



Cisco Unified Communications Solution Reference Network Design (SRND)

Based on Cisco Unified CallManager Release 4.x
October 30, 2008

Americas 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: 408 527-0883

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS.

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

CCDE, CCENT, Cisco Eos, Cisco Lumin, Cisco Nexus, Cisco StadiumVision, Cisco TelePresence, Cisco WebEx, the Cisco logo, DCE, and Welcome to the Human Network are trademarks; Changing the Way We Work, Live, Play, and Learn and Cisco Store are service marks; and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MeetingPlace Chime Sound, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0809R)

Cisco Unified Communications SRND (Based on Cisco Unified CallManager 4.x)
Copyright © 2005 Cisco Systems, Inc. All rights reserved.



CONTENTS

Preface xxi

CHAPTER 1

Introduction 1-1

- Overview of Cisco Unified Communications 1-2
 - Cisco IP Network Infrastructure 1-3
 - Quality of Service 1-4
 - Call Processing Agent 1-4
 - Communication Endpoints 1-5
 - Conferencing, Messaging, and Collaboration Capabilities 1-6
 - Security 1-6

CHAPTER 2

IP Telephony Deployment Models 2-1

- Single Site 2-2
 - Benefits of the Single-Site Model 2-3
 - Best Practices for the Single-Site Model 2-4
- Multisite WAN with Centralized Call Processing 2-4
 - Best Practices for the Centralized Call Processing Model 2-7
 - Remote Site Survivability 2-8
 - Voice Over the PSTN as a Variant of Centralized Call Processing 2-10
 - VoPSTN Using AAR 2-12
 - VoPSTN Using Dial Plan 2-13
- Multisite WAN with Distributed Call Processing 2-14
 - Benefits of the Distributed Call Processing Model 2-16
 - Best Practices for the Distributed Call Processing Model 2-16
 - Call Processing Agents for the Distributed Call Processing Model 2-17
- Clustering Over the IP WAN 2-17
 - WAN Considerations 2-18
 - Intra-Cluster Communications 2-19
 - Failover Between Subscriber Servers 2-20
 - Cisco Unified CallManager Publisher 2-20
 - Call Detail Records (CDR) 2-20
 - Delay Testing 2-20
 - Error Rate 2-21
 - Troubleshooting 2-21

Local Failover Deployment Model	2-22
Cisco Unified CallManager Provisioning for Local Failover	2-24
Gateways for Local Failover	2-25
Voicemail for Local Failover	2-25
Music on Hold and Media Resources for Local Failover	2-25
Remote Failover Deployment Model	2-26
Design Considerations for Section 508 Conformance	2-27

CHAPTER 3

Network Infrastructure 3-1

What's New in This Chapter	3-3
LAN Infrastructure	3-4
LAN Design for High Availability	3-4
Campus Access Layer	3-4
Campus Distribution Layer	3-7
Campus Core Layer	3-9
Network Services	3-10
Power over Ethernet (PoE)	3-19
Category 3 Cabling	3-20
IBM Type 1A and 2A Cabling	3-20
LAN Quality of Service (QoS)	3-21
Traffic Classification	3-22
Interface Queuing	3-24
Bandwidth Provisioning	3-24
Impairments to IP Communications if QoS is Not Employed	3-24
WAN Infrastructure	3-25
WAN Design and Configuration	3-25
Deployment Considerations	3-25
Guaranteed Bandwidth	3-27
Best-Effort Bandwidth	3-27
WAN Quality of Service (QoS)	3-28
Traffic Prioritization	3-29
Link Efficiency Techniques	3-31
Traffic Shaping	3-32
Resource Reservation Protocol (RSVP)	3-35
RSVP Principles	3-35
RSVP and QoS in WAN Routers	3-38
RSVP Application ID	3-42
RSVP Design Best Practices	3-44
Bandwidth Provisioning	3-44

Provisioning for Bearer Traffic	3-46
Additional Considerations for Bearer Traffic with RSVP	3-49
Provisioning for Call Control Traffic with Centralized Call Processing	3-53
Provisioning for Call Control Traffic with Distributed Call Processing	3-57
Wireless LAN Infrastructure	3-58
WLAN Design and Configuration	3-58
Wireless Infrastructure Considerations	3-58
Wireless AP Configuration and Design	3-61
Wireless Security	3-62
WLAN Quality of Service (QoS)	3-64
Traffic Classification	3-64
Interface Queuing	3-65
Bandwidth Provisioning	3-65

CHAPTER 4

Gateways 4-1

What's New in This Chapter	4-1
Understanding Cisco Gateways	4-2
Cisco Access Analog Gateways	4-2
Cisco Access Digital Trunk Gateways	4-2
Gateway Selection	4-2
Core Feature Requirements	4-3
Gateway Protocols	4-3
Gateway Protocols and Core Feature Requirements	4-6
DTMF Relay	4-6
Supplementary Services	4-7
Cisco Unified CallManager Redundancy	4-10
Call Survivability	4-12
Site-Specific Gateway Requirements	4-12
QSIG Support	4-19
Fax and Modem Support	4-20
Gateway Support for Fax Pass-Through and Cisco Fax Relay	4-20
Gateway Support for Modem Pass-Through	4-21
Supported Platforms and Features	4-22
Platform Protocol Support	4-23
Gateway Combinations and Interoperability of Features	4-24
Feature Support Between Similar Gateways	4-25
Gateway Configuration Examples	4-26
Cisco IOS Gateway Configuration	4-26
Cisco VG248 Configuration	4-26

Cisco Unified CallManager Configuration for Cisco IOS Gateways	4-27
Clock Sourcing for Fax and Modem Pass-Through	4-27
T.38 Fax Relay	4-28
Loose Gateway Controlled with Named Service Event (NSE)	4-28
Gateway Controlled with Capability Exchange Through H.245 or Session Description Protocol (SDP)	4-29
Call-Agent-Controlled T.38 with H.323 Annex D	4-30
Gateways for Video Telephony	4-31
Routing Inbound Calls from the PSTN	4-33
Routing Outbound Calls to the PSTN	4-34
Automated Alternate Routing (AAR)	4-35
Least-Cost Routing	4-37
ISDN B-Channel Binding, Rollover, and Busy Out	4-38
Inbound Calls	4-38
Outbound Calls	4-39
Configuring the Gateways in Cisco Unified CallManager	4-40
Call Signaling Port Numbers	4-40
Call Signaling Timers	4-41
Voice Gateways Bearer Capabilities	4-41

CHAPTER 5

Cisco Unified CallManager Trunks 5-1

H.323 Trunks	5-2
Gatekeeper Controlled Trunks	5-2
Non-Gatekeeper Controlled Trunks	5-3
Gatekeeper Trunk Redundancy, Resilience, and Load Balancing	5-4
H.323 Trunks with Media Termination Points	5-8
H.323 Operation in Cisco Unified CallManager	5-9
SIP Trunks	5-11

CHAPTER 6

Media Resources 6-1

Voice Termination	6-2
Medium and High Complexity Mode	6-2
Flex Mode	6-2
DSP Resources for Voice Termination	6-3
Audio Conferencing	6-7
Audio Conferencing Resources	6-7
Transcoding	6-10
Transcoding Resources	6-11
Media Termination Point (MTP)	6-12

Re-Packetization of a Stream	6-12
H.323 Supplementary Services	6-12
H.323 Outbound Fast Start	6-12
Named Telephony Events (RFC 2833)	6-13
When are Media Termination Points Required for Named Telephony Events?	6-13
Configuration of DTMF on SIP and H.323 Gateways	6-15
CTI Route Points	6-16
MTP Resources	6-16
Annunciator	6-17
Cisco RSVP Agent	6-18
Cisco IP Voice Media Streaming Application	6-19
Hardware and Software Capacities	6-20
PVDMS	6-20
Cisco 2800 and 3800 Series Platforms	6-21
Network Modules	6-21
Calculating DSP Requirements for the NM-HDV	6-22
General Design Guidelines	6-22
Media Resource Groups and Lists	6-22
Deployment Models	6-23
IP PSTN Access	6-26

CHAPTER 7
Music on Hold 7-1

Deployment Basics of MoH	7-2
Unicast and Multicast MoH	7-2
Co-resident and Standalone MoH Servers	7-3
Fixed and Audio File MoH Sources	7-3
MoH Server as Part of the Cisco Unified CallManager Cluster	7-4
Basic MoH and MoH Call Flows	7-5
Basic MoH	7-5
User and Network Hold	7-6
Unicast and Multicast MoH Call Flows	7-8
MoH Configuration Considerations and Best Practices	7-8
Codec Selection	7-8
Multicast Addressing	7-8
MoH Audio Sources	7-9
Using Multiple Fixed or Live Audio Sources	7-9
Unicast and Multicast in the Same Cisco Unified CallManager Cluster	7-10
Redundancy	7-11
Quality of Service (QoS)	7-11

Hardware and Capacity Planning for MoH Resources	7-12
Server Platform Limits	7-12
Resource Provisioning and Capacity Planning	7-14
Implications for MoH With Regard to IP Telephony Deployment Models	7-15
Single-Site Campus (Relevant to All Deployments)	7-15
Centralized Multisite Deployments	7-15
Call Admission Control and MoH	7-16
Multicast MoH from Branch Router Flash	7-17
Distributed Multisite Deployments	7-19
Clustering Over the WAN	7-20
Detailed Unicast and Multicast MoH Call Flows	7-20

CHAPTER 8

Call Processing 8-1

Cisco Unified CallManager Cluster Guidelines	8-2
Hardware Platforms	8-2
Cisco Unified CallManager Cluster Services	8-3
Intracenter Communications	8-4
Publisher	8-5
Call Processing Subscriber	8-6
TFTP Server	8-10
CTI Manager	8-10
IP Voice Media Streaming Application	8-11
Voice Activity Detection	8-11
Extension Mobility	8-11
Cisco Unified CallManager Platform Capacity Planning	8-12
Capacity Calculations	8-13
Cisco CallManager Capacity Tool	8-14
Gatekeeper Design Considerations	8-18
Hardware Platform Selection	8-18
Gatekeeper Redundancy	8-18
Hot Standby Router Protocol (HSRP)	8-19
Gatekeeper Clustering (Alternate Gatekeeper)	8-21
Directory Gatekeeper Redundancy	8-24
Interoperability of Cisco Unified CallManager and CallManager Express	8-27
Multisite IP Telephony Deployments Using Cisco Unified CallManager and CME, with H.323 Trunks and an IP-to-IP Gateway	8-28

CHAPTER 9

Call Admission Control 9-1

Best Practices Summary	9-2

Call Admission Control Principles	9-3
Topology-Unaware Call Admission Control	9-3
Topology-Aware Call Admission Control	9-7
Special Considerations for MPLS Networks	9-11
Call Admission Control Elements	9-11
Cisco Unified CallManager Static Locations	9-12
Cisco IOS Gatekeeper Zones	9-15
Cisco IOS Gatekeeper and IP-to-IP Gateway with RSVP	9-17
Via-Zone Gatekeeper	9-18
Design Best Practices	9-19
Redundancy	9-20
Configuration Guidelines	9-21
Call Admission Control Design	9-25
Simple Hub-and-Spoke Topologies	9-25
Centralized Cisco Unified CallManager Deployments	9-26
Distributed Cisco Unified CallManager Deployments	9-27
Two-Tier Hub-and-Spoke Topologies	9-29
Centralized Cisco Unified CallManager Deployments	9-30
Distributed Cisco Unified CallManager Deployments	9-33
Simple MPLS Topologies	9-34
Centralized Cisco Unified CallManager Deployments	9-36
Distributed Cisco Unified CallManager Deployments	9-39
Generic Topologies	9-41
Case Study	9-42

CHAPTER 10

Dial Plan 10-1

Planning Considerations	10-2
Dialed Pattern Recognition	10-2
On-Net versus Off-Net Dialing	10-3
Abbreviated Dialing	10-3
Avoiding Overlap of Extension Dialing	10-4
Dialing String Length	10-4
Uniform On-Net Dial Plan	10-4
Variable Length On-Net Dial Plan	10-6
On-Net and Off-Net Access Codes	10-7
Plan Ahead	10-7
Dial Plan Elements	10-7
User Input on SCCP Phones	10-8
Call Routing in Cisco Unified CallManager	10-9

Route Patterns	10-11
Route Lists	10-14
Route Groups	10-14
Route Group Devices	10-15
Calling Privileges in Cisco Unified CallManager	10-15
Partitions	10-16
Calling Search Spaces	10-17
Digit Manipulation in Cisco Unified CallManager	10-21
Automated Alternate Routing	10-22
Establish the Phone Number of the Destination	10-22
Prefix the Required Access Codes	10-23
Voicemail Considerations	10-24
Select the Proper Dial Plan and Route	10-24
Special Considerations for Sites Located Within the Same Local Dialing Area	10-25
Device Mobility	10-26
Extension Mobility	10-27
Hunt Lists and Line Groups	10-29
Hunt Pilot	10-31
Hunt List	10-32
Line Group	10-32
Hunt Group Logout	10-33
Line Group Devices	10-33
Time-of-Day Routing	10-34
Call Routing in Cisco IOS with H.323 Dial Peers	10-35
Call Routing in Cisco IOS with a Gatekeeper	10-38
Centralized Gatekeeper Configuration	10-42
Distributed Gatekeeper Configuration	10-44
Distributed Gatekeeper Configuration with Directory Gatekeeper	10-45
Calling Privileges in Cisco IOS with H.323 Dial Peers	10-47
Digit Manipulation in Cisco IOS with H.323 Dial Peers	10-49
Design Considerations	10-51
Design Guidelines for Multisite Deployments	10-52
Choosing a Dial Plan Approach	10-54
Deploying Uniform On-Net Dial Plans	10-56
Inter-Site Calls Within a Cluster	10-57
Outgoing PSTN and IP WAN Calls	10-57
Emergency Calls	10-57
Incoming Calls	10-57
Voicemail Calls	10-58
Deploying Variable-Length On-Net Dial Plans with Partitioned Addressing	10-58

Inter-Site Calls Within a Cluster	10-61
Outgoing PSTN and IP WAN Calls	10-61
Incoming Calls	10-63
Voicemail Calls	10-63
Deploying Variable-Length On-Net Dial Plans with Flat Addressing	10-63
Inter-Site Calls Within a Cluster	10-66
Outgoing PSTN and IP WAN Calls	10-67
Incoming Calls	10-70
Voicemail Calls	10-70
Special Considerations for Deployments Without Site Codes	10-70
Building Classes of Service for Cisco Unified CallManager	10-72
Classes of Service with the Traditional Approach	10-72
Classes of Service with the Line/Device Approach	10-76
Guidelines for the Line/Device Approach	10-80
Extension Mobility Considerations for the Line/Device Approach	10-81
Building Classes of Service in Cisco IOS with H.323	10-84
Deploying Call Coverage	10-87
Deploying Call Coverage in a Multisite Centralized Call Processing Model	10-88
Deploying Call Coverage in a Multisite Distributed Call Processing Model	10-89
Hunt Pilot Scalability	10-90

CHAPTER 11

Emergency Services 11-1

Planning for 911 Functionality	11-1
Public Safety Answering Point (PSAP)	11-2
911 Network Service Provider	11-2
Interface Points into the Appropriate 911 Networks	11-3
Interface Type	11-4
Dynamic ANI (Trunk Connection)	11-4
Static ANI (Line Connection)	11-6
Emergency Response Location Mapping	11-6
Emergency Location Identification Number Mapping	11-7
Nomadic Phone Considerations	11-9
Cisco Emergency Responder	11-9
Emergency Call String	11-10
Gateway Considerations	11-11
Gateway Placement	11-11
Gateway Blocking	11-11
Answer Supervision	11-12
Cisco Emergency Responder Considerations	11-13

Emergency Responder Version Compatibility with Cisco Unified CallManager	11-13
Device Mobility Across Call Admission Control Locations	11-13
Default Emergency Response Location	11-13
Soft Clients	11-14
Test Calls	11-14
PSAP Callback to Shared Directory Numbers	11-14
Multi-Cluster Considerations	11-14
Single Cisco ER Group	11-15
Multiple Cisco ER Groups	11-16
Emergency Call Routing within a Cisco ER Cluster	11-18
Scalability Considerations for Cisco ER Clustering	11-19
ALI Formats	11-19

CHAPTER 12

Third-Party Voicemail Design 12-1

SMDI	12-1
Cisco Messaging Interface	12-2
Cisco VG248	12-2
Considerations When Using FXS Ports	12-3
Digital Set Emulation	12-4
Dual PBX Integration	12-5
Centralized Voicemail	12-6
Positive Disconnect Supervision	12-11
Summary of Third-Party Voicemail Integration	12-12

CHAPTER 13

Cisco Unity 13-1

Messaging Deployment Models	13-2
Single-Site Messaging	13-2
Centralized Messaging	13-2
Distributed Messaging	13-3
Messaging Failover	13-4
Messaging System Infrastructure Components	13-5
Port Groups (Separate Integrations)	13-5
Managing Bandwidth	13-6
Cisco Unity and Unity Connection Native Transcoding Operation	13-8
Voice Port Integration with a Cisco Unified CallManager Cluster	13-9
Voice Port Integration with Dedicated Cisco Unified CallManager Backup Servers	13-11
Centralized Messaging and Centralized Call Processing	13-12
Distributed Messaging with Centralized Call Processing	13-13

Combined Messaging Deployment Models	13-15
Centralized Messaging with Clustering Over the WAN	13-16
Distributed Messaging with Clustering Over the WAN	13-18
Cisco Unity Messaging Failover	13-19
Cisco Unity Failover and Clustering Over the WAN	13-19
Centralized Messaging with Multiple Cisco Unified CallManager Servers	13-20

CHAPTER 14
Cisco Unity Express 14-1

Deployment Models for Cisco Unity Express	14-2
Best Practices for Deploying Cisco Unity Express	14-5

CHAPTER 15
Cisco Unified MeetingPlace Integration 15-1

MeetingPlace Server Recommendations	15-2
Deployment Models	15-2
MeetingPlace Components	15-5
MeetingPlace Audio Server	15-5
MeetingPlace H.323/SIP IP Gateway	15-5
Sizing a MeetingPlace Deployment	15-7
Sizing Voice Conferencing Usage	15-8
Sizing Web Conferencing Usage	15-8
Video Integration and Sizing the MCU	15-9
MeetingPlace Network Infrastructure	15-10
Connectivity Between MeetingPlace and IP Telephony Components	15-11
Quality of Service (QoS)	15-11
Traffic Classification	15-11
Call Admission Control and Bandwidth Provisioning	15-12
Rate and Codec Selection for Audio and Video	15-12
MeetingPlace Web Session Network Utilization	15-13
Jitter	15-14
Domain Name System (DNS)	15-14
Network Time Protocol (NTP)	15-14
Demilitarized Zone (DMZ) Requirements	15-14
Interoperability Protocols	15-17
Internet Protocol (IP) Network	15-17
Protocols Supported by the MeetingPlace Audio Server	15-17
Protocols Supported by Other MeetingPlace Components	15-18
Public Switched Telephone Network (PSTN)	15-21
Digital Trunks	15-21

Conferencing	15-23
Audio Conferencing	15-23
Call Flow	15-23
Meeting Type	15-24
Port Management	15-25
Scheduling	15-25
Audio Conference Cascading	15-25
Dialing into an Audio-Only Conference Using Video Endpoints	15-26
Web Conferencing	15-26
MeetingPlace Web Server	15-26
SQL Database	15-27
Meeting Type	15-28
Web Conference Cascading	15-28
Segmented Meetings	15-28
Video Conferencing	15-29
Voice Link	15-29
MeetingPlace Video	15-29
MCU Configuration	15-30
Enhanced Media Processor (EMP) Requirement	15-31
Port Management	15-31
Scheduling	15-32
Attending a Video Conference	15-32
Meeting Type	15-33
Video Conference Call Flow	15-34
Video Conference Cascading	15-37
Gatekeeper and Dial Plan	15-37
Dynamic H.323 Addressing in Cisco Unified CallManager	15-37
Cisco Unified CallManager Redundancy Groups and H.323 Clients	15-38
MCU Registration	15-38
MeetingPlace	15-38
Reservationless Single Number Access (RSNA)	15-38
Redundancy and Load Balancing	15-39
MeetingPlace Audio Server	15-39
MeetingPlace H.323/SIP IP Gateway	15-40
MeetingPlace Web	15-41
Cisco Unified CallManager	15-41
MeetingPlace Video	15-42
MCU	15-42

CHAPTER 16**Cisco Unified MeetingPlace Express 16-1**

Deployment Models	16-1
Single Site	16-2
Multisite WAN with Centralized Call Processing	16-3
Multisite WAN with Distributed Call Processing	16-4
Clustering Over the WAN	16-5
Security and DMZ Deployment	16-5
Call Admission Control, Bandwidth, and QoS	16-7
Call Admission Control	16-7
Bandwidth Considerations for Web Applications and Screen Sharing	16-7
QoS	16-8
Directory Integration	16-9
H.323 and SIP Direct Integration	16-9
H.323 Gateway	16-9
SIP Trunk	16-10
Gatekeeper Integration	16-11
Design Considerations	16-11
Redundancy	16-12
Cisco Unified MeetingPlace Express Server Redundancy	16-12
Physical Network Redundancy	16-12
Redundancy Using a Gatekeeper	16-12
Redundancy Using H.323 Gateway Integration	16-13
Redundancy Using SIP Trunk Integration	16-13
Other Important Design Considerations	16-14

CHAPTER 17**IP Video Telephony 17-1**

IP Video Telephony Solution Components	17-1
Administration Considerations	17-2
Protocols	17-2
Regions	17-3
Locations	17-6
Retry Video Call as Audio	17-7
Wait for Far-End to Send TCS	17-9
Multipoint Conferencing	17-11
SCCP MCU Resources	17-14
Media Resource Groups and Lists	17-15
H.323 MCU Resources	17-16
Sizing the MCU	17-18

IVR for Dial-In Conference	17-18
Gatekeepers	17-19
Supported Gatekeeper Platforms	17-22
Endpoint Gatekeepers	17-23
Provisioning H.323 Clients	17-24
Provisioning H.323 MCUs	17-29
Provisioning H.320 Gateways	17-30
Gatekeeper Zone Configuration	17-31
Summary of Endpoint Gatekeepers	17-39
Migrating from Cisco Unified CallManager 4.0	17-42
Applications	17-42
CTI Applications	17-43
Cisco Emergency Responder	17-43
Cisco Unified CallManager Assistant	17-43
Cisco IP Interactive Voice Response and Cisco IP Contact Center	17-44
Cisco Attendant Console	17-44
Cisco Personal Assistant	17-44
Cisco IP SoftPhone and Cisco IP Communicator	17-45
Collaboration Solutions	17-45
T.120 Application Sharing	17-45
Cisco Unified MeetingPlace	17-45
Wireless Networking Solutions	17-45
Cisco Aironet 802.11b Networking Solutions	17-45
Cisco Unified Wireless IP Phone 7920	17-46
XML Services	17-46

CHAPTER 18

LDAP Directory Integration 18-1

What is Directory Integration?	18-2
Directory Access for IP Telephony Endpoints	18-3
Directory Integration with Cisco Unified CallManager 4.x	18-6
Cisco Unified CallManager 4.x Directory Architecture	18-6
The Cisco Customer Directory Configuration Plugin	18-9
Security Considerations	18-10
Multilevel Administration Account Security	18-10
Adding Cisco Unified CallManager Servers to a Domain	18-11
Comparison with the Cisco Unified CallManager 5.0 Approach	18-12
Best Practices for Directory Integration	18-13
Planning the Directory Integration	18-13
Preparing the Directory for Integration	18-15

Integrating Cisco Unified CallManager with the Directory	18-19
Maintaining the Directory Integration	18-22

CHAPTER 19

IP Telephony Migration Options 19-1

Phased Migration	19-1
Parallel Cutover	19-2
The Need for QSIG in Multiple-Site Enterprises	19-3
Summary	19-5

CHAPTER 20

Voice Security 20-1

General Security	20-1
Security Policy	20-1
Security in Layers	20-2
Secure Infrastructure	20-3
Physical Security	20-4
IP Addressing	20-4
Phone Security	20-5
PC Port on the Phone	20-5
Gratuitous ARP	20-6
PC Voice VLAN Access	20-7
Web Access	20-8
Video Capabilities	20-9
Settings Access	20-9
Phone Authentication and Encryption	20-10
Access Security	20-11
Voice and Video VLANs	20-11
Switch Port	20-12
Port Security: MAC CAM Flooding	20-13
Port Security: Prevent Port Access	20-14
Port Security: Prevent Rogue Network Extensions	20-14
DHCP Snooping: Prevent Rogue DHCP Server Attacks	20-15
DHCP Snooping: Prevent DHCP Starvation Attacks	20-17
DHCP Snooping: Binding Information	20-19
Requirement for Dynamic ARP Inspection	20-20
Quality of Service	20-24
Access Control Lists	20-24
VLAN Access Control Lists	20-24
Router Access Control Lists	20-26

Gateways and Media Resources	20-28
Putting Firewalls Around Gateways	20-29
Firewalls	20-31
Routed ASA and PIX	20-33
Transparent ASA and PIX	20-34
ASA and PIX Configuration Example	20-35
FWSM Routed Mode	20-36
FWSM Transparent Mode	20-37
FWSM Configuration Example	20-37
Data Center	20-39
Applications Servers	20-39
Cisco Security Agent on the Cisco Unified CallManager and Application Servers	20-39
Unmanaged Cisco Security Agent	20-39
Managed Cisco Security Agent	20-40
Antivirus	20-40
General Server Guidelines	20-41
Deployment Examples	20-42
Lobby Phone Example	20-42
Firewall Deployment Example (Centralized Deployment)	20-44
Conclusion	20-45

CHAPTER 21

IP Telephony Endpoints 21-1

Analog Gateways	21-2
Analog Interface Module	21-2
Low-Density Analog Interface Module	21-2
High-Density Analog Interface Module	21-3
Supported Platforms and Cisco IOS Requirements for Analog Interface Modules	21-4
Cisco Communication Media Module (CMM)	21-5
WS-X6624-FXS Analog Interface Module	21-5
Cisco VG224 Gateway	21-5
Cisco VG248 Gateway	21-5
Cisco ATA 186 and 188	21-6
Cisco Unified IP Phones	21-6
Cisco Basic IP Phones	21-6
Cisco Business IP Phones	21-6
Cisco Manager IP Phones	21-6
Cisco Executive IP Phones	21-7
Cisco Unified IP Phone Expansion Module 7914	21-7
Software-Based Endpoints	21-7

Cisco IP Communicator	21-7
Maximum IP Communicator Configuration Limits	21-8
Codec Selection	21-8
Call Admission Control	21-8
Cisco IP SoftPhone	21-9
Maximum Cisco IP SoftPhone Configuration Limits	21-10
Codec Selection	21-10
Call Admission Control	21-11
Wireless Endpoints	21-12
Site Survey	21-12
Authentication	21-12
Capacity	21-13
Phone Configuration	21-14
Roaming	21-14
AP Call Admission Control	21-15
Cisco IP Conference Station	21-16
Video Endpoints	21-16
SCCP Video Endpoints	21-16
Cisco Unified Video Advantage	21-16
Cisco IP Video Phone 7985G	21-19
Codecs Supported by Cisco Unified Video Advantage and Cisco IP Video Phone 7985G	21-20
Third-Party SCCP Video Endpoints	21-20
QoS Recommendations	21-21
Cisco VG224 and VG248	21-21
Cisco ATA 186 and IP Conference Station	21-22
Cisco ATA 188 and IP Phones	21-22
Software-Based Endpoints	21-26
Cisco Unified Wireless IP Phone 7920	21-30
Video Telephony Endpoints	21-32
Cisco Unified Video Advantage with a Cisco Unified IP Phone	21-33
Cisco IP Video Phone 7985G	21-33
Sony and Tandberg SCCP Endpoints	21-34
H.323 Video Endpoints	21-34
Endpoint Features Summary	21-35

CHAPTER 22

Device Mobility 22-1

Need for Device Mobility	22-2
Device Mobility Feature	22-3
Dial Plan Design Considerations	22-7

Device Mobility Considerations for Building Classes of Service	22-7
Traditional Approach	22-8
Line/Device Approach	22-9
Choosing a Dial Plan Model	22-11
Uniform On-Net Dialing Using the Line/Device Approach	22-11
Variable Length On-Net Dialing with Partitioned Addressing Using the Line/Device Approach	22-13
Variable Length On-Net Dialing with Flat Addressing Using the Line/Device Approach	22-15
Design Guidelines for Using a VPN	22-16

APPENDIX A

Recommended Hardware and Software Combinations A-1

GLOSSARY

INDEX



Preface

Last revised on: October 30, 2008

This document provides design considerations and guidelines for deploying a Cisco Unified Communications System based on Cisco Unified CallManager 4.x releases.

This document focuses on the following components of the Cisco Unified Communications System:

- Cisco Unified CallManager
- Cisco Unified Video Advantage
- Cisco Unified MeetingPlace and Cisco Unified MeetingPlace Express

This document should be used in conjunction with other documentation available at the following locations:

- For more information about the Cisco Unified Communications System:
<http://www.cisco.com/go/unified-techinfo>
<http://www.cisco.com>
- For more information about Cisco Unified CallManager:
http://www.cisco.com/en/US/products/sw/voicesw/ps556/tsd_products_support_series_home.html
<http://www.cisco.com>
- For other Solution Reference Network Design (SRND) documents:
<http://www.cisco.com/go/designzone>

New or Changed Information for This Release



Note

Unless stated otherwise, the information in this document applies to Cisco Unified CallManager 4.0 and subsequent 4.x releases. Any differences between the various releases of Cisco Unified CallManager are specifically noted in the text.

The following chapters are either new in the current release of this document, or they contain information that has changed significantly from previous releases of this document.

- [Network Infrastructure, page 3-1](#)
- [Gateways, page 4-1](#)

Within each chapter, new and revised information is listed in a section titled *What's New in This Chapter*.

Revision History

This document may be updated at any time without notice. You can obtain the latest version of this document online at

<http://www.cisco.com/go/designzone>

Visit this Cisco.com website periodically and check for documentation updates by comparing the revision date on the front title page of your copy with the revision date of the online document.

The following table lists the revision history for this document.

Revision Date	Comments
October 30, 2008	Updated content as indicated in New or Changed Information for This Release , page xxi.
February 13, 2008	This document was updated to correct some errors in various chapters.
July, 2006	This document was updated and modified to apply to Cisco Unified CallManager releases 4.0 through 4.2.
February, 2006	This document was updated to include a new chapter on voice security as well as a few other minor changes.
August, 2005	This document was updated and modified to apply to both Cisco CallManager releases 4.0 and 4.1.
April, 2005	Initial version of this document for Cisco CallManager Release 4.1.

Obtaining Documentation, Obtaining Support, and Security Guidelines

For information on obtaining documentation, obtaining support, providing documentation feedback, security guidelines, and also recommended aliases and general Cisco documents, see the monthly *What's New in Cisco Product Documentation*, which also lists all new and revised Cisco technical documentation, at:

<http://www.cisco.com/en/US/docs/general/whatsnew/whatsnew.html>



CHAPTER 1

Introduction

Last revised on: February 13, 2008

The Cisco Unified Communications System delivers fully integrated communications by enabling data, voice, and video to be transmitted over a single network infrastructure using standards-based Internet Protocol (IP). Leveraging the framework provided by Cisco IP hardware and software products, the Cisco Unified Communications System delivers unparalleled performance and capabilities to address current and emerging communications needs in the enterprise environment. The Cisco Unified Communications family of products is designed to optimize feature functionality, reduce configuration and maintenance requirements, and provide interoperability with a wide variety of other applications. The Cisco Unified Communications System provides this capability while maintaining a high level of availability, quality of service (QoS), and security for your network.

The Cisco Unified Communications System incorporates and integrates the following major communications technologies:

- **IP telephony**
IP telephony refers to technology that transmits voice communications over a network using IP standards. Cisco Unified Communications includes a wide array of hardware and software products such as call processing agents, IP phones (both wired and wireless), voice messaging systems, video devices, and many special applications.
- **Customer contact center**
Cisco IP Contact Center products are a combination of strategy and architecture that promote efficient and effective customer communications across a globally capable network by enabling organizations to draw from a broader range of resources to service customers. They include access to a large pool of agents and multiple channels of communication as well as customer self-help tools.
- **Video telephony**
The Cisco Unified Video Advantage products enable real-time video communications and collaboration using the same IP network and call processing agent as Cisco Unified Communications. With Cisco Unified Video Advantage, making a video call is now as easy as dialing a phone number.
- **Rich-media conferencing**
Cisco Conference Connection and Cisco Unified MeetingPlace enhance the virtual meeting environment with a integrated set of IP-based tools for voice, video, and web conferencing.

- Third-party applications

Cisco works with leading-edge companies to provide the broadest selection of innovative third-party IP communications applications and products focused on critical business needs such as messaging, customer care, and workforce optimization.

The remainder of this document focuses on system design considerations for the following components of the Cisco Unified Communications System:

- Cisco Unified CallManager
- Cisco Unified Video Advantage
- Cisco Unified MeetingPlace

For information about other aspects of the Cisco Unified Communications System, such as Cisco IP Contact Center, refer to the documentation available at the following locations:

<http://www.cisco.com/go/designzone>

<http://www.cisco.com/go/unified-techinfo>


You can also find additional documentation for the Cisco Unified Communications family of products at the following location:

<http://www.cisco.com>

Overview of Cisco Unified Communications

Cisco IP Telephony is the leading converged network IP telephony solution for organizations that want to increase productivity and reduce the costs associated with managing and maintaining separate voice and data networks. The flexibility and sophisticated functionality of the Cisco IP network infrastructure provides the framework that permits rapid deployment of emerging applications such as desktop IP Telephony, unified messaging, video telephony, desktop collaboration, enterprise application integration with IP phone displays, and collaborative IP contact centers. These applications enhance productivity and increase enterprise revenues.

Figure 1-1 illustrates a typical IP Telephony deployment employing the Cisco IP network infrastructure, with Cisco Unified CallManager as the call processing agent.

Figure 1-1 Typical IP Telephony Deployment

The foundation architecture for Cisco IP Telephony includes the following major components (see Figure 1-1):

- [Cisco IP Network Infrastructure, page 1-3](#)
- [Quality of Service, page 1-4](#)
- [Call Processing Agent, page 1-4](#)
- [Communication Endpoints, page 1-5](#)
- [Conferencing, Messaging, and Collaboration Capabilities, page 1-6](#)
- [Security, page 1-6](#)

Cisco IP Network Infrastructure

The network infrastructure includes public switched telephone network (PSTN) gateways, analog phone support, and digital signal processor (DSP) farms. The infrastructure can support multiple client types such as hardware phones, software phones, and video devices. The infrastructure also includes the

interfaces and features necessary to integrate legacy PBX, voicemail, and directory systems. Typical products used to build the infrastructure include Cisco voice gateways (non-routing, routing, and integrated), Cisco IOS and Catalyst switches, and Cisco routers.

For more information about the IP network infrastructure, see the chapter on [Network Infrastructure](#), page 3-1.

Quality of Service

Voice, as a class of IP network traffic, has strict requirements concerning packet loss, delay, and delay variation (also known as jitter). To meet these requirements for voice traffic, Cisco Unified Communications includes Quality of Service (QoS) features such as traffic classification, queuing, traffic shaping, compressed Real-Time Transport Protocol (cRTP), and Transmission Control Protocol (TCP) header compression.

The QoS components for Cisco Unified Communications are provided through the rich IP traffic management, queuing, and shaping capabilities of the Cisco IP network infrastructure. Key elements of this infrastructure that enable QoS for IP Telephony include:

- Traffic marking
- Enhanced queuing services
- Link fragmentation and interleaving (LFI)
- Compressed RTP (cRTP)
- Low-latency queuing (LLQ)
- Link efficiency
- Traffic shaping
- Call admission control (bandwidth allocation)

For more information about QoS, see the various sections on QoS in the chapter on [Network Infrastructure](#), page 3-1.

Call Processing Agent

Cisco Unified CallManager is the core call processing software for Cisco IP Telephony. It builds call processing capabilities on top of the Cisco IP network infrastructure. Cisco Unified CallManager software extends enterprise telephony features and capabilities to packet telephony network devices such as IP phones, media processing devices, voice gateways, and multimedia applications.

You can deploy the call processing capabilities of Cisco Unified CallManager according to one of the following models, depending on the size, geographical distribution, and functional requirements of your enterprise:

- Single-site call processing model

In the single-site model, each site or campus has its own Cisco Unified CallManager or cluster of Cisco Unified CallManagers to perform call processing functions. No voice traffic travels over the IP WAN; instead, external calls or calls to remote sites use the public switched telephone network (PSTN).

- Multisite WAN model with centralized call processing

In the multisite WAN model with centralized call processing, the Cisco Unified CallManager cluster resides at the main (or central) campus, and communication with remote branch offices normally takes place over the IP WAN. If either the central site or the IP WAN is down, the remote sites can continue to have service through a feature called Survivable Remote Site Telephony (SRST). The remote sites can also place calls over the PSTN if the IP WAN is temporarily oversubscribed, and you can interconnect a number of central sites through intercluster trunks.

- Multisite WAN model with distributed call processing

In the multisite WAN model with distributed call processing, each site has its own Cisco Unified CallManager cluster for call processing. Communication between sites normally takes place over the IP WAN, with the PSTN serving as a backup voice path. With this model, you can interconnect any number of sites across the IP WAN.

- Clustering over the IP WAN

You may deploy a single Cisco Unified CallManager cluster across multiple sites that are connected by an IP WAN with QoS features enabled. To provide call processing redundancy, you can deploy backup servers either locally at each site or at a remote site across the WAN. Clustering over the WAN is well suited as a disaster recovery plan for business continuance sites or as a single solution for small or medium sites.

For information on how to apply these deployment models in designing your Cisco IP Telephony network, see [IP Telephony Deployment Models, page 2-1](#).

Communication Endpoints

A communication endpoint is a user instrument such as a desk phone or even a software phone application that runs on a PC. In the IP environment, each phone has an Ethernet connection. IP phones have all the functions you expect from a telephone, as well as more advanced features such as the ability to access World Wide Web sites.

In addition to various models of desktop Cisco Unified IP Phones, IP Telephony endpoints include the following devices:

- Software-based endpoints

Cisco IP Communicator and Cisco Unified Personal Communicator are desktop applications that turn your computer into a full-featured IP phone with the added advantages of call tracking, desktop collaboration, and one-click dialing from online directories. Cisco IP Communicator is a software-based application that delivers enhanced telephony support through the PC. It is designed to meet diverse customer needs by serving as a supplemental telephone when traveling, as a telecommuting device, or as a primary desktop telephone. Cisco Unified Personal Communicator integrates a wide variety of communications applications and services into a single desktop computer application to provide quick and easy access to powerful communications tools for voice, video, web conferencing, call management, directories, and presence information.

- Video telephony endpoints

Video telephony capability is now fully integrated with Cisco Unified CallManager. In addition, Cisco Unified Video Advantage brings video telephony functionality to Cisco Unified IP Phones and the Cisco IP Communicator softphone application. The video telephony solution consists of a Windows-based application and USB camera. Users make calls from their Cisco Unified IP Phones using the familiar phone interface, and calls are displayed with video on their PCs without requiring any extra button pushes or mouse clicks.

- Wireless endpoints

The Cisco Wireless IP Phone 7920 extends the Cisco family of IP Phones from 10/100 Ethernet to 802.11 wireless LAN (WLAN). The Cisco Wireless IP Phone 7920 provides multiple line appearances with functionality similar to other Cisco Unified IP Phones. In addition, the Cisco Wireless IP Phone 7920 provides enhanced WLAN security and Quality of Service (QoS) for operation in 802.11b networks, and it provides support for XML-based data access and services.

For more information about the various types of endpoints, see [IP Telephony Endpoints, page 21-1](#).

Conferencing, Messaging, and Collaboration Capabilities

Cisco Unified Communications supports the following additional features and applications to provide conferencing, voice messaging, and multi-media collaboration functionality:

- Conferencing

Cisco Unified CallManager can interface with a number of other Cisco software and hardware devices to provide complete conferencing capability, including annunciator functions and music on hold. For more information about designing and provisioning conference capabilities in Cisco Unified CallManager, see [Media Resources, page 6-1](#).

- Voice messaging

Cisco Unified CallManager is capable of integrating with third-party voicemail systems and with Cisco Unity and Unity Connection to provide a full range of voicemail and voice messaging functionality. For more information about integrating third-party voicemail systems with Cisco Unified CallManager, see [Third-Party Voicemail Design, page 12-1](#). For more information about integrating Cisco Unity and Unity Connection with Cisco Unified CallManager, see [Cisco Unity, page 13-1](#).

- Video telephony

Video telephony capability is now fully integrated into Cisco Unified CallManager, and there are also new video endpoints available from Cisco and its strategic partners. Video calls and conferences are now just as easy to make as a voice call on an IP phone. For more information about video capability with Cisco Unified CallManager, see [IP Video Telephony, page 17-1](#).

- Multimedia collaboration

Cisco Unified MeetingPlace is a complete rich-media conferencing application that integrates voice, video, and web conferencing capabilities to make remote meetings as natural and effective as face-to-face meetings. For more information about integrating MeetingPlace with Cisco Unified CallManager, see [Cisco Unified MeetingPlace Integration, page 15-1](#).

Security

Security for a Cisco Unified Communications deployment involves the following main considerations, among others:

- Physical security for restricting physical access to important application servers and network components
- Network access security to prevent hostile logins or attacks
- Security measures for Cisco Unified CallManager, endpoint devices, and various directories and databases

- Mechanisms for defining calling privileges for various classes of users
- Careful network design and management to enhance security

For more information about security for IP Telephony networks, see [Voice Security, page 20-1](#).



CHAPTER 2

IP Telephony Deployment Models

Last revised on: February 13, 2008

This chapter describes the IP Telephony deployment models for Cisco Unified CallManager 4.2. For design guidance with earlier releases of Cisco Unified CallManager, refer to the IP Telephony Solution Reference Network Design (SRND) documentation available at

<http://www.cisco.com/go/designzone>

Each Cisco Unified Communications solution is based on one of the following main deployment models described in this chapter:

- [Single Site, page 2-2](#)

The single-site model for IP telephony consists of a call processing agent located at a single site and a LAN or metropolitan area network (MAN) to carry voice traffic throughout the site. Calls beyond the LAN or MAN use the public switched telephone network (PSTN). If an IP WAN is incorporated into the single-site model, it is for data traffic only; no telephony services are provided over the WAN.

Use this model for a single campus or site with less than 30,000 lines.

- [Multisite WAN with Centralized Call Processing, page 2-4](#)

The multisite WAN model with centralized call processing consists of a single call processing agent that provides services for many sites and uses the IP WAN to transport voice traffic between the sites. The IP WAN also carries call control signaling between the central site and the remote sites.

Use this model for a main site with many smaller remote sites that are connected via a QoS-enabled WAN but that do not require full features and functionality during a WAN outage.

- [Multisite WAN with Distributed Call Processing, page 2-14](#)

The multisite WAN model with distributed call processing consists of multiple independent sites, each with its own call processing agent connected to an IP WAN that carries voice traffic between the distributed sites. The IP WAN in this model does not carry call control signaling between the sites because each site has its own call processing agent.

Use this model for a large central site with more than 30,000 lines or for a deployment with more than six large sites (more than 30,000 lines total) interconnected via a QoS-enabled WAN.

- [Clustering Over the IP WAN, page 2-17](#)

This model deploys a single Cisco Unified CallManager cluster across multiple sites that are connected by an IP WAN with QoS features enabled.

Use this model for a deployment with a maximum of six large sites (maximum of 30,000 lines total) interconnected via a QoS-enabled WAN.

**Note**

Other sections of this document assume that you understand the concepts involved with these deployment models, so please become thoroughly familiar with them before proceeding.

In addition, the section on [Design Considerations for Section 508 Conformance, page 2-27](#), presents guidelines for designing your IP telephony network to provide accessibility to users with disabilities, in conformance with U.S. Section 508.

**Note**

For the most recent information on recommended hardware platforms and software releases, frequently refer to the documentation at

http://www.cisco.com/en/US/products/sw/voicesw/ps556/products_device_support_tables_list.html


Single Site

The single-site model for IP telephony consists of a call processing agent located at a single site, or campus, with *no* telephony services provided over an IP WAN. An enterprise would typically deploy the single-site model over a LAN or metropolitan area network (MAN), which carries the voice traffic within the site. In this model, calls beyond the LAN or MAN use the public switched telephone network (PSTN).

The single-site model has the following design characteristics:

- Single Cisco Unified CallManager or Cisco Unified CallManager cluster.
- Maximum of 30,000 Skinny Client Control Protocol (SCCP) IP phones or SCCP video endpoints per cluster.
- Maximum of 500 H.323 devices (gateways, MCUs, trunks, and clients) per Cisco Unified CallManager cluster.
- PSTN for all external calls.
- Digital signal processor (DSP) resources for conferencing, transcoding, and media termination point (MTP).
- Voicemail or unified messaging components.
- Capability to integrate with legacy private branch exchange (PBX) and voicemail systems.
- H.323 clients, MCUs, and H.323/H.320 gateways that require a gatekeeper to place calls must register with a Cisco IOS Gatekeeper (Cisco IOS Release 12.3(8)T or greater). Cisco Unified CallManager then uses an H.323 trunk to integrate with the gatekeeper and provide call routing and bandwidth management services for the H.323 devices registered to it. Multiple Cisco IOS Gatekeepers may be used to provide redundancy.
- MCU resources are required for multipoint video conferencing. Depending on conferencing requirements, these resources may be either SCCP or H.323, or both.
- H.323/H.320 video gateways are needed to communicate with H.320 videoconferencing devices on the public ISDN network.
- High-bandwidth audio (for example, G.711, G.722, or Cisco Wideband Audio) between devices within the site.
- High-bandwidth video (for example, 384 kbps or greater) between devices within the site. The Cisco Unified Video Advantage Wideband Codec, operating at 7 Mbps, is also supported.

[Figure 2-1](#) illustrates the model for an IP telephony network within a single campus or site.

Figure 2-1 Single-Site Deployment Model

Benefits of the Single-Site Model

A single infrastructure for a converged network solution provides significant cost benefits and enables IP telephony to take advantage of the many IP-based applications in the enterprise. Single-site deployment also allows each site to be completely self-contained. There is no dependency for service in the event of an IP WAN failure or insufficient bandwidth, and there is no loss of call processing service or functionality.

In summary, the main benefits of the single-site model are:

- Ease of deployment
- A common infrastructure for a converged solution
- Simplified dial plan
- No transcoding resources required, due to the use of only G.711 codecs

Best Practices for the Single-Site Model

Follow these guidelines and best practices when implementing the single-site model:


- Provide a highly available, fault-tolerant infrastructure based on a common infrastructure philosophy. A sound infrastructure is essential for easier migration to IP telephony, integration with applications such as video streaming and video conferencing, and expansion of your IP telephony deployment across the WAN or to multiple Cisco Unified CallManager clusters.
- Know the calling patterns for your enterprise. Use the single-site model if most of the calls from your enterprise are within the same site or to PSTN users outside your enterprise.
- Use G.711 codecs for all endpoints. This practice eliminates the consumption of digital signal processor (DSP) resources for transcoding, and those resources can be allocated to other functions such as conferencing and Media Termination Points (MTPs).
- Use Media Gateway Control Protocol (MGCP) gateways for the PSTN if you do *not* require H.323 functionality. This practice simplifies the dial plan configuration. H.323 might be required to support specific functionality such as support for Signaling System 7 (SS7) or Non-Facility Associated Signaling (NFAS).
- Implement the recommended network infrastructure for high availability, connectivity options for phones (in-line power), Quality of Service (QoS) mechanisms, and security. (See [Network Infrastructure, page 3-1](#).)
- Follow the provisioning recommendations listed in the chapter on [Call Processing, page 8-1](#).

Multisite WAN with Centralized Call Processing

The multisite WAN model with centralized call processing consists of a single call processing agent that provides services for many sites and uses the IP WAN to transport IP telephony traffic between the sites. The IP WAN also carries call control signaling between the central site and the remote sites. [Figure 2-2](#) illustrates a typical centralized call processing deployment, with a Cisco Unified CallManager cluster as the call processing agent at the central site and an IP WAN with QoS enabled to connect all the sites. The remote sites rely on the centralized Cisco Unified CallManager cluster to handle their call processing. Applications such as voicemail and Interactive Voice Response (IVR) systems are typically centralized as well to reduce the overall costs of administration and maintenance.

**Note**

In each solution for the centralized call processing model presented in this document, the various sites connect to an IP WAN with QoS enabled.

Figure 2-2 Multisite Deployment Model with Centralized Call Processing

The multisite model with centralized call processing has the following design characteristics:

- Single Cisco Unified CallManager or Cisco Unified CallManager cluster.
- Maximum of 30,000 Skinny Client Control Protocol (SCCP) IP phones or SCCP video endpoints per cluster.
- Maximum of 500 H.323 devices (gateways, MCUs, trunks, and clients) per Cisco Unified CallManager cluster.
- PSTN for all external calls.
- Digital signal processor (DSP) resources for conferencing, transcoding, and media termination point (MTP).
- Voicemail or unified messaging components.
- Capability to integrate with legacy private branch exchange (PBX) and voicemail systems.
- H.323 clients, MCUs, and H.323/H.320 gateways that require a gatekeeper to place calls must register with a Cisco IOS Gatekeeper (Cisco IOS Release 12.3(8)T or greater). Cisco Unified CallManager then uses an H.323 trunk to integrate with the gatekeeper and provide call routing and bandwidth management services for the H.323 devices registered to it. Multiple Cisco IOS Gatekeepers may be used to provide redundancy.
- MCU resources are required for multipoint video conferencing. Depending on conferencing requirements, these resources may be either SCCP or H.323, or both, and may all be located at the central site or may be distributed to the remote sites if local conferencing resources are required.

- H.323/H.320 video gateways are needed to communicate with H.320 videoconferencing devices on the public ISDN network. These gateways may all be located at the central site or may be distributed to the remote sites if local ISDN access is required.
- High-bandwidth audio (for example, G.711, G.722, or Cisco Wideband Audio) between devices in the same site, and low-bandwidth audio (for example, G.729 or G.728) between devices in different sites.
- High-bandwidth video (for example, 384 kbps or greater) between devices in the same site, and low-bandwidth video (for example, 128 kbps) between devices at different sites. The Cisco Unified Video Advantage Wideband Codec, operating at 7 Mbps, is recommended only for calls between devices at the same site.
- Minimum of 768 kbps or greater WAN link speeds. Video is *not* recommended on WAN connections that operate at speeds lower than 768 kbps.
- Cisco Unified CallManager locations provide call admission control, and automated alternate routing (AAR) is also supported for video calls.
- Maximum of 500 locations per Cisco Unified CallManager cluster.
- Survivable Remote Site Telephony (SRST) is *not* supported for video. If the WAN connection fails, SCCP video endpoints located at the remote sites become audio-only devices, and H.323 video devices fail. An active SCCP video call between two devices at remote sites will continue with both video and audio channels, but it will not be able to activate additional features such as call transfer. Active H.323 video calls fail. In SRST mode, new calls are always audio-only. Any call that goes over the WAN will be dropped if the WAN connection fails.

Connectivity options for the IP WAN include:

- Leased lines
- Frame Relay
- Asynchronous Transfer Mode (ATM)
- ATM and Frame Relay Service Inter-Working (SIW)
- Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN)
- Voice and Video Enabled IP Security Protocol (IPSec) VPN (V3PN)

Routers that reside at the WAN edges require quality of service (QoS) mechanisms, such as priority queuing and traffic shaping, to protect the voice traffic from the data traffic across the WAN, where bandwidth is typically scarce. In addition, a call admission control scheme is needed to avoid oversubscribing the WAN links with voice traffic and deteriorating the quality of established calls. For centralized call processing deployments, the *locations* construct within Cisco Unified CallManager provides call admission control. (Refer to the section on [Cisco Unified CallManager Static Locations, page 9-12](#), for more information on locations.)

A variety of Cisco gateways can provide the remote sites with PSTN access. When the IP WAN is down, or if all the available bandwidth on the IP WAN has been consumed, users at the remote sites can dial the PSTN access code and place their calls through the PSTN. The Survivable Remote Site Telephony (SRST) feature on Cisco IOS gateways provides call processing at the branch offices in the event of a WAN failure.

Best Practices for the Centralized Call Processing Model

Follow these guidelines and best practices when implementing the multisite WAN model with centralized call processing:

- Minimize delay between Cisco Unified CallManager and remote locations to reduce voice cut-through delays (also known as clipping).
- Use the locations mechanism in Cisco Unified CallManager to provide call admission control into and out of remote branches. See the chapter on [Call Admission Control, page 9-1](#), for details on how to apply this mechanism to the various WAN topologies.
- The locations mechanism works across multiple servers in Cisco Unified CallManager Release 3.1 and later. This configuration can support a maximum of 30,000 IP phones when Cisco Unified CallManager runs on the largest supported server.
- The number of IP phones and line appearances supported in Survivable Remote Site Telephony (SRST) mode at each remote site depends on the branch router platform, the amount of memory installed, and the Cisco IOS release. (For the latest SRST platform and code specifications, refer to the SRST documentation at Cisco.com.) Generally speaking, however, the choice of whether to adopt a centralized call processing or distributed call processing approach for a given site depends on a number of factors such as:
 - IP WAN bandwidth or delay limitations
 - Criticality of the voice network
 - Feature set needs
 - Scalability
 - Ease of management
 - Cost

If a distributed call processing model is deemed more suitable for the customer's business needs, the choices include installing a local Cisco Unified CallManager server or running Cisco Unified CallManager Express on the branch router.

Remote Site Survivability


When deploying IP Telephony across a WAN with the centralized call processing model, you should take additional steps to ensure that data and voice services at the remote sites are highly available. [Table 2-1](#) summarizes the different strategies for providing high availability at the remote sites. The choice of one of these strategies may depend on several factors, such as specific business or application requirements, the priorities associated with highly available data and voice services, and cost considerations.

Table 2-1 *Strategies for High Availability at the Remote Sites*

Strategy	High Availability for Data Services?	High Availability for Voice Services?
Redundant IP WAN links in branch router	Yes	Yes
Redundant branch router platforms + Redundant IP WAN links	Yes	Yes
Survivable Remote Site Telephony (SRST) only	No	Yes
Data-only ISDN backup + SRST	Yes	Yes
Data and voice ISDN backup	Yes	Yes (see rules below)

The first two solutions listed in [Table 2-1](#) provide high availability at the network infrastructure layer by adding redundancy to the IP WAN access points, thus maintaining IP connectivity between the remote IP phones and the centralized Cisco Unified CallManager at all times. These solutions apply to both data and voice services, and are entirely transparent to the call processing layer. The options range from adding a redundant IP WAN link at the branch router to adding a second branch router platform with a redundant IP WAN link.

The third solution listed in [Table 2-1](#), Survivable Remote Site Telephony (SRST), provides high availability for voice services only, by providing a subset of the call processing capabilities within the remote office router and enhancing the IP phones with the ability to “re-home” to the call processing functions in the local router if a WAN failure is detected. [Figure 2-3](#) illustrates a typical call scenario with SRST.

Figure 2-3 Survivable Remote Site Telephony (SRST) Application Scenario

74353

Under normal operations shown in the left part of [Figure 2-3](#), the branch office connects to the central site via an IP WAN, which carries data traffic, voice traffic, and call signaling. The IP phones at the branch office exchange call signaling information with the Cisco Unified CallManager cluster at the central site and place their calls across the IP WAN. The branch router or gateway forwards both types of traffic (call signaling and voice) transparently and has no knowledge of the IP phones.

If the WAN link to the branch office fails, or if some other event causes loss of connectivity to the Cisco Unified CallManager cluster, the branch IP phones re-register with the branch router. The branch router queries the IP phones for their configuration and uses this information to build its own

configuration automatically. The branch IP phones can then make and receive calls either internally or through the PSTN. The phone displays the message “Unified CM fallback mode,” and some advanced Cisco Unified CallManager features are unavailable and are grayed out on the phone display.

When WAN connectivity to the central site is reestablished, the branch IP phones automatically re-register with the Cisco Unified CallManager cluster and resume normal operation. The branch router deletes its information about the IP phones and reverts to its standard routing or gateway configuration.

The last two solutions in [Table 2-1](#) use an ISDN backup link to provide survivability during WAN failures. The two deployment options for ISDN backup are:

- Data-only ISDN backup

With this option, ISDN is used for data survivability only, while SRST is used for voice survivability. Note that you should configure an access control list on the branch router to prevent Skinny Client Control Protocol (SCCP) traffic from entering the ISDN interface, so that signaling from the IP phones does not reach the Cisco Unified CallManager at the central site.

- Data and voice ISDN backup

With this option, ISDN is used for both data and voice survivability. In this case, SRST is not used because the IP phones maintain IP connectivity to the Cisco Unified CallManager cluster at all times. However, Cisco recommends that you use ISDN to transport data and voice traffic only if all of the following conditions are true:

- The bandwidth allocated to voice traffic on the ISDN link is the same as the bandwidth allocated to voice traffic on the IP WAN link.
- The ISDN link bandwidth is fixed.
- All the required QoS features have been deployed on the router's ISDN interfaces. Refer to the chapter on [Network Infrastructure, page 3-1](#), for more details on QoS.

Voice Over the PSTN as a Variant of Centralized Call Processing

The centralized call processing deployment model can be adapted so that inter-site voice media is sent over the PSTN instead of the WAN. With this configuration, the signaling (call control) of all telephony endpoints is still controlled by the central Cisco Unified CallManager cluster, therefore this Voice over the PSTN (VoPSTN) model variation still requires a QoS-enabled WAN with appropriate bandwidth configured for the signaling traffic.

You can implement VoPSTN in one of the following ways:

- Using the automated alternate routing (AAR) feature. (For more information on AAR, see the section on [Automated Alternate Routing, page 10-22](#).)
- Using a combination of dial plan constructs in both Cisco Unified CallManager and the PSTN gateways.

VoPSTN can be an attractive option in deployments where IP WAN bandwidth is either scarce or expensive with respect to PSTN charges, or where IP WAN bandwidth upgrades are planned for a later date but the IP telephony system is already being deployed.



Note

Because of its very nature, the VoPSTN deployment model variation offers basic voice functionality that is a reduced subset of the Cisco Unified CallManager feature set.

In particular, regardless of the implementation choice, the system designer should address the following issues, among others:

- Centralized voicemail requires:
 - A telephony network provider that supports redirected dialed number identification service (RDNIS) end-to-end for all locations that are part of the deployment. RDNIS is required so that calls redirected to voicemail carry the redirecting DN, to ensure proper voicemail box selection.
 - If the voicemail system is accessed through an MGCP gateway, the voicemail pilot number must be a fully qualified E.164 number.
- The Extension Mobility feature is limited to IP phones contained within a single branch site.
- All on-net (intra-cluster) calls will be delivered to the destination phone with the same call treatment as an off-net (PSTN) call. This includes the quantity of digits delivered in the call directories such as Missed Calls and Received Calls.
- Each inter-branch call generates two independent call detail records (CDRs): one for the call leg from the calling phone to the PSTN, and the other for the call leg from the PSTN to the called phone.
- There is no way to distinguish the ring type for on-net and off-net calls.
- All destination phones require a fully qualified Direct Inward Dial (DID) PSTN number that can be called directly. Non-DID DNs cannot be reached directly from a different branch site.
- With VoPSTN, music on hold (MoH) is limited to cases where the holding party is co-located with the MoH resource. If MoH servers are deployed at the central site, then only calls placed on hold by devices at the central site will receive the hold music.
- Transfers to a destination outside the branch site will result in the hairpinning of the call through the branch's gateway. Traffic engineering of the branch's gateway resources must be adjusted accordingly.
- Call forwarding of any call coming into the branch's gateway to a destination outside the branch site will result in hairpinning of the call through the gateway, thus using two trunk ports. This behavior applies to:
 - Calls forwarded to a voicemail system located outside the branch
 - Calls forwarded to an on-net abbreviated dialing destination located in a different branch

The gateway port utilization resulting from these call forwarding flows should be taken into account when sizing the trunks connecting the branch to the PSTN.

- Conferencing resources must be co-located with the phone initiating the conference.
- VoPSTN does not support applications that require streaming of IP audio from the central site (that is, not traversing a gateway). These applications include, but are not limited to:
 - Centralized music on hold (MoH) servers
 - Interactive Voice Response (IVR)
 - CTI-based applications
- Use of the Attendant Console outside of the central site can require a considerable amount of bandwidth if the remote sites must access large user account directories without caching them.
- Because all inter-branch media (including transfers) are sent through the PSTN, the gateway trunk group must be sized to accommodate all inter-branch traffic, transfers, and centralized voicemail access.
- Cisco recommends that you do not deploy shared lines across branches, such that the devices sharing the line are in different branches.

In addition to these general considerations, the following sections present recommendations and issues specific to each of the following implementation methods:

- [VoPSTN Using AAR, page 2-12](#)
- [VoPSTN Using Dial Plan, page 2-13](#)

VoPSTN Using AAR

This method consists of configuring the Cisco Unified CallManager dial plan as in a traditional centralized call processing deployment, with the automated alternate routing (AAR) feature also properly configured. AAR provides transparent re-routing over the PSTN of inter-site calls when the locations mechanism for call admission control determines that there is not enough available WAN bandwidth to accept an additional call.

To use the PSTN as the primary (and only) voice path, you can configure the call admission control bandwidth of each location (branch site) to be 1 kbps, thus preventing *all* calls from traversing the WAN. With this configuration, all inter-site calls trigger the AAR functionality, which automatically re-routes the calls over the PSTN.

The AAR implementation method for VoPSTN offers the following benefits:

- An easy migration path to a complete IP telephony deployment. When bandwidth becomes available to support voice media over the WAN, the dial plan can be maintained intact, and the only change needed is to update the location bandwidth value for each site.
- Support for some supplementary features, such as callback on busy.

In addition to the general considerations listed for VoPSTN, the following design guidelines apply to the AAR implementation method:

- AAR functionality must be configured properly.
- As a general rule, supported call initiation devices include IP phones, gateways, and line-side gateway-driven analog phones.
- Inter-branch calls can use AAR only if the destination devices are IP phones or Cisco Unity ports.
- Inter-branch calls to other endpoints must use a fully qualified E.164 number.
- All on-net, inter-branch calls will display the message, "Network congestion, rerouting."
- If destination phones become unregistered (for example, due to WAN connectivity interruption), AAR functionality will not be invoked and abbreviated dialing will not be possible. If the destination phone has registered with an SRST router, then it can be reached by directly dialing its PSTN DID number.
- If originating phones become unregistered (for example, due to WAN connectivity interruption), they will go into SRST mode. To preserve abbreviated dialing functionality under these conditions, configure the SRST router with an appropriate set of translation rules.
- Shared lines within the same branch should be configured in a partition included only in that branch's calling search spaces. Inter-site access to the shared line requires one of the following:
 - The originating site dials the DID number of the shared line.
 - If inter-site abbreviated dialing to the shared line is desired, use a translation pattern that expands the user-dialed abbreviated string to the DID number of the shared line.



Note

In this case, direct dialing of the shared line's DN from another branch would trigger multiple AAR-based PSTN calls.

VoPSTN Using Dial Plan

This method relies on a specific dial plan configuration within Cisco Unified CallManager and the PSTN gateways to route all inter-site calls over the PSTN. The dial plan must place IP phone DN's at each site into a different partition, and their calling search space must provide access only to the site's internal partition and a set of route patterns that point to the local PSTN gateway.

Abbreviated inter-site dialing can still be provided via a set of translations at each branch site, one for each of the other branch sites. These translations are best accomplished with H.323 gateways and translation rules within Cisco IOS.

The dial plan method for implementing VoPSTN offers the following benefits:

- Easier configuration because AAR is not needed.
- Abbreviated dialing automatically works even under WAN failure conditions on either the originating or destination side, because the Cisco IOS translation rules within the H.323 gateway are effective in SRST mode.


In addition to the general considerations listed for VoPSTN, the following design guidelines apply to the dial plan implementation method:

- There is no support for supplementary features such as callback on busy.
- Some CTI-based applications do not support overlapping extensions (that is, two or more phones configured with the same DN, although in different partitions).
- There is no easy migration to a complete IP telephony deployment because the dial plan needs to be redesigned.

Multisite WAN with Distributed Call Processing

The multisite WAN model with distributed call processing consists of multiple independent sites, each with its own call processing agent connected to an IP WAN that carries voice traffic between the distributed sites. Figure 2-4 illustrates a typical distributed call processing deployment.

Figure 2-4 Multisite Deployment Model with Distributed Call Processing



141851

Each site in the distributed call processing model can be one of the following:

- A single site with its own call processing agent, which can be either:
 - Cisco Unified CallManager
 - Cisco Unified CallManager Express
 - Other IP PBX
- A centralized call processing site and all of its associated remote sites
- A legacy PBX with Voice over IP (VoIP) gateway

The multisite model with distributed call processing has the following design characteristics:

- Maximum of 30,000 Skinny Client Control Protocol (SCCP) IP phones or SCCP video endpoints per cluster.
- Maximum of 500 H.323 devices (gateways, MCUs, trunks, and clients) per Cisco Unified CallManager cluster.
- PSTN for all external calls.
- Digital signal processor (DSP) resources for conferencing, transcoding, and media termination point (MTP).
- Voicemail or unified messaging components.
- Capability to integrate with legacy private branch exchange (PBX) and voicemail systems.
- H.323 clients, MCUs, and H.323/H.320 gateways that require a gatekeeper to place calls must register with a Cisco IOS Gatekeeper (Cisco IOS Release 12.3(8)T or greater). Cisco Unified CallManager then uses an H.323 trunk to integrate with the gatekeeper and provide call routing and bandwidth management services for the H.323 devices registered to it. Multiple Cisco IOS Gatekeepers may be used to provide redundancy. Cisco IOS Gatekeepers may also be used to provide call routing and bandwidth management between the distributed Cisco Unified CallManager clusters. In most situations, Cisco recommends that each Cisco Unified CallManager cluster have its own set of endpoint gatekeepers and that a separate set of gatekeepers be used to manage the intercluster calls. It is possible in some circumstances to use the same set of gatekeepers for both functions, depending on the size of the network, complexity of the dial plan, and so forth. (For details, see [Gatekeepers, page 17-19](#).)
- MCU resources are required in each cluster for multipoint video conferencing. Depending on conferencing requirements, these resources may be either SCCP or H.323, or both, and may all be located at the regional sites or may be distributed to the remote sites of each cluster if local conferencing resources are required.
- H.323/H.320 video gateways are needed to communicate with H.320 videoconferencing devices on the public ISDN network. These gateways may all be located at the regional sites or may be distributed to the remote sites of each cluster if local ISDN access is required.
- High-bandwidth audio (for example, G.711, G.722, or Cisco Wideband Audio) between devices in the same site, but low-bandwidth audio (for example, G.729 or G.728) between devices in different sites.
- High-bandwidth video (for example, 384 kbps or greater) between devices in the same site, but low-bandwidth video (for example, 128 kbps) between devices at different sites. The Cisco Unified Video Advantage Wideband Codec, operating at 7 Mbps, is recommended only for calls between devices at the same site. Note that the Cisco VT Camera Wideband Video Codec is not supported over intercluster trunks.
- Minimum of 768 kbps or greater WAN link speeds. Video is *not* recommended on WAN connections that operate at speeds lower than 768 kbps.
- Call admission control is provided by Cisco Unified CallManager locations for calls between sites controlled by the same Cisco Unified CallManager cluster, and by the Cisco IOS Gatekeeper for calls between Cisco Unified CallManager clusters (that is, intercluster trunks). Automated alternate routing (AAR) is also supported for both intra-cluster and inter-cluster video calls.

An IP WAN interconnects all the distributed call processing sites. Typically, the PSTN serves as a backup connection between the sites in case the IP WAN connection fails or does not have any more available bandwidth. A site connected only through the PSTN is a standalone site and is not covered by the distributed call processing model. (See [Single Site, page 2-2](#).)

Connectivity options for the IP WAN include:

- Leased lines
- Frame Relay
- Asynchronous Transfer Mode (ATM)
- ATM and Frame Relay Service Inter-Working (SIW)
- Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN)
- Voice and Video Enabled IP Security Protocol (IPSec) VPN (V3PN)

Benefits of the Distributed Call Processing Model

The multisite WAN model with distributed call processing provides the following benefits:

- PSTN call cost savings when using the IP WAN for calls between sites.
- Use of the IP WAN to bypass toll charges by routing calls through remote site gateways, closer to the PSTN number dialed. This practice is known as tail-end hop-off (TEHO).
- Maximum utilization of available bandwidth by allowing voice traffic to share the IP WAN with other types of traffic.
- No loss of functionality during IP WAN failure because there is a call processing agent at each site.
- Scalability to hundreds of sites.

Best Practices for the Distributed Call Processing Model

A multisite WAN deployment with distributed call processing has many of the same requirements as a single site or a multisite WAN deployment with centralized call processing. Follow the best practices from these other models in addition to the ones listed here for the distributed call processing model. (See [Single Site, page 2-2](#), and [Multisite WAN with Centralized Call Processing, page 2-4](#).)

Gatekeeper or Session Initiation Protocol (SIP) proxy servers are among the key elements in the multisite WAN model with distributed call processing. They each provide dial plan resolution, with the gatekeeper also providing call admission control. A gatekeeper is an H.323 device that provides call admission control and E.164 dial plan resolution.

The following best practices apply to the use of a gatekeeper:

- Use a Cisco IOS gatekeeper to provide call admission control into and out of each site.
- To provide high availability of the gatekeeper, use Hot Standby Router Protocol (HSRP) gatekeeper pairs, gatekeeper clustering, and alternate gatekeeper support. In addition, use multiple gatekeepers to provide redundancy within the network. (See [Gatekeeper Design Considerations, page 8-18](#).)
- Size the platforms appropriately to ensure that performance and capacity requirements can be met.
- Use only one type of codec on the WAN because the H.323 specification does not allow for Layer 2, IP, User Data Protocol (UDP), or Real-time Transport Protocol (RTP) header overhead in the bandwidth request. (Header overhead is allowed only in the payload or encoded voice part of the packet.) Using one type of codec on the WAN simplifies capacity planning by eliminating the need to over-provision the IP WAN to allow for the worst-case scenario.
- Gatekeeper networks can scale to hundreds of sites, and the design is limited only by the WAN topology.

For more information on the various functions performed by gatekeepers, refer to the following sections:

- For gatekeeper call admission control, see [Call Admission Control, page 9-1](#).
- For gatekeeper scalability and redundancy, see [Call Processing, page 8-1](#).
- For gatekeeper dial plan resolution, see [Dial Plan, page 10-1](#).

SIP devices provide resolution of E.164 numbers as well as SIP uniform resource identifiers (URIs) to enable endpoints to place calls to each other. Cisco Unified CallManager supports the use of E.164 numbers only.

The following best practices apply to the use of SIP proxies:

- Provide adequate redundancy for the SIP proxies.
- Ensure that the SIP proxies have the capacity for the call rate and number of calls required in the network.
- Planning for call admission control is outside the scope of this document.

Call Processing Agents for the Distributed Call Processing Model

Your choice of call processing agent will vary, based on many factors. The main factors, for the purpose of design, are the size of the site and the functionality required.

For a distributed call processing deployment, each site has its own call processing agent. The design of each site varies with the call processing agent, the functionality required, and the fault tolerance required. For example, in a site with 500 phones, a Cisco Unified CallManager cluster containing two servers can provide one-to-one redundancy, with the backup server being used as a publisher and Trivial File Transfer Protocol (TFTP) server.

The requirement for IP-based applications also greatly affects the choice of call processing agent because only Cisco Unified CallManager provides the required support for many Cisco IP applications.

[Table 2-2](#) lists recommended call processing agents.

Table 2-2 Recommended Call Processing Agents

Call Processing Agent	Recommended Size	Comments
Cisco Unified CallManager Express (CME)	Up to 240 phones	• For small remote sites • Capacity depends on Cisco IOS platform
Cisco Unified CallManager	50 to 30,000 phones	• Small to large sites, depending on the size of the Cisco Unified CallManager cluster • Supports centralized or distributed call processing
Legacy PBX with VoIP gateway	Depends on PBX	• Number of IP WAN calls and functionality depend on the PBX-to-VoIP gateway protocol and the gateway platform

Clustering Over the IP WAN

You may deploy a single Cisco Unified CallManager cluster across multiple sites that are connected by an IP WAN with QoS features enabled. This section provides a brief overview of clustering over the WAN. For further information, refer to the chapter on [Call Processing, page 8-1](#).

Clustering over the WAN can support two types of deployments:

- [Local Failover Deployment Model, page 2-22](#)

Local failover requires that you place the Cisco Unified CallManager subscriber and backup servers at the same site, with no WAN between them. This deployment model is ideal for two to four sites with Cisco Unified CallManager

- [Remote Failover Deployment Model, page 2-26](#)

Remote failover allows you to deploy the backup servers over the WAN. Using this deployment model, you may have up to eight sites with Cisco Unified CallManager subscribers being backed up by Cisco Unified CallManager subscribers at another site.

You can also use a combination of the two deployment models to satisfy specific site requirements. For example, two main sites may each have primary and backup subscribers, with another two sites containing only a primary server each and utilizing either shared backups or dedicated backups at the two main sites.

The key advantages of clustering over the WAN are:

- Single point of administration for users for all sites within the cluster
- Feature transparency
- Shared line appearances
- Extension mobility within the cluster
- Unified dial plan

These features make this solution ideal as a disaster recovery plan for business continuance sites or as a single solution for up to eight small or medium sites.

WAN Considerations

For clustering over the WAN to be successful, you must carefully plan, design, and implement various characteristics of the WAN itself. The Intra-Cluster Communication Signaling (ICCS) between Cisco Unified CallManager servers consists of many traffic types. The ICCS traffic types are classified as either priority or best-effort. Priority ICCS traffic is marked with IP Precedence 3 (DSCP 24 or PHB CS3). Best-effort ICCS traffic is marked with IP Precedence 0 (DSCP 0 or PHB BE). The various types of ICCS traffic are described in [Intra-Cluster Communications, page 2-19](#), which also provides further guidelines for provisioning. The following design guidelines apply to the indicated WAN characteristics:

- Delay

The maximum one-way delay between any Cisco Unified CallManager servers for all priority ICCS traffic should not exceed 20 ms, or 40 ms round-trip time (RTT). Delay for other ICCS traffic should be kept reasonable to provide timely database and directory access. Measuring the delay is covered in [Delay Testing, page 2-20](#). Propagation delay between two sites introduces 6 microseconds per kilometer without any other network delays being considered. This equates to a theoretical maximum distance of approximately 3000 km for 20 ms delay or approximately 1860 miles. These distances are provided only as relative guidelines and in reality will be shorter due to other delay incurred within the network.

- Jitter

Jitter is the varying delay that packets incur through the network due to processing, queue, buffer, congestion, or path variation delay. Jitter for the IP Precedence 3 ICCS traffic must be minimized using Quality of Service (QoS) features.

- Packet loss and errors

The network should be engineered to provide sufficient prioritized bandwidth for all ICCS traffic, especially the priority ICCS traffic. Standard QoS mechanisms must be implemented to avoid congestion and packet loss. If packets are lost due to line errors or other “real world” conditions, the ICCS packet will be retransmitted because it uses the TCP protocol for reliable transmission. The retransmission might result in a call being delayed during setup, disconnect (teardown), or other supplementary services during the call. Some packet loss conditions could result in a lost call, but this scenario should be no more likely than errors occurring on a T1 or E1, which affect calls via a trunk to the PSTN/ISDN.

- Bandwidth

Provision the correct amount of bandwidth between each server for the expected call volume, type of devices, and number of devices. This bandwidth is in addition to any other bandwidth for other applications sharing the network, including voice and video traffic between the sites. The bandwidth provisioned must have QoS enabled to provide the prioritization and scheduling for the different classes of traffic. The general rule of thumb for bandwidth is to over-provision and under-subscribe.

- Quality of Service

The network infrastructure relies on QoS engineering to provide consistent and predictable end-to-end levels of service for traffic. Neither QoS nor bandwidth alone is the solution; rather, QoS-enabled bandwidth must be engineered into the network infrastructure.

Intra-Cluster Communications

In general, intra-cluster communications means all traffic between servers. There is also a real-time protocol called Intra-Cluster Communication Signaling (ICCS), which provides the communications with the Cisco CallManager Service process that is at the heart of the call processing in each server or node within the cluster.

The intra-cluster traffic between the servers consists of the following:

- Database traffic from the MS-SQL database that provides the main configuration information. The SQL database is replicated from the publisher server to all other servers in the cluster using best-effort. The SQL traffic may be re-prioritized in line with Cisco QoS recommendations to a higher priority data service (for example, IP Precedence 1 if required by the particular business needs). An example of this is extensive use of Extension Mobility, which relies on SQL database configuration.
- ICCS real-time traffic, which consists of signaling, call admission control, and other information regarding calls as they are initiated and completed. ICCS uses a Transmission Control Protocol (TCP) connection between all servers that have the Cisco CallManager Service enabled. The connections are a full mesh between these servers. Because only eight servers may have the Cisco CallManager Service enabled in a cluster, there may be up to seven connections on each server. This traffic is priority ICCS traffic and is marked dependant on release and service parameter configuration.
- CTI Manager real-time traffic is used for CTI devices involved in calls or for controlling or monitoring other third-party devices on the Cisco Unified CallManager servers. This traffic is marked as priority ICCS traffic and exists between the Cisco Unified CallManager server with the CTI Manager and the Cisco Unified CallManager server with the CTI device.

Failover Between Subscriber Servers

With Cisco Unified CallManager Release 4.2, the device configuration records are cached during the initialization or boot-up time. The effect is that Cisco Unified CallManager might take a longer time to initialize, but any failover or failback for all devices is not affected by the delay in accessing the publisher database.

Cisco Unified CallManager Publisher

The publisher replicates a read-only copy of the master database to all other servers in the cluster. If changes are made in the publisher's master database during a period when another server in the cluster is unreachable, the publisher will replicate the updated database when communications are re-established. During any period when the publisher is unreachable or offline, no changes can be made to the configuration database. All subscriber databases are read-only and may not be modified. Most normal operations of the cluster will not be affected during this period, including the following:

- Call processing
- Failover
- Installation registration of previously configured devices

There are some features and functions that require access to the master database on the publisher because they make modifications to records and therefore need write access. The publisher is the only server in a Cisco Unified CallManager cluster that has a read and write configuration database. The main features and functions that require access to the publisher for write access include:

- Configuration additions, changes, and deletions
- Extension Mobility and Device Mobility
- User speed dials
- Cisco Unified CallManager User page options requiring the database
- Cisco Unified CallManager software upgrades
- Call Forward All changes
- Message Waiting Indicator (MWI) state

Other services or applications might also be affected, and their ability to function without the publisher should be verified when deployed.

Call Detail Records (CDR)

Call detail records, when enabled, are collected by each subscriber and uploaded to the publisher periodically. During a period that the publisher is unreachable, the CDRs are stored on the subscriber's local hard disk. When connectivity is re-established to the publisher, all outstanding CDRs are uploaded to the publisher.

Delay Testing

The maximum round-trip time (RTT) between any two servers must not exceed 40 ms at any time. This time limit must include all delays in the transmission path between the two servers. Verifying the round trip delay using the **ping** utility on the Cisco Unified CallManager server will not provide an accurate result. The ping is sent as a best-effort tagged packet and is not transported using the same QoS-enabled path as the ICCS traffic. Therefore, Cisco recommends that you verify the delay by using the closest

network device to the Cisco Unified CallManager servers, ideally the access switch to which the server is attached. Cisco IOS provides an extended ping capable of setting the Layer 3 type of service (ToS) bits to make sure the ping packet is sent on the same QoS-enabled path that the ICCS traffic will traverse. The time recorded by the extended ping is the round-trip time (RTT), or the time it takes to traverse the communications path and return. The maximum RTT between any two Cisco Unified CallManager servers is 40 ms, and therefore the maximum one-way delay would be 20 ms. This delay is critical to the performance of the call processing capabilities of the Cisco Unified CallManager cluster and must be strictly enforced.

The following example shows a Cisco IOS extended ping with the ToS bit (IP Precedence) set to 3:

```
Access_SW#ping
Protocol [ip]:
Target IP address: 10.10.10.10
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: y
Source address or interface:
Type of service [0]: 3
Set DF bit in IP header? [no]:
Validate reply data? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Sweep range of sizes [n]:
Type escape sequence to abort.
Sending 5, 100-byte Intelligent Contact ManagementP Echos to 10.10.10.10, timeout is 2
seconds:
!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
```

Error Rate

The expected error rate should be zero. Any errors, dropped packets, or other impairments to the IP network can have an impact on the call processing performance of the cluster. This may be noticeable by delay in dial tone, slow key or display response on the IP phone, or delay from off-hook to connection of the voice path. Although Cisco Unified CallManager will tolerate random errors, they should be avoided to avoid impairing the performance of the cluster.

Troubleshooting

If the Cisco Unified CallManager subscribers in a cluster are experiencing impairment of the ICCS communication due to higher than expected delay, errors, or dropped packets, some of the following symptoms might occur:

- IP phones, gateways, or other devices on a remote Cisco Unified CallManager server within the cluster might temporarily be unreachable.
- Calls might be disconnected or might fail during call setup.
- Users might experience longer than expected delays before hearing dial tone.
- Busy hour call completions (BHCC) might be low.

- The ICCS (SDL session) might be reset or disconnected. The following example shows a Cisco Unified CallManager SDL trace of a remote server VO30-7835-8 going Out of Service and the devices that were reachable by that server being removed as “available” destinations:


```
RemoteCMOutOfService: Ip address: VO30-7835-8 remoteClusterId
VO30-7835-1-Cluster|<CLID::VO30-7835-1-Cluster><NID::VO30-7835-2>
|Delete entries from SsManagerTable, now this table has 75
entries|<CLID::VO30-7835-1-Cluster><NID::VO30-7835-2><CT::0,0,0,0.0><IP::><DEV::>
|Delete entries from FeatActTable, now this table has 70
entries|<CLID::VO30-7835-1-Cluster><NID::VO30-7835-2><CT::0,0,0,0.0><IP::><DEV::>
|Digit analysis: Remove remote pattern /5000020 , PID:
7:34:1|<CLID::VO30-7835-1-Cluster><NID::VO30-7835-2><CT::0,0,0,0.0><IP::><DEV::>
|Digit analysis: Remove remote pattern /5000066 , PID:
7:34:2|<CLID::VO30-7835-1-Cluster><NID::VO30-7835-2><CT::0,0,0,0.0><IP::><DEV::>
.
.
.
|Digit analysis: Remove remote pattern /5001002 , PID:
7:34:106|<CLID::VO30-7835-1-Cluster><NID::VO30-7835-2><CT::0,0,0,0.0><IP::><DEV::>
```

In summary, perform the following tasks to troubleshoot ICCS communication problems:

- Verify the delay between the servers.
- Check all links for errors or dropped packets.
- Verify that QoS is correctly configured.
- Verify that sufficient bandwidth is provisioned for the queues and across the WAN to support all the traffic.

Local Failover Deployment Model

The local failover deployment model provides the most resilience for clustering over the WAN. Each of the sites in this model contains at least one primary Cisco Unified CallManager subscriber and one backup subscriber. This configuration can support up to four sites. The maximum number of phones and other devices will be dependant on the quantity and type of servers deployed. The maximum total number of IP phones for all sites is 30,000. (See [Figure 2-5](#).)

Figure 2-5 Example of Local Failover Model

Observe the following guidelines when implementing the local failover model:

- Configure each site to contain at least one primary Cisco Unified CallManager subscriber and one backup subscriber.
- Configure Cisco Unified CallManager *groups* and *device pools* to allow devices within the site to register with only the servers at that site under all conditions.
- Cisco highly recommends that you replicate key services (TFTP, DNS, DHCP, LDAP, and IP Phone Services), all media resources (conference bridges and music on hold), and gateways at each site to provide the highest level of resiliency. You could also extend this practice to include a voicemail system at each site.
- Under a WAN failure condition, sites without access to the publisher database will lose some functionality, as follows:
 - System administration at the local site will not be able to add, modify, or delete any part of the configuration.
 - Extension mobility users will not be able to log in or log out of the IP phones.
 - Changes to Call Forward All will not be allowed.

- Under WAN failure conditions, calls made to phone numbers that are not currently communicating with the subscriber placing the call, will result in either a fast-busy tone or a call forward (possibly to voicemail, dependant on the location of the phone number being forwarded to). During this condition, users should manually dial those numbers via the PSTN.
- Every 10,000 busy hour call attempts (BHCA) between sites that are clustered over the WAN requires 900 kbps of bandwidth for Intra-Cluster Communication Signaling (ICCS). This is a minimum bandwidth requirement, and bandwidth is allocated in multiples of 900 kbps. The ICCS traffic types are classified as either priority or best-effort. Priority ICCS traffic is marked with IP Precedence 3 (DSCP 24 or PHB CS3). Best-effort ICCS traffic is marked with IP Precedence 0 (DSCP 0 or PHB BE).
- The minimum recommended bandwidth between sites that are clustered over the WAN is 1.544 Mbps. This amount allows for the minimum of 900 kbps for ICCS and 644 kbps for database and other inter-server traffic.
- A maximum round-trip time (RTT) of 40 ms is allowed between any two servers in the Cisco Unified CallManager cluster. This time equates to a 20 ms maximum one-way delay, or a transmission distance of approximately 1860 miles (3000 km) under ideal conditions.
- If remote branches using centralized call processing are connected to the main sites via clustering over the WAN, pay careful attention to the configuration of call admission control to avoid oversubscribing the links used for clustering over the WAN.
 - If the bandwidth is not limited on the links used for clustering over the WAN (that is, if the interfaces to the links are OC-3s or STM-1s and there is no requirement for call admission control), then the remote sites may be connected to any of the main sites because all the main sites should be configured as location <None>. This configuration still maintains hub-and-spoke topology for purposes of call admission control.
 - If you are using the Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) feature, all sites in Cisco Unified CallManager locations and the remote sites may register with any of the main sites.
 - If bandwidth is limited between the main sites, call admission control must be used between sites, and all remote sites must register with the main site that is configured as location <None>. This main site is considered the hub site, and all other remote sites and clustering-over-the-WAN sites are spokes sites.
- During a software upgrade, all servers in the cluster should be upgraded during the same maintenance period, using the standard upgrade procedures outlined in the software release notes.

Cisco Unified CallManager Provisioning for Local Failover

Provisioning of the Cisco Unified CallManager cluster for the local failover model should follow the design guidelines for capacities outlined in the chapter on [Call Processing, page 8-1](#). If voice or video calls are allowed across the WAN between the sites, then you must configure Cisco Unified CallManager *locations* in addition to the default location for the other sites, to provide call admission control between the sites. If the bandwidth is over-provisioned for the number of devices, it is still best practice to configure call admission control based on locations. If the locations-based call admission control rejects a call, automatic failover to the PSTN can be provided by the automated alternate routing (AAR) feature.

To improve redundancy and upgrade times, Cisco recommends that you enable the Cisco Trivial Transfer Protocol (TFTP) service on two Cisco Unified CallManager servers. Cisco also recommends that you do not enable more than one server with the Cisco TFTP service at a location remote to publisher, over a low-speed WAN link. More than two TFTP servers can be deployed in a cluster, however this configuration can result in an extended period for rebuilding all the TFTP files on all TFTP servers.

You can run the TFTP service on either a publisher or a subscriber server, depending on the site and the available capacity of the server. The TFTP server option must be correctly set in the DHCP servers at each site. If DHCP is not in use or if the TFTP server is manually configured, you should configure the correct address for the site. During an upgrade of the cluster or a restart of the Cisco TFTP service, the TFTP servers that are remote from the publisher will take longer to upgrade and rebuild all the configuration files. This time depends on the latency between the TFTP server and the publisher and on the quantity of devices configured in the database.

Other services, which may affect normal operation of Cisco Unified CallManager during WAN outages, should also be replicated at all sites to ensure uninterrupted service. These services include DHCP servers, DNS servers, corporate directories, and IP phone services. On each DHCP server, set the DNS server address correctly for each location.

IP phones may have shared line appearances between the sites. The ICCS bandwidth provisioned between the sites allows for the additional ICCS traffic that shared line appearances generate. During a WAN outage, call control for each line appearance is segmented, but call control returns to a single Cisco Unified CallManager server once the WAN is restored. During the WAN restoration period, there is additional traffic between the two sites. If this situation occurs during a period of high call volume, the shared lines might not operate as expected during that period. This situation should not last more than a few minutes, but if it is a concern, you can provision additional prioritized bandwidth to minimize the effects.

Gateways for Local Failover

Normally, gateways should be provided at all sites for access to the PSTN. The device pools should be configured to register the gateways with the Cisco Unified CallManager servers at the same site. Partitions and calling search spaces should also be configured to select the local gateways at the site as the first choice for PSTN access and the other site gateways as a second choice for overflow. Take special care to ensure emergency service access at each site.

You can centralize access to the PSTN gateways if access is not required during a WAN failure and if sufficient additional bandwidth is configured for the number of calls across the WAN. For E911 requirements, additional gateways might be needed at each site.

Voicemail for Local Failover

Cisco Unity or other voicemail systems can be deployed at all sites and integrated into the Cisco Unified CallManager cluster. This configuration provides voicemail access even during a WAN failure and without using the PSTN. Using Voice Mail Profiles, you can allocate the correct voicemail system for the site to the IP phones in the same location. You can configure a maximum of four voicemail systems per cluster that use the SMDI protocol, that are attached directly to the COM port on a subscriber, and that use the Cisco Messaging Interface (CMI).


Music on Hold and Media Resources for Local Failover

Music on hold (MoH) servers and other media resources such as conference bridges should be provisioned at each site, with sufficient capacity for the type and number of users. Through the use of media resource groups (MRGs) and media resource group lists (MRGLs), media resources are provided by the on-site resource and are available during a WAN failure.

Remote Failover Deployment Model

The remote failover deployment model provides flexibility for the placement of backup servers. Each of the sites contains at least one primary Cisco Unified CallManager subscriber and may or may not have a backup subscriber. This model allows for a deployment of up to eight sites, with IP phones and other devices normally registered to a local subscriber when using 1:1 redundancy and the 50/50 load balancing option described in the chapter on [Call Processing, page 8-1](#). Backup subscribers are located across the WAN at one or more of the other sites. (See [Figure 2-6](#).)

Figure 2-6 Remote Failover Model with Four Sites



When implementing the remote failover model, observe all guidelines for the local failover model (see [Local Failover Deployment Model, page 2-22](#)), with the following modifications:

- Configure each site to contain at least one primary Cisco Unified CallManager subscriber and an optional backup subscriber as desired.
- You may configure Cisco Unified CallManager *groups* and *device pools* to allow devices to register with servers over the WAN.
- Signaling or call control traffic requires bandwidth when devices are registered across the WAN with a remote Cisco Unified CallManager server in the same cluster. This bandwidth might be more than the ICCS traffic and should be calculated using the bandwidth provisioning calculations for signalling, as described in [Bandwidth Provisioning, page 3-44](#).

Design Considerations for Section 508 Conformance

Regardless of which deployment model you choose, you should consider designing your IP telephony network to make the telephony features more accessible to users with disabilities, in conformance with Section 255 of the Telecommunications Act and U.S. Section 508.

Observe the following basic design guidelines when configuring your IP telephony network to conform to Section 508:

- Enable Quality of Service (QoS) on the network.
- Configure only the G.711 codec for phones that will be connected to a terminal teletype (TTY) device or a Telephone Device for the Deaf (TDD). Although low bit-rate codecs such as G.729 are acceptable for audio transmissions, they do not work well for TTY/TDD devices if they have an error rate higher than 1% Total Character Error Rate (TCER).
- Configure TTY/TDD devices for G.711 across the WAN, if necessary.
- Enable (turn ON) Echo Cancellation for optimal performance.
- Voice Activity Detection (VAD) does not appear to have an effect on the quality of the TTY/TDD connection, so it may be disabled or enabled.
- Configure the appropriate *regions* and *device pools* in Cisco Unified CallManager to ensure that the TTY/TDD devices always use G.711 codes.
- Connect the TTY/TDD to the IP telephony network in either of the following ways:
 - Direct connection (Recommended method)

Plug a TTY/TDD with an RJ-11 analog line option directly into a Cisco FXS port. Any FXS port will work, such as the one on the Cisco VG248, Catalyst 6000, Cisco ATA 188 module, or any other Cisco voice gateway with an FXS port. Cisco recommends this method of connection.
 - Acoustic coupling

Place the IP phone handset into a coupling device on the TTY/TDD. Acoustic coupling is less reliable than an RJ-11 connection because the coupling device is generally more susceptible to transmission errors caused by ambient room noise and other factors.
- If stutter dial tone is required, use an analog phone in conjunction with an FXS port on the Cisco VG248 or ATA 188.



CHAPTER 3

Network Infrastructure

Last revised on: October 30, 2008


This chapter describes the requirements of the network infrastructure needed to build an IP telephony system in an enterprise environment. [Figure 3-1](#) illustrates the roles of the various devices that form the network infrastructure, and [Table 3-1](#) summarizes the features required to support each of these roles.

IP telephony places strict requirements on IP packet loss, packet delay, and delay variation (or jitter). Therefore, you need to enable most of the Quality of Service (QoS) mechanisms available on Cisco switches and routers throughout the network. For the same reasons, redundant devices and network links that provide quick convergence after network failures or topology changes are also important to ensure a highly available infrastructure

The following sections describe the network infrastructure features as they relate to:

- [LAN Infrastructure, page 3-4](#)
- [WAN Infrastructure, page 3-25](#)
- [Wireless LAN Infrastructure, page 3-58](#)

Figure 3-1 Typical Campus Network Infrastructure



77290

Table 3-1 Required Features for Each Role in the Network Infrastructure

Infrastructure Role	Required Features
Campus Access Switch	• In-Line Power • Multiple Queue Support • 802.1p and 802.1Q • Fast Link Convergence
Campus Distribution or Core Switch	• Multiple Queue Support • 802.1p and 802.1Q • Traffic Classification • Traffic Reclassification
WAN Aggregation Router (Site that is at the hub of the network)	• Multiple Queue Support • Traffic Shaping • Link Fragmentation and Interleaving (LFI) • Link Efficiency • Traffic Classification • Traffic Reclassification • 802.1p and 802.1Q
Branch Router (Spoke site)	• Multiple Queue Support • LFI • Link Efficiency • Traffic Classification • Traffic Reclassification • 802.1p and 802.1Q
Branch or Smaller Site Switch	• In-Line Power • Multiple Queue Support • 802.1p and 802.1Q

What's New in This Chapter

[Table 3-2](#) lists the topics that are new in this chapter or that have changed significantly from previous releases of this document.

Table 3-2 New or Changed Information Since the Previous Release of This Document

New or Revised Topic	Described in:
Different versions of software on Trivial File Transfer Protocol (TFTP) servers	Multi-Cluster Campus TFTP Services, page 3-15

LAN Infrastructure

Campus LAN infrastructure design is extremely important for proper IP telephony operation on a converged network. Proper LAN infrastructure design requires following basic configuration and design best practices for deploying a highly available network. Further, proper LAN infrastructure design requires deploying end-to-end QoS on the network. The following sections discuss these requirements:

- [LAN Design for High Availability, page 3-4](#)
- [LAN Quality of Service \(QoS\), page 3-21](#)

LAN Design for High Availability

Properly designing a LAN requires building a robust and redundant network from the top down. By structuring the LAN as a layered model (see [Figure 3-1](#)) and developing the LAN infrastructure one step of the model at a time, you can build a highly available, fault tolerant, and redundant network. Once these layers have been designed properly, you can add network services such as DHCP and TFTP to provide additional network functionality. The following sections examine the infrastructure layers and network services:


- [Campus Access Layer, page 3-4](#)
- [Campus Distribution Layer, page 3-7](#)
- [Campus Core Layer, page 3-9](#)
- [Network Services, page 3-10](#)

For more information on campus design, refer to the *Gigabit Campus Network Design* white paper at [http://www.cisco.com/warp/public/cc/so/neso/lnso/cpso/gcnd_wp.pdf](http://www.cisco.com/warp/public/cc/so/neso/lnso/cpso/gcnd/wp.pdf)

Campus Access Layer

The access layer of the Campus LAN includes the portion of the network from the desktop port(s) to the wiring closet switch.

Proper access layer design starts with assigning a single IP subnet per virtual LAN (VLAN). Typically, a VLAN should not span multiple wiring closet switches; that is, a VLAN should have presence in one and only one access layer switch (see [Figure 3-2](#)). This practice eliminates topological loops at Layer 2, thus avoiding temporary flow interruptions due to Spanning Tree convergence. However, with the introduction of standards-based IEEE 802.1w Rapid Spanning Tree Protocol (RSTP) and 802.1s Multiple Instance Spanning Tree Protocol (MISTP), Spanning Tree can converge at much higher rates. More importantly, confining a VLAN to a single access layer switch also serves to limit the size of the broadcast domain. There is the potential for large numbers of devices within a single VLAN or broadcast domain to generate large amounts of broadcast traffic periodically, which can be problematic. A good rule of thumb is to limit the number of devices per VLAN to about 512, which is equivalent to two Class C subnets (that is, a 23-bit subnet masked Class C address). Typical access layer switches include the stackable Cisco Catalyst 2950, 3500XL, 3550, and 3750, as well as the Cisco 3560 and the larger, higher-density Catalyst 4000 and 6000 switches.

Figure 3-2 Access Layer Switches and VLANs for Voice and Data

When you deploy voice, Cisco recommends that you enable two VLANs at the access layer: a native VLAN for data traffic (VLANs 10, 11, 30, 31, and 32 in [Figure 3-2](#)) and a voice VLAN under Cisco IOS or Auxiliary VLAN under CatOS for voice traffic (represented by VVIDs 110, 111, 310, 311, and 312 in [Figure 3-2](#)).

Separate voice and data VLANs are recommended for the following reasons:

- Address space conservation and voice device protection from external networks
Private addressing of phones on the voice or auxiliary VLAN ensures address conservation and ensures that phones are not accessible directly via public networks. PCs and servers are typically addressed with publicly routed subnet addresses; however, voice endpoints should be addressed using RFC 1918 private subnet addresses.
- QoS trust boundary extension to voice devices
QoS trust boundaries can be extended to voice devices without extending these trust boundaries and, in turn, QoS features to PCs and other data devices.
- Protection from malicious network attacks
VLAN access control, 802.1Q, and 802.1p tagging can provide protection for voice devices from malicious internal and external network attacks such as worms, denial of service (DoS) attacks, and attempts by data devices to gain access to priority queues via packet tagging.
- Ease of management and configuration
Separate VLANs for voice and data devices at the access layer provide ease of management and simplified QoS configuration.

To provide high-quality voice and to take advantage of the full voice feature set, access layer switches should provide support for:

- 802.1Q trunking and 802.1p for proper treatment of Layer 2 CoS packet marking on ports with phones connected
- Multiple egress queues to provide priority queuing of RTP voice packet streams

- The ability to classify or reclassify traffic and establish a network trust boundary
- Inline power capability (Although inline power capability is not mandatory, it is highly recommended for the access layer switches.)
- Layer 3 awareness and the ability to implement QoS access control lists (These features are required if you are using certain IP telephony endpoints, such as a PC running a softphone application, that cannot benefit from an extended trust boundary.)

Spanning Tree Protocol (STP)

To minimize convergence times and maximize fault tolerance at Layer 2, enable the following STP features:

- PortFast
Enable PortFast on all access ports. The phones, PCs, or servers connected to these ports do not forward bridge protocol data units (BPDUs) that could affect STP operation. PortFast ensures that the phone or PC, when connected to the port, is able to begin receiving and transmitting traffic immediately without having to wait for STP to converge.
- Root guard or BPDU guard
Enable root guard or BPDU guard on all access ports to prevent the introduction of a rogue switch that might attempt to become the Spanning Tree root, thereby causing STP re-convergence events and potentially interrupting network traffic flows. Ports that are set to **errdisable** state by BPDU guard must either be re-enabled manually or the switch must be configured to re-enable ports automatically from the errdisable state after a configured period of time.
- UplinkFast and BackboneFast
Enable these features where appropriate to ensure that, when changes occur on the Layer 2 network, STP converges as rapidly as possible to provide high availability. When using stackable switches such as the Catalyst 2950, 3550, or 3750, enable Cross-Stack UplinkFast (CSUF) to provide fast failover and convergence if a switch in the stack fails.
- UniDirectional Link Detection (UDLD)
Enable this feature to reduce convergence and downtime on the network when link failures or misbehaviors occur, thus ensuring minimal interruption of network service. UDLD detects, and takes out of service, links where traffic is flowing in only one direction. This feature prevents defective links from being mistakenly considered as part of the network topology by the Spanning Tree and routing protocols.



Note

With the introduction of RSTP 802.1w, features such as PortFast and UplinkFast are not required because these mechanisms are built in to this standard. If RSTP has been enabled on the Catalyst switch, these commands are not necessary.

Campus Distribution Layer

The distribution layer of the Campus LAN includes the portion of the network from the wiring closet switches to the next-hop switch, and it is the first Layer-2-to-Layer-3 traversal in the LAN. Distribution layer switches typically include Layer 3-enabled Catalyst 4000 and 6000 switches and the Catalyst 3750 for smaller deployments.

At the distribution layer, it is important to provide redundancy to ensure high availability, including redundant links between the distribution layer switches (or routers) and the access layer switches. To avoid creating topological loops at Layer 2, use Layer 3 links for the connections between redundant Distribution switches when possible.

Hot Standby Router Protocol (HSRP)

HSRP should also be enabled at the distribution layer to ensure that all routers are made redundant and that, in the event of a failure, another router can take over. HSRP configuration should incorporate the following:

- Standby track

The **standby track** command indicates that the HSRP should monitor a particular interface(s). If the interface goes down, then the HSRP priority of the box is reduced, typically forcing a failover to another device. This command is used in conjunction with the **standby preempt** command.

- Standby preempt

This command ensures that, when a device's priority becomes higher than all the other HSRP-configured devices in the standby group, that device will take over as the active Layer 3 router for the HSRP standby address.


HSRP should also be configured in such a way as to load-balance traffic between both HSRP routers. To provide load balancing, configure each HSRP device as the active HSRP router for one VLAN or interface, and configure the standby router for another VLAN or interface. Evenly distributing active and standby VLANs between both HSRP devices ensures load-balancing. Devices on one VLAN use the active HSRP device as their default gateway, and devices on another VLAN use the same HSRP device as a standby default gateway only if the other HSRP device fails. This type of configuration prevents all network traffic from being sent to a single active router and enables other HSRP devices to help carry the load.

[Figure 3-3](#) shows an example of an HSRP-enabled network. In this figure, the two Catalyst 6500 switches (6500-SW1 and 6500-SW2) have been configured with multiple VLAN interfaces. Assuming that there are no link failures in the network, 6500-SW1 is the standby HSRP router for VLAN 110 (the voice VLAN for Group A phones) and is the active HSRP router for VLAN 10 (the data VLAN) and for VLAN 120 (the voice VLAN for Group B phones). 6500-SW2 is configured in reverse; it is the active HSRP router for VLAN 110 and the standby HSRP router for VLAN 10 and VLAN 120. As configured, both switches are actively in use, and the load can be distributed between the two by evenly distributing all Layer 2 VLANs between them. Each switch is also configured to track its local VLAN 200 interface and, in the event of a VLAN 200 link failure, the other switch will preempt and become the active router for all VLANs. Likewise, if either switch fails, the other switch will handle the traffic for all three VLANs.

The PCs and phones at the access layer in [Figure 3-3](#) have been configured with default gateways that correspond to the HSRP addresses for each of the HSRP groups. Devices in voice VLANs 110 and 120 are pointing to 10.100.10.1 and 10.100.20.1, respectively, as the default gateways, which correspond to the HSRP addresses for the VLAN 110 and 120 interfaces on both switches. Devices in data VLAN 10 are pointing to 64.100.10.1 as the default gateway, which corresponds to the HSRP address of the VLAN 10 interface on both switches. Note that, while traffic flowing from the access layer to the

distribution layer will be distributed between the two switches (as long as there are no failures), no mechanism exists to ensure distribution on the return path. Traffic returning from the core layer and destined for the access layer will follow the shortest and/or least costly routed path.

Figure 3-3 HSRP Network Configuration Example with Standby Preempt and Standby Track



Example 3-1 and Example 3-2 list the configurations for the two Catalyst 6500 switches shown in Figure 3-3.

Example 3-1 Configuration for 6500-SW1

```
interface Vlan 10
description Data VLAN 10
ip address 64.100.10.11 255.255.255.0
standby preempt
standby ip 64.100.10.1
standby track Vlan 200

interface Vlan110
description Voice VLAN 110
ip address 10.100.10.11 255.255.255.0
standby preempt
standby ip 10.100.10.1
standby track Vlan 200
standby priority 95

interface Vlan120
description Voice VLAN 120
ip address 10.100.20.11 255.255.255.0
standby preempt
standby ip 10.100.20.1
standby track Vlan 200
```

114399

Example 3-2 Configuration for 6500-SW2

```
interface Vlan 10
  description Data VLAN 10
  ip address 64.100.10.12 255.255.255.0
  standby preempt
  standby ip 64.100.10.1
  standby track Vlan 200
  standby priority 95

interface Vlan110
  description Voice VLAN 110
  ip address 10.100.10.12 255.255.255.0
  standby preempt
  standby ip 10.100.10.1
  standby track Vlan 200

interface Vlan120
  description Voice VLAN 120
  ip address 10.100.20.11 255.255.255.0
  standby preempt
  standby ip 10.100.20.1
  standby track Vlan 200
  standby priority 95
```

How quickly HSRP converges when a failure occurs depends on how the HSRP hello and hold timers are configured. By default, these timers are set to 3 and 10 seconds respectively, which means that an hello packet will be sent between the HSRP standby group devices every 3 seconds and that the standby device will become active when an hello packet has not been received for 10 seconds. You can lower these timer settings to speed up the failover or preemption; however, to avoid increased CPU usage and unnecessary standby state flapping, do not set the hello timer below one (1) second or the hold timer below 4 seconds. Note that, if you are using the HSRP tracking mechanism and the tracked link fails, then the failover or preemption occurs immediately regardless of the hello and hold timers.

Routing Protocols

Configure Layer 3 routing protocols, such as Open Shortest Path First (OSPF) and Enhanced Interior Gateway Routing Protocol (EIGRP), at the distribution layer to ensure fast convergence, load balancing, and fault tolerance. Use parameters such as routing protocol timers, path or link costs, and address summaries to optimize and control convergence times as well as to distribute traffic across multiple paths and devices. Cisco also recommends using the **passive-interface** command to prevent routing neighbor adjacencies via the access layer. These adjacencies are typically unnecessary, and they create extra CPU overhead and increased memory utilization because the routing protocol keeps track of them. By using the **passive-interface** command on all interfaces facing the access layer, you prevent routing updates from being sent out on these interfaces and, therefore, neighbor adjacencies are not formed.

Campus Core Layer

The core layer of the Campus LAN includes the portion of the network from the distribution routers or Layer 3 switches to one or more high-end core Layer 3 switches or routers. Layer 3-capable Catalyst 6000 switches are the typical core layer devices, and these core switches can provide connectivity between numerous campus distribution layers.

At the core layer, it is again very important to provide the following types of redundancy to ensure high availability:

- Redundant link or cable paths

Redundancy here ensures that traffic can be rerouted around downed or malfunctioning links.

- Redundant devices

Redundancy here ensures that, in the event of a device failure, another device in the network can continue performing tasks that the failed device was doing.

- Redundant device sub-systems

This type of redundancy ensures that multiple power supplies and Supervisor engines are available within a device so that the device can continue to function in the event that one of these components fails.

Routing protocols at the core layer should again be configured and optimized for path redundancy and fast convergence. There should be no STP in the core because network connectivity should be routed at Layer 3. Finally, each link between the core and distribution devices should belong to its own VLAN or subnet and be configured using a 30-bit subnet mask.

Data Center and Server Farm

Typically, Cisco Unified CallManager cluster servers, including media resource servers, reside in a data center or server farm environment. In addition, centralized gateways and centralized hardware media resources such as conference bridges, DSP or transcoder farms, and media termination points are located in the data center or server farm. Because these servers and resources are critical to voice networks, Cisco recommends distributing all Cisco Unified CallManager cluster servers, centralized voice gateways, and centralized hardware resources between multiple physical switches and, if possible, multiple physical locations within the campus. This distribution of resources ensures that, given a hardware failure (such as a switch or switch line card failure), at least some servers in the cluster will still be available to provide telephony services. In addition, some gateways and hardware resources will still be available to provide access to the PSTN and to provide auxiliary services. Besides being physically distributed, these servers, gateways, and hardware resources should be distributed among separate VLANs or subnets so that, if a broadcast storm or denial of service attack occurs on a particular VLAN, not all voice connectivity and services will be disrupted.

Network Services

The deployment of an IP Communications system requires the coordinated design of a well structured, highly available, and resilient network infrastructure as well as an integrated set of network services including Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), Trivial File Transfer Protocol (TFTP), and Network Time Protocol (NTP).

Domain Name System (DNS)

DNS enables the mapping of host names and network services to IP addresses within a network or networks. DNS server(s) deployed within a network provide a database that maps network services to hostnames and, in turn, hostnames to IP addresses. Devices on the network can query the DNS server and receive IP addresses for other devices in the network, thereby facilitating communication between network devices.

Relying on DNS, however, can be problematic. If the DNS server becomes unavailable and a network device is relying on that server to provide a hostname-to-IP-address mapping, communication can and will fail. For this reason, do not rely on DNS for communication between Cisco Unified CallManager and the IP telephony endpoints.

Configure Cisco Unified CallManager(s), gateways, and endpoint devices to use IP addresses rather than hostnames. Cisco does *not* recommend configuration of DNS parameters such as DNS server addresses, hostnames, and domain names. Likewise, configuring HOSTS files is discouraged because the

administration of these files in a large IP telephony network with thousands of endpoints and servers can be extremely time-consuming and inefficient. If you eliminate DNS configuration within the IP telephony network, telephony devices and applications do not have to rely on the DNS server.

You must, however, configure the LMHOSTS file. This file provides a mechanism for resolving or mapping server hostnames or NetBios names to IP addresses. The LMHOSTS file must contain a list of server names and corresponding IP address. There should be an entry for every server within the cluster as well as an entry with **127.0.0.1 localhost** (loopback entry). You should copy this file to every server in the cluster. The LMHOSTS file is located in the directory path C:\WINNT\system32\drivers\etc.

[Example 3-3](#) shows a typical LMHOSTS file for a cluster with six servers.

Example 3-3 LMHOSTS File

```
127.0.0.1      localhost      ! The local host entry
64.101.1.7     ccml-hq-1
64.101.2.7     ccml-hq-2
64.101.1.8     ccml-hq-3
64.101.2.8     ccml-hq-4
64.101.1.21    ccml-moh-1
64.101.2.21    ccml-moh-2
```

There are some situations in which configuring and using DNS might be unavoidable. For example, if Network Address Translation (NAT) is required for communication between the IP phones and Cisco Unified CallManager in the telephony network, DNS is required to ensure proper mapping of NAT translated addresses to network host devices. Likewise, some IP telephony disaster recovery network configurations rely on DNS to ensure proper failover of the network during failure scenarios by mapping hostnames to secondary backup site IP addresses.

If either of these two situations exists and DNS must be configured, you must deploy DNS servers in a redundant fashion so that a single DNS server failure will not prevent network communication between IP telephony devices. By providing DNS server redundancy in the event of a single DNS server failure, you ensure that devices relying on DNS to communicate on the network can still receive hostname-to-IP-address mappings from a backup or secondary DNS server.



Note

DNS names within the cluster are resolved only at system initialization (that is, when a server is booted up). As a result, in order for a server within the cluster to resolve a DNS name that has been changed on the DNS server, the Cisco Unified CallManager service must be restarted on all servers within the cluster.

Dynamic Host Configuration Protocol (DHCP)

DHCP is used by hosts on the network to obtain initial configuration information, including IP address, subnet mask, default gateway, and TFTP server address. DHCP eases the administrative burden of manually configuring each host with an IP address and other configuration information. DHCP also provides automatic reconfiguration of network configuration when devices are moved between subnets. The configuration information is provided by a DHCP server located in the network, which responds to DHCP requests from DHCP-capable clients.

You should configure IP Communications endpoints to use DHCP to simplify deployment of these devices. Any RFC 2131 compliant DHCP server can be used to provide configuration information to IP Communications network devices. When deploying IP telephony devices in an existing data-only network, all you have to do is add DHCP voice scopes to an existing DHCP server for these new voice devices. Because IP telephony devices are configured to use and rely on a DHCP server for IP configuration information, you must deploy DHCP servers in a redundant fashion. At least two DHCP

servers should be deployed within the telephony network such that, if one of the servers fails, the other can continue to answer DHCP client requests. You should also ensure that DHCP server(s) are configured with enough IP subnet addresses to handle all DHCP-reliant clients within the network.

DHCP Option 150

IP telephony endpoints can be configured to rely on DHCP Option 150 to identify the source of telephony configuration information, available from a server running the Trivial File Transfer Protocol (TFTP).

In the simplest configuration, where a single TFTP server is offering service to all deployed endpoints, Option 150 is delivered as a single IP address pointing to the system's designated TFTP server. The DHCP scope can also deliver two IP addresses under Option 150, for deployments where there are two TFTP servers within the same cluster. The phone would use the second address if it fails to contact the primary TFTP server, thus providing redundancy. To achieve both redundancy and load sharing between the TFTP servers, you can configure Option 150 to provide the two TFTP server addresses in reverse order for half of the DHCP scopes.



Note

If the primary TFTP server is available but is not able to grant the requested file to the phone (for example, because the requesting phone is not configured on that cluster), the phone will not attempt to contact the secondary TFTP server.

Cisco highly recommends using a direct IP address (that is, not relying on a DNS service) for Option 150 because doing so eliminates dependencies on DNS service availability during the phone boot-up and registration process.



Note

Even though IP phones support a maximum of two TFTP servers under Option 150, you could configure a cluster with more than two TFTP servers. For instance, if a Cisco Unified CallManager system is clustered over a WAN at three separate sites, three TFTP servers could be deployed (one at each site). Phones within each site could then be granted a DHCP scope containing that site's TFTP server within Option 150. This configuration would bring the TFTP service closer to the endpoints, thus reducing latency and ensuring failure isolation between the sites (one site's failure would not affect TFTP service at another site). However, when the configuration file is modified, the publisher must send a new copy of the database to each TFTP server in the cluster. This propagation of the database consumes publisher CPU resources and can slow performance if more than two TFTP servers are deployed in the cluster.

DHCP Lease Times

Configure DHCP lease times as appropriate for the network environment. Given a fairly static network in which PCs and telephony devices remain in the same place for long periods of time, Cisco recommends longer DHCP lease times (for example, one week). Shorter lease times require more frequent renewal of the DHCP configuration and increase the amount of DHCP traffic on the network. Conversely, networks that incorporate large numbers of mobile devices, such as laptops and wireless telephony devices, should be configured with shorter DHCP lease times (for example, one day) to prevent depletion of DHCP-managed subnet addresses. Mobile devices typically use IP addresses for short increments of time and then might not request a DHCP renewal or new address for a long period of time. Longer lease times will tie up these IP addresses and prevent them from being reassigned even when they are no longer being used.

Cisco Unified IP Phones adhere to the conditions of the DHCP lease duration as specified in the DHCP server's scope configuration. Once half the lease time has expired since the last successful DHCP server acknowledgment, the IP phone will request a lease renewal. This DHCP client Request, once acknowledged by the DHCP server, will allow the IP phone to retain use of the IP scope (that is, the IP

address, default gateway, subnet mask, DNS server (optional), and TFTP server (optional)) for another lease period. If the DHCP server becomes unavailable, an IP phone will not be able to renew its DHCP lease, and as soon as the lease expires, it will relinquish its IP configuration and will thus become unregistered from Cisco Unified CallManager until a DHCP server can grant it another valid scope.

In centralized call processing deployments, if a remote site is configured to use a centralized DHCP server (through the use of a DHCP relay agent such as the IP Helper Address in Cisco IOS) and if connectivity to the central site is severed, IP phones within the branch will not be able to renew their DHCP scope leases. In this situation, branch IP phones are at risk of seeing their DHCP lease expire, thus losing the use of their IP address, which would lead to service interruption. Given the fact that phones attempt to renew their leases at half the lease time, DHCP lease expiration can occur as soon as half the lease time since the DHCP server became unreachable. For example, if the lease time of a DHCP scope is set to 4 days and a WAN failure causes the DHCP server to be unavailable to the phones in a branch, those phones will be unable to renew their leases at half the lease time (in this case, 2 days). The IP phones could stop functioning as early as 2 days after the WAN failure, unless the WAN comes back up and the DHCP server is available before that time. If the WAN connectivity failure persists, all phones see their DHCP scope expire after a maximum of 4 days from the WAN failure.

This situation can be mitigated by one of the following methods:

- Set the DHCP scope lease to a long duration (for example, 8 days or more).

This method would give the system administrator a minimum of half the lease time to remedy any DHCP reachability problem. Long lease durations also have the effect of reducing the frequency of network traffic associated with lease renewals.

- Configure co-located DHCP server functionality (for example, run a DHCP server function on the branch's Cisco IOS router).

This approach is immune to WAN connectivity interruption. One effect of such an approach is to decentralize the management of IP addresses, requiring incremental configuration efforts in each branch. (See [DHCP Network Deployments, page 3-13](#), for more information.)

DHCP Network Deployments

There are two options for deploying DHCP functionality within an IP telephony network:

- Centralized DHCP Server

Typically, for a single-site campus IP telephony deployment, the DHCP server should be installed at a central location within the campus. As mentioned previously, redundant DHCP servers should be deployed. If the IP telephony deployment also incorporates remote branch telephony sites, as in a centralized multisite Cisco Unified CallManager deployment, a centralized server can be used to provide DHCP service to devices in the remote sites. This type of deployment requires that you configure the **ip helper-address** on the branch router interface. Keep in mind that, if redundant DHCP servers are deployed at the central site, both servers' IP addresses must be configured as **ip helper-address**. Also note that, if branch-side telephony devices rely on a centralized DHCP server and the WAN link between the two sites fails, devices at the branch site will be unable to send DHCP requests or receive DHCP responses.



Note

By default, **service dhcp** is enabled on the Cisco IOS device and does not appear in the configuration. Do not disable this service on the branch router because doing so will disable the DHCP relay agent on the device, and the **ip helper-address** configuration command will not work.

- Centralized DHCP Server and Remote Site Cisco IOS DHCP Server

When configuring DHCP for use in a centralized multisite Cisco Unified CallManager deployment, you can use a centralized DHCP server to provide DHCP service to centrally located devices. Remote devices could receive DHCP service from a locally installed server or from the Cisco IOS router at the remote site. This type of deployment ensures that DHCP services are available to remote telephony devices even during WAN failures. [Example 3-4](#) lists the basic Cisco IOS DHCP server configuration commands.

Example 3-4 Cisco IOS DHCP Server Configuration Commands

```
! Activate DHCP Service on the IOS Device

service dhcp

! Specify any IP Address or IP Address Range to be excluded from the DHCP pool

ip dhcp excluded-address <ip-address>|<ip-address-low> <ip-address-high>

! Specify the name of this specific DHCP pool, the subnet and mask for this
! pool, the default gateway and up to four TFTP

ip dhcp pool <dhcp-pool name>
  network <ip-subnet> <mask>
  default-router <default-gateway-ip>
  option 150 ip <tftp-server-ip-1> ...

! Note: IP phones use only the first two addresses supplied in the option 150
! field even if more than two are configured.
```

Cisco Unified CallManager DHCP Sever (Standalone versus Co-Resident Server)

Typically DHCP servers are dedicated machine(s) in most network infrastructures, and they run in conjunction with the DNS and Windows Internet Naming Service (WINS) services used by that network. In some instances, given a small Cisco Unified CallManager deployment with no more than 1000 devices registering to the cluster, you may run DHCP on a Cisco Unified CallManager server to support those devices. However, if the server experiences high CPU load, you should move DHCP to a standalone server. If more than 1000 devices are registered to the cluster, DHCP must not be run on a Cisco Unified CallManager server and must be run on a dedicated or standalone server(s).

Trivial File Transfer Protocol (TFTP)

Within a Cisco Unified CallManager system, endpoints (such as IP phones running the SCCP protocol) rely on a TFTP-based process to acquire configuration information. The endpoints request a configuration file whose name is based on the requester's MAC address (for example, for an IP phone with MAC address ABCDEF123456, the file name would be SEPABCDEF123456.cnf.xml). The configuration file includes the version of software that the phone must run and a list of Cisco Unified CallManager servers with which the phone should register.

If the configuration file instructs the phone to run a software file other than the one it currently uses, the phone will request the new version of software from the TFTP server. The phone goes through this process once per software upgrade.

Centralized call processing deployments require remote phones to download configuration files and phone software through the branch's WAN link. When scheduled maintenance involves the downloading of new software, download times are a function of the number of phones requiring upgrades, the file size, and the WAN link's bandwidth and traffic utilization.

For example, in Cisco Unified CallManager Release 4.2, the size of a phone configuration file is approximately 3250 bytes, and the combined size of the required software load files (P00308000300.loads, P00308000300.sbn, and P00308000300.sb2) for a Cisco Unified IP Phone 7960 is 830,845 bytes. If a branch has a WAN bandwidth of 256 Kbps available to download the software, a single phone would require about 26 seconds to download new software during an upgrade. If that same branch has 10 phones requiring the new software, the download process would take about 4.5 minutes.

Multi-Cluster Campus TFTP Services

In multi-cluster systems, it is possible to have a single subnet or VLAN containing phones from multiple clusters. In this situation, the TFTP server whose address is provided to all phones in the subnet or VLAN must answer the file transfer requests made by each phone, regardless of which cluster contains the phone. Therefore, this centralized TFTP server must have access to files created and managed by other clusters.

To provide this file access, each cluster's TFTP server must be configured to create and manage configuration files on the centralized TFTP server's drive. Perform the configuration by using the alternate file location entry under each TFTP server's configuration (with the exception of the centralized TFTP server).

With Cisco Unified CallManager Release 3.2 and later, Cisco TFTP servers cache the IP phone configuration files in RAM by default. For those files to be written to a centralized TFTP server, you must disable (turn off) file caching by setting the following service parameters as indicated on each TFTP server configured to write to the centralized TFTP server:

- Enable Caching of Configuration Files: **False** (required)
- Enable Caching of Constant and Bin Files at Startup: **False** (recommended)

If the TFTP server receives a request for a file that it does not have (such as a configuration file created and maintained by the TFTP server of a different cluster), it will search for that file in a list of alternate file locations. The centralized TFTP server must be configured to search through the subdirectories associated with the other clusters.

For centralized TFTP configurations, ensure that the main TFTP server exists in the cluster that runs the highest version of Cisco Unified CallManager. For example, if you are using a centralized TFTP server between a compatible Cisco CallManager 3.x cluster and a Cisco Unified CallManager 4.x cluster, ensure that the central TFTP server exists in the Cisco Unified CallManager 4.x cluster.

If the centralized TFTP server exists in the cluster that runs the lower version of Cisco Unified Communications Manager, all phones use the locale files from this centralized TFTP server. These older locale files can create display issues for phones registered to clusters running a higher version of Cisco Unified CallManager because the newer localized phrases are not included in the locale files that are served from the main cluster's TFTP server.

Example 3-5 Alternate TFTP File Locations

A large campus system is deployed using three clusters, and each cluster contains a TFTP server. TFTP1, the TFTP server for Cluster1, is configured as the centralized TFTP server for the campus. The other clusters and TFTP servers are named in sequence as TFTP2 for Cluster2 and TFTP3 for Cluster3. In all subnets, the DHCP scope provides TFTP1's IP address as Option 150.

First, TFTP2 and TFTP3 are configured to write their configuration files to TFTP1's drive, each in a different subdirectory, as follows:

- TFTP2's alternate file location is set to: \\TFTP1_IP\\Program Files\\Cisco\\TFTPpath\\TFTP2
- TFTP3's alternate file location is set to: \\TFTP1_IP\\Program Files\\Cisco\\TFTPpath\\TFTP3

Second, TFTP1 is configured to search in the alternate file locations as follows:

- Alternate File Location 1: c:\Program Files\Cisco\TFTPpath\TFTP2
- Alternate File Location 2: c:\Program Files\Cisco\TFTPpath\TFTP3



Note In this example, TFTP1_IP represents the IP address of TFTP1. Also, TFTP1 requires that Windows NT subdirectories be created manually for TFTP2 and TFTP3.

Cisco recommends that you use Universal Naming Convention (UNC) paths (in the format \\<IP_address>\<Full_path_to_folder>) to point a TFTP server to alternate file locations. Cisco does not recommend creating non-default NT "shares" or using DNS names. Also, ensure that all clusters meet the proper login ID, password, and security privileges (workgroup, domain, or directory-based) for the Cisco TFTP service.

TFTP Server Redundancy

Option 150 allows up to two IP addresses to be returned to phones as part of the DHCP scope. The phone tries the first address in the list, and it tries the subsequent address only if it cannot establish communications with the first TFTP server. This address list provides a redundancy mechanism that enables phones to obtain TFTP services from another server even if their primary TFTP server has failed.


As illustrated in [Figure 3-4](#), two TFTP servers can be configured in a cluster, and each can create and manage separate lists of the same configuration files. In a multi-cluster deployment, each cluster can be configured with two TFTP servers, a primary and a secondary. The primary TFTP servers can be configured to write files to a centralized primary TFTP server; likewise, the secondary TFTP servers can be configured to write files to a centralized secondary TFTP server. This creates two separate groups (primary and secondary) of TFTP servers configured to ensure redundancy, each with a member serving as the centralized server.

Example 3-6 TFTP Server Redundancy

If we wanted to provide TFTP redundancy for the case described in [Example 3-5](#), we could configure each cluster with two TFTP servers. All primary TFTP servers would be configured to write their configuration files to TFTP1_P, while all the secondary TFTP servers would write theirs to TFTP1_S, as follows:

- TFTP2_P's alternate file location is set to: \\TFTP1_P\Program Files\Cisco\TFTPpath\TFTP2.
- TFTP3_P's alternate file location is set to: \\TFTP1_P\Program Files\Cisco\TFTPpath\TFTP3.
- TFTP2_S's alternate file location is set to: \\TFTP1_S\Program Files\Cisco\TFTPpath\TFTP2.
- TFTP3_S's alternate file location is set to: \\TFTP1_S\Program Files\Cisco\TFTPpath\TFTP3.

Both TFTP1_P and TFTP1_S must be configured as in [Example 3-5](#) to search through the list of alternate file locations.

Figure 3-4 TFTP Server Redundancy with Centralized TFTP Servers for All Clusters**TFTP Load Sharing**

The preceding sections explain how to use one TFTP server at a time to service phones from multiple clusters. For this approach, Cisco recommends that you grant different ordered lists of TFTP servers to different subnets to allow for load balancing. For example:

- In subnet 10.1.1.0/24: Option 150: TFTP1_P, TFTP1_S
- In subnet 10.1.2.0/24: Option 150: TFTP1_S, TFTP1_P

Under normal operations, a phone in subnet 10.1.1.0/24 will request TFTP services from TFTP1_P, while a phone in subnet 10.1.2.0/24 will request TFTP services from TFTP1_S. If TFTP1_P fails, then phones from both subnets will request TFTP services from TFTP1_S.

Load balancing avoids having a single TFTP server hot spot, where all phones from multiple clusters rely on the same server for service. TFTP load balancing is especially important when phone software loads are transferred, such as during a Cisco Unified CallManager upgrade, because more files of larger size are being transferred, thus imposing a bigger load on the TFTP server.

Network Time Protocol (NTP)

NTP allows network devices to synchronize their clocks to a network time server or network-capable clock. NTP is critical for ensuring that all devices in a network have the same time. When troubleshooting or managing a telephony network, it is crucial to synchronize the time stamps within all error and security logs, traces, and system reports on devices throughout the network. This synchronization enables administrators to recreate network activities and behaviors based on a common timeline. Billing records and call detail records (CDRs) also require accurate synchronized time.

Cisco Unified CallManager NTP Time Synchronization

Time synchronization is especially critical on Cisco Unified CallManager servers. You should configure automatic NTP time synchronization on all Cisco Unified CallManager servers within the network by performing the following steps:

Step 1 Configure the NTP.conf file.

The NTP.conf file is located in the C:\WINNT\ directory. The file must be configured with a list of one or more NTP Time servers that can be queried for the time. The file can also be configured to receive NTP Time server updates via NTP broadcasts on the local LAN segment. There must be a broadcast-capable NTP Time server available for Cisco Unified CallManager to receive the time via broadcast messages. [Example 3-7](#) illustrates both methods of configuring the NTP.conf file.

Step 2 Configure the NTP Service to start automatically at boot-up.

Under the Services application in Microsoft Windows on each Cisco Unified CallManager server, configure the NTP Service to start automatically at system boot-up.

Step 3 Use the Services application to start or restart the NTP Service on each Cisco Unified CallManager server.**Note**

If the NTP Service does not appear in the Microsoft Services control panel application on the Cisco Unified CallManager server, install it manually by running C:\Program Files\Cisco\Xntp>**install.bat**.

Example 3-7 NTP.conf Configuration Files

```
server 64.100.21.254
server 64.200.40.10
driftfile %windir%\ntp.drift
```

Or

```
broadcastclient
driftfile %windir%\ntp.drift
```

The driftfile referenced in [Example 3-7](#) is automatically updated via the NTP Service, based on information in the NTP messages received from the NTP Time server.

**Note**

By default, NTP updates will not occur if the clock offset or the adjustment between the NTP server and the NTP client is larger than 1000 seconds. To adjust this default behavior, you can add the **tinker panic** *<number of seconds>* command to the NTP.conf file, where the number of seconds determines the amount of slip time that can occur. Setting this value to 0 disables the panic feature, and a clock offset of any value will be accepted.

Cisco IOS and CatOS NTP Time Synchronization

Time synchronization is also important for other devices within the network. Cisco IOS routers and Catalyst switches should be configured to synchronize their time with the rest of the network devices via NTP. This is critical for ensuring that debug, syslog, and console log messages are time-stamped appropriately. Troubleshooting telephony network issues is simplified when a clear timeline can be drawn for events that occur on devices throughout the network.

[Example 3-8](#) illustrates the configuration of NTP time synchronization on Cisco IOS and CatOS devices.

Example 3-8 Cisco IOS and CatOS NTP Configuration

Cisco IOS configuration:

```
ntp server 64.100.21.254
```

CatOS configuration:

```
set ntp server 64.100.21.254
set ntp client enable
```

To ensure proper NTP time synchronization on routers and switches, it might be necessary to configure time zones using the **clock timezone** command (in Cisco IOS) and/or the **set timezone** command (in CatOS).

Power over Ethernet (PoE)

PoE (or inline power) is 48 Volt DC power provided over standard Ethernet unshielded twisted-pair (UTP) cable. Instead of using wall power, IP phones and other inline powered devices (PDs) such as the Aironet Wireless Access Points can receive power provided by inline power-capable Catalyst Ethernet switches or other inline power source equipment (PSE). Inline power is enabled by default on all inline power-capable Catalyst switches.

Deploying inline power-capable switches with uninterrupted power supplies (UPS) ensures that IP phones continue to receive power during power failure situations. Provided the rest of the telephony network is available during these periods of power failure, then IP phones should be able to continue making and receiving calls. You should deploy inline power-capable switches at the campus access layer within wiring closets to provide inline-powered Ethernet ports for IP phones, thus eliminating the need for wall power.

Cisco PoE is delivered on the same wire pairs used for data connectivity (pins 1, 2, 3, and 6). If existing access switch ports are not capable of inline power, you can use a power patch panel to inject power into the cabling. (In this case pins 4, 5, 7, and 8 are used.) Additionally, power injectors may be used for specific deployment needs.



Caution

The use of power injectors or power patch panels can damage some devices because power is always applied to the Ethernet pairs. PoE switch ports automatically detect the presence of a device that requires PoE before enabling it on a port-by-port basis.

In addition to Cisco PoE inline power, Cisco now supports the IEEE 802.3af PoE standard. Currently, only some access switches and phones comply with 802.3af. Over time, all phones and switches will support 802.3af PoE. The Catalyst 6500, 4500, and 3750 are currently capable of supporting 802.3af. For information about which Cisco Unified IP Phones support the 802.3af PoE standard, see the [Endpoint Features Summary, page 21-35](#).

Category 3 Cabling

The use of Category 3 cabling is supported for IP Communications under the following conditions:

- Phones with a PC port and a PC attached to it (Cisco Unified IP Phones 7971, 7970, 7961, 7960, 7941, 7940, 7911, and 7910+SW) should be set to 10 Mb, full-duplex.

This setting requires hard-coding the upstream switch port, the phone switch and PC ports, and the PC NIC port to 10 Mb, full-duplex. No ports should be set to AUTO negotiate. If desired, you can hard-code the phone's PC port to 10 Mb half-duplex, thereby forcing the PC's NIC to negotiate to 10 Mb half-duplex (assuming the PC's NIC is configured to AUTO negotiate). This configuration is acceptable as long as the uplink between the phone and the upstream switch port is set to 10 Mb full-duplex.

- Phones with no PC ports and with 10 Mb switch ports (Cisco Unified IP Phones 7902, 7905, and 7910) should be allowed to auto-negotiate to 10 Mb, half-duplex.

Because these phones support only 10 Mb Ethernet and their ports cannot be manually configured, the upstream switch port should be set to either AUTO negotiate or 10 Mb, half-duplex. In both cases, these phones will negotiate to 10 Mb, half-duplex.

- Phones with a PC port but no PC attached to it (Cisco Unified IP Phones 7971, 7970, 7961, 7960, 7941, 7940, 7912, 7911, and 7910+SW) can be allowed to negotiate to 10 Mb, half-duplex.

If you leave these phones with the default switch port configuration of AUTO negotiate and configure the upstream switch port to 10 Mb, half-duplex, these phones will revert to 10Mb, half-duplex.

**Note**

The Cisco Unified IP Phone 7912 should not be used with Category 3 cable when a PC is attached because the switch and PC ports on this phone cannot be forced to 10 Mb, full duplex.

IBM Type 1A and 2A Cabling

The use of IBM Cabling System (ICS) or Token Ring shielded twisted-pair type 1A or 2A cabling is supported for IP Communications under the following conditions:

- Cable lengths should be 100 meters or less.
- Adapters without impedance matching should be used for converting from universal data connector (UDC) to RJ-45 Ethernet standard.

**Note**

There are only two twisted pairs in the Token Ring cables. Therefore, inline power for IP phones can be supported, but mid-span power insertion cannot (with Cisco Inline Power and 802.3af) because it requires more than two pairs.

Running data over the network is not always a sufficient test of the quality of the cable plant because some non-compliance issues might not be apparent. Therefore, customers might want to perform a cable plant survey to verify that their type 1A and 2A cabling installation is compliant with Ethernet standards.

For more information about the use of IBM cabling, refer to the Product Bulletin *Shielded Twisted-Pair Cabling Support for Cisco Fast Ethernet Products*, available at


<http://www.cisco.com>

LAN Quality of Service (QoS)

Until recently, quality of service was not an issue in the enterprise campus due to the asynchronous nature of data traffic and the ability of network devices to tolerate buffer overflow and packet loss. However, with new applications such as voice and video, which are sensitive to packet loss and delay, buffers and not bandwidth are the key QoS issue in the enterprise campus.

Figure 3-5 illustrates the typical oversubscription that occurs in LAN infrastructures.

Figure 3-5 Data Traffic Oversubscription in the LAN



This oversubscription, coupled with individual traffic volumes and the cumulative effects of multiple independent traffic sources, can result in the egress interface buffers becoming full instantaneously, thus causing additional packets to drop when they attempt to enter the egress buffer. The fact that campus switches use hardware-based buffers, which compared to the interface speed are much smaller than those found on WAN interfaces in routers, merely increases the potential for even short-lived traffic bursts to cause buffer overflow and dropped packets.

Applications such as file sharing (both peer-to-peer and server-based), remote networked storage, network-based backup software, and emails with large attachments, can create conditions where network congestion occurs more frequently and/or for longer durations. Some of the negative effects of recent worm attacks have been an overwhelming volume of network traffic (both unicast and broadcast-storm based), increasing network congestion. If no buffer management policy is in place, loss, delay, and jitter performance of the LAN may be affected for all traffic.

Another situation to consider is the effect of failures of redundant network elements, which cause topology changes. For example, if a distribution switch fails, all traffic flows will be reestablished through the remaining distribution switch. Prior to the failure, the load balancing design shared the load between two switches, but after the failure all flows are concentrated in a single switch, potentially causing egress buffer conditions that normally would not be present.

For applications such as voice, this packet loss and delay results in severe voice quality degradation. Therefore, QoS tools are required to manage these buffers and to minimize packet loss, delay, and delay variation (jitter).

The following types of QoS tools are needed from end to end on the network to manage traffic and ensure voice quality:

- Traffic classification

Classification involves the marking of packets with a specific priority denoting a requirement for class of service (CoS) from the network. The point at which these packet markings are trusted or not trusted is considered the trust boundary. Trust is typically extended to voice devices (phones) and not to data devices (PCs).

- Queuing or scheduling

Interface queuing or scheduling involves assigning packets to one of several queues based on classification for expedited treatment throughout the network.

- Bandwidth provisioning

Provisioning involves accurately calculating the required bandwidth for all applications plus element overhead.

The following sections discuss the use of these QoS mechanisms in a campus environment:

- [Traffic Classification, page 3-22](#)
- [Interface Queuing, page 3-24](#)
- [Bandwidth Provisioning, page 3-24](#)
- [Impairments to IP Communications if QoS is Not Employed, page 3-24](#)

Traffic Classification

It has always been an integral part of the Cisco network design architecture to classify or mark traffic as close to the edge of the network as possible. Traffic classification is an entrance criterion for access into the various queuing schemes used within the campus switches and WAN interfaces. The IP phone marks its voice control signaling and voice RTP streams at the source, and it adheres to the values presented in [Table 3-3](#). As such, the IP phone can and should classify traffic flows.

[Table 3-3](#) lists the traffic classification requirements for the LAN infrastructure.

Table 3-3 Traffic Classification Guidelines for Various Types of Network Traffic

Application	Layer-3 Classification			Layer-2 Classification
	IP Precedence (IPP)	Per-Hop Behavior (PHB)	Differentiated Services Code Point (DSCP)	Class of Service (CoS)
Routing	6	CS6	48	6
Voice Real-Time Transport Protocol (RTP)	5	EF	46	5
Videoconferencing	4	AF41	34	4
Streaming video	4	CS4	32	4
Call signaling ¹	3	CS3 (currently) AF31 (previously)	24 (currently) 26 (previously)	3

Table 3-3 Traffic Classification Guidelines for Various Types of Network Traffic (continued)

Application	Layer-3 Classification			Layer-2 Classification
	IP Precedence (IPP)	Per-Hop Behavior (PHB)	Differentiated Services Code Point (DSCP)	Class of Service (CoS)
Transactional data	2	AF21	18	2
Network management	2	CS2	16	2
Scavenger	1	CS1	8	1
Best effort	0	0	0	0

1. The recommended DSCP/PHB marking for call control signaling traffic has been changed from 26/AF31 to 24/CS3. A marking migration is planned within Cisco to reflect this change, however many products still mark signaling traffic as 26/AF31. Therefore, in the interim, Cisco recommends that both AF31 and CS3 be reserved for call signaling.

For more information about traffic classification, refer to the *Enterprise QoS Solution Reference Network Design (SRND)*, available at

<http://www.cisco.com/go/designzone>

Traffic Classification for Video Telephony

The main classes of interest for IP Video Telephony are:

- Voice
Voice is classified as CoS 5 (IP Precedence 5, PHB EF, or DSCP 46).
- Videoconferencing
Videoconferencing is classified as CoS 4 (IP Precedence 4, PHB AF41, or DSCP 34).
- Call signaling
Call signaling for voice and videoconferencing is now classified as CoS 3 (IP Precedence 3, PHB CS3, or DSCP 24) but was previously classified as PHB AF31 or DSCP 26.

Cisco highly recommends these classifications as *best practices* in a Cisco Unified Communications network.

The voice component of a call can be classified in one of two ways, depending on the type of call in progress. A voice-only (or normal) telephone call would have the media classified as CoS 5 (IP Precedence 5 or PHB EF), while the audio channel of a video conference would have the media classified as CoS 4 (IP Precedence 4 or PHB AF41). All the Cisco IP Video Telephony products adhere to the Cisco Corporate QoS Baseline standard, which requires that the audio and video channels of a video call both be marked as CoS 4 (IP Precedence 4 or PHB AF41). The reasons for this recommendation include, but are not limited to, the following:

- To preserve lip-sync between the audio and video channels
- To provide separate classes for audio-only calls and video calls

The signaling class is applicable to all voice signaling protocols (such as SCCP, MGCP, and so on) as well as video signaling protocols (such as SCCP, H.225, RAS, CAST, and so on). These protocols are discussed in more detail in the section on [Software-Based Endpoints](#), page 21-7.

Given the recommended classes, the first step is decide where the packets will be classified (that is, which device will be the first to mark the traffic with its QoS classification). There are essentially two places to mark or classify traffic:

- On the originating endpoint — the classification is then trusted by the upstream switches and routers

- On the switches and/or routers — because the endpoint is either not capable of classifying its own packets or is not trustworthy to classify them correctly

Interface Queuing

After packets have been marked with the appropriate tag at Layer 2 (CoS) and Layer 3 (DSCP or PHB), it is important to configure the network to schedule or queue traffic based on this classification, so as to provide each class of traffic with the service it needs from the network. By enabling QoS on campus switches, you can configure all voice traffic to use separate queues, thus virtually eliminating the possibility of dropped voice packets when an interface buffer fills instantaneously.

Although network management tools may show that the campus network is not congested, QoS tools are still required to guarantee voice quality. Network management tools show only the average congestion over a sample time span. While useful, this average does not show the congestion peaks on a campus interface.

Transmit interface buffers within a campus tend to congest in small, finite intervals as a result of the bursty nature of network traffic. When this congestion occurs, any packets destined for that transmit interface are dropped. The only way to prevent dropped voice traffic is to configure multiple queues on campus switches. For this reason, Cisco recommends always using a switch that has at least two output queues on each port and the ability to send packets to these queues based on QoS Layer 2 and/or Layer 3 classification. Cisco Catalyst 6000, 4000, 3750, 35XX, and 2950 switches all support two or more output queues per port.

Bandwidth Provisioning

In the campus LAN, bandwidth provisioning recommendations can be summarized by the motto, *Over provision and under subscribe*. This motto implies careful planning of the LAN infrastructure so that the available bandwidth is always considerably higher than the load and there is no steady-state congestion over the LAN links.

The addition of voice traffic onto a converged network does not represent a significant increase in overall network traffic load; the bandwidth provisioning is still driven by the demands of the data traffic requirements. The design goal is to avoid extensive data traffic congestion on any link that will be traversed by telephony signaling or media flows. Contrasting the bandwidth requirements of a single G.711 voice call (approximately 86 kbps) to the raw bandwidth of a FastEthernet link (100 Mbps) indicates that voice is not a source of traffic that causes network congestion in the LAN, but rather it is a traffic flow to be protected from LAN network congestion.

Impairments to IP Communications if QoS is Not Employed

If QoS is not deployed, packet drops and excessive delay and jitter can occur, leading to impairments of the telephony services. When media packets are subjected to drops, delay, and jitter, the user-perceivable effects include clicking sound, harsh-sounding voice, extended periods of silence, and echo.

When signaling packets are subjected to the same conditions, user-perceivable impairments include unresponsiveness to user input (such as delay to dial tone), continued ringing upon answer, and double dialing of digits due to the user's belief that the first attempt was not effective (thus requiring hang-up and redial). More extreme cases can include endpoint re-initialization, call termination, and the spurious activation of SRST functionality at branch offices (leading to interruption of gateway calls).

These effects apply to all deployment models. However, single-site (campus) deployments tend to be less likely to experience the conditions caused by sustained link interruptions because the larger quantity of bandwidth typically deployed in LAN environments (minimum links of 100 Mbps) allows for some residual bandwidth to be available for the IP Communications system.

In any WAN-based deployment model, traffic congestion is more likely to produce sustained and/or more frequent link interruptions because the available bandwidth is much less than in a LAN (typically less than 2 Mbps), so the link is more easily saturated. The effects of link interruptions impact the users, whether or not the voice media traverses the packet network.

WAN Infrastructure

Proper WAN infrastructure design is also extremely important for proper IP telephony operation on a converged network. Proper infrastructure design requires following basic configuration and design best practices for deploying a WAN that is as highly available as possible and that provides guaranteed throughput. Furthermore, proper WAN infrastructure design requires deploying end-to-end QoS on all WAN links. The following sections discuss these requirements:

- [WAN Design and Configuration, page 3-25](#)
- [WAN Quality of Service \(QoS\), page 3-28](#)
- [Resource Reservation Protocol \(RSVP\), page 3-35](#)
- [Bandwidth Provisioning, page 3-44](#)

WAN Design and Configuration

Properly designing a WAN requires building fault-tolerant network links and planning for the possibility that these links might become unavailable. By carefully choosing WAN topologies, provisioning the required bandwidth, and approaching the WAN infrastructure as another layer in the network topology, you can build a fault-tolerant and redundant network. The following sections examine the required infrastructure layers and network services:

- [Deployment Considerations, page 3-25](#)
- [Guaranteed Bandwidth, page 3-27](#)
- [Best-Effort Bandwidth, page 3-27](#)

Deployment Considerations

WAN deployments for voice networks may be hub-and-spoke or an arbitrary topology. A hub-and-spoke topology consists of a central hub site and multiple remote spoke sites connected into the central hub site. In this scenario, each remote or spoke site is one WAN-link hop away from the central or hub site and two WAN-link hops away from all other spoke sites. An arbitrary topology may contain multiple WAN links and any number of hops between the sites. In this scenario there may be many different paths to the same site or there may be different links used for communication with some sites compared to other sites. The simplest example is three sites, each with a WAN link to the other two sites, forming a triangle. In that case there are two potential paths between each site to each other site.

Topology-unaware call admission control requires the WAN to be hub-and-spoke, or a spoke-less hub in the case of MPLS VPN. This topology ensures that call admission control, provided by Cisco Unified CallManager's locations or a gatekeeper, works properly in keeping track of the bandwidth available between any two sites in the WAN. In addition, multiple hub-and-spoke deployments can be interconnected via WAN links.

Topology-aware call admission control may be used with either hub-and-spoke or an arbitrary WAN topology. This form of call admission control requires parts of the WAN infrastructure to support Resource Reservation Protocol (RSVP). For details, see [Resource Reservation Protocol \(RSVP\)](#), [page 3-35](#), and [Call Admission Control](#), [page 9-1](#).

For more information about centralized and distributed multisite deployment models as well as Multiprotocol Label Switching (MPLS) implications for these deployment models, see the chapter on [IP Telephony Deployment Models](#), [page 2-1](#).

WAN links should, when possible, be made redundant to provide higher levels of fault tolerance. Redundant WAN links provided by different service providers or located in different physical ingress/egress points within the network can ensure backup bandwidth and connectivity in the event that a single link fails. In non-failure scenarios, these redundant links may be used to provide additional bandwidth and offer load balancing of traffic on a per-flow basis over multiple paths and equipment within the WAN. Topology-unaware call admission control normally requires redundant paths to be over-provisioned and under-subscribed to allow for failures that reduce the available bandwidth between sites without the call admission control mechanism being aware of those failures or the reduction in bandwidth. Topology-aware call admission control is able to adjust dynamically to many of the topology changes and allows for efficient use of the total available bandwidth.

Voice and data should remain converged at the WAN, just as they are converged at the LAN. QoS provisioning and queuing mechanisms are typically available in a WAN environment to ensure that voice and data can interoperate on the same WAN links. Attempts to separate and forward voice and data over different links can be problematic in many instances because the failure of one link typically forces all traffic over a single link, thus diminishing throughput for each type of traffic and in most cases reducing the quality of voice. Furthermore, maintaining separate network links or devices makes troubleshooting and management difficult at best.

Because of the potential for WAN links to fail or to become oversubscribed, Cisco recommends deploying non-centralized resources as appropriate at sites on the other side of the WAN. Specifically, media resources, DHCP servers, voice gateways, and call processing applications such as Survivable Remote Site Telephony (SRST) and Cisco Unified CallManager Express (CME) should be deployed at non-central sites when and if appropriate, depending on the site size and how critical these functions are to that site. Keep in mind that de-centralizing voice applications and devices can increase the complexity of network deployments, the complexity of managing these resources throughout the enterprise, and the overall cost of a the network solution; however, these factors can be mitigated by the fact that the resources will be available during a WAN link failure.

When deploying voice in a WAN environment, Cisco recommends that you use the lower-bandwidth G.729 codec for any voice calls that will traverse WAN links because this practice will provide bandwidth savings on these lower-speed links. Furthermore, media resources such as MoH should be configured to use multicast transport mechanism when possible because this practice will provide additional bandwidth savings.

Finally, recommendation G.114 of the International Telecommunication Union (ITU) states that the one-way delay in a voice network should be less than or equal to 150 milliseconds. It is important to keep this in mind when implementing low-speed WAN links within a network. Topologies, technologies, and physical distance should be considered for WAN links so that one-way delay is kept at or below this 150-millisecond recommendation. For deployments that use clustering over the WAN, the one-way delay for signaling traffic between clusters should not exceed 20 milliseconds (see [Clustering Over the IP WAN](#), [page 2-17](#)).

Guaranteed Bandwidth

Because voice is typically deemed a critical network application, it is imperative that bearer and signaling voice traffic always reaches its destination. For this reason, it is important to choose a WAN topology and link type that can provide guaranteed dedicated bandwidth. The following WAN link technologies can provide guaranteed dedicated bandwidth:

- Leased Lines
- Frame Relay
- Asynchronous Transfer Mode (ATM)
- ATM/Frame-Relay Service Interworking
- Multiprotocol Label Switching (MPLS)
- Cisco Voice and Video Enabled IP Security VPN (IPSec V3PN)

These link technologies, when deployed in a dedicated fashion or when deployed in a private network, can provide guaranteed traffic throughput. All of these WAN link technologies can be provisioned at specific speeds or bandwidth sizes. In addition, these link technologies have built-in mechanisms that help guarantee throughput of network traffic even at low link speeds. Features such as traffic shaping, fragmentation and packet interleaving, and committed information rates (CIR) can help ensure that packets are not dropped in the WAN, that all packets are given access at regular intervals to the WAN link, and that enough bandwidth is available for all network traffic attempting to traverse these links.

Best-Effort Bandwidth

There are some WAN topologies that are unable to provide guaranteed dedicated bandwidth to ensure that network traffic will reach its destination, even when that traffic is critical. These topologies are extremely problematic for voice traffic, not only because they provide no mechanisms to provision guaranteed network throughput, but also because they provide no traffic shaping, packet fragmentation and interleaving, queuing mechanisms, or end-to-end QoS to ensure that critical traffic such as voice will be given preferential treatment.

The following WAN network topologies and link types are examples of this kind of best-effort bandwidth technology:

- The Internet
- DSL
- Cable
- Satellite
- Wireless

In most cases, none of these link types can provide the guaranteed network connectivity and bandwidth required for critical voice and voice applications. However, these technologies might be suitable for personal or telecommuter-type network deployments. At times, these topologies can provide highly available network connectivity and adequate network throughput; but at other times, these topologies can become unavailable for extended periods of time, can be throttled to speeds that render network throughput unacceptable for real-time applications such as voice, or can cause extensive packet losses and require repeated retransmissions. In other words, these links and topologies are unable to provide guaranteed bandwidth, and when traffic is sent on these links, it is sent best-effort with no guarantee that it will reach its destination. For this reason, Cisco recommends that you do *not* use best-effort WAN topologies for voice-enabled networks that require enterprise-class voice services and quality.

**Note**

There are some new QoS mechanisms for DSL and cable technologies that can provide guaranteed bandwidth; however, these mechanisms are not typically deployed by service providers, and these services are still significantly oversubscribed.

WAN Quality of Service (QoS)

Before placing voice and video traffic on a network, it is important to ensure that there is adequate bandwidth for all required applications. Once this bandwidth has been provisioned, voice priority queuing must be performed on all interfaces. This queuing is required to reduce jitter and possible packet loss if a burst of traffic oversubscribes a buffer. This queuing requirement is similar to the one for the LAN infrastructure.

Next, the WAN typically requires additional mechanisms such as traffic shaping to ensure that WAN links are not sent more traffic than they can handle, which could cause dropped packets.

Finally, link efficiency techniques can be applied to WAN paths. For example, link fragmentation and interleaving (LFI) can be used to prevent small voice packets from being queued behind large data packets, which could lead to unacceptable delays on low-speed links.

The goal of these QoS mechanisms is to ensure reliable, high-quality voice by reducing delay, packet loss, and jitter for the voice traffic. [Table 3-4](#) lists the QoS features and tools required for the WAN infrastructure to achieve this goal.

Table 3-4 QoS Features and Tools Required to Support IP Telephony for each WAN Technology and Link Speed

WAN Technology	Link Speed: 56 kbps to 768 kbps	Link Speed: Greater than 768 kbps
Leased Lines	• Multilink Point-to-Point Protocol (MLP) • MLP Link Fragmentation and Interleaving (LFI) • Low Latency Queuing (LLQ) • Optional: Compressed Real-Time Transport Protocol (cRTP)	• LLQ
Frame Relay (FR)	• Traffic Shaping • LFI (FRF.12) • LLQ • Optional: cRTP • Optional: Voice-Adaptive Traffic Shaping (VATS) • Optional: Voice-Adaptive Fragmentation (VAF)	• Traffic Shaping • LLQ • Optional: VATS
Asynchronous Transfer Mode (ATM)	• TX-ring buffer changes • MLP over ATM • MLP LFI • LLQ • Optional: cRTP (requires MLP)	• TX-ring buffer changes • LLQ

Table 3-4 QoS Features and Tools Required to Support IP Telephony for each WAN Technology and Link Speed

WAN Technology	Link Speed: 56 kbps to 768 kbps	Link Speed: Greater than 768 kbps
Frame Relay and ATM Service Inter-Working (SIW)	• TX-ring buffer changes • MLP over ATM and FR • MLP LFI • LLQ • Optional: cRTP (requires MLP)	• TX-ring buffer changes • MLP over ATM and FR • LLQ
Multiprotocol Label Switching (MPLS)	• Same as above, according to the interface technology • Class-based marking is generally required to remark flows according to service provider specifications	• Same as above, according to the interface technology • Class-based marking is generally required to remark flows according to service provider specifications


The following sections highlight some of the most important features and techniques to consider when designing a WAN to support both voice and data traffic:

- [Traffic Prioritization, page 3-29](#)
- [Link Efficiency Techniques, page 3-31](#)
- [Traffic Shaping, page 3-32](#)

Traffic Prioritization

In choosing from among the many available prioritization schemes, the major factors to consider include the type of traffic involved and the type of media on the WAN. For multi-service traffic over an IP WAN, Cisco recommends low-latency queuing (LLQ) for all links. This method supports up to 64 traffic classes, with the ability to specify, for example, priority queuing behavior for voice and interactive video, minimum bandwidth class-based weighted fair queuing for voice control traffic, additional minimum bandwidth weighted fair queues for mission critical data, and a default best-effort queue for all other traffic types.

Figure 3-6 shows an example prioritization scheme.

Figure 3-6 Optimized Queuing for VoIP over the WAN

Cisco recommends the following prioritization criteria for LLQ:

- The criterion for *voice* to be placed into a priority queue is the differentiated services code point (DSCP) value of 46, or a per-hop behavior (PHB) value of EF.
- The criterion for *video conferencing* traffic to be placed into a priority queue is a DSCP value of 34, or a PHB value of AF41. However, due to the larger packet sizes of video traffic, these packets should be placed in the priority queue only on WAN links that are faster than 768 Kbps. Link speeds below this value require packet fragmentation, but packets placed in the priority queue are not fragmented, thus smaller voice packets could be queued behind larger video packets. For links speeds of 768 Kbps or lower, video conferencing traffic should be placed in a separate class-based weighted fair queue (CBWFQ).



Note One-way video traffic, such as the traffic generated by streaming video applications for services such as video-on-demand or live video feeds, should always use a CBWFQ scheme because that type of traffic has a much higher delay tolerance than two-way video conferencing traffic

- As the WAN links become congested, it is possible to starve the *voice control* signaling protocols, thereby eliminating the ability of the IP phones to complete calls across the IP WAN. Therefore, voice control protocols, such as H.323, MGCP, and Skinny Client Control Protocol (SCCP), require their own class-based weighted fair queue. The entrance criterion for this queue is a DSCP value of 24 or a PHB value of CS3.



Note Cisco has begun to change the marking of voice control protocols from DSCP 26 (PHB AF31) to DSCP 24 (PHB CS3). However many products still mark signaling traffic as DSCP 26 (PHB AF31); therefore, in the interim, Cisco recommends that you reserve both AF31 and CS3 for call signaling.

- In some cases, certain data traffic might require better than best-effort treatment. This traffic is referred to as *mission-critical data*, and it is placed into one or more queues that have the required amount of bandwidth. The queuing scheme within this class is first-in-first-out (FIFO) with a minimum allocated bandwidth. Traffic in this class that exceeds the configured bandwidth limit is placed in the default queue. The entrance criterion for this queue could be a Transmission Control Protocol (TCP) port number, a Layer 3 address, or a DSCP/PHB value.
- All remaining traffic can be placed in a default queue for best-effort treatment. If you specify the keyword **fair**, the queuing algorithm will be weighted fair queuing (WFQ).

Link Efficiency Techniques

The following link efficiency techniques improve the quality and efficiency of low-speed WAN links.

Compressed Real-Time Transport Protocol (cRTP)

You can increase link efficiency by using Compressed Real-Time Transport Protocol (cRTP). This protocol compresses a 40-byte IP, User Datagram Protocol (UDP), and RTP header into approximately two to four bytes. cRTP operates on a per-hop basis. Use cRTP on a particular link only if that link meets *all* of the following conditions:

- Voice traffic represents more than 33% of the load on the specific link.
- The link uses a low bit-rate codec (such as G.729).
- No other real-time application (such as video conferencing) is using the same link.

If the link fails to meet any one of the preceding conditions, then cRTP is not effective and you should not use it on that link. Another important parameter to consider before using cRTP is router CPU utilization, which is adversely affected by compression and decompression operations.

cRTP on ATM and Frame Relay Service Inter-Working (SIW) links requires the use of Multilink Point-to-Point Protocol (MLP).

Note that cRTP compression occurs as the final step before a packet leaves the egress interface; that is, after LLQ class-based queueing has occurred. Beginning in Cisco IOS Release 12.(2)2T and later, cRTP provides a feedback mechanism to the LLQ class-based queueing mechanism that allows the bandwidth in the *voice* class to be configured based on the compressed packet value. With Cisco IOS releases prior to 12.(2)2T, this mechanism is not in place, so the LLQ is unaware of the compressed bandwidth and, therefore, the *voice* class bandwidth has to be provisioned as if no compression is taking place. [Table 3-5](#) shows an example of the difference in *voice* class bandwidth configuration given a 512-kbps link with G.729 codec and a requirement for 10 calls.

Note that [Table 3-5](#) assumes 24 kbps for non-cRTP G.729 calls and 10 kbps for cRTP G.729 calls. These bandwidth numbers are based on voice payload and IP/UDP/RTP headers only. They do not take into consideration Layer 2 header bandwidth. However, actual bandwidth provisioning should also include Layer 2 header bandwidth based on the type WAN link used.

Table 3-5 *LLQ Voice Class Bandwidth Requirements for 10 Calls with 512 kbps Link Bandwidth and G.729 Codec*

Cisco IOS Release	With cRTP Not Configured	With cRTP Configured
Prior to 12.2(2)T	240 kbps	240 kbps ¹
12.2(2)T or later	240 kbps	100 kbps

1. 140 kbps of unnecessary bandwidth must be configured in the LLQ *voice* class.

It should also be noted that, beginning in Cisco IOS Release 12.2(13)T, cRTP can be configured as part of the voice class with the Class-Based cRTP feature. This option allows cRTP to be specified within a class, attached to an interface via a service policy. This new feature provides compression statistics and bandwidth status via the **show policy interface** command, which can be very helpful in determining the offered rate on an interface service policy class given the fact that cRTP is compressing the IP/RTP headers.


For additional recommendations about using cRTP with a Voice and Video Enabled IPsec VPN (V3PN), refer to the V3PN documentation available at

<http://www.cisco.com/go/designzone>

Link Fragmentation and Interleaving (LFI)

For low-speed links (less than 768 kbps), use of link fragmentation and interleaving (LFI) mechanisms is required for acceptable voice quality. This technique limits jitter by preventing voice traffic from being delayed behind large data frames, as illustrated in [Figure 3-7](#). The two techniques that exist for this purpose are Multilink Point-to-Point Protocol (MLP) LFI (for Leased Lines, ATM, and SIW) and FRF.12 for Frame Relay.

Figure 3-7 Link Fragmentation and Interleaving (LFI)



Voice-Adaptive Fragmentation (VAF)

In addition to the LFI mechanisms mentioned above, voice-adaptive fragmentation (VAF) is another LFI mechanism for Frame Relay links. VAF uses FRF.12 Frame Relay LFI; however, once configured, fragmentation occurs only when traffic is present in the LLQ priority queue or when H.323 signaling packets are detected on the interface. This method ensures that, when voice traffic is being sent on the WAN interface, large packets are fragmented and interleaved. However, when voice traffic is not present on the WAN link, traffic is forwarded across the link unfragmented, thus reducing the overhead required for fragmentation.

VAF is typically used in combination with voice-adaptive traffic shaping (see [Voice-Adaptive Traffic Shaping \(VATS\)](#), [page 3-34](#)). VAF is an optional LFI tool, and you should exercise care when enabling it because there is a slight delay between the time when voice activity is detected and the time when the LFI mechanism engages. In addition, a configurable deactivation timer (default of 30 seconds) must expire after the last voice packet is detected and before VAF is deactivated, so during that time LFI will occur unnecessarily. VAF is available in Cisco IOS Release 12.2(15)T and later.

Traffic Shaping

Traffic shaping is required for multiple-access, non-broadcast media such as ATM and Frame Relay, where the physical access speed varies between two endpoints and several branch sites are typically aggregated to a single router interface at the central site.

[Figure 3-8](#) illustrates the main reasons why traffic shaping is needed when transporting voice and data on the same IP WAN.


Figure 3-8 Traffic Shaping with Frame Relay and ATM

Figure 3-8 shows three different scenarios:

1. Line speed mismatch

While the central-site interface is typically a high-speed one (such as T1 or higher), smaller remote branch interfaces may have significantly lower line speeds, such as 64 kbps. If data is sent at full rate from the central site to a slow-speed remote site, the interface at the remote site might become congested and degrade voice performance.

2. Oversubscription of the link between the central site and the remote sites

It is common practice in Frame Relay or ATM networks to oversubscribe bandwidth when aggregating many remote sites to a single central site. For example, there may be multiple remote sites that connect to the WAN with a T1 interface, yet the central site has only a single T1 interface. While this configuration allows the deployment to benefit from statistical multiplexing, the router interface at the central site can become congested during traffic bursts, thus degrading voice quality.


3. Bursting above Committed Information Rate (CIR)

Another common configuration is to allow traffic bursts above the CIR, which represents the rate that the service provider has guaranteed to transport across its network with no loss and low delay. For example, a remote site with a T1 interface might have a CIR of only 64 kbps. When more than

64 kbps worth of traffic is sent across the WAN, the provider marks the additional traffic as "discard eligible." If congestion occurs in the provider network, this traffic will be dropped with no regard to traffic classification, possibly having a negative affect on voice quality.

Traffic shaping provides a solution to these issues by limiting the traffic sent out an interface to a rate lower than the line rate, thus ensuring that no congestion occurs on either end of the WAN. Figure 3-9 illustrates this mechanism with a generic example, where R is the rate with traffic shaping applied.

Figure 3-9 Traffic Shaping Mechanism



Voice-Adaptive Traffic Shaping (VATS)

VATS is an optional dynamic mechanism that shapes traffic on Frame Relay permanent virtual circuits (PVCs) at different rates based on whether voice is being sent across the WAN. The presence of traffic in the LLQ voice priority queue or the detection of H.323 signaling on the link causes VATS to engage. Typically, Frame Relay shapes traffic to the guaranteed bandwidth or CIR of the PVC at all times. However, because these PVCs are typically allowed to burst above the CIR (up to line speed), traffic shaping keeps traffic from using the additional bandwidth that might be present in the WAN. With VATS enabled on Frame Relay PVCs, WAN interfaces are able to send at CIR when voice traffic is present on the link. However, when voice is not present, non-voice traffic is able to burst up to line speed and take advantage of the additional bandwidth that might be present in the WAN.

When VATS is used in combination with voice-adaptive fragmentation (VAF) (see [Link Fragmentation and Interleaving \(LFI\)](#), page 3-32), all non-voice traffic is fragmented and all traffic is shaped to the CIR of the WAN link when voice activity is detected on the interface.

As with VAF, exercise care when enabling VATS because activation can have an adverse effect on non-voice traffic. When voice is present on the link, data applications will experience decreased throughput because they are throttled back to well below CIR. This behavior will likely result in packet drops and delays for non-voice traffic. Furthermore, after voice traffic is no longer detected, the deactivation timer (default of 30 seconds) must expire before traffic can burst back to line speed. It is important, when using VATS, to set end-user expectations and make them aware that data applications will experience slowdowns on a regular basis due to the presence of voice calls across the WAN. VATS is available in Cisco IOS Release 12.2(15)T and later.

For more information on the Voice-Adaptive Traffic Shaping and Fragmentation features and how to configure them, refer to the documentation at

http://www.cisco.com/en/US/products/sw/iosswrel/ps1839/products_feature_guide09186a00801541c6.html

Resource Reservation Protocol (RSVP)

The Resource Reservation Protocol (RSVP) is the first significant industry-standard protocol for dynamically setting up end-to-end QoS across a heterogeneous network. RSVP, which runs over IP, was first introduced by the IETF in RFC 2205, and it enables an application to reserve network bandwidth dynamically. Using RSVP, applications can request a certain level of QoS for a data flow across a network. Because of its distributed and dynamic nature, RSVP is capable of reserving bandwidth across any network topology, therefore it can be used to provide topology-aware call admission control for voice and video calls.

This section focuses on the RSVP protocol principles and its interactions with the WAN infrastructure, specifically the QoS aspects, while the motivation and the mechanisms for call admission control based on RSVP are described in the chapter on [Call Admission Control, page 9-1](#).

This section covers the following specific topics:


- [RSVP Principles, page 3-35](#)
- [RSVP and QoS in WAN Routers, page 3-38](#)
- [RSVP Application ID, page 3-42](#)
- [RSVP Design Best Practices, page 3-44](#)

RSVP Principles

RSVP performs resource reservation for a given data flow across a network by defining signaling messages that are exchanged between the source and destination devices for the data flow and that are processed by intermediate routers along the path. The RSVP signaling messages are IP packets with the protocol number in the IP header set to 46, and they are routed through the network according to the existing routing protocols.

Not all routers on the path are required to support RSVP because the protocol is designed to operate transparently across RSVP-unaware nodes. On each RSVP-enabled router, the RSVP process intercepts the signaling messages and interacts with the QoS manager for the router interfaces involved in the data flow in order to "reserve" bandwidth resources. When the available resources are not sufficient for the data flow anywhere along the path, the routers signal the failure back to the application that originated the reservation request.

The principles of RSVP signaling can be explained by using the example shown in [Figure 3-10](#). In this diagram, an application wishes to reserve network resources for a data stream flowing from Device 1, whose IP address is 10.10.10.10, to Device 2, whose IP address is 10.60.60.60.

Figure 3-10 Example of RSVP Path and Resv Message Flow

1471853

The following steps describe the RSVP signaling process for the example in [Figure 3-10](#):

1. The application residing on Device 1 originates an RSVP message called Path, which is sent to the same destination IP address as the data flow for which a reservation is requested (that is, 10.60.60.60) and is sent with the "router alert" option turned on in the IP header. The Path message contains, among other things, the following objects:
 - The "session" object, consisting of destination IP address, protocol number, and UDP/TCP port, which is used to identify the data flow in RSVP-enabled routers.
 - The "sender Tspec" (traffic specification) object, which characterizes the data flow for which a reservation will be requested. This typically translates into a token bucket model that specifies a data rate and a burst size (or bucket depth).
 - The "P Hop" (or previous hop) object, which contains the IP address of the router interface that last processed the Path message. In this example, the P Hop is initially set to 10.10.10.10 by Device 1.
2. By means of the "router alert" option, the Path message is intercepted by the CPU of the RSVP-aware router identified as 10.20.20.20 in [Figure 3-10](#), which sends it to the RSVP process. RSVP creates a path state for this data flow, storing the values of the session, sender Tspec, and P Hop objects contained in the Path message. Then it forwards the message downstream, after having replaced the P Hop value with the IP address of its outgoing interface (10.20.20.20 in this example).

3. Similarly, the Path message is intercepted by the CPU of the following RSVP-aware router, identified as 10.30.30.30 in [Figure 3-10](#). After creating the path state and changing the P Hop value to 10.30.30.30, this router also forwards the message downstream.
4. The Path message now arrives at the RSVP-unaware router identified as 10.40.40.40 in [Figure 3-10](#). Because RSVP is not enabled on this router, it just routes this message according to the existing routing protocols like any other IP packet, without any additional processing and without changing the content of any of the message objects.
5. Therefore, the Path message gets to the RSVP-aware router identified as 10.50.50.50, which processes the message, creates the corresponding path state, and forwards the message downstream. Notice that the P Hop recorded by this router still contains the IP address of the last RSVP-aware router along the network path, or 10.30.30.30 in this example.
6. The RSVP Receiver at Device 2 receives the Path message with a P Hop value of 10.50.50.50, and it can now initiate the actual reservation by originating a message called Resv. For this reason, RSVP is known as a receiver-initiated protocol. The Resv message carries the reservation request hop-by-hop from the receiver to the sender, along the reverse paths of the data flow for the session. At each hop, the IP destination address of the Resv message is the IP address of the previous-hop node, obtained from the path state. Hence, in this case Device 2 sends the Resv message with a destination IP address of 10.50.50.50. The Resv message contains, among other things, the following objects:
 - The "session" object, which is used to identify the data flow.
 - The "N Hop" (or next hop) object, which contains the IP address of the node that generated the message. In this example, the N Hop is initially set to 10.60.60.60 by Device 2.
7. When RSVP-aware router 10.50.50.50 receives the Resv message for this data flow, it matches it against the path state information using the received session object, and it verifies if the reservation request can be accepted based on the following criteria:
 - Policy control — Is this user and/or application allowed to make this reservation request?
 - Admission control — Are there enough bandwidth resources available on the relevant outgoing interface to accommodate this reservation request?
8. In this case, we assume that both policy and admission control are successful on 10.50.50.50, which means that the bandwidth provided by the Tspec in the path state for this session is reserved on the outgoing interface (in the same direction as the data flow, that is from Device 1 to Device 2), and a corresponding "reservation state" is created. Now router 10.50.50.50 can send a Resv message upstream by sending it as a unicast IP packet to the destination IP address stored in the P Hop for this session, which was 10.30.30.30. The N Hop object is also updated with the value of 10.50.50.50.
9. The Resv message now transits through the RSVP-unaware router identified as 10.40.40.40, which will route it toward its destination of 10.30.30.30 like any other IP packet. This mechanism allows RSVP signaling to work across a heterogeneous network where some nodes are not RSVP-enabled.
10. The RSVP-aware router identified as 10.30.30.30 receives the Resv message and processes it according to the mechanisms described in steps 7 and 8. Assuming policy and admission control are successful also at this hop, the bandwidth is reserved on the outgoing interface and a Resv message is sent to the previous hop, or 10.20.20.20 in this example.
11. After a similar process within the router identified as 10.20.20.20, the Resv finally reaches the RSVP sender, Device 1. This indicates to the requesting application that an end-to-end reservation has been established and that bandwidth has been set aside for this data flow in all RSVP-enabled routers across the network.

This example shows how the two main RSVP signaling messages, Path and Resv, travel across the network to establish reservations. Several other messages are defined in the RSVP standard to address error situations, reservation failures, and release of resources. In particular, the ResvErr message is used

to signal failure to reserve the requested resources due to either policy control or admission control somewhere along the network. If, for example, admission control had failed at node 10.50.50.50 in [Figure 3-10](#), this node would have sent a ResvErr message back to Device 2, specifying the cause of the failure, and the application would have been notified.

Another important aspect of the RSVP protocol is that it adopts a soft-state approach, which means that for each session both the path state and the reservation state along the network need to be refreshed periodically by the application by sending identical Path and Resv messages. If a router does not receive refresh messages for a given session for a certain period of time, it deletes the corresponding state and releases the resources reserved. This allows RSVP to react dynamically to network topology changes or routing changes due to link failures. The reservations simply start flowing along the new routes based on the routing protocol decisions, and the reservations along the old routes time-out and are eventually deleted.

**Note**

This section focuses on providing an overview of RSVP principles and mechanisms. For more information on protocol behavior and extensions, complete message formats, and interactions with other protocols, refer to the numerous RFC documents related to RSVP, available at <http://www.ietf.org>.

RSVP and QoS in WAN Routers

RSVP has been supported in Cisco routers for many years, however most configurations recommended in this document are based on the RSVP Scalability Enhancements feature, which was first introduced in Cisco IOS Release 12.2(2)T.

By issuing the following Cisco IOS command in interface configuration mode on each Cisco IOS router interface, you can enable RSVP and define the maximum amount of bandwidth that it can control:

```
ip rsvp bandwidth [interface-kbps] [single-flow-kbps]
```

The *interface-kbps* parameter specifies the upper limit of bandwidth that RSVP can reserve on the given interface, while the *single-flow-kbps* parameter provides an upper bandwidth limit for each individual reservation (so that flows with higher bandwidth requests will be rejected even if there is bandwidth available on the interface).


**Note**

When RSVP is enabled on a router interface, all other interfaces in the router will drop RSVP messages unless they are also enabled for RSVP. To avoid dropping RSVP messages, enable RSVP on all interfaces through which you expect RSVP signaling to transit. If call admission control is not desired on an interface, set the bandwidth value to 75% of the interface bandwidth.

Within Cisco IOS, RSVP can be configured to operate according to two different models: the Integrated Services (IntServ) model, described in RFC 2210, or the Integrated Services/Differentiated Services (IntServ/DiffServ) model, described in RFC 2998. Both RFC documents are available on the IETF website at

<http://www.ietf.org>


[Figure 3-11](#) shows the difference between these two approaches from the perspective of a Cisco IOS router.

Figure 3-11 The Two RSVP Operation Models: IntServ and IntServ/DiffServ

The IntServ Model

As shown on the left side of [Figure 3-11](#), RSVP in the IntServ model involves both the control plane and the data plane. In the control plane, RSVP admits or denies the reservation request. In the data plane, it classifies the data packets, polices them based on the traffic description contained in the RSVP messages, and queues them in the appropriate queue. The classification that RSVP performs is based on the 5-tuple consisting of the source IP address, source port, destination IP address, destination port, and protocol number. In this model, all data packets transiting through the router must be intercepted by RSVP so that RSVP can inspect the 5-tuple and look for a match among the established reservations. If a match is found, the packets are scheduled and policed by RSVP according to the reservation's traffic specification.

As shown in [Figure 3-12](#), when you combine the IntServ model with Low Latency Queuing (LLQ), the usable bandwidth is divided between RSVP and the predefined LLQ queues. RSVP controls the entrance criteria to the RSVP reserved bandwidth, while policy maps control the entrance criteria for the predefined queues.

Figure 3-12 Combining the IntServ Model with LLQ

To use the IntServ operation model on a Cisco IOS router, use the following commands in interface configuration mode:

```
ip rsvp resource-provider wfq [interface | pvc]
no ip rsvp data-packet classification
```

When these commands are active, RSVP admits or rejects new reservations, not only based on the upper bandwidth limit defined within the **ip rsvp bandwidth** command, but also based on the actual bandwidth resources available. For example, if there are LLQ classes with bandwidth statements, these amounts are deducted from the bandwidth pool that can be assigned to RSVP reservations. While LLQ classes statically allocate bandwidth at configuration time, RSVP does not allocate any amount until a reservation request is received. Therefore, it is important to ensure that an appropriate percentage of the available interface bandwidth is *not* allocated to LLQ classes, so that it can be used by RSVP as reservation requests are received.

Because the total maximum bandwidth that can be assigned to QoS mechanisms on a link is equal to 75% of the link speed, if you want to reserve 33% of the link bandwidth for RSVP-admitted flows, you have to make sure that the bandwidth assigned to LLQ classes does not exceed $(75 - 33) = 42\%$ of the link bandwidth.

Because RSVP is in control of assigning packets to the various queues within this model, it is possible to define a mechanism for RSVP to know whether or not to place flows in the Priority Queue (PQ) by using the following Cisco IOS command in interface configuration mode:

```
ip rsvp pq-profile [r [b [p-to-r]]]
```

RSVP uses the parameters *r*, *b*, and *p-to-r* to determine if the flow being signaled for is a voice flow that needs PQ treatment. These parameters represent the following values:

- *r* = the average traffic rate in bytes per second
- *b* = the maximum burst of a flow in bytes
- *p-to-r* = the ratio of peak rate to average rate, expressed as a percentage

If the traffic characteristics specified by the RSVP messages for a certain flow are less than or equal to the parameters in the command, then RSVP will direct the flow into the PQ. If no parameters are provided with the command, the following values, representing the largest of the commonly used voice codecs (G.711), are used as default:

- $r = 12288$ bytes per second
- $b = 592$ bytes
- $p\text{-to-}r = 110\%$


The IntServ/DiffServ Model

As shown on the right side of [Figure 3-11](#), RSVP in the IntServ/DiffServ model involves only the control plane performing admission control but does not involve the data plane. This means that the call admission control function is separate from the scheduling and policing functions, which can be performed by the Low Latency Queuing (LLQ) algorithm according to predefined class maps, policy maps, and service policies.

With the IntServ/DiffServ model, it is therefore possible to add RSVP call admission control to a network that is already using a Differentiated Services approach to QoS. RSVP admits or rejects calls based on a preconfigured bandwidth amount, but the actual scheduling is based on the pre-existing LLQ criteria such as the DSCP value of each packet.

The entire usable bandwidth (75% of the link speed) can be assigned to LLQ classes, as shown in [Figure 3-13](#), as it normally is today. The policy maps define the traffic that is admitted into each queue. RSVP is typically configured to admit flows up to the amount of bandwidth defined for priority traffic, but keep in mind that RSVP in this model does not adjust the scheduling, so any traffic admitted by RSVP in excess of the predefined priority queue may be dropped or remapped to other lower-priority queues.

If all applications that send priority traffic are RSVP-enabled, you may configure the RSVP bandwidth to match the size of the priority queue. If, on the other hand, there are non-RSVP applications that also need to send priority traffic (such as Cisco Unified CallManager static locations or a gatekeeper), as shown in [Figure 3-13](#), the priority queue is divided into priority traffic that is controlled by non-RSVP mechanisms and priority traffic that is controlled by RSVP. The combined non-RSVP and RSVP admission control mechanisms must not use more bandwidth than is allocated to ensure that the priority queue is never over-subscribed.

Figure 3-13 LLQ Bandwidth Allocation with RSVP

To use the IntServ/DiffServ operation model on a Cisco IOS router, use the following commands in interface configuration mode:

```
ip rsvp resource-provider none
ip rsvp data-packet classification none
```

When these commands are active, RSVP admits or rejects new reservations uniquely based on the upper bandwidth limits defined within the **ip rsvp bandwidth** command, independently from the actual bandwidth resources available on the interface. Once admitted, the RSVP flows are subject to the same scheduling rules as all other non-RSVP traffic (for example, LLQ class and policy maps). Therefore, it is important to ensure that the RSVP-enabled traffic is marked with the appropriate DSCP value and that the bandwidth of the corresponding PQ or CBWFQ queues is provisioned to accommodate both RSVP-enabled traffic and all other traffic.

In this operating model, the **ip rsvp pq-profile** command is inactive because RSVP does not control the scheduling function.

RSVP Application ID

An application identity (app-id) is an RSVP object that can be inserted into the policy element of an RSVP message. This object is described in RFC 2872. This policy object helps to identify the application and associate it with the RSVP reservation request, thus allowing routers along the path to make appropriate decisions based on the application information.

The need for an app-id arises because RSVP is used to support multiple applications such as voice and video.

Without using an app-id, there is only one bandwidth value that is configurable per interface in RSVP. RSVP will admit requests until this bandwidth limit is reached. It does not differentiate between the requests and is not aware of the type of application for which the bandwidth is requested. As a result of this, it is quite possible for RSVP to exhaust the allowed bandwidth by admitting requests representing

just one type of application, thus causing all subsequent requests to be rejected due to unavailable bandwidth. In this way, a few video calls could prevent all or most of the voice calls from being admitted. For example, if an organization allocates 1000 units to RSVP, RSVP might exhaust a majority of this amount by admitting two 384-kbps video calls, thus leaving very little bandwidth for voice calls.

The solution to this problem is to configure separate bandwidth limits for individual applications or classes of traffic. Limiting bandwidth per application requires that an RSVP local policy matching the application bandwidth limit be applied to the router interface and that each reservation request flag the application to which it belongs so that it may be admitted against the appropriate bandwidth limit.

The app-id is not a single piece of information but multiple variable-length strings. As is described in RFC 2872, the object may include the following attributes:

- An identifier of the application (APP). This attribute is required.
- Global unique identifier (GUID). Optional.
- Version number of the application (VER). This attribute is required.
- Sub-application identifier (SAPP). An arbitrary number of sub-application elements can be included. Optional.

For example:

- APP = AudioStream
- GUID = CiscoSystems
- VER = 5.0.1.0
- SAPP = (not specified)

How Cisco Unified CallManager Uses the Application ID

To support RSVP Application ID functionality, Cisco Unified CallManager has two cluster-wide service parameters that define the Application ID used to tag audio and video call reservations using RSVP:

- RSVP Audio Application ID (Default is "AudioStream")
- RSVP Video Application ID (Default is "VideoStream")

It is possible to change these service parameters, but Cisco recommend that you leave them at their default values unless you require the ability to differentiate one cluster's reservations from another using the same link.

How Voice Calls are Tagged

When a voice call is made between locations with an RSVP policy, the resulting reservations for the audio stream will be tagged with the RSVP Audio Application ID.

How Video Calls are Tagged

When a video call is made between locations with an RSVP policy, the resulting reservations for the audio stream will be tagged with the RSVP Audio Application ID and the reservations for the video stream will be tagged with the RSVP Video Application ID.

Application ID Call Admission Control Model

As mentioned in the chapter on [Call Admission Control](#), page 9-1, the call admission control model supported by Application ID differs from the one supported by "static" locations. Because the audio stream of a video call is marked with the RSVP Audio Application ID, it is possible to guarantee a

minimum number of voice calls and allow them to take over the entire available bandwidth. Video calls are allowed up to a certain maximum bandwidth but may be denied when the entire bandwidth is consumed by voice calls that were established first.


RSVP Design Best Practices

When deploying RSVP in the IP WAN in conjunction with Cisco Unified CallManager, observe the following design best practices:

- Cisco recommends that you use the IntServ/DiffServ model if either of the following statements is true:
 - The only traffic destined for the Priority Queue (PQ) in the IP WAN interfaces is RSVP-enabled traffic.
 - All the non-RSVP traffic destined for the PQ can be deterministically limited to a certain amount by an out-of-band call admission control mechanism (such as Cisco Unified CallManager locations or a Cisco IOS gatekeeper).
- If all the PQ traffic is RSVP-enabled, the value specified in the **ip rsvp bandwidth** command and the **priority** command should match once Layer 2 overhead is taken into account.
- If RSVP is enabled on one or more interfaces of a router, all interfaces through which you expect RSVP signaling to transit should also be enabled for RSVP to ensure that RSVP messages do not get dropped. If call admission control is not desired on an interface, set the bandwidth value to 75% of the interface bandwidth.
- If some PQ traffic is not RSVP-enabled, you must ensure that the sum of the values specified in the **ip rsvp bandwidth** command and in the out-of-band call admission control mechanism do not exceed the bandwidth value specified in the **priority** command.
- Enable RSVP Application ID support if you need to limit the maximum amount of bandwidth used by video calls. Application ID Support is introduced in Cisco IOS Release 12.4(6)T. For more information, see [RSVP Application ID, page 3-42](#).
- Enable RSVP at the edge of the network, including the router WAN interfaces on both sides of the WAN link.
- Enable RSVP at all possible WAN congestion points, including redundant links of different speeds.
- Ensure symmetric routing on load-balanced MPLS WAN links.
- Bundle Interfaces, including MLPPP, ATM-IMA, and FRF.16, should have the RSVP bandwidth configured to the size of one physical link.
- RSVP is currently not available on Tunnel Interfaces.
- RSVP is currently not available on the Catalyst Switching Platforms.

Bandwidth Provisioning

Properly provisioning the network bandwidth is a major component of designing a successful IP network. You can calculate the required bandwidth 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 it should not exceed approximately 75% of the total available bandwidth for the link. This 75% rule assumes that some bandwidth is required for overhead traffic, such as routing and Layer 2 keep-alives. [Figure 3-14](#) illustrates this bandwidth provisioning process.

Figure 3-14 Link Bandwidth Provisioning

In addition to using no more than 75% of the total available bandwidth for data, voice, and video, the total bandwidth configured for all LLQ priority queues should typically not exceed 33% of the total link bandwidth. Provisioning more than 33% of the available bandwidth for the priority queue can be problematic for a number of reasons. First, provisioning more than 33% of the bandwidth for voice can result in increased CPU usage. Because each voice call will send 50 packets per second (with 20 ms samples), provisioning for large numbers of calls in the priority queue can lead to high CPU levels due to high packet rates. In addition, if more than one type of traffic is provisioned in the priority queue (for example, voice and video), this configuration defeats the purpose of enabling QoS because the priority queue essentially becomes a first-in, first-out (FIFO) queue. A larger percentage of reserved priority bandwidth effectively dampens the QoS effects by making more of the link bandwidth FIFO. Finally, allocating more than 33% of the available bandwidth can effectively starve any data queues that are provisioned. Obviously, for very slow links (less than 192 kbps), the recommendation to provision no more than 33% of the link bandwidth for the priority queue(s) might be unrealistic because a single call could require more than 33% of the link bandwidth. In these situations, and in situations where specific business needs cannot be met while holding to this recommendation, it may be necessary to exceed the 33% rule.

From a traffic standpoint, an IP telephony call consists of two parts:

- The voice and video bearer streams, which consists of Real-Time Transport Protocol (RTP) packets that contain the actual voice samples.
- The call control signaling, which consists of packets belonging to one of several protocols, according to the endpoints involved in the call (for example, H.323, MGCP, SCCP, or (J)TAPI). Call control functions are, for instance, those used to set up, maintain, tear down, or redirect a call.

Bandwidth provisioning should include not only the bearer traffic but also the call control traffic. In fact, in multisite WAN deployments, the call control traffic (as well as the bearer traffic) must traverse the WAN, and failure to allocate sufficient bandwidth for it can adversely affect the user experience.

The next three sub-sections describe the bandwidth provisioning recommendations for the following types of traffic:

- Voice and video bearer traffic in all multisite WAN deployments (see [Provisioning for Bearer Traffic, page 3-46](#))
- Call control traffic in multisite WAN deployments with centralized call processing (see [Provisioning for Call Control Traffic with Centralized Call Processing, page 3-53](#))
- Call control traffic in multisite WAN deployments with distributed call processing (see [Provisioning for Call Control Traffic with Distributed Call Processing, page 3-57](#))

Provisioning for Bearer Traffic


The section describes bandwidth provisioning for the following types of traffic:

- [Voice Bearer Traffic, page 3-46](#)
- [Video Bearer Traffic, page 3-48](#)

Voice Bearer Traffic

As illustrated in [Figure 3-15](#), a voice-over-IP (VoIP) packet consists of the payload, IP header, User Datagram Protocol (UDP) header, Real-Time Transport Protocol (RTP) header, and Layer 2 Link header. At the default packetization rate of 20 ms, VoIP packets have a 160-byte payload for G.711 or a 20-byte payload for G.729. When Secure Real-Time Transport Protocol (SRTP) encryption is used, the payload for each packet is increased by 4 bytes. At the default packetization rate of 20 ms, SRTP VoIP packets have a 164-byte payload for G.711 or a 24-byte payload for G.729. The IP header is 20 bytes, the UDP header is 8 bytes, and the RTP header is 12 bytes. The link header varies in size according to the Layer 2 media used.

Figure 3-15 Typical 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 packets per second). [Table 3-6](#) details the bandwidth per VoIP flow at a default packet rate of 50 packets per second (pps). [Table 3-6](#) does not include Layer 2 header overhead and does not take into account any possible compression schemes, such as compressed Real-Time Transport Protocol (cRTP). You can use the Service Parameters menu in Cisco Unified CallManager Administration to adjust the packet rate.

[Table 3-6](#) lists the bandwidth consumed by the voice payload and IP header only, at a default packet rate of 50 packets per second (pps) and at a rate of 33.3 pps for both non-encrypted and encrypted payloads.

Table 3-6 *Bandwidth Consumption for Voice Payload and IP Header Only*

CODEC	Sampling Rate	Voice Payload in Bytes	Packets per Second	Bandwidth per Conversation
G.711	20 ms	160	50.0	80.0 kbps
G.711 (SRTP)	20 ms	164	50.0	81.6 kbps
G.711	30 ms	240	33.3	74.7 kbps
G.711 (SRTP)	30 ms	244	33.3	75.8 kbps
G.729A	20 ms	20	50.0	24.0 kbps
G.729A (SRTP)	20 ms	24	50.0	25.6 kbps
G.729A	30 ms	30	33.3	18.7 kbps
G.729A (SRTP)	30 ms	34	33.3	19.8 kbps

A more accurate method for provisioning is to include the Layer 2 headers in the bandwidth calculations. [Table 3-7](#) lists the amount of bandwidth consumed by voice traffic when the Layer 2 headers are included in the calculations.


Table 3-7 *Bandwidth Consumption with Layer 2 Headers Included*

CODEC	Header Type and Size						
	Ethernet 14 Bytes	PPP 6 Bytes	ATM 53-Byte Cells with a 48-Byte Payload	Frame Relay 4 Bytes	MLPPP 10 Bytes	MPLS 4 Bytes	WLAN 24 Bytes
G.711 at 50.0 pps	85.6 kbps	82.4 kbps	106.0 kbps	81.6 kbps	84.0 kbps	81.6 kbps	89.6 kbps
G.711 (SRTP) at 50.0 pps	87.2 kbps	84.0 kbps	106.0 kbps	83.2 kbps	85.6 kbps	83.2 kbps	N/A
G.711 at 33.3 pps	78.4 kbps	76.3 kbps	84.8 kbps	75.7 kbps	77.3 kbps	75.7 kbps	81.1 kbps
G.711 (SRTP) at 33.3 pps	79.5 kbps	77.4 kbps	84.8 kbps	76.8 kbps	78.4 kbps	76.8 kbps	N/A
G.729A at 50.0 pps	29.6 kbps	26.4 kbps	42.4 kbps	25.6 kbps	28.0 kbps	25.6 kbps	33.6 kbps
G.729A (SRTP) at 50.0 pps	31.2 kbps	28.0 kbps	42.4 kbps	27.2 kbps	29.6 kbps	27.2 kbps	N/A
G.729A at 33.3 pps	22.4 kbps	20.3 kbps	28.3 kbps	19.7 kbps	21.3 kbps	19.7 kbps	25.1 kbps
G.729A (SRTP) at 33.3 pps	23.5 kbps	21.4 kbps	28.3 kbps	20.8 kbps	22.4 kbps	20.8 kbps	N/A

While it is possible to configure the sampling rate above 30 ms, doing so usually results in very poor voice quality. As illustrated in [Figure 3-16](#), as sampling size increases, the number of packets per second decreases, resulting in a smaller impact to the CPU of the device. Likewise, as the sample size increases, IP header overhead is lower because the payload per packet is larger. However, as sample size increases,

so does packetization delay, resulting in higher end-to-end delay for voice traffic. The trade-off between packetization delay and packets per second must be considered when configuring sample size. While this trade-off is optimized at 20 ms, 30 ms sample sizes still provide a reasonable ratio of delay to packets per second; however, with 40 ms sample sizes, the packetization delay becomes too high.

Figure 3-16 Voice Sample Size: Packets per Second vs. Packetization Delay



Video Bearer Traffic

For audio, it is relatively easy to calculate a percentage of overhead per packet given the sample size of each packet. For video, however, it is nearly impossible to calculate an exact percentage of overhead because the payload varies depending upon how much motion is present in the video (that is, how many pixels changed since the last frame).

To resolve this inability to calculate the exact overhead ratio for video, Cisco recommends that you add 20% to the call speed regardless of which type of Layer-2 medium the packets are traversing. The additional 20% gives plenty of headroom to allow for the differences between Ethernet, ATM, Frame Relay, PPP, HDLC, and other transport protocols, as well as some cushion for the bursty nature of video traffic. (See [Table 3-8](#).)

Table 3-8 Recommended Bandwidth for Various Video Call Speeds

Call Speed Requested by Endpoint	Actual Layer-2 Bandwidth Required
128 kbps	153.6 kbps
256 kbps	307.2 kbps
384 kbps	460.8 kbps
512 kbps	614.4 kbps
768 kbps	921.6 kbps
1.5 Mbps	1.766 Mbps
2.048 Mbps	2.458 Mbps
7 Mbps	8.4 Mbps

Note that the values in [Table 3-8](#) represent the maximum burst speed of the call, with some additional amount for a cushion. The average speed of the call is typically much less than these values. The concepts of media channels and bandwidth usage are critical to understanding what values to use when configuring call admission control.

Additional Considerations for Bearer Traffic with RSVP

RSVP was built to support requesting a Quality of Service (QoS) for any traffic flow, not just voice or video, across a wide range of Layer 2 technologies. To accomplish this, RSVP must be able to describe in detail the traffic flow for which it is requesting QoS so that the intermediate routers can make admittance decisions correctly.

The RSVP Sender generates Path messages containing the stream's bandwidth requirements, in the form of a traffic specification (Tspec). The Tspec, as described in RFC 2210, contains the following attributes to describe the traffic flow in detail:

- Token Bucket Rate [r] — The average traffic rate in bytes per second
- Token Bucket Size [b] — The maximum burst of a flow in bytes
- Peak Data Rate [p] — The peak traffic rate in bytes per second
- Minimum Policed Unit [m]
- Maximum Packet Size [M]

Calculating RSVP Bandwidth Values for Use with Cisco Unified CallManager

At the time Cisco Unified CallManager instructs the Cisco RSVP Agent to make the initial reservation for the call flow, the endpoints that are involved in the call have not fully exchanged their codec capabilities. Without this information, Cisco Unified CallManager must rely on the region settings to determine how to describe the traffic flow. The size of the traffic flow is a function of two things, the codec bit-rate and the sampling rate (or packets per second). The region settings contain the codec but do not describe the sampling rate. The preferred sampling rates for G.729 and G.711 audio codecs are defined in the following clusterwide service parameters:

- Preferred G711 millisecond packet size: 20 ms by default
- Preferred G729 millisecond packet size: 20 ms by default

However, the codec sampling rate is negotiated for every call and might not be the preferred setting because it is not supported on one or more of the endpoints. To avoid having to increase the reservation size once the capabilities are fully exchanged, possibly causing a post-ring failure, this initial reservation is for the worst-case scenario, or the smallest packet size, of that codec. Once media capabilities have been exchanged between the endpoints, then the reservation is revised to the correct bandwidth allocation. In most cases, the default sampling rate is used, resulting in the reservation being reduced.

**Note**

Cisco Unified CallManager does not include SRTP overhead or the Layer 2 overhead in the RSVP Reservation. When compared to the RSVP bandwidth value, the priority queue must be over-provisioned. (See [Table 3-7](#) and [Table 3-8](#).)

Voice Bearer Traffic

Inter-region call with audio codec set to G729:

- Initial request: 40 kbps, using a 10 ms worst-case scenario
- Updated request: 24 kbps, using the preferred sample size of 20 ms

Inter-region call with audio codec set to G711:

- Initial request: 96 kbps, using a 10 ms worst-case scenario
- Updated request: 80 kbps, using the preferred sample size of 20 ms

Video Bearer Traffic

As with the audio stream, the initial reservation for the video stream will rely on the region settings because the endpoint codec capabilities will not be fully negotiated at the time of the reservation. The region settings for video calls include the bandwidth for the audio stream. (See [IP Video Telephony, page 17-1](#), for more information.) Because the audio stream has its own reservation, the final reservation for the video stream will be the region setting minus the audio codec bit-rate. However, because these codecs have not been fully negotiated, the video stream reservation will be for the worst-case scenario, which assumes no audio stream. Once media capabilities have been exchanged between the endpoints, then the reservation will be revised to the correct bandwidth allocation.

Because video is inherently bursty, it is necessary to add some overhead to the stream requirements. (See [Video Bearer Traffic, page 3-48](#), for more information.) Cisco Unified CallManager uses the stream bandwidth to determine how to calculate the overhead, as follows:

- If the stream is < 256 kbps, then the overhead will be 20%
- If the stream is $\geq$ 256 kbps, then the overhead will be 7%

Inter-region video call, with G.729 audio codec and video setting of 384 kbps:

- Initial request: $384 * 1.07 = 410$ kbps
- Updated request: $(384 - 8) * 1.07 = 402$ kbps

Inter-region video call, with G.711 audio codec and video setting of 384 kbps:

- Initial request: $384 * 1.07 = 410$ kbps
- Updated request: $(384 - 64) * 1.07 = 342$ kbps

Configuration Recommendation

Because the initial reservation will be larger than the actual packet flow, over-provisioning the RSVP and LLQ bandwidth is required to ensure that the desired number of calls can complete.

When provisioning the RSVP bandwidth value for N calls, Cisco recommends that the Nth value be the worst-case bandwidth to ensure that the Nth call gets admitted.

For example:

- To provision four G.729 streams:
 $(3 * 24) + 40 = 112 \text{ kbps}$
- To provision four G.711 streams:
 $(3 * 80) + 96 = 336 \text{ kbps}$
- To provision four 384 kbps video streams (G.729 audio)
 $(3 * (384 - 8) + 384) * 1.07 = 1618 \text{ kbps}$
- To provision four 384 kbps video streams (G.711 audio)
 $(3 * (384 - 64) + 384) * 1.07 = 1438 \text{ kbps}$

Configuring Cisco IOS Application ID Support

RSVP Application ID feature support was introduced in Cisco IOS Release 12.4(6)T, and that is the minimum release required for the following examples.

Combined Priority Queue

To utilize the functionality allowed in Cisco Unified CallManager's implementation of Application ID support (that is, allowing voice calls to consume all the bandwidth available in the priority queue), we must modify the previous recommendations that voice and video priority queues be kept separate. (See [Application ID Call Admission Control Model, page 3-43](#).) To use this functionality, you should combine both the voice and video match criteria into one class-map. Because the requirements are to match either voice traffic or video traffic, be sure to make the class-map match criteria **match-any** instead of **match-all**, as follows:

```
class-map match-any IPC-RTP
  match ip dscp ef
  match ip dscp af41 af42
```

Configure the priority queue to support both the voice and video traffic. The following example configuration allocates 33% of the link bandwidth to the priority queue:

```
policy-map Voice-Policy
  class IPC-RTP
    priority percent 33
```

Mapping Application ID to RSVP Policy Identities

The RSVP Local Policy provides the mechanism for controlling a reservation based on an Application ID. Application IDs are mapped to RSVP Local Policies through the **ip rsvp policy identity** command. RSVP Local Policy identities are defined globally and are available to each interface for policy enforcement. Each identity can have one policy locator defined to match an Application ID.

To give the user as much flexibility as possible in matching application policy locators to local policies, the RSVP local policy command line interface (CLI) accepts application ID match criteria in the form of Unix-style regular expressions for the policy locator. Regular expressions are already used in the CLI for existing Cisco IOS components such as Border Gateway Protocol (BGP). Refer to the following documentation for more information on how regular expressions are used in Cisco IOS:

- *Access and Communication Servers Command Reference*

http://www.cisco.com/en/US/docs/ios/11_0/access/command/reference/arbook.html

- *Using Regular Expressions in BGP*
http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a0080094a92.shtml
- *Regex Engine Performance Enhancement*
http://www.cisco.com/en/US/docs/ios/12_3t/12_3t4/feature/guide/gt_rexpe.html

RSVP Policy Identities for Matching the Default Cisco Unified CallManager Application IDs

```
ip rsvp policy identity rsvp-video policy-locator .*VideoStream.*
ip rsvp policy identity rsvp-voice policy-locator .*AudioStream.*
```

Interface RSVP Local Policies

Whether configuring Application ID support or not, for an interface to support RSVP, you must configure the **ip rsvp bandwidth <value>** command on that interface. This value cannot be exceeded by any one RSVP reservation or the sum of RSVP reservations on that interface, regardless of Application ID support. In fact, if a reservation passes the local policy check, it still must pass the interface RSVP bandwidth check before it is reserved.

Local policies based on Application ID are applied to an interface using the **ip rsvp policy local identity** command.

For reservations that match its policy locator value, a local policy has the ability to perform the following functions:

- Define the maximum amount of bandwidth the reservations can reserve as a group or as a single sender
- Forward or not forward RSVP messages
- Accept or not accept RSVP messages
- Define the maximum bandwidth the group or sender can reserve

For example, to limit the amount of video bandwidth to 384 kbps on a Serial T1, use the following commands:

```
interface Serial0/0/1:0
ip rsvp bandwidth 506
ip rsvp policy local identity rsvp-video
    maximum bandwidth group 384
forward all
```

There is also a catch-all local policy called the default local policy. This local policy will match any RSVP reservation that did not match the other RSVP local policies configured on the link. The default local policy can be used to match reservations that are not tagged with an Application ID or reservations that are tagged with an Application ID that you want to treat as untagged traffic.

Example

The following example supports both voice and video calls using the model discussed in [How Cisco Unified CallManager Uses the Application ID, page 3-43](#). The voice calls are guaranteed 352 kbps of bandwidth while video calls are limited to 154 kbps of bandwidth. Voice calls can use all of the available RSVP bandwidth.

```
interface Serial0/0/1:0
ip address 10.2.101.5 255.255.255.252
service-policy output Voice-Policy
ip rsvp bandwidth 506
ip rsvp data-packet classification none
ip rsvp resource-provider none
ip rsvp policy local identity rsvp-voice
```

```
maximum bandwidth group 506
forward all
ip rsvp policy local identity rsvp-video
maximum bandwidth group 154
forward all
ip rsvp policy local default
no accept all ! Will not show in the configuration
no forward all! Will not show in the configuration
```

In this example, if an RSVP reservation is received that does not have an Application ID or its Application ID does not match the two configured options, the reservation will fail. This configuration works if RSVP traffic originates only from Cisco RSVP Agents controlled by Cisco Unified CallManager. However, if there is intercluster RSVP traffic via an IP-IP gateway or if RSVP messages from a controller other than Cisco Unified CallManager are traversing this link, then the default local policy should be configured to accept and forward the reservations and a maximum bandwidth value should be configured on the policy. Note that it is possible to oversubscribe the RSVP bandwidth via the use of multiple RSVP local policies (if the sum of the policies is greater than the RSVP interface bandwidth), but reservations then become first-come, first-serve.

Provisioning for Call Control Traffic with Centralized Call Processing

In a centralized call processing deployment, the Cisco Unified CallManager cluster and the applications (such as voicemail) are located at the central site, while several remote sites are connected through an IP WAN. The remote sites rely on the centralized Cisco Unified CallManagers to handle their call processing.

The following considerations apply to this deployment model:

- Each time a remote branch phone places a call, the control traffic traverses the IP WAN to reach the Cisco Unified CallManager at the central site, even if the call is local to the branch.
- The signaling protocols that may traverse the IP WAN in this deployment model are SCCP (encrypted and non-encrypted), H.323, MGCP, and CTI-QBE. All the control traffic is exchanged between a Cisco Unified CallManager at the central site and endpoints or gateways at the remote branches.

As a consequence, the area where you must provision bandwidth for control traffic lies between the branch routers and the WAN aggregation router at the central site.

The control traffic that traverses the WAN in this scenario can be split into two categories:

- Quiescent traffic, which consists of keep-alive messages periodically exchanged between the branch endpoints (phones and gateways) and Cisco Unified CallManager, regardless of call activity. This traffic is a function of the quantity of endpoints.
- Call-related traffic, which consists of signaling messages exchanged between the branch endpoints and the Cisco Unified CallManager at the central site when a call needs to be set up, torn down, forwarded, and so forth. This traffic is a function of the quantity of endpoints and their associated call volume.

To obtain an estimate of the generated call control traffic, it is therefore necessary to make some assumptions regarding the average number of calls per hour made by each branch IP phone. In the interest of simplicity, the calculations in this section assume an average of 10 calls per hour per phone.



Note

If this average number does not satisfy the needs of your specific deployment, you can calculate the recommended bandwidth by using the [Advanced Formulas, page 3-55](#).

Given the assumptions made, and initially considering the case of a remote branch with no signaling encryption configured, the recommended bandwidth needed for call control traffic can be obtained from the following formula:

Equation 1: Recommended Bandwidth Needed for SCCP Control Traffic without Signaling Encryption.

$$\text{Bandwidth (bps)} = 265 * (\text{Number of IP phones and gateways in the branch})$$

Equation 1 and all other formulas within this section include a 25% over-provisioning factor. Control traffic has a bursty nature, with peaks of high activity followed by periods of low activity. For this reason, assigning just the minimum bandwidth required to a control traffic queue can result in undesired effects such as buffering delays and, potentially, packet drops during periods of high activity. The default queue depth for a Class-Based Weighted Fair Queuing (CBWFQ) queue in Cisco IOS equals 64 packets. The bandwidth assigned to this queue determines its servicing rate. Assuming that the bandwidth configured is the average bandwidth consumed by this type of traffic, it is clear that, during the periods of high activity, the servicing rate will not be sufficient to "drain" all the incoming packets out of the queue, thus causing them to be buffered. Note that, if the 64-packet limit is reached, any subsequent packets are either assigned to the best-effort queue or are dropped. It is therefore advisable to introduce this 25% over-provisioning factor to absorb and smooth the variations in the traffic pattern and to minimize the risk of a temporary buffer overrun. This method is equivalent to increasing the servicing rate of the queue.

If encryption is configured, the recommended bandwidth is affected because encryption increases the size of signaling packets exchanged between Cisco Unified CallManager and the endpoints. The following formula takes into account the impact of signaling encryption:

Equation 2: Recommended Bandwidth Needed for SCCP Control Traffic with Signaling Encryption.

$$\text{Bandwidth with signaling encryption (bps)} = 415 * (\text{Number of IP phones and gateways in the branch})$$

If we now take into account the fact that the smallest bandwidth that can be assigned to a queue on a Cisco IOS router is 8 kbps, we can summarize the values of minimum and recommended bandwidth for various branch office sizes, as shown in [Table 3-9](#).

Table 3-9 *Recommended Layer 3 Bandwidth for Call Control Traffic With and Without Signaling Encryption*

Branch Office Size (Number of IP Phones and Gateways)	Recommended Bandwidth for SCCP Control Traffic (no encryption)	Recommended Bandwidth for SCCP Control Traffic (with encryption)
1 to 10	8 kbps	8 kbps
20	8 kbps	9 kbps
30	8 kbps	13 kbps
40	11 kbps	17 kbps
50	14 kbps	21 kbps
60	16 kbps	25 kbps
70	19 kbps	29 kbps
80	21 kbps	33 kbps
90	24 kbps	38 kbps
100	27 kbps	42 kbps
110	29 kbps	46 kbps

Table 3-9 Recommended Layer 3 Bandwidth for Call Control Traffic With and Without Signaling Encryption (continued)

Branch Office Size (Number of IP Phones and Gateways)	Recommended Bandwidth for SCCP Control Traffic (no encryption)	Recommended Bandwidth for SCCP Control Traffic (with encryption)
120	32 kbps	50 kbps
130	35 kbps	54 kbps
140	37 kbps	58 kbps
150	40 kbps	62 kbps

**Note**

Table 3-9 assumes 10 calls per hour per phone.

Advanced Formulas

The previous formulas presented in this section assume an average call rate per phone of 10 calls per hour. However, this rate might not correspond to your deployment if the call patterns are significantly different (for example, with call center agents at the branches). To calculate call control bandwidth requirements in these cases, use the following formulas, which contain an additional variable (CH) that represents the average calls per hour per phone:

Equation 3: Recommended Bandwidth Needed for SCCP Control Traffic for a Branch with No Signaling Encryption.

$$\text{Bandwidth (bps)} = (53 + 21 * \text{CH}) * (\text{Number of IP phones and gateways in the branch})$$

Equation 4: Recommended Bandwidth Needed for SCCP Control Traffic for a Remote Branch with Signaling Encryption.

$$\text{Bandwidth with signaling encryption (bps)} = (73.5 + 33.9 * \text{CH}) * (\text{Number of IP phones and gateways in the branch})$$

**Note**

Equations 3 and 4 are based on the default SCCP keep-alive period of 30 seconds.

Considerations for Shared Line Appearances

Calls placed to shared line appearances, or calls sent to line groups using the Broadcast distribution algorithm, have two net effects on the bandwidth consumed by the system:

- Because all the phones on which the line is configured ring simultaneously, they represent a load on the system corresponding to a much higher calls-per-hour (CH) value than the CH of the line. The corresponding bandwidth consumption is therefore increased. The network infrastructure's bandwidth provisioning requires adjustments when WAN-connected shared line functionality is deployed. The CH value employed for Equations 3 and 4 must be increased according to the following formula:

$$\text{CHS} = \text{CHL} * (\text{Number line appearances}) / (\text{Number of lines})$$

Where CHS is the shared-line calls per hour to be used in Equations 3 and 4, and CHL is the calls-per-hour rating of the line. For example, if a site is configured with 5 lines making an average of 6 calls per hour but 2 of those lines are shared across 4 different phones, then:

$$\text{Number of lines} = 5$$

Number of line appearances = (2 lines appear on 4 phones, and 3 lines appear on only one phone) = $(2 * 4) + 3 = 11$ line appearances

CHL = 6

CHS = $6 * (11 / 5) = 13.2$

- Because each of the ringing phones requires a separate signaling control stream, the quantity of packets sent from Cisco Unified CallManager to the same branch is increased in linear proportion to the quantity of phones ringing. Because Cisco Unified CallManager is attached to the network through a 100 Mbps interface, it can instantaneously generate a very large quantity of packets that must be buffered while the queuing mechanism is servicing the signaling traffic. The servicing speed is limited by the WAN interface's effective information transfer speed, which is typically two orders of magnitude smaller than 100 Mbps.

This traffic may overwhelm the queue depth of the central site's WAN router. By default, the queue depth available for each of the classes of traffic in Cisco IOS is 64. In order to prevent any packets from being dropped before they are queued for the WAN interface, you must ensure that the signaling queue's depth is sized to hold all the packets from at least one full shared-line event for each shared-line phone. Avoiding drops is paramount in ensuring that the call does not create a race condition where dropped packets are retransmitted, causing system response times to suffer.

Therefore, the quantity of packets required to operate shared-line phones in Cisco Unified CallManager 4.x is 14 packets per shared-line phone.

For example, in Cisco Unified CallManager 4.1 with 5 phones sharing the same line, the queue depth for the signaling class of traffic must be adjusted to a minimum of 70. [Table 3-10](#) provides recommended queue depths based on the quantity of shared line appearances within a branch site.

Table 3-10 Recommended Queue Depth per Branch Site

Number of Shared Line Appearances	Queue Depth (Packets) with SCCP and Cisco Unified CallManager 4.x
5	70
10	140
15	210
20	280
25	350

When using a Layer 2 WAN technology such as Frame Relay, you must make this adjustment on the circuit corresponding to the branch where the shared-line phones are located.

If you are using a Layer 3 WAN technology such as MPLS, there might be a single signaling queue servicing multiple branches. In this case, you must make an adjustment for the total of all branches serviced. For example, with Cisco Unified CallManager 4.1, if 10 branches each require 5 phones sharing a line, the central site's WAN router must be adjusted to have a signaling queue depth of 700.

Provisioning for Call Control Traffic with Distributed Call Processing

In distributed call processing deployments, several sites are connected through an IP WAN. Each site contains a Cisco Unified CallManager cluster and can follow either the single-site model or the centralized call processing model. A gatekeeper may be used for call admission control between sites.

The following considerations apply to this deployment model:

- The signaling protocol used to place a call across the WAN is H.323 or SIP.
- Control traffic is exchanged between the Cisco IOS gatekeeper and the Cisco Unified CallManager clusters at each site, as well as between the Cisco Unified CallManager clusters themselves.

Therefore, bandwidth for control traffic must be provisioned on the WAN links between Cisco Unified CallManagers as well as between each Cisco Unified CallManager and the gatekeeper. Because the topology is limited to hub-and-spoke, with the gatekeeper typically located at the hub, the WAN link that connects each site to the other sites usually coincides with the link that connects the site to the gatekeeper.

The control traffic that traverses the WAN belongs to one of the following categories:

- Quiescent traffic, which consists of registration messages periodically exchanged between each Cisco Unified CallManager and the gatekeeper
- Call-related traffic, which in turn consists of two types of traffic:
 - Call admission control traffic, exchanged between the Cisco Unified CallManagers and the infrastructure control point (such as a gatekeeper) before a call can be set up and after it has been torn down.
 - Signaling traffic associated with a media stream and exchanged over an intercluster trunk when a call needs to be set up, torn down, forwarded, and so on.

Because the total amount of control traffic depends on the number of calls that are set up and torn down at any given time, it is necessary to make some assumptions about the call patterns and the link utilization. The WAN links that connect each of the spoke sites to the hub site are normally provisioned to accommodate different types of traffic (for example, data, voice, and video). Using a traditional telephony analogy, we can view the portion of the WAN link that has been provisioned for voice as a number of *virtual tie lines*.

Assuming an average call duration of 2 minutes and 100 percent utilization of each virtual tie line, we can deduce that each tie line carries a volume of 30 calls per hour. This assumption allows us to obtain the following formula that expresses the recommended bandwidth for call control traffic as a function of the number of virtual tie lines.

Equation 5: Recommended Bandwidth Based on Number of Virtual Tie Lines.

$$\text{Recommended Bandwidth (bps)} = 116 * (\text{Number of virtual tie lines})$$

If we take into account the fact that 8 kbps is the smallest bandwidth that can be assigned to a queue on a Cisco IOS router, we can deduce that a minimum queue size of 8 kbps can accommodate the call control traffic generated by *up to 70 virtual tie lines*. This amount should be sufficient for most large enterprise deployments.

Wireless LAN Infrastructure

Wireless LAN infrastructure design becomes important when IP telephony is added to the wireless LAN (WLAN) portions of a converged network. With the addition of wireless IP telephony endpoints such as the Cisco Unified Wireless IP Phone 7920, voice traffic has moved onto the WLAN and is now converged with the existing data traffic there. Just as with wired LAN and wired WAN infrastructure, the addition of voice in the WLAN requires following basic configuration and design best-practices for deploying a highly available network. In addition, proper WLAN infrastructure design requires understanding and deploying QoS on the wireless network to ensure end-to-end voice quality on the entire network. The following sections discuss these requirements:

- [WLAN Design and Configuration, page 3-58](#)
- [WLAN Quality of Service \(QoS\), page 3-64](#)

For more information about WLAN design, refer to the *Cisco Wireless LAN SRND* guide, available at <http://www.cisco.com/go/designzone>

For more information about the Cisco Wireless IP Phone 7920, refer to the *Cisco Unified Wireless IP Phone 7920 Design and Deployment Guide*, available at <http://www.cisco.com/go/designzone>

WLAN Design and Configuration

Properly designing a WLAN requires, first and foremost, ensuring that the existing wired network is deployed in a highly available, fault-tolerant and redundant manner. Next, an understanding of wireless technology is required. Finally, by configuring and deploying wireless access points (APs) and wireless telephony endpoints in an effective way, you can build a flexible, secure, redundant, and highly scalable network.

The following sections examine the WLAN infrastructure layers and network services:

- [Wireless Infrastructure Considerations, page 3-58](#)
- [Wireless AP Configuration and Design, page 3-61](#)
- [Wireless Security, page 3-62](#)

Wireless Infrastructure Considerations

The following sections provide guidelines and best practices for designing the WLAN infrastructure:

- [VLANs, page 3-58](#)
- [Roaming, page 3-59](#)
- [Wireless Channels, page 3-59](#)
- [Wireless Interference, page 3-60](#)
- [Multicast on the WLAN, page 3-60](#)

VLANs

Just as with a wired LAN infrastructure, when deploying voice in a wireless LAN, you should enable at least two virtual LANs (VLANs) at the Access Layer. The Access Layer in a wireless LAN environment includes the access point (AP) and the first-hop access switch. On the AP and access switch, you should configure both a native VLAN for data traffic and a voice VLAN (under Cisco IOS) or Auxiliary VLAN

(under CatOS) for voice traffic. This voice/auxiliary VLAN should be separate from all the other wired voice VLANs in the network. In addition, as with voice endpoints on wired LANs, wireless voice endpoints should be addressed using RFC 1918 private subnet addresses. When deploying a wireless infrastructure, Cisco also recommends configuring a separate management VLAN for the management of WLAN APs. This management VLAN should not have a WLAN appearance; that is, it should not have an associated service set identifier (SSID) and it should not be directly accessible from the WLAN.

Roaming

Another very important consideration for wireless infrastructure is wireless endpoint roaming. When wireless devices roam at Layer 2, they keep their IP address and network configuration. For this reason, roaming can occur fairly quickly (in 100 to 400 ms). All that is required is re-authentication, if Cisco LEAP or Extensible Authentication Protocol (EAP) is used, and the passing of Inter-Access Point Protocol (IAPP) messages between the last AP and the new AP to indicate that the endpoint has roamed. Layer 2 roaming is typically unnoticeable to the end user.

When devices roam at Layer 3, they move from one AP to another AP across native VLAN boundaries. The Cisco Catalyst 6500 Series Wireless LAN Services Module (WLSM) allows the Cisco Wireless IP Phone 7920 to roam at Layer 3 while still maintaining an active call. The Cisco Wireless IP Phone 7920 can roam at Layer 3 using Static WEP or Cisco Centralized Key Management (Cisco CKM) protocols. Cisco CKM enables the Cisco Wireless IP Phone 7920 to achieve full Layer 3 mobility while using either WEP 128 or TKIP encryption. Seamless Layer 3 roaming occurs only when the client is roaming within the same mobility group. For details about the Cisco WLSM and Layer 3 roaming, refer to the product documentation available at

<http://www.cisco.com>

If you are using 802.1x authentication in the wireless LAN, Cisco CKM is recommended to minimize roaming downtime. Whether roaming at Layer 2 or Layer 3, device downtime decreases from 300-400 ms to 100 ms or less. Cisco CKM also takes some of the load off the Access Control Server (ACS) by reducing the number of authentication requests that must be sent to the ACS.



Note

If Cisco Catalyst 4000 Series switches are used as Layer 3 devices at the distribution layer, a minimum of a Supervisor Engine 2+ (SUP2+) or Supervisor Engine 3 (SUP3) module is required. The Supervisor Engine 1 or 2 (SUP1 or SUP2) modules can cause roaming delays. The Cisco Catalyst 2948G, 2948G-GE-TX, 2980G, 2980G-A, and 4912 switches are also known to introduce roaming delays. Cisco does not recommend using these switches in a wireless voice network.

Wireless Channels

Wireless endpoints and APs communicate via radios on particular channels. When communicating on one channel, wireless endpoints typically are unaware of traffic and communication occurring on other non-overlapping channels.

Optimal channel configuration for 2.4 GHz 802.11b requires a minimum of five-channel separation between configured channels to prevent interference or overlap between channels. In North America, with allowable channels of 1 to 11, channels 1, 6, and 11 are the three usable non-overlapping channels for APs and wireless endpoint devices. However, in Europe where the allowable channels are 1 to 13, multiple combinations of five-channel separation are possible. Multiple combinations of five-channel separation are also possible in Japan, where the allowable channels are 1 to 14.

AP coverage should be deployed so that no (or minimal) overlap occurs between APs configured with the same channel. Same channel overlap should typically occur at 19 dBm separation. However, proper AP deployment and coverage on non-overlapping channels require a minimum overlap of 15% to 20%.

This amount of overlap ensures smooth roaming for wireless endpoints as they move between AP coverage cells. Overlap of less than 15% to 20% can result in slower roaming times and poor voice quality.

Deploying wireless devices in a multi-story building such as an office high-rise or hospital introduces a third dimension to wireless AP and channel coverage planning. The 2.4 GHz wave form of 802.11b can pass through floors and ceilings as well as walls. For this reason, not only is it important to consider overlapping cells or channels on the same floor, but it is also necessary to consider channel overlap between adjacent floors. With only three channels, proper overlap can be achieved only through careful three-dimensional planning.


Note

Careful deployment of APs and channel configuration within the wireless infrastructure are imperative for proper wireless network operation. For this reason, Cisco requires that a complete and thorough site survey be conducted before deploying wireless networks in a production environment. The survey should include verifying non-overlapping channel configurations, AP coverage, and required data and traffic rates; eliminating rogue APs; and identifying and mitigating the impact of potential interference sources.

Wireless Interference

Interference sources within a wireless environment can severely limit endpoint connectivity and channel coverage. In addition, objects and obstructions can cause signal reflection and multipath distortion. Multipath distortion occurs when traffic or signaling travels in more than one direction from the source to the destination. Typically, some of the traffic arrives at the destination before the rest of the traffic, which can result in delay and bit errors in some cases. You can reduce the affects of multipath distortion by eliminating or reducing interference sources and obstructions, and by using diversity antennas so that only a single antenna is receiving traffic at any one time. Interference sources should be identified during the site survey and, if possible, eliminated. At the very least, interference impact should be alleviated by proper AP placement and the use of location-appropriate directional or omni-directional diversity radio antennas.

Possible interference sources include:

- Other APs on overlapping channels
- Other 2.4 GHz appliances, such as 2.4 GHz cordless phones, personal wireless network devices, sulphur plasma lighting systems, microwave ovens, rogue APs, and other WLAN equipment that takes advantage of the license-free operation of the 2.4 GHz band
- Metal equipment, structures, and other metal or reflective surfaces such as metal I-beams, filing cabinets, equipment racks, wire mesh or metallic walls, fire doors and fire walls, concrete, and heating and air conditioning ducts
- High-power electrical devices such as transformers, heavy-duty electric motors, refrigerators, elevators, and elevator equipment

Multicast on the WLAN

Cisco does not recommend forwarding multicast traffic on a WLAN with voice devices because of the following reasons:

- Multicast packets are buffered on the AP if devices associated with the AP are in power-save mode. When devices such as the Cisco Unified Wireless IP Phone 7920 go into power-save mode, an AP buffers all multicast packets until the next active interval for these devices. This buffering causes packet delay and affects all devices associated with the AP even if they are not in power-save mode. This condition can be extremely problematic for real-time multicast applications such as music on hold and streaming video.

- Multicast packets on the WLAN are unacknowledged and are not retransmitted if lost or corrupted. AP and wireless endpoint devices use acknowledgements on the link layer to ensure reliable delivery. When packets are not received or acknowledged, they are retransmitted. This retransmission does not occur for multicast traffic on the WLAN. Given that wireless networks have a higher instance of bit error than wired networks, this lack of retransmission results in a larger number of lost packets when compared to a wired LAN.

Before enabling multicast applications on the wireless network, Cisco recommends testing these applications to ensure that performance and behavior are acceptable.

For additional considerations with multicast traffic, see [Music on Hold, page 7-1](#).

Wireless AP Configuration and Design

Proper AP selection, deployment, and configuration are essential to ensure that the wireless network handles voice traffic in a way that provides high-quality voice to the end users.

AP Selection

Cisco recommends the following APs for deploying wireless voice:

- Aironet 350 Series AP
- Aironet 1100 Series AP
- Aironet 1130 Series AP
- Aironet 1200 Series AP
- Aironet 1240 Series AP
- Aironet 1300 Series AP
- Airespace 1000 Series AP

For these APs, Cisco IOS Release 12.3(4) JA or later is recommended.

AP Deployment

When deploying Cisco access points (APs), do not associate more than 15 to 25 devices to a single AP at any given time. This number will vary depending on usage profiles. The number of devices on an AP affects the amount of time each device has access to the medium. As the number of devices increases, the traffic contention increases. Associating more than 15 to 25 devices to an AP can result in poor AP performance and slower response times for associated devices.

While there is no specific mechanism to ensure that no more than 25 devices are associated to a single AP, system administrators can manage device-to-AP ratios by conducting periodic site surveys and analyzing user and device traffic patterns. If additional devices and users are added to the network in a particular area, additional site surveys should be conducted to determine whether additional APs are required to handle the number of endpoints that need to access the network.

AP Configuration

When deploying wireless voice, observe the following specific AP configuration requirements:

- Enable Address Resolution Protocol (ARP) caching.

ARP caching is required on the AP because it enables the AP to answer ARP requests for the wireless endpoint devices without requiring the endpoint to leave power-save or idle mode. This feature results in extended battery life for the wireless endpoint devices.

- Match the transmit power on the AP to that on the wireless voice endpoints.

When possible, the transmit power on the AP and the voice endpoints should match. Matching transmit power on the AP and voice endpoints helps eliminate the possibility of one-way audio traffic. If the transmit power configuration on the APs must vary, the transmit power on all voice endpoints should be configured to match the AP with the highest transmit power.



Note Beginning with version 1.0(8) of the Cisco Wireless IP Phone 7920 firmware, the phone will take advantage of the Dynamic Transmit Power Control (DTPC) feature by automatically adjusting its transmit power based on the Limit Client Power (mW) setting of the current AP.

- Set the data rate to 11 Mbps.

Configuring the maximum 11 Mbps data rate ensures the best level of throughput for voice devices and the largest number of active calls per AP.

- Manually configure the RF channel selection. (Do *not* use the option **Search for Least Congested Channel**).

To control wireless network channels and eliminate channel overlap, it is important to configure a channel number manually on each AP based on its location.

- Assign a Service Set Identifier (SSID) to each VLAN configured on the AP.

SSIDs enable endpoints to select the wireless VLAN they will use for sending and receiving traffic. These wireless VLANs and SSIDs map to wired VLANs. For voice endpoints, this mapping ensures priority queuing treatment and access to the voice VLAN on the wired network.

- Enable **QoS Element for Wireless Phones** on the AP.

This feature ensures that the AP will provide QoS Basic Service Set (QBSS) information elements in beacons. The QBSS element provides an estimate of the channel utilization on the AP, and Cisco wireless voice devices use it to help make roaming decisions and to reject call attempts when loads are too high. Beginning with Cisco IOS Release 12.3(7)JA, the AP also provides 802.11e clear channel assessment (CCA) QBSS in beacons. The CCA-based QBSS values reflect true channel utilization.

- Configure two QoS policies on the AP, and apply them to the VLANs and interfaces.

Configure a voice policy and a data policy with default classifications for the respective VLANs to ensure that voice traffic is given priority queuing treatment. (See [Interface Queuing, page 3-65](#), for more information).

Wireless Security

Another important consideration for a wireless infrastructure is security. Wireless endpoints, including wireless phones, can connect to the wireless network using one of the following security mechanisms:

- Extensible Authentication Protocol-Flexible Authentication via Secure Tunneling (EAP-FAST)

A standards-based security protocol that first requires the establishment of an authenticated tunnel between the wireless client and an Authentication, Authorization, and Accounting (AAA) server using a Protected Access Credential (PAC). Next the wireless endpoint authenticates across the tunnel using a user name and password to authenticate with the network via 802.1X. Once this authentication occurs, traffic to and from the wireless device is encrypted using TKIP or WEP. Using the 802.1X authentication method requires an EAP-compliant Remote Authentication Dial-In User Service (RADIUS) authentication server such as the Cisco Secure Access Control Server (ACS), which provides access to a user database for authenticating the wireless devices.

- **Wi-Fi Protected Access (WPA)**

A standards-based security protocol requires the wireless endpoint to provide a user name and password to authenticate with the network. Once this authentication occurs using either 802.1X or Pre-Shared Key (PSK), traffic to and from the wireless device is encrypted using Temporal Key Integrity Protocol (TKIP). Using the 802.1X authentication method requires an EAP-compliant Remote Authentication Dial-In User Service (RADIUS) authentication server such as the Cisco Secure Access Control Server (ACS), which provides access to a user database for authenticating the wireless devices.

- **Cisco LEAP**

Cisco LEAP requires the wireless endpoint to provide a user name and password to authenticate with the network. Once this authentication occurs, a dynamic key is generated, and traffic to and from the wireless device is encrypted. This method requires an EAP-compliant Remote Authentication Dial-In User Service (RADIUS) authentication server such as the Cisco Secure Access Control Server (ACS), which provides access to a user database for authenticating the wireless devices.

- **Static Wired Equivalent Privacy (WEP)**

Static WEP requires the exchange of a statically configured 40-bit or 128-bit character key between the wireless endpoint and the AP. If the keys match, the wireless device is given access to the network. Be aware that there are known weaknesses in the WEP encryption algorithm. These weaknesses, coupled with the complexity of configuring and maintaining static keys, can make this security mechanism undesirable in many cases.

Authentication and ACS Deployment Models

Extensible Authentication Protocol (EAP) is the preferred method of wireless device authentication (especially voice devices) because it provides the most secure and robust mechanism for access to the network and voice VLAN(s). Because an EAP-compliant RADIUS server is required, Cisco recommends the use of Cisco Secure ACS for Windows Server Version 3.1 or later.

When deploying EAP-FAST, WPA, or Cisco LEAP for wireless authentication and encryption, carefully consider the placement of the ACS within the network, and select one of the following ACS deployment models:

- **Centralized ACS**

ACS server or servers are located in a centralized place within the network and are used to authenticate all wireless devices and users within the network.

- **Remote ACS**

In networks where remote locations are separated from the central site by low-speed or congested WAN links, an ACS server can be located at the remote site and remote wireless devices or users can be authenticated by this server locally, thus eliminating the potential for delayed authentication via a centralized ACS across the WAN link.

- **Local and fallback RADIUS server on the Cisco AP**

In networks where remote locations are separated from a central site by low-speed WAN links, local wireless devices can authenticate against local Cisco IOS APs. APs running Cisco IOS Release 12.2(11)JA or later can authenticate users and devices locally without relying on an external ACS. A single AP can support up to 50 users with this functionality. This feature can be used in lieu of a centralized or local ACS, or in the case of a WAN or ACS failure in which the remote site users are unable to contact a local ACS or the central site ACS.

When choosing a deployment model for the ACS, it is imperative to make authentication services redundant so that the ACS does not become a single point of failure when wireless devices attempt to access the network. For this reason, each ACS server should replicate its database to a secondary server. Furthermore, it is always a good idea to provide a local ACS or an on-AP RADIUS server at remote sites to ensure that remote wireless devices can still authenticate in the event of a WAN failure.

In addition to ACS server placement, it is also important to consider the implications of user database location in relation to the ACS server. Because the ACS server must access the user database to authenticate wireless devices, the location of the user database affects the amount of time the authentication will take. If the user database is a Microsoft Active Directory (AD) server located on the network, the ACS must send an authentication request to the AD server and wait for a response. To ensure the fastest response times for wireless voice endpoints attempting to authenticate to the network, Cisco recommends defining users locally on the ACS server. Remote databases have unknown response times and can adversely affect authentication times.

WLAN Quality of Service (QoS)

Just as QoS is necessary for LAN and WAN wired network infrastructure in order to ensure high voice quality, QoS is also required for wireless LAN infrastructure. Because of the bursty nature of data traffic and the fact that real-time traffic such as voice is sensitive to packet loss and delay, QoS tools are required to manage wireless LAN buffers, limit radio contention, and minimize packet loss, delay, and delay variation.

However, unlike most wired networks, wireless networks are a shared medium, and wireless endpoints do not have dedicated bandwidth for sending and receiving traffic. While wireless endpoints can mark traffic with 802.1p CoS, DSCP, and PHB, the shared nature of the wireless network means limited admission control and access to the network for these endpoints.

Wireless QoS involves the following main areas of configuration:

- [Traffic Classification, page 3-64](#)
- [Interface Queuing, page 3-65](#)
- [Bandwidth Provisioning, page 3-65](#)

Traffic Classification

As with wired network infrastructure, it is important to classify or mark pertinent wireless traffic as close to the edge of the network as possible. Because traffic marking is an entrance criterion for queuing schemes throughout the wired and wireless network, marking should be done at the wireless endpoint device whenever possible. Marking or classification by wireless network devices should be identical to that for wired network devices, as indicated in [Table 3-3](#).

In accordance with traffic classification guidelines for wired networks, the Cisco Wireless IP Phone 7920 marks voice media traffic or RTP traffic with DSCP 46 (or PHB EF) and voice signaling traffic (SCCP) with DSCP 24 (or PHB CS3). Once this traffic is marked, it can be given priority or better than best-effort treatment and queuing throughout the network. All wireless voice devices should be capable of marking traffic in this manner. All other traffic on the wireless network should be marked as best-effort or with some intermediary classification as outlined in wired network marking guidelines.

Interface Queuing

Once marking has occurred, it is necessary to enable the wired network APs and devices to provide QoS queuing so that voice traffic types are given separate queues to reduce the chances of this traffic being dropped or delayed as it traversed the wireless LAN. Queuing on the wireless network occurs in two directions, upstream and downstream. Upstream queuing concerns traffic traveling from the wireless endpoint up to the AP and from the AP up to the wired network. Downstream queuing concerns traffic traveling from the wired network to the AP and down to the wireless endpoint.

Unfortunately, there is little upstream queuing available in a wireless network. While wireless devices such as the Cisco Wireless IP Phone 7920 can provide queuing upstream as the packets leave the device, there is no mechanism in place to provide queuing among all clients on the wireless LAN because wireless networks are a shared medium. Therefore, although voice media packets might receive priority treatment leaving the wireless endpoint, these packets must contend with all the other packets that other wireless devices may be attempting to send. For this reason, it is extremely important to follow the guideline of no more than 15 to 25 wireless clients per AP. Going beyond the upper limit of this guideline can result in additional voice packet delay and jitter.

As for downstream QoS, Cisco APs currently provide up to eight queues for downstream traffic being sent to wireless clients. The entrance criterion for these queues can be based on a number of factors including DSCP, access control lists (ACLs), and VLAN. Although eight queues are available, Cisco recommends using only two queues when deploying wireless voice. All voice media and signaling traffic should be placed in the highest-priority queue, and all other traffic should be placed in the best-effort queue. This ensures the best possible queuing treatment for voice traffic.

To set up this two-queue configuration, create two QoS policies on the AP. Name one policy **voice** and configure it with the class of service **Voice <10 ms Latency (6)** as the **Default Classification for all packets on the Vlan**. Name the other policy **data** and configure it with the class of service **Best Effort (0)** as the **Default Classification for all packets on the Vlan**. Then assign the **data** policy to the incoming and outgoing radio interface for the data VLAN(s), and assign the **voice** policy to the incoming and outgoing radio interfaces for the voice VLAN(s). With the QoS policies applied at the VLAN level, the AP is not forced to examine every packet coming in or going out to determine the type of queuing it should receive. This configuration ensures that all voice media and signaling are given priority queuing treatment in a downstream direction.

Bandwidth Provisioning

Another QoS requirement for wireless networking is the appropriate provisioning of bandwidth. Bandwidth provisioning involves the bandwidth between the wired and wireless networks as well as the number of simultaneous voice calls that an AP can handle. Wireless APs typically connect to the wired network via a 100 Mbps link to an Access Layer switch port. While the ingress Ethernet port on the AP can receive traffic at 100 Mbps, the maximum throughput on an 802.11b wireless network is 11 Mbps. After taking into account the half-duplex nature of the wireless medium and the overhead of wireless headers, the practical throughput on the 802.11b wireless network is about 7 Mbps. This mismatch in throughput between the wired and wireless network can result in packet drops when traffic bursts occur in the network.

Rather than allowing traffic bursts to send excessive traffic toward the AP only to have it dropped by the AP, it is a good idea to rate-limit or police this traffic to a rate that the wireless network can handle. Forcing the AP to drop excessive traffic causes increased CPU utilization and congestion at the AP. Instead, limiting the traffic rate to 7 Mbps on the link between the wired Access Layer switch and the wireless AP ensures that traffic is dropped at the Access Layer switch, thus removing the burden from the AP. For more information on rate-limiting traffic sent to the AP, see the QoS recommendations in the

section on the [Cisco Unified Wireless IP Phone 7920, page 21-30](#). Note that, depending on the wireless network deployment, the practical throughput might be less than 7 Mbps, especially if more than the recommended number of devices are associated to a single AP.

Based on wireless voice network testing, Cisco has determined that a single 802.11b wireless AP can support up to seven (7) active G.711 voice streams or eight (8) active G.729 voice streams.

**Note**


A call between two phones associated to the same AP counts as two active voice streams.

If these limits are exceeded, voice quality will suffer and voice calls might be dropped. While there is no true call admission control mechanism or method for provisioning wireless bandwidth for voice traffic, the Cisco Wireless IP Phone 7920 can provide a simplified version of call admission control or bandwidth provisioning based on channel utilization information received from APs on the network. This information can be sent by the AP to the phone via a beacon that includes the QoS Basic Service Set (QBSS). The higher the QBSS element value, the higher the channel utilization and the less likely the channel and AP can provide sufficient bandwidth for additional wireless voice devices. If the QBSS element value is above the maximum threshold, then any calls attempted by the wireless IP phone will be rejected with a "Network Busy" message. In addition, the wireless IP phone considers the QBSS element in its roaming algorithm and will not roam to an AP that is sending beacons with a QBSS element above the maximum threshold.

**Note**

Beginning with Cisco IOS Release 12.3(7)JA, the AP sends 802.11e CCA-based QBSS. These QBSS values represent true channel utilization for a particular AP.

The QBSS information element is sent by the AP only if **QoS Element for Wireless Phones** has been enable on the AP. (Refer to [Wireless AP Configuration and Design, page 3-61](#).)



CHAPTER 4

Gateways

Last revised on: October 30, 2008

Gateways provide a number of methods for connecting an IP telephony network to the Public Switched Telephone Network (PSTN), a legacy PBX, or key systems. Gateways range from specialized, entry-level and stand-alone voice gateways to high-end, feature-rich integrated router and Cisco Catalyst gateways.

This chapter explains important factors to consider when selecting a Cisco gateway to provide the appropriate protocol and feature support for your IP Telephony network. The main topics discussed in this chapter include:

- [Understanding Cisco Gateways, page 4-2](#)
- [Gateway Selection, page 4-2](#)
- [QSIG Support, page 4-19](#)
- [Fax and Modem Support, page 4-20](#)
- [Gateways for Video Telephony, page 4-31](#)

What's New in This Chapter

[Table 4-1](#) lists the topics that are new in this chapter or that have changed significantly from previous releases of this document.

Table 4-1 *New or Changed Information Since the Previous Release of This Document*

New or Revised Topic	Described in:
Video gateways	Gateways for Video Telephony, page 4-31

Understanding Cisco Gateways

Cisco access gateways allow Cisco Unified CallManager to communicate with non-IP telecommunications devices. There are two types of Cisco access gateways, analog and digital.

Cisco Access Analog Gateways

There are two categories of Cisco access analog gateways, trunk gateways and station gateways.

- Access analog station gateways

Analog station gateways connect Cisco Unified CallManager to Plain Old Telephone Service (POTS) analog telephones, interactive voice response (IVR) systems, fax machines, and voice mail systems. Station gateways provide Foreign Exchange Station (FXS) ports.

- Access analog trunk gateways

Analog trunk gateways connect Cisco Unified CallManager to PSTN central office (CO) or PBX trunks. Trunk gateways provide Foreign Exchange Office (FXO) ports for access to the PSTN, PBXs, or key systems, and E&M (recEive and transMit, or ear and mouth) ports for analog trunk connection to a legacy PBX. Whenever possible, use digital gateways to minimize any answer and disconnect supervision issues. Analog Direct Inward Dialing (DID) and Centralized Automatic Message Accounting (CAMA) are also available for PSTN connectivity.

Cisco Access Digital Trunk Gateways

A Cisco access digital trunk gateway connects Cisco Unified CallManager to the PSTN or to a PBX via digital trunks such as Primary Rate Interface (PRI), Basic Rate Interface (BRI), or T1 Channel Associated Signaling (CAS). Digital T1 PRI trunks may also be used to connect to certain legacy voice mail systems.

Gateway Selection

When selecting an IP telephony gateway, consider the following factors:

- [Core Feature Requirements, page 4-3](#)
- [Gateway Protocols, page 4-3](#)
- [Gateway Protocols and Core Feature Requirements, page 4-6](#)
- [Site-Specific Gateway Requirements, page 4-12](#)

Core Feature Requirements

Gateways used in IP telephony applications must meet the following core feature requirements:

- Dual tone multifrequency (DTMF) relay capabilities

DTMF relay capability, specifically out-of-band DTMF, separates DTMF digits from the voice stream and sends them as signaling indications through the gateway protocol (H.323, SCCP, MGCP, or SIP) signaling channel instead of as part of the voice stream or bearer traffic. Out-of-band DTMF is required when using a low bit-rate codec for voice compression because the potential exists for DTMF signal loss or distortion.

- Supplementary services support

Supplementary services are typically basic telephony functions such as hold, transfer, and conferencing.

- Fax/modem support

Fax over IP enables interoperability of traditional analog fax machines with IP telephony networks. The fax image is converted from an analog signal and is carried as digital data over the packet network. For more information, see [Fax and Modem Support, page 4-20](#)

- Cisco Unified CallManager redundancy support

Cisco Unified Communications is based on a distributed model for high availability. Cisco Unified CallManager clusters provide for Cisco Unified CallManager redundancy. The gateways must support the ability to “re-home” to a secondary Cisco Unified CallManager in the event that a primary Cisco Unified CallManager fails. Redundancy differs from call survivability in the event of a Cisco Unified CallManager or network failure.

Refer to the gateway product documentation to verify that any IP Telephony gateway you select for an enterprise deployment can support the preceding core requirements. Additionally, every IP Telephony implementation has its own site-specific feature requirements, such as analog or digital access, DID, and capacity requirements (see [Site-Specific Gateway Requirements, page 4-12](#)).

Gateway Protocols

Cisco Unified CallManager (Release 3.1 and later) supports the following gateway protocols:

- H.323
- Media Gateway Control Protocol (MGCP)

Cisco Unified CallManager Release 4.0 and later supports Session Initiation Protocol (SIP) on the trunk side. The SIP trunk implementation has been enhanced in Cisco Unified CallManager Release 5.0 to support more features.

Cisco Unified IP Phones use Skinny Client Control Protocol (SCCP), which is a lighter-weight protocol. SCCP uses a master/slave model, while H.323 is a peer-to-peer model. MGCP also follows a master/slave model.

Protocol selection depends on site-specific requirements and the installed base of equipment. For example, most remote branch locations have Cisco 2600XM, 2800, 3700, or 3800 Series routers installed. These routers support H.323 and MGCP 0.1 with Cisco IOS Release 12.2.11(T) and Cisco Unified CallManager Release 3.1 or later. For gateway configuration, MGCP might be preferred to H.323 due to simpler configuration. On the other hand, H.323 might be preferred over MGCP because of the robustness of the interfaces supported.

Simplified Message Desk Interface (SMDI) is a standard for integrating voice mail systems to PBXs or Centrex systems. Connecting to a voice mail system via SMDI and using either analog FXS or digital T1 PRI would require either SCCP or MGCP protocol because H.323 devices do not identify the specific line being used from a group of ports. Use of H.323 gateways for this purpose means the Cisco Message Interface cannot correctly correlate the SMDI information with the actual port or channel being used for an incoming call.

In addition, the Cisco Unified CallManager deployment model being used can influence gateway protocol selection. (Refer to the chapter on [IP Telephony Deployment Models](#), page 2-1.)

[Table 4-2](#) shows which gateways support a given protocol. Each of these protocols follows a slightly different methodology to provide support for the core gateway requirements. [Gateway Protocols and Core Feature Requirements](#), page 4-6, describes how each protocol provides these feature requirements.

Table 4-2 Supported Gateway Protocols and Cisco Unified Communications Gateways

Cisco Gateway	MGCP 0.1	H.323	SCCP	SIP
Cisco 3800	Yes, beginning with Cisco IOS Release 12.3.11T	Yes, beginning with Cisco IOS Release 12.3.11T	Yes, beginning with Cisco IOS Release 12.3.11T	Yes, SIP trunk
Cisco 2800	Yes, beginning with Cisco IOS Release 12.3.8T4	Yes, beginning with Cisco IOS Release 12.3.8T4	Yes, beginning with Cisco IOS Release 12.3.8T4	Yes, SIP trunk
Cisco 3700	Yes Supported with: - Analog FXS/FXO - T1 CAS (E&M Wink Start; Delay Dial only) - T1/E1 PRI	Yes	DSP farm in Cisco IOS Release 12.2.13T	Yes, SIP trunk
Communication Media Module (CMM)	Yes Supported with: - T1 CAS FXS - T1/E1 PRI - FXS	Yes	No	Yes
Catalyst 6000 WS-X6608-x1 Gateway Module and FXS Module WS-X6624	Yes Supported with: - T1 CAS E&M - T1 CAS FXS - T1/E1 PRI - FXS with WS-6624	No	No	No
VG224	Yes, FXS only. Also supports conferencing and transcoding for VG224 beginning with Cisco IOS Release 12.3(T).	Yes, FXS only	Yes, beginning with Cisco IOS Release 12.4(2)T	Yes, SIP trunk

Table 4-2 Supported Gateway Protocols and Cisco Unified Communications Gateways (continued)

Cisco Gateway	MGCP 0.1	H.323	SCCP	SIP
VG248	No	No	Yes ¹	No
Cisco ATA 188	Yes, FXS only	Yes, FXS only	Yes, FXS only	Yes, third-party SIP phone
Cisco AS5350 Cisco AS5400	No	Yes	No	Yes, SIP trunk
Cisco AS5850	No	Yes	No	Yes, SIP trunk
Cisco 5300	No	Yes	No	Yes, SIP trunk
Cisco 3640 and 3660	Yes Supported with: - Analog FXS/FXO - T1 CAS (E&M Wink Start; Delay Dial only) - T1/E1 PRI	Yes	DSP farm in Cisco IOS Release 12.2.13T	Yes, SIP trunk
Cisco 2600 and 2600XM ²	Yes Supported with: - Analog FXS/FXO - T1 CAS (E&M Wink Start; Delay Dial only) - T1/E1 PRI	Yes	DSP farm in Cisco IOS Release 12.2.13T	Yes, SIP trunk
Cisco 1751 and 1760	Yes	Yes	Yes, conferencing and transcoding	Yes, SIP trunk
VG200 ³	Yes Supported with: - Analog FXS/FXO - T1 CAS (E&M Wink Start; Delay Dial only) - T1/E1 PRI	Yes	Yes (DSP farm)	No
Cisco 7200	No	Yes	No	Yes, SIP trunk
Catalyst 4000 WS-X4604-GWY Gateway Module	Yes	Yes	No	No
Cisco ICS7750-MRP	No	Yes	No	No
Cisco ICS7750-ASI	No	Yes	No	No
DE-30+, DT-24+ ⁴	Yes	No	No	No
Cisco 827-V4 ⁴	No	Yes, supported for FXS	No	No

1. The VG248 is not a true gateway in that it uses Skinny Client Control Protocol (SCCP) instead of H.323, MGCP, or SIP.

2. For IP Telephony applications, use Cisco 2800 Series Routers.
3. The VG200 is no longer available for purchase and has been replaced by the Cisco 2800 Router. Existing models of the VG200 can still be used in an IP Telephony deployment.
4. These models have reached end of life.

**Note**

Prior to deployment, check the Cisco IOS software release notes to confirm feature or interface support.

Gateway Protocols and Core Feature Requirements

This section describes how each protocol (SCCP, H.323, MGCP, and SIP) supports the following gateway feature requirements:

- [DTMF Relay, page 4-6](#)
- [Supplementary Services, page 4-7](#)
- [Cisco Unified CallManager Redundancy, page 4-10](#)

DTMF Relay

Dual-Tone Multifrequency (DTMF) is a signaling method that uses specific pairs of frequencies within the voice band for signals. A 64 kbps pulse code modulation (PCM) voice channel can carry these signals without difficulty. However, when using a low bite-rate codec for voice compression, the potential exists for DTMF signal loss or distortion. An out-of-band signaling method for carrying DTMF tones across a Voice over IP (VoIP) infrastructure provides an elegant solution for these codec-induced symptoms.

SCCP Gateways

The SCCP gateways, such as the Cisco VG248, carry DTMF signals out-of-band using Transmission Control Protocol (TCP) port 2002. Out-of-band DTMF is the default gateway configuration mode for the VG248.

H.323 Gateways

The H.323 gateways, such as the Cisco 3700 series products, can communicate with Cisco Unified CallManager using the enhanced H.245 capability for exchanging DTMF signals out-of-band. The following is an example out-of-band DTMF configuration on a Cisco IOS gateway:

```
dial-peer voice 100 voip
destination-pattern 555...
session target ipv4:10.1.1.1
CODEC g729ar8
dtmf-relay h245-alphanumeric
preference 0
```

MGCP Gateway

The Cisco IOS-based VG224, 2600XM, 2800, 3700, and 3800 platforms use MGCP for Cisco Unified CallManager communication. Within the MGCP protocol is the concept of *packages*. The MGCP gateway loads the DTMF package upon start-up. The MGCP gateway sends *symbols* over the

control channel to represent any DTMF tones it receives. Cisco Unified CallManager then interprets these signals and passes on the DTMF signals, out-of-band, to the signaling endpoint. The global configuration command for DTMF relay is:

```
mgcp dtmf-relay CODEC all mode out-of-band
```

You must enter additional configuration parameters in the Cisco Unified CallManager MGCP gateway configuration interface.

The Catalyst 6000, DE-30+, and DT-24+ all support MGCP with Cisco Unified CallManager Release 3.1 and later. DTMF relay is enabled by default and does not need additional configuration.

SIP Gateway

The Cisco IOS-based VG224, 2600XM, 2800, 3700, 3800 platforms can use SIP for Cisco Unified CallManager communication. They support various methods for DTMF, but only the following two methods can be used to communicate with Cisco Unified CallManager:

- Named Telephony Events (NTE), or RFC 2833
- Unsolicited SIP Notify (UN)

The following example shows a configuration for NTE:

```
dial-peer voice 100 voip
destination-pattern 555...
session target ipv4:10.1.1.1
session protocol sipv2
dtmf-relay rtp-nte
```

The following example shows a configuration for UN:

```
dial-peer voice 100 voip
destination-pattern 555...
session target ipv4:10.1.1.1
session protocol sipv2
dtmf-relay sip-notify
```

For more details on DTMF method selection, see the chapter on [Media Resources](#), page 6-1.

Supplementary Services

Supplementary services provide user functions such as hold, transfer, and conferencing. These are considered fundamental requirements of any voice installation. Each gateway evaluated for use in an IP telephony network should provide support for supplementary services natively, without the use of a software media termination point (MTP).

SCCP Gateways

The Cisco VG224, VG248, and ATA 188 gateways provide full supplementary service support. The SCCP gateways use the Gateway-to-Cisco Unified CallManager signaling channel and SCCP to exchange call control parameters.

H.323 Gateways

H.323v2 implements Open/Close LogicalChannel and the emptyCapabilitySet features. The use of H.323v2 by H.323 gateways, beginning in Cisco IOS Release 12.0(7)T and Cisco Unified CallManager Release 3.0 and later, eliminates the requirement for an MTP to provide supplementary services. With


Cisco Unified CallManager Release 3.1 and later, a transcoder is allocated dynamically only if required during a call to provide access to G.711-only devices while still maintaining a G.729 stream across the WAN. Full support for H.323v2 is available in Cisco IOS Release 12.1.1T.

Once an H.323v2 call is set up between a Cisco IOS gateway and an IP phone, using the Cisco Unified CallManager as an H.323 proxy, the IP phone can request to modify the bearer connection. Because the Real-Time Transport Protocol (RTP) stream is directly connected to the IP phone from the Cisco IOS gateway, a supported voice codec can be negotiated.

Figure 4-1 and the following steps illustrate a call transfer between two IP phones:

1. If IP Phone 1 wishes to transfer the call from the Cisco IOS gateway to Phone 2, it issues a transfer request to Cisco Unified CallManager using SCCP.
2. Cisco Unified CallManager translates this request into an H.323v2 CloseLogicalChannel request to the Cisco IOS gateway for the appropriate SessionID.
3. The Cisco IOS gateway closes the RTP channel to Phone 1.
4. Cisco Unified CallManager issues a request to Phone 2, using SCCP, to set up an RTP connection to the Cisco IOS gateway. At the same time, Cisco Unified CallManager issues an OpenLogicalChannel request to the Cisco IOS gateway with the new destination parameters, but using the same SessionID.
5. After the Cisco IOS gateway acknowledges the request, an RTP voice bearer channel is established between Phone 2 and the Cisco IOS gateway.


Figure 4-1 H.323 Gateway Supplementary Service Support



MGCP Gateway

The MGCP gateways provide full support for the hold, transfer, and conference features through the MGCP protocol. Because MGCP is a master/slave protocol with Cisco Unified CallManager controlling all session intelligence, Cisco Unified CallManager can easily manipulate MGCP gateway voice connections. If an IP telephony endpoint (for example, an IP phone) needs to modify the session (for example, transfer the call to another endpoint), the endpoint would notify Cisco Unified CallManager using SCCP. Cisco Unified CallManager then informs the MGCP gateway, using the MGCP User Datagram Protocol (UDP) control connection, to terminate the current RTP stream associated with the Session ID and to start a new media session with the new endpoint information. [Figure 4-2](#) illustrates the protocols exchanged between the MGCP gateway, endpoints, and Cisco Unified CallManager.

Figure 4-2 MGCP Gateway Supplementary Service Support



77300

SIP Gateway

The Cisco Unified CallManager SIP trunk interface to Cisco IOS SIP gateways supports supplementary services such as hold, blind transfer, and attended transfer. The support for supplementary services is achieved via SIP methods such as INVITE and REFER. For more details, refer to the following documentation:

- *Cisco Unified CallManager 5.0 System Guide*, available at <http://www.cisco.com>
- *Cisco IOS SIP Configuration Guide*, available at http://www.cisco.com/en/US/products/ps6441/products_installation_and_configuration_guides_list.html

Cisco Unified CallManager Redundancy

An integral piece of the IP telephony architecture is the provisioning of low-cost, distributed PC-based systems to replace expensive and proprietary legacy PBX systems. This distributed design lends itself to the robust fault tolerant architecture of clustered Cisco Unified CallManagers. Even in its most simplistic form (a two-system cluster), a secondary Cisco Unified CallManager should be able to pick up control of all gateways initially managed by the primary Cisco Unified CallManager.

SCCP Gateways

Upon boot-up, the Cisco VG224, VG248, and ATA 188 gateways are provisioned with Cisco Unified CallManager server information. When these gateways initialize, a list of Cisco Unified CallManagers is downloaded to the gateways. This list is prioritized into a primary Cisco Unified CallManager and secondary Cisco Unified CallManager. In the event that the primary Cisco Unified CallManager becomes unreachable, the gateway registers with the secondary Cisco Unified CallManager.

H.323 Gateways

Using several enhancements to the **dial-peer** and **voice class** command sets in Cisco IOS Release 12.1(2)T, Cisco H.323 gateways support redundant Cisco Unified CallManagers. A new command, **H.225 tcp timeout <seconds>**, has been added. This command tracks the time it takes for the H.323 gateway to establish an H.225 control connection for H.323 call setup. If the H.323 gateway cannot establish an H.225 connection to the primary Cisco Unified CallManager, it tries a second Cisco Unified CallManager defined in another **dial-peer** statement. The H.323 gateway shifts to the **dial-peer** statement with next highest **preference** setting. The following commands allow you to configure Cisco Unified CallManager redundancy for a H.323 gateway:

```
dial-peer voice 101 voip
  destination-pattern 1111
  session target ipv4:10.1.1.101
  preference 0
  voice class h323 1
dial-peer voice 102 voip
  destination-pattern 1111
  session target ipv4:10.1.1.102
  preference 1
  voice class h323 1
voice class h323 1
  h225 tcp timeout <1-30 sec>
```

MGCP Gateway

MGCP gateways also have the ability to fail over to a secondary Cisco Unified CallManager in the event of communication loss with the primary Cisco Unified CallManager. When the failover occurs, active calls are preserved.

Within the MGCP gateway configuration file, the primary Cisco Unified CallManager is identified using the **call-agent <hostname>** command, and a list of secondary Cisco Unified CallManager is added using the **ccm-manager redundant-host** command. Keepalives with the primary Cisco Unified CallManager are through the MGCP application-level keepalive mechanism, whereby the MGCP gateway sends an empty MGCP notify (NTFY) message to Cisco Unified CallManager and waits for an acknowledgement. Keepalive with the backup Cisco Unified CallManagers is through the TCP keepalive mechanism.

If the primary Cisco Unified CallManager becomes available at a later time, the MGCP gateway can “re-home,” or switch back to the original Cisco Unified CallManager. This re-homing can occur either immediately, after a configurable amount of time, or only when all connected sessions have been released. This is enabled through the following global configuration commands:

```
ccm-manager redundant-host <hostname1 | ipaddress1 > <hostname2 | ipaddress2>
[no] call-manager redundancy switchback [immediate|graceful|delay <delay_time>]
```

SIP Gateway

Redundancy with Cisco IOS SIP gateways can be achieved similarly to H.323. If the SIP gateway cannot establish a connection to the primary Cisco Unified CallManager, it tries a second Cisco Unified CallManager defined under another dial-peer statement with a higher preference.

By default the Cisco IOS SIP gateway transmits the SIP INVITE request 6 times to the Cisco Unified CallManager IP address configured under the dial-peer. If the SIP gateway does not receive a response from that Cisco Unified CallManager, it will try to contact the Cisco Unified CallManager configured under the other dial-peer with a higher preference.

Cisco IOS SIP gateways wait for the SIP 100 response to an INVITE for a period of 500 ms. By default, it can take up to 3 seconds for the Cisco IOS SIP gateway to reach the backup Cisco Unified CallManager. You can change the SIP INVITE retry attempts under the **sip-ua** configuration by using the command **retry invite <number>**. You can also change the period that the Cisco IOS SIP gateway waits for a SIP 100 response to a SIP INVITE request by using the command **timers trying <time>** under the **sip-ua** configuration.

One other way to speed up the failover to the backup Cisco Unified CallManager is to configure the command **monitor probe icmp-ping** under the **dial-peer** statement. If Cisco Unified CallManager does not respond to an Internet Control Message Protocol (ICMP) echo message (ping), the dial-peer will be shut down. This command is useful only when the Cisco Unified CallManager is not reachable. ICMP echo messages are sent every 10 seconds.

The following commands enable you to configure Cisco Unified CallManager redundancy on a Cisco IOS SIP gateway:

```
sip-ua
  retry invite <number>
  timers trying <time>

dial-peer voice 101 voip
  destination-pattern 2...
  session target ipv4:10.1.1.101
  preference 0
  monitor probe icmp-ping
  session protocol sipv2
```

```
dial-peer voice 102 voip
destination-pattern 2...
session target ipv4:10.1.1.102
preference 1
monitor probe icmp-ping
session protocol sipv2
```

Call Survivability

Prior to Cisco Unified CallManager 4.2, call survivability was available only with MGCP gateways. If the signaling component of a call disappeared, the media connection was preserved until call termination, thereby allowing completion of the call.

Cisco Unified CallManager 4.2 introduces a new H.323 feature called *quiet clear*. In addition, Cisco IOS Release 12.4(4)XC introduces H.323 VoIP call preservation enhancements for WAN link failures. Both Cisco Unified CallManager and the H.323 gateway must be configured appropriately in order to allow call survivability.

For details on configuring the H.323 gateway, refer to the *Cisco IOS H.323 Configuration Guide*, available at

<http://www.cisco.com>

Site-Specific Gateway Requirements

Each IP Telephony implementation has its own site-specific requirements. The following questions can help you with IP Telephony gateway selection:

- Is the PSTN (or PBX) access analog or digital?
- What type of analog (FXO, FXS, E&M, DID, CAMA) or digital (T1, E1, CAS, CCS) interface is required for the PSTN or PBX?
- If the PSTN access is digital, what type of signaling is required (T1 CAS, Q.931 PRI, E1 CAS, or R2)?
- What type of signaling does the PBX currently use?
 - FXO or FXS: loop start or ground start
 - E&M: wink-start, delay-start, or immediate-start
 - E&M: type I, II, III, IV, or V
 - T1: CAS, Q.931 PRI (User-Side or Network-Side), QSIG, DPNSS, or Proprietary d-channel (CCS) signaling
 - E1: CAS, R2, Q.931 PRI (User-Side or Network-Side), QSIG, DPNSS, Proprietary d-channel (CCS) signaling
- What type of framing (SF, ESF, or G.704) and line encoding (B8ZS, AMI, CRC-4, or HDB3) does the PBX currently use?
- Does the PBX require passing proprietary signaling? If so, which time slot is the signaling passed on, and is it HDLC-framed?
- What is the required capacity of the gateway; that is, how many channels are required? (Typically, if 12 or more voice channels are required, then digital is more cost effective than an analog solution.)

- Is Direct Inward Dialing (DID) required? If so, specify analog or digital.
- Is Calling Line ID (CLID) needed?
- Is Calling Name needed?
- What types of fax and modem support are required?
- What types of voice compression are required?
- What types of supplementary services are required?
- Will the PBX provide clocking, or will it expect the Cisco gateway to provide clocking?
- Is rack space available for all needed gateways, routers, and switches?

**Note**

Direct Inward Dial (DID) refers to a private branch exchange (PBX) or Centrex feature that permits external calls to be placed directly to a station line without use of an operator.

**Note**

Calling Line Identification (CLI, CLID, or ANI) refers to a service available on digital phone networks to display the calling number to the called party. The central office equipment identifies the phone number of the caller, enabling information about the caller to be sent along with the call itself. CLID is synonymous with Automatic Number Identification (ANI).

Cisco Unified Communications gateways are able to inter-operate with most major PBX vendors, and they are EIA/TIA-464B compliant.

The site-specific and core gateway requirements are a good start to help narrow the possible choices. Once you have defined the required features, you can make a gateway selection for each of the pertinent configurations, whether they are single-site enterprise deployments of various sizes and complexities or multisite enterprise deployments.

The following tables summarize the features and interface types supported by the various Cisco gateway models.

**Note**

In the following tables, the Cisco IOS and Cisco Unified CallManager release numbers refer to the minimum release that can support the listed feature on a particular gateway platform. For specific recommendations about the preferred software release for each hardware platform, refer to the documentation at

http://www.cisco.com/en/US/products/sw/voicesw/ps556/products_device_support_tables_list.html

Cisco Analog Gateways

Table 4-3 lists supported interface types for Cisco analog gateways using H.323 or Session Initiation Protocol (SIP), and Table 4-4 lists supported interface types for Cisco analog gateways using Media Gateway Control Protocol (MGCP).

Table 4-3 Supported Analog H.323 and SIP Features

Cisco Gateway	Interface Type					
	FXS	FXO	E&M	FXO, Battery Reversal	Analog DID	CAMA 911
3800 Series	Yes	Yes	Yes	Yes	Yes	Yes
2800 Series	Yes	Yes	Yes	Yes	Yes	Yes

Table 4-3 Supported Analog H.323 and SIP Features (continued)

Cisco Gateway	Interface Type					
	FXS	FXO	E&M	FXO, Battery Reversal	Analog DID	CAMA 911
3700 Series	Yes	Yes	Yes	Yes	Yes	Yes
Communication Media Module (CMM) 24FXS	Yes	N/A	N/A	N/A	N/A	N/A
CMM-6T1/E1	N/A	N/A	N/A	N/A	N/A	N/A
6608 and 6624	N/A	N/A	N/A	N/A	N/A	N/A
VG224	Yes	N/A	N/A	N/A	N/A	N/A
VG248	No	No	No	No	No	No
Analog Telephone Adapter (ATA)	Yes	No	No	No	No	No
3600 Series	Yes	Yes	Yes	Yes	Yes	12.2.11T
2600 Series	Yes	Yes	Yes	Yes	Yes	12.2.11T
1751 and 1760	Yes	Yes	Yes	Yes	Yes	Yes
VG200	Yes	Yes	Yes	No	Yes	No
7x00 family	N/A	N/A	N/A	N/A	N/A	N/A
ICS 7750	Yes	Yes	Yes	Yes	Yes	No
Catalyst 4000 Access Gateway Module (AGM)	Yes	Yes	No	No	No	No
827-4V ¹	Yes	No	No	No	No	No

1. This model has reached end of life.

Table 4-4 Supported Analog MGCP Features

Cisco Gateway	Interface Type					
	FXS	FXO	E&M	FXO, Battery Reversal	Analog DID	CAMA 911
3800 Series	Yes	Yes	No	Yes	No	No
2800 Series	Yes	Yes	No	Yes	No	No
3700 Series	Yes	Yes	No	Yes	No	No
Communication Media Module (CMM) 24FXS	Yes	N/A	N/A	N/A	N/A	N/A
CMM-6T1/E1	N/A	N/A	N/A	N/A	N/A	N/A
6608 and 6624	Yes	No	No	No	No	No
VG224	Yes	No	No	No	No	No
VG248	No	No	No	No	No	No
Analog Telephone Adapter (ATA)	Yes	N/A	N/A	N/A	N/A	N/A

Table 4-4 Supported Analog MGCP Features (continued)

Cisco Gateway	Interface Type					
	FXS	FXO	E&M	FXO, Battery Reversal	Analog DID	CAMA 911
3600 Series	Yes	Yes	No	Yes	No	No
2600 Series	Yes	Yes	No	Yes	No	No
1751 and 1760	Yes	Yes	No	Yes	No	No
VG200	Yes	Yes	No	Yes	No	No
7x00 family	N/A	N/A	N/A	N/A	N/A	N/A
ICS 7750	Yes	Yes	No	No	No	No
Catalyst 4000 Access Gateway Module (AGM)	Yes	Yes	No	No	No	No
827-4V ¹	No	No	N/A	N/A	N/A	N/A

1. This model has reached end of life.

Cisco Digital Gateways

Table 4-5 through Table 4-8 list supported interface types for Cisco digital gateways using H.323 or Session Initiation Protocol (SIP). Table 4-9 lists supported interface types for Cisco digital gateways using Media Gateway Control Protocol (MGCP).

Table 4-5 Supported Digital H.323 and SIP Features for BRI, T1 CAS, T1 FGB, T1 FGD, and T1 QSIG

Cisco Gateway	Interface Type							
	BRI (TE, User side)	BRI (NT, Network side)	BRI QSIG (Net3)	BRI Phones	T1 CAS (Robbed bit)	T1 FGB	T1 FGD	T1 QSIG
3800 Series	Yes	Yes	Yes	No	Yes	No	Yes	Yes
2800 Series	Yes	Yes	Yes	No	Yes	No	Yes	Yes
3700 Series	Yes	Yes	Yes	No	Yes	No	Yes	Yes
Communication Media Module (CMM) 24FXS	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
CMM-6T1/E1	N/A	N/A	N/A	N/A	Yes	No	No	Yes
6608 and 6624	N/A	N/A	N/A	N/A	No	No	No	No
VG224	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
VG248	No	No	No	No	No	No	No	No
Analog Telephone Adapter (ATA)	No	No	No	No	No	No	No	No
3600 Series	Yes	Yes	Yes	No	Yes	No	Yes	Yes
2600 Series	Yes	Yes	Yes	No	Yes	No	Yes	Yes
1751 and 1760	No	Yes	Yes	No	Yes	No	No	Yes
VG200	Yes	Yes	No	No	Yes	No	Yes	No
7x00 family	N/A	N/A	N/A	N/A	Yes	No	Yes	Yes

Table 4-5 Supported Digital H.323 and SIP Features for BRI, T1 CAS, T1 FGB, T1 FGD, and T1 QSIG (continued)

Cisco Gateway	Interface Type							
	BRI (TE, User side)	BRI (NT, Network side)	BRI QSIG (Net3)	BRI Phones	T1 CAS (Robbed bit)	T1 FGB	T1 FGD	T1 QSIG
ICS 7750	Yes	Yes	No	No	Yes	No	Yes	No
Catalyst 4000 Access Gateway Module (AGM)	Yes	No	Yes	No	Yes	No	Yes	Yes
827-4V ¹	No	No	No	No	No	No	No	No

1. This model has reached end of life.

Table 4-6 Supported Digital H.323 and SIP Features for T1 PRI SL-1, 4ESS, and 5ESS

Cisco Gateway	Interface Type					
	T1 PRI (User, DMS-100)	T1 PRI (Network, SL-1)	T1 PRI (User, 4ESS)	T1 PRI (Network, 4ESS)	T1 PRI (User, 5ESS)	T1 PRI (Network, 5ESS)
3800 Series	Yes	Future	Yes	Yes	Yes	Yes
2800 Series	Yes	Future	Yes	Yes	Yes	Yes
3700 Series	Yes	Future	Yes	Future	Yes	Future
Communication Media Module (CMM) 24FXS	N/A	N/A	N/A	N/A	N/A	N/A
CMM-6T1/E1	Yes	Yes	Yes	Yes	Yes	Yes
6608 and 6624	No	No	No	No	No	No
VG224	N/A	N/A	N/A	N/A	N/A	N/A
VG248	No	No	No	No	No	No
Analog Telephone Adapter (ATA)	No	No	No	No	No	No
3600 Series	Yes	Future	Yes	Yes	Yes	Yes
2600 Series	Yes	Future	Yes	Yes	Yes	Yes
1751 and 1760	Yes	Future	Yes	Future	Yes	Future
VG200	Yes	No	Yes	No	Yes	No
7x00 family	Yes	Future	Yes	Future	Yes	Future
ICS 7750	Yes	No	Yes	No	Yes	No
Catalyst 4000 Access Gateway Module (AGM)	Yes	Future	Yes	Future	Yes	Future
827-4V ¹	No	No	No	No	No	No

1. This model has reached end of life.

Table 4-7 **Supported Digital H.323 and SIP Features for T1 PRI NI2, NFAS, and Network Specific Facilities (NSF) Service**

Cisco Gateway	Interface Type					
	T1 PRI (User, NI2)	T1 PRI (Network, NI2)	T1 PRI NFAS (User, DMS-100)	T1 PRI NFAS (User, 4ESS)	T1 PRI NFAS (User, 5ESS)	T1 PRI (Megacom or SDN, 4ESS)
3800 Series	Yes	Yes	Yes	Yes	Yes	Yes
2800 Series	Yes	Yes	Yes	Yes	Yes	Yes
3700 Series	Yes	Yes	Yes	Yes	Yes	Yes
Communication Media Module (CMM) 24FXS	N/A	N/A	N/A	N/A	N/A	N/A
CMM-6T1/E1	Yes	Yes	Yes	Future	Future	No
6608 and 6624	No	No	No	No	No	No
VG224	N/A	N/A	N/A	N/A	N/A	N/A
VG248	No	No	No	No	No	No
Analog Telephone Adapter (ATA)	No	No	No	No	No	No
3600 Series	Yes	Yes	Yes	Yes	Yes	Yes
2600 Series	Yes	Yes	Yes	Yes	Yes	Yes
1751 and 1760	Yes	Yes	No	No	No	No
VG200	Yes	Yes	No	No	No	No
7x00 family	Yes	Yes	No	No	No	No
ICS 7750	Yes	Yes	Yes	Yes	Yes	No
Catalyst 4000 Access Gateway Module (AGM)	Yes	Yes	Future	Future	Future	Future
827-4V ¹	No	No	No	No	No	No

1. This model has reached end of life.

Table 4-8 **Supported Digital H.323 and SIP Features for E1 and J1**

Cisco Gateway	Interface Type						
	E1 CAS	E1 MELCAS	E1 R2	E1 PRI (User side, Net5)	E1 PRI (Network side, Net5)	E1 QSIG	J1
3800 Series	Yes	Yes	Yes	Yes	Yes	Yes	Yes
2800 Series	Yes	Yes	Yes	Yes	Yes	Yes	Yes
3700 Series	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Communication Media Module (CMM) 24FXS	N/A	N/A	N/A	N/A	N/A	N/A	N/A
CMM-6T1/E1	No	No	Yes	Yes	Yes	Yes	N/A
6608 and 6624	No	No	No	No	No	No	No

Table 4-8 Supported Digital H.323 and SIP Features for E1 and J1 (continued)

Cisco Gateway	Interface Type						
	E1 CAS	E1 MELCAS	E1 R2	E1 PRI (User side, Net5)	E1 PRI (Network side, Net5)	E1 QSIG	J1
VG224	N/A	N/A	N/A	N/A	N/A	N/A	N/A
VG248	No	No	No	No	No	No	No
Analog Telephone Adapter (ATA)	No	No	No	No	No	No	No
3600 Series	Yes	Yes	Yes	Yes	Yes	Yes	Yes
2600 Series	Yes	Yes	Yes	Yes	Yes	Yes	Yes
1751 and 1760	No	No	Yes	Yes	Yes	Yes	No
VG200	No	Yes	Yes	Yes	Yes	No	Yes
7x00 family	Yes	No	Yes	Yes	Yes	Yes	No
ICS 7750	No	No	Yes	Yes	Yes	No	No
Catalyst 4000 Access Gateway Module (AGM)	No	No	Yes	Yes	Yes	Yes	No
827-4V ¹	No	No	No	No	No	No	No

1. This model has reached end of life.

Table 4-9 Supported Digital MGCP Features

Cisco Gateway	Interface Type					
	BRI ¹	T1 CAS (E&M)	T1 PRI	T1 QSIG	E1 PRI	E1 QSIG
3800 Series	12.4(2)T	Yes ²	Yes ²	Yes ²	Yes ²	Yes ²
2800 Series	12.4(2)T	Yes ²	Yes ²	Yes ²	Yes ²	Yes ²
3700 Series	12.4(2)T	Yes ²	Yes ²	Yes ²	Yes ²	Yes ²
Communication Media Module (CMM) 24FXS	N/A	N/A	N/A	N/A	N/A	N/A
CMM-6T1/E1	N/A	Yes	Yes	Yes	Yes	Yes
6608	N/A	Yes	Yes	Yes	Yes	Yes
6624	N/A	N/A	N/A	N/A	N/A	N/A
VG224	N/A	N/A	N/A	N/A	N/A	N/A
VG248	N/A	N/A	N/A	N/A	N/A	N/A
Analog Telephone Adapter (ATA)	N/A	N/A	N/A	N/A	N/A	N/A
3600 Series	12.4(2)T	Yes ²	Yes ²	Yes ²	Yes ²	Yes ²
2600 Series	12.4(2)T	Yes ²	Yes ²	Yes ²	Yes ²	Yes ²
1751 and 1760	12.3(14)T	Yes	Yes	Yes	Yes	Yes
VG200	No	Yes	Yes	Yes	Yes	Yes

Table 4-9 Supported Digital MGCP Features (continued)

Cisco Gateway	Interface Type					
	BRI ¹	T1 CAS (E&M)	T1 PRI	T1 QSIG	E1 PRI	E1 QSIG
7x00 family	N/A	No	No	No	No	No
ICS 7750	12.3.7T	Yes	Yes	Yes	Yes	Yes
Catalyst 4000 Access Gateway Module (AGM)	No	Yes	Yes	Yes	Yes	Yes
827-4V ³	N/A	N/A	N/A	N/A	N/A	N/A

1. Cisco IOS Release 12.4(2)T supports BRI MGCP with the following hardware: NM-HDV2, NM-HD-XX and on-board H-WIC slots. BRI MGCP is also supported on older Cisco IOS releases with NM-1V/2V hardware.
2. AIM-VOICE-30 modules require Cisco IOS Release 12.2.13T to support MGCP.
3. This model has reached end of life.

QSIG Support

QSIG is a suite of international standards designed to provide flexibility in connecting PBX equipment to a corporate network. Among its other features, QSIG provides an open, standards-based method for interconnecting PBX equipment from different vendors.

ECMA QSIG is currently supported in H.323 gateways in PBX-to-PBX mode. The H.323 gateways provide full QSIG feature transparency for QSIG information elements. Basic call setup and teardown are supported using H.323 QSIG gateways, as summarized in [Table 4-10](#).

Table 4-10 QSIG Support on H.323 Gateways

Platform	Media	Minimum Cisco IOS Software Release Required
Cisco 3800	BRI and T1/E1 QSIG	12.3.11T
Cisco 2800 Series	BRI and T1/E1 QSIG	12.3.8T4
Cisco 3700	T1/E1 QSIG	12.2.8T
Cisco AS5350	T1/E1	12.2.2T
Cisco AS5400		
Cisco 5300	T1/E1	12.0.7T
Cisco 2600 and 3600 Series	BRI and T1/E1 QSIG	12.1.2T
Cisco 1751 and 1760	BRI	12.2(8)YH
	T1/E1 QSIG	12.2(4)YB
Cisco 7200	T1/E1 QSIG	12.1.2T

For more information on QSIG support on Cisco IOS gateways, refer to

http://www.cisco.com/en/US/products/sw/iosswrel/ps1834/products_feature_guide09186a00800800b0.html#xtocid116542

Prior to Cisco Unified CallManager Release 3.3, basic PRI functionality is all that is provided whenever a PBX is connected to a gateway using QSIG via H.323 and calls are made between phones on the PBX and IP phones attached to the Cisco Unified CallManager. This basic functionality, which includes only the Calling Line Identifier (CLID) and Direct Inward Dialed (DID) number, is provided by the gateway terminating the QSIG protocol rather than by Cisco Unified CallManager.

For Cisco Unified CallManager to support QSIG functionality, QSIG must be back-hauled directly to Cisco Unified CallManager. This support is provided in Cisco Unified CallManager Release 3.3 and later, in conjunction with MGCP gateways such as the Catalyst 6608, 2600XM Series, and 3640/60 Series.

Fax and Modem Support

This section describes the fax and modem support available with Cisco Unified CallManager and Cisco voice gateways. This section first presents brief overviews of fax and modem support on Cisco voice gateways, followed by a listing of supported platforms and example configuration files.

Gateway Support for Fax Pass-Through and Cisco Fax Relay

Fax over IP enables interoperability of traditional analog fax machines with IP Telephony networks. The fax image is converted from an analog signal and is carried as digital data over the packet network.

In its original form, fax data is digital. However, to transmit across a traditional PSTN, it is modulated and converted to analog. Fax over IP reverses this analog conversion, transmitting digital data over the packet network and then reconverts the digital data to analog for the receiving fax machine.

Most Cisco voice gateways currently support two methods to transmit fax traffic across the IP network:

- Cisco Fax Relay — In fax relay mode, the gateways terminate the T.30 or T.38 fax signaling.
- Fax Pass-Through — In fax pass-through mode, the gateways do not distinguish a fax call from a voice call.

Cisco Fax Relay mode is the preferred method to transmit fax traffic. However, if a specific gateway does not support Cisco fax relay, it supports fax pass-through.

Best Practices

The following recommendations and guidelines can assist you in best implementing fax support on Cisco voice gateways:

- When using QoS, make every effort to minimize the following:
 - Packet loss
 - Delay
 - Delay variation (jitter)

For detailed information about implementing QoS in a Cisco Unified Communications network, refer to the *Cisco Network Infrastructure Enterprise Quality of Service Design* guide, available at

<http://www.cisco.com/go/designzone>

- The following tips can help ensure the integrity of the fax calls:
 - Use call admission control (CAC) to ensure that calls are not admitted if they exceed the specified total bandwidth limit.

- Disable call waiting on all dedicated modem and fax ports.
- For best performance, verify that you have Cisco Fax Relay on both the originating and terminating gateways. If two Cisco IOS gateways have differing transports, they will negotiate to use Cisco Fax Relay.

The only non-IOS gateway that does not support Cisco Fax Relay is the Cisco Digital Access DT-24/DE-30+. If you connect this gateway to a Cisco IOS gateway, you should configure both gateways to use fax pass-through mode.

- Ensure that constant packet delay on the network does not exceed 1 second and that delay variation (jitter) does not exceed 240 milliseconds.
- To improve performance in networks with a high frequency of out-of-order packet arrival, disable Error Correction Mode (ECM) on the fax machines.
- Most fax machines appear to accept packet drop in the range of 0.4% to 0.6% without slowing down to the next speed. However, in a network with packet drop in the range of 0.8% to 1%, you should disable ECM.
- You can disable ECM on the gateway itself rather than disabling it on multiple fax machines. However, if packet drops occur, the fax image quality might deteriorate. Therefore, you should disable ECM only after considering whether you want to risk compromising image quality rather than experiencing longer call durations or dropped calls. You should also monitor and evaluate the network to identify and resolve the cause of the dropped packets.

Gateway Support for Modem Pass-Through

In general, there are two mechanisms for supporting modem sessions over an IP network using voice gateways:

- Modem pass-through
- Modem relay

Currently, both modem relay and modem pass-through are supported on Cisco voice gateways.

Modem pass-through is the transport of modem signals through a packet network using pulse code modulation (PCM) encoded packets and a G.711 codec. Modem pass-through requires the ability of the gateways to discriminate between modem signals and voice signals and take appropriate action. When the gateway detects the modem signal, it disables the following services:

- Echo cancellation (EC)
- Voice activity detection (VAD)

In modem pass-through mode, the gateways do not distinguish a modem call from a voice call. The communication between the two modems is carried in-band in its entirety over a "voice" call. The modem traffic is transparently carried over a QoS-enabled IP infrastructure, and at no point is the data demodulated within the IP network.

Modem upspeed is similar to pass-through in the sense that the modem call is carried in-band over the "voice" call. The difference is that the gateways are, to some extent, aware of the modem call when the upspeed feature is used. Although relay mechanisms are not employed, the gateways do recognize the modem tone, automatically change the "voice" codec to G.711 (the upspeed portion), and turn off VAD and echo cancellation (EC) for the duration of the call.

Currently, this upspeed feature is not supported on any Cisco IOS platform except the Cisco AS5300 via Cisco IOS Release 12.1.3T. For Cisco 2600XM, 3700, VG224, and Catalyst 4000 Access Gateway Module (AGM) platforms, the modem upspeed feature will be supported in a future Cisco IOS release. For these platforms, you can configure **no vad** on the dial peer until the modem upspeed feature becomes available.

The modem upspeed feature is also supported on the Catalyst 6000 gateway modules.

Best Practices

Observe the following recommended best practices to ensure optimum performance of modem traffic transported over an IP infrastructure:

- Ensure that the IP network is enable for Quality of Service (QoS) and that you adhere to all of the recommendations for providing QoS in the LAN, MAN, and WAN environments. Every effort should be made to minimize the following parameters:
 - Packet loss — Fax and modem traffic requires an essentially loss-free transport. A single lost packet will result in retransmissions.
 - Delay
 - Delay variation (jitter)

For more information, refer to the *Cisco Network Infrastructure Enterprise Quality of Service Design* guide, available at

<http://www.cisco.com/go/designzone>

- Use call admission control (CAC) to ensure that calls are not admitted if they exceed the specified total bandwidth limit.
- Use G.711 for all calls involving a modem. If one of the gateways does not support modem relay, modem pass-through will be negotiated (G.711 only). If modems are used, the best-practice recommendation is to use G.711 for all calls.
- Do not use the IP network to connect modems that will be used to troubleshoot or diagnose problems with the IP network. In this case, the modems used to troubleshoot the devices that compose the IP infrastructure should be connected to a plain old telephone service (POTS).
- Where possible, use a single signaling protocol and gateway family to minimize interoperability issues.
- Disable call waiting on all dedicated modem and fax ports.

V.90 Support

Currently, Cisco equipment supports only V.34 modems. Although V.90 modems will function on existing hardware, and speeds higher than V.34 speeds can be achieved, full V.90 support cannot be guaranteed.

Supported Platforms and Features

The following Cisco platforms support fax and modem features:

Analog Gateways

Cisco IOS Gateways:

- 2600XM and 2691 (with FXS)
- 2800 (with FXS)

- 3725 and 3745 (with FXS)
- 3800 (with FXS)
- VG200 (with FXS)
- VG224
- 1751 and 1760
- Communication Media Module (CMM) FXS card

Non-IOS Gateways:

- VG248
- ATA 188
- 6624

Digital Gateways

Cisco IOS Gateways:

- 2600XM and 2691
- 2800
- 3725 and 3745
- 3800
- VG200
- VG224
- 1751 and 1760
- 7200 and 7500
- AS5300, 5350, 5400, and 5850
- Communication Media Module (CMM)

Non-IOS Gateways:

- 6608



Note

Fax and modem support was tested on the above platforms using Cisco IOS Release 12.3(1) on the Cisco IOS gateways and Release 1.2.1 of the Cisco VG248 Analog Phone Gateway.

Platform Protocol Support

Common call control protocols used today in enterprise solutions include H.323, Session Initiation Protocol (SIP), Media Gateway Control Protocol (MGCP), and Skinny Client Control Protocol (SCCP). Not all Cisco voice platforms support all of these protocols or all of the fax and modem features, thus raising interoperability issues. Additional interoperability issues occur when mixing Cisco IOS gateways, such as the Cisco 2600XM or the Cisco 3700 Series, with non-IOS gateways such as the VG248. This section lists the combinations of gateways that provide support for interoperability of fax, modem, and protocol features.

At a high level, Cisco IOS Release 12.3(1) – Load 47 on the Cisco 6608 and Load 41 on the Cisco 6624 – and Release 1.2.1 on the VG248 do support interoperability of Cisco fax relay, modem pass-through, and voice features. Prior to Cisco IOS Release 12.2(11)T1, only voice and Cisco fax relay were supported between Cisco IOS and non-IOS voice platforms because incompatibility of the pass-through Named Service Event (NSE) scheme prevented modem pass-through from interoperating.

Some of the common combinations of protocols in a network include MGCP and H.323, SCCP and H.323, and SCCP and MGCP. Common voice gateways included the Cisco VG224, VG248, 2600XM, 2800, 3700, 3800, 5300, and Catalyst 6000.

Table 4-11 lists the protocol combinations that currently support fax and modem interoperability.

Table 4-11 Fax and Modem Features Supported with Various Combinations of Call Control Protocols

Protocol Combinations	Modem Relay	Modem Pass-Through	T.38 Fax Relay	Cisco Fax Relay	Fax Pass-Through
Cisco Unified CallManager using MGCP combined with Cisco Unified CallManager using H.323 or SIP	Yes	Yes	No	Yes	Yes
Cisco Unified CallManager using MGCP combined with Cisco Unified CallManager using MGCP	Yes	Yes	No	Yes	Yes
SCCP combined with Cisco Unified CallManager using H.323 or SIP	Yes	Yes	No	Yes	Yes
SCCP combined with Cisco Unified CallManager using MGCP	Yes	Yes	No	Yes	Yes
Cisco Unified CallManager using H.323 combined with H.323 or SIP	Yes	Yes	Yes	Yes	Yes
Cisco Unified CallManager using SIP combined with H.323 or SIP	Yes	Yes	Yes	Yes	Yes




Note


Cisco ATA 188, VG248 and Catalyst 6000 platforms currently do not support T.38 fax relay. When these platforms connect to Cisco AS5350 or AS5400 gateways, only fax pass-through is supported for fax applications.

Gateway Combinations and Interoperability of Features

The most frequent questions about fax and modem interoperability arise from combining a Cisco IOS gateway (such as a Cisco 2800 or 3800) with a non-IOS gateway (such as a Cisco VG248), as illustrated in Figure 4-3.

Figure 4-3 Configuration Combining a Cisco IOS and Non-IOS Gateway

The second most common source of questions about fax and modem interoperability arise in configurations using only Cisco IOS gateways, as illustrated in [Figure 4-4](#).

Figure 4-4 Configuration Using Only Cisco IOS Gateways

The answer is basically the same for both scenarios: Prior to Cisco IOS Load 47 on the 6608 and Release 1.2.1 on the VG248, only voice and Cisco fax relay are supported, while fax and modem pass-through are not supported because of NSE incompatibility. With Cisco IOS Load 47 or later on the 6608, Load 41 or later on the 6624, and Release 1.2.1 on the VG248, all three platforms can interoperate with Cisco IOS gateways for voice, Cisco fax relay, and modem pass-through, regardless of call control protocol. The NSE pass-through scheme is independent of call control protocol because it operates in the bearer path instead of the signaling path.

Feature Support Between Similar Gateways

[Table 4-12](#) lists the fax and modem features supported between gateways of the same general type, such as between the Cisco VG248 and 6608, between 2600XM and 3700, or between 2600XM and AS5300. In these cases, as long as both platforms support a given feature, those platforms will interoperate.

Table 4-12 Fax and Modem Feature Support on Gateways of the Same Type

Gateway Type	Fax Pass-Through	Cisco Fax Relay	T.38 Fax Relay	Modem Pass-Through	Modem Relay
Cisco IOS gateways	Supported	Supported (except on 5350 and 5400)	Supported	Supported	Supported (only on NM-HDV)
Non-IOS gateways	Supported	Supported (except on ATA 188)	N/A	Supported (except on ATA 188)	N/A

Gateway Configuration Examples

This section provides listings of example gateway configurations for fax and modem support.

Cisco IOS Gateway Configuration

```
H.323
!
! Cisco fax relay is ON by default
!(except for 5350/5400, where Cisco fax relay is not supported)
!
dial-peer voice 1000 voip
 destination-pattern 1T
 session target ipv4:10.10.10.1
 modem passthrough mode nse codec g711ulaw
!
!

MGCP
!
ccm-manager mgcp
mgcp
mgcp call-agent 10.10.10.1 service-type mgcp version 0.1
mgcp modem passthrough voip mode nse
mgcp fax t38 inhibit
!
dial-peer voice 100 pots
 application mgcpapp
 port 1/0/0
!
```

Cisco VG248 Configuration

Cisco VG248 (VGC10d8002407)		
Advanced settings		
Allow last good configuration (enabled)		
SRST policy (disabled)		
SRST provider ()		
Call preservation (enabled: no timeout)		
Media receive timeout (disabled)		
Busy out off hook ports (disabled)		
DTMF tone dur ----- 100ms)		
Echo cancelli	Passthrough signalling	e: use DSP)
Passthrough s	-----)	
Hook flash ti	legacy	default>)
Hook flash re	IOS mode	
Fax relay max ----- 14400 bps)		
Fax relay playout delay (default: 300)		

```

-----
|                               Cisco VG248 (VGC10d8002407)                               |
-----
| Advanced settings                                                         |
|-----|
| Allow last good configuration (enabled)                                  |
| SRST policy (disabled)                                                  |
| SRST provider ()                                                        |
| Call preservation (enabled: no timeout)                                |
| Media receive timeout (disabled)                                        |
| Busy out off hook ports (disabled)                                      |
| DTMF tone duration (default: 100ms)                                    |
| Echo cancelling policy (alternate: use DSP)                            |
| Passthrough signalling (IOS mode)                                    |
| Hook flash timer (<country default>)                                  |
| Hook flash reject period (none)                                        |
| Fax relay maximum speed (default: 14400 bps)                            |
| Fax relay playout delay (default: 300)                                  |
|-----|
-----

```

Cisco Unified CallManager Configuration for Cisco IOS Gateways

Perform the following steps in Cisco Unified CallManager to configure it for the Cisco IOS gateways (such as the Cisco 6608 and 6624).

-
- Step 1** In Cisco Unified CallManager Administration, select **Device > Gateway** to display the **Find/List Gateways** window.
 - Step 2** Search for the gateway you want to modify (if it already exists), or click on **Add a New Gateway** to add a new gateway to the Cisco Unified CallManager database.
 - Step 3** After selecting the appropriate type of gateway (for example, Cisco Catalyst 6000), click on **Fax Relay Enable** to enable Cisco fax relay.
 - Step 4** Using the **NSE Type** drop-down list box, select **IOS Gateways** for modem pass-through.
 - Step 5** Click **Update** to save the changes.
 - Step 6** Reset the gateway to apply the changes.
-

This configuration supports voice, Cisco fax relay, and modem pass-through between Cisco VG248, 6608, 6624, and IOS gateways, with the exception of Cisco AS5350 and AS540 gateways (which do not support Cisco fax relay). The configuration also supports a V.34 modem connection in pass-through mode. V.90 modem connections are not guaranteed but are possible, depending on amount of network jitter and clock sync.

Clock Sourcing for Fax and Modem Pass-Through

The clock signal plays a critical role in enabling fax and modem pass-through to work correctly. The gateway clock must synchronize with the PSTN clock, where Stratum clocking is provided. Without this clock synchronization, fax and (especially) modem pass-through will not work. To synchronize the clocks correctly, enter the following configuration in the T1 controller. (In this example, the T1 controller is the voice gateway that connects to the PSTN.)

```

!
controller T1 0
  framing esf
  linecode b8zs
  clock source line
  channel-group 1 timeslots 1-24 speed 64
!

```

Also enter this configuration in all other interfaces connected to the PSTN.

T.38 Fax Relay

T.38 fax relay is not supported on Cisco ATA 188, VG248, 6608, and 6624 gateways, but it is supported on most of the high-performance Cisco IOS voice platforms such as the Cisco 2800 and 3800 Series Routers. When operating in either H.323 or SIP mode, these platforms do not support MGCP.

You can configure T.38 fax relay in any of the following ways:

- [Loose Gateway Controlled with Named Service Event \(NSE\), page 4-28](#)
- [Gateway Controlled with Capability Exchange Through H.245 or Session Description Protocol \(SDP\), page 4-29](#)
- [Call-Agent-Controlled T.38 with H.323 Annex D, page 4-30](#)

Loose Gateway Controlled with Named Service Event (NSE)

This configuration uses static T.38 configuration on the dial-peer, as illustrated in the following Cisco IOS gateway configuration example:

H.323

```

!
dial-peer voice 1000 voip
  destination-pattern 1T
  session target ipv4:10.10.10.1
  modem passthrough mode nse codec g711ulaw
  fax protocol t38
!
!

```

MGCP

```

!
ccm-manage mgcp
mgcp
mgcp call-agent 10.10.10.1 service-type mgcp version 0.1
mgcp modem passthrough voip mode nse
no mgcp fax t38 inhibit
!
dial-peer voice 100 pots
  application mgcpapp
  port 1/0/0
!
!

```

Gateway Controlled with Capability Exchange Through H.245 or Session Description Protocol (SDP)

The following characteristics apply to this method of configuring T.38 fax relay:

- T.38 capability is exchanged between gateways. A Named Service Event (NSE) message is sent on the RTP stream from the terminating gateway to signal the originating gateway to switch to T.38 fax relay upon detection of a fax tone. Because the NSE message is sent on the RTP stream, it is transparent to call control signaling.
- Cisco Unified CallManager cannot support this capability exchange with MGCP. Therefore, you must use a configuration command to force T.38 fax relay even though T.38 capability is not exchanged.
- There are three fallback mechanism to choose from:
 - Cisco fax relay (default)
 - Fax pass-through
 - None

The following example illustrates this type of configuration:

H.323

```
!
dial-peer voice 1000 voip
 destination-pattern 1T
 session target ipv4:10.10.10.1
 modem passthrough mode nse codec g711ulaw
!
! To enable T.38 fax relay and fall back to Cisco fax relay when
! T.38 fax negotiation fails. This is the default case.
fax protocol t38 fallback cisco
!
dial-peer voice 1001 voip
 destination-pattern 2T
 session target ipv4:10.10.10.2
 modem passthrough mode nse codec g711ulaw
!
! To enable T.38 fax relay and fall back to fax passthrough when
! T.38 fax negotiation fails.
fax protocol t38 nse fallback pass-through
!
dial-peer voice 1002 voip
 destination-pattern 3T
 session target ipv4:10.10.10.3
 modem passthrough mode nse codec g711ulaw
!
! This CLI is needed when talking to MGCP endpoint where CA/GK
! doesn't support T.38 fax relay such as CCM.
fax protocol t38 nse force fallback none
!
!
```

MGCP

```
!
ccm-manage mgcp
mgcp
mgcp call-agent 10.10.10.1 service-type mgcp version 0.1
mgcp modem passthrough voip mode nse
no mgcp fax t38 inhibit
!
```

```

! This CLI is needed when CA doesn't support T.38 fax relay
mgcp fax t38 gateway force
!
dial-peer voice 100 pots
  application mgcpapp
  port 1/0/0
!
!

```

In topologies that employ the Cisco VG248 and 6608 or 6624, use the following Cisco IOS commands:

```

fax protocol t38 [nse [force]] fallback [cisco | none]
modem passthrough nse codec {g711ulaw|g711alaw}

```

These two commands enable Cisco IOS gateways to interoperate with the VG248 for Cisco fax relay and modem pass-through as well as with other Cisco IOS gateways for T.38 fax relay and modem pass-through.

Call-Agent-Controlled T.38 with H.323 Annex D

The following characteristics apply to this method of configuring T.38 fax relay:

- The call control agent (for example, Cisco Unified CallManager) controls the T.38 fax relay, and the gateways operate in passive mode.
- No NSE messages are sent from gateway to gateway.
- In this type of configuration, the T.38 fax relay is *not* transparent to the call control protocol. The call agent performs the protocol translation between H.323 and SIP.
- This method of configuring T.38 fax relay is available with Cisco IOS Release 12.3(1). The Cisco BTS 10200 Softswitch also supports this method.
- The Cisco Voice Media Streaming Application does not support T.38, but Cisco IOS media termination points (MTPs) do. Therefore, ensure that Cisco IOS MTPs are correctly prioritized in the media resource group list (MRGL).

The following example illustrates this type of configuration:

H.323

```

!
dial-peer voice 1000 voip
  destination-pattern 1T
  session target ipv4:10.10.10.1
  modem passthrough mode nse codec g711ulaw
!
! To enable T.38 fax relay.
fax protocol t38
!
!

```

MGCP

```

!
ccm-manager mgcp
mgcp
mgcp call-agent 10.10.10.1 service-type mgcp version 0.1
!
! T.38 fax relay is ON by default. HOWEVER, Unified CM doesn't
! support CA controlled mode. This is the configuration for
! talking to BTS.

```

```
!  
dial-peer voice 100 pots  
  application mgcpapp  
  port 1/0/0  
!
```

Gateways for Video Telephony

Cisco offers voice gateway functionality in a variety of forms, such as standalone devices, modules that integrate into Cisco IOS Routers, or line cards that integrate into Cisco Catalyst Ethernet Switches. These gateways support multiple VoIP protocols (such as H.323, MGCP, SIP, and SCCP), multiple port interface types (such as FXS, FXO, E&M, T1/E1-CAS, T1/E1-PRI, ISDN BRI, and so on), and a myriad of advanced VoIP features. They also offer a rich set of management and troubleshooting interfaces.

Cisco IP/VC gateways are scalable and offer a robust video gateway solution for large networks. The IP/VC gateways have the following characteristics:

- They support only H.323 and H.320.
- They are standalone devices that cannot be integrated into Cisco IOS Routers or Cisco Catalyst Switches.
- They support only T1/E1-PRI, ISDN BRI, and V.35 interface types.
- They support only G.711, G.728, G.723, and G.722; they do not support G.729 audio.
- They support the H.245 Empty Capabilities Set (ECS).
- They do not support many of the manageability and troubleshooting capabilities inherent in Cisco Voice Gateways.

Cisco IOS routers also support H.320 protocol, enabling video gateway capability in addition to voice. The Cisco IOS video gateways now add support for the H.26x family of video codecs and bonding to provide channel aggregation for video conferencing.



Deploying Cisco IOS gateways for voice and video can have the following advantages:

- Single gateway for voice and video
- Single circuit for PSTN from service provider

Telephony customers can choose to deploy the Cisco IP/VC Video gateways or Cisco IOS gateways for video. Smaller locations with existing Cisco IOS voice gateways can add the additional video gateway functionality at these distributed locations, while the larger locations can use the IP/VC and dedicated voice gateways to better achieve scalability.

When using a single ISDN circuit for voice and H.320 videoconferencing, these lines are shared, as depicted in [Figure 4-5](#).

With separate Video gateways and voice gateways, PSTN lines are not shared as depicted in [Figure 4-6](#).

Figure 4-5 *Traditional PBX Sharing PSTN Lines Between Voice and H.320 Videoconferencing***Figure 4-6** *Cisco Unified CallManager System with Separate PSTN Lines for Voice and IP Video Telephony*

With separate voice and video gateways, the route plans must also be separate for both inbound and outbound calls. For inbound calls, there is no way to have a single Direct Inward Dial (DID) extension for a user who wants to be able to receive both voice and video calls. Typically, each user will already have a DID for voice calls. When you introduce video into the scenario, users will have to be dialed some other way, such as via a second DID number or by dialing the main number of the video gateway and then entering the users video extension when prompted by the Interactive Voice Response (IVR). For outbound calls, there is no way to have a single PSTN access code for both voice and video calls. Typically, users will already have a well-known access code for voice (such as 9 in most US enterprises), but when you introduce video into the scenario, they will have to dial some other access code to place outbound video calls.

Another consideration for deploying two types of gateways is the placement of those gateways. Typically, enterprises have many PSTN gateway resources consolidated at their central site(s), and each branch office has some local gateway resources as well. For instance, Cisco Catalyst 6500 gateways may be deployed at the central site with several T1/E1 circuits connected to them, while Cisco Integrated Services Routers (ISRs) may be deployed at each branch office with either analog or digital trunks to the local CO. When video is introduced into this scenario, the customer must also determine the number

of PSTN circuits they will need for video and where the video gateways will be placed. For instance, will they deploy only a few IP/VC 3500 Series gateways at the central site, or will they also deploy them at each branch office?

Finally, consider how calls will be routed across the IP network to a remote gateway for the purpose of providing toll bypass, and how calls will be re-routed over the PSTN in the event that the IP network is unavailable or does not have enough bandwidth to complete the call. More specifically, do you want to invoke automated alternate routing (AAR) for video calls?

Routing Inbound Calls from the PSTN

Use one of the following methods to route inbound calls from the PSTN:

- Assign at least two different directory numbers to each video-enabled device in the Cisco Unified CallManager cluster, with one line for audio and another line for video. With this method, the outside (PSTN) caller must dial the correct number to enable video.
- For video calls, have outside callers dial the main number of the video gateway. Cisco Unified Videoconferencing gateways offer an integrated IVR that prompts the caller to enter the extension number of the party they are trying to reach. Cisco Unified CallManager will then recognize that it is a video call when ringing the destination device. This method relieves the caller from having to remember two different DID numbers for each called party, but it adds an extra step to dialing an inbound video call.



Note The outside video endpoints must support DTMF in order to enter the extension of the called party at the IVR prompt.

The following example illustrates the second method:

A user has a Cisco Unified IP Phone 7960 attached to a PC running Cisco Unified Video Advantage. The extension of the IP Phone is 51212, and the fully qualified DID number is 1-408-555-1212. To reach the user from the PSTN for a voice-only call, people simply dial the DID number. The CO sends calls to that DID number through T1-PRI circuit(s) connected to a Cisco Voice Gateway. When the call is received by the gateway, Cisco Unified CallManager knows that the gateway is capable of audio only, so it negotiates only a single audio channel for that call. Conversely, for people to reach the user from the PSTN for a video call, they must dial the main number of the video gateway and then enter the user's extension. For example, they might dial 1-408-555-1000. The CO would send calls to that number through the T1-PRI circuit(s) connected to a Cisco Unified Videoconferencing 3500 Series video gateway. When the call is received by the gateway, an IVR prompt asks the caller to enter the extension of the person they are trying to reach. When the caller enters the extension via DTMF tones, Cisco Unified CallManager knows that the gateway is capable of video, so it negotiates both audio and video channels for that call.

Gateway Digit Manipulation

The Cisco Unified Videoconferencing 3500 Series Gateways cannot manipulate digits for calls received from the PSTN. It takes the exact number of digits passed to it in the Q.931 Called Party Number field and sends them all to Cisco Unified CallManager. Therefore, Cisco Unified CallManager must manipulate the digits in order to match the directory number (DN) of the destination device. For instance, if the circuit from the CO switch to the gateway is configured to pass 10 digits but the extension

of the called party is only five digits, Cisco Unified CallManager must strip off the leading five digits before attempting to find a matching DN. You can implement this digit manipulation in one of the following ways:

- By configuring the Significant Digits field on the H.323 gateway device or on the H.225 gatekeeper-controlled trunk that carries the incoming calls from the IP/VC gateway

This method enables you to instruct Cisco Unified CallManager to pay attention to only the least-significant N digits of the called number. For example, setting the Significant Digits to 5 will cause Cisco Unified CallManager to ignore all but the last 5 digits of the called number. This is the easiest approach, but it affects all calls received from that gateway. Thus, if you have variable-length extension numbers, this is not the recommended approach.

- By configuring a translation pattern and placing it in the calling search space of the H.323 gateway device or of the H.225 gatekeeper-controlled trunk that carries the incoming calls from the IP/VC gateway

This method enables Cisco Unified CallManager to match calls to the full number of digits received, to modify the called number, and then to continue performing digit analysis on the resulting modified number. This approach is slightly more complex than the preceding method, but it is more flexible and enables you to use a finer granularity for matching calls and for specifying how they will be modified.

Routing Outbound Calls to the PSTN

Use one of the following methods to route outbound calls to the PSTN:

- Assign different access codes (that is, different route patterns) for voice and video calls. For example, when the user dials 9 followed by the PSTN telephone number they are trying to reach, it could match a route pattern that directs the call out a voice gateway. Similarly, the digit 8 could be used for the route pattern that directs calls out a video gateway.
- Assign at least two different directory numbers on each video-enabled device in the Cisco Unified CallManager cluster, with one line for audio and another line for video. The two lines can then be given different calling search spaces. When users dial the access code (9, for example) on the first line, it could be directed out a voice gateway, while dialing the same access code on the second line could direct the call out a video gateway. This method alleviates the need for users to remember two different access codes but requires them to press the correct line on their phones when placing calls.

Gateway Service Prefixes

The Cisco Unified Videoconferencing Gateways use service prefixes to define the speed for outbound calls. When you configure a service prefix in the gateway, you must choose one of the following speeds:

- Voice-only
- 128 kbps
- 256 kbps
- 384 kbps
- 768 kbps
- Auto (dynamically determined; supports any call speed in the range of 128 kbps to 768 kbps)

**Note**

Each of the above speeds represents a multiple of 64 kbps. For 56-kbps dialing, there is a check-box on the service prefix configuration page to restrict each channel to 56 kbps. Therefore, a 128-kbps service with restricted mode enabled would result in a 112-kbps service; a 384 kbps service with restricted mode enabled would result in a 336-kbps service; and so on.

Calls from an IP endpoint toward the PSTN must include the service prefix at the beginning of the called number in order for the gateway to decide which service to use for the call. Optionally, you can configure the default prefix to be used for calls that do not include a service prefix at the beginning of the number. This method can become quite complex because users will have to remember which prefix to dial for the speed of the call they wish to make, and you would have to configure multiple route patterns in Cisco Unified CallManager (one for each speed). Fortunately, the Auto speed enables you to minimize this effort. If the majority of your calls are made using 64 kbps per channel (for example, 128 kbps, 384 kbps, 512 kbps, 768 kbps, and so on), you could use the Auto service in that case. You would then need to create only one other service for the rare case in which someone makes a call using 56 kbps per channel (for example, 112 kbps, 336 kbps, and so on).

Cisco recommends that you always use a # character in your service prefixes because the gateway recognizes the # as an end-of-dialing character. By placing this character in the service prefix, you block people from attempting to use the gateway for toll fraud by dialing the main number of the gateway, reaching the IVR, and then dialing out to an off-net number. The # can either be at the beginning (recommended) or the end of the service prefix. For example, if your access code to reach the PSTN is 8 for video calls, Cisco recommends that you configure the service prefix as #8 or 8#. Or, if you have two service prefixes as described above, you might use #80 for the Auto 64-kbps service and #81 for the Auto 56-kbps service.

The ramification of using a service prefix is that Cisco Unified CallManager must prepend the service prefix to the called number when sending calls to the IP/VC gateway. Because forcing users to dial the # would not be very user-friendly, Cisco recommends that you configure Cisco Unified CallManager to prepend the # to the dialed number. For example, if the access code to dial a video call to the PSTN is 8, you could configure a route pattern as 8.@ in Cisco Unified CallManager, and in the route pattern configuration you would configure the called number translation rule to prepend #8 whenever that route pattern is dialed. Or, if you have two service prefixes as described above, you might use 80.@ for the Auto 64-kbps service (prefixing # to the called number) and 81.@ for the Auto 56-kbps service (prefixing # to the called number).

Automated Alternate Routing (AAR)

When the IP network does not have enough bandwidth available to process a call, Cisco Unified CallManager uses its call admission control mechanism to determine what to do with the call. As described in the chapter on [IP Video Telephony, page 17-1](#), Cisco Unified CallManager performs one of the following actions with the call, depending on how you have configured it:

- Fail the call, playing busy tone to the caller and displaying a Bandwidth Unavailable message on the caller's screen
- Retry the video call as an audio-only call
- Use automated alternate routing (AAR) to re-route the call over an alternative path, such as a PSTN gateway

The first two options are covered in the chapter on [IP Video Telephony, page 17-1](#), and this section covers the AAR option.

To provide AAR for voice or video calls, you must configure the calling and called devices as members of an AAR group and configure an External Phone Number Mask for the called device. The External Phone Number Mask designates the fully qualified E.164 address for the user's extension, and the AAR group indicates what digits should be prepended to the External Phone Number Mask of the called device in order for the call to route successfully over the PSTN.

For example, assume that user A is in the San Jose AAR group and user B is in the San Francisco AAR group. User B's extension is 51212, and the External Phone Number Mask is 6505551212. The AAR groups are configured to prepend 91 for calls between the San Jose and San Francisco AAR groups. Thus, if user A dials 51212 and there is not enough bandwidth available to process the call over the IP WAN between those two sites, Cisco Unified CallManager will take user B's External Phone Number Mask of 6505551212, prepend 91 to it, and generate a new call to 916505551212 using the AAR calling search space for user A.

This same logic applies to video calls as well, with one additional step in the process. For video-capable devices, there is field called Retry Video Call as Audio. As described in the chapter on [IP Video Telephony, page 17-1](#), if this option is enabled (checked), Cisco Unified CallManager does not perform AAR but retries the same call (that is, the call to 51212) as a voice-only call instead. If this option is disabled (unchecked), Cisco Unified CallManager performs AAR. By default, all video-capable devices in Cisco Unified CallManager have the Retry Video Call as Audio option enabled (checked). Therefore, to provide AAR for video calls, you must disable (uncheck) the Retry Video Call as Audio option. Additionally, if a call admission control policy based on Resource Reservation Protocol (RSVP) is being used between locations, the RSVP policy must be set to Mandatory for both the audio and video streams.

Furthermore, Cisco Unified CallManager looks at only the called device to determine whether the Retry Video Call as Audio option is enabled or disabled. So in the scenario above, user B's phone would have to have the Retry Video Call as Audio option disabled in order for the AAR process to take place.

Finally, devices can belong to only one AAR group. Because the AAR groups determine which digits to prepend, AAR groups also influence which gateway will be used for the rerouted call. Depending on your choice of configuration for outbound call routing to the PSTN, as discussed in the previous section, video calls that are rerouted by AAR might go out a voice gateway instead of a video gateway. Therefore, carefully construct the AAR groups and the AAR calling search spaces to ensure that the correct digits are prepended and that the correct calling search space is used for AAR calls.

While these considerations can make AAR quite complex to configure in a large enterprise environment, AAR is easier to implement when the endpoints are strictly of one type or the other (such as IP Phones for audio-only calls and systems such as the Tandberg T-1000 dedicated for video calls). When endpoints are capable of both audio and video calls (such as Cisco Unified Video Advantage or a Cisco IP Video Phone 7985G), the configuration of AAR can quickly become unwieldy. Therefore, Cisco recommends that large enterprise customers who have a mixture of voice and video endpoints give careful thought to the importance of AAR for each user, and use AAR only for select video devices such as dedicated videoconference rooms or executive video systems. [Table 4-13](#) lists scenarios when it is appropriate to use AAR with various device types.

Table 4-13 *When to Use AAR with a Particular Device Type*

Device Type	Device is used to call:	Enable AAR?	Comments
IP Phone	Other IP Phones and video-capable devices	Yes	Even when calling a video-capable device, the source device is capable of audio-only, thus AAR can be configured to route calls out a voice gateway.
IP Phone with Cisco Unified Video Advantage, or Cisco IP Video Phone 7985G	Other video-capable devices only	Yes	Because the device is used strictly for video calls, you can configure the AAR groups accordingly.
	IP Phones and other video-capable devices	No	It will be difficult to configure the AAR groups to route audio-only calls differently than video calls.
Sony or Tandberg SCCP endpoint	Other video-capable devices only	Yes	Because the device is used strictly for video calls, you can configure the AAR groups accordingly.
	IP Phones and other video-capable devices	No	It will be difficult to configure the AAR groups to route audio-only calls differently than video calls.
H.323 or SIP client	Other video-capable devices only	Yes	Because the device is used strictly for video calls, you can configure the AAR groups accordingly.
	IP Phones and other video-capable devices	No	It will be difficult to configure the AAR groups to route audio-only calls differently than video calls

Least-Cost Routing

Least-cost routing (LCR) and tail-end hop-off (TEHO) are very popular in VoIP networks and can be used successfully for video calls as well. In general, both terms refer to a way of configuring the call routing rules so that calls to a long-distance number are routed over the IP network to the gateway closest to the destination, in an effort to reduce toll charges. (For Cisco Unified CallManager Release 4.1, LCR basically means the same thing as TEHO.) Cisco Unified CallManager supports this feature through its rich set of digit analysis and digit manipulation capabilities, including:

- Partitions and calling search spaces
- Translation patterns
- Route patterns and route filters
- Route lists and route groups

Configuring LCR for video calls is somewhat more complicated than for voice calls, for the following reasons:

- Video calls require their own dedicated gateways, as discussed previously in this chapter
- Video calls require much more bandwidth than voice calls

With respect to dedicated gateways, the logic behind why you might or might not decide to use LCR for video calls is very similar to that explained in the section on [Automated Alternate Routing \(AAR\)](#), [page 4-35](#). Due to the need to have different types of gateways for voice and video, it can become quite

complex to configure all the necessary partitions, calling search spaces, translation patterns, route patterns, route filters, route lists, and route groups needed for LCR to route voice calls out one gateway and video calls out another.

With respect to bandwidth requirements, the decision to use LCR depends on whether or not you have enough available bandwidth on your IP network to support LCR for video calls to/from a given location. If the current bandwidth is not sufficient, then you have to determine whether the benefits of video calls are worth the cost of either upgrading your IP network to make room for video calls or deploying local gateways and routing calls over the PSTN. For example, suppose you have a central site with a branch office connected to it via a 1.544-Mbps T1 Frame Relay circuit. The branch office has twenty video-enabled users in it. A 1.544-Mbps T1 circuit can handle at most about four 384-kbps video calls. Would it really make sense in this case to route video calls up to the central site in order to save on toll charges? Depending on the number of calls you want to support, you might have to upgrade your 1.544-Mbps T1 circuit to something faster. Is video an important enough application to justify the additional monthly charges for this upgrade? If not, it might make more sense to deploy an IP/VC video gateway at the branch office and not bother with LCR. However, placing local IP/VC gateways at each branch office is not inexpensive either, so ultimately you must decide how important video-to-PSTN calls are to your business. If video is not critical, perhaps it is not worth upgrading the bandwidth or buying video gateways but, instead, using the Retry Video Call as Audio feature to reroute video calls as voice-only calls if they exceed the available bandwidth. Once a call is downgraded to voice-only, local gateway resources and bandwidth to perform LCR become more affordable and easier to configure.

ISDN B-Channel Binding, Rollover, and Busy Out

H.320 video uses multiple ISDN channels bound together to achieve the speeds needed to pass full-motion video. One of the problems with this bonding mechanism is that, when an inbound ISDN video call is received, the gateway does not know how many channels will be requested for that call until after it accepts the call and the source device indicates how many additional channels are required. If there are not enough B-Channels to satisfy the request, the call is disconnected. Therefore, careful traffic engineering is required to minimize the possibility that this situation will occur. Essentially, you want to ensure that there are always enough B-Channels available to handle the next call that might come in.

This B-Channel issue occurs in two cases:

- Inbound calls from the PSTN to the IP network
- Outbound calls from the IP network to the PSTN

Inbound Calls

For inbound calls, consider the following scenario:

A company has a Cisco 3526 IP/VC Gateway with an ISDN PRI circuit connecting it to a central office (CO) switch. The ISDN PRI circuit in this case offers 23 B-Channels. A video call is received from the PSTN at 384 kbps. This call takes six B-Channels, leaving 17 available. A second and third 384-kbps call are received on the line while the first one is still active. These each take an additional six channels, leaving five channels available. When the fourth 384-kbps call is received, the gateway will answer the call but, recognizing that it does not have enough B-Channels available (it only has five left but the call requires six), it will disconnect (by sending a Q.931 RELEASE COMPLETE with "16: Normal Call Clearing" as the reason). The caller attempting to make the fourth call will not know why the call failed and might redial the number repeatedly, trying to make the call work.

On Cisco Unified Videoconferencing gateways, you can minimize your chances of running into this issue by configuring the gateway to send a request to the CO to busy-out the remaining B-Channels (in this example, five channels) whenever the gateway reaches a certain threshold of utilization (configured as a percentage of total bandwidth).

In addition, you can have the CO provision multiple ISDN circuits in a trunk group. When the first circuit reaches the busy-out threshold, calls will roll over to the next PRI in the group. The Cisco 3540 IP/VC Gateway offers two ISDN PRI connections and supports bonding channels across both ports. For example, port 1 might have only five channels available while port 2 is sitting idle and, therefore, has 23 channels available. By taking the five channels from port 1 and one channel from port 2 and bonding them together, the fourth 384-kbps call can succeed. This leaves 22 channels available on controller 2, and at some point additional inbound calls would reach the busy-out threshold again. At that point the remaining channels on port 2 will be busied out, and all further inbound calls will be rejected with cause code "Network Congestion." Cisco Unified Videoconferencing gateways cannot bond channels across different gateways or across different Cisco 3540 gateway models in the same Cisco 3544 chassis, so two ports is the maximum that you can bond together. The CO switch can still roll calls over to a third or fourth PRI in the trunk group (most COs support trunk groups of up to 6 circuits), but you cannot bond channels between PRI number one and PRI number three, for example, as you can between PRI number one and PRI number two.

The busy-out logic described above depends on the assumption that all calls take place at the same speed. Suppose, for example, that two 384-kbps calls are active on a port and a 128-kbps call came in. This call would take only two channels, using a total of 14 channels for the three calls ($6+6+2 = 14$) and leaving nine channels available on the circuit. However, if the busy-out threshold is set at 18 channels (assuming that all calls would take place at 384-kbps), only four channels are still available under this busy-out threshold. If another 384 kbps call comes in at this point, the call will fail because the remaining four channels are not enough to support the call. Also, because the busy-out threshold of 18 channels has not been reached yet (only 14 channels are used), the circuit is not busied out and calls will not roll over to the next circuit. This condition will persist until one of the existing calls is disconnected. To avoid such situations, it is important to try to standardize on a single call speed for all calls.

Outbound Calls

Outbound calls encounter the same potential situations as inbound calls, but the way in which the busy-out occurs is different. The Cisco 3500 Series IP/VC Gateways support messages called Resource Availability Indicator and Resource Availability Confirm (RAI/RAC). The RAI/RAC messages are defined under the H.225 RAS specification and are used by the gateways to tell the gatekeeper that they are full and to no longer route any more calls to them. When the gateway reaches the busy-out threshold, it sends an RAI message with a status of True to the gatekeeper. True means "Do not send me any more calls;" False means "I am available." The gateway sends an RAI=False as soon as it is no longer at its busy-out threshold. The busy-out threshold for outbound calls is separate from the busy-out threshold for inbound calls, and you can configure them differently so that inbound calls will roll over to the next available circuit but outbound calls will still be accepted, or vice versa. For example, you could configure the RAI threshold to 12 channels but the ISDN busy-out threshold to 18 channels. When two 384 kbps are active, outbound calls will roll over to the next available gateway, but a third 384-kbps inbound call could still be received. An equally efficient method of achieving outbound call busy-out failover is to use Cisco Unified CallManager's route group and route list construct, as described in the following section, instead of the RAI/RAC method.

Configuring the Gateways in Cisco Unified CallManager

You can configure an IP/VC gateway in either of the following ways in Cisco Unified CallManager:

- Configure it as an H.323 gateway, and Cisco Unified CallManager will route calls directly to the gateway.
- Configure an H.225 gatekeeper-controlled trunk to the gatekeeper, and route calls to the gateway through the gatekeeper.

If you have only one gateway, it is probably easier to configure it directly in Cisco Unified CallManager instead of going through a trunk to get to it. If you have multiple gateways for load balancing and redundancy, you can either configure them all in Cisco Unified CallManager and place them into a route group(s) and route list, or configure an H.225 trunk to the gatekeeper and rely on RAI/RAC between the gateways and the gatekeeper to tell Cisco Unified CallManager which gateway it should send a given call to.

For inbound calls from the PSTN to Cisco Unified CallManager, the Cisco Unified Videoconferencing gateways can either register with a gatekeeper or be configured with the IP addresses of up to three Cisco Unified CallManager servers to which they should send all inbound call requests. This method is known as peer-to-peer mode. Either way, the goal is have all inbound calls received by the gateways sent to Cisco Unified CallManager so that Cisco Unified CallManager can decide how to route the calls. See [Gatekeepers, page 17-19](#), for more details on how to configure the gatekeeper to route calls from the gateways to Cisco Unified CallManager.

Call Signaling Port Numbers

By default, the Cisco Unified Videoconferencing Gateways listen on TCP port 2720 instead of the well-known port 1720. However, also by default, Cisco Unified CallManager sends H.323 calls to port 1720. You can change the port that the gateway listens on or you can change the port that Cisco Unified CallManager sends to in the H.323 gateway device configuration in Cisco Unified CallManager. Either way, both sides have to match in order for outbound calls to the gateway to succeed.

In the inbound direction, when configured to operate in peer-to-peer mode, the Cisco Unified Videoconferencing Gateways will send the call to Cisco Unified CallManager on port 1720. When configured to register with a gatekeeper, Cisco Unified CallManager uses a randomly generated port number for all gatekeeper-controlled trunks. This method enables Cisco Unified CallManager to have multiple trunks to the same gatekeeper. This port number is included in the Registration Request (RRQ) from Cisco Unified CallManager to the gatekeeper, so the inbound H.225 setup message from the gateway to Cisco Unified CallManager will be sent to this port number. However, if the gateway is configured directly in Cisco Unified CallManager as an H.323 gateway device, Cisco Unified CallManager will ignore the fact that the call came in on the TCP port of the H.225 trunk and will instead match the source IP address to the H.323 gateway device configured in its database. If it does not find a matching device, Cisco Unified CallManager will treat the call as if it came in on the trunk.

In the outbound direction, if Cisco Unified CallManager uses a gatekeeper-controlled H.225 trunk to reach the gateway, the gatekeeper will tell Cisco Unified CallManager which TCP port to use to reach the gateway. If the gateway is configured in Cisco Unified CallManager as an H.323 gateway device (that is, peer-to-peer mode), then Cisco Unified CallManager must be configured to send calls either to port 2720 (default) or to 1720 (if the listening port on the gateway has been modified).

Call Signaling Timers

Due to the delay inherent in H.320 bonding, video calls can take longer to complete than voice calls. Several timers in Cisco Unified CallManager are tuned, by default, to make voice calls process as fast as possible, and they can cause video calls to fail. Therefore, you must modify the following timers from their default values in order to support H.320 gateway calls:

- H.245TCSTimeout
- Media Exchange Interface Capability Timer
- Media Exchange Timer

Cisco recommends that you increase each of these timers to 25 by modifying them under the Service Parameters in Cisco Unified CallManager Administration. Note that these are cluster-wide service parameters, so they will affect calls to all types of H.323 devices, including voice calls to existing H.323 Cisco Voice Gateways.

Voice Gateways Bearer Capabilities

H.323 calls use the H.225/Q.931 Bearer Capabilities Information Element (bearer-caps) to indicate what type of call is being made. A voice-only call has its bearer-caps set to "speech" or "3.1 KHz Audio," while a video call has its bearer-caps set to "Unrestricted Digital Information." Cisco voice gateways, some legacy PBXs, and most cellular carriers do not support Unrestricted Digital Information bearer-caps. Therefore, calls to a voice gateway might fail if Cisco Unified CallManager attempts the call as a video call.

Cisco Unified CallManager decides which bearer-caps to set, based on the following factors:

- Whether the calling and/or called devices are video-capable
- Whether the region in Cisco CallManager is configured to allow video for calls between those devices

For example, consider a network in which a video-capable device (such as a Cisco Unified IP Phone with a VT Advantage client associated to it) is configured in the same region as a Cisco voice gateway. When the user dials 9 to access an outside line, Cisco Unified CallManager determines that the calling device is video-capable and that the region is set to allow 384 kbps of video bandwidth.

Cisco Unified CallManager then sets the bearer-caps to Unrestricted Digital Information for that call. But because the call is to a Cisco voice gateway, the gateway rejects the call with the cause code "Incompatible Destination." This problem will occur in any network that uses H.323 voice gateways and that has IP Phones associated with Cisco Unified Video Advantage. From the user's perspective, things will appear to work fine before installing Cisco Unified Video Advantage, but calls to the PSTN will fail as soon as the user plugs a PC running Cisco Unified Video Advantage into the IP Phone.

This situation exists only on calls to H.323 voice gateways. If the Cisco voice gateway uses MGCP to communicate with Cisco Unified CallManager, the problem will not occur because Cisco Unified CallManager does not support video on its MGCP protocol stack and because, in MGCP mode, Cisco Unified CallManager has complete control over the D-Channel signaling to the PSTN. Likewise, if the Cisco voice gateway uses SIP to communicate with Cisco Unified CallManager, the problem will not occur then either because Cisco Unified CallManager does not support video on its SIP protocol stack, but even if it did, the gateway would simply need to ignore the video capabilities passed in Cisco Unified CallManager's outgoing Session Description Protocol (SDP) advertisement.

To prevent this situation, configure the bearer-caps on all Cisco H.323 voice gateways by using the **bearer-caps** command under the **voice-port** configuration mode, as illustrated in the following example:

```
gateway#configure terminal
gateway(config)#voice-port 1/0:23
gateway(config-voiceport)#bearer-caps speech
```



CHAPTER 5

Cisco Unified CallManager Trunks

Last revised on: February 13, 2008

This chapter discusses design considerations for Cisco Unified Callmanager Release 4.2, but much of the discussion also applies to Cisco Unified CallManager releases 4.1, 4.0, and 3.3.

The use of the term *trunk* in Cisco Unified CallManager is different than the traditional meaning for trunk used when discussing circuit switched systems. The definition of trunk is expanded beyond physical connections such as a PRI, T1, or analog trunk, to also include virtual connections between two entities. The virtual connection defines a TCP/IP session between the two entities that is used for call signaling, and associated with the connection is a path over a TCP/IP network used for the media. The signaling path and the media path may terminate on different sets of endpoints and also may take different paths through the TCP/IP network. In some cases the trunk may define an entity that is the destination where the VoIP media stream of the call ultimately terminates, or it may define an entity that provides an intermediate routing function. In the first case, there must exist a separate configuration for each pair of endpoints, and the relationship between the two entities is peer-to-peer. In the latter case, a single trunk may be defined to specify a gatekeeper that provides a routing function between all entities registered to that gatekeeper. This relationship is one-to-many.

In Cisco Unified CallManager, trunks may be configured with one of two protocols: H.323 or SIP. The following sections introduce the characteristics of these trunks.

H.323 Protocol

There are four different configurations in Cisco Unified CallManager that may be used to communicate to other entities via the H.323 protocol suite:

- H.323 gateway (See [Gateways, page 4-1.](#))

This type of configuration is used to define a peer-to-peer relationship between the cluster and another H.323 gateway or endpoint without gatekeeper control. It is not a trunk under the definition used by Cisco Unified CallManager but is mentioned here for completeness of the discussion. This configuration may *not* be used to connect two Cisco Unified CallManager clusters.

- H.225 trunk, gatekeeper controlled (See [Gatekeeper Controlled Trunks, page 5-2.](#))

This trunk configuration defines communication between a Cisco Unified CallManager cluster and one or more H.323 gateways, endpoints, or other clusters under control of an H.323 gatekeeper.

- Intercluster trunk, non-gatekeeper controlled (See [Non-Gatekeeper Controlled Trunks, page 5-3.](#))

This trunk configuration defines communication between a single pair of Cisco Unified CallManager clusters using an H.323-based protocol. There is no gatekeeper associated with this trunk.

- Intercluster trunk, gatekeeper controlled (See [Gatekeeper Controlled Trunks, page 5-2.](#))

This configuration allows multiple Cisco Unified CallManager clusters to interconnect via H.323 with a single trunk definition under control of a gatekeeper.

SIP Protocol

Cisco Unified CallManager classifies SIP entities as either line side or trunk side. Entities that are integrated via SIP trunks include gateways, voicemail systems, presence servers, other clusters, SIP proxies, and so forth. Cisco Unified CallManager provides only one type of SIP trunk. (See [SIP Trunks, page 5-11.](#))

SIP standards do not define the term trunk; it is terminology that is specific to the Cisco Unified CallManager implementation. A SIP Trunk configures members of a cluster to behave as a SIP User Agent, and the cluster expects the other SIP entity defined in the configuration to also act as a SIP User Agent.

Support for SIP trunks was first added in Cisco Unified CallManager Release 4.0.

H.323 Trunks

The following types of H.323 trunks can be configured in a Cisco Unified CallManager cluster:

- H.225 trunk (gatekeeper controlled)
- Intercluster trunk (non-gatekeeper controlled)
- Intercluster trunk (gatekeeper controlled)

All three types of H.323 trunks may be used to connect Cisco Unified CallManager clusters (Release 3.2 and later) to each other, but only the H.225 gatekeeper controlled trunk may be used to connect a Cisco Unified CallManager cluster to other H.323 endpoints. When any of these three types of trunks are used to communicate between clusters, extensions to the H.225 protocol are used to provide additional feature transparency between the clusters. The use of these protocol extensions is invoked automatically via autodetection.

Protocol Autodetection

This feature provides the ability to determine, on a call-by-call basis, if the calling device is from Cisco Unified CallManager Release 3.2 or later. When a call is received on an H.323 trunk, Cisco Unified CallManager looks for an H.225 User-to-User Information Element (UUIE) that indicates if the other end is another Cisco Unified CallManager. Only Cisco Unified CallManager 3.2 or later will send this UUIE. When the UUIE is present, then both clusters will utilize the Cisco-specific H.225 protocol extensions. If no UUIE is found, the use of the Cisco extensions are dependant on the configured trunk type.

Gatekeeper Controlled Trunks

You may want to provision a gatekeeper for any of the following reasons:

- To reduce the administration of the full-mesh configuration that might be required for larger numbers of H.323 endpoints.

With a single Cisco gatekeeper, it is possible to have 100 clusters all registering a single trunk each, with all clusters being able to call each other. With non-gatekeeper controlled trunks, this same topology would require 99 trunks configured in each cluster.

- To provide gatekeeper-based call admission control. (See [Call Admission Control, page 9-1](#), for more information.)
- To improve failover times.

If a subscriber server in a cluster becomes unreachable, there will be a five-second (default) timeout while the call is attempted. If an entire cluster is unreachable, the number of attempts before either call failure or rerouting over the PSTN will depend on the number of remote servers defined for the trunk and on the number of trunks in the route list or route group. If there are many remote servers and many non-gatekeeper controlled trunks, the call delay can become excessive.

With a gatekeeper controlled trunk, you configure only one trunk that can then communicate via the gatekeeper with all other clusters registered to the gatekeeper. If a cluster or subscriber becomes unreachable, the gatekeeper automatically directs the call to another subscriber in the cluster or rejects the call if no other possibilities exist, thus allowing the call to be rerouted over the PSTN (if required) with little incurred delay.

If a gatekeeper is used, then the following two types of trunks are available:

- H.225 trunk (gatekeeper controlled)

Always use this configuration to communicate with any combination Cisco Unified CallManager (Release 3.2 or later) or other H.323 endpoints that are all controlled by a common set of gatekeeper(s). Cisco Unified CallManager will use the H.225-based Intercluster Trunk Protocol between clusters and use standard H.225 to other H.323 endpoints. The autodetection feature will invoke the appropriate behavior.

Do not use this trunk for connection to any version of Cisco CallManager prior to Release 3.2 because earlier versions will not send the UUIE and the connection will use the standard H.225 protocol.

- Intercluster trunk (gatekeeper controlled)

This trunk must be used only for communication to Cisco Unified CallManagers. The trunk always utilizes the Intercluster Trunk Protocol, which will not be understood by a standard H.323 endpoint. Configure this type of trunk only when interoperability is required with versions of Cisco CallManager prior to Release 3.2.

Non-Gatekeeper Controlled Trunks

If no gatekeeper is used, then the only choice for cluster-to-cluster trunking is an intercluster trunk (non-gatekeeper controlled).

This trunk type is the simplest and is used for connecting to other Cisco Unified CallManager clusters in either a multi-cluster single campus or a distributed call processing deployment. This trunk does not use a gatekeeper for call admission control, although it may use locations configured in Cisco Unified CallManager if bandwidth control is required. This type of trunk must be used only for communication with other Cisco Unified CallManager clusters because it utilizes the Intercluster Trunk Protocol.

When defining this type of trunk, you may define up to three remote Cisco Unified CallManager servers in the same destination cluster. The trunk will automatically load-balance across all defined servers. In the remote cluster, it is important to configure a corresponding intercluster trunk (non-gatekeeper controlled) that has a Cisco Unified CallManager group containing the same servers that were defined as remote Cisco Unified CallManager servers in the first cluster. A similar configuration is required in each Cisco Unified CallManager cluster connected by the intercluster trunks.

For example, if Cluster 1 has a trunk to Cluster 2, and Cluster 2 has a trunk to Cluster 1, the following configurations would be needed:

- Cluster 1
 - Servers B, C, and D are configured as members of the Cisco Unified CallManager group defined in the device pool associated with the trunk to Cluster 2.
 - The non-gatekeeper controlled trunk has Cluster 2's remote servers D, E, and F configured.
- Cluster 2
 - Servers D, E, and F are configured as members of the Cisco Unified CallManager group defined in the device pool associated with the trunk to Cluster 1.
 - The non-gatekeeper controlled trunk has Cluster 1's remote servers B, C, and D configured.

Gatekeeper Trunk Redundancy, Resilience, and Load Balancing

Redundancy can be achieved in several ways, depending on the requirements of the design. The simplest method is to configure a gatekeeper controlled trunk and assign up to three subscribers in the Cisco Unified CallManager group associated with the device pool assigned to that trunk. This configuration will cause all servers to register with the same gatekeeper in the same zone with the same technology prefix. However, the H.323 trunk name that is used for the h323_id will have a suffix of "_n" where *n* is the node number in the cluster. This ID is automatically generated and cannot be changed. You configure a single trunk, but the gatekeeper registers multiple trunks, one from each subscriber in the Cisco Unified CallManager group.

If you have additional redundancy requirements, it is possible to configure another gatekeeper controlled trunk with a different name and different subscribers in the Cisco Unified CallManager group, but with all the other parameters identical to the first trunk. This second trunk will cause additional subscribers to register with the gatekeeper.

Cisco recommends assigning device pools that contain a Cisco Unified CallManager group consisting of the two servers that make up the standard subscriber pair. (See [Call Processing Subscriber](#), page 8-6, for more information on subscriber redundancy.) For complete redundancy in a full cluster, four trunks would be needed using four different device pools, resulting in eight subscribers registering with the gatekeeper. (The same result could be achieved with three trunks and larger Cisco Unified CallManager groups.)

During registration, several parameters are passed between Cisco Unified CallManager and the gatekeeper. Cisco Unified CallManager uses an ephemeral User Datagram Protocol (UDP) port for gatekeeper Registration Admission Status (RAS) messages on a per-trunk basis. This port would normally be UDP 1719. However, Cisco Unified CallManager must be able to determine which trunk is the destination for a particular RAS message, therefore it uses a range of UDP ports and assigns them dynamically.


During the registration process, a trunk registers the following information for the other subscribers in its Cisco Unified CallManager group:

- H.225 call signaling port
- h323_id
- CanMapAlias support
- Technology prefix
- H.225 call signaling address

If the recommended clustered gatekeepers are used, the gatekeeper will return a list of alternate gatekeeper addresses that may be used if the primary gatekeeper fails or does not have sufficient available resources.


Figure 5-1 shows a cluster of gatekeepers that use Gatekeeper Update Protocol (GUP) to communicate. (See the chapter on [Call Processing](#), page 8-1, for more information on gatekeepers.)

Figure 5-1 Gatekeeper Cluster




If an H.323 trunk has only a single subscriber in its Cisco Unified CallManager group, there will be only one connection between the configured gatekeeper in Cisco Unified CallManager and the gatekeeper cluster, as illustrated in Figure 5-2.


Figure 5-2 H.323 Trunk with a Single Cisco Unified CallManager Subscriber



If there are multiple subscribers in the Cisco Unified CallManager group associated with the trunk, additional connections will be established between the Cisco Unified CallManager cluster and the gatekeeper cluster, as illustrated in Figure 5-3.

Figure 5-3 H.323 Trunk with Multiple Cisco Unified CallManager Subscribers

This approach provides redundancy for subscriber failures as well as gatekeeper failures after registration because the alternate gatekeeper is communicated when the trunk registers. This approach does not, however, provide redundancy if the configured gatekeeper is unavailable at initial registration or following a reset because the list of alternate gatekeepers is dynamic and not stored in the database. To provide an additional level of redundancy as well as load balancing, an additional gatekeeper from the gatekeeper cluster is configured in Cisco Unified CallManager. For example, if the original trunk is registered with Element 2, the additional gatekeeper could be configured as Element 4, as illustrated in Figure 5-4.

Figure 5-4 Additional Gatekeeper Configured for Load Balancing and Additional Redundancy

The Cisco Unified CallManager configuration for the example in Figure 5-4 would contain the following components:

- Two gatekeepers for Element 2 and Element 4
- Two H.323 trunks defined with a Cisco Unified CallManager group containing subscriber servers A and B

Using this approach, the Cisco Unified CallManager cluster will still be able to register when either Element 2 or Element 4 is not reachable during initial registration (that is, during power-up or trunk reset).

Load balancing of calls inbound to the Cisco Unified CallManager cluster is done automatic by default because the gatekeeper randomly selects one of the subscribers registered within the zone. If this is not the desired behavior, you can use the **gw-priority** configuration command in the gatekeeper to modify this default behavior, as illustrated in Example 5-1.

Example 5-1 Using the gw-priority Command to Direct Calls to a Particular Trunk

```
gatekeeper
zone local SJC cisco.com 10.0.1.10
zone prefix SJC 1408..... gw-priority 10 sjc-trunk_2
zone prefix SJC 1408..... gw-priority 9 sjc-trunk_3
zone prefix SJC 1408..... gw-default-priority 0
gw-type-prefix 1#* default-technology
arq reject-unknown-prefix
no shutdown
endpoint ttl 60
```

In [Example 5-1](#), the H.323 trunk was configured as sjc-trunk in Cisco Unified CallManager, and the “_2” and “_3” suffixes are appended automatically by the Cisco Unified CallManager subscribers to indicate which node number they are in the cluster. Therefore, this example uses node 2 as the first choice, which should be the highest-priority Cisco Unified CallManager in the CallManager group for this trunk. Node 3 is the second choice in this case.

The use of **gw-default-priority 0** is optional. It was used in this example to disable the use of any other trunk that might accidentally be configured to register in this zone.

Outbound calls from a Cisco Unified CallManager cluster can be load-balanced in any of the following ways:

- A single H.323 trunk in a route group will always use the highest-priority subscriber available in the Cisco Unified CallManager group. Lower-priority subscribers are used only when higher-priority subscribers are unavailable.
- Multiple H.323 trunks in a circular route group provide equal call loading across all H.323 trunks in the group.

The following examples list how to configure load balancing in various scenarios.

All calls originate from a single subscriber in a cluster:

- Single H.323 trunk in a route group

Calls spread across four primary subscribers in a cluster:

- Four H.323 trunks with four Cisco Unified CallManager groups, all contained within a circular route group
- Cisco Unified CallManager groups defined as follows:
 - Subscriber A, Subscriber B
 - Subscriber C, Subscriber D
 - Subscriber E, Subscriber F
 - Subscriber G, Subscriber H

Where subscribers A, C, E, and G are all primaries and subscribers B, D, F, and H are the backups.

Calls spread across eight subscribers in a cluster:

- Eight H.323 trunks with eight different Cisco Unified CallManager groups, each containing only one subscriber, and all contained within a circular route group
- Cisco Unified CallManager groups defined as follows:
 - Subscriber A
 - Subscriber B
 - Subscriber C

- Subscriber D
- Subscriber E
- Subscriber F
- Subscriber G
- Subscriber H

H.323 Trunks with Media Termination Points

Media termination points (MTPs) are generally not required for normal operation of the H.323 trunk. They are, however, required for communication with devices that are H.323 Version 1 or that do not support the Empty Capabilities Set (ECS) for supplementary services.

To test whether or not an MTP is required, use the following simple procedure:

1. Place a call from a phone via the H.323 trunk to the other device. This call should work normally.
2. Place the call on hold, then resume it. If the call drops, then it is highly likely that an MTP is required to ensure interoperability between Cisco Unified CallManager and the other device.

MTPs are very useful for terminating media streams from other devices that make calls over the H.323 trunk and for re-originating the media streams with the same voice payload; however, in such cases the IP address is changed to that of the MTP. With this fact in mind, you can utilize MTPs in the following scenarios:

- If the phones, gateways, and other devices within your enterprise all use RFC 1918 private addresses, you might still want to connect to other systems on a public network without using Network Address Translation (NAT) for all your voice and video devices. If the Cisco Unified CallManager subscriber that communicates to the public network is using a public IP address, the signaling will be routed. If all MTPs are also using public addresses, the media from the devices with RFC 1918 addresses will be terminated on the MTP and then originated again, but this time with a public address that is routable on the public network. This approach allows tens of thousands of devices with RFC 1918 addresses to communicate with the public network. This same method can be used to conceal the real IP addresses of devices in an enterprise network when communicating with other enterprises or service providers.
- Trust boundaries can be established to traverse firewalls or to allow access through an access control list (ACL). Normally, for media to traverse a firewall, you could either use an Application Layer Gateway (ALG) or fix-up to provide access dynamically for the media streams or you could allocate a wide range of addresses and ports for use by all voice devices that need to communicate across the firewall. All calls that use the H.323 trunk and traverse a firewall or ACL will have media that is sourced from the MTP(s), which may use either a single IP address or a small range of IP addresses.

With both of these methods, if the **MTP Required** box is checked, the default behavior is to allow calls on the H.323 trunks even if MTP resources are unavailable or exhausted. This default behavior might result in no voice path for the call, but the behavior can be changed by setting the Cisco CallManager service parameter **Fail Call if MTP allocation fails** under the H.323 section to **True**.

H.323 Operation in Cisco Unified CallManager

This section provides information on how the H.323 protocol is used and implemented in Cisco Unified CallManager, and it explains how and why certain features work the way they do.

The most important point to understand is which subscribers run the call signaling daemons. These daemons are pieces of code that make and receive H.323 calls. They are usually referred to as H.225 daemons, or H.225Ds. H.225 is part of the H.323 protocol and is mainly responsible for call control. H.245 is the other major component of H.323 that is responsible for the media control of a call.

The subscribers listed in the Cisco Unified CallManager group for a particular H.323 device determine which subscribers run the daemons and when. This point is a very important because calls sent to an incorrect subscriber might be rejected or handled by a different H.225D. For example, this situation would occur if a Cisco IOS H.323 gateway is configured with dial peers that send calls to subscriber C in a Cisco Unified CallManager cluster but the Cisco Unified CallManager group for that gateway has only subscribers A and B in its list. In such a case, the call will fail or be handled by an H.323 trunk daemon if one happens to be configured on the subscriber.

The following scenarios describe where and when H.225Ds are created on subscribers:

- H.323 client

The H.225D is active on only the highest-priority subscriber available in the Cisco Unified CallManager group associated with the H.323 client.

If the H.323 client is gatekeeper controlled, the RasAggregator device registers from only the highest-priority subscriber available in the Cisco Unified CallManager group associated with the gatekeeper controlled H.323 client.

The RasAggregator is a special device that registers in gatekeeper zones for the purpose of providing two specific features:

- If H.323 clients use DHCP, they cannot be used with a Cisco Unified CallManager using DNS unless they support Dynamic DNS. With the RasAggregator, Cisco Unified CallManager can obtain the IP address of a specific H.323 client that is registered with the gatekeeper whenever a call is placed. The gatekeeper registration is done using standard RAS ARQ messages that contain the E.164 address of the H.323 client. The gatekeeper resolves the E.164 address and provides the IP address back to Cisco Unified CallManager in an ACF message.
- The RasAggregator also ensures that all calls by the H.323 clients are made through Cisco Unified CallManager and not directly between the clients themselves, thus ensuring that dialing rules and codec restrictions are enforced.


- H.323 gateway

The H.225D is active on all subscribers in the Cisco Unified CallManager group associated with the H.323 gateway.

- H.323 trunks

The H.225D is active on all subscribers in the Cisco Unified CallManager group associated with the H.323 trunk. A RAS daemon registers the trunk with the gatekeeper from all subscribers in the associated Cisco Unified CallManager group.

When an incoming H.323 call is made to a subscriber in a Cisco Unified CallManager cluster, various decisions are made to determine if the call is accepted or rejected and which H.225D will receive the call if it is accepted. [Figure 5-5](#) shows how this process works.

Figure 5-5 Process for Determining if an H.323 Call is Accepted or Rejected

132369

Cisco Unified CallManager H.323 protocol includes the following additional features:

- Protocol Auto Detect

This feature provides the ability to determine, on a call-by-call basis, if the calling device is from Cisco Unified CallManager Release 3.2 or later. Whenever a call is received, Cisco Unified CallManager looks for an H.225 User-to-User Information Element (UUIE) that indicates if the other end is another Cisco Unified CallManager. If it is, it will always use the Intercluster Trunk Protocol. If no UUIE is found, it will use the configured protocol for that device. This feature enables an H.225 gatekeeper controlled trunk to switch between Intercluster Trunk Protocol and H.225 on a call-by-call basis, allowing a mixture of Cisco Unified CallManager clusters and other H.323 devices to use the gatekeeper. Intercluster Trunk Protocol is the same as H.225 except for several differences that enable specific features to work correctly between Cisco Unified CallManager clusters.

- Tunneled Q.SIG or H.323 Annex M1

With the release of Cisco Unified CallManager 4.1(3), this feature can be enabled on all H.323 trunks. It allows specific H.323 Annex M1 features to be implemented between Cisco Unified CallManager clusters and other verified systems that also support H.323 Annex M1. These features include:

- Path replacement
- Message waiting indication (MWI)
- Callback

- Alternate Endpoints

When registering with a gatekeeper that supports this feature, such as a Cisco Multimedia Conference Manager (MCM) Gatekeeper, Cisco Unified CallManager can inform the gatekeeper of alternate destinations for calls to the H.323 trunk. These alternate endpoints or destinations are sent to the calling device by the gatekeeper when this H.323 trunk is called. They are the other subscribers listed in the Cisco Unified CallManager group associated with the H.323 trunk that registers with the gatekeeper.

- Alternate Gatekeeper

When an H.323 trunk registers with a gatekeeper that supports this feature (for example, a Cisco gatekeeper cluster), Cisco Unified CallManager is dynamically informed about other gatekeepers that can process registrations, call admission requests, and other RAS functions in the event that this gatekeeper fails or exhausts its own resources.

- CanMapAlias

When an H.323 trunk sends an admission request (ARQ) to the gatekeeper, it might receive a different E.164 number in the admission confirmation message (ACF), indicating that the original called number should be replaced with this new one. This feature requires a route server using Gatekeeper Transaction Message Protocol (GKTMP) to communicate with Cisco gatekeepers.



Note CanMapAlias is supported for the called number only.

- Bandwidth Requests

H.323 trunks can update the gatekeeper with bandwidth information to indicate a change in the requested bandwidth allocated to a specific call. This feature is disabled by default and is controlled by setting the Cisco CallManager service parameter **BRQ Enabled** to **True**, under the H.323 section. This feature is especially important when video is used on an H.323 trunk because the original bandwidth request is for the maximum amount allowed. Enabling this feature ensures that call admission control uses the actual bandwidth negotiated during call setup.

SIP Trunks

Support for SIP trunks was first added in Cisco Unified Callmanager Release 4.0.

A SIP trunk configuration causes members of a Cisco Unified CallManager cluster to behave as a SIP User Agent, and the cluster expects the other SIP entity defined in the configuration to act as a SIP User Agent also. The SIP trunk is used to configure communication to any SIP entity that can act as a SIP User Agent, which includes SIP proxies, SIP gateways, SIP redirect servers, or other Cisco Unified CallManager clusters. Specifically, it also includes Cisco Unity and Cisco Unified MeetingPlace SIP interfaces.

In Cisco Unified CallManager 4.2 and earlier releases, SIP trunks have the following characteristics:

- Each SIP trunk defined in the system must use a unique incoming port number.
- Each call using a SIP trunk requires a media termination point (MTP) resource.
- The codec for all calls established over a given trunk is restricted to either G.711 mu-law or a-law.
- Q.SIG Tunneling using Annex M1 is not supported.
- The trunk may be used to connect only to devices that implement early media support. In other words, Cisco Unified CallManager will initiate calls by sending an Invite message with Session Description Protocol (SDP), and the far end must be capable of accepting this Invite message.
- Secure Real-Time Transport Protocol (SRTP) is not supported.
- Video calls are not supported.
- TLS is not supported.


**Note**

Refer to the *Cisco Unified Communications SRND for Cisco Unified CallManager 5.0* to understand which of the above restriction no longer apply for that release.

SIP Trunk Redundancy

One of the main benefits of using SIP trunks for intercluster trunking is call survivability. Unlike H.323 trunks in Cisco Unified CallManager Release 3.3 and later, SIP trunks can point to only a single IP address or DNS Server (SRV) record. To provide failover and load balancing for intercluster SIP trunks that are not capable of DNS SRV, configure multiple SIP trunks. In addition, these SIP trunks must be members of route groups and route lists.

It is important to note that Cisco Unified CallManager accepts calls only from a SIP device whose IP address matches one of the destination addresses of the configured SIP trunks. In addition, the incoming port number of the SIP messages has to match the port number configured for that SIP trunk. As a consequence, Cisco recommends that you configure as many SIP trunks with destination addresses as needed to match all IP addresses of any far-end SIP devices that can potentially place an inbound call. This method is not desirable for deployments with more than two Cisco Unified CallManager clusters, and it is better to use SIP trunks with DNS SRV if there are more than two clusters. [Figure 5-6](#) shows the call flow for an intercluster SIP trunk call using DNS SRV.

Figure 5-6 Call Flow for Intercluster SIP Trunk Using DNS SRV

Note: The DNS A Lookup has been removed from this call flow

Figure 5-6 illustrates the following steps in the call flow:

1. The IP phone in Cluster1 calls 87522001.
2. The call matches a route pattern of 8752XXXX that is pointing to SIP Trunk with DNS SRV of cluster2.foo.com. CCM3 in Cluster1 is the node handling this call because the SIP trunk is registered to it. CCM3 sends a DNS SRV lookup for cluster2.foo.com
3. The DNS server replies with two records: CCM-A.cluster2.foo.com and CCM-B.cluster2.foo.com. Because CCM-A.cluster2.foo.com has a higher priority, the call is attempted to this Cisco Unified CallManager. Before sending the SIP Invite, another DNS lookup is done for CCM-A.cluster2.foo.com.
4. CCM3 sends a SIP Invite to 87522001@cluster2.foo.com with destination address set to the IP address of CCM-A.
5. Cisco Unified CallManager interprets this call as a local call because the host portion of the uniform resource identifier (URI) matches the Cluster FQDN enterprise parameter. Cluster2 does not have any SIP trunk configured with a destination of CCM3, so it does a DNS SRV lookup for all domains configured under the SIP trunks with DNS SRV. In this case, the example shows a single trunk with a DNS SRV destination of cluster1.foo.com

6. The DNS server returns two entries, and one of them matches the source IP address of the invite. The cluster accepts the call and extends it to extension 87522001.



CHAPTER 6

Media Resources

Last revised on: February 13, 2008

A media resource is a software-based or hardware-based entity that performs media processing functions on the data streams to which it is connected. Media processing functions include mixing multiple streams to create one output stream (conferencing), passing the stream from one connection to another (media termination point), converting the data stream from one compression type to another (transcoding), echo cancellation, signaling, termination of a voice stream from a TDM circuit (coding/decoding), packetization of a stream, streaming audio (annunciation), and so forth.

Use this chapter to determine which of the media resources described below are needed for your deployment. Also determine whether the resources that are required can be provided by software-based features or whether it is necessary to provision digital signal processors (DSPs) to implement the resources. Although the resources are discussed in individual sections, the same basic resources (DSPs and Cisco IP Voice Media Streaming Application) may be shared to implement higher-level functions.

This chapter focuses on the following features:

- [Voice Termination, page 6-2](#)
- [Audio Conferencing, page 6-7](#)
- [Transcoding, page 6-10](#)
- [Media Termination Point \(MTP\), page 6-12](#)
- [Annunciator, page 6-17](#)
- [Cisco RSVP Agent, page 6-18](#)
- [Cisco IP Voice Media Streaming Application, page 6-19](#)

For more details about music on hold, see [Music on Hold, page 7-1](#).

For details about the dependencies on hardware and software, see the section on [Hardware and Software Capacities, page 6-20](#).

You can control media resources on Cisco Unified CallManager through the use of media resource groups and media resource group lists. You can create pools of resources to control the specific hardware or software that is used. Cisco recommends that you use pools to group resources based on physical location. For design guidelines based on the various call processing models, see the section on [General Design Guidelines, page 6-22](#).

Voice Termination

Voice termination applies to a call that has two call legs, one leg on a time-division multiplexing (TDM) interface and the second leg on a Voice over IP (VoIP) connection. The TDM leg must be terminated by hardware that performs coding/decoding and packetization of the stream. This termination function is performed by digital signal processor (DSP) resources residing in the same hardware module, blade, or platform. All DSP hardware on Cisco TDM gateways is capable of terminating voice streams, and certain hardware is also capable of performing other media resource functions such as conferencing or transcoding (see [Audio Conferencing, page 6-7](#) and [Transcoding, page 6-10](#)).

[Table 6-2](#) through [Table 6-6](#) show the number of calls that each hardware platform can support, which is determined by the type of DSP chipset and the number of DSPs on the hardware. The hardware has either fixed DSP resources that cannot be upgraded or changed, or it has modular DSP resources that can be upgraded. For modular (upgradeable) hardware, [Table 6-2](#) through [Table 6-6](#) also list the maximum number of DSPs per hardware module.

The number of supported calls depends on the computational complexity of the codec used for a call and also on the complexity mode configured on the DSP. Cisco IOS enables you to configure a complexity mode on the hardware module. Some hardware platforms have two complexity modes, medium complexity and high complexity, while other hardware platforms have medium and high complexity as well as flex mode.

Medium and High Complexity Mode

To determine how many calls a module can support, locate the module in [Table 6-2](#) through [Table 6-6](#), and determine the number of DSPs it can provide as well as the desired codec type. For example, an NM-HD-2VE module with three C2510 DSPs configured in flex mode can support eight G.729 calls on each DSP, for a total of 24 calls using flex mode and G.729 codecs. Using G.711 codecs in flex mode, the same hardware could support 48 calls.

As indicated in [Table 6-1](#), if a codec is supported in medium complexity mode, it is also supported in high complexity mode, albeit with fewer supported calls.

You can configure each DSP separately as either medium complexity, high complexity, or flex mode (C5510 only). The DSP treats all calls according to its configured complexity, regardless of the actual complexity of the codec of the call. A resource with configured complexity equal or higher than the actual complexity of the incoming call must be available, or the call will fail. For example, if a call requires a high-complexity codec but the DSP resource is configured for medium complexity mode, the call will fail. However, if a medium-complexity call is attempted on a DSP configured for high complexity mode, then the call will succeed and Cisco IOS will allocate a high-complexity mode resource.

To determine the maximum number of calls supported, find the appropriate row in [Table 6-2](#) through [Table 6-6](#) that contains the desired hardware. Check the medium and high complexity columns in [Table 6-1](#) to verify which complexity modes are capable of handling the desired codecs. Then read the maximum number of supported calls for each DSP at the desired complexity mode.

Flex Mode

Flex mode, available only on hardware platforms that use the C5510 chipset, eliminates the requirement to specify the codec complexity at configuration time. A DSP in flex mode accepts a call of any supported codec type, as long as it has available processing power. The overhead of each call is tracked dynamically via a calculation of processing power in millions of instructions per second (MIPS).

Cisco IOS performs a MIPS calculation for each call received and subtracts MIPS credits from its budget whenever a new call is initiated. The number of MIPS consumed by a call depends on the codec of the call, as shown in the Flex Mode column of [Table 6-1](#). The DSP will allow a new call as long as it has remaining MIPS credits greater than or equal to the MIPS required for the incoming call. The Flex Mode column of [Table 6-1](#) groups the supported codecs by the number of MIPS for a single call (either 15, 30, or 40 MIPS per call) and shows the MIPS budget available for various hardware.

Flex mode has an advantage when calls of multiple codecs must be supported on the same hardware because flex mode can support more calls than when the DSPs are configured as medium or high complexity. However, flex mode does allow oversubscription of the resources, which introduces the risk of call failure if all resources are used. With flex mode it is possible to have fewer DSP resources than with physical TDM interfaces.

For example, each DSP has a budget of 240 MIPS, for a total budget of 720 MIPS per NM-HD-2VE module. For an NM-HDV2 module, the budget per DSP is also 240 MIPS, but refer to [Table 6-2](#) to determine the total number of MIPS available based on your choice and the number of PVDMS.

Compared to medium or high complexity mode, flex mode has the advantage of supporting the most G.711 calls per DSP. In medium complexity mode a DSP can support 8 G.711 calls, while flex mode supports 16 G.711 calls.

DSP Resources for Voice Termination

[Table 6-2](#) through [Table 6-6](#), organized by DSP chipset, provide information on DSP support by platform, DSP density, and the number of voice terminations (or calls) supported per DSP. [Table 6-1](#) lists the codecs supported by the hardware modules in each complexity mode.

Table 6-1 Codecs Supported by Each Complexity Mode

Medium Complexity	High Complexity	Flex Mode
G.711 (a-law, mu-law)	G.711 (a-law, mu-law)	<i>At 15 MIPS per call:</i>
Fax/modem pass-through	Fax/modem pass-through	• G.711 (a-law, mu-law)
Clear channel	Clear channel	• Fax/modem pass-through
G.726 (32K, 24K, 16K)	G.726 (32K, 24K, 16K)	• Clear channel
GSM-FR	GSM-FR	<i>At 30 MIPS per call:</i>
Fax relay	Fax relay	• G.726 (32K, 24K, 16K)
G.729 (a, ab)	G.729	• GSM-FR
	G.729 (a, b, ab)	• Fax relay
	G.728	• G.729
	G.723.1 (32K, 24K, 16K)	• G.729 (a, b, ab)
	G.723.1a (5.3K, 6.3K)	<i>At 40 MIPS per call:</i>
	GSM-EFR	• G.728
	Modem relay	• G.723.1 (32K, 24K, 16K)
		• G.723.1a (5.3K, 6.3K)
		• GSM-EFR
		• Modem relay

Hardware based on the C5510 chipset supports medium and high complexity modes as well as flex mode, as indicated in [Table 6-2](#).

Table 6-2 *DSP Resources on Cisco IOS Hardware Platforms with C5510 Chipset*

Hardware Module or Chassis	DSP Configuration	Maximum Number of Voice Terminations (Calls) per DSP and per Module		
		Medium Complexity (8 calls per DSP)	High Complexity (6 calls per DSP)	Flex Mode ¹ (240 MIPS per DSP)
VG-224	Fixed at 4 DSPs	N/A	24 calls per platform Supported codecs: • G.711 (a-law, mu-law) • G.729a	N/A
NM-HD-1V ²	Fixed at 1 DSP	4 calls per NM	4 calls per NM	240 MIPS per NM
NM-HD-2V	Fixed at 1 DSP	8 calls per NM	6 calls per NM	240 MIPS per NM
NM-HD-2VE	Fixed at 3 DSPs	24 calls per NM	18 calls per NM	720 MIPS per NM
NM-HDV2	1 to 4 of:	Calls per PVDM:	Calls per PVDM:	MIPS per PVDM:
NM-HDV2-2T1/E1	PVDM2-8 ³ (½ DSP)	4	3	120
NM-HDV2-1T1/E1	PVDM2-16 (1 DSP)	8	6	240
	PVDM2-32 (2 DSPs)	16	12	480
	PVDM2-48 (3 DSPs)	24	18	720
	PVDM2-64 (4 DSPs)	32	24	960
2801	1 to 2 of:	Calls per PVDM:	Calls per PVDM:	MIPS per PVDM:
2811	PVDM2-8 ³ (½ DSP)	4	3	120
	PVDM2-16 (1 DSP)	8	6	240
	PVDM2-32 (2 DSPs)	16	12	480
	PVDM2-48 (3 DSPs)	24	18	720
	PVDM2-64 (4 DSPs)	32	24	960
2821	1 to 3 of:	Calls per PVDM:	Calls per PVDM:	MIPS per PVDM:
2851	PVDM2-8 ³ (½ DSP)	4	3	120
	PVDM2-16 (1 DSP)	8	6	240
	PVDM2-32 (2 DSPs)	16	12	480
	PVDM2-48 (3 DSPs)	24	18	720
	PVDM2-64 (4 DSPs)	32	24	960
3825	1 to 4 of:	Calls per PVDM:	Calls per PVDM:	MIPS per PVDM:
3845	PVDM2-8 ³ (½ DSP)	4	3	120
	PVDM2-16 (1 DSP)	8	6	240
	PVDM2-32 (2 DSPs)	16	12	480
	PVDM2-48 (3 DSPs)	24	18	720
	PVDM2-64 (4 DSPs)	32	24	960

1. In flex mode, the maximum number of supported calls depends on the number of MIPS used per call (see [Table 6-1](#)).
2. With the NM-HD-1V module, the number of voice terminations (calls) is limited by the number of physical ports on the module.
3. The PVDM2-8 has a single half-capacity C5510.

Hardware that is based on the C5421 chipset may have DSPs configured as either medium complexity or high complexity. [Table 6-3](#) lists the call density per DSP, and [Table 6-1](#) lists the codecs supported in each complexity mode.

Table 6-3 *DSP Resources on Cisco IOS Hardware Platforms with C5421 Chipset*

Hardware Module	DSP Configuration	Maximum Number of Calls per DSP and per Module	
		Medium Complexity (8 calls per DSP)	High Complexity (8 calls per DSP)
NM-HDA-4FXS	Fixed at 2 DSPs	16 calls per NM	8 calls per NM
	<i>or</i> Fixed at 1 of DSP-HDA-16 (4 DSPs)	16 calls per NM	16 calls per NM
AIM-VOICE-30 AIM-ATM-VOICE-30	Fixed at 4 DSPs	30 or 60 calls per AIM	16 or 30 calls per AIM

Hardware that is based on the C549 chipset may have DSPs configured as either medium complexity or high complexity. [Table 6-4](#) lists the call density per DSP, and [Table 6-1](#) lists the codecs supported in each complexity mode.

Table 6-4 *DSP Resources on Cisco IOS Hardware Platforms with C549 Chipset*

Hardware Module	DSP Configuration	Maximum Number of Calls per DSP and per Module	
		Medium Complexity (4 calls per DSP)	High Complexity (2 calls per DSP)
NM-HDV NM-HDV-FARM	1 to 5 of PVDM-12 (3 DSPs per PVDM-12)	12, 24, 36, 48, or 60 calls per NM	6, 12, 18, 24, or 30 calls per NM
1751 ¹ 1760	1 to 2 of: PVDM-256K-4 (1 DSP) PVDM-256K-8 (2 DSPs) PVDM-256K-12 (3 DSPs) PVDM-256K-16HD (4 DSPs) PVDM-256K-20HD (5 DSPs)	Calls per NM: 4 or 8 8 or 16 12 or 24 16 or 32 20	Calls per NM: 2 or 4 4 or 8 6 or 12 8 or 16 10
PA-VXA-1TE1-24+	Fixed at: 7 DSPs	Calls per PA: 28	Calls per PA: 14
PA-VXA-1TE1-30+	8 DSPs	32	16
PA-VXB-2TE1+	12 DSPs	48	24
PA-VXC-2TE1+	30 DSPs	120	60
PA-MCX-2TE1 PA-MCX-4TE1 PA-MCX-8TE1	Fixed (No on-board DSPs)	Depends on PA-VX(x) ²	Depends on PA-VX(x) ²

1. The 1751 supports a maximum of 8 DSPs (32 channels), and you can order these modules in multiples of 2 PVDMs as long as the total is less than 32 channels. The part number indicates the number of channels.
2. The multichannel port adaptors utilize unused DSPs from PA-VXAs, PA-VXBs, or PA-VXCs across the mix backplane.

Hardware that is based on the C542 chipset supports the following codecs:

- G.711 (a-law, mu-law)
- Fax/modem pass-through
- Clear channel
- G.726 (32K, 24K, 16K)
- GSM-FR
- Fax relay
- G.729
- G.729 (a, b, ab)
- G.728
- G.723.1 (32K, 24K, 16K)
- G.723.1a (5.3K, 6.3K)
- GSM-EFR
- Modem relay

Table 6-5 lists the call density per DSP.

Table 6-5 *DSP Resources on Cisco IOS Hardware Platforms with C542 Chipset*

Hardware Module ¹	DSP Configuration	Maximum Number of Calls per DSP and per Module
NM-1V	Fixed at 2 DSPs	1 call per DSP 2 calls per NM
NM-2V	Fixed at 4 DSPs	1 call per DSP 4 calls per NM

1. These modules do not have any complexity mode but support all codecs equally.

Table 6-6 lists non-IOS hardware for DSP resources. All non-IOS hardware platforms have a fixed configuration of DSPs, as indicated in Table 6-6.

Table 6-6 *DSP Resources on Non-IOS Hardware Platforms*

Hardware Module or Platform	DSP Configuration	Maximum Number of Calls per DSP and per Module	Codecs Supported
WS-6608-T1 WS-6608-E1	Fixed at 64 of C549 (8 DSPs per port)	2 calls per DSP 256 calls per module ¹	G.711 a-law, mu-law G.729a
WS-6624-FXS	Fixed at 12 of C549	2 calls per DSP 24 calls per module	G.711 a-law, mu-law G.729a
VG-248	Fixed at 12 of C5409	4 calls per DSP 48 calls per platform	G.711 a-law, mu-law G.729a
WS-SVC-CMM-ACT	Fixed at 4 of Broadcom 1500	32 calls per DSP 128 calls per module	G.711 (10-30 ms) G.729 (10-60 ms) G.723 (30-60 ms)

Table 6-6 *DSP Resources on Non-IOS Hardware Platforms (continued)*

Hardware Module or Platform	DSP Configuration	Maximum Number of Calls per DSP and per Module	Codecs Supported
WS-SVC-CMM-6T1	Fixed at 12 of C5441	15 calls per DSP 144 calls per module	G.711 (10, 20, 30 ms) G.729 (10, 20, 30, 40, 50, 60 ms)
WS-SVC-CMM-6E1	Fixed at 12 of C5441	15 calls per DSP 180 calls per module	G.711 (10, 20, 30 ms) G.729 (10, 20, 30, 40, 50, 60 ms)
WS-SVC-CMM-24FXS	Fixed at 3 of C5441	15 calls per DSP 24 calls per module	G.711 a-law, mu-law G.729 G.729a
ATA-188 ²	Fixed at 1 of Komodo 3880	2 calls per platform	G.711 a-law, mu-law G.729

1. A maximum of 192 calls are possible with T1 and 240 calls with E1, based on the number of physical ports. If none of the DSPs are configured for T1 or E1, then a maximum of 256 DSP resources are available.
2. The ATA module does not define complexity. It supports G.711, G.729, and G.723 only.

Audio Conferencing

A conference bridge is a resource that joins multiple participants into a single call. It can accept any number of connections for a given conference, up to the maximum number of streams allowed for a single conference on that device. There is a one-to-one correspondence between media streams connected to a conference and participants connected to the conference. The conference bridge mixes the streams together and creates a unique output stream for each connected party. The output stream for a given party is the composite of the streams from all connected parties minus their own input stream. Some conference bridges mix only the three loudest talkers on the conference and distribute that composite stream to each participant (minus their own input stream if they are one of the talkers).

Audio Conferencing Resources

A hardware conference bridge has all the capabilities of a software conference bridge. In addition, some hardware conference bridges can support multiple low bit-rate (LBR) stream types such as G.729, GSM, or G.723. This capability enables some hardware conference bridges to handle mixed-mode conferences. In a mixed-mode conference, the hardware conference bridge transcodes G.729, GSM, and G.723 streams into G.711 streams, mixes them, and then encodes the resulting stream into the appropriate stream type for transmission back to the user. Some hardware conference bridges support only G.711 conferences.

All conference bridges that are under the control of Cisco Unified CallManager use Skinny Client Control Protocol (SCCP) to communicate with Cisco Unified CallManager.

Cisco Unified CallManager allocates a conference bridge from a conferencing resource that is registered with the Cisco Unified CallManager cluster. Both hardware and software conferencing resources can register with Cisco Unified CallManager at the same time, and Cisco Unified CallManager can allocate and use conference bridges from either resource. Cisco Unified CallManager does not distinguish between these types of conference bridges when it processes a conference allocation request.

The number of individual conferences that may be supported by the resource varies, and the maximum number of participants in a single conference varies, depending on the resource.

The following types of conference bridge resources may be used on a Cisco Unified CallManager system:

- [Software Audio Conference Bridge \(Cisco IP Voice Media Streaming Application\)](#), page 6-8
- [Hardware Audio Conference Bridge \(Cisco NM-HDV2, NM-HD-1V/2V/2VE, 2800 Series, and 3800 Series Routers\)](#), page 6-8
- [Hardware Audio Conference Bridge \(Cisco WS-SVC-CMM-ACT\)](#), page 6-9
- [Hardware Audio Conference Bridge \(Cisco NM-HDV and 1700 Series Routers\)](#), page 6-9
- [Hardware Audio Conference Bridge \(Cisco Catalyst WS-X6608-T1 and WS-X6608-E1\)](#), page 6-9
- [Built-in Conference](#), page 6-10

Software Audio Conference Bridge (Cisco IP Voice Media Streaming Application)

A software unicast conference bridge is a standard conference mixer that is capable of mixing G.711 audio streams and Cisco Wideband audio streams. Any combination of Wideband or G.711 a-law and mu-law streams may be connected to the same conference. The number of conferences that can be supported on a given configuration depends on the server where the conference bridge software is running and on what other functionality has been enabled for the application. The Cisco IP Voice Media Streaming Application is a resource that can also be used for several functions, and the design must consider all functions together (see [Cisco IP Voice Media Streaming Application](#), page 6-19).

Hardware Audio Conference Bridge (Cisco NM-HDV2, NM-HD-1V/2V/2VE, 2800 Series, and 3800 Series Routers)

DSPs that are configured through Cisco IOS as conference resources will load firmware into the DSPs that is specific to conferencing functionality only, and these DSPs cannot be used for any other media feature.

The following guidelines and considerations apply to these DSP resources:

- Based on the C5510 DSP chipset, the NM-HDV2 and the router chassis use the PVDM2 modules for providing DSPs.
- DSPs on PVDM2 hardware are configured individually as either voice termination, conferencing, media termination, or transcoding, so that DSPs on a single PVDM may be used as different resource types. Allocate DSPs to voice termination first, then to other functionality as needed.
- The NM-HDV2 has 4 slots that will accept PVDM2 modules in any combination. The other network modules have fixed numbers of DSPs.
- A conference based on these DSPs allows a maximum of 8 participants. When a conference begins, all 8 positions are reserved at that time.
- The PVDM2-8 is listed as having $\frac{1}{2}$ a DSP because it has a DSP that has half the processing capacity of the PVDM2-16. For example, if the DSP on a PVDM2-8 is configured for G.711, it can provide $(0.5 * 8)$ bridges/DSP = 4 conference bridges.
- Use [Table 6-1](#) and [Table 6-2](#) to determine how many DSPs may be provisioned with specific hardware.
- A DSP farm configuration in Cisco IOS specifies which codecs may be accepted for the farm. A DSP farm that is configured for conferencing and G.711 provides 8 conferences. When configured to accept both G.711 and G.729 calls, a single DSP provides 2 conferences because it is also reserving its resources for performing transcoding of streams.
- The I/O of an NM-HDV2 is limited to 400 streams, so ensure that the number of conference resources allocated does not cause this limit to be exceeded. If G.711 conferences are configured, then no more than 48 DSPs should be allocated per NM because $(48 * 8)$ participants = 384 streams.

If you configure all conferencing for both G.711 and G.729 codecs, then each DSP provides only 2 conferences of 8 participants each. In this case, it is possible to populate the NM fully and configure it with 16 DSPs so there would be 256 streams.

- Conferences cannot natively accept calls utilizing the GSM codec. A transcoder must be provided separately for these calls to participate in a conference.
- Any PVDM2-based hardware, such as the NM-HDV2, may be used simultaneously in a single chassis for voice termination but may not be used simultaneously for other media resource functionality. The DSPs based on PVDM-256K and PVDM2 have different DSP farm configurations, and only one may be configured in a router at a time.

Hardware Audio Conference Bridge (Cisco WS-SVC-CMM-ACT)

The following guidelines and considerations apply to this DSP resource:

- DSPs on this hardware are configured individually as either voice termination, conferencing, media termination, or transcoding, so that DSPs on a single module may be used as different resource types. Allocate DSPs to voice termination first.
- This Cisco Catalyst-based hardware provides DSP resources that may provide conference bridges of up to 32 participants per bridge.
- Each module contains 4 DSPs that are individually configurable. Each DSP can support 32 conference bridges.
- The G.711 and G.729 codecs are supported on these conference bridges without extra transcoder resources. However, transcoder resources would be necessary if other codecs are used.

Hardware Audio Conference Bridge (Cisco NM-HDV and 1700 Series Routers)

The following guidelines and considerations apply to these DSP resources:

- This hardware utilizes the PVDM-256K type modules that are based on the C549 DSP chipset.
- Conferences using this hardware provide bridges that allow up to 6 participants in a single bridge.
- The resources are configured on a per-DSP basis as conference bridges.
- The NM-HDV may have up to 4 PVDM-256K modules, while the Cisco 1700 Series Routers may have 1 or 2 PVDM-256K modules.
- Each DSP provides a single conference bridge that can accept G.711 or G.729 calls.
- The Cisco 1751 is limited to 5 conference calls per chassis, and the Cisco 1760 can support 20 conference calls per chassis.
- Any PVDM2-based hardware, such as the NM-HDV2, may be used simultaneously in a single chassis for voice termination but may not be used simultaneously for other media resource functionality. The DSPs based on PVDM-256K and PVDM2 have different DSP farm configurations, and only one may be configured in a router at a time.

Hardware Audio Conference Bridge (Cisco Catalyst WS-X6608-T1 and WS-X6608-E1)

The following guidelines and considerations apply to these DSP resources:

- This hardware has 8 DSPs that are physically associated to each port, and there are 8 ports per card. Configuration of the DSPs is at the port level, so that all DSPs associated to a port perform the same function.
- Conference bridges may have up to 32 participants, and each port supports 32 conference bridges.
- For conferences with G.711 or G.723 there may be 32 conferences per port. If G.729 calls are used, there may be 24 conferences per port.

Built-in Conference

Some phone models have a built-in conference resource that allows a three-way conference. This bridge is invoked only by the Barge feature and is not used as a general conferencing resource. For details on which phones have this bridge, see [IP Telephony Endpoints, page 21-1](#). This bridge accepts only G.711 calls.

Transcoding

A transcoder is a device that converts an input stream from one codec into an output stream that uses a different codec. It may also connect two streams that utilize the same codec but with a different sampling rate. In a Cisco Unified CallManager system, the typical use of a transcoder is to convert between a G.711 voice stream and the low bit-rate compressed voice stream G.729a. The following cases determine when transcoder resources are needed:

- Single codec for the entire system

When a single codec is configured for all calls in the system, then no transcoder resources are required. The G.711 codec is supported by all vendors. A single-site deployment usually has no need for conserving bandwidth, and a single codec can be used. In this scenario, G.711 is the most common choice.

- Multiple codecs in use in the system, and all endpoints are capable of all codec types

The most common reason for multiple codecs is to use G.711 for LAN calls to maximize the call quality and to use a low-bandwidth codec to maximize bandwidth efficiency for calls that traverse a WAN with limited bandwidth. Cisco recommends using G.729a as the low-bandwidth codec because it is supported on all Cisco Unified IP Phone models as well as most other Cisco Unified Communications devices, therefore it can eliminate the need for transcoding. Although Cisco Unified CallManager allows configuration of other low-bandwidth codecs between regions, the current phone models do not support those codecs and would require transcoders. They would require one transcoder for a call to a gateway and two transcoders if the call is to another IP phone. The use of transcoders is avoided if all devices support and are configured for both G.711 and G.729 because the devices will use the appropriate codec on a call-by-call basis.

- Multiple codecs in use in the system, and some endpoints support or are configured for G.711 only

This condition exists when G.729a is used in the system but there are devices that do not support this codec, or a device with G.729a support may be configured to not use it. In this case, a transcoder is also required. Devices from some third-party vendors may not support G.729. Another example where G.729 is supported but might not be configured is with Cisco Unity. Cisco Unity has support for accepting calls with G.729a, but the codec is implemented in software and is CPU-intensive. Because as few as 10 simultaneous calls can cause significant CPU utilization, many deployments choose to disable G.729 on Cisco Unity and off-load the transcoding function to a dedicated transcoding resource external to the Unity server. If your system includes Cisco Unity, determine whether Unity will accept G.729a calls or whether it will be configured for G.711 only.

**Note**

Cisco Unified MeetingPlace Express currently supports G.711 only. In an environment where G.729 is configured for a call into Cisco Unified MeetingPlace Express, transcoder resources are required.

To finalize the design, it is necessary to know how many transcoders are needed and where they will be placed. If multiple codecs are needed, it is necessary to know how many endpoints do not support all codecs, where those endpoints are located, what other groups will be accessing those resources, how many maximum simultaneous calls these device must support, and where those resources are located in the network.

Transcoding Resources

DSP resources are required to perform transcoding. Those DSP resources can be located in the voice modules and the hardware platforms for transcoding that are listed in the following sections.

Hardware Transcoder (Cisco NM-HDV2, NM-HD-1V/2V/2VE, 2800, and 3800 Series Routers)

The following guidelines and considerations apply to these DSP resources:

- Transcoding is available between G.711 mu-law or a-law and G.729a or G.729ab. Each DSP can support 8 sessions.
- Cisco Unified IP Phones use only the G.729a variants of the G.729 codecs. The default for a new DSP farm profile is G.729a/G.729ab/G.711u/G.711a. Because a single DSP can provide only one function at a time, the maximum sessions configured on the profile should be specified in multiples of 8 to prevent wasted resources.
- Transcoding is also available between G.711mu-law/G.711a-law and G.729/G.729b, but is not usually used in a Cisco Unified CallManager system. Each DSP can support 6 sessions.
- Refer to the section on [DSP Resources for Voice Termination, page 6-3](#), to determine numbers of DSPs available on a particular platform or network module.

Hardware Transcoder (Cisco WS-SVC-CMM-ACT)

The following guidelines and considerations apply to this DSP resource:

- Transcoding is available between G.711 mu-law or a-law and G.729, G.729b, or G.723.
- There are 4 DSPs per ACT that may be allocated individually to DSP pools.
- The CCM-ACT can have 16 transcoded calls per DSP or 64 per ACT. The ACT reports resources as streams rather than calls, and a single transcoded call consists of two streams.

Hardware Transcoder (Cisco NM-HDV and 1700 Series Routers)

The following guidelines and considerations apply to these DSP resources:

- This hardware utilizes the PVDM-256K type modules that are based on the C549 DSP chipset.
- The NM-HDV may have up to 4 PVDM-256K modules, and the Cisco 1700 series routers may have 1 or 2 PVDM-256K modules.
- NM-HDV and NM-HDV2 modules may be used simultaneously in a single chassis for voice termination but may not be used simultaneously for other media resource functionality. Only one type of DSP farm configuration may be active at one time (either the NM-HDV or the NM-HDV2) for conferencing, MTP, or transcoding.
- Transcoding is supported from G.711 mu-law or a-law to any of G.729, G.729a, G.729b, G.729ab, or GSM codecs.
- Each DSP can provide 2 transcoding sessions.
- The Cisco 1751 has a chassis limit of 16 sessions, and the Cisco 1760 has a chassis limit of 20 sessions.

Hardware Transcoder (Cisco WS-X6608)

The following guidelines and considerations apply to this DSP resource:

- DSPs are allocated to functions at a port level. Each port can provide 24 transcoding sessions.
- There are 8 ports per blade.
- Transcoding is available between G.711 mu-law or a-law and G.729a, G.729ab, G.729, or G.729b.

A transcoder is also capable of performing the same functionality as a media termination point (MTP). In cases where transcoder functionality and MTP functionality are both needed, a transcoder is allocated by the system. If MTP functionality is required, Cisco Unified CallManager will allocate either a transcoder or an MTP from the resource pool, and the choice of resource will be determined by the media resource groups, as described in the section on [Media Resource Groups and Lists, page 6-22](#).

Media Termination Point (MTP)

A media termination point (MTP) is an entity that accepts two full-duplex G.711 streams. It bridges the media streams together and allows them to be set up and torn down independently. The streaming data received from the input stream on one connection is passed to the output stream on the other connection, and vice versa. MTPs have many possible uses, as described in the following sections.

Re-Packetization of a Stream

An MTP can be used to transcode a-law to mu-law and vice versa, or it can be used to bridge two connections that utilize different packetization periods (different sample sizes).

H.323 Supplementary Services

MTPs can be used to extend supplementary services to H.323 endpoints that do not support the H.323v2 OpenLogicalChannel and CloseLogicalChannel request features of the Empty Capabilities Set (ECS). This requirement occurs infrequently. Cisco H.323 endpoints support ECS, and most third party endpoints have support as well. When needed, an MTP is allocated and connected into a call on behalf of an H.323 endpoint. Once inserted, the media streams are connected between the MTP and the H.323 device, and these connections are present for the duration of the call. The media streams connected to the other side of the MTP can be connected and disconnected as needed to implement features such as hold, transfer, and so forth.

When an MTP is required on an H.323 call and none is available, the call will proceed but will not be able to invoke supplementary services.

H.323 Outbound Fast Start

H.323 defines a feature called Fast Start, which reduces the number of packets exchanged during a call setup, thereby reducing the amount of time for media to be established. It is useful when two devices that are utilizing H.323 have high network latency between them because the time to establish media is dependant on that latency. Cisco Unified CallManager distinguishes between inbound and outbound fast start based on the direction of the call setup, and the distinction is important because the MTP

requirements are not equal. For inbound fast start, no MTP is required. Outbound calls on an H.323 trunk do require an MTP when fast start is enabled. Frequently it is only inbound calls that are problematic, and it is possible to use inbound fast start to solve the issue without also enabling outbound fast start.

Named Telephony Events (RFC 2833)

DTMF tones can be used during a call to signal to a far-end device for purposes of navigating a menu system, entering data, or other manipulation. They are processed differently than DTMF tones sent during a call setup as part of the call control.

Named Telephony Events (NTEs) defined by RFC 2833 are a method of sending DTMF from one endpoint to another after the call media has been established. The tones are sent as packet data using the already established RTP stream and are distinguished from the audio by the RTP packet type field. For example, the audio of a call can be sent on a session with an RTP packet type that identifies it as G.711 data, and the DTMF packets are sent with an RTP packet type that identifies them as NTEs. The consumer of the stream utilizes the G.711 packets and the NTE packets separately.

When are Media Termination Points Required for Named Telephony Events?

For *Cisco Unified CallManager 4.x*, only SIP trunks are supported, and they require that an MTP be allocated for all calls using the SIP trunk. The SIP trunk has a configuration parameter, "MTP required," which is selected by default and cannot be changed.

Cisco Unified CallManager 5.0 adds support for SIP on line devices and also removes the requirement that MTPs be allocated for all calls using SIP trunks. In Cisco Unified CallManager 5.0, an MTP is required when two endpoints do not have a method in common for sending DTMF between them or when the system configuration specifies that one must be allocated. The following description expands on the details and applies only to Cisco Unified CallManager 5.0.

Verify the types of endpoints planned for the system based on the following rules:

1. Two non-SIP endpoints do not require MTPs.

All Cisco Unified Communications endpoints other than SIP send DTMF to Cisco Unified CallManager via various signaling paths, and Cisco Unified CallManager forwards the DTMF between dissimilar endpoints. For example, an IP phone may use SCCP messages to Cisco Unified CallManager to send DTMF, which then gets sent to an H.323 gateway via H.245 signaling events. The two endpoints have a common method by sending DTMF to Cisco Unified CallManager. If no SIP endpoints are used, stop here.

2. Two Cisco SIP endpoints do not require MTPs.

All Cisco SIP endpoints support NTEs, therefore MTPs can be avoided entirely. DTMF is sent directly between the endpoints using NTE. If all endpoints are Cisco SIP devices, then no MTPs are required for DTMF conversion.

3. A combination of a SIP endpoint and a non-SIP endpoint might require MTPs.

To determine the support for NTE in your devices, see [Table 6-7](#). RFC 2833 is not limited to SIP and can be supported in devices with other call control protocols. For example, Cisco Unified IP Phones that run either an SCCP or SIP stack can support NTEs in both modes. Some devices support DTMF via multiple methods (for example, a Cisco Unified IP Phone 7960 with an SCCP stack can send NTEs to the other device or SCCP to Cisco Unified CallManager). Other devices such as the Cisco Wireless IP Phone 7920 can send only SCCP, and still others can send only NTE (such as the Cisco Unified IP Phone 7960 with a SIP stack).

Cisco Unified CallManager has the ability to allocate MTPs dynamically on a call-by-call basis, based on the capabilities of the pair of endpoints. Use [Table 6-7](#) to determine whether MTPs should be provisioned.

4. Dynamic allocation of MTPs is only for outbound calls.

A SIP trunk that does not require MTPs will use dynamic allocation as described in rule 3. The SIP trunk is capable of accepting calls with or without Session Description Protocol (SDP) in the Invite, therefore it will not allocate an MTP for an inbound call. Only outbound calls may have an MTP allocated as described above. If a particular SIP trunk is used for inbound calls only, then it not necessary to allocate any MTP resources on the system for receiving calls on that trunk.

5. MTPs may be forced by the configuration.

In Cisco Unified CallManager 5.0, the SIP trunk parameter **MTP Required** is not selected by default, and the field is unlocked. When a SIP trunk is defined for a Cisco device (such as a SIP gateway), the device should be configured for NTE support. The default setting for MTP Required should normally be used, and MTPs are allocated only when needed.

There are other reasons for forcing an MTP. The establishment of a SIP dialog behaves differently when an MTP is allocated before call setup starts. SIP uses Session Description Protocol (SDP) for establishing the parameters of a session, and SDP is embedded into SIP messages. When an MTP is forced, the SDP is sent with the Invite message that initiates the call. When an MTP is not forced, it is allocated (if needed) after the Invite message. The Invite message does not contain SDP, and it is sent at a later point in the call setup. If the far-end device supports only Invite messages with embedded SDP, then the MTP Required parameter must be checked (enabled).

When this configuration parameter forces an MTP, the codec used on the call leg from the MTP to the far-end SIP device must be G.711 mu-law or a-law. There is an additional parameter on the trunk configuration that becomes unlocked when MTP Required is selected. This parameter allows a choice of which variant of G.711 to use. When an MTP is not forced by this configuration parameter, then the normal mechanisms for determining the codec apply.



Note

When an MTP is placed into a media stream by configuration of **MTP Required**, the following MTP resources may be used: the Cisco IP Voice Media Streaming Application or the CMM-ACT modules with Cisco IOS Release 12.4(6)T or later. This limitation is a result of endpoints that send DTMF by multiple methods (for example, a Cisco Unified IP Phone 7960 sends both SCCP and NTE events simultaneously). Only the MTPs listed here currently handle this situation.

Table 6-7 **Endpoint Support for DTMF Methods**

Endpoint Protocol Stack: Endpoints	DTMF Method		
	SCCP	NTE	KPML ¹
SCCP Stack:	Yes	No	No
12SP+, 30 VOIP,			
7910, 7920, 7935, 7936			
VG248,			
DPA-7610, DPA-7630,			
CTI Port, First Party Control			

Table 6-7 *Endpoint Support for DTMF Methods (continued)*

Endpoint Protocol Stack: Endpoints	DTMF Method		
	SCCP	NTE	KPML ¹
SCCP Stack: 7902, 7905, 7912, 7940, 7941, 7960, 7961, 7970 Future new phone models	Yes	Yes	No
SIP Stack: 7905, 7911, 7912, 7940, 7941, 7960, 7961, 7970 Future new phone models	No	Yes	Yes for: 7911, 7941, 7961, 7970 No for: 7905, 7912, 7940, 7960

1. Key Press Markup Language (KPML)

**Note**

IP Phones play DTMF to the end user when DTMF is received via SCCP, but they do not play tones received by NTE. However, there is no requirement to send DTMF to another end user. It is necessary only to consider the endpoints that originate calls combined with endpoints that might need DTMF, such as PSTN gateways, application servers, and so forth.

Example 6-1 *Call Flow that Requires an MTP for NTE Conversion*

Assume the example system has CTI route points with first-party control (the CTI port terminates the media), which integrate to a system that uses DTMF to navigate an IVR menu. If all phones in the system are running SCCP, then no MTP is required. In this case Cisco Unified CallManager controls the CTI port and receives DTMF from the IP phones via SCCP. Cisco Unified CallManager provides DTMF conversion.

However, if there are phones running a SIP stack, an MTP is required. NTEs are part of the media stream, therefore Cisco Unified CallManager does not receive them. An MTP is invoked into the media stream and has one call leg that uses SCCP, and the second call leg uses NTEs. The MTP is under SCCP control of Cisco Unified CallManager and performs the NTE-to-SCCP conversion under the control of Cisco Unified CallManager.

Configuration of DTMF on SIP and H.323 Gateways

Configure either **sip-notify** or **rtp-nte** as a method under a SIP dial peer. The best configuration of a SIP dial peer will depend on the mix of endpoints that exist in the system (see [Table 6-7](#)). When Unsolicited Notify is used on the gateway, that dial peer may be used by endpoints that support SCCP without the need for MTP resources. Endpoints that support only NTE must invoke an MTP to use the gateway.

Conversely, if NTE is configured on the SIP dial peer, then any endpoint that can use NTE can send DTMF directly to the gateway without need for MTPs. An endpoint that supports only SCCP will need to invoke an MTP.

SIP gateways have the ability to send DTMF using NTEs, Unsolicited Notify, or audio tones in the media stream. Unsolicited Notify is a Cisco proprietary method that sends a SIP Notify message with an event that contains the DTMF tone. This method is also supported on Cisco Unified CallManager.

The following example lists a SIP gateway configuration for Named Telephony Events:

```
dial-peer voice 10 voip
dtmf-relay rtp-nte
```

Cisco IOS SIP gateways do not currently support Key Press Markup Language (KPML).

H.323 gateways have support for H.245 Alphanumeric, H.245 Signal, NTE, and audio in the media stream. The NTE option must not be used because it is not supported on Cisco Unified CallManager for H.323 at this time. The preferred option is H.245 signal. MTPs are required for establishing calls to an H.323 gateway if the other endpoint does not have signaling capability in common with Cisco Unified CallManager. For example, a Cisco Unified IP Phone 7960 running the SIP stack supports only NTEs, so an MTP is needed with an H.323 gateway.

The following configuration is recommended for DTMF methods on an H.323 gateway:

```
dial-peer voice 10 voip
dtmf-relay h.245-signal
```



Note

SIP devices may also support another method of sending DTMF, called Key Press Markup Language (KPML). Although many Cisco phones do support KPML, it is not yet widely supported. Cisco IOS SIP gateways do not support KPML as of Cisco IOS Release 12.4(4)T, but they will soon be adding support. At that time, KPML may be used in place of Unsolicited Notifies on the gateway.

CTI Route Points

CTI Route Points that have first-party control of a phone call will participate in the media stream of the call and require an MTP to be inserted. When the CTI has third-party control of a call such that the media passes through a device that is controlled by the CTI, then the requirement for an MTP is dependant on the capabilities of the controlled device.

MTP Resources

The following types of devices are available for use as an MTP:

Software MTP (Cisco IP Voice Media Streaming Application)

A software MTP is a device that is implemented by installing the Cisco IP Voice Media Streaming Application on a server. When the installed application is configured as an MTP application, it registers with a Cisco Unified CallManager node and informs Cisco Unified CallManager of how many MTP resources it supports. A software MTP device supports only G.711 streams. The IP Voice Media Streaming Application is a resource that may also be used for several functions, and the design guidance must consider all functions together (see [Cisco IP Voice Media Streaming Application, page 6-19](#)).

Software MTP (Based on Cisco IOS)

- The capability to provide a software-based MTP on the router is available beginning with Cisco IOS Release 12.3(11)T for the Cisco 3800 Series Routers and Release 12.3(8)T4 for other router models.
- This MTP allows configuration of any of the following codecs, but only one may be configured at a given time: G.711 mu-law and a-law, G.729a, G.729, G.729ab, G.729b, GSM, and pass-through. Some of these are not pertinent to a Cisco Unified CallManager implementation.

- The router configuration permits up to 500 individual streams, which support 250 transcoded sessions. This number of G.711 streams generates 5 Mbytes of traffic.

Hardware MTP (Cisco NM-HDV2, NM-HD-1V/2V/2VE, 2800 and 3800 Series Routers)

- This hardware uses the PVDM-2 modules for providing DSPs.
- Each DSP can provide 16 G.711 mu-law or a-law MTP sessions or 6 G.729, G.729b, or GSM MTP sessions.

Hardware MTP (Cisco WS-SVC-CMM-ACT)

- This module has 4 DSPs that may be configured individually.
- Each DSP can support 128 G.729, G.729b, or GSM MTP sessions or 256 G.711 mu-law or a-law MTP sessions.

Hardware MTP (Catalyst WS-X6608-T1 and WS-X6608-E1)

- Codec support is G.711 mu-law or a-law, G.729, G.720b, or GSM.
- Configuration is done at the port level. Eight ports are available per module.
- Each port configured as an MTP resources provides 24 sessions.

Annunciator

An annunciator is a software function of the Cisco IP Voice Media Streaming Application that provides the ability to stream spoken messages or various call progress tones from the system to a user. It is capable of sending multiple one-way RTP streams to devices such as Cisco IP phones or gateways, and it uses SCCP messages to establish the RTP stream. The device must be capable of SCCP to utilize this feature. Tones and announcements are predefined by the system. The announcements support localization and also may be customized by replacing the appropriate .wav file. The annunciator is capable of supporting G.711 a-law and mu-law, G.729, and Wideband codecs without any transcoding resources.

The following features require an annunciator resource:

- Cisco Multilevel Precedence Preemption (MLPP)

This feature has streaming messages that it plays in response to the following call failure conditions.

- Unable to preempt due to an existing higher-precedence call.
- A precedence access limitation was reached.
- The attempted precedence level was unauthorized.
- The called number is not equipped for preemption or call waiting.

- Integration via SIP trunk

SIP endpoints have the ability to generate and send tones in-band in the RTP stream. Because SCCP devices do not have this ability, an annunciator is used in conjunction with an MTP to generate or accept DTMF tones when integrating with a SIP endpoint. The following types of tones are supported:

- Call progress tones (busy, alerting, and ringback)
- DTMF tones

- Cisco IOS gateways and intercluster trunks

These devices require support for call progress tone (ringback tone).

- System messages

During the following call failure conditions, the system plays a streaming message to the end user:

- A dialed number that the system cannot recognize
- A call that is not routed due to a service disruption
- A number that is busy and not configured for preemption or call waiting

- Conferencing

During a conference call, the system plays a barge-in tone to announce that a participant has joined or left the bridge.

An annunciator is automatically created in the system when the Cisco IP Voice Media Streaming Application is activated on a server. If the Media Streaming Application is deactivated, then the annunciator is also deleted. A single annunciator instance can service the entire Cisco Unified CallManager cluster if it meets the performance requirements (see [Annunciator Performance, page 6-18](#)); otherwise, you must configure additional annunciators for the cluster. Additional annunciators can be added by activating the Cisco IP Voice Media Streaming Application on other servers within the cluster.

The annunciator registers with a single Cisco Unified CallManager at a time, as defined by its device pool. It will automatically fail over to a secondary Cisco Unified CallManager if a secondary is configured for the device pool. Any announcement that is playing at the time of an outage will not be maintained.

An annunciator is considered a media device, and it can be included in media resource groups (MRGs) to control which annunciator is selected for use by phones and gateways.

Annunciator Performance

By default, the annunciator is configured to support 48 simultaneous streams, which is the maximum recommended for an annunciator running on the same server (co-resident) with the Cisco Unified CallManager service. If the server has only 10 Mbps connectivity, lower the setting to 24 simultaneous streams.

A standalone server without the Cisco CallManager service can support up to 255 simultaneous announcement streams, and a high-performance server with dual CPUs and a high-performance disk system can support up to 400 streams. You can add multiple standalone servers to support the required number of streams.

Cisco RSVP Agent

In order to provide topology-aware call admission control, Cisco Unified CallManager invokes one or two RVSP Agents during the call setup to perform an RSVP reservation across the IP WAN. These agents are MTP or transcoder resources that have been configured to provide RSVP functionality. RSVP resources are treated the same way as regular MTPs or transcoders from the perspective of allocation of an MTP or transcoder resource by Cisco Unified CallManager.

The Cisco RSVP Agent feature was first introduced in Cisco IOS Release 12.4(6)T. For details on RSVP and Cisco RSVP Agents, refer to the chapter on [Call Admission Control, page 9-1](#).

Cisco IP Voice Media Streaming Application

The Cisco IP Voice Media Streaming Application provides the following resources in software:

- Music on Hold (MoH)
- Annunciator
- Software conference bridge
- Media termination point (MTP)

When the Media Streaming Application is activated, one of each of the above resources is automatically configured. If the annunciator, software conference bridge, or MTP are not needed, Cisco recommends that you disable them by disabling the Run Flag service parameter for the Cisco IP Voice Media Streaming Application.

Give careful consideration to situations that require multiple resources and to the load they place on the Media Streaming Application. Each resource has a service parameter that controls the maximum number of connections it can handle, with an associated default setting. You may run limited amounts of all four resources on the same server if you have not changed the default settings. However, if your deployment requires more than the default number of any one resource, configure that resource to run on its own dedicated server (without any other resources or the Cisco CallManager service running on that server).

The annunciator is the only media resource that is available only through the IP Voice Media Streaming Application. Conferencing, MTPs, and music on hold (MoH) can all reside on servers other than the Cisco Unified CallManager servers. Cisco highly recommends that you disable the MTP and conferencing resources on Cisco Unified CallManager and that you provide external dedicated resources for this functionality.

Cisco also strongly recommends that you run the IP Voice Media Streaming Application on a server other than the publisher or any Cisco Unified CallManager server that provides call processing. The increase in CPU load for media resources might adversely impact call processing performance, and security issues can arise because User Datagram Protocol (UDP) traffic must be received on the Cisco Unified CallManager server.

Hardware and Software Capacities

This section provides data on capacities of the network modules and chassis that carry DSPs, the capacities of the chassis to carry network modules, and software dependencies of the hardware.

PVDMs

Table 6-8 and Table 6-9 show the number of DSPs that can reside on the two models of PVDMs or on fixed-configuration network modules. The PVDM2-xx modules are newer than the PVDM-256K-xx modules, and the two types are not interchangeable.

Table 6-8 *Number of DSPs per PVDM-256K Module*

Module	Number of DSPs
PVDM-256K-4	1 DSP
PVDM-256K-8	2 DSPs
PVDM-256K-12	3 DSPs
PVDM-256K-16HD	4 DSPs
PVDM-256K-20HD	5 DSPs

Table 6-9 *Number of DSPs per PVDM2 Module or Fixed-Configuration Hardware*

Hardware Module or Chassis	Number of DSPs
PVDM2-8	½ DSP
PVDM2-16	1 DSP
PVDM2-32	2 DSPs
PVDM2-48	3 DSPs
PVDM2-64	4 DSPs
NM-HD-1V	1 DSP
NM-HD-2V	
NM-HD-2VE	3 DSPs

Table 6-10 specifies minimum versions of Cisco IOS software needed to support media resource functionality on these hardware platforms and network modules.

Table 6-10 *PVDM2 Slots Available and Cisco IOS Versions Required for Media Support*

Chassis or Network Module	Number of PVDM2 Slots	Minimum Cisco IOS Release for Media
2801	2	12.3(11)T
2811	2	12.3(8)T4
2821 or 2851	3	12.3(8)T4
3825 or 3845	4	12.3(11)T
NM-HDV2	4	

Cisco 2800 and 3800 Series Platforms

Although the Cisco 2800 and 3800 Series Routers all have two AIM slots, they do not support the AIM-VOICE-30 or AIM-ATM-VOICE-30 cards because the functionality of these cards is provided instead by PVDM2 modules that are installed on the motherboard.

Network Modules

You can install the NM-HDV2, NM-HD-xx, and NM-HDV modules in the Cisco IOS platforms listed in [Table 6-11](#), up to the maximum numbers indicated.

All three families of modules in Table 6-11 may be installed in a single chassis. However, the conferencing and transcoding features cannot be used simultaneously on both the NM-HDV family and either of the other two families (NM-HD-xx or NM-HDV2). In addition, the NM-HDV (TI-549), NM-HD-xx, and NM-HDV2 (TI-5510) cannot be used simultaneously for conferencing and transcoding within a single chassis.

You can mix NM-HDV and NM-HDV-FARM modules in the same chassis. Not all chassis can be completely populated by these modules. Table 6-11 shows the maximum number of modules that each type of hardware platform can support.

Table 6-11 *Number of Module Slots per Platform Type*

Cisco IOS Platform	Number of Slots
2691, 2811, 2821, 2851	1
3620 ¹ , 3725, 3825	2
3640	3
3745, 3845	4
3660	6

1. Although it has two NM slots, the Cisco 3620 Router supports only a single NM-HDV module.

Table 6-12 *Modules Supported by Platform Type*

Cisco IOS Platform	Modules Supported by Platform		
	NM-HDV2	NM-HD-1V NM-HD-2V NM-HD-2VE	NM-HDV NM-HDV-FARM
VG2001 26002 36203 3640	No	No	Yes
3660	No	Yes	Yes
2600XM, 2691, 3725, 3745 2811, 2821, 2851 3825, 3845	Yes	Yes	Yes

**Note**

The Cisco VG200, 2620, 2621, and 3620 do not support the NM-HDV-FARM, nor do they support MTP, conferencing, or transcoding. The Cisco 2801 has no NM slot.

Calculating DSP Requirements for the NM-HDV

There are specific situations where the NM-HDV may not be fully populated. The sampling rate is not normally changed from the system defaults. If you do not require the sampling rate to be changed, then this issue may be ignored.

For sample rates of 20, 30, 40, or 60 ms with voice activity detection (VAD) enabled or disabled (or 10 ms with VAD enabled), it is possible to configure an NM-HDV or NM-HDV-FARM with a full complement of 5 PVDMs, giving 60 usable DSP resources.

For a sampling rate of 10 ms with VAD disabled, it is not possible to utilize all DSPs on a fully populated NM-HDV. The following additional calculation is required to ensure that the packet rate does not exceed 6600 packets per second (pps), which is the capacity of the NM-HDV:

$$100 \text{ pps (number of voice terminations)} + 600 \text{ pps (number of conferences)} + 200 \text{ pps (number of transcoding sessions)} < 6600 \text{ pps}$$

General Design Guidelines

The Cisco Unified CallManager constructs of media resource groups (MRGs) and media resource group lists (MRGLs) are used to control how the resources described in this chapter are organized and accessed. This section discusses considerations for how to utilize these constructs effectively. It also discusses specific considerations for the various Cisco Unified CallManager deployment models.

Media Resource Groups and Lists

Media resource groups and lists provide a method to control how resources are allocated that could include rights to resources, location of resources, or resource type for specific applications. This section assumes you have an understanding of media resource groups, and it highlights the following design considerations:

- The system defines a default media resource group that is not visible in the user interface and that all resources are members of when they are created. Consumers of media resources use resources first from any media resource group (MRG) or media resource group list (MRGL) that their configuration specifies. If the required resource is not available, the default MRG is searched for the resource. For simple deployments, the default MRG alone may be used.
- When using MRGs to control access to resources, it is necessary to move the resources out of the default MRG by explicitly configuring them in some other group. If the desired effect is for resources to be available only as a last resort for all calls, then the resources may remain in the default group. Also, if no control over resources is necessary, they may remain in the default group.
- An MRG may contain multiple types of resources, and the appropriate resource will be allocated from the group based on the feature needed. MTPs and transcoders are a special case because a transcoder may also be used as an MTP.
- One use of MRGs is to group resources of similar types. Conference bridge resources vary in the number of participants they support, and MRGs could be used to group the conference resources by conference bridge size.

- Use media resource groups (MRGs) and media resource group lists (MRGLs) to provide sharing of resources across multiple Cisco Unified CallManagers. If you do not use MRGs and MRGLs, the resources are available to a single Cisco Unified CallManager only.
- You can also use MRGs and MRGLs to separate resources based on geographical location, thereby conserving WAN bandwidth whenever possible.
- MRGLs will use MRGs in the order that they are listed in the configuration. If one MRG does not have the needed resource, the next MRG is searched. If all MRGs are searched and no resource is found, the search terminates.
- The algorithm for allocating like resources from an MRG attempts to spread the load across the like resources. When a resource has been used, a pointer for that MRG is incremented to the next device. A device may be present in more than one MRG, which will affect the pointers of all groups in which the device is a member. When MTPs are required and transcoders exist in the same group, the MTPs are always allocated until all MTPs are used, and then transcoders are used as MTPs. When resources of differing capacities exist in the same group, the load sharing attempts to allocate resources based on the capacity. The system will spread the load across resources, but because of the above factors, it frequently will not be round-robin in behavior.
- Cisco Unified CallManager Administration displays the devices in an MRG in alphabetical order, but they are allocated based on their order in the configuration database. You cannot change this order. If you want media resources to be allocated in a specific order, then create a separate MRG for each individual resource and use MRGLs to specify the order of allocation.
- Ensure that the media resources themselves have configurations that prevent further invocation of other media resources. For example, if an MTP is inserted into a call and the codec configured on that MTP does not match the one needed by Cisco Unified CallManager for the call, then a transcoder may also be invoked. A frequent mistake is to configure an MTP for G.729 or G.729b when Cisco Unified CallManager needs G.729a.

Deployment Models

This section examines where and when the MTP and transcoding resources are used within the following three enterprise IP Telephony deployment models plus a fourth application scenario:

- [Single-Site Deployments, page 6-23](#), consist of one or more call processing agents within a single site, with no voice traffic carried over the IP WAN.
- [Multisite WAN Deployments with Centralized Call Processing, page 6-24](#), consist of a single call processing agent that provides service for many sites connected through an IP WAN.
- [Multisite WAN Deployments with Distributed Call Processing, page 6-25](#), consist of call processing agents residing at each of several remote sites connected through an IP WAN.
- [IP PSTN Access, page 6-26](#), is another scenario that requires MTP resources, and it applies to all of the preceding deployment models.

Single-Site Deployments

In a single-site deployment, there is no need for transcoding because there are no low-speed links to justify the use of a low bit-rate (LBR) codec. Some MTP resources might be required in the presence of a significant number of devices that are not compliant with H.323v2, such as older versions of Microsoft NetMeeting or certain video devices. MTP resources may be required for DTMF conversion if SIP endpoints are present (see [Named Telephony Events \(RFC 2833\), page 6-13.](#))


Multisite WAN Deployments with Centralized Call Processing

In a centralized call processing deployment, the Cisco Unified CallManager cluster and the applications (such as voice mail and IVR) are located at the central site, while several remote sites are connected through an IP WAN. The remote sites rely on the centralized Cisco Unified CallManagers to handle their call processing.

Because WAN bandwidth is typically limited, calls are configured to use a low bit-rate codec such as G.729 when traversing the WAN. (See [Figure 6-1](#).)

Voice compression between IP phones is easily configured through the use of *regions* and *locations* in Cisco Unified CallManager. A region defines the type of compression (for example, G.711 or G.729) used by the devices in that region, and a location specifies the total amount of bandwidth available for calls to and from devices at that location.

Figure 6-1 Transcoding for the WAN with Centralized Call Processing



Cisco Unified CallManager uses media resource groups (MRGs) to enable sharing of MTP and transcoding resources among the Cisco Unified CallManager servers within a cluster. In addition, when using an LBR codec (for example, G.729a) for calls that traverse different regions, the transcoding resources are used only if one (or both) of the endpoints is unable to use the LBR codec.

In [Figure 6-1](#), Cisco Unified CallManager knows that a transcoder is required and allocates one based on the MRGL and/or MRG of the device that is using the higher-bandwidth codec. In this case it is the VM/UM server that determines which transcoder device is used. This behavior of Cisco Unified CallManager is based on the assumption that the transcoder resources are actually located close to the higher-bandwidth device. If this system was designed so that the transcoder for the VM/UM server was located at the remote site, then G.711 would be sent across the WAN, which would defeat the intended design. As a result, if there are multiple sites with G.711-only devices, then each of these sites would need transcoder resources when an LBR is run on the WAN.

The placement of other resources is also important. For example, if a conference occurs with three phones at a remote site and the conference resource is located in the central (call processing) site, then three media streams are carried over the WAN. If the conference resource were local, then the calls would not traverse the WAN. It is necessary to consider this factor when designing the bandwidth and call admission control for your WAN.

Multisite WAN Deployments with Distributed Call Processing


In distributed call processing deployments, several sites are connected through an IP WAN. Each site contains a Cisco Unified CallManager cluster that can, in turn, follow the single-site model or the centralized call processing model. A gatekeeper may be used for call admission control between sites.

Because WAN bandwidth is typically limited, calls between sites may be configured to use an LBR codec (such as G.729a) when traversing the WAN. H.323v2 intercluster trunks are used to connect Cisco Unified CallManager clusters. Cisco Unified CallManager also supports compressed voice call connections through the MTP service if a hardware MTP is used. (See [Figure 6-2](#).)

A distributed call processing deployment might need transcoding and MTP services in the following situations:

- With current versions of Cisco applications, it is possible and recommended to avoid the use of transcoding resources. There might be specific instances where G.711 on a specific device cannot be avoided.
- Some endpoints (for example, video endpoints) do not support the H.323v2 features.

Figure 6-2 Intercluster Call Flow with Transcoding



Cisco Unified CallManager uses media resource groups (MRGs) to enable sharing of MTP and transcoding resources among the Cisco Unified CallManager servers within a cluster. In addition, for calls across intercluster trunks, MTP and transcoding resources are used only when needed, thus eliminating the need to configure the MTP service for applications that do not support LBR codecs.


The following characteristics apply to distributed call processing deployments:

- Only the intercluster calls that require transcoding will use the MTP service. For example, if both endpoints of a call are capable of using a G.729 codec, no transcoding resources will be used.
- Sharing MTP resources among servers within a cluster provides more efficient resource utilization.

IP PSTN Access

Another application scenario for MTP and transcoding resources involves a service provider that provides its customers with access to an IP PSTN instead of the traditional PSTN. In such a scenario, the gatekeepers reside in the service provider network. In order to simplify the dial plan, each customer is required to use an MTP to anchor its calls, so that the individual IP addresses assigned to the endpoints can be hidden. The service provider's central office can then relay through the traditional PSTN and/or provide IP connectivity to other customers. [Figure 6-3](#) illustrates this deployment model.

Figure 6-3 IP PSTN Access Example



77307

Note that the customer site in [Figure 6-3](#) can use any of the previous three deployment models: single site, multisite WAN with centralized call processing, or multisite WAN with distributed call processing.

The H.323 trunk from the customer site to the IP PSTN must be configured with MTP so that the endpoint IP addresses remain masked. Thus, all external calls use the MTP resources. However, MTP resources can be shared within the Cisco Unified CallManager cluster to achieve more efficient use of the resources. This technique of utilizing MTPs for address hiding may be used for SIP trunks as well.



CHAPTER 7

Music on Hold

Last revised on: February 13, 2008

Music on hold (MoH) is an integral feature of the Cisco Unified Communications system. This feature provides music to callers when their call is placed on hold, transferred, parked, or added to an ad-hoc conference. Implementing MoH is relatively simple but requires a basic understanding of unicast and multicast traffic, MoH call flows, configuration options, server behavior and requirements. This chapter describes how to design and provision MoH resources for a Cisco Enterprise IP Telephony deployment.

Cisco Unified CallManager provides access to a variety of media resources. A media resource is a software-based or hardware-based entity that performs some media processing function on the voice data streams that are connected to it. Media processing functions include mixing multiple streams to create one output stream, passing the stream from one connection to another, or transcoding the data stream from one compression type to another.

Cisco Unified CallManager allocates and uses the following types of media resources:

- Media termination point (MTP) resources
- Transcoding resources
- Unicast conferencing resources
- Annunciator resources
- Music on hold resources

For more information about media resources in general, see the chapter on [Media Resources, page 6-1](#).

This chapter examines the following design aspects of the MoH feature:

- [Deployment Basics of MoH, page 7-2](#)
- [Basic MoH and MoH Call Flows, page 7-5](#)
- [MoH Configuration Considerations and Best Practices, page 7-8](#)
- [Hardware and Capacity Planning for MoH Resources, page 7-12](#)
- [Implications for MoH With Regard to IP Telephony Deployment Models, page 7-15](#)
- [Detailed Unicast and Multicast MoH Call Flows, page 7-20](#)

Deployment Basics of MoH

For callers to hear music on hold, Cisco Unified CallManager must be configured to support the MoH feature. The MoH feature has two main requirements:

- An MoH server to provide the MoH audio stream sources
- Cisco Unified CallManager configured to use the MoH streams provided by the MoH server when a call is placed on hold

The integrated MoH feature allows users to place on-net and off-net users on hold with music streamed from a streaming source. This source makes music available to any on-net or off-net device placed on hold. On-net devices include station devices and applications placed on hold, consult hold, or park hold by an interactive voice response (IVR) or call distributor. Off-net users include those connected through Media Gateway Control Protocol (MGCP) and H.323 gateways. The MoH feature is also available for plain old telephone service (POTS) phones connected to the Cisco IP network through Foreign Exchange Station (FXS) ports. The integrated MoH feature includes media server, database administration, call control, media resource manager, and media control functional areas. The MoH server provides the music resources and streams.

You can configure the MoH feature via the Cisco Unified CallManager Administration interface. When an end device or feature places a call on hold, Cisco Unified CallManager connects the held device to an MoH media resource. Essentially, Cisco Unified CallManager instructs the end device to establish a connection to the MoH server. When the held device is retrieved, it disconnects from the MoH resource and resumes normal activity.

Unicast and Multicast MoH

Cisco Unified CallManager supports two types of MoH transport mechanisms:

- Unicast
- Multicast

Unicast MoH consists of streams sent directly from the MoH server to the endpoint requesting an MoH audio stream. A unicast MoH stream is a point-to-point, one-way audio Real-Time Transport Protocol (RTP) stream between the server and the endpoint device. Unicast MoH uses a separate source stream for each user or connection. As more endpoint devices go on hold via a user or network event, the number of MoH streams increases. Thus, if twenty devices are on hold, then twenty streams of RTP traffic are generated over the network between the server and these endpoint devices. These additional MoH streams can potentially have a negative effect on network throughput and bandwidth. However, unicast MoH can be extremely useful in those networks where multicast is not enabled or where devices are not capable of multicast, thereby still allowing an administrator to take advantage of the MoH feature.

Multicast MoH consists of streams sent from the MoH server to a multicast group IP address that endpoints requesting an MoH audio stream can join as needed. A multicast MoH stream is a point-to-multipoint, one-way audio RTP stream between the MoH server and the multicast group IP address. Multicast music on hold conserves system resources and bandwidth because it enables multiple users to use the same audio source stream to provide music on hold. Thus, if twenty devices are on hold, then potentially only a single stream of RTP traffic is generated over the network. For this reason, multicast is an extremely attractive technology for the deployment of a service such as MoH because it greatly reduces the CPU impact on the source device and also greatly reduces the bandwidth consumption for delivery over common paths. However, multicast MoH can be problematic in situations where a network is not enabled for multicast or where the endpoint devices are not capable of handling multicast.

For information about IP multicast network design, refer to the *IP Multicast SRND* document, available online at

<http://www.cisco.com/go/designzone>

Recommended Unicast/Multicast Gateways

The following recommended gateways support both unicast and multicast MoH:

- Cisco 6624 and 6608 gateway modules with MGCP and Cisco Unified CallManager Release 3.3(3) or later
- Cisco Communication Media Module (CMM) with MGCP or H.323 and Cisco Unified CallManager Release 4.0, Cisco IOS Release 12.2(13)ZP3 or later, and Catalyst OS Release 8.1(1) or later
- Cisco 2600, 2800, 3600, 3700, and 3800 Series Routers with MGCP or H.323 and Cisco IOS Release 12.2(8)T or later

Co-resident and Standalone MoH Servers

The MoH feature requires the use of a server that is part of a Cisco Unified CallManager cluster. You can configure the MoH server in either of the following ways:

- Co-resident deployment

In a co-resident deployment, the MoH feature runs on any server (either publisher or subscriber) in the cluster that is also running the Cisco Unified CallManager software. Because MoH shares server resources with Cisco Unified CallManager in a co-resident configuration, this type of configuration drastically reduces the number of simultaneous streams that an MoH server can send.
- Standalone deployment

A standalone deployment places the MoH feature on a dedicated server within the Cisco Unified CallManager cluster. The sole function of this dedicated server is to send MoH streams to devices within the network. A standalone deployment allows for the maximum number of streams from a single MoH server.

Fixed and Audio File MoH Sources

You can set up the source for MoH in any of the following ways:

- MoH from an audio file on the Cisco Unified CallManager or MoH server
 - Unicast MoH from an audio file
 - Multicast MoH from an audio file
- MoH from a fixed music source (via sound card)
 - Unicast MoH from a fixed source
 - Multicast MoH from a fixed source

MoH can be generated from an audio file stored on the MoH server. Audio files must be in one of the following formats:

- G.711 A-law or mu-law
- G.729 Annex A
- Wideband

You can create these files with the Cisco MoH Audio Translator service, which transcodes and formats audio source files (such as .wav or .mp3 files) into the appropriate MoH source file for the specified codec type(s). The MoH server requests these files based on the audio sources configured. When an MoH event occurs, the configured audio source file is streamed to the requesting device on hold.

If recorded or live audio is needed, MoH can be generated from a fixed source. For this type of MoH, a sound card is required. The fixed audio source is connected to the audio input of the local sound card.

This mechanism enables you to use radios, CD players, or any other compatible sound source. The stream from the fixed audio source is transcoded in real-time to support the codec that was configured through Cisco Unified CallManager Administration. The fixed audio source can be transcoded into G.711 (A-law or mu-law), G.729 Annex A, and Wideband, and it is the only audio source that is transcoded in real-time.

The following sound cards are supported for a fixed or live audio source:

- Griffin Technologies iMic USB

A USB sound device supported in Cisco CallManager Release 3.3(5) and later, with Microsoft Windows 2000 (OS 2000 version 2.7 or later). This device is supported on all Cisco MCS-78xxH or MCS-78xxI servers with 3.0 GHz or greater processor.

- Telex P-800 USB

A USB sound device supported in Cisco CallManager Release 3.3(3) with Microsoft Windows 2000 (OS 2000 version 2.5). This device is supported on all Cisco MCS-78xxH or MCS-78xxI servers with 2.2 GHz or greater processor.

**Note**

The Telex P-800 USB card is no longer available. While existing P-800 cards are still supported as indicated above, the Griffin iMic USB card should be used in new deployments.

**Note**

Prior to using a fixed audio source to transmit music on hold, you should consider the legalities and the ramifications of re-broadcasting copyrighted audio materials. Consult your legal department for potential issues.

MoH Server as Part of the Cisco Unified CallManager Cluster

The MoH feature requires that each MoH server must be part of a Cisco Unified CallManager cluster. All MoH servers must share their configurations with the publisher server and participate in the database replication schema. Specifically, the MoH server must share the following information (configured through Cisco Unified CallManager Administration) by means of the database:

- Audio sources — The number and identity of all configured MoH audio sources
- Multicast or unicast — The transport nature (multicast or unicast) configured for each of these sources
- Multicast address — The multicast base IP address of those sources configured to stream as multicast

The MoH server becomes part of the Cisco Unified CallManager cluster and participates in the database replication automatically. To configure a standalone MoH server, start with a normal Cisco Unified CallManager installation on that server, then disable the Cisco CallManager service (on the standalone MoH server only) and enable the Cisco IP Voice Media Streaming Application.

Basic MoH and MoH Call Flows

This section describes the basic operation of MoH as implemented in Cisco Unified CallManager, as well as typical call flow scenarios.


Basic MoH

The basic operation of MoH in a Cisco Unified Communications environment consists of a holder and a holdee. The *holder* is the endpoint user or network application placing a call on hold, and the *holdee* is the endpoint user or device placed on hold.

The MoH stream that an endpoint receives is determined by a combination of the User Hold MoH Audio Source of the device placing the endpoint on hold (holder) and the configured media resource group list (MRGL) of the endpoint placed on hold (holdee). The User Hold MoH Audio Source configured for the holder determines the audio file that will be streamed when the holder puts a call on hold, and the holdee's configured MRGL indicates the resource or server from which the holdee will receive the MoH stream.

In simplest terms, the holder's configuration determines which audio file to play, and the holdee's configuration determines which resource or server will play that file. As illustrated by the example in [Figure 7-1](#), if phones A and B are on a call and phone B (holder) places phone A (holdee) on hold, phone A will hear the MoH audio source configured for phone B (Audio-source2). However, phone A will receive this MoH audio stream from the MRGL (resource or server) configured for phone A (MRGL A).

Figure 7-1 User Hold Audio Source and Media Resource Group List (MRGL)



Because the configured MRGL determines the server from which a unicast-only device will receive the MoH stream, you must configure unicast-only devices with an MRGL that points to a unicast MoH resource or media resource group (MRG). Likewise, a device capable of multicast should be configured with an MRGL that points to a multicast MRG.

MoH Configuration Settings

You can configure the settings for MRGLs and User and Network Hold Audio Sources in several places within Cisco Unified CallManager Administration, and you can configure different (and potentially conflicting) settings in each place.

To determine which User and Network Audio Source configuration setting to apply in a particular case, Cisco Unified CallManager interprets these settings for the *holder* device in the following priority order:

1. Directory or line setting (Devices with no line definition, such as gateways, do not have this level.)
2. Device setting
3. Common Profile setting
4. Cluster-wide default setting

When attempting to determine the audio source for a particular holder, Cisco Unified CallManager first looks at the User (or Network) Audio Source configured at the directory or line level. If this level is not defined, Cisco Unified CallManager looks at the User (or Network) Audio Source configured on the holder device. If this level is not defined, Cisco Unified CallManager looks at the User (or Network) Audio Source configured for the common profile of the holder device. If this level is not defined, then Cisco Unified CallManager looks at the cluster-wide default audio source ID configured under the Cisco Unified CallManager system parameters. (By default, this audio source ID is set to 1 for both User and Network Hold Audio Sources, which is the SampleAudioSource.)

Cisco Unified CallManager also interprets the MRGL configuration settings of the *holdee* device in the following priority order:

1. Device setting
2. Device pool setting
3. System default MoH resources


When attempting to determine the MRGL for a particular holdee, Cisco Unified CallManager looks at the MRGL configured at the device level. If this level is not defined, Cisco Unified CallManager looks at the MRGL configured for the device pool of the holdee device. If this level is not defined, then Cisco Unified CallManager uses the system default MoH resources. System default MoH resources are those resources not assigned to any MRGL, and they are always unicast.

User and Network Hold

There are two basic types of user hold:

- User on hold at an IP phone or other endpoint device
- User on hold at the PSTN, where MoH is streamed to the gateway


Figure 7-2 shows these two types of call flows. If phone A is in a call with phone B and phone A (holder) pushes the Hold softkey, then a music stream is sent from the MoH server to phone B (holdee). The music stream can be sent to holdees within the IP network or holdees on the PSTN, as is the case if phone A places phone C on hold. In the case of phone C, the MoH stream is sent to the voice gateway interface and converted to the appropriate format for the PSTN phone. When phone A presses the Resume softkey, the holdee (phone B or C) disconnects from the music stream and reconnects to phone A.

Figure 7-2 Basic User Hold Example

Network hold includes the following types:

- Call transfer
- Call Park
- Conference setup
- Application-based hold

Figure 7-3 shows the call transfer call flow. When phone A receives a call from PSTN phone C (step 1), phone A answers the call and then transfers it to phone B (step 2). During the transfer process, phone C receives an MoH stream from the MoH server via the gateway (step 3). After phone A completes the transfer action, phone C disconnects from the music stream and gets redirected to phone B, the transfer destination. This process is the same for other network hold operations such as call park and conference setup.

Figure 7-3 Basic Network Hold Example for Call Transfer

Unicast and Multicast MoH Call Flows

MoH operation is quite similar to normal phone call flows, with the MoH server acting as an endpoint device to which the holdee device can connect or disconnect as required. However, there are distinct differences between unicast and multicast MoH call flow behavior. A unicast MoH call flow is initiated by a message from Cisco Unified CallManager to the MoH server. This message tells the MoH server to send an audio stream to the holdee device's IP address. On the other hand, a multicast MoH call flow is initiated by a message from Cisco Unified CallManager to the holdee device. This message instructs the endpoint device to join the multicast group address of the configured multicast MoH audio stream.

For a detailed look at MoH call flows, see the section on [Detailed Unicast and Multicast MoH Call Flows, page 7-20](#).

MoH Configuration Considerations and Best Practices

This section highlights some MoH configuration considerations and best practice to help you design a robust MoH solution.

Codec Selection

If you need multiple codecs for MoH deployment, configure them in the IP Voice Streaming Media App service parameter under Cisco CallManager Service Parameters Configuration. Select the desired codec types from the Supported MoH Codecs list under the Clusterwide Parameters section. By default, only G.711 mu-law is selected. To select another codec type, click on it in the scrollable list. For multiple selections, hold down the CTRL key and use the mouse to select multiple codecs from the scrollable list. After making your selection, click the Update button.



Note

If you are using the G.729 codec for MoH audio streams, be aware that this codec is optimized for speech and it provides only marginal audio fidelity for music.

Multicast Addressing

Proper IP addressing is important for configuring multicast MoH. Addresses for IP multicast range from 224.0.1.0 to 239.255.255.255. The Internet Assigned Numbers Authority (IANA), however, assigns addresses in the range 224.0.1.0 to 238.255.255.255 for public multicast applications. Cisco strongly discourages using public multicast addresses for music on hold. Instead, Cisco recommends that you configure multicast MoH audio sources to use IP addresses in the range 239.1.1.1 to 239.255.255.255, which is reserved for administratively controlled applications on private networks.

Furthermore, you should configure multicast audio sources to increment on the IP address and not the port number, for the following reasons:

- IP phones placed on hold join multicast IP addresses, not port numbers.

Cisco IP phones have no concept of multicast port numbers. Therefore, if all the configured codecs for a particular audio stream transmit to the same multicast IP address (even on different port numbers), all streams will be sent to the IP phone even though only one stream is needed. This has the potential of saturating the network with unnecessary traffic because the IP phone is capable of receiving only a single MoH stream.

- IP network routers route multicast based on IP addresses, not port numbers.

Routers have no concept of multicast port numbers. Thus, when it encounters multiple streams sent to the same multicast group address (even on different port numbers), the router forwards all streams of the multicast group. Because only one stream is needed, network bandwidth is over-utilized and network congestion can eventually result.

MoH Audio Sources

Audio sources are shared among *all* MoH servers in the Cisco Unified CallManager cluster. You can configure up to 51 unique audio sources per cluster (50 audio file sources and one fixed/live source via a sound card). For exceptions to this limit, refer to the sections on [Using Multiple Fixed or Live Audio Sources, page 7-9](#) and [Multicast MoH from Branch Router Flash, page 7-17](#).

Using Multiple Fixed or Live Audio Sources

Each MoH server is capable of streaming only one fixed audio source. In most cases, if multiple fixed or live audio sources are needed, a separate MoH server is required for each source. However, it is possible to provide multiple fixed-source MoH audio streams by using external non-MoH servers or devices that are capable of streaming multicast from fixed or live sources.


For each external source, you must configure the MoH server with an audio source that has the same multicast IP address and port number as that of the audio source stream being multicast by the external source server or device. In addition, you should block this configured (non-external) audio source from traversing the WAN by setting its maximum hop count to one (1) or by using access control lists (ACLs) to prevent the packets from streaming further than the local subnet.

[Figure 7-4](#) shows an example of an external live source being used as an MoH stream. In this figure, the MoH server is streaming a multicast audio source to 239.1.1.1 (on RTP port 16384), and this stream has been limited to a maximum hop of one, thus ensuring that it will not travel beyond the local MoH server's subnet. At the same time, the media server is multicasting an audio stream derived from a live radio station feed. This stream is also using 239.1.1.1 as its multicast address and 16384 as the RTP port number, but this stream has a hop count or Time to Live (TTL) greater than one to ensure that it is able to reach phone B when phone A presses the Hold softkey.



Note

Multicast Time to Live (TTL) values are decremented (or they expire) only when packets traverse Layer 3 interfaces.

Figure 7-4 External Live Audio Source Example**Note**

Using live radio broadcasts as multicast audio sources can have legal ramifications. Consult your legal department for potential issues.

Numerous streams can be multicast from one or more external media servers, by configuring additional audio sources on multiple MoH servers and then sourcing audio streams from the external servers using the same multicast group addresses configured on the MoH servers. However, because a combination of the user/network hold audio source of the holder and the MRGL of the holdee determines the MoH stream that an endpoint device hears, it can become difficult to predict which particular stream an endpoint will receive in an environment with many overlapping multicast group addresses. For this reason, Cisco recommends that you configure only a single multicast audio source on each MoH server. This recommendation ensures that the audio source an endpoint receives is uniquely identifiable by a single combination of user/network hold audio source and MRGL.

Unicast and Multicast in the Same Cisco Unified CallManager Cluster

In some cases, administrators might want to configure a single Cisco Unified CallManager cluster to handle both unicast and multicast MoH streams. This configuration might be necessary because the telephony network contains devices or endpoint that do not support multicast or because some portions of the network are not enabled for multicast.

Use one of the following methods to enable a cluster to support both unicast and multicast MoH audio streams:

- Deploy separate MoH servers, with one server configured as a unicast MoH server and the second server configured as a multicast MoH server.
- Configure separate media resource groups (MRGs) for the same MoH server, with one MRG configured to use multicast for audio streams and the second MRG configured to use unicast.

In either case, you must configure at least two MRGs and at least two media resource group lists (MRGLs). Configure one unicast MRG and one unicast MRGL for those endpoints requiring unicast MoH. Likewise, configure one multicast MRG and one multicast MRGL for those endpoints requiring multicast MoH.

When deploying separate MoH servers, configure one server without multicast enabled (unicast-only) and configure a second MoH server with multicast enabled. Assign the unicast audio resource of the unicast-only MoH server and the multicast audio resource of the multicast MoH server to the unicast and multicast MRGs, respectively. Ensure that the **Use Multicast for MoH Audio** box is checked for the multicast MRG but not for the unicast MRG. Also assign these unicast and multicast MRGs to their respective MRGLs. In this case, an MoH stream is unicast or multicast based on the server from which it is served and on whether the MRG is configured to use multicast.

When deploying a single MoH server for both unicast and multicast MoH, configure the server and its audio source for multicast. Assign this same audio source to both the unicast MRG and the multicast MRG, and check the **Use Multicast for MoH Audio** box for the multicast MRG. In this case, an MoH stream is unicast or multicast based solely on whether the MRG is configured to use multicast.

**Note**

When configuring the unicast MRG, do not be confused by the fact that the audio resource you are adding to this MRG has [Multicast] appended to the end of the resource name even though you are adding it to the unicast MRG. This label is simply an indication that the resource is capable of being multicast, but the **Use Multicast for MoH Audio** box determines whether the resource will be sent as unicast or multicast.

In addition, you must configure individual devices or device pools to use the appropriate MRGL. You can place all unicast devices in a device pool or pools and configure those device pools to use the unicast MRGL. Likewise, you can place all multicast devices in a device pool or pools and configure those device pools to use the multicast MRGL. Optionally, you can configure individual devices to use the appropriate unicast or multicast MRGL. Also configure a User Hold Audio Source and Network Hold Audio Source for each device pool, individual device, or (in the case of phone devices) individual lines or directory numbers to determine the appropriate audio source to stream.

When choosing a method for deploying both multicast and unicast MoH in the same cluster, an important factor to consider is the number of servers required. When using a single MoH server for both unicast and multicast, fewer MoH servers are required throughout the cluster. Deploying separate multicast and unicast MoH servers will obviously require more servers within the cluster.

Redundancy

Cisco recommends that you configure and deploy multiple MoH servers for completely redundant MoH operation. If the first MoH server fails or becomes unavailable because it no longer has the resources required to service requests, the second server can provide continued MoH functionality. For proper redundant configuration, assign resources from at least two MoH servers to each MRG in the cluster.

Resources in the MRG are used in the order listed. When a device requests an MoH audio resource, Cisco Unified CallManager attempts to stream the first MoH resource in the MRG to the device. If the first resource is unavailable (due to server failure or lack of resources), Cisco Unified CallManager then attempts to use the next MoH resource in the MRG.

In environments where both multicast and unicast MoH are required, be sure to provide redundancy for both transport types to ensure MoH redundancy for all endpoints in the network.

Quality of Service (QoS)

Convergence of data and voice on a single network requires adequate QoS to ensure that time-sensitive and critical real-time applications such as voice are not delayed or dropped. To ensure proper QoS for voice traffic, the streams must be marked, classified, and queued as they enter and traverse the network.

to give the voice streams preferential treatment over less critical traffic. MoH servers automatically mark audio stream traffic the same as voice bearer traffic, with a Differentiated Services Code Point (DSCP) value of 46 or Per Hop Behavior (PHB) value of EF (ToS of 0xB8). Therefore, as long as QoS is properly configured on the network, MoH streams will receive the same classification and priority queueing treatment as voice RTP media traffic.

Call signaling traffic between MoH servers and Cisco Unified CallManager servers is marked with a DSCP value of 26 or PHB value of AF31 (ToS of 0x68) by default. However, in order to conform with Cisco QoS marking recommendations, this traffic should be marked with a DSCP value of 24 or PHB value of CS3 (ToS of 0x60) to ensure that it is properly classified and queued within the network. The default traffic marking can be changed to the appropriate marking by changing the *IP Type of Service to Cisco CallManager* service parameter value from 0x68 to 0x60 under the Cisco IP Voice Media Streaming App Service Parameter configuration page.

Hardware and Capacity Planning for MoH Resources

As with all media resources, capacity planning is crucial to make certain that the hardware, once deployed and configured, can support the anticipated call volume of the network. For this reason, it is important to be aware of the hardware capacity for MoH resources and to consider the implications of multicast and unicast MoH in relation to this capacity.

Server Platform Limits

Table 7-1 lists the server platforms and the maximum number of simultaneous MoH sessions each can support. Ensure that network call volumes do not exceed these limits because, once MoH sessions have reached these limits, additional load could result in poor MoH quality, erratic MoH operation, or even loss of MoH functionality.

Table 7-1 Maximum Number of MoH Sessions per Server Platform Type

Server Platform	Codecs Supported	MoH Sessions Supported
MCS 7815	G.711 (A-law and mu-law)	Co-resident server: 40 MoH sessions Standalone MoH server: 200 MoH sessions
MCS 782x (All models)	G.729a	
MCS 7830 (All models)	Wideband audio	
SPE-310		
HP DL320		
IBM xSeries 33x (All models)		
MCS 7835 (All models)	G.711 (A-law and mu-law)	Coresident server: 100 MoH sessions Standalone MoH server: 250 MoH sessions ¹
MCS 7845 (All models)	G.729a	
HP DL380	Wideband audio	
IBM xSeries 34x (All models)		

1. You can configure a maximum of 51 unique audio sources per Cisco Unified CallManager cluster.

The following MoH Server Configuration parameters affect MoH server capacity:

- **Maximum Half Duplex Streams**

This parameter determines the number of devices that can be placed on unicast MoH. By default this value is set to 250.

The Maximum Half Duplex Streams parameter should be set to the value derived from the following formula:

$$(\text{Server and deployment capacity}) - ((\text{Number of multicast MoH sources}) * (\text{Number of MoH codecs enabled}))$$

For example:

MCS-7835 standalone MoH server	Multicast MoH audio sources	MoH codecs enables (G.711 mu-law and G.729)	Maximum half-duplex streams
250	- (12	* 2)	= 226

Therefore, in this example, the Maximum Half Duplex Streams parameter would be configured with a value of no more than 226.

The value of this parameter should never be set higher than the capacities indicated in [Table 7-1](#), based on the platform and deployment type (co-resident or standalone).

- **Maximum Multicast Connections**

This parameter determines the number of devices that can be placed on multicast MoH. By default this value is set to 30.

The Maximum Multicast Connections parameter should be set to a number that ensures that all devices can be placed on multicast MoH if necessary. Although the MoH server can generate only a finite number of multicast streams (maximum of 204), large numbers of held devices can join each multicast stream, thus it is necessary to set this parameter to an artificially large number. Typically multicast traffic is accounted for based on the number of streams being generated; however, Cisco Unified CallManager maintains a count of the actual number of devices placed on multicast MoH or joined to each multicast MoH stream. This method is different than the way multicast traffic is normally tracked.

Failure to configure these parameters properly could lead to under-utilization of MoH server resources or failure of the server to handle the network load.



Note

The maximum session limits listed in [Table 7-1](#) apply to unicast, multicast, or simultaneous unicast and multicast sessions. The limits represent the recommended maximum sessions a platform can support, irrespective of the transport mechanism.

Resource Provisioning and Capacity Planning

When provisioning for co-resident or standalone MoH server configurations, network administrators should consider the type of transport mechanism used for the MoH audio streams. If using unicast MoH, each device on hold will require a separate MoH stream. However, if using multicast MoH and only a single audio source, then only a single MoH stream is required for each configured codec type, no matter how many devices of that type are on hold.

For example, given a cluster with 30,000 phones and a 2% hold rate (only 2% of all endpoint devices are on hold at any given time), 600 MoH streams or sessions would be required. Given a unicast-only MoH environment, two standalone MoH servers (MCS 7835 or 7845) and one co-resident MoH server (MCS 7835 or 7845) would be required to handle this load, as shown by the following calculation:

$$[(250 \text{ sessions per MCS 7835 or 7845 standalone server}) * (2 \text{ standalone servers})] + [(100 \text{ sessions per MCS 7835 or 7845 co-resident server}) * (1 \text{ co-resident server})] = 600 \text{ sessions}$$

By comparison, a multicast-only MoH environment with 36 unique MoH audio streams, for example, would require only one co-resident MoH server (MCS 7815, 782x, or 7830), as shown by the following calculation:

$$(40 \text{ sessions per MCS 7815, 782x, or 7830 co-resident server}) * (1 \text{ co-resident server}) > 36 \text{ sessions}$$

These 36 unique multicast streams could be provisioned in any one of the following ways:

- 36 unique audio sources streamed using a single codec
- 18 unique audio sources streamed using only 2 codecs
- 12 unique audio sources streamed using only 3 codecs
- 9 unique audio source streamed using all 4 codecs

As these examples show, multicast MoH can provide a considerable savings in server resources over unicast MoH.

In the preceding examples, the 2% hold rate is based on 30,000 phones and does not take into account gateways or other endpoint devices in the network that are also capable of being placed on hold. You should consider these other devices when calculating a hold rate because they could potentially be placed on hold just as the phones can.

The preceding calculations also do not provide for MoH server redundancy. If an MoH server fails or if more than 2% of the users go on hold at the same time, there are no other MoH resources in this scenario to handle the overflow or additional load. Your MoH resource calculations should include enough extra capacity to provide for redundancy.

**Note**

Because you can configure only 51 unique audio sources per Cisco Unified CallManager cluster and because there are only four possible codecs for MoH streams, the maximum number of multicast streams per MoH server is 204.

Implications for MoH With Regard to IP Telephony Deployment Models

The various IP Telephony deployment models introduce additional considerations for MoH configuration design. Which deployment model you choose can also affect your decisions about MoH transport mechanisms (unicast or multicast), resource provisioning, and codecs. This section discusses these issues in relation to the various deployment models.

For more detailed information about the deployment models, see the chapter on [IP Telephony Deployment Models](#), page 2-1.

Single-Site Campus (Relevant to All Deployments)

Single-site campus deployments are typically based on a LAN infrastructure and provide sufficient bandwidth for large amounts of traffic. Because bandwidth is typically not limited in a LAN infrastructure, Cisco recommends the use of the G.711 (A-law or mu-law) codec for all MoH audio streams in a single-site deployment. G.711 provides the optimal voice and music streaming quality in an IP Telephony environment.

MoH server redundancy should also be considered. In the event that an MoH server becomes overloaded or is unavailable, configuring multiple MoH servers and assigning them in preferred order to MRGs ensures that another server can take over and provide the MoH streams.

With the increasing diversity of network technologies, in a large single-site campus it is likely that some endpoint devices will be unable to support multicast. For this reason, you might have to deploy both unicast and multicast MoH resources. For example, wireless IP phones do not support multicast due to the behavior of wireless technology. Thus, when deploying wireless IP phones, you have to configure both multicast and unicast MoH.

To ensure that off-net calls and application-handled calls receive expected MoH streams when placed on hold, configure all gateways and other devices with the appropriate MRGLs and audio sources, or assign them to appropriate device pools and common profiles.

Centralized Multisite Deployments

Multisite IP telephony deployments with centralized call processing typically contain WAN connections to multiple non-central sites. These WAN links usually cause bandwidth and throughput bottlenecks. To minimize bandwidth consumption on these links, Cisco recommends the use of the G.729 codec for all MoH audio streams traversing the WAN. Because the G.729 codec is optimized for voice and not music applications, you should use G.729 only across the WAN, where the bandwidth savings far outweighs the lower quality afforded by G.729 for MoH transport. Likewise, because multicast traffic provides significant bandwidth savings, you should always use multicast MoH when streaming audio to endpoints across the WAN.

If the sound quality of an MoH stream becomes an issue when using the G.729 codec across the WAN, you can use the G.711 codec for MoH audio streams across the WAN while still using G.729 for voice calls. In order to send MoH streams across the WAN with the G.711 codec but voice calls across the WAN with the G.729 codec, place all MoH servers in a Cisco Unified CallManager region by themselves, and configure that region to use G.711 between itself and all other regions. Thus, when a call is placed between two phones on either side of a WAN, the G.729 codec is used between their

respective regions. However, when the call is placed on hold by either party, the MoH audio stream is encoded using G.711 because G.711 is the configured codec to use between the MoH server's region and the region of the phone placed on hold.

Call Admission Control and MoH

Call admission control (CAC) is required when IP telephony traffic is traveling across WAN links. Due to the limited bandwidth available on these links, it is highly possible that voice media traffic might get delayed or dropped. The Cisco Unified CallManager locations-based call admission control mechanism enables you to configure each location in the IP telephony environment to accept or allow only a certain number of calls across the WAN link to other locations, thus preventing oversubscription of WAN bandwidth and delayed or lost voice packets. By specifying a bandwidth value for the WAN link, you can limit the number of calls based on the speed of the link. If the limit is reached or exceeded, Cisco Unified CallManager rejects all other calls that are attempted across the link. For additional information, see [Call Admission Control, page 9-1](#).


Cisco Unified CallManager locations-based call admission control is capable of tracking unicast MoH streams traversing the WAN but not multicast MoH streams. Thus, even if WAN bandwidth has been fully subscribed, a multicast MoH stream will not be denied access to the WAN by call admission control. Instead, the stream will be sent across the WAN, likely resulting in poor audio stream quality and poor quality on all other calls traversing the WAN. To ensure that multicast MoH streams do not cause this over-subscription situation, you should over-provision the QoS configuration on all downstream WAN interfaces by configuring the low-latency queuing (LLQ) voice priority queue with additional bandwidth. Because MoH streams are uni-directional, only the voice priority queues of the downstream interfaces (from the central site to remote sites) must be over-provisioned. Add enough bandwidth for every unique multicast MoH stream that might traverse the WAN link. For example, if there are four unique multicast audio streams that could potentially traverse the WAN, then add 96 kbps to the voice priority queue ($4 * 24 \text{ kbps per G.729 audio stream} = 96 \text{ kbps}$).

[Figure 7-5](#) shows an example of call admission control and MoH in a centralized multisite deployment. For this example, assume that IP phone C is in a call with a PSTN phone (phone B). At this point, no bandwidth has been consumed on the WAN. When phone C pushes the Hold softkey (step 1), phone B receives an MoH stream from the central-site MoH server via the WAN, thereby consuming bandwidth on the link. Whether or not this bandwidth is taken into consideration by call admission control depends on the type of MoH stream. If multicast MoH is streamed, then call admission control will not consider the 24 kbps being consumed (therefore, QoS on the downstream WAN interfaces should be provisioned accordingly). However, if unicast MoH is streamed, call admission control will subtract 24 kbps from the available WAN bandwidth (step 2).



Note

The preceding example might seem to imply that unicast MoH should be streamed across the WAN. However, this is merely an example used to illustrate locations-based call admission control with MoH and is not intended as a recommendation or endorsement of this configuration. As stated earlier, multicast MoH is the recommended transport mechanism for sending MoH audio streams across the WAN.

Figure 7-5 Locations-Based Call Admission Control and MoH

Multicast MoH from Branch Router Flash

Beginning with Cisco IOS Release 12.2(15)ZJ and SRST Release 3.0, MoH can be multicast in a remote or branch site via the branch router's flash. Multicast MoH from a Cisco IOS router's flash enhances the MoH feature for the following reasons:

- The branch gateway or router can provide multicast MoH when it is in SRST mode and the branch devices have lost connectivity to the central-site Cisco Unified CallManager.
- This configuration can eliminate the need to forward MoH across the WAN to remote branch sites by providing locally sourced MoH even when the WAN is up and the phones are controlled by Cisco Unified CallManager.

[Example 7-1](#) illustrates the commands to use in the Cisco IOS router configuration (under the SRST section) to enable multicast MoH from the router flash.

Example 7-1 Enabling Multicast MoH from Branch Router Flash

```
SRST-router(config)#call-manager-fallback
SRST-router(config-cm-fallback)#ip source-address 10.1.1.1
SRST-router(config-cm-fallback)#moh music-on-hold.au
SRST-router(config-cm-fallback)#multicast moh 239.192.240.1 port 16384 route 10.1.1.254
```

In [Example 7-1](#), the name of the audio file on the router flash is `music-on-hold.au`, and the configured multicast address and port number are `239.192.240.1` and `16384` respectively. The optional **route** command indicates a source interface address for the multicast stream. If no route option is specified, the multicast stream will be sourced from the configured SRST default address as specified by the **ip source-address** command under the SRST configuration. Note that you can stream only a single audio file from flash and that you can use only a single multicast address and port number per router.

When the branch router is operating in SRST mode, it can stream multicast MoH to all analog and digital ports within the chassis, thereby providing MoH to analog phones and PSTN callers. At this time, IP phones in SRST mode cannot receive multicast MoH from the SRST router's flash and will receive tone-on-hold instead.

**Note**

An SRST license is required regardless of whether the SRST functionality will actually be used. The license is required because the configuration for streaming MoH from branch router flash is done under the SRST configuration mode and, even if SRST functionality will not be used, at least one **max-ephones** and one **max-dn** must be configured. These configuration commands are required in addition to the commands listed in [Example 7-1](#).


Once configured, the router will continue to stream the MoH stream from flash even when not in SRST mode. When the branch router is not operating in SRST mode, it can multicast MoH from the flash to all local devices, including IP phones. The branch router's configuration for non-SRST multicast MoH from flash is the same as for the SRST configuration. (See [Example 7-1](#).) However, which multicast address you configure on the router depends on the intended operation. If you want multicast MoH from flash only in SRST mode (for example, if MoH received by remote devices is sourced from the central MoH server when not in SRST mode), then the multicast address and port number configured on the router should not overlap with any of the central-site MoH server audio sources. Otherwise, remote devices might continue to receive MoH from the local router flash, depending on the configured user/network hold audio sources.

If you always want multicast MoH from the branch router flash, then you must configure the central-site server with an audio source that has the same multicast IP address and port number as configured on the branch router. In this scenario, because the multicast MoH audio stream is always coming from the router's flash, it is not necessary for the central site MoH server audio source to traverse the WAN.

To prevent the central site audio stream(s) from traversing the WAN, use one of the following methods:

- Configure a maximum hop count
Configure the central-site MoH audio source with a maximum hop count (or TTL) low enough to ensure that it will not stream further than the central-site LAN.
- Configure an access control list (ACL) on the WAN interface
Configure an ACL on the central-site WAN interface to disallow packets destined to the multicast group address(es) from being sent out the interface.
- Disable multicast routing on the WAN interface
Do not configure multicast routing on the WAN interface, thus ensuring that multicast streams are not forwarded into the WAN.

[Figure 7-6](#) illustrates streaming multicast MoH from the flash of a remote router when it is not in SRST mode. After phone A places phone C on hold, phone C receives multicast MoH from the local SRST router. In this figure, the MoH server is streaming a multicast audio source to 239.192.240.1 (on RTP port 16384), however this stream has been limited to a maximum hop of one (1) to ensure that it will not travel off the local MoH server's subnet and across the WAN. At the same time, the branch office SRST router/gateway is multicasting an audio stream from flash. This stream is also using 239.192.240.1 as its multicast address and 16384 as the RTP port number. When phone A presses the Hold softkey, phone C receives the MoH audio stream sourced by the SRST router.

Figure 7-6 Multicast MoH from Branch Router Flash

When using this method for delivering multicast MoH, configure all devices within the Cisco Unified CallManager cluster to use the same user hold and network hold audio source and configure all branch routers with the same multicast group address and port number. Because the user or network hold audio source of the holder is used to determine the audio source, if you configure more than one user or network hold audio source within the cluster, there is no way to guarantee that a remote holdee will always receive the local MoH stream. For example, suppose a central-site phone is configured with an audio source that uses group address 239.192.254.1 as its user and network hold audio source. If this phone places a remote device on hold, the remote device will attempt to join 239.192.254.1 even if the local router flash MoH stream is sending to multicast group address 239.192.240.1. If instead all devices in the network are configured to use the user/network hold audio source with multicast group address 239.192.240.1 and all branch routers are configured to multicast from flash on 239.192.240.1, then every remote device will receive the MoH from its local router's flash.

In networks with multiple branch routers configured to stream multicast MoH from flash, it is possible to have more than 51 unique MoH audio sources in a cluster. Each branch site router can multicast a unique audio file from flash, although all routers must multicast this audio on the same multicast group address. In addition, the central-site MoH server can multicast a MoH stream on this same multicast group address. Thus, if there are 100 branch sites each multicasting an audio file from flash, then the cluster can contain 101 unique MoH audio sources (100 branch streams and one central-site stream). If you want more than one unique audio stream in the central site, you can stream fixed/live sources from additional MoH servers or from external media servers (as described in [Using Multiple Fixed or Live Audio Sources](#), page 7-9), but you should not configure more than one audio source per server.

Distributed Multisite Deployments

Multisite IP telephony deployments with distributed call processing typically contain WAN or MAN connections between the sites. These lower-speed links usually cause bandwidth and throughput bottlenecks. To minimize bandwidth consumption on these links, Cisco recommends use of the G.729

codec for all MoH audio streams traversing them. Because the G.729 codec is optimized for voice and not music applications, you should use G.729 only across the WAN/MAN links, where the bandwidth savings far outweighs the lower quality afforded by G.729 for MoH transport.

Unlike with centralized multisite deployments, in situations where G.711 might be required for MoH audio streams traveling across a WAN, MoH audio streams cannot be forced to G.711 in a distributed multisite deployment. Even when MoH servers are placed in a separate Cisco Unified CallManager region and the G.711 codec is configured between this region and the intercluster or SIP trunk's region, the codec of the original voice call is maintained when a call between the two clusters is placed on hold by either phone. Because these intercluster calls are typically encoded using G.729 for bandwidth savings, a MoH stream from either cluster will also be encoded using G.729.

Furthermore, multicast MoH is not supported for calls between Cisco Unified CallManager clusters (intercluster calls). Therefore, you must configure at least one unicast MoH resource in each Cisco Unified CallManager cluster if you want MoH on the intercluster or SIP trunk.

Proper multicast address management is another important design consideration in the distributed intercluster environment. All MoH audio source multicast addresses must be unique across all Cisco Unified CallManager clusters in the deployment to prevent possible overlap of streaming resources throughout the distributed network.

Clustering Over the WAN


As its name suggests, clustering-over-the-WAN deployments also contain the same type of lower-speed WAN links as other multisite deployments and therefore are subject to the same requirements for G.729 codec, multicast transport mechanism, and solid QoS for MoH traffic traversing these links.

In addition, you should deploy MoH server resources at each side of the WAN in this type of configuration. In the event of a WAN failure, devices on each side of the WAN will be able to continue to receive MoH audio streams from their locally deployed MoH server. Furthermore, proper MoH redundancy configuration is extremely important. The devices on each side of the WAN should point to an MRGL whose MRG has a priority list of MoH resources with at least one local resource as the highest priority. Additional MoH resources should be configured for this MRG in the event that the primary server becomes unavailable or is unable to process requests. At least one other MoH resource in the list should point to an MoH resource on the remote side of the WAN in the event that resources at the local side of the WAN are unavailable.

Detailed Unicast and Multicast MoH Call Flows

The following sections provide detailed illustrations and explanations of unicast and multicast MoH call flows.

[Figure 7-7](#) illustrates a typical multicast call flow. As shown in the diagram, when the Hold softkey is pressed at phone A, Cisco Unified CallManager instructs both phone A and phone B to Close Receive Channel and Stop Media Transmission. This action effectively stops the RTP two-way audio stream. Next, Cisco Unified CallManager tells phone B (the holdee) to Start Multicast Media Reception from multicast group address 239.192.240.1. The phone then issues an Internet Group Management Protocol (IGMP) V2 Membership Report message indicating that it is joining this group.

Figure 7-7 Detailed Multicast MoH Call Flow


Meanwhile, the MoH server has been sourcing RTP audio to this multicast group address and, upon joining the multicast group, phone B begins receiving the MoH stream. Once phone A presses the Resume softkey, Cisco Unified CallManager instructs phone B to Stop Multicast Media Reception. Phone B then sends an IGMP V2 Leave Group message to 224.0.0.2 to indicate that the multicast stream is no longer needed. This effectively ends the MoH session. Next, Cisco Unified CallManager sends a series of Open Receive Channel messages to phones A and B, just as would be sent at the beginning of a phone call between the two phones. Soon afterwards, Cisco Unified CallManager instructs both phones to Start Media Transmission to each other's IP addresses. The phones are once again connected via an RTP two-way audio stream.

**Note**

The call flow diagrams in [Figure 7-7](#) and [Figure 7-8](#) assume that an initial call is up between phones A and B, with a two-way RTP audio stream. These diagrams are representative of call flows and therefore include only the pertinent traffic required for proper MoH operation. Thus, keep-alives, acknowledgements, and other miscellaneous traffic have been eliminated to better illustrate the interaction. The initial event in each diagram is the Hold softkey action performed by phone A.

Figure 7-8 depicts a unicast MoH call flow. In this call flow diagram, when the Hold softkey is pressed at phone A, Cisco Unified CallManager instructs both phone A and phone B to Close Receive Channel and Stop Media Transmission. This action effectively stops the RTP two-way audio stream. Up to this point, unicast and multicast MoH call flows behave exactly the same.

Figure 7-8 Detailed Unicast MoH Call Flow



Next, Cisco Unified CallManager tells phone B (the holdee) to Open Receive Channel. (This is quite different from the multicast case, where Cisco Unified CallManager tells the holdee to Start Multicast Media Reception.) Then Cisco Unified CallManager tells the MoH server to Start Media Transmission to the IP address of phone B. (This too is quite different behavior from the multicast MoH call flow, where the phone is prompted to join a multicast group address.) At this point, the MoH server is sending a one-way unicast RTP music stream to phone B. When phone A presses the Resume softkey, Cisco Unified CallManager instructs the MoH server to Stop Media Transmission and instructs phone B to Close Receive Channel, effectively ending the MoH session. As with the multicast scenario, Cisco Unified CallManager sends a series of Open Receive Channel messages and Start Media Transmissions messages to phones A and B with each other's IP addresses. The phones are once again connected via an RTP two-way audio stream.



CHAPTER 8

Call Processing

Last revised on: February 13, 2008

This chapter provides design guidance for scalable and resilient call processing systems with Cisco Unified CallManager 4.2. This chapter also discusses how to choose the appropriate hardware and deployment scenario for Cisco Unified CallManager based on specific requirements that include the following:

- Scale — The number of users, gateways, applications, and so forth
- Performance — The call rate
- Resilience — The amount of redundancy

This chapter focuses on the following topics:

- [Cisco Unified CallManager Cluster Guidelines, page 8-2](#)

This section discusses the minimum hardware requirements for Cisco Unified CallManager. This section also explains the various feature services that can be enabled on the Cisco Unified CallManager server and the purpose of each.

- [Cisco Unified CallManager Platform Capacity Planning, page 8-12](#)

This section provides guidelines for using the Cisco CallManager Capacity Tool, which must be used when planning an IP Telephony deployment. The Cisco CallManager Capacity Tool provides guidance on resources used on a Cisco Unified CallManager server, based on certain deployment requirements.

- [Gatekeeper Design Considerations, page 8-18](#)

This section explains how gatekeepers can be used in an IP Telephony deployment. Cisco Gatekeeper may also be paired with other standby gatekeepers or may be clustered for higher performance and resilience. Gatekeepers may also be used for call routing and call admission control.

- [Interoperability of Cisco Unified CallManager and CallManager Express, page 8-27](#)

This section explains the H.323 integration between Cisco Unified CallManager and Cisco Unified CallManager Express in a distributed call processing deployment.

Cisco Unified CallManager Cluster Guidelines

The Cisco Unified CallManager architecture enables a group of physical servers to work together as a single IP PBX system. This grouping of servers is known as a *cluster*. A cluster of Cisco Unified CallManager servers may be distributed across an IP network, within design limitations, allowing for spatial redundancy and, hence, resilience to be designed into the IP Communications system.

This section describes the various functions performed by the servers that form a Cisco Unified CallManager cluster, and it provides guidelines for deploying the servers in ways that achieve the desired scale, performance, and resilience.

Hardware Platforms

Cisco Unified CallManager clusters utilize various types of servers, depending on the scale, performance, and redundancy required. They range from non-redundant, single-processor servers to highly redundant, multi-processor units.

Table 8-1 lists the general types of servers you can use in a cluster, along with their main characteristics.

Table 8-1 *Types of Cisco Unified CallManager Servers*

Server Type	Cisco Server Model	Characteristics
Standard server (not high availability)	MCS 7815 or equivalent	• Single processor • Single power supply • Non-RAID hard disk
High-availability standard server	MCS 7825 or equivalent	• Single processor • Multiple power supplies • Single SCSI RAID hard disk array
High-performance server	MCS 7835 and MCS 7845 or equivalent	• Multiple processors • Multiple power supplies • Multiple SCSI RAID hard disk arrays

For a complete list of currently supported hardware configurations, refer to the documentation available at

<http://www.cisco.com/go/swonly>

The server platforms have particular memory requirements, which are documented in *Product Bulletin 2864, Physical Memory Recommendations For Cisco CallManager Version 4.0 and Later*, available at

http://www.cisco.com/en/US/products/sw/voicesw/ps556/prod_bulletin0900aecd80284099.html

Servers should be deployed in an environment that provides high availability, not only for the IP network but also for power and cooling. Servers should be powered from an uninterruptible power supply (UPS) if building power does not have the required availability. Servers with dual power supplies could also be plugged into two different power sources to avoid the failure of one power circuit causing the server to fail.

Connectivity to the IP network should also ensure maximum performance and availability. The Cisco Unified CallManager servers should be connected to the Ethernet at 100 Mbps full-duplex. If 100 Mbps is not available on smaller deployments, then use 10 Mbps full-duplex. Many servers also include the capability of using Gigabit Ethernet, which is also an option. Ensure that servers are connected to the network using full-duplex, which can be achieved with 10 Mbps and 100 Mbps by hard-coding the switch port and the server NIC. For 1000 Mbps, Cisco recommends using Auto/Auto for speed and duplex configuration on both the NIC and the switch port.

**Note**

A mismatch will occur if either the server port or the Ethernet switch port is left in Auto mode and the other port is configured manually. The best practice is to configure both the server port and the Ethernet switch port manually, with the exception of Gigabit Ethernet ports which should be set to Auto/Auto.

Server platforms with dual Ethernet NICs can support NIC Teaming. This feature will depend on the manufacturer, and it allows a server to be connected to the Ethernet via two NICs and, hence, two cables. In some circumstances, this feature can be useful to provide additional levels of redundancy; however, careful design of the IP infrastructure is required prior to successful deployment. For additional information regarding the installation of teaming drivers on Cisco and Hewlett-Packard (HP) server platforms, refer to the documentation at

http://www.cisco.com/en/US/products/hw/voiceapp/ps378/prod_installation_guide09186a00803a0032.html

General Clustering Guidelines

The following guidelines apply to all Cisco Unified CallManager clusters:

**Note**

A cluster may contain a mix of server platforms, but all servers in the cluster must run the same Cisco Unified CallManager software release.

- Under normal circumstances, place all members of the cluster within the same LAN or MAN. Cisco does not recommend placing all members of a cluster on the same VLAN or switch.
- For redundancy, the members of the cluster should be deployed in the following manner to minimize the impact of any failures in the infrastructure or building:
 - Different access switches connected to the same distribution or core switch
 - Different access switches attached to different distribution or core switches
 - Different buildings within the same LAN or MAN
- If the cluster spans an IP WAN, follow the guidelines for clustering over an IP WAN as specified in the section on [Clustering Over the IP WAN](#), page 2-17.

Cisco Unified CallManager Cluster Services

Within a Cisco Unified CallManager cluster, there are servers that provide unique services. Each of these services can coexist with others on the same physical server. For example, in a small system it is possible to have a single server be a database publisher, backup subscriber, music on hold (MoH) server, TFTP server, CTI Manager, and Conference Bridge. As the scale and performance requirements of the cluster increase, many of these services should be moved to a single, dedicated physical server.


A Cisco Unified CallManager 4.2 cluster may contain as many as 20 servers, of which a maximum of eight may run the Cisco CallManager service that provides call processing. The other servers may be configured as a dedicated database publisher, dedicated Trivial File Transfer Protocol (TFTP) server, or music on hold (MoH) servers. Media streaming applications (conference bridge or media termination point) may also be installed on a separate server that registers with the cluster.

When deploying a cluster with Cisco MCS 7815 or equivalent servers, there is a minimum limit of two servers in a cluster: one as the publisher, TFTP server, and backup call processing server, and the other as the primary call processing server. A maximum of 300 phones is supported in this configuration on a Cisco MCS 7815. When deploying a two-server cluster with higher-capacity servers, Cisco recommends that you do not exceed 1250 users in the cluster. Above 1250 users, a dedicated publisher and separate servers for primary and secondary call processing services is recommended, thus increasing the number of servers in a cluster.

It is also possible to deploy a single-server cluster with an MCS 7825 or greater servers. With an MCS 7825 or equivalent server, the limit is 500 users; with a higher-availability server, the single-server cluster should not exceed 1000 users. In a single-server configuration, there is no redundancy unless Survivable Remote Site Telephony (SRST) is also deployed to provide service during periods when the Cisco Unified CallManager is not available. Cisco does not recommend a single-server deployment for production environments. The load balancing option is not available when the publisher is a backup call processing subscriber.

Figure 8-1 illustrates a typical Cisco Unified CallManager cluster.

Figure 8-1 Typical Cisco Unified CallManager Cluster



Intracuster Communications


There are two primary kinds of communication within a Cisco Unified CallManager cluster (intracuster communications). (See Figure 8-2.) The first is a mechanism for distributing the database that contains all the device configuration information (see Database Replication in Figure 8-2). The configuration database is stored on a publisher server, and a read-only copy is replicated to the subscriber members of the cluster. Changes made on the publisher are communicated to the subscriber databases, ensuring that the configuration is consistent across the members of the cluster, as well as facilitating spatial redundancy of the database.

The second type of Intracuster communication is the propagation and replication of run-time data such as registration of devices, locations bandwidth, and shared media resources (see ICCS in Figure 8-2). This information is shared across all members of a cluster running the Cisco Unified CallManager Service, and it assures the optimum routing of calls between members of the cluster and associated gateways.

Lightweight Directory Access Protocol (LDAP) directory information is also replicated between all servers in a cluster. The LDAP directory is replicated by the publisher to all other servers. Cisco Unified CallManager may be integrated into a corporate LDAP directory, such as Microsoft Active Directory or Netscape Directory. The replication is then dependant on the integration method deployed and is outside the scope of this document. For more information on directory integration, refer to the chapter on [LDAP Directory Integration, page 18-1](#).

Figure 8-2 Intracluster Communications

Intra-Cluster Communication Signaling (ICCS) SQL and directory replication



Publisher

The publisher is a required server in all clusters, and there can currently be only one per cluster. This server is the first to be installed and provides the database services to all other members in the cluster. The publisher server is the only server that has read and write access to the configuration database. When configuration changes are made, other members of the cluster have a read-only copy of the database. On larger systems with more than 1250 users, Cisco recommends a dedicated publisher to prevent administrative operations from affecting the telephony users. A dedicated publisher does not have any call processing services or TFTP services running on the server. Other servers run the TFTP and Cisco Unified CallManager services.

Servers in the cluster will attempt to use the publisher's database when initializing. If the publisher is not available, they will use the local read-only copy from their hard drives.

If the system is running but the publisher is not available, the following operations will not be available:

- Call forwarding changes
- Configuration changes
- Extension Mobility login or logout operations, and Device Mobility

Extension Mobility will not function without the publisher because it requires access to the read/write database. Therefore, Cisco recommends running this service on the publisher only. Device Mobility will not update the device pool settings for the phone when roaming to another site. The phone uses its home device pool settings if the publisher is not available.

The choice of hardware platform for the publisher is based on the scale and performance of the cluster. Cisco recommends that the publisher have the same performance capability as the call processing subscribers. Ideally the publisher should also be a high-availability server to minimize the impact of a hardware failure.

Call Processing Subscriber

When installing the Cisco Unified CallManager software, you can define two types of servers, publisher and subscriber. These terms are used to define the database relationship during installation. Once the software is installed, only the database and network services are enabled. All subscribers will subscribe to the publisher to obtain a read-only copy of the database information.

A call processing subscriber is a server that has the Cisco CallManager Service enabled. A single server license is required to enable this service on a subscriber. The Cisco CallManager Service cannot be enabled on a server if the publisher is not available because the publisher acts as a licensing server and distributes the licenses needed to activate the Cisco CallManager Service. Once this service is enabled, the server is able to perform call processing functions. Devices such as phones, gateways, and media resources can register and make calls only to servers with this service enabled.

Cisco Unified CallManager 4.2 supports up to eight servers in a cluster with the Cisco CallManager Service enabled.

Depending on the redundancy scheme chosen (see [Call Processing Redundancy, page 8-6](#)), the call processing subscriber will be either a primary (active) subscriber or a backup (standby) subscriber. In the load-balancing option, the subscriber can be both a primary and backup subscriber. When planning the design of a cluster, you should generally dedicate the call processing subscribers to this function. In larger-scale or higher-performance clusters, the call processing service should not be enabled on the publisher and TFTP server. Call processing subscribers normally operate in either dedicated pairs or shared pairs, depending on the redundancy scheme adopted. One-to-one redundancy uses dedicated pairs, while two-to-one redundancy uses two pairs of servers that share one server from each pair (the backup server).

The choice of hardware platform depends on the scale, performance, redundancy, and cost of the servers. Scale and performance are covered in the section on [Cisco Unified CallManager Platform Capacity Planning, page 8-12](#), and redundancy is covered in the section on [Call Processing Redundancy, page 8-6](#).

Call Processing Redundancy

With Cisco Unified CallManager 4.2, you can choose from the following redundancy configurations:

- Two to one (2:1) — For every two primary subscribers, there is one shared backup subscriber.
- One to one (1:1) — For every primary subscriber, there is a backup subscriber.

The 1:1 redundancy scheme allows upgrades with only the failover periods impacting the cluster. The 1:1 redundancy scheme enables you to upgrade the cluster using the following method:

-
Step 1	Upgrade the publisher server.
Step 2	Upgrade dedicated TFTP servers one at a time; reboot the server and wait for configuration files to be rebuilt before upgrading any further servers in the cluster.
Step 3	Upgrade dedicated music on hold (MoH) and other media resource servers one at a time.
Step 4	Upgrade the backup subscribers one at a time. This step will impact some users if 50/50 load balancing is implemented.
Step 5	Fail-over the primary subscribers to their backups, and stop the Cisco CallManager Service on the primaries.
Step 6	Upgrade the primaries one at a time, then re-enable the Cisco CallManager Service.
-

With this upgrade method, there is no period (except for the failover period) when devices are registered to subscriber servers that are running different versions of the Cisco Unified CallManager software. This factor can be important because the Intra-Cluster Communication Signaling (ICCS) protocol that communicates between subscribers can detect a different software version and shut down communications to that subscriber.

The 2:1 redundancy scheme allows for fewer servers in a cluster, but it can potentially result in an outage during upgrades.

**Note**

You must use 1:1 redundancy when 10,000 or more IP phones are registered on the two primary subscribers because there cannot be more than 10,000 backup registrations on a single backup subscriber.

**Note**

Before you do an upgrade, Cisco recommends that you back up the Cisco Unified CallManager and Call Detail Record (CDR) database to an external network directory using the Disaster Recovery Framework. This practice will prevent any loss of data if the upgrade fails.

Call Processing Subscriber Redundancy

The following figures illustrate typical cluster configurations to provide call processing redundancy with Cisco Unified CallManager.

Figure 8-3 Basic Redundancy Schemes






Figure 8-3 illustrates the two basic redundancy schemes available. In each case the backup server must be capable of handling the capacity of at least a single primary call processing server failure. In the 2:1 redundancy scheme, the backup might have to be capable of handling the failure of a single call processing server or potentially both primary call processing servers, depending on the requirements of a particular deployment. Sizing the capacity of the servers and choosing the hardware platforms is covered in the section on [Cisco Unified CallManager Platform Capacity Planning, page 8-12](#).

Figure 8-4 1:1 Redundancy Configuration Options

114952

In [Figure 8-4](#) the five options shown all indicate 1:1 redundancy. Option 1 is used for clusters supporting less than 1250 users. Options 2 through 5 illustrate increasingly scalable clusters. The exact scale will depend on the hardware platforms chosen or required.

Note that the illustrations show only publisher and call processing subscribers.

Figure 8-5 2:1 Redundancy Configuration Options

114953

Load Balancing

Normally a backup server has no devices registered to it unless its primary is unavailable. This model allows for:

- Easier troubleshooting — Because all call processing takes place on primary servers, obtaining traces and alert notifications becomes easier.
- Less configuration — Because all the devices are registered to the primary server, the need to define additional Cisco Unified CallManager redundancy groups or device pools for the various devices can be reduced by 50%.

The 1:1 redundancy scheme enables you to balance distribution of the devices over the primary and backup server pairs. With load balancing, you can move up to half of the device load from the primary to the secondary subscriber by using the Cisco Unified CallManager redundancy groups and device pool settings. This model allows for:

- Load sharing — The call processing load is distributed on multiple servers, which can provide faster response time.
- Faster failover and failback — Because all devices (such as IP phones, CTI ports, gateways, trunks, voicemail ports, and so forth) are distributed across all active subscribers, only some of the devices fail over to the secondary subscriber if the primary subscriber fails. In this way, you can reduce by 50% the impact of any server becoming unavailable.

To plan for 50/50 load balancing, calculate the capacity of a cluster without load balancing, and then distribute the load across the primary and backup subscribers based on devices and call volume. To allow for failure of the primary or the backup server, the total load on the primary and secondary subscribers should not exceed that of a single subscriber server.

TFTP Server

The TFTP server performs two main functions:

- The serving of files for services such as MoH, configuration files for devices such as phones and gateways, binary files for the upgrade of phones as well as some gateways, and various security files.
- Generation of configuration and security files. Most files generated by the Cisco TFTP service are signed and in some cases encrypted before being available for download.

The TFTP service can be enabled on any server in the cluster. However, in a cluster with more than 1250 users, other services might be impacted by configuration changes that can cause the TFTP service to regenerate configuration files. Therefore, Cisco recommends that you dedicate a specific server to the TFTP service in a cluster with more than 1250 users, with Extension Mobility, or with other features that cause configuration changes.

The TFTP server is used by phones and MGCP gateways to obtain configuration information. There is no restriction on the number of servers that can have TFTP service enabled, however Cisco recommends deploying 2 TFTP servers for a large cluster, thus providing redundancy for TFTP service. More than 2 TFTP servers can be deployed in a cluster, but this can result in an extended period for rebuilding of all TFTP files on all TFTP servers. When configuring the TFTP options using DHCP or statically, you can normally define an IP address array (more than one IP address) for a TFTP server. Therefore, you can assign half of the devices to use TFTP server A as the primary and TFTP server B as the backup, and the other half to use TFTP server B as the primary and TFTP server A as the backup. To improve performance on dedicated TFTP servers, you can set service parameters to increase the number of simultaneous TFTP sessions allowed on the server.

When upgrading a Cisco Unified CallManager cluster, Cisco highly recommends that you upgrade the TFTP servers after the publisher and before any other server, also allowing additional time following the upgrade for the TFTP server to rebuild all the configuration files. Either use the typical Cisco TFTP - BuildDuration time or use the real-time monitoring tool to monitor the Cisco TFTP - DeviceBuildCount until it stops incrementing. This upgrade order ensures that any new binaries and configuration changes are available before the upgrade of other services in the cluster. If you are manually adding a specific binary or firmware load for a phone or gateway, be sure to copy the file to each TFTP server in the cluster.

By default, Cisco Unified CallManager Release 3.3 and later versions cache the configuration files in memory and do not store them on the hard drive of the TFTP server. This default setting can be changed to place the configuration files on the hard drive of the TFTP server, but doing so will impact TFTP performance. Therefore, Cisco recommends that you do not change this default setting.

Cisco recommends that you use the same hardware platform for the TFTP servers as used for the call processing subscribers.

CTI Manager

CTI Manager is required in a cluster for applications that use TAPI or JTAPI Computer Telephony Integration (CTI). The CTI Manager acts as a broker between the CTI application and the Cisco Unified CallManager service. It provides authentication of the application and enables control or monitoring of authorized devices. The CTI application communicates with a primary CTI Manager and, in the event of a failure, will switch to a backup CTI Manager. The CTI Manager should be enabled only on call processing subscribers, thus allowing for a maximum of eight CTI Managers in a cluster. Cisco recommends that you load-balance CTI applications across the various CTI Managers in the cluster to provide maximum resilience, performance, and redundancy.

Generally, it is good practice to associate devices that will be controlled or monitored by an application with the same server pair used for the CTI Manager. For example, an interactive voice response (IVR) application requires four CTI ports. They would be provisioned as follows, assuming the use of 1:1 redundancy and 50/50 load balancing:

- Two CTI Ports would have a Cisco Unified CallManager redundancy group of server A as the primary and server B as the backup (or secondary). The other two ports would have a Cisco Unified CallManager redundancy group of server B as the primary and server A as the backup.
- The IVR application would be configured to use the CTI Manager on server A as the primary and server B as the backup.

The above example allows for redundancy in case of failure of the CTI Manager on server A and also allows for the IVR call load to be spread across two servers. This approach also minimizes the impact of a Cisco Unified CallManager server failure.

IP Voice Media Streaming Application

Media resources such as conferencing and music on hold may be provided by IP Voice Media Streaming Application services running on the same physical server as the Cisco Unified CallManager service.

Media resources include:

- Music on Hold (MoH) — Provides multicast or unicast music to devices that are placed on hold or temporary hold, transferred, or added to a conference. (See [Music on Hold, page 7-1](#).)
- Annunciator service — Provides announcements in place of tones to indicate incorrectly dialed numbers or call routing unavailability. (See [Annunciator, page 6-17](#).)
- Conference bridges — Provide software-based conferencing for ad-hoc and meet-me conferences. (See [Audio Conferencing, page 6-7](#).)
- Media termination point (MTP) services — Provide features for H.323 clients, H.323 trunks, and Session Initiation Protocol (SIP) trunks. (See [Media Termination Point \(MTP\), page 6-12](#).)

Because of the additional processing and network requirements for media, it is essential to follow all guidelines for running media resources within a cluster. Generally, Cisco recommends non-dedicated servers for multicast MoH and the annunciator, but a dedicated media resource server is recommended for large-scale software-based conferencing and MTP unless those services are within the design guidelines detailed in the chapters on [Media Resources, page 6-1](#), and [Music on Hold, page 7-1](#).

Voice Activity Detection

Cisco also recommends that you leave voice activity detection (VAD) disabled within the cluster. VAD is disabled by default in the Cisco CallManager service parameters, and you should disable it on H.323 dial peers by using the **no vad** command.

Extension Mobility

When Extension Mobility is used in a cluster, the expected capacity and performance can impact the choice of the publisher hardware platform. As a user logs in or logs out from a phone, the configuration must be updated in the configuration database, the TFTP service has to regenerate the configuration file, and then the device is reset to make the change. Most of this takes place on the publisher.

Currently supported limits for Extension Mobility are as follows:

- Cisco Unified CallManager releases 3.3 and 4.0
 - 33 sequential logins/logouts per minute
- Cisco Unified CallManager release 4.1 and later 4.x releases
 - 33 sequential logins/logouts per minute
 - A Cisco MCS 7845 publisher can support 75 sequential logins/logouts per minute.

Cisco Unified CallManager Platform Capacity Planning

Many types of devices can register with Cisco Unified CallManager; for example, IP phones, voicemail ports, CTI (TAPI or JTAPI) devices, gateways, and DSP resources such as transcoding and conferencing. Each of these devices requires resources from the server platform with which it is registered. The required resources can include memory, processor usage, and disk I/O. Each device then consumes additional server resources during transactions, which are normally in the form of calls. For example, a device that makes only 6 calls per hour consumes fewer resources than a device making 12 calls per hour.

The recommendations provided in this section are based on calculations made using the Cisco Unified CallManager Capacity Calculator with default trace levels and CDRs enabled. Higher levels of performance can be achieved by disabling, reducing, or reconfiguring other functions that are not directly related to processing calls. Increasing some of these functions can also have an impact on the call processing capabilities of the system. These functions include tracing, call detail recording, highly complex dial plans, and other services that are co-resident on the server. Highly complex dial plans can include multiple line appearances as well as large numbers of partitions, calling search spaces, route patterns, translations, route groups, hunt groups, pickup groups, route lists, extensive use of call forwarding, co-resident services, and other co-resident applications. All of these functions can consume additional resources within the Cisco Unified CallManager server.

To improve system performance, the following techniques can provide useful options:

- Install additional certified memory in the server, up to the maximum supported for the particular platform. Doubling the RAM in MCS 7825 and MCS 7835 or equivalent servers is recommended in large configurations for that server class. Verification using a performance monitor will indicate if this memory upgrade is required. As the server approaches maximum utilization of physical memory, the operating system will start to swap to disk. This swapping is a sign that additional physical memory should be installed. For specific server memory recommendations, refer to *Product Bulletin 2864, Physical Memory Recommendations For Cisco CallManager Version 4.0 and Later*, available at http://www.cisco.com/en/US/products/sw/voicesw/ps556/prod_bulletin0900aecd80284099.html
- On supported hardware, make sure the battery-backed write cache (BBWC) is enabled with 50% write capability. This feature might require the addition of a battery to the Small Computer System Interface (SCSI) controller card.
- On MCS-7845 or equivalent servers, set the trace files location to the F: drive. This setting is a Cisco Unified CallManager service parameter.
- Reformat the trace file partition to use smaller block sizes. (Refer to the release notes for the version of the operating system you have installed.) A utility is provided for this purpose.
- A Cisco Unified CallManager cluster with a very large dial plan containing many gateways, route patterns, translation patterns, and partitions, can take an extended amount of time to initialize when the Cisco CallManager Service is first started. If the system does not initialize within the default

time, there are service parameters that can be modified to allow additional time for the configuration to initialize. For details on the service parameters, refer to the online help for Service Parameters in Cisco Unified CallManager Administration.

The following guidelines apply to Cisco Unified CallManager Release 4.2:

- Within a cluster, you may enable a maximum of 8 servers with the Cisco Unified CallManager Service. Other servers may be used for more dedicated functions such as TFTP, publisher, music on hold, and so forth.
- You can configure a maximum of 800 CTI connections or associations per standard server, or a maximum of 3200 per cluster if they are equally balanced among the servers.
- You can configure a maximum of 2500 CTI connections or associations per high-performance server, or a maximum 10,000 per cluster if they are equally balanced among the servers.
- Each cluster can support a maximum of 30,000 SCCP phones.
- Each cluster can support a maximum of 600 H.323 devices (gateways, trunks, and clients), digital MGCP devices, and SIP trunks

Capacity Calculations

The Cisco CallManager Capacity Tool for Cisco Unified CallManager Release 4.x enables you to calculate the capacity of the system for various configurations. The capacity planning tool is currently available to all www.cisco.com users with a login account. If your system does not meet the following guidelines, or if you consider the system to be more complex and would like to verify the capacity but you do not have access to the Cisco CallManager Capacity Tool, please contact your Cisco Systems Engineer (SE) or the Cisco Technical Assistance Center (TAC).

The Cisco CallManager Capacity Tool is available (with proper login authentication) at

<http://www.cisco.com/partner/WWChannels/technologies/resources/CallManager/>

If your system meets the following requirements, you should not need to verify your configuration with the Cisco CallManager Capacity Tool:

- System contains less than 25% of the maximum number of users for the server platforms
- Average number of lines per phone does not exceed 1.1
- Average busy hour call attempts (BHCA) per user is below 4 calls per hour
- Cluster security is not enabled
- No more than 20% trunking (5 users per trunk) either via gateways or trunks
- No more than 5% voicemail ports (20 users per voicemail port)
- No more than 20 MoH streams per server
- No CTI, JTAPI, or TAPI devices
- No more than 5% conference bridges (20 users per port on a bridge)
- No transcoders
- Only MTPs are used to support the maximum calls required for trunking
- Less than 20 locations
- System contains nothing else that is not defined above and is not an IP Phone, IP Communicator, gateway, media resource, voicemail port, or trunk

The maximum number of users Cisco Unified CallManager can support depends on the server platform, as indicated in [Table 8-2](#).

Table 8-2 *Maximum Number of Devices per Server Platform*

Server Platform Characteristics	Maximum Users per Server ¹	High-Availability Server ²	High-Performance Server
Cisco MCS-7845 (All supported models)	7500	Yes	Yes
Cisco MCS-7835 (All supported models)	2500	Yes	No
Cisco MCS-7825 (All supported models)	1000	No	No
Cisco MCS-7815 (All supported models) ³	300	No	No

1. A platform that is not a high-availability server can support a maximum of 500 IP Phones in a non-redundant installation.
2. A high-availability server supports redundancy for both the power supplies and the hard disks.
3. MCS-7815 servers support only N+1 redundancy and may not be a member of a cluster.

For the latest information on supported platforms, third-party platforms, and specific hardware configurations, refer to the online documentation at

<http://www.cisco.com/go/swonly>



Note

The maximum supported non-redundant installation on platforms that are not highly available is 500 IP phones.

Cisco CallManager Capacity Tool

The Cisco CallManager Capacity Tool requires various pieces of information to provide a calculation of the minimum size and type of servers required for a system. The information includes the type and quantity of devices, such as IP phones, gateways, and media resources. For each device type, the capacity tool also requires the average BHCA and the average utilization time. For example, if all IP phones make an average of 5 calls per hour and the average call lasts 3 minutes, then the BHCA is 5 and the utilization is 0.25. (Five calls of 3 minutes each equals 15 minutes per hour on the phone, which is 0.25 hour.)

In addition to the device information, the capacity tool also requires information regarding the dial plan, such as route patterns and translation patterns. When all the details have been entered, the capacity tool will calculate how many primary servers of the desired server type are required, as well as the number of clusters if the required capacity exceeds a single cluster.

The capacity tool replaces the previous mechanisms known as device weights, BHCA multipliers, call type multipliers, and dial plan weights.

When you use the Cisco CallManager Capacity Tool, the results will indicate the highest capacity of the various resources calculated for the configuration. The resources include hard limits, memory, processor utilization, and disk I/O activity. If you add more devices, it might look as if no more capacity is being used because the additional configuration is using resources that are below the current highest capacity indicated.

For example, if 2500 third-party controlled SCCP IP phones were added to the Cisco CallManager Capacity Tool, it would show that a single MCS 7845 subscriber is required and is at 100% capacity. If an additional 1000 phones are added, the results still indicate a single MCS 7845 subscriber is required and is at 100% capacity. This result is due to the fact that CTI capacity for this server is at 100%, and

additional IP phones do not add CTI requirements. Until the number of other devices exceeds one of the other limits, it will appear as if no additional capacity is being used. This is normal and expected behavior for the Cisco CallManager Capacity Tool.

Use the guidelines in the following sections to provide information for the capacity calculator.

Phone Calculations

The main Telephony section (not the Contact Center section) of the Capacity Tool lists the following phone types:

- SCCP Phones (Non-secure)
- Secure SCCP Phones

For each phone type, you can enter quantity, BHCA, utilization, and line appearance values, which are further explained in the following sections.

Quantity

This value is the total number of IP phones of each type that will be configured in the cluster. The quantity includes all Cisco 7900 Series IP Phones, VG248 ports, VG224, IP Communicator, and other third-party SCCP endpoint devices. The quantity must include all configured phones even if they are not active.

BHCA

This value is the average BHCA of all phones of each type. If you have shared lines across multiple phones, the BHCA should include one call for each and every phone that shares that line. Shared lines across multiple phone types will affect the BHCA for each type. (That is, one call to a shared line will be calculated as multiple calls, one to each phone that rings.) If you have different groups of phones that generate different BHCAs, use the following method to provide a BHCA value for use in the Cisco CallManager Capacity Tool.

For example, assume there are two classes of user with the following characteristics:

- 100 phones at 20 BHCA = 2,000 BHCA total
- 5000 phones at 4 BHCA = 20,000 BHCA total

The total BHCA for all phone devices in this case is 22,000.

Then divide the total BHCA by the total number of phone devices, to yield:

$$\text{Average BHCA per phone device} = 22,000 / 5,100 = 4.31 \text{ BHCA}$$

Utilization

This value is the average call utilization per phone, and it includes all the time that a call is on the phone. The actual utilization of a phone can exceed 100% if that phone allows multiple calls per line or has multiple line appearances that also are frequently utilized. Utilization is measured as a percentage of an hour. For example, a phone that is on a call for 3 minutes in the busiest hour is considered to be 5% utilized. The same method for calculating BHCA can also be used to calculate the average utilization for various groups of phone. If the phone has shared lines, then only the expected utilization for that phone should be calculated and not the actual utilization for the line that is shared. The Cisco CallManager Capacity Tool currently accepts an average utilization for all phones.

Lines Appearances

This value is the average number of lines for all the phones. If a phone has multiple appearances of the same DN in different partitions, it will be counted as multiple line appearances. Shared lines should be counted as one line, but ensure that the correct BHCA and utilization are calculated. Multiple calls per

line do not increase the line appearance count, but they do affect the BHCA and utilization calculations. For example, if it is normal for a phone to have two calls, one active and the other on hold for various amounts of time, then the utilization will be higher for that device because two calls are actually connected.

Gateways

Quantity of Gateways

The quantity of gateways is split into several entries because it depends on gateway type, as follows:

- MGCP T1/E1 gateways

This value is the total number of gateways that need to be configured in the Cisco Unified CallManager database. For example, a Cisco IOS MGCP gateway can have multiple T1s or E1s, but they would be added as a single gateway. A Cisco WS-6608 module would be added as a gateway per port that is configured as a T1 or E1 gateway, up to the maximum of eight per module.

- MGCP analogue gateways

This value is the total number of analogue gateways to be added in the Cisco Unified CallManager database. Generally, an entire module (WS-6624 or Cisco Unified CallManager) or hardware platform (Cisco IOS Router platform) would be added as a single gateway.

- H.323 gateways

An entire module or hardware platform would be added as a single gateway. This quantity does not include H.323 gateways that are not defined in Cisco Unified CallManager but are used via H.323 trunks.

Number of DS0s

This value is the total number of DS0s or analogue ports supported by each type of gateway. The number of DS0s is split between the different types of gateways as follows:

- T1 CAS:
 $24 * (\text{Number of T1 CAS spans})$
- T1 (E1) PRI:
 $(23 \text{ or } 30) * (\text{Total number of PRIs})$
- H.323 gateways
 $(\text{Total number of DS0s}) / (\text{Number of calls supported on all digital, analogue, or IP interfaces})$

BHCA

The BHCA is the average for all DS0s or analogue ports on a gateway during the busiest hour. The method for calculating this average is the same as the method used for the phone BHCA.

Utilization

This value is the average utilization of all DS0s or analogue ports during the busiest hour.

EM Profiles

Extension Mobility (EM) profiles contain line appearances that should also be included in the phone calculations. They will not affect the quantity of devices, but they will increase the average number of line appearances per phone. The BHCA and utilization of an EM user should already have been calculated for the phone they will log into.

H.323 and SIP Trunks

Quantity of Trunks

This value is the total number of trunks that will be configured in the Cisco Unified CallManager database. A gatekeeper controlled trunk is not affected by the possible number of destinations and, therefore, will be counted once for each configured gatekeeper controlled trunk. This situation is also the case for Session Initiation Protocol (SIP) trunks using a SIP proxy.

Quantity of Calls

This value is the total number of concurrent calls expected to be allowed on all trunks. The number of allowed calls will generally be controlled by locations or gatekeeper call admission control. Keep in mind that regions and codecs can also affect the total number of calls allowed.

Utilization

This value is the average utilization of all calls across all trunks. It is a percentage of the busiest hour, where 75% utilization means that calls will be active for 45 minutes per hour.

MTP Requirements

MTP requirements must also be considered separately in the calculator. If you have MTP required on an H.323 trunk, then an MTP resource is required for every concurrent call on that trunk. MTP resources are always required for SIP trunks.

CTI Route Points, Ports, and Third-Party Controlled Lines

CTI route point calculations include the total quantity, average BHCA, and utilization values. Any assigned directory numbers should be added to the Dial Plan section of the tool.

CTI ports can be used in a variety of ways by the applications controlling them. For simplicity they have been divided into two groups:

- Simple call or redirect call

A simple call is a call either to or from the port to another device, with no supplementary service being invoked. It is the typical call scenario where one party calls another, the called party answers, the call proceeds, and the parties hang up.

A redirect is a call to a CTI port that is not actually answered. Instead, it is forwarded to another destination or redirected. This is even simpler than a simple call because no call is connected, so no media is connected to the CTI port.

- Transfer or conference

These call types are typical of an IVR or auto-attendant application.

A transfer type call is characterized as a call that is connected to a CTI port, then at some point the CTI port transfers the caller without consultation (or blind transfer) to another destination.

A conference is a variation of the transfer, whereby a call connected to the CTI port is either transferred with consultation (three-way call) or is conferenced with another party for the duration of the call.

Third-Party Controlled Lines

Each application that requires third-party control or monitoring of a device or line must be counted and entered in this field of the calculator. If more than one application is monitoring or controlling the same device, it will need to be counted more than once.

Gatekeeper Design Considerations

A single Cisco IOS gatekeeper can provide call routing and call admission control for up to 100 Cisco Unified CallManager clusters in a distributed call processing environment. Multiple gatekeepers can be configured to support thousands of Cisco Unified CallManager clusters. You can also implement a hybrid Cisco Unified CallManager and toll-bypass network by using Cisco IOS gatekeepers to provide communication and call admission control between the H.323 gateways and Cisco Unified CallManager.

Gatekeeper call admission control is a policy-based scheme requiring static configuration of available resources. The gatekeeper is not aware of the network topology, so it is limited to hub-and-spoke topologies.

The Cisco 2600, 2800, 3600, 3700, 3800, and 7200 Series routers all support the gatekeeper feature. You can configure Cisco IOS gatekeepers in a number of different ways for redundancy, load balancing, and hierarchical call routing. This section considers the design requirements for building a gatekeeper network, but it does not deal with the call admission control or dial plan resolution aspects, which are covered in the chapters on [Call Admission Control, page 9-1](#), and [Dial Plan, page 10-1](#), respectively.

For additional information regarding gatekeepers, refer to the *Cisco IOS H.323 Configuration Guide*, available at

http://www.cisco.com/en/US/products/sw/iosswrel/ps5207/products_configuration_guide_book09186a00801fcee1.html

Hardware Platform Selection

The choice of gatekeeper platform is based on the number of calls per second and the number of concurrent calls. A higher number of calls per second requires a more powerful CPU, such as a Cisco 3700, 3800, or 7200 Series Router. A higher number of concurrent calls requires more memory. For the latest information on platform selection, contact your Cisco partner or Cisco Systems Engineer (SE).

Gatekeeper Redundancy

With gatekeepers providing all call routing and admission control for intercluster communications, redundancy is required. Prior to Cisco Unified CallManager Release 3.3, the only method for providing gatekeeper redundancy was Hot Standby Router Protocol (HSRP); however, beginning with Cisco Unified CallManager Release 3.3, gatekeeper clustering and redundant gatekeeper trunks are also available as methods of providing gatekeeper redundancy. The following sections describe these methods.

**Note**

Cisco recommends that you use gatekeeper clustering to provide gatekeeper redundancy whenever possible. Use HSRP for redundancy only if gatekeeper clustering is not available in your software feature set.

Hot Standby Router Protocol (HSRP)


Hot Standby Router Protocol (HSRP) is the only option for gatekeeper redundancy with Cisco Unified CallManager prior to Release 3.3. HSRP does not provide the features required to build a redundant and scalable gatekeeper network, therefore it should be used only in Cisco Unified CallManager environments prior to Release 3.3.

The following guidelines apply to HSRP:

- Only one gatekeeper is active at a time.
 - The standby gatekeeper does not process any calls unless the primary fails.
 - No load balancing features are available.
- All gatekeepers must reside in the same subnet or location.
- No previous state information is available after failover.
- After a failover, the standby gatekeeper is not aware of the calls that are already active, so over-subscription of the bandwidth is possible.
- Failover time can be substantial because the endpoints have to re-register with the HSRP standby gatekeeper before calls can be placed. The failover time depends on the settings of the registration timers.

Figure 8-6 show a network configuration using HSRP for gatekeeper redundancy.

Figure 8-6 Gatekeeper Redundancy Using HSRP



Example 8-1 shows the configuration for Gatekeeper 1 and Example 8-2 shows the configuration for Gatekeeper 2 in Figure 8-6. Both configurations are identical except for the HSRP configuration on the Ethernet interface.

Example 8-1 Configuration for Gatekeeper 1

```
interface Ethernet0/0
ip address 10.1.10.2 255.255.255.0
 standby ip 10.1.10.1
 standby priority 110

gatekeeper
zone local GK-Site1 customer.com 10.1.10.1
zone local GK-Site2 customer.com
zone prefix GK-Site1 408.....
zone prefix GK-Site2 212.....
bandwidth interzone default 160
gw-type-prefix 1#* default-technology
arq reject-unknown-prefix
no shutdown
```

Example 8-2 Configuration for Gatekeeper 2

```
interface Ethernet0/0
ip address 10.1.10.3 255.255.255.0
 standby ip 10.1.10.1

gatekeeper
zone local GK-Site1 customer.com 10.1.10.1
zone local GK-Site2 customer.com
zone prefix GK-Site1 408.....
zone prefix GK-Site2 212.....
bandwidth interzone default 160
gw-type-prefix 1#* default-technology
arq reject-unknown-prefix
no shutdown
```

The following notes also apply to [Example 8-1](#) and [Example 8-2](#):

- Each router has **standby** commands configured for HSRP and to identify the virtual IP address shared by each. Gatekeeper 1 is configured as the primary with the command **standby priority 110**.
- Each Cisco Unified CallManager cluster has a local zone configured on each router to support Cisco Unified CallManager trunk registrations. Note that the IP address defined on the first zone should match the virtual IP address used by HSRP.
- A zone prefix is configured for each zone on both routers, allowing inter-zone and inter-cluster call routing.
- Bandwidth statements are configured on each router for both sites. Cisco recommends that you use the **bandwidth interzone** command because the **bandwidth total** command does not work in some configurations.
- The **gw-type-prefix** command is configured on both routers, allowing all locally unresolved calls to be forwarded to a device registered with a technology prefix of 1# in the local zone. In this example, all Cisco Unified CallManager trunks have been configured to register with a 1# prefix.
- The **arq reject-unknown-prefix** command is configured on both routers to guard against potential call routing loops across redundant Cisco Unified CallManager trunks.

For additional and advanced HSRP information, refer to the online documentation at the following locations:

- <http://www.cisco.com/en/US/docs/internetworking/case/studies/cs009.html>
- http://www.cisco.com/en/US/tech/tk648/tk362/technologies_q_and_a_item09186a00800a9679.shtml
- http://www.cisco.com/en/US/tech/tk648/tk362/technologies_tech_note09186a0080094afd.shtml

Gatekeeper Clustering (Alternate Gatekeeper)


Gatekeeper clustering (alternate gatekeeper) enables the configuration of a "local" gatekeeper cluster, with each gatekeeper acting as primary for some Cisco Unified CallManager trunks and an alternate for others. Gatekeeper Update Protocol (GUP) is used to exchange state information between gatekeepers in a local cluster. GUP tracks and reports CPU utilization, memory usage, active calls, and number of registered endpoints for each gatekeeper in the cluster. Load balancing is supported by setting thresholds for any of the following parameters in the GUP messaging:

- CPU utilization
- Memory utilization
- Number of active calls
- Number of registered endpoints

With the support of gatekeeper clustering (alternate gatekeeper) and Cisco Unified CallManager Release 3.3 or later, stateful redundancy and load balancing is available. Gatekeeper clustering provides the following features:

- Local and remote clusters
- Up to five gatekeepers in a local cluster
- Gatekeepers in local clusters can be located in different subnets or locations
- No failover delay (Because the alternate gatekeeper is already aware of the endpoint, it does not have to go through the full registration process.)
- Gatekeepers in a cluster pass state information and provide load balancing

Figure 8-7 shows three sites with Cisco Unified CallManager distributed call processing and three distributed gatekeepers configured in a local cluster.

Figure 8-7 Gatekeeper Clustering

In [Figure 8-7](#), Gatekeeper 2 is the backup for Gatekeeper 1, Gatekeeper 3 is the backup for Gatekeeper 2, and Gatekeeper 1 is the backup for Gatekeeper 3.

[Example 8-3](#) shows the configuration for Gatekeeper 1 (SJC), and [Example 8-4](#) shows the configuration for Gatekeeper 2 (CHC). The configuration for Gatekeeper 3 (NYC) is not shown because it is very similar to the other two.

Example 8-3 Gatekeeper Clustering Configuration for Gatekeeper 1

```
gatekeeper
zone local SJC cisco.com 10.1.1.1
zone local CHC_GK1 cisco.com
zone local NYC_GK1 cisco.com
!
```

```

zone cluster local SJC_Cluster SJC
  element SJC_GK2 10.1.2.1 1719
  element SJC_GK3 10.1.3.1 1719
!
zone cluster local CHC_Cluster CHC_GK1
  element CHC 10.1.2.1 1719
  element CHC_GK3 10.1.3.1 1719
!
zone cluster local NYC_Cluster NYC_GK1
  element NYC 10.1.3.1 1719
  element NYC_GK2 10.1.2.1 1719
!
zone prefix SJC 40852.....
zone prefix NYC_GK1 21251.....
zone prefix CHC_GK1 72067.....
gw-type-prefix 1#* default-technology
load-balance cpu 80 memory 80
bandwidth interzone SJC 192
bandwidth interzone NYC_GK1 160
bandwidth interzone CHC_GK1 160
arq reject-unknown-prefix
no shutdown

```

Example 8-4 Gatekeeper Clustering Configuration for Gatekeeper 2

```

gatekeeper
zone local CHC cisco.com 10.1.2.1
zone local SJC_GK2 cisco.com
zone local NYC_GK2 cisco.com
!
zone cluster local CHC_Cluster CHC
  element CHC_GK3 10.1.3.1 1719
  element CHC_GK1 10.1.1.1 1719
!
zone cluster local SJC_Cluster SJC_GK2
  element SJC 10.1.1.1 1719
  element SJC_GK3 10.1.3.1 1719
!
zone cluster local NYC_Cluster NYC_GK2
  element NYC_GK1 10.1.1.1 1719
  element NYC 10.1.3.1 1719
!
zone prefix SJC_GK2 40852.....
zone prefix NYC_GK2 21251.....
zone prefix CHC 72067.....
gw-type-prefix 1#* default-technology
load-balance cpu 80 memory 80
bandwidth interzone CHC_Voice 160
bandwidth interzone SJC_Voice2 192
bandwidth interzone NYC_Voice3 160
arq reject-unknown-prefix
no shutdown

```

The following notes also apply to [Example 8-3](#) and [Example 8-4](#):

- Each Cisco Unified CallManager cluster has a local zone configured to support Cisco Unified CallManager trunk registrations.
- A cluster is defined for each local zone, with backup zones on the other gatekeepers listed as elements. Elements are listed in the order in which the backups are used.
- A zone prefix is configured for each zone to allow inter-zone and inter-cluster call routing.

- The **gw-type-prefix** command allows all locally unresolved calls to be forwarded to a device registered with a technology prefix of 1# in the local zone. In this example, all Cisco Unified CallManager trunks have been configured to register with a 1# prefix.
- The **load-balance cpu 80 memory 80** command limits CPU and memory usage. If the router hits either limit, all new requests are denied and the first backup in the list is used until utilization drops below the threshold.
- Bandwidth statements are configured for each site. Cisco recommends that you use the **bandwidth interzone** command because the **bandwidth total** command does not work in some configurations.
- The **arq reject-unknown-prefix** command guards against potential call routing loops across redundant Cisco Unified CallManager trunks.

All gatekeepers in the cluster display all Cisco Unified CallManager trunk registrations. For trunks that use the gatekeeper as a primary resource, the flag field is blank. For trunks that use another gatekeeper in the cluster as their primary gatekeeper, the flag field is set to A (alternate). Having all endpoints registered as primary or alternate allows all calls to be resolved locally without having to send a location request (LRQ) to another gatekeeper.

[Example 8-5](#) shows the output from the **show gatekeeper endpoints** command on Gatekeeper 1 (SJC).

Example 8-5 Output for Gatekeeper Endpoints

```

                                GATEKEEPER ENDPOINT REGISTRATION
                                =====
CallSignalAddr  Port  RASignalAddr  Port  Zone Name      Type      Flags
-----
10.1.1.12      1307  10.1.1.12     1254  SJC             VOIP-GW
H323-ID: SJC-to-GK-trunk_1
10.1.1.12      4422  10.1.1.12     4330  SJC             VOIP-GW
H323-ID: SJC-to-GK-trunk_2
10.1.2.12      4587  10.1.2.12     4330  CHC_GK1         VOIP-GW    A
H323-ID: CHC-to-GK-trunk_1
10.1.3.21      2249  10.1.3.21     1245  NYC_GK1         VOIP-GW    A
H323-ID: NYC-to-GK-trunk_1
Total number of active registrations = 4

```

Directory Gatekeeper Redundancy


You can implement directory gatekeeper redundancy by using HSRP or by configuring multiple identical directory gatekeepers. When a gatekeeper is configured with multiple remote zones using the same zone prefix, the gatekeeper can use either of the following methods:

- Sequential LRQs (default)
Redundant remote zones (matching zone prefixes) are assigned a cost, and LRQs are sent to the matching zones in order based on the cost values. Using sequential LRQs saves WAN bandwidth by not blasting LRQs to all matching gatekeepers.
- LRQ Blast
LRQs are sent to redundant zones (matching zone prefixes) simultaneously. The first gatekeeper to respond with an Location Confirm (LCF) is the one that is used.

Cisco recommends that you use multiple active directory gatekeepers with sequential LRQs, thus allowing directory gatekeepers to be placed in different locations. Using HSRP requires both directory gatekeepers to be located in the same subnet, and only one gatekeeper can be active at any time.

Figure 8-8 illustrates a Cisco Unified CallManager distributed call processing environment with two active directory gatekeepers.

Figure 8-8 Redundant Directory Gatekeepers



Example 8-6 and Example 8-7 show the configurations for the two directory gatekeepers in Figure 8-8.

Example 8-6 Configuration for West Directory Gatekeeper

```
gatekeeper
zone local DGKW customer.com 10.1.10.1
zone remote SJC customer.com 10.1.1.1
zone remote CHC customer.com 10.1.2.1
zone remote NYC customer.com 10.1.3.1
zone prefix SJC 408.....
```

```

zone prefix CHC 720.....
zone prefix NYC 212.....
lrq forward-queries
no shutdown

```

Example 8-7 Configuration for East Directory Gatekeeper

```

gatekeeper
zone local DGKE customer.com 10.1.12.1
zone remote SJC customer.com 10.1.1.1
zone remote CHC customer.com 10.1.2.1
zone remote NYC customer.com 10.1.3.1
zone prefix SJC 408.....
zone prefix CHC 720.....
zone prefix NYC 212.....
lrq forward-queries
no shutdown

```

The following notes also apply to [Example 8-6](#) and [Example 8-7](#):

- Both directory gatekeepers are configured exactly the same.
- A local zone is configured for the directory gatekeeper.
- Remote zones are configured for each remote gatekeeper.
- Zone prefixes are configured for both remote zones for inter-zone call routing. The wildcard (*) could be used in the zone prefix to simplify configuration, but the use of dots (.) is more specific. Calls are not routed to the DGK zone, so a prefix is not required for it.
- The **lrq forward-queries** command allows the directory gatekeeper to forward an LRQ received from another gatekeeper.



Note

Directory gatekeepers do not contain any active endpoint registrations and do not supply any bandwidth management.

[Example 8-8](#), [Example 8-9](#), and [Example 8-10](#) show the configurations for Gatekeepers 1 to 3 in [Figure 8-8](#).

Example 8-8 Configuration for Gatekeeper 1 (SJC)

```

zone local SJC customer.com 10.1.1.1
zone remote DGKW customer.com 10.1.10.1
zone remote DGKE customer.com 10.1.12.1
zone prefix SJC 408.....
zone prefix DGKW .....
zone prefix DGKE .....
bandwidth remote 192
gw-type-prefix 1# default-technology
arq reject-unknown-prefix
no shutdown

```

Example 8-9 Configuration for Gatekeeper 2 (CHC)

```

gatekeeper
zone local GK-CHC customer.com 10.1.2.1
zone remote DGKE customer.com 10.1.12.1
zone remote DGKW customer.com 10.1.10.1

```

```

zone prefix CHC 720.....
zone prefix DGKE .....
zone prefix DGKW .....
bandwidth remote 160
gw-type-prefix 1# default-technology
arq reject-unknown-prefix
no shutdown

```

Example 8-10 Configuration for Gatekeeper 3 (NYC)

```

gatekeeper
zone local NYC customer.com 10.1.3.1
zone remote DGKE customer.com 10.1.12.1
zone remote DGKW customer.com 10.1.10.1
zone prefix NYC 212.....
zone prefix DGKE .....
zone prefix DGKW .....
bandwidth remote 160
gw-type-prefix 1# default-technology
arq reject-unknown-prefix
no shutdown

```

The following notes also apply to [Example 8-8](#), [Example 8-9](#), and [Example 8-10](#):

- Each Cisco Unified CallManager cluster has a local zone configured to support Cisco Unified CallManager trunk registrations.
- Remote zones are configured for each directory gatekeeper.
- Zone prefixes are configured for the local zone and both remote zones for inter-zone call routing. Both directory gatekeeper prefixes are 10 dots. Sequential LRQs are used by default when matching zone prefixes are configured. The gatekeeper sends an LRQ to the directory gatekeeper with the lowest cost; if there is no response, the gatekeeper tries the second directory gatekeeper.
- The **bandwidth remote** command is used to limit bandwidth between the local zone and any other remote zone.
- The **gw-type-prefix** command allows all locally unresolved calls to be forwarded to a device registered with a technology prefix of 1# in the local zone. In this example, all Cisco Unified CallManager trunks have been configured to register with a 1# prefix.
- The **arq reject-unknown-prefix** command guards against potential call routing loops across redundant Cisco Unified CallManager trunks.

Interoperability of Cisco Unified CallManager and CallManager Express

This section explains the requirements for interoperability and internetworking of Cisco Unified CallManager with Cisco Unified CallManager Express (CME, formerly called Cisco IOS Telephony Services, or ITS) using H.323 protocol in a multisite IP telephony deployment. This section highlights the recommended deployments between phones controlled by Cisco Unified CallManager and phones controlled by CME.

Cisco CME 3.4 supports Cisco IP SCCP Phones 7902, 7905, 7910, 7912, 7920, 7935, 7936, 7940, 7960, 7970, 7971, and Cisco IP Communicator. Cisco Unified CallManager 4.2 supports CME with SIP phones.

The following section provides guidelines for achieving interoperability between Cisco Unified CallManager and Cisco Unified CallManager Express:

- [Multisite IP Telephony Deployments Using Cisco Unified CallManager and CME, with H.323 Trunks and an IP-to-IP Gateway, page 8-28](#)

For more information about CME, refer to the product documentation for Cisco Unified CallManager Express, available online at

<http://www.cisco.com>

Multisite IP Telephony Deployments Using Cisco Unified CallManager and CME, with H.323 Trunks and an IP-to-IP Gateway


An IP-to-IP gateway is a separate router that provides a proxy (or front end) for a system that does not support H.450, such as Cisco Unified CallManager. The IP-to-IP gateway can be placed between Cisco Unified CallManager and a CME router, and it provides H.323-to-H.323 call connections to terminate and re-originate calls for transfers and forwards to endpoints that do not support H.450.

Cisco Unified CallManager Express can be directly integrated with Cisco Unified CallManager using an H.323 interface without the IP-to-IP gateway. However, because Cisco Unified CallManager does not support H.450 specifications, supplementary services such call transfer and call forward initiated from Cisco Unified CallManager phones will require the media to be hair-pinned (over the WAN in some cases) through Cisco Unified CallManager even though the Cisco Unified CallManager phones might not be involved in the call. The IP-to-IP gateway ensure that this media hair-pinning does not occur over the WAN.

The IP-to-IP gateway can also act as the PSTN gateway for the remote Cisco CME systems as well as for Cisco Unified CallManager. A separate PSTN gateway is not required in such cases.

The IP-to-IP gateway must run a Cisco IOS release that is compliant with Cisco CME 3.4 and that supports H.450, such as Cisco IOS Release 12.4(6)T or later with the IP VOICE feature set.

[Figure 8-9](#) shows an IP Telephony deployment with Cisco Unified CallManager connected to Cisco CME through an IP-to-IP gateway.

Figure 8-9 Multisite IP Telephony Deployment with Cisco Unified CallManager, CME, and an IP-to-IP Gateway

Best Practices

Follow these guidelines and best practices when using the deployment model illustrated in [Figure 8-9](#):

- Configure a gatekeeper controlled H.225 trunk between Cisco Unified CallManager and the IP-to-IP gateway, with **Media Termination Point Required** checked and **Wait For Far End H.245 Terminal Capability Set** unchecked.
- In Cisco Unified CallManager, set the service parameter **Send H225 user info message to H225 info for Call Progress Tone**.
- Use the Cisco Unified CallManager dial plan configuration (route patterns, route lists, and route groups) to send calls to the H.225 trunk that connects to the IP-to-IP gateway.
- Register Cisco CME and the IP-to-IP gateway as H.323 gateways on the gatekeeper.
- A media termination point (MTP) is required for H.450 to perform Empty Capability Set (ECS) signaling conversion. The MTP should be configured on the IP-to-IP gateway and registered to Cisco Unified CallManager. An MTP can cause media to be hair-pinned over the WAN in certain rare cases. For more details on media termination points, see the chapter on [Media Resources](#), [page 6-1](#).
- Configure the **allow-connection h323 to h323** command on the IP-to-IP gateway to allow H.323-to-H.323 call connections. Remote CME routers need not have this command enabled.

**Note**

When using H.323 trunks between CME and Cisco Unified CallManager, configure the **allow-connection h323 to h323** command on CME to allow H.323-to-H.323 hair-pinned call connections. The Cisco Unified CallManager Express auto-detection feature (starting with CME 3.1) disables all H.450-based communications on the interface connected to Cisco Unified CallManager. To complete the call transfer or forward between Cisco Unified CallManager and CME phones, an H.323-to-H.323 connection is required.

- Define VoIP dial peers on the IP-to-IP gateway to route calls to the Cisco Unified CallManager and CME endpoints.
- Define VoIP dial peers on CME to forward calls, destined for Cisco Unified CallManager endpoints, to the IP-to-IP gateway.
- Supplementary services such as transfer and call forward will result in calls being media hair-pinned when the two endpoints reside in the same CME branch location.

**Note**

When multiple PSTN connections exist (one for Cisco Unified CallManager and one for CME), fully attended transfer between a Cisco Unified CallManager endpoint and a CME endpoint to a PSTN endpoint will fail. The recommendation is to use blind transfer when using multiple PSTN connections, and it is configured under **telephony-service** as **transfer-system full-blind**.

[Example 8-11](#) lists a sample configuration for the IP-to-IP gateway.

Example 8-11 Configuration for IP-to-IP Gateway

```
voice service voip
  allow-connections h323 to h323
  supplementary-service h450.2
  supplementary-service h450.3
  supplementary-service h450.12
  h323
  emptycapability
  h225 id-passthru
  h225 connect-passthru
  h245 passthru tcsnonstd-passthru
dial-peer voice 1 voip          /* To Cisco Unified CallManager endpoints */
  destination-pattern xxxx
  session target ipv4:y.y.y.y
  dtmf-relay h245-alphanumeric
  codec g729r8
  no vad
dial-peer voice 1 voip          /* To Cisco Unified CallManager endpoints */
  destination-pattern zzzz
  session target ras            /* "ras" if gatekeeper is used, otherwise "ipv4:a.b.c.d" */
  dtmf-relay h245-alphanumeric
  codec g729r8
  no vad
```

Example 8-12 lists a sample configuration for the in this H.323 deployment model.

Example 8-12 Configuration for Cisco CME 3.4 with H.323

```
voice service voip
  h323
interface FastEthernet0/1
  ip address 10.10.10.23 255.255.255.0
  h323-gateway voip interface
  h323-gateway voip id cme ipaddr 10.10.10.30 1719
dial-peer voice 1 voip          /* To Cisco Unified CallManager endpoints */
  destination-pattern xxxx
  session target ras
  session transport tcp
  codec g729r8                 /* Voice class can also be used */
  no vad
telephony-service
  ip source-address 10.10.10.22 port 2000
  create cnf-files
  keepalive 45
  max-conferences 8 gain -6
  moh music-on-hold.au
  transfer-system full-blind /* Used with multiple PSTN connections */
  transfer-pattern .T
```





CHAPTER 9

Call Admission Control

Last revised on: February 13, 2008

The call admission control function is an essential component of any IP telephony system that involves multiple sites connected through an IP WAN. In order to better understand what call admission control does and why it is needed, consider the example in [Figure 9-1](#).

Figure 9-1 Why Call Admission Control is Needed



As shown on the left side of [Figure 9-1](#), traditional TDM-based PBXs operate within circuit-switched networks, where a circuit is established each time a call is set up. As a consequence, when a legacy PBX is connected to the PSTN or to another PBX, a certain number of physical trunks must be provisioned. When calls have to be set up to the PSTN or to another PBX, the PBX selects a trunk from those that are available. If no trunks are available, the call is rejected by the PBX and the caller hears a network-busy signal.

Now consider the IP telephony system shown on the right side of [Figure 9-1](#). Because it is based on a packet-switched network (the IP network), no circuits are established to set up an IP telephony call. Instead, the IP packets containing the voice samples are simply routed across the IP network together with other types of data packets. Quality of Service (QoS) is used to differentiate the voice packets from

the data packets, but bandwidth resources, especially on IP WAN links, are not infinite. Therefore, network administrators dedicate a certain amount of "priority" bandwidth to voice traffic on each IP WAN link. However, once the provisioned bandwidth has been fully utilized, the IP telephony system must reject subsequent calls to avoid oversubscription of the priority queue on the IP WAN link, which would cause quality degradation for all voice calls. This function is known as call admission control, and it is essential to guarantee good voice quality in a multisite deployment involving an IP WAN.

To preserve a satisfactory end-user experience, the call admission control function should always be performed during the call setup phase so that, if there are no network resources available, a message can be presented to the end-user or the call can be rerouted across a different network (such as the PSTN).

This chapter discusses the following main topics:

- [Best Practices Summary, page 9-2](#)

This section summarizes the key best practices, recommendations, and notes about call admission control for the readers who are already familiar with the principles and mechanisms described in the remainder of this chapter.

- [Call Admission Control Principles, page 9-3](#)

This section defines the two fundamental approaches to call admission control in an IP-based telephony system: topology-aware and topology-unaware call admission control.

- [Call Admission Control Elements, page 9-11](#)

This section describes the call admission control mechanisms available through the various components of a Cisco IP Communications system, such as Cisco Unified CallManager locations, Cisco IOS gatekeeper, RSVP, and the IP-to-IP gateway.

- [Call Admission Control Design, page 9-25](#)

This section shows how to apply and combine the mechanisms described in the previous sections, based on the IP WAN topology (simple hub-and-spoke, two-tier hub-and-spoke, MPLS, or other topologies) and also based on the Cisco Unified CallManager deployment model adopted.

Best Practices Summary

This section briefly summarizes the best practices for providing call admission control in various Cisco Unified CallManager deployments. The remainder of this chapter explains these best practices in more detail.

The following recommendations apply to deployments with a single Cisco Unified CallManager cluster:

- For simple hub-and-spoke topologies, use Cisco Unified CallManager static locations. Leave the hub site devices in the <None> location.
- For Multiprotocol Label Switching (MPLS) topologies, use Cisco Unified CallManager static locations, with devices at every site (including the central site) assigned to a location.

The following recommendations apply to deployments with multiple Cisco Unified CallManager clusters:

- For simple hub-and-spoke topologies, use Cisco IOS gatekeeper zones between sites where Cisco Unified CallManager clusters reside.
- For two-tier hub-and-spoke topologies where Cisco Unified CallManager clusters are located at the first and second level hub sites, use Cisco IOS gatekeeper zones for the links between first- and second-level hub sites and use Cisco Unified CallManager static locations for the links between second-level hub sites and spoke sites.

- For MPLS topologies with no dual links, use Cisco Unified CallManager static locations, with every site in a location and with no gatekeeper zones. Leave intercluster trunks in the <None> location unless an MTP is required. You may use a gatekeeper for intercluster call routing, but it is not needed for call admission control.
- For generic topologies where Cisco Unified CallManager clusters are located at each site, use RSVP-enabled IP-to-IP gateways across clusters.

Call Admission Control Principles

As mentioned previously, call admission control is a function of the call processing agent in an IP-based telephony system, so in theory there could be as many call admission control mechanisms as there are IP-based telephony systems. However, most of the existing call admission control mechanisms fall into one of the following two main categories:

- Topology-unaware call admission control — Based on a static configuration within the call processing agent
- Topology-aware call admission control — Based on communication between the call processing agent and the network about the available resources

The remainder of this section first analyzes the principles of topology-unaware call admission control and its limitations, then it presents the principles of topology-aware call admission control.


Topology-Unaware Call Admission Control

We define as topology-unaware call admission control any mechanism that is based on a static configuration within a call processing agent or IP-based PBX, aimed at limiting the number of simultaneous calls to or from a remote site connected via the IP WAN.


As shown in [Figure 9-2](#), most of these mechanisms rely on the definition of a logical "site" entity, which generally corresponds to a geographical branch office connected to the enterprise IP WAN.

After assigning all the devices located at each branch office to the corresponding site entity, the administrator usually configures a maximum number of calls (or a maximum amount of bandwidth) to be allowed in or out of that site.

Each time a new call needs to be established, the call processing agent checks the sites to which the originating and terminating endpoints belong, and verifies whether there are available resources to place the call (in terms of number of calls or amount of bandwidth for both sites involved). If the check succeeds, the call is established and the counters for both sites are decremented. If the check fails, the call processing agent can decide how to handle the call based on a pre-configured policy. For example, it could send a network-busy signal to the caller device, or it could attempt to reroute the call over a PSTN connection.

Figure 9-2 Principles of Topology-Unaware Call Admission Control

Because of their reliance on static configurations, topology-unaware call admission control mechanisms can generally be deployed only in networks with a relatively simple IP WAN topology. In fact, most of these mechanisms mandate a simple hub-and-spoke topology or a simple MPLS-based topology (where the MPLS service is provided by a service provider), as shown in Figure 9-3.

Figure 9-3 Domain of Applicability of Topology-Unaware Call Admission Control

In a hub-and-spoke network or MPLS-based network such as those shown in Figure 9-3, each spoke site is assigned to a "site" within the call processing agent, and the number of calls or amount of bandwidth for that "site" is configured to match the bandwidth available for voice (and/or video) on the IP WAN link that connects the spoke to the IP WAN.


Notice the absence of redundant links from the spoke sites to the hub site and of links directly connecting two spoke sites. The next section explains why such links create problems for topology-unaware call admission control.

Limitations of Topology-Unaware Call Admission Control

In today's enterprise networks, high availability is a common requirement, and it often translates into a desire to provide redundancy for the IP WAN network connectivity.

When considering the IP WAN topology in a typical enterprise network, you are likely to encounter a number of characteristics that complicate the assumption of a pure hub-and-spoke topology. Figure 9-4 shows several of these network characteristics in a single diagram. Obviously, only the largest enterprise networks present all these characteristics at once, but it is highly likely that most IP WAN networks feature at least one of them.

Figure 9-4 Topology Characteristics of Typical Enterprise Networks




As explained in the section on [Call Admission Control Design, page 9-25](#), it is sometimes possible to adapt a topology-unaware call admission control mechanism to a complex network topology, but there are limitations in terms of when this approach can be used and what behavior can be achieved. For example, consider the simple case of a branch site connected to a hub site via the IP WAN, where redundancy is a network requirement. Typically, redundancy can be achieved in one of the following ways:

- A single router with a primary and a backup link to the IP WAN
- A single router with two active WAN links in a load-balancing configuration
- Two router platforms, each connected to the IP WAN, with load-balanced routing across them

The examples [Figure 9-5](#) attempt to apply a topology-unaware call admission control mechanism to the case of a single router with a primary and backup link and the case of a single router with two active load-balanced links. (The case of two router platforms has the same call admission control implications as the latter example.)

Figure 9-5 *Topology-Unaware Call Admission Control in Presence of Dual Links*



For the first example in [Figure 9-5](#), branch office A is normally connected to the IP WAN via a primary link, whose Low Latency Queuing (LLQ) bandwidth is provisioned to allow a maximum of 10 simultaneous calls. When this primary link fails, a smaller backup link becomes active and preserves the connectivity to the IP WAN. However, the LLQ bandwidth of this backup link is provisioned to allow only up to 2 simultaneous calls.

In order to deploy a topology-unaware call admission control mechanism for this branch office, we must define a "site" A in the call processing agent and configure it for a certain number of calls (or amount of bandwidth). If we choose to use 10 calls as the maximum for site A, the backup link can be overrun during failures of the primary link, thereby causing bad voice quality for all active calls. If, on the other hand, we choose 2 calls as the maximum, we will not be able to use the bandwidth provisioned for the remaining 8 calls when the primary link is active.

Now consider branch office B, which has two active links connecting it to the IP WAN. Each of these links is provisioned to allow a maximum of 10 simultaneous calls, and the routing protocol automatically performs load-balancing between them. When deploying a topology-unaware call admission control mechanism for this branch office, we must define a "site" B in the call processing agent and configure it for a certain number of calls (or amount of bandwidth). Similar to the case of branch office A, if we choose to add up the capacity of the two links and use 20 calls as the maximum for site B, there is a potential to overrun the LLQ on one of the two links during failures of the other one. For example, if link #2 fails, the system still allows 20 simultaneous calls to and from site B, which are now all routed via link #1, thus overrunning it and causing poor voice quality for all calls. On the other hand, if site B is configured for a maximum of 10 simultaneous calls, the available LLQ bandwidth is never fully utilized under normal conditions (when both links are operational).

These two simple examples show how IP WAN bandwidth provisioning in real enterprise networks is often too complex to be summarized in statically configured entries within the call processing agent. Deploying topology-unaware call admission control in such networks forces the administrator to make assumptions, develop workarounds, or accept sub-optimal use of network resources.

The optimal way to provide call admission control in the presence of a network topology that does not conform to a simple hub-and-spoke is to implement topology-aware call admission control, as described in the following section.

**Note**

Some IP telephony systems augment classic topology-unaware call admission control with a feedback mechanism based on observed congestion in the network, which forces calls through the PSTN when voice quality deteriorates. This approach is still not equivalent to true topology-aware call admission control because it is performed after the calls have already been established and because the call processing agent still does not have knowledge of exactly where congestion is occurring. As mentioned at the beginning of the chapter, in order to be effective, call admission control must be performed before the call is set up.

Topology-Aware Call Admission Control

We define as topology-aware call admission control any mechanism aimed at limiting the number of simultaneous calls across IP WAN links that can be applied to any network topology and can dynamically adjust to topology changes.

To accomplish these goals, topology-aware call admission control must rely on real-time communications about the availability of network resources between a call processing agent (or IP-based PBX) and the network. Because the network is a distributed entity, real-time communications require a signaling protocol.


The Resource Reservation Protocol (RSVP) is the first significant industry-standard signaling protocol that enables an application to reserve bandwidth dynamically across an IP network. Using RSVP, applications can request a certain amount of bandwidth for a data flow across a network (for example, a voice call) and can receive an indication of the outcome of the reservation based on actual resource availability.

In the specific case of call admission control for voice or video calls, an IP-based PBX can synchronize the call setup process with RSVP reservations between the two remote sites and can make a routing decision based on the outcome of the reservations. Because of its distributed and dynamic nature, RSVP is capable of reserving bandwidth across any network topology, thus providing a real topology-aware call admission control mechanism.

To better understand the basic principles of how RSVP performs bandwidth reservation in a network, consider the simple example depicted in [Figure 9-6](#). This example does not analyze the exact message exchanges and protocol behaviors, but rather focus on the end results from a functionality perspective. For more information on the RSVP message exchanges, see [RSVP Principles, page 3-35](#).


Assume that RSVP is enabled on each router interface in the network shown in [Figure 9-6](#) and that the numbers shown in the circles represent the amount of available RSVP bandwidth remaining on each interface.

Figure 9-6 Sample Network to Show RSVP Principles



Now consider an RSVP-enabled application that wants to reserve a certain amount of bandwidth for a data stream between two devices. This scenario is depicted in Figure 9-7, which shows a particular data stream that requires 24 kbps of bandwidth from Device 1 to Device 2.

Figure 9-7 RSVP Signaling for a Successful Reservation




The following considerations apply to [Figure 9-7](#):

- RSVP does not perform its own routing; instead it uses underlying routing protocols to determine where it should carry reservation requests. As routing changes paths to adapt to topology changes, RSVP adapts its reservations to the new paths wherever reservations are in place.
- The RSVP protocol attempts to establish an end-to-end reservation by checking for available bandwidth resources on all RSVP-enabled routers along the path from Device 1 to Device 2. As the RSVP messages progress through the network, the available RSVP bandwidth gets decremented by 24 kbps on the outbound router interfaces, as shown in [Figure 9-7](#).
- The available bandwidth on all outbound interfaces is sufficient to accept the new data stream, so the reservation succeeds and the application is notified.
- RSVP reservations are unidirectional (in this case, the reservation is established from Device 1 to Device 2, and not vice versa). In the presence of bidirectional applications such as voice and videoconferencing, two reservations must be established, one in each direction.
- RSVP provides transparent operation through router nodes that do not support RSVP. If there are any routers along the path that are not RSVP-enabled, they simply ignore the RSVP messages and pass them along like any other IP packet, and a reservation can still be established. (See [RSVP Principles, page 3-35](#), for details on protocol messages and behaviors.) However, in order to have an end-to-end QoS guarantee, you have to ensure that there is no possibility of bandwidth congestion on the links controlled by the non-RSVP routers.

After a reservation has been successfully established between Device 1 and Device 2, now assume that another application requests a 24-kbps reservation between Device 3 and Device 4, as depicted in [Figure 9-8](#).

Figure 9-8 *RSVP Signaling for an Unsuccessful Reservation*



The following considerations apply to [Figure 9-8](#):


- The RSVP protocol attempts to establish an end-to-end reservation by checking for available bandwidth resources on all RSVP-enabled routers along the path from Device 3 to Device 4. As the RSVP messages progress through the network, the available RSVP bandwidth gets decremented by 24 kbps on the outbound router interfaces, as shown in [Figure 9-8](#).
- In this example, the available bandwidth on R5's outbound interface toward R6 is not sufficient to accept the new data stream, so the reservation fails and the application is notified. The available RSVP bandwidth on each outbound interface along the path is then restored to its previous value.
- The application can then decide what to do. It could abandon the data transfer or decide to send it anyway with no QoS guarantees, as best-effort traffic.

We can now apply the topology-aware call admission control approach based on RSVP to the examples of dual-connected branch offices A and B introduced in the previous section.

As shown in [Figure 9-9](#), branch office A has a primary link with an LLQ provisioned for 10 calls, while the backup link can accommodate only 2 calls. With this approach, RSVP is configured on both router interfaces so that the RSVP bandwidth matches the LLQ bandwidth. Branch A is also configured within the call processing agent to require RSVP reservations for all calls to or from other branches. Now calls are admitted or rejected based on the outcome of the RSVP reservations, which automatically follow the path determined by the routing protocol. Under normal conditions (when the primary link is active), up to 10 calls will be admitted; during failure of the primary link, only up to 2 calls will be admitted.

Policies can typically be set within the call processing agent to determine what to do in the case of a call admission control failure. For example, the call could be rejected, rerouted across the PSTN, or sent across the IP WAN as a best-effort call with a different DSCP marking.

Figure 9-9 Topology-Aware Call Admission Control for Dual Links



Similar considerations apply to branch B, connected to the IP WAN via two load-balanced links, as shown on the right side of [Figure 9-9](#). RSVP is enabled on each of the two router interfaces, with a bandwidth value that matches the LLQ configuration (in this case, enough bandwidth for 10 calls). Branch B is also configured within the call processing agent to request RSVP reservations for calls to or from other branches. Again, calls are admitted or rejected based on the actual bandwidth available along

the path chosen by the routing protocol. So in a case of perfectly even load-balancing across the two links, up to 20 calls could be admitted under normal conditions (when both links are operational); if one of the two links fails, only up to 10 calls would be admitted.

In the case that one of the two links failed while more than 10 calls were active, some calls would fail to re-establish a reservation on the new path. At this point, the call processing agent would be notified and could react based on the configured policy (for example, by dropping the extra calls or by remarking them as best-effort calls).

In conclusion, topology-aware call admission control allows administrators to protect call quality with any network topology, to automatically adjust to topology changes, and to make optimal use of the network resources under all circumstances.

Special Considerations for MPLS Networks

From the call admission control perspective, a network based on MPLS differs from one based on traditional Layer 2 WAN Services with respect to support for RSVP in the "hub" of the network. Hub sites of traditional Layer 2 wide-area networks consist, in most cases, of an enterprise-controlled router that can be enabled to participate in RSVP. Because the entire network (cloud) is the "hub site" in MPLS networks, there is no enterprise-controlled hub location to enable RSVP. (For more information, see [Simple MPLS Topologies, page 9-34](#).) Therefore, to provide topology-aware call admission control in an MPLS environment, the Customer Edge (CE) devices of the network must be configured for RSVP support.

Because RSVP must be enabled on the CE, control of this equipment is important. If this equipment is not under the control of the enterprise, you must work with your service provider to determine if they will enable RSVP on your WAN interface and if that implementation will support advanced features such as RSVP application ID.

RSVP messages will transparently pass across the RSVP-unaware MPLS cloud, so this does not pose a problem with end-to-end RSVP capability. Configuring RSVP on the CE WAN interface will ensure that its priority queue will not be overrun. Because RSVP reservations are unidirectional, the following rules must be observed to protect the priority queue on the Provider Edge (PE) router when RSVP is not enabled in the MPLS cloud:

- The media streams must be the same size in both directions.
- The media has to be symmetrically routed.

If your MPLS network does not comply with these rules, contact your local Cisco account team for further assistance before implementing RSVP.

Call Admission Control Elements

There are several mechanisms that perform the call admission control function in a Cisco IP Communications system. This section provides design and configuration guidelines for all of these mechanisms, according to their category:

- Topology-unaware mechanisms
 - [Cisco Unified CallManager Static Locations, page 9-12](#)
 - [Cisco IOS Gatekeeper Zones, page 9-15](#)
- Topology-aware mechanisms
 - [Cisco IOS Gatekeeper and IP-to-IP Gateway with RSVP, page 9-17](#)

**Note**

Cisco Unified CallManager 5.0 introduces topology-aware call admission control by extending the concept of *locations*, which already existed in previous releases. Therefore, to avoid confusion, this document refers to the existing topology-unaware mechanism as *static locations* and to the new topology-aware mechanism as *RSVP-enabled locations*. Cisco Unified CallManager 4.2 provides support for static locations only.

Cisco Unified CallManager Static Locations

Cisco Unified CallManager provides a simple mechanism known as static locations for implementing call admission control across the IP WAN in centralized call processing deployments. When you configure a device in Cisco Unified CallManager, the device can be assigned to a location. A bandwidth pool is allocated to each location for calls to devices belonging to any other location. (Calls within the same location do not need call admission control.) Each time a call is set up between any two locations, the required amount of bandwidth is subtracted from the corresponding location bandwidth pools, until there is no more bandwidth available in the pools. At that point, call admission control rejects subsequent calls to or from the given location.

By default, Cisco Unified CallManager assumes that each device's physical location matches the configured location. So if a device is moved from one physical location to another, the system administrator must perform a manual update on its location configuration. An alternative is to enable the Device Mobility feature, introduced in Cisco Unified CallManager Release 4.2. This feature enables Cisco Unified CallManager to assign a device dynamically to a device pool (and hence to a location) based on its IP address, and it is described in detail in the chapter on [Device Mobility, page 22-1](#).

**Note**

The Device Mobility feature is not available in Cisco Unified CallManager Release 5.0.

Each device is in location <None> by default. Location <None> is a special unnamed location that has unlimited audio and video bandwidth. If the devices at a given site are configured in the <None> location, all calls to and from that devices at that site will be admitted by the call admission control mechanism.

Cisco Unified CallManager allows you to define a voice and video bandwidth pool for each location. If the location's audio and video bandwidth are configured as **Unlimited**, there will be unlimited bandwidth available for that location, and every audio or video call to or from that location will be permitted by Cisco Unified CallManager. On the other hand, if the bandwidth values are set to a finite number of kilobits per second (kbps), Cisco Unified CallManager will allow calls in and out of that location as long as the aggregate bandwidth used by all active calls is less than or equal to the configured values. If the video bandwidth for the location is configured as **None**, every video call to or from that location is denied but the video calls within the same location are not affected.

For video calls, the video location bandwidth takes into account both the video and the audio portions of the call. Therefore, for a video call, no bandwidth is deducted from the audio bandwidth pool.

The devices that can specify membership in a location include:

- IP phones
- CTI ports
- H.323 clients
- CTI route points
- Conference bridges

- Music on hold (MoH) servers
- Gateways
- Trunks

The static locations call admission control mechanism also takes into account the mid-call changes of call type. For example, if an inter-site video call is established, Cisco Unified CallManager will subtract the appropriate amount of video bandwidth from the respective locations. If this video call changes to an audio-only call via a transfer to a device that is not capable of video, Cisco Unified CallManager will return the allocated bandwidth to the video pool and allocate the appropriate amount of bandwidth from the audio pool. Calls that change from audio to video will cause the opposite change of bandwidth allocation.

Table 9-1 lists the amount of bandwidth requested by the static locations algorithm for various call speeds. For an audio call, Cisco Unified CallManager counts the media bit rates plus the Layer 3 overhead. For example, a G.711 audio call consumes 80 kbps allocated from the location's audio bandwidth pool. For a video call, Cisco Unified CallManager counts only the media bit rates for both the audio and video streams. For example, for a video call at a speed of 384 kbps, Cisco Unified CallManager will allocate 384 kbps from the video bandwidth pool.

Table 9-1 Amount of Bandwidth Requested by the Static Locations Algorithm

Call Speed	Static Location Bandwidth Value
G.711 audio call (64 kbps)	80 kbps
G.729 audio call (8 kbps)	24 kbps
128-kbps video call	128 kbps
384-kbps video call	384 kbps
512-kbps video call	512 kbps
768-kbps video call	768 kbps

Figure 9-10 shows the configuration for the location Branch 1, with 256 kbps of available audio bandwidth and 384 kbps of available video bandwidth. The Branch 1 location can support up to three G.711 audio calls (at 80 kbps per call) or ten G.729 audio calls (at 24 kbps per call), or any combination of both that does not exceed 256 kbps. The location can also support different numbers of video calls depending on the video and audio codecs being used. For example, one video call requesting 384 kbps of bandwidth or three video calls with each requesting 128 kbps of bandwidth.

Figure 9-10 Defining a Location in Cisco Unified CallManager

The screenshot shows the Cisco Unified CallManager Administration web interface. At the top is a navigation bar with links: System, Route Plan, Service, Feature, Device, User, Application, and Help. Below this is the header 'Cisco Unified CallManager Administration' with the tagline 'For Cisco Unified Communications' and the Cisco Systems logo. The main content area is titled 'Location Configuration'. On the right side of this area are three links: 'Add a New Location', 'Back to Find/List Locations', and 'Dependency Records'. The configuration is for 'Location: Branch 1', with a status of 'Insert completed'. There are four buttons: 'Copy', 'Update', 'Delete', and 'Resync Bandwidth'. Below these are three sections: 'Location Information' with a text field for 'Location Name*' containing 'Branch 1'; 'Audio Calls Information' with a radio button for 'Unlimited' and a text field for '256 kbps', accompanied by a note: 'If the audio quality is poor or choppy, lower the bandwidth setting. For ISDN use multiples of 56 kbps or 64 kbps.'; and 'Video Calls Information' with radio buttons for 'None' and 'Unlimited', and a text field for '384 kbps'. A footnote at the bottom left states '* indicates required item'. A vertical text '190317' is on the right edge of the screenshot.

**Note**

Call admission control does not apply to calls between devices within the same location.

When a call is placed from one location to the other, Cisco Unified CallManager deducts the appropriate amount of bandwidth from both locations. For example, a G.729 call between two locations causes Cisco Unified CallManager to deduct 24 kbps from the available bandwidth at both locations. When the call has completed, Cisco Unified CallManager returns the bandwidth to the affected locations. If there is not enough bandwidth at either branch location, the call is denied by Cisco Unified CallManager and the caller receives the network busy tone. If the calling device is an IP phone with a display, that device also displays the message "Not Enough Bandwidth."

When an inter-site call is denied by call admission control, Cisco Unified CallManager can automatically reroute the call to the destination via the PSTN connection by means of the Automated Alternate Routing (AAR) feature. For detailed information on the AAR feature, see [Automated Alternate Routing](#), page 10-22.

**Note**

AAR is invoked only when the locations-based call admission control denies the call due to a lack of network bandwidth. AAR is not invoked when the IP WAN is unavailable or other connectivity issues cause the called device to become unregistered with Cisco Unified CallManager. In such cases, the calls are redirected to the target specified in the Call Forward No Answer field of the called device.

Cisco IOS Gatekeeper Zones

A Cisco IOS gatekeeper can provide call routing and call admission control between devices such as Cisco Unified CallManager, Cisco Unified CallManager Express, or H.323 gateways connected to legacy PBXs. It uses the H.323 Registration Admission Status (RAS) protocol to communicate with these devices and route calls across the network.

Gatekeeper call admission control is a policy-based scheme requiring static configuration of available resources. The gatekeeper is not aware of the network topology, so it is limited to simple hub-and-spoke topologies. Refer to the section on [Call Admission Control Design, page 9-25](#), for detailed topology examples.

The Cisco 2600, 3600, 3700, 2800, 3800, and 7200 Series routers all support the gatekeeper feature. You can configure Cisco IOS gatekeepers in a number of different ways for redundancy, load balancing, and hierarchical call routing. This section focuses on the call admission control aspect of the gatekeeper feature. For redundancy and scalability considerations, refer to the section on [Gatekeeper Design Considerations, page 8-18](#). For call routing considerations, refer to [Call Routing in Cisco IOS with a Gatekeeper, page 10-38](#).

The call admission control capabilities of a Cisco IOS gatekeeper are based on the concept of gatekeeper *zones*. A zone is a collection of H.323 devices, such as endpoints, gateways, or Multipoint Control Units (MCUs), that register with a gatekeeper. There can be only one active gatekeeper per zone, and you can define up to 100 local zones on a single gatekeeper. A local zone is a zone that is actively handled by that gatekeeper – that is, all H.323 devices assigned to that zone register with that gatekeeper.

When multiple gatekeepers are deployed in the same network, a zone is configured as a local zone on only one gatekeeper. On the other gatekeepers, that zone is configured as a remote zone. This configuration instructs the gatekeeper to forward calls destined for that zone to the gatekeeper that "owns it" (that is, the gatekeeper on which that zone is configured as a local zone).

Use the **bandwidth** command to manage the number of calls that the gatekeeper will allow, thus providing call admission control functionality. This command has several options, but the most relevant are the following:

- The **interzone** option controls the amount of bandwidth for all calls into or out of a given local zone.
- The **total** option controls the amount of bandwidth for all calls into, out of, or within a given local zone.
- The **session** option controls the amount of bandwidth per call for a given local zone.
- The **remote** option controls the total amount of bandwidth to or from all remote zones.

The bandwidth value deducted by the gatekeeper for every active call is double the bit-rate of the call, excluding Layer 2, IP, and RTP overhead. For example, a G.711 audio call that uses 64 kbps would be denoted as 128 kbps in the gatekeeper, and a 384-kbps video call would be denoted as 768 kbps.

[Table 9-2](#) shows the bandwidth values used by the gatekeeper feature for some of the most popular call speeds.

Table 9-2 Gatekeeper Bandwidth Settings for Various Call Speeds

Call Speed	Gatekeeper Bandwidth Value
G.711 audio call (64 kbps)	128 kbps
G.729 audio call (8 kbps)	16 kbps
128-kbps video call	256 kbps
384-kbps video call	768 kbps


Table 9-2 Gatekeeper Bandwidth Settings for Various Call Speeds (continued)

Call Speed	Gatekeeper Bandwidth Value
512-kbps video call	1024 kbps
768-kbps video call	1536 kbps

**Note**

Bandwidth calculations for the call Admission Request (ARQ) do not include compressed Real-Time Transport Protocol (cRTP) or any other transport overhead. See [Bandwidth Provisioning, page 3-44](#), for details on how to provision interface queues.

To better understand the application of the **bandwidth** commands in a real network, consider the example shown in [Figure 9-11](#).

Figure 9-11 Example of Cisco IOS Gatekeeper bandwidth Commands

Assuming that all calls are voice-only calls using the G.711 codec, and given the configuration commands shown in [Figure 9-11](#), the following statements hold true:

- The maximum amount of bandwidth requested by any device in zone A for a single call is 128 kbps, which means that calls trying to use codecs with a higher bit-rate than 64 kbps will be rejected.
- The maximum amount of bandwidth used by all calls involving devices in zone A (either within the zone or with other zones) is 384 kbps, which means that there can be at most three active calls involving devices in zone A.
- The maximum amount of bandwidth used by all calls between devices in zone B and devices in any other zone is 256 kbps, which means that there can be at most two active calls between devices in zone B and devices in zones A, C, and D.
- The maximum amount of bandwidth used by all calls between devices registered with gatekeeper GK 1 and devices registered with any other gatekeeper is 512 kbps, which means that there can be at most four active calls between devices in zones A and B and devices in zones C and D.

Cisco IOS Gatekeeper and IP-to-IP Gateway with RSVP

The Cisco Multiservice IP-to-IP Gateway (also referred to as IP-IP gateway or IPIPGW) can be used to ease the restriction of hub-and-spoke topologies for the IP WAN connections between Cisco Unified CallManager clusters and/or H.323 gateways.


This Cisco IOS feature provides a mechanism to enable H.323 Voice over IP (VoIP) and videoconferencing calls from one IP network to another. The main purpose of the IP-IP gateway is to provide a control point and a demarcation for VoIP and video calls traversing administrative domains. This gateway performs most of the same functions of a PSTN-to-IP gateway, but typically joins two IP call legs rather than a PSTN and an IP call leg.

The most interesting feature of the IP-IP gateway in an enterprise IP Communications environment is that it can generate an RSVP reservation for each call that traverses it. As described in the section on [Topology-Aware Call Admission Control, page 9-7](#), RSVP is a network-based signaling protocol that provides a topology-aware call admission control mechanism that does not require a hub-and-spoke topology but works with any network topology.

As a consequence, you can perform call admission control over any IP WAN topology by inserting two IP-IP gateways in the call flow and enabling RSVP between them. [Figure 9-12](#) shows a basic example where two sites, A and B, each have a Cisco Unified CallManager cluster and are connected via an IP WAN that has an arbitrary topology. An IP-IP gateway is also located at each site, and the two Cisco Unified CallManager clusters are configured so that all inter-site calls are routed via a trunk that points to the local IP-IP gateway. When a call is set up between Site A and Site B, the following events occur:

- Cisco Unified CallManager at Site A sets up a call through an H.323 trunk to Site A's IP-IP gateway. (This is call leg 1 in the figure.)
- Site A's IP-IP gateway attempts to establish another call to Site B's IP-IP gateway, but first it uses RSVP to allocate bandwidth resources along the IP WAN path.
- If the RSVP reservation is successful, call leg 2 is established between the two IP-IP gateways.
- Site B's IP-IP gateway generates another call to Site B's Cisco Unified CallManager cluster. (This is call leg 3 in the figure.)

Figure 9-12 Simple Example of the IP-to-IP Gateway for RSVP Call Admission Control



The example in [Figure 9-12](#) is a simple scenario in which all calls between Cisco Unified CallManager clusters are routed via a pair of IP-IP gateways. However, in many real-world cases this approach might not prove scalable or flexible enough. For these cases, Cisco IOS gatekeepers can be used to provide a wider range of communication options between Cisco Unified CallManager clusters, H.323 gateways, H.323 videoconferencing endpoints, and IP-IP gateways.

**Note**

All scenarios involving IP-IP gateways described in this section apply to calls between different Cisco Unified CallManager clusters. Cisco does not recommend inserting IP-IP gateways for calls between endpoints registered to the same Cisco Unified CallManager cluster.


Via-Zone Gatekeeper

Traditional Cisco IOS gatekeeper functionality has been extended to accommodate for IP-IP gateways through the concept of a *via-zone gatekeeper*. A via-zone gatekeeper differs from legacy gatekeepers in how it uses LRQ and ARQ messages for call routing. Using via-zone gatekeepers will maintain normal gatekeeper functionality and extend it with additional features. Legacy gatekeepers examine incoming LRQs based on the called number, and more specifically the dialedDigits field in the destinationInfo portion of the LRQ. Via-zone gatekeepers look at the origination point of the LRQ before looking at the called number. If an LRQ comes from a gatekeeper listed in the via-zone gatekeeper's remote zone configurations, the gatekeeper checks to see that the zone remote configuration contains an **invia** or **outvia** keyword. If the configuration contains these keywords, the gatekeeper uses the new via-zone behavior; if not, it uses legacy behavior.

For ARQ messages, the gatekeeper determines if an **outvia** keyword is configured on the destination zone. If the **outvia** keyword is configured and the zone named with the **outvia** keyword is local to the gatekeeper, the call is directed to an IP-IP gateway in that zone by returning an ACF pointing to the IP-IP gateway. If the zone named with the **outvia** keyword is remote, the gatekeeper sends a location request to the outvia gatekeeper rather than the remote zone gatekeeper. The **invia** keyword is not used in processing the ARQ.

Figure 9-13 shows an example of how IP-IP gateways and via-zone gatekeepers can be used in conjunction with Cisco Unified CallManager clusters and legacy gatekeepers to provide call routing and call admission control. The following considerations apply to this scenario:

- Site A's Cisco Unified CallManager clusters use the Site A gatekeeper to route calls directly between them.
- The Site A gatekeeper sends all calls directed to Site B's E.164 numbers to the Site A via-zone gatekeeper.
- The Site A via-zone gatekeeper inserts an IP-IP gateway for all calls coming from or destined to the Site A gatekeeper.
- The Site A IP-IP gateway attempts RSVP reservations before sending calls toward Site B's IP-IP gateway.
- The Cisco Unified CallManager clusters, gatekeepers, and the IP-IP gateway at Site B are configured in a similar fashion to their counterparts at Site A.

Figure 9-13 IP-to-IP Gateway for RSVP Using Via-Zone Gatekeepers

Design Best Practices

When deploying IP-IP gateways in conjunction with Cisco Unified CallManager in order to enable RSVP call admission control in the IP WAN, observe these design best practices:

- When configuring trunks in Cisco Unified CallManager for either voice or video communication with other Cisco Unified CallManager clusters through one or more IP-IP gateways, use gatekeeper-controlled H.225 trunks. When using Cisco IOS Release 12.4(6)T or later and Cisco Unified CallManager Release 4.1 or later, MTP resources are no longer required to invoke supplementary services such as hold/resume, transfer, and conference across IP-IP gateways. In order to ensure interoperability, you must configure the following items:
 - On Cisco Unified CallManager Administration trunk configuration page, leave the **Media Termination Point required** field unchecked (default configuration), and also uncheck the field **Wait for Far End H.245 Terminal Capability Set**.
 - In Cisco Unified CallManager Administration, under the Advanced Service Parameters page for Cisco Unified CallManager, set the **Send H225 User Info Message** field to **H225 Info For Call Progress Tone**.
 - On the IP-IP gateways, configure the following Cisco IOS commands to ensure interoperability with Cisco Unified CallManager when invoking supplementary services:

```
voice service voip
  h323
    emptycapability
    h245 passthru tcsnonstd-passthru
```

- The MTP resources are preferred in some deployments to provide a proxy functionality and to terminate the signaling and media streams on behalf of endpoint devices. If the MTP resources are required, Cisco recommends that you place the MTP resources in the same site as the IP-IP gateway to avoid further IP WAN bandwidth usage when calling across the intercluster trunks. These MTP resources may be software-based (for example, on a Cisco MCS server or on a Cisco IOS router) or hardware-based (for example, on a Catalyst 6500 with a Cisco Communications Media Module or on a Cisco IOS router with an NM-HDV network module). Refer to the chapter on [Media Resources](#), page 6-1, for a complete list of available MTP resources. However, when MTPs are used, media packets will transit through the initial MTP resources for the entire duration of the call, with the

potential for hairpinning in the case of subsequent call transfers. Note that video calls will not be established across clusters through IP-IP gateways if the **Media Termination Point required** option is checked on the H.225 trunks (because MTPs do not support video calls).

- Configure the IP-IP gateway as an H.323 gateway in Cisco Unified CallManager only when you want all intercluster calls to use the IP-IP gateway. In this case, the IP-IP gateway can still use a gatekeeper to resolve the remote destination.
- Configure gatekeeper-controlled intercluster trunks in Cisco Unified CallManager when you want to use a gatekeeper to resolve intercluster calls and decide whether they need to go through an IP-IP gateway or be routed directly. This approach is more flexible and allows for greater scalability.
- Compatibility with Cisco CallManager 3.3(2) or later on the IP-IP gateway is available on Cisco IOS Release 12.3(1) or later. Cisco recommends that you use Cisco IOS Release 12.4(6)T or later.
- Keep the gatekeeper and via-zone gatekeeper functions separate by running them on different router platforms. Each IP-IP gateway should have a dedicated via-zone gatekeeper.
- You can run the via-zone gatekeeper function and the IP-IP gateway function on the same router platform (co-resident); however, be aware of the scalability requirements described in the section on [Redundancy, page 9-20](#).
- Do not use IP-IP gateways for calls between endpoints controlled by the same Cisco Unified CallManager cluster.
- Use the following options under the dial-peer configuration when enabling RSVP reservations for the IP-IP gateway:

```
req-qos guaranteed-delay audio
req-qos guaranteed-delay video
acc-qos guaranteed-delay audio
acc-qos guaranteed-delay video
```

This configuration ensures that for each voice or video call, the IP-IP gateway will request an RSVP reservation using the guaranteed delay service. The fact that both the requested QoS and the acceptable QoS specify this RSVP service means that the RSVP reservation is mandatory for the call to succeed (that is, if the reservation cannot be established, the call will fail). For more information on configuration details, see [Configuration Guidelines, page 9-21](#).

Redundancy

Redundancy and scalability can be provided by registering multiple IP-IP gateways with the same via-zone gatekeeper and in the same via-zone. The via-zone gatekeeper will automatically distribute the incoming calls between all the IP-IP gateways registered in the same via-zone, using a round-robin algorithm.

When an IP-IP gateway fails, it loses its registration to the via-zone gatekeeper, and the gatekeeper removes it from the list of available resources.

It is also possible to manually configure maximum-load thresholds within the IP-IP gateway so that a certain IP-IP gateway stops being selected for new calls when more than a certain percentage of its circuits are in use, and it becomes available again when the circuits in use drop below a certain percentage. The Cisco IOS commands used for this configuration are:

- On the IP-IP gateway:


```
ip circuit max-calls max-call-number
```

The above command specifies an aggregated session capacity of the IP-IP gateway, in terms of call legs. The default value is 1000 reserved call legs. Any call being handled by the IP-IP gateway costs two sessions from the available IP circuits, one session for the inbound call leg and another session for the outbound call leg. Like a regular H.323 gateway, the IP-IP gateway will automatically send its session capacity information to the via-zone gatekeeper using the H.323 version 4 protocol.

- On the gatekeeper:

```
endpoint resource-threshold onset onset-threshold abatement abatement-threshold
```

The above command makes the via-zone gatekeeper monitor the call volume in each of its gateways, including the IP-IP gateway. The via-zone gatekeeper does the active call counting when the gateway reports its session capacity information in an admission request (ARQ) or disengaged request (DRQ) message. If the active call capacity usage in a particular gateway is above the high-water mark (Range = 1 to 99; Default = 90), the via-zone gatekeeper will stop sending calls to that gateway. If the gateway's active call volume falls below the low-water mark (Range = 1 to 99; Default = 70), the via-zone gatekeeper will resume sending calls to the gateway. These thresholds are global values and affect all gateways registered with a given gatekeeper.

With both of the above commands configured, the via-zone gatekeeper is able to calculate the current session capacity utilization on the IP-IP gateway, which can prevent the via-zone gatekeeper from sending a call to the IP-IP gateway without enough capacity resources. Otherwise, the gatekeeper call admission control would fail, with an admission reject (ARJ) or location reject (LRJ) message being returned to the originating device.


For more details on these commands, refer to the Cisco IOS command reference documentation available at

<http://www.cisco.com>

Configuration Guidelines

This section presents a simple configuration example based on the network diagram shown in [Figure 9-14](#). This section is not intended as an exhaustive command reference guide, but rather as a collection of guidelines that can prove helpful in the most common deployment scenarios. Complete information on how to configure IP-IP gateways and via-zone gatekeepers is provided in the online documentation for the Cisco Multiservice IP-to-IP Gateway, available at

<http://www.cisco.com>

Figure 9-14 Configuration Example for IP-IP Gateway with Via-Zone Gatekeeper

For the network shown in [Figure 9-14](#), assume that there are two Cisco Unified CallManager clusters at Site A: cluster A1, with phone extensions 1xxx, and cluster A2, with phone extensions 2xxx. There are also two Cisco Unified CallManager clusters at Site B: cluster B1, with phone extensions 3xxx, and cluster B2, with phone extensions 4xxx.

The following subsections show the relevant configurations for devices located at Site A, so that calls within Site A are routed directly between Cisco Unified CallManager clusters (using Site A's legacy gatekeeper) while calls to Site B are routed through the two IP-IP gateways (using the respective legacy gatekeepers and via-zone gatekeepers).

Cisco Unified CallManager

Both cluster A1 and cluster A2 use a gatekeeper-controlled intercluster trunk, which is an intercluster trunk (ICT) without MTP required and which points to Site A's legacy gatekeeper.

A [34]XXX route pattern points to the ICT through a route list and route group construct, so as to reach Site B's clusters through the gatekeepers and the IP-IP gateways.

Another route pattern (2XXX for cluster A1 and 1XXX for cluster A2) allows cluster A1 and A2 to communicate with each other through the gatekeeper, by pointing to the ICT through a route list and route group construct.

Legacy Gatekeeper

The legacy gatekeeper at Site A routes calls between clusters A1 and A2 directly, while it sends all calls for Site B (extensions 3xxx and 4xxx) to Site A's via-zone gatekeeper. [Example 9-1](#) shows the relevant configuration.

Example 9-1 Legacy Gatekeeper Configuration for Site A

```
gatekeeper
zone local CCM-A1 customer.com 10.10.10.10
zone local CCM-A2 customer.com
zone remote A-VIAGK customer.com 10.10.100.10
zone prefix CCM-A1 1...
zone prefix CCM-A2 2...
zone prefix A-VIAGK 3...
```

```

zone prefix A-VIAGK 4...
gw-type-prefix 1#* default-technology
arq reject-unknown-prefix
no shutdown

```

Via-zone Gatekeeper

The via-zone gatekeeper at Site A sends calls directed to Site B's Cisco Unified CallManager clusters (extensions 3xxx and 4xxx) to Site B's via-zone gatekeeper and invokes an IP-IP gateway for calls to or from Site B. Calls directed to Site A's clusters are routed to Site A's legacy gatekeeper without invoking an IP-IP gateway. [Example 9-2](#) shows the relevant configuration.

Example 9-2 Via-Zone Gatekeeper Configuration for Site A

```

gatekeeper
zone local A-VIAGK customer.com 10.10.100.10
zone remote CCM-A1 customer.com 10.10.10.10
zone remote CCM-A2 customer.com 10.10.10.10
zone remote B-VIAGK customer.com 10.20.200.20 invia A-VIAGK outvia A-VIAGK
zone prefix B-VIAGK 3...
zone prefix B-VIAGK 4...
zone prefix CCM-A1 1...
zone prefix CCM-A2 2...
gw-type-prefix 1#* default-technology
arq reject-unknown-prefix lrq forward-queries
no shutdown

```

The following considerations apply to the configuration shown in [Example 9-2](#):

- The **invia** and **outvia** keywords in the command line related to the B-VIAGK remote zone trigger the via-zone gatekeeper behavior for that zone. This means that, for all calls destined to or coming from the B-VIAGK remote zone, the via-zone gatekeeper will invoke an IP-IP gateway resource registered in the A-VIAGK local zone.
- The absence of **invia** and **outvia** keywords in the command line related to the CCM-A1 and CCM-A2 remote zones means that the standard gatekeeper behavior is applied, and no IP-IP gateway is invoked for calls to or from these zones.

IP-IP Gateway

The IP-IP gateway at Site A requests RSVP reservations for voice and video calls directed toward Site B's Cisco Unified CallManager clusters (extensions 3xxx and 4xxx) but not for calls directed toward Site A's Cisco Unified CallManager clusters (extensions 1xxx and 2xxx). [Example 9-3](#) shows the relevant configuration.

Example 9-3 IP-IP Gateway Configuration for Site A

```

voice service voip
  allow-connections h323 to h323
  h323
    emptycapability
    h245 passthru tcsnonstd-passthru
!
gateway
!
interface FastEthernet0/1
  ip address 10.10.1.10 255.255.255.0
  ip rsvp bandwidth 200

```

```

ip rsvp data-packet classification none
ip rsvp resource-provider none
h323-gateway voip interface
h323-gateway voip id A-VIAGK ipaddr 10.10.100.10
h323-gateway voip h323-id A-IPIPGW
h323-gateway voip bind srcaddr 10.10.1.10
h323-gateway voip tech-prefix 1#
!
dial-peer voice 5 voip
  session target ras
  incoming called-number [3-4]...
  codec transparent
!
dial-peer voice 10 voip
  destination-pattern [3-4]...
  session target ras
  req-qos guaranteed-delay audio
  req-qos guaranteed-delay video
  acc-qos guaranteed-delay audio
  acc-qos guaranteed-delay video
  codec transparent
!
dial-peer voice 15 voip
  session target ras
  incoming called-number [1-2]...
  req-qos guaranteed-delay audio
  req-qos guaranteed-delay video
  acc-qos guaranteed-delay audio
  acc-qos guaranteed-delay video
  codec transparent
!
dial-peer voice 20 voip
  destination-pattern [1-2]...
  session target ras
  codec transparent

```

The following considerations apply to the configuration shown in [Example 9-3](#):

- The **emptycapability** command enables the H.245 Empty Capabilities Set (ECS) between Cisco Unified CallManager and the IP-IP gateway to invoke supplementary services for the established calls.
- The **req-qos guaranteed-delay [audio | video]** commands specify that the IP-IP gateway should request a guaranteed-delay RSVP reservation for voice and video calls that use dial-peer 10 or 15.
- The **acc-qos guaranteed-delay [audio | video]** commands specify that the minimum acceptable QoS level for voice and video calls is also a guaranteed-delay RSVP reservation. This means that, if the RSVP request fails, the call will also fail, which equates to making the RSVP reservation mandatory. To configure the IP-IP gateway so that the RSVP reservation is optional (so that the call succeeds even if the reservation fails), use the commands **acc-qos best-effort [audio | video]** instead.
- Voice calls with a transparent codec for RSVP should be used only when the codecs are the same on both call legs of the IP-IP gateway.

Call Admission Control Design

This section describes how to apply the call admission control mechanisms to the various Cisco Unified CallManager deployment models and to the following IP WAN topologies:


- [Simple Hub-and-Spoke Topologies, page 9-25](#)
- [Two-Tier Hub-and-Spoke Topologies, page 9-29](#)
- [Simple MPLS Topologies, page 9-34](#)
- [Generic Topologies, page 9-41](#)

For each topology, these sections present different sets of design considerations based on the Cisco Unified CallManager deployment model adopted.

Simple Hub-and-Spoke Topologies

Figure 9-15 depicts a simple hub-and-spoke topology, also known as a star topology. In this type of network topology, all sites (called *spoke sites*) are connected via a single IP WAN link to a central site (called the *hub site*). There are no direct links between the spoke sites, and every communication between them must transit through the hub site.

Figure 9-15 **A Simple Hub-and-Spoke Topology**



The design considerations in this section apply to simple hub-and-spoke topologies that use traditional Layer 2 IP WAN technologies such as:

- Frame Relay
- ATM
- Frame Relay/ATM Service Interworking
- Leased Lines

For IP WAN deployments based on the MPLS technology, refer to the section on [Simple MPLS Topologies, page 9-34](#).

The remainder of this section contains design best practices for simple hub-and-spoke topologies according to the Cisco Unified CallManager deployment model adopted:

- [Centralized Cisco Unified CallManager Deployments, page 9-26](#)
- [Distributed Cisco Unified CallManager Deployments, page 9-27](#)


One or more Cisco Unified CallManager clusters are located at the hub site, but only phones and gateways are located at the spoke sites.

A Cisco Unified CallManager cluster or Cisco Unified CallManager Express is located at each site.

Centralized Cisco Unified CallManager Deployments

In multisite WAN deployments with centralized call processing in a simple hub-and-spoke topology, use Cisco Unified CallManager static *locations* for implementing call admission control. Figure 9-16 shows an example of how to apply this mechanism to such a topology.

Figure 9-16 Call Admission Control for Simple Hub-and-Spoke Topologies Using Static Locations



Follow these guidelines when using static locations for call admission control:

- Configure a separate location in Cisco Unified CallManager for each spoke site.
- Configure the appropriate bandwidth limits for voice and video calls for each site according to the types of codecs used at that site. (See [Table 9-1](#) for recommended bandwidth settings.)
- Assign all devices at each spoke site to the appropriate location.
- Leave devices at the hub site in the <None> location.
- If you move a device to another location, change its location configuration as well.
- Cisco Unified CallManager supports up to 500 locations.
- If you require automatic rerouting over the PSTN when the WAN bandwidth is not sufficient, configure the automated alternate routing (AAR) feature on Cisco Unified CallManager. (See [Automated Alternate Routing](#), page 10-22.)
- If multiple Cisco Unified CallManager clusters are located at the same hub site, leave the intercluster trunk devices in the <None> location. You may use a gatekeeper for dial plan resolution. However, gatekeeper call admission control is not necessary in this case because all IP WAN links are controlled by the locations algorithm.


**Note**

If one or more sites have dual connectivity to the IP WAN, you need to configure the location bandwidth according to the worst-case scenario (that is, only the smallest of the two IP WAN links is active), or upgrade to Cisco Unified CallManager 5.0 and deploy RSVP-enabled locations. See [Limitations of Topology-Unaware Call Admission Control, page 9-5](#), for more information.

Distributed Cisco Unified CallManager Deployments

For distributed call processing deployments on a simple hub-and-spoke topology, you can implement call admission control with a Cisco IOS gatekeeper. In this design, the call processing agent (which could be a Cisco Unified CallManager cluster, Cisco Unified CallManager Express, or an H.323 gateway) registers with the Cisco IOS gatekeeper and queries it each time the agent wants to place an IP WAN call. The Cisco IOS gatekeeper associates each call processing agent with a zone that has specific bandwidth limitations. Thus, the Cisco IOS gatekeeper can limit the maximum amount of bandwidth consumed by IP WAN voice calls into or out of a zone.

[Figure 9-17](#) illustrates call admission control with a gatekeeper. In brief, when the call processing agent wants to place an IP WAN call, it first requests permission from the gatekeeper. If the gatekeeper grants permission, the call processing agent places the call across the IP WAN. If the gatekeeper denies the request, the call processing agent can try a secondary path (the PSTN, for example) or can simply fail the call.

Figure 9-17 Call Admission Control for Hub-and-Spoke Topologies Using a Gatekeeper

153172

Follow these guidelines when deploying call admission control with a gatekeeper:

- In Cisco Unified CallManager, configure an H.225 gatekeeper-controlled trunk if you have a mixed environment with Cisco Unified CallManager Express and H.323 gateways.
- In Cisco Unified CallManager, configure an intercluster gatekeeper-controlled trunk if you have an environment exclusively based on Cisco Unified CallManager clusters.
- Ensure that the zone configured in Cisco Unified CallManager matches the correct gatekeeper zone for the site.
- Each Cisco Unified CallManager subscriber listed in the device pool's CallManager Redundancy Group registers a gatekeeper-controlled trunk with the gatekeeper. (Maximum of three.)
- Calls are load-balanced across the registered trunks in the Cisco Unified CallManager cluster.
- Cisco Unified CallManager supports multiple gatekeepers and trunks.
- You can place the trunk in a route group and route list construct to provide automatic PSTN failover. (See [Dial Plan](#), page 10-1, for more details.)
- Configure a separate zone in the gatekeeper for each site supporting Cisco Unified CallManagers, Cisco Unified CallManager Express, or H.323 gateways.

- Use the **bandwidth interzone** command on the gatekeeper to control bandwidth between Cisco Unified CallManager clusters, Cisco Unified CallManager Express servers, and H.323 devices registered directly with the gatekeeper. (See [Table 9-2](#) for bandwidth settings by codec type.)
- A single Cisco IOS gatekeeper can support up to 100 zones or sites.
- You can provide gatekeeper redundancy by using gatekeeper clustering (alternate gatekeeper) or Cisco Hot Standby Router Protocol (HSRP). Use HSRP only if gatekeeper clustering is not available in your software feature set.


**Note**

If one or more sites have dual connectivity to the IP WAN and you want to take full advantage of the bandwidth available on both links, Cisco recommends that you deploy topology-aware call admission control using RSVP-enabled IP-IP gateways, as described in the section on [Generic Topologies](#), [page 9-41](#). See also [Limitations of Topology-Unaware Call Admission Control](#), [page 9-5](#), for more information.

Two-Tier Hub-and-Spoke Topologies

[Figure 9-18](#) depicts a two-tier hub-and-spoke topology. This type of network topology consists of sites at three hierarchical levels: the first-tier hub site, the second-tier hub sites, and the spoke sites. A group of spoke sites are connected to a single second-tier hub site, and each second-tier hub site is in turn connected to the single first-tier hub site. As in the simple hub-and-spoke topology, there are no direct links between the spoke sites, and every communications between them must transit through the second-tier hub site. Similarly, there are no direct links between the second-tier hub sites, and all communications between them must transit through the first-tier hub site.

Figure 9-18 *A Two-Tier Hub-and-Spoke Topology*



The design considerations in this section apply to two-tier hub-and-spoke topologies that use traditional Layer 2 IP WAN technologies such as:

- Frame Relay
- ATM
- Frame Relay/ATM Service Interworking
- Leased Lines

For IP WAN deployments based on the MPLS technology, refer to the section on [Simple MPLS Topologies](#), page 9-34.

The remainder of this section contains design best practices for two-tier hub-and-spoke topologies according to the Cisco Unified CallManager deployment model adopted:

- [Centralized Cisco Unified CallManager Deployments](#), page 9-30

One or more Cisco Unified CallManager clusters are located at the first-tier hub site, but only phones and gateways are located at the second-tier hub sites and the spoke sites.


- [Distributed Cisco Unified CallManager Deployments](#), page 9-33

Cisco Unified CallManager clusters are located at the first-tier hub site and at the second-tier hub sites, while only endpoints and gateways are located at the spoke sites.

Centralized Cisco Unified CallManager Deployments

Figure 9-19 depicts a single centralized Cisco Unified CallManager cluster deployed in a two-tier hub-and-spoke IP WAN topology. In this scenario, the Cisco Unified CallManager cluster is located at the first-tier hub site, while all second-tier hub and spoke sites contain only endpoints and gateways.


Figure 9-19 Two-Tier Hub-and-Spoke Topology with Centralized Cisco Unified CallManager



This type of topology poses a problem in terms of call admission control. The Cisco Unified CallManager static locations algorithm is the only available mechanism for call admission control in a centralized call processing deployment, and it is designed to work for simple hub-and-spoke topologies. This section describes an attempt to mitigate the issues generated by such topologies, but you should keep in mind that this workaround has limitations, and the recommended approaches to solve this problem are either to change the topology to a simple hub-and-spoke or to upgrade to Cisco Unified CallManager 5.0 and deploy RSVP-enabled locations.

Figure 9-20 shows a possible approach to enable a Cisco Unified CallManager 4.2 cluster in a centralized call processing deployment to support a two-tier hub-and-spoke topology.

Figure 9-20 Two-Tier Hub-and-Spoke in a Purely Centralized Deployment



When using this approach, adhere to the following guidelines:

- Leave the devices at the first-tier hub site (Central Site) in the <None> location.
- Place devices at each spoke site in their own location, and define the location bandwidth according to how many calls you want to allow in or out of each site. In the example of Figure 9-20, the location bandwidth for the three spoke sites (branches 1, 2 and 3) is set to 48 kbps, or enough for two calls using the G.729 codec.

- Provision the priority queue bandwidth in the LLQ configuration for the links between the second-tier hub site and the spoke sites so that it matches the spoke site location bandwidth. In this example, this bandwidth setting is 48 kbps.
- Place devices at each second-tier hub site in their own location, and define the location bandwidth according to how many calls you want to allow in or out of each site. Keep in mind that this bandwidth accounts for all calls entering or leaving the second-tier hub site, so it includes calls toward the underlying spoke sites as well as calls toward the first-tier hub site. In the example of [Figure 9-20](#), the location bandwidth for the second-tier hub site shown (Regional Site) is set to 192 kbps, or enough for eight calls using the G.729 codec.
- Provision the priority queue bandwidth in the LLQ configuration for the links between the first-tier hub site and each second-tier hub site so that it matches the sum of the location bandwidth of the second-tier hub site itself plus all the location bandwidths of the spoke sites connected to that second-tier hub site. In this example, this bandwidth setting is $192 + 48 + 48 + 48 = 336$ kbps.

**Note**

For simplicity, this example shows LLQ bandwidth provisioning based on Layer 3 values (that is, 24 kbps for a G.729 call). In reality, Layer 2 overhead also must be taken into account when provisioning the LLQ queues. Refer to [Table 3-7](#) in the section on [Bandwidth Provisioning](#), page 3-44, for a complete list of bandwidth values for the various Layer 2 WAN technologies.

As shown by the example in [Figure 9-20](#), this approach relies on over-provisioning the priority queue bandwidth between the first-tier and the second-tier hubs to compensate for the fact that Cisco Unified CallManager locations are unaware that groups of spoke sites are actually connected via a certain second-tier hub site. As a consequence, the configuration must account for bandwidth as if communications between any two sites has to transit through the first-tier hub site.

While this approach protects voice quality under all circumstances, it also involves the following design considerations and caveats:

- As the number of spoke sites increases per second-tier hub site, the amount of priority queue bandwidth that must be provisioned between the second-tier sites and the first-tier hub site also increases significantly.
- The bandwidth utilization is suboptimal because the calculations are based on the worst-case scenario. For example, in [Figure 9-20](#), if each of the three branches has two active calls to the Regional Site, Cisco Unified CallManager will allow only two additional calls from the Regional Site to the Central Site, even if the actual bandwidth available in the priority queue would allow for 14 additional calls.
- Call completion across the IP WAN cannot be guaranteed even if the actual bandwidth resources would allow for it. For example, in [Figure 9-20](#), if the Regional Site has eight active calls to the Central Site, Cisco Unified CallManager will not allow any of the branches to call the Regional Site, even if the actual bandwidth available would allow for two calls from each branch. On the other hand, in the same situation the branches will be allowed to call the Central Site.

**Note**

Cisco recommends that you do not place all spoke sites connected to a given second-tier hub site in the same location as the second-tier hub site. This solution cannot guarantee voice quality in all scenarios and is *not* supported by Cisco.


As mentioned previously, the workaround illustrated in this section has several limitations. For a better solution, Cisco recommends that you implement one of the following changes:

- Change the network topology to a simple hub-and-spoke by connecting all spoke sites directly to the first-tier hub site. Use Cisco Unified CallManager static locations to perform call admission control, as described in the section on [Simple Hub-and-Spoke Topologies](#), page 9-25.
- Change the network topology to an MPLS-based network. Use Cisco Unified CallManager static locations to perform call admission control, as described in the section on [Simple MPLS Topologies](#), page 9-34.
- Upgrade to Cisco Unified CallManager 5.0. Use Cisco Unified CallManager RSVP-enabled locations to perform call admission control, as described in the *Cisco Unified Communications SRND Based on Cisco Communications Manager 5.x*, available at <http://www.cisco.com/go/designzone>

Distributed Cisco Unified CallManager Deployments

To provide call admission control in deployments that use a two-tier hub-and-spoke topology, with Cisco Unified CallManagers at the first-tier and second-tier hub sites, you can combine the static locations and gatekeeper zone mechanisms as illustrated in [Figure 9-21](#).

Figure 9-21 Combining the Locations and Gatekeeper Mechanisms for Call Admission Control



Follow these recommendations when combining gatekeeper zones with static locations for call admission control:

- Use call admission control based on static locations for sites with no local Cisco Unified CallManager (that is, the spoke sites).

- Use gatekeeper-based call admission control between Cisco Unified CallManager clusters (that is, between the first-tier hub site and the second-tier hub sites).
- For each site without a local Cisco Unified CallManager, configure a location for that site in the Cisco Unified CallManager cluster supporting the site.
- Configure the appropriate bandwidth limits for voice and video calls at each site according to the type of codec used at that site. (See [Table 9-1](#) and [Table 9-2](#) for bandwidth settings.)
- Assign each device configured in Cisco Unified CallManager to a location. If you move a device to another location, change its location configuration as well.
- Cisco Unified CallManager supports up to 500 locations.
- Each Cisco Unified CallManager cluster registers a gatekeeper-controlled trunk with the gatekeeper.
- On the gatekeeper, configure a zone for each Cisco Unified CallManager cluster, and use the **bandwidth interzone** command to control the number of calls to and from each cluster.

**Note**


If one or more sites have dual connectivity to the IP WAN, you need to configure the location bandwidth according to the worst-case scenario (that is, only the smallest of the two IP WAN links is active), or upgrade to Cisco Unified CallManager 5.0 and deploy RSVP-enabled locations. See [Limitations of Topology-Unaware Call Admission Control, page 9-5](#), for more information.

Simple MPLS Topologies

[Figure 9-22](#) shows an IP WAN (from a service provider) based on the Multiprotocol Label Switching (MPLS) technology. The main design difference between traditional Layer 2 WAN services offered by service providers and services based on MPLS is that, with MPLS, the IP WAN topology does not conform to a hub-and-spoke but instead provides "full-mesh" connectivity between all sites.

This topology difference means that, from an IP routing perspective on the enterprise side of the network, each site is one IP hop away from all other sites. Thus, there is no need to transit through a hub site to reach any other site. In fact, there is no concept of a "hub site." All sites are considered equal, and the only difference between them is the amount of bandwidth that they are allowed to use across the IP WAN.

Figure 9-22 *MPLS IP WAN from a Service Provider, and Its Topology Equivalent*




Based on these considerations, it is easy to see that, from a call admission control perspective, a service-provider IP WAN service based on MPLS is in reality equivalent to a hub-and-spoke topology without a hub site. (See [Figure 9-22](#).) In fact, the network itself could be considered as the hub site, while all the enterprise sites (including the headquarters, or central site) are equivalent to spoke sites. These differences have implications on how to perform call admission control, which are described in the remainder of this section.

An exception to the above considerations that is worth mentioning here is represented by multisite deployments where an MPLS-based WAN co-exists with an IP WAN based on a traditional Layer 2 technology, such as Frame Relay or ATM. Such a scenario could occur, for example, in the case of network migration phases, company mergers, or various other situations.

As shown in [Figure 9-23](#), integrating a hub-and-spoke IP WAN based on a traditional Layer 2 technology (such as Frame Relay) with an MPLS-based IP WAN results in a network topology that is neither a simple hub-and-spoke nor a full-mesh, but rather is equivalent to a two-tier hub-and-spoke.

In this case the first-tier hub site is represented by the MPLS network, the second-tier hub sites are represented by the MPLS-based sites as well as the MPLS-enabled Frame Relay hub site, and the spoke sites are represented by the Frame Relay spoke sites. Therefore, for design considerations on such deployments, refer to the section on [Two-Tier Hub-and-Spoke Topologies](#), [page 9-29](#).

Figure 9-23 Co-existence of MPLS Sites and Frame Relay Sites, and the Topology Equivalent



The remainder of this section contains design best practices for MPLS-based topologies according to the Cisco Unified CallManager deployment model adopted:

- [Centralized Cisco Unified CallManager Deployments](#), [page 9-36](#)

One or more Cisco Unified CallManager clusters are located at only one site, while only endpoints and gateways are located at all other sites.

- [Distributed Cisco Unified CallManager Deployments](#), [page 9-39](#)

Cisco Unified CallManager clusters are located at multiple sites, while endpoints and gateways are located at all other sites.

**Note**

This section focuses on enterprise deployments where the MPLS WAN service is provided by a service provider. In cases where the MPLS network is deployed by the enterprise itself, call admission control can be performed effectively if one of the following two conditions is satisfied: (1) routing in the MPLS network is configured so that it is equivalent to a hub-and-spoke, or (2) bandwidth in the core of the MPLS network is heavily over-provisioned so that congestion can occur only at the edge.

**Note**


If one or more sites have dual connectivity to the IP WAN, you need to configure the location bandwidth according to the worst-case scenario (that is, only the smallest of the two IP WAN links is active), or upgrade to Cisco Unified CallManager 5.0 and deploy RSVP-enabled locations. See [Limitations of Topology-Unaware Call Admission Control, page 9-5](#), for more information.

Centralized Cisco Unified CallManager Deployments

In multisite WAN deployments with centralized call processing in an MPLS topology, use Cisco Unified CallManager static *locations* for implementing call admission control.

In a hub-and-spoke WAN topology (for example, Frame Relay or ATM), each link to and from a branch site terminates at the central site. For example, in a Frame Relay network, all permanent virtual circuits (PVCs) from the branch routers are aggregated at the central site's head-end router. In such a scenario, there is no need to apply call admission control to devices at the central site because the bandwidth accounting occurs at the branch ends of the WAN links. Therefore, within the Cisco Unified CallManager locations configuration, devices at the central site are left in the <None> location, while devices at each branch are placed in their appropriate location to ensure proper call admission control.


With an MPLS WAN network, all branches are deemed to be adjacent at Layer 3, thus they do not have to rely on the central site for connectivity. [Figure 9-24](#) illustrates a spoke-to-spoke call between two branch sites in this type of deployment.

Figure 9-24 Spoke-to-Spoke Calls in an MPLS Deployment

Also, in an MPLS WAN, the link connecting the central site to the WAN does not aggregate every branch's WAN link. By placing all the central site devices in their own call admission control location (that is, not in the <None> location), this configuration requires that call admission control be performed on the central site link independently of the branch links. (See [Figure 9-25](#).)

**Note**

Some devices such as trunks do not terminate media and are normally left in the <None> location. However, to avoid errors in call admission control when an MTP is required on a trunk, the trunk must be assigned to a location other than <None> and any MTP in the trunk's MRGL must be physically located at the site associated with that location. This configuration is required because an MTP cannot be assigned a location directly, so it inherits the location of the device that selected it.


Figure 9-25 Calls to and from the Hub in an MPLS Deployment

When all the available bandwidth for a particular site has been utilized, you can provide automatic failover to the PSTN by using the automated alternate routing (AAR) feature within Cisco Unified CallManager. (For more information on AAR, see [Automated Alternate Routing, page 10-22](#).)

Distributed Cisco Unified CallManager Deployments


In multisite deployments where a Cisco Unified CallManager cluster is present at more than one site without any branch locations and the sites are linked through an MPLS WAN, a gatekeeper can provide dial-plan resolution as well as call admission control between the sites, with each site being placed in a different gatekeeper zone. This is the same mechanism adopted for hub-and-spoke topologies based on Layer 2 WAN technologies. (See [Figure 9-26](#).)

Figure 9-26 Gatekeeper Call Admission Control in a Distributed Deployment with MPLS



In deployments where branch sites are required, a gatekeeper can be used for dial-plan resolution between clusters, but a gatekeeper is not recommended for call admission control.

When calls occur between branches belonging to different clusters, the audio path is established between the two branches directly, with no media transiting through each cluster's central site. Therefore, call admission control is required only on the WAN links at the two branches. (See [Figure 9-27](#).)

Figure 9-27 Multiple Clusters Connected by Intercluster Trunks (ICTs)

As in the centralized Cisco Unified CallManager deployments, devices that terminate media at each site (including the central sites for each cluster) must be placed in an appropriately configured location.

Note that the intercluster trunks are purely signaling devices, and there is no media transiting through them. Therefore, all intercluster trunks must be left in location <None>. The exception is when the trunk requires an MTP, in which case the trunk and MTP should both be in the location of the site in which they reside.

When all the available bandwidth for a particular site has been used, you can provide automatic failover to the PSTN by using a combination of the following two methods:


- The route list and route group construct for calls across multiple Cisco Unified CallManager clusters
- The automated alternate routing (AAR) feature for calls within a Cisco Unified CallManager cluster (For more information on AAR, see [Automated Alternate Routing](#), page 10-22.)

Generic Topologies

In the context of this chapter, a generic topology is a network topology that cannot be reduced to a simple or two-tier hub-and-spoke or to a simple MPLS-based network.

As [Figure 9-28](#) illustrates, a generic topology can present full-mesh features, hub-and-spoke features, partial-mesh features, or possibly all of them combined in a single network. It may also present dual connections between sites, as well as multiple paths from one site to another.

Figure 9-28 *A Generic Topology*



The complex nature of these networks requires the adoption of topology-aware call admission control mechanisms based on RSVP. In particular, these mechanisms can properly control bandwidth in presence of any of the following topology aspects:

- Remote sites dual-homed to different hub sites
- Multiple IP WAN links between any two sites, either in a primary/backup configuration or in an active/active load-balanced configuration
- Redundant hubs or data centers with a dedicated connection
- Fully-meshed core networks
- Multiple equal-cost IP paths between any two sites
- Multi-tiered architectures

Cisco Unified CallManager 5.0 introduces support for topology-aware call admission control using RSVP-enabled locations, therefore Cisco recommends upgrading to this release for any deployment using generic topologies.

However, even with Cisco Unified CallManager 4.x releases, you can introduce topology-aware call admission control for calls across different clusters by using the Cisco IP-to-IP Gateway, as described in the section on [Cisco IOS Gatekeeper and IP-to-IP Gateway with RSVP](#), page 9-17.

Depending on the deployment details, it might prove unfeasible to add IP-IP gateway functionality to all sites within a large network. However, assuming that certain parts of the network can be simplified to a hub-and-spoke topology (either by altering the network connectivity or by making assumptions based on the available bandwidth), it is possible to combine the topology-unaware call admission control mechanisms such as Cisco Unified CallManager static locations and Cisco IOS gatekeeper zones with topology-aware call admission control provided by RSVP across IP-to-IP gateways in the same network.

The remainder of this section presents a detailed case study based on a hypothetical large customer network whose topology cannot be simplified to a hub-and-spoke. This example combines the following call admission control mechanisms to provide an end-to-end solution:


- Cisco Unified CallManager locations
- Cisco IOS Gatekeeper
- RSVP (using the Cisco IP-to-IP Gateway)

Case Study

[Figure 9-29](#) shows the high-level topology of a hypothetical large customer network. The customer sites are divided into three main regions: the Americas (AMER), Europe-Middle East-Africa (EMEA), and Asia-Pacific (AP). Within each of these regions, the network topology is a three-tiered hub-and-spoke,

with a regional hub and several country hubs, each of which can have one or two lower levels of smaller sites. The core network interconnects the three regions across an arbitrary network topology, possibly containing multiple equal-cost paths between any two destinations.

Figure 9-29 Hypothetical Customer Network



This complex topology illustrates where to apply the various call admission control mechanisms and how to combine them. To provide a better understanding of the overall solution, this section analyzes different portions of the network separately.

First, consider the AMER region. As illustrated in [Figure 9-29](#), there are two country hubs that connect to the region hub: the US hub site, located in Dallas, and the Canada hub site, located in Toronto.

[Figure 9-30](#) shows how to implement call admission control for the US portion of the network.

Call Admission Control for the First Two Levels of the Customer Network in the US



A Cisco IOS gatekeeper is placed at the US hub site in Dallas. A number of first-level spoke sites (up to 100 per gatekeeper) can connect to this hub site. In the example shown in [Figure 9-30](#), these sites are Los Angeles (LAX) and Boston (BOS). By defining a local zone for each site in the gatekeeper (USHub, USWest, and USEast in the example), you can perform call admission control on the links between the first-level spoke sites and the US hub site using the Cisco IOS gatekeeper command **bandwidth interzone zone zone-name bandwidth-value**. This command ensures that the bandwidth used by voice and video calls in and out of each zone does not exceed the configured value.


In this figure, the first-level spoke sites (LAX and BOS) each host a Cisco Unified CallManager cluster. Because Cisco Unified CallManager uses its own call admission control mechanism (the locations construct), it is possible to add a number of second-level spoke sites (up to 500 per cluster) to each of the first-level spoke sites. In this example, two additional sites are controlled by each Cisco Unified CallManager cluster: San Francisco (SFO) and Seattle (SEA) for the Los Angeles cluster and New York (JFK) and Washington, DC (IAD), for the Boston cluster. The bandwidth on the links between these second-level spoke sites and their respective first-level site is therefore controlled by the Cisco Unified CallManager locations.

**Note**

As far as the gatekeeper is concerned, the second-level spoke sites are considered to belong to the same zone as their "parent" first-level spoke site because they are controlled by the same Cisco Unified CallManager cluster. This analysis is correct because calls from any of the second-level spoke sites directed to an on-net destination not controlled by the Cisco Unified CallManager cluster will have to traverse the link between the "parent" first-level spoke site and the country hub site.

Figure 9-31 shows the Canada (CND) sites, which provide another example of how call admission control is implemented at this level in the network.

Figure 9-31 Call Admission Control for the First Level of the Customer Network in Canada




Another Cisco IOS gatekeeper is placed at the Canada hub site in Toronto, and a number of spoke sites connect to the hub site. In the example, the two spoke sites are Vancouver and Montreal. Note that these spoke sites do not contain a Cisco Unified CallManager cluster but do contain a variety of H.323-based devices such as voice gateways, video endpoints, Multipoint Control Units (MCUs), and H.320 video gateways. For each site, a local zone (CNDHub, CNDWest, and CNDEast in the example) is defined within the gatekeeper. The bandwidth on the links between the Canada hub site and the spoke sites can again be controlled using the Cisco IOS gatekeeper command **bandwidth interzone zone zone-name bandwidth-value**. This command ensures that the bandwidth used by voice and video calls in and out of each zone does not exceed the configured value.

Because the H.323 devices at the spoke sites cannot perform call admission control, it is not possible to add a second level of spoke sites in this case. However, because all these devices can also be controlled by Cisco Unified CallManager, it would be possible to add a Cisco Unified CallManager cluster to the first-level spoke sites and distribute these H.323 devices across multiple second-level spoke sites.

Also note that, within the sites controlled by the same gatekeeper, it is possible to mix and match Cisco Unified CallManager spoke sites and H.323 spoke sites.

Now consider the next level up in the network hierarchy. Figure 9-32 illustrates the call admission control mechanisms used between the AMER region hub site in Chicago and the two country hub sites for the US and Canada, located in Dallas and Toronto respectively.

Figure 9-32 Call Admission Control for the Second Level of the Customer Network (Region Hubs)



This part of the network is also a hub-and-spoke, but in this case the approach used to perform call admission control is different from that used between the country hub sites and the first-level spoke sites.

As discussed earlier, the country hub sites contain a Cisco IOS gatekeeper that controls bandwidth to the first-level spoke sites using local gatekeeper zones and the **bandwidth interzone** Cisco IOS command. The AMER region hub site contains a gatekeeper used as a directory gatekeeper, whose main purpose is to handle call routing between the country gatekeepers.

To control bandwidth usage on the links between the country hub sites and the region hub site, you can use the **bandwidth remote bandwidth-value** Cisco IOS command on the two gatekeepers located at the country hub sites. This command ensures that the bandwidth used by voice and video calls between all the local zones and all non-local zones does not exceed the configured value. Because all calls to

destinations that do not belong to a local zone must traverse the link between the country hub site and the region hub site, you can use the **bandwidth remote** command to perform call admission control effectively on this link.

In Figure 9-32, also notice the presence of a via-zone gatekeeper at the AMER region hub site. The directory gatekeeper is configured (configuration not shown in Figure 9-32) to route calls to destinations in other regions (such as AP and EMEA) to the AMER via-zone gatekeeper.

Next consider the core section of the network, which connects the three region hub sites. Figure 9-33 and Figure-35 illustrate how call routing and call admission control are performed across the core network.

Figure 9-33 Call Routing Across the Core of the Customer Network Using Via-Zone Gatekeepers




Figure 9-33 shows partial configurations of the three via-zone gatekeepers located in each of the three region hub sites. Each region's via-zone gatekeeper is configured so that all calls destined to or coming from the other two regions will cause the gatekeeper to insert an IP-to-IP gateway into the call.

For example, the configuration for the AMER region's via-zone gatekeeper defines the following zones:

- A single local zone called AMER-ViaGK
- The remote zone AMER-DGK, used to reach the AMER directory gatekeeper and, hence, all destinations within the AMER region

- The remote zone EMEA-DGK, used to reach destinations in the EMEA region
- The remote zone AP-DGK, used to reach destinations in the AP region


The definitions of the EMEA-DGK and AP-DGK remote zones also specify that the local zone AMER-ViaGK should be used as the **invia** and **outvia** zone to reach these destinations. This means that all calls destined to or coming from these remote zones will cause the AMER via-zone gatekeeper to insert an IP-to-IP gateway into the call. The specific IP-to-IP gateway resource will be selected from among all those registered with the gatekeeper in the local zone AMER-ViaGK.

**Note**

Inserting an IP-to-IP gateway into a call effectively splits the call into two call legs. In this example, because an IP-to-IP gateway is inserted at each region hub site, a call between an endpoint in the AMER region and one in the EMEA region would consist of three call legs: the first leg from the AMER endpoint to the AMER IP-to-IP gateway, the second leg between the AMER IP-to-IP gateway and the EMEA IP-to-IP gateway, and the third leg between the EMEA IP-to-IP gateway and the EMEA endpoint.

Figure 9-34 illustrates the call admission control mechanism used in the core network.

Figure 9-34 RSVP-Based Call Admission Control Across the Core of the Customer Network



126691

The IP-to-IP gateways are inserted into the call flow by the via-zone gatekeepers whenever the call needs to traverse the core network. Each IP-to-IP gateway is in turn configured to use the via-zone gatekeeper to route its calls and to use RSVP call admission control for the calls it originates and terminates across the core network. Because RSVP reservations are unidirectional, each voice call generates two RSVP reservations, one for each direction.

Figure 9-34 shows a call between the AMER IP-to-IP gateway and the EMEA IP-to-IP gateway, assuming that multiple paths exist between the two regional hubs across the core network. If the routes taken are asymmetrical, the two RSVP reservations will travel along different paths. However, for a given reservation, the RSVP messages will always travel along the same path because RSVP uses reverse hop-by-hop routing to establish the reservation. RSVP ensures that the bandwidth used by voice and video calls across the core network does not exceed the configured values on each router's interface.

In summary, end-to-end call admission control is achieved by combining different techniques in different portions of the network.

Referring to the network depicted in Figure 9-29, for example, a call between two IP phones located in San Francisco in the US and Liverpool in the UK would use the following mechanisms:

- Cisco Unified CallManager locations for the link between San Francisco and Los Angeles
- Cisco IOS gatekeeper **bandwidth inter-zone** command for the link between Los Angeles and Dallas
- Cisco IOS gatekeeper **bandwidth remote** command for the link between Dallas and Chicago
- RSVP between Chicago and Amsterdam (the two regional hub sites)
- Cisco IOS gatekeeper **bandwidth remote** command for the link between Amsterdam and London (the country hub site for the UK)
- Cisco IOS gatekeeper **bandwidth inter-zone** command for the link between London and Manchester (the hub site for England)
- Cisco Unified CallManager locations for the link between Manchester and Liverpool



CHAPTER 10

Dial Plan

Last revised on: February 13, 2008

The dial plan is one of the key elements of an IP Telephony system, and an integral part of all call processing agents. Generally, the dial plan is responsible for instructing the call processing agent on how to route calls. Specifically, the dial plan performs the following main functions:

- **Endpoint addressing**
Reachability of internal destinations is provided by assigning directory numbers (DNs) to all endpoints (such as IP phones, fax machines, and analog phones) and applications (such as voicemail systems, auto attendants, and conferencing systems)
- **Path selection**
Depending on the calling device, different paths can be selected to reach the same destination. Moreover, a secondary path can be used when the primary path is not available (for example, a call can be transparently rerouted over the PSTN during an IP WAN failure).
- **Calling privileges**
Different groups of devices can be assigned to different classes of service, by granting or denying access to certain destinations. For example, lobby phones might be allowed to reach only internal and local PSTN destinations, while executive phones could have unrestricted PSTN access.
- **Digit manipulation**
In some cases, it is necessary to manipulate the dialed string before routing the call; for example, when rerouting over the PSTN a call originally dialed using the on-net access code, or when expanding an abbreviated code (such as 0 for the operator) to an extension.
- **Call coverage**
Special groups of devices can be created to handle incoming calls for a certain service according to different rules (top-down, circular hunt, longest idle, or broadcast).

This chapter examines the following main aspects of dial plan:

- [Planning Considerations, page 10-2](#)
This section analyzes the thought process involved in planning an IP Telephony dial plan, ranging from the number of digits used for internal extensions to the overall architecture of a company's internal dial plan. (Prerequisite: Some familiarity with dial plans in general.)
- [Dial Plan Elements, page 10-7](#)

This section provides detailed explanations of the elements of a Cisco Unified Communications dial plan. Covered topics include call routing logic, calling privileges, and digit manipulation techniques for various Cisco products. (Prerequisite: A working knowledge of Cisco Unified CallManager and Cisco IOS is recommended.)

- [Design Considerations, page 10-51](#)

This section contains design and deployment guidelines related to multisite IP Telephony networks, endpoint addressing methods, approaches to building classes of service, and call coverage functionality. (Prerequisite: A working knowledge of Cisco Unified CallManager and Cisco IOS is recommended.)

For more details, refer to the *Cisco Unified CallManager System Guide*, the *Cisco IOS Voice, Video, and Fax Configuration Guide, Release 12.2*, and other product documentation available at

<http://www.cisco.com>

Planning Considerations

The dial plan is the most fundamental attribute of a telephony system. It is at the very core of the user experience because it defines the rules that govern how a user reaches any destination. These rules include:

- Extension dialing — how many digits must be dialed to reach an extension on the system
- Extension addressing — how many digits are used to identify extensions
- Dialing privileges — allowing or not allowing certain types of calls
- Path selection — for example, using the IP network for on-net calls, or using one carrier for local PSTN calls and another for international calls
- Automated selection of alternate paths in case of network congestion — for example, using the local carrier for international calls if the preferred international carrier cannot handle the call
- Blocking of certain numbers— for example, pay-per-minute calls
- Transformation of the called number — for example, retaining only the last five digits of a call dialed as a ten-digit number
- Transformation of the calling number — for example, replacing a caller's extension with the office's main number when calling the PSTN

A dial plan suitable for an IP telephony system is not fundamentally different from a dial plan designed for a traditional TDM telephony system; however, an IP-based system presents the dial plan architect with some new possibilities. For example, because of the flexibility of IP-based technology, telephony users in separate sites who used to be served by different, independent TDM systems can now be included in one, unified IP-based system. These new possibilities afforded by IP-based systems require some rethinking of the way we look at dial plans. This section examines some of the elements that the system planner must consider to properly establish the requirements that drive the design of the dial plan.

Dialed Pattern Recognition

Digit strings dialed by a user on a telephone generally follow patterns. For instance, many enterprises implement a five-digit abbreviated dialing pattern for calls made within the same office location. Also, many enterprises rely on a single-digit access code to represent outside dialing, followed by some

quantity of digits to reach a local PSTN number or a long-distance PSTN number (for example, 9 followed by seven digits to reach a local number, or 9 followed by 1 and ten digits to reach a long-distance destination).

The system administrator must plan the system's recognition of these patterns to ensure that the system will act promptly upon detection of a string that corresponds to a predetermined pattern so that users experience no (or minimal) post-dialing delay.

In Cisco Unified CallManager 4.2, you can implement pattern recognition by configuring route patterns, translation patterns, phone DNs, and so forth. With each digit dialed by the user, the phone sends a signaling message to Cisco Unified CallManager, which performs the incremental work of recognizing a matching pattern. As each key press from the user input is collected, Cisco Unified CallManager's digit analysis provides appropriate user feedback, such as:

- Playing dial tone when the phone first goes off-hook
- Stopping dial tone once a digit has been dialed
- Providing secondary dial tone if an appropriate sequence of digits has been dialed, such as when the off-net access code 9 is dialed

Once digit dialing is completed, Cisco Unified CallManager provides user feedback in the form of call progress tones, such as ringback tone if the destination is in the alerting stage or reorder tone if the destination is invalid.

On-Net versus Off-Net Dialing

Calls that originate and terminate on the same telephony network are considered to be on-network (or on-net). By contrast, if a call originates in company A and terminates at company B, it probably has to be routed through different telephony networks: first company A's network, followed by the PSTN, and finally into company B's network. From the caller's perspective, the call was routed off-network (or off-net); from the called party's perspective, the call originated off-net.

In TDM systems, the on-net boundaries of a telephony system are established by the PBX or Centrex system, and they typically do not extend outside of a single site. When they do, they typically do not include sites not immediately on the periphery of a large system hub.

One of the key attributes of IP telephony is its ability to expand the boundaries of calls that can be considered on-net. For instance, telephony users in an enterprise with six dispersed branch offices might be used to reaching one colleague with abbreviated dialing (for example, four-digit dialing) if the called party is located in the same site but dialing a full PSTN number to reach another colleague located in a different site. With an IP-based system where all users are served by the same IP network, it now becomes economically feasible to unite all six branches under a four-digit abbreviated dialing plan, with the IP network as the preferred path and automated overflow to the PSTN as a secondary path if the IP network is congested.

Abbreviated Dialing

Consider an extension with direct inward dial (DID) capability, which can be reached directly from the PSTN. An off-net PSTN caller has to dial the fully qualified PSTN number (for example, 1 415 555 1234) to reach a DID extension. An on-net caller might, however, prefer the ability to reach that same extension by simply dialing the last few digits of the DID number. In a four-digit abbreviated dial plan, the on-net caller would dial only 1234 in this example to reach the same extension.

Avoiding Overlap of Extension Dialing

A telephony system must be configured so that any extension can be reached in an unambiguous manner. To accomplish this goal, the dial plan must satisfy the following requirements:

- All on-net extension dialing must be globally unique. For instance, in a system using an abbreviated four-digit on-net dial plan, there cannot be an extension 1000 in site A and another extension 1000 in site B if the requirement is to reach either of them by dialing only four digits from site C.
- There cannot be any partial overlap between different dial strings.
 - For instance, if 9 is used as an off-net access code in a four-digit abbreviated dial plan (for example, to make PSTN calls), there cannot be any extensions in the 9XXX range. Attempting to do so would create situations where calls are not routed immediately. For example, if a user dialed 9141, the system would have to wait for either more digits (if the user were dialing 91415551234, for example) or the expiration of the interdigit timeout before routing the call to extension 9141. Likewise, if an operator code is used (for example, 0), the entire 0XXX extension range would have to be excluded from a four-digit uniform dial plan.
 - There cannot be overlapping strings of different length. For example, a system with extensions 1000 and 10000 would force users to wait for the interdigit timeout when they dial 1000.

Dialing String Length

The number of dialable extensions determines the quantity of digits needed to dial extensions. For example, a four-digit abbreviated dial plan cannot accommodate more than 10000 extensions (from 0000 to 9999). If 0 and 9 are reserved as operator code and off-net access code, respectively, the number range is further reduced to 8000 (1000 to 8999).

Uniform On-Net Dial Plan

A dial plan can be designed so that all extensions within the system are reached in a uniform way; that is, a fixed quantity of digits is used to reach a given extension from any on-net origination point. Uniform dialing is desirable because of the simplicity it presents to users. A user does not have to remember different ways to dial a number when calling from various on-net locations.

For example, if phone A is reached by dialing 1234 from any on-net location, whether the calling phone is in the same office or at a different site, the enterprise's dial plan is deemed uniform.

When the enterprise consists of few sites, this approach can be used with few complications. The larger the enterprise, in terms of number of extensions and sites, the more of the following challenges it faces in designing a uniform dial plan:

- The number of extensions can exceed the range afforded by the quantity of digits being considered for the dial plan. For instance, if more than 8000 extensions are required (considering the exclusions of the 0XXX and 9XXX ranges), the system may require that an abbreviated dial plan use more than four digits.
- Matching on-net abbreviated extensions to DID numbers means that, when a new DID range is obtained from a local exchange carrier, it cannot conflict with the pre-existing on-net abbreviated dial ranges. For example, if the DID range of 415 555 1XXX exists in a system using a four-digit uniform abbreviated dial plan, and DID range 650 556 1XXX is also being considered, it might be desirable to increase the quantity of digits for on-net dialing to five. In this example, the five-digit on-net ranges 51XXX and 61XXX would not overlap.

- Most systems require the exclusion of certain ranges due to off-net access codes and operator dialing. In such a system where 9 and 0 are reserved codes, no dial plan (uniform or not) could accommodate on-net extension dialing that begins with 9 or 0. This means that DID ranges could not be used if they would force the use of 9 or 0 as the first digit in the dial plan. For instance, in a five-digit abbreviated dial plan, the DID range 415 559 XXXX (or any subset thereof) could not be used. In this example, alternatives include increasing the length of the abbreviated dialing to six or more digits, or avoiding any DID range whose last five digits start with 9.

Once a given quantity of digits has been selected and the requisite ranges have been excluded (for example, ranges beginning with 9 or 0), the remaining dialing space has to be divided between all sites.

Most systems require that two ranges be excluded, thus leaving eight different possibilities for the leading digit of the dial range. [Table 10-1](#) exemplifies the distribution of the dialing space for a typical four-digit uniform dial plan.

Table 10-1 *Distribution of Digits in a Typical Four-Digit Uniform Dial Plan*

Range	Use	DID Ranges	Non-DID Ranges
0XXX	Excluded; 0 is used as off-net access code		
1XXX	Site A extensions	418 555 1XXX	N/A
2XXX	Site B extensions	919 555 2XXX	N/A
3XXX	Site C extensions	415 555 30XX	3[1-9]XX
4[0-4]XX	Site D extensions	613 555 4[0-4]XX	N/A
4[5-9]XX	Site E extensions	450 555 4[5-9]XX	N/A
5XXX	Site A extensions	418 555 5XXX	N/A
6XXX	Site F extensions	514 555 6[0-8]XX	69XX
7XXX	Future	XXX XXX 7XXX	7XXX
8XXX	Future	XXX XXX 8XXX	8XXX
9XXX	Excluded; 9 is used as off-net access code		

For the example in [Table 10-1](#), the various sites were assigned numbers in the following ways:

- Site A, the company headquarters, requires more than 1000 extensions, so two entire ranges of numbers have been retained (1XXX and 5XXX). Note that the corresponding DID ranges must also be retained from the site's local exchange carrier.
- Site B has been assigned an entire range (2XXX), allowing for up to 1000 extensions.
- Site C was also assigned an entire range, but it has been split between 100 DID extensions (415 555 30XX) and up to 900 non-DID extensions. If growth requires more extensions for DID, any unassigned numbers from the non-DID range could be used.
- Sites D and E were each assigned 500 numbers from the 4XXX range. Note that their corresponding DID ranges must match each of the site's respective portions of the 4XXX range. Because the DID ranges are for different sites (probably from different PSTN service providers), more coordination effort is required to split ranges between sites. As the quantity of sites assigned within a given range increases, it becomes increasingly difficult (sometimes impossible) to make full use of an entire range.

- Site F's range is split between 900 DID numbers (6[0-8]XX) and 100 non-DID numbers (69XX).
- The ranges 7XXX and 8XXX are reserved for future use.

When implementing a new dial plan, one of the main desires of any planner is to avoid having to change phone numbers. In addition, the extension ranges of any existing phone systems may have overlapped without any problems in the past, but they could be incompatible with a uniform dial plan.

Variable Length On-Net Dial Plan

Systems with many sites or overlapping site extension ranges can benefit from the use of a variable-length dial plan with the following characteristics:

- Within a site, the system retains the use of abbreviated dialing for calls to on-net extensions (for example, four-digit dialing).
- Between sites, users dial an access code followed by a site code and the destination's on-net extension.
- Off-net calls require an access code followed by a PSTN number.

The use of access and site codes (see [Table 10-2](#)) enables the on-net dial plan to differentiate between extensions that would overlap if a uniform abbreviated dial plan were implemented.

Table 10-2 *Typical Use of Site Codes*

Site Code	Range	Use	DID Ranges	Non-DID Ranges
1	1XXX	Site A extensions	418 555 10XX	1[1-9]XX
2	1XXX	Site B extensions	919 555 1XXX	N/A
3	1XXX	Site C extensions	907 555 1XXX	N/A

In [Table 10-2](#), sites A, B, and C are independently assigned the four-digit range 1XXX. For calls from site A to site B under the old telephony system, the calls had to be routed as off-net calls. With the new system, these calls can be dialed as on-net calls.

From site A, users simply dial 1234 to reach extension 1234. But from site B, the dial plan must accommodate the ability to reach extension 1234 at site A without conflicting with site B's own extension 1234. Therefore, each site is assigned a site code.

From site B, merely dialing the combination of site A's code with the desired extension is not feasible: in this case because 11234 would partially overlap with site B's extension 1123, thus causing interdigit timeout issues. If, instead, we assign 8 as an inter-site on-net access code, this would allow site B to dial 81234 to reach site A's extension 1234.

The following factors determine the quantity of digits required to dial an on-net, off-site extension:

- One digit for the inter-site access code
- N digits for the site code, where N is chosen to satisfy the quantity of site codes required (For example, if a system has 13 sites, then a minimum of two digits are required for the site code.)
- The quantity of digits required by the destination site's local dial plan

For example, a system with 75 sites which each use four-digit abbreviated dialing would require a format of 8 + SS + XXXX, where 8 is the on-net access code, SS is a two-digit site code, and XXXX is a four-digit extension number, giving a total of seven digits.

On-Net and Off-Net Access Codes

It is common practice in most enterprise telephony systems to dedicate a digit (for example, 9) as an off-net access code to steer calls to an off-net destination. In the variable-length on-net dial plan, another steering digit (for example, 8) is also required as an on-net access code to dial calls to on-net extensions at other sites. These two access codes, along with the use of an operator access code (for example, 0), implicitly exclude three of the ten possible leading digits of any dialed string. This restriction might not prove convenient for either of the following reasons:

- The users would be required to know the difference between on-net and off-net destinations, and to select the proper access code.
- The exclusion of three entire dialing ranges can become too restrictive or can conflict with some pre-assigned extension ranges. For instance, if a site already uses an abbreviated dialing range beginning with 8, the use of that same digit as an access code would require a change.

In systems where the same off-net access code (for example, 9) is already in use by all sites, it might be desirable to use the same code for both off-net and on-net off-site destinations. This approach has two main implications:

- To avoid partial overlap and interdigit timeout situations, the quantity of digits expected after the access code should be uniform.
- The telephony system must be able to recognize all on-net numbers dialed as off-net numbers and to route them over the IP network. This task can be simple for small systems with only one Cisco Unified CallManager cluster but complex for large systems with multiple Cisco Unified CallManager clusters.

Plan Ahead

Implementing an IP-based system might require changing certain existing user practices. Although it is preferable to plan a new system so that the implementation is as transparent as possible to users, dialing habits might have to be adapted to accommodate the integration of multiple sites that used to be on separate telephony systems. For instance, adapting to a new global, enterprise-wide dial plan might require changing the way a user reaches another user at a different site, the quantity of digits used to make intra-site calls and, in some cases, the extension numbers. To avoid exposing users to multiple generations of dial plan changes, try to anticipate expansion of the enterprise, which could result in the addition of sites in different dialing regions, an increase in the required number of on-net extensions, PSTN renumbering (for example, an area code split), or business expansion into different countries.

Dial Plan Elements

This section provides design and configuration guidelines for the following dial plan elements within a Cisco Unified Communications system:

- [User Input on SCCP Phones, page 10-8](#)
- [Call Routing in Cisco Unified CallManager, page 10-9](#)
- [Calling Privileges in Cisco Unified CallManager, page 10-15](#)
- [Digit Manipulation in Cisco Unified CallManager, page 10-21](#)
- [Automated Alternate Routing, page 10-22](#)
- [Device Mobility, page 10-26](#)


- [Extension Mobility, page 10-27](#)
- [Hunt Lists and Line Groups, page 10-29](#)
- [Call Routing in Cisco IOS with H.323 Dial Peers, page 10-35](#)
- [Calling Privileges in Cisco IOS with H.323 Dial Peers, page 10-47](#)
- [Digit Manipulation in Cisco IOS with H.323 Dial Peers, page 10-49](#)

User Input on SCCP Phones

IP phones using SCCP report every single user input event to Cisco Unified CallManager immediately. For instance, as soon as the user goes off-hook, a signaling message is sent from the phone to the Cisco Unified CallManager server with which it is registered. The phone can be considered to be a terminal, where all decisions resulting from the user input are made by the Cisco Unified CallManager server's configured dial plan.

As other user events are detected by the phone, they are relayed to Cisco Unified CallManager individually. A user who goes off-hook and then dials 1000 would trigger five individual signaling events from the phone to Cisco Unified CallManager. All the resulting feedback provided to the user, such as screen messages, playing dial tone, secondary dial tone, ring back, reorder, and so forth, are commands issued by Cisco Unified CallManager to the phone in response to the dial plan configuration. (See [Figure 10-1](#).)

Figure 10-1 User Input and Feedback for SCCP Phones



It is neither required nor possible to configure dial plan information on IP phones running SCCP. All dial plan functionality is contained in the Cisco Unified CallManager cluster, including the recognition of dialing patterns as user input is collected.

If the user dials a pattern that is denied by Cisco Unified CallManager, reorder tone is played to the user as soon as that pattern becomes the best match in Cisco Unified CallManager's digit analysis. For instance, if all calls to the pay-per-minute Numbering Plan Area (or area code) 976 are denied, reorder tone would be sent to the user's phone as soon as the user dials 91976.

Call Routing in Cisco Unified CallManager


All dialing destinations configured in Cisco Unified CallManager are added to its internal call routing table as patterns. These destinations include IP phone lines, voicemail ports, route patterns, translation patterns, and CTI route points.

When a number is dialed, Cisco Unified CallManager uses closest-match logic to select which pattern to match from among all the patterns in its call routing table. In practice, when multiple potentially matching patterns are present, the destination pattern is chosen based on the following criteria:

- It matches the dialed string, and
- Among all the potentially matching patterns, it matches the fewest strings other than the dialed string.

For example, consider the case shown in [Figure 10-2](#), where the call routing table includes the patterns 1XXX, 12XX, and 1234.

Figure 10-2 Cisco Unified CallManager Call Routing Logic Example



When user A dials the string 1200, Cisco Unified CallManager compares it with the patterns in its call routing table. In this case, there are two potentially matching patterns, 1XXX and 12XX. Both of them match the dialed string, but 1XXX matches a total of 1000 strings (from 1000 to 1999) while 12XX matches only 100 strings (from 1200 to 1299). Therefore, 12XX is selected as the destination of this call.

When user B dials the string 1212, there are three potentially matching patterns, 1XXX, 12XX and 121X. As mentioned above, 1XXX matches 1000 strings and 12XX matches 100 strings. However, 121X matches only 10 strings; therefore it is selected as the destination of the call.

When user C dials the string 1234, there are three potentially matching patterns, 1XXX, 12XX, and 1234. As mentioned above, 1XXX matches 1000 strings and 12XX matches 100 strings. However, 1234 matches only a single string (the dialed string); therefore it is selected as the destination of this call.




Note

Whenever a directory number (DN) is configured in Cisco Unified CallManager Release 4.0 and later, it is placed in the call routing table, regardless of whether the respective device (for example, an IP phone) is registered or not. This behavior differs from previous Cisco Unified CallManager versions,

where patterns were added to the call routing table only if the respective devices were registered. An implication of this modified behavior is that it is no longer possible to rely on secondary, identical patterns to provide failover capabilities to applications when the application (and hence the primary pattern) is unregistered. Because the primary pattern is now permanently in the call routing table, the secondary pattern will never be matched. It is, however, possible to achieve the same effect by using the Forward No Answer field on primary patterns such as CTI route points because this field accepts wildcards in Cisco Unified CallManager Release 4.0 and later.

Cisco Unified CallManager automatically "knows" how to route calls to internal destinations within the same cluster. For external destinations such as PSTN gateways, H.323 gatekeepers, or other Cisco Unified CallManager clusters, you have to use the external route construct (described in the following sections) to configure explicit routing. This construct is based upon a three-tiered architecture that allows for multiple layers of call routing as well as digit manipulation. Cisco Unified CallManager searches for a configured route pattern that matches the external dialed string and uses it to select a corresponding route list, which is a prioritized list of the available paths for the call. These paths are known as route groups and are very similar to trunk groups in traditional PBX terminology. Figure 10-3 depicts the three-tiered architecture of the Cisco Unified CallManager external route construct.

Figure 10-3 External Route Pattern Architecture



The following sections describe the individual elements of the external route construct in Cisco Unified CallManager:

- [Route Patterns](#), page 10-11
- [Route Lists](#), page 10-14

- [Route Groups, page 10-14](#)
- [Route Group Devices, page 10-15](#)

Route Patterns

Route patterns are strings of digits and wildcards, such as 9.[2-9]XXXXXX, configured in Cisco Unified CallManager to route calls to external entities. The route pattern can point directly to a gateway for routing calls or point to a route list, which in turn points to a route group and finally to a gateway.

Cisco strongly recommends that you use the complete route pattern, route list, and route group construct because it provides the greatest flexibility for call routing, digit manipulation, route redundancy, and future dial plan growth.

The @ Wildcard

- The @ wildcard is a special macro function that expands into a series of patterns representing the entire national numbering plan for a certain country. For example, configuring a single unfiltered route pattern such as 9.@ with the North American Numbering Plan really adds 166 individual route patterns to the Cisco Unified CallManager internal dial plan database.
- It is possible to configure Cisco Unified CallManager to accept other national numbering plans. Once this is done, the @ wildcard can be used for different numbering plans even within the same Cisco Unified CallManager cluster, depending on the value selected in the Numbering Plan field on the Route Pattern configuration page. For more information, please refer to the *Cisco Unified CallManager Dial Plan Deployment Guide*, available at http://www.cisco.com/en/US/products/sw/voicesw/ps5629/prod_maintenance_guides_list.html
- The @ wildcard can be practical in several small and medium deployments, but it can become harder to manage and troubleshoot in large deployments because adopting the @ wildcard forces the administrator to use route filters to block certain patterns (see [Route Filters, page 10-11](#)).

Route Filters

- Use route filters only with the @ route pattern to reduce the number of route patterns created by the @ wildcard. A route filter applied to a pattern not containing the @ wildcard has no effect on the resulting dial plan.
- The logical expression you enter with the route filter can be up to 1024 characters, excluding the NOT-SELECTED fields.
- As you increase the number of logical clauses in a route filter, the refresh time of the configuration page also increases and can become unacceptably long.
- For large-scale deployments, use explicit route patterns rather than the @ wildcard and route filters. This practice also facilitates management and troubleshooting because all patterns configured in Cisco Unified CallManager are easily visible from the Route Pattern configuration page.

International and Variable-Length Route Patterns

- International destinations are usually configured using the ! wildcard, which represents any quantity of digits. For example, in North America the route pattern 9.011! is typically configured for international calls. In most European countries, the same result is accomplished with the 0.00! route pattern.
- The ! wildcard is also used for deployments in countries where the dialed numbers can be of varying lengths. In such cases, Cisco Unified CallManager does not know when the dialing is complete and will wait for 15 seconds (by default) before sending the call. You can reduce this delay in any of the following ways:
 - Reduce the T302 timer (Service Parameter TimerT302_msec) to indicate end of dialing, but do not set it lower than 4 seconds to prevent premature transmission of the call before the user is finished dialing.
 - Configure a second route pattern followed by the # wildcard (for example, 9.011!# for North America or 0.00!# for Europe), and instruct the users to dial # to indicate end of dialing. This action is analogous to hitting the "send" button on a cell phone.

Overlap Sending and Overlap Receiving

In countries whose national numbering plan is not easily defined with static route patterns, you can configure Cisco Unified CallManager for overlap sending and overlap receiving.

Overlap sending means that Cisco Unified CallManager keeps collecting digits as they are dialed by the end users, and passes them on to the PSTN as they are dialed. To enable overlap sending in Cisco Unified CallManager Release 4.0 or later, check the Allow Overlap Sending box on the Route Pattern Configuration page. (In earlier Cisco Unified CallManager releases, overlap sending is enabled by setting the SendingCompleteIndicator service parameter to False.) The route pattern needs only to include the PSTN access code (for example, "9." in North America or "0." in many European countries).

Overlap receiving means that Cisco Unified CallManager receives the dialed digits one-by-one from a PRI PSTN gateway, and it then waits for completion of the dialed string before attempting to route the call to an internal destination. To enable overlap receiving in Cisco Unified CallManager Release 3.3(3) or later, set the OverlapReceivingFlagForPRI service parameter to True. (In earlier Cisco Unified CallManager releases, the parameter name was OverlapReceivingForPriFlag.)

Digit Manipulation in Route Patterns

- Configure digit manipulation only in the route list's view of its member route group and not in the route pattern.
- Digit manipulation in the route group completely overrides any digit manipulation done in the route pattern.
- If you configure digit manipulation in the route pattern, the Call Detail Record (CDR) records the dialed number after the digit manipulation has occurred. If you configure digit manipulation only in the route group, the CDR records the actual dialed number prior to the digit manipulation.
- Similarly, if you configure digit manipulation in the route pattern, the IP phone display of the calling party will show the manipulated number. If you configure digit manipulation only in the route group, the manipulations will be transparent to the end user.

Calling Line ID

- The calling line ID presentation can be enabled or disabled on the gateway and also can be manipulated in the route pattern, based on site requirements.
- If you select the option Use Calling Party's External Phone Number Mask, then the external call uses the calling line ID specified for the IP phone placing the call. If you do not select this option, then the mask placed in the Calling Party Transform Mask field is used to generate the calling party ID.

Urgent Priority

- The Urgent Priority checkbox is often used to force immediate routing of certain calls as soon as a match is detected, without waiting for the T302 timer to expire. For example, in North America, if the patterns 9.911 and 9.[2-9]XXXXXX are configured and a user dials 9911, Cisco Unified CallManager has to wait for the T302 timer before routing the call because further digits may cause the 9.[2-9]XXXXXX to match. If Urgent Priority is enabled for the 9.911 route pattern, Cisco Unified CallManager makes its routing decision as soon as the user has finished dialing 9911, without waiting for the T302 timer.
- It is important to note that the Urgent Priority checkbox forces the T302 timer to expire as soon as the configured pattern is the best match for the dialed number. This does not mean that the urgent pattern has a higher priority than other patterns; the closest-match logic described in the section on [Call Routing in Cisco Unified CallManager, page 10-9](#), still applies.

For example, assume the route pattern 1XX is configured as urgent and the pattern 12! is configured as a regular route pattern. If a user dials 123, Cisco Unified CallManager will not make its routing decision as soon as it receives the third digit because even though 1XX is an urgent pattern, it is not the best match (10 total patterns matched by 12! versus 100 patterns matched by 1XX). Cisco Unified CallManager will have to wait for inter-digit timeout before routing the call because the pattern 12! allows for more digits to be input by the user.

Consider another example, where pattern 12[2-5] is marked as urgent and 12! is configured as a regular pattern. If the user dials 123, the pattern 12[2-5] is the best match (4 total patterns matched by 12[2-5] versus 10 patterns matched by 12!). Because the urgent-priority pattern is the best match, the T302 timer is aborted and no further user input is expected. Cisco Unified CallManager routes the call using pattern 12[2-5].

Call Classification

- Calls using this route pattern can be classified as on-net or off-net calls. This route pattern can be used to prevent toll fraud by prohibiting off-net to off-net call transfers or by tearing down a conference bridge when no on-net parties are present. (Both of these features are controlled by Service Parameters within Cisco Unified CallManager Administration.)
- When the "Allow device override" box is enabled, the calls are classified based on the call classification settings on the associated gateway or trunk.

Forced Account Codes (FAC)

- The Forced Account Codes (FAC) checkbox is used to restrict the outgoing calls when using a particular route pattern. If you enable FAC through route patterns, users must enter an authorization code to reach the intended recipient of the call.
- When a user dials a number that is routed through a FAC-enabled route pattern, the system plays a tone that prompts for the authorization code. To complete the call, the user authorization code must meet or exceed the level of authorization that is specified to route the dialed number.
- Only the authorization name appears in the call detail records (CDR); the authorization code does not appear in the CDR.
- The FAC feature is not available if the "Allow overlap sending" checkbox is enabled.

Client Matter Codes (CMC)

- The Client Matter Code (CMC) checkbox is used to track calls to certain numbers when using a particular route pattern. (For example, a company can use it to track calls to certain clients.)
- If you enable CMC for a route pattern, users must enter a code to reach the intended destination.
- When a user dials a number that is routed through a CMC-enabled route pattern, the system plays a tone that prompts for the code. The user must enter the correct code in order to complete the call.
- Client Matter Codes appear in the call detail records so that they can be used by the CDR analysis and reporting tool to generate reports for client billing and accounting.
- The CMC feature is not available if the "Allow overlap sending" checkbox is enabled.
- If both CMC and FAC are enabled, the user dials a number, enters the FAC when prompted to do so, and then enters the CMC at the next prompt.

Route Lists

A route list is a prioritized list of eligible paths (route groups) for an outbound call. Typically, a route list is associated with a remote location, and multiple route patterns may point to it. A typical use of a route list is to specify two paths for a remote destination, where the first-choice path is across the IP WAN and the second-choice path is through the local PSTN gateways.

Route lists have the following characteristics:

- Multiple route patterns may point to the same route list.
- A route list is a prioritized list of route groups that function as alternate paths to a given destination. For example, you can use a route list to provide least-cost routing, where the primary route group in the list offers a lower cost per call and the secondary route group is used only if the primary is unavailable due to an "all trunks busy" condition or insufficient IP WAN resources.
- Each route group in the route list can have its own digit manipulation. For example, if the route pattern is 9.@ and a user dials 9 1 408 555 4000, the IP WAN route group can strip off the 9 1 while the PSTN route group may strip off just the 9.
- Multiple route lists can contain the same route group. The digit manipulation for the route group is associated with the specific route list that points to the route group.
- If you are performing several digit manipulations in a route pattern or a route group, the order in which the transformations are performed can impact the resulting E.164 address. Cisco Unified CallManager performs the following major types of digit manipulations in the order indicated:

1. Discarding digits
2. Called party transformations
3. Prefixing digits

Route Groups

Route groups control and point to specific devices, which are typically gateways (MGCP or H.323), H.323 trunks to a gatekeeper or remote Cisco Unified CallManager cluster, or SIP trunks to a SIP proxy. (In Cisco CallManager Release 3.2 and earlier, the role of the H.323 trunk was performed by the Anonymous Device gateway and by H.323 gateways configured using the Intercluster Trunk protocol.)

Cisco Unified CallManager sends calls to the devices according to the distribution algorithm assigned. Cisco Unified CallManager supports top-down and circular algorithms.

Route Group Devices

The route group devices are the endpoints accessed by route groups, and they typically consist of gateways or H.323 trunks to a gatekeeper or to remote Cisco Unified CallManagers. You can configure the following types of devices in Cisco Unified CallManager:

- Media Gateway Control Protocol (MGCP) gateways
- H.323 gateways
- H.225 trunk, gatekeeper controlled — trunk to standard H.323 gateways, via a gatekeeper
- Intercluster trunk, not gatekeeper controlled — direct trunk to another Cisco Unified CallManager cluster
- Intercluster trunk, gatekeeper controlled — trunk to other Cisco Unified CallManager clusters and/or H.323 gateways, via a gatekeeper
- SIP trunk — trunk to a SIP proxy (available with Cisco Unified CallManager Release 4.0 or later)

**Note**

Both the H.225 and intercluster trunk (gatekeeper controlled) will automatically discover if the other endpoint is a standard H.323 gateway or a Cisco Unified CallManager and will select H.225 or Intercluster Trunk protocol accordingly.


Calling Privileges in Cisco Unified CallManager

Dialing privileges are configured in order to control which types of calls are allowed (or prevented) for a particular endpoint (such as phones, gateways, or CTI applications). All calls handled by Cisco Unified CallManager are subjected to the dialing privileges implemented through the configuration of the following elements:

- [Partitions, page 10-16](#)
- [Calling Search Spaces, page 10-17](#)

A *partition* is a group of directory numbers (DNs) with similar accessibility, and a *calling search space* defines which partitions are accessible to a particular device. A device can call only those DNs located in the partitions that are part of its calling search space.


As illustrated in [Figure 10-4](#), items that can be placed in partitions all have a dialable pattern, and they include phone lines, route patterns, translation patterns, CTI route group lines, CTI port lines, voicemail ports, and Meet-Me conference numbers. Conversely, items that have a calling search space are all devices capable of dialing a call, such as phones, phone lines, gateways, and applications (via their CTI route groups or voicemail ports).

Figure 10-4 Partitions and Calling Search Spaces

Partitions

The dial plan entries that you may place in a partition include IP phone directory numbers, translation patterns, route patterns, CTI route points, and voicemail ports. As described in the section on [Call Routing in Cisco Unified CallManager, page 10-9](#), if two or more dial plan entries (directory numbers, route patterns, or so forth) overlap, Cisco Unified CallManager selects the entry with the closest match (most specific match) to the dialed number. In cases where two dial plan entries match the dialed pattern equally, Cisco Unified CallManager selects the dial plan entry that appears first in the calling search space of the device making the call.

For example, consider [Figure 10-5](#), where route patterns 1XXX and 23XX are part of Partition_A and route patterns 12XX and 23XX are part of Partition_B. The calling search space of the calling device lists the partitions in the order Partition_A:Partition_B. If the user of this device dials 2345, Cisco Unified CallManager selects route pattern 23XX in Partition_A as the matching entry because it appears first in the calling device's calling search space. However, if the user dials 1234, Cisco Unified CallManager selects route pattern 12XX in Partition_B as the matching entry because it is a closer match than 1XXX in Partition_A. Remember that the partition order in a calling search space is used exclusively as a tie-breaker in case of equal matches based on the closest-match logic.

Figure 10-5 *Impact of Partition Order on the Matching Logic***Note**


When multiple equal-precision matches occur in the same partition, Cisco Unified CallManager selects the entry that is listed first in its local dial plan database. Since you cannot configure the order in which the dial plan database lists dial plan entries, Cisco strongly recommends that you avoid any possibility of equal-precision matches coexisting within the same partition because the resulting dial plan logic is not predictable in such cases.

Beginning with Cisco Unified CallManager Release 4.1, partitions can be activated or deactivated based on the time and date. You can activate or deactivate partitions by first configuring time periods and schedules within Cisco Unified CallManager Administration and then assigning a specific time schedule to each partition. Outside of the times and days specified by the schedule, the partition is inactive, and all patterns contained within it are ignored by the Cisco Unified CallManager call routing engine. For more information on this feature, see [Time-of-Day Routing, page 10-34](#).

Calling Search Spaces

A calling search space defines which partitions are accessible to a particular device. Devices that are assigned a certain calling search space can access only the partitions listed in that calling search space. Attempts to dial a DN in a partition outside that calling search space will fail, and the caller will hear a busy signal.

If you configure a calling search space both on an IP phone line and on the device (phone) itself, Cisco Unified CallManager concatenates the two calling search spaces and places the line's calling search space in front of the device's calling search space, as shown in [Figure 10-6](#).

Figure 10-6 Concatenation of Line and Device Calling Search Spaces for IP Phones**Note**

In versions of Cisco Unified CallManager prior to Release 4.2, the device calling search space remains the same as the device is moved to different parts of the network. With Cisco Unified CallManager 4.2, the device calling search space can be determined dynamically based on where in the network the phone is physically located, as determined by the phone's IP address. See [Device Mobility](#), page 10-26, for more details.

If the same route pattern appears in two partitions, one contained in the line's calling search space and one contained in the device's calling search space, then according to the rules described in the section on [Partitions](#), page 10-16, Cisco Unified CallManager selects the route pattern listed first in the concatenated list of partitions (in this case, the route pattern associated with the line's calling search space).

For recommendations on how to set the line and device calling search spaces, refer to the sections on [Building Classes of Service for Cisco Unified CallManager](#), page 10-72, and [Classes of Service with the Line/Device Approach](#), page 10-76.

The maximum length of the combined calling search space (device plus line) is 1024 characters, including separator characters between each partition name. (For example, the string “partition_1:partition_2:partition_3” contains 35 characters.) Thus, the maximum number of partitions in a calling search space varies, depending on the length of the partition names. Also, because the calling search space clause combines the calling search space of the device and that of the line, the maximum character limit for an individual calling search space is 512 (half of the combined calling search space clause limit of 1024 characters).

Therefore, when you are creating partitions and calling search spaces, keep the names of partitions short relative to the number of partitions that you plan to include in a calling search space. For more details on configuring calling search spaces, refer to the *Cisco Unified CallManager Administration Guide*, available online at

<http://www.cisco.com>

Before you configure any partitions or calling search spaces, all DNs reside in a special partition named <None>, and all devices are assigned a calling search space also named <None>. When you create custom partitions and calling search spaces, any calling search space you create also contains the <None> partition, while the <None> calling search space contains *only* the <None> partition.

**Note**

Any dial plan entry left in the <None> partition is implicitly reachable by *any* device making a call. Therefore, to avoid unexpected results, Cisco strongly recommends that you do not leave dial plan entries in the <None> partition.

**Note**

Cisco strongly recommends that you do not leave any calling search space defined as <None>. Doing so can introduce dial plan behavior that is difficult to predict.

Call-Forward Calling Search Spaces

While the main calling search space configured on the line is concatenated with the device calling search space, the calling search spaces configured for the various types of call forward (Forward All, Forward Busy, Forward No Answer, Forward No Coverage, and Forward Unregistered) are standalone values not concatenated with any other calling search space.

Call Forward settings (except Forward All) can be configured separately for internal or external call types. For example, a user might want to have their phone Call Forward No Answer to voicemail for external callers but to a cell phone number if the caller is a co-worker calling from another Cisco Unified CallManager phone on the network. You can implement this behavior by using different configurations for the Internal and External Call Forward settings.

When Forward All is initiated from an IP phone running SCCP, user input is simultaneously compared to the patterns allowed in the configured Forward All calling search space. If an invalid destination is entered by the phone user, reorder tone is heard.

When the Forward All calling search space is left as <None>, the results are difficult to predict and depend on the Cisco Unified CallManager release. Therefore, Cisco recommends the following best practices when configuring call-forward calling search spaces:

- Always provision the Forward All, Forward Busy, Forward No Answer, Forward No Coverage, and Forward Unregistered calling search spaces with a value other than <None>. This practice avoids confusion and facilitates troubleshooting because it enables the network administrator to know exactly which calling search space is being used for forwarded calls.
- Configure the Call Forward Busy and Call Forward No Answer calling search spaces with values that allow them to reach the DN's for the voicemail pilot and voicemail ports but not external PSTN numbers.
- Configure the Call Forward All calling search space according to your company's policy. Many companies choose to restrict forwarded calls to internal numbers only, to prevent users from forwarding their IP phone lines to a long-distance number and dialing their local IP phone number from the PSTN to bypass long-distance toll charges on personal calls.

Cisco Unified CallManager 4.2 introduces the Call Forward Unregistered (CFUR) feature as a way to reroute calls placed to a temporarily unregistered destination phone. The configuration of CFUR consists of two main elements:

- Destination selection

When the DN is unregistered, calls can be rerouted to either of the following destinations:

- Voicemail

Calls can be sent to voicemail by selecting the voicemail checkbox and configuring the CFUR calling search space to contain the partition of the voicemail pilot number.

- A directory number used to reach the phone through the PSTN

This approach is preferred when a phone is located within a site whose WAN link is down. If the site is equipped with Survivable Remote Site Telephony, the phone (and its co-located PSTN gateway) will re-register with the co-located SRST router. The phone is then able to receive calls placed to its PSTN DID number.

In this case, the appropriate CFUR destination is the corresponding PSTN DID number of the original destination DN. Configure this PSTN DID in the destination field, along with applicable access codes and prefixes (for example, 9 1 415 555 1234).

- Calling search space

Cisco Unified CallManager attempts to route the call to the configured destination number by using the called DN's CFUR calling search space. The CFUR calling search space is configured on the target phone and is used by all devices calling the unregistered phone. This means that all calling devices will use the same combination of route pattern, route list, route group, and gateway to place the call, and that all CFUR calls to a given unregistered device will be routed through the same unique gateway regardless of where the calling phone is located. Cisco recommends that you select a centralized gateway as the egress point to the PSTN for CFUR calls and that you configure the CFUR calling search space to route calls to the CFUR destination to this centralized gateway.

The Call Forward Unregistered functionality can result in telephony routing loops if a phone is unregistered while the gateway associated with the phone's DID number is still under control of Cisco Unified CallManager, as is the case if a phone is simply disconnected from the network. In such a case, the initial call to the phone would prompt the system to attempt a first CFUR call to the phone's DID through the PSTN. The resulting incoming PSTN call would in turn trigger another CFUR attempt to reach the same phone's DN, triggering yet another CFUR call from the central PSTN gateway through the PSTN. This cycle could repeat itself until system resources are exhausted.

The new service parameter `MaximumForwardUnRegisteredHopsToDn` controls the maximum number of CFUR calls that are allowed for a DN at the same time. The default value of 0 means the counter is disabled. If any DNs are configured to reroute CFUR calls through the PSTN, loop prevention is required. Configuring this service parameter to a value of 1 would stop CFUR attempts as soon as a single call is placed through the CFUR mechanism. This setting would also allow only one call to be forwarded to voicemail, if CFUR is so configured. Configuring this service parameter to a value of 2 would allow for up to two simultaneous callers to reach the voicemail of a DN whose CFUR setting is configured for voicemail, while also limiting potential loops to two for DNs whose CFUR configuration sends calls through the PSTN.

**Note**

Extension Mobility DNs should not be configured to send Call Forward Unregistered calls to the PSTN DID associated with the DN. The DNs of Extension Mobility profiles in the logged-out state are deemed to be unregistered, therefore any calls to the PSTN DID number of a logged-out DN would trigger a routing loop. To ensure that calls made to Extension Mobility DNs in the logged-out state are sent to voicemail, ensure that their corresponding Call Forward Unregistered parameters are configured to send calls to voicemail.

Digit Manipulation in Cisco Unified CallManager

The following tools provide digit manipulation capability in Cisco Unified CallManager:


- The external route construct (route patterns, route lists, and route groups)
- Translation patterns

The external route construct can be used to introduce some digit manipulation while routing calls to external devices, and it is described in the section on [Call Routing in Cisco Unified CallManager](#), page 10-9.

Translation patterns are the most powerful tool in Cisco Unified CallManager to manipulate digits for any type of call. They follow the same general rules and use the same wildcards as route patterns. As with route patterns, you assign a translation pattern to a partition. However, when the dialed digits match the translation pattern, Cisco Unified CallManager does not route the call to an outside entity such as a gateway; instead, it performs the translation first and then routes the call again, this time using the calling search space configured within the translation pattern.

Translation patterns can be used for a variety of applications, as shown by the example in [Figure 10-7](#).

Figure 10-7 Application Example for Translation Patterns



In this example, the administrator wishes to provide users with an operator service that is reached by dialing 0, while also maintaining a fixed-length internal numbering plan. The IP phones are configured with the Phone_css calling search space, which contains the Translations_pt partition (among others). A translation pattern 0 is defined in this partition, and the configured Called Party Transform Mask instructs Cisco Unified CallManager to replace the dialed string (0) with the new string 2001, which corresponds to the DN of the operator phone. A second lookup (of 2001 this time) is forced through the call routing engine, using the Internal_css calling search space, and the call can now be extended to the real operator DN of 2001, which resides in the AllPhones_pt partition.



Note

When a dialed number is manipulated using a translation pattern, the translated number is recorded in the call detail record (CDR). However, when the digit manipulation occurs within a route list, the CDR will show the originally dialed number, not the translated one. The Placed Calls directory on the IP phone always shows the string as it was dialed by the user.

Automated Alternate Routing

The automated alternate routing (AAR) feature enables Cisco Unified CallManager to establish an alternate path for the voice media when the preferred path between two endpoints within the same cluster runs out of available bandwidth, as determined by the locations mechanism for call admission control.

The AAR feature applies primarily to centralized call processing deployments. For instance, if a phone in branch A calls a phone in branch B and the available bandwidth for the WAN link between the branches is insufficient (as computed by the locations mechanism), AAR can reroute the call through the PSTN. The audio path of the call would be IP-based from the calling phone to its local (branch A) PSTN gateway, TDM-based from that gateway through the PSTN to the branch B gateway, and IP-based from the branch B gateway to the destination IP phone.

AAR can be transparent to the users. You can configure AAR so that users dial only the on-net (for example, four-digit) directory number of the called phone and no additional user input is required to reach the destination through the alternate network (such as the PSTN).

**Note**

AAR does not support CTI route points as the origin or the destination of calls. Also, AAR is incompatible with the Extension Mobility feature when users roam across different sites. Refer to [Extension Mobility, page 10-27](#), for more details.

You must provide the following main elements for AAR to function properly:

- [Establish the Phone Number of the Destination, page 10-22](#)
- [Prefix the Required Access Codes, page 10-23](#)
- [Select the Proper Dial Plan and Route, page 10-24](#)

**Note**

Beginning with Cisco Unified CallManager Release 4.1.3, automated alternate routing (AAR) can be applied to voicemail hunt group members.

Establish the Phone Number of the Destination

The rerouting of calls requires using a destination directory number (DN) that can be routed through the alternate network (for example, the PSTN). Prior to Cisco Unified CallManager 4.2, AAR uses the dialed digits to establish the on-cluster destination of the call and then combines them with the called party's External Phone Number Mask. The combination of these two elements must yield a fully qualified number that can be routed by the alternate network.

In Cisco Unified CallManager 4.2, two new features are available to determine the AAR destination:

- The AAR configuration allows for calls to be directed to the voicemail pilot number if you select the voicemail checkbox. This choice does not rely on the numbers originally dialed by the caller, but routes the call according to the Voice Mail profile configuration.

**Note**

By default, the directory number configuration retains the AAR leg of the call in the call history, which ensures that the AAR forward to the voice messaging system will select the proper voice mailbox. If you choose **Remove this destination from the call forwarding history**, the AAR leg of the call is not present in the call history, which would prevent the automated voice mailbox selection and would offer the caller the generic voicemail greeting.

- The AAR Destination Mask allows the destination phone number to be determined independently of the External Phone Number mask. For example, if the Caller ID policy for a company required a phone's external phone number mask to be the main directory number of an office (for example, 415 555 1000), the AAR destination mask could be set to 415 555 1234 to provide AAR with the phone's specific PSTN number.

For example, assume phone A in San Francisco (DN = 2345) dials an on-net DN (1234) configured on phone B located in New York. If locations-based call admission control denies the call, AAR retrieves the External Phone Number Mask of the New York phone (212555XXXX) and uses it to derive the fully qualified number (2125551234) that can be routed on the PSTN.

The PSTN routing of a call from San Francisco to New York requires a "1" as a prefix to the phone number. Cisco recommends that you do not include this prefix as part of the External Phone Number Mask of phones because it would be displayed as part of the Calling Party Identification (CallerID) for any calls made by the phones to an off-net destination. Instead, Cisco recommends that you add the "1" as part of the AAR group configuration.

For deployments that cover multiple countries within the same Cisco Unified CallManager cluster, keep in mind that the external phone number mask needs to be configured in such a way that the destination phone can be reached from the same country or from a different country by simply prefixing digits. This means that any in-country prefixes (such as 0 in many countries) should not be included in the external phone number mask if they are not part of the E.164 address.

In order to better understand this, consider an example where a Cisco Unified CallManager cluster has a site in London (United Kingdom), one in Paris (France), and one in Nice (France). The E.164 address of the DID range in Paris is +33145678XXX, but these extensions are usually reached as 0145678XXX when calling from within the French PSTN.

When somebody in the London office wishes to dial the Paris office via the PSTN, the dialed string is 90033145678XXX. However, when somebody in the Nice office wishes to dial the Paris office via the PSTN, the dialed string is 00145678XXX. Therefore, in this case the external phone number mask for the Paris office phones must be set to 145678XXX and not to the usual French national number of 0145678XXX because, if the 0 were included in the mask, it would not be possible to obtain the string 90033145678XXX by simply prefixing additional digits.

Prefix the Required Access Codes

The destination number might require a prefix for an off-net access code (for example, 9) to be routed properly by the origination branch's dial plan. Furthermore, if the point of origin is located in a different area code (also known as Numbering Plan Area, or NPA), then a prefix of "1" might be required as part of the dialed string. When configuring AAR, you place the DNs in AAR groups. For each pair of AAR groups, you can then configure prefix digits to add to the DNs for calls between the two groups, including prefix digits for calls originating and terminating within the same AAR group.

As a general rule, place DNs in the same AAR group if they share all of the following characteristics:

- A common off-net access code (for example, 9)
- A common PSTN dialing structure for inter-area calls (for example, 1-NPA-NXX-XXXX in North America)
- A common External Phone Number Mask format

For example, assume that both the San Francisco and New York sites share all of the preceding characteristics. We could place the DNs for San Francisco and New York into a single AAR group (for example, US) and configure the group such that AAR calls placed within this same AAR group are

prefixed with 91. For phone A in San Francisco to reach phone B in New York (at 212 555 1234), the AAR group configuration prefixes 91 to the dialed string, yielding a completed string of 91 212 555 1234.

For deployments that cover multiple countries, you will typically need at least one AAR group per country. Considering the example introduced in the previous section, two AAR groups could be defined: the UK AAR group (assigned to all DN in London) and the France AAR group (assigned to all DN in Paris and Nice). The UK AAR group configuration prefixes 90033 for calls to the France AAR group, while the France AAR group simply prefixes 00 for calls within the same AAR group.

Voicemail Considerations

With Cisco Unified CallManager 4.2, AAR can direct calls to voicemail. The voicemail pilot number is usually dialed without the need for an off-net access code (if the voicemail pilot number is a fully qualified on-net number, such as 8 555 1000). When AAR is configured to send calls to voicemail, the AAR group mechanism will still prefix the configured access code(s). This configuration requires the creation of an AAR group to be used by all DN whose desired AAR destination is voicemail (for example, `vmail_aar_grp`). Ensure that the configuration for this voicemail AAR group uses no prefix numbers when receiving calls from other AAR group DNs.

For example: Assume that DNs located in sites San Francisco and New York are configured with AAR group US, which prefixes 91 to calls made between any two DNs in the group. If a DN in San Francisco is configured to send AAR calls to voicemail (for example, 8 555 1000), a call would be placed to 9185551000, which would result in a failed call. Instead, the San Francisco DN should be configured with AAR group `vmail_aar_grp`. The prefix digits for calls from AAR group US to AAR group `vmail_aar_grp` are `<none>`, as shown in the following table. The call will be placed successfully to 85551000.

AAR Group		Destination	
		US	<code>vmail_aar_grp</code>
Source	US	91	<code><none></code>
	<code>vmail_aar_grp</code>	91	<code><none></code>



Note

In versions of Cisco Unified CallManager prior to Release 4.2, the AAR group configuration of a DN remains the same as the device is moved to different parts of the network. With Cisco Unified CallManager 4.2, the AAR group can be determined dynamically based on where in the network the phone is physically located, as determined by the phone's IP address. See [Device Mobility, page 10-26](#), for more details.

Select the Proper Dial Plan and Route

AAR calls should egress through a gateway within the same location as the calling phone, thus causing the completed dial string to be sent through the origination site's dial plan. To ensure that this is the case, select the appropriate AAR calling search space on the device configuration page in Cisco Unified CallManager Administration. Configure the off-net dial plan entries (for example, route patterns) in the AAR calling search space to point to co-located gateways and to remove the access code before presenting the call to the PSTN.

For example, phones at the San Francisco site can be configured with an AAR calling search space that permits long distance calls dialed as 91-NPA-NXX-XXXX but that delivers them to the San Francisco gateway with the access code (9) stripped.

**Note**

If you have configured additional route patterns to force on-net internal calls dialed as PSTN calls, ensure that these patterns are not matched by the AAR feature. Refer to [Design Guidelines for Multisite Deployments, page 10-52](#), for more details.

**Note**

To avoid denial of re-routed calls due to call admission control, AAR functionality requires the use of a LAN as the IP path between each endpoint and its associated gateway to the PSTN. Therefore, AAR dial plans cannot rely on centralized gateways for PSTN access.


**Note**

In versions of Cisco Unified CallManager prior to Release 4.2, the AAR calling search space configuration of a device remains the same as the device is moved to different parts of the network. With Cisco Unified CallManager 4.2, the AAR calling search space can be determined dynamically based on where in the network the phone is physically located, as determined by the phone's IP address. See [Device Mobility, page 10-26](#), for more details.

Special Considerations for Sites Located Within the Same Local Dialing Area

In some instances, the AAR dial string might have to be modified locally to allow for local area dialing. For example, assume two separate sites located in New York share the same area code of 212. (See [Figure 10-8](#).) In this case, a number dialed as 91 212 555 1234 would have to be transformed to 9 555 1234.

This transformation is best done with a site-specific translation pattern of 91212.555XXXX (to strip the pre-dot digits and prepend 9). This translation pattern should be placed in a member partition of the AAR calling search space for the New York sites only; the San Francisco site still needs to reach this same destination as 91 212 555 1234. This translation pattern should also be placed in the New York sites' dial plan to provide proper routing of locally reachable numbers dialed as long distance calls. The New York sites' dial plan is responsible for accepting 9 555 1234 as a valid string and transforming it to 555 1234 before delivering the call to the PSTN.

Figure 10-8 Dialed Number Transformations for AAR Calls Between Sites**Note**

The AAR functionality is not triggered upon detection that the destination phone is unreachable. Therefore, WAN failures do not trigger AAR functionality.

Device Mobility

Cisco Unified CallManager 4.2 introduces new functionality designed to enhance the mobility of devices within an IP network (for example, a phone initially configured for use in San Francisco is physically moved to New York). Although the device still registers with the same Cisco Unified CallManager cluster, it now will adapt some of its behavior based on the new site where it is located. Those changes are triggered by the IP subnet in which the phone is located.

From a dial-plan perspective, the functionality of the following four main configuration parameters can be modified due to the physical location of the phone. For these parameters to be modified, the device must be deemed as roaming outside its home physical location but within its home device mobility group.

- Device calling search space

The roaming device pool's Device Mobility calling search space is used instead of the device calling search space configured on the device's configuration page. For example, if a device is roaming from San Francisco to New York, the Device Mobility calling search space of the New York device pool is used as the roaming phone's device calling search space. If you use the line/device approach to classes of service, this approach will establish the path taken for PSTN calls, routing them to the local New York gateway.

- AAR calling search space

The roaming device pool's AAR calling search space is used instead of the AAR calling search space configured on the device's configuration page. For example, if a device is roaming from San Francisco to New York, the AAR calling search space of the New York device pool is used as the roaming phone's AAR calling search space. This calling search space will establish the path taken for outgoing AAR PSTN calls, routing them to the local New York gateway.

- DN's AAR group

For incoming AAR calls, the AAR group assigned to a DN is retained, whether or not the DN's host phone is roaming. This ensures that the reachability characteristics established for the AAR destination number are retained.

For outgoing AAR calls, the calling DN's AAR group uses the roaming device pool's AAR group instead of the AAR group selected on the DN's configuration page. Note that this AAR group will be applied to all DNs on the roaming device. For example, all DNs on a device roaming from New York to Paris (assuming both locations are in the same Device Mobility group) would inherit the AAR group configured for outgoing calls in the Paris device pool. This AAR group would be applied to all DNs on the roaming device and would allow for the appropriate prepending of prefix digits to AAR calls made from DNs on the roaming phone.

The chapter on [Device Mobility, page 22-1](#), explains the details of this feature.


Extension Mobility

The Extension Mobility feature enables a user to log in to an IP phone and automatically apply his or her profile to that phone, including extension number, speed dials, message waiting indicator (MWI) status, and calling privileges. This mechanism relies on the creation of a device profile associated with each Extension Mobility user. The device profile is effectively a virtual IP phone on which you can configure one or more lines and define calling privileges, speed dials, and so on.

When an IP phone is in the logged-out state, (that is, no Extension Mobility user has logged into it), the phone characteristics are determined by the device configuration page and the line configuration page(s). When a user logs in to an IP phone, the device configuration does not change, but the existing line configuration is saved in the Cisco Unified CallManager database and is replaced by the line configuration of the user's device profile.

One of the key benefits of Extension Mobility is that users can be reached at their own extensions regardless of where they are located, provided that they can log in to an IP phone controlled by the same Cisco Unified CallManager cluster. When Extension Mobility is applied to multisite deployments with centralized call processing, this capability is extended to multiple sites geographically separated from each other.

However, if you combine the Extension Mobility feature with the AAR feature described in the section on [Automated Alternate Routing, page 10-22](#), some limitations exist. Consider the example shown in [Figure 10-9](#), where Extension Mobility and AAR are deployed in a centralized call processing Cisco Unified CallManager cluster with one site in San Jose and one in New York.

Figure 10-9 Extension Mobility and AAR

In this example, assume that an Extension Mobility user who is normally based in San Jose has a DN of 1000 and a DID number of (408) 555-1000. That user's external phone number mask (or the AAR mask in Cisco Unified CallManager 4.2) is therefore configured as 4085551000. The user now moves to the New York site and logs in. Also, assume that the IP WAN bandwidth between San Jose and New York has been entirely utilized.

When the user in San Jose with extension 1001 tries to call 1000, AAR is triggered and, based on the AAR calling search space of the calling party and the AAR groups of both parties, a new call to 914085551000 is attempted by the San Jose phone. This call uses the San Jose gateway to access the PSTN, but because the DID (408) 555-1000 is owned by that same gateway, the PSTN sends the call back to it. The San Jose gateway tries to complete the call to the phone with extension 1000, which is now in New York. Because no bandwidth is available to New York, the AAR feature is invoked again, and one of the following two scenarios will occur:

- If the gateway's AAR calling search space contains external PSTN route patterns, this is the beginning of a loop that eventually uses all the PSTN trunks at the San Jose site.
- If, on the other hand, the gateway's AAR calling search space contains only internal numbers, the call fails and the caller hears a fast-busy tone. In this case, one PSTN call is placed and one is received, so two PSTN trunks are utilized on the San Jose gateway for the duration of the call setup.

**Tip**

To prevent routing loops such as the one described here, always configure all calling search spaces on the gateway configuration pages to include only internal destinations and no route patterns pointing to route lists or route groups containing that same gateway.

This example highlights the fact that Extension Mobility leverages the dynamic aspect of Cisco IP Communications and, therefore, requires that the call routing between sites use the IP network. Because the E.164 numbers defined in the PSTN are static and the PSTN network is unaware of the movements of the Extension Mobility users, the AAR feature, which relies on the PSTN for call routing, cannot be used to reach Extension Mobility users who move to a site other than their home site.

**Note**

However, if the Extension Mobility user moves to a remote site that belongs to the same AAR group as his or her home site, he or she can use the AAR feature to place calls to other sites when the available IP WAN bandwidth is not sufficient. This is because the path of such a call is determined by the AAR calling search space of the phone from which the call originates. This AAR calling search space does not change when users log in/out of Extension Mobility, and it should be configured to use the visited remote site's gateway.

**Tip**

Configure unregistered Extension Mobility profile DNs to send calls to voicemail. See [Call-Forward Calling Search Spaces](#), page 10-19, for details.

Hunt Lists and Line Groups

The *hunt pilot* is typically used for call coverage, or distributing a call through a list of phones. For call distribution, you can use a hunt construct. This hunt construct is based upon a three-tiered architecture, similar to that used to route external calls, that allows for multiple layers of call routing as well as digit manipulation.

Cisco Unified CallManager searches for a configured hunt pilot that matches an incoming called number and uses it to select a corresponding hunt list, which is a prioritized list of the available paths for the call. These paths are known as *line groups*. [Figure 10-10](#) depicts the three-tiered architecture of the hunt construct in Cisco Unified CallManager Release 4.1 or later.

**Note**

In Cisco CallManager Release 3.3 and earlier, call coverage functionality was provided by hunt groups, which are controlled by the Telephony Call Dispatcher (TCD) service, also used by the Cisco Attendant Console. Cisco Unified CallManager Release 4.0 introduced hunt pilots, hunt lists, and line groups, but in that release the hunt pilot structure was combined with the route pattern structure and the hunt list was combined with the route list. Starting with Cisco Unified CallManager Release 4.1, these structures are independent. [Table 10-3](#) summarizes a feature comparison of the hunt lists and line groups in Cisco Unified CallManager Releases 4.0 and 4.1 and hunt groups using the Attendant Console in Cisco Unified CallManager Releases 3.3 and earlier.

Table 10-3 Comparison of Features Supported by Route/Hunt List, Hunt Pilot, and Hunt Groups

Feature	Hunt Groups in Cisco CallManager Release 3.3 and Earlier	Route/Hunt List in Cisco Unified CallManager Release 4.0	Hunt Pilot in Cisco Unified CallManager Release 4.1	Hunt Pilot in Cisco Unified CallManager Release 4.2
Skinny Client Control Protocol (SCCP) endpoints	Yes	Yes (line groups)	Yes	Yes
Gateways and trunks (off-net destinations)	No	Yes (route groups)	No	No
Top-down algorithm	Yes	Yes	Yes	Yes
Circular algorithm	Yes	Yes	Yes	Yes
Longest-idle algorithm	Yes	Yes	Yes	Yes
Broadcast algorithm	Yes	Yes	Yes	Yes

Table 10-3 Comparison of Features Supported by Route/Hunt List, Hunt Pilot, and Hunt Groups (continued)

Feature	Hunt Groups in Cisco CallManager Release 3.3 and Earlier	Route/Hunt List in Cisco Unified CallManager Release 4.0	Hunt Pilot in Cisco Unified CallManager Release 4.1	Hunt Pilot in Cisco Unified CallManager Release 4.2
Hunt options	No	Yes	Yes	Yes
Ring-no-answer reversion	No	Yes	Yes	Yes
Performance monitoring (PerfMon)	No	Yes	Yes	Yes
SCCP voicemail port (Cisco Unity)	No	Yes (line groups)	Yes	Yes
Simplified Message Desk Interface (SMDI) voicemail systems	Yes	Yes	Yes	Yes
Queuing	Yes	No	No	No
Linked hunt groups and hunt pilot	Yes	No	Yes	Yes
Log out of hunt groups	No	No	No	Yes

For comparison purposes, [Figure 10-10](#) illustrates the hunt pilot architecture in Cisco Unified CallManager 4.1 and above, while [Figure 10-11](#) shows the hunt pilot architecture in Cisco Unified CallManager 4.0.



Figure 10-10 Three-Tiered Architecture for the Hunt Construct in Cisco Unified CallManager Releases 4.1 and 4.2

Figure 10-11 Hunt Construct in Cisco Unified CallManager Release 4.0

Hunt Pilot

Hunt pilots are strings of digits and wildcards similar to route patterns, such as 9.[2-9]XXXXXX, configured in Cisco Unified CallManager to route calls to directory numbers. The hunt pilot points directly to a hunt list. Hunt lists point to line groups, which finally point to SCCP endpoints.

Beginning with Cisco Unified CallManager Release 4.1, calls can be redirected to a final destination when the hunting fails because of one or both of the following reasons:

- All hunting options have been exhausted and the call still is not answered.
- A time-out period has expired.

This call redirection is configured in the Hunt Forward Settings section of the Hunt Pilot configuration page, and the destination for this redirect can be either of the following options:

- A specific pattern in the internal call routing table of Cisco Unified CallManager
- Personal preferences, which point to the Call Forward No Coverage settings for the originally called number when hunting on behalf of that number fails

For example, you can implement the personal preferences option by configuring a user's phone so that the Forward No Answer field redirects the call to a hunt pilot, in order to search for someone else who can answer the call. If the call hunting fails, either because all the hunting options were exhausted or because a time-out period expired, the call can be sent to a destination personalized for the person who was originally called. For example, if you set the Forward No Coverage field within the person's DN configuration page to the voicemail number, the call will be sent to that person's voicemail box if hunting fails.



Note

Cisco Unified CallManager Release 4.0 does not support call redirection.

The following considerations apply to calls handled by hunt pilots:

- Call Pickup and Group Call Pickup are not supported on calls distributed by a hunt pilot. A member of the line group cannot pick up the hunt pilot call offered to another member in the line group, even if they belong to the same call pickup group.
- Calls distributed by the hunt pilot number do not support separate call-forward treatment configured on a directory number in a line group. Thus, the Immediate Divert (iDivert) softkey or any call-forward configurations on the directory number do not work for calls distributed by the hunt pilot. Only the call-forward conditions available as hunt options in a line group configuration apply to the hunt pilot call. However, the iDivert softkey or call-forward configurations work for all incoming calls other than those distributed by the hunt pilot.
- The hunt pilot can distribute calls to any of its line group members, even if the members and the hunt pilot reside in different partitions. A call distributed by the hunt pilot overrides all the partitions and calling search space restrictions.

Hunt List

A hunt list is a prioritized list of eligible paths (line groups) for call coverage. Hunt lists have the following characteristics:

- Multiple hunt pilots may point to the same hunt list.
- A hunt list is a prioritized list of line groups that function as alternate sets of phones which are offered a call placed to the hunt pilot number. For example, you can use a hunt list to attempt to find a taker for the call within a set of phones at a particular site. If the call is not taken, then the hunt list attempts to offer the call through a second line group pointing to phones at a second site.
- Hunt lists do not do any digit manipulation.
- Multiple hunt lists can contain the same line group.

Line Group

Line group members are user extension numbers that are controlled by Cisco Unified CallManager. Thus, when the call is being distributed through the line group members, Cisco Unified CallManager is in control of the call. Hunt options can be applied to the call when it is not answered or if the extension is busy or unregistered.

Line groups control the order in which the call is distributed, and they have the following characteristics:

- Line groups point to specific extensions, which are typically IP phone extensions or voicemail ports.
- A single extension may be present in multiple line groups.
- Computer Telephony Integration (CTI) ports and CTI route points cannot be added within a line group. Therefore, calls cannot be distributed to endpoints controlled through CTI applications such as Cisco Customer Response Solution (CRS), IP Interactive Voice Response (IP IVR), and so forth.
- Cisco Unified CallManager distributes calls to the devices according to the distribution algorithm assigned. Cisco Unified CallManager supports the following algorithms:
 - Top-down
 - Circular
 - Longest idle time
 - Broadcast

- In the event of No-Answer, Busy, or Not-Available, line groups redirect a call distributed to an extension based on the hunt options. Cisco Unified CallManager supports the following hunt options:
 - Try next member, then try next group in hunt list.
 - Try next member, but do not go to next group.
 - Skip remaining members and go directly to next group.
 - Stop hunting.

Hunt Group Logout

In Cisco Unified CallManager 4.2, a user can log out of a hunt group by activating the HLog softkey. Once activated, this function effectively makes all lines configured on the phone act as though they are not part of any hunt group. The phone displays "Logged out of Hunt Group." If a line group contains a shared line, all instances of the shared line that are on devices in the logged-out state will *not* ring; conversely, all instances of the shared line on devices in the logged-in state will ring.

Lines that are not part of any hunt groups will continue to ring normally, no matter the state of the HLog function.

The HLog function can be activated from Cisco Unified CallManager Administration. By default, the HLog softkey is not configured on the softkey templates. Once added to a phone's softkey template, the HLog button appears in the display when the phone is in the connected, off-hook, or on-hook state.

The Hunt Group Logoff Notification service parameter provides the option of audible ring tones when calls that come in from a line group arrive at the phone in the logged-off state. The Hunt Group Logoff Notification service parameter is in the Clusterwide Parameters (Device - Phone) section of the Service Parameters Configuration window. For enabling the function, specify a valid ring tone file (which must be present on the TFTP server under the C:\Program Files\Cisco\TFTPPath directory). If an invalid file name is provided, no tone will be played.

For more information about hunt algorithms and hunt options, refer to the *Cisco Unified CallManager Administration Guide*, available at

<http://www.cisco.com>

Line Group Devices

The line group devices are the endpoints accessed by line groups, and they can be of any of the following types:

- Any Skinny Client Control Protocol (SCCP) endpoints, such as Cisco Unified IP Phones, VG248, or ATA 188
- Voicemail ports (for Cisco Unity)
- H.323 clients
- FXS extensions attached to an MGCP gateway

Time-of-Day Routing


Cisco Unified CallManager Release 4.1 introduces the time-of-day (ToD) routing feature. To use this feature, configure the following elements:

- Time period
- Time schedule

The time period allows you to configure start and end times for business hours. The start and end times indicate the times during which the calls can be routed. In addition to these times, you can set the event to repeat itself on a weekly or yearly basis. Moreover, you can also configure non-business hours by selecting "No business hours" from the Start Time and End Time options. All incoming calls will be blocked when this option is selected.

A time schedule is a group of specific time periods assigned to the partition. It determines whether the partition is active or inactive during the specified time periods. A matching/dialing pattern can be reached only if the partition in which the dialing pattern resides is active.

As illustrated in [Figure 10-12](#), two hunt pilots with the same calling pattern (8000) are configured in two partitions (namely, RTP_Partition and SJC_Partition). Each of these partitions is assigned a time schedule, which contains a list of defined time periods. For example, RTP phones can be reached using Hunt Pilot 1 from 8:00 AM to 12:00 PM EST (GMT - 5.00) Monday through Friday as well as 8:00 AM to 5:00 PM on Sundays. In the same way, SJC phones can be reached using Hunt Pilot 2 from 8:00 AM to 5:00 PM PST (GMT - 8.00) Monday through Friday and 8:00 AM to 5:00 PM on Saturdays. Both of the hunt pilots in this example are inactive on July 4th.

Figure 10-12 Time-of-Day Routing

For the example in Figure 10-12, an incoming call to the hunt pilot (8000) on Wednesday at 3:00 PM will be forwarded to the SJC phones, while a person calling the hunt pilot on July 4th will get a fast busy tone unless there is another pattern that matches 8000.

Call Routing in Cisco IOS with H.323 Dial Peers

The call routing logic on Cisco IOS routers using the H.323 protocol relies on the dial peer construct. Dial peers are similar to static routes; they define where calls originate and terminate and what path the calls take through the network. Dial peers are used to identify call source and destination endpoints and to define the characteristics applied to each call leg in the call connection. Attributes within the dial peer determine which dialed digits the router collects and forwards to telephony devices.

For a detailed description of dial peers and their configuration, refer to *Configuring Dial Plans, Dial Peers, and Digit Manipulation*, part of the *Cisco IOS Voice, Video, and Fax Configuration Guide*, Release 12.2, available at

<http://www.cisco.com>

One of the keys to understanding call routing with dial peers is the concept of incoming versus outgoing call legs and, consequently, of incoming versus outgoing dial peers. Each call passing through the Cisco IOS router is considered to have two call legs, one entering the router and one exiting the router. The call leg entering the router is the *incoming call leg*, while the call leg exiting the router is the *outgoing call leg*.

Call legs can be of two main types:

- Traditional TDM telephony call legs, connecting the router to the PSTN, analog phones, or PBXs
- IP call legs, connecting the router to other gateways, gatekeepers, or Cisco Unified CallManagers


For all calls going through the router, Cisco IOS associates one dial peer to each call leg. Dial peers are also of two main types, according to the type of call leg with which they are associated:

- POTS dial peers, associated with traditional TDM telephony call legs
- VoIP dial peers, associated with IP call legs

Figure 10-13 shows the following examples of different types of calls going through a Cisco IOS router:

- Call 1 is from another H.323 gateway across the IP network to a traditional PBX connected to the router (for example, via a PRI interface). For this call, an incoming VoIP dial peer and an outgoing POTS dial peer are selected.
- Call 2 is from an analog phone connected to an FXS port on the router to a Cisco Unified CallManager cluster across the IP network. For this call, an incoming POTS dial peer and an outgoing VoIP dial peer are selected by the router.
- Call 3 is from an IP phone controlled by Cisco Unified CallManager Express or SRST to a PSTN interface on the router (for example, a PRI interface). For this call, an automatically generated POTS dial peer (corresponding to the **ephone** configured on the router) and an outgoing POTS dial peer are selected.

Figure 10-13 Incoming and Outgoing Dial Peers



To match incoming call legs to incoming dial peers, the router selects a dial peer by matching the information elements in the setup message (called number/DNIS and calling number/ANI) with four configurable dial peer attributes. The router attempts to match these items in the following order:

1. Called number with **incoming called-number**
2. Calling number with **answer-address**
3. Called number with **destination-pattern**
4. Incoming voice port with configured voice port

The router must match only one of these conditions. It is not necessary for all the attributes to be configured in the dial peer or that every attribute match the call setup information; only one condition must be met for the router to select a dial peer. The router stops searching as soon as one dial peer is matched, and the call is routed according to the configured dial peer attributes. Even if there are other dial peers that would match, only the first match is used.

How the router selects an outbound dial peer depends on whether **direct-inward-dial** (DID) is configured in the inbound POTS dial peer:

- If DID is not configured in the inbound POTS dial peer, the router performs two-stage dialing and collects the incoming dialed string digit-by-digit. As soon as one dial peer matches the destination pattern, the router immediately places the call using the configured attributes in the matching dial peer.
- If DID is configured in the inbound POTS dial peer, the router uses the full incoming dial string to match the destination pattern in the outbound dial peer. With DID, the setup message contains all the digits necessary to route the call, so no additional digit collection is required. If more than one dial peer matches the dial string, all of the matching dial peers are used to form a *hunt group*. The router attempts to place the outbound call leg using all of the dial peers in the hunt group until one is successful.

By default, dial peers in a hunt group are selected according to the following criteria, in the order listed:

1. Longest match in phone number

This method selects the destination pattern that matches the greatest number of dialed digits. For example, if one dial peer is configured with a dial string of 345.... and a second dial peer is configured with 3456789, the router would first select 3456789 because it has the longest explicit match of the two dial peers.

2. Explicit preference

This method uses the priority configured with the **preference** dial peer command. The lower the preference number, the higher the priority. The highest priority is given to the dial peer with preference order 0. If the same preference is defined in multiple dial peers with the same destination pattern, a dial peer is selected randomly.

3. Random selection

In this method, all destination patterns are weighted equally.

You can change this default selection order or choose different methods for hunting dial peers by using the **dial-peer hunt** global configuration command. An additional selection criterion is *least recent use*, which selects the destination pattern that has waited the longest since being selected (equivalent to *longest idle* for Cisco Unified CallManager line groups).

Observe the following best practices when configuring H.323 dial peers on a Cisco IOS router:

- To ensure that incoming PSTN calls are directly routed to their destination based on the DNIS information, create a default POTS dial peer with the **direct-inward-dial** attribute, as follows:
- When using the router as an H.323 gateway connected to a Cisco Unified CallManager cluster, provide redundancy by configuring at least two VoIP dial peers with the same destination pattern pointing to two different Cisco Unified CallManager servers. Use the **preference** attribute to select the priority order between primary and secondary Cisco Unified CallManager servers. The following example shows the use of the **preference** attribute:

```
dial-peer voice 999 pots
  incoming called-number .
  direct-inward-dial
  port 1/0:23
```

```
dial-peer voice 100 voip
  preference 1
```

```
!--- Make this the first choice dial peer.
```

```
ip precedence 5
destination-pattern 1...
session target ipv4:10.10.10.2
```

```
!--- This is the address of the primary Cisco Unified CallManager.
```

```
dtmf-relay h245-alpha
```

```
dial-peer voice 101 voip
  preference 2
```

```
!--- This is the second choice.
```

```
ip precedence 5
destination-pattern 1...
session target ipv4:10.10.10.3
```

```
!--- This is the address of the secondary Cisco Unified CallManager.
```

```
dtmf-relay h245-alpha
```

Call Routing in Cisco IOS with a Gatekeeper

An H.323 gatekeeper is an optional node that manages endpoints (such as H.323 terminals, gateways, and Multipoint Control Units (MCUs), as well as Cisco Unified CallManager Express and Cisco Unified CallManager clusters) in an H.323 network, providing them with call routing and call admission control functions. The endpoints communicate with the gatekeeper using the H.323 Registration Admission Status (RAS) protocol.

Endpoints attempt to register with a gatekeeper on startup. When they want to communicate with another endpoint, they request admission to initiate a call using a symbolic alias for the endpoint, such as an E.164 address or an email address. If the gatekeeper decides that the call can proceed, it returns a destination IP address to the originating endpoint. This IP address might not be the actual address of the destination endpoint, but instead it might be an intermediate address, such as the address of an IP-to-IP gateway or a gatekeeper that routes call signaling.

For more details about the H.323 protocol and the message exchange between H.323 endpoints and gatekeepers, refer to the *Cisco IOS H.323 Configuration Guide*, available at

<http://www.cisco.com>

The Cisco 2600, 3600, 3700, 2800, 3800, and 7200 Series routers all support the gatekeeper feature. You can configure Cisco IOS gatekeepers in a number of different ways for redundancy, load balancing, and hierarchical call routing. This section focuses on the call routing capabilities of the gatekeeper feature. For redundancy and scalability considerations, refer to [Gatekeeper Redundancy, page 8-18](#); for call admission control considerations, refer to [Cisco IOS Gatekeeper Zones, page 9-15](#).

Call routing in Cisco IOS gatekeeper is based on the following types of information:

- Statically configured information, such as zone prefixes and default technology prefixes
- Dynamic information, such as E.164 addresses and technology prefixes provided by H.323 devices during the registration phase
- Per-call information, such as called number and technology prefix

A zone is a collection of H.323 devices (such as endpoints, gateways, or MCUs) that register with a gatekeeper. There can be only one active gatekeeper per zone, and you can define up to 100 local zones on a single gatekeeper.

When an H.323 endpoint registers with the gatekeeper, it is assigned to a zone and it can optionally register one or more E.164 addresses for which it is responsible, as well as a technology prefix that specifies which kinds of calls it can handle (for example, voice, video, fax, and so on).

For each zone, you can configure one or more *zone prefixes* on the gatekeeper. Zone prefixes are strings that contain digits and wildcards and are used by the gatekeeper to facilitate call routing decisions. The following characters are allowed in a zone prefix string:

- All numbers between 0 and 9, which match a single specific digit
- The dot (.) wildcard, which matches one digit between 0 and 9
- The * wildcard, which matches one or more digits between 0 and 9

To understand the gatekeeper call routing behavior, it is helpful to consider the message parsing logic. [Figure 10-14](#) illustrates the parsing logic for an Admission Request (ARQ). To initiate a call, an endpoint sends an Admission Request (ARQ) to the gatekeeper. The ARQ contains either an H.323 ID or the E.164 address of the destination, or called party, as well as the E.164 address or H.323 ID of the source, or calling party.

If the ARQ contains the E.164 address (with Cisco Unified CallManager, the ARQ always contains an E.164 address), the ARQ may or may not contain a technology prefix. If the ARQ contains a technology prefix, the gatekeeper strips it from the called number. If the ARQ does not contain a technology prefix, the gatekeeper uses the default technology prefix if one is configured (see the **gw-type-prefix** command in the section on [Centralized Gatekeeper Configuration, page 10-42](#)). The technology prefix thus obtained is stored in memory, and the gatekeeper continues with the call routing algorithm.


Next, the gatekeeper tries to match the called number with one of the configured zone prefixes. Longest-match is used if multiple potential matches exist. If no zone prefix can be matched, and if the gatekeeper is configured to accept calls with an unknown prefix, the gatekeeper then assumes that the destination zone is equal to the source zone.

At this point, the gatekeeper searches in the chosen destination zone for a registered E.164 address that matches the called number. If there is a match, the gatekeeper can send an Admission Confirm (ACF), provided that the requested bandwidth for the call is available and that the called endpoint is registered with the gatekeeper. The ACF will contain the IP address of the destination endpoint. If the bandwidth is unavailable or the called endpoint is not registered, the gatekeeper returns an Admission Reject (ARJ) to the calling endpoint.

If there is no matching E.164 address registered in the destination zone, the gatekeeper will use the previously stored technology prefix to choose a gateway registered in that zone as the destination for the call. The same considerations regarding bandwidth availability and endpoint registration dictate whether the gatekeeper will send an ACF or an ARJ to the calling endpoint.


Upon receipt of an ACF from the gatekeeper, the source endpoint can send a setup message directly to the destination endpoint by using the IP address returned in the ACF.

Figure 10-14 Gatekeeper Address Resolution for an ARQ



90616

Figure 10-15 illustrates the parsing logic for a Location Request (LRQ). LRQ messages are exchanged between gatekeepers and are used for inter-zone (remote zone) calls. For example, gatekeeper A receives an ARQ from a local zone gateway requesting call admission for a remote zone device. Gatekeeper A then sends an LRQ message to gatekeeper B. Gatekeeper B replies to the LRQ message with either a Location Confirm (LCF) or Location Reject (LRJ) message, depending on whether it is configured to admit or reject the inter-zone call request and whether the requested resource is registered.

Figure 10-15 Gatekeeper Address Resolution for an LRQ

90617

Traditional Cisco IOS gatekeeper functionality has been extended to accommodate for IP-to-IP gateways through the concept of a *via-zone gatekeeper*. (Refer to [Cisco IOS Gatekeeper and IP-to-IP Gateway with RSVP, page 9-17](#), for some deployment examples.)

A *via-zone gatekeeper* differs from legacy gatekeepers in how it uses LRQ and ARQ messages for call routing. Using *via-zone gatekeepers* will maintain normal clusters and functionality. Legacy gatekeepers examine incoming LRQs based on the called number, and more specifically the `dialedDigits` field in the `destinationInfo` portion of the LRQ. *Via-zone gatekeepers* look at the origination point of the LRQ before looking at the called number. If an LRQ comes from a gatekeeper listed in the *via-zone gatekeeper's* remote zone configurations, the gatekeeper checks to see that the zone remote configuration contains an **invia** or **outvia** keyword. If the configuration contains either of these keywords, the gatekeeper uses the new *via-zone* behavior; if not, it uses legacy behavior.


For ARQ messages, the gatekeeper determines if an **outvia** keyword is configured on the destination zone. If the **outvia** keyword is configured and the zone named with the **outvia** keyword is local to the gatekeeper, the call is directed to an IP-IP gateway in that zone by returning an ACF pointing to the IP-IP

gateway. If the zone named with the **outvia** keyword is remote, the gatekeeper sends a location request to the **outvia** gatekeeper rather than the remote zone gatekeeper. The **invia** keyword is not used in processing the ARQ.

Centralized Gatekeeper Configuration

A single gatekeeper can support call routing between clusters and call admission control for up to 100 Cisco Unified CallManager clusters. [Figure 10-16](#) illustrates a distributed call processing environment with two Cisco Unified CallManager clusters and a single, centralized gatekeeper.

Figure 10-16 Centralized Gatekeeper Supporting Two Clusters



[Example 10-1](#) shows the gatekeeper configuration for the example in [Figure 10-16](#).

Example 10-1 Configuration for Centralized Gatekeeper

```
gatekeeper
zone local GK-Site1 customer.com 10.1.10.100
zone local GK-Site2 customer.com
zone prefix GK-Site1 408.....
zone prefix GK-Site2 212.....
bandwidth interzone GK-Site1 160
bandwidth interzone GK-Site2 160
gw-type-prefix 1#* default-technology
arq reject-unknown-prefix
no shutdown
```

The following notes also apply to [Figure 10-16](#):

- Each Cisco Unified CallManager cluster has a local zone configured to support Cisco Unified CallManager trunk registrations.
- A zone prefix is configured for each zone to allow inter-zone and inter-cluster call routing.

- Bandwidth statements are configured for each site. Cisco recommends that you use the **bandwidth interzone** command because using the **bandwidth total** command can cause issues in some configurations. Bandwidth is measured in kilobits per second (kbps).
- The **gw-type-prefix 1# default-technology** command allows all locally unresolved calls to be forwarded to a device registered with a technology prefix of 1# in the local zone. In this example, all Cisco Unified CallManager trunks have been configured to register with a 1# prefix.

Technology prefixes indicate the type of call being made. The specific values used as technology prefixes are arbitrary and are defined by the network administrator. The same values should be used consistently throughout the entire deployment.

Technology prefixes are sent as a prefix to the E.164 address (phone number) to indicate whether the call is voice, video, or some other type. The # symbol is generally used to separate the prefix from the E.164 number. If a prefix is not included, the default technology prefix is used to route the call. There can be only one default technology prefix for the entire deployment.


Cisco IOS gateways automatically add a technology prefix to outbound calls if the gateway has a prefix configured. The gateway also automatically strips the prefix from incoming H.323 calls. Cisco Unified CallManager can register with the gatekeeper using a technology prefix, as specified on the configuration page for gatekeeper-controlled H.323 trunks. However, this technology prefix is not automatically added to outgoing calls to the gatekeeper, and is not automatically stripped from inbound calls to Cisco Unified CallManager. You can use translation patterns and significant digits to manipulate the called number so as to add or strip the technology prefix as needed.

- The **arq reject-unknown-prefix** command guards against potential call routing loops across redundant Cisco Unified CallManager trunks.

Distributed Gatekeeper Configuration

Gatekeepers can be distributed to conserve bandwidth or to provide local call routing for H.323 gateways in case of a WAN failure. Figure 10-17 illustrates a distributed call processing environment with two clusters and two gatekeepers.

Figure 10-17 Distributed Gatekeepers Supporting Two Clusters



Example 10-2 shows the gatekeeper configuration for Site 1 in Figure 10-17.

Example 10-2 Gatekeeper Configuration for Site 1

```
gatekeeper
zone local GK-Site1 customer.com 10.1.10.100
zone remote GK-Site2 customer.com 10.1.11.100
zone prefix GK-Site1 408.....
zone prefix GK-Site2 212.....
bandwidth remote 160
gw-type-prefix 1#* default-technology
arq reject-unknown-prefix
no shutdown
```

The following notes apply to Example 10-2:

- A local zone is configured for registration of local Cisco Unified CallManager cluster trunks.
- A remote zone is configured for routing calls to the Site 2 gatekeeper.
- Zone prefixes are configured for both zones for inter-zone call routing.
- The **bandwidth remote** command is used to limit bandwidth between the local zone and any other remote zone.
- The **gw-type-prefix 1# default-technology** command allows all locally unresolved calls to be forwarded to a device registered with a technology prefix of 1# in the local zone. In this example, all Cisco Unified CallManager trunks have been configured to register with a 1# prefix.

- The **arq reject-unknown-prefix** command guards against potential call routing loops across redundant Cisco Unified CallManager trunks.

Example 10-3 shows the gatekeeper configuration for Site 2 in Figure 10-17.

Example 10-3 Gatekeeper Configuration for Site 2

```
gatekeeper
zone local GK-Site2 customer.com 10.1.11.100
zone remote GK-Site1 customer.com 10.1.10.100
zone prefix GK-Site2 212.....
zone prefix GK-Site1 408.....
bandwidth remote 160
gw-type-prefix 1#* default-technology
arq reject-unknown-prefix
no shutdown
```

The following notes apply to Example 10-3:


- A local zone is configured for registration of local Cisco Unified CallManager cluster trunks.
- A remote zone is configured for routing calls to the Site 1 gatekeeper.
- Zone prefixes are configured for both zones for inter-zone call routing.
- The **bandwidth remote** command is used to limit bandwidth between the local zone and any other remote zone.
- The **gw-type-prefix 1# default-technology** command allows all locally unresolved calls to be forwarded to a device registered with a technology prefix of 1# in the local zone. In this example, all Cisco Unified CallManager trunks have been configured to register with a 1# prefix.
- The **arq reject-unknown-prefix** command guards against potential call routing loops across redundant Cisco Unified CallManager trunks.

Distributed Gatekeeper Configuration with Directory Gatekeeper

Because there is no gatekeeper protocol available to update gatekeeper routing tables, use of a directory gatekeeper can help make distributed gatekeeper configurations more scalable and more manageable. Implementing a directory gatekeeper makes gatekeeper configurations at each site simpler and moves most of the configuration for inter-zone communication into the directory gatekeeper.

Without a directory gatekeeper, you would have to add an entry in every gatekeeper on the network every time you add a new zone on one of the gatekeepers. However, with a directory gatekeeper, you can add the new zone in the local gatekeeper and the directory gatekeeper only. If the local gatekeeper cannot resolve a call request locally, it forwards that request to the directory gatekeeper with a matching zone prefix.

Figure 10-18 illustrates a Cisco Unified CallManager distributed call processing environment with distributed gatekeepers for local call routing and a directory gatekeeper to provide call routing between gatekeepers.

Figure 10-18 Distributed Gatekeepers with a Directory Gatekeeper

Example 10-4 shows the gatekeeper configuration for Site 1 in Figure 10-18. Because the Site 1 and Site 2 gatekeeper configurations are almost identical in this example, only Site 1 is covered here.

Example 10-4 Gatekeeper Configuration for Site 1, with Directory Gatekeeper

```
gatekeeper
zone local GK-Sitel customer.com 10.1.10.100
zone remote DGK customer.com 10.1.10.101
zone prefix GK-Sitel 408.....
zone prefix DGK .....
bandwidth remote 160
gw-type-prefix 1# default-technology
arq reject-unknown-prefix
no shutdown
```

The following notes also apply to Example 10-4:

- A local zone is configured for registration of local Cisco Unified CallManager cluster trunks.
- A remote zone is configured for the directory gatekeeper.
- Zone prefixes are configured for both zones for inter-zone call routing.
- The directory gatekeeper zone prefix is configured with 10 dots. This pattern matches any unresolved 10-digit dial strings. Multiple zone prefixes can be configured for a single zone, allowing matches on different length dial strings. A wildcard (*) can also be used for a directory gatekeeper zone prefix, but this method can cause call routing issues in some instances.
- The **bandwidth remote** command is used to limit bandwidth between the local zone and any other remote zone.
- The **gw-type-prefix 1# default-technology** command allows all locally unresolved calls to be forwarded to a device registered with a technology prefix of 1# in the local zone. In this example, all Cisco Unified CallManager trunks have been configured to register with a 1# prefix.

- The **arq reject-unknown-prefix** command guards against potential call routing loops across redundant Cisco Unified CallManager trunks.

Example 10-5 shows the directory gatekeeper configuration for the example in Figure 10-18.

Example 10-5 Directory Gatekeeper Configuration

```
gatekeeper
zone local DGK customer.com 10.1.10.101
zone remote GK-Site1 customer.com 10.1.10.100
zone remote GK-Site2 customer.com 10.1.11.100
zone prefix GK-Site1 408*
zone prefix GK-Site2 212*
lrq forward-queries
no shutdown
```

The following notes also apply to Example 10-5:

- A local zone is configured for the directory gatekeeper.
- Remote zones are configured for each remote gatekeeper.
- Zone prefixes are configured for both remote zones for inter-zone call routing. The wildcard (*) is used in the zone prefix to simplify configuration. Calls will not be routed to the DGK zone, so no prefix is required for it.
- The **lrq forward-queries** command allows the directory gatekeeper to forward an LRQ received from another gatekeeper.

Calling Privileges in Cisco IOS with H.323 Dial Peers

Use the class of restriction (COR) functionality to implement calling privileges with Cisco IOS-based systems using H.323, including H.323 gateways, SRST, and Cisco Unified CallManager Express. This functionality provides flexibility in network design, allows administrators to block calls for all users (for example, calls to 900 numbers), and applies different calling privileges to call attempts from different originators (for example, it allows some users but not others to call international numbers).

The fundamental mechanism at the center of the COR functionality relies on the definition of incoming and outgoing *corlists* that are associated to existing dial peers, where the concepts of incoming and outgoing are relative to the Cisco IOS router (as in the case of dial peers). Each corlist is defined to include a number of members, which are simply tags previously defined within Cisco IOS.

When a call goes through the router, an incoming dial peer and an outgoing dial peer are selected based on the Cisco IOS dial peer routing logic. If corlists are associated with the selected dial peers, the following additional check is performed before extending the call:


- If the members of the outgoing corlist associated with the outgoing dial peer are a subset of the members of the incoming corlist associated with the incoming dial peer, the call is permitted.
- If the members of the outgoing corlist associated with the outgoing dial peer are *not* a subset of the members of the incoming corlist associated with the incoming dial peer, the call is rejected.

If no corlist statements are applied to some dial peers, the following properties apply:

- When no incoming corlist is configured on a dial-peer, the default incoming corlist is used. The default incoming corlist has the highest possible priority, and it therefore allows this dial-peer to access all other dial-peers, regardless of their outgoing corlist.
- When no outgoing corlist is configured on a dial-peer, the default outgoing corlist is used. The default outgoing corlist has the lowest possible priority, and it therefore allows all other dial-peers to access this dial-peer, regardless of their incoming corlist.

This behavior is best illustrated with an example as shown in Figure 10-19, where one VoIP dial-peer and two POTS dial-peers are defined.

Figure 10-19 Example of COR Behavior



The VoIP dial-peer is associated with the c1 incoming corlist, with members A, B, and C. You can think of members of the incoming corlist as "keys."

The first POTS dial-peer has a destination-pattern of 1.. and is associated with the c2 outgoing corlist, with members A and B. The second POTS dial-peer has a destination-pattern of 2.. and is associated with the c3 outgoing corlist, with members A, B, and D. You can think of members of the outgoing corlists as "locks."

For the call to succeed, the incoming corlist of the incoming dial-peer must have all the "keys" needed to open all the "locks" of the outgoing corlist of the outgoing dial-peer.

In the example shown in Figure 10-19, a first VoIP call with destination 100 is received by the router. The Cisco IOS call routing logic matches the incoming call leg with the VoIP dial-peer and the outgoing call leg with the first POTS dial-peer. The COR logic is then applied; because the c1 incoming corlist has all the keys needed for the c2 outgoing corlist locks (A and B), the call succeeds.

A second VoIP call with destination 200 is then received by the router. The Cisco IOS call routing logic matches the incoming call leg with the VoIP dial-peer and the outgoing call leg with the second POTS dial-peer. The COR logic is then applied; because the c1 incoming corlist is missing one "key" for the c3 outgoing corlist (D), the call is rejected.

Follow these steps when configuring the COR functionality in Cisco IOS:

-
Step 1	Define "tags" to be used as corlist members with the command dial-peer cor custom .
Step 2	Define corlists with the command dial-peer cor list <i>corlist-name</i> .
Step 3	Associate corlists with existing VoIP or POTS dial-peers by using the command corlist {incoming outgoing} <i>corlist-name</i> within the dial-peer configuration.
-

With Cisco IOS Release 12.2(8)T and later, you can apply the COR functionality to SRST-controlled IP phones. Because IP phones register with the SRST router dynamically, SRST has no prior knowledge of the individual phones until they lose connectivity to the Cisco Unified CallManager cluster. Therefore, the COR feature configuration for SRST is based on the phone DN's instead. When the IP phones register with the SRST router, they communicate their DN to it, thus allowing the SRST router to assign them to the appropriate corlist.

Configure COR for IP phones controlled by SRST by using the command **cor {incoming | outgoing}** *corlist-name* {*corlist-number starting-number – ending-number* | **default**} within the **call-manager-fallback** configuration mode.

The following limitations apply to this command:

- The maximum number of **cor {incoming | outgoing}** statements under **call-manager-fallback** is 5 (plus the default statement) in SRST version 2.0, available on Cisco IOS Release 12.2(8)T or later.
- The maximum number of **cor {incoming | outgoing}** statements under **call-manager-fallback** is 20 (plus the default statement) in SRST version 3.0, available on Cisco IOS Release 12.3(4)T or later.

The COR functionality can also be deployed with Cisco Unified CallManager Express, using Cisco IOS Release 12.2(8)T and later. Because the individual IP phones are specifically configured within Cisco Unified CallManager Express, you can apply corlists directly to the IP phones themselves by using the command **cor {incoming | outgoing}** *corlist-name* within the **ephone-dn dn-tag** configuration mode of each IP phone.

Refer to the section on [Building Classes of Service in Cisco IOS with H.323](#), page 10-84, for an example of how to apply all these concepts in practice.

For more details on configuration of Cisco SRST and Cisco Unified CallManager Express, refer to the *Cisco SRST 3.0 System Administrator Guide* and the *Cisco Unified CallManager Express 3.1 System Administrator Guide* available at

<http://www.cisco.com>

Digit Manipulation in Cisco IOS with H.323 Dial Peers

In Cisco IOS routers running H.323, digit manipulation is performed via voice translation profiles, which are used to manipulate the calling number (ANI) or called number (DNIS) digits for a voice call or to change the numbering type of a call.

Voice translation profiles are available starting with Cisco IOS Release 12.2(11)T, and they are used to convert a telephone number into a different number before the call is matched to an incoming dial peer or before the call is forwarded by the outgoing dial peer. For example, within your company you might dial a five-digit extension to reach an employee at another site. If the call is routed through the PSTN to reach the other site, the originating gateway must use voice translation profiles to convert the five-digit extension into the ten-digit format that is recognized by the central office switch.

You configure voice translation profiles by using the **voice translation-rule** and **voice translation-profile** Cisco IOS commands, which use regular expressions to define the digit strings to be modified. You then specify if the manipulations should be associated to calling numbers, called numbers, or redirected called numbers. After you define a voice translation profile, you can apply it to any of the following elements:

- All incoming POTS call legs that terminate on a specific voice port
- All incoming VoIP call legs into the router
- Outgoing call legs associated with a specific VoIP or POTS dial peer
- All incoming or outgoing call legs that terminate on the IP phones controlled by SRST
- Incoming call legs for calls originated by all IP phones controlled by SRST

**Note**

Voice translation profiles using the **voice translation-rule** command replace and enhance the functionality previously provided by the **translation-rule** command. The syntax of the new command differs from that used by the old command. For more details, refer to **voice translation-rule** in the *Cisco IOS Voice Command Reference*, Release 12.2(11)T or later, available at <http://www.cisco.com>.

A typical application of voice translation profiles is to enable the preservation of on-net inter-site dialing habits from a branch site even when the IP WAN is unavailable and the router is running in SRST mode. For example, consider a simple deployment with a central site in San Jose and three remote sites in San Francisco, New York, and Dallas. [Table 10-4](#) shows the DID ranges and the internal site codes for this example.

Table 10-4 Example of DID Ranges and Site Codes for Translation Rule Application

	San Jose	San Francisco	New York	Dallas
DID Range	(408) 555-1XXX	(415) 555-1XXX	(212) 555-1XXX	(972) 555-1XXX
Site Code	1	2	3	4

Inter-site calls are normally placed across the IP WAN by dialing the on-net access code 8 followed by the one-digit site code and the four-digit extension of the called party. To preserve these dialing habits even when the IP WAN is down and Cisco SRST is active, the internal numbers must be converted back into the E.164 numbers before being sent to the PSTN. A sample configuration for the San Francisco router is as follows:

```
voice translation-rule 1
  rule 1 /^81/ /91408555/
  rule 2 /^83/ /91212555/
  rule 3 /^84/ /91972555/

voice translation-profile on-net-xlate
  translate called 1

call-manager-fallback
  translation-profile outgoing on-net-xlate

dial-peer voice 2 pots
  destination-pattern 91[2-9]..[2-9].....
  port 1/0:0
  direct-inward-dial
  forward-digits 11
```

With this configuration, when the San Francisco site is in SRST mode and a user dials 831000, the router will match **rule 2** of **voice translation-rule 1** and translate the called number to 912125551000. This new number will then be used to match the outgoing dial peer (**dial-peer voice 2**).

For a detailed description of dial peers and their configuration, refer to *Configuring Dial Plans, Dial Peers, and Digit Manipulation*, part of the *Cisco IOS Voice, Video, and Fax Configuration Guide, Release 12.2*, available at

<http://www.cisco.com>

For a thorough explanation of regular expression syntax in Cisco IOS, refer to the document on *Regular Expressions*, available at

http://www.cisco.com/en/US/docs/ios/termserv/configuration/guide/tsv_reg_express_ps6441_TSD_Products_Configuration_Guide_Chapter.html

Design Considerations

This section presents the following dial plan design considerations for multisite deployments:

- [Design Guidelines for Multisite Deployments, page 10-52](#), covers guidelines and best practices that apply to all multisite deployment models.
- [Choosing a Dial Plan Approach, page 10-54](#), presents the various approaches to organizing a dial plan for uniform versus variable-length on-net dialing and, for this second option, partitioned versus flat addressing.
- The following sections analyze in detail three dial plan approaches and provide configuration guidelines for each:
 - [Deploying Uniform On-Net Dial Plans, page 10-56](#)
 - [Deploying Variable-Length On-Net Dial Plans with Partitioned Addressing, page 10-58](#)
 - [Deploying Variable-Length On-Net Dial Plans with Flat Addressing, page 10-63](#)
- The following sections present two alternative ways of configuring classes of service within Cisco Unified CallManager:
 - [Building Classes of Service for Cisco Unified CallManager, page 10-72](#)
 - [Classes of Service with the Line/Device Approach, page 10-76](#)
- [Building Classes of Service in Cisco IOS with H.323, page 10-84](#), explains how to implement classes of service within a Cisco IOS router running the H.323 protocol.
- [Deploying Call Coverage, page 10-87](#), provides guidelines and best practices for implementing call coverage functionality using hunt lists and line groups with Cisco Unified CallManager.

Design Guidelines for Multisite Deployments

The following guidelines and best practices apply in common to all multisite IP Telephony deployments. For deployments that involve more than one Cisco Unified CallManager cluster, also refer to the section on [Additional Considerations for Distributed Call Processing Deployments](#), page 10-53.

- To prevent routing loops, make sure the calling search spaces of all PSTN gateways do not include any partitions that contain external route patterns.
- When choosing DID ranges with your Local Exchange Carrier (LEC), try to select them so that no overlap occurs within a site. For example, if you use four-digit dialing within a site and you need two blocks of 1000 DIDs, the blocks (408)555-1XXX and (408)999-1XXX would overlap when reduced to four-digit numbers and would introduce additional complexity due to inbound and outbound translations.
- Allow for multiple ways of dialing emergency numbers. For example, in North America, configure both 911 and 9.911 as emergency route patterns within Cisco Unified CallManager.
- When automated alternate routing (AAR) is deployed, ensure that the external phone number mask configured on the IP phones (or the AAR Destination Mask configured on the DN in Cisco Unified CallManager 4.2) is compatible with all the prefixes added by the various AAR groups. For example, in a multi-national deployment, do not include national access codes such as 0 in the mask unless they are part of the global E.164 address.
- You can force calls to on-net destinations, but dialed as PSTN calls, to be routed within the cluster by adding translation patterns that match the E.164 DID ranges for each site and that manipulate the digits so they match the destination's internal extension. For example, if a DN is reachable on-net by dialing 1234, but someone within the system dials this same destination as 9 1 415 555 1234, you can force the call to be on-net by creating a translation pattern 9 1 415 555.1XXX, removing all pre-dot digits, and routing the call on-net to the resulting number. However, remember to provision AAR accordingly, using one of the following methods, so that automated PSTN failover is possible when the IP WAN is out of bandwidth:
 - Configure the AAR calling search space to exclude the partition containing the "force on-net" translation patterns but to include a partition containing the regular route patterns pointing to the PSTN.
 - Configure the AAR group to prefix digits with a special character such as *, and configure additional route patterns such as *9.!, and *9.!!# (or *0.!, and *0.!!#) in the regular partition.

The second approach is preferable because it does not require additional calling search spaces to be defined for AAR. The additional * route patterns also provide a good troubleshooting and testing tool that allows for overriding the "force on-net" configuration and placing calls over the PSTN even when AAR is not invoked, or to force a call through the PSTN when the destination number is located in a branch under SRST mode.

- Within a centralized call processing cluster with N sites, you can implement Tail-End Hop-Off (TEHO) using one of the following methods:
 - TEHO with centralized failover
This method involves configuring a set of N route patterns in a global partition, with each pattern pointing to a route list that has the appropriate remote site route group as the first choice and the central site route group as the second choice.
 - TEHO with local failover
This method involves configuring N sets of N route patterns in site-specific partitions, with each pattern pointing to a route list that has the appropriate remote site route group as the first choice and the local site route group as the second choice.

While the second approach allows for an optimal failover scenario when the remote gateway or the IP WAN is unavailable, it also introduces a high level of complexity into the dial plan because it requires a minimum of N^2 route patterns and N^2 route lists, as opposed to the N route patterns and N route lists needed with the first approach.

- When appropriate for your national numbering plan, you may configure an additional translation pattern at each site to catch local PSTN calls dialed as long-distance calls and to translate them into the proper abbreviated form. This translation pattern should be accessible only from phones located within the site. Such a configuration also helps simplify the AAR configuration (see [Special Considerations for Sites Located Within the Same Local Dialing Area, page 10-25](#)).
- Do not use the multilevel precedence and preemption (MLPP) feature to assign higher priority to emergency calls. An emergency-related call might not appear as such to the IP Telephony system, and you would risk terminating an existing emergency call to place another call to the main emergency service routing number. For example, an emergency situation might prompt someone to place a call to a regular ten-digit number to reach a medical professional. Preemption of this call would abort the ongoing emergency communication and could delay handling of the emergency. Also, incoming calls from emergency service personnel would be at risk of preemption by MLPP.

**Note**

A Cisco Unified CallManager cluster with a very large dial plan containing many gateways, route patterns, translation patterns, and partitions can take an extended amount of time to initialize when the Cisco CallManager Service is first started. If the system does not initialize within the default time, there are service parameters that you can increase to allow additional time for the configuration to initialize. For details on the service parameters, refer to the online help for Service Parameters in Cisco Unified CallManager Administration.

Additional Considerations for Distributed Call Processing Deployments

In addition to the considerations made in the previous section, observe the following best practices when designing a dial plan for a distributed call processing deployment (that is, any multisite deployment involving multiple Cisco Unified CallManager clusters):

- Avoid splitting DID ranges across multiple Cisco Unified CallManager clusters. This practice would make intercluster routing very difficult because summarization would not be possible. Each DID range should belong to a single Cisco Unified CallManager cluster.
- Avoid splitting devices within a remote site between multiple Cisco Unified CallManager 4.x clusters using call admission control based on locations. Locations-based call admission control is significant only within a cluster, and having devices belong to different clusters at the same remote site would lead to poor utilization of the IP WAN bandwidth because you would have to split the available bandwidth between the clusters. Each remote site should belong to a single Cisco Unified CallManager cluster.
- Use gatekeeper-controlled intercluster trunks to route calls between Cisco Unified CallManager clusters. This practice enables you to add or modify clusters easily in your network without reconfiguring all other clusters.
- Implement redundancy in the connection between Cisco Unified CallManager and the gatekeeper by using a gatekeeper cluster and by assigning the intercluster trunk to a device pool that uses a Cisco Unified CallManager group with multiple servers configured.
- When sending calls to the gatekeeper, expand the called number to the full E.164 address. This practice simplifies PSTN failover when the IP WAN is not available because no additional digit manipulation is required to reroute the call via a PSTN gateway. Also, this practice eliminates the need to configure the local (calling) Cisco Unified CallManager with dial length information for each remote site.

- Within the gatekeeper, configure one zone per Cisco Unified CallManager cluster. For each cluster/zone, add zone prefix statements to match all DN ranges owned by that cluster.
- You can implement Tail-End Hop-Off (TEHO) across multiple Cisco Unified CallManager clusters by following these guidelines:
 - Add specific route patterns for the relevant E.164 ranges to the source (originating) Cisco Unified CallManager cluster, and point them to a route list that has the IP WAN route group as the first choice and a local PSTN route group as the second choice
 - Within the Cisco IOS gatekeeper configuration, add zone prefix statements for all the relevant E.164 ranges and point them to the appropriate Cisco Unified CallManager cluster
 - Ensure that the intercluster trunk calling search space in the destination Cisco Unified CallManager cluster includes partitions featuring route patterns that match the local PSTN numbers, and that digit manipulation is applied as needed (for example, stripping the area code before sending the call to the PSTN).

For details on how to configure a Cisco IOS gatekeeper for distributed call processing deployments, see [Gatekeeper Design Considerations, page 8-18](#).

Choosing a Dial Plan Approach

As introduced in [Planning Considerations, page 10-2](#), there are two main approaches to a dial plan for internal destinations within an IP Telephony system:


- Uniform on-net dial plan, where each internal destination is dialed in the same way regardless of whether the caller is located in the same site or in a different site.
- Variable-length on-net dial plan, where internal destinations are dialed differently within a site than across sites. Typically, this approach uses four or five-digit abbreviated dialing for calls within a site and either full E.164 addresses or an on-net access code followed by a site code and the extension for calls across sites.

To help you decide which approach is best suited for your needs, consider the following high-level design questions:

- How many sites will eventually be served by the IP Telephony system?
- What are the calling patterns between sites or branches?
- What do users dial within a site and to reach another site?
- Are there any calling restrictions applied to on-net inter-site calls?
- What transport network (PSTN or IP WAN) will be used for most inter-site calls?
- What (if any) CTI applications are being used?
- Is there a desire for a standardized on-net dialing structure using site codes?

Uniform on-net dial plans are the easiest to design and configure; however, they work best for small to medium deployments, and they become impractical when the number of sites and users increases. They are described and analyzed in detail in the section on [Deploying Uniform On-Net Dial Plans, page 10-56](#).

Variable-length on-net dial plans are more scalable but also more complex to design and configure. [Figure 10-20](#) shows the typical requirements for a large-scale deployment using the variable-length on-net dial plan approach.

Figure 10-20 Typical Dialing Requirements for Large Multisite Deployments

With Cisco Unified CallManager, there are two main methods for implementing a variable-length on-net dial plan:

- Partitioned addressing

The internal extensions are placed in different partitions according to the site where they are located. This method is typically based on full E.164 addresses for inter-site calls and is analyzed in detail in the section on [Deploying Variable-Length On-Net Dial Plans with Partitioned Addressing](#), page 10-58.

- Flat addressing

All internal extensions are placed in the same partition. This method is typically based on on-net site codes for inter-site calls and is analyzed in detail in the section on [Deploying Variable-Length On-Net Dial Plans with Flat Addressing](#), page 10-63. In some cases it is possible to use this approach even when using full E.164 addresses for inter-site calls, as described in the section on [Special Considerations for Deployments Without Site Codes](#), page 10-70.


Deploying Uniform On-Net Dial Plans

You can implement a uniform on-net dial plan by following these guidelines:

- Uniquely identify all phones with an abbreviated extension.
- Place all the phone DNs in a single partition.
- At each site, place PSTN route patterns in one or more site-specific partitions, according to the chosen class-of-service approach.

Figure 10-21 shows an example implementation for a single Cisco Unified CallManager cluster deployment.

Figure 10-21 Example of Uniform On-Net Dial Plan Deployment



Use this approach if both of the following conditions apply:

- The DID ranges available do not overlap when considering the number of digits chosen to identify internal extensions.
- The number of sites covered by the IP Telephony system is not expected to grow significantly over time.

The following sections analyze implementation details and best practices for different types of calls within the framework of uniform on-net dial plans:

- [Inter-Site Calls Within a Cluster, page 10-57](#)
- [Outgoing PSTN and IP WAN Calls, page 10-57](#)
- [Incoming Calls, page 10-57](#)
- [Voicemail Calls, page 10-58](#)

Inter-Site Calls Within a Cluster

Because all internal DNs are directly reachable from every device's calling search space, all on-net calls (intra-site and inter-site) are automatically enabled and need no additional configuration within Cisco Unified CallManager.

Outgoing PSTN and IP WAN Calls

PSTN calls are enabled via the site-specific partitions and route patterns, so that emergency and local calls can be routed via the local branch gateway. Long-distance and international calls may be routed via the same branch gateway (as shown in [Figure 10-21](#)) or via a centralized gateway, depending on company policy. This second option would simply require an additional route list per site, with the first-choice route group pointing to the central site gateway and, optionally, a second-choice route group pointing to the local branch gateway.

Abbreviated dialing to another Cisco Unified CallManager cluster or Cisco Unified CallManager Express is also possible via a gatekeeper. For these IP WAN calls, Cisco recommends that you expand the abbreviated string to the full E.164 via a translation pattern before sending it to the gatekeeper.

Emergency Calls

If Cisco Emergency Responder is used for managing emergency calls, the partition containing the CTI route point used to send calls to Cisco Emergency Responder should be part of the calling search space of all phones in all branches instead of the site-specific 911 patterns as illustrated in [Figure 10-21](#). Cisco Emergency Responder will be able to identify the calling phone because there is no duplicity of DNs allowed in the internal partition. For more information on Cisco Emergency Responder considerations, refer to the chapter on [Emergency Services, page 11-1](#), and to the Cisco Emergency Responder product documentation available at

<http://www.cisco.com>

Incoming Calls

Incoming PSTN calls simply require stripping the excess digits in order to match the extension length configured in Cisco Unified CallManager. This can be done within the gateway configuration or, alternatively, via a translation pattern included in the gateway's calling search space.

Voicemail Calls

Because every extension is unique within the system, the extension itself can be used to configure voicemail boxes within the voicemail system. No translations are necessary to send calls to the voicemail system or to enable a Message Waiting Indicator (MWI) within Cisco Unified CallManager.

However, when users access the voicemail system from the PSTN, they need to be trained to enter their eight-digit extension to access their voicemail boxes.

Deploying Variable-Length On-Net Dial Plans with Partitioned Addressing

You can implement a variable-length on-net dial plan with partitioned addressing by placing the abbreviated DN's for phones at each site in a different partition (one partition per site), and by relying on a set of translation patterns placed in a global partition (one translation pattern per site) to perform the inter-site call routing.

This method enables you to satisfy the key requirement of supporting abbreviated dialing (typically four or five digits) within a site and full E.164 dialing between sites, at the price of higher complexity in the dial plan.



Note

These deployments are also known as *overlapping dial plans* or *dial plans with overlapping extensions* because the abbreviated DN's defined at the various sites usually overlap with each other.

Table 10-5 illustrates the basic relationship between calling search spaces and partitions at each site, without taking into account additional elements required to implement classes of service.

Table 10-5 **Calling Search Spaces and Partitions for Variable-Length Dial Plans with Partitioned Addressing**

Calling Search Space	Partition	Partition Contents
Site1_css	Site1Phones_pt	Site 1 phone DN's (abbreviated form)
	Site1PSTN_pt	PSTN route patterns for Site 1 (more partitions may be needed based on classes of service)
	Translations_pt	Translation patterns for inter-site calling within the cluster
...	...	...
SiteN_css	SiteNPhones_pt	Site N phone DN's (abbreviated form)
	SiteNPSTN_pt	PSTN route patterns for Site N (more partitions may be needed based on classes of service)
	Translations_pt	Translation patterns for inter-site calling within the cluster

Use this approach if one or more of the following conditions apply:

- A global on-net numbering plan using site codes is *not* desired

**Note**


While this approach can also be used in the presence of an on-net internal numbering plan using site codes, for this scenario Cisco recommends that you follow the flat addressing approach described in the section on [Deploying Variable-Length On-Net Dial Plans with Flat Addressing, page 10-63](#), because it enables you to greatly simplify the dial plan structure and to better manage and troubleshoot the system.

- Policy restrictions must be applied to on-net inter-site calls (that is, some or all users are not allowed to dial other sites on-net).
- Inter-site calls are always routed over the PSTN (as described in the section on [Voice Over the PSTN as a Variant of Centralized Call Processing, page 2-10](#)).
- CTI-based applications are not used across sites.

**Note**

Some CTI-based applications (such as Cisco Emergency Responder), as well as Cisco Attendant Console, do not support overlapping extensions and cannot be deployed with the partitioned addressing approach. Other applications, such as Cisco Personal Assistant and Cisco Unified CallManager Assistant, require additional dial plan configuration that becomes more complex and might be impaired in presence of overlapping extensions. If you are planning to deploy these applications, Cisco recommends that you choose the flat addressing approach described in the next section.

To better understand how to deploy the partitioned addressing approach, consider the hypothetical customer network shown in [Figure 10-22](#). This network has a main site (San Jose) and a number of smaller branch sites (New York and Dallas) in the United States, and a similar topology with a main site (London) and several smaller branch sites (Paris and Milan) in Europe. Given the user count, the administration needs, and the network topology, two Cisco Unified CallManager clusters with centralized call processing are deployed, one in the United States and one in Europe. A Cisco IOS gatekeeper cluster is used to provide E.164 address resolution and call admission control between the two clusters. It has also been decided that a variable-length on-net dial plan is required, with four-digit dialing within each site (the 1XXX extension range is chosen at each site) and full E.164 dialing between sites.

Figure 10-22 Example of a Large Multisite Deployment


Using the United States (US) cluster from [Figure 10-22](#) as an example, the following sections analyze implementation details and best practices for different types of calls within the framework of the partitioned addressing approach:

- [Inter-Site Calls Within a Cluster](#), page 10-61
- [Outgoing PSTN and IP WAN Calls](#), page 10-61
- [Incoming Calls](#), page 10-63
- [Voicemail Calls](#), page 10-63

Inter-Site Calls Within a Cluster

Figure 10-23 shows an example configuration for inter-site calls within the US cluster.

Figure 10-23 Inter-Site Calls Within a Cluster for the Partitioned Addressing Method




To provide connectivity between sites and partitions, use translation patterns according to the following guidelines:

- Define one translation pattern for each site, and place them all in the **Translations_pt** partition.
- Each pattern should match a site's E.164 address range, including the PSTN access code (9 in the example).
- The resulting called numbers after translation should match the site's extensions (1XXX in the example).
- The calling search space where the call is sent after translation should include the partition that contains the destination site's IP phone DN's.

Outgoing PSTN and IP WAN Calls

Depending on how the various types of PSTN calls need to be routed (centralized gateways versus distributed gateways), the configuration may vary. In the example shown in Figure 10-24, all domestic PSTN calls are routed via the local branch gateway. International calls are first routed through the gatekeeper to intercept calls toward the sites controlled by the Europe cluster, and then through the local PSTN gateway.

Figure 10-24 Outgoing PSTN and IP WAN Calls for the Partitioned Addressing Method.



114729

As Figure 10-24 illustrates, for outgoing PSTN and IP WAN calls, there are no special considerations due to the partitioned nature of the internal addressing.



Note

Figure 10-24 uses the traditional approach for building classes of service. However, the line/device approach can also be combined with the partitioned addressing approach. Class-of-service approaches and endpoint addressing approaches are independent and interchangeable.

Incoming Calls

To dispatch incoming PSTN or IP WAN calls to the appropriate extension, you can reuse the translation patterns in the Translations_pt partition mentioned previously. It is sufficient to assign all PSTN gateways a calling search space that contains only the Translations_pt partition, making sure that the gateways also prefix a 9 or a 91 to the incoming dialed number to match the already defined translation patterns. This approach assumes that the PSTN delivers the fully qualified number of each phone to the gateways, so that if an incoming dialed number is prefixed with 91, Cisco Unified CallManager will be presented with a number that matches the global translation patterns located in the Translations_pt partition.

Voicemail Calls

Voicemail integration requires special attention to the following requirements with the partitioned addressing approach:

- Voicemail boxes must have unique identifiers. This means that the IP phone extension cannot be used as a voicemail box, and some sort of digit manipulation is needed to obtain a unique number.
- Message waiting indicator (MWI) messages from the voicemail system must be able to reach the right IP phone even in the presence of non-unique extensions.

The first issue is addressed in Cisco Unified CallManager by the use of the Voice Mail Box Mask field in the Voice Mail Profile Configuration page. This parameter, when configured, is used to communicate with the voicemail system and to uniquely identify the user. For example, the Voice Mail Box Mask parameter can be set to the full E.164 number associated with the user.

The second issue is addressed by reusing the translation patterns in the on-cluster partition. If the voicemail system has been configured with the full E.164 numbers, it is sufficient to prepend 9 to the E.164 number in order to match the translation patterns previously defined in the Translations_pt partition and to ensure proper inter-site communication. In this way, the MWI messages coming from the voicemail system with the full E.164 number will be translated to the appropriate extension in the specific partition. For example, the voicemail ports can be configured with a calling search space that contains a VM_Translations_pt partition with a single translation pattern that prefixes 9 to the E.164 numbers dialed by the voicemail system. The calling search space of this translation pattern contains the Translations_pt partition, which then provides access to all the extensions via the previously defined translation patterns.



Note

This approach requires the configuration of two service parameters within Cisco Unified CallManager. The MultiTenantMwiMode parameter within the Cisco CallManager service must be set to True, and the ValidateDNs parameter within the Cisco Messaging Interface (CMI) service must be set to False.

Deploying Variable-Length On-Net Dial Plans with Flat Addressing

You can implement a variable-length on-net dial plan with flat addressing by defining phone DN's as unique strings containing an on-net access code, a site code, and the extension (for example, 8-123-1000). You can place all these DN's in the same global partition, thus enabling inter-site calls using the site code, and you can define translation patterns in site-specific partitions (one translation pattern and one partition per site) to enable abbreviated dialing within a site.

This internal structure can be hidden from the end users by configuring the Line Text Label parameter within the Directory Number configuration page with the four or five-digit number that users are accustomed to dialing within the site. The external phone number mask should also be provisioned with the corresponding PSTN number in order to enable AAR and to give users a visual indication of their own DID number on the IP phone display.

Table 10-6 illustrates the basic relationship between calling search spaces and partitions at each site, without taking into account additional elements required to implement classes of service:

Table 10-6 *Calling Search Spaces and Partitions for Variable-Length Dial Plans with Flat Addressing*

Calling Search Space	Partition	Partition Contents
Site1_css	Site1Translations_pt	Translation patterns for Site 1's abbreviated dialing
	Site1PSTN_pt	PSTN route patterns for Site 1 (more partitions may be needed based on classes of service)
	Internal_pt	All IP phone DNs (unique form)
...	...	...
SiteN_css	SiteNTranslations_pt	Translation patterns for Site N's abbreviated dialing
	SiteNPSTN_pt	PSTN route patterns for Site N (more partitions may be needed based on classes of service)
	Internal_pt	All IP phone DNs (unique form)

Use this approach if one or more of the following conditions apply:

- No dialing restrictions are needed for on-net inter-site calls.
- A global on-net numbering plan using site codes is desired.
- Inter-site calls are normally routed over the IP WAN.
- CTI-based applications are used across sites.



Note

If dialing restrictions need to be applied to on-net inter-site calls, or an on-net numbering plan using site codes is not desired, refer to the section on [Special Considerations for Deployments Without Site Codes](#), page 10-70, for a variant of this approach that can accommodate these needs.

The following considerations apply to this approach:

- The destination numbers of intra-site four-digit calls get expanded to the unique internal DN on the IP phone display.
- The Placed Calls directory will display the original four-digit string as dialed by the user.
- Calling number, and numbers in the Missed Calls and Received Calls directories, appear as the unique internal DN.
- If you wish to preserve the four-digit dialing feature when the IP WAN is unavailable and the branch phones are in SRST mode, you need to apply a translation rule to the **call-manager-fallback** configuration within the SRST router.

- When the branch phones are in SRST mode, the Line Text Label that masks the unique internal DN as a four-digit number on the IP phone display is not available, so the users will see their full internal DN appear instead.

To better understand how to deploy the flat addressing approach, consider again the hypothetical customer network shown in [Figure 10-22](#). In this case, it has been decided that a variable-length on-net dial plan is required, with four-digit dialing within each site (the 1XXX extension range is chosen at each site) and inter-site dialing with an eight-digit string consisting of an on-net access code (8 in this example), a three-digit site code, and the four-digit extension. The three-digit site code is derived from the NANP area code for the sites located in the United States, and from the E.164 country code followed by a site identifier for the sites located in Europe. [Table 10-7](#) summarizes the site codes chosen.

Table 10-7 **Site Codes for the Customer Network in [Figure 10-22](#)**

	San Jose	New York	Dallas	London	Paris	Milan
Site Code	408	212	972	442	331	392


Using the US cluster from this example, the following sections analyze implementation details and best practices for different types of calls within the framework of the flat addressing approach:

- [Inter-Site Calls Within a Cluster, page 10-66](#)
- [Outgoing PSTN and IP WAN Calls, page 10-67](#)
- [Incoming Calls, page 10-70](#)
- [Voicemail Calls, page 10-70](#)
- [Special Considerations for Deployments Without Site Codes, page 10-70](#)

Inter-Site Calls Within a Cluster

Figure 10-25 shows an example configuration for inter-site calls within the US cluster.

Figure 10-25 *Inter-Site Calls Within a Cluster for the Flat Addressing Method*



To provide connectivity between sites and partitions, use the following guidelines:

- Place all unique DNs, including the on-net access code 8, in a global partition (named Internal_pt in this example).
- Create one partition per site, each containing a translation pattern that expands four-digit numbers into the fully qualified eight-digit number for that site, thus enabling abbreviated dialing within the site.
- For each site, include both the Internal_pt partition and the local translation partition in the phone's calling search space.

The inclusion of the on-net access code in the DN configured in Cisco Unified CallManager enables you to place all internal extensions in a partition directly accessible by all phones, and at the same time ensures that all call directories on the IP phones are populated with numbers that can be directly redialed.



Note

You must ensure that the on-net access code and site code combinations do not overlap with the local abbreviated dialing range at any site.

Outgoing PSTN and IP WAN Calls

Depending on how the various types of PSTN calls need to be routed (centralized gateways versus distributed gateways), the configuration may vary.

To provide on-net connectivity for inter-site calls to the Europe (EU) cluster, the following options are possible:

Option 1: Eight-Digit Numbers Only


This option relies on a single route pattern that matches all eight-digit ranges (8XXXXXXX) and points to a route list or route group that contains only a gatekeeper-controlled intercluster trunk. The gatekeeper is configured to use the site codes as zone prefixes.

This solution is simple and easy to maintain because no information is needed about other clusters' site codes or E.164 ranges. However, no automatic PSTN failover is provided when the IP WAN is unavailable, and users are expected to redial manually using the PSTN access code and the E.164 address of the destination.

Option 2: Eight-Digit Numbers and E.164 Addresses with Centralized PSTN Failover

This option, illustrated in [Figure 10-26](#), uses a global set of translation patterns that match the Europe eight-digit ranges and translate them into the corresponding E.164 numbers. The translation patterns use the central site's calling search space (in this case, San Jose), and the call then matches the international PSTN route pattern within the central site's PSTN partition. At each site, the international PSTN route pattern points to a route list with the IP WAN route group as a first choice and the local PSTN route group as a second choice. The gatekeeper is configured to use the E.164 ranges as zone prefixes.

Figure 10-26 Outgoing PSTN and IP WAN Calls for the Flat Addressing Method with Centralized PSTN Failover for IP WAN Calls



**Note**

The configuration example in [Figure 10-26](#) assumes that the line/device approach to building classes of service is being used, but the same considerations apply when using the traditional approach.


This solution requires a little more configuration and maintenance than that outlined in Option 1 because it requires that you configure and maintain information about other clusters' site codes and E.164 ranges. On the other hand, it provides automatic PSTN failover when the IP WAN is unavailable. PSTN failover is provided using the central site's gateway only, so the IP WAN bandwidth usage is not optimal.

Also observe that calls to the Europe sites dialed as PSTN calls will automatically be placed on-net if the IP WAN is available, with an automatic PSTN failover that in this case uses the local gateway.

Option 3: Eight-Digit Numbers and E.164 Addresses with Distributed PSTN Failover

This option, illustrated in [Figure 10-27](#), uses a set of translation patterns per site. Each set matches the Europe eight-digit ranges and translates them into the corresponding E.164 numbers. The translation patterns use the originating site's calling search space, and the call then matches the international PSTN route pattern within the originating site's PSTN partition. At each site, the international PSTN route pattern points to a route list with the IP WAN route group as a first choice and the local PSTN route group as a second choice. The gatekeeper is configured to use the E.164 ranges as zone prefixes.

Figure 10-27 *Outgoing PSTN and IP WAN Calls for the Flat Addressing Method with Distributed PSTN Failover for IP WAN Calls*



This solution requires a lot more configuration and maintenance than that outlined in Option 2 because it requires that you configure and maintain information about other clusters' site codes and E.164 ranges, and that you repeat the configuration for each remote site within the cluster. On the other hand, it provides automatic PSTN failover when the IP WAN is unavailable, using the local site's gateway so that the IP WAN bandwidth usage is optimal.

Also in this case, as for Option 2, calls to the Europe sites dialed as PSTN calls will automatically be placed on-net if the IP WAN is available, with an automatic PSTN failover using the local gateway.

Incoming Calls

Incoming PSTN calls require that the E.164 number be manipulated to obtain the eight-digit internal number in order to reach the destination phone. You can implement this requirement in any of the following ways:

- Configure the Num Digits and Prefix Digits fields within the Gateway Configuration page in Cisco Unified CallManager to strip and then prefix the needed digits.
- If you have configured translation patterns to force on-net inter-site calls within the cluster, you can reuse these patterns by simply prefixing the PSTN access code to the incoming called number on the gateway.
- If you are using an H.323 gateway, you can use translation rules within the gateway to manipulate the digits before sending the call to Cisco Unified CallManager.

The third approach has the advantage that the translation rules you configured can be reused to provide incoming PSTN connectivity to the IP phones when the branch is in SRST mode.

Voicemail Calls


Because every eight-digit extension is unique within the system, the extension itself can be used to configure voicemail boxes within the voicemail system. No translations are necessary to send calls to the voicemail system or to enable Message Waiting Indicators (MWIs) within Cisco Unified CallManager. Note that users are required to use their eight-digit on-net number when prompted for the mailbox number.

Special Considerations for Deployments Without Site Codes

This scenario is a variant of the flat addressing approach that does not rely on the definition of an on-net numbering plan based on site codes. In this scenario, intra-site calls are still dialed as four-digit numbers, while inter-site calls are dialed as regular PSTN calls and are then intercepted and routed across the IP WAN by Cisco Unified CallManager.

To implement this mechanism, follow these guidelines, illustrated in [Figure 10-28](#):

- Define the phone DNs as the full E.164 addresses and place them all in the same partition (named `Internal_pt` in this example).
- Configure the phones' calling search spaces so that they do not have direct access to the `Internal_pt` partition. Instead, they point to a global partition named `Global_Translations_pt`, which contains one translation pattern per site or DID range.
- Configure the translation patterns to strip the PSTN access code and country code (for example, 91) and send the call to the `Internal_pt` partition via the `Hidden_css` calling search space.
- Abbreviated four-digit dialing within each site can still be provided via site-specific partitions containing translation patterns (one partition and one translation pattern per site or DID range).

Figure 10-28 Variable-Length Dial Plans with Flat Addressing Without Site Codes

This configuration variant allows you to establish dialing restrictions for inter-site calls. For example, if you must prevent a certain group of users from calling other sites, configure their calling search space so that it does not contain the Global_Translations_pt partition. However, keep in mind the following factors when choosing this approach:

- Additional complexity is introduced in the form of translation patterns.
- Because you are effectively forcing on-net PSTN calls, remember to configure AAR so that calls can still be placed across the PSTN when the IP WAN bandwidth is not sufficient. Refer to [Design Guidelines for Multisite Deployments, page 10-52](#), for details.
- The placed-calls directory displays digit strings as they were dialed by the user. For example, if the user dialed 1000 and a call was placed to phone 4085551000, the placed-calls directory would display 1000, thus allowing for direct redialing of the number without having to edit the dial string.
- The missed-calls and received-calls directories display phone numbers as they appeared when the call was offered to the phone. The calling number thus offered depends on the Calling Number configuration parameters of the translation patterns and the line configuration of the calling phone.
- The configuration shown in [Figure 10-28](#) assumes that PSTN calls are dialed in the same way in all sites, which is usually true for deployments contained within a single country. International deployments require one additional set of translation patterns per country in order to intercept inter-site calls.

- International deployments also introduce the complexity of having to deal with variable-length internal DNs because E.164 addresses have different lengths in different countries (and even within a single country in some cases).

Building Classes of Service for Cisco Unified CallManager

Cisco Unified CallManager offers two main approaches to defining and applying classes of service to users and devices: the traditional approach and the line/device approach. The fundamental elements addressed for each approach include the types of calls to be allowed (for example, local, national, or international) and the path taken by the calls (for example, IP network, local gateway, or central gateway). Both elements depend on the calling search space configuration. The following sections describe the two main approaches used in Cisco Unified CallManager systems to implement classes of service. Both approaches are based on the fundamental functionality of the line and device calling search space.

**Note**


In versions of Cisco Unified CallManager prior to Release 4.2, the device calling search space configuration of a phone remains the same as the device is moved to different parts of the network. With Cisco Unified CallManager 4.2, the device calling search space can be determined dynamically based on where in the network the phone is physically located, as determined by the phone's IP address. See [Device Mobility, page 10-26](#), for more details.

Classes of Service with the Traditional Approach

With Cisco Unified CallManager, you can define classes of service for IP Telephony users by combining partitions and device calling search spaces with external route patterns, as follows:

- Place external route patterns in partitions associated with the destinations they can call. While you could place all route patterns in a single partition, you can achieve more refined call restriction policies by associating the route patterns with partitions according to the destinations they can call. For example, if you place local and international route patterns in the same partition, then all users can reach both local and international destinations, which might not be desirable. Cisco recommends that you group route patterns in partitions according to the reachability policies for the various classes of service.
- Configure each calling search space to be able to reach only the partitions associated with its call restriction policy. For example, configure the local calling search space to point to the internal and local partitions, so that users assigned to this calling search space can place only internal and local calls.
- Assign these calling search spaces to the phones by configuring them on the Cisco Unified CallManager device pages. In this way, all lines configured on the device automatically receive the same class of service.

[Figure 10-29](#) shows a simple example for a single-site deployment.

Figure 10-29 Basic Example of Classes of Service Using the Traditional Approach

114720

With this approach, the device calling search space performs two distinct logical functions:


- Path selection

The calling search space contains specific partitions, which in turn contain specific route patterns that point to specific PSTN gateways through route lists and their associated route groups.

- Class of service

By selectively including certain partitions and not others in the device calling search space, you effectively apply calling restrictions to certain groups of users.

As a consequence, when you apply this approach to a multisite deployment with centralized call processing, you have to replicate partitions and calling search spaces for each site because for each site you have to create classes of service and, at the same time, route the PSTN calls out of the local branch gateways, as illustrated in [Figure 10-30](#).

Figure 10-30 Calling Search Spaces and Partitions Needed with the Traditional Approach

Follow these additional guidelines when applying this dial plan approach to a multisite deployment with centralized call processing:

- To facilitate on-net dialing between sites, place all IP phone DN's in an on-cluster or internal partition that can be reached from the calling search spaces of all sites. Note that this is not possible if the IP phone DN's overlap. For more details on dial plans that use overlapping extensions, refer to [Deploying Variable-Length On-Net Dial Plans with Partitioned Addressing](#), page 10-58.
- Give each remote site its own set of partitions and route patterns. The number of partitions per remote site depends on the number of calling restriction policies associated with the route patterns.
- Give each site its own set of calling search spaces for its IP phones. The calling search spaces point to the on-cluster partition as well as to the appropriate local route pattern partitions.
- Depending on your company's call forwarding restriction policy, you can reuse one of the site-specific device calling search spaces for the Forward All calling search space.

As a general rule, use the following formulas to calculate the total number of calling search spaces and partitions needed:

Total Partitions = (Number of classes of service) * (Number of sites) + (1 Partition for all IP phone DNs)

Total Calling Search Spaces = (Number of classes of service) * (Number of sites)

**Note**

These values represent the minimum number of partitions and calling search spaces required. You may need additional partitions and calling search spaces for special devices and applications, as well as for on-net patterns for other call processing agents.

Extension Mobility Considerations with the Traditional Approach

When using the Extension Mobility feature, the dialing restrictions of a phone are a function of the logged-in (or logged-out) status of the phone. Typically, a logged-out phone should be restricted to calling other phones and services (such as 911), but access to local or toll calls through the PSTN are restricted. Conversely, a phone where a user is logged-in should allow calls according to that user's dialing privileges and should route those calls to the appropriate gateway (for example, a co-located branch gateway for local calls).

With the traditional approach for building classes of service, consider the following guidelines to apply calling restrictions when using Extension Mobility:

- At each site, configure the device calling search space for all IP phones to point to only PSTN emergency services (using the local gateway).
- Configure the line calling search spaces for IP phones used for Extension Mobility in a logged-out state to point to internal numbers only.
- For each Extension Mobility user, configure the line calling search space within the device profile to point to internal numbers and the additional PSTN route patterns allowed for their specific class of service (again, using an appropriate gateway according to company policy).

Note that, when an Extension Mobility user who is normally based in Site 1 logs into an IP phone in Site 2, the path selection for PSTN calls will change as follows:

- Emergency calls will be correctly routed using Site 2's PSTN gateway because the emergency services are provided by the device calling search space of the IP phone at Site 2.
- All other PSTN calls will be routed according to the Extension Mobility user's profile (more specifically, the line calling search space configured in the device profile). Typically, this means that these PSTN calls will traverse two WAN links and use Site 1's gateway to access the PSTN.

You can use one of the following methods to modify this behavior and ensure that PSTN calls are always routed via the local PSTN gateway even when Extension Mobility users roam across different sites:

- Include local PSTN route patterns in the device calling search space and remove them from the line calling search space within the device profile. This method ensures that local PSTN calls will be routed via the co-located branch gateway, but it also means that users will be able to dial these calls even without logging into the IP phones. Long-distance and international calls will still be routed according to the Extension Mobility user's device profile, so this solution is suitable only if these calls are usually routed via a centralized gateway.
- Define multiple device profiles for each user, one for each site to which they usually roam. Each device profile is configured so that its line calling search space points to PSTN route patterns that use the local gateway for that site. This method might prove burdensome to configure and manage if a significant number of users roam to a significant number of sites.

- Implement the line/device approach described in the next section on [Classes of Service with the Line/Device Approach](#), page 10-76.

**Note**

When Cisco Emergency Responder is used, the site-specific calling search space configured on the device should include the partition containing the 911 CTI route point pointing to Cisco Emergency Responder. This same partition can also contain a translation pattern 9.911 pointing to the same 911 CTI route point, to allow users to dial 9911 when trying to reach emergency services.

Classes of Service with the Line/Device Approach

The traditional approach outlined in the preceding section can result in a large number of partitions and calling search spaces when applied to large multisite deployments with centralized call processing. This configuration is required because the device calling search space is used to determine both the path selection (which PSTN gateway to use for external calls) and the class of service.


It is possible to significantly decrease the total number of partitions and calling search spaces needed by dividing these two functions between the line calling search space and the device calling search space, in what is called the *line/device approach*.

Keeping in mind how the line calling search space and the device calling search space for each given IP phone are combined within Cisco Unified CallManager, and how the line calling search space partitions appear first in the resulting calling search space (see [Calling Privileges in Cisco Unified CallManager](#), page 10-15), you can apply the following general rules to the line/device approach:

- Use the device calling search space to provide call routing information (for example, which gateway to select for PSTN calls).
- Use the line calling search space to provide class-of-service information (for example, which calls to allow).

To better understand how to apply these rules, consider the example shown in [Figure 10-31](#), where the device calling search space contains a partition with route patterns to all PSTN numbers, including international numbers. The route patterns point to a PSTN gateway via the route list and route group construct.

Figure 10-31 Key Concepts in the Line/Device Approach



At the same time, the line calling search space contains a partition with a single translation pattern that matches international numbers and that has been configured as a blocked pattern.

The resulting calling search space therefore contains two identical patterns matching international numbers, with the blocked pattern in the line calling search space appearing first. The result is that international calls from this line will be blocked.

It is possible to use route patterns instead of translation patterns to block calls within the line calling search space. To configure a blocked route pattern, first create a "dummy" gateway with an unused IP address and place it into a "dummy" route list and route group construct. Then point the route pattern to the dummy route list. The main difference between using a route pattern and a translation pattern to block calls is the end-user experience when trying to dial a blocked number, as follows:


- When using a translation pattern, the end users will be able to dial the entire number and only then will they hear a fast-busy tone.
- When using a route pattern, the end users will hear a fast-busy tone as soon as the number they are dialing can no longer match any allowed pattern.

Follow these guidelines to implement the line/device approach in a multisite deployment with centralized call processing:

- Create an unrestricted calling search space for each site and assign it to the phone's device calling search space. This calling search space should contain a partition featuring route patterns that route the calls to the appropriate gateway for the phone's location (for example, a co-located branch gateway for emergency services and a centralized gateway for long-distance calls).
- Create calling search spaces containing partitions featuring blocked translation/route patterns for those types of calls not part of the user's dialing privileges, and assign them to the user's lines. For instance, if a user has access to all types of calls except international, that user's line (or lines) should be configured with a calling search space that blocks the 9.011! route pattern.

Figure 10-32 shows an example of how these guidelines can apply to a multisite deployment with N sites.

Figure 10-32 Calling Search Spaces and Partitions Needed with the Line/Device Approach



This approach has the significant advantage that only a single site-specific, unrestricted calling search space is required for each site (that is, one per branch). Because the dialing privileges are implemented through the use of blocked route patterns (which are not site-dependent), the same set of blocking calling search spaces can be used in all branches.

Consequently, you can use the following formulas to calculate the total number of calling search spaces and partitions needed:

Total Partitions = (Number of classes of service) + (Number of sites) + (1 Partition for all IP phone DNs)

Total Calling Search Spaces = (Number of classes of service) + (Number of sites)

**Note**

These values represent the minimum numbers of partitions and calling search spaces required. You may need additional partitions and calling search spaces for special devices and applications, as well as for on-net patterns for other call processing agents.

**Note**

If Cisco Emergency Responder is used, the 911 CTI route pattern and 9.911 translation pattern can be placed in the global On-Cluster partition.

When applied to centralized call processing deployments with large numbers of sites, the line/device approach drastically reduces the number of partitions and calling search spaces needed. For example, a deployment with 100 remote sites and 4 classes of service requires at least 401 partitions and 400 calling search spaces with the traditional approach but only 105 partitions and 104 calling search spaces with the line/device approach.

However, the line/device approach relies on the fact that you can globally identify the types of PSTN calls that must be restricted for certain classes of service (for example, local, long-distance, and international calls). If the national numbering plan of your country does not allow this global identification of the different types of calls, the efficiency of this approach (in terms of configuration savings) is lower than that indicated in the formulas above.

For example, in France the numbering plan is based on five area codes, from 01 to 05 (plus the 06 area code for cellular phones), followed by eight digits for the subscriber number. The key characteristic is that each PSTN destination is reached by dialing exactly the same number (for example, 01XXXXXXXX for Paris numbers, 04XXXXXXXX for Nice numbers, and so on), whether calling from the same local area or from a different area. This means that it is not possible to block access to long-distance calls with a single partition and a single route pattern because the concept of "long-distance call" changes depending on the area where the calling party is located. (For example, a call to 014455667788 is a local call if the caller is in Paris but a long-distance call if the caller is in Nice or Lyon.)

In such cases, you will have to configure additional sets of blocking calling search spaces and partitions, one for each area within which local and long distance calls are dialed the same way. In the example of France, you would have to define five additional blocking calling search spaces and partitions, one for each area code, as shown in [Table 10-8](#):

Table 10-8 The Line/Device Approach Applied to the French National Numbering Plan

Calling Search Space	Partition	Blocked Route Patterns
Internal_css	BlockAllNational_pt	0.0[1-6]XXXXXXXX
	BlockIntl_pt	0.00!, 0.00!#
Local01_css	BlockLD01_pt	0.0[2-6]XXXXXXXX
	BlockIntl_pt	0.00!, 0.00!#
Local02_css	BlockLD02_pt	0.0[13-6]XXXXXXXX
	BlockIntl_pt	0.00!, 0.00!#

Table 10-8 The Line/Device Approach Applied to the French National Numbering Plan

Calling Search Space	Partition	Blocked Route Patterns
Local03_css	BlockLD03_pt	0.0[124-6]XXXXXXXX
	BlockIntl_pt	0.00!, 0.00!#
Local04_css	BlockLD04_pt	0.0[1-356]XXXXXXXX
	BlockIntl_pt	0.00!, 0.00!#
Local05_css	BlockLD05_pt	0.0[1-46]XXXXXXXX
	BlockIntl_pt	0.00!, 0.00!#
LD_css	BlockIntl_pt	0.00!, 0.00!#
Intl_css	NoBlock_pt	none

Guidelines for the Line/Device Approach

Consider the following guidelines when using the line/device approach:

- The Forward Busy and Forward No Answer calling search spaces are not concatenated with the line or device calling search space.
- Set the Forward Busy and Forward No Answer calling search spaces to a global calling search space that can reach the voicemail pilot number and ports.
- Set the Forward All calling search space(s) according to your company's policy:
 - If forwarded calls must have unrestricted privileges, set the Forward All calling search space to match the site-specific device calling search space. (For additional considerations with Extension Mobility, see [Extension Mobility Considerations for the Line/Device Approach, page 10-81](#).)
 - If forwarded calls must be restricted to internal numbers only, set the Forward All calling search space to a global calling search space that can reach only internal numbers.
 - If forwarded calls must have some intermediate restriction (such as access to local PSTN numbers but not international numbers), the line/device approach might become less efficient because additional site-specific calling search spaces would be required. As a result, the total quantity of configured calling search spaces in this case would be similar to that of the traditional approach described above.
- Remember that, for this approach to work, the blocked patterns configured within the line calling search space must be at least as specific as the route patterns configured within the device calling search space. Wherever possible, Cisco recommends that you configure the blocked patterns as more specific than the routed ones, to avoid any possibility of error. Use extra care when dealing with the @ wildcard because the patterns defined within it are very specific.
- AAR is triggered when on-net DN's are dialed. Access to these DN's can be controlled by the same processes described previously. AAR uses a different calling search space for rerouted calls. In most cases, the AAR calling search space can be the same as the site-specific, unrestricted device calling search space because it can never be dialed directly by end users.

**Note**

The priority order between line and device is reversed for CTI devices (CTI route points and CTI ports). For these devices, the partitions in the device calling search space are placed before the line calling search space in the resulting calling search space. Therefore, the line/device approach cannot be applied to CTI devices such as Cisco IP SoftPhone unless you are careful not to rely solely on the concatenation order for pattern selection but instead ensure that the desired blocked pattern's precision is greater in all cases than that of the permitted pattern(s).


Extension Mobility Considerations for the Line/Device Approach

When using the Extension Mobility feature, the line/device approach provides a natural way to implement the dialing restrictions of a phone as a function of the logged-in (or logged-out) status of the phone. Typically, a logged-out phone should be restricted to calling other phones and services (such as 911), but access to local or toll calls through the PSTN are restricted. Conversely, a phone where a user is logged-in should allow calls according to that user's dialing privileges and should route those calls to the appropriate gateway (for example, a co-located branch gateway for local calls).

With the line/device approach for building classes of service, you can simply apply the same rules described in the previous section to the Extension Mobility device profile construct. Consider the following guidelines to apply calling restrictions when using Extension Mobility:

- At each site, configure the device calling search space for all IP phones to point to a site-specific partition that contains all possible PSTN route patterns and that routes them appropriately (for example, using the local gateway for emergency and local calls and a centralized gateway for long distance calls).
- Configure the line calling search space for all IP phones (or the line calling search space for the default logout device profile) to point to a global calling search space featuring blocked translation/route patterns that block all calls except those to be allowed when no user is logged in (for example, internal extensions and emergency services).
- For each Extension Mobility user, configure the line calling search space within the device profile to point to a global calling search space featuring blocked translation/route patterns to selectively block the PSTN calls that are not to be allowed for their specific class of service (for example, block only international calls). If some users must have unrestricted calling privileges, assign them to a line calling search space featuring an empty partition.

The key advantage of using the line/device approach for extension mobility is that, in a multisite deployment with centralized call processing, appropriate call routing is ensured even when users log in to IP phones located at branch sites different from their home site, as illustrated in [Figure 10-33](#).

Figure 10-33 Extension Mobility with the Line/Device Approach

As described previously in this chapter, the line calling search space configured within the device profile replaces the line calling search space configured on the physical IP phone when a user logs in through Extension Mobility. Because the call routing is correctly determined by the device calling search space, the login operation is used merely to "unlock" the phone by removing the phone's line calling search space (which contains blocked patterns) and replacing it with the device profile's line calling search space (which does not contain blocked patterns in this simplified example).

Because all the call routing is done within the device calling search space while the line calling search space only introduces blocked patterns, whenever a user logs in at a different site from their home site, they will automatically inherit all the local dialing habits for that site. For example, assume that phone DNs are defined as eight-digit numbers, but four-digit dialing is provided within each site by local translation patterns. In this case, a user roaming to a different site will not be able to dial a colleague at the home site by using only four digits because the four-digit number will now be translated according to the rules of the host site where the user logged in.


In summary, when you use the line/device approach for Extension Mobility, end-users have to adopt the dialing behavior of the site where they logged in.

Call Forwarding Considerations

When applying the Line/Device calling search space approach to a centralized call processing environment with Extension Mobility, you should be aware of the call forwarding behavior if users need to be allowed to forward all their calls to external PSTN numbers.

In [Figure 10-34](#), an Extension Mobility user is normally based in Site 1, with a device profile allows the user to place unrestricted PSTN calls and to forward all incoming calls to any PSTN number.

Figure 10-34 *Call Forwarding Considerations for Extension Mobility with the Line/Device Approach*



As described in the section on [Call-Forward Calling Search Spaces](#), [page 10-19](#), the Forward All calling search space is not concatenated with the line or the device calling search spaces and therefore needs to be set to Site1_all, which includes all PSTN routes using the Site 1 gateway.

When this user moves to Site 2 and logs into phone D, the user's device profile applies its line calling search space and Forward All calling search space(s) to the physical device. For direct PSTN calls, the calling search space used is the concatenation of the line and device calling search space, and phone D's device calling search space (Site2_all) correctly points to the Site 2 gateway.

If the user now configures the phone to forward all calls to a PSTN number, any forwarded call will use the Site1_all calling search space, which still points to the Site 1 gateway. This condition results in the following behavior:

- Incoming PSTN calls enter the IP network at the Site 1 gateway and are hairpinned back into the PSTN within the same gateway.
- Calls originating from Site 1 phones (such as phone A) are correctly forwarded to the PSTN via the Site 1 gateway.
- Calls originating from Site 2 phones (such as phone C) traverse the WAN to Site 1 and access the PSTN via the Site 1 gateway. The same behavior applies to calls originating from other sites within the same Cisco Unified CallManager cluster.

Keep this behavior in mind when designing the network and training the users.

Building Classes of Service in Cisco IOS with H.323

The following scenarios require you to define classes of service within Cisco IOS routers running the H.323 protocol:

- Cisco Unified CallManager multisite deployments with centralized call processing
- Cisco Unified CallManager Express deployments

Under normal conditions in Cisco Unified CallManager multisite deployments with centralized call processing, classes of service are implemented using partitions and calling search spaces within Cisco Unified CallManager. However, when IP WAN connectivity is lost between a branch site and the central site, Cisco SRST takes control of the branch IP phones, and all the configuration related to partitions and calling search spaces is unavailable until IP WAN connectivity is restored. Therefore, it is desirable to implement classes of service within the branch router when running in SRST mode.


Similarly, in Cisco Unified CallManager Express deployments, the router needs a mechanism to implement classes of service for the IP phones.

For both of these applications, define classes of service in Cisco IOS routers by using the class of restriction (COR) functionality (refer to [Calling Privileges in Cisco IOS with H.323 Dial Peers](#), page 10-47, for details on COR).

You can adapt the COR functionality to replicate the Cisco Unified CallManager concepts of partitions and calling search spaces by following these main guidelines:

- Define tags for each type of call that you want to distinguish.
- Assign "basic" outgoing corlists (equivalent to partitions), containing a single member tag, to the respective POTS dial peers that route each type of call.
- Assign "complex" incoming corlists (equivalent to calling search spaces), containing subsets of the member tags, to the IP phones that belong to the various classes of service.

[Figure 10-35](#) illustrates an implementation example based on SRST, where the IP phone with DN 2002 is configured to have unrestricted PSTN access, the IP phone with DN 2001 is configured to have only local PSTN access, and all other IP phones are configured to have access only to internal numbers and emergency services.

Figure 10-35 Building Classes of Service for Cisco SRST using COR

The following steps provide examples and guidelines for implementing a Cisco IOS solution like the one shown in Figure 10-35.

- Step 1** Using the **dial-peer cor custom** command, define meaningful tags for the various types of calls (Emergency, VMail, Local, LD, Intl):

```
dial-peer cor custom
  name Emergency
  name VMail
  name Local
  name LD
  name Intl
```

- Step 2** Using the **dial-peer cor list** command, define basic corlists to be used as partitions, each containing a single tag as a member:

```
dial-peer cor list EmPt
  member Emergency

dial-peer cor list VMailPt
  member VMail

dial-peer cor list LocalPt
  member Local

dial-peer cor list LDPT
  member LD
```

```
dial-peer cor list IntlPt
member Intl
```

- Step 3** Using the **dial-peer cor list** command, define more complex corlists to be used as calling search spaces, each containing a subset of the tags as members according to the classes of service needed:

```
dial-peer cor list InternalCSS
member Emergency
member VMail
```

```
dial-peer cor list LocalCSS
member Emergency
member VMail
member Local
```

```
dial-peer cor list LDCSS
member Emergency
member VMail
member Local
member LD
```

```
dial-peer cor list IntlCSS
member Emergency
member VMail
member Local
member LD
member Intl
```

- Step 4** Using the command **corlist outgoing** *corlist-name*, configure the basic “partition” corlists as outgoing corlists assigned to the corresponding POTS dial peers:

```
dial-peer voice 100 pots
corlist outgoing EmPt
destination-pattern 911
no digit-strip
direct-inward-dial
port 1/0:23
```

```
dial-peer voice 101 pots
corlist outgoing VMailPt
destination-pattern 914085551234
forward-digits 11
direct-inward-dial
port 1/0:23
```

```
dial-peer voice 102 pots
corlist outgoing LocalPt
destination-pattern 9[2-9].....
forward-digits 7
direct-inward-dial
port 1/0:23
```

```
dial-peer voice 103 pots
corlist outgoing LDPT
destination-pattern 91[2-9]..[2-9].....
forward-digits 11
direct-inward-dial
port 1/0:23
```

```
dial-peer voice 104 pots
corlist outgoing IntlPt
destination-pattern 9011T
prefix-digits 011
direct-inward-dial
```



```
port 1/0:23
```

- Step 5** Using the **cor incoming** command within the **call-manager-fallback** configuration mode, configure the complex corlists acting as "calling search spaces" to be incoming corlists assigned to the various phone DNs:

```
call-manager-fallback
 cor incoming InternalCSS default
 cor incoming LocalCSS 1 3001 - 3003
 cor incoming LDCSS 2 3004
 cor incoming IntlCSS 3 3010
```

When deploying COR for SRST, keep in mind the following limitations:

- In SRST version 2.0, available on Cisco IOS Release 12.2(8)T or later, the maximum number of **cor incoming** statements allowed under **call-manager-fallback** is 5 (plus the default statement).
- In SRST version 3.0, available on Cisco IOS Release 12.3(4)T or later, the maximum number of **cor incoming** statements allowed under **call-manager-fallback** is 20 (plus the default statement).

Therefore, if the phone DNs of users with non-default privileges are not consecutive and the SRST site is relatively large, you might have to reduce the number of classes of service in SRST mode to accommodate all the DNs without exceeding these limitations.

Although the preceding example is based on Cisco SRST, the same concepts can be applied to a Cisco Unified CallManager Express deployment, except for the following considerations:

- With Cisco Unified CallManager Express, the corlist expressing the class of service (equivalent to a calling search space) can be assigned directly to the individual IP phones by using the **cor {incoming | outgoing} corlist-name** command under the **ephone-dn dn-tag** configuration mode.
- According to COR general rules, all IP phones for which no corlist is configured have unrestricted access to all dial peers, regardless of their outgoing corlist. Cisco Unified CallManager Express has no mechanism equivalent to the **cor incoming corlist-name default** command, which applies default restrictions to all phones.

Deploying Call Coverage

Call coverage functionality is a key feature in many IP telephony deployments. Many customer-focused service companies have to route customer calls to the appropriate service representatives expeditiously. This section focuses on design guidelines for using the hunting mechanism based on hunt pilots, hunt lists, and line groups in Cisco Unified CallManager Release 4.1 to manage call distribution, and it covers the following main topics:

- [Deploying Call Coverage in a Multisite Centralized Call Processing Model, page 10-88](#)
- [Deploying Call Coverage in a Multisite Distributed Call Processing Model, page 10-89](#)




Note

Call coverage functionality does not offer call queues per se, and the caller is presented with ringback tone until a destination is found for the call. To provide prompting, music on hold, and so forth, Cisco offers many contact center technologies such as the Cisco Unified Customer Voice Portal (CVP). For more information on the contact center technologies available from Cisco, refer to the documentation at <http://www.cisco.com/go/designzone>.

Deploying Call Coverage in a Multisite Centralized Call Processing Model

Figure 10-36 shows an example of configuring hunt lists for a multisite centralized call processing deployment. This example assumes that the hunt pilot call will be distributed first through the remote office operators. If the call is not answered or is rejected due to call admission control, the call will then be routed to central-site operators or voicemail.

Figure 10-36 Call Coverage Between Multiple Sites in a Centralized Call Processing Deployment



In centralized IP telephony systems, features such as Automated Alternate Routing (AAR) and Survivable Remote Site Telephony (SRST) may be enabled for high availability. Consider the following guidelines when deploying call coverage functionality with AAR or SRST features enabled:

- Automated Alternate Routing (AAR)

The line group members can be assigned in different locations and regions. Call admission control implemented through locations works as expected. However, a call being distributed from a hunt pilot will not use AAR to reroute a call if Cisco Unified CallManager blocks the call to one of its line group members due to insufficient WAN bandwidth. Instead, Cisco Unified CallManager distributes the call to the next available member or next available line group.



Note

Calls distributed by a hunt pilot can use AAR in Cisco Unified CallManager Release 4.1(3) and above. However, Cisco strongly recommends that you use AAR only when using voicemail ports within line groups.


- Survivable Remote Site Telephony (SRST)
 - When Cisco Unified CallManager receives a call for the hunt pilot, and if some of its line group members are at the remote sites operating in SRST mode, then Cisco Unified CallManager skips those members and distributes the call to the next available line group member. From the perspective of Cisco Unified CallManager, the members operating in SRST mode are unregistered, and hunt pilot calls are not forwarded to unregistered members.
 - When a router operating in SRST mode receives a call for the hunt pilot, call coverage functionality is unavailable. The call fails if no configuration is added to reroute the call to a registered and available extension. You can use the **alias** or the **default-destination** command under the **call-manager-fallback** mode in Cisco IOS to reroute the call destined for the hunt pilot to an operator extension or to voicemail.

Deploying Call Coverage in a Multisite Distributed Call Processing Model

Beginning with Cisco Unified CallManager Release 4.1, route groups can no longer be added to hunt lists. Thus, a hunt list cannot be used to send the calls to other clusters or to a remote gateway. But the hunt option settings in Hunt Pilot, introduced in Cisco Unified CallManager Release 4.1, can be used to match a route pattern that in turn points to gateways or trunks.

Figure 10-37 shows an example of configuring hunt lists for a distributed call processing deployment with an intercluster trunk. This example assumes that the hunt pilot call is first distributed within Cluster A. If the call is not answered, it is rerouted to Cluster B for call distribution using the Forward Hunt No Answer setting, which matches a route pattern. The route pattern, in turn, points to an intercluster trunk to Cluster B.

Figure 10-37 Call Coverage Between Clusters in a Distributed Call Processing Deployment



**Tip**

In distributed call processing deployments, load sharing of incoming hunt pilot calls can be managed using Cisco VoIP gateways and gatekeepers. In the event that the call is not answered within one cluster, it can overflow to another cluster for service. Calls can also be sent through gateways or trunks to IVR treatment. Tool Command Language (TCL) IVR applications can be implemented on Cisco IOS gateways.

Guidelines

If calls are distributed across multiple clusters in a distributed call processing deployment, the route patterns should be properly configured to account for any digit transformations that are done on the outbound or inbound route group devices. If digit transformations are not done, then the configured route patterns and hunt pilot should be the same on all clusters, otherwise the calls will not be distributed appropriately.

Hunt Pilot Scalability

Cisco recommends using the following guidelines when deploying call coverage using top-down, circular, and longest-idle algorithms:

- The Cisco Unified CallManager cluster supports a maximum of 15,000 hunt list devices.
- The hunt list devices may be a combination of 1500 hunt lists with 10 IP phones in each hunt list, or a combination of 750 hunt lists with 20 IP phones in each hunt list. However, an increase in a number of hunt lists can require increasing the dial plan initialization timer specified in the Cisco CallManager service parameters. Cisco recommends setting the dial plan initialization timer to 600 seconds if 1500 hunt lists are configured.

**Note**

When using the broadcast algorithm for call coverage, the number of hunt list devices is limited by the number of busy hour call attempts (BHCA). Note that a BHCA of 10 on a hunt pilot pointing to a hunt list or hunt group containing 10 phones and using the broadcast algorithm is equivalent to 10 phones with a BHCA of 10.

- Cisco recommends having a maximum of 35 directory numbers in a single line group configured to send the calls simultaneously to all DNs. Additionally, the number of broadcast line groups depends on the BHCC. If there are multiple broadcast line groups in a Cisco Unified CallManager system, the number of maximum directory numbers in a line group must be less than 35. The number of busy hour call attempts (BHCA) for all the broadcast line groups should not exceed 35 calls set up per second



CHAPTER 11

Emergency Services

Last revised on: February 13, 2008

Emergency services are of great importance in the proper deployment of a voice system. This chapter presents a summary of the following major design considerations essential to planning for emergency calls:

- [Planning for 911 Functionality, page 11-1](#)
- [Gateway Considerations, page 11-11](#)
- [Cisco Emergency Responder Considerations, page 11-13](#)

This chapter presents some information specific to the 911 emergency networks as deployed in Canada and the United States. Many of the concepts discussed here are adaptable to other locales. Please consult with your local telephony network provider for appropriate implementation of emergency call functionality.

In the United States, some states have already enacted legislation covering the 911 functionality required for users in a multi-line telephone system (MLTS). The National Emergency Number Association (NENA) has also produced the *Technical Information Document on Model Legislation Enhanced 911 for Multi-Line Telephone Systems*, available online at

http://www.nena.org/media/files/MLTS_ModLeg_Nov2000.pdf

The Federal Communications Commission (FCC) has also drafted a proposed new section to title 47, part 68, section 319, which is available at

<http://www.apcointl.org/about/pbx/worddocs/mltspart68.doc>

This chapter assumes that you are familiar with the generic 911 functionality available to residential PSTN users in North America. For more information on the subject, refer to the following URL for a description of the current state of E911 services in North America:

<http://www.nena.org/florida/Directory/911Tutorial%20Study%20Guide.pdf>

Planning for 911 Functionality

This section highlights some of the functionality requirements for lifeline calls in multi-line telephone systems (MLTS). In the context of this section, lifeline calls are 911 calls serviced by the North American public switched telephone network (PSTN).

When planning an MLTS deployment, first establish all of the physical locations where phone services are needed. The locations can be classified as follows:

- Single building deployments, where all users are located in the same building.
- Single campus deployments, where the users are located in a group of buildings situated in close proximity.
- Multisite deployments, where users are distributed over a wide geographical area and linked to the telephony call processing site via WAN connectivity.

The locations, or type of deployment, affect the criteria used to design and implement 911 services. The following sections describe the key criteria, along with typical and exceptional situations for each. When analyzing and applying these criteria, consider how they are affected by the phone locations in your network.

Public Safety Answering Point (PSAP)

The public safety answering point (PSAP) is the party responsible for answering the 911 call and arranging the appropriate emergency response, such as sending police, fire, or ambulance teams. The physical location of the phone making the 911 call is the primary factor in determining the appropriate PSAP for answering that call. Generally, each building is serviced by one local PSAP.

To determine the responsible PSAP for a given location, contact a local public safety information service such as the local fire marshal or police department. Also, the phone directory of the local exchange carrier usually lists the agency responsible for servicing 911 calls in a given area.

Typical Situation

- For a given street address, there is only one designated PSAP.
- For a given street address, all 911 calls are routed to the same PSAP.

Exceptional Situation

- The physical size of the campus puts some of the buildings in different PSAP jurisdictions.
- Some of the 911 calls need to be routed to an on-net location (campus security, building security).

911 Network Service Provider

After identifying the responsible PSAPs, you must also identify the 911 network service providers to which each PSAP is connected. It is commonly assumed that PSAPs receive 911 phone calls from the PSTN, but that is not the case. Instead, 911 calls are carried over dedicated, regionally significant networks, and each PSAP is connected to one or more such regional networks. In the majority of cases, the incumbent Local Exchange Carrier (LEC) is the 911 network service provider for a PSAP. Some exceptions include military installations, university campuses, federal or state parks, or other locations where the public safety responsibility falls outside the jurisdiction of the local authorities and/or where a private network is operated by an entity other than a public local exchange carrier.

If you are in doubt about the 911 network service provider for a given PSAP, contact the PSAP directly to verify the information.

Typical Situation

- For a given street address, the 911 network service provider is the incumbent Local Exchange Carrier (LEC). For a location served by Phone Company X, the corresponding PSAP is also served by Phone Company X.
- All 911 calls are routed directly to an off-net location, or all 911 calls are routed directly to an on-net location.

Exceptional Situation

- The local exchange carrier (LEC) through which the MLTS interfaces to the PSTN is *not* the same LEC that serves as 911 network service provider to the PSAP. (For example, the phone system is served by Phone Company X, but the PSAP is connected to Phone Company Y.) This situation might require either a special arrangement between the LECs or special, dedicated trunks between the phone system and the PSAP's 911 network service provider.
- Some LECs may not accept 911 calls on their networks. If this is the case, the only two options are to change LECs or to establish trunks (dedicated to 911 call routing) connected to a LEC that can route 911 calls to the appropriate PSAPs.
- Some (or all) of the 911 calls have to be routed to an on-net location such as campus security or building security. This situation can easily be accommodated during the design and implementation phases, but only if the destination of 911 calls for each phone has been properly planned and documented.

Interface Points into the Appropriate 911 Networks

For larger telephony systems, 911 connectivity might require many interface points. Typically, more than one E911 selective router is used within a LEC's territory, and these routers usually are *not* interconnected.

For example, an enterprise with a large campus could have the following situation:

- Building A located in San Francisco
- Building B located in San Jose
- San Francisco Police Department and San Jose Police Department are the appropriate PSAPs
- San Francisco Police Department and San Jose Police Department are served by the same 911 network service provider
- However, San Francisco Police Department and San Jose Police Department are served by different E911 selective routers operated by that same 911 network service provider!

This type of situation would require two separate interface points, one per E911 selective router. The information pertaining to the E911 selective router territories is generally kept by the incumbent LEC, and the local account representative for that LEC should be able to provide an enterprise customer with the pertinent information. Many LECs also provide the services of 911 subject matter experts who can consult with their own account representatives on the proper mapping of 911 access services.

Typical Situation

- For single-site deployments or campus deployments, there is usually only one PSAP for 911 calls.
- If access to only one PSAP is required, then only one interface point is required. Even if access to more than one PSAP is required, they might be reachable from the same E911 selective router, through the same centralized interface. If the enterprise's branch sites are linked via a WAN

(centralized call processing), it is desirable to give each location its own local (that is, located inside each branch office) access to 911 to prevent 911 isolation during WAN failure conditions where Survivable Remote Site Telephony (SRST) operation is activated.

Exceptional Situation

- The physical size of the campus puts some of the buildings in different PSAP jurisdictions, *and*
- Some of the 911 calls have to be routed to different E911 selective routers, through different interface points.



Note

Some of the information required to establish the geographical territories of PSAPs and E911 selective routers is available online or from various competitive local exchange carrier (CLEC) information web sites. (For example, <https://clec.att.com/clec/hb/shell.cfm?section=782> provides some valuable data about the territory covered by SBC/Pacific Bell.) However, Cisco strongly recommends that you obtain proper confirmation of the appropriate interface points from the LEC prior to the design and implementation phases of 911 call routing.

Interface Type

In addition to providing voice communications, the interfaces used to present 911 calls to the network must also provide identification data about the calling party.

Automatic Number Identification (ANI) refers to the E.164 number of the calling party, which is used by networks to route a 911 call to the proper destination. This number is also used by the PSAP to look up the Automatic Location Identification (ALI) associated with a call.

911 calls are source-routed, which means that they are routed according to the calling number. Even though different locations are all dialing the same number (911), they will reach different PSAPs based on their location of origin, which is represented by the ANI (calling number).

You can implement 911 call functionality with either of the following interface types:

- Dynamic ANI assignment
- Static ANI assignment

While dynamic ANI assignment scales better (because it supports multiple ANIs) and lends itself to all but the smallest of applications, static ANI assignment can be used in a wider variety of environments, from the smallest to the largest systems.

Dynamic ANI (Trunk Connection)

The dynamic aspect of ANI refers to the fact that a system has many phones sharing access to the 911 network across the same interface, and the ANI transmitted to the network might need to be different for each call.

There are two main types of dynamic ANI interfaces:

- Integrated Services Digital Network Primary Rate Interface (ISDN-PRI, or simply PRI)
- Centralized Automatic Message Accounting (CAMA).

PRI

This type of interface usually connects a telephony system to a PSTN Class 5 switch. The calling party number (CPN) is used at call setup time to identify the E.164 number of the calling party.

Most LECs treat the CPN differently when a call is made to 911. Depending upon the functionality available in the Class 5 switch and/or upon LEC or government policy, the CPN may not be used as the ANI for 911 call routing. Instead, the network may be programmed to use the listed directory number (LDN) or the bill-to number (BTN) for ANI purposes.

If the CPN is not used for ANI, then 911 calls coming from a PRI interface all look the same to the 911 network because they all have the same ANI, and they are all routed to the same destination (which might not be the appropriate one).

Some LECs offer a feature to provide CPN transparency through a PRI interface for 911 calls. With this feature, the CPN presented to the Class 5 switch at call setup is used as ANI to route the call. The feature name for this functionality varies, depending on the LEC. (For example, SBC calls it Inform 911 in California.)

**Note**

The CPN *must* be a routable E.164 number, which means that the CPN must be entered in the routing database of the associated E911 selective router.

**Note**

For Direct Inward Dial (DID) phones, the DID number could be used as the ANI for 911 purposes, but only if it is properly associated with an Emergency Service Number in the 911 service provider's network. For non-DID phones, use another number. (See [Emergency Location Identification Number Mapping, page 11-7](#), for more information.)

Many Class 5 switches are connected to E911 selective routers through trunks that do not support more than one area code. In such cases, if PRI is used to carry 911 calls, then the only 911 calls that will be routed properly are those whose CPN (or ANI) have the same Numbering Plan Area (NPA) as the Class 5 switch.

Example

An MLTS is connected to a Class 5 switch in area code 514 (NPA = 514). If the MLTS were to send a 911 call on the PRI trunk, with a CPN of **450.555.1212**, the Class 5 switch would send the call to the E911 selective router with an ANI of **514.555.1212** (instead of the correct **450.555.1212**), yielding inappropriate routing and ALI lookup.

To use PRI properly as a 911 interface, the system planner must ensure that the CPN will be used for ANI and must properly identify the range of numbers (in the format NPA NXX TNTN) acceptable on the link. For example, if a PRI link is defined to accept ANI numbers within the range 514 XXX XXXX, then only calls that have a Calling Party Number with NPA = 514 will be routed appropriately.

CAMA

Centralized Automatic Message Accounting (CAMA) trunks also allow the MLTS to send calls to the 911 network, with the following differences from the PRI approach:

- CAMA trunks are connected directly into the E911 selective router. Extra mileage charges may apply to cover the distance between the E911 selective router and the MLTS gateway point.
- CAMA trunks support 911 calls only. The capital and operational expenses associated with the installation and operation of CAMA trunks support 911 traffic only.

- CAMA trunks for the MLTS market may be limited to a fixed area code, and the area code is typically implied (that is, not explicitly sent) in the link protocol. The connection assumes that all calls share the same deterministic area code, therefore only 7 or 8 digits are sent as ANI.

**Note**

Cisco supports CAMA-based 911 functionality through the VIC-2CAMA, VIC-2FXO, and VIC-4FXO trunk cards.

Static ANI (Line Connection)

Static ANI provides a line (rather than a trunk) connection to the PSTN, and the ANI of the line is associated with all 911 calls made on that line, regardless to the CPN of the calling phone. A plain old telephone service (POTS) line is used for this purpose.

POTS lines are one of the simplest and most widely supported PSTN interfaces. A POTS line usually comes fully configured to accept 911 calls. In addition, the existing E911 infrastructure supports 911 calls from POTS lines very well.

The POTS approach has the following attributes:

- The operational costs associated with a POTS line are low.
- The POTS line can even serve as a backup line in case of power failure.
- The POTS line number can be used as the callback number entered into the ALI database.
- POTS lines represent the lowest cost 911 support for locations where user density does not justify local PRI or CAMA access into the PSTN.
- POTS lines are ubiquitous in PSTN installations.

All outgoing 911 calls through this type of interface are treated the same by the E911 network, and the tools that enable Cisco Unified CallManager to control the ANI presented to the E911 network (such as calling party transformation masks) are irrelevant because the ANI can be only the POTS line's number.

Emergency Response Location Mapping

The National Emergency Number Association (NENA) has recently proposed model legislation to be used by state and federal agencies in enacting the rules that govern 911 in enterprise telephony systems. One of the concepts in the NENA proposal is that of the emergency response location (ERL), which is defined as:

A location to which a 911 emergency response team may be dispatched. The location should be specific enough to provide a reasonable opportunity for the emergency response team to quickly locate a caller anywhere within it.

Rather than having to identify each phone's location individually, the requirement allows for the grouping of phones into a "zone," the ERL. The maximum size of the ERL may vary, depending upon local implementation of the legislation, but we will use 7000 square feet (sq ft) as a basis for discussion in this section. (The concepts discussed here are independent of the maximum ERL size that may be allowed in any given state or region.)

An emergency location identification number (ELIN) is associated with each ERL. The ELIN is a fully qualified E.164 number, used to route the call within the E911 network. The ELIN is sent to the E911 network for any 911 call originating from the associated ERL. This process allows more than one phone to be associated with the same fully qualified E.164 number for 911 purposes, and it can be applied to DID and non-DID phones alike.

**Note**

This document does not attempt to present the actual requirements of any legislation. Rather, the information and examples presented here are for the purposes of discussion only. The system planner is responsible for verifying the applicable local requirements.

For example, assume a building has a surface area of 70,000 sq ft and 100 phones. In planning for 911 functionality, the building can be divided into 10 zones (ERLs) of 7000 sq ft each, and each phone can be associated with the ERL where it is located. When a 911 call is made, the ERL (which could be the same for multiple phones) is identified by sending the associated ELIN to the PSAP. If the phones were evenly distributed in this example, each group of 10 phones would have the same ERL and, therefore, the same ELIN.

The various legislations define a minimum number of phones (for example, 49) and a minimum surface area (for example, 40,000 sq ft) below which the requirements for MLTS 911 are not applicable. But even if the legislation does not require 911 functionality for a given enterprise, it is always best practice to provision for it.

Emergency Location Identification Number Mapping

In general, you must associate a single fully qualified E.164 number, known as the emergency location identification number (ELIN), with each ERL. (However, if using Cisco Emergency Responder, you can configure more than one ELIN per ERL.) The ELIN is used to route the call across the E911 infrastructure and is used by the PSAP as the index into the ALI database.

ELINs must meet the following requirements:

- They must be routable across the E911 infrastructure. (See the examples in the section on [Interface Type, page 11-4](#).) If an ELIN is not routable, 911 calls from the associated ERL will, at best, be handled according to the default routing programmed in the E911 selective router.
- Once the ERL-to-ELIN mapping of an enterprise is defined, the corresponding ALI records must be established with the LEC so that the ANI and ALI database records serving the PSAP can be updated accurately.

The ELIN mapping process can be one of the following, depending on the type of interface to the E911 infrastructure for a given ERL:

- Dynamic ANI interface

With this type of interface, the calling party number identification passed to the network is controlled by the MLTS. The telephony routing table of the MLTS is responsible for associating the correct ELIN with the call, based on the calling phone's ERL. Within Cisco Unified CallManager, the calling party number can be modified by using transformation masks for calls made to 911. For example, all phones located in a given ERL can share the same calling search space that lists a partition containing a translation pattern (911) and a calling party transformation mask that would replace the phone's CPN with the ELIN for that location.

- Static ANI interface

With this type of interface, the calling party number identification passed to the network is controlled by the PSTN. This is the case if the interface is a POTS line. The ELIN is the phone number of the POTS line, and no further manipulation of the phone's calling party identification number is possible.

PSAP Callback

The PSAP might have to reach the caller after completion of the initial conversation. The PSAP's ability to call back relies on the information that it receives with the original incoming call.

The delivery of this information to the PSAP is a two-part process:

1. The Automatic Number Identification (ANI) is first sent to the PSAP. The ANI is the E.164 number used to route the call. In our context, the ANI received at the PSAP is the ELIN that the MLTS sent.
2. The PSAP then uses the ANI to query a database and retrieve the Automatic Location Identification (ALI). The ALI provides the PSAP attendant with information such as:
 - Caller's name
 - Address
 - Applicable public safety agency
 - Other optional information, which could include callback information. For example, the phone number of the enterprise's security service could be listed, to aid in the coordination of rescue efforts.

Typical Situation

- The ANI information is used for PSAP callback, which assumes that the ELINs are dialable numbers.
- The ELINs are PSTN numbers associated with the MLTS. If someone calls the ELIN from the PSTN, the call will terminate on an interface controlled by the MLTS.
- It is the responsibility of the MLTS system administrator to program the call routing so that calls made to any ELIN in the system will ring a phone (or multiple phones) in the immediate vicinity of the associated ERL.
- Once the ERL-to-ELIN mapping is established, it needs be modified only when there are changes to the physical situation of the enterprise. If phones are simply added, moved, or deleted from the system, the ERL-to-ELIN mapping and its associated ANI/ALI database records need not be changed.

Exceptional Situation

- Callback to the immediate vicinity of the originating ERL may be combined with (or even superseded by) routing the callback to an on-site emergency desk, which will assist the PSAP in reaching the original caller and/or provide additional assistance with the emergency situation at hand.
- The situation of the enterprise could change, for example, due to area code splits, city or county service changes requiring a new distribution of the public safety responsibilities, new buildings being added, or any other change that would affect the desired routing of a call for 911 purposes. Any of these events could require changes in the ERL-to-ELIN mapping and the ANI/ALI database records for the enterprise.

Nomadic Phone Considerations

All discussions in this chapter thus far have relied upon the assumption that the phone locations are static. If, however, phones are moved across ERL boundaries, then 911 calls from the newly relocated phone will not be routed correctly. Because it is now physically located in a different ERL, the phone should use the ELIN of its current ERL. If the configuration is not changed in the Cisco Unified CallManager database, the following events will occur:

- The ELIN of the previous ERL will be used to route calls on the E911 infrastructure.
- The egress point from the IP network to the E911 infrastructure might be incorrect.
- The callback functionality provided to the PSAP might reach the wrong destination.
- The ALI information provided to the PSAP might result in the dispatching of emergency response personnel to the wrong location.
- The location-based call admission control for the phone might not properly account for the WAN bandwidth usage of the phone, yielding possible over-subscription or under-subscription of WAN bandwidth resources.

The only way to remedy this situation is to manually update the phone's configuration (including its calling search space and location information) in Cisco Unified CallManager to reflect its new physical location.

Cisco Emergency Responder

Ease of administration for moves, adds, and changes is one of the key advantages of IP telephony technology. To provide for moves, adds, and changes that automatically update 911 information without user intervention, Cisco has developed a product called the Cisco Emergency Responder (Cisco ER).

Cisco ER provides the following primary functionality:

- Dynamic association of a phone to an ERL, based on the detected physical location of the phone.
- Dynamic association of the ELIN to the calling phone, for callback purposes. In contrast to non-ER scenarios outlined in preceding sections, Cisco ER enables the callback to ring the exact phone that initiated the 911 call.
- On-site notification to designated parties (by pager, web page, or phone call) to inform them that there is an emergency call in progress. The pager and web page notifications include the calling party name and number, the ERL, and the date and time details associated with the call. Phone notification provides the information about the calling number from which the emergency call was placed.

For more information on Cisco ER, refer to the section on [Cisco Emergency Responder Considerations, page 11-13](#), and to the Cisco ER product documentation available online at

http://www.cisco.com/en/US/products/sw/voicesw/ps842/tsd_products_support_series_home.html

The key functionality of Cisco ER relies on the detection of the phone's location by discovery of the network port (Layer 2 port, such as a Fast Ethernet switched port) from which the phone made the 911 call. The discovery mechanism relies on two main assumptions:

- The wired infrastructure of the enterprise is well established and does not change sporadically.
- The infrastructure is available for Cisco ER to browse; that is, Cisco ER can establish Simple Network Management Protocol (SNMP) sessions to the underlying network infrastructure and can scan the network ports for the discovery of connected phones.

Once Cisco ER discovers the originating port for the call, it associates the call with the pre-established ERL for the location of that port. This process also yields an association with a pre-established ELIN for the location and the selection of the appropriate egress point to the E911 infrastructure, based on the originating ERL.

With Cisco ER, the ERL-to-ELIN mapping process described in the preceding sections still applies, but with one variation: without Cisco ER, each ERL is associated with only one ELIN, but Cisco ER allows for the use of two or more ELINs per ERL. The purpose of this enhancement is to cover the specific case of more than one 911 call originating from a given ERL within the same general time period, as illustrated by the following examples.

Example 1

- Phone A and phone B are both located within ERL X, and ERL X is associated with ELIN X.
- Phone A makes a 911 call at 13:00 hours. ELIN X is used to route the call to PSAP X, and PSAP X answers and releases the call. Then, at 13:15 hours, phone B makes a 911 call. ELIN X is again used to route the call to PSAP X.
- PSAP X, after releasing the call from phone B, decides to call back phone A for further details pertaining to phone A's original call. The PSAP dials ELIN X, and gets phone B (instead of the desired phone A).

To work around this situation, Cisco ER allows you to define a pool of ELINs for each ERL. This pool provides for the use, in a round-robin fashion, of a distinct ELIN for each successive call. With the definition of two ELINs for ERL X in our example, we now have the situation described in Example 2.

Example 2

- Phone A and phone B are both located within ERL X. ERL X is associated with both ELIN X1 and ELIN X2.
- Phone A makes a 911 call at 13:00 hours. ELIN X1 is used to route the call to PSAP X, and PSAP X answers and releases the call. Then, at 13:15 hours, phone B makes a 911 call, and ELIN X2 is used to route this call to PSAP X.
- PSAP X, after releasing the call from phone B, decides to call back phone A for further details pertaining to phone A's original call. The PSAP dials ELIN X1 and gets phone A.

Of course, if a third 911 call were made but there were only two ELINs for the ERL, the situation would allow for callback functionality to properly reach only the last two callers in the sequence.

Emergency Call String

It is highly desirable to configure a dial plan so that the system easily recognizes emergency calls, whether an access code (for example, 9) is used or not. The emergency string for North America is generally 911. Cisco strongly recommends that you configure the system to recognize both the strings 911 and 9911.

Cisco also strongly recommends that you explicitly mark the emergency route patterns with Urgent Priority so that Cisco Unified CallManager does *not* wait for the inter-digit timeout (Timer T.302) before routing the call.

Other emergency call strings may be supported concurrently on your system. Cisco highly recommends that you provide your telephony system users with training on the selected emergency call strings.

Also, it is highly desirable that users be trained to react appropriately if they dial the emergency string by mistake. In North America, 911 may be dialed in error by users trying to access a long distance number through the use of 9 as an access code. In such a case, the user should remain on the line to

confirm that there is no emergency, and therefore no need to dispatch emergency personnel. Cisco ER's on-site notification capabilities can help in identifying the phone at the origin of such spurious 911 calls by providing detailed accounts of all calls made to 911, including calls made by mistake.

Gateway Considerations

Consider the following factors when selecting the gateways to handle emergency calls for your system:

- [Gateway Placement, page 11-11](#)
- [Gateway Blocking, page 11-11](#)
- [Answer Supervision, page 11-12](#)
- [Answer Supervision, page 11-12](#)

Gateway Placement

Within the local exchange carrier (LEC) networks, 911 calls are routed over a locally significant infrastructure based on the origin of the call. The serving Class 5 switches are connected either directly to the relevant PSAP for their location or to an E911 selective router, which itself is connected to a group of PSAPs significant for its region.

With Cisco's IP-based enterprise telephony architecture, it is possible to route calls on-net to gateways that are remotely situated. As an example, a phone located in San Francisco could have its calls carried over an IP network to a gateway situated in San Jose, and then sent to the LEC's network.

For 911 calls, it is critical to choose the egress point to the LEC network so that emergency calls are routed to the appropriate local PSAP. In the example above, a 911 call from the San Francisco phone, if routed to a San Jose gateway, could not reach the San Francisco PSAP because the San Jose LEC switch receiving the call does not have a link to the E911 selective router serving the San Francisco PSAP. Furthermore, the San Jose area 911 infrastructure would not be able to route the call based on a San Francisco calling party number.

As a rule of thumb, route 911 calls to a gateway physically co-located with the originating phone. Contact the LEC to explore the possibility of using a common gateway to aggregate the 911 calls from multiple locations. Be aware that, even if the 911 network in a given region lends itself to using a centralized gateway for 911 calls, it might be preferable to rely on gateways co-located with the calling phones to prevent 911 call routing from being impacted during WAN failures.

Gateway Blocking

It is highly desirable to protect 911 calls from "all trunks busy" situations. If a 911 call needs to be connected, it should be allowed to proceed even if other types of calls are blocked due to lack of trunking resources. To provide for such situations, you can dedicate an explicit trunk group just for 911 calls.

It is acceptable to route emergency calls exclusively to an emergency trunk group. Another approach is to send emergency calls to the same trunk group as the regular PSTN calls (if the interface permits it), with an alternative path to a dedicated emergency trunk group. This latter approach allows for the most flexibility.

As an example, we can point emergency calls to a PRI trunk group, with an alternate path (reserved exclusively for emergency calls) to POTS lines for overflow conditions. If we put 2 POTS lines in the alternate trunk group, we are guarantying that a minimum of two simultaneous 911 calls can be routed in addition to any calls that were allowed in the main trunk group.

If the preferred gateway becomes unavailable, it may be acceptable to overflow emergency calls to an alternate number so that an alternate gateway is used. For example, in North America calls dialed as 911 could overflow to an E.164 (non-911) local emergency number. This approach does not take advantage of the North American 911 network infrastructure (that is, there is no selective routing, ANI, or ALI services), and it should be used only if it is acceptable to the applicable public safety authorities and only as a last resort to avoid blocking the emergency call due to a lack of network resources.

Answer Supervision

Under normal conditions, calls made to an emergency number should return answer supervision upon connection to the PSAP. The answer supervision may, as with any other call, trigger the full-duplex audio connection between the on-net caller and the egress interface to the LEC's network.

With some North American LECs, answer supervision might not be returned when a "free" call is placed. This may be the case for some toll-free numbers (for example, 800 numbers). In exceptional situations, because emergency calls are considered "free" calls, answer supervision might not be returned upon connection to the PSAP. You can detect this situation simply by making a 911 test call. Upon connection to the PSAP, if audio is present, the call timer should record the duration of the ongoing call; if the call timer is absent, it is very likely that answer supervision was not returned. If answer supervision is not returned, Cisco highly recommends that you contact the LEC and report this situation because it is most likely not the desired functionality.

If this situation cannot be rectified by the Local Exchange Carrier, it would be advisable to configure the egress gateway *not* to require answer supervision when calls are placed to the LEC's network, and to cut through the audio in both directions so that progress indicator tones, intercept messages, and communications with the PSAP are possible even if answer supervision is not returned.

By default, Cisco IOS-based H.323 gateways must receive answer supervision in order to connect audio in both directions. To forego the need for answer supervision on these gateways, use the following commands:

- **progress_ind alert enable 8**

This command provides the equivalent of receiving a progress indicator value of 8 (in-band information now available) when alerting is received. This command allows the POTS side of the gateway to connect audio toward the origin of the call.

- **voice rtp send-recv**

This command allows audio cut-through in both the backward and forward directions before a Connect message is received from the destination switch. This command affects all Voice over IP (VoIP) calls when it is enabled.

Be advised that, in situations where answer supervision is not provided, the call detail records (CDRs) will not accurately reflect the connect time or duration of 911 calls. This inaccuracy can impede the ability of a call reporting system to document the relevant statistics properly for 911 calls.

In all cases, Cisco highly recommends that you test 911 call functionality from all call paths and verify that answer supervision is returned upon connection to the PSAP.

Cisco Emergency Responder Considerations

Device mobility brings about special design considerations for emergency calls. Cisco Emergency Responder (Cisco ER) can be used to track device mobility and to adapt the system's routing of emergency calls based on a device's dynamic physical location.

Emergency Responder Version Compatibility with Cisco Unified CallManager

Emergency Responder 1.3(1) is required for compatibility with Cisco Unified CallManager 5.0. For a full description of the compatibility between Emergency Responder and Cisco Unified CallManager software versions, refer to the *Cisco Emergency Responder Administration Guide 1.3(1)*, available at

<http://www.cisco.com>

Device Mobility Across Call Admission Control Locations

In a centralized call processing deployment, Cisco ER cannot fully support device movement across call admission control locations because Cisco Unified CallManager does not know about device movements. For example, if you physically move a phone from Branch A to Branch B but the phone's call admission control location remains the same (for example, Location_A), then it is possible that calls made to 911 from that phone would be blocked due to call admission control denial if all available bandwidth to Location_A is in use for other calls. This call blocking occurs even if the phone, now in location B, is physically co-located with the gateway used to connect to the PSAP for location B.

For the same reasons, Cisco ER cannot support device movement across gatekeeper-controlled call admission control zones. However, Cisco ER can fully support device movements within a call admission control location.

In centralized call processing deployments, Cisco ER automatically supports device movement within branches. However, if a device is moved between branches, manual intervention is required to adapt the device's location and region parameters before Cisco ER can fully support 911 calls.

Default Emergency Response Location

If Cisco ER cannot directly determine the physical location of a phone, it assigns a default emergency response location (ERL) to the call. The default ERL points all such calls to a specific PSAP. Although there is no universal recommendation as to where calls should be sent when this situation occurs, it is usually desirable to choose a PSAP that is centrally located and that offers the largest public safety jurisdiction. It is also advisable to populate the ALI records of the default ERL's emergency location identification numbers (ELINs) with contact information for the enterprise's emergency numbers and to offer information about the uncertainty of the caller's location. In addition, it is advisable to mark those ALI records with a note that a default routing of the emergency call has occurred.

Soft Clients

In cases where soft clients such as Cisco IP Communicator are used within the enterprise, Cisco ER can provide device mobility support. However, if the soft client is used outside the boundaries of the enterprise (for example, VPN access from a home office or hotel), Cisco ER will not be able to determine the location of the caller. Furthermore, it is unlikely that the Cisco system would have a gateway properly situated to allow sending the call to the appropriate PSAP for the caller's location.

It is a matter of enterprise policy to allow or not to allow the use of soft clients for 911 calls. It is highly advisable to disallow 911 calls by policy for soft clients that can roam across the internet. Nevertheless, if such a user were to call 911, the best-effort system response would be to route the call to either an on-site security force or a large PSAP close to the system's main site.

The following paragraph is an example notice that you could issue to users to warn them that emergency call functionality is not guaranteed to soft client users:

Emergency calls should be placed from phones that are located at the site for which they are configured (for example, your office). A local safety authority might not answer an emergency call placed from a phone that has been removed from its configured site. If you must use this phone for emergency calls while away from your configured site, be prepared to provide the answering public safety authority with specific information regarding your location. Use a phone that is locally configured to the site (for example, your hotel phone or your home phone) for emergency calls when traveling or telecommuting.

Test Calls

For any enterprise telephony system, it is a good idea to test 911 call functionality, not only after the initial installation, but regularly, as a preventive measure.

The following suggestions can help you carry out the testing:

- Contact the PSAP to ask for permission before doing any tests, and provide them with the contact information of the individuals making the tests.
- During each call, indicate that it is *not* an actual emergency, just a test.
- Confirm the ANI and ALI that the call taker has on their screen.
- Confirm the PSAP to which the call was routed.
- Confirm that answer supervision was received by looking at the call duration timer on the IP phone. An active call timer is an indication that answer supervision is working properly.

PSAP Callback to Shared Directory Numbers

Cisco ER handles the routing of inbound calls made to emergency location identification numbers (ELINs). In cases where the line from which a 911 call was made is a shared directory number, the PSAP callback will cause all shared directory number appearances to ring. Any of the shared appearances can then answer the call, which means that it may not be the phone from which the 911 call originated.

Multi-Cluster Considerations

Enterprise telephony systems based on multiple Cisco Unified CallManager clusters can benefit from the functionality of Cisco Emergency Responder (Cisco ER).

The *Cisco Emergency Responder Administration Guide* provides detailed descriptions of the terms used herein, as well as the background information required to support the following discussion. Of specific interest is the chapter on *Planning for Cisco Emergency Responder*. This documentation is available at


<http://www.cisco.com>

Single Cisco ER Group

A single Emergency Responder group can be deployed to handle emergency calls from two or more Cisco Unified CallManager clusters. The design goal is to ensure that any phone's emergency call is routed to the Cisco ER group, which will assign an ELIN and route the call to the appropriate gateway based on the phone's location.

One advantage of using a single Cisco ER group is that all ERLs and ELINs are configured into a single system. A phone registered on any cluster will be located by the single Cisco ER group because that group is responsible for polling all of the system's access switches. [Figure 11-1](#) illustrates a single Cisco ER group interfaced with two Cisco Unified CallManager clusters.

Figure 11-1 A Single Cisco ER Group Connected to Two Cisco Unified CallManager Clusters



The single Cisco ER group in [Figure 11-1](#) interfaces with the following components:

- Each Cisco Unified CallManager cluster via SNMP, to collect information about their respective configured phones

- All of the enterprise's switches via SNMP, so that any cluster's phone, connected to any switch, can be located. This connection is not required if the phone locations are being identified based on IP subnets. For details on configuring IP subnet-based ERLs, refer to the chapter on Configuring Cisco Emergency Responder in the *Cisco Emergency Responder Administration Guide*, available at <http://www.cisco.com>
- Each Cisco Unified CallManager cluster via JTAPI, to allow for the call processing required by any phone that dials 911 – for example, identification of the calling phone's ERL, assignment of the ELIN, redirection of the call to the proper gateway (based on the calling phone's location), and the handling of the PSAP callback functionality

The version of the JTAPI interface used by Cisco Emergency Responder is determined by the version of the Cisco Unified CallManager software to which it is connected. At system initialization, Cisco ER interrogates the Cisco Unified CallManager cluster and loads the appropriate JTAPI Telephony Service Provider (TSP). Because there can be only one version of JTAPI TSP on the Cisco ER server, all Cisco Unified CallManager clusters to which a single Cisco ER group is interfaced *must* run the same version of Cisco Unified CallManager software.

For some deployments, this software version requirement might present some difficulties. For instance, during a Cisco Unified CallManager upgrade, different clusters will be running different versions of software, and some of the clusters will be running a version of JTAPI that is not compatible with the version running on the Cisco ER servers. When this situation occurs, emergency calls from the cluster running a version of JTAPI different than that of the Cisco ER group might receive the call treatment provided by the Call Forward Busy settings of the emergency number's CTI Route Point.

When considering if a single Cisco ER group is appropriate for multiple Cisco Unified CallManager clusters, apply the following guidelines:

- Make Cisco Unified CallManager upgrades during an acceptable maintenance window when emergency call volumes are as low as possible (for example, after hours, when system use is at a minimum).
- Use a single Cisco ER group only if the quantity and size of the clusters allow for minimizing the amount of time when dissimilar versions of JTAPI are in use during software upgrades.

For example, a deployment with one large eight-server cluster in parallel with a small two-server cluster could be considered for use with a single Cisco ER group. In this case, it would be best to upgrade the large cluster first, thus minimizing the number of users (those served by the small cluster) that might be without Cisco ER service during the maintenance window of the upgrade. Furthermore, the small cluster's users can more appropriately be served by the temporary static routing of emergency calls in effect while Cisco ER is not reachable because they can be identified by the single ERL/ELIN assigned to all non-ER calls made during that time.

**Note**

Emergency Responder version 1.3(1) is required if any of the Cisco Unified CallManager clusters are running Cisco Unified CallManager Release 4.2 or 5.0.

Multiple Cisco ER Groups

Multiple Cisco ER groups can also be deployed to support multi-cluster systems. In this case, each ER group interfaces with the following components:

- A Cisco Unified CallManager cluster via the following methods:
 - SNMP, to collect information about its configured phones
 - JTAPI, to allow for the call processing associated with redirection of the call to the proper gateway or, in the case of roaming phones, the proper Cisco Unified CallManager cluster

- The access switches (via SNMP) to which most of the phones associated with the Cisco Unified CallManager of the Cisco ER group are most likely to be connected

This approach allows Cisco Unified CallManager clusters to run different versions of software because each is interfaced to a separate Cisco ER group.

To allow phones to roam between various parts of the network and still be tracked by Cisco ER, you might have to configure the Cisco ER groups into a Cisco ER cluster. For details on Cisco ER clusters and groups, refer to the chapter on *Planning for Cisco Emergency Responder* in the *Cisco Emergency Responder Administration Guide*, available at

<http://www.cisco.com>

Figure 11-2 presents a sample topology illustrating some of the basic concepts behind Cisco ER clustering.

Figure 11-2 Multiple Cisco ER Groups




Figure 11-2 illustrates the following topology:

- Cisco ER group A is interfaced to Cisco Unified CallManager cluster A to access switches A1 and A2, and it is deemed to be the home Cisco ER group of all phones registered to Cisco Unified CallManager cluster A.
- Likewise, Cisco ER group B is interfaced to Cisco Unified CallManager cluster B to access switches B1 and B2, and it is deemed to be the home Cisco ER group of all phones registered to Cisco Unified CallManager cluster B.

**Note**

Emergency Responder 1.3(1) requires that all ER groups in an ER cluster run the same version of software. Therefore, if any of the Cisco Unified CallManager clusters are using Cisco Unified CallManager Release 4.2 or 5.0, then all of the ER groups must use Emergency Responder 1.3(1).

Phone Movements Within the Tracking Domain of a Cisco ER Group

The emergency call processing for phones moving between access switches controlled by the same home Cisco ER group is the same as the processing done for a deployment with a single Cisco Unified CallManager cluster. For example, a phone moving between access switches A1 and A2 remains registered with Cisco Unified CallManager cluster A, and its location is determined by Cisco ER group A both before and after the move. The phone is still under full control of Cisco ER group A, for both the discovery of the phone by Cisco Unified CallManager cluster A and the determination of the phone's location by switch A2. The phone is therefore not considered to be an unlocated phone.

Phone Movements Between the Various Tracking Domains of a Cisco ER Cluster

A Cisco ER cluster is essentially a collection of Cisco ER groups that share location information through a Lightweight Directory Access Protocol (LDAP) database. Each group shares the location of any phone it finds on an access switch or in an IP subnet. However, any phone found in the Cisco ER group's own Cisco Unified CallManager cluster is deemed to be *unknown*, and its information is not shared.

Cisco ER groups also share information about phones that cannot be located within a Cisco ER group's tracking domain (in switches or IP subnets) but which are known to be registered in the group's associated Cisco Unified CallManager cluster. Such phones are deemed *unlocated*.

If a phone is roaming between access switches monitored by different Cisco ER groups, those groups must be configured in a Cisco ER cluster so they can exchange information about the phone's location. For example, phone A3 is registered with Cisco Unified CallManager cluster A, but it is connected to an access switch controlled by Cisco ER group B. Cisco ER group A is aware that phone A3 is registered with Cisco Unified CallManager cluster A, but group A cannot locate phone A3 in any of the site A switches. Therefore, phone A3 is deemed *unlocated* by Cisco ER group A.

Cisco ER group B, on the other hand, has detected the presence of phone A3 in one of the switches that it monitors. Because the phone is not registered with Cisco Unified CallManager cluster B, phone A3 is advertised through the Cisco ER LDAP database as an *unknown* phone.

Because the two Cisco ER groups are communicating through an LDAP database, they can determine that Cisco ER group B's *unknown* phone A3 is the same as Cisco ER group A's *unlocated* phone A3.

The Unlocated Phone page in Cisco ER group A will display the phone's host name along with the remote Cisco ER group (in this, case Cisco ER group B).

Emergency Call Routing within a Cisco ER Cluster

Cisco ER clustering also relies on route patterns that allow emergency calls to be redirected between pairs consisting of a Cisco Unified CallManager cluster and a Cisco ER. For more details, refer to the section on *Creating Route Patterns for Inter-Cisco Emergency Responder-Group Communications* in the *Cisco Emergency Responder Administration Guide*, available at

<http://www.cisco.com>

If phone A3 places an emergency call, the call signaling flow will be as follows:

1. Phone A3 sends the emergency call string to Cisco Unified CallManager cluster A for processing.
2. Cisco Unified CallManager cluster A sends the call to Cisco ER group A for redirection.

3. Cisco ER group A determines that phone A3 is located in Cisco ER group B's tracking domain, so it redirects the call to a route pattern that points to Cisco Unified CallManager cluster B.
4. Cisco Unified CallManager cluster A sends the call to Cisco Unified CallManager cluster B.
5. Cisco Unified CallManager cluster B sends the call to Cisco ER group B for redirection.
6. Cisco ER group B identifies the ERL and ELIN associated with phone A3's location and redirects the call to Cisco Unified CallManager cluster B. The calling number is transformed into the ELIN associated with the ERL of phone A3, and the called number is modified to route the call to the proper gateway.
7. Cisco Unified CallManager cluster B routes the call according to the new called number information obtained from Cisco ER group B.
8. Cisco Unified CallManager cluster B sends the call out the gateway toward the Emergency PSTN network.

Scalability Considerations for Cisco ER Clustering

In a Cisco ER cluster, the quantity of phones roaming outside the tracking domain of their home Cisco ER group is a scalability factor that you must keep within the limits set forth in the section on *Network Hardware and Software Requirements* in the *Cisco Emergency Responder Administration Guide 1.2(3)*, available at:

http://www.cisco.com/en/US/products/sw/voicesw/ps842/prod_maintenance_guides_list.html

With the Cisco MCS 7845 server platform and version 1.2(3) of the Cisco ER software, a Cisco ER cluster can support a maximum of 3000 roaming phones. For deployments that have to exceed this limit (for instance, large campus deployments with multiple Cisco Unified CallManager clusters), phone movement can be tracked by IP subnets. By defining the IP subnets in each of the Cisco ER groups and by assigning each ERL with one ELIN per Cisco ER group, you can virtually eliminate roaming phones because all phones in the campus will be part of the tracking domain of their respective Cisco ER group.

ALI Formats

In multi-cluster configurations, there might be instances where the physical locations of ERLs and ELINs defined in a single Cisco ER group span the territory of more than one phone company. This condition can lead to situations where records destined for different phone companies have to be extracted from a common file that contains records for multiple LECs.

Cisco ER exports this information in ALI records that conform to National Emergency Number Association (NENA) 2.0, 2.1, and 3.0 formats. However, many service providers do not use NENA standards. In such cases, you can use the ALI Formatting Tool (AFT) to modify the ALI records generated by Cisco ER so that they conform to the formats specified by your service provider. That service provider can then use the reformatted file to update their ALI database.

The ALI Formatting Tool (AFT) enables you to perform the following functions:

- Select a record and update the values of the ALI fields. AFT allows you to edit the ALI fields to customize them to meet the requirements of various service providers. Your service provider can then read the reformatted ALI files and use them to update their ELIN records.
- Perform bulk updates on multiple ALI records. Using the bulk update feature, you can apply common changes to all the records that you have selected, to one area code, or to one area code and one city code.

- Selectively export ALI records based on area code, city code, or a four-digit directory number. By selecting to export all the ALI records in an area code, for example, you can quickly access all the ELIN records for each service provider, thereby easily supporting multiple service providers.

Given the flexibility of the AFT, a single Cisco ER group can export ALI records in multiple ALI database formats. For a Cisco ER group serving a Cisco Unified CallManager cluster with sites in the territories of two LECs, the basic approach is as follows:

1. Obtain an ALI record file output from Cisco Emergency Responder in standard NENA format. This file contains the records destined for multiple LECs.
2. Make a copy of the original file for each required ALI format (one copy per LEC).
3. Using the AFT of the first LEC (for example, LEC-A), load a copy of the NENA-formatted file and delete the records of all the ELINs associated with the other LECs. The information to delete can usually be identified by NPA (or area code).
4. Save the resulting file in the required ALI format for LEC-A, and name the file accordingly.
5. Repeat steps 3 and 4 for each LEC.

For more information about the ALI formatting tools, refer to the online documentation available at http://www.cisco.com/en/US/products/sw/voicesw/ps842/prod_maintenance_guides_list.html

For LECs not listed at this URL, the output from Cisco Unified CallManager can be formatted using standard text file editing tools, such as spreadsheet programs and standard text editors.



CHAPTER 12

Third-Party Voicemail Design

Last revised on: February 13, 2008

This chapter discusses various options for deploying third-party voicemail systems with Cisco Unified CallManager.



Note

This chapter does not discuss how to size a voicemail system for ports and/or storage. For this type of information, contact your voicemail vendor, who should be better able to discuss the individual requirements of their own system based upon specific traffic patterns.

There are many voicemail vendors, and it is not uncommon for customers to want to continue to use an existing voicemail system when deploying Cisco Unified CallManager. With this requirement in mind, Cisco provides support for the industry standard voicemail protocol known as Simplified Message Desk Interface (SMDI). SMDI is a serial protocol that provides all the necessary call information required for a voicemail system to answer calls appropriately.

There are also other options for integrating Cisco Unified CallManager to voicemail systems, such as digital set emulation, Microsoft TAPI, QSIG, and so forth. Each method has its own pros and cons, and the method you employ will largely depend on how your voicemail system is integrated to your current PBX.

This section covers the following aspects of integrating third-party voicemail systems with Cisco Unified CallManager:

- [SMDI, page 12-1](#)
- [Digital Set Emulation, page 12-4](#)
- [Dual PBX Integration, page 12-5](#)
- [Centralized Voicemail, page 12-6](#)
- [Positive Disconnect Supervision, page 12-11](#)
- [Summary of Third-Party Voicemail Integration, page 12-12](#)

SMDI

Cisco Unified CallManager supports use of SMDI through either of the following methods:


- [Cisco Messaging Interface, page 12-2](#)
- [Cisco VG248, page 12-2](#)

Cisco Messaging Interface

The Cisco Messaging Interface (CMI) is a Cisco CallManager service that should be run only on the publisher server. This service intercepts calls destined for voicemail and generates appropriate SMDI messages, which are then delivered to one of the server's Component Object Model (COM) ports. The CMI service is compatible with any MGCP gateway that supports analog FXS ports or T1 CAS E&M; however, the WS-X6624 and VG224 modules are two of only three gateways that support positive disconnect supervision (see [Positive Disconnect Supervision, page 12-11](#)) and are therefore the only gateways currently recommended for use with the CMI service.

Figure 12-1 illustrates the use of SMDI through the CMI service in Cisco Unified CallManager.

Figure 12-1 SMDI via Cisco Unified CallManager




Through the CMI, Cisco Unified CallManager supports integration with virtually any voicemail system that can provide SMDI with analog FXS ports, including (but not limited to) the following:

- Octel 100, 200/300, and 250/350
- Intuity Audix
- Siemens PhoneMail
- Centigram/BayPoint (OnePoint and NuPoint Messenger)
- Lyrinx ECS
- IBM Message Center

Cisco VG248

The Cisco VG248 is an SCCP gateway that supports 48 analog FXS ports and generates SMDI locally (that is, it runs independent of the CMI service). As with the WS-X6624 and VG224 modules, the VG248 also supports positive disconnect supervision.

Figure 12-2 illustrates the use of SMDI with the VG248.

Figure 12-2 SMDI via the VG248

Voicemail integration through the VG248 provides the following features and advantages:

- Multiple SMDI links per Cisco Unified CallManager
- SMDI failover capability
- Independence from the location of the voicemail system

The VG248 is also capable of supporting two other serial protocols that are sometimes used for voicemail integration: NEC Message Center Interface (MCI) and Ericsson MD110 proprietary protocol.

Considerations When Using FXS Ports

If your voicemail system is equipped with analog FXS ports, use the following Cisco gateways to integrate with the voicemail system:

- WS-X6624

Use this module when there is a slot available for it in the Catalyst 6500 chassis.

- VG224

Use this gateway when there is no physical Catalyst 6500 chassis slot available and when automatic failover of the serial port is not deemed necessary.

- VG248

Use this gateway when full failover is required for the serial port as well as voice ports, when serial protocols other than SMDI are required (such as NEC MCI or Ericsson MD110), or when no Catalyst 6500 chassis slot is available.

Digital Set Emulation

Digital set emulation (DSE) is another method of integrating PBXs to voicemail systems. In this mode, the voicemail ports look like proprietary digital handsets to the PBX. This integration method has the following advantages over the use of SMDI with analog FXS ports:

- The circuit is completely digital for both the voice path and signaling of call information.
- There is no out-of-band signaling because the voice and signaling are both transported over the same physical circuit.
- Overall quality of the call is generally better.

Digital PBX Adapter (DPA)

Cisco developed the Digital PBX Adapter (DPA) specifically to integrate Cisco Unified CallManager with third-party voicemail systems via digital set emulation. The DPA essentially functions as a number of digital PBX extensions on one side, with IP connectivity on the other. The DPA enables you to maintain your existing voicemail system, as well as its interfaces, while connecting to Cisco Unified CallManager.

The DPA is available in two variants:

- DPA 7630 for emulating the Avaya Definity G3 7400 series digital telephone set
- DPA 7610 for emulating the Nortel Meridian 1 2616 digital telephone set




Note

The DPA will work with Octel Aria 250/350 or Serenade 200/300 voicemail systems only when using either Lucent/Avaya or Nortel digital set emulation.

Figure 12-3 illustrates the DPA integrating an Octel voicemail system with Cisco Unified CallManager.

Figure 12-3 DPA Integrating Octel Voicemail with Cisco Unified CallManager



Dual PBX Integration

Dual PBX integration is a useful option for enterprises that want to maintain existing voicemail services while migrating from their current PBX to IP Telephony.




Note

Most voicemail vendors do not support this scenario due to its complex nature, but some will provide support on a "site-specific" basis if requested. Consult with your voicemail vendor before attempting to implement this solution.

The Cisco VG248 has inherent multiplexing capabilities that enable it to provide dual integration. The VG248 can combine information from an existing serial link with its own link, and then present a single serial stream to the voicemail system. (See [Figure 12-4](#).)


Figure 12-4 Dual Integration via the VG248 and SMDI



The VG248 works with any voicemail system that has SMDI capability and analog FXS ports. The following prerequisites are required prior to implementation, assuming a dual integration is required:

- Uniform dial plan
- Transfer and reconnect sequences
- Connectivity between the PBX and Cisco Unified CallManager

The Cisco DPA also has the capability to provide dual integration through digital set emulation, as illustrated in [Figure 12-5](#).

Figure 12-5 Dual Integration via DPA

The DPA works with Octel digital set emulation. The following prerequisites are required prior to implementation, assuming a dual integration is required:


- Uniform dial plan
- Transfer and reconnect sequences
- Connectivity between the PBX and Cisco Unified CallManager

Centralized Voicemail


In a centralized voicemail deployment, two or more PBXs share a single voicemail system. The sharing is achieved by integrating the voicemail system to only one PBX and then utilizing an inter-PBX private networking protocol to extend voicemail services to remote subscribers. The networked PBXs look and act like one large PBX to the voicemail system. Various PBX manufacturers have developed proprietary protocols that enable the delivery of such services as well as providing feature transparency to subscribers across a large network (for example, Avaya DCS, Nortel MCDN, Siemens CorNet, Alcatel ABC, NEC CCIS, and Fujitsu FIPN).

The primary motivation for using a centralized voicemail system stems from the desire to offer voicemail services to IP Telephony subscribers from the existing voicemail system so that the subscribers do not have to learn a new telephony user interface (TUI).

Some voicemail systems are capable of supporting multiple PBXs (dual PBX integration) via protocols such as Simple Messaging Desktop Interface (SMDI). Other solutions have been introduced, such as the Cisco Digital PBX Adapter (DPA), which also allow for dual integration. In some circumstances, these solutions are either not possible because the voicemail vendor may choose not to support this configuration, or a dual integration is simply not technically possible because the voicemail system cannot support dissimilar PBX integrations simultaneously. In such circumstances, a centralized voicemail deployment provides an alternative solution to dual integration. (See [Figure 12-6](#).)


Figure 12-6 Centralized Voicemail with Cisco Unified CallManager and QSIG

If you want to use an existing voicemail system, consider the make and model of that system. If the voicemail system in question is from the same manufacturer as the PBX system, then full voicemail functionality is typically available to Cisco Unified CallManager subscribers. See [Figure 12-7](#) for an example of a Nortel system and [Figure 12-8](#) for an Avaya system.

Figure 12-7 Nortel M1 Centralized Voicemail with Meridian Mail or CallPilot

The system in [Figure 12-7](#) has the following characteristics:


- Voicemail services are available to all subscribers.
- Voicemail is hosted on the message center PINX.
- QSIG MWI works only with Meridian Mail or CallPilot.

Figure 12-8 Avaya G3 Centralized Voicemail with Intuity Audix

The system in [Figure 12-8](#) has the following characteristics:

- Voicemail services are available to all subscribers.
- Voicemail is hosted on the message center PINX.
- QSIG MWI works only with Avaya Intuity Audix.


If the voicemail system originates from a different manufacturer than the PBX system, then some functionality might not be passed back through the QSIG trunk to Cisco Unified CallManager. In this instance, the Cisco Digital PBX Adapter (DPA) may be used specifically to deliver MWI directly to Cisco Unified CallManager. See [Figure 12-9](#) for an example of a Nortel system and [Figure 12-10](#) for an Avaya system.

Figure 12-9 Nortel M1 Centralized Voicemail with Octel Aria or Serenade

The system in Figure 12-9 has the following characteristics:

- Voicemail services are available to all subscribers.
- Station-invoked message center functionality is not mapped to QSIG MWI.
- The Cisco DPA is used for MWI to Cisco Unified CallManager only.

Figure 12-10 Avaya G3 Centralized Voicemail with Octel Aria or Serenade



The system in Figure 12-10 has the following characteristics:

- Voicemail services are available to all subscribers.
- Leave Word Calling (LWC) is not mapped to QSIG MWI.
- The Cisco DPA is used for MWI to Cisco Unified CallManager only.

Note that the term *centralized voicemail* does not refer to the voicemail system itself. Centralized voicemail is a function of the PBX's inter-PBX networking protocol (either a proprietary protocol such as Avaya DCS, Nortel MCDN, or Siemens CorNet or a standards-based protocol such as QSIG or DPNSS), which is needed to deliver the voicemail features.

The following important terms and concepts apply to centralized voicemail:

- Message center Private Integrated Services Network Exchange (PINX) — This is the PBX that is "hosting" the voicemail system (the PBX directly connected to the voicemail system).
- Subscriber PINX — This is the PBX that is "remote" from the voicemail system (not directly connected to the voicemail system).


A centralized voicemail configuration requires a suitable inter-PBX networking protocol such as QSIG. This protocol must also deliver the following minimum level of feature support:

- Message Waiting Indication (MWI)
- Transfer — Needed to ensure that the correct calling and called party IDs are delivered to the voicemail system.
- Divert — Needed to ensure that the correct calling and called party IDs are delivered to the voicemail system.

Other features might also be required, depending on how the voicemail system will be used. For example, if the voicemail system is also serving as an automated attendant, then the Path Replacement feature is needed to prevent calls from hair-pinning.

Not all PBXs are capable of serving as the message center PINX. In this case, consider relocating the voicemail system to Cisco Unified CallManager and have Cisco Unified CallManager act as the message center PINX, with the PBX acting as the subscriber PINX. (See Figure 12-11.)

Figure 12-11 Centralized Voicemail with Cisco Unified CallManager Acting as Message Center PINX



Support

Cisco cannot guarantee that another vendor's product will act in a particular manner, nor can Cisco specify what is required in terms of configuration changes or upgrades to another vendor's product. It is the responsibility of the customer to ask these questions and seek confirmation directly from the supplier and/or vendor of each product.

Cisco can assist you in determining which particular questions to ask your supplier and/or vendor, such as: "What do I have to do to my PBX to enable remote PBX users, connected via QSIG, to have a mailbox as well as full access to all voicemail features such as MWI?"

To help with PBX interoperability, Cisco has tested a number of different PBXs with Cisco Unified CallManager and has documented these tests in the form of Application Notes. These documents, while not a guarantee of success, do provide some level of guidance in terms of features supported as well as configuration details for both Cisco Unified CallManager and the PBX. Application Notes for Cisco Unified CallManager have already been written for the leading PBXs, and they cover the scenario of centralized voicemail with Cisco Unified CallManager acting as the message center PINX. The Application Notes are available at

<http://www.cisco.com/go/interoperability>



Note

It is not feasible for Cisco to test other vendors' PBXs acting as the message center PINX. Cisco has neither the facilities nor the expertise to configure these systems, therefore customers must request this information directly from their supplier and/or vendor.

Summary

- Centralized voicemail is a function of the inter-PBX networking protocol, not the voicemail system itself.
- Not all PBXs can act as the message center PINX. Customers *must* confirm this feature with their PBX supplier and/or vendor; Cisco cannot provide or support this feature on the PBX.
- Cisco Unified CallManager can act as the message center PINX, thus providing customers with an alternative if their PBX cannot perform this function.
- Confirm if Path Replacement is needed. Cisco Unified CallManager Release 4.1 and later supports this feature.

Positive Disconnect Supervision

Positive disconnect supervision is a signal sent from a PBX port to the voicemail system to indicate that the far-end device has gone on-hook. This signal typically takes the form of a drop in loop current for approximately 600 ms, causing the voicemail system to terminate the session.

Without this signal, the voicemail system would be unaware that the far-end device has gone on-hook and would continue to record whatever supervisory tones the PBX provides under this condition. (For example, some PBXs play dial tone while others play busy tone.) The voicemail system would continue to record these tones until the maximum message time has expired. (For example, if the mailbox has a limit is 3 minutes per message and a caller hangs up after 30 seconds, then the voicemail system would continue to record such tones for another 2.5 minutes in the absence of positive disconnect supervision.) This unnecessary recording can be annoying to subscribers and can impact system performance by increasing disk usage as well as causing higher port usage times. Some voicemail systems are able to deal with this scenario by monitoring for known tones and then deleting them, but system performance is still impacted in this case.

A similar issue exists when subscribers call into their mailboxes to check for messages. If a user simply hangs up without disconnect supervision, the voicemail system will stay on the line waiting for a valid response before any activity timers eventually expire. In this scenario, the main impact is from the additional port usage time incurred.

For these reasons, positive disconnect supervision must be provided by the analog ports connecting to the voicemail system.

Summary of Third-Party Voicemail Integration

There are other methods for connecting voicemail systems to Cisco Unified CallManager (such as Microsoft TAPI and PRI ISDN trunks in conjunction with SMDI), but these methods are uncommon. The vast majority of third-party voicemail integrations use the Cisco VG248 or the Digital PBX Adapter (DPA), therefore they are the recommended solutions.

The method of integrating Cisco Unified CallManager to the voicemail system depends on how the voicemail system is currently integrated to the PBX. If analog ports are currently in use, then the Cisco VG248 or the VG224 provides the best integration method. However, if Avaya or Nortel digital set emulation is currently in use, then the Cisco DPA can provide a solution at lower cost by integrating with your current voicemail system without the need to re-engineer it for analog FXS ports.

**Note**

Cisco does not test or certify any third-party voicemail systems. Within the industry, it is generally considered to be the responsibility of the voicemail vendor to test and/or certify their products with various PBX systems. Cisco does, of course, test its interfaces to such equipment and will support these interfaces regardless of which third-party voicemail system is connected.



CHAPTER 13

Cisco Unity

Last revised on: February 13, 2008

This chapter focuses on the design aspects of integrating Cisco Unity and Cisco Unity Connection with Cisco Unified CallManager. The design topics covered in this chapter apply to both voicemail and unified messaging configurations.

Additionally, this chapter discusses design issues for deploying Cisco Unity with Microsoft Exchange 2000 or 2003 or Lotus Notes Domino message stores and Microsoft Windows 2000 or 2003. This chapter does *not* cover deployments or upgrades from Microsoft NT 4.0 and/or Exchange 5.5. Cisco Unity Connection has no dependencies on an external message store.

For additional design information about Cisco Unity and Unity Connection, including integrations with other non-Cisco messaging systems, refer to the *Cisco Unity Design Guide*, available at

<http://www.cisco.com>

This chapter covers the following Cisco Unity and Unity Connection design topics:

- [Messaging Deployment Models, page 13-2](#)
- [Messaging System Infrastructure Components, page 13-5](#)
- [Port Groups \(Separate Integrations\), page 13-5](#)
- [Managing Bandwidth, page 13-6](#)
- [Cisco Unity and Unity Connection Native Transcoding Operation, page 13-8](#)
- [Voice Port Integration with a Cisco Unified CallManager Cluster, page 13-9](#)
- [Voice Port Integration with Dedicated Cisco Unified CallManager Backup Servers, page 13-11](#)
- [Centralized Messaging and Centralized Call Processing, page 13-12](#)
- [Distributed Messaging with Centralized Call Processing, page 13-13](#)
- [Combined Messaging Deployment Models, page 13-15](#)
- [Centralized Messaging with Clustering Over the WAN, page 13-16](#)
- [Distributed Messaging with Clustering Over the WAN, page 13-18](#)
- [Cisco Unity Messaging Failover, page 13-19](#)
- [Cisco Unity Failover and Clustering Over the WAN, page 13-19](#)
- [Centralized Messaging with Multiple Cisco Unified CallManager Servers, page 13-20](#)

Cisco Unified CallManager Release 4.0 introduced *line groups*, *hunt lists*, and *hunt pilots*, which can impact the operation of existing voice ports. Before upgrading to Cisco Unified CallManager 4.x from Cisco Unified CallManager versions older than Release 4.0, refer to the appropriate *Cisco Unified CallManager Administration Guide* for your 4.x release, available at

<http://www.cisco.com>

All of the deployment models and design considerations discussed in this chapter are fully supported by the Cisco Unified CallManager 4.x releases.

While Cisco Unified CallManager 4.2 supports SIP trunks, Unity and Unity Connection are not supported for SIP trunk integrations with this release.

Messaging Deployment Models

Cisco Unity supports three primary messaging deployment models:

- [Single-Site Messaging, page 13-2](#)
- Multisite WAN deployment with [Centralized Messaging, page 13-2](#)
- Multisite WAN deployment with [Distributed Messaging, page 13-3](#)

Cisco Unity Connection supports single-site messaging and centralized messaging with a single messaging server.

Deployments involving both Cisco Unified CallManager and Cisco Unity or Unity Connection use one call processing model for Cisco Unified CallManager and one messaging model for Cisco Unity. The messaging deployment model is independent from the type of call processing model deployed.

In addition to the three messaging deployment models, Cisco Unity also supports messaging failover. (See [Messaging Failover, page 13-4](#).) Cisco Unity Connection does not currently support failover. Because Cisco Unity Connection uses only a single server without any networking between servers, it supports only the single-site and centralized messaging models.

Cisco Unity supports voicemail only and unified messaging in all messaging models. Cisco Unity Connection supports voicemail only in single-site and centralized messaging models.

Single-Site Messaging

In this model, the messaging systems and messaging infrastructure components are all located at the same site, on the same highly available LAN. The site can be either a single site or a campus site interconnected via high-speed metropolitan area networks (MANs). All clients of the messaging system are also located at the single (or campus) site. The key distinguishing feature of this model is that there are no remote clients.

Centralized Messaging

In this model, similar to the single-site model, all the messaging system and messaging infrastructure components are located at the same site. The site can be one physical site or a campus site interconnected via high-speed MANs. However, unlike the single-site model, centralized messaging clients can be located both locally and remotely.

Because messaging clients may be either local or remote from the messaging system, special design considerations apply to the following graphical user interface (GUI) clients: ViewMail for Outlook (VMO) with Microsoft Exchange and Lotus Notes with Domino Unified Communications Services (DUC) for Cisco. The use of the telephone record and playback (TRaP) and message streaming features should be restricted. Remote clients should not use TRaP and should be configured to download messages before playback because the additional bandwidth usage placed on WAN links cannot be controlled by call admission control. Additionally, because different features and operations for local and remote clients can cause user confusion, TRaP should be disabled on the voice ports and GUI clients should be configured to download messages and not use TRaP, regardless of whether the client is local or remote. Cisco Unity inbox accessed via Cisco Unity Personal Assistant (CPCA) should be allowed only for local clients. The Cisco Unity telephone user interface (TUI) operates the same way for both local and remote clients.

Distributed Messaging

In distributed messaging, the messaging systems and messaging infrastructure components are co-located in a distributed fashion. There can be multiple locations, each with its own messaging system and messaging infrastructure components. All client access is local to each messaging system, and the messaging systems share a messaging backbone that spans all locations. Message delivery from the distributed messaging systems occurs via the messaging backbone through a hub-and-spoke type of message routing infrastructure. No messaging infrastructure components should be separated by a WAN from the messaging system they service. Distributed messaging is essentially multiple, single-site messaging models with a common messaging backbone.


The distributed messaging model has the same design criteria as centralized messaging with regard to local and remote GUI clients, TRaP, and message downloads.

Because Cisco Unity Connection supports only a single server, distribution of messaging servers is not possible with it. Therefore, Cisco Unity Connection does not support the distributed messaging model.

Messaging Failover

All three messaging deployment models support messaging failover. You can implement local messaging failover as illustrated in Figure 13-1. With local failover, both the primary and secondary Cisco Unity servers are located at the same site on the same highly available LAN. This design is only for Cisco Unity; Cisco Unity Connection does not currently support failover.

Figure 13-1 Local Failover of Cisco Unity Messaging



The following list summarize the supported combinations of messaging and call processing models. Cisco Unity and Cisco Unified CallManager support all of the following combinations of messaging and call processing deployment models. Cisco Unity Connection supports all combinations except the distributed messaging models.

- Single-site messaging and single-site call processing
- Centralized messaging and centralized call processing
- Distributed messaging and centralized call processing
- Centralized messaging and distributed call processing
- Distributed messaging and distributed call processing


Refer to the Cisco Unity and Cisco Unity Connection design guides available at <http://www.cisco.com> for further details on site classification and a detailed analysis of supported combinations of messaging and call processing deployment models.

Messaging System Infrastructure Components

Cisco Unity interacts with various network resources, include Dynamic Domain Name Server (DDNS), directory servers, and message stores. (See [Figure 13-2](#).) Cisco Unity supports both Microsoft Exchange and IBM Lotus Domino as message stores. Each of these message store types has different infrastructure components (refer to the appropriate deployment type in the *Cisco Unity Design Guide*, available at <http://www.cisco.com>).

Instead of viewing Cisco Unity as a single monolithic device, it is better to think of it as a system of interdependent components. Each messaging system infrastructure component in a Cisco Unity messaging system is required for normal operation, and it is critical that all of these components be on the same highly available LAN. (In most cases they will be physically co-located.) Any WAN links between these components can introduce delays that will impact the operation of Cisco Unity. These delays will manifest themselves as long delays and periods of silence during TUI operation. For more information, refer to the *Network and Infrastructure Considerations* chapter in the *Cisco Unity Design Guide*, available from <http://www.cisco.com>.

Figure 13-2 Cisco Unity Messaging System Infrastructure Components (Exchange Specific)



[Figure 13-2](#) is a logical representation of the message system infrastructure components. Some of these components can be located on the same server. In the case of IBM Lotus Domino, the message store and directory (Names.nsf) are located on the same server. Microsoft Windows, the global catalog server, and domain controller may also be on the same server. You can also use multiple instances of each component, as in message store clustering, which Cisco Unity supports for Microsoft Exchange 2000 or 2003 and Lotus Domino. All messaging system infrastructure components must be located on the same highly available LAN as the Cisco Unity server(s) that they service.

When deploying Cisco Unity Connection, it is possible to locate the automatic speech recognition (ASR) service on a separate server. In this deployment environment, the Cisco Unity Connection server and the ASR server should reside in the same location. Cisco Unity Connection does not have the same infrastructure requirements or dependencies as Cisco Unity.

Port Groups (Separate Integrations)

Unity Connection introduces the concept of *port groups*, also known as *separate integrations* in Cisco Unity 4.2. In earlier releases of Unity, under the Cisco Unity Telephony Integration Manager (UTIM) you were able to configure *multiple clusters* of the same type of integration. It is important to understand

how port groups and separate integrations are different than multiple clusters under a single integration. With the single-integration multiple-cluster approach, the Unity database associated users to an integration, not a cluster. Because there was only a single integration, all subscribers were associated with that single integration regardless of which cluster they were on. For message waiting indication (MWI), Unity then had to broadcast MWI updates out all ports under the single integration, including to clusters where the subscriber was not present. With the new addition of separate integrations in Unity (4.2) and port groups in Unity Connection, subscribers are associated with the specific integration to which they belong. MWI lights no longer have to be broadcast out all ports but are sent only to the specific ports that are associated with the particular subscriber.

Dual integrations occur when you integrate more than one type of system to a Unity or Unity Connection server, such as when integrating a legacy phone system and a Cisco Unified CallManager phone system. Multiple integrations occur when you integrate Cisco Unity 4.1 or earlier versions to more than one Cisco Unified CallManager cluster via the same type of integration, such as when integrating via SCCP to two or more Cisco Unified CallManager clusters.

With Unity 4.2, separate integrations are configured via the UTIM, while for Unity Connection port groups are configured under the System Administrator.


With Cisco Unified CallManager 4.2, SIP integrations to Cisco Unity and Unity Connection are not currently supported. For SIP integrations with Cisco Unified CallManager, refer to the *Cisco Unified Communications SRND for Cisco Communications Manager 5.x*, available at

<http://www.cisco.com/go/designzone>

Managing Bandwidth

Cisco Unified CallManager provides a variety of features for managing bandwidth. Through the use of regions, locations, and even gatekeepers, Cisco Unified CallManager can ensure that the number of voice calls going over a WAN link does not oversubscribe the existing bandwidth and cause poor voice quality. Cisco Unity and Unity Connection rely on Cisco Unified CallManager to manage bandwidth and to route calls. If you deploy Cisco Unity or Unity Connection in an environment where calls or voice ports might cross WAN links, these calls will be transparent to gatekeeper-based call admission control. This situation occurs any time the Cisco Unity or Unity Connection server is servicing either distributed clients (distributed messaging (Unity only) or distributed call processing) or when Cisco Unified CallManager is remotely located (distributed messaging (Unity only) or centralized call processing). Cisco Unified CallManager provides regions and locations for call admission control.

Figure 13-3 uses a small centralized messaging and centralized call processing site to illustrate how regions and locations work together to manage available bandwidth. For a more detailed discussion of regions and locations, refer to the chapter on [Call Admission Control](#), page 9-1.

Figure 13-3 *Locations and Regions*

In [Figure 13-3](#), regions 1 and 2 are configured to use G.711 for intra-region calls and G.729 for inter-region calls. Locations 1 and 2 are both set to 24 kbps. Location bandwidth is budgeted only in the case of inter-location calls.

An intra-region (G.711) call would not be budgeted against the available bandwidth for the location. For example, when extension 100 calls extension 101, this call is not budgeted against the 24 kbps of available bandwidth for Location 1. However, an inter-region call using G.729 is budgeted against both bandwidth allocations of 24 kbps for Location 1 and Location 2. For example, when extension 100 calls extension 200, this call would be connected but any additional (simultaneous) inter-region calls would receive reorder (busy) tone.

Impact of Non-Delivery of RDNIS on Voicemail Calls Routed via AAR

Automated alternate routing (AAR), a Cisco Unified CallManager feature, can route calls over the PSTN when the WAN is oversubscribed. However, when calls are rerouted over the PSTN, Redirected Dialed Number Information Service (RDNIS) can be affected. Incorrect RDNIS information can impact voicemail calls that are rerouted over the PSTN by AAR when Cisco Unity or Unity Connection is remote from its messaging clients. If the RDNIS information is not correct, the call will not reach the voicemail box of the dialed user but will instead receive the auto-attendant prompt, and the caller might be asked to reenter the extension number of the party they wish to reach. This behavior is primarily an issue when the telephone carrier is unable to ensure RDNIS across the network. There are numerous reasons why the carrier might not be able to ensure that RDNIS is properly sent. Check with your carrier

to determine if they provide guaranteed RDNIS deliver end-to-end for your circuits. The alternative to using AAR for oversubscribed WANs is simply to let callers hear reorder tone in an oversubscribed condition.

Cisco Unity and Unity Connection Native Transcoding Operation

By default, the Cisco Unity or Unity Connection server performs transcoding automatically. Currently, Cisco Unified CallManager and Cisco Unity support only G.729 and G.711 for the Skinny Client Control Protocol (SCCP) TAPI Service Provider (TSP) voice ports, and other codecs are supported with legacy integration using Intel or Dialogic voice boards. Cisco Unity native transcoding does not use external hardware transcoders but instead uses the server's main CPU. For SCCP integrations only, the native transcoding must be disabled (turned off) on the Cisco Unity server via a registry setting in order for Cisco Unified CallManager to assign hardware transcoders to voice port calls. The registry setting is called "Audio - Enable G.729a codec support," and the tool for setting it is the *Advanced Settings Tool* available at <http://www.CiscoUnityTools.com>. (To disable native transcoding on Cisco Unity Connection, the configuration of the messaging codec is done differently, as described at the end of this section.)

By default, there is no codec registry key, and therefore native transcoding is enabled (on). The Advanced Settings Tool adds a new registry key that enables you to select either one of the two available codecs. Cisco Unity will then "advertise" to Cisco Unified CallManager that it supports only one codec. If a call terminating or originating on the voice ports is using a different codec than the type configured for the Cisco Unity server, Cisco Unified CallManager will assign an external transcoding resource for the call. For detailed information on how to configure transcoding resources on Cisco Unified CallManager, refer to the chapter on [Media Resources](#), page 6-1.


When the Advanced Settings Tool is used to enable only one codec, the Cisco Unity server system prompts must be the same as the codec used. By default, the system prompts are G.711. If the codec is set to G.711, the system prompts will work correctly. However, if G.729 is selected, you will have to change the system prompts. Refer to the *Cisco Unity Administration Guide* on <http://www.cisco.com> for details on how to change the system prompts. If the system prompts are not the same as the codec selected by the registry, then the end users will hear unintelligible system prompts.

To change how Cisco Unity Connection advertises which codecs it supports, the configuration is done differently than for Cisco Unity. Instead of using the Cisco Unity Tools Depot, the configuration changes are made from the Cisco Unity Connection Administration page. Under the Telephony Integrations heading, select the Port Groups link. On the Port Groups page, you can change the configuration to advertise either G.711 only, G.729 only, or both, but "prefer" either G.711 or G.729. In the preferred setting, Cisco Unity Connection will advertise that it supports both protocols (native transcoding) but will prefer to use the one specified over the other. With native transcoding disabled, Cisco Unified CallManager will provide a transcoder resource when needed instead of using the CPU on the Unity or Unity Connection server.

Voice Port Integration with a Cisco Unified CallManager Cluster

When deploying Cisco Unity in a single-site messaging environment, integration with the Cisco Unified CallManager cluster occurs through the SCCP voice ports. Design considerations must include proper deployment of the voice ports among the Cisco Unified CallManager subscribers so that, in the event of a subscriber failure (Cisco Unified CallManager failover), users and outside calls can continue to access voice messaging. (See [Figure 13-4](#).)

Figure 13-4 Cisco Unity Server(s) Integrated with a Cisco Unified CallManager Cluster (No Dedicated Backup Servers)



The Cisco Unified CallManager cluster in [Figure 13-4](#) employs 1:1 server redundancy and 50/50 load balancing. During normal operations, each subscriber server is active and handles up to 50% of the total server call processing load. In the event of a subscriber server failure, the remaining subscriber server takes up the load of the failed server.

This configuration uses two groups of voicemail ports, with each group containing one-half of the total number of licensed voice ports. One group is configured so that its primary server is Sub1 and its secondary (backup) server is Sub2. The second group is configured so that Sub2 is the primary server and Sub1 is the backup.

Make sure that MWI-only ports or any other special ports are equally distributed between the two groups. During the configuration of the voice ports, pay special attention to the naming convention. When configuring the two groups of ports in the Cisco Unity Telephony Integration Manager (UTIM) utility, make sure that the device name prefix is unique for each group and that you use the same device name when configuring the voicemail ports in Cisco Unified CallManager Administration. The device name prefix is unique for each group of ports in this example, with group Sub1 using CiscoUM1 as the device name prefix and Sub2 using CiscoUM2 in this example.

For additional design information on the ratio of inbound to outbound voicemail ports (for MWI, message notification, and TRaP), refer to the *Cisco Unity Design Guide* available at <http://www.cisco.com>.

**Note**

The device name prefix is unique for each group of ports and must match the same naming convention for the voicemail ports configured in Cisco Unified CallManager Administration.

In Cisco Unified CallManager Administration, half of the ports in this example are configured to register using the unique device name prefix of CiscoUM1, and the other half are configured to register using the unique device prefix. (See [Table 13-1](#).) When the ports register with Cisco Unified CallManager, half will be registered with subscriber server Sub1, and the other half will be registered with Sub2, as shown in [Table 13-1](#).

Table 13-1 Voicemail Port Configuration in Cisco Unified CallManager Administration

Device Name	Description	Device Pool	SCCP Security Profile	Status	IP Address
CiscoUM1-VI1	Unity1	Default	Standard Profile	Registered with sub1	1.1.2.9
CiscoUM1-VI2	Unity1	Default	Standard Profile	Registered with sub1	1.1.2.9
CiscoUM1-VI3	Unity1	Default	Standard Profile	Registered with sub1	1.1.2.9
CiscoUM1-VI4	Unity1	Default	Standard Profile	Registered with sub1	1.1.2.9
CiscoUM2-VI1	Unity1	Default	Standard Profile	Registered with sub2	1.1.2.9
CiscoUM2-VI2	Unity1	Default	Standard Profile	Registered with sub2	1.1.2.9
CiscoUM2-VI3	Unity1	Default	Standard Profile	Registered with sub2	1.1.2.9
CiscoUM2-VI4	Unity1	Default	Standard Profile	Registered with sub2	1.1.2.9

**Note**

The naming convention used for the voicemail ports in Cisco Unified CallManager Administration must match the device name prefix used in Cisco UTIM, otherwise the ports will fail to register.


Cisco Unified CallManager 4.0 introduced numerous changes to the configuration of SCCP voicemail ports with regard to how they hunt and forward calls. For information on how these changes impact the voicemail port operation and configuration, refer to the Cisco Unified CallManager 4.0 documentation available at

<http://www.cisco.com>

Voice Port Integration with Dedicated Cisco Unified CallManager Backup Servers

This Cisco Unified CallManager cluster configuration allows each subscriber server to operate at a call processing load higher than 50%. Each primary subscriber server has either a dedicated or shared backup server. (See Figure 13-5.) During normal operation the backup server processes no calls, in the event of failure or maintenance of a Subscriber server, the backup server will then take the full load of that server.

Figure 13-5 Cisco Unity Server(s) Integrated with a Single Cisco Unified CallManager Cluster with Backup Subscriber Server(s)



Configuration of the voicemail ports in this case is similar to the 50/50 load-balanced cluster. However instead of configuring the voice ports to use the opposite subscriber server as the secondary server, the individual shared or dedicated backup server is used. In the Cisco Unified CallManager clustered with a shared backup server, both of the secondary ports for the subscribers servers are configured to use the single backup server.


The voice port names (device name prefix) must be unique for each Cisco UTIM group and must be the same as the device names used on the Cisco Unified CallManager server.

To configure the voicemail ports on Cisco Unity, use the UTIM tool. For Cisco Unity Connection, use the Telephony Integration section of the Unity Connection Administration console. For details, refer to the Cisco Unity or Cisco Unity Connection administration guides available at <http://www.cisco.com>.

Centralized Messaging and Centralized Call Processing

In centralized messaging, the Cisco Unity server is co-located with the Cisco Unified CallManager cluster. With centralized call processing, subscribers may be located either remotely and/or locally to the cluster and messaging server(s). (See [Figure 13-6](#).) When remote users access resources at the central site (such as voice ports, IP phones, or PSTN gateways, as in Tail-End Hop-Off (TEHO)), these calls are transparent to gatekeeper call admission control. Therefore, regions and locations must be configured in Cisco Unified CallManager for call admission control. (See [Managing Bandwidth, page 13-6](#).) When making inter-region calls to IP phones or MGCP gateways, IP phones automatically select the inter-region codec that has been configured. Native transcoding should be disabled so that the Cisco Unity voice ports use Cisco Unified CallManager transcoding resources for calls that transverse the WAN (inter-region).

Figure 13-6 Centralized Messaging with Centralized Call Processing



In [Figure 13-6](#), regions 1 and 2 are configured to use G.711 for intra-region calls and G.729 for inter-region calls. Native transcoding has been disabled on the Cisco Unity server.

As [Figure 13-6](#) shows, when a call is made from extension 200 to the voicemail ports in Region 1, the inter-region G.729 codec is used at the endpoint but the RTP stream is transcoded to use G.711 on the voice ports. Native transcoding on the Cisco Unity server has been disabled in this example. Cisco Unified CallManager transcoding resources must be located at the same site as the voicemail system.


Hairpinning

Another issue to consider is hairpinning (or tromboning) of voice calls over multiple Cisco Unity voice ports. Hairpinning is not a concern in environments that use only SCCP TSP voice ports, but it is a concern with dual-integration environments, where hairpinning can occur between the voice ports of the legacy system and the SCCP TSP voice ports.

For more information about dual integration, refer to the *Cisco Unity Administration Guide*, available at <http://www.cisco.com>

Distributed Messaging with Centralized Call Processing

Distributed messaging means that there are multiple messaging systems distributed within the telephony environment, and each messaging system services only local messaging clients. This model differs from centralized messaging, where clients are both local and remote from the messaging system. [Figure 13-7](#) illustrates the distributed messaging model with centralized call processing. As with other multisite call processing models, the use of regions and locations is required to manage WAN bandwidth. For this model, you must also disable Cisco Unity native transcoding.

Figure 13-7 Distributed Messaging with Centralized Call Processing

For the configuration in Figure 13-7, transcoder resources must be local to each Cisco Unity message system site. Regions 1 and 2 are configured to use G.711 for intra-region calls and G.729 for inter-region calls. Native transcoding has been disabled on the Cisco Unity servers.

Voice messaging ports for both Cisco Unity servers must be assigned the appropriate region and location by means of calling search spaces and device pools configured on the Cisco Unified CallManager server. In addition, to associate telephony users with a specific group of voicemail ports, you must configure Cisco Unified CallManager voicemail profiles. For details on configuring calling search spaces, device pools, and voicemail profiles, refer to the applicable version of the *Cisco Unified CallManager Administration Guide*, available at

<http://www.cisco.com>


The messaging systems are "networked" together to present a single messaging system to both inside and outside users. For information about Cisco Unity Networking for the distributed Unity servers, Refer to the *Networking in Cisco Unity Guide*, available at

http://www.cisco.com/en/US/products/sw/voicesw/ps2237/products_feature_guides_list.html

Combined Messaging Deployment Models

It is possible to combine messaging models in the same deployment, provided that the deployment adheres to all the guidelines listed in the preceding sections. [Figure 13-8](#) shows a user environment in which both centralized and distributed messaging are employed simultaneously.

Figure 13-8 Combined Deployment Models



[Figure 13-8](#) shows the combination of two messaging models. Regions 1 and 3 use centralized messaging with centralized call processing, while Region 2 uses distributed messaging with centralized call processing. All regions are configured to use G.711 for intra-regions calls and G.729 for inter-region calls.

In [Figure 13-8](#), centralized messaging and centralized call control are used between the Central Site and Site3. The messaging system at the Central Site provides messaging services for clients at both the Central Site and Site3. Site2 uses the distributed messaging model with centralized call control. The messaging system (Unity2) located at Site2 provides messaging services for only those users located within Site2. In this deployment, both models adhere to their respective design guidelines as presented


in this chapter. Transcoding resources are located locally to each messaging system site, and they support clients who access messaging services from a remote site (relative to the messaging system), as in the case of a Site2 user leaving a message for a Central Site user.

Centralized Messaging with Clustering Over the WAN

This section addresses Cisco Unity design issues for deploying centralized messaging with Cisco Unified CallManager clustering over the WAN with local failover. In the case of a WAN failure with this model, all remote messaging sites will lose voicemail capability until the WAN is restored. (See [Figure 13-9](#).)

Clustering over the WAN supports local failover. With local failover, each site has a backup subscriber server physically located at the site. This section focuses on deploying Cisco Unity centralized messaging with local failover for clustering over the WAN.

For additional information, refer to the section on [Clustering Over the IP WAN, page 2-17](#).

Figure 13-9 Cisco Unity Centralized Messaging and Clustering Over the WAN with Local Failover


The minimum bandwidth required between cluster sites is a T1 line (1,544 MHz). This amount of bandwidth can support signaling and database traffic for up to 10,000 busy hour call attempts (BHCA), but it does not include the required media bandwidth.

Clustering over the WAN with Cisco CallManager Release 3.3(3) and earlier supports up to four sites per cluster, but Cisco Unified CallManager Release 4.1 and higher supports up to eight sites. Cisco Unity supports up to the maximum number of sites in either case. The voicemail ports are configured only at the site where the Cisco Unity messaging system is located (see Figure 13-9). Voicemail ports do not register over the WAN to the remote site(s). Messaging clients at the other site(s) access all voicemail resources from the primary site. There is no benefit to configuring voice ports over the WAN to any of the remote sites because, in the event of a WAN failure, remote sites would lose access to the centralized messaging system. Because of bandwidth consideration, the voicemail ports should have TRaP disabled and all messaging clients should download voicemail messages to their local PCs (unified messaging only).

Distributed Messaging with Clustering Over the WAN

Local failover sites that also have Cisco Unity messaging server(s) deployed would have voice ports registered to the local Cisco Unified CallManager subscriber server(s), similar to the centralized messaging model. For information about configuring the voice ports, see [Voice Port Integration with a Cisco Unified CallManager Cluster](#), page 13-9, and [Voice Port Integration with Dedicated Cisco Unified CallManager Backup Servers](#), page 13-11.

Figure 13-10 Cisco Unity Distributed Messaging and Clustering over the WAN with Local Failover



In a purely distributed messaging implementation with clustering over the WAN, each site in the cluster would have its own Cisco Unity messaging server with messaging infrastructure components. If not all of the sites have local Cisco Unity messaging systems but some sites have local messaging clients using a remote messaging server(s), this deployment would be a combination model with both distributed

messaging and centralized messaging. (See [Combined Messaging Deployment Models, page 13-15](#).) In the event of a WAN failure in this model, all remote sites that use centralized messaging will lose voicemail capability until the WAN is restored.

Each site that does not have a local messaging server must use a single messaging server for all of its messaging clients, but all such sites do not have to use the same messaging server. For example, suppose Site1 and Site2 each have a local messaging server. Site3 can then have all of its clients use (register with) the messaging server at Site2, while Site4 can have all of its clients use the messaging server at Site1. Transcoder resources are required at sites that have local Cisco Unity messaging server(s).

As with other distributed call processing deployments, calls going between these sites are transparent to gatekeeper call admission control, therefore you must configure regions and locations in Cisco Unified CallManager to provide call admission control. (See [Managing Bandwidth, page 13-6](#).)

The distributed Cisco Unity servers may also be networked digitally. For more information on this topic, refer to the *Cisco Unity Networking Guide*, available at <http://www.cisco.com>. The networking guides are specific to the particular messaging store deployed.

Cisco Unity Messaging Failover

Cisco Unity failover provides voicemail survivability in the case of a Cisco Unity server failure. (See [Figure 13-1](#).) With Cisco Unity local failover, both the primary and secondary Unity messaging servers reside at the same physical location, and the messaging infrastructure components are co-located with the primary server. Optionally, messaging infrastructure components (such as messaging store servers, domain controller and global catalog (DC/GC) servers, and DNS servers) can also have redundant components that may be physically co-located with the Cisco Unity secondary server.

Cisco Unity failover is supported by all of the messaging deployment models. Cisco Unity Connection currently does not support failover.

For information on configuring Cisco Unity failover, refer to the *Cisco Unity Failover Configuration and Administration Guide*, available at

<http://www.cisco.com>

Cisco Unity Failover and Clustering Over the WAN


When deploying Cisco Unity local failover with clustering over the WAN, apply the same design practices described in [Centralized Messaging with Clustering Over the WAN, page 13-16](#), and [Distributed Messaging with Clustering Over the WAN, page 13-18](#). The voice ports from the primary Cisco Unity server should not cross the WAN during normal operation.

[Figure 13-11](#) depicts Cisco Unity local failover. Note that the primary and secondary Cisco Unity servers are both physically located at the same site. Cisco Unity failover supports up to the maximum number of remote sites available with clustering over the WAN for Cisco Unified CallManager.



Note

Unity Connection does not currently support failover.

Figure 13-11 Cisco Unity Local Failover and Clustering Over the WAN


For information on configuring Cisco Unity failover, refer to the *Cisco Unity Failover Configuration and Administration Guide*, available at

<http://www.cisco.com>

Centralized Messaging with Multiple Cisco Unified CallManager Servers

Cisco Unity and Unity Connection support port groups, which allow subscribers on either of the Unity product servers to be associated with a port group and ultimately an integration. When a subscriber is associated with a particular port group, the MWI function for the user occurs only on the MWI ports belonging to that port group. Cisco Unity supports up to nine port groups, while Cisco Unity Connection supports up to seven port groups. For details, refer to the appropriate Cisco Unity or Cisco Unity Connection administration guides available at <http://www.cisco.com>.

Figure 13-12 Integrating Cisco Unity or Unity Connection with Multiple Cisco Unified CallManager Clusters



Messaging clients at both Cluster 1 and Cluster 2 sites use the Cisco Unity or Unity Connection messaging infrastructure physically located at Cluster 1. (See [Figure 13-12](#).)



CHAPTER 14

Cisco Unity Express

Last revised on: February 13, 2008

This chapter discusses Cisco Unity Express as a distributed voicemail and automated attendant (AA) solution for Cisco Unified CallManager. Cisco Unity Express is an entry-level voicemail system for offices of up to 250 people (250 mailboxes), and it integrates physically as a blade within the branch office router at each site.

Use Cisco Unity Express as a distributed voicemail solution if any of the following conditions apply to your Cisco Unified CallManager network deployment:

- Survivability of voicemail and AA access must be ensured regardless of WAN availability.
- Available WAN bandwidth is insufficient to support voicemail calls traversing the WAN to a central voicemail server.
- There is limited geographic coverage of the AA or branch site PSTN phone numbers published to the local community, and these numbers cannot be dialed to reach a central AA server without incurring toll charges.
- The likelihood is high that a PSTN call into a branch office will be transferred from the branch AA to a local extension in the same office.
- Management philosophy allows remote locations to select their own voicemail and AA technology.



Note

Interoperability with Cisco Unified CallManager requires a minimum of Cisco Unity Express Release 1.1 and Cisco Unified CallManager Release 3.3.3 or later 3.3-based release. Cisco Unity Express 2.0 provides interoperability with Cisco Unified CallManager Release 4.0, Cisco Unity Express 2.1 provides interoperability with Cisco Unified CallManager Release 4.1, and Cisco Unity Express 2.3 provides interoperability with Cisco Unified CallManager Release 4.2.

For more information about Cisco Unity Express, refer to the product documentation available at

<http://www.cisco.com/go/cue>

Deployment Models for Cisco Unity Express

Cisco Unity Express is supported with all of the Cisco Unified CallManager deployment models, including single-site deployments, multi-site WAN with centralized call processing, and multi-site WAN with distributed call processing. [Figure 14-1](#) shows a centralized call processing deployment incorporating Cisco Unity Express, and [Figure 14-2](#) shows a distributed call processing deployment.

Cisco Unity Express can also be deployed as a co-resident voicemail solution with Cisco Unified CallManager Express. Cisco Unity Express sites controlled by Cisco Unified CallManager Express, as well as other sites controlled by Cisco Unified CallManager, can be interconnected in the same network. Although Cisco Unity Express can integrate with either Cisco Unified CallManager or Cisco Unified CallManager Express, it cannot integrate with both simultaneously.

For additional information on supported deployment models with Cisco Unified CallManager Express, refer to the appropriate Cisco Unified CallManager Express design documentation available at

<http://www.cisco.com>

Figure 14-1 Cisco Unity Express in a Centralized Call Processing Deployment





Figure 14-2 Cisco Unity Express in a Distributed Call Processing Deployment

The most likely deployment model to use Cisco Unity Express is the multi-site WAN with centralized call processing, where Cisco Unity Express provides distributed voicemail at the smaller remote offices and a central Cisco Unity system provides voicemail to the main campus and larger remote sites.

The following characteristics and guidelines apply to Cisco Unity Express in either a centralized or distributed Cisco Unified CallManager deployment:

- A single Cisco Unity Express can be integrated with a single Cisco Unified CallManager cluster.
- Cisco Unity Express integrates with Cisco Call Manager using a JTAPI application and Computer Telephony Integration (CTI) Quick Buffer Encoding (QBE) protocol. CTI ports and CTI route points control the Cisco Unity Express voicemail and automated attendant applications.
- Cisco Unity Express provides voicemail functionality to Cisco IP Phones running Skinny Client Control Protocol (SCCP). Cisco Unity Express 2.3 also provides support for Session Initiation Protocol (SIP) IP phones.
- The following CTI route points are defined on Cisco Unified CallManager for Cisco Unity Express:
 - Automated attendant entry point (Cisco Unity Express can contain up to five distinct AAs and may therefore require up to five different route points.)
 - Voicemail pilot number
 - Greeting management system (GMS) pilot number (Optional; if the GMS is not used, then this route point need not be defined.)
- The following CTI ports are defined on Cisco Unified CallManager for Cisco Unity Express:
 - A 12 or 25-mailbox system (4 ports)
 - A 50-mailbox AIM-CUE system (6 ports)

- A 100-mailbox NM-CUE system (8 ports)
- A 250-mailbox NM-CUE-EC system (16 ports)
- Each Cisco Unity Express site has 250 mailboxes or less. For deployments that require more than 250 mailboxes, consider using Cisco Unity or other voicemail solutions.
- Each Cisco Unity Express mailbox can be associated with a maximum of two different extensions, if needed.
- The automated attendant function for any office deployed with Cisco Unity Express can be local to the office (using the AA application in Cisco Unity Express) or centralized (using Cisco Unity Express for voicemail only).
- Beginning with Cisco Unity Express Release 2.0, a single Cisco Unity Express can be networked only with other Cisco Unity Expresses or with Cisco Unity. Thus, a Cisco Unity Express subscriber can forward or send messages to another remote Cisco Unity Express or Cisco Unity subscriber.
- Cisco Unity Express allows you to specify up to three Cisco Unified CallManagers for failover. If IP connectivity to all three Cisco Unified CallManagers is lost, Cisco Unity Express switches to Survivable Remote Site Telephony (SRST) call control, thus providing AA call answering service as well as mailbox access to IP phones and PSTN calls coming into the branch office.
- Cisco Unity Express automated attendant supports dial-by-extension and dial-by-name functions. The dial-by-extension operation enables a caller to transfer a call to any user endpoint in the network. The dial-by-name operation uses the directory database internal to Cisco Unity Express and does not interact with external LDAP or Active Directory databases.

Figure 14-3 shows the protocols involved in the call flow between Cisco Unified CallManager and Cisco Unity Express.

Figure 14-3 Protocols Used Between Cisco Unity Express and Cisco Unified CallManager




Figure 14-3 illustrates the following signaling and media flows:

- Phones are controlled via SCCP from Cisco Unified CallManager.
- Cisco Unity Express is controlled via JTAPI (CTI-QBE) from Cisco Unified CallManager.
- The Message Waiting Indicator (MWI) on the phone is affected by Cisco Unity Express communicating a change of mailbox content to Cisco Unified CallManager via CTI-QBE, and by Cisco Unified CallManager in turn sending a MWI message to the phone to change the state of the lamp.

- The voice gateway communicates via either H.323, SIP, or MGCP to Cisco Unified CallManager.
- Real-Time Transport Protocol (RTP) stream flows carry the voice traffic between endpoints.

Figure 14-4 shows the protocols involved in the call flow between the SRST router and Cisco Unity Express when the WAN link is down.

Figure 14-4 Protocols Used Between Cisco Unity Express and the SRST Router




Figure 14-4 illustrates the following signaling and media flows:

- Phones are controlled via SCCP from the SRST router.
- Cisco Unity Express communicates with the SRST router via an internal SIP interface.
- Currently, MWI changes are not supported in SRST mode. Voice messages can be sent and retrieved, as during normal operation, but MWI lamp state on the phone remains unchanged until the phone registers again with Cisco Unified CallManager. At that time, all MWI lamp states are automatically resynchronized with the current state of the users' Cisco Unity Express voicemail boxes.
- When not located on the same router, the voice gateway communicates via H.323 to the SRST router.
- RTP stream flows carry the voice traffic between endpoints.

Best Practices for Deploying Cisco Unity Express

Use the following guidelines and best practices when deploying the Cisco Unity Express.

- Ensure that the IP phones having Cisco Unity Express as their voicemail destination are located on the same LAN segment as the router hosting Cisco Unity Express.
- If uninterrupted AA and voicemail access is required for a site deployed with Cisco Unity Express, ensure that Cisco Unity Express, SRST, and the PSTN voice gateway are all located at the same physical site. Hot Standby Router Protocol (HSRP) or other redundant router configurations are not currently supported with Cisco Unity Express.

- Each mailbox can be associated with a primary extension number and a primary E.164 number. Typically, this number is the direct-inward-dial (DID) number that PSTN callers use. If the primary E.164 number is configured to any other number, use Cisco IOS translation patterns to match either the primary extension number or primary E.164 number so that the correct mailbox can be reached during SRST mode.
- Each Cisco Unity Express site must be associated with a CTI route point for voicemail and one for AA (if licensed and purchased), and you must configure the same number of CTI route points as Cisco Unity Express ports licensed. Ensure that the number of sites with Cisco Unity Express does not exceed the CTI scalability guidelines presented in the chapter on [Call Processing, page 8-1](#).
- Cisco Unity Express is associated with a JTAPI user on Cisco Unified CallManager. Although a single JTAPI user can be associated with multiple Cisco Unity Expresses in a system, Cisco recommends associating each dedicated JTAPI user in Cisco Unified CallManager with a single Cisco Unity Express.
- Calls into Cisco Unity Express use G.711 only. Cisco recommends using a local transcoder to convert the G.729 calls traversing the WAN into G.711 calls. You can configure Cisco Unified CallManager regions with the G.711 voice codec for intra-region calls and the G.729 voice codec for inter-region calls.
- If transcoding facilities are not available at the Cisco Unity Express site, provision enough bandwidth for the required number of G.711 voicemail calls over the WAN. Configure the Cisco Unified CallManager regions with the G.711 voice codec for calls between the IP phones and Cisco Unity Express devices (CTI ports and CTI route points).
- The CTI ports and CTI route points can be defined in specific locations. Cisco recommends using location-based call admission control between Cisco Unified CallManager and Cisco Unity Express. RSVP may also be used.
- Ensure proper Quality of Service (QoS) and bandwidth for signaling traffic that traverses the WAN between Cisco Unity Express and Cisco Unified CallManager. Provision 20 kbps of bandwidth for CTI-QBE signaling for each Cisco Unity Express site. See the chapter on [Network Infrastructure, page 3-1](#), for more details.
- The CTI-QBE signaling packets from Cisco Unified CallManager to Cisco Unity Express are marked with a DSCP value of AF31 (0x68). However, CTI-QBE signaling packets from Cisco Unity Express to Cisco Unified CallManager are not marked. Cisco recommends using access control lists (ACLs) to mark these CTI-QBE packets to ensure proper QoS. [Example 14-1](#) shows a sample configuration for achieving proper QoS for CTI-QBE signaling packets. Cisco CallManager uses TCP port 2748 for CTI-QBE signaling.

Example 14-1 QoS Configuration for CTI-QBE Signaling Packets

```
access-list 101 permit tcp host a.b.c.d any eq 2748
!
class-map match-all cti-qbe
  match access-group 101
!
policy-map cti-qbe
  class cti-qbe
    set dscp af31
    bandwidth 20
!
interface Serial0/1
  service-policy output cti-qbe
```




CHAPTER 15

Cisco Unified MeetingPlace Integration

Last revised on: February 13, 2008

This chapter describes the technical and design issues for incorporating Cisco Unified MeetingPlace into an existing Cisco Unified Communications network. The fundamental network infrastructure and IP Telephony design considerations for MeetingPlace remain the same as the IP Telephony design considerations for Cisco Unified CallManager, and this chapter assumes that you have basic knowledge and experience with Cisco Unified Communications.

This chapter focuses mainly on the integration and design issues for combining both Cisco Unified MeetingPlace and IP Telephony in one converged network. The time-division multiplexing (TDM) connection from MeetingPlace to the public switched telephone network (PSTN) is not a consideration for this setup because the PSTN access is typically provided through Cisco Unified CallManager by means of a voice gateway.

This chapter does not cover other MeetingPlace components that do not affect the integration, such as the MeetingPlace components for Outlook, Lotus Notes, Email (Simple Mail Transfer Protocol, or SMTP), Directory Services, and Instant Messaging. Also, specific product-level information about

MeetingPlace (such as feature descriptions and configuration options) is beyond the scope of this document. For more information about MeetingPlace, refer to the product documentation available at <http://www.cisco.com>.

MeetingPlace Server Recommendations

Table 15-1 lists the recommended Cisco Media Convergence Server (MCS) models to use in various deployment scenarios.

Table 15-1 MCS Deployment Recommendations

MeetingPlace Components on the Server	Up to 480 Voice User Licenses	More than 480 Voice User Licenses
All bundles (including H.323/SIP IP Gateway, Web User Interface, Email Gateway)	MCS 7835	MCS 7845
Outlook or Notes		
Directory Services		
Video Gateway		
Web Conferencing	Add one MCS 7845 for each 200 web user licenses	
Instant Messaging (IM) Gateway	Add one MCS 7835	

The following guidelines also apply to MCS deployments:

- Deployments with up to 50 Web Conferencing user licenses can run web conferencing on the same MCS with the bundled software and other options.
- Deployments with more than 50 Web Conferencing user licenses should move Outlook or Notes and the Web Conferencing Schedule, Find, and Attend functions to an MCS 7845 dedicated to Web Conferencing.
- Deployments with more than 100 Web Conferencing user licenses should not use Microsoft Desktop Engine (MSDE) 2000 because it is limited to eight concurrent connections for scheduling and web conferences. SQL 2000 is required and must be provided by the customer. Large installations (more than 500 user licenses) should run SQL 2000 on a dedicated server.
- Deployments with more than 200 voice user licenses (typically more than 10,000 records) require a dedicated MCS 7835 for Directory Integration.


Deployment Models

This section covers the major deployment models used to deploy MeetingPlace with Cisco Unified CallManager. For the purpose of this section, assume that the MeetingPlace system is placed at a major site that includes a Cisco Unified CallManager cluster.

Figure 15-1 shows Cisco Unified MeetingPlace 5.3 integrated with a comprehensive Cisco IP Communications topology. The Cisco Unified CallManager cluster (running Cisco Unified CallManager 4.0 or above) includes a video deployment (both SCCP and H.323 video endpoints) with the video conferencing capability provided by a Multipoint Control Unit (MCU). The H.323-to-H.320 video gateway handles video calls to the PSTN. Through the MeetingPlace H.323/SIP

IP Gateway, Cisco Unified CallManager fully utilizes the rich-media features provided by MeetingPlace with its audio, video, and web conferencing solution, which external users can also access from the Internet.

Figure 15-1 Example Topology for Integrating MeetingPlace with IP Telephony



The following sections describe the major deployment models for integrating MeetingPlace with Cisco Unified CallManager.

Single Site

The single-site model for IP Telephony and MeetingPlace consists of a call processing agent located at a single site and a LAN or metropolitan area network (MAN) to carry voice, video, and collaboration traffic throughout the site. Conference participants beyond the LAN or MAN use the public switched telephone network (PSTN). If an IP WAN is incorporated into the single-site model, it is for data and web collaboration traffic only; no telephony services are provided over the WAN.

For a more detailed explanation of the IP Telephony single-site deployment model, see the chapter on [IP Telephony Deployment Models](#), page 2-1.

MeetingPlace should be deployed within the data center to provide maximum resilience for the conferencing and collaboration system. Cisco recommends using Media Gateway Control Protocol (MGCP) voice gateways in a single-site model; however, an H.323/H.320 Cisco Unified Videoconferencing Gateway is required to support external video participants. The recommended method for providing external access is to use either a toll-free service or dedicated direct inward dial (DID) numbers for the MeetingPlace pilot numbers via Cisco Unified CallManager gateways, but you

also have the option to use dedicated PSTN T1 or E1 lines directly connected to MeetingPlace. Typical systems include a single toll-free number, a single DID or central office (CO) number, and an internal dialing number provided for users to dial into an audio conference session. Unique DID numbers can be used for dedicated crisis management or unique meeting IDs that are always available (24/7/365) for other applications.

Multisite WAN with Centralized Call Processing

The multisite WAN model with centralized call processing consists of a single call processing agent that provides services for many sites and uses the IP WAN to transport voice and /or video traffic between the sites. The IP WAN also carries call control signaling between the central site and the remote sites.

For a more detailed explanation of the IP Telephony multisite deployment model with centralized call processing, see the chapter on [IP Telephony Deployment Models, page 2-1](#).

MeetingPlace should be deployed at the main site to provide centralized conferencing and collaboration services for all sites. When considering the number of conference participants from each site, plan for the required WAN bandwidth and/or the number of PSTN calls to the central-site MeetingPlace. If there is insufficient bandwidth on the WAN, users will have to redial via the PSTN using either a toll-free service or a dedicated CO or DID number. If remote sites enter Survivable Remote Site Telephony (SRST) mode or switch to Cisco Unified CallManager Express for connectivity loss to the central Cisco Unified CallManager cluster, then the only access to MeetingPlace conferencing is via the PSTN, and video support is not automatically included. The web collaboration traffic, however, could still use a backup data path if one is available.

Multisite WAN with Distributed Call Processing

The multisite WAN model with distributed call processing consists of multiple independent sites, each with its own call processing agent connected to an IP WAN that carries voice and video traffic between the distributed sites.

For a more detailed explanation of the IP Telephony multisite deployment model with distributed call processing, see the chapter on [IP Telephony Deployment Models, page 2-1](#).

MeetingPlace may be deployed at one, some, or all of the distributed call processing sites. Access to the various MeetingPlace systems in different Cisco Unified CallManager clusters is by means of intercluster trunks and/or the PSTN. Web collaboration is always via the IP WAN. For maximum flexibility, implement a uniform dial plan to ensure that MeetingPlace is accessible by any user on any Cisco Unified CallManager cluster using a MeetingPlace connected to the same or any other Cisco Unified CallManager cluster. MeetingPlace also provides a feature called *multiserver meetings*, which are cascaded conferences that can conserve bandwidth or PSTN calls when multiple users at multiple sites are in the same conference. (For more information on multiserver meetings, see the chapter on [Conferencing, page 15-23](#).)

Clustering Over the IP WAN

This model deploys a single Cisco Unified CallManager cluster across multiple sites that are connected by an IP WAN with QoS features enabled.

For a more detailed explanation of clustering over the WAN, see the chapter on [IP Telephony Deployment Models, page 2-1](#).

With clustering over the WAN, you can deploy multiple MeetingPlace Audio Servers at different Cisco Unified CallManager sites. (See the information on *Dual Conference Servers* in the section on [Disaster Recovery, page 15-39](#).) The user profile information can be synchronized between the servers by using the MeetingPlace Directory Service Integration option.

On Cisco Unified CallManager, you can configure route groups and route lists so that, if one of the MeetingPlace Audio Servers fails, the calls will still be routed to another MeetingPlace system. The meeting information does not transfer between servers, so the administrator must manually upload the meeting information from the failed system onto another system for use.

If there is a WAN outage, each server is still able to operate and serve its local users. For remote users to join the meeting on the same audio server, you can configure route groups and route lists on Cisco Unified CallManager to reroute the calls.

With multiple audio servers, Cisco highly recommends turning off the Vanity ID feature (which allows users to choose and set their own meeting IDs) in order to support failover mode and reduce the probability that meeting IDs would conflict during uploads.

MeetingPlace Components

This section describes the following major components that affect the design of a MeetingPlace system:


- [MeetingPlace Audio Server, page 15-5](#)
- [MeetingPlace H.323/SIP IP Gateway, page 15-5](#)

MeetingPlace Audio Server

The MeetingPlace Audio Server (8112 or 8016) provides digital signal processor (DSP) resources for audio conferencing capability. Additionally, the MeetingPlace server acts as a scheduling module for web, audio, and video conferences. Based on the reservation, the MeetingPlace server will control how many users can join the conference as well as the duration of the meeting.

MeetingPlace H.323/SIP IP Gateway

The MeetingPlace Audio Server was originally designed for TDM connectivity, via T1/E1 trunks, either to a service provider or behind a PBX. To incorporate it into an existing Cisco Unified Communications network, the MeetingPlace H.323/SIP IP Gateway is required to link the two networks. Although it is called a gateway, it is not a hardware gateway but a software application that resides on an MCS server and is used to communicate between the MeetingPlace Audio Server and Cisco Unified CallManager. (See [Figure 15-2](#).)

Figure 15-2 MeetingPlace H.323/SIP IP Gateway Used to Interface

The MeetingPlace H.323/SIP IP Gateway supports H.323 and SIP simultaneously. For details on the protocols supported by the MeetingPlace H.323/SIP IP Gateway, refer to the section on [Interoperability Protocols](#), page 15-17.

MeetingPlace audio servers can support mixed TDM and IP connections. The capacities for a mixed environment depend on the combination of both TDM ports and IP cards (user licenses) in the system.

Although the MeetingPlace H.323/SIP IP Gateway can coexist on the same Cisco Media Convergence Server (MCS) with other MeetingPlace applications (such as MeetingPlace Web, SMTP Email, Video Integration, and so forth), in a pure IP setup Cisco recommends that you install the MeetingPlace H.323/SIP IP Gateway independently on a single MCS so that audio communications will not be affected if any problems occur with the web or video portions. However, for small deployments the MeetingPlace H.323/SIP IP Gateway can coexist with the MeetingPlace Web Conferencing Service, and the performance and capacity of the web conference should not be impacted significantly because it does not use much of the MCS CPU resources.

A single MeetingPlace Audio Server can be connected to multiple MeetingPlace H.323/SIP IP Gateways. If one MeetingPlace H.323/SIP IP Gateway fails, calls in progress are dropped, and new calls are routed to the secondary MeetingPlace H.323/SIP IP Gateway. For outdialing from the MeetingPlace Audio Server, the server chooses the MeetingPlace H.323/SIP IP Gateway with the least activity. For more information, refer to the section on [Redundancy and Load Balancing](#), page 15-39.

Each MeetingPlace H.323/SIP IP Gateway can support only one Cisco Unified CallManager without using a gatekeeper, and multiple Cisco Unified CallManagers with a gatekeeper. Each MeetingPlace Audio Server can connect to a maximum of 16 MeetingPlace H.323/SIP IP Gateways for the purposes of redundancy, load balancing, and support for multiple Cisco Unified CallManager clusters.

An Audio Server can support up to a maximum of 16 GWSIM MCS servers with MeetingPlace applications installed, so multiple IP gateways can be supported (until the combined MCS application server count = 16).

**Note**

Do *not* use Network Address Translation (NAT) between the MeetingPlace Audio Server and the MeetingPlace H.323/SIP IP Gateway.

The MeetingPlace H.323/SIP IP Gateway can communicate with Cisco Unified CallManager in one or more of the following ways:

- Cisco Unified CallManager communicates directly with the MeetingPlace H.323/SIP IP Gateway via SIP

To implement this method, configure the connection to the MeetingPlace H.323/SIP IP Gateway as a SIP Trunk in Cisco Unified CallManager.

- Cisco Unified CallManager communicates directly with the MeetingPlace H.323/SIP IP Gateway via H323

To implement this method, configure the connection to the MeetingPlace H.323/SIP IP Gateway as an H.323 gateway in Cisco Unified CallManager.

- Cisco Unified CallManager communicates with the MeetingPlace H.323/SIP IP Gateway through a gatekeeper

In this method, Cisco Unified CallManager and the MeetingPlace H.323/SIP IP Gateway register with a gatekeeper, and the gatekeeper handles the call routing.

Sizing a MeetingPlace Deployment

Sizing a Cisco Unified MeetingPlace deployment involves the following major considerations:

- Voice user licenses, or audio ports
- Web conferencing licenses
- Video conferencing IP/VC MCU sizing

When sizing a MeetingPlace system, always start with the 8100 Audio Server, and use the most accurate estimate available for the *average conferencing minutes per month*. This estimate can be derived from current billing information obtained from the conferencing service provider(s). For example, if the billing information contains a yearly total conferencing minutes, then you can either divide the total by 12 or select a peak month to use as the estimate. User surveys and feedback can also be helpful in deriving the average monthly usage.

In addition, you should add some growth factor (typically 10% to 30%) for at least the first year and possibly subsequent years. If you are incorporating the Cisco Unified MeetingPlace Outlook Integration option (which enables end users to easily schedule, notify, and attend voice, web, or video meetings), then your growth factor might be larger (possibly 30% to 50%) based on projected usage.

Use the following general guidelines when sizing a MeetingPlace solution:

- The actual conferencing usage from current service provider bills is the best measurement.
- Always add a growth percentage (typically 10% to 30%) for at least the first year before applying the sizing formulas.
- If you do not know the current conferencing usage (rare), then you can apply the following assumptions:
 - Statistically, every 20 telephony users need at least one audio conference port (license) for MeetingPlace audio service. (For example, 2500 users / 20 = 125 voice user licenses required.)
 - Each user averages 100 conferencing minutes per month. (For example, 5500 users * 100 = 550,000 minutes per month.)

You can also use these assumptions to compare to the billing data.

Sizing Voice Conferencing Usage

The general equation for calculating the number of required voice user licenses is:

$$\text{Number of MeetingPlace voice user licenses} = \frac{(\text{Average conferencing minutes per month} + \text{growth factor})}{\text{Baseline}}$$

Round any fractions up to the next whole number.

Table 15-2 lists the baseline values to use in this equation.

Table 15-2 Baseline Values for Calculating the Number of Required User Licenses

Average Minutes per Month	Baseline (Minutes per User License)
50,000 to 300,000	2,000
300,000 to 700,000	2,500
700,000 to 1,000,000	3,000
1,000,000 to 2,000,000	3,500
Above 2,000,000	4,000

For example, if you estimate the average monthly usage of your system to be 528,000 conferencing minutes, the required number of user licenses would be:

$$\text{Number of MeetingPlace voice user licenses} = 528,000 * (1 + 20\%) / 2500 = 254$$

Sizing Web Conferencing Usage

You can calculate the number of required web user licenses based on either a ratio of the total voice user licenses or service provider billing information about web conference usage. Typical deployments use 25% to 50% as many web conferencing licenses as voice conferencing licenses (25% is most common), but some enterprises deploy equal number of voice and web user licenses.

For example, a deployment with 240 voice user licenses would need 60 web user licenses to provide 25% coverage for web users. One MCS 7845 could handle the web services for this system, while also providing for growth up to 200 web user licenses.

Web conferencing directly affects the number of Cisco MCS 7800 Series servers that are required. An MCS 7835 co-located with other MeetingPlace integration modules can accommodate approximately 50 web sessions, while an MCS 7845 can accommodate up to approximately 200 web sessions (depending on the actual type of usage at a particular peak time). Therefore, sizing the web conferencing resources is critical to the overall design and placement of the solution in the enterprise data environment.

Web conferencing usage typically experiences faster growth rates than voice conferencing usage, and it depends on the amount of web collaboration used within the enterprise. For web conferencing, try to estimate growth for the next two to five years (typically 20% to 30% growth per year).

Video Integration and Sizing the MCU

MeetingPlace Video Integration is a system-wide software module, and no user licenses are required for video conferencing. The number of supported video users depends on the number of H.323 ports available on the MCU for MeetingPlace to control. The IP/VC 3511 MCU provides 15 ports (either H.323 or SCCP but not both), and the 3540 MCU provides 30, 60, or 100 ports (either H.323, SCCP, or a mixture of both). (See [Table 15-3](#).)

MeetingPlace Video Integration allows users to schedule video conferences with only a single IP/VC MCU. MeetingPlace currently does not support multiple MCUs or cascaded video meetings. You may deploy other third-party MCUs to support other videoconferencing applications that do not integrate with the MeetingPlace Video Integration solution.

MeetingPlace Video can reside on only one MeetingPlace Web MCS in the system, which can be either an internal or external web server (but not both). Video meetings are enabled only on the web server where MeetingPlace Video is installed, and they are not supported on any other web server in the system. Voice and web conferences are still supported on all servers. The MeetingPlace Video Integration software can reside on the same MCS as the Web Conferencing module, along with other MeetingPlace integration modules (such as Outlook, Notes, Directory Services, and so forth), depending on the expected number of web user licenses and the total number of web servers.

Because it can reside on only one web server, the MeetingPlace Video solution does not currently provide redundancy.

You can use the same Cisco Unified Videoconferencing 3540 or 3511 MCU for MeetingPlace video integration with other SCCP video telephony endpoints, and a single IP/VC 3540 MCU can support both ad-hoc video telephony calls (SCCP controlled) and MeetingPlace H.323 conferences. In such a deployment, you allocate some of the ports on the MCU to support SCCP for video telephony and the rest of the ports to support H.323 for MeetingPlace. This dual mode is not supported on the IP/VC 3511 MCU. To support both modes with the IP/VC 3511 MCU, you would have to deploy one 3511 MCU for video telephony and a separate 3511 MCU for MeetingPlace Video Integration. (See [Table 15-3](#).)

Table 15-3 Possible MCU Port Allocations for SCCP and H.323

IP/VC-3511-MCU (15 ports) ¹		IP/VC-3540-MC03A (30 ports)		IP/VC-3540-MC06A (60 ports)		IP/VC-3540-MC10A (100 ports)	
H.323	SCCP	H.323	SCCP	H.323	SCCP	H.323	SCCP
100%	0%	100%	0%	100%	0%	100%	0%
0%	100%	50%	50%	75%	25%	70%	30%
		0%	100%	50%	50%	50%	50%
				25%	75%	30%	70%
				0%	100%	0%	100%


1. This MCU supports 16 ports for SCCP at 128 or 384 kbps.

MeetingPlace Network Infrastructure

This section explains the requirements of the network infrastructure needed to build a MeetingPlace conferencing solution in an existing Cisco Unified Communications enterprise environment. The requirements in this section must be used in conjunction with the network infrastructure requirements described in the chapter on [Network Infrastructure](#), page 3-1.

[Figure 15-3](#) illustrates a typical network configuration for a MeetingPlace conferencing solution.

Figure 15-3 Typical MeetingPlace Conferencing Solution in an Enterprise IP Telephony Network



The following sections list the main requirements for the network that connects the MeetingPlace Audio Server and its components to the IP Telephony infrastructure:

- [Connectivity Between MeetingPlace and IP Telephony Components](#), page 15-11
- [Quality of Service \(QoS\)](#), page 15-11

Connectivity Between MeetingPlace and IP Telephony Components

Adhere to the following practices when connecting MeetingPlace to the IP Telephony network:

- Manually configure switch ports and MeetingPlace components to be 100 MB full duplex or 1000 MB full duplex.
- Physically co-locate the MeetingPlace components.
- Build redundant network connections between the MeetingPlace infrastructure and other IP Telephony products to ensure survival during WAN failure conditions.

Quality of Service (QoS)

This section discusses the following QoS mechanisms as they relate to MeetingPlace:

- [Traffic Classification, page 15-11](#)
- [Call Admission Control and Bandwidth Provisioning, page 15-12](#)
- [Jitter, page 15-14](#)

Traffic Classification

Traffic classification is a keystone for voice quality in congested networks. Voice packets must be classified or differentiated from data packets and must be handled with high priority. Some of the voice traffic is marked at its source, but other traffic needs to be classified as close as possible to the entry points of the network.

The MeetingPlace Audio Server has the following traffic classification characteristics:

- It does not support Layer 2 Class of Service (CoS) marking.
- It does not support Layer 3 voice signaling marking.
- It supports marking Real-Time Transport Protocol (RTP) packets at the source, if enabled.

The MeetingPlace Audio Server allows one of the following Layer 3 settings to be applied to RTP packets:

- IP Precedence
- Type of Service (ToS)
- Differentiated Services Code Point (DSCP, or DiffServ)



Note

The IP Precedence setting overwrites the DSCP value. Therefore, if you want to use the DSCP setting, you must manually configure the IP Precedence and ToS settings as **unused**. If the DSCP field is set to the desired value and IP Precedence and ToS are set to 0, RTP traffic will not be marked from the MeetingPlace Audio Server.

There is no mechanism within the MeetingPlace Audio Server to mark the voice signaling traffic at its source. The traffic between all MeetingPlace components must be marked as soon as it enters the network. The traffic between MeetingPlace components is usually from the Gateway System Integrity Manager (GWSIM) application, which is loaded on an MCS 8100 Series server and is always present with the MeetingPlace integration software modules (Web, Outlook, Notes, Video, and so forth). When this traffic traverses the WAN, it must be given the same priority as the other voice signaling traffic.

Cisco recommends that you use packet markings that are consistent with other devices. Refer to the chapter on [Network Infrastructure, page 3-1](#), for detailed information.

Call Admission Control and Bandwidth Provisioning

The MeetingPlace Audio Server and other MeetingPlace components do not have any call admission control mechanisms of their own. The MeetingPlace Audio Server can accept IP calls or conference requests until all ports or resources have been exhausted. This behavior can result in degraded voice quality for all audio streams across a given link if sufficient bandwidth is not available.

Call admission control should be implemented in deployments involving bandwidth-restricted links such as WAN links. Refer to the chapter on [Network Infrastructure, page 3-1](#), for detailed information regarding call admission control.

Rate and Codec Selection for Audio and Video

Audio

In an audio- only conference, the audio codec selected for a voice call depends on the regions configuration in Cisco Unified CallManager. The codec specified in the regions configuration in Cisco Unified CallManager must match the audio codec for which the MeetingPlace Audio Server is configured. Most commonly used codecs are G.711 A-law or mu-law and G.729. MeetingPlace supports both types of codecs, but only G.711 is enabled by default. If G.729 support is also desired, you should enable it during Audio Server initialization.

Video

Cisco Unified CallManager Release 5.0 supports H.261, H.263, and H.264 codecs. For a video conference, the video rate can be configured in the following places:

- MCU service view
- Cisco Unified CallManager region
- MeetingPlace user profile

The maximum video rate in a MeetingPlace user profile is set to 384 kbps. The video rate specified in the Cisco Unified CallManager region configuration will take effect only if the bandwidth requested or the video rate configured in the MCU or MeetingPlace user profile exceeds the region configuration.

For example, consider the following video bandwidth settings:

- Cisco Unified CallManager regions: 512 kbps
- MCU service view: 384 kbps
- MeetingPlace user profile: 256 kbps


In this case, if a video user dials into the MCU to join the conference, the video rate selected will be 384 kbps which corresponds to the setting on the MCU service view. The video bandwidth setting in the MeetingPlace user profile will not take effect at all because MeetingPlace Video is not involved in the call flow.

In the case of outdialing to a user, the MeetingPlace Video application will pass the video bandwidth value (256 kbps) in the MeetingPlace user profile for this user to the MCU. The MCU compares this value with its own service view value (384 kbps) and selects the lowest bandwidth value (in this case, 256 kbps) for the video conference and extends a call request to the Cisco Unified CallManager video telephony endpoint or the H.323 video endpoint. The video bandwidth selected for the call depends on the video endpoint. In the case of a Skinny Client Control Protocol (SCCP) video endpoint, the

bandwidth is selected according to the video rate settings on the MCU service view and the Cisco Unified CallManager region. (In this case 384 kbps is selected.) However, in the case of an H.323 video endpoint, the bandwidth selection is based on all three video rate settings. (In this case, 256 kbps is selected.)

Figure 15-4 shows the algorithm for selecting the video bandwidth.

Figure 15-4 Algorithm for Selecting the Video Bandwidth



The Cisco Unified CallManager region bandwidth setting (512 kbps in our example) takes effect only when the MCU or MeetingPlace Video application requests more bandwidth than the region setting (in this case, more than 512 kbps) and the endpoint dials into the video meeting. If the endpoint user employs the outdial feature to connect, then the MeetingPlace profile bandwidth setting is used.

MeetingPlace Web Session Network Utilization

MeetingPlace Web sessions generate the following amounts of network traffic:

- A first-time participant downloads approximately 1 MB of data and might be asked to accept a security notice and close the browser to reset.
- Each participant downloads between 0.5 to 0.75 MB of data at the start of the meeting.
- The average baseline utilization is about 600 Bytes/second per participant.
- In application mode, the amount of data sent when the presenter changes the view depends on the amount of color and graphics being displayed. The following general guidelines apply in such cases:
 - Low-complexity presentations send about 10 kB of data per participant.
 - Medium-complexity presentations send about 50 kB of data per participant.

- High-complexity presentations send about 430 kB of data per participant.
- Most presentations are in the low to medium range.

Cisco Unified MeetingPlace Web Conferencing will transmit 24-bit color if the bandwidth is available and will automatically reduce it to 8-bit color for slower connections. This determination is made automatically by attempting to transmit at the higher rate and reducing that rate if data is significantly delayed or lost.

Jitter

Variations in the packet delay, known as jitter, can seriously affect the voice quality. A jitter buffer smooths out variations in network delay. The MeetingPlace Audio Server contains the following parameters for setting the jitter buffer:

- Jitter Buffer Minimum Size — The jitter buffer automatically adapts to changing jitter values, but this parameter defines a minimum value for the jitter buffer.
- Jitter Buffer Optimization — This value controls how quickly the jitter buffer can react to network jitter.

In most cases, you should not change the default values of these parameters. These values may be adjusted if voice quality issues are encountered.

Domain Name System (DNS)

The MeetingPlace Audio Server does not support DNS internally. Cisco recommends assigning static IP addresses to all MeetingPlace Audio Server components and MCS components in your implementation. The Audio Server attempts to use IP addresses when accessing other servers. MeetingPlace components (such as MeetingTime System Administration client) may use host names to contact the MeetingPlace Audio Server. MeetingPlace components may also use host names to communicate with other Cisco Unified Communications products.

Network Time Protocol (NTP)


The MeetingPlace Audio Server uses NTP to synchronize its clock with the network time server. The MeetingPlace Audio Server may be configured with up to three NTP servers. MeetingPlace windows-based components can use the w32time.exe service to act as the NTP client or server. Cisco highly recommends that you use only one NTP server across the entire enterprise network so that the clocks are synchronized on all components.

Demilitarized Zone (DMZ) Requirements

MeetingPlace components placed in a DMZ allow external users access to meetings. For external users to access web conferences, a separate MeetingPlace Web server can be placed in the DMZ. Internal users can still have confidential meetings using an internal MeetingPlace Web server.

In a DMZ deployment, internal users have full web access and MeetingPlace capabilities (schedule, find, attend, access attachments, and record), while external users are allowed only to attend meetings. When the meeting starts, the internal users are redirected to the external web server.

Figure 15-5 illustrates this type of deployment.

Figure 15-5 MeetingPlace Web Conferencing with a DMZ for External Users

If MeetingPlace components in a DMZ are separated from internal MeetingPlace components by security policies such as a firewall, it will be necessary to open certain Transmission Control Protocol (TCP) or User Datagram Protocol (UDP) ports through the firewall for integration to work properly.

No ports need to be opened from the DMZ to the internal network, but the following ports from the internal network to the DMZ are needed:

- TCP 80 or 443 for web conferencing
- TCP 5003 bidirectionally for communication between the Audio Server and Web server
- TCP 5005 for attachment and recording
- TCP 1503, optionally, for Microsoft NetMeeting
- TCP 1627, optionally, for higher performance with web conferencing

The following ports are open between the DMZ and the Internet:

- TCP 80 or 443
- TCP 1503, optionally, for Microsoft NetMeeting
- TCP 1627, optionally, for higher performance with web conferencing

The DNS service must be split (one internal DNS server and one external DNS server) for the following reasons:

- Ease of use — With a single click, internal and external users can attend a meeting.
- Security — External users cannot find out information about the internal network.
- Increased performance — This deployment removes the load of internal requests from the external DNS server.

The internal DNS server resolves queries from inside the corporate network only, while the external DNS server holds the publicly addressable entities for the corporation's domain. The DNS servers must map both URLs and IP addresses.


Note

NAT is not supported between the MeetingPlace Audio Server and other MeetingPlace components due to the IP address being embedded in the data.

Table 15-4 lists all the TCP and UDP ports used by the MeetingPlace Audio Server and its components.

Table 15-4 TCP and UDP Ports Used by the MeetingPlace Audio Server

TCP or UDP Port	Protocols and/or Applications
TCP 21	FTP
TCP 23	Telnet
TCP 25	Between the MeetingPlace Email Gateway and the Simple Mail Transfer Protocol (SMTP) server
TCP 161	Simple Network Management Protocol (SNMP)
TCP 389	Between the MeetingPlace Directory Services Gateway and the enterprise corporate directory
TCP 443	Secure Socket Layer (SSL)
TCP 1443	MeetingPlace database
TCP 1627	Direct inbound access to host a meeting with WebShare
TCP 3336	Cisco Unified MeetingPlace Video integration with the IP/VC MCU
TCP 5001	Cisco MeetingTime (Opened from client to Audio Server)
TCP 5003	Cisco Unified MeetingPlace Gateway System Integrity Manager (SIM)
TCP 5005	File transfer between an Audio Server and Cisco Unified MeetingPlace Web or a gateway. (Opened from client to server)
UDP 53	DNS
UDP 123	Network Time Protocol (NTP)

Interoperability Protocols

This section describes the communication protocols used by the MeetingPlace Audio Server and its components to achieve interoperability with other Cisco products. Interoperability with the Cisco Unified MeetingPlace conferencing solution involves integration with up to two networks:

- IP Network

The MeetingPlace components (MeetingPlace H.323/SIP IP Gateway, MeetingPlace Web server, and MeetingPlace Video application) integrate directly with other products, and the MeetingPlace components then communicate with the MeetingPlace Audio Server.

- Public Switched Telephony Network (PSTN)

The MeetingPlace Audio Server integrates directly with other products through a TDM interface on the Audio Server.

Internet Protocol (IP) Network

This section describes the protocols used to integrate MeetingPlace components with an IP network. The protocols are of the following types:

- [Protocols Supported by the MeetingPlace Audio Server, page 15-17](#)
- [Protocols Supported by Other MeetingPlace Components, page 15-18](#)

Protocols Supported by the MeetingPlace Audio Server

The MeetingPlace Audio Server uses the following protocols or applications to integrate with an IP Communications system:

GWSIM

The Gateway System Integrity Manager (GWSIM) is a proprietary application that serves as an interface between the MeetingPlace Audio Server and its components. The Cisco Unified MeetingPlace GWSIM is based on a client/server architecture, which provides a means of exchanging information with the remote MeetingPlace components. The MeetingPlace Audio Server manages all the meeting information using a System Integrity Manager (SIM). The SIM uses the IP network to exchange all signaling and conferencing information with the GWSIM agent installed on the MeetingPlace components. This communication uses a proprietary Cisco protocol using TCP ports 5001 and 5003, and it is commonly referred to as GWSIM communication.

Real-Time Transport Protocol (RTP)

The MeetingPlace Audio Server uses User Datagram Protocol (UDP) for transmission of real-time audio directly to the endpoints (such as IP phones, IP/VC MCU, and so forth). The Audio Server, by default, is enabled to use only the G.711 codec for RTP. However, you can change this default to the desired codec by using the command line interface.

The MeetingPlace Audio Server requires use of the MeetingPlace H.323/SIP IP Gateway in order to support H.323 and SIP.

Protocols Supported by Other MeetingPlace Components

This section describes the protocols used by MeetingPlace components, other than the MeetingPlace Audio Server, to integrate with IP Communications networks.

H.323

H.323 networks can be integrated with MeetingPlace for a feature-rich conferencing solution. As mentioned earlier, because the MeetingPlace Audio Server does not natively support H.323, the MeetingPlace H.323/SIP IP Gateway is required to provide H.323 protocol support. The Cisco Unified MeetingPlace H.323/SIP IP Gateway supports H.323 and communicates directly with the other H.323 elements in the network to enable conferencing with the Audio Server.

The MeetingPlace H.323/SIP IP Gateway integrates directly with Cisco Unified CallManager or Cisco Unified CallManager Express using H.323. The number of simultaneous H.323 calls that a MeetingPlace H.323/SIP IP Gateway can handle depends on the maximum IP port limits of the Audio Server. Optionally, a gatekeeper can also be used for address resolution and bandwidth control when integrating with Cisco Unified CallManager or CallManager Express.

The MeetingPlace H.323/SIP IP Gateway uses the following TCP and UDP ports for the indicated purposes:


- H.323
 - Call setup for H.225 uses static TCP port 1720.
 - Call setup for H.245 uses a random TCP port in the address range 1024 to 65535.
 - RTP voice streams use random UDP ports in the range 5000 to 65535.
- Gatekeepers
 - Require static UDP port 1719 for Registration Admission Status (RAS).

**Note**

The MeetingPlace H.323/SIP IP Gateway does not support H.323 Fast Start. However, even if it receives an inbound H.323 fast-start call, the call is completed using normal H.323 signaling procedures.

Figure 15-6 shows the interoperability of MeetingPlace with Cisco Unified CallManager and CallManager Express.

Figure 15-6 *Integration with Cisco Unified CallManager and CallManager Express Using H.323*



SIP

The MeetingPlace H.323/SIP IP Gateway also provides the ability to interconnect Session Initiation Protocol (SIP) networks to the Cisco Unified MeetingPlace Audio Server. The MeetingPlace H.323/SIP IP Gateway integrates directly with a SIP proxy server or SIP gateway and converts the SIP messages to GWSIM messages, thus allowing SIP endpoints to use the MeetingPlace Audio Server for conferencing.


Cisco Unified CallManager Release 4.2 can be integrated directly with the MeetingPlace H.323/SIP IP Gateway using a SIP trunk on Cisco Unified CallManager. Using a SIP trunk between Cisco Unified CallManager and the MeetingPlace H.323/SIP IP Gateway has three significant requirements:

- The use of media termination points (MTPs) is required.
- Due to the requirement of statically enabled MTPs, only the G.711 codec is supported.
- UDP Transport must be used, and TCP is not supported. (Outgoing Transport Type must be set to UDP in the trunk configuration in Cisco Unified CallManager.)

The MeetingPlace H.323/SIP IP Gateway uses the following UDP ports for SIP:

- Call setup uses static UDP port 5060.
- RTP voice streams use random UDP ports in the range 5000 to 65535.

Figure 15-7 shows the interoperability of MeetingPlace with Cisco Unified CallManager.


Figure 15-7 SIP Integration**Note**

The MeetingPlace H.323/SIP IP Gateway can support H.323 and SIP connections simultaneously.

XML Application

The MeetingPlace Audio Server can handle only audio communications. To support video conferencing, it requires an external MeetingPlace Video application that integrates directly with the IP/VC Multipoint Control Unit (MCU). Cisco Unified MeetingPlace Video Integration communicates with the Cisco Unified Videoconferencing MCU using XML messaging through TCP port 3336.

Figure 15-8 shows the interoperability of MeetingPlace with the IP/VC MCU:

Figure 15-8 Video Integration

Public Switched Telephone Network (PSTN)

The Cisco Unified MeetingPlace Audio Server connects to the PSTN either directly or through a PBX using digital (T1 or E1) or analog trunks. The MeetingPlace Audio Server is also reachable by PSTN via Cisco Unified CallManager with a voice gateway connected.

Digital Trunks

Cisco Unified MeetingPlace supports T1 and E1 digital trunks, as described in the following sections.

T1

Cisco Unified MeetingPlace supports the following protocols for T1 digital trunks:

- T1-CAS (loop start, wink start, or ground start)
- T1-PRI (AT&T PRI, Nortel PRI, or Bell PRI)

T1 Smart Blades support T1 directly connected from either the PSTN or PBXs. The multi-access blades (MA-4 or MA-16) support T1-PRI or E1-PRI digital connections to a PBX or to the PSTN. The framing for the digital lines can be either Extended Superframe (ESF) or D4 framing. The digital lines can use either Binary 8-Zero Substitution (B8ZS) or Alternate Mark Inversion (AMI) coding.



Note

Cisco recommends using ESF framing and B8ZS coding.

Digital connections usually provide either E&M or ground start (GST) signaling. On T1 circuits configured for E&M signaling, MeetingPlace can receive only dialed number information – either Direct Inward Dial (DID/DDI) or Dialed Number Identification Service (DNIS). MeetingPlace can use dialed number information to connect the caller directly to a meeting or to determine the MeetingPlace services to which the caller has access.

Cisco Unified MeetingPlace also supports fractional T1 services.

E1

Cisco Unified MeetingPlace supports the following protocols for E1 digital trunks:

- Euro-ISDN (ETSI 300-102)
- QSIG (ECMA version) — Channels are numbered 1 to 30
- QSIG (ETSI version) — Channels are numbered 1 to 15 and 17 to 31



Note

The Cisco Unified MeetingPlace system supports only E1-PRI protocols; it does *not* support E1-CAS protocols.

Each E1 protocol allows software configuration of signaling options. The signaling options must match with the options configured on the PBX or PSTN switch.



Note

Cisco does not support mixing protocols on the MeetingPlace Audio Server, except in combination with IP ports. For example, a Cisco Unified MeetingPlace Audio Server cannot have both T1 and E1 ports configured on the same system, but it can have T1 (either PRI or CAS) and IP ports, or E1 and IP ports. Also, a Cisco Unified MeetingPlace Audio Server cannot have both T1-CAS and T1-PRI ports configured on the same system. (See [Table 15-5](#) through [Table 15-7](#).)

Table 15-5 Port Capacities for a Mixed System with T1 and IP Ports

Number of IP ports	0	96	192	240	480	576	600	960
Maximum number of T1 DS0s	1152	960	768	576	576	394	192	0
Total ports	1152	1056	960	816	1056	960	792	960

Table 15-6 Port Capacities for a Mixed System with T1-PRI and IP Ports

Number of IP ports	0	120	400	480	960
Maximum number of PRI ports	736	368	368	368	0
Total ports	736	488	768	848	960

Table 15-7 Port Capacities for a Mixed System with E1 and IP Ports

Number of IP ports	0	120	384	480	960
Maximum number of E1 ports	960	480	480	480	0
Total ports	960	600	864	960	960

For more information, refer to the *Getting Started Guide for Cisco Unified MeetingPlace Audio Server* and the *Configuration Guide for Cisco Unified MeetingPlace Audio Server*, both available at

http://www.cisco.com/en/US/products/sw/ps5664/ps5669/tsd_products_support_series_home.html

Conferencing

Cisco Unified MeetingPlace supports the following types of conferencing:

- [Audio Conferencing, page 15-23](#)
- [Web Conferencing, page 15-26](#)
- [Video Conferencing, page 15-29](#)

Audio Conferencing



There are two platforms for the MeetingPlace Audio Server:

- Cisco Unified MeetingPlace 8106 Audio Server
 - Maximum of 480 IP ports in increments of 24 or 30 user licenses
 - Maximum of 536 T1-CAS ports in increments of 24 user licenses
 - Maximum of 480 T1-PRI or E1 ports in increments of 24 user licenses
- Cisco Unified MeetingPlace 8112 — Supports up to 960 IP ports
 - Maximum of 960 IP ports in increments of 24 or 30 user licenses
 - Maximum of 1152 T1-CAS ports in increments of 24 user licenses
 - Maximum of 960 T1-PRI or E1 ports in increments of 24 user licenses

The maximum number of simultaneous meetings equals the number of ports divided by two. The servers support up to 550 sessions (participants) per meeting. However, you can cascade up to three Audio Servers to achieve a total of 1650 audio sessions in a single meeting.

Call Flow

The media request (signaling) is handled by the MeetingPlace H.323/SIP IP Gateway on behalf of the MeetingPlace Audio Server, while the media stream flows directly between the MeetingPlace Audio Server and IP phones. [Figure 15-9](#) illustrates the call flow for inbound calls, and [Figure 15-10](#) illustrates the call flow for outbound calls.

Figure 15-9 Audio Dial-in Call Flow for MeetingPlace Integrated with Cisco Unified CallManager**Figure 15-10** Audio Outdial Call Flow from Web Interface

Meeting Type

MeetingPlace supports two main types of audio meetings:

- Standard (scheduled) meeting

While scheduling this type of meeting, you can specify the meeting parameters, such as meeting ID, number of participants, and so forth. When participants join, they are connected directly to the main meeting. Immediate meetings are basically scheduled meetings in terms of the resource reservation, with the only difference being the start time (starting immediately rather than in the future).

- Reservationless meeting

Any profile user can start a reservationless meeting from the phone if that feature is enabled. The user must sign in with a profile ID and password, which starts the meeting. The database also stores information about who started the meeting.

Reservationless mode is both a system-wide parameter and a user group parameter. It can be turned on for the whole system and off for a particular set of end users, but not the reverse. This parameter will change the prompts that play when participants join the meetings.

Both types of meetings can reside on the same MeetingPlace Audio Server, and they share the same pool of conference ports. Enabling the reservationless mode does not affect the capability to schedule the standard type of meetings simultaneously. However, reservationless meetings automatically set the user's profile number as the meeting ID. Therefore, if the reservationless mode is enabled, user profile numbers become reserved and cannot be assigned manually as meeting IDs for standard scheduled meetings.

Port Management

The MeetingPlace Audio Server provides the following types of ports:

- Access ports

These ports are reserved for scheduling, attending, and listening to recorded meetings.

- Conference ports

These ports are the ones in use or reserved for meetings. Conference ports also include the following types of special-usage ports:

- Contingency ports are those conference ports reserved to handle call transfers. The system keeps them in reserve, making it possible for meeting participants to reach a contact person or attendant for assistance during a meeting and for the system manager to dial into meetings.
- Floating ports are configured to handle unexpected meeting attendance. They can float between meetings, taking up the slack when an extra person attends a meeting that is already full.

- Overbook ports

These ports do not exist physically. They are software ports configured to allow users to schedule more ports than are actually available, based on the assumption that there are usually some reserved but unused ports in some scheduled conferences. Because these ports do not exist physically, there is no limit to the number that can be configured.

Scheduling

The audio conference can be scheduled through any of the following interfaces:

- MeetingPlace Web User Interface (UI)
- Email Calendar Integration (Outlook or Lotus Notes)
- MeetingTime client
- Cisco Unified IP Phone XML Service

Audio Conference Cascading

Conference cascading is also known as a *multiserver meeting*. This feature provides a virtual link between different MeetingPlace systems so that users on each server can communicate with each other as if they were in the same meeting. When users schedule multiserver meetings, they use the web schedule on the primary server to select the secondary servers required for the meeting. At the start of the meeting, the primary server places a call to the secondary servers. The secondary servers then add the primary server to the meeting, just as they would with any participant who dialed into their system.

You can cascade up to three Audio Servers for a single meeting.

Cisco recommends configuring Network Time Protocol (NTP) servers on all MeetingPlace Audio Servers to synchronize time across all the servers.

For more information on multiserver meetings, refer to the *Administration Guide for Cisco Unified MeetingPlace Audio Server*, available at

http://www.cisco.com/en/US/products/sw/ps5664/ps5669/prod_maintenance_guides_list.html

Dialing into an Audio-Only Conference Using Video Endpoints

Video endpoints can join audio-only conferences by dialing into the MeetingPlace Audio Server directly. MeetingPlace supports both out-of-band and in-band DTMF digits from the endpoints. When using the outdial feature from the Web server, a user must enter the video endpoint extension in the audio endpoint phone number box.



Tip

For Polycom H.323 devices, users must press the # key first, then a small telephone keypad appears on the screen. Polycom devices send in-band DTMF digits, which work only when using the G.711 codec.

Web Conferencing

Web conferencing involves the following MeetingPlace components:

- [MeetingPlace Web Server, page 15-26](#)
- [SQL Database, page 15-27](#)

MeetingPlace Web Server

The MeetingPlace Web application provides the following primary functions:

- MeetingPlace Web User Interface
This feature provides scheduling, meeting room, meeting controls, reference center, and account management.
- MeetingPlace Web Conferencing
This feature provides collaboration, whiteboard, application sharing, file uploads, and more.
- Converting audio file
If MeetingPlace Web is installed with the audio service option, it first converts MeetingPlace voice (.mpv) files to .wav format, then you can choose to convert them to one of the other supported audio formats, such as Windows Media (.wma), RealAudio (.ra or .rm), and MP3. These files can be downloaded and made available to other applications such as a custom website or a CD.
- Synchronized audio-web recording
The .wav file is used by MeetingPlace Web to create a synchronized audio/web recording if the optional voice and web recording is enabled on the system.

MeetingPlace Web requires a Media Convergence Server (MCS) with the IIS service running, and it uses a Microsoft SQL database that is automatically synchronized with the MeetingPlace 8100 Audio Server master database.

The Cisco MCS 7835 server supports up to 50 simultaneous user sessions and the MCS 7845 supports up to 200 simultaneous user sessions, but there is a limit of 150 user sessions per meeting. In addition, there is a limit of 100 simultaneous web conferences for application sharing or 150 for presentation mode on a single Web server. (Presentation mode involves an uploaded presentation that is converted to JPEG format and downloaded individually to the local browser cache as each participant joins the web conference.)

Deployments with more than 100 Web Conferencing user licenses should not use Microsoft Desktop Engine (MSDE) 2000 because it is limited to eight concurrent connections for scheduling and web conferences. SQL 2000 is required and must be provided by the customer. Large installations (more than 500 user licenses) should run SQL 2000 on a dedicated server.

To increase performance in terms of capacity and load balancing, you can group MeetingPlace Web servers into clusters. The MeetingPlace Web cluster can support a maximum of six Web servers in either internal or external configurations.

**Note**

Do not use any Cisco or third-party web balancing products with the MeetingPlace solution. MeetingPlace Web has the correct web balancing built into the product, and other balancers will not work with this particular application.

Clients can connect to MeetingPlace Web via either HTTP or HTTP over Secure Socket Layer (HTTPS). The following destination ports are used on the MeetingPlace Web server:

- TCP port 5003 is used to communicate to the MeetingPlace Audio Server via GWSIM.
- TCP port 5005 is used for attachment and recording.
- TCP port 80 or 443 is used for scheduling and displaying web pages.
- TCP port 1627 (tunnel through port 80 if not open) or 443 is used for web conferencing when Secure Socket Layer (SSL) is deployed on that server. You must supply the SSL certificates for deployment on the MeetingPlace Web server.

SQL Database

The Microsoft SQL database on the Web server includes profile information as well as meeting information. The type of meeting information depends on whether the Web server is internal (hosting all internal and external meetings) or external (hosting only external meetings). The primary database is on the Audio Server. When a meeting is scheduled, the meeting information is also copied to the appropriate Web server to facilitate certain functions without having to go through the Audio Server for all requests.

Beginning with MeetingPlace 5.3, the database contains the following data:

- Cached MeetingPlace data
- Web server/site configuration data
- Tables used by Microsoft Outlook and Lotus Notes to translate new, shorter Click-To-Attend links into their longer equivalents
- Most English and localized strings used in the Web user interface
- Any strings that the system administrator has customized
- WebPolls along with their results

Microsoft Outlook and Lotus Notes do not use the database directly. They make requests to MeetingPlace Web to store Click-To-Attend links and then to attend meetings using those links.

MeetingPlace Video does not use the SQL database at all.

**Note**

Cisco Security Agent currently is *not* supported for all the MeetingPlace servers, but the Cisco-approved virus protection programs are supported.

Meeting Type

The MeetingPlace Web application supports the following types of web conferences:

- **Open Forum Meeting**
Each attendee can speak and listen equally.
- **Lecture-Style Meeting**
The majority of attendees can enter as listeners by default, with one or more designated speakers. Listeners cannot speak during a meeting.
- **Immediate Meeting**
This meeting is scheduled using the default system parameters. Users are not allowed to specify their own meeting parameters with this scheduling method.
- **Reservationless Meeting**
This type of meeting has no requirement for resource reservation and meeting ID assignment in advance. The MeetingPlace Audio Server must be set to reservationless mode before this type of meeting can be initiated. Attendees in a reservationless meeting are placed in a waiting room until the meeting moderator logs into the meeting. Attendees in the waiting room cannot speak with each other.
- **Continuous Meeting**
Only a System Manager can set up a continuous meeting (via the MeetingTime client) that is always available (24/7/365). These meeting ports are dedicated for this meeting only and are not available to the Audio system scheduling tasks.
- **Reserve All Ports**
This feature is used by system administrators to busy-out all ports on the MeetingPlace Audio Server to provide a maintenance window.

Web Conference Cascading

Web conferences cannot be cascaded; all the participants must be on the same MeetingPlace Web server.

Segmented Meetings

For external users to attend the web conference via the Internet, Cisco recommends deploying another external MeetingPlace Web server in a demilitarized zone (DMZ). This arrangement will keep the internal network confidential inside the firewall. For details on this type of deployment, see [Demilitarized Zone \(DMZ\) Requirements, page 15-14](#).

Video Conferencing

Although the MeetingPlace solution supports both H.323 and SIP, currently only H.323 video endpoints are supported for video conferences, which are handled by the MCU. All of the video conferences created by MeetingPlace are H.323 conferences, so MeetingPlace Video Integration requires H.323 resources on the MCU. In this solution, MeetingPlace plays the role of controlling application while the MCU provides the real resources and does bridging. In order to communicate with audio-only participants on the MeetingPlace Audio Server, a voice link is set up between the MCU and Audio Server, acting as a single audio participant of the video conference on the MCU.

Once MeetingPlace controls the MCU, all the H.323 ports are controlled, preventing conferences from being created directly on the MCU. MeetingPlace checks on a regular basis to ensure availability of the H.323 ports.

If the MCU is configured to allocate some resources for SCCP conferences as well, then they will still be usable for Cisco Unified CallManager to set up SCCP ad-hoc video conferences.

Voice Link

When a video endpoint participates in a conference, a special link called a *voice link* is established between the MeetingPlace Audio Server and the MCU. It acts as an audio participant dialing into the video conference on the MCU from the MeetingPlace Audio Server. All the audio endpoints of the conference connect to the MeetingPlace Audio Server, while all the video endpoints connect to the MCU. They communicate with each other via the voice link. The voice link is not established until the first video participant joins the conference. When the first video user joins the meeting, the MCU signals the MeetingPlace Video application, which causes the MeetingPlace Audio server to initiate the voice link to the MCU.

If there are multiple MeetingPlace H.323/SIP IP Gateways with various E.164 numbers used for the same Audio Server, all of these E.164 numbers must be entered in the MeetingPlace Video configuration page. The E.164 number being used depends on which MeetingPlace H.323/SIP IP Gateway that Audio Server has chosen for outdialing. The MeetingPlace Video application must be aware of all the possible E.164 numbers for the Audio Server in order to accept the voice link, no matter which number is chosen.

To utilize all the MeetingPlace features for pure video conferencing with no audio-only endpoints, the voice link still is required to connect the two systems because this link supports some advanced features on the MeetingPlace Audio Server, such as scheduling and recording the web and/or audio portion of the video conference.

Video recording is not supported on MeetingPlace Video Integration. MeetingPlace can be used to record the audio and web portions of video conferences. Third-party recording systems may be used to record the video sessions.

MeetingPlace Video

MeetingPlace Video is an application installed on an MCS that acts as a conference authorizer to the MCU. When the MCU receives a request to create a video conference or admit additional participants into an existing video conference, the MCU sends that request to MeetingPlace Video for approval.

One MeetingPlace Video application can communicate with only a single MeetingPlace Web server, and it must be installed on a machine running MeetingPlace Web 5.3. If there are multiple MeetingPlace Web servers being attached to a MeetingPlace Audio Server, only one of the Web servers can have MeetingPlace Video installed or enabled. Other MeetingPlace Web servers cannot provide the video feature for users.

If users outside the firewall need the video features, MeetingPlace Video must be enabled on the Web server located in the DMZ. Web collaboration for external users is supported via the Internet, but at this time the assumption is that external telephone and video participants will use the PSTN (or ISDN) to access gateways connected either to Cisco Unified CallManager or directly to MeetingPlace.

With MeetingPlace Video running in the DMZ, only external video meetings can be scheduled. Requests to schedule internal video meetings will fail.

If a meeting is scheduled with a request for video ports, the web conference for that meeting will be held on the MeetingPlace Web server where the MeetingPlace Video application is installed.

MCU Configuration

Cisco Unified MeetingPlace video conferencing can be integrated with Cisco Unified Videoconferencing MCUs. The MCU must be configured to authorize MeetingPlace to take control of scheduling and managing the conferences as a third-party agent. In the conference management configuration on the MCU, the **External conference authentication policy** must be set to **Authorize**.


Inbound calls can be routed to the MCU in either of the following ways:

- Through Cisco Unified CallManager
On Cisco Unified CallManager, configure the route pattern for the MCU service prefix to route to the MCU directly (defined as a gateway).
- Through a gatekeeper
On Cisco Unified CallManager, you can configure the route pattern for the MCU service prefix to route to the H.225 trunk to the gatekeeper.

Outbound calls from the MCU always require a gatekeeper for the following reasons:

- The H.323 service code defined on the MCU must be registered to the gatekeeper so that the call can be routed to the gatekeeper when a caller dials in to the conference.
- A gatekeeper is needed for address resolution because the MCU itself does not resolve the address from the H.323 video terminal.

The MCU is able to register with the gatekeeper either as a gateway or as an MCU (see [Figure 15-11](#)). Usually the MCU is registered as a gateway, but for MeetingPlace the MCU should register as an MCU. If the MCU is registered as an MCU instead of a gateway, it acts more like a terminal, with the service prefix registered to the gatekeeper as an E.164 number.

Figure 15-11 Two Methods of Registering the MCU with a Gatekeeper

A unique service code (prefix) must be configured on the MCU to be used by MeetingPlace. The protocol must be H.323. The service prefix must be the same as the prefix configured in the route pattern in Cisco Unified CallManager for routing video conference calls to the MCU. For more details, see [Video Conference Call Flow, page 15-34](#).

MeetingPlace Video requires at least one conference view to be in the designated service code, and the view must be a single-panel view. The administrator can provide the second view to take advantage of MeetingPlace Web options for switching between voice-activated (VA) and continuous-presence (CP) viewing modes. (Cisco recommends using an Enhanced Media Processor module if you want to run CP mode.) The second view must be a multi-panel view.

Enhanced Media Processor (EMP) Requirement

The Cisco Unified Videoconferencing Enhanced Media Processor (EMP) is a powerful DSP resource module that can be installed on the MCU for more complicated computing and processing. Although the EMP is not required for most video conferencing capability (either voice activated or continuous presence for H.323), it is still strongly recommended because it provides better quality with the latest MeetingPlace Video software.

To run continuous presence on SCCP video devices, EMP is always required. From the MCU's perspective, the conference is still an H.323 conference instead of an SCCP conference.



Note

MeetingPlace video integration does not support the Rate Matching Module and the Data Collaboration Server on Cisco Unified Videoconferencing MCU platforms.

Port Management

MeetingPlace Video provides the following types of ports:

- Video Conference Ports

These ports are assigned dynamically based on the setup of the service code configured on the MCU. MeetingPlace Video periodically verifies if the resources (maximum number of video ports and conferences) have been modified in the MCU, and it updates the MeetingPlace Audio Server accordingly.

- Video Floating Ports

These ports are allocated from the pool of video conference ports. They enable the system administrator to reserve a number of video ports for conferences that require more video ports than originally scheduled. A minimum number of floating ports is required to facilitate Voice Links. The formula for calculating this minimum number of floating ports is documented in the *Administrator's Guide for Cisco Unified MeetingPlace Video Integration*, available at

http://www.cisco.com/en/US/products/sw/ps5664/ps5669/prod_maintenance_guides_list.html

- Video Overbook Ports

These ports do not exist physically. They are software ports configured to allow users to schedule more ports than are actually available, based on the assumption that there are usually some reserved but unused ports in some scheduled conferences. Because these ports do not exist physically, there is no limit to the number that can be configured.

Scheduling

The MeetingPlace video conference can be scheduled through the following interfaces only:

- Web interface
- Outlook or Notes interface
- MeetingTime

Currently the Voice User Interface does not support video scheduling capability.

The video conference is not set up until the first video participant joins the conference. Users cannot get the video resources on the fly. The conference must be scheduled in advance or it must be reservationless.

The video conference is created using the designated service code defined on the MCU. Only one service code is allowed for MeetingPlace video integration. That service code controls the behavior and default parameters of all the MeetingPlace video conferences.

The dial-in number for the each created conference is unique and has the following format:

Designated service code (configured on MCU) + *Meeting ID* (assigned by MeetingPlace)

Either immediate or reservationless meetings can be created.

Attending a Video Conference

Video endpoints can join the meeting in either of the following ways:

- Outdial from MeetingPlace

Users join the conference first via MeetingPlace Web or MeetingPlace Outlook. The endpoint address is part of the user's profile and will appear in the video endpoint address box. The user then clicks to have MeetingPlace outdial to the video endpoint. This is the easier way for the video endpoints to join a meeting.

- Dial into the MCU

The number to dial is formatted as described in the section on [Scheduling, page 15-32](#). This method is different than the typical user experience for audio-only endpoints because the prompts do not ask the video user to enter the meeting ID, as is the case when dialing into the MeetingPlace Audio Server. Because video conferencing does not currently use authentication, meetings with the following characteristics will not allow video dial-in:

- Password-required meeting
- Invitation-only or profile-only meeting (vs. public/anyone meeting)

For these types of meetings, users join from the MeetingPlace Web interface or Outlook.

Outdialing to video endpoints is not automatic for video conferences because the endpoints cannot be defined when the meeting is scheduled. After the meeting has started, users can click on the profile (phone number) to outdial.

Meeting Type

All video conferences are created on an ad-hoc basis using one of the following types:

- Standard reserved or scheduled meeting
- Continuous meeting

As with a scheduled meeting, the video conference is created when the first video user joins the conference. The video conference does not terminate until the entire meeting (audio and video) has been terminated by the MeetingPlace platform. The voice link between the Audio Server and the MCU will terminate 5 minutes after the last video participant leaves the conference.

- Immediate meeting (system without reservationless meetings enabled)

Even if MeetingPlace does not have reservationless mode enabled, users still can start the immediate type of meetings via the Web interface as long as the current conference number and video ports in use do not exceed the capacity in the MCU.

- Reservationless meeting

This type of meeting does not reserve any video ports in advance. Users may join this type of meeting as long as there are video conference slots and floating ports available. If a video user joins the reservationless meeting before it starts, the corresponding video conference will be created but the user will be put in the waiting room first until the conference initiator joins the meeting from the audio side. While in waiting mode, users cannot communicate with each other and will see a visual message stating that the meeting is not yet in session.

- Lecture-style meeting

There are two types of participants in these meetings: speakers and listeners. Video users are always treated as listeners when they join, and they have no way to enable themselves as speakers even if they are really scheduled as speakers. The only way to join the lecture-style conference as a speaker is join the conference as an audio participant rather than a video participant. The speaker has the ability to choose starting the meeting with participants in the waiting room (muted and video blocked), as listeners (muted but able to see the image), or as an open floor (can communicate with audio and video immediately).

- Multiserver meeting

See the section on [Video Conference Cascading, page 15-37](#).

Video Conference Call Flow

When the first video participant joins a conference, MeetingPlace Video opens an XML control channel to the MCU and initiates an outdial from the audio server to the MCU, joining the audio conference to the video conference. The link behaves as an audio participant in both conferences.

Figure 15-12 illustrates the process involved in creating the voice link when the first conference participant initiates the video conference. Before a user can initiate the video conference, the MCU must first be registered to the gatekeeper as an MCU, and the MCU must also be defined in Cisco Unified CallManager as an H.323 gateway.

Figure 15-12 Video Conference Initiation Process




Figure 15-12 illustrates the following steps in the video conference initiation process:

1. The first video participant dials 8151234 (where 1234 is the meeting ID created by MeetingPlace), and the call is routed to Cisco Unified CallManager.
2. In Cisco Unified CallManager, the call matches route pattern 815XXXXX, which points to the H.323 gateway that is defined in Cisco Unified CallManager to represent the MCU.
3. The call is routed to the MCU.
4. The MCU sends authorization to MeetingPlace Video to verify the meeting and participant information.
5. The MCU also registers the newly generated conference number (8151234) with the gatekeeper as an additional E.164 number.
6. MeetingPlace Video notifies the Audio Server to create the voice link.
7. The Audio Server dials 8151234 (the MCU) to join the video conference.

8. The call is directed to Cisco Unified CallManager, and again the same route pattern for this number is matched to route the call.
9. This call from the Audio Server is routed to the MCU.
10. The voice link is created between the Audio Server and the MCU.
11. The audio and video streams for the first video participant are set up between the endpoint and the MCU.

For the second and later video participants, the call flow follows steps 1, 2, 3, 4, and 11.

For video outdialing, the request is passed to the MCU from the MeetingPlace Video Integration (see [Figure 15-14](#)). For dial-in to the MCU from endpoints, the call flow is identical to the normal dial-in video call flow, with MeetingPlace getting involved only to check the user authentication (see [Figure 15-15](#)).

Figure 15-13 Call Flow for Video Conference Initiation (Voice Link Creation)






Figure 15-14 Call Flow for Video Conference Outdialing**Figure 15-15** Call Flow for Video Conference Dial-In

Video Conference Cascading

Currently video cannot be cascaded between MCUs in the MeetingPlace implementation. Different MeetingPlace Audio Servers can be cascaded for the same conference, which is referred to as a *multiserver meeting*, with the various MeetingPlace systems communicating with each other via the voice links. With a multiserver meeting, all participants are on the Web server associated with the primary Audio Server. Video participation can occur if the Web server used has video integration installed.

For more information on multiserver meetings, refer to the *Administration Guide for Cisco Unified MeetingPlace Audio Server*, available at

http://www.cisco.com/en/US/products/sw/ps5664/ps5669/prod_maintenance_guides_list.html

Gatekeeper and Dial Plan

A gatekeeper is a required component when H.323 video terminals are integrated with an IP Communications system. The gatekeeper mainly provides address resolution for the H.323 endpoints when they initiate calls.

With Cisco IOS Release 12.3(8)T or later, Cisco Unified CallManager can register itself to a via-zone gatekeeper as an IP-to-IP gateway (a new software feature in Cisco voice gatekeepers). Also, it can register H.323 endpoints with aliases (dynamic addresses) rather than with static IP addresses. These new features greatly influence H.323 video design with regard to MeetingPlace Video Integration, as described in the following sections.

When using multiple Cisco Unified CallManager clusters with a centralized gatekeeper providing intercluster trunk call admission control, the recommendation is to separate the gatekeepers used for video endpoints from those used for intercluster trunks. This practice allows for a more flexible dial plan. If the video zone and the intercluster trunk zone are configured on the same gatekeeper, the zone prefix for the intercluster trunk zone cannot overlap any endpoint numbers.

For more detailed information regarding the Gatekeeper and the Dial Plan, see the chapter on [Dial Plan](#), page 10-1.

Dynamic H.323 Addressing in Cisco Unified CallManager

In Cisco Unified CallManager 4.0, H.323 clients are configured with static IP addresses. In Cisco Unified CallManager 4.1 and higher, H.323 clients can also be configured with E.164 addresses. E.164 addressing simplifies H.323 client administration for deployments where H.323 clients are configured for Dynamic Host Configuration Protocol (DHCP). E.164 addressing works in conjunction with a via-zone gatekeeper, which can be configured to route every call to Cisco Unified CallManager (acting as an IP-to-IP gateway) in the **outvia** zone (which is the Cisco Unified CallManager zone).

With this configuration, Cisco Unified CallManager can automatically create a special trunk, called the RasAggregator trunk, and register it to the gatekeeper. As the name suggests, this trunk is used for Registration Admission Status (RAS) and to administer and manage all the dynamically addressed H.323 devices defined in Cisco Unified CallManager as an aggregated device registered to the same gatekeeper. Whenever there are calls involving those H.323 devices, the gatekeeper communicates with Cisco Unified CallManager through this RasAggregator trunk.

Cisco Unified CallManager Redundancy Groups and H.323 Clients

The Cisco Unified CallManager redundancy group defined in the device pool of an H.323 client is important because it affects how calls may be placed using Cisco Unified CallManager and, if configured incorrectly, can cause calls to fail. The following rules apply when configuring redundancy groups:

- If an H.323 endpoint uses peer-to-peer mode instead of a gatekeeper, the H.323 client must be configured to use the subscribers in the same priority order as listed in the redundancy group for that H.323 client device on Cisco Unified CallManager.
- For each Cisco Unified CallManager redundancy group, you must configure a unique zone per gatekeeper.

MCU Registration

If the MCU registers with the gatekeeper as a gateway, its service prefix (for example, 85) is registered as a technology prefix. If an incoming call to the gatekeeper has a matching technology prefix in the called number, the gatekeeper will first try to look for the IP-to-IP gateway that is registered with that particular prefix in the via-zone. Therefore, when an H.323 endpoint dials into the MCU (using an access code of 851234, for example), the dialed number matches the technology prefix of 85 but the call fails because the IP-to-IP gateway (which is Cisco Unified CallManager) is actually registered with a technology prefix of 1#.

To avoid unwanted matches with the technology prefix, Cisco recommends registering the MCU as an MCU rather than a gateway. Registering the MCU as an MCU eliminates the technology prefix registration, thereby preventing the gatekeeper from finding an IP-to-IP gateway with that particular technology prefix.

Cisco also recommends configuring the MCU to register the conference IDs so that individual E.164 numbers are registered with the gatekeeper whenever a conference ID is created.

MeetingPlace

Because MeetingPlace registers with the gatekeeper as a terminal, you can put the MeetingPlace H.323/SIP IP Gateway in one zone with all other H.323 video endpoints. The RasAggregator trunk from Cisco Unified CallManager also registers to this same zone.

Reservationless Single Number Access (RSNA)

The Reservationless Single Number Access (RSNA) feature enables multiple MeetingPlace Audio Servers to appear as one server to the user community. This feature is designed primarily for reservationless meetings. With this feature, users can dial a single directory number to access multiple Audio Servers deployed in one location. Users dial in and enter their reservationless ID, and the system directs the calls to the appropriate server for that meeting. This feature also enables users in a remote location to dial the local server.

RSNA relies on SIP endpoints that support the REFER method. Cisco Unified CallManager 5.0 supports REFER, thus allowing RSNA to be implemented. Previous versions of Cisco Unified CallManager did not support REFER or RSNA.

For more details, refer to the MeetingPlace product documentation, available at

http://www.cisco.com/en/US/products/sw/ps5664/ps5669/tsd_products_support_series_home.html

Redundancy and Load Balancing

This section describes redundancy and load balancing considerations for the following MeetingPlace components:

- [MeetingPlace Audio Server, page 15-39](#)
- [MeetingPlace H.323/SIP IP Gateway, page 15-40](#)
- [MeetingPlace Web, page 15-41](#)
- [Cisco Unified CallManager, page 15-41](#)
- [MeetingPlace Video, page 15-42](#)
- [MCU, page 15-42](#)

MeetingPlace Audio Server

If the MeetingPlace Audio Server fails, calls in progress will be dropped. If the Audio Server is unreachable by the MeetingPlace H.323/SIP IP Gateway, the MeetingPlace H.323/SIP IP Gateway will immediately reject any incoming calls. If the MeetingPlace H.323/SIP IP Gateway is down, Cisco Unified CallManager will be unable to establish an H.323 connection with the MeetingPlace H.323/SIP IP Gateway. There is no automatic failover mechanism available for the MeetingPlace Audio Server.

Disaster Recovery

A disaster recovery plan provides for orderly restoration of the database, computing, and services. The plan usually includes the following provisions:

- Redundancy (spare parts of hardware)
- Data and software backups
- Alternate emergency locations
- Operations split across multiple sites

A component called the MeetingPlace Network Backup Gateway enables system managers to transfer multiple copies of the MeetingPlace database (audio configuration files and scheduled meeting information) from the Audio Server to a designated storage server on the network. Files are encrypted for security. A maximum of three Network Backup Gateways per system can be implemented, but only one of them can make the database transfer at a time.

Redundancy

If two or more MeetingPlace Audio Servers are deployed, any of the following plans for redundancy can be implemented:

- Shadow Server

A shadow server is a redundant MeetingPlace Audio Server that is not active until the primary Audio Server fails. During the disaster, the shadow server must be switched from shadow mode to active mode using the command line interface (CLI). Once active, the shadow server can store only limited information, such as profile information, meeting information (past, present, and future), and

participant information. Recording, attachments, hardware and network configuration, and automatic software upgrades are not available on the shadow server. The shadow server has the following network connectivity requirements:

- Less than 250 ms round trip delay
- Less than 1% packet loss
- Minimum bandwidth of 384 kbps
- Dual Conference Servers

In this plan, two MeetingPlace Audio Servers are active production servers and each one can overflow to the other. Profiles are synchronized automatically between the servers through the Directory Services. Automatic failover between the two Audio Servers does *not* occur. Meetings have to be uploaded to the remaining server during an emergency, and users must rejoin the meetings manually. On Cisco Unified CallManager, you can configure route groups and route lists so that if one MeetingPlace Audio Servers fails, the calls will be routed to the other active MeetingPlace Audio Server.
- Continuous Meeting Server

In this plan, a MeetingPlace Audio Server is implemented with a set of critical meetings pre-created and always available. This deployment is suitable for any crisis management application, with features such as blast outdial to crisis management teams.

MeetingPlace H.323/SIP IP Gateway

The following redundancy and load balancing considerations apply to the MeetingPlace H.323/SIP IP Gateway.

Redundancy

If the MeetingPlace H.323/SIP IP Gateway fails, calls in progress will be dropped.

Two or more MeetingPlace H.323/SIP IP Gateways can connect to the same MeetingPlace Audio Server. Dial-in redundancy can be implemented in the following ways:

- Configure all of the MeetingPlace H.323/SIP IP Gateways on Cisco Unified CallManager and put them in the same route group. If the link between Cisco Unified CallManager and the MeetingPlace H.323/SIP IP Gateway fails, Cisco Unified CallManager chooses the next MeetingPlace H.323/SIP IP Gateway in the current route group. This same method also works for load balancing.
- Have all the MeetingPlace H.323/SIP IP Gateways register to an H.323 gatekeeper in a specific zone, then the gatekeeper can take care of the failover. Configure **gw-priority** on the gatekeeper to specify the chosen priority of each MeetingPlace H.323/SIP IP Gateway.

For outdial, the MeetingPlace Audio Server uses the following algorithm:

- If all the MeetingPlace H.323/SIP IP Gateways are functioning correctly with no outdial failures at all, the Audio Server will choose the least busy gateway for outdial, thus load balancing between them.
- If any of the MeetingPlace H.323/SIP IP Gateways has a failure, the Audio Server will skip that one and do load balancing on the remained gateways that do not have failures.
- If all the servers have had failures, then the Audio Server will choose the one with the oldest failure time and will keep using that same gateway until it fails again.

Load Balancing

For dial-in load balancing, you can use either of the following same approaches as for redundancy:

- Use a route group in Cisco Unified CallManager, but choose **Circular** instead of **Top down** for the Distribution Algorithm. Cisco Unified CallManager will distribute the outbound calls between the gateways in round-robin fashion.
- Use a gatekeeper. For purposes of load balancing, do not configure **gw-priority** to specify a preferred gateway. Instead, let the gatekeeper distribute the calls in round-robin fashion to all the registered gateways.

As explained in the previous section, load balancing for outdial is done only among those MeetingPlace H.323/SIP IP Gateways that never have outdial failures.

MeetingPlace Web

The following redundancy and load balancing considerations apply to the MeetingPlace Web.

Redundancy

If a web-conferencing server fails during a web conference, all users are disconnected from the web portion of the meeting, and the current web conference cannot continue. Users can rejoin their web conference on another active web server. If the SQL server database is on the failed server, all web conferencing servers in the cluster become nonfunctional and users will be unable to conduct web conferences until the database is restored.

Load Balancing

MeetingPlace Web servers can be grouped into clusters to increase performance for load balancing (to distribute the meeting requests to different web servers) and to increase the capability of handling more meetings. The MeetingPlace Web cluster can contain a maximum of six web servers in either internal or external configurations.

The MeetingPlace WebConnect feature provides integration between multiple internal and external MeetingPlace systems through a single URL. It can schedule across multiple Audio Servers. If there is not enough capacity on the first Audio Server, it will try the second, and so on.

Cisco Unified CallManager

The following redundancy and load balancing considerations apply to Cisco Unified CallManager.

Redundancy

If Cisco Unified CallManager fails, MeetingPlace calls in progress using that Cisco Unified CallManager will be dropped, and users must rejoin the meeting. For dial-in from endpoints, redundancy is achieved with redundancy groups configured in the Cisco Unified CallManager cluster.

For outdialing from MeetingPlace through a MeetingPlace H.323/SIP IP Gateway, each MeetingPlace H.323/SIP IP Gateway can communicate with only one Cisco Unified CallManager. Therefore, the following guidelines apply to redundancy:

- If the MeetingPlace H.323/SIP IP Gateway is configured to communicate with Cisco Unified CallManager directly as a gateway, redundancy can be implemented only by connecting two or more MeetingPlace H.323/SIP IP Gateways to the same MeetingPlace Audio Server.

- If the MeetingPlace H.323/SIP IP Gateway is registered to a gatekeeper, then the gatekeeper will handle the failover situation, routing calls to an active Cisco Unified CallManager.

Load Balancing

For dial-in from endpoints to MeetingPlace, load balancing for calls is achieved through Cisco Unified CallManager redundancy groups and route group configurations.

For outdialing from MeetingPlace through a MeetingPlace H.323/SIP IP Gateway, the same guidelines apply for load balancing as for redundancy.

MeetingPlace Video

The following redundancy and load balancing considerations apply to the MeetingPlace Video.

Redundancy

Only one MeetingPlace Video application can be implemented in a MeetingPlace 5.3 system. MeetingPlace Video must be installed on a server running MeetingPlace Web 5.3. MeetingPlace Video can communicate with only one MeetingPlace Web server.

If a meeting is scheduled with video ports requested, the web conference for that meeting will be held on the MeetingPlace Web server where MeetingPlace Video is installed. If that server fails, the meeting will roll over to another server after a period equal to five times the Load Stats Poll Period configured on the site administration UI page. (The default is one minute.) If a conference with scheduled video ports has rolled over to another web server, no video features will be provided in the GUI. Users will not be able to outdial, but they can still choose to dial in.

Load Balancing

Because only one MeetingPlace Video application can be active at a time, there is no load balancing for MeetingPlace Video.

MCU

Currently with MeetingPlace 5.3 video integration, you can have only one MCU per MeetingPlace Audio Server. Therefore, there is no redundancy or load balancing for the MCU.



CHAPTER 16

Cisco Unified MeetingPlace Express

Last revised on: February 13, 2008

Cisco Unified MeetingPlace Express is a rich-media appliance for voice and web collaboration. This chapter addresses system-level design considerations. For detailed product information, refer to the product documentation available on

<http://www.cisco.com>

Deployment Models

This section describes design recommendations for integrating Cisco Unified MeetingPlace Express with the following major IP Telephony deployment models:


- [Single Site, page 16-2](#)
- [Multisite WAN with Centralized Call Processing, page 16-3](#)
- [Multisite WAN with Distributed Call Processing, page 16-4](#)
- [Clustering Over the WAN, page 16-5](#)

In addition to these deployment models, this section covers deployments based on a demilitarized zone (DMZ). The DMZ deployment is treated separately and is applicable to each of the IP Telephony deployment models listed above. (See [Security and DMZ Deployment, page 16-5](#).)

Single Site

In a single-site deployment, all call processing is local and all participants are local as well. Bandwidth is usually not a concern, as it would be in multisite configurations. In this model, the Cisco Unified MeetingPlace Express server connects to Cisco Unified CallManager through a gatekeeper, H.323 trunk, or SIP trunk. (See [Figure 16-1](#).)

Figure 16-1 Single-Site Deployment




The general design considerations presented in the remaining sections of this chapter apply to this basic deployment model.

Multisite WAN with Centralized Call Processing

In a multisite WAN deployment with centralized call processing, the Cisco Unified MeetingPlace Express server is usually located in the central site, where all call processing occurs. (See [Figure 16-2](#).)

Figure 16-2 Multisite WAN Deployment with Centralized Call Processing




In this deployment, participants in the central site have no special considerations, but participants at remote sites require the following additional considerations:

- Cisco Unified MeetingPlace Express supports only G.711. Transcoders at the central site must be used for all remote calls other than G.711.
- Web collaboration requires significant bandwidth that must be provisioned. (See [Bandwidth Considerations for Web Applications and Screen Sharing](#), page 16-7.)
- The same QoS and call admission control mechanisms used for the existing IP Telephony deployment are required. For purposes of call admission control, the Cisco Unified MeetingPlace Express server should be treated as a telephony endpoint in the central site.

Multisite WAN with Distributed Call Processing

In a multisite WAN deployment with distributed call processing, the Cisco Unified MeetingPlace Express server is usually located in a single site local to Cisco Unified CallManager, similar to a multisite WAN deployment with centralized call processing. [Figure 16-3](#) depicts an example of a multisite WAN deployment with distributed call processing.

Figure 16-3 Multisite WAN Deployment with Distributed Call Processing



This deployment differs from a multisite WAN deployment with centralized call processing because the call processing is distributed to multiple sites. These sites may use Cisco Unified CallManager, Unified CallManager Express, or other methods of call processing. Because Cisco Unified MeetingPlace Express is a single server with no cascading, mirroring, or redundancy capabilities, a single active server at a single site would be typical. An additional server as a secondary active server may be deployed at a separate site but would not allow for a true failover environment. See [Redundancy](#), page 16-12, for more information.

To have a multisite redundant and distributed collaboration system, consider Cisco Unified MeetingPlace as an option. See the chapter on [Cisco Unified MeetingPlace Integration](#), page 15-1, for more information.

Design Considerations

In this deployment, participants local to the Cisco Unified MeetingPlace Express server have no special considerations, but participants remote from the Cisco Unified MeetingPlace Express server require the following additional considerations:

- Cisco Unified MeetingPlace Express supports only G.711. Transcoders at the site containing the Cisco Unified MeetingPlace Express server must be used for all remote calls other than G.711.
- Web collaboration requires significant bandwidth that must be provisioned. (See [Bandwidth Considerations for Web Applications and Screen Sharing, page 16-7](#).)
- The same QoS and call admission control mechanisms used for the existing IP Telephony deployment are required. For purposes of call admission control, the Cisco Unified MeetingPlace Express server should be treated as a telephony endpoint in its local site.
- The Cisco Unified CallManager cluster at the remote site (Site 2 in [Figure 16-3](#)) can directly route calls to Cisco Unified MeetingPlace Express located at the central site (Site 1) by defining an H.323 gateway or SIP trunk directly on the remote cluster. This practice is not recommended due to its potential impact on call admission control. Routing calls directly to MeetingPlace and not through the central cluster via an intercluster trunk or through gatekeeper can cause calls not to be accounted for in call admission control measurements.

Clustering Over the WAN

In a deployment with clustering over the WAN, the Cisco Unified CallManager cluster is split across one or more locations separated by high-capacity and high-speed WAN or MAN links. Cisco Unified MeetingPlace Express does not have any server-to-server redundancy and cannot fully utilize the redundant nature of this model. An additional server as a secondary active server may be deployed at the alternate site but would not allow for a true failover environment. See [Redundancy, page 16-12](#), for more information.

Placement and design considerations would be identical to the previous deployment model, [Multisite WAN with Distributed Call Processing, page 16-4](#).

To have a multisite redundant and distributed collaboration system, consider Cisco Unified MeetingPlace as an option. See the chapter on [Cisco Unified MeetingPlace Integration, page 15-1](#), for more information.

Security and DMZ Deployment


Cisco Unified MeetingPlace Express may be placed in a demilitarized zone (DMZ), although there is no mechanism to allow internal and external applications because it is a single server. All internal and external web collaboration and telephony calls must reach the Cisco Unified MeetingPlace Express server in the DMZ.

External telephony calls would come into a gateway inside the firewall and be treated the same as internal calls being sent into the DMZ to connect to the Cisco Unified MeetingPlace Express server.

External web collaboration would come into the DMZ from outside the firewall, and internal from inside.

[Figure 16-4](#) illustrates communication between the Cisco Unified MeetingPlace Express server in the DMZ and the internal and external networks.

Figure 16-4 DMZ Deployment



Design Considerations for DMZ Deployment

- If port 1935 is not configured in the firewall access control list (ACL), Real Time Messaging Protocol (RTMP) traffic will tunnel over port 80 with a small performance hit.
- If Network Time Protocol (NTP) is to be used, connectivity internally or externally to an NTP source over port 123 is needed.
- Secure Shell (SSH) access over port 22 to the server from the internal network might be needed. Cisco Technical Assistance Center (TAC) might require SSH access to support Cisco Unified MeetingPlace Express properly.
- Real-Time Transport Protocol (RTP) and Real-Time Transport Control Protocol (RTCP) ports are required for voice termination on the Cisco Unified MeetingPlace Express server. With Cisco firewall devices, RTP ports will be opened dynamically to allow the correct traffic flows through the firewall based on the H.323 or SIP call setup. Therefore, it is not necessary to open a range of ports on the firewall.

Call Admission Control, Bandwidth, and QoS

This section describes important call admission control, bandwidth, and QoS considerations for Cisco Unified MeetingPlace Express deployments.

Call Admission Control

Cisco Unified MeetingPlace Express should be treated as a local voice gateway for purposes of call admission control. For additional guidance on call admission control, see the chapter on [Call Admission Control, page 9-1](#).

Bandwidth Considerations for Web Applications and Screen Sharing

Screen sharing consumes large amounts of bandwidth, on the order of 2 Mbps per user in the worst case. Because of the excessive bandwidth and bursty nature of screen sharing, the following design considerations must be observed.

Design Considerations

Use of 100 Mbps for LAN connectivity will limit the number of concurrent web collaboration sessions. When possible, use 1 Gbps connections.

Users at remote sites that cause web collaboration to traverse a WAN will require special consideration. The client flash session bandwidth or room bandwidth setting for these users should be lowered, reducing the load across the WAN. Because web collaboration data is delivered unicast, the largest burst of data should be multiplied by the number of clients at a remote site. For example, assume a remote site has 100 users, 10 of which are on a web collaboration session at any one time. If bursts of 2 Mbps occur in the data from the remote server to each user, 20 Mbps bursts can be experienced across the WAN connection.

When WAN links become congested from excessive web collaboration data or other sources, the degradation of all traffic is compounded by packet loss, retransmissions, and increased latency. Sustained congestion will have a sustained degrading impact on all remote collaboration sessions.

The following settings on the client web collaboration sessions control the rate at which the participant receives data as well as the rate at which the presenter sends data:

- **Modem** — Bandwidth limited to 64 kbps
- **DSL** — Bandwidth limited to 640 kbps
- **LAN** — Bandwidth up to 3,000 kbps (Bandwidth bursts above 3,000 kbps are possible if high-resolution images or photos are shared. Sharing normal to complex presentations or document should not generate bursts above 3,000 kbps unless large complex images are embedded.)

Bandwidth settings are not automatically adjusted when congestion occurs; they must be adjusted manually. Bandwidth settings default to LAN and must be set at the initialization of each collaboration session. A new session is set to LAN regardless of previous settings.

Bandwidth settings can be adjusted in the following locations in the web collaboration client screen:

- My Connection Speed
 - Limits the speed data is delivered to an individual user from the Cisco Unified MeetingPlace Express server.

- Limits the speed data is sent from the presenter to the Cisco Unified MeetingPlace Express server.
- Presenter speed does not limit the rate at which participants receive data.
- Optimize Room Bandwidth
 - Limits the rate data is delivered to all users.
 - Currently has data burst issues and might impact WAN bandwidth utilization.

**Note**

The setting for My Connection Speed on the client web interface for the presenter and participant are independent of each other. Presenter speed does not limit the rate at which participants receive data. If the presenter is set to send data to Cisco Unified MeetingPlace Express at Modem speed (70 kbps) and a participant is set to receive data at LAN speed (3,000 kbps), the participant will still receive data up to 3,000 kbps.

**Note**

While the Optimize Room Bandwidth setting does reduce the rate data is sent to users, bursts in the delivery of that data still occur and can congest WAN links. Changing My Connection Speed on the participant's client eliminates bursts and delivers data to the client at a steady rate. The burst issues with the Optimize Room Bandwidth setting will be addressed in future releases.

The client resolution setting also impacts the bandwidth used by limiting the bits per image. A meeting room set at 640x480 pixels will typically generate less than one-third of the traffic compared to a setting of 1280x1024 pixels.

Design Recommendations Summary

- Connect both Cisco Unified MeetingPlace Express server interfaces to the LAN with 1 Gbps connections.
- For meetings involving remote users, have the remote users set their My Connection Speed to **DSL**. If congestion issues occur, lower the setting to **Modem**, or lower the screen resolution settings.

QoS

Quality of Service (QoS) considerations encompass the following topics.

Traffic Classification or Markings

Traffic classification or markings from the Cisco Unified MeetingPlace Express server are as follows:

- SIP, H.323, and call control — Not marked
- Voice media (RTP) — Marked EF (DSCP 46)
- Web traffic (HTTP and HTTPS) — Not marked
- Flash web collaboration (RTMP) — Not marked

Voice and Call Control Traffic

Voice and call control traffic should be classified with the standard classifications as described in the chapter on [Network Infrastructure, page 3-1](#). If the connection between the Cisco Unified MeetingPlace Express server and the gatekeeper or Cisco Unified CallManager extends over a WAN connection, call control traffic should be re-marked to CS3 (DSCP 24).

Web Interface Traffic

No prioritization is recommended for the main web interface. Web traffic to the Cisco Unified MeetingPlace Express server should be treated the same as traffic to other internal web application servers.

Flash Web Collaboration Traffic

Because of the large amounts of bandwidth consumed and the bursty nature of this data, as described in the previous section, prioritization of web collaboration traffic across the WAN is difficult. If you attempt to prioritize the web collaboration traffic, classify it separately and lower than other prioritized data.

Directory Integration

Cisco Unified MeetingPlace Express integrates directly with Cisco Unified CallManager 4.x. For third-party directory integration, a plug-in is required. Cisco Unified MeetingPlace Express supports Microsoft Active Directory, Sun One LDAP, and Netscape/iPlanet LDAP.

Cisco Unified MeetingPlace Express also integrates directly with Cisco Unified CallManager 5.x via Cisco AVVID XML Layer (AXL) Simple Object Access Protocol (SOAP). Third-party directory integration in a Cisco Unified CallManager 5.x environment requires Cisco Unified MeetingPlace Express to point to Cisco Unified CallManager only.

For more information, refer to the *Administrator's Configuration and Maintenance Guide for Cisco MeetingPlace Express*, available at

<http://www.cisco.com>

H.323 and SIP Direct Integration

Integration with Cisco Unified CallManager can be accomplished by several methods. This section examines two of those methods, H.323 and SIP direct integration, while the third (H.323 via gatekeeper) is examined in the next section.

H.323 Gateway

The simplest integration option is to define Cisco Unified MeetingPlace Express as an H.323 gateway in Cisco Unified CallManager. You must also configure route patterns in Cisco Unified CallManager that point to the defined gateway. Cisco Unified CallManager 4.1, 4.2, and 5.0, as well as Cisco Unified CallManager Express, all support this option.

SIP Trunk

Using a SIP trunk to connect to Cisco Unified MeetingPlace Express has some conditions and version dependencies. The following current combinations are supported between Cisco Unified CallManager and Cisco Unified MeetingPlace Express.

Cisco Unified MeetingPlace Express 1.1.1 and Cisco Unified CallManager 5.x

With this combination, a SIP trunk can be defined directly from Cisco Unified CallManager to Cisco Unified MeetingPlace Express. The following design rules apply to this integration:

- A separate SIP Security Profile must be created and associated with the SIP trunk to Cisco Unified MeetingPlace Express. This security profile must have the transport set to UDP. The default setting is TCP, which will not work with Cisco Unified MeetingPlace Express.
- The SIP trunk configuration must have **MTP Required** checked and must have MTP resources available. MTPs will be invoked for each call to Cisco Unified MeetingPlace Express for the entire duration of the call. This requirement does not apply to Cisco Unified MeetingPlace Express 1.1.2 and higher.

Cisco Unified MeetingPlace Express 1.1.2 and Cisco Unified CallManager 5.x

With this combination, a SIP trunk can be defined directly from Cisco Unified CallManager to Cisco Unified MeetingPlace Express. The following design rules apply to this integration:

- Cisco Unified CallManager 5.0.4 or higher is required for Cisco Unified MeetingPlace Express 1.1.2. Prior Cisco Unified CallManager 5.0 versions are not supported.
- A separate SIP Security Profile must be created and associated with the SIP trunk to Cisco Unified MeetingPlace Express. This security profile must have the transport set to UDP. The default setting is TCP, which will not work with Cisco Unified MeetingPlace Express.
- The SIP trunk does not require MTPs because Cisco Unified MeetingPlace Express 1.1.2 supports enhanced SIP capabilities. Do not check the MTP Required option on the SIP trunk.
- Ensure that MTPs are available in the associated media resource group list. Some endpoints may cause an MTP to be invoked dynamically when completing a call to Cisco Unified MeetingPlace Express.

Cisco Unified MeetingPlace Express 1.1.x and Cisco Unified CallManager 4.1 and 4.2

With this combination, a SIP trunk can be defined directly from Cisco Unified CallManager to Cisco Unified MeetingPlace Express. The following design rules apply to this integration:

- Outgoing transport must be set to UDP. The default setting is TCP, which will not work with Cisco Unified MeetingPlace Express.
- The SIP trunk configuration must have **MTP Required** checked and must have MTP resources available. MTPs will be invoked for each call to Cisco Unified MeetingPlace Express for the entire duration of the call. This requirement will not be removed in future releases of Cisco Unified MeetingPlace Express due to the requirements of Cisco Unified CallManager 4.x.

Gatekeeper Integration

Cisco Unified MeetingPlace Express can be integrated into a gatekeeper environment in accordance with the design considerations presented in this section. Integration as an H.323 gateway or via SIP trunk directly from Cisco Unified CallManager is still an option in a gatekeeper environment. The gatekeeper registration method is as a gateway, using a single E.164 address with the registration. The following design considerations must be taken into account when designing for use in a gatekeeper environment.

Design Considerations



Note

Cisco Unified MeetingPlace Express will not register into a specific zone. The default, or first, local zone is always used when multiple local zones exist.

To force Cisco Unified MeetingPlace Express to use a specific local zone, use the **no zone** command on all zones where you do not want Cisco Unified MeetingPlace Express to register. For example, the following commands force Cisco Unified MeetingPlace Express (with address 10.20.110.50) to register to testzone2:

```
gatekeeper
zone local mp2-gk1 mp2.com 10.20.105.50
zone local testzone2 mp2.com
no zone subnet mp2-gk1 10.20.110.50/32 enable
```



Note

The Cisco Unified MeetingPlace Express direct meeting dial-in feature requires additional gatekeeper configuration.

Cisco Unified MeetingPlace Express registers as a single E.164 address, so extensions for direct meeting dial-in cannot be reached through a gatekeeper without adding manual entries to the gatekeeper by one of the methods shown in the following examples.

Option 1: Use Zone Prefix to Route (Recommended)

This example routes extensions 1000 through 1009 to the Cisco Unified MeetingPlace Express server.

```
gatekeeper
zone local mp2-gk1 mp2.com 10.20.105.50
zone prefix mp2-gk1 100.
gw-type-prefix 1#* default-technology gw ipaddr 10.20.110.50 1720
```

Option 2: Use Static E.164 Addresses (Not Recommended)

```
gatekeeper
alias static 10.20.110.50 1720 gkid mp2-gk1 gateway voip ras 10.20.110.50 62675 e164
1005 e164 1008 e164 1007 e164 1006
```

```
show gatekeeper endpoints
```

CallSignalAddr	Port	RASSignalAddr	Port	Zone Name	Type	Flags
-----	----	-----	----	-----	----	-----
10.20.110.50	1720	10.20.110.50	62675	mp2-gk1	UNKN-GW	S
E164-ID: 1000						
H323-ID: MeetingPlace						

```
E164-ID: 1005 (static)
E164-ID: 1008 (static)
E164-ID: 1007 (static)
E164-ID: 1006 (static)
```

Redundancy

This section describes the following types of redundancy:

- [Cisco Unified MeetingPlace Express Server Redundancy, page 16-12](#)
- [Physical Network Redundancy, page 16-12](#)
- [Redundancy Using a Gatekeeper, page 16-12](#)
- [Redundancy Using H.323 Gateway Integration, page 16-13](#)
- [Redundancy Using SIP Trunk Integration, page 16-13](#)

Cisco Unified MeetingPlace Express Server Redundancy

Cisco Unified MeetingPlace Express is a standalone server with no redundancy built in, except for some server components which are themselves redundant. Mirrored drives, dual power supplies, redundant fans, and so forth, give a high level of internal server redundancy and availability. Other methods outside of the server may be incorporated to enhance redundancy characteristics for calls sent from or delivered to Cisco Unified MeetingPlace Express.

Although Cisco Unified MeetingPlace Express is a standalone server, multiple Cisco Unified MeetingPlace Express servers can be deployed using the same Cisco Unified CallManager directory. This configuration enables users to log into an alternate server, via the common directory, and quickly reschedule meetings if the server they were originally using should fail. Two or more fully licensed Cisco Unified MeetingPlace Express servers are need for this option.

Physical Network Redundancy

Cisco Unified MeetingPlace Express has two Ethernet ports that are utilized for different purposes, one for call control, web interface, and media, and the other for web collaboration. While traffic rerouting should occur, these ports cannot be relied upon for redundancy. If one port on the server or switch (or the cable between them) fails, the server might not operate properly.

Redundancy Using a Gatekeeper

Inbound Calls

The only other redundancy with respect to inbound calls is gatekeeper redundancy, which would apply to outbound calls as well. You can implement gatekeeper redundancy by either of the following methods:

- Alternate gatekeeper

In this method, gatekeepers are set up in a redundant gatekeeper cluster. Cisco Unified MeetingPlace Express registers with its defined gatekeeper. Upon registration, the gatekeeper informs Cisco Unified MeetingPlace Express of an alternate gatekeeper in the cluster for use in the event of a failure.

- Gatekeeper Hot Standby Router Protocol (HSRP)

In this method, two gatekeepers share a single HSRP IP address. In the case of failure, Cisco Unified MeetingPlace Express registers to the secondary gatekeeper using the same gatekeeper IP address.

**Note**

Because the alternate gatekeeper information is delivered during registration and is not permanently stored in Cisco Unified MeetingPlace Express, booting or rebooting of the server when the primary gatekeeper is inaccessible will prevent Cisco Unified MeetingPlace Express from registering with the alternate gatekeeper.

Detailed information about gatekeeper redundancy can be found in the section on [Gatekeeper Redundancy, page 8-18](#).

Outbound Calls

Multiple Cisco Unified CallManager servers in a single cluster can register to a gatekeeper for redundancy. The gatekeeper will choose an active Cisco Unified CallManager to complete an outdial from Cisco Unified MeetingPlace Express.

Gatekeeper redundancy as described in the previous section ([Inbound Calls, page 16-12](#)) applies to outbound calls as well.

Redundancy Using H.323 Gateway Integration

Inbound Calls

If you integrate Cisco Unified MeetingPlace Express as an H.323 gateway, inbound calls will be sent from an available Cisco Unified CallManager server within a cluster. If the server fails, inbound calls will automatically be sent from another server in the cluster.

Outbound Calls

Cisco Unified MeetingPlace Express allows definition of multiple outbound Cisco Unified CallManagers for outbound call resolution. Should one fail, the next one is used.

Redundancy Using SIP Trunk Integration

Inbound Calls

If you integrate Cisco Unified MeetingPlace Express using a SIP trunk, inbound calls will be sent from an available Cisco Unified CallManager server within a cluster. If the server fails, inbound calls will automatically be sent from another server in the cluster.

Outbound Calls

Cisco Unified MeetingPlace Express allows definition of multiple outbound Cisco Unified CallManagers for outbound call resolution via SIP trunk. Should one fail, the next one is used.

Other Important Design Considerations

This section lists other important design considerations not covered previously in this chapter.

Network Connectivity

- Cisco Unified MeetingPlace Express uses two network interface cards (NICs) for connectivity. At this time, both NICs must be placed in the same subnet with a common default gateway. Connectivity issues will arise if the NICs are placed in separate subnets.
- Changing the IP addresses of the Cisco Unified MeetingPlace Express server must be done through the **net** command by connecting to the server via SSH and using the **mpxadmin** login. Changing the IP address through other means will not properly change the configuration of Cisco Unified MeetingPlace Express and will cause configuration issues.

IP Telephony Gateways

- Gateways used with Cisco Unified MeetingPlace Express must be configured with the standard recommendations set forth in the chapter on [Gateways, page 4-1](#).
- The main gateway issue that may impact Cisco Unified MeetingPlace Express is echo cancellation. Cisco recommends increasing the tail allocation on the gateway to 128 ms if echo issues are encountered.



CHAPTER 17


IP Video Telephony

Last revised on: February 13, 2008

Cisco introduced its IP Video Telephony solution in Cisco Unified CallManager Release 4.0. Video is fully integrated into Cisco Unified CallManager, and there are also many video endpoints available from Cisco and its strategic partners. Cisco Unified Video Advantage is just as easy to deploy, manage, and use as a Cisco Unified IP Phone.

IP Video Telephony Solution Components

The Cisco IP Video Telephony solution consists of Cisco Unified CallManager 4.2, Cisco Unified Videoconferencing 3500 Series Multipoint Control Units (MCUs) for both H.323 and Skinny Client Control Protocol (SCCP) conference calls, Cisco Unified Videoconferencing 3500 Series H.320 Gateways, Cisco IOS H.323 Gatekeeper, Cisco Unified Video Advantage, Cisco IP Video Phone 7985, Sony and Tandberg SCCP endpoint solutions, and the existing range of H.323 products from partners such as Polycom, Tandberg, Sony, and others. (See [Figure 17-1](#).)

Figure 17-1 IP Video Telephony Components

119450

Administration Considerations

This section discusses the following configuration elements in Cisco Unified CallManager Administration that pertain to Video Telephony:

- [Protocols, page 17-2](#)
- [Regions, page 17-3](#)
- [Locations, page 17-6](#)
- [Retry Video Call as Audio, page 17-7](#)
- [Wait for Far-End to Send TCS, page 17-9](#)

Protocols

Cisco Unified CallManager supports a large number of protocols. Any device can call any other device, but video is supported only on SCCP and H.323 devices. Specifically, video is not supported in the following protocols in Cisco Unified CallManager Release 4.2:

- Computer Telephony Integration (CTI) applications (TAPI and JTAPI)
- Media Gateway Control Protocol (MGCP)
- Session Initiation Protocol (SIP)

Therefore, Cisco Unified CallManager currently supports the types of calls listed in [Table 17-1](#).

Table 17-1 *Types of Calls Supported in Cisco Unified CallManager Release 4.2*

Calling Device Type	Called Device Type				
	SCCP	H.323	MGCP	TAPI/JTAPI	SIP
SCCP	Audio and video	Audio and video	Audio only	Audio only	Audio only
H.323	Audio and video	Audio and video	Audio only	Audio only	Audio only
MGCP	Audio only	Audio only	Audio only	Audio only	Audio only
TAPI/JTAPI	Audio only	Audio only	Audio only	Audio only	Audio only
SIP	Audio only	Audio only	Audio only	Audio only	Audio only

Table 17-2 lists the audio and video algorithms and protocols currently supported in Cisco Unified CallManager.

Table 17-2 *Capabilities Supported in Cisco Unified CallManager Release 4.2*

H.323	SCCP
H.261	H.261
H.263+	H.263+
	H.264
G.711 A-law and mu-law	G.711 A-law and mu-law
G.723.1	G.723.1
G.728	G.728
G.729, G.729a, G.729b, and G.729ab	G.729, G.729a, G.729b, and G.729ab
G.722	G.722
Far-end camera control (Supported by Cisco CallManager but not by all endpoints)	Far-end camera control (Supported by Cisco CallManager but not by all endpoints)
Out-of-band DTMF (H.245 alphanumeric)	Out-of-band DTMF
	Cisco Wideband Audio
	Cisco VT Camera Wideband Video Codec

Regions

When configuring a region, you set two fields in Cisco Unified CallManager Administration: the Audio Codec and the Video Bandwidth. The audio setting specifies a codec type, while the video setting specifies the amount of bandwidth you want to allow. However, even though the notation is different, the Audio Codec and Video Bandwidth fields actually perform similar functions. The Audio Codec field defines the maximum bit-rate allowed for audio-only calls as well as for the audio channel in video calls. For instance, if you set the Audio Codec for a region to G.711, Cisco Unified CallManager allocates 64 kbps as the maximum bandwidth allowed for the audio channel for that region. In this case, Cisco Unified CallManager will permit calls using either G.711, G.722, G.728, or G.729. However, if

you set the Audio Codec to G.729, Cisco Unified CallManager allocates only 8 kbps as the maximum amount of bandwidth allowed for the audio channel, and it will permit calls using only G.729 because G.728, G.711, and G.722 all take more than 8 kbps.

**Note**

If both endpoints support G.711 and G.722, then G.722 will be negotiated by default.

**Note**

The Audio Codec setting also applies to the audio channel of video calls.

The Video Bandwidth field defines the maximum bit-rate allowed for the video channel of the call. However, for historical continuity with the practices used in traditional videoconferencing products, the value used in this field also includes the bandwidth of the audio channel. For instance, if you want to allow calls at 384 kbps using G.711 audio, you would set the Video Bandwidth field to 384 kbps and not 320 kbps.

In summary, the Audio Codec field defines the maximum bit-rate used for audio-only calls and for the audio channel of video calls, while the Video Bandwidth field defines the maximum bit-rate allowed for video calls and should include the audio portion of the call.

Choosing the correct audio codec bandwidth limit is very important because each device supports only certain audio codecs, as shown in [Table 17-3](#). (For the most recent list of codecs supported by a particular endpoint, refer to the product documentation for that endpoint.)

Table 17-3 *Types of Audio Codecs Supported by Endpoint Devices*


Codec Type	Cisco 7900 Series IP Phones	SCCP Sony and Tandberg Endpoints	Typical H.323 Endpoints	IP/VC 3500 Series Gateways	IP/VC 3500 Series MCUs
G.729	Yes	Yes, depending on the model	No	No	Yes (with transcoder)
G.728	No	Yes, depending on the model	Yes	Yes (with transcoder)	Yes (with transcoder)
G.711	Yes	Yes	Yes	Yes	Yes
G.722	No	Yes	Yes	Yes (with transcoder)	Yes (with transcoder)
Cisco Wideband Audio	Yes	No	No	No	No

As [Table 17-3](#) indicates, if you set the region to G.729, not all videoconferencing devices are able to support this type of codec. For example, calls between a Cisco Unified Video Advantage endpoint and a Tandberg endpoint would fail, or Cisco Unified CallManager would allocate an audio transcoding resource for the call.

Transcoders do not currently support the pass-through capability, so the call would connect as audio-only and would be transcoded between G.729 and G.711. To avoid this situation without using Cisco IOS Enhanced Transcoders, you would have to set the region to use G.711 instead. However, a region set for G.711 would also use G.711 for audio calls between two IP phones, which you might not want due to the increased consumption of bandwidth over the WAN.

If you want to use G.729 for audio-only calls to conserve bandwidth and to use G.711 for video calls, then you should configure one region to use G.711 for video endpoints that do not support G.729 and a separate region (or regions) to use G.729 for IP phones. (See [Figure 17-2](#).) This method increases the number of regions needed but provides the desired codec and bandwidth allocations.

Figure 17-2 Using G.711 for Video Calls and G.729 for Audio-Only Calls



Note

It is possible to configure a pair of regions to prohibit video. If two video-capable devices in that region pair try to call each other, they will connect as audio-only unless Retry Video Call as Audio is not checked, in which case AAR rerouting logic will take over.

[Table 17-4](#) lists some example configurations and their outcomes.

Table 17-4 Scenarios for Various Region Settings

Region Setting	Setting of Retry Video as Audio	Result
Region allows video	Enabled	Video calls allowed
Region allows video	Disabled	Video calls allowed
Region does not allow video	Enabled	Video calls will proceed as audio
Region does not allow video	Disabled	If AAR is not configured, video calls fail (with busy tone and "Bandwidth Unavailable" message displayed)

The Video Bandwidth field accepts values in the range of 1 to 8128 kbps. However, to allow for compatibility with H.323 and H.320 videoconferencing devices, Cisco recommends that you always enter values for this field in increments of either 56 or 64 kbps. Therefore, valid values for this field include 112 kbps, 128 kbps, 224 kbps, 256 kbps, 336 kbps, 384 kbps, and so forth.

When the call speed requested by the endpoint exceeds the bandwidth value configured for the region, Cisco Unified CallManager automatically negotiates the call down to match the value allowed in the region setting. For instance, assume that an H.323 endpoint calls another H.323 endpoint at 768 kbps, but the region is set to allow a maximum of 384 kbps. The incoming H.225 setup request from the calling party would indicate that the call speed is 768 kbps, but Cisco Unified CallManager would change that value to 384 kbps in the outgoing H.225 setup message to the called party. Thus, the called endpoint

would think that it was a 384-kbps call to begin with, and the call would be negotiated at that rate. The calling endpoint would show the requested bandwidth as 768 kbps, but the negotiated bandwidth would be 384 kbps.

However, if you set the Video Bandwidth to "None" in the region, Cisco Unified CallManager will either terminate the call (and send an H.225 Release Complete message back to the calling party) or will allow the call to pass as an audio-only call instead, depending on whether or not the called device has the Retry Video Call as Audio option enabled. (See [Retry Video Call as Audio](#), page 17-7.)

Locations

When configuring static locations, you also set two fields in Cisco Unified CallManager Administration: the Audio Bandwidth and the Video Bandwidth. Unlike regions however, the Audio Bandwidth for static locations applies only to audio-only calls, while the Video Bandwidth applies to both the audio and video channels of video calls. The audio and video bandwidth are kept separate because, if both types of calls shared a single allocation of bandwidth, then it is very likely that audio calls would take all the available bandwidth and leave no room for any video calls, or vice versa. Also, separate bandwidth pools for audio and video correspond to the way queues are configured in the switches and routers in the network, which typically have a priority queue for voice traffic and a separate priority queue or a Class-Based Weighted Fair Queue for video traffic. (See [WAN Quality of Service \(QoS\)](#), page 3-28, for more details.)

Both the Audio Bandwidth and the Video Bandwidth fields offer three options: None, Unlimited, or a field that accepts numeric values. However, the values entered in these fields use two different calculation models. For the Audio Bandwidth field, the values entered should include the Layers 3 to 7 overhead required for the call. For instance, if you want to permit a single G.729 call to or from a location, you would enter the value of 24 kbps. For a G.711 call, you would enter the value of 80 kbps. The Video Bandwidth field, by contrast, should be entered without the overhead included. For instance, for a 128-kbps call you would enter 128 kbps, and for a 384-kbps call you would enter 384 kbps. As with the values used in the Video Bandwidth field for regions, Cisco recommends that you always use increments of 56 kbps or 64 kbps for the Video Bandwidth field for locations as well.

For example, assume that a company has a three-site network. The San Francisco location has a 1.544-Mbps T1 circuit connecting it to the San Jose main campus. The system administrator wants to allow four G.729 voice calls and one 384-kbps (or two 128-kbps) video calls to/from that location. The Dallas location has two 1.544-Mbps T1 circuits connecting it to the San Jose main campus, and the administrator wants to allow eight G.711 voice calls and two 384-kbps video calls to/from that location. For this example, the administrator would set the San Francisco and Dallas locations to the following values:

Location	Number of Audio Calls Desired	Audio Bandwidth Field Value	Number of Video Calls Desired	Video Bandwidth Field Value
San Francisco	4 using G.729	96 kbps (4 * 24 kbps)	1 at 384 kbps	384 kbps
Dallas	8 using G.711	640 kbps (8 * 80 kbps)	2 at 384 kbps	768 kbps

When the call speed requested by the endpoint exceeds the value configured for the location, Cisco Unified CallManager will *not* automatically negotiate the call speed (as it does for regions) to match the value allowed in the location setting. Instead, the call will be rejected or will be retried as an audio-only call (if the Retry Video as Audio setting is enabled on the called device). Therefore, you should always set the region's video bandwidth to a value lower than the location's video bandwidth value. For example, if you have two regions (region A and region B) and you set the video bandwidth

between those two regions to 768 kbps but the devices in region A are in a location that has a video bandwidth set to 384 kbps, then all calls between those two regions will fail or will result in audio-only calls (depending on the Retry Video Call as Audio setting).

Retry Video Call as Audio

This check-box is available on all endpoint types that support video, including Cisco Unified IP Phones 7940, 7941, 7960, 7961, 7970, 7971, and Cisco IP Video Phone 7985, as well as Sony or Tandberg SCCP endpoints, and all H.323 devices (clients, gateways and all types of H.323 trunks). When this option is activated (checked), if there is not enough bandwidth to reach the device (for example, if the Cisco Unified CallManager regions or locations do not allow video for that call), then Cisco Unified CallManager will retry the call as an audio-only call. When this option is deactivated (unchecked), Cisco Unified CallManager will not retry the call as audio-only but instead will either fail the call or reroute the call by whatever automated alternate routing (AAR) path is configured. By default, this retry option is enabled (checked).


This feature applies to the following scenarios only:

- The region is configured not to allow video.
- The location is configured not to allow video, or the requested video speed exceeds the available video bandwidth for that location.
- For calls between Cisco Unified CallManager clusters, the requested video speed exceeds the gatekeeper's zone bandwidth limits.

The Retry Video Call as Audio option takes effect only on the terminating (called) device, thus allowing the flexibility for the calling device to have different options (retry or AAR) for different destinations.


If the video call fails due to bandwidth limitations but automated alternate routing (AAR) is enabled, Cisco Unified CallManager will attempt to reroute the failed call as a video call to the AAR destination. If AAR is not enabled, the failed call will result in a busy tone and an error message being sent to the caller. (See [Figure 17-3](#).) Depending on the type of device that is calling, the failed call will result in one of the following conditions:


- If the calling device is an SCCP endpoint with an LCD screen (such as a Sony or Tandberg SCCP endpoint or many models of Cisco Unified IP Phones), the caller will hear busy tone and will see the message "Bandwidth Unavailable" displayed on the device.
- If the calling device is an SCCP endpoint without an LCD screen (such as a Cisco Unified IP Phone 7902), the caller will hear busy tone.
- If the calling device is an H.323 device or a PSTN device connected through a gateway, the caller will hear busy tone and Cisco Unified CallManager will send the appropriate error message (such as a Q.931 Network Congestion cause code) back to the H.323 or MGCP device.

Figure 17-3 Possible Scenarios for a Video Call

See the chapter on [Call Admission Control, page 9-1](#), for further details on the use of AAR.

[Figure 17-4](#) illustrates the steps of a call between two clusters using non-gatekeeper controlled intercluster trunks.

Figure 17-4 Call Flow Between Two Clusters Using Non-Gatekeeper Controlled Intercluster Trunks



119464

Wait for Far-End to Send TCS

This check-box is available on all H.323 devices, including H.323 clients, H.323 gateways, and H.225 gatekeeper-controlled trunks. This feature pertains to the H.245 capabilities-exchange phase of H.323 calls. When this feature is enabled, Cisco Unified CallManager waits for the remote H.323 device to send its Terminal Capabilities Set (TCS) to Cisco Unified CallManager before Cisco Unified CallManager will send its TCS to the H.323 device. When the option is disabled, Cisco Unified CallManager does not wait but sends its TCS to the remote H.323 device immediately.

By default, the Wait for Far-End to Send TCS option is enabled (checked). However, you must uncheck (disable) it in the following circumstances:

- When the H.323 device communicating with Cisco Unified CallManager is also waiting for the far-end to send its TCS

In this case, a stalemate occurs because neither side sends its TCS, and the H.245 connection times out after a few seconds. Examples of devices that also wait for the far-end to send TCS are some H.323 routed-mode gatekeepers, H.320 gateways, H.323 proxies (or IP-to-IP gateways) and some H.323 multipoint conference bridges. These devices wait for the far-end to send TCS for the same reason Cisco Unified CallManager does: because they are waiting for both sides of the connection to send their TCSs before forwarding them to the other side.

- When communicating with another Cisco Unified CallManager cluster over an intercluster trunk

**Note**

For intercluster trunks and gatekeeper-controlled intercluster trunks, the Wait for Far-End to Send TCS option is always disabled and cannot be enabled.

In many scenarios, Cisco Unified CallManager performs the role of a software switch connecting two endpoint devices (such as two H.323 clients that are trying to call each other). In such cases, it is best if Cisco Unified CallManager waits until both devices have sent their TCS messages so that it knows the capabilities of each device and can therefore make the most intelligent decision about what TCS to send back to each party (depending on region and location configurations, among other things). In these cases, the Wait for Far-End to Send TCS feature should be enabled.


However, some other H.323 devices (such as an H.320 gateway, which connects an H.323 device to an H.320 device) also perform the function of connecting two or more parties together. The gateway also prefers to wait until both ends send their TCS messages, so that it can make the most intelligent choice about how to set up the call. A stalemate could result if Cisco Unified CallManager and the gateway both wait for the other side to send their TCSs. To avoid this stalemate situation, disable (uncheck) the Wait for Far-End to Send TCS feature.

For instance, consider the following call scenarios illustrated in [Figure 17-5](#):

- Scenario 1 — Cisco Unified Video Advantage calls an H.320 endpoint
- Scenario 2 — An H.323 client calls an H.320 endpoint

In both of these scenarios, Wait for Far-End to Send TCS feature can be left at its default setting of enabled (checked)

Figure 17-5 Scenarios with the Wait for Far-End to Send TCS Feature Enabled (Checked)




In Scenario 1 from [Figure 17-5](#), Cisco Unified CallManager already knows the capabilities of the Cisco Unified Video Advantage client because SCCP devices provide Cisco Unified CallManager with their media capabilities during registration. But Cisco Unified CallManager does not know the capabilities of the H.320 gateway until the gateway sends its TCS to Cisco Unified CallManager during the H.245 phase of the call. Likewise, the H.320 gateway does not know what TCS to send to Cisco Unified CallManager until the H.320 endpoint sends its TCS to the gateway. In this case it is better to leave the Wait for Far-End to Send TCS feature enabled because the H.320 endpoint will send its TCS to the gateway, the gateway will send its TCS to Cisco Unified CallManager, and Cisco Unified CallManager will then have the TCSs from both endpoints with which to make a decision.

Figure 17-6 shows the following call scenarios, in which the call will fail unless the Wait for Far-End to Send TCS feature is disabled:

- Scenario 1 — Cisco Unified Video Advantage calls another Cisco Unified Video Advantage in a remote cluster via the ISDN network
- Scenario 2 — An H.323 client calls another H.323 client in the remote cluster via the ISDN network

Figure 17-6 Scenarios with the Wait for Far-End to Send TCS Feature Disabled (Unchecked)



In both scenarios from Figure 17-6, there will be a stalemate because both Cisco Unified CallManagers will wait to receive the TCSs from the gateways, and the gateways will both wait to receive the TCS from the ISDN side as well. The call will time-out after a few seconds and fail. From the perspective of the users, the caller will hear ringback tone indicating that the call is progressing and the called party will hear a ring indicating an incoming call. When the called party attempts the answer the call, the H.245 phase will fail due to the stalemate, and the call will then fail by hanging up on both parties.

To work around this issue in scenarios such as these, Cisco recommends that you disable (uncheck) the Wait for Far-End to Send TCS option on the device representing the H.320 gateway in Cisco Unified CallManager. This device could be an H.225 gatekeeper-controlled trunk or an H.323 gateway device, depending on how you have configured Cisco Unified CallManager to reach the H.320 gateway.

However, if the Wait for Far-End to Send TCS option is disabled, there is a possibility that the initial capabilities exchanged will not work for the remote device. For instance, the Cisco Unified CallManager region might be set to 768-kbps video but the H.320 device might only support 384 kbps, or the selected audio codec might not work for the remote party. In such cases, the initially negotiated logical channels might have to be torn down and reopened with the correct speed and codec. Many legacy H.323 and H.320 devices do not handle this situation well and will disconnect the call when Cisco Unified CallManager sends a CloseLogicalChannel message to renegotiate the channel to a different value. Thus, you must be careful where and when you disable the Wait for Far-End to Send TCS option.

Multipoint Conferencing

Whenever three or more parties want to engage in the same video call together, a Multipoint Control Unit (MCU) is required. An MCU consists of the following main components:

- Multipoint Controller (MC)

- Multipoint Processor (MP)

The MC handles all aspects of call setup and teardown for the conference, including media negotiation, call signaling, and choosing which MP to use for the call. The MP processes all the audio and video packets. The MC controls the MP, and one MC can control multiple MPs. The MP can be either software-based or hardware-based. Software-based MPs are typically not capable of advanced transcoding, transrating (multiple speeds), or composition features.

Since 1999, Cisco has offered the IP/VC 3500 Series H.323 Multipoint Conference Units (MCUs). The first product in this family was the Cisco Unified Videoconferencing 3510. This model is no longer available for sale, and is not compatible with Cisco Unified CallManager. In 2002, Cisco introduced the IP/VC 3511 and 3540 models. These models offer significantly improved features and scalability unavailable on the older 3510 model.

In 2003, Cisco introduced software version 3.2+ on the IP/VC 3511 and 3540 models, which adds support for the Skinny Client Control Protocol (SCCP). SCCP support is not available on the IP/VC 3510. In addition, there are three types of MPs available in the IP/VC MCUs:

- A software-based MP built into each MCU
- The Rate Matching (RM) module, a dedicated software-based module for the IP/VC 3540 chassis
- The Enhanced Media Processor (EMP), a hardware-based solution available as a dedicated module for the IP/VC 3540 chassis or as an integrated component in the IP/VC 3511 model

Cisco Unified CallManager Release 4.2 supports the IP/VC 3511 and 3540 models in SCCP and H.323 modes. Each protocol offers different features and is used for different reasons, so each of these MCUs is equipped to run both protocols. The IP/VC 3511 can be configured to run in either SCCP mode or H.323 mode, while the IP/VC 3540 model can be configured to run both protocols simultaneously and divide the total number of available MP resources between the two.

Regardless of signaling protocol, the MCU provides the same basic function of receiving the audio and video streams from each participant and sending those streams back out to all other participants in some sort of combined view. There are two types of views in a multipoint video conference:

- Voice-activated (switched)
- Continuous presence

Voice Activation

Voice-activated conferences take in the audio and video streams of all the participants, decide which participant is the dominant speaker, and send only the dominant speaker's video stream back out to all other participants. The participants then see a full-screen image of the dominant speaker (and the current speaker sees the previous dominant speaker). The audio streams from all participants are mixed together, so everyone hears everyone else, but only the dominant speaker's video is displayed.

You can use any of the following methods to select the dominant speaker:

- Voice activation mode

Using this mode, the MCU automatically selects the dominant speaker by determining which conference participant is speaking the loudest and the longest. To determine loudness, the MCU calculates the strength of the voice signal for each participant. As conditions change throughout the conversation, the MCU automatically selects a new dominant speaker and switches the video to display that participant. A hold timer prevents the video from switching too hastily. To become the dominant speaker, a participant has to speak for a specified number of seconds and be more dominant than all other participants.

- Manual selection of the dominant speaker through the MCU's web-based conference control user interface
The conference controller (or chairperson) can log onto the MCU's web page, highlight a participant, and select that person as the dominant speaker. This action disables voice activity detection, and the dominant speaker remains constant until the chairperson either selects a new dominant speaker or re-enables voice activation mode.
- Configuring the MCU to cycle through the participant list automatically, one participant at a time
With this method, the MCU stays on each participant for a configured period of time and then switches to the next participant in the list. The conference controller (or chairperson) can turn this feature on and off (re-enable voice activation mode) via the web interface.

Continuous Presence

Continuous-presence conferences display some or all of the participants together in a composite view. The view can display from 2 to 16 squares (participants) in a variety of different layouts. Each layout offers the ability to make one of the squares voice-activated, which is useful if there are more participants in the conference than there are squares to display them all in the composite view. For instance, if you are using a four-way view but there are five participants in the call, only four of them will be displayed at any given time. You can make one of the squares in this case voice-activated so that participants 4 and 5 will switch in and out of that square, depending on who is the dominant speaker. The participants displayed in the other three squares would be fixed, and all of the squares can be manipulated via the conference control web-based user interface.



Note

Continuous presence requires the use of an Enhanced Media Processor (EMP) in the IP/VC MCU.

MP Resources

For both types of conferences, the type of MP resource determines which video formats, transrating, and transcoding capabilities the MCU can support. If endpoints connect to the conference at different speeds, a transrating-capable MP is required. The RM and EMP modules are both capable of transrating between speeds. If a transrating-capable MP is not available, the MCU will send out flow-control messages to all the endpoints, instructing them to lower their transmit speed to match the maximum receive rate of the slowest endpoint. For example, if three participants are connected in a 384-kbps conference and a fourth participant joins at 128 kbps, the MCU will send flow-control messages to the other three participants instructing them to lower their transmit rates to match the 128-kbps participant. This method causes all participants to suffer degraded quality because one participant is less capable. A transrating-capable MP would, instead, convert the 128-kbps stream to 384 kbps, and vice versa, so that each participant can enjoy the best quality their connection allows.

For continuous-presence conferences, a transrating-capable MP is also very important. The software-based MP built into the MCU combines all the input streams and sends the resulting combination back out to each participant. For instance, if four participants are connected in a continuous-presence conference at 384 kbps using G.711 audio, each participant will transmit 320 kbps of video and 64 kbps of audio into the MCU. The MCU will take these four input video streams and combine them into the four-way composite view. The MCU will then transmit 1280 kbps of video back to each endpoint, along with the mixed 64 kbps of audio, for a total of 1344 kbps per endpoint. This method is known as Asynchronous Continuous Presence, and it can have a negative impact on bandwidth requirements, call admission control mechanisms, and interoperability with certain devices.



Note

Cisco strongly advises against the use of Asynchronous Continuous Presence.

With the RM or EMP modules, the MCU is capable of transrating each input stream before combining them, so that the total output bandwidth matches the input bandwidth. For instance, if the MCU is using the four-quadrant layout and each participant transmits 320 kbps of video and 64 kbps of audio into the MCU, the MCU would essentially transrate each input stream down to 80 kbps, combine them so that the resulting four-quadrant view is 320 kbps of video ($4 * 80$ kbps), combine this video with the mixed 64 kbps audio, and transmit the final combination back to each participant. This method is known as Synchronous Continuous Presence. Cisco strongly recommends that all continuous-presence conferences use the Synchronous Continuous Presence mode. However, use of this mode means that each MCU must have a transrating-capable MP (such as an RM or EMP) available, which does increase the cost of the MCU.

**Note**

For an H.323 client with a built-in MCU, Cisco Unified CallManager does not allow the H.323 client to generate a second call, thereby negating the functionality of the built-in MCU.

SCCP MCU Resources

As stated previously, the Cisco Unified Videoconferencing 3511 and 3540 MCUs both offer support for SCCP beginning with software version 3.2+ for those models and with Cisco Unified CallManager Release 4.0. When configured in SCCP mode, Cisco Unified CallManager provides the MC function while the MCU provides the MP function. The SCCP MCU is controlled completely by Cisco Unified CallManager.

Only the following events invoke SCCP MCU resources:

- The user of an SCCP endpoint (such as an IP Phone or a Tandberg endpoint) presses the Conf, Join, or cBarge softkeys to invoke an ad-hoc conference
- The user of an SCCP endpoint (such as an IP Phone or a Tandberg endpoint) presses the MeetMe softkey to invoke a reservationless meet-me conference.


Participants in either of these types of conferences can include any type of endpoint (that is, video and non-video devices using any signaling protocol that Cisco Unified CallManager supports via any supported gateway type); however, only SCCP endpoints can invoke the SCCP MCU resources. In other words, an H.323 video endpoint cannot invoke an SCCP MCU resource, but an SCCP video endpoint can invoke the resource and then join an H.323 video participant to the call. For example, the user at the SCCP endpoint could press the Conf softkey, dial the directory number of an H.323 client, and then press the Conf softkey again to complete the transaction. The H.323 client will be joined as a participant on the SCCP MCU conference.

However, for ad-hoc conferences initiated via the Conf, Join, or cBarge softkey, the signaling protocol used by the other participants must support the ability to be placed on hold and then have their audio and video channels redirected to the MCU. For H.323 devices (H.323 clients, H.323 gateways, H.320 gateways, and all types of H.323 trunks) Cisco Unified CallManager uses the Empty Capabilities Set (ECS) method defined in the H.245 specification to achieve this functionality. If the H.323 endpoint does not support receiving an ECS message from Cisco Unified CallManager, it will react by hanging up or possibly even crashing and/or rebooting. To work around this problem, you can enable (check) the "MTP Required" option on the H.323 device but assign it a media resource group list (MRGL) that does not contain any MTP devices, and then set the Cisco CallManager Service Parameter "Fail Call if MTP Allocation Fails" to False. (For details, see [Media Resource Groups and Lists, page 17-15](#).) This configuration will cause the softkeys to be greyed out on the phone, disabling the user from invoking any supplementary services with that endpoint, including placing it on hold, conferencing it into an existing call, joining it with an existing call, or barging into an existing call containing that endpoint.

For reservationless conferences via the MeetMe softkey, the signaling protocol used by the other endpoints does not have to support being placed on hold and transferred. For these types of conferences, each endpoint dials the MeetMe dial-in number arranged by the SCCP client who initiated the conference.

Figure 17-7 illustrates how H.323 and SCCP endpoints can participate in the same SCCP conference. In this example, the conference was initiated by an SCCP endpoint using the Conf softkey to invite the three members.

Figure 17-7 Ad-Hoc Conference Between SCCP and H.323 Endpoints



SCCP conferences support voice-activated mode as well as continuous presence. Furthermore, SCCP conferences support the integrated software-based MP of the MCU and the Rate Matching (RM) module as well as the Enhanced Media Processor (EMP) module.

Media Resource Groups and Lists

Cisco Unified CallManager uses media resource groups (MRGs) and media resource group lists (MRGLs) to determine which specific conference bridge resource to use for a given endpoint. How you group the resources is completely at your discretion, but it is typically done either by geographical placement (so that all endpoints at a given site use the MCU closest to them) or by endpoint type (so that video-capable endpoints use a video-capable MCU while audio-only endpoints use a different conference bridge resource). When a user of an SCCP device activates the Conf, Join, or MeetMe softkey, Cisco Unified CallManager uses the MRGL of the initiating endpoint to determine which conference bridge should be used.

Cisco Unified CallManager applies the following criteria, in the order listed, to select which conference bridge resource to use:

1. The priority order in which the media resource groups (MRGs) are listed in the media resource group list (MRGL)
2. Within the selected MRG, the resource that has been used the least

Because of this selection process, the video-capable MCU must be at the top of the MRGL for the video-capable SCCP endpoint in order for that MCU to be selected when the user activates the Conf, Join, or MeetMe softkey on that endpoint. However, some endpoints are not dedicated video endpoints. For example, the IP Phone used in conjunction with Cisco Unified Video Advantage might be used for audio-only calls most of the time and for video calls only occasionally. Thus, if you place the MCU at the top of the MRGL for that phone, the MCU will always be chosen even for audio-only conferences

that do not involve any video-capable participants. In this scenario, the MCU resources might be wasted on audio-only conferences and not be available to satisfy the request for a video conference when it occurs.

Therefore, Cisco recommends that you carefully select which users you want to give access to the video-capable MCU resources through MRGLs. The following considerations can help you make that selection:

- Is the endpoint a dedicated video device (such as a Sony or Tandberg SCCP endpoint), or is it an IP Phone that will only occasionally have a Cisco Unified Video Advantage client associated to it?
- Does the user require video when invoking SCCP conferences, or will audio-only suffice?
- How much are you willing to spend on MCU resources to outfit your network with enough resources to satisfy the SCCP conferencing requirements?

The answers to these selection criteria will be different for every company. One company might deem multipoint video to be a critical feature and be willing to spend the money necessary to outfit their network with enough MCU resources so that, even if a few ports get wasted on audio-only conferences, there will still be resources available to satisfy all of their video conferences. Another company might choose to enable video resources only for their Sony and Tandberg endpoints and assign audio-only conference bridge resources to Cisco Unified Video Advantage users. A third company might choose to enable video resources for all Sony and Tandberg SCCP endpoints and for select Cisco Unified Video Advantage users (perhaps managers and above), but allocate audio-only resources for the rest of the user population.

H.323 MCU Resources

When configured in H.323 mode, the MCU provides the MC function and behaves like an H.323 peer to Cisco Unified CallManager. H.323 MCU conferences can be invoked in a number of different ways, but they all fall into two major categories:


- Scheduled
- Reservationless

A scheduled conference uses a scheduling application to reserve the MCU resources in advance of the call. The scheduling function typically is provided through a web-based user interface such as Cisco Unified MeetingPlace, MagicSoft VCWizard, RADVision VCS, Tandberg TMS, or FVC.COM Click to Meet. The scheduling application usually generates an invitation that provides the user with the date and time of the conference, the number of ports reserved for the conference, and the dial-in information. Alternatively, the scheduling system can be configured to dial out to some or all of the participants at the beginning of the conference.

For a reservationless conference, the MCU has a certain number of resources that are available on demand. To create a conference, a user simply dials into the MCU at any time. If that user is the first participant to dial in, the MCU dynamically creates a new conference using the settings defined in the service template. (For more information on service templates, see [Service Templates and Prefixes](#), [page 17-17](#).) Subsequent users who dial into the same conference number are joined to that conference.

Any type of endpoint can create and participate in scheduled and reservationless H.323 conferences. For instance, an SCCP endpoint can dial into the H.323 MCU to create a reservationless conference just as well as an H.323 endpoint can.

[Figure 17-8](#) illustrates how H.323 and SCCP endpoints can participate in the same H.323 conference. In this example, the conference was initiated by an SCCP endpoint that dialed into the H.323 MCU to create a new reservationless conference, and the other two parties subsequently dialed into that conference.

Figure 17-8 *SCCP and H.323 Endpoints in a Reservationless Conference*

H.323 conferences support both voice-activated and continuous-presence modes. Furthermore, H.323 conferences support all of the MP types, including the integrated software-based MP of the MCU, the Rate Matching (RM) module, and the Enhanced Media Processor (EMP) module.

Service Templates and Prefixes

A service on the MCU defines the settings that pertain to each conference. You can have different services defined for the different types of conferences. Each service defines, at a minimum, the following settings:

- Speed of the conference (the video bit-rate)
This setting might include multiple speeds if a transrating-capable MP is used.
- Minimum and/or maximum number of participants
The minimum number defines how many ports will be reserved at the beginning of the conference. The maximum number defines the maximum number of participants that the MCU will allow to join this conference.
- Video codec type (H.261, H.263, or H.264)
- Frame rate (15 or 30 frames per second)
- Resolution (QCIF or CIF)
- MP resource (Auto, MP, RM, or EMP)
- Video layout to be displayed (voice activated or continuous presence)
A conference can include multiple layouts or even dynamic layouts that change as the number of participants in the conference increases or decreases.
- H.323 or SCCP
When the "SCCP service" check-box is enabled (checked), the service is used for SCCP conferences. When this box is disabled (unchecked), the service is used for H.323 conferences.

For H.323 services, each service is assigned a service prefix that is dialed by the endpoints to reach that particular service. The service prefix forms the leading digits of the conference number, and the trailing digits define the conference ID. This format allows multiple conferences to run simultaneously using the same service prefix. For instance, you can have a service prefix of 555, while the full dial string of the conference is seven digits. This scheme would allow for conference numbers in the range of 5550000 through 5559999, with four digits for the conference ID. The user must dial the full string to access the conference. When the call is received by the MCU, the MCU parses the dialed digits to try to match a

service prefix. Once it determines which service prefix is being dialed, the MCU then uses the remaining digits as the conference ID. If the conference ID does not exist yet, the MCU creates a new reservationless conference with that ID. If the conference already exists, the user is added to that conference.

SCCP services must also have a service prefix defined, but the digits themselves do not mean anything because users do not "dial" into an SCCP service. The prefix is used only in the SCCP registration messages between Cisco Unified CallManager and the SCCP MCU resource. Users either use the Conf, Join, or cBarge softkeys to access the SCCP MCU conference (ad hoc), or they dial a MeetMe number assigned by Cisco Unified CallManager to join the conference (reservationless). Therefore, the digits you specify for the SCCP service prefix do not matter and can be any digits you wish, such as 999999 for instance. This prefix is not exposed outside of the SCCP signaling between the MCU and Cisco Unified CallManager (that is, it cannot be dialed, it is not included in the MCU's registration with the gatekeeper, and so forth).

Sizing the MCU

As mentioned previously, the current MCU models are the IP/VC 3511 and 3540. The IP/VC 3511 MCU supports a fixed number of ports, while the IP/VC 3540 MCU is a modular system that can accept various sizes of modules. When calculating the total number of ports available, you must also consider the number of sessions that the Audio Transcoder Daughter Card and the Rate Matching (RM) module or Enhanced Media Processor (EMP) module can support. Therefore, consider the following factors when calculating the size of the MCU:

- The number of ports that the MCU can support
This value depends on the speed of the conference; as the speed increases, the number of supported ports decreases.
- The number of ports that the Audio Transcoding Daughter Card can support
This value depends on which audio codecs are used in the conference.
- The number of conferences that the RM or EMP module can support
This value depends on how many participants need to be transrated and how many views are in use in the conference.

For specific information about the number of ports supported, refer to the product documentation for your MCU hardware, available on Cisco.com. Due to the almost infinite number of possible variations, it is very difficult to provide any concrete design guidance in this document. Many customers end up with a mixture of SCCP ad-hoc conferences, H.323 reservationless conferences, and H.323 scheduled conferences. The MCUs must be sized to accommodate all of those types of conferences at the correct speeds and video layouts. Needless to say, this can become quite complex to determine. Please consult with your Cisco sales representative for assistance on sizing the MCUs for your particular environment.

IVR for Dial-In Conference

Dial-in conferences typically use an interactive voice response (IVR) system to prompt users to enter the conference ID and the password (if one is configured) of the conference they want to join. You can use either of the following types of IVRs with the Cisco Unified Videoconferencing 3500 Series MCUs:

- The IVR built into the MCU
- Cisco Unified IP-IVR

The built-in IVR of the MCU has the following characteristics:

- Can prompt only for the password of the conference
It cannot prompt the user for the conference ID first. In other words, the users must dial the specific conference number they wish to join, then they are prompted for the password of that conference.
- Supports both in-band and out-of-band (H.245 alpha-numeric) DTMF
- Cannot be customized to provide more flexible menus or functionality

The only thing that can be customized is the recorded audio file that is played to the user.

If you want to have a single dial-in number and then prompt the user for the conference ID, you can use Cisco Unified IP-IVR in conjunction with the MCU.

Cisco Unified IP-IVR has the following characteristics:

- Can prompt for the conference ID and the password (among other things)
- Supports only out-of-band DTMF
That is, the calling device must support an out-of-band DTMF method, such as H.245 alpha-numeric on H.323 devices. These out-of-band DTMF messages are then relayed by Cisco Unified CallManager to the Cisco IP IVR server. If the calling device supports only in-band DTMF tones, the Cisco IP IVR server will not recognize them and the calling device will be unable to enter the conference.
- Can be highly customized to provide more flexible menus and other advanced functionality
Customizations can include such things as verifying the user's account against a back-end database before permitting that user to enter into the conference, or queuing the participants until the chairperson joins.

**Note**

Because Cisco Unified IP-IVR supports only out-of-band signaling, it will not work with H.323 endpoints that use in-band DTMF tones.

With Cisco Unified IP-IVR, users dial a CTI route point that routes the call to the Cisco Unified IP-IVR server instead of dialing a route pattern that routes directly to the MCU. After collecting the DTMF digits of the conference ID, the Cisco Unified IP-IVR then transfers the call to the route pattern that routes the call to the MCU. This transfer operation requires that the calling device supports having its media channels closed and reopened to a new destination. For example, an H.323 video device that calls the Cisco Unified IP-IVR will initially negotiate an audio channel to the Cisco Unified IP-IVR server and then, after entering the appropriate DTMF digits, it will be transferred to the MCU, at which point Cisco Unified CallManager will invoke the Empty Capabilities Set (ECS) procedure described earlier in this chapter to close the audio channel between the endpoint and the Cisco Unified IP-IVR server and open new logical channels between the endpoint and the MCU. If the H.323 video endpoint does not support receiving an ECS from Cisco Unified CallManager, it will react by disconnecting the call or, worse, crashing and/or rebooting.


Gatekeepers

Prior to the introduction of video support in Cisco Unified CallManager, H.323 videoconferencing networks relied on gatekeepers to perform device registration management, call routing, and bandwidth control. The Cisco IOS Gatekeeper, formerly known as the Multimedia Conference Manager (MCM), provides these functions. However, most gatekeepers, including Cisco's, offer only rudimentary call

routing capabilities compared to what might be expected in a typical enterprise-class PBX. When used to route H.323 video calls, Cisco Unified CallManager supplements the basic gatekeeper features and provides full enterprise-class PBX functionality to H.323 video calls.

Cisco Unified CallManager and the gatekeeper work as a team to manage H.323 video endpoints. The gatekeeper handles all Registration, Admission, and Status (RAS) signaling, while Cisco Unified CallManager handles all of the H.225 call signaling and H.245 media negotiations. Therefore, you have to deploy gatekeepers along with the Cisco Unified CallManager servers if RAS signaling procedures are required for the H.323 endpoints in your network, as illustrated in Figure 17-9.

Figure 17-9 Cisco Unified CallManager and IOS Gatekeeper Provide RAS Signaling for H.323 Endpoints



RAS signaling is required any time either of the following conditions exists:

- The endpoint does not use a fixed IP address.

If the endpoint uses a static IP address, Cisco Unified CallManager does not require RAS procedures to locate the endpoint. Instead, the endpoint is provisioned in Cisco Unified CallManager Administration with its static IP address, and calls to that H.323 client's directory number are routed directly to that static IP address. If the endpoint does not use a static IP address, then Cisco Unified CallManager must query the gatekeeper to obtain the endpoint's current IP address each time Cisco Unified CallManager extends a call to the endpoint.

- The endpoint requires RAS procedures to place calls to E.164 addresses.

Most H.323 videoconferencing endpoints are capable of dialing another endpoint directly only when dialing by IP address (that is, the user enters the IP address of the destination endpoint in dotted-decimal format and then pushes the call button). However, if the user dials an E.164-formatted number (a numeric value not in the dotted-decimal format of an IP address) or an H.323-ID (in the format of *username* or *username@domain*), most endpoints today provide only one way to resolve these types of destinations – by a RAS query to their gatekeeper. A growing number of endpoints, however, can be configured so that, for any call to an E.164 address, they skip any RAS procedures and instead send an H.225 SETUP message directly to a specified IP address. This method of operation is known as peer-to-peer mode. Tandberg H.323 endpoints are one example that use this mode, in which you can either configure a gatekeeper address for them to register with, or configure the IP address of the Cisco Unified CallManager server they should use. In the latter case, the endpoint sends all calls directly to the specified IP address, bypassing the need for RAS procedures with any gatekeeper.

In addition to managing RAS procedures for H.323 video endpoints, gatekeepers also continue to play an important role in managing dial plan resolution and bandwidth restrictions between Cisco Unified CallManager clusters in large multisite distributed call processing environments. A gatekeeper can also integrate with large numbers of H.323 VoIP gateways within the organization, or it can act as a session border controller between an enterprise IP Telephony network and a service provider VoIP transport network.

Therefore, as it pertains to Cisco IP Video Telephony deployments, the Cisco IOS Gatekeeper can perform one or both of the following roles:

- Endpoint gatekeeper

An endpoint gatekeeper is configured to manage all RAS procedures for calls to, from, and between H.323 clients, MCUs, and H.320 video gateways. The endpoint gatekeeper directs all such calls to the appropriate Cisco Unified CallManager cluster so that Cisco Unified CallManager can perform all of the H.225 call routing and H.245 media negotiations.

- Infrastructure gatekeeper

An infrastructure gatekeeper is configured to manage all dial plan resolution and bandwidth restrictions (call admission control) between Cisco Unified CallManager clusters, between a Cisco Unified CallManager cluster and a network of H.323 VoIP gateways, or between a Cisco Unified CallManager cluster and a service provider's H.323 VoIP transport network.

In Cisco Unified CallManager Release 4.0, the endpoint gatekeeper and the infrastructure gatekeeper had to run on separate routers, and each endpoint gatekeeper could service only a single Cisco Unified CallManager cluster. If multiple Cisco Unified CallManager clusters existed within the enterprise, a separate endpoint gatekeeper had to be deployed for each Cisco Unified CallManager cluster. With Cisco Unified CallManager Release 4.1 or above, it is possible to combine these roles on a single gatekeeper, using it as an endpoint gatekeeper for one or more Cisco Unified CallManager clusters and as the infrastructure gatekeeper for managing calls between clusters or between a cluster and other H.323 VoIP networks. However, for the following reasons (among others), Cisco recommends that you still separate these roles onto two or more gatekeepers:

- Scalability

Depending on the Cisco IOS router platform you choose to deploy and your estimated busy hour call volume, you might need several gatekeepers to handle the load.

- Geographical resiliency

Putting all of your eggs into one basket may not be wise in a large, multi-national VoIP network. Having gatekeepers placed throughout your network (typically by geography) can provide better fault isolation in the event of a gatekeeper failure.

- Incompatibilities


Some configuration aspects of the gatekeeper are global in nature (they pertain to all endpoints registered with that gatekeeper). For example, the command **arq reject-unknown-prefix**, which may be useful in some H.323 VoIP transport environments, conflicts with the use of the **gw-type-prefix <prefix> default-technology** command, which is used in endpoint gatekeepers to route calls to Cisco Unified CallManager. While Cisco IOS does not stop you from configuring both commands on the same gatekeeper, the **arq reject-unknown-prefix** command takes precedence and, therefore, calls to unknown numbers will be rejected instead of being routed to Cisco Unified CallManager. In this case, you would have to use one gatekeeper for the H.323 VoIP transport network and another gatekeeper for the Cisco Unified CallManager cluster(s).

Another example of incompatibility can occur in the way you configure the gatekeeper for redundancy. Most Cisco H.323 voice devices, including Cisco Voice Gateways and Cisco Unified CallManager, support the H.323v3 Alternate Gatekeeper feature, which would allow you to configure the gatekeepers as a gatekeeper cluster using the Gatekeeper Update Protocol

(GUP) to keep in sync with each other. However, many H.323 video endpoints do not support Alternate Gatekeeper, so the gatekeepers must be configured to use Hot Standby Routing Protocol (HSRP) for redundancy. You cannot mix and match these two redundancy methods on the same gatekeeper. In this case, you might decide to use a gatekeeper cluster for those endpoints that support Alternate Gatekeeper and an HSRP pair of gatekeepers for those that do not.

Figure 17-10 illustrates a network scenario with two Cisco Unified CallManager clusters. Each cluster consists of SCCP and H.323 clients, H.323 MCUs, and H.320 gateways. To manage the RAS aspects of the H.323 clients, MCUs, and H.320 gateways, an endpoint gatekeeper is deployed with each cluster. A separate infrastructure gatekeeper manages dial plan resolution and bandwidth between the clusters. Gatekeeper redundancy is not shown in the figure, although each of these gatekeepers may actually be multiple gatekeepers configured for either Alternate Gatekeeper or HSRP-based redundancy.

Figure 17-10 Two Cisco Unified CallManager Clusters with Required Gatekeepers



Supported Gatekeeper Platforms

To act as an endpoint gatekeeper with Cisco Unified CallManager 4.1 and above, the Cisco IOS Gatekeeper must run Cisco IOS Release 12.3(11)T or greater. For minimum Cisco IOS release requirements on the infrastructure gatekeeper, see [Recommended Hardware and Software Combinations, page A-1](#).

The following router platforms support the Cisco IOS Gatekeeper:

- Cisco 2600XM Series and 2691
- Cisco 2800 Series
- Cisco 3640, 3640A, 3660
- Cisco 3725 and 3745
- Cisco 3825 and 3845
- Cisco 7200 Series, 7301, and 7400 Series

To determine which release and feature set you should use for your router platform, use the *Cisco Feature Navigator* (requires a Cisco.com login account), available at

<http://tools.cisco.com/ITDIT/CFN/jsp/index.jsp>

More information is also available at

<http://www.cisco.com/web/psa/products/index.html?c=268438303>

That document states that Cisco IOS Release 12.3(11)T provides integrated voice and video services, therefore it is the recommended minimum release.

Endpoint Gatekeepers

An endpoint gatekeeper is required any time both of the following conditions are met:

- The cluster contains H.323 clients, H.323 MCUs, or H.320 gateways (collectively referred to as H.323 endpoints). If none of these types of endpoints exists (for example, if all clients are SCCP endpoints and there are no MCUs or H.320 gateways), then an endpoint gatekeeper is not needed.
- And either of the following conditions is true:
 - The H.323 endpoints require RAS procedures to initiate calls to E.164 addresses. As mentioned earlier, a growing number of devices are capable of peer-to-peer call signaling, in which case there is no need for those devices to register with a gatekeeper.
 - The H.323 endpoints do not use static IP addresses.

The role of the endpoint gatekeeper is simply to handle the RAS aspects of communications with the endpoints, providing a place for these H.323 endpoints to register. The endpoint gatekeeper responds to all call requests made to, from, or between these endpoints by directing the call to the appropriate Cisco Unified CallManager server(s) so that Cisco Unified CallManager can perform all of the call routing and bandwidth control functions. To accomplish this call routing and bandwidth control, you configure Cisco Unified CallManager to register H.323 trunk(s) with the gatekeeper and configure the gatekeeper to route calls to those trunks for all calls to, from, or within that zone.

Cisco Unified CallManager Release 4.1 introduced a new type of H.323 trunk called the RASAggregator trunk. This type of trunk is used for all H.323 client, H.323 MCU, or H.320 gateway zones, while the gatekeeper-controlled intercluster trunk and gatekeeper-controlled H.225 trunk from previous Cisco Unified CallManager releases are used to integrate with infrastructure gatekeepers. If you have deployed H.323 video endpoints based on the recommendations documented in the *IP Video Telephony SRND for Cisco Unified CallManager 4.0*, you should modify your configuration to use the new RASAggregator trunk so that you can take advantage of its flexibility. This configuration change should be carefully planned and done at a time when it is convenient for the administrator so as to minimize disruption to the existing H.323 video endpoints. For a description of the procedure for migrating from a Cisco Unified CallManager 4.0 deployment, see [Migrating from Cisco Unified CallManager 4.0, page 17-42](#).

Provisioning H.323 Clients

H.323 clients are provisioned much the same way as other phones are, in that you create a new phone (model type = H.323 Client), assign a directory number to it, and assign it a calling search space, device pool, and so forth. You configure the H.323 clients in Cisco Unified CallManager in one of the following ways. The method you use depends on whether or not the client uses a static IP address and whether or not the client requires RAS procedures to dial E.164 addresses.

- Gatekeeper controlled

This type of configuration is used for clients that do not have a static IP address assigned to them (they use a DHCP-assigned address) and that require RAS procedures to dial E.164 addresses. A RASAggregator trunk is used to communicate to and from these clients. (See [Figure 17-11](#) and [Figure 17-12](#).)

- Non-gatekeeper controlled, asynchronous

This type of configuration is used for clients that have a static IP address assigned to them but that require RAS procedures to dial E.164 addresses. While Cisco Unified CallManager can signal directly to them without the need of a gatekeeper to resolve their IP addresses, they are not able to signal directly to Cisco Unified CallManager but instead must query the gatekeeper to resolve the E.164 address they are trying to dial (thus, asynchronous communications). To support these types of clients, you must have at least one gatekeeper-controlled client defined in Cisco Unified CallManager for each zone on the gatekeeper, even if all the clients actually use static IP addresses. In this case, the non-gatekeeper controlled client may be a "dummy" client that does not actually exist. Its purpose is merely to create the RASAggregator trunk so that the gatekeeper will be able to route calls from the clients to Cisco Unified CallManager. (See [Figure 17-13](#) and [Figure 17-14](#).)

- Non-gatekeeper controlled, synchronous

This type of configuration is used for clients that have a static IP address and are also capable of peer-to-peer signaling (that is, they do not require RAS procedures to dial E.164 numbers). Cisco Unified CallManager signals directly to them, and they signal directly to Cisco Unified CallManager (thus, synchronous communications). No gatekeeper or RASAggregator trunk is needed for this type of client. (See [Figure 17-15](#) and [Figure 17-16](#).)

[Figure 17-11](#) through [Figure 17-16](#) illustrate the call signaling flows used in these three scenarios.



Figure 17-11 Call to Gatekeeper-Controlled Client from Cisco Unified CallManager**Figure 17-12** Call from Gatekeeper-Controlled Client to Cisco Unified CallManager

Figure 17-13 *Call to Non-Gatekeeper Controlled Client from Cisco Unified CallManager (Asynchronous)*




Figure 17-14 *Call from Non-Gatekeeper Controlled Client to Cisco Unified CallManager (Asynchronous)*




Figure 17-15 Call to Non-Gatekeeper Controlled Client from Cisco Unified CallManager (Synchronous)





Figure 17-16 Call from Non-Gatekeeper Controlled Client to Cisco Unified CallManager (Synchronous)



Gatekeeper-Controlled Clients

When you configure an H.323 client as gatekeeper-controlled, you may enter any alpha-numeric string (such as a descriptive name) in the Device Name field, check the **Gatekeeper-controlled** box, and fill in the following fields:

- Device Pool

The device pool you want the client to use. All H.323 clients (whether gatekeeper-controlled or non-gatekeeper controlled) that are registered in the same zone must use the same device pool. If you accidentally assign different device pools across the endpoints, Cisco Unified CallManager will register multiple RASAggregator trunks within the zone, and an inbound call might be rejected by Cisco Unified CallManager if the call is directed to the wrong RASAggregator trunk.

- Gatekeeper

A drop-down list of gatekeeper IP addresses. You must define the gatekeeper in Cisco Unified CallManager before configuring any gatekeeper-controlled H.323 clients.

- Technology Prefix

The technology prefix used by the RASAggregator trunk to register in the client zone on the gatekeeper. This technology prefix must match what is configured as the default technology prefix on the gatekeeper. All gatekeeper-controlled H.323 clients that are registered in the same zone must use the same technology prefix. If you accidentally assign different technology prefixes across the endpoints, Cisco Unified CallManager will register multiple RASAggregator trunks within the zone, and an inbound call might be rejected by Cisco Unified CallManager if the call is directed to the wrong RASAggregator trunk. Cisco recommends that you use **1#** for this prefix.

- Zone Name

The (case-sensitive) name of the client zone as configured in the gatekeeper. All gatekeeper-controlled H.323 clients that are registered in the same zone must use the same zone name. If you accidentally assign different zone names (remember, the field is case sensitive) across the endpoints, Cisco Unified CallManager will attempt to register multiple RASAggregator trunks with the gatekeeper (but the one with the incorrect zone name will fail to register), and an inbound call might be rejected by Cisco Unified CallManager if the call is directed to the wrong RASAggregator trunk.

Also, you must set the Cisco CallManager service parameter **Send Product ID and Version ID to True**. This parameter allows the RASAggregator trunk to register with the gatekeeper as an H323-GW, so that the gatekeeper can direct all H.323 calls to, from, or within the client zone to the RASAggregator trunk.

Non-Gatekeeper Controlled Clients

When provisioning an H.323 client as non-gatekeeper controlled, you must enter the static IP address of the client into the Device Name field and leave all of the settings under the Gatekeeper-controlled section blank (unchecked). Cisco Unified CallManager then uses the static IP address to reach the client any time a call is extended to its directory number.

If the client is configured to use peer-to-peer mode, then no further configuration is required. If the client requires RAS procedures to place calls to E.164 addresses, then you must also configure a dummy gatekeeper-controlled H.323 client in order to create the RASAggregator trunk, by filling in the following fields:

- Device Name

A descriptive name that identifies this client as a dummy client used for the purpose of creating the RASAggregator trunk for the client zone.

- Device Pool

The device pool you chose when configuring the non-gatekeeper controlled H.323 client(s). If the device pool assigned to the dummy client is different than that assigned to the real clients, inbound calls from the real clients might be rejected by Cisco Unified CallManager.

- Gatekeeper

A drop-down list of gatekeeper IP addresses. You must define the gatekeeper in Cisco Unified CallManager before configuring the dummy gatekeeper-controlled H.323 client.

- Technology Prefix

The technology prefix used by the RASAggregator trunk to register in the client zone on the gatekeeper. This technology prefix must match what is configured as the default technology prefix on the gatekeeper. Cisco recommends that you use **1#** for this prefix.

- Zone Name

The (case-sensitive) name of the client zone as configured in the gatekeeper.

Also, you must set the Cisco CallManager service parameter **Send Product ID and Version ID** to **True**. This parameter allows the RASAggregator trunk to register with the gatekeeper as an H323-GW, so that the gatekeeper can direct all H.323 calls to, from, or within the client zone to the RASAggregator trunk.

Provisioning H.323 MCUs

H.323 MCUs are provisioned in Cisco Unified CallManager as H.323 gateways, and then route patterns are configured to extend calls to these devices. When provisioning an H.323 gateway, you must enter the static IP address and TCP signaling port of the MCU into the Device Name field.

Cisco Unified CallManager then uses the static IP address and TCP port to reach the MCU any time a call matches the route pattern(s) associated with it.



Note

The Cisco Unified Videoconferencing 3500 Series MCUs do not listen on TCP port 1720 by default. (The IP/VC 3500 Series MCUs listen on port 2720 by default.) You must verify which TCP port they are listening on, and either change it to 1720 or provision the correct port in Cisco Unified CallManager.

If the MCU is configured to use peer-to-peer mode, then no further configuration is required. (Cisco Unified Videoconferencing MCUs do not currently support peer-to-peer mode, but some third-party MCUs do.) If the MCU requires RAS procedures to place calls to E.164 addresses, then you must also configure a dummy gatekeeper-controlled H.323 client in order to create the RASAggregator trunk, by filling in the following fields:

- Device Name

A descriptive name that identifies this client as a dummy client used for the purpose of creating the RASAggregator trunk for the MCU zone.

- Device Pool

The device pool you chose when configuring the H.323 gateway representing the MCU. If the device pool assigned to the dummy client is different than that assigned to the H.323 gateway device representing the MCU, inbound calls from the MCU might be rejected by Cisco Unified CallManager.

- Gatekeeper

A drop-down list of gatekeeper IP addresses. You must define the gatekeeper in Cisco Unified CallManager before configuring the dummy gatekeeper-controlled H.323 client.

- Technology Prefix

The technology prefix used by the RASAggregator trunk to register in the MCU zone on the gatekeeper. This technology prefix must match what is configured as the default technology prefix on the gatekeeper. Cisco recommends that you use **1#** for this prefix.

- Zone Name

The (case-sensitive) name of the MCU zone as configured in the gatekeeper.

Also, you must set the Cisco CallManager service parameter **Send Product ID and Version ID** to **True**. This parameter allows the RASAggregator trunk to register with the gatekeeper as an H323-GW, so that the gatekeeper can direct all H.323 calls to, from, or within the MCU zone to the RASAggregator trunk.

MCU Service Prefixes

H.323 MCUs can use either E.164 addresses or technology prefixes (also referred to as service prefixes in the MCU) as the dial-in number(s) to reach reservationless or scheduled H.323 conferences running on them. Cisco recommends that you configure the MCUs to use E.164 addresses by setting the MCU Mode to **MCU** instead of **Gateway** in the MCU administration screens. If the **MCU** setting is not available on the model of MCU you are using, then you must use the following special configuration to properly route calls placed from other H.323 endpoints to the MCU:

If the MCU is configured in **Gateway** mode or is another vendor's MCU that (for whatever reason) requires its conference IDs to register as technology prefixes instead of as E.164 addresses, then the service prefix(s) of the MCU must begin with a # character. For example, if the MCU's service prefix is 8005551212, then you must provision the service prefix on the MCU as #8005551212. Thus, when other H.323 endpoints dial 8005551212, the gatekeeper will not find a matching technology prefix registered and will instead route the call to the RASAggregator trunk that is registered with the default technology prefix in the zone of the endpoint that is placing the call.

Cisco Unified CallManager must then prepend the # character to the beginning of the called number before extending the call to the MCU. This character is prepended on the route pattern(s) associated with the H.323 gateway representing the MCU. Calls to the MCU from SCCP clients will therefore also have this # character prepended to the calling number.

If the MCU is configured in **MCU** mode or is another vendor's MCU that uses E.164 addresses for its conference IDs, then you do not have to prepend the # character. Also note that, if the MCU uses peer-to-peer mode and hence does not need to register its technology prefixes with any gatekeeper, then this situation does not apply and you do not have to prepend a # character.

Provisioning H.320 Gateways

As with H.323 MCUs, H.320 gateways are provisioned in Cisco Unified CallManager as H.323 gateways, and then route patterns are configured to extend calls to these devices. When provisioning an H.323 gateway, you must enter the static IP address and TCP signaling port of the H.320 gateway into the Device Name field. Cisco Unified CallManager then uses the static IP address and TCP port to reach the gateway any time a call matches the route pattern(s) associated with it.



Note

The Cisco Unified Videoconferencing 3500 Series Gateways do not listen on TCP port 1720 by default. (The IP/VC 3500 Series Gateways listen on port 1820 by default.) You must verify which TCP port they are listening on, and either change it to 1720 or provision the correct port in Cisco Unified CallManager.

If the gateway is configured to use peer-to-peer mode, then no further configuration is required. If the gateway requires RAS procedures to place calls to E.164 addresses, then you must also configure a dummy gatekeeper-controlled H.323 client in order to create the RASAggregator trunk, by filling in the following fields:

- **Device Name**
A descriptive name that identifies this client as a dummy client used for the purpose of creating the RASAggregator trunk for the gateway zone.
- **Device Pool**
The device pool you chose when configuring the H.323 gateway representing the H.320 gateway. If the device pool assigned to the dummy client is different than that assigned to the gateway, inbound calls from the gateway might be rejected by Cisco Unified CallManager.
- **Gatekeeper**
A drop-down list of gatekeeper IP addresses. You must define the gatekeeper in Cisco Unified CallManager before configuring the dummy gatekeeper-controlled H.323 client.
- **E.164**
This field requires an entry. Ensure that it is "non-dialable" within Cisco Unified CallManager.
- **Technology Prefix**
The technology prefix used by the RASAggregator trunk to register in the gateway zone on the gatekeeper. This technology prefix must match what is configured as the default technology prefix on the gatekeeper. Cisco recommends that you use **1#** for this prefix.
- **Zone Name**
The (case-sensitive) name of the gateway zone as configured in the gatekeeper.

Also, you must set the Cisco CallManager service parameter **Send Product ID and Version ID** to **True**. This parameter allows the RASAggregator trunk to register with the gatekeeper as an H323-GW, so that the gatekeeper can direct all H.323 calls to, from, or within the gateway zone to the RASAggregator trunk.

Gateway Service Prefixes

H.320 gateways use technology prefixes (also referred to as service prefixes in the gateway) as the prefix that users should dial to reach an ISDN destination. For calls to route correctly, you must configure the service prefix(s) of the gateway to begin with a # character. For example, if the gateway's service prefix that clients dial to reach an ISDN number is 9, then you must provision the service prefix on the gateway as #9. In this way, when H.323 clients dial 9 plus the PSTN number (such as 918005551212), the gatekeeper will not find a matching technology prefix registered and will instead route the call to the Cisco Unified CallManager trunk that is registered with the default technology prefix.

Cisco Unified CallManager must then prepend the # character to the beginning of the called number before extending the call to the gateway. Note that, if the gateway uses peer-to-peer mode and hence does not need to register its technology prefixes with any gatekeeper, then this situation does not apply and you do not have to prepend a # character.

Gatekeeper Zone Configuration

The preceding sections discuss how to provision the endpoints in Cisco Unified CallManager Administration. You must also configure the endpoint gatekeeper(s) with the appropriate zone definitions. You must configure a zone for each type of endpoint (client, MCU, or gateway) and, optionally, for each device pool associated with these endpoints in Cisco Unified CallManager.

Each zone is configured to route all calls placed to, from, or within the zone to the RASAggregator trunk registered in that zone. You configure the zones on the endpoint gatekeeper by using the following command syntax:

```
zone local <zone_name> <domain_name> <ip_address> invia <zone_name>
outvia <zone_name> enable-intrazone
```

The command argument **invia** applies to calls placed to the zone from any other zone, **outvia** applies to calls placed from the zone to any other zone, and **enable-intrazone** applies to calls placed within the zone. The following sections illustrate the use of these commands.

Client Zones

The number of client zones you have to configure within each endpoint gatekeeper depends on the following factors:

- The device pools to which the H.323 clients are associated

The device pool determines which Cisco Unified CallManager servers are primary, secondary, and tertiary servers for each H.323 client. If you assign all H.323 clients to the same device pool, then you need to define only a single client zone in the endpoint gatekeeper. In other words, for each device pool used by H.323 clients, you must configure a separate client zone in the gatekeeper.

- Whether the endpoint gatekeeper provides services for a single Cisco Unified CallManager cluster or multiple Cisco Unified CallManager clusters

Each client zone is configured to route calls to a particular RASAggregator trunk. Therefore, if one endpoint gatekeeper is used to service multiple Cisco Unified CallManager clusters, then you must define a separate client zone for each cluster that the gatekeeper services.

To illustrate, the following three examples show how client zones may be configured. [Example 17-1](#) shows a single client zone defined for a single Cisco Unified CallManager cluster in which all H.323 clients are associated with the same device pool. [Example 17-2](#) shows a single Cisco Unified CallManager cluster in which the H.323 clients are divided between two different device pools. [Example 17-3](#) shows two Cisco Unified CallManager clusters in which the H.323 clients are divided between two different device pools in each cluster.



Note

Some of the commands shown in the following examples are the default values applied in the Cisco IOS Gatekeeper and, therefore, would not have to be configured explicitly, nor would they appear in the running configuration. They are included here for thoroughness but are marked by a ! at the beginning of the command line.

Example 17-1 Client Zone for a Single Cisco Unified CallManager Cluster and Single Device Pool

```
gatekeeper
zone local clients domain.com invia clients outvia clients enable-intrazone
gw-type-prefix 1# default-technology
no use-proxy clients default inbound-to terminal
no use-proxy clients default outbound-from terminal
! no arg reject-unknown-prefix
endpoint ttl 60
no shutdown
```

Example 17-2 Client Zones for a Single Cisco Unified CallManager Cluster and Two Device Pools

```
gatekeeper
zone local dp1-clients domain.com invia dp1-clients outvia dp1-clients enable-intrazone
```



```

zone local dp2-clients domain.com invia dp2-clients outvia dp2-clients enable-intrazone
gw-type-prefix 1# default-technology
no use-proxy dp1-clients default inbound-to terminal
no use-proxy dp1-clients default outbound-from terminal
no use-proxy dp2-clients default inbound-to terminal
no use-proxy dp2-clients default outbound-from terminal
! no arg reject-unknown-prefix
endpoint ttl 60
no shutdown

```

Example 17-3 Client Zones for Two Cisco Unified CallManager Clusters with Two Device Pools per Cluster

```

gatekeeper
zone local clstr1-dp1-clients domain.com invia clstr1-dp1-clients outvia
clstr1-dp1-clients enable-intrazone
zone local clstr1-dp2-clients domain.com invia clstr1-dp2-clients outvia
clstr1-dp2-clients enable-intrazone
zone local clstr2-dp1-clients domain.com invia clstr2-dp1-clients outvia
clstr2-dp1-clients enable-intrazone
zone local clstr2-dp2-clients domain.com invia clstr2-dp2-clients outvia
clstr2-dp2-clients enable-intrazone
gw-type-prefix 1# default-technology
no use-proxy clstr1-dp1-clients default inbound-to terminal
no use-proxy clstr1-dp1-clients default outbound-from terminal
no use-proxy clstr1-dp2-clients default inbound-to terminal
no use-proxy clstr1-dp2-clients default outbound-from terminal
no use-proxy clstr2-dp1-clients default inbound-to terminal
no use-proxy clstr2-dp1-clients default outbound-from terminal
no use-proxy clstr2-dp2-clients default inbound-to terminal
no use-proxy clstr2-dp2-clients default outbound-from terminal
! no arg reject-unknown-prefix
endpoint ttl 60
no shutdown

```

Disabling The Use of Proxy

The Cisco IOS Gatekeeper, formerly known as the Cisco Multimedia Conference Manager (MCM), previously offered an H.323 proxy function that has been at End of Life (EOL) for some time and is not compatible with Cisco Unified CallManager, but the commands in the gatekeeper to use a proxy for all calls to and from terminals (clients) are still enabled by default. You must disable this function for each client zone by using the following command syntax:

```

gatekeeper
no use-proxy <zone_name> default [inbound-to | outbound-from] terminals

```

The Cisco MCM proxy was replaced by a new solution called the Cisco IOS Multiservice IP-to-IP Gateway and the associated via-zone-enabled Cisco IOS Gatekeeper. This document does not discuss the IP-to-IP Gateway, but Cisco Unified CallManager Release 4.1 and above leverages the via-zone and IP-to-IP gateway constructs by registering its RASAggregator trunks with the gatekeeper, effectively mimicking an IP-to-IP gateway so that the gatekeeper will route all invia, outvia, and enable-intrazone calls to the RASAggregator trunk as if it were an IP-to-IP gateway.

Client Zone Prefixes

For H.323 client zones, there is no need to configure zone prefixes or technology prefixes of any kind, except for the default technology prefix. Instead, the **invia**, **outvia**, **enable-intrazone**, and **gw-type-prefix <1#> default-technology** commands ensure that all calls placed are routed to the RASAggregator trunk associated with the zone in which the call originated.

MCU Zones

The number of MCU zones you have to configure within each endpoint gatekeeper depends on the following factors:

- The device pools to which the MCUs are associated

The device pool determines which Cisco Unified CallManager servers are primary, secondary, and tertiary servers for each MCU. If you assign all MCUs to the same device pool, then you need to define only a single MCU zone in the endpoint gatekeeper. In other words, for each device pool used by MCUs, you must configure a separate MCU zone in the gatekeeper.

- Whether the endpoint gatekeeper provides services for a single Cisco Unified CallManager cluster or multiple Cisco Unified CallManager clusters

Each MCU zone is configured to route calls to a particular RASAggregator trunk. Therefore, if one endpoint gatekeeper is used to service multiple Cisco Unified CallManager clusters, then you must define a separate MCU zone for each cluster that the gatekeeper services.

To illustrate, the following three examples show how MCU zones may be configured. [Example 17-4](#) shows a single MCU zone defined for a single Cisco Unified CallManager cluster in which all MCUs are associated with the same device pool. [Example 17-5](#) shows a single Cisco Unified CallManager cluster in which the MCUs are divided between two different device pools. [Example 17-6](#) shows two Cisco Unified CallManager clusters in which the MCUs are divided between two different device pools.



Note

Some of the commands shown in the following examples are the default values applied in the Cisco IOS Gatekeeper and, therefore, would not have to be configured explicitly, nor would they appear in the running configuration. They are included here for thoroughness but are marked by a ! at the beginning of the command line.

Example 17-4 MCU Zone for a Single Cisco Unified CallManager Cluster and Single Device Pool

```
gatekeeper
zone local MCUs domain.com invia MCUs outvia MCUs enable-intrazone
gw-type-prefix 1# default-technology
! no use-proxy MCUs default inbound-to [MCU | gateway]
! no use-proxy MCUs default outbound-from [MCU | gateway]
! no arq reject-unknown-prefix
endpoint ttl 60
no shutdown
```

Example 17-5 MCU Zones for a Single Cisco Unified CallManager Cluster and Two Device Pools

```
gatekeeper
zone local dp1-MCUs domain.com invia dp1-MCUs outvia dp1-MCUs enable-intrazone
zone local dp2-MCUs domain.com invia dp2-MCUs outvia dp2-MCUs enable-intrazone
gw-type-prefix 1# default-technology
! no use-proxy dp1-MCUs default inbound-to [MCU | gateway]
! no use-proxy dp1-MCUs default outbound-from [MCU | gateway]
! no use-proxy dp2-MCUs default inbound-to [MCU | gateway]
! no use-proxy dp2-MCUs default outbound-from [MCU | gateway]
! no arq reject-unknown-prefix
endpoint ttl 60
no shutdown
```

Example 17-6 MCU Zones for Two Cisco Unified CallManager Clusters with Two Device Pools per Cluster

```

gatekeeper
zone local clstr1-dp1-MCUs domain.com invia clstr1-dp1-MCUs outvia clstr1-dp1-MCUs
enable-intrazone
zone local clstr1-dp2-MCUs domain.com invia clstr1-dp2-MCUs outvia clstr1-dp2-MCUs
enable-intrazone
zone local clstr2-dp1-MCUs domain.com invia clstr2-dp1-MCUs outvia clstr2-dp1-MCUs
enable-intrazone
zone local clstr2-dp2-MCUs domain.com invia clstr2-dp2-MCUs outvia clstr2-dp2-MCUs
enable-intrazone
gw-type-prefix 1# default-technology
! no use-proxy clstr1-dp1-MCUs default inbound-to [MCU | gateway]
! no use-proxy clstr1-dp1-MCUs default outbound-from [MCU | gateway]
! no use-proxy clstr1-dp2-MCUs default inbound-to [MCU | gateway]
! no use-proxy clstr1-dp2-MCUs default outbound-from [MCU | gateway]
! no use-proxy clstr2-dp1-MCUs default inbound-to [MCU | gateway]
! no use-proxy clstr2-dp1-MCUs default outbound-from [MCU | gateway]
! no use-proxy clstr2-dp2-MCUs default inbound-to [MCU | gateway]
! no use-proxy clstr2-dp2-MCUs default outbound-from [MCU | gateway]
! no arq reject-unknown-prefix
endpoint ttl 60
no shutdown

```

Disabling The Use of Proxy

By default, the Cisco IOS Gatekeeper is set to not use a proxy for calls to and from MCUs or gateways. However, if you have enabled the use of proxy for those types of endpoints, you must disable it for each MCU zone by using the following command syntax:

```

gatekeeper
no use-proxy <zone_name> default [inbound-to | outbound-from] [MCU | gateway]

```

If your MCU is registering as an MCU, then use the **MCU** argument at the end of the **no use-proxy** command; if your MCU is registering as a gateway, then use the **gateway** argument instead.

MCU Zone Prefixes

For H.323 MCU zones, there is no need to configure zone prefixes or technology prefixes of any kind, except for the default technology prefix. Instead, the **invia**, **outvia**, **enable-intrazone**, and **gw-type-prefix <1#> default-technology** commands ensure that all calls placed are routed to the RASAggregator trunk associated with the zone in which the call originated.

If your MCUs are registering their service prefixes as technology prefixes instead of E.164 addresses, use the special configuration described previously for prepending a # character to the MCU's service prefixes (see [MCU Service Prefixes](#), page 17-30). Due to the way the Cisco IOS Gatekeeper selects a via-zone for calls to a technology prefix, when the endpoint dials the service prefix of the MCU, the call will fail if the gatekeeper finds a matching technology prefix registered. You must ensure that the client does not dial the # character, so that the gatekeeper will not find a matching technology prefix and will instead route the call to the RASAggregator trunk associated with the zone in which the call originated.

H.320 Gateway Zones

The number of H.320 gateway zones you have to configure within each endpoint gatekeeper depends on the following factors:

- The device pools to which the H.320 gateways are associated

The device pool determines which Cisco Unified CallManager servers are primary, secondary, and tertiary servers for each H.320 gateway. If you assign all gateways to the same device pool, then you need to define only a single gateway zone in the endpoint gatekeeper. In other words, for each device pool used by H.320 gateways, you must configure a separate gateway zone in the gatekeeper.

- Whether the endpoint gatekeeper provides services for a single Cisco Unified CallManager cluster or multiple Cisco Unified CallManager clusters

Each gateway zone is configured to route calls to a particular RASAggregator trunk. Therefore, if one endpoint gatekeeper is used to service multiple Cisco Unified CallManager clusters, then you must define a separate gateway zone for each cluster that the gatekeeper services.

To illustrate, the following three examples show how gateway zones may be configured. [Example 17-7](#) shows a single gateway zone defined for a single Cisco Unified CallManager cluster in which all H.320 gateways are associated with the same device pool. [Example 17-8](#) shows a single Cisco Unified CallManager cluster in which the gateways are divided between two different device pools. [Example 17-9](#) shows two Cisco Unified CallManager clusters in which the gateways are divided between two different device pools.



Note

Some of the commands shown in the following examples are the default values applied in the Cisco IOS Gatekeeper and, therefore, would not have to be configured explicitly, nor would they appear in the running configuration. They are included here for thoroughness but are marked by a ! at the beginning of the command line.

Example 17-7 Gateway Zone for a Single Cisco Unified CallManager Cluster and Single Device Pool

```
gatekeeper
zone local gateways domain.com invia gateways outvia gateways enable-intrazone
gw-type-prefix 1# default-technology
! no use-proxy gateways default inbound-to gateway
! no use-proxy gateways default outbound-from gateway
! no arq reject-unknown-prefix
endpoint ttl 60
no shutdown
```

Example 17-8 Gateway Zones for a Single Cisco Unified CallManager Cluster and Two Device Pools

```
gatekeeper
zone local dp1-gateways domain.com invia dp1-gateways outvia dp1-gateways enable-intrazone
zone local dp2-gateways domain.com invia dp2-gateways outvia dp2-gateways enable-intrazone
gw-type-prefix 1# default-technology
! no use-proxy dp1-gateways default inbound-to gateway
! no use-proxy dp1-gateways default outbound-from gateway
! no use-proxy dp2-gateways default inbound-to gateway
! no use-proxy dp2-gateways default outbound-from gateway
! no arq reject-unknown-prefix
endpoint ttl 60
no shutdown
```

Example 17-9 Gateway Zones for Two Cisco Unified CallManager Clusters with Two Device Pools per Cluster

```

gatekeeper
zone local clstr1-dp1-gateways domain.com invia clstr1-dp1-gateways outvia
clstr1-dp1-gateways enable-intrazone
zone local clstr1-dp2-gateways domain.com invia clstr1-dp2-gateways outvia
clstr1-dp2-gateways enable-intrazone
zone local clstr2-dp1-gateways domain.com invia clstr2-dp1-gateways outvia
clstr2-dp1-gateways enable-intrazone
zone local clstr2-dp2-gateways domain.com invia clstr2-dp2-gateways outvia
clstr2-dp2-gateways enable-intrazone
gw-type-prefix 1# default-technology
! no use-proxy clstr1-dp1-gateways default inbound-to gateway
! no use-proxy clstr1-dp1-gateways default outbound-from gateway
! no use-proxy clstr1-dp2-gateways default inbound-to gateway
! no use-proxy clstr1-dp2-gateways default outbound-from gateway
! no use-proxy clstr2-dp1-gateways default inbound-to gateway
! no use-proxy clstr2-dp1-gateways default outbound-from gateway
! no use-proxy clstr2-dp2-gateways default inbound-to gateway
! no use-proxy clstr2-dp2-gateways default outbound-from gateway
! no arq reject-unknown-prefix
endpoint ttl 60
no shutdown

```

Disabling The Use of Proxy

By default, the Cisco IOS Gatekeeper is set to not use a proxy for calls to and from gateways. However, if you have enabled the use of proxy for those types of endpoints, you must disable it for each H.320 gateway zone by using the following command syntax:

```

gatekeeper
no use-proxy <zone_name> default [inbound-to | outbound-from] gateway

```

Gateway Zone Prefixes

There is no need to configure zone prefixes of any kind for H.320 gateway zones. Instead, the **invia**, **outvia**, **enable-intrazone**, and **gw-type-prefix <1#> default-technology** commands ensure that all calls placed are routed to the RASAggregator trunk associated with the zone in which the call originated.

You must also use the special configuration described previously for prepending a # character to the gateway's service prefixes (see [Gateway Service Prefixes, page 17-31](#)). Due to the way the Cisco IOS Gatekeeper selects a via-zone for calls to a technology prefix, when the endpoint dials the service prefix of the gateway, the call will fail if the gatekeeper finds a matching technology prefix registered. You must ensure that the client does not dial the # character, so that the gatekeeper will not find a matching technology prefix and will instead route the call to the RASAggregator trunk associated with the zone in which the call originated.

Zone Subnets

As mentioned previously, the H.323 specification permits a single gatekeeper to manage multiple zones. However, the gatekeeper needs a way to decide which zone an endpoint should be placed in when it receives a Registration Request (RRQ) from that device. The RRQ message contains a Gatekeeper Identifier field that enables the endpoint to indicate the zone in which it would like to register. However, many H.323 video endpoints do not populate this field, and if the gatekeeper has multiple zones defined, it will not know which zone to place the endpoint into. Therefore, you must use of the **zone subnet** command to tell the gatekeeper which zone to associate with the endpoint. This command defines which

IP addresses or IP address ranges are permitted to register in each zone. The command syntax requires that you enter a network mask. Therefore, you can specify either a particular host address by entering a 32-bit (/32) network mask or a range of addresses by specifying a smaller network mask.

Because MCUs, H.320 gateways, and Cisco Unified CallManager servers typically use fixed IP addresses but H.323 clients can use DHCP addresses, Cisco recommends that you define **zone subnet** commands only for the MCU and gateway zones but leave the client zones open so that any IP address is permitted in them. Note that you must also permit the Cisco Unified CallManager servers to register in the MCU and gateway zones, as illustrated in [Example 17-10](#).

**Note**

Some of the commands shown in the following example are the default values applied in the Cisco IOS Gatekeeper and, therefore, would not have to be configured explicitly, nor would they appear in the running configuration. They are included here for thoroughness but are marked by a ! at the beginning of the command line.

Example 17-10 Defining Zone Subnets

```
gatekeeper
no zone subnet MCUs default enable
zone subnet MCUs [MCUs_IP_addr]/32 enable
zone subnet MCUs [RASAggregators_IP_addr]/32 enable
no zone subnet gateways default enable
zone subnet gateways [gateways_IP_addr]/32 enable
zone subnet gateways [RASAggregators_IP_addr]/32 enable
! zone subnet clients default enable
no zone subnet clients [MCUs_IP_addr]/32 enable
no zone subnet clients [gateways_IP_addr]/32 enable
```

The configuration in [Example 17-10](#) explicitly permits the MCU and the RASAggregator for the MCU zone to register in the MCU zone, and it explicitly permits the gateway and RASAggregator for the gateway zone to register in the gateway zone. It also explicitly denies the MCU and gateway from registering in the client zone, while implicitly permitting all other IP addresses (including the RASAggregator for the client zone) to register in the client zone.

Endpoint Time to Live

Endpoints send lightweight Registration Requests (RRQs) to their gatekeeper periodically to maintain their registration status. The frequency with which they send these RRQs is referred to as the Time to Live (TTL) value. The endpoint may specify the TTL it wishes to use in the body of its RRQs. The gatekeeper may then honor the endpoint's requested TTL value by echoing it in the Registration Confirm (RCF) response or, alternatively, may override the endpoint's request by specifying a different TTL value in the RCF.

If the TTL value is not specified in the RRQ, the gatekeeper should specify one in its RCF response. The endpoint should then honor the TTL specified by the gatekeeper. The Cisco IOS Gatekeeper honors all TTL values specified by the endpoints. However, many H.323 video endpoints do not specify a TTL value in their RRQs. In such cases, the Cisco IOS Gatekeeper defaults to specifying a TTL value of 1800 seconds (30 minutes). The Cisco IOS Gatekeeper will flush the endpoint's registration after three TTL intervals have passed without receiving any messages from the endpoint ($3 * 30 \text{ minutes} = 90 \text{ minutes}$).

A large TTL value can cause problems with H.323 clients that do not use static IP addresses. For example, with the default TTL value of 1800 seconds, if you disconnect the client from the network and move it to another location in which it receives a different DHCP address, it will fail to register with the gatekeeper (Registration Reject (RRJ) cause value "duplicate alias") until three TTL intervals have passed, and the gatekeeper will flush that endpoint's original registration.

Therefore, Cisco recommends that you consider reducing the TTL value to as low a number as possible without causing any negative effect on your network. The Cisco IOS Gatekeeper permits you to set the TTL value anywhere in the range of 60 seconds to 3600 seconds. In most cases, 60 seconds should work well. However, if your gatekeeper is already heavily utilized, adjusting the TTL from the default of 1800 seconds to 60 seconds might cause it to become overwhelmed.

Use the following command syntax to set the TTL value:

```
gatekeeper
endpoint ttl <seconds>
```

Summary of Endpoint Gatekeepers

This section summarizes some key points to remember about endpoint gatekeepers and provides some example configurations that combine techniques used in the previous examples.

- Configure a separate zone in the endpoint gatekeeper for each type of endpoint (clients, MCUs, and H.320 gateways). If the endpoints are associated with multiple device pools, configure multiple zones for each type of endpoint.
- Configure a RASAggregator trunk to register in each zone. This trunk is automatically created when you configure gatekeeper-controlled H.323 clients in Cisco Unified CallManager Administration. However, for non-gatekeeper controlled H.323 clients, H.323 MCUs, and H.320 gateways, you must configure a dummy gatekeeper-controlled H.323 client in order to create the RASAggregator trunk for that zone.
- Set the service parameter **Send Product ID and Version ID** to **True** in order for the RASAggregator trunk to register with the gatekeeper as an IP-to-IP gateway. This setting enables the RASAggregator to be selected by the gatekeeper for all calls placed to, from, or within each zone due to the use of the **invia**, **outvia**, **enable-intrazone**, and **gw-type-prefix <1#> default-technology** commands applied to each local zone definition.
- You do not have to associate any zone prefixes for any of the endpoint zones. No matter what the endpoint dials, the gatekeeper should not find a matching zone prefix or technology prefix but should instead route the call to the RASAggregator trunk associated with the zone from which the call originated. To avoid having the gatekeeper accidentally match the dialed number to the technology prefix of your MCUs or gateways, mask all MCU and gateway service prefixes with a # character, and then prepend the # character in the route pattern associated with that MCU or gateway.
- Configure zone subnets if any of the H.323 endpoints do not support the ability to specify the Gatekeeper Identifier (name of the zone) with which they wish to register.
- Disable the use of the old MCM Proxy for all zones.
- Set the endpoint registration Time to Live (TTL) to as low of a value as you can without creating undo stress on the gatekeeper. In extreme cases where the gatekeeper is serving hundreds of endpoint registrations, setting the TTL to 60 seconds might cause an unmanageable amount of RAS traffic. In smaller environments, setting it to 60 seconds should work well.

[Example 17-11](#) shows a configuration for an endpoint gatekeeper servicing a single Cisco Unified CallManager cluster in which a single device pool is used to service all H.323 video endpoint types.

**Note**

Some of the commands shown in the following examples are the default values applied in the Cisco IOS Gatekeeper and, therefore, would not have to be configured explicitly, nor would they appear in the running configuration. They are included here for thoroughness but are marked by a ! at the beginning of the command line.

Example 17-11 Endpoint Gatekeeper Configuration for a Single Cluster and a Single Device Pool

```
gatekeeper
zone local clients domain.com invia clients outvia clients enable-intrazone
zone local MCUs domain.com invia MCUs outvia MCUs enable-intrazone
zone local gateways domain.com invia gateways outvia gateways enable-intrazone
! zone subnet clients default enable
no zone subnet clients [MCUs_IP_addr]/32 enable
no zone subnet clients [gateways_IP_addr]/32 enable
no zone subnet MCUs default enable
zone subnet MCUs [MCUs_IP_addr]/32 enable
zone subnet MCUs [RASAggregators_IP_addr]/32 enable
no zone subnet gateways default enable
zone subnet gateways [gateways_IP_addr]/32 enable
zone subnet gateways [RASAggregators_IP_addr]/32 enable
no use-proxy clients inbound-to terminals
no use-proxy clients outbound-from terminals
! no use-proxy MCUs inbound-to [MCU | gateway]
! no use-proxy MCUs outbound-from [MCU | gateway]
! no use-proxy gateways inbound-to gateway
! no use-proxy gateways outbound-from gateway
gw-type-prefix 1# default-technology
! no arg reject-unknown-prefix
endpoint ttl 60
no shutdown
```

Example 17-12 shows a configuration for an endpoint gatekeeper servicing two Cisco Unified CallManager clusters. Each cluster has two different device pools for its H.323 video endpoints.

Example 17-12 Endpoint Gatekeeper Configuration for Two Clusters and a Two Device Pools

```
gatekeeper
zone local clstr1-dp1-clients domain.com invia clstr1-dp1-clients outvia
clstr1-dp1-clients enable-intrazone
zone local clstr1-dp1-MCUs domain.com invia clstr1-dp1-MCUs outvia clstr1-dp1-MCUs
enable-intrazone
zone local clstr1-dp1-gateways domain.com invia clstr1-dp1-gateways outvia
clstr1-dp1-gateways enable-intrazone
zone local clstr1-dp2-clients domain.com invia clstr1-dp2-clients outvia
clstr1-dp2-clients enable-intrazone
zone local clstr1-dp2-MCUs domain.com invia clstr1-dp2-MCUs outvia clstr1-dp2-MCUs
enable-intrazone
zone local clstr1-dp2-gateways domain.com invia clstr1-dp2-gateways outvia
clstr1-dp2-gateways enable-intrazone
zone local clstr2-dp1-clients domain.com invia clstr2-dp1-clients outvia
clstr2-dp1-clients enable-intrazone
zone local clstr2-dp1-MCUs domain.com invia clstr2-dp1-MCUs outvia clstr2-dp1-MCUs
enable-intrazone
zone local clstr2-dp1-gateways domain.com invia clstr2-dp1-gateways outvia
clstr2-dp1-gateways enable-intrazone
zone local clstr2-dp2-clients domain.com invia clstr2-dp2-clients outvia
clstr2-dp2-clients enable-intrazone
```



```

zone local clstr2-dp2-MCUs domain.com invia clstr2-dp2-MCUs outvia clstr2-dp2-MCUs
enable-intrazone
zone local clstr2-dp2-gateways domain.com invia clstr2-dp2-gateways outvia
clstr2-dp2-gateways enable-intrazone
! zone subnet clstr1-dp1-clients default enable
no zone subnet clstr1-dp1-clients [clstr1-dp1 MCUs_IP_addr]/32 enable
no zone subnet clstr1-dp1-clients [clstr1-dp2 MCUs_IP_addr]/32 enable
no zone subnet clstr1-dp1-clients [clstr2-dp1 MCUs_IP_addr]/32 enable
no zone subnet clstr1-dp1-clients [clstr2-dp2 MCUs_IP_addr]/32 enable
no zone subnet clstr1-dp1-clients [clstr1-dp1 gateways_IP_addr]/32 enable
no zone subnet clstr1-dp1-clients [clstr1-dp2 gateways_IP_addr]/32 enable
no zone subnet clstr1-dp1-clients [clstr2-dp1 gateways_IP_addr]/32 enable
no zone subnet clstr1-dp1-clients [clstr2-dp2 gateways_IP_addr]/32 enable
! zone subnet clstr1-dp2-clients default enable
no zone subnet clstr1-dp2-clients [clstr1-dp1 MCUs_IP_addr]/32 enable
no zone subnet clstr1-dp2-clients [clstr1-dp2 MCUs_IP_addr]/32 enable
no zone subnet clstr1-dp2-clients [clstr2-dp1 MCUs_IP_addr]/32 enable
no zone subnet clstr1-dp2-clients [clstr2-dp2 MCUs_IP_addr]/32 enable
no zone subnet clstr1-dp2-clients [clstr1-dp1 gateways_IP_addr]/32 enable
no zone subnet clstr1-dp2-clients [clstr1-dp2 gateways_IP_addr]/32 enable
no zone subnet clstr1-dp2-clients [clstr2-dp1 gateways_IP_addr]/32 enable
no zone subnet clstr1-dp2-clients [clstr2-dp2 gateways_IP_addr]/32 enable
! zone subnet clstr2-dp1-clients default enable
no zone subnet clstr2-dp1-clients [clstr1-dp1 MCUs_IP_addr]/32 enable
no zone subnet clstr2-dp1-clients [clstr1-dp2 MCUs_IP_addr]/32 enable
no zone subnet clstr2-dp1-clients [clstr2-dp1 MCUs_IP_addr]/32 enable
no zone subnet clstr2-dp1-clients [clstr2-dp2 MCUs_IP_addr]/32 enable
no zone subnet clstr2-dp1-clients [clstr1-dp1 gateways_IP_addr]/32 enable
no zone subnet clstr2-dp1-clients [clstr1-dp2 gateways_IP_addr]/32 enable
no zone subnet clstr2-dp1-clients [clstr2-dp1 gateways_IP_addr]/32 enable
no zone subnet clstr2-dp1-clients [clstr2-dp2 gateways_IP_addr]/32 enable
zone subnet clstr2-dp2-clients default enable
no zone subnet clstr2-dp2-clients [clstr1-dp1 MCUs_IP_addr]/32 enable
no zone subnet clstr2-dp2-clients [clstr1-dp2 MCUs_IP_addr]/32 enable
no zone subnet clstr2-dp2-clients [clstr2-dp1 MCUs_IP_addr]/32 enable
no zone subnet clstr2-dp2-clients [clstr2-dp2 MCUs_IP_addr]/32 enable
no zone subnet clstr2-dp2-clients [clstr1-dp1 gateways_IP_addr]/32 enable
no zone subnet clstr2-dp2-clients [clstr1-dp2 gateways_IP_addr]/32 enable
no zone subnet clstr2-dp2-clients [clstr2-dp1 gateways_IP_addr]/32 enable
no zone subnet clstr2-dp2-clients [clstr2-dp2 gateways_IP_addr]/32 enable
no zone subnet clstr1-dp1-MCUs default enable
zone subnet clstr1-dp1-MCUs [clstr1-dp1 MCUs_IP_addr]/32 enable
zone subnet clstr1-dp1-MCUs [clstr1-dp1 RASAggregators_IP_addr]/32 enable
no zone subnet clstr1-dp2-MCUs default enable
zone subnet clstr1-dp2-MCUs [clstr1-dp2 MCUs_IP_addr]/32 enable
zone subnet clstr1-dp2-MCUs [clstr1-dp2 RASAggregators_IP_addr]/32 enable
no zone subnet clstr2-dp1-MCUs default enable
zone subnet clstr2-dp1-MCUs [clstr2-dp1 MCUs_IP_addr]/32 enable
zone subnet clstr2-dp1-MCUs [clstr2-dp1 RASAggregators_IP_addr]/32 enable
no zone subnet clstr2-dp2-MCUs default enable
zone subnet clstr2-dp2-MCUs [clstr2-dp2 MCUs_IP_addr]/32 enable
zone subnet clstr2-dp2-MCUs [clstr2-dp2 RASAggregators_IP_addr]/32 enable
no zone subnet clstr1-dp1-gateways default enable
zone subnet clstr1-dp1-gateways [clstr1-dp1 gateways_IP_addr]/32 enable
zone subnet clstr1-dp1-gateways [clstr1-dp1 RASAggregators_IP_addr]/32 enable
no zone subnet clstr1-dp2-gateways default enable
zone subnet clstr1-dp2-gateways [clstr1-dp2 gateways_IP_addr]/32 enable
zone subnet clstr1-dp2-gateways [clstr1-dp2 RASAggregators_IP_addr]/32 enable
no zone subnet clstr2-dp1-gateways default enable
zone subnet clstr2-dp1-gateways [clstr2-dp1 gateways_IP_addr]/32 enable
zone subnet clstr2-dp1-gateways [clstr2-dp1 RASAggregators_IP_addr]/32 enable
no zone subnet clstr2-dp2-gateways default enable
zone subnet clstr2-dp2-gateways [clstr2-dp2 gateways_IP_addr]/32 enable
zone subnet clstr2-dp2-gateways [clstr2-dp2 RASAggregators_IP_addr]/32 enable

```

```

no use-proxy clstr1-dp1-clients inbound-to terminals
no use-proxy clstr1-dp1-clients outbound-from terminals
no use-proxy clstr1-dp2-clients inbound-to terminals
no use-proxy clstr1-dp2-clients outbound-from terminals
no use-proxy clstr2-dp1-clients inbound-to terminals
no use-proxy clstr2-dp1-clients outbound-from terminals
no use-proxy clstr2-dp2-clients inbound-to terminals
no use-proxy clstr2-dp2-clients outbound-from terminals
! no use-proxy clstr1-dp1-MCUs inbound-to [MCU | gateway]
! no use-proxy clstr1-dp1-MCUs outbound-from [MCU | gateway]
! no use-proxy clstr1-dp2-MCUs inbound-to [MCU | gateway]
! no use-proxy clstr1-dp2-MCUs outbound-from [MCU | gateway]
! no use-proxy clstr2-dp1-MCUs inbound-to [MCU | gateway]
! no use-proxy clstr2-dp1-MCUs outbound-from [MCU | gateway]
! no use-proxy clstr2-dp2-MCUs inbound-to [MCU | gateway]
! no use-proxy clstr2-dp2-MCUs outbound-from [MCU | gateway]
! no use-proxy clstr1-dp1-gateways inbound-to gateway
! no use-proxy clstr1-dp1-gateways outbound-from gateway
! no use-proxy clstr1-dp2-gateways inbound-to gateway
! no use-proxy clstr1-dp2-gateways outbound-from gateway
! no use-proxy clstr2-dp1-gateways inbound-to gateway
! no use-proxy clstr2-dp1-gateways outbound-from gateway
! no use-proxy clstr2-dp2-gateways inbound-to gateway
! no use-proxy clstr2-dp2-gateways outbound-from gateway
gw-type-prefix 1# default-technology
! no arq reject-unknown-prefix
endpoint ttl 60
no shutdown

```

Migrating from Cisco Unified CallManager 4.0

The process of migrating from the design methodology documented in the *Cisco IP Video Telephony SRND for Cisco Unified CallManager 4.0* to the design methodology presented in this chapter consists of the following fundamental steps:

1. Reconfigure your H.323 clients in Cisco Unified CallManager Administration to take advantage of the new gatekeeper-controlled H.323 client feature.
2. Remove any H.225 gatekeeper-controlled trunks that were used for the H.323 client, MCU, and H.320 gateway zones, and replace them with RASAggregator trunks by creating dummy gatekeeper-controlled H.323 clients.
3. Reconfigure your endpoint gatekeeper(s) to use the new via-zone definition constructs, removing any zone prefixes that might be applied to the H.323 client, MCU, and gateway zones.
4. Modify all MCU and H.320 gateway service prefixes as necessary to include a # character at the beginning of the prefix.

Each of these steps will cause service disruptions to existing H.323 video endpoints, so ensure that these configuration changes are carefully planned and done at a time when the administrator will be able to minimize disruption to the existing H.323 video endpoints.

Applications

Cisco IP Communications provides an expanding portfolio of applications that extend the features of Cisco Unified CallManager and provide advanced capabilities and integration with other communication media. Many of these applications can be used in conjunction with IP Video Telephony

devices, even if they do not specifically support video. For instance, Cisco Unified CallManager Release 4.1 does not support the negotiation of video channels for CTI applications using the TAPI/JTAPI protocols, but that does not necessarily preclude using a CTI application in conjunction with a video call. This section reviews some of the Cisco and third-party applications and discusses whether or not they can be used to provide advanced call treatment for video calls.

CTI Applications

The following applications are based on the Computer Telephony Integration (CTI) interface.

Cisco Emergency Responder

Cisco Emergency Responder (ER) routes emergency (911) calls to the correct Public Safety Answering Point (PSAP). It also provides the PSAP with the correct calling line ID of the originating device so that the PSAP can respond to the correct physical location of the incident and call the party back in the event that the call is disconnected. Cisco ER uses JTAPI to integrate with Cisco Unified CallManager. Emergency calls are routed to Cisco ER via a CTI route point, then Cisco ER decides which PSAP to forward the call to and what calling line ID to display. Cisco ER tracks each endpoint on the network to determine its physical location by using Simple Network Management Protocol (SNMP) and Cisco Discovery Protocol (CDP) to discover the physical port and specific Cisco Catalyst Ethernet switch to which the endpoint is connected. If CDP is not available, Cisco ER can be configured to locate endpoints by their IP subnet instead. Cisco ER then correlates this information with the physical location of the switch and stores the information in its database.

Both Cisco Unified Video Advantage and the Cisco IP Video Phone 7985 support CDP for the purpose of Cisco ER discovery. Cisco Unified Video Advantage does not send CDP messages to the switch directly but relies on its associated Cisco Unified IP Phone for this support. Therefore, if a Video Telephony user dials 911, Cisco ER is able to route the call to the correct PSAP.

Because Sony and Tandberg SCCP endpoints do not support CDP, Cisco ER must track these endpoints by their IP subnet. Cisco ER is therefore able to route the call to the correct PSAP.

Because H.323 videoconferencing clients do not support CDP, Cisco ER must track them by their IP subnet. Cisco ER is therefore able to route the call to the correct PSAP. However, the H.323 device must support the Empty Capabilities Set (ECS) procedure in order to have its call routed by Cisco ER. If the H.323 endpoint does not support receiving an ECS from Cisco Unified CallManager, calls to 911 that are handled by Cisco ER will fail.

Cisco Unified CallManager Assistant

Cisco Unified CallManager Assistant enables administrative assistants to provide coverage for the managers with whom they are associated. Unified CM Assistant uses JTAPI to integrate with Cisco Unified CallManager. While Unified CM Assistant is not specifically video-capable, there is nothing stopping Unified CM Assistant from being used with phones that are video-enabled. Once Unified CM Assistant has handled the call and the call is transferred to the final destination device, the two devices in the call communicate with each other directly and could establish video channels at that point. For example, if a video-capable endpoint dialed the directory number of a Manager, and the Assistant covered the call using the Unified CM Assistant application, there might not be any video established during the initial handling of the call. But once the Assistant transferred the caller to the Manager, Cisco Unified CallManager could then negotiate video channels. However, H.323 devices must support the Empty Capabilities Set (ECS) procedure in order to interoperate with

Unified CM Assistant. If the H.323 endpoint does not support receiving an ECS from Cisco Unified CallManager, calls that are intercepted by Unified CM Assistant will fail when the Assistant attempts to transfer the call to the Manager.

Cisco IP Interactive Voice Response and Cisco IP Contact Center

Cisco Unified IP Interactive Voice Response (Unified IP IVR) and Cisco Unified Contact Center Enterprise (Unified CCE) use JTAPI to integrate with Cisco Unified CallManager. If a video-capable device calls into an IVR application (such as a help desk), the communication is audio-only while the caller is connected to the application server (that is, while the caller browses the IVR menu or waits in queue for a help-desk member to take the call). However, once the IVR application transfers the call to its final destination, video channels can be negotiated at that time. H.323 devices must support the Empty Capabilities Set (ECS) procedure in order to interoperate with Cisco Unified IP-IVR and Unified CCE. If the H.323 endpoint does not support receiving an ECS from Cisco Unified CallManager, calls that are intercepted by Cisco Unified IP-IVR or Unified CCE will fail when the application attempts to transfer the caller to the final destination.

IVR applications often use DTMF tones to select options in the IVR menu. An alternative is speech recognition, which enables the caller to speak commands to the IVR server instead of pressing keys on the phone. Because Cisco Unified IP-IVR and Unified CCE both use JTAPI to integrate with Cisco Unified CallManager, they pass DTMF tones through out-of-band signaling messages. Many H.323 devices on the market today use in-band DTMF tones, and these H.323 clients would not be able to use DTMF to navigate a Unified IP IVR or Unified CCE menu. However, these H.323 clients could use speech recognition if the IVR server is enabled for it. Video-capable devices such as Cisco Unified Video Advantage, Sony or Tandberg SCCP devices, and any H.323 endpoint that uses H.245 alphanumeric out-of-band signaling for DTMF, can navigate the IVR menus using DTMF tones.

Cisco Attendant Console

Cisco Attendant Console uses JTAPI to integrate with Cisco Unified CallManager. The Attendant Console is used as a front-office device to handle incoming calls. Although Attendant Console does not specifically support video, video channels can be negotiated once the call is transferred to its final destination. However, H.323 devices must support the Empty Capabilities Set (ECS) procedure in order to interoperate with the Attendant Console. If the H.323 endpoint does not support receiving an ECS from Cisco Unified CallManager, calls that are intercepted by an Attendant Console will fail when the attendant attempts to transfer the caller to the final destination.

Cisco Personal Assistant

Cisco Personal Assistant (PA) provides two main functions:

- Dialing assistance either through DTMF tones or voice commands
- User-programmed call routing preferences based on the time of day, caller ID information, and other factors

For both of these functions, no video is established while the call is connected to the PA server, but video channels can be negotiated once PA transfers the call to its final destination. H.323 devices must support the Empty Capabilities Set (ECS) procedure in order to interoperate with PA. If the H.323 endpoint does not support receiving an ECS from Cisco Unified CallManager, calls that are intercepted by PA will fail when the application attempts to transfer the caller to the final destination.

Cisco IP SoftPhone and Cisco IP Communicator

Cisco IP SoftPhone uses TAPI to integrate with Cisco Unified CallManager and can be configured either as a standalone softphone or as a software interface to control an associated SCCP hardware phone. Cisco IP SoftPhone does not specifically support video, but it can be used in conjunction with an IP Phone that has a Cisco Unified Video Advantage client associated to it. Cisco IP SoftPhone cannot be used to control a Sony or Tandberg SCCP device.

Cisco IP Communicator differs from IP SoftPhone in that it is an SCCP client and therefore acts like a Cisco 7970 Series IP Phone. If you use Cisco Unified Video Advantage 2.0 and Cisco IP Communicator 2.0 on the same PC, these two applications can "associate" with each other, thus eliminating the need for a physical IP Phone.

Collaboration Solutions

The following technologies are sometimes used to provide video communications between endpoints.

T.120 Application Sharing

Some videoconferencing endpoints use the T.120 protocol to share documents, whiteboards, and text among participants in a conference. Cisco Unified CallManager does not support negotiating a T.120 channel. Instead of T.120, Cisco recommends using web-based collaboration solutions such as Cisco MeetingPlace or other third-party collaboration solutions such as Placeware, Web-Ex, or IBM E-Collaborate.

Cisco Unified MeetingPlace

Cisco Unified MeetingPlace combines a high-end audio and video conferencing solution with a web-based front end for scheduling and participating in conferences. For more information, refer to the chapter on [Cisco Unified MeetingPlace Integration, page 15-1](#).

Wireless Networking Solutions

Because video is so bandwidth intensive, Cisco does *not* recommend using a shared wireless medium such as 802.11b for video endpoints.

Cisco Aironet 802.11b Networking Solutions

The Tandberg T-1000 model endpoint offers a PCMCIA interface in which you can install a Cisco Aironet PCMCIA 802.11b Wireless Adapter. However, you should take care to ensure that the video endpoints do not share the wireless bandwidth with any production IP Phones because video will consume much of the bandwidth and make it difficult to support video, audio, and data all on the same wireless medium.

If you use Cisco Unified Video Advantage with a hardware IP Phone, Cisco Unified Video Advantage relies on a physical Ethernet connection to the IP Phone with which it is associated. It is not uncommon for a user to have both a physical Ethernet interface and an Aironet 802.11b Wireless Adapter installed in the same PC. This configuration could cause problems for Cisco Unified Video Advantage if the wireless interface happens to be the preferred path to the network because Cisco Unified Video

Advantage will not associate over that interface. Cisco recommends that you always make the physical Ethernet interface the preferred path. Also, when users connect to the PC port on the back of the IP Phone, instruct them to disable their Aironet Adapter to keep it from accidentally taking precedence.

Cisco Unified Wireless IP Phone 7920

The Cisco Unified Wireless IP Phone 7920 does not support video. Nothing prevents a video endpoint from calling a Cisco Unified Wireless IP Phone 7920; however, it will negotiate as an audio-only call. The Wireless IP Phone user can then put the call on hold, transfer it, or conference it. If the caller is an H.323 video endpoint, it must support the Empty Capabilities Set (ECS) procedure in order for these supplementary services to work.

XML Services

Currently there are no XML applications specifically written for the Cisco Unified Video Advantage client solution, Cisco IP Video Phone 7985, or the Sony or Tandberg SCCP endpoints. However, with the exception of the Sony endpoints, these endpoints do support XML applications. Cisco VT Advantage uses a Cisco Unified IP Phone, so any XML application supported on those phone models should also work with VT Advantage.

The Tandberg SCCP endpoints support XML, but not all XML applications currently work on the Tandberg endpoints. For example, Cisco Extension Mobility and the Berbee InformaCast product are two popular XML applications that currently do not work on the Sony or Tandberg SCCP endpoints. Support for these applications will require firmware upgrades to the endpoints as well as changes in Cisco Unified CallManager Administration in some cases.



CHAPTER 18

LDAP Directory Integration

Last revised on: February 13, 2008

Directories are specialized databases that are optimized for a high number of reads and searches, and occasional writes and updates. Directories typically store data that does not change often, such as employee information, user privileges on the corporate network, and so on.

Another aspect of directories is that they are extensible, which means that the type of information stored in them can be modified and extended. The term *directory schema* refers to the type of stored information and the rules it obeys.

The Lightweight Directory Access Protocol (LDAP) provides applications with a standard method for accessing and potentially modifying the information stored in the directory. This capability enables companies to centralize all user information in a single repository available to several applications, with a remarkable reduction in maintenance costs through the ease of adds, moves, and changes.

This chapter covers the main design principles for integrating a Cisco IP Communications system based on Cisco Unified CallManager 4.x with a corporate LDAP directory. The main topics include:

- [What is Directory Integration?, page 18-2](#)

This section analyzes the various requirements for integration with a corporate LDAP directory in a typical enterprise IT organization.

- [Directory Access for IP Telephony Endpoints, page 18-3](#)

This section describes the technical solution to enable directory access for Cisco Unified Communications endpoints and provides design best-practices around it.

- [Directory Integration with Cisco Unified CallManager 4.x, page 18-6](#)

This section describes the technical solutions and provides design best-practices for directory integration with Cisco Unified CallManager 4.x.

The considerations presented in this chapter apply to Cisco Unified CallManager 4.x as well as the following applications bundled with it: Extension Mobility, Cisco Unified CallManager Assistant, WebDialer, Bulk Administration Tool, and Real-Time Monitoring Tool.

For all other Cisco voice applications, refer to the respective product documentation available at:

<http://www.cisco.com>

In particular, for Cisco IP Contact Center refer to the *Cisco Cisco Unified Contact Center Enterprise Edition SRND* and the *Cisco Cisco Unified Contact Center Express SRND*, both available at

<http://www.cisco.com/go/designzone>


For Cisco Unity, refer to the *Cisco Unity Design Guide* and to the following white papers: *Cisco Unity Data and the Directory*, *Active Directory Capacity Planning*, and *Cisco Unity Data Architecture and How Cisco Unity Works*, also available at

<http://www.cisco.com>

What is Directory Integration?

Integration between voice applications and a corporate LDAP directory is a common task for many enterprise IT organizations. However, the exact scope of the integration varies from company to company, and it can translate to one or more specific and independent requirements, as shown in Figure 18-1.

Figure 18-1 Various Requirements for Directory Integration



For example, one common requirement is to enable user lookups (sometimes called the "white pages" service) from IP phones or other voice and/or video endpoints, so that users can dial a contact directly after looking up their number in the directory.

Another requirement is to provision users automatically from the corporate directory into the user database of voice and/or video applications. This method avoids having to add, remove, or modify core user information manually each time a change occurs in the corporate directory.

Often authentication of end-users and administrators of the voice and/or video applications using the corporate directory credentials is also required. This method enables the IT department to deliver single log-on functionality and reduces the number of passwords that each user needs to maintain across different corporate applications.

Each of these requirements can be satisfied by a Cisco IP Communications system using different mechanisms according to the Cisco Unified CallManager version used, as summarized in Table 18-1.

Table 18-1 **Directory Requirements and Cisco Solutions**

Requirement	Cisco Solution	Cisco Unified CallManager 4.x Feature	Cisco Unified CallManager 5.0 Feature
User lookup for endpoints	Directory access	Cisco Unified IP Phone Services SDK	Cisco Unified IP Phone Services SDK
User provisioning	Directory integration	Cisco Customer Directory Configuration Plugin	LDAP Synchronization
Authentication for IP Telephony end users	Directory integration	Cisco Customer Directory Configuration Plugin	LDAP Authentication
Authentication for IP Telephony application administrators	Directory integration	Cisco Customer Directory Configuration Plugin + Cisco Multilevel Administration	LDAP Authentication

As shown in [Table 18-1](#), within the context of a Cisco IP Communications system, the term *directory access* refers to mechanisms and solutions that satisfy the requirement of user lookups for IP Telephony endpoints, while the term *directory integration* refers to mechanisms and solutions that satisfy the requirements of user provisioning and authentication (for both end users and administrators).

The remainder of this chapter describes how to address these requirements in a Cisco IP Communications system based on Cisco Unified CallManager Release 4.x. For a detailed description of directory integration solutions with Cisco Unified Communications Manager Release 5.x, refer to the *Cisco Unified Communications SRND Based on Cisco Unified Communications Manager 5.x*, available at

<http://www.cisco.com/go/designzone>

Directory Access for IP Telephony Endpoints


This section describes how to configure corporate directory access to any LDAP-compliant directory server to perform user lookups from Cisco Unified Communications endpoints (such as Cisco Unified IP Phones). The guidelines contained in this section apply regardless of whether Cisco Unified CallManager or other IP Telephony applications have been integrated with a corporate directory for user provisioning and authentication.

Cisco Unified IP Phones equipped with a display screen can search a user directory when a user presses the Directories button on the phone. The IP Phones use Hyper-Text Transfer Protocol (HTTP) to send requests to a web server. The responses from the web server must contain some specific Extensible Markup Language (XML) objects that the phone can interpret and display.

By default, Cisco Unified IP Phones are configured to perform user lookups against Cisco Unified CallManager's embedded database. However, it is possible to change this configuration so that the lookup is performed on a corporate LDAP directory. In this case, the phones send their HTTP requests to an external web server that operates as a proxy and translates these requests into LDAP queries against the corporate directory. The LDAP responses are then encapsulated in the appropriate XML objects and sent back to the phones via HTTP.

Figure 18-2 illustrates this mechanism in a deployment where Cisco Unified CallManager has not been integrated with the corporate directory. Note that, in this scenario, Cisco Unified CallManager is not involved in the message exchange related to the user lookup.

Figure 18-2 *Directory Access for Cisco Unified IP Phones Using the Cisco Unified IP Phone Services SDK*



153280

In the example shown in Figure 18-2, the web server proxy function is provided by the Cisco LDAP Search Component Object Model (COM) server, which is included in the Cisco Unified IP Phone Services Software Development Kit (SDK) version 3.3(4) or later. You can download the latest Cisco Unified IP Phone Services SDK from Cisco Developer Support Central at

http://www.cisco.com/en/US/products/svcs/ps3034/ps5408/ps5418/serv_home.html

The IP Phone Services SDK can be installed on a Microsoft Windows web server running IIS 4.0 or later, but it cannot be installed on a Cisco Unified CallManager server. The SDK includes some sample scripts to provide simple directory lookup functionality.

To set up a corporate directory lookup service using the IP Phone Services SDK, perform the following steps:

-
- Step 1** Modify one of the sample scripts to point to your corporate LDAP directory, or write your own script using the LDAP Search COM Programming Guide provided with the SDK.
- Step 2** In Cisco Unified CallManager, configure the URL Directories parameter (under **System > Enterprise Parameters**) to point to the URL of the script on the external web server.
- Step 3** Reset the phones to make the changes take effect.
-

**Note**

If you want to offer the service only to a subset of users, configure the URL Directories parameter directly within the Phone Configuration page instead of the Enterprise Parameters page.

In conclusion, the following design considerations apply to directory access with the Cisco Unified IP Phone Services SDK:

- User lookups are supported against any LDAP-compliant corporate directory.
- When querying Microsoft Active Directory, you can perform lookups against the Global Catalog by pointing the script to a Global Catalog server and specifying port 3268 in the script configuration. This method typically results in faster lookups.
- There is no impact on Cisco Unified CallManager for this functionality, and only minimal impact on the LDAP directory server.
- The sample scripts provided with the SDK allow only a minimal amount of customization (for example, you can prefix a digit string to all returned numbers). For a higher degree of manipulation, you will have to develop custom scripts, and a programming guide is included with the SDK to aid in writing the scripts.
- This functionality does not entail provisioning or authentication of Cisco Unified CallManager users with the corporate directory.

Directory Integration with Cisco Unified CallManager 4.x

This section describes the mechanisms and best practices for directory integration with Cisco Unified CallManager 4.x to allow for user provisioning and authentication with a corporate LDAP directory. This section covers the following topics:

- [Cisco Unified CallManager 4.x Directory Architecture, page 18-6](#)
- [The Cisco Customer Directory Configuration Plugin, page 18-9](#)
- [Security Considerations, page 18-10](#)
- [Adding Cisco Unified CallManager Servers to a Domain, page 18-11](#)
- [Comparison with the Cisco Unified CallManager 5.0 Approach, page 18-12](#)

Cisco Unified CallManager 4.x Directory Architecture

Cisco Unified CallManager 4.x uses an embedded Microsoft SQL database to store system and device configuration data, such as dial plan information, phone and gateway configurations, and media resource utilization. It also uses an embedded LDAP directory to store user and application profiles, such as devices controlled by a user, computer telephony integration (CTI) user parameters, and personal address book entries.

Both the SQL database and the LDAP directory run on every Cisco Unified CallManager 4.x server within a cluster, and replication agreements are automatically set up between servers. The publisher server contains the master copy of both the SQL database and the LDAP directory, and it handles replication to all subscriber servers, which contain read-only copies of both repositories.

**Note**

The examples and recommendations in this chapter are based on Cisco Unified CallManager 4.x, which introduced some new features and enhancements. Some behaviors might differ and some features might not be available if you are running an older version of Cisco CallManager.

To store application-specific information in an LDAP directory, Cisco Unified CallManager 4.x adopts an approach that is valid both when using the embedded directory and when integrating with a corporate directory.

Cisco Unified CallManager utilizes a subset of the Cisco registered object identifies (OIDs) when it extends the schema. The OID values are 1.2.840.113548.3.1.4.xx for attributes and 1.2.840.113548.3.2.4.xx for object classes.


Because different directory vendors typically use different User object models with several additional, nonstandard attributes, Cisco Unified CallManager 4.x uses only the standard LDAPv3 core attributes from the User object. The User object is then augmented with an auxiliary class, `ciscoocUser`, which contains the following attributes:

- `ciscoatGUID`
This attribute uniquely identifies a user within the directory.
- `ciscoatUserProfile`
This attribute is used by earlier versions of Cisco Unified CallManager 4.x and other applications. It is still present for backward compatibility.

- ciscoatUserProfileString

This attribute is a distinguished name pointer to another object in the directory, which contains the user's application-specific profile. This approach minimizes the impact on the core User object, and all the application-specific information can be stored in a separate organizational unit (OU) within the directory, usually called the Cisco subtree, CISCOBASE, or Cisco Directory Information Tree (DIT). [Figure 18-3](#) illustrates this process.

Figure 18-3 Approach to Storing Application-Specific User Information in the Directory with Cisco Unified CallManager 4.x



The object pointed to by the `ciscoatUserProfileString` attribute belongs to a structural object class called `ciscoocUserProfile`. The main purpose of this object is to store some specific details for the user, including the user's locale, any Cisco IP Manager Assistant (IPMA) assistants for the user, and pointers to various specific profile objects for all Cisco applications that integrate with the directory. The application profile used by Cisco Unified CallManager 4.x belongs to the auxiliary class called `ciscoCCNocAppProfile`, and it is where Cisco Unified CallManager 4.x stores the user's extension mobility PIN, the list of devices controlled by this user, information on whether this user is permitted to use CTI applications, and so on. Both of these profile objects are created by Cisco Unified CallManager 4.x under the "Cisco" subtree.



Note

The list of devices associated with a user is stored in the directory as a multi-valued attribute. If you are using the embedded Cisco Unified CallManager 4.x directory (known as DC Directory), the maximum number of devices that can be associated with a user is 2,000 (or 2,500 if all subscribers in the cluster are dual-CPU servers). However, Microsoft Active Directory limits the number of values that can be stored in a multi-valued attribute. The total size of the record should never exceed 7,000 bytes, which allows for approximately 850 attribute values.

If you are integrating Cisco Unified CallManager 4.x with Microsoft Active Directory (AD), use the following method to calculate the maximum number of devices, device profiles, or combination of both that can be associated with a single user:

Total record size = $9 * (\text{Total number of devices associated with a user}) + (\text{Size in bytes of the name of each associated device profile})$.

This calculated record size must be less than 7,000 bytes.

The following examples illustrate the use of this calculation method.

Example 18-1 Devices Only (Device Profiles Not Used)

Number of devices = 500.

Total record size = $9 * 500 = 4,500$ bytes.

Because the calculated record size is less than 7,000, this user's device associations are within the acceptable limit for AD.

Example 18-2 Devices and Device Profiles

Number of devices = 400.

Number of device profiles = 339.

Each profile name is 10 characters (bytes).

Total record size = $(9 * 400) + (339 * 10) = 3600 + 3390 = 6,990$ bytes.

Because the calculated record size is less than 7,000, this user's device associations are within the acceptable limit for AD.

The total number of associations for this user = $400 + 339 = 739$.

Example 18-3 Devices and Device Profiles

Number of devices = 400.

Number of device profiles = 400.

Each profile name is 10 characters (bytes).

Total record size = $(9 * 400) + (400 * 10) = 3600 + 4000 = 7,600$ bytes.

This user's device associations exceed the limit of 7,000 bytes and would not be an acceptable directory entry in AD. To bring this user's device associations within the record size limit for AD, reduce the number of devices and device profiles associated with this user.

The Cisco Customer Directory Configuration Plugin


To integrate Cisco Unified CallManager with an external LDAP directory, run the Cisco Customer Directory Configuration Plugin, which is bundled with Cisco Unified CallManager (**Applications > Install Plugins**). This plugin serves three main purposes:

- It extends the corporate directory schema to accommodate the application-specific objects and attributes.
- It populates the "Cisco" subtree with the configuration objects needed by Cisco Unified CallManager.
- It configures Cisco Unified CallManager to use the corporate directory and disables its embedded directory.

Usually, running the plugin locally on Cisco Unified CallManager performs the schema update. However, starting with Cisco Unified CallManager Release 4.0, a new option exists to create the LDAP Data Interchange Format (LDIF) files separately so that the schema update can be performed directly on the corporate directory's Schema Master server using the LDIF files. This option allows different groups of people to perform the relevant parts of the work and reduces the need to update over the network when Cisco Unified CallManager is not local to the Schema Master server.

After the plugin has been run, Cisco Unified CallManager effectively uses the corporate directory to store user preferences. If the Cisco IP Telephony endpoints have also been enabled to access the corporate directory, as described in the preceding section, the resulting scenario is that shown in [Figure 18-4](#).

Figure 18-4 Message Exchange for Cisco IP Phone Corporate Directory Access When Cisco Unified CallManager Is Integrated with the Corporate Directory



Security Considerations

Starting with Cisco Unified CallManager Release 4.1, the communication between Cisco Unified CallManager and the embedded directory is set by default to use LDAP over Secure Socket Layer (SSL), also known as LDAPS.

When integrating with a corporate directory, it is also possible to enable LDAP over SSL, thus ensuring that all sensitive LDAP data goes over a secure connection. The LDAP over SSL option can be configured as part of the Cisco Customer Directory Plugin, and it requires a certificate shared with the corporate directory and issued by the same Certificate Authority.

When LDAP over SSL is enabled for Cisco Unified CallManager, the following additional Cisco IP Communications applications will also communicate with the directory over this secure channel:

- User pages within Cisco Unified CallManager Administration
- Cisco Multi-Level Administration (MLA)
- Cisco IP Phone Options pages
- Extension Mobility application — SSL is used for communications between the Extension Mobility application, running on a Cisco Unified CallManager server, and the corporate directory. However, SSL is not used for communications between the IP phones and the Extension Mobility application, which use HTTP instead.
- Cisco CTI Manager
- Serviceability and Cisco Real-Time Monitoring Tool (RTMT)
- Cisco CDR Analysis and Reporting (CAR)
- Cisco IP Manager Assistant (IPMA) service
- Cisco Bulk Administration Tool (BAT)

Multilevel Administration Account Security

Prior to Cisco Unified CallManager Release 4.2, the Multilevel Administration (MLA) function performed caching of the authentication status of users that had an MLA association. When MLA was integrated with an off-cluster Microsoft Active Directory (AD), the directory's password security features (such as password aging) could cause interruption of correct MLA functionality. In Cisco Unified CallManager 4.2, a new set of controls was added to allow control over the caching function of MLA so that it would not be affected by these password security features.

The password caching can allow a user to authenticate successfully in some instances when the directory has disabled or otherwise blocked an account since the cache entry was made. The new feature also avoids this behavior.

In Cisco Unified CallManager 4.2, the User Cache Flush Timeout parameter is enabled with a default value of 5 minutes. The parameter determines the lifetime of the cached authentication status. The AVVID XML Layer (AXL) and Simple Object Access Protocol (SOAP) interfaces on Cisco Unified CallManager are also affected by this parameter because they use the same authentication mechanism as MLA. The parameter for caching can be set to disable caching. When using applications enabled for SOAP or AXL, be aware that their performance will be impacted because they are usually polling-type applications that require frequent authentication. Therefore, Cisco recommends that you do *not* disable the caching.

Adding Cisco Unified CallManager Servers to a Domain

Adding the Cisco Unified CallManager servers to a Microsoft Windows domain differs substantially from integrating Cisco Unified CallManager with an external directory. Although these operations do not exclude each other, they are completely independent and have different implications:

- Adding Cisco Unified CallManager servers to a Microsoft Windows Active Directory (AD) domain could cause domain policies to be applied to the Windows 2000 Server operating system, and the addition affects only the management of the Cisco Unified CallManager server itself.
- Integrating Cisco Unified CallManager with an external directory (such as Microsoft Active Directory or Netscape Directory Server) causes Cisco Unified CallManager to store all its user information and preferences in that directory, but it does not affect management of the Cisco Unified CallManager server itself.

Cisco recommends that you keep Cisco Unified CallManager servers as workgroup servers. However, if you want to add the servers to a domain, avoid applying domain policies to the server that could interfere with its normal operation.


For additional recommendations about adding servers to a domain, refer to the latest Cisco Unified CallManager product documentation available online at

<http://www.cisco.com>

Comparison with the Cisco Unified CallManager 5.0 Approach

The approach to directory integration changes significantly with Cisco Unified CallManager Release 5.0, and this section highlights the main differences with the current Cisco Unified CallManager 4.x approach. Figure 18-5 shows high-level functional diagrams of the two approaches to achieve user provisioning and authentication.

Figure 18-5 Directory Integration Approaches in Cisco Unified CallManager 4.x and 5.0



Cisco Unified CallManager Release 4.x uses an embedded LDAP directory to store user-related information. Directory integration is enabled by extending the corporate directory schema, shutting down the embedded directory, and using the corporate directory to store the application-specific data related to the users. Because the corporate directory is effectively used as the back-end storage repository for user information, this method satisfies the requirement for both user provisioning and user authentication. Any changes to user data in the corporate directory are immediately picked up by Cisco Unified CallManager because it accesses the same data store.

However, this approach has an impact on the corporate directory in terms of schema extensions and additional data, and it also introduces dependencies between the real-time functionality of the IP Communications system and the availability of the directory. When connectivity is lost or the directory becomes unavailable, Cisco Unified CallManager is unable to access all user-related configurations, which impacts applications such as Extension Mobility, Attendant Console, and IP Contact Center Express. In this approach, the user provisioning and user authentication functions have to be enabled at the same time because they rely on the same integration process. In addition, using the corporate directory as the storage repository for application-specific data also imposes limitations on the day-to-day maintenance operations on the corporate directory itself, as described in the remainder of this chapter.

By contrast, the approach to directory integration adopted by Cisco Unified CallManager Release 5.0 relies on two separate components to satisfy the user provisioning and user authentication requirements independently. User provisioning is performed with a one-way synchronization of user data from the corporate directory to the Cisco Unified CallManager embedded database. The synchronization uses standard LDAPv3 and can be triggered manually or scheduled periodically to ensure that changes are incorporated into the Cisco IP Communications system. This solution avoids the need to write anything to the corporate directory, and it does not require any schema extensions.

User authentication is enabled independently from user provisioning, and it provides authentication of end-user passwords against the corporate directory credentials. With this approach, the Cisco IP Communications system preserves all of its real-time functionality even when the corporate directory is unavailable or unreachable.

For more information on directory integration with Cisco Unified Communications Manager 5.x, refer to the *Cisco Unified Communications SRND Based on Cisco Unified Communications Manager 5.x*, available on Cisco.com at

<http://www.cisco.com/go/designzone>

Best Practices for Directory Integration

The directory integration process involves several components and services in your network and therefore should be carefully planned and implemented. This section addresses following topics:

- [Planning the Directory Integration, page 18-13](#)
- [Preparing the Directory for Integration, page 18-15](#)
- [Integrating Cisco Unified CallManager with the Directory, page 18-19](#)
- [Maintaining the Directory Integration, page 18-22](#)



Note

Because the majority of known Cisco Unified CallManager integrations are performed with Microsoft Active Directory (AD), this chapter focuses on best practices for AD. However, most recommendations and best practices mentioned in this section also apply to the other directory product supported by Cisco Unified CallManager, the Sun/iPlanet Netscape Directory Server.

Planning the Directory Integration

Because the directory is an enterprise-wide resource that is used by a potentially large number of applications and end users, it is essential to plan the integration carefully to minimize the impact on all other applications.



Tip

Before starting the integration, ensure that your organization's directory team is involved in the planning, design, and implementation phases.

As mentioned previously in this chapter, integrating Cisco Unified CallManager and other applications with an external directory involves extending the directory schema. Schema extension is a delicate operation. For example, in the case of Microsoft Windows 2000 Active Directory, schema changes cannot be undone. You should take the following precautions to avoid damaging the directory:

- Review the planned schema changes with your organization's directory team. This practice should be part of your organization's change control procedures.
- Create a replica of the production directory in a lab setup, and test the integration against it.
- Back up the production directory, both data as well as schema, before integration, and make sure a workable back-out plan exists to enable successful restoration of the data and schema if it becomes necessary.
- Plan and perform the schema extension during off-peak hours to minimize the impact on other applications and end users.

Although the preceding list of precautions might cause concern at first, in practice the schema extension rarely causes problems that would require it to be backed out. However, no matter how safe an operation is known to be, you should not add risk by skipping any possible precautions to ensure that you can recover from unplanned issues.

Another important consideration is that, as soon as the voice applications have been integrated with the directory, they rely on it for their correct operation, and inability to reach the directory server can negatively affect the voice system.

For example, if the directory suddenly becomes unavailable, end users cannot log into the Cisco Unified CallManager User Options web page to configure their preferences; Extension Mobility users, Attendant Console operators, and Unified Contact Center Express agents cannot log in or out; and the dial-by-name function is unavailable.

To avoid these problems, you should design your directory infrastructure so that it is highly available to all Cisco voice applications. You can use any of the following methods to achieve this high availability:

- Leverage the directory replication mechanism to place a directory server or servers in the same location as the Cisco voice applications.
- Use a server load-balancing mechanism, such as Cisco IOS software server load balancing (SLB), to provide server redundancy in a specific campus or data center and to ensure that local servers are accessed by preference.
- Use Domain Name System (DNS) domain names instead of specific domain controller host names when configuring the directory plugin.

With redundant servers, the first name returned by DNS might be the name of a server that is not as local to Cisco Unified CallManager as others returned later in the response. Also, if your DNS server has the round-robin feature enabled, by design it rotates the order in which addresses are returned in the response. Depending on mechanisms such as client-side DNS cache timeout, along with other possible clients querying for the same domain in the interim, Cisco Unified CallManager could run two consecutive operations against two different domain controllers (DCs). In addition to the locality problem already mentioned, using DNS redundancy could keep objects created in the first operation from being found by a search on a different DC by a later query if the directory has not replicated in the meantime. Therefore, before choosing to use DNS to make the implementation redundant, be sure that these issues do not affect your deployment.

Also note that DNS is needed for proper LDAP referral; Cisco Unified CallManager must be able to resolve the host names of any of the DCs returned in an LDAP referral.

**Note**

Although Microsoft Windows 2000 DNS has a feature that returns local resources first (called LocalNetPriority), it is based on inspecting the requesting client's classful IP address. Therefore, it is of limited use in subnetted networks. Microsoft Knowledge Base article 177883 documents this feature (see <http://support.microsoft.com/>). If you are not using Windows 2000 DNS, you should check to see which features in your chosen implementation might alleviate some of these issues. These recommendations are based on the assumptions that Cisco Unified CallManager is set to use DNS and that it is the same DNS infrastructure used by your Active Directory.

Preparing the Directory for Integration

With Microsoft Windows 2000 Active Directory, the domain controllers have to be set to allow schema changes. This requirement applies only to the domain controller that acts as the Schema Master (the location where change takes place), and it is fully documented in Microsoft Knowledge Base article 285172, which is available at

<http://support.microsoft.com>

If you are integrating Cisco Unified CallManager with the Microsoft Windows 2000 Active Directory, and if Microsoft Exchange 2000 needs to coexist in the same forest, you must perform an additional preparation step. Cisco Unified CallManager uses the labeledURI attribute specified by the iNetOrgPerson class, as defined in RFC 2798. Microsoft currently defines this attribute differently for Exchange 2000, which causes a naming clash with the Cisco Unified CallManager schema. The problem is documented in Microsoft Knowledge Base article 314649 (available from <http://support.microsoft.com>). You can obtain the iNetOrgPerson kit from

<http://msdn.microsoft.com/en-us/library/ms808546.aspx>

**Caution**

Remember to back up your directory before extending the schema. Make sure you have tested the restore mechanism you intend to use *before* you need it.

To extend the directory schema for Cisco Unified CallManager, run the Cisco Customer Directory Configuration plugin from the publisher server for the cluster, and follow these best practices:

- Always use the **Custom** setup type when running the plugin. The Express setup type is appropriate only for integration with a standalone domain used exclusively for Cisco Unified CallManager and other Cisco voice applications. You should never use the Express setup when integrating with an existing domain.
- Select only the **Install Schema on the Schema Master** option in the plugin configuration screen.
- Try to ensure that the Schema Master server is relatively local to Cisco Unified CallManager or that a high-speed connection exists between them. If neither is possible, then consider creating just the LDIF files with the plugin and using them directly to update the schema on the Schema Master.
- In this phase, provide the plugin with credentials (distinguished name and password) for a user who is a member of the Schema Admins group in Active Directory. These credentials are used only for the schema extension, not for normal Cisco Unified CallManager operation.

**Note**

If you have a large AD forest or a complex topology, schema changes might take some time to propagate to all domains and all domain controllers in the forest. Allow enough time for this propagation to happen before continuing with the preparation process, or else force replication if required.


As soon as the schema has been extended, you can decide where to create the "Cisco" OUs (that is, the subtrees) that will be used by the Cisco Unified CallManager clusters and the other Cisco voice applications. There must be one organizational unit (OU) for each Cisco Unified CallManager cluster to be integrated.

For a deployment with a single domain AD or Netscape directory server, placement is not critical. The OU can effectively be placed anywhere in the tree.

In a multiple-domain AD forest, the root domain is often kept free of users and resources and is used as a placeholder domain, so the "Cisco" subtrees would typically reside in the child domains. In this type of multiple-domain topology, domains can be created based on geographic boundaries. Therefore, it is unlikely that every location has a local domain controller for each domain. To reduce replication traffic across the network, domain controllers are usually placed only where needed. With this in mind, it is best to try to place the Cisco OU for a given cluster in the domain containing the majority of users serviced by that cluster.

Figure 18-6 shows a multiple-domain, single-tree AD forest in which the "Cisco" OUs for two Unified CallManager clusters have been created in the two respective child domains, emea.vse.lab and amer.vse.lab.

Figure 18-6 Multiple-Domain, Single-Tree AD Forest




In Figure 18-6, each cluster services the corresponding geography in a centralized call processing model, thus ensuring that its user data stored in AD is also local. This design saves having to retrieve the information from a DC that is probably non-local, and it reduces the number of LDAP referrals required to find the relevant information in a search.

Cisco does not recommend having a Cisco Unified CallManager cluster service users in different domains because response times while user data is being retrieved might be less than optimal if domain controllers for all included domains are not local. However, this scenario is not a common one because the reasons for creating a multiple-domain AD (namely, geography, bandwidth, or organizational structure) are usually the same reasons for needing multiple clusters.

Currently clusters cannot span trees within an AD forest because they would not have a contiguous namespace, which is a requirement for LDAP referrals. A cluster can exist within a domain or a single tree, even in a multiple-tree forest (as described previously), but all the users for a specific cluster must be contained in the same namespace, as shown in Figure 18-7.

Figure 18-7 Cisco Unified CallManager Clusters Must Be Contained Within a Single Tree (Contiguous Namespace)



The User Search Base is another important element used by Cisco Unified CallManager when integrating with a directory. The User Search Base refers to the root of the subtree used by Cisco Unified CallManager to search for users who can potentially be associated with devices within the cluster. (The section on [Integrating Cisco Unified CallManager with the Directory](#), page 18-19, discusses how to set this parameter.)

You should create a special user account that Cisco Unified CallManager and the other Cisco voice applications can use to access and manage the directory. There should be one account per Cisco Unified CallManager cluster because this arrangement allows each account to be granted specific permissions only when needed and allows for easier administration on a per-cluster basis without the risk of affecting other parts of the enterprise. In the examples in this chapter, this account is called the CCM Directory Manager, but the name you choose for this user account may be different.

Each CCM Directory Manager account should be granted at least the following permissions within your directory:


- **Read/Write/Create all child objects/Delete all child objects** privileges on the respective "Cisco" OU subtrees. The rights must be set to apply to **This object and all child objects** for both the object and the properties. In AD, you can set this privilege by using the advanced options for security within Active Directory Users and Computers (ADUC). The default is to apply to the object only, so you will have to change the rights.

- **Read** privileges on all the OUs contained at and below the User Search Base. You can set this privilege just at the User Search Base level, as long as inheritance is not blocked lower in the tree.
- **Read/Write** privileges on the `ciscoatGUID`, `ciscoatUserProfile`, and `ciscoatUserProfileString` attributes for all User objects contained below the User Search Base. In AD, you can set this privilege by using the advanced options for security within ADUC.

**Tip**

In AD, to set the permissions for the `ciscoatGUID`, `ciscoatUserProfile`, and `ciscoatUserProfileString` attributes for all User objects within the User Search Base, select the CCM Directory Manager user from the Advanced security window for the organizational unit (OU) at the root of the User Search Base. (In this chapter, the user is called CCM Directory Manager, but your user name may be different.) Then click **View/Edit** and go to the **Properties** tab of the new window, as shown in Figure 18-8. From the **Apply onto** drop-down menu, select **User objects**, and then scroll down to the `ciscoatGUID`, `ciscoatUserProfile`, and `ciscoatProfileString` attributes. Allow **Write** permissions for all of them.

Figure 18-8 Setting Permissions for the User Account in Active Directory

**Tip**

While creating the CCM Directory Manager account, set the **Password never expires** option. When you want to change the password, run the CCMPwdChanger utility from Cisco Unified CallManager. This method updates the password in AD, updates the registry in Cisco Unified CallManager, and updates directory initialization files.

Integrating Cisco Unified CallManager with the Directory

After you have prepared the directory as described in the preceding section, you can perform the integration by running the Cisco Customer Directory Configuration plugin again. The two key concepts to consider at this point are the User Search Base and the User Creation Base.

**Note**

The User Creation Base was introduced in Cisco Unified CallManager Release 4.0. In previous versions of Cisco Unified CallManager, the User Search Base is also used for user creation.

As mentioned in the preceding section, the User Search Base parameter represents the root of the subtree used by Cisco Unified CallManager for all user searches.

The User Creation Base parameter tells Cisco Unified CallManager where to create the following system accounts, which are needed by some applications and the features bundled with them:

- CCM Administrator, used by Cisco Unified CallManager Multilevel Administration Access (MLA)
- CCM SysUser, used by CallBack and Extension Mobility
- IPMA SysUser, used by Cisco IP Manager Assistant

The User Creation Base must be contained within the User Search Base because Cisco Unified CallManager has to be able to search for the system account users before authenticating them.

To set the User Search Base, look at where the users serviced by the cluster are located, and set the User Search Base to the first level that includes all of them. The lower you set the User Search Base, the better response times and performance you achieve because searches will not have to follow many referrals and cross slow WAN links to get to remote domain controllers. Also, user data not needed by the cluster does not have to be parsed in the search.

In a single-domain AD forest (or standalone Netscape Directory), set the User Search Base to the lowest organizational unit (OU) that contains all potential users for the Cisco Unified CallManager cluster (for example: ou=AVVID Users, dc=vse, dc=lab). This OU might even be the root of the domain (for example: dc=vse, dc=lab, or o=avvid.lab) if users are spread across a set of OUs directly under this one.

In a multiple-domain AD forest, try to keep the users for a specific Cisco Unified CallManager cluster within a single domain, and follow the guidelines described previously. If a single domain is not possible because users are spread across multiple domains, set the User Search Base to the lowest point in the tree containing all domains with users serviced by the Cisco Unified CallManager cluster. In structures in which serviced child domains are under the top-level domain, the User Search Base must be set at the root of the entire AD forest. In all cases, though, try to ensure that a domain controller for each serviced domain is collocated with Cisco Unified CallManager, or that the network is sufficiently resilient and fast to allow remote searches with no greater performance degradation than occurs with local searches.

**Note**


Although the User Creation Base must exist within the User Search Base, take care to ensure that no existing Active Directory user policies are applied to the system accounts created there. A simple way to protect the system accounts from user policies is to put the accounts in a sub-OU and block inheritance of the Group Policy Object (GPO) to that OU.

It is possible to integrate multiple Cisco Unified CallManager clusters with the same AD forest. However, due to the potentially large number of combinations of customer AD configurations and Cisco voice applications that can be deployed together with Cisco Unified CallManager and that also use the directory, you must request specific support from the Cisco engineering team before proceeding with a multi-cluster integration. Contact your local Cisco account representative to initiate the support request. When deploying Cisco voice applications other than Cisco Unified CallManager (such as the CDR


Analysis and Reporting tool, Multilevel Administration Access, Unified Contact Center Enterprise, Unified Contact Center Express, and so forth), additional limitations may apply. Refer to the respective online documentation and release notes for details on these other products

When integrating multiple Cisco Unified CallManager clusters with the same AD domain (or standalone Netscape Directory), you can share the system accounts across clusters by specifying the same User Creation Base, as shown in [Figure 18-9](#).

Figure 18-9 Setting the User Search Base and User Creation Base in a Single AD Domain



When integrating multiple Cisco Unified CallManager clusters with different AD domains within a forest, Cisco recommends that you define a different User Creation Base for each cluster, setting it to an OU within the relevant domain, as shown in [Figure 18-10](#).

Figure 18-10 Setting the User Search Base and User Creation Base in a Multiple-Cluster, Multiple-Domain Forest**Tip**

You can prevent system and service accounts from appearing in Cisco Unified CallManager Administration by adding the string **CiscoPrivateUser** to the user's Description field. The CCM Administrator, CCM SysUser, and IPMA SysUser accounts have this field set by default, but you can safely add the description to the CCM Directory Manager account as well. Use Microsoft ADSIEdit (Active Directory Service Interfaces, available as a part of the Windows 2000 Support Tools) or any other LDAP tool to update the Description field.

After you have decided how to set the User Search Base and User Creation Base, you can run the Cisco Customer Directory Configuration plugin again on the Cisco Unified CallManager publisher server within your cluster. Follow these best practices when performing this step:

- Select the **Custom** plugin setup type, and select the **Configure Active Directory** and the **Enable CallManager Integration with Active Directory** options only.
- Specify the host name of a domain controller located within the same LAN as Cisco Unified CallManager, or use the domain name if you are using DNS and all domain controllers for this domain are located in the same LAN as Cisco Unified CallManager. Refer to [Planning the Directory Integration, page 18-13](#), for more details on how to provide high availability for the directory integration.

After completing these steps on the publisher server, set the passwords for the three system accounts by using either the CCMPwdChanger tool bundled with Cisco Unified CallManager (open a DOS window by selecting **Start > Run**, type **cmd**, enter **CCMPwdChanger**, and press **Enter**) or your corporate

directory's interface (for example, ADUC in the case of Active Directory). Cisco recommends that you set the password policy for these users so that their passwords never expire and are not set to change on first logon. This password policy is also one reason to avoid having any GPOs applied to these accounts.

If you enforce an expiration policy, Cisco Unified CallManager will stop working when the password expires, and there will be no warning to alert you that the problem is an expired password. If you have policies that require password changes every three months, for example, you should run the CCMPwdChanger tool every three months.

After completing all the preceding steps, you can run the Cisco Customer Directory Configuration plugin on every subscriber server within the Cisco Unified CallManager cluster.

**Note**

When you perform the integration, no data migration occurs between the Cisco Unified CallManager embedded directory and the corporate directory. If you want to migrate users and profiles that you had configured in the embedded directory, Cisco has developed some migration scripts that can help you in this task. To obtain these scripts, contact your Cisco account team or channel partner. Note that the scripts are provided as-is and without support.

Maintaining the Directory Integration

After you have integrated Cisco Unified CallManager with an external directory, you should set the user and password management procedures and policies accordingly.

Use the corporate directory's interface (or supported API) for the following administrative operations:

- Adding a user
- Removing a user
- Setting and changing the core user attributes (such as display name, department, address, and password)

In addition, use Cisco Unified CallManager Administration for the following administrative operations:

- Configuring CallManager-specific user attributes (such as PIN and user locale)
- Associating a user with a device (such as an IP phone or a CTI port)

By default, you cannot use Cisco Unified CallManager Administration to add or remove users. You also cannot modify any of their core user attributes, such as name and phone number. If you want to enable adding and removing users in Cisco Unified CallManager Administration, you can modify the UMDirectoryConfiguration.ini file on the Cisco Unified CallManager servers, as described in *Installing the Cisco Customer Directory Configuration Plugin for Cisco CallManager Release 4.0(1)*, available at

<http://www.cisco.com>

This functionality, provided for your convenience, does not replace your existing user/directory management tools. Be aware that this functionality is limited; Cisco expects that you typically will add or delete users by using other available tools.

Even if Cisco Unified CallManager is integrated with Microsoft Active Directory, you still cannot set or modify user passwords through Cisco Unified CallManager Administration because Active Directory does not allow passwords to be set through clear-text LDAP. To change directory passwords in a secure fashion, you can use either the CCMPwdChanger tool bundled with Cisco Unified CallManager or the management interface provided by the directory vendor.

Even after modifying the UMDirectoryConfiguration.ini file on Cisco Unified CallManager, you still have to give the CCM Directory Manager account sufficient permissions to create and delete users in AD.

When integrating multiple Cisco Unified CallManager clusters with the same directory, keep in mind that you cannot associate the same user with devices in different clusters. Each user must be associated with a single, specific Cisco Unified CallManager cluster at any point in time. (You can, of course, move users from one cluster to another by simply disassociating them from devices in the first cluster and associating them to devices in the second cluster.)

For user-initiated password changes and the setting of preferences, instruct your users to do the following:

- Use the directory application's interface to change their passwords. In the case of Microsoft Active Directory, password changes are done through the users' Windows workstations or by the administrator using the management tools.
- Use the Cisco Unified CallManager User Options web page to change PINs and Cisco Unified CallManager preferences (such as speed dials and call-forward-all numbers).

**Note**

Although the Cisco Unified CallManager User Options web page allows users to change their passwords even after integration with AD, Cisco does not recommend this practice because users probably will not realize that they are changing their Windows passwords at the same time. Also, the communication between the client workstation and the Cisco Unified CallManager server uses HTTP, so the password would travel across the network in clear text. You can remove the **Change your Password** option from the Cisco Unified CallManager User Options web page simply by removing the relevant code from the active server page (ASP).

As far as adds, moves, and changes are concerned, be aware that the following operations are not supported when Cisco Unified CallManager is integrated with the directory:

- Changing a username (in the case of AD, this is the sAMAccountName)
- Moving a user from one OU to another
- Moving or renaming the "Cisco" OU

However, to work around these restrictions, you can manually delete the user's CallManager-specific attributes (such as the profile in the "Cisco" OU and the data in the ciscoatGUID, ciscoatUserProfile, and ciscoatUserProfileString attributes for the user in question). Then the user can be renamed or moved using the directory management tools before being re-added as a Cisco Unified CallManager subscriber. Although it is more cumbersome, this procedure preserves the user's file ownership and security principles in the directory application.

Cisco Unified CallManager Upgrades

Because the schema might change with every major Cisco Unified CallManager release, you should run the Cisco Customer Directory Integration plugin after every upgrade.

If you have integrated multiple Cisco Unified CallManager clusters with the same directory, you need to extend the schema only once. If the clusters are running different Cisco Unified CallManager releases, you should extend the schema from within the cluster that is running the most recent release.



CHAPTER 19

IP Telephony Migration Options

Last revised on: February 13, 2008

This chapter describes the following main methods for migrating to an IP Telephony system (or any other phone system, for that matter):

- [Phased Migration, page 19-1](#)
- [Parallel Cutover, page 19-2](#)

Neither method is right or wrong, and both depend upon individual customer circumstances and preferences to determine which option is most suitable.

This chapter also discusses [The Need for QSIG in Multiple-Site Enterprises, page 19-3](#).

Phased Migration

This approach typically entails a small initial IP Telephony deployment that is connected to the main corporate PBX. The choice of which signaling protocol to use is determined by the required features and functionality as well as by the cost of implementation. Cisco Unified CallManager can support either regular PSTN-type PRI or QSIG PRI as well as H.323 and SIP. Of these options, T1/E1 QSIG provides the highest level of feature transparency between the two systems.

PSTN-type PRI provides for basic call connectivity as well as Automatic Number Identification (ANI). In some instances, the protocol also supports calling name information, as illustrated in [Figure 19-1](#).

The diagram illustrates a Unified Communications Network Architecture. On the left, a **Cisco Unified CM** (represented by a blue box with a white 'M' and a handset icon) is connected to a blue switch. This switch is connected to three IP phones (labeled 'IP'). The switch is also connected to a **Gateway** (represented by a blue circle with a white 'V' and a handset icon). The Gateway is connected to a **PBX** (represented by a blue box with a white handset icon) via a **T1/E1 Trunk or Analog** line. The PBX is connected to three more IP phones. The PBX is also connected to the **PSTN** (Public Switched Telephone Network), represented by a cloud icon. A callout box from the trunk line provides details on the supported protocols:

T1/E1 CAS or Analog	T1/E1 PRI PSTN Protocol
• Voice	• Voice • Calling/Called # • Calling Name (maybe)

114432

Cisco Unified CallManager, Release 3.3 and later, incorporates the International Standards Organization (ISO) variant of QSIG. The QSIG protocol allows for additional feature transparency between PBXs from different vendors, over and above the features that can be obtained from PSTN-type PRI, and it is therefore more appropriate for large enterprises that are already operating complex networks. (Refer to [The Need for QSIG in Multiple-Site Enterprises, page 19-3.](#))

The Cisco San Jose campus, consisting of some 23,000 subscribers housed in approximately 60 buildings, was migrated to IP Telephony in this manner and took just over one year from start to finish. We converted one building per weekend. All subscribers in the selected building were identified, and their extensions were deleted from the PBX on a Friday evening. At the same time, additions were made to the PBX routing tables so that anyone dialing those extension numbers would then be routed over the correct PRI trunk for delivery to Cisco Unified CallManager. During the weekend, new extensions were created in Cisco Unified CallManager for the subscribers, and new IP phones were delivered to their appropriate locations, ready for use by Monday morning. This process was simply repeated for each building until all subscribers had been migrated.

This approach begins with implementation of a complete IP Telephony infrastructure that is redundant, highly available, QoS-enabled, and equipped with Ethernet ports that are powered inline. Once the infrastructure is complete, the IP Telephony application can then be deployed. All IP phones and gateways can be fully configured and deployed so that subscribers have two phones on their desk simultaneously, an IP phone as well as a PBX phone. This approach provides the opportunity to test the

system as well as allowing subscribers time to familiarize themselves with their new IP phones. Outbound-only trunks can also be connected to the IP Telephony system, giving subscribers the opportunity to use their new IP phones to place external calls as well as internal calls.

Once the IP Telephony system is fully deployed, you can select a time slot for bringing the new system into full service by transferring the inbound PSTN trunks from the PBX to the IP Telephony gateways. You can also leave the PBX in place until such a time as you are confident with the operation of the IP Telephony system, at which point the PBX can be decommissioned.

A parallel cutover has the following advantages over a phased migration:


- If something unexpected occurs, the parallel cutover provides a back-out plan that allows you to revert to the PBX system by simply transferring the inbound PSTN trunks from the IP Telephony gateways back to the PBX.
- The parallel cutover allows for verification of the IP Telephony database configuration before the system carries live PSTN traffic. This scenario can be run for any length of time prior to the cutover of the inbound PSTN trunks from the PBX to the IP Telephony gateways, thereby ensuring correct configuration of all subscriber information, phones, gateways, the dial plan, and so forth.
- Training can be carried out at a more relaxed pace by allowing subscribers to explore and use the IP Telephony system at their own leisure before the cutover of the inbound PSTN trunks.
- The system administrator does not have to make special provisions for "communities of interest." With a phased approach, you have to consider maintaining the integrity of call pick-up groups, hunt groups, shared lines, and so forth. These associations can be easily accounted for when moving the complete site in a parallel cutover.

One disadvantage of the parallel cutover is that it requires the IP Telephony solution to be fully funded from the beginning because the entire system must be ready prior to bringing it into service. With a phased migration, on the other hand, you can purchase individual components of the system when they are needed, and this approach does not prevent you from starting with a small trial system prior to moving to full deployment.

The Need for QSIG in Multiple-Site Enterprises

While some enterprises consist of only one location, others consist of many sites, some of which may potentially be spread over large distances. PBX networks for multisite enterprises are usually connected using T1 or E1 PRI trunks (depending on location) running a proprietary protocol such as Avaya DCS, Nortel MCDN, Siemens CorNet, NEC CCIS, Fujitsu FIPN, or Alcatel ABC, among others. These proprietary networking protocols enable the PBXs to deliver a high level of feature transparency between subscribers.

QSIG was developed to enable the interconnection of PBXs from different vendors, thereby allowing similar levels of feature transparency. Cisco first added QSIG to Cisco Unified CallManager Release 3.3 to enable Cisco Unified CallManager to be part of a large enterprise network. (See [Figure 19-2](#).)

Figure 19-2 QSIG Used Between Cisco Unified CallManager and a PBX

QSIG as implemented in Cisco Unified CallManager Release 4.1 supports the following features:

- Basic call
- Direct Inward Dialing (DID)
- Calling number
- Called number
- Connected name
- Transfer (by join)
- Message Waiting Indication (MWI)
- Divert (by forward-switching)
- Calling name restriction
- Calling number restriction
- Divert (by re-route)
- Divert (responding to “check restriction” request)
- Alerting name (on-ringing)
- Path Replacement
- Callback — Call Completion Busy Subscriber (CCBS) and Call Completion No Reply (CCNR)

By supporting QSIG, Cisco Unified CallManager can be introduced into a large enterprise network while also maintaining feature transparency between subscribers. PBX locations can then be converted to IP Telephony whenever convenient.

However, unless you already have QSIG enabled on your PBX or have a specific need for its additional features and functionality, the cost of upgrading the PBX might be hard to justify if it will be retired within a short period of time. For example, why spend \$30,000 on enabling the PBX for QSIG if you plan to retire the PBX in two or three months?

Summary

Although both methods of migration work well and neither method is right or wrong, the parallel cutover method usually works best in most cases. In addition, large enterprises can improve upon either migration method by using QSIG to enable Cisco Unified CallManager to become part of the enterprise network.

Cisco has a lab facility dedicated to testing interoperability between Cisco Unified CallManager and PBX systems. The results of that testing are made available as Application Notes, which are posted at

<http://www.cisco.com/go/interoperability>

The Application Notes are updated frequently, and new documents are continuously added to this website. Check the website often to obtain the latest information.



CHAPTER 20

Voice Security

Last revised on: February 13, 2008

This chapter presents guidelines and recommendations for securing IP Telephony networks. Following the guidelines in this chapter does not guarantee a secure environment, nor will it prevent all penetration attacks on a network. You can achieve reasonable security by establishing a good security policy, following that security policy, staying up-to-date on the latest developments in the hacker and security communities, and maintaining and monitoring all systems with sound system administration practices.

The security guidelines presented in this chapter pertain specifically to IP Telephony technology and the voice network. For more information on data network security, refer to the Cisco SAFE Blueprint documentation available at

http://www.cisco.com/en/US/netsol/ns744/networking_solutions_program_home.html

This chapter addresses centralized but not distributed call processing, and it includes clustering over the WAN but not local failover mechanisms such as Survivable Remote Site Telephony (SRST). This chapter assumes that all remote sites have a redundant link to the head-end or local call-processing backup in case of head-end failure. The interaction between Network Address Translation (NAT) and IP Telephony, for the most part, is not addressed here. This chapter also assumes that all networks are privately addressed and do not contain overlapping IP addresses.

General Security

This section covers general security features and practices that can be used to protect the voice data within a network.

Security Policy

This chapter presumes that your enterprise already has a security policy in place. Cisco Systems recommends that you do not deploy any technology without an associated security policy. The security policy defines which data in your network is sensitive so that it can be protected properly when transported throughout the network. Having this security policy helps you define the security levels required for the types of data traffic that are on your network. Each type of data may or may not require its own security policy.

If no security policy exists for data on the company network, you should create one before enabling any of the security recommendations in this chapter. Without a security policy, there is no way of telling if the security that is enabled in a network is doing what it is designed to accomplish. Without a security policy, there is also no systematic way of enabling security for all the applications and types of data that run in a network.

**Note**

While it is important to adhere to the security guidelines and recommendations presented in this chapter, they alone are not sufficient to constitute a security policy for your company. You must define a corporate security policy before implementing any security technology.

This chapter details the features and functionality of a Cisco Systems network that are available to protect the voice data on a network. It is up to the security policy to define which data to protect, how much protection is needed for that type of data, and which security techniques to use to provide that protection.

One of the more difficult issues with a security policy that includes IP Telephony is combining the security policies that usually exist for both the data network and the traditional voice network. Ensure that all aspects of the integration of the voice data onto the network are secured at the correct level for your security policy or corporate environment.

The basis of a good security policy is defining how important your data is within the network. Once you have ranked the data according to its importance, you can decide how the security levels should be established for each type of data. You can then achieve the correct level of security by using both the network and application features.


In summary, you can use the following process to define a security policy:

- Define the data that is on the network.
- Define the importance of that data.
- Apply security based on the importance of the data.

Security in Layers

This chapter starts at the phone port that a user can plug a PC into, and it works its way through the network from the phone to the access switch, to the distribution layer, into the core, and then into the data center. (See [Figure 20-1](#).) We build layer upon layer of security, starting at the access port into the network itself. With each feature or function that is discussed, we list advantages and disadvantages that need to be taken into account from the standpoint of your corporate security policy.

For example, [Figure 20-1](#) shows both the advantage and disadvantage of using an IP Telephony network. The voice products can be placed anywhere in a network because they use IP to connect to all those devices. This feature gives a network designer the ability to place the devices where it is both physically and logically easy to deploy IP Telephony applications. But because of this ease of deployment, the security complexity increases because the IP Telephony devices can be placed anywhere in a network as long as they have connectivity.

Figure 20-1 **Layers of Security**

Secure Infrastructure

As the IP Telephony data crosses a network, that data is only as safe and secure as the devices that are transporting the data. Depending on the security level that is defined in your security policy, the security of the network devices might have to be improved or they might already be secure enough for the transportation of IP Telephony traffic.

There are many best-practices within a data network that, if used, will increase the entire security of your network. For example, instead of using Telnet (which sends passwords in clear text) to connect to any of the network devices, use Secure Shell (SSH, the secure form of Telnet) so that an attacker would not be able to see a password in clear text. There are many documents on the Cisco.com website that talk about overall security within a network. Use that documentation along with your security policy to determine what security the infrastructure needs.

Video Infrastructure

The Cisco IOS feature sets that provide the gatekeeper functionality (that is, the IP/H323 and EnterprisePlus/H323 MCU feature sets) support only Telnet and not Secure Shell (SSH). Cisco recommends that you use access control lists (ACLs) to control who is permitted to connect to the routers using Telnet, and that you always try to connect to the gatekeeper from a host that is in a secure segment of the network, because user names and passwords are sent over Telnet in clear text.

Cisco Unified Videoconferencing 3500 Series MCUs and H.320 gateways support Telnet, FTP, HTTP, and SNMP. These IP/VC devices do not support TACACS or RADIUS authentication, and only a limited number of administrative accounts can be configured locally in the device. User names and passwords are sent via clear text in all Telnet, FTP, HTTP, and SNMP communications. Cisco recommends that you access these devices from a host that is in a secure segment of the network. You should also use firewalls, access control lists, Cisco Authentication Proxy, and other Cisco security tools to help protect these devices from unauthorized access.

The following links list just a few of the security documents available on Cisco.com:

- Best Practices for Cisco Switches (login authentication required)
http://cisco.com/en/US/partner/products/hw/switches/ps663/products_tech_note09186a0080094713.shtml
- SAFE: A Security Blueprint for Enterprise Networks
http://www.cisco.com/en/US/prod/collateral/wireless/wirelssw/ps1953/product_implementation_design_guide09186a00800a3016.pdf

Physical Security

Just as a traditional PBX is usually locked in a secure environment, the IP network should be treated in a similar way. Each of the devices that carries IP Telephony traffic is really part of an IP PBX, and normal general security practices should be used to control access to those devices. Once a user or attacker has physical access to one of the devices in a network, all kinds of problems could occur. Even if you have excellent password security and the user or attacker cannot get into the network device, that does not mean that they cannot cause havoc in a network by simply unplugging the device and stopping all traffic.

For more information on general security practices, refer to the documentation at the following locations:

- http://www.cisco.com/en/US/netsol/ns744/networking_solutions_program_home.html
- http://www.cisco.com/web/about/ac123/iqmagazine/archives/q2_2005/addressing_network_security.html

IP Addressing

IP addressing can be critical for controlling the data that flows in and out of the logically separated IP Telephony network. The more defined the IP addressing is within a network, the easier it becomes to control the devices on the network.

As stated in other sections of this document (see [Campus Access Layer, page 3-4](#)), you should use IP addressing based on RFC 1918. This method of addressing allows deployment of a IP Telephony system into a network without redoing the IP addressing of the network. Using RFC 1918 also allows for better control in the network because the IP addresses of the voice endpoints are well defined and easy to understand. If the voice endpoints are all addressed within a 10.x.x.x. network, access control lists (ACLs) and tracking of data to and from those devices are simplified.

Advantages

If you have a well defined IP addressing plan for your voice deployments, it becomes easier to write ACLs for controlling the IP Telephony traffic and it also helps with firewall deployments.

Using RFC 1918 enables you easily to deploy one VLAN per switch, which is a best-practise for campus design, and also enables you to keep the Voice VLAN free of any Spanning Tree Protocol (STP) loops.

If deployed correctly, route summarization could help to keep the routing table about the same as before the voice deployment, or just slightly larger.


Disadvantages

Routing tables could become large if not designed correctly or if route summarization is not used.

Phone Security

Cisco Unified IP Phones contain built-in features to increase security on a IP Telephony network. These features can be enabled or disabled on a phone-by-phone basis to increase the security of an IP Telephony deployment. Depending on the placement of the phones, a security policy will help determine if these features need to be enabled and where they should be enabled. (See [Figure 20-2](#).)

Figure 20-2 Security at the Phone Level



Before attempting to configure the security features on a phone, check the documentation at the following link to make sure the features are available on that particular phone model:

http://www.cisco.com/en/US/products/sw/voicesw/ps556/products_administration_guide_chapter09186a00803fe693.html

PC Port on the Phone

The phone has the ability to turn on or turn off the port on the back of the phone, to which a PC would normally be connected. This feature can be used as a control point to access the network if that type of control is necessary.

Depending on the security policy and placement of the phones, the PC port on the back of any given phone might have to be disabled. Disabling this port would prevent a device from plugging into the back of the phone and getting network access through the phone itself. A phone in a common area such as a lobby would typically have its port disabled. Most companies would not want someone to get into the network on a non-controlled port because physical security is very weak in a lobby. Phones in a normal work area might also have their ports disabled if the security policy requires that no device should ever

get access to the network through a phone PC port. Depending on the model of phone deployed, Cisco Unified CallManager can disable the PC port on the back of the phone. Before attempting to enable this feature, check the documentation at the following link to verify that this feature is supported on your particular model of Cisco Unified IP Phone:

http://www.cisco.com/en/US/products/hw/phones/ps379/tsd_products_support_series_home.html

Advantages

Disabling the phone PC port allows for phones to be placed in areas where network access from the phone should not be allowed. It controls access to the network that otherwise would be there if the PC port on the back of the phone was enabled.

Disadvantages

For each person who needs to have network access and is approved for access, a separate Ethernet port would be required to provide that person with network access if the PC port on the phone is disabled. A person could still unplug the ethernet jack from the phone and attempt to plug it into another device.

For Cisco Unified Video Advantage to operate properly, the PC port and Video Capabilities must both be enabled. The other settings can safely be disabled.

Gratuitous ARP

Just like any other data device on the network, the phones are vulnerable to traditional data attacks. The phones have features to prevent some of the common data attacks that can occur on a corporate network. One such feature is Gratuitous ARP (Gratuitous Address Resolution Protocol, or GARP). This feature helps to prevent man-in-the-middle (MITM) attacks to the phone. A MITM attack involves an attacker who tricks an end station into believing that he is the router and tricks the router into believing that he is the end station. This scheme makes all the traffic between the router and the end station travel through the attacker, thus enabling the attacker to log all of the traffic or inject new traffic into the data conversation.

Gratuitous ARP helps protect the phones from having an attacker capture the signaling and RTP voice streams from the phone if the attacker was able to get onto the voice segment of the network. This feature protects only the phones; it does not protect the rest of the infrastructure from a Gratuitous ARP attack. This feature is of less importance if you are running a Cisco infrastructure because the switch port provides features that protect both the phones and the network gear. For a description of these switch port features see the section on [Switch Port](#), page 20-12.


Advantages

The Gratuitous ARP feature protects the phone from a traditional MITM attack on the signaling and RTP voice streams that are sourced from the phone to the network.

Disadvantages

The downstream signaling and RTP voice streams coming from another phone or coming across the network are not protected by this feature in the phone. Only the data coming from the phone that has this feature enabled is protected. (See [Figure 20-3](#).)

If the default gateway is running Hot Standby Router Protocol (HSRP), if the HSRP configuration uses the burned-in MAC address rather than the virtual MAC address for the default gateway, and if the primary router fails-over to a secondary router that has a new MAC address, the phones could maintain the old MAC address of the default gateway. This scenario could cause an outage for up to 40 minutes. Always use the virtual MAC address in an HSRP environment to avoid this potential problem.

Figure 20-3 Gratuitous ARP Protects the Phone that Has It but Not Other Traffic

As shown in Figure 20-3, the traffic from the phone that has Gratuitous ARP is protected, but the attacker could still see the traffic coming from another endpoint because that endpoint might not have the ability to protect the data flow.


PC Voice VLAN Access

Because there are two VLANs from the switch to the phone, the phone needs to protect the voice VLAN from any unwanted access. The phones can prevent unwanted access into the voice VLAN from the back of the phone. A feature called PC Voice VLAN Access prevents any access to the voice VLAN from the PC port on the back of the phone. When disabled, this feature does not allow the devices plugged into the PC port on the phone to "jump" VLANs and get onto the voice VLAN by sending 802.1q tagged information destined for the voice VLAN to the PC port on the back of the phone. The feature operates one of two ways, depending on the phone that is being configured. On the more advanced phones, the phone will block any traffic destined for the voice VLAN that is sent into the PC port on the back of the phone. In the example shown in Figure 20-4, if the PC tries to send any voice VLAN traffic (with an 802.1q tag of 200 in this case) to the PC port on the phone, that traffic will be blocked. The other way this feature can operate is to block all traffic with an 802.1q tag (not just voice VLAN traffic) that comes into the PC port on the phone.

Currently, 802.1q tagging from an access port is not normally used. If that feature is a requirement for the PC plugged into the port on the phone, you should use a phone that allows 802.1q tagged packets to pass through the phone.

Before attempting to configure the PC Voice VLAN Access feature on a phone, check the documentation at the following link to make sure the feature is available on that particular phone model:

http://www.cisco.com/en/US/products/hw/phones/ps379/tsd_products_support_series_home.html

Figure 20-4 Blocking Traffic to the Voice VLAN from the Phone PC Port**Advantages**

The PC Voice VLAN Access feature prevents attackers from sending uncontrolled data into the voice VLAN through the PC port on the back of the phone.

Disadvantages

If the device plugged into the phone is normally allowed to send 802.1q tagged packets, then these packets will be dropped. Most end stations are not allowed to perform this function at the access layer. If that function is considered normal operation within a network, this feature would not allow that function to work.

Web Access

Each Cisco Unified IP Phone has a web server built into it to help with debugging and remote status of the phone for management purposes. The web server also enables the phones to receive applications pushed from Cisco Unified CallManager to the phones. Access to this web server can be enabled or disabled on a phone by means of the Web Access feature in the Cisco Unified CallManager configuration. This setting can be global, or it could be enabled or disabled on a phone-by-phone basis.

Advantages

With Web Access enabled on the phones, the phones can be used to assist in debugging issues with a phone or within the network. If Web Access from the phone is disabled, users or an attacker cannot get information from the phone about the IP Telephony network.

Disadvantages

If Web Access from the phone is disabled, debugging network or IP Telephony issues can be more difficult. If the web server is globally disable but it is needed to help with debugging, then the administrator for Cisco Unified CallManager will have to enable this feature on the phones. The ability to get to this web page can be controlled by an ACL in the network, leaving network operators with the capability to get to the web page when needed.

With the Web Access feature disabled, the phones will be unable to receive applications pushed to them from Cisco Unified CallManager.

Video Capabilities

For Cisco Unified Video Advantage to operate properly, the PC port and Video Capabilities must both be enabled. The other settings can safely be disabled. The Device Security Mode operates as specified, even when Cisco Unified Video Advantage is in use, but Cisco Unified Video Advantage itself does not support authentication or encryption of the Cisco Audio Session Tunnel (CAST) protocol or its RTP media traffic. When an IP Phone is in Authenticated mode, the Skinny Client Control Protocol (SCCP) signaling between the phone and Cisco Unified CallManager will be authenticated, but the CAST signaling between the phone and Cisco Unified Video Advantage will not be authenticated. Likewise, when a phone is in Encrypted mode, the audio stream between the phones will be encrypted, but the video streams between the Cisco Unified Video Advantage clients will not be encrypted. Users should be notified that the video channel is not encrypted even though an icon on the phone appears to indicate that they are in an encrypted call.

Advantages

The PC port and Video Capabilities are required for Cisco Unified Video Advantage to function correctly.

Disadvantages

Enabling these features could possibly allow communication to the phone from the PC if ACLs are not used to protect the phone.

Settings Access

Each Cisco Unified IP Phone has a network settings page that lists many of the network elements and detailed information that is needed for the phone to operate. This information could be used by an attacker to start a reconnaissance on the network with some of the information that is displayed on the phone's web page. For example, an attacker could look at the settings page to determine the default gateway, the TFTP server, and the Cisco Unified CallManager IP address. Each of these pieces of information could be used to gain access to the voice network or to attack a device in the voice network.

This access can be disabled on a phone-by-phone basis (see [Figure 20-5](#)) to prevent end users or attackers from obtaining the additional information such as Cisco Unified CallManager IP address and TFTP server information.

For more information on the phone settings page, refer to *Cisco Unified IP Phone Authentication and Encryption for Cisco Unified CallManager*, available at

http://www.cisco.com/en/US/products/sw/voicesw/ps556/prod_maintenance_guides_list.html

Figure 20-5 Phone Configuration Page in Cisco Unified CallManager

Product Specific Configuration	
Disable Speakerphone	☐
Disable Speakerphone and Headset	☐
Forwarding Delay*	Disabled
PC Port*	Disabled
Settings Access*	Disabled
Gratuitous ARP*	Disabled
PC Voice VLAN Access*	Disabled
Video Capabilities*	Disabled
Auto Line Select*	Disabled
Web Access*	Disabled

148492

Advantages

With access to the phone settings page disabled, end users and potential attackers are not able to see detailed information about the network and the IP Telephony information used by the voice system. With this feature disabled, some of the information that would be protected includes the IP address of the phone and the Cisco Unified CallManager to which the phone is registered.

Disadvantages

With access to the phone settings page disabled, end users lose the ability to change many of the settings on the phone that they would normally be able to control, such as speaker volume, contrast, and ring type. It might not be practical to use this security feature because of the limitations it places on end users with respect to the phone interface. However, access to the phone settings page does not have to be lost if the administrator chooses to restrict access rather than disable it.

Phone Authentication and Encryption

Cisco Unified CallManager can be configured to provide multiple levels of security to the phones within a voice system, if those phones support those features. Depending on your security policy, phone placement, and phone support, the security can be configured to fit the needs of your company.

For information on which Cisco Unified IP Phone models support specific security features, refer to the documentation available at

http://www.cisco.com/en/US/products/hw/phones/ps379/tsd_products_support_series_home.html

To enable security on the phones and in the Cisco Unified CallManager cluster, refer to the *Cisco Unified CallManager Security Guide*, available at

http://www.cisco.com/en/US/products/sw/voicesw/ps556/prod_maintenance_guides_list.html

Advantages

When the security features are properly configured in Cisco Unified CallManager, all supported phones will have the following capabilities:

- Integrity — Does not allow TFTP file manipulation but does allow Transport Layer Security (TLS) signaling to the phones when enabled.
- Authentication — The image for the phone is authenticated from Cisco Unified CallManager to the phone, and the device (phone) is authenticated to Cisco Unified CallManager. All signaling messages between the phone and Cisco Unified CallManager are verified as being sent from the authorized device.
- Encryption — For supported devices, signaling and media can be encrypted to prevent eavesdropping.
- Secure Real-time Transport Protocol (SRTP) — Is supported to Cisco IOS MGCP gateways and, of course, phone-to-phone. Cisco Unity also supports SRTP for voicemail.

Disadvantages

Cisco Unified CallManager supports authentication, integrity, and encryption for calls between two Cisco Unified IP Phones within a single cluster where no media services are used. However, Cisco Unified CallManager does not provide authentication, integrity, or encryption for all devices or phones. To determine if your device supports these features, refer to the documentation available at

http://www.cisco.com/en/US/products/hw/phones/ps379/tsd_products_support_series_home.html

In addition, auto-registration does not work if you configure the cluster for mixed mode, which is required for device authentication. You cannot implement signaling or media encryption if device authentication does not exist in the cluster – that is, if you do not install and configure the Cisco Certificate Trust List (CTL) client. Application Layer Gateways (ALGs) that allow IP Telephony traffic to traverse firewalls and Network Address Translation (NAT) also do not work with signaling encryption. Not all gateways, phones, or conference are supported with encrypted media.

Access Security

This section covers security features at the Access level that can be used to protect the voice data within a network.

Voice and Video VLANs

Before the phone has its IP address, the phone determines which VLAN it should be in by means of the Cisco Discovery Protocol (CDP) negotiation that takes place between the phone and the switch. This negotiation allows the phone to send packets with 802.1q tags to the switch in a "voice VLAN" so that the voice data and all other data coming from the PC behind the phone are separated from each other at Layer 2. Voice VLANs are not required for the phones to operate, but they provide additional separation from other data on the network.

Cisco Unified Video Advantage is a client application running on a PC, but it is also associated with an IP Phone. The PC will likely reside in the data VLAN while the phone will likely reside in the voice VLAN. To associate with the IP Phone, Cisco Unified Video Advantage uses the Cisco Audio Session Tunnel (CAST) protocol, which operates over TCP/IP. Therefore, Cisco Unified Video Advantage will have to communicate through whatever Layer 3 router is configured to route IP packets between the voice and data VLANs. If there are any access control lists or firewalls configured between those VLANs, they will have to be modified to permit the CAST protocol to operate. Fortunately, CAST uses

TCP port 4224 in both directions, making this task easier. Cisco Unified Video Advantage communicates with the IP Phone but not with Cisco Unified CallManager, except when Cisco Unified Video Advantage periodically checks with the TFTP server (which could be co-resident on one or more of the Cisco Unified CallManager servers) to download any software updates. Therefore, you must also permit the TFTP protocol between the data VLAN and the TFTP server(s).

Sony and Tandberg SCCP endpoints do not support Cisco Discovery Protocol (CDP) or 802.1Q VLAN ID tagging. Therefore, in a typical environment, these devices will reside in the data VLAN, unless the port has been configured manually to use the voice VLAN as the native VLAN. The Sony and Tandberg endpoints communicate with the TFTP server to download their configurations, with Cisco Unified CallManager for SCCP signaling, and with other endpoints for RTP audio/video media channels. Therefore, the TFTP protocol must be permitted between the data VLAN and the TFTP server(s), TCP port 2000 must be permitted between the data VLAN and the Cisco Unified CallManager server(s), and UDP ports for RTP media must be permitted between the data and voice VLANs.

H.323 clients, Multipoint Control Units (MCUs), and gateways communicate with Cisco Unified CallManager using the H.323 protocol. Cisco Unified CallManager H.323 trunks (such as H.225 and intercluster trunk variants as well as the RASAggregator trunk type) use a random port range rather than the well-known TCP port 1720. Therefore, you must permit a wide range of TCP ports between these devices and the Cisco Unified CallManager servers. MCUs and gateways are considered infrastructure devices, and they typically reside within the datacenter adjacent to the Cisco Unified CallManager servers. H.323 clients, on the other hand, typically reside in the data VLAN.

Cisco Unified Videoconferencing 3500 Series MCUs configured to run in SCCP mode communicate with the TFTP server(s) to download their configuration, with the Cisco Unified CallManager servers for signaling, and with other endpoints for RTP media traffic. Therefore, TFTP must be permitted between the MCU and the TFTP server(s), TCP port 2000 must be permitted between the MCUs and the Cisco Unified CallManager server(s), and UDP ports for RTP media must be permitted between the MCUs and the voice, data, and gateway VLANs.

Advantages

Voice VLANs can be assigned automatically from the switch to the phone, thus allowing for Layer 2 and Layer 3 separations between voice data and all other data on a network. A voice VLAN also allows for a different IP addressing scheme because the separate VLAN can have a separate IP scope at the Dynamic Host Configuration Protocol (DHCP) server.

Applications use the CDP messaging from the phones to assist in locating phones during an emergency phone call. The location of the phone will be much more difficult to determine if CDP is not enabled on the access port to which that phone is attached.

Disadvantages

There is a possibility that information could be gathered from the CDP messaging that would normally go to the phone, and that information could be used to discover some of the network. Not all devices that can be used for voice or video with Cisco Callmanager are able to use CDP to assist in discovering the voice VLAN.


Switch Port

There are many security features within a Cisco switch infrastructure that can be used to secure a data network. This section describes some of the features that can be used in Cisco Access Switches to protect the IP Telephony data within a network. (See [Figure 20-6](#).) This section does not cover all of the security features available for all of the current Cisco switches, but it does list the most common security feature

used across many of the switches that Cisco manufactures. For additional information on the security features available on the particular Cisco gear deployed within your network, refer to the appropriate product documentation available at

<http://www.cisco.com>

Figure 20-6 A Typical Access Layer Design to Which the Phones Attach



Port Security: MAC CAM Flooding

A classic attack on a switched network is a MAC content-addressable memory (CAM) flooding attack. This type of attack floods the switch with so many MAC addresses that the switch does not know which port an end station or device is attached to. When the switch does not know which port a device is attached to, it broadcasts the traffic destined for that device to the entire VLAN. In this way, the attacker is able to see all traffic that is coming to all the users in a VLAN.

To disallow malicious MAC flooding attacks from hacker tools such as macof, limit the number of MAC addresses allowed to access individual ports based on the connectivity requirements for those ports. Malicious end-user stations can use macof to originate MAC flooding from random-source to random-destination MAC addresses, both directly connected to the switchport or through the IP phone. The macof tool is very aggressive and typically can fill a Cisco Catalyst switch content-addressable memory (CAM) table in less than ten seconds. The flooding of subsequent packets that remain unlearned because the CAM table is filled, is as disruptive and insecure as packets on a shared Ethernet hub for the VLAN that is being attacked.

Either port security or dynamic port security can be used to inhibit a MAC flooding attack. A customer with no requirement to use port security as an authorization mechanism would want to use dynamic port security with the number of MAC addresses appropriate to the function attached to a particular port. For example, a port with only a workstation attached to it would want to limit the number of learned MAC addresses to one. A port with a Cisco Unified IP Phone and a workstation behind it would want to set the number of learned MAC addresses to two (one for the IP phone itself and one for the workstation behind the phone) if a workstation is going to plug into the PC port on the phone. This setting in the past has been three MAC addresses, used with the older way of configuring the port in trunk mode. If you use the multi-VLAN access mode of configuration for the phone port, this setting will be two MAC addresses, one for the phone and one for the PC plugged into the phone. If there will be no workstation on the PC port, then the number of MAC addresses on that port should be set to one. These configurations are for a multi-VLAN access port on a switch. The configuration could be different if the port is set to trunk mode (not the recommended deployment of an access port with a phone and PC).

Port Security: Prevent Port Access

Prevent all port access except from those devices designated by their MAC addresses to be on the port. This is a form of device-level security authorization. This requirement is used to authorize access to the network by using the single credential of the device's MAC address. By using port security (in its non-dynamic form), a network administrator would be required to associate MAC addresses statically for every port. However, with dynamic port security, network administrators can merely specify the number of MAC addresses they would like the switch to learn and, assuming the correct devices are the first devices to connect to the port, allow only those devices access to that port for some period of time.

The period of time can be determined by either a fixed timer or an inactivity timer (non-persistent access), or it can be permanently assigned. The feature to permanently assign a MAC address on a Cisco 6000 switch is called *autoconfigure*, and on the Cisco 4500, 2550, 2750, or 2950 switches the feature is called *sticky*. In both cases, the MAC address learned will remain on the port even in the event of a reload or reboot of the switch.


Persistent assignment of MAC addresses via *autoconfigure* or *sticky* can be cleared only by a command. The most common default behavior seen across the Cisco Catalyst switching platforms currently is the non-persistent behavior; the only behavior prior to Cisco CatOS Release 7.6(1) was persistent. No provision is made for device mobility by static port security or persistent dynamic port security. Although it is not the primary requirement, MAC flooding attacks are implicitly prevented by port security configurations that aim to limit access to certain MAC addresses.

From a security perspective, there are better mechanisms for both authenticating and authorizing port access based on userid and/or password credentials rather than using MAC address authorization. MAC addresses alone can easily be spoofed or falsified by most operating systems.

Port Security: Prevent Rogue Network Extensions

Prevent rogue network extensions via hub or wireless access points (APs). Because it limits the number of MAC addresses to a port, port security can also be used as a mechanism to inhibit user extension to the IT-created network. For example, if a user plugs a wireless AP into a user-facing port or data port on a phone with port security defined for a single MAC address, the wireless AP itself would occupy that MAC address and not allow any devices behind it to access the network. (See [Figure 20-7](#).) Generally, a configuration appropriate to stop MAC flooding is also appropriate to inhibit rogue access.

Figure 20-7 Limited Number of MAC Addresses Prevents Rogue Network Extensions



Advantages

Port security prevents an attacker from flooding the CAM table of a switch and from turning any VLAN into a hub that transmits all received traffic to all ports. It also prevents unapproved extensions of the network by adding hubs or switches into the network.

Disadvantages

If the number of MAC addresses is not defined correctly, there is a possibility of denying access to the network or error-disabling the port and removing all devices from the network.

Configuration Example



Note This example configuration is based on a switch running correct code levels to support these features and not running trunk mode to the phone.

The following example illustrates the Cisco IOS commands to configure an access port with dynamic port security, running a phone with a device plugged into the data port on the phone:

```
switchport access vlan 10
switchport mode access
switchport voice vlan 20
switchport port-security
switchport port-security maximum 2
switchport port-security violation restrict
switchport port-security aging time 2
switchport port-security aging type inactivity
```

The commands in the preceding example perform the following functions:

- **switchport port-security *x/x* enable**

This command enables the port security on the specified module/port.

- **switchport port-security violation restrict**

This command is the recommended configuration. The default is to disable the port. If you **restrict** the port, it will learn up to the maximum number of MAC addresses and then stop learning any new MAC addresses. If you have the port on the default setting of **disable** and the maximum number of MAC addresses is exceeded, the port will error-disable and turn off power to the phone. The default timer for the port to re-enable is 5 minutes. Depending on your security policy, it might be preferable to restrict the port and not shut down the phone by disabling the port.

- **switchport port-security aging time 2**

This command sets the amount of time that the MAC address will remain on the port without any traffic from that MAC address. Because of CDP communication between some switches and the phone, the recommended minimum time is 2 minutes.

- **switchport port-security aging type inactivity**

This command defines the type of aging that will be used on the port to time-out the learned MAC address.

DHCP Snooping: Prevent Rogue DHCP Server Attacks


Dynamic Host Configuration Protocol (DHCP) Snooping prevents a non-approved DHCP or rogue DHCP server from handing out IP addresses on a network by blocking all replies to a DHCP request unless that port is allowed to reply. Because most phone deployments use DHCP to provide IP addresses

to the phones, you should use the DHCP Snooping feature in the switches to secure DHCP messaging. Rouge DHCP servers can attempt to respond to the broadcast messages from a client to give out incorrect IP addresses, or they can attempt to confuse the client that is requesting an address.


When enabled, DHCP Snooping treats all ports in a VLAN as untrusted by default. An untrusted port is a user-facing port that should never make any reserved DHCP responses. If an untrusted DHCP-snooping port makes a DHCP server response, it will be blocked from responding. Therefore, rogue DHCP servers will be prevented from responding. However, legitimately attached DHCP servers or uplinks to legitimate servers must be trusted.

Figure 20-8 illustrates the normal operation of a network-attached device that requests an IP address from the DHCP server.

Figure 20-8 Normal Operation of a DHCP Request



However, an attacker can request not just a single IP address but all of the IP addresses that are available within a VLAN. (See Figure 20-9.) This means that there would be no addresses for a legitimate device trying to get on the network, and without an IP address the phone cannot connect to Cisco Unified CallManager.

Figure 20-9 An Attacker Can Take All Available IP Addresses on the VLAN**Advantaged**


DHCP Snooping prevents unapproved DHCP servers from being on a network.

Disadvantages

Incorrect configurations of this feature can deny IP addresses to approved users.

DHCP Snooping: Prevent DHCP Starvation Attacks

DHCP address scope starvation attacks from tools such as Gobbler are used to create a DHCP denial-of-service (DoS) attack. Because the Gobbler tool makes DHCP requests from different random source MAC addresses, you can prevent it from starving a DHCP address space by using port security to limit the number of MAC addresses. (See [Figure 20-10](#).) However, a more sophisticated DHCP starvation tool can make the DHCP requests from a single source MAC address and vary the DHCP payload information. With DHCP Snooping enabled, untrusted ports will make a comparison of the source MAC address to the DHCP payload information and fail the request if they do not match.

Figure 20-10 Using DHCP Snooping to Prevent DHCP Starvation Attacks**Advantages**

DHCP Snooping prevents any single device from capturing all the IP addresses in any given scope.

Disadvantages

Incorrect configurations of this feature can deny IP addresses to approved users.

Configuration Example

The following example illustrates the Cisco IOS commands to configure an access port with DHCP Snooping, running a phone with a device plugged into the data port on the phone:

- Global commands


```
ip dhcp snooping vlan 10, 20
no ip dhcp snooping information option
ip dhcp snooping
```
- Interface commands


```
no ip dhcp snooping trust (Default)
ip dhcp snooping limit rate 10 (pps)
ip dhcp snooping trust
```

The global commands in the preceding example perform the following functions:

- **ip dhcp snooping vlan 10, 20**
This command specifies which VLANs have DHCP Snooping enabled.
- **No ip dhcp snooping information option**
This command should be used so that Option 82 information is not required to lease a DHCP address. The Option 82 information must be supported by the DHCP server, but most enterprise servers do not support this feature. Option 82 is supported in Cisco IOS DHCP servers.
- **ip dhcp snooping**
This command enables DHCP Snooping at the global level on the switches.

The interface commands in the preceding example perform the following functions:

- **no ip dhcp snooping trust**
This command sets the interface to not trust any information coming into the port from a DHCP server.

- **ip dhcp snooping limit rate 10**

This command sets the default rate limit that is configured on the interface when DHCP Snooping is initially configured. This value can be changed depending on your security policy.

- **ip dhcp snooping trust**

You must use this command on the port through which DHCP information will be sent from a DHCP server. If you do not trust the port from which the DHCP information is coming, then none of the devices will ever receive a DHCP address. At least one port (access port or trunk port) with the DHCP server on it must be configured to allow this information to get to the clients. This command can also be used to trust any device connected to a port that has a static IP address and that will not use DHCP to get an IP address. Note that the uplink port to the DHCP server, or the trunk port to the DHCP server, will also have to be trusted.

DHCP Snooping: Binding Information

Another function of DHCP Snooping is to record the DHCP binding information for untrusted ports that successfully get IP addresses from the DHCP servers. The binding information is recorded in a table on the Cisco Catalyst switch. The DHCP binding table contains the IP address, MAC address, lease length, port, and VLAN information for each binding entry. The binding information from DHCP Snooping remains in effect for the length of the DHCP binding period set by the DHCP server (that is, the DHCP lease time). The DHCP binding information is used to create dynamic entries for Dynamic ARP Inspection (DAI) to limit ARP responses for only those addresses that are DHCP-bound. The DHCP binding information is also used by the IP source guard to limit sourcing of IP packets to only those addresses that are DHCP-bound.

The following examples show binding information from DHCP Snooping.

- Displaying binding information for Cisco IOS:

```
show ip dhcp snooping binding
-----
MacAddress      IpAddress      Lease(sec)      Type            VLAN  Interface
-----
00:03:47:B5:9F:AD  10.120.4.10    193185          dhcp-snooping   10    FastEthernet3/18
```

- Displaying binding information for Cisco CatOS:

```
ngcs-6500-1> (enable) show dhcp-snooping bindings
MacAddress      IpAddress      Lease(sec)      VLAN      Port
-----
00-10-a4-92-bf-dd  10.10.10.21    41303           10        2/5
```

There is a maximum limit to the number of binding table entries that each type of switch can store for DHCP Snooping. (Refer to the product documentation for your switch to determine this limit.) If you are concerned about the number of entries in your switch's binding table, you can reduce the lease time on the DHCP scope so that the entries in the binding table time-out sooner. The entries remain in the DHCP binding table until the lease runs out. In other words, the entries remain in the DHCP Snooping binding table as long as the DHCP server thinks the end station has that address. They are not removed from the port when the workstation or phone is unplugged.

If you have a Cisco Unified IP Phone plugged into a port and then move it to a different port, you might have two entries in the DHCP binding table with the same MAC and IP address on different ports. This behavior is considered normal operation.

Requirement for Dynamic ARP Inspection

Dynamic Address Resolution Protocol (ARP) Inspection (DAI) is a feature used on the switch to prevent Gratuitous ARP attacks on the devices plugged into the switch and on the router. Although it is similar to the Gratuitous ARP feature mentioned previously for the phones, Dynamic ARP protects all the devices on the LAN, and it is not just a phone feature.

In its most basic function, Address Resolution Protocol (ARP) enables a station to bind a MAC address to an IP address in an ARP cache, so that the two stations can communicate on a LAN segment. A station sends out an ARP request as a MAC broadcast. The station that owns the IP address in that request will give an ARP response (with its IP and MAC address) to the requesting station. The requesting station will cache the response in its ARP cache, which has a limited lifetime. The default ARP cache lifetime for Microsoft Windows is 2 minutes; for Linux, the default lifetime is 30 seconds; and for Cisco IP phones, the default lifetime is 40 minutes.


ARP also makes the provision for a function called Gratuitous ARP. Gratuitous ARP (GARP) is an unsolicited ARP reply. In its normal usage, it is sent as a MAC broadcast. All stations on a LAN segment that receive a GARP message will cache this unsolicited ARP reply, which acknowledges the sender as the owner of the IP address contained in the GARP message. Gratuitous ARP has a legitimate use for a station that needs to take over an address for another station on failure.

However, Gratuitous ARP can also be exploited by malicious programs that want to illegitimately take on the identity of another station. When a malicious station redirects traffic to itself from two other stations that were talking to each other, the hacker who sent the GARP messages becomes the man-in-the-middle. Hacker programs such as ettercap do this with precision by issuing "private" GARP messages to specific MAC addresses rather than broadcasting them. In this way, the victim of the attack does not see the GARP packet for its own address. Ettercap also keeps its ARP poisoning in effect by repeatedly sending the private GARP messages every 30 seconds.

Dynamic ARP Inspection (DAI) is used to inspect all ARP requests and replies (gratuitous or non-gratuitous) coming from untrusted (or user-facing) ports to ensure that they belong to the ARP owner. The ARP owner is the port that has a DHCP binding which matches the IP address contained in the ARP reply. ARP packets from a DAI trusted port are not inspected and are bridged to their respective VLANs.

Using DAI

Dynamic ARP Inspection (DAI) requires that a DHCP binding be present to legitimize ARP responses or Gratuitous ARP messages. If a host does not use DHCP to obtain its address, it must either be trusted or an ARP inspection access control list (ACL) must be created to map the host's IP and MAC address. (See [Figure 20-11](#).) Like DHCP Snooping, DAI is enabled per VLAN, with all ports defined as untrusted by default. To leverage the binding information from DHCP Snooping, DAI requires that DHCP Snooping be enabled on the VLAN prior to enabling DAI. If DHCP Snooping is not enabled before you enable DAI, none of the devices in that VLAN will be able to use ARP to connect to any other device in their VLAN, including the default gateway. The result will be a self-imposed denial of service to any device in that VLAN.

Figure 20-11 Using DHCP Snooping and DAI to Block ARP Attacks

Because of the importance of the DHCP Snooping binding table to the use of DAI, it is important to back up the binding table. The DHCP Snooping binding table can be backed up to bootflash, File Transfer Protocol (FTP), Remote Copy Protocol (RCP), slot0, and Trivial File Transfer Protocol (TFTP). If the DHCP Snooping binding table is not backed up, the Cisco Unified IP Phones could lose contact with the default gateway during a switch reboot. For example, assume that the DHCP Snooping binding table is not backed up and that you are using Cisco Unified IP Phones with a power adapter instead of line power. When the switch comes back up after a reboot, there will be no DHCP Snooping binding table entry for the phone, and the phone will not be able to communicate with the default gateway unless the DHCP Snooping binding table is backed up and loads the old information before traffic starts to flow from the phone.

Advantages

DAI prevents an attacker from running ARP-based attacks in a network to disrupt or sniff the traffic between people who are adjacent to the attacker at Layer 2.

Disadvantages

Incorrect configurations of this feature can deny network access to approved users. If a device has no entry in the DHCP Snooping binding table, then that device will not be able to use ARP to connect to the default gateway and therefore will not be able to send traffic. If you use static IP addresses, those addresses will have to be entered manually into the DHCP Snooping binding table. If you have devices that do not use DHCP again to obtain their IP addresses when a link goes down (some UNIX or Linux machines behave this way), then you must back up the DHCP Snooping binding table.

Configuration Example

The following example illustrates the Cisco IOS commands to configure access ports with DHCP Snooping and Dynamic ARP Inspection:

- Global commands


```
ip dhcp snooping vlan 10,20 (required)
no ip dhcp snooping information option (required without option 82 dhcp server)
ip dhcp snooping (required)
ip arp inspection vlan 10,20
ip dhcp snooping database tftp://172.26.168.10/tftpboot/cisco/ngcs-dhcpdb
```

- Interface commands

```
ip dhcp snooping trust
ip arp inspection trust
no ip arp inspection trust (default)
ip arp inspection limit rate 15 (pps)
```

The global commands in the preceding example perform the following functions:

- **ip arp inspection vlan 10,20**

This command specifies which VLANs have Dynamic ARP Inspection (DAI) enabled.

- **ip arp inspection trust**

As with **ip dhcp snooping trust**, this command allows a trusted device such as a router to reply to ARP messages. This command must be configured on the port for your router, otherwise the router will not be able to respond to any ARP requests because the router will not be in the DHCP Snooping binding table.

- **no ip arp inspection trust**

This is the default setting for every port in the VLAN. Trust must be enabled.

- **ip arp inspection limit rate 15 (pps)**

This command sets the global default value for the maximum number of packets per second (pps) that are allow for ARP messages on the interfaces. If this value is exceeded, then the interface will be disabled. If this behavior is a concern, you can increase or decrease the limit or set it to **none**.

- **ip dhcp snooping database tftp://172.26.168.10/tftpboot/cisco/ngcs-dhcpdb**

This command backs up the DHCP Snooping binding table to the TFTP server. The DHCP Snooping binding table can be backed up to bootflash, FTP, RCP, slot0, and TFTP.

The interface commands in the preceding example perform the following functions:

- **no ip arp inspection trust**

This command enables DAI on the port and checks all ARPs against the DHCP Snooping binding table.

- **ip arp inspection limit rate 15 (pps)**

This command specifies the maximum number of packets per second (pps) that are allow for ARP messages on the interface. If the interface sees more than the specified number of ARP messages in a second, it will disable the port. Depending on your security policy, the default value (15 pps) might be the preferred setting. If you do not want to disable the phone when the port receives more than 15 ARP messages in a second, you can set the rate limit to **none**, which will allow the phone to stay up.


IP Source Guard

Beyond ARP spoofing, an attacker can also do IP address spoofing. This method is commonly use to perform DoS attacks on a second party by sending packets through a third party, thus masking the identity of the attacking system. A simple example of this occurs when an attacker pings a third-party system while sourcing the IP address of the second party that is being attacked. The ping response will be directed to the second party from the third-party system. Aggressive SYN-flooding originating from spoofed IP addresses is another common type of attack used to overwhelm a server with TCP half-sessions.

The IP Source Guard (IPSG) feature, when invoked, dynamically creates an ACL based on the contents of the DHCP Snooping binding table. This ACL ensures that traffic is sourced from the IP address issued at DHCP binding time and prevents any traffic from being forwarded by other spoofed addresses. While

DHCP Snooping is a prerequisite for IP Source Guard, DAI is not. However, Cisco recommend that you enable DAI in addition to IP Source Guard to prevent ARP-poisoning man-in-the-middle attacks in addition to IP address spoofing. (See [Figure 20-12](#).)

Figure 20-12 Using IP Source Guard to Prevent Address Spoofing



With IP address spoofing, an attacker can impersonate a valid address either by manually changing an address or by running a program design to do address spoofing, such as hping2. Internet worms can use spoofing techniques to disguise their origins.

Configuration Example

The following example illustrates the Cisco IOS commands to configure access ports with IP Source Guard:

- Commands that must be enabled before you enable IP Source Guard

```
ip dhcp snooping vlan 4,104
no ip dhcp snooping information option
ip dhcp snooping
```

- Interface command — This command enables the IP Source Guard without DHCP Option 82.

```
ip verify source vlan dhcp-snooping
```

Additional Information

For additional information on network security, refer to the Cisco documentation at the following locations:

- http://www.cisco.com/en/US/products/hw/switches/ps708/products_white_paper0900aecd8015f0ae.shtml
- http://www.cisco.com/en/US/prod/collateral/wireless/wirelssw/ps1953/product_implementation_design_guide09186a00800a3016.pdf

Quality of Service

Quality of Service (QoS) is a vital part of any security policy for an enterprise network. Even though most people think of QoS as setting the priority of traffic in a network, it also controls the amount of data that is allowed into the network. In the case of Cisco switches, that control point is at the port level when the data comes from the phone to the Ethernet switch. The more control applied at the edge of the network at the access port, the fewer problems will be encountered as the data aggregates in the network.

As mentioned previously in the lobby phone example, you can provide enough flow control of the traffic at the access port level to prevent any attacker from launching a denial-of-service (DoS) attack from that port in the lobby. The configuration for that example was not as aggressive as it could be because the QoS configuration allowed traffic sent to the port to exceed the maximum rate, but the traffic was remarked to the level of scavenger class. Given a more aggressive QoS policy, any amount of traffic that exceeded that maximum limit of the policy could just be dropped at the port, and that "unknown" traffic would never make it into the network. QoS should be enabled across the entire network to give the IP Telephony data high priority from end to end.

For more information on QoS, refer to the chapter on [Network Infrastructure, page 3-1](#), and the *Enterprise QoS Solution Reference Network Design (SRND) Guide* available at

<http://www.cisco.com/go/designzone>

Advantages

QoS can be used to control not only the priority of the traffic in the network but also the amount of traffic that can travel through any specific interface. Cisco Smartports templates have been created to assist in deploying voice QoS in a network at the access port level.

Disadvantages

If QoS settings are outside the standard Cisco Smartports template, the configuration can be complex and difficult to manage in large IP Telephony deployments.

Access Control Lists

This section covers access control lists (ACLs) and their uses in protecting voice data.

VLAN Access Control Lists

You can use VLAN access control lists (ACLs) to control data that flows on a network. Cisco switches have the capability of controlling Layers 2 to 4 within a VLAN ACL. Depending on the types of switches in a network, VLAN ACLs can be used to block traffic into and out of a particular VLAN. They can also be used to block intra-VLAN traffic to control what happens inside the VLAN between devices.

If you plan to deploy a VLAN ACL, you should verify which ports are needed to allow the phones to function with each application used in your IP Telephony network. Normally any VLAN ACL would be applied to the VLAN that the phones use. This would allow control at the access port, as close as possible to the devices that are plugged into that access port.

Refer to the following product documentation for information on configuring VLAN ACLs:

- Cisco Catalyst 3750 Switches

http://www.cisco.com/en/US/products/hw/switches/ps5023/products_installation_and_configuration_guides_list.html

- Cisco Catalyst 4500 Series Switches
http://www.cisco.com/en/US/products/hw/switches/ps4324/tsd_products_support_configure.html
- Cisco Catalyst 6500 Series Switches running Cisco IOS
http://www.cisco.com/en/US/products/hw/switches/ps708/products_installation_and_configuration_guides_list.html
- Cisco 6500 Series Switches running Cisco CatOS
http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_configure.html

The following example represents a VLAN ACL that allows only the traffic for a Cisco 7960 IP Phone to boot and function in a VLAN. (Inline comments indicate the purpose of each line of the ACL.) This example VLAN ACL is for the ports used with Cisco Unified CallManager Release 4.1. The example uses the following IP address ranges:

- Phones are in the range 10.0.20.*
- Servers are in the range 10.0.10.*
- Gateways are in the range 10.0.30.*
- Default gateways are 10.0.10.2 and 10.0.10.3
- DNS server IP address is 10.0.40.3

**Note**

The ports do change when either the application is updated or the OS is updated, or both. This note applies to all the IP Telephony devices in the network, including phones. To obtain the latest list of ports used by a product, refer to the appropriate documentation for the version of the product that is running on your network. Port usage documentation is available at http://www.cisco.com/en/US/products/sw/voicesw/ps556/prod_maintenance_guides_list.html

```
20 permit udp host 10.0.10.2 eq 1985 any
30 permit udp host 10.0.10.3 eq 1985 any
!permit HSRP from the routers
40 permit udp any any eq bootpc
50 permit udp any any eq bootps
!permit DHCP activity
60 permit udp 10.0.10.0 0.0.0.255 range 49152 65535 10.0.20.0 0.0.0.255 eq tftp
70 permit udp 10.0.20.0 0.0.0.255 range 1024 5000 10.0.10.0 0.0.0.255 range 49152 65535
80 permit udp 10.0.10.0 0.0.0.255 range 49152 65535 10.0.20.0 0.0.0.255 range 1024 5000
!permit the tftp traffic from the tftp server and phone
90 permit udp 10.0.10.0 0.0.0.255 range 49152 65535 host 10.0.40.3 eq domain
100 permit udp host 172.19.244.2 eq domain 10.0.10.0 0.0.0.255 range 49152 65535
!permit DNS to and from the phone
110 permit tcp 10.0.10.0 0.0.0.255 range 49152 65535 10.0.20.0 0.0.0.255 eq 2000
120 permit tcp 10.0.20.0 0.0.0.255 eq 2000 10.0.10.0 0.0.0.255 range 49152 65535
!permit signaling to and from the phone.
130 permit udp 10.0.10.0 0.0.0.255 range 16384 32767 10.0.10.0 0.0.0.255 range 16384 32767
140 permit udp 10.0.0.0 0.0.255.255 range 16384 32767 10.0.10.0 0.0.0.255 range 16384
32767
150 permit udp 10.0.10.0 0.0.0.255 range 16384 32767 10.0.0.0 0.0.255.255 range 16384
43767
!permit all phones to send udp to each other
160 permit tcp 10.0.10.0 0.0.0.255 range 49152 65535 10.0.20.0 0.0.0.255 eq www
170 permit tcp 10.0.20.0 0.0.0.255 eq www 10.0.10.0 0.0.0.255 range 49152 65535
180 permit tcp 10.0.20.0 0.0.0.255 range 49152 65535 10.0.10.0 0.0.0.255 eq www
190 permit tcp 10.0.10.0 0.0.0.255 eq www 10.0.20.0 0.0.0.255 range 49152 65535
!permit web access to and from the phone
200 permit Intelligent Contact ManagementP any any
```

```
!allow all icmp - phone to phone, gateway to phone, and NMS to phone
220 permit udp 10.0.30.0 0.0.0.255 rang 16384 327676 10.0.10.0 0.0.0.255 rang 16384 32767
!permit udp to the gateways in the network for pstn access
```

As this example ACL illustrates, the more well-defined the IP addresses are in a network, the easier it is to write and deploy an ACL.

For more details on how to apply VLAN ACLs, refer to the following documentation:

- Cisco Catalyst 3750 Switch
http://www.cisco.com/en/US/products/hw/switches/ps5023/products_installation_and_configuration_guides_list.html
- Cisco Catalyst 4500 Series Switches
http://www.cisco.com/en/US/products/hw/switches/ps4324/tsd_products_support_configure.html
- Cisco Catalyst 6500 Series Switches with Cisco CatOS
http://www.cisco.com/en/US/products/hw/switches/ps708/products_installation_and_configuration_guides_list.html
- Cisco Catalyst 6500 Series Switches with Cisco IOS
http://www.cisco.com/en/US/products/hw/switches/ps708/tsd_products_support_configure.html

Advantages


ACLs provide the ability to control the network traffic in and out of a VLAN as well as the ability to control the traffic within the VLAN.

Disadvantages

VLAN ACLs are very difficult to deploy and manage at an access-port level that is highly mobile. Because of these management issues, care should be taken when deploying VLAN ACLs at the access port in the network.

Router Access Control Lists

As with VLAN ACLs, routers have the ability to process both inbound and outbound ACLs by port. The first Layer 3 device is the demarcation point between voice data and other types of data when using voice and data VLANs, where the two types of data are allowed to send traffic to each other. Unlike the VLAN ACLs, router ACLs are not deployed in every access device in your network. Rather, they are applied at the edge router, where all data is prepared for routing across the network. This is the perfect location to apply a Layer 3 ACL to control which areas the devices in each of the VLANs have the ability to access within a network. Layer 3 ACLs can be deployed across your entire network to protect devices from each other at points where the traffic converges. (See [Figure 20-13](#).)

Figure 20-13 Router ACLs at Layer 3

There are many types of ACLs that can be deployed at Layer 3. For descriptions and examples of the most common types, refer to *Configuring Commonly Used IP ACLs*, available (with Cisco partner login required) at

http://cisco.com/en/US/partner/tech/tk648/tk361/technologies_configuration_example09186a0080100548.shtml

Depending on your security policy, the Layer 3 ACLs can be as simple as not allowing IP traffic from the non-voice VLANs to access the voice gateway in the network, or the ACLs can be detailed enough to control the individual ports and the time of the day that are used by other devices to communicate to IP Telephony devices. Assuming that there are no software phones, ACLs can be written to block all traffic (by IP address or IP range) to Cisco Unified CallManagers, voice gateways, phones, and any other voice application that is being used for voice-only services. This method simplifies the ACLs at Layer 3 compared to the ACLs at Layer 2 or VLAN ACLs.

This example uses the following IP address ranges:

- Phones are in the range 10.0.20.*
- IP Telephony servers are in the range 10.0.10.*
- Gateways are in the range 10.0.30.*
- All other devices in the network are in the range 192.168.*.*

```
10 deny ip 192.168.0.0 0.0.255.255 10.0.10.0 0.0.0.255
!deny all non voice devices to the voip servers
20 deny 192.168.0.0 0.0.255.255 10.0.30.0 0.0.0.255
!deny all non voice devices to the voip gateways
30 deny 192.168.0.0 0.0.255.255 10.0.20.0 0.0.0.255
!deny all non voice devices to communicate with the phones ip addresses
```

Advantages

ACLs are much easier to manage and deploy at Layer 3. Layer 3 is one of the first opportunities to apply control to voice data and other non-voice data in the network.

Disadvantages

As the ACLs become more granular and detailed, any changes in port usage in a network could break not only voice but also other applications in the network.

If there are software phones in the network, if web access to the phone is allowed, or if you use the Attendant Console or other applications that need access to the voice VLAN subnets, the ACLs are much more difficult to deploy and control.


Gateways and Media Resources

Gateways and media resources are devices that convert an IP Telephony call into a PSTN call. When an outside call is placed, the gateway or media resource is one of the few places within an IP Telephony network to which all the voice RTP streams flow.

Because IP Telephony gateways and media resources can be placed almost anywhere in a network, securing an IP Telephony gateway or media resource might be considered more difficult than securing other devices, depending on your security policy. However, depending on which point trust is established in the network, the gateways and media resources can be quite easy to secure. Because of the way the gateways and media resources are controlled by Cisco Unified CallManager, if the path that the signaling takes to the gateway or media resource is in what is considered a secure section of the network, a simple ACL can be used to control signaling to and from the gateway or media resource. If the network is not considered secure between the gateways (or media resources) and where the Cisco Unified CallManagers are located (such as when a gateway is located at a remote branch), the infrastructure can be used to build IPSec tunnels to the gateways and media resources to protect the signaling. In most networks there would most likely be a combination of the two approaches (ACL and IPSec) to secure those devices.

For H.323 videoconferencing devices, an ACL can be written to block port 1720 for H.225 trunks from any H.323 client in the network. This method would block users from initiating an H.225 session with each other directly. Cisco devices might use different ports for H.225, so refer to the product documentation for your equipment to see which port is used. If possible, change the port to 1720 so that only one ACL is needed to control signaling.

Because we use QoS at the edge of the network, if an attacker can get into the voice VLAN and determine where the gateways and media resources are, QoS at the port would limit how much data the attacker would be able to send to the gateway or media resource. (See [Figure 20-14](#).)

Figure 20-14 Securing Gateways and Media Resources with IPSec, ACLs, and QoS

Some gateways and media resources support Secure RTP (SRTP) to the gateways and media resources from the phones, if the phone is enabled for SRTP. To determine if a gateway or media resource supports SRTP, refer to the appropriate product documentation at

<http://www.cisco.com>


For more information on IPSec tunnels, refer to the *Site-to-Site IPSec VPN Solution Reference Network Design (SRND)*, available at

<http://www.cisco.com/go/designzone>

Putting Firewalls Around Gateways

Some very interesting issues arise from placing firewalls between a phone making a call and the gateway to the PSTN network. Stateful firewalls look into the signaling messages between Cisco Unified CallManager, the gateway, and the phone, and they open up a pinhole for the RTP streams to allow the call to take place. To do the same thing with a normal ACL, the entire port ranges used by the RTP streams would have to be open to the gateway.

There are two ways to deploy gateways within a network: behind a firewall and in front of a firewall. If you place the gateway behind a firewall, all the media from the phones that are using that gateway have to flow through the firewall, and additional CPU resources are required to run those streams through the firewall. In turn, the firewall adds control of those streams and protects the gateway from denial-of-service attacks. (See Figure 20-15.)

Figure 20-15 Gateway Placed Behind a Firewall

The second way to deploy the gateway is outside the firewall. Because the only type of data that is ever sent to the gateway from the phones is RTP streams, the access switch's QoS features control the amount of RTP traffic that can be sent to that gateway. The only thing that Cisco Unified CallManager sends to the gateway is the signaling to set up the call. If the gateway is put in an area of the network that is trusted, the only communication that has to be allowed between Cisco Unified CallManager and the gateway is that signaling. (See [Figure 20-15](#).) This method of deployment decreases the load on the firewall because the RTP streams are not going through the firewall.

Advantages

Unlike an ACL, most firewall configurations will open only the RTP stream port that Cisco Unified CallManager has told the phone and the gateway to use between those two devices as long as the signaling goes through the firewall. The firewall also has additional features for DoS attacks and Cisco Intrusion Detection System (IDS) signatures to look at interesting traffic and determine if any attackers are doing something they should not be doing.

Disadvantages

As stated in the section on [Firewalls](#), [page 20-31](#), when a firewall is looking at all the signaling and RTP streams from phones to a gateway, capacity could be an issue. Also, if data other than voice data is running through the firewall, CPU usage must be monitored to make sure that the firewall does not affect the calls that are running through the firewall.

In active or standby mode, the lowest setting for the failover timer is 3 seconds for the Cisco Adaptive Security Appliance (ASA) and Cisco Private Internet Exchange (PIX), and the minimum timer setting for the Cisco Firewall Services Module (FWSM) is also 3 seconds. The firewalls will fail-over quickly once the standby unit has determined that it needs to take over. The state of the data running through the primary firewall will be passed to the failover unit if stateful failover is configured, so everything that was taking place before the failover will be maintained. But there is a possibility that, in the case of a complete failure of the primary unit or connectivity to that unit, at least 3 seconds of traffic would not pass to the gateway for the ASA or PIX, or 3 seconds of traffic for the FWSM. In other words, if there is some type of failure that forces the firewalls to fail-over, there would be at least a 3-second stoppage of RTP streams.

Firewalls

Firewalls can be used in conjunction with ACLs to protect the voice servers and the voice gateways from devices that are not allowed to communicate with IP Telephony devices. Because of the dynamic nature of the ports used by IP Telephony, having a firewall does help to control opening up a large range of ports needed for IP Telephony communications. Given the complexities that firewalls introduce into a network design, you must take care in placing and configuring the firewalls and the devices around the firewalls to allow the traffic that is considered correct to pass while blocking the traffic that needs to be blocked.

IP Telephony networks have unique data flows. The phones use a client/server model for signaling for call setup, and Cisco Unified CallManager controls the phones through that signaling. The data flows for the IP Telephony RTP streams are more like a peer-to-peer network, and the phones or gateways talk directly to each other via the RTP streams. If the signaling flows do not go through the firewall so that the firewall can inspect the signaling traffic, the RTP streams could be blocked because the firewall will not know which ports need to be opened to allow the RTP streams for a conversation.

A firewall placed in a correctly designed network can force all the data through that device, so capacities and performance need to be taken into account. Performance includes the amount of latency, which can be increased by a firewall if the firewall is under high load or even under attack. The general rule in an IP Telephony deployment is to keep the CPU usage of an FWSM, ASA, or PIX to less than 60% for normal usage. If the CPU runs over 60%, it increases the chance of impacting IP phones, call setup, and registration. If the CPU usage stays at a sustained level above 60%, the registered IP phones will be affected, quality of calls in progress will degrade, and call setup for new calls will suffer. In the worst case, if the sustained CPU usage stays above 60%, phones will start to unregister. When this happens, they will attempt to re-register with Cisco Unified CallManager, thus increasing the load on the firewalls even more. If this were to happen, the effect would be a rolling blackout of phones unregistering and attempting to re-register with Cisco Unified CallManager. Until the CPU usage of the firewall decreases to under 60% sustained load, this rolling blackout would continue and most (if not all) of the phones would be affected. If you are currently using a Cisco firewall in your network, you should monitor the CPU usage carefully when adding IP Telephony traffic to your network so that you do not adversely affect that traffic.


There are many ways to deploy firewalls. This section concentrates on the ASA, PIX, and FWSM in the active/standby mode in both routed and transparent scenarios. Each of the configurations in this section is in single-context mode within the voice sections of the firewall configurations.

All of the Cisco firewalls can run in either multiple-context or single-context mode. In single-context mode, the firewall is a single firewall that controls all traffic flowing through it. In multiple-context mode, the firewalls can be turned into many virtual firewalls. Each of these contexts or virtual firewalls have their own configurations and can be controlled by different groups or administrators. Each time a new context is added to a firewall, it will increase the load and memory requirements on that firewall. When you deploy a new context, make sure that the CPU requirements are met so that voice RTP streams are not adversely affected.

Functionality Differences of the ASA or PIX and the FWSM

Figure 20-16 shows a logical representation of redundant firewalls in a network. Placement is the same for both routed configurations and transparent configurations.

Figure 20-16 Redundant Routed or Transparent Firewalls



Both the Cisco Adaptive Security Appliance (ASA) and the Cisco Private Internet Exchange (PIX) operate in a different manner than the Cisco Firewall Services Modules (FWSM). Within an ASA or PIX, as long as there are no ACLs on a more trusted interface, all traffic from that interface is trusted and allowed out to a less-trusted interface. (See Figure 20-17.) For example, all traffic from the inside or data center interface of an ASA will be allowed to go out to the outside interface of the ASA. Once any ACL is applied to the more trusted interface on a ASA/PIX, all other traffic is denied (DENY) and the firewall will then function very much like the FWSM. (See Figure 20-18.)

Figure 20-17 Functionality of a Cisco ASA or PIX





Figure 20-18 **Functionality of a Cisco FWSM****Overall Advantages of Firewalls**

A firewall provides a security control point in the network for applications that run over the network. A firewall also provides dynamic opening of ports for IP Telephony conversations if that traffic is running through the firewall.

Using its Application Layer Gateway (ALG) capability, the firewall can inspect the traffic that runs through it to determine if that traffic is really the type of traffic that the firewall is expecting. For example, does the HTTP traffic really look like HTTP traffic, or is it an attack? If it is an attack, then drop that packet and do not allow that packet to get to the HTTP server behind the firewall.

Overall Disadvantages of Firewalls

Not all IP Telephony application servers or applications are supported through a firewall. Some applications that are not supported with firewalls or with an ALG in the firewall include Cisco Unity voicemail servers, Attendant Console, Cisco Unified Contact Center Enterprise, and Cisco Unified Contact Center Express. ACLs can be written for these applications to allow traffic to flow through a firewall.

SCCP video is not supported through any of the firewalls currently shipping (up to and including versions of ASA/PIX 7.1.x and FWSM 3.1.x). SIP used on *Cisco Unified CallManager 5.x* is not currently supported through the firewalls using ALGs. If SIP signaling or media is required to traverse the firewalls, ACLs must be used until support is available.

Versions of Cisco FWSM prior to version 3.0 do not support SCCP fragmentation. If an SCCP packet is fragmented from a phone, from Cisco Unified CallManager, or from a gateway to another IP Telephony device, the fragmented packet will not be allowed through the FWSM. In cases where fragmentation occurs with an FWSM running version 2.x code, an ACL should be used without the ALG feature of the firewall for the signaling traffic. This configuration will allow the signaling traffic through the FWSM but will not do packet inspection as the signaling goes through the firewall.

To determine if the applications running on your network are supported with the version of firewall in the network or if ACLs have to be written, refer to the appropriate application documentation available at

<http://www.cisco.com>

Routed ASA and PIX

The ASA or PIX firewall in routed mode acts as a router between connected networks, and each interface requires an IP address on a different subnet. In single-context mode, the routed firewall supports Open Shortest Path First (OSPF) and Routing Information Protocol (RIP) in passive mode. Multiple-context mode supports static routes only. Cisco recommends using the advanced routing capabilities of the

upstream and downstream routers instead of relying on the security appliance for extensive routing needs. For more information on the routed mode, refer to the *Cisco Security Appliance Command Line Configuration Guide*, available at

http://www.cisco.com/en/US/products/ps6120/products_installation_and_configuration_guides_list.html

Advantages

The routed ASA or PIX firewall supports QoS, NAT, and VPN termination to the box, which are not supported in the transparent mode (see [Transparent ASA and PIX](#), page 20-34).

Figure 20-16 shows the logical placement of firewalls for both routed and transparent configurations in active standby mode. With the routed configuration, each interface on the ASA or PIX would have an IP address. In the transparent mode, there would be no IP address on the interfaces other than the IP address to manage the ASA or PIX remotely.

Disadvantages

Unlike with transparent mode, the device can be seen in the network and, because of that, it can be a point of attack. Placing a routed ASA or PIX firewall in a network changes the network routing because some of the routing can be done by the firewall. IP addresses must also be available for all the interfaces on the firewall that are going to be used, so changing the IP addresses of the routers in the network might also be required. If a routing protocol or RSVP is to be allowed through the ASA or PIX firewall, then an ACL will have to be put on the inside (or most trusted) interface to allow that traffic to pass to the outside (or lesser trusted) interfaces. That ACL must also define all other traffic that will be allowed out of the most trusted interface.

Transparent ASA and PIX

The ASA or PIX firewall can be configured to be a Layer 2 firewall (also known as "bump in the wire" or "stealth firewall"). In this configuration, the firewall does not have an IP address (other than for management purposes), and all of the transactions are done at Layer 2 of the network. Even though the firewall acts as a bridge, Layer 3 traffic cannot pass through the security appliance unless you explicitly permit it with an extended access list. The only traffic allowed without an access list is Address Resolution Protocol (ARP) traffic.

Advantages

This configuration has the advantage that an attacker cannot see the firewall because it is not doing any dynamic routing. Static routing is required to make the firewall work even in transparent mode.

This configuration also makes it easier to place the firewall into an existing network because routing does not have to change for the firewall. It also makes the firewall easier to manage and debug because it is not doing any routing within the firewall. Because the firewall is not processing routing requests, the performance of the firewall is usually somewhat higher with **inspect** commands and overall traffic than the same firewall model and software that is doing routing.

Disadvantages

With transparent mode, you are unable to use NAT on the firewall. If you are going to pass data for routing, you will also have to define the ACLs both inside and outside the firewall to allow traffic, unlike with the same firewall in routed mode. Cisco Discovery Protocol (CDP) traffic will not pass through the device even if it is defined. Each directly connected network must be on the same subnet. You cannot share interfaces between contexts; if you plan on running multiple-context mode, you will have to use additional interfaces. You must define all non-IP traffic, such as routing protocols, with an ACL to allow

that traffic through the firewall. QoS is not supported in transparent mode. Multicast traffic can be allowed to go through the firewall with an extended ACL, but it is not a multicast device. In transparent mode, the firewall does not support VPN termination other than for the management interface.

If a routing protocol or RSVP is to be allowed through the ASA or PIX firewall, then an ACL will have to be put on the inside (or most trusted) interface to allow that traffic to pass to the outside (or lesser trusted) interfaces. That ACL must also define all other traffic that will be allowed out of the most trusted interface.

For more information on the transparent mode, refer to the *Cisco Security Appliance Command Line Configuration Guide*, available at

http://www.cisco.com/en/US/products/ps6120/products_installation_and_configuration_guides_list.html

ASA and PIX Configuration Example

The following configuration example lists the ports and **inspect** commands that are used to make the firewalls work with voice for ASA and PIX software Release 7.04. This is an example only, and you should review the ports list from all the applications that are used in your network before deploying any firewall. This configuration example shows only the voice sections.

```
!
!
object-group service remote-access tcp
  description remote access
  !Windows terminal
  port-object range 3389 3389
  !VNC
  port-object range 5800 5800
  !VNC
  port-object range 5900 5900
  port-object range 8080 8080
  port-object eq ssh
  !SSH
  port-object eq ftp-data
  !FTP data transport
  port-object eq www
  !HTTP Access
  port-object eq ftp
  !FTP
  port-object eq https
  !HTTPS Access
object-group service voice-protocols-tcp tcp
  description TCP voice protocols
  CTI/QBE
  port-object range 2428 2428
  !SIP communication
  port-object eq ctigbe
  !SCCP
  port-object range 2000 2000
  !Secure SCCP
  port-object range 2443 2443
object-group service voice-protocols-udp udp
  !TFTP
  port-object eq tftp
  !MGCP Signaling
  port-object range 2427 2427
  !DNS
  port-object eq domain
  !RAS
```

```

port-object range 1719 1719
!SIP

!Object Group applied for remote-access
access-list OUTSIDE extended permit tcp any any object-group remote-access
!Object Group applied for voice-protocols-tcp
access-list OUTSIDE extended permit tcp any any object-group voice-protocols-tcp
!Object Group applied for voice-protocols-udp
access-list OUTSIDE extended permit udp any any object-group voice-protocols-udp
! Object Group applied for remote-access
access-list inside_access_in extended permit tcp any any object-group remote-access
! Object Group applied for voice-protocols-tcp
access-list inside_access_in extended permit tcp any any object-group voice-protocols-tcp
! Object Group applied for voice-protocols-udp
access-list inside_access_in extended permit udp any any object-group voice-protocols-udp

!Failover config
ip address 172.19.245.3 255.255.255.248 standby 172.19.245.4
failover
failover lan unit primary
failover lan interface failover GigabitEthernet0/3
failover polltime unit 1 holdtime 3
!Lowest and fastest setting for failover
failover polltime interface 3
failover link failover_state GigabitEthernet0/2
failover interface ip failover 192.168.1.1 255.255.255.0 standby 192.168.1.2
failover interface ip failover_state 192.168.0.1 255.255.255.0 standby 192.168.0.2

!
!Default inspection with inspections enabled
class-map inspection_default
  match default-inspection-traffic
!
!
policy-map global_policy
  class inspection_default
    inspect h323 h225
    inspect h323 ras
    inspect skinny
    inspect sip
    inspect tftp
    inspect mgcp

```

FWSM Routed Mode

In routed mode, the FWSM is considered to be a router hop in the network. It performs NAT between connected networks and can use OSPF or passive RIP (in single-context mode). Routed mode supports up to 256 interfaces per context or in single mode, with a maximum of 1000 interfaces divided between all contexts.

Advantages

As a routed device in your network, the FWSM supports routing and all other features that are not available in transparent mode.

Disadvantages

Unlike the transparent mode, the routed device is visible in the network and, because of that, it can be a point of attack. To place the device in a network, IP addressing and routing must be change.

FWSM Transparent Mode

In transparent mode, the FWSM acts like a "bump in the wire," or a "stealth firewall," and is not a router hop. The FWSM connects the same network on its inside and outside interfaces, but each interface must be on a different VLAN. No dynamic routing protocols or NAT are required. However, like routed mode, transparent mode also requires ACLs to allow traffic to pass through. Transparent mode can also optionally use EtherType ACLs to allow non-IP traffic. Transparent mode supports only two interfaces, an inside interface and an outside interface.

You might use a transparent firewall to simplify your network configuration. Transparent mode is also useful if you want the firewall to be invisible to attackers. You can also use a transparent firewall for traffic that would otherwise be blocked in routed mode. For example, a transparent firewall can allow multicast streams using an EtherType ACL.

Advantages

This configuration has the advantage that an attacker cannot see the firewall because it is not doing any routing. This configuration also makes it easier to place the firewall into an existing network because routing does not have to change for the firewall. It also makes the firewall easier to manage and debug because it is not doing any routing within the firewall. You can also bridge non-IP traffic and IP multicast traffic, static ARP inspection, and MAC move detection and static MAC.

Disadvantages

To avoid loops when you use failover in transparent mode, you must use switch software that supports Bridge Port Data Unit (BPDU) forwarding, and you must configure the FWSM to allow BPDUs.

Transparent mode does not support NAT, dynamic routing, or a unicast reverse path forwarding (RPF) check. There is no NAT 0 with the FWSM in transparent mode.

FWSM Configuration Example

The following configuration example lists the ports and **inspect** commands that are used to make the firewalls work with voice for FWSM software Release 2.3.x. This is only an example, and you should review the ports list from all the applications that are used in your network before deploying any firewall. This configuration example shows only the voice sections.

```
fixup protocol h323 H225 1720
!Enable fixup h3232 h225

fixup protocol h323 ras 1718-1719
!Enable fixup h323 RAS

fixup protocol mgcp 2427
!Enable fixup mgcp

fixup protocol skinny 2000
!Enable fixup

fixup protocol tftp 69
!Enable fixup

object-group service VoiceProtocols tcp
  description Unified CM Voice protocols
  port-object eq ctigbe
  port-object eq 2000
  port-object eq 3224
  port-object eq 2443
```

```

    port-object eq 2428
    port-object eq h323
!Defining the ports for TCP voice

object-group service VoiceProtocolsUDP udp
    description UDP based Voice Protocols
    port-object range 2427 2427
    port-object range 1719 1719
    port-object eq tftp
!Defining the ports for UDP voice

object-group service RemoteAccess tcp
    description Remote Acces
    port-object range 3389 3389
    port-object range 5800 5809
    port-object eq ssh
    port-object range 5900 5900
    port-object eq www
    port-object eq https
!Defining remote access TCP ports

access-list inside_nat0_outbound extended permit ip any any
!

access-list phones_access_in extended permit tcp any any object-group RemoteAccess log
notifications interval 2
access-list phones_access_in extended permit tcp any any object-group VoiceProtocols log
notifications interval 2
access-list phones_access_in extended permit udp any any object-group VoiceProtocolsUDP
log notifications interval 2
access-list phones_access_in extended deny ip any any log notifications interval 2
access-list outside_access_in extended permit tcp any any object-group VoiceProtocols log
notifications interval 2
access-list outside_access_in extended permit tcp any any object-group RemoteAccess log
notifications interval 2
access-list outside_access_in extended permit udp any any object-group VoiceProtocolsUDP
log notifications interval 2
!Access lists applying the object groups defined above for inside and outside interfaces

access-list outside_access_in extended deny ip any any log notifications interval 2
access-list inside_access_in extended deny ip any any
!Deny all other traffic

access-list phones_nat0_outbound extended permit ip any any
!

failover
failover lan unit primary
failover lan interface fln vlan 4050
failover polltime unit 1 holdtime 5
failover polltime interface 15
!Failover config - 15 seconds
failover interface-policy 50%
failover link fln vlan 4051
failover interface ip fln 1.1.1.1 255.255.255.252 standby 1.1.1.2
failover interface ip fln 1.1.1.5 255.255.255.252 standby 1.1.1.6
nat (inside) 0 access-list inside_nat0_outbound_V1
access-group outside_access_in in interface outside
access-group inside_access_in in interface inside

```

Data Center

Within the data center, the security policy should define what security is needed for the IP Telephony applications servers. Because the Cisco Unified Communications servers are based on IP, the security that you would put on any other time-sensitive data within a data center could be applied to those servers as well.

If clustering over the WAN is being used between data centers, any additional security that is applied both within and between those data centers has to fit within the maximum round-trip time that is allowed between nodes in a cluster. If your current security policy for application servers within your network covers the IP Telephony servers from Cisco, then you should use that security. You can also use any of the infrastructure security that is already deployed.

To design appropriate data center security for your data applications, Cisco recommends following the guidelines presented in the *Data Center Networking: Server Farm Security SRND (Server Farm Security in the Business Ready Data Center Architecture)*, available at

<http://www.cisco.com/go/designzone>

Applications Servers

For a list of the Cisco Unified CallManager security features and how to enable them, refer to the *Cisco Unified CallManager Security Guide*, available at

http://www.cisco.com/en/US/products/sw/voicesw/ps556/prod_maintenance_guides_list.html

Before enabling any of the Cisco Unified CallManager security features, verify that they will satisfy the security requirements specified in your enterprise security policy for these types of devices in a network.

Cisco Security Agent on the Cisco Unified CallManager and Application Servers

The Cisco Security Agent is used on most of the application servers that Cisco uses to provide IP Telephony and IP Telephony services. The Cisco Security Agent software is Host Intrusion Prevention software that looks at the behavior of the traffic to and from the server, and the way the applications are running on that server, to determine if everything is working correctly. If something is considered abnormal, then the Cisco Security Agent software prevents that activity from happening. For example, if a virus is trying to install a software package on a Cisco Unified CallManager and that is something that has never happened before, the virus would be prevented from installing. Antivirus software is still needed on the server because Cisco Security Agent cannot clean a server that has been infected; it can only prevent the infection.

Unmanaged Cisco Security Agent

Cisco developed a default Cisco Security Agent policy for its servers that allows all the correct things needed for that IP Telephony server to function, while also preventing known and unknown attacks from affecting the IP Telephony servers. At a minimum, you should install and run this unmanaged version Cisco Security Agent. This software protects the application and the operating system from both viruses and worm attacks. To get the maximum protection from these types of intrusions, ensure that the newest version of the Cisco Security Agent software is always installed on your servers. With the unmanaged

agent installed on your servers, you will be able to see the logs of attacks on only the system where the agent is installed. You will have to log into each system to check the log files that might be there because of some type of alarm that has occurred.

Advantages

The unmanaged Cisco Security Agent protects each system from both known and unknown attacks, worms, and viruses.

Disadvantages

When run in unmanaged mode, Cisco Security Agent does not correlate alarms, and you have to access each system individually to see the log files on that system. When you upgrade the unmanaged Cisco Security Agent, you have to install the new client and usually reboot the system to make the client settings take effect. Cisco Security Agent cannot clean a system if that system does become infected for some reason. You also have to run an antivirus software on the system to help maintain security and protect the system.

Managed Cisco Security Agent

Managed Cisco Security Agent works in the same way as the unmanaged version, but some additional benefits come along with the management console. If you run the managed system, you will receive all the alarms from all of the systems on one console. This feature also enables you to be notified of any anomaly by being emailed or paged if an anomaly reaches a critical level.

**Note**

Managed Cisco Security Agent capability is not currently available for Cisco Unified CallManager 5.0

Advantages

Managed Cisco Security Agent provides the same protection as the unmanaged system, but it also gives you control of the agent. This control allows for the correlation of events, global reporting back to the management console, and upgrades to the Cisco Security Agent configuration of the agent without reloading the system when you have an update.

Disadvantages

A separate piece of software is required on a separate server for global monitoring and configuration of the managed agent. Cisco Security Agent cannot clean a system if that system does become infected for some reason. You must also run an antivirus software on the system to help maintain security and protect the system.

Antivirus

Approved antivirus software should be run on all of the IP Telephony servers and IP Telephony application servers that have been approved to run that software. As with any other server in a network, antivirus software helps protect the Cisco Unified CallManager server from being infected with a worm or a virus that could impact call processing. The Cisco Security Agent should not be the only preventative software installed on the system because Cisco Security Agent cannot clean the system of infection. Only an antivirus software can clean an infected system.

Additional information about running antivirus software on the Cisco Unified CallManager servers is available at

http://www.cisco.com/en/US/products/sw/voicesw/ps556/tsd_products_support_series_home.html

Advantages

Antivirus software helps prevent the application server from being infected and suffering performance degradation.

Disadvantages

Managing the antivirus software involves some overhead. In addition, you must ensure that the version of software is approved for installation on Cisco Unified CallManager and the other IP Telephony application servers.

General Server Guidelines

Your Cisco Unified CallManager and other IP Telephony application servers should not be treated as normal servers. Anything you do while configuring the system could affect calls that are trying to be places or that are in progress. As with any other business-class application, major configuration changes should be done within maintenance windows to keep from disrupting phone conversations.

Standard security policies for application servers might not be adequate for IP Telephony servers. Unlike email servers or web servers, voice servers will not allow you to refresh a screen or re-send a message. The voice communications are real-time events. Any security policy for IP Telephony servers should ensure that work that is not related to configuring or managing the voice systems is not done on the IP Telephony servers at any time. Activities that might be considered normal on application servers within a network (for example, surfing the internet) should not take place on the IP Telephony servers.

In addition, Cisco provides a well defined patch system for the IP Telephony servers, and it should be applied based on the patch policy within your IT organization. You should not patch the system normally using the OS vendor's patch system unless it is approved by Cisco Systems. All patches should be downloaded from Cisco or from the OS vendor as directed by Cisco Systems, and applied according to the patch installation process.

Additional information on how to increase the hardening of the OS for Cisco Unified CallManager is listed in the C:\Utils\SecurityTemplates directory on your Cisco Unified CallManager server. You should use the OS hardening techniques if your security policy requires you to lock down the OS even more than what is provided in the default installation.

Various software patches are available at the following location:

<http://www.cisco.com/kobayashi/sw-center/sw-voice.shtml>

**Note**

A Cisco.com login account is required for access to this link.

The above site also contains a notification tool that will email you when a critical patch must be applied to an IP Telephony server.

Advantages

General server security practices help to mitigate viruses and worms if the application server is treated like a PBX and not like other application servers.

Disadvantages

When the additional security features are configured, some of the Cisco Unified CallManager functionality might be reduced. Also, additional care is needed during upgrades because some of the services that are disabled for the additional security will have to be enabled for a successful upgrade.

Deployment Examples

This section presents examples of what could be done from a security perspective for a lobby phone and a firewall deployment. A good security policy should be in place to cover deployments similar to these types.

Lobby Phone Example

The example in this section illustrates one possible way to configure a phone and a network for use in an area with low physical security, such as a lobby area. None of the features in this example are required for a lobby phone, but if your security policy states more security is needed, then you could use the features listed in this example.

Because you would not want anyone to gain access to the network from the PC port on the phone, you should disable the PC port on the back of the phone to limit network access (see [PC Port on the Phone, page 20-5](#)). You should also disable the settings page on the phone so that potential attackers cannot see the IP addresses of the network to which the lobby phone is connected (see [Settings Access, page 20-9](#)). The disadvantage of not being able to change the settings on the phone usually will not matter for a lobby phone.

Because there is very little chance that a lobby phone will be moved, you could use a static IP address for that phone. A static IP address would prevent an attacker from unplugging the phone and then plugging into that phone port to get a new IP address (see [IP Addressing, page 20-4](#)). Also, if the phone is unplugged, the port state will change and the phone will no longer be registered with Cisco Unified CallManager. You can track this event in just the lobby phone ports to see if someone is trying to attach to the network.

Using static port security for the phone and not allowing the MAC address to be learned would mean that an attacker would have to change his MAC address to that of the phone, if he were able to discover that address. Dynamic port security could be used with an unlimited timer to learn the MAC address (but never unlearn it), so that it would not have to be added. Then the switchport would not have to be changed to clear that MAC address unless the phone is changed. The MAC address is listed in a label on the bottom of the phone. If listing the MAC address is considered a security issue, the label can be removed and replaced with a "Lobby Phone" label to identify the device. (See [Switch Port, page 20-12](#).)

A single VLAN could be used and Cisco Discovery Protocol (CDP) could be disabled on the port so that attackers would not be able to see any information from the Ethernet port about that port or switch to which it is attached. In this case, the phone would not have a CDP entry in the switch for E911 emergency calls, and each lobby phone would need either a label or an information message to local security when an emergency number is dialed.

A static entry in the DHCP Snooping binding table could be made because there would be no DHCP on the port (see [DHCP Snooping: Prevent Rogue DHCP Server Attacks, page 20-15](#)). Once the static entry is in the DHCP Snooping binding table, Dynamic ARP Inspection could be enabled on the VLAN to keep the attacker from getting other information about one of the Layer 2 neighbors on the network (see [Requirement for Dynamic ARP Inspection, page 20-20](#)).

With a static entry in the DHCP Snooping binding table, IP Source Guard could be used (see [IP Source Guard, page 20-22](#)). If an attacker got the MAC address and the IP address and then started sending packets, only packets with the correct IP address could be sent.

A VLAN ACL could be written to allow only the ports and IP addresses that are needed for the phones to operate (see [VLAN Access Control Lists, page 20-24](#)). The following example contains a very small ACL that can be applied to a port at Layer 2 or at the first Layer 3 device to help control access into the network (see [Router Access Control Lists, page 20-26](#)). This example is based on a Cisco 7960 IP Phone being used in a lobby area, without music on hold to the phone or HTTP access from the phone.

This example uses the following IP address ranges:

- The lobby phone has an IP address of 10.0.40.5
- The Cisco Unified CallManager cluster uses the address range 10.0.20.*
- The DNS server has an IP address of 10.0.30.2
- The HSRP routers have IP addresses 10.0.10.2 and 10.0.10.3
- Other phones in the network use IP addresses in the range 10.0.*.*

```
10 permit icmp any any
! Allow all icmp - phone to phone, gateway to phone and NMS to phone

20 permit udp host 10.0.10.2 eq 1985 any
!Allow HSRP information in, do not allow out

30 permit udp host 10.0.10.3 eq 1985 any
! Allow in from HSRP neighbor

40 permit udp host 10.0.40.5 range 49152 65535 10.0.20.0 0.0.0.255 eq tftp
! Using ip host from ephemeral port range from phone to the TFTP server port 69 (start of tftp)

50 permit udp 10.0.20.0 0.0.0.255 range 1024 5000 host 10.0.40.5 range 49152 65535
!Using IP subnet from TFTP server with ephemeral port range to ip host and ephemeral port range for phone

60 permit udp host 10.0.40.5 range 49152 65535 10.0.20.0 0.0.0.255 range 1024 5000
! Using host from phone to TFTP server with ephemeral port range to ip range and ephemeral port range for TFTP (continue the TFTP conversation)

70 permit udp host 10.0.40.5 range 49152 65535 host 10.0.30.2 eq domain
! Using IP host and ephemeral port range from phone to DNS server host

80 permit udp host 10.0.30.2 eq domain host 10.0.40.5 range 49152 65535
! Using IP from DNS server to phone host ip and ephemeral port range

90 permit tcp 10.0.40.5 range 49152 65535 10.0.20.0 0.0.0.255 eq 2000
! Using IP host and ephemeral port range from phone to Unified CM cluster for SCCP

100 permit tcp 10.0.20.0 0.0.0.255 eq 2000 host 10.0.40.5 range 49152 65535
! Using IP range and SCCP port to phone IP host and ephemeral port range

110 permit udp 10.0.0.0 0.0.255.255 range 16384 32767 host 10.0.40.5 range 16384 32767
! Using IP range and ephemeral port range from all phones or gateways outside a vlan to the IP host to phone

120 permit udp host 10.0.40.5 range 16384 32767 10.0.0.0 0.0.255.255 range 16384 43767
! Using IP host and ephemeral port range from vlan to all other phones or gateways

130 permit udp host 172.19.244.3 range 1024 5000 host 10.0.40.5 eq snmp
!From IP host of NMS server and ephemeral port range (Different for Windows vs Sun) to IP host of phones and SNMP port (161)
```

```
140 permit udp host 10.0.40.5 eq snmp host 172.19.244.3 range 1024 5000
! From IP host of phone with SNMP port (161) to IP host of NMS server and ephemeral port
range
```

Basic QoS Example for a Lobby Phone

Set the voice stream to G.729 and use QoS to limit the amount of traffic that can be sent into the port (see [Quality of Service, page 20-24](#)). Even if the QoS maximum is exceeded, the traffic will be reset to CS1, or Scavenger Class, which is the lowest priority traffic in most enterprise networks.


```
CAT2970(config)#mls qos map policed-dscp 0 24 46 to 8
! Excess traffic marked 0 or CS3 or EF will be remarked to CS1
CAT2970(config)#
CAT2970(config)#class-map match-all LOBBY-VOICE
CAT2970(config-cmap)# match access-group name LOBBY-VOICE
CAT2970(config-cmap)#class-map match-all LOBBY-SIGNALING
CAT2970(config-cmap)# match access-group name LOBBY-SIGNALING
CAT2970(config-cmap)#exit
CAT2970(config)#
CAT2970(config)#policy-map LOBBY-PHONE
CAT2970(config-pmap)#class LOBBY-VOICE
CAT2970(config-pmap-c)# set ip dscp 46 ! Lobby phone VoIP is marked to DSCP EF
CAT2970(config-pmap-c)# police 48000 8000 exceed-action policed-dscp-transmit
! Out-of-profile Lobby voice traffic (g.729) is marked down to Scavenger (CS1)
CAT2970(config-pmap-c)#class LOBBY-SIGNALING
CAT2970(config-pmap-c)# set ip dscp 24 ! Signaling is marked to DSCP CS3
CAT2970(config-pmap-c)# police 32000 8000 exceed-action policed-dscp-transmit
! Out-of-profile Signaling traffic is marked down to Scavenger (CS1)
CAT2970(config-pmap-c)#class class-default
CAT2970(config-pmap-c)# set ip dscp 0
CAT2970(config-pmap-c)# police 56000 8000 exceed-action policed-dscp-transmit
! Out-of-profile data traffic is marked down to Scavenger (CS1)
CAT2970(config-pmap-c)# exit
CAT2970(config-pmap)#exit
CAT2970(config)#
CAT2970(config)#interface GigabitEthernet0/1
CAT2970(config-if)# service-policy input LOBBY-PHONE ! Applies policy to int
CAT2970(config-if)#exit
CAT2970(config)#
CAT2970(config)#ip access list extended LOBBY-VOICE
CAT2970(config-ext-nacl)# permit udp any any range 16384 32767 ! VoIP ports
CAT2970(config-ext-nacl)#
CAT2970(config-ext-nacl)#ip access list extended LOBBY-SIGNALING
CAT2970(config-ext-nacl)# permit tcp any any range 2000 2002 ! SCCP ports
CAT2970(config-ext-nacl)#end
CAT2970#
```

Firewall Deployment Example (Centralized Deployment)

The example in this section is one way that firewalls could be deployed within the data center, with Cisco Unified CallManagers behind them (see [Figure 20-19](#)). In this example, the Cisco Unified CallManagers are in a centralized deployment, single cluster with all the phones outside the firewalls. Because the network in this deployment already contained firewalls that are configured in routed mode within the corporate data center, the load was reviewed before the placement of gateways was determined. After reviewing the average load of the firewall, it was decided that all the RTP streams would not transverse the firewall in order to keep the firewalls under the 60% CPU load (see [Putting Firewalls Around Gateways, page 20-29](#)). The gateways are placed outside the firewalls, and ACLs

within the network are used to control the TCP data flow to and from the gateways from the Cisco Unified CallManagers. An ACL is also written in the network to control the RTP streams from the phones because the IP addresses of the phones are well defined (see [IP Addressing, page 20-4](#)). The voice applications servers are placed within the demilitarized zone (DMZ), and ACLs are used at the firewalls to control access to and from the Cisco Unified CallManagers and to the users in the network. This configuration will limit the amount of RTP streams through the firewall using inspects, which will minimize the impact to the firewalls when the new voice applications are added to the existing network.

Figure 20-19 Firewall Deployment Example



Conclusion

This chapter did not cover all of the security that could be enabled to protect the voice data within your network. The techniques presented here are just a subset of all the tools that are available to network administrators to protect all the data within a network. On the other hand, even these tools do not have to be enabled within a network, depending on what level of security is required for the data within the network overall. Choose your security methods wisely. As the security within a network increases, so do the complexity and troubleshooting problems. It is up to each enterprise to define both the risks and the requirements of its organization and then to apply the appropriate security within the network and on the devices attached to that network.



CHAPTER 21

IP Telephony Endpoints

Last revised on: February 13, 2008

This chapter summarizes various types of IP Telephony endpoints along with their features and QoS recommendations. The IP Telephony endpoints can be categorized into the following major types:

- [Analog Gateways, page 21-2](#)
- [Cisco Unified IP Phones, page 21-6](#)
- [Software-Based Endpoints, page 21-7](#)
- [Wireless Endpoints, page 21-12](#)
- [Cisco IP Conference Station, page 21-16](#)
- [Video Endpoints, page 21-16](#)

These sections provide detailed information about each endpoint type. In addition, the section on [QoS Recommendations, page 21-21](#), lists generic QoS configurations, and the [Endpoint Features Summary, page 21-35](#), lists all the endpoint features.

The following list summarizes high-level recommendations for selecting IP Telephony endpoints:

- For low-density analog connections, use the Cisco Analog Telephone Adapter (ATA) or low-density analog interface module.
- For medium to high-density analog connections, use the high-density analog interface module, Cisco Communication Media Module (CMM) with 24-FXS port adapter, Catalyst 6500 24-FXS analog interface module, Cisco VG224, or Cisco VG248.
- For telephony users with limited call features who generate small amounts of traffic, use the Cisco Unified IP Phones 7902G, 7905G, 7906G, 7911G, 7912G, or 7912G-A.
- For transaction-type telephony users who generate a medium amount of traffic, use Cisco Unified IP Phones 7940G, 7941G, or 7941G-GE.
- For managers and administrative assistants who generate medium to heavy telephony traffic, use Cisco Unified IP Phones 7960G, 7961G, or 7961G-GE.
- For executives with extensive call features who generate high amounts of telephony traffic, use Cisco Unified IP Phones 7970G or 7971G-GE.
- For mobile workers and telecommuters, use the Cisco IP communicator.
- For users who need a mobile IP phone, use the Cisco Unified Wireless IP Phone 7920.

- For making video calls, use Cisco Unified Video Advantage associated with a Cisco Unified IP Phone, the Cisco IP Video Phone 7985G, or Sony and Tandberg SCCP endpoints.
- For formal conferencing environments, use the Cisco Unified IP Conference Station 7936.

Analog Gateways

Analog gateways include router-based analog interface modules, Cisco Communication Media Module (CMM) with 24-FXS port adapter, Catalyst 6500 24-FXS analog interface module, Cisco VG224, Cisco VG248, and Cisco Analog Telephone Adaptor (ATA) 186 and 188. An analog gateway is usually used to connect analog devices, such as fax machines, modems, TDD/TTYs, and analog phones, to the VoIP network so that the analog signal can be packetized and transmitted over the IP network.

Analog Interface Module

Cisco router-based analog interface modules include low-density interface modules (NM-1V, NM-2V, NM-HD-1V, NM-HD-2V, NM-HD-2VE, NM-HDV2, NM-HDV2-1T1/E1, and NM-HDV2-2T1/E1) and high-density interface modules (NM-HDA-4FXS and EVM-HD-8FXS/DID). Cisco analog interface modules connect the PSTN and other legacy telephony equipment, including PBXs, analog telephones, fax machines, and key systems, to Cisco multiservice access routers. Cisco analog interface modules are best suited for connecting low- to high-density analog devices to the IP network with limited call features.

Low-Density Analog Interface Module

The low-density analog interface modules include the NM-1V, NM-2V, NM-HD-1V, NM-HD-2V, NM-HD-2VE, NM-HDV2, NM-HDV2-1T1/E1, and NM-HDV2-2T1/E1. The NM-1V and NM-2V contain one or two interface cards (VIC). The interface cards include: two-port FXS VIC (VIC-2FXS); two-port FXO VIC (VIC-2FXO, VIC-2FXO-M1/M2/M3, and VIC-2FXO-EU); two-port Direct Inward Dial VIC (VIC-2DID); two-port E&M VIC (VIC-2E/M); two-port Centralized Automated Message Accounting VIC (VIC-2CAMA); and two-port BRI VIC (VIC-2BRI-S/T-TE and VIC-2BRI-NT/TE). The NM-1V and NM-2V can serve up to two and four FXS connections, respectively.

**Note**

The NM-1V and NM-2V are not supported on the Cisco 2800 and 3800 Series platforms. On the Cisco 2800 and 3800 Series platforms, the voice interface cards are supported in the on-board High-Speed WIC slots, including the VIC-2DID, VIC4-FXS/DID, VIC2-2FXO, VIC-2-4FXO, VIC2-2FXS, VIC2-2E/M, and VIC2-2BRI-NT/TE.

The NM-HD-1V and NM-HD-2V contain one and two VICs, respectively. The NM-HD-2VE contains two VICs or two voice/WAN interface cards (VWIC), or a combination of one VIC and one VWIC. The NM-HD-1V, NM-HD-2V, and NM-HD-2VE can serve up to 4, 8, and 8 FXS or FXO connections, respectively. The NM-HDV2, NM-HDV2-1T1/E1, and NM-HDV2-2T1/E1 can be fitted with either digital T1/E1 or analog/BRI voice interface cards, with up to 4 FXS or FXO connections. The difference among these three interface modules is that the NM-HDV2-1T1/E1 has one built-in T1/E1 port while the NM-HDV2-2T1/E1 has two built-in T1/E1 ports.

The voice interface cards include: 2-port and 4-port FXS VICs (VIC2-2FXS and VIC-4FXS/DID); 2-port and 4-port FXO VICs (VIC2-2FXO and VIC2-4FXO); 2-port Direct Inward Dial VIC (VIC-2DID); 2-port E&M VIC (VIC2-2E/M); and 2-port BRI VIC (VIC2-2BRI-NT/TE). The voice/WAN interface cards include: 1-port and 2-port RJ-48 multiflex trunk (MFT) T1/E1 VWICs for

both voice and WAN connections (VWIC-1MFT-T1, VWIC-2MFT-T1, VWIC-2MFT-T1-DI, VWIC-1MFT-E1, VWIC-2MFT-E1, VWIC-2MFT-E1-DI, VWIC-1MFT-G703, VWIC-2MFT-G703, VWIC2-1MFT-T1/E1, VWIC2-2MFT-T1/E1, VWIC2-1MFT-G703, and VWIC2-2MFT-G703). The G.703 interface cards are primarily for data connectivity but can in some cases be configured to support voice applications.

High-Density Analog Interface Module

The high-density analog interface module includes the NM-HDA-4FXS and EVM-HD-8FXS/DID. The NM-HDA-4FXS has four on-board FXS ports and room for two expansion modules from the following options:

- EM-HDA-8FXS: An 8-port FXS interface card
- EM-HDA-4FXO/EM2-HDA-4FXO: A 4-port FXO interface card

The NM-HDA-4FXS provides up to 12 analog ports (4 FXS and 8 FXO) with four built-in FXS ports and two EM-HDA-4FXO or EM2-HDA-4FXO extension modules, or 16 analog ports (12 FXS and 4 FXO) with four built-in FXS ports and one EM-HDA-8FXS, and one EM-HDA-4FXO or EM2-HDA-4FXO extension module. A configuration using two 8-port FXS extension modules is not supported. The NM-HDA also has a connector for a daughter module (DSP-HDA-16) that provides additional DSP resources to serve an additional 8 high-complexity calls or 16 medium-complexity calls.



Note

The EM2-HDA-4FXO supports the same density and features as the EM-HDA-FXO, but it provides enhanced features including longer loop length support of up to 15,000 feet and improved performance under poor line conditions when used in ground-start signaling mode.

The EVM-HD-8FXS/DID provides eight individual ports on the baseboard module and can be configured for FXS or DID signaling. In addition, the EVM-HD-8FXS/DID has room for two expansion modules from the following options:

- EM-HDA-8FXS: An 8-port FXS interface card
- EM-HDA-6FXO: A 6-port FXO interface card
- EM-HDA-3FXS/4FXO: A 3-port FXS and 4-port FXO interface card
- EM-4BRI-NT/TE: A 4-port BRI interface card

These extension modules can be used in any combination and provide for configurations of up to 24 FXS ports per EVM-HD-8FXS/DID.

Supported Platforms and Cisco IOS Requirements for Analog Interface Modules

The supported platforms for Cisco analog interface modules are the Cisco 2600, 2800, 3600, 3700, and 3800 Series. [Table 21-1](#) lists the maximum number of interface modules supported per platform, and [Table 21-2](#) lists the minimum Cisco IOS software version required.

Table 21-1 Maximum Number of Analog Interface Modules Supported per Platform

Platform	Maximum Number of Interface Modules Supported				
	NM-1V, -2V	NM-HDA-4FXS	EVM-HD	NM-HD-1V, -2V, -2VE	NM-HDV2, -1T1/E1, -2T1/E1
Cisco 2600XM	1	1	No	1	1
Cisco 2691	1	1	No	1	1
Cisco 3640	3	3	No	3	No
Cisco 3660	6	6	No	6	No
Cisco 3725	2	2	No	2	2
Cisco 3745	4	4	No	4	4
Cisco 2811	No	1	1	1	1
Cisco 2821	No	1	1	1	1
Cisco 2851	No	1	1	1	1
Cisco 3825	No	2	1	2	2
Cisco 3845	No	4	2	4	4

Table 21-2 Minimum Cisco IOS Requirements for Analog Interface Modules

Platform	Minimum Cisco IOS Software Release Required				
	NM-1V, -2V	NM-HDA-4FXS	EVM-HD	NM-HD-1V, -2V, -2VE	NM-HDV2, -1T1/E1, -2T1/E1
Cisco 2600XM	12.2(8)T	12.2(8)T	No	12.3.4T	12.3(7)T
Cisco 2691	12.2(8)T	12.2(8)T	No	12.3.4T	12.3(7)T
Cisco 3640	12.0(1)T or later	12.2(8)T or later	No	12.3.4T	No
Cisco 3660	12.0(1)T or later	12.2(8)T or later	No	12.3.4T	No
Cisco 3725	12.2(8)T or later	12.2(8)T	No	12.3.4T	12.3(7)T
Cisco 3745	12.2(8)T or later	12.2(8)T	No	12.3.4T	12.3(7)T
Cisco 2811	No	12.3.8T4	12.3.8T4	12.3.8T4	12.3.8T4
Cisco 2821	No	12.3.8T4	12.3.8T4	12.3.8T4	12.3.8T4
Cisco 2851	No	12.3.8T4	12.3.8T4	12.3.8T4	12.3.8T4
Cisco 3825	No	12.3(11)T	12.3(11)T	12.3(11)T	12.3(11)T
Cisco 3845	No	12.3(11)T	12.3(11)T	12.3(11)T	12.3(11)T

Cisco Communication Media Module (CMM)

The Cisco CMM is a line card that provides high-density analog, T1, and E1 gateway connections for Catalyst 6000 and Cisco 7600 Series switches. The Cisco CMM can serve up to 72 FXS connections. The CMM operates as either an MGCP or H.323 gateway, and it provides Survivable Remote Site Telephony (SRST) service for up to 480 IP phones.

Cisco CMM can contain the following interface port adapters: 24-port FXS analog port adapter (WS-SVC-CMM-24FXS), 6-port T1 interface port adapter (WS-SVC-CMM-6T1), 6-port E1 interface port adapter (WS-SVC-CMM-6E1), and conference/transcoding port adapter (WS-SVC-CMM-ACT). [Table 21-3](#) lists the minimum software requirements for the compatible port adapters.

Table 21-3 Software Requirements for CMM Port Adapters

	WS-SVC-CMM-24FXS	WS-SVC-CMM-6T1	WS-SVC-CMM-6E1	WS-SVC-CMM-ACT
Cisco IOS Release	12.3(8)XY	12.3(8)XY	12.3(8)XY	12.3(8)XY
CatOS Release	7.3(1)	7.3(1)	7.3(1)	7.6.8
Native IOS Release	12.1(15)E	12.1(14)E	12.1(13)E	12.1(13)E
Maximum number of port adapters per CMM	3	3	3	4

WS-X6624-FXS Analog Interface Module

The Cisco WS-X6624-FXS analog interface module is an MGCP-based device for connecting high-density analog devices to the IP telephony network, and it provides 24 analog ports.



Note

The WS-X6624 FXS analog interface module is no longer available for sale.

Cisco VG224 Gateway

The Cisco VG224 analog gateway is a Cisco IOS high-density 24-port gateway for connecting analog devices to the IP Telephony network. In Cisco IOS Release 12.4(2)T and later, the Cisco VG224 can act as a Skinny Client Control Protocol (SCCP), Media Gateway Control Protocol (MGCP), or H.323 endpoint with Cisco Unified CallManager and re-home to a Survivable Remote Site Telephone (SRST) router in failover scenarios. The Cisco VG224 supports Cisco Unified CallManager Release 3.1 and later. The Cisco VG224 also supports modem pass-through, modem relay, fax pass-through, and fax relay.

Cisco VG248 Gateway

The Cisco VG248 is a high-density, 48 port, Skinny Client Control Protocol (SCCP) gateway for connecting analog devices such as analog phones, fax machines, modems and speakerphones to an enterprise Cisco Unified CallManager (Release 3.1 and later) and voice network. The Cisco VG248 also supports Cisco Unified CallManager integration with legacy voicemail systems and PBXs compatible with Simplified Message Desk Interface (SMDI), NEC Message Center Interface (MCI), or Ericsson voicemail protocols. The Cisco VG248 supports failover to Survivable Remote Site Telephone (SRST).

Cisco ATA 186 and 188

The Cisco Analog Telephone Adaptor (ATA) 186 or 188 can connect two analog devices to the IP telephony network, and it is the best suited for low-density analog devices connecting to the IP network.

The difference between the Cisco ATA 186 and 188 is that the former has only one 10 Base-T Ethernet connection while the later has an integrated Ethernet switch providing two 10/100 Base-T Ethernet connections for itself and a co-located PC or other Ethernet-based device. Cisco Unified CallManager 4.2 has native SCCP support for the Cisco ATA 186 or 188.

The Cisco ATA 186 and 188 can be configured in any of the following ways:

- Cisco ATA web configuration page
- Cisco ATA voice configuration menu
- Configuration file downloaded from the TFTP server

Cisco Unified IP Phones

The Cisco IP phone portfolio includes basic IP phones, business IP phones, manager IP Phones, and executive IP phones. For a complete list of supported features of Cisco basic IP phone models, see the [Endpoint Features Summary, page 21-35](#).

Cisco Basic IP Phones

The Cisco basic IP phone is best suited for low-traffic users with limited call features and budget requirements. The Cisco basic IP phones include Cisco Unified IP Phone 7902G, 7905G, 7906G, 7911G, and 7912G.

**Note**

The initial version of the Cisco Unified IP Phone 7912G is no longer available for sale. The initial version of the Cisco Unified IP Phone 7912G is now replaced by the Cisco Unified IP Phone 7912G-A, which offers identical features but has an enhanced Ethernet switch.

Cisco Business IP Phones

The Cisco business IP phone is best suited for the transaction-type worker with medium telephony traffic use and extensive call features, such as speakers, headset, and so forth. The business IP phones include Cisco Unified IP Phone 7940G, 7941G, and 7941G-GE.

Cisco Manager IP Phones

The Cisco manager IP phone is best suited for managers and administrative assistants with medium to heavy telephony traffic use and extensive call features such as speakers, headset, and so forth. The business IP phones include Cisco Unified IP Phone 7960G, 7961G, and 7961G-GE.

Cisco Executive IP Phones

The Cisco executive IP phone is best suited for the executive high-traffic user with extensive call features. The executive IP phones include the Cisco Unified IP Phone 7970G and 7971G-GE.

**Note**

In addition to using the inline power from the access switch or local wall power, a Cisco Unified IP Phone can also be supplied power by a power injector, Promax. Promax connects Cisco Unified IP Phones to Cisco switches that do not support inline power or to non-Cisco switches. Promax is compatible with all Cisco Unified IP Phones, and it supports both Cisco PoE and IEEE 802.3af PoE. It has two 10/100/1000 Base-T Ethernet ports. One Ethernet port connects to the switch access port and the other connects to the Cisco Unified IP Phone.

Cisco Unified IP Phone Expansion Module 7914

The Cisco Unified IP Phone Expansion Module 7914 is for administrative assistants and others who need to determine the status of a number of lines beyond the current line capability of the phone.

The Cisco Unified IP Phone Expansion Module 7914 extends the capability of the Cisco Unified IP Phone 7960G, 7961G, 7961G-GE, 7970G, or 7971G-GE with additional buttons and an LCD. The Cisco Unified IP Phone Expansion Module 7914 provides 14 buttons per module, and the Cisco Unified IP Phones 796xG and 797xG can support up to two Cisco Unified IP Phone Expansion Modules. If the IP phone uses Cisco inline power or IEEE802.3af PoE, then the Cisco Unified IP Phone Expansion Module 7914 requires the use of an external power adaptor (CP-PWR-CUBE-3).

Software-Based Endpoints

Software-based endpoints include Cisco IP Communicator and Cisco IP SoftPhone. A software-based endpoint is an application installed on a client PC, and it registers with (and is controlled by) Cisco Unified CallManager.

Cisco IP Communicator

Cisco IP Communicator is a Microsoft Windows-based application that endows computers with the functionality of IP phones. This application enables high-quality voice calls on the road, in the office, or from wherever users can access the corporate network. It is an ideal solution for remote users and telecommuters. Cisco IP Communicator is easy to deploy and features some of the latest technology and advancements available with IP communications today. This section summarizes the following design considerations that apply when using Cisco IP Communicator with Cisco Unified CallManager:

- [Maximum IP Communicator Configuration Limits, page 21-8](#)
- [Codec Selection, page 21-8](#)
- [Call Admission Control, page 21-8](#)


Maximum IP Communicator Configuration Limits

Because Cisco IP Communicator is an SCCP standalone device, the design guidelines for IP phones in the various IP Telephony deployment models still hold true for the Cisco IP Communicator. Refer to the chapter on [IP Telephony Deployment Models, page 2-1](#), for details.

Codec Selection

The Cisco IP Communicator supports G.711 and G.729a codecs. The codec selection can be done by configuring the region in which Cisco IP Communicator is located. Cisco recommends G.729a low-bandwidth codec configurations in deployments with telecommuters connecting their Cisco IP Communicator across the WAN. The Cisco IP Communicator also has a low-bandwidth codec overriding capability in a G.711 region, which can be enabled by checking the Optimize for Low Bandwidth option in the Audio setting window (see [Figure 21-1](#)). In this case, the Cisco IP Communicator will use a G.729 codec to set up a call with another phone in the same region. Cisco IP Communicator can make video calls when it is used together with Cisco Unified Video Advantage. For more information on the video codec support of Cisco Unified Video Advantage, see the section on [Video Endpoints, page 21-16](#).

Figure 21-1 Cisco IP Communicator Audio Setting



Call Admission Control

Call admission control ensures that there is enough bandwidth available to process IP phone calls over the network. While there are several mechanisms for implementing call admission control, the Cisco IP Communicator uses the *locations* mechanism configured in Cisco Unified CallManager for centralized call processing deployments. Refer to the chapter on [Call Admission Control, page 9-1](#), for more information on call admission control with Cisco Unified CallManager locations.

With Cisco Unified CallManager 4.2, the Cisco IP Communicator can automatically update its location setting as it moves from one location to another. Cisco Unified CallManager performs the call admission control based on the device's current location setting and subtracts a certain amount of bandwidth for the call from the current location's bandwidth pool. For more information on device mobility, refer to the chapter on [Device Mobility](#), page 22-1.

Cisco IP SoftPhone

This section summarizes the following design considerations that apply when using Cisco IP SoftPhone with Cisco Unified CallManager:


- [Maximum Cisco IP SoftPhone Configuration Limits](#), page 21-10
- [Codec Selection](#), page 21-10
- [Call Admission Control](#), page 21-11

The information in this section applies explicitly to Cisco IP SoftPhone Release 1.3. For specific details about Cisco IP SoftPhone configuration and features, refer to the *Cisco IP Softphone Administrator Guide (1.3)*, available online at

http://www.cisco.com/en/US/products/sw/voicesw/ps1860/tsd_products_support_eol_series_home.html

Figure 21-2 illustrates that the Cisco IP SoftPhone application can either monitor or control the hardware IP phone associated with it. In the case of a third-party controlled phone, the Cisco IP SoftPhone acts as a virtual extension of the desktop IP phone. The Cisco IP SoftPhone application is able to see and handle incoming and outgoing calls for the hardware phone. From the perspective of provisioning devices and CTI resources, each user with this configuration would be provisioned as a third-party controlled IP phone. As a CTI port, the Cisco IP SoftPhone is a dedicated line to process calls directly to the client machine without the additional control and monitoring of a desktop phone.

Figure 21-2 Cisco IP SoftPhone Device Association Options



The Cisco IP SoftPhone is not able to run a CTI port and a third-party controlled phone with the same directory number (DN) at the same time. As shown in [Figure 21-2](#), the user is able to have extension 8110 as either a CTI port or a control for the desktop phone.

For more information on device and resource provisioning, refer to the chapter on [Call Processing](#), [page 8-1](#).

Maximum Cisco IP SoftPhone Configuration Limits

Regardless of the device limits allowed per server, there are maximum limits on the number CTI devices you can configure in Cisco Unified CallManager. The CTI device limits as they apply to Cisco IP SoftPhone are as follows:

- Maximum of 800 Cisco IP SoftPhones per Cisco Media Convergence Server (MCS) 7825 or 7835; maximum of 3,200 Cisco IP SoftPhones per cluster of MCS 7825s or 7835s
- Maximum of 2,500 Cisco IP SoftPhones per Cisco Media Convergence Server (MCS) 7845; maximum of 10,000 Cisco IP SoftPhones per cluster of MCS 7845s


The following assumptions apply to the preceding maximum Cisco IP SoftPhone limits:

- Each Cisco IP SoftPhone is configured with one line appearance.
- Each Cisco IP SoftPhone is processing an estimated six or fewer busy hour call attempts (BHCA).
- No other CTI applications requiring CTI devices are configured in the Cisco Unified CallManager cluster.

Codec Selection

Cisco IP SoftPhone supports G.711, G.723, and G.729a codecs. Cisco recommends G.729a low-bandwidth codec configurations in deployments with telecommuters connecting their Cisco IP SoftPhones across the WAN.

Because Cisco Unified CallManager does not support the G.723 codec, Cisco IP Softphone has two bandwidth codec settings to use: G.711 is the default setting, with a user-configurable option to select the low-bandwidth codec setting of G.729 on the TAPI Service Provider (TSP) client (see [Figure 21-3](#)). Refer to the chapter on [Network Infrastructure](#), [page 3-1](#), for details on provisioning network bandwidth.

Figure 21-3 Cisco IP Softphone Audio Setting

Cisco IP SoftPhone users with low-bandwidth connections across the WAN should consider selecting this low-bandwidth G.729 codec setting.

Call Admission Control

Call admission control ensures that there is enough bandwidth available to process IP phone calls over the network. While there are several mechanisms for implementing call admission control, the Cisco IP SoftPhone uses the *locations* mechanism configured in Cisco Unified CallManager for centralized call processing deployments. Refer to the chapter on [Call Admission Control, page 9-1](#), for more information on call admission control with Cisco Unified CallManager locations.

With Cisco Unified CallManager 4.2, the Cisco IP SoftPhone can automatically update its location setting as it moves from one location to another. Cisco Unified CallManager performs the call admission control based on the device's current location setting and subtracts a certain amount of bandwidth for the call from the current location's bandwidth pool. For more information on device mobility, refer to the chapter on [Device Mobility, page 22-1](#).

Wireless Endpoints

Cisco wireless endpoints use a wireless LAN (WLAN) infrastructure via wireless access points (APs) to provide telephony functionality and features. This type of endpoint is ideal for environments with the need for mobile users within a area where traditional wired phones are undesirable or problematic. (Refer to [Wireless LAN Infrastructure](#), page 3-58, for more information about wireless network design.)

The Cisco Unified Wireless IP Phone 7920 is a hardware-based phone with a built-in radio antenna that enables 802.11b wireless LAN connectivity to the network. These phones register with Cisco Unified CallManager using Skinny Client Control Protocol (SCCP), just like the hardware-based phones and Cisco IP Communicator. For more information, refer to the *Cisco Unified Wireless IP Phone 7920 Design and Deployment Guide*, available at

<http://www.cisco.com/go/designzone>

Site Survey

Before deploying the Cisco Unified Wireless IP Phone 7920, you must perform a complete site survey to determine the appropriate number and location of APs required to provide radio frequency (RF) coverage. Your site survey should take into consideration which types of antennas will provide the best coverage, as well as where sources of RF interference might exist. A site survey requires the use of the Site Survey tool on the Cisco Unified Wireless IP Phone 7920 (accessed via **Menu > Network Config > Site Survey**) and the Aironet Client Utility Site Survey Tool used with a Cisco Aironet NIC card on a laptop or PC. Additional third-party tools can also be used for site surveys; however, Cisco highly recommends that you conduct a final site survey using the Cisco Unified Wireless IP Phone 7920 because each endpoint or client radio can behave differently depending on antenna sensitivity and survey application limitations.

Authentication

To connect to the wireless network, the Cisco Unified Wireless IP Phone 7920 must first use one of the following authentication methods to associate and communicate with the AP:

- Extensible Authentication Protocol-Flexible Authentication via Secure Tunneling (EAP-FAST)

This method allows the Cisco Unified Wireless IP Phone 7920 to be authenticated to the AP via 802.1X with a user name and password once a secure authenticated tunnel is established between the client and an EAP-compliant Remote Authentication, Authorization, and Accounting server via Protected Access Credential (PAC). Upon authentication, traffic to and from the wireless device is encrypted using TKIP or WEP. Using the 802.1X authentication method and the PAC authenticated tunnel exchange requires an EAP-compliant Remote Authentication Dial-In User Service (RADIUS) authentication server such as the Cisco Secure Access Control Server (ACS), which provides access to a user database.

- Wi-Fi Protected Access (WPA)

This method allows the Cisco Unified Wireless IP Phone 7920 to be authenticated to the AP via 802.1X with a user name and password. Upon authentication, traffic to and from the wireless device is encrypted using Temporal Key Integrity Protocol (TKIP). Using the 802.1X authentication method requires an EAP-compliant Remote Authentication Dial-In User Service (RADIUS) authentication server such as the Cisco Secure Access Control Server (ACS), which provides access to a user database.

- Wi-Fi Protected Access Pre-Shared Key (WPA-PSK)

This method allows the Cisco Unified Wireless IP Phone 7920 to be authenticated to the AP via the configuration of a shared key on the Cisco Unified Wireless IP Phone 7920 and the AP. Upon authentication, traffic to and from the wireless device is encrypted using TKIP. This method of authentication is not recommended for enterprise deployments.

- Cisco Centralized Key Management (Cisco CKM)

This method allows the Cisco Unified Wireless IP Phone 7920 to be authenticated to the AP via 802.1x with a user name and password. Upon authentication, traffic to and from the wireless device is encrypted using either WEP 128 or TKIP. The 802.1X authentication method requires an EAP-compliant RADIUS authentication server such as the Cisco ACS, which provides access to a user database for the initial authentication request. Subsequent authentication requests are validated via the wireless domain service (WDS) at the AP, which shortens re-authentication times and ensures fast, secure roaming.

- Cisco LEAP

This method allows the Cisco Unified Wireless IP Phone 7920 and AP to be authenticated mutually based on a user name and password. Upon authentication, the dynamic key is generated and used for encrypting traffic between the Cisco Unified Wireless IP Phone 7920 and the AP. A LEAP-compliant Radius authentication server, such as the Cisco Secure Access Control Server (ACS), is required to provide access to the user database.

- Shared Key

This method involves the configuration of static 10 (40-bit) or 26 (128-bit) character keys on the Cisco Unified Wireless IP Phone 7920 and the AP. This method is AP-based authentication in which access to the network is gained if the device has a matching key.

- Open Authentication

This method requires no exchange of identifying information between the Cisco Wireless IP Phone 7920 and the AP. Cisco does *not* recommend this method because it provides no secure exchange of voice or signaling, and it allows any rogue device to associate to the AP.

Capacity

Each AP can support a maximum of seven active G.711 voice streams or eight G.729 streams. If these numbers are exceeded, poor quality can result due to dropped or delayed voice packets or dropped calls. AP rates set lower than 11 Mbps will result in lower call capacity per AP.



Note

A call between two phones associated to the same AP counts as two active voice streams.

Based on these active call capacity limits, and using Erlang ratios, you can calculate the number of Cisco Unified Wireless IP Phone 7920s that each AP can support. For example, given a typical user-to-call capacity ratio of 3:1, a single AP can support 21 to 24 Cisco Unified Wireless IP Phone 7920s, depending on whether the codec used is G.711 or G.729. However, this number does not take into consideration the possibility that other Cisco Unified Wireless IP Phone 7920s could roam to this AP, so a lower number of phones per AP is more realistic.

The number of APs per VLAN or Layer 2 subnet should also be considered. To optimize memory and performance on the APs, Cisco recommends deploying no more than 30 APs on a single VLAN or subnet. This recommendation, taken with typical user-to-call capacity ratios, limits the number of Cisco Unified Wireless IP Phone 7920s per Layer 2 subnet to approximately 500 (or about 15 to 17 Cisco Unified Wireless IP Phone 7920s per AP).

These capacities were calculated with voice activity detection (VAD) disabled and a packetization sample size of 20 milliseconds (ms). VAD is a mechanism for conserving bandwidth by not sending RTP packets while no speech is occurring during the call. However, enabling or disabling VAD is a global cluster-wide configuration parameter on Cisco Unified CallManager. (It is referred to as Silence Suppression in Cisco Unified CallManager.) Thus, if VAD is enabled for the Cisco Unified Wireless IP Phone 7920, then it will be enabled for all devices in the Cisco Unified CallManager cluster. Cisco recommends leaving VAD (Silence Suppression) *disabled* to provide better overall voice quality.

At a sampling rate of 20 ms, a voice call will generate 50 packets per second (pps) in either direction. Cisco recommends setting the sample rate to 20 ms for almost all cases. By using a larger sample size (for example, 30 or 40 ms), you can increase the number of simultaneous calls per AP, but a larger end-to-end delay will result. In addition, the percentage of acceptable voice packet loss within a wireless environment decreases dramatically with a larger sample size because more of the conversation is missing when a packet is lost. For more information about voice sampling size, see [Bandwidth Provisioning, page 3-44](#).

Phone Configuration

You can configure the Cisco Unified Wireless IP Phone 7920 either through the phone's keypad or with the 7920 Configuration Utility running on a PC that is attached to the phone via a USB cable. In either case, you must configure the following parameters:

- **Network Configuration**

Configure either the DHCP server address or static settings such as IP address, subnet mask, default gateway, TFTP server, and DNS server, as appropriate for the network. These settings can be found on the Cisco Unified Wireless IP Phone 7920 under **Menu > Profiles > Network Profile**.

- **Wireless Configuration**

Configure the service set identifier (SSID) for the voice VLAN and the authentication type, including the WEP key and/or LEAP user name and password when appropriate. These settings can be found on the Cisco Unified Wireless IP Phone 7920 under **Menu > Profiles > Network Profile**.

Roaming

Currently the Cisco Unified Wireless IP Phone 7920 is able to roam at Layer 2 (within the same VLAN or subnet) and still maintain an active call.

Layer 2 roaming occurs in the following situations:

- During the initial boot-up of the Cisco Unified Wireless IP Phone 7920, the phone roams to a new AP for the first time.
- If the Cisco Unified Wireless IP Phone 7920 receives no beacons or responses from the AP to which it is currently associated, the phone assumes that the current AP is unavailable and it attempts to roam and associate with a new AP.

- The Cisco Unified Wireless IP Phone 7920 maintains a list of eligible AP roam targets. If conditions change on the current AP, the phone consults the list of available AP roam targets. If one of the roam targets is determined to be a better choice, then the phone attempts to roam to the new AP.
- If the configured SSID or authentication type on the Cisco Unified Wireless IP Phone 7920 is changed, the phone must roam to re-associate with an AP.

In trying to determine eligible AP roam targets for Layer 2 roaming, the wireless IP phone uses the following variables to determine the best AP to associate with:

- **Relative Signal Strength Indicator (RSSI)**
Used by the wireless IP phone to determine the signal strength and quality of available APs within an RF coverage area. The phone will attempt to associate with the AP that has the highest RSSI value and matching authentication/encryption type.
- **QoS Basic Service Set (QBSS)**
Enables the AP to communicate channel utilization information to the wireless phone. The phone will use the QBSS value to determine if it should attempt to roam to another AP, because APs with high channel utilization might not be able to handle VoIP traffic effectively.

Layer 2 roaming times for the wireless IP phone depend on the type of authentication used. If static WEP keys are used for authentication between the phone and the AP, Layer 2 roaming occurs in less than 100 ms. If LEAP (with local Cisco Secure ACS authentication) is used, Layer 2 roaming occurs in 200 to 400 ms. Using Cisco Centralized Key Management (Cisco CKM) can reduce roaming time to less than 100 ms.

When devices roam at Layer 3, they move from one AP to another AP across native VLAN boundaries. The Cisco Catalyst 6500 Series Wireless LAN Services Module (WLSM) allows the Cisco Unified Wireless IP Phone 7920 to roam at Layer 3 while still maintaining an active call. The Cisco Wireless IP Phone 7920 can roam at Layer 3 by using Static WEP or Cisco CKM protocols. Cisco CKM enables the phone to achieve full Layer 3 mobility while using either WEP 128 or TKIP encryption. Seamless Layer 3 roaming occurs only when the client is roaming within the same mobility group. For details about the Cisco WLSM and Layer 3 roaming, refer to the product documentation available at

<http://www.cisco.com>

If you are using 802.1x authentication in the wireless LAN, Cisco CKM is recommended to minimize roaming downtime. Whether the roaming is at Layer 2 or Layer 3, device downtime decreases from 300-400 ms to 100 ms or less. Cisco CKM also takes some of the load off the ACS server by reducing the number of authentication requests that must be sent to the ACS.

AP Call Admission Control

Call admission control mechanisms in Cisco Unified CallManager or in a gatekeeper can control WAN bandwidth utilization and provide QoS for existing calls, but both mechanisms are applied only at the beginning of a call. For calls between static devices, this type of call admission control is sufficient. However, for a call between two mobile wireless devices such as the Cisco Unified Wireless IP Phone 7920, there must also be a call admission control mechanism at the AP level because the wireless devices may roam from one AP to another.

The AP mechanism for call admission control is QBSS, which is the beacon information element that enables the AP to communicate channel utilization information to the wireless IP phone. As previously mentioned, this QBSS value helps the phone determine whether it should roam to another AP. A lower QBSS value indicates that the AP is a good candidate to roam to, while a higher QBSS value indicates that the phone should not roam to this AP.

While this QBSS information is useful, it does not guarantee that calls will retain proper QoS during roaming. When a Cisco Unified Wireless IP Phone 7920 is associated to an AP with a high QBSS, the AP will prevent a call from being initiated or received by rejecting the call setup and sending a Network Busy message to the initiating phone. However, once a call is set up between a wireless IP phone and another endpoint, the phone may roam and associate with an AP with a high QBSS, thus resulting in an oversubscription of the available bandwidth on that AP.

Cisco IP Conference Station

The Cisco IP Conference Station combines conference room speaker-phone technology with Cisco IP Communications technology. The Cisco IP Conference Station is best suited for use in conferencing environments providing 360-degree room coverage.

The Cisco Unified IP Conference Station 7936 has an external speaker and three built-in microphones. The Cisco Unified IP Conference Station 7936 requires Cisco CallManager Release 3.3(3) SR3 or later. The Cisco Unified IP Conference Station 7936 also features a pixel-based LCD display with backlighting, and optional extension microphones can be connected to it for extended microphone coverage in larger rooms.

Video Endpoints

Cisco Unified CallManager Release 4.2 supports the following types of video-enabled endpoints:

- Cisco Unified Video Advantage associated with a Cisco Unified IP Phone 7940, 7941, 7960, 7961, 7970, or 7971 running Skinny Client Control Protocol (SCCP)
- Cisco IP Video Phone 7985
- Tandberg 2000 MXP, 1500 MXP, 1000 MXP, 770 MXP, 550 MXP, T-1000, or T-550 models running SCCP
- Sony PCS-1, PCS-TL30, or PCS-TL50 models running SCCP
- H.323 clients (Polycom, Sony, PictureTel, Tandberg, VCON, VTEL, Microsoft NetMeeting, and others)

SCCP Video Endpoints

SCCP video endpoints register directly with Cisco Unified CallManager and download their configurations via Trivial File Transfer Protocol (TFTP). They support many features and supplementary services, including hold, transfer, conference, park, pickup and group pickup, music on hold, shared line appearances, mappable softkeys, call forwarding (busy, no answer, and unconditional), and much more.

Cisco Unified Video Advantage

Cisco Unified Video Advantage is a Windows-based application and USB camera that you can install on a Windows 2000 or Windows XP personal computer. When the PC is physically connected to the PC port on a Cisco Unified IP Phone 7940, 7941, 7960, 7961, 7970, or 7971 running the Skinny Client Control Protocol, the Cisco Unified Video Advantage application "associates" with the phone, thus enabling

users to operate their phones as they always have but now with the added benefit of video. In Cisco Unified Video Advantage Release 2.0, this association can also be to Cisco IP Communicator running on the same PC.

The system administrator can control which IP Phones allow this association to take place by toggling the **Video Capabilities: Enabled/Disabled** setting on the IP Phone configuration page in Cisco Unified CallManager Administration. When this feature is enabled, an icon representing a camera appears in the bottom right-hand corner of the IP Phone display. By default, Cisco Unified Video Advantage is disabled. You can also use the Bulk Administration Tool to modify this setting on many phones at once. Note that the **PC Port: Enabled/Disabled** setting must also be enabled for Cisco Unified Video Advantage to work with a hardware IP Phone; however, the **PC Access to Voice VLAN** setting does not have to be enabled.

To achieve the association with a hardware IP Phone, Cisco Unified Video Advantage installs a Cisco Discovery Protocol (CDP) driver onto the Ethernet interface of the PC. CDP enables the PC and the hardware IP Phone to discover each other automatically, which means that the user does not have to configure anything on the PC or the hardware IP Phone in order for Cisco Unified Video Advantage to work. The user can, therefore, plug the PC into any hardware IP Phone that is video-enabled and automatically associates with it. (See Figure 21-4.)

Cisco Unified Video Advantage 2.0 does not rely on CDP to discover the presence of Cisco IP Communicator running on the same PC. Instead, it listens for a private Windows message sent from the IP Communicator process. Once IP Communicator is discovered, the association process works exactly like it does for a hardware IP Phone. (See Figure 21-5.)



Note

When you install Cisco Unified Video Advantage, the CDP packet drivers install on all Ethernet interfaces of the PC. If you add a new network interface card (NIC) or replace an old NIC with a new one, you must reinstall Cisco Unified Video Advantage so that the CDP drivers also install on the new NIC.

Figure 21-4 Cisco Unified Video Advantage Operational Overview




Figure 21-4 illustrates the following events:

1. The IP Phone and PC exchange Cisco Discovery Protocol (CDP) messages. The phone begins listening for PC association packets on TCP port 4224 from the IP address of its CDP neighbor.

2. The PC initiates association messages to the phone over TCP/IP. Association packets are routed up to the Layer-3 boundary between VLANs. Firewalls and/or access control lists (ACLs) must permit TCP port 4224.
3. The phone acts as an SCCP proxy between Cisco Unified Video Advantage and Cisco Unified CallManager. Cisco Unified CallManager tells the phone to open video channels for the call, and the phone proxies those messages to the PC.
4. The phone sends/receives audio, and the PC sends/receives video. Both audio and video traffic are marked DSCP AF41. Video traffic uses UDP port 5445.

Figure 21-5 Cisco IP Communicator Associating with Cisco Unified Video Advantage




Figure 21-5 illustrates the following events:

1. Cisco IP Communicator sends a private Windows message to Cisco Unified Video Advantage. The message includes the IP address of Cisco IP Communicator and the port number for CAST messages.
2. Cisco Unified Video Advantage initiates CAST messages to Cisco IP Communicator over TCP/IP. CAST messages do not leave the PC because it is a connected address.
3. Cisco IP Communicator acts as an SCCP proxy between Cisco Unified Video Advantage and Cisco Unified CallManager. Cisco Unified CallManager tells IP Communicator to open video channels for the call, and IP Communicator proxies those messages to Cisco Unified Video Advantage via CAST protocol.
4. Cisco IP Communicator sends/receives audio, and Cisco Unified Video Advantage sends/receives video. Audio traffic is marked DSCP EF, and video traffic is marked DSCP 0. The switch port must be set to use an access control list (ACL) to re-mark the traffic instead of trusting the DSCP values, otherwise Cisco Unified Video Advantage packets will not have QoS.

When a call is made using Cisco Unified Video Advantage, the audio is handled by the IP Phone while the video is handled by the PC. There is no synchronization mechanism between the two devices, so QoS is essential to minimize jitter, latency, fragmented packets, and out-of-order packets.

If a hardware IP Phone is used, the IP Phone resides in the voice VLAN, while the PC resides in the data VLAN, which means that there must be a Layer-3 routing path between the voice and data VLANs in order for the association to occur. If there are access control lists (ACLs) or firewalls between these

VLANs, they must be configured to permit the association protocol (which uses TCP port 4224 in both directions) to pass. If Cisco IP Communicator is used, this communication happens internal to the PC, and there are no Layer-3 boundaries to cross.

Cisco Unified Video Advantage supports Differentiated Services Code Point (DSCP) traffic classifications. Cisco Unified CallManager specifies the DSCP value in the SCCP messages it sends to the phone. When the IP Phone makes an audio-only call, it marks its SCCP control traffic as DSCP CS3 and its audio RTP media traffic as DSCP EF. However, when the IP Phone makes a video call, it marks its SCCP control traffic as DSCP CS3 and its audio RTP media traffic as DSCP AF41, and the Cisco Unified Video Advantage application marks its video RTP media traffic as DSCP AF41 as well. Both the IP Phone and the Cisco Unified Video Advantage application mark their "association" protocol messages as DSCP CS3 because it is considered to be signaling traffic and is grouped with all other signaling traffic such as SCCP.

**Note**

Cisco Unified CallManager Release 4.0 added security features to the Cisco Unified IP Phone 7970 and 7971 to enable it to use Transport Layer Security (TLS) and Secure RTP (SRTP) to authenticate and encrypt signaling and audio media traffic. The association protocol does not use this authentication or encryption, nor are the video RTP media streams encrypted. However, the SCCP signaling and the audio RTP media streams are encrypted if they are so configured.

**Note**

Do not set the voice VLAN equal to the data VLAN because doing so can cause issues with connectivity.

Cisco Unified Video Advantage, like any other application that runs on a PC, does have an impact on system performance, which you should take into consideration. Cisco Unified Video Advantage 1.0 supports two types of video codecs: H.263 and the Cisco VT Camera Wideband Video Codec. Cisco Unified Video Advantage 2.0 also supports two types of codecs: H.263 and H.264. The Cisco VT Camera Wideband Video Codec places the least demand on the PC but the most demand on the network. H.263 places a lesser demand on the network but a higher demand on the PC. Finally, H.264 places the least demand on the network but the highest demand on the PC. Therefore, if your network has plenty of available bandwidth, you can use the Cisco VT Camera Wideband Video Codec and save on PC CPU and memory resources.

The H.263 and H.264 codecs support a range of speeds up to 1.5 Mbps. In summary, customers must balance PC performance with network utilization when deploying Cisco Unified Video Advantage.

System Requirements

For detailed PC requirements, refer to the *Cisco Unified Video Advantage Data Sheet*, available at

http://www.cisco.com/en/US/prod/collateral/video/ps7190/ps5662/product_data_sheet0900aecd8044de04.html

Cisco IP Video Phone 7985G

The Cisco IP Video Phone 7985G is a personal desktop video phone. Unlike Cisco Unified Video Advantage, which is an application that runs on a PC, the Cisco IP Video Phone 7985G is a standalone phone with integrated video features. The phone has an 8.4 inch color LCD screen and an embedded video camera for making video calls. The phone supports up to eight line appearances, has two 10/100 BaseT Ethernet connections, and has buttons for Directories, Messages, Settings, and Services. Like other Cisco Unified IP Phones, the Cisco IP Video Phone 7985G uses CDP to learn VLAN and CoS information from the attached switch to use in 802.1p/q markings.

Codecs Supported by Cisco Unified Video Advantage and Cisco IP Video Phone 7985G

Table 21-4 lists the codecs supported by Cisco Unified Video Advantage and Cisco IP Video Phone 7985G

Table 21-4 **Codecs Supported by Cisco Unified Video Advantage and Cisco IP Video Phone 7985G**

Codec or Feature	Cisco Unified Video Advantage	Cisco IP Video Phone 7985G
H.264	Yes in Release 2.0	Yes
H.263	Yes	Yes
H.261	No	Yes
G.711	Yes	Yes
G.722	No	Yes
G.722.1	No	No
G.723.1	No	No
G.728	No	No
G.729	Yes	Yes
Cisco Wideband	Yes in Release 1.0	No
Maximum Bandwidth	7 Mbps for Release 1.0 and 1.5 Mbps for Release 2.0	768 kbps
Video Resolution	CIF, QCIF	NTSC: 4SIF, SIF PAL: 4CIF, QCIF, SQCIF

Third-Party SCCP Video Endpoints

Two manufacturers of video endpoints, Sony and Tandberg, currently have the following products that support the Cisco Skinny Client Control Protocol (SCCP):

- Sony PCS-1 and PCS-TL50
- Tandberg T-1000 and T-550

SCCP on both the Sony and Tandberg endpoints is modeled after SCCP on the Cisco Unified IP Phone 7940. Most features found on the Cisco Unified IP Phone 7940 user interface are also supported on the Sony endpoints as well as the Tandberg endpoints, including multiple line appearances, softkeys, and buttons for Directories, Messages, Settings, Services, and so forth. The Sony and Tandberg endpoints also support the Option 150 field in DHCP to discover the IP address of the TFTP server, and they download their configurations from the TFTP server. However, software upgrades of the Sony and Tandberg endpoints are not done via TFTP. Instead, the customer must manually upgrade each endpoint using tools provided by the vendor. (Tandberg uses an FTP method, while Sony uses FTP or a physical memory stick.) The Sony and Tandberg endpoints register with up to three Cisco Unified CallManager servers and will fail-over to its secondary or tertiary servers if its primary server becomes unreachable.

While the Sony and Tandberg endpoints support softkey functionality similar to that of the Cisco Unified IP Phone 7940 and 7960, the exact feature support differs between vendor and model. Check the manufacturers' documentation for supported features. Features that are currently known to be missing on some platforms include:

- Messages button
- Directories (placed calls, received calls, missed calls, and corporate directory)
- Settings and Services buttons
- Some XML services (such as Extension Mobility and Berbee's InformaCast)

Because the Sony and Tandberg endpoints use SCCP, dialing a video call from an endpoint is similar to dialing an audio call from a Cisco Unified IP Phone. If users are familiar with Cisco Unified IP Phones, they should also find the Sony and Tandberg endpoints very intuitive to use. The main difference in the user interface is that the Sony and Tandberg endpoints do not have a button keypad or a handset like those on a phone. Instead, the remote control is used to access features and to dial numbers on the Sony and Tandberg endpoints.

**Note**

Sony and Tandberg endpoints do not support the Cisco Discovery Protocol (CDP) or IEEE 802.Q/p. Therefore, you must manually configure the VLAN ID and Quality of Service trust boundary on the ethernet switch to which they are attached. (For more details, see [Network Infrastructure, page 3-1.](#))

Codecs Supported by Sony and Tandberg SCCP Endpoints

Codec support for the third-party SCCP endpoints varies by vendor, model, and software version. Check the vendors' product documentation for the supported codecs.

QoS Recommendations

This sections provides the basic QoS guidelines and configurations for the Cisco Catalyst switches most commonly deployed with IP Telephony endpoints. For more details, refer to the *Quality of Service* design guide at

<http://www.cisco.com/go/designzone>

Cisco VG224 and VG248

Analog gateways are trusted endpoints. For Cisco VG224 and VG248 gateways, configure the switch to trust the DSCP value of the VG248 packets. The following sections list the commands to configure the most common Cisco Catalyst switches for the Cisco VG224 and VG248 analog gateways.

**Note**

In the following sections, *vlan_id* is the voice VLAN ID and *dvlan_id* is the data VLAN ID.

Cisco 2950

```
CAT2950(config)#interface interface-id
CAT2950(config-if)#mls qos trust dscp
CAT2950(config-if)#switchport mode access
CAT2950(config-if)#switchport access vlan vlan_id
```

**Note**

The **mls qos trust dscp** command is available only with Enhanced Image (EI).

Cisco 2970 or 3750

```
CAT2970(config)#mls qos
CAT2970(config)#interface interface-id
CAT2970(config-if)#mls qos trust dscp
CAT2970(config-if)#switchport mode access
CAT2970(config-if)#switchport access vlan vvlan_id
```

Cisco 3550

```
CAT3550(config)#mls qos
CAT3550(config)#interface interface-id
CAT3550(config-if)#mls qos trust dscp
CAT3550(config-if)#switchport mode access
Cat3550(config-if)#switchport access vlan vvlan_id
```

Cisco 4500 with SUP3, IV, or V

```
CAT4500(config)#qos
CAT4500(config)#interface interface-id
CAT4500(config-if)#qos trust dscp
CAT4500(config-if)#switchport mode access
CAT4500(config-if)#switchport access vlan vvlan_id
```

Cisco 6500

```
CAT6500>(enable)set qos enable
CAT6500>(enable)set port qos 2/1 vlan-based
CAT6500>(enable)set vlan vvlan_id mod/port
CAT6500>(enable)set port qos mod/port trust trust-dscp
```

Cisco ATA 186 and IP Conference Station

Because the Cisco Analog Telephone Adaptor (ATA) 186 and IP Conference Station are trusted endpoints, their QoS configurations are identical to those described in the section on [Cisco VG224 and VG248, page 21-21](#).

Cisco ATA 188 and IP Phones

For the Cisco Analog Telephone Adaptor (ATA) 188 and IP Phones, Cisco recommends segregating the voice VLAN from the data VLAN. For the Cisco ATA 186, 7902, 7905, 7906, 7910, and IP Conference Station, Cisco still recommends configuring voice and data VLAN segregation and an auxiliary voice VLAN. In this way, the same access-layer configurations can be used with different IP phone models and ATAs, and end-users can plug their IP phones or ATAs into different access ports on the switch and get the same treatment. For the Cisco ATA 186, 7902, 7905, 7906, 7910, and IP Conference Stations, the command to override the CoS value of the frames from the attached PC has no effects because these devices do not have a PC connected to them.

The following sections list the configuration commands for IP phones on the most commonly deployed Cisco Catalyst switches.

Cisco 2950

```

CAT2950(config)#
CAT2950(config)#class-map VVLAN
CAT2950(config-cmap)# match access-group name VVLAN
CAT2950(config-cmap)#class-map VLAN
CAT2950(config-cmap)# match access-group name DVLAN
CAT2950(config-cmap)#exit
CAT2950(config)#
CAT2950(config)#policy-map IPPHONE-PC
CAT2950(config-pmap)# class VVLAN
CAT2950(config-pmap-c0# set ip dscp 46
CAT2950(config-pmap-c)# police 1000000 8192 exceed-action-drop
CAT2950(config-pmap)# class DVLAN
CAT2950(config-pmap-c0# set ip dscp 0
CAT2950(config-pmap-c)# police 5000000 8192 exceed-action-drop
CAT2950(config-pmap-c)#exit
CAT2950(config-pmap)#exit
CAT2950(config)#
CAT2950(config)#interface interface-id
CAT2950(config-if)#mls qos trust device cisco-phone
CAT2950(config-if)#mls qos trust cos
CAT2950(config-if)#switchport mode access
CAT2950(config-if)#switchport voice vlan vvlan_id
CAT2950(config-if)#switchport access vlan dvlan_id
CAT2950(config-if)#service-policy input IPPHONE-PC
CAT2950(config-if)#exit
CAT2950(config)#
CAT2950(config)#ip access-list standard VVLAN
CAT2950(config-std-nacl)# permit voice_IP_subnet wild_card_mask
CAT2950(config-std-nacl)#exit
CAT2950(config)#ip access-list standard DVLAN
CAT2950(config-std-nacl)# permit data_IP_subnet wild_card_mask
CAT2950(config-std-nacl)#end

```

**Note**

The **mls qos map cos-dscp** command is available only with Enhanced Image (EI). With Standard Image (SI), this command is not available and the default CoS-to-DSCP mapping is as follows:

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

Cisco 2970 or 3750

```

CAT2970(config)# mls qos map cos-dscp 0 8 16 24 34 46 48 56
CAT2970(config)# mls qos map policed-dscp 0 24 to 8
CAT2970(config)#
CAT2970(config)#class-map match-all VVLAN-VOICE
CAT2970(config-cmap)# match access-group name VVLAN-VOICE
CAT2970(config-cmap)#
CAT2970(config-cmap)#class-map match-all VVLAN-CALL-SIGNALING
CAT2970(config-cmap)# match access-group name VVLAN-CALL-SIGNALING
CAT2970(config-cmap)#
CAT2970(config-cmap)#class-map match-all VVLAN-ANY
CAT2970(config-cmap)# match access-group name VVLAN-ANY
CAT2970(config-cmap)#
CAT2970(config-cmap)# policy-map IPPHONE-PC
CAT2970(config-pmap)#class VVLAN-VOICE
CAT2970(config-pmap-c)# set ip dscp 46
CAT2970(config-pmap-c)# police 128000 8000 exceed-action drop

```

```

CAT2970(config-pmap-c)# class VVLAN-CALL-SIGNALING
CAT2970(config-pmap-c)# set ip dscp 24
CAT2970(config-pmap-c)# police 32000 8000 exceed-action policed-dscp-transmit
CAT2970(config-pmap-c)# class VVLAN-ANY
CAT2970(config-pmap-c)# set ip dscp 0
CAT2970(config-pmap-c)# police 32000 8000 exceed-action policed-dscp-transmit
CAT2970(config-pmap-c)# class class-default
CAT2970(config-pmap-c)# set ip dscp 0
CAT2970(config-pmap-c)# police 5000000 8000 exceed-action policed-dscp-transmit
CAT2970(config-pmap-c)# exit
CAT2970(config-pmap)# exit
CAT2970(config)#
CAT2970(config)#
CAT2970(config)#interface interface-id
CAT2970(config-if)# switchport voice vlan vvlan_id
CAT2970(config-if)# switchport access vlan dvlan_id
CAT2970(config-if)# mls qos trust device cisco-phone
CAT2970(config-if)# service-policy input IPPHONE-PC
CAT2970(config-if)# exit
CAT2970(config)#
CAT2970(config)#ip access list extended VVLAN-VOICE
CAT2970(config-ext-nacl)# permit udp Voice_IP_Subnet Subnet_Mask any range 16384 32767
dscp ef
CAT2970(config-ext-nacl)# exit
CAT2970(config)#ip access list extended VVLAN-CALL-SIGNALING
CAT2970(config-ext-nacl)# permit tcp Voice_IP_Subnet Subnet_Mask any range 2000 2002 dscp
cs3
CAT2970(config-ext-nacl)# permit tcp Voice_IP_Subnet Subnet_Mask any range 2000 2002 dscp
Af31
CAT2970(config-ext-nacl)# exit
CAT2970(config)#ip access list extended VVLAN-ANY
CAT2970(config-ext-nacl)# permit ip Voice_IP_Subnet Subnet_Mask any
CAT2970(config-ext-nacl)# end
CAT2970#

```

Cisco 3550

```

CAT3550(config)# mls qos map cos-dscp 0 8 16 24 34 46 48 56
CAT3550(config)# mls qos map policed-dscp 0 24 26 46 to 8
CAT3550(config)#class-map match-all VOICE
CAT3550(config-cmap)# match ip dscp 46
CAT3550(config-cmap)#class-map match-any CALL SIGNALING
CAT3550(config-cmap)# match ip dscp 26
CAT3550(config-cmap)# match ip dscp 24
CAT3550(config-cmap)#
CAT3550(config-cmap)#class-map match-all VVLAN-VOICE
CAT3550(config-cmap)# match vlan vvlan_id
CAT3550(config-cmap)# match class-map VOICE
CAT3550(config-cmap)#
CAT3550(config-cmap)#class-map match-all VVLAN-CALL-SIGNALING
CAT3550(config-cmap)# match vlan vvlan_id
CAT3550(config-cmap)# match class-map CALL SIGNALING
CAT3550(config-cmap)#
CAT3550(config-cmap)#class-map match-all ANY
CAT3550(config-cmap)# match access-group name ACL_Name
CAT3550(config-cmap)#
CAT3550(config-cmap)# class-map match-all VVLAN-ANY
CAT3550(config-cmap)# match vlan vvlan_id
CAT3550(config-cmap)# match class-map ANY
CAT3550(config-cmap)#
CAT3550(config-cmap)#class-map match-all DVLAN-ANY
CAT3550(config-cmap)# match vlan dvlan_id
CAT3550(config-cmap)# match class-map ANY

```

```

CAT3550(config-cmap)#
CAT3550(config-cmap)#policy-map IPPHONE-PC
CAT3550(config-pmap-c)# class VVLAN-VOICE
CAT3550(config-pmap-c)# set ip dscp 46
CAT3550(config-pmap-c)# police 128000 8000 exceed-action drop
CAT3550(config-pmap-c)#
CAT3550(config-pmap-c)#class VVLAN-CALL-SIGNALING
CAT3550(config-pmap-c)# set ip dscp 24
CAT3550(config-pmap-c)# police 32000 8000 exceed-action policed-dscp-transmit
CAT3550(config-pmap-c)#
CAT3550(config-pmap-c)#class VVLAN-ANY
CAT3550(config-pmap-c)# set ip dscp 0
CAT3550(config-pmap-c)# police 32000 8000 exceed-action policed-dscp-transmit
CAT3550(config-pmap-c)#
CAT3550(config-pmap-c)#class DVLAN-VOICE
CAT3550(config-pmap-c)# set ip dscp 0
CAT3550(config-pmap-c)# police 5000000 8000 exceed-action policed-dscp-transmit
CAT3550(config-pmap-c)#exit
CAT3550(config-pmap)#exit
CAT3550(config)#interface interface-id
CAT3550(config-if)# switchport voice vlan vvlan_id
CAT3550(config-if)# switchport access vlan dvlan_id
CAT3550(config-if)# mls qos trust device cisco-phone
CAT3550(config-if)# service-policy input IPPHONE-PC
CAT3550(config-if)# exit
CAT3550(config)#
CAT3550(config)#ip access list standard ACL_ANY
CAT3550(config-std-nacl)# permit any
CAT3550(config-std-nacl)# end
CAT3550#

```

Cisco 4500 with SUP3, IV, or V

```

CAT4500(config)# qos map cos 5 to dscp 46
CAT4500(config)# qos map cos 0 24 26 46 to dscp 8
CAT4500(config)#
CAT4500(config)#class-map match-all VVLAN-VOICE
CAT4500(config-cmap)# match access-group name VVLAN-VOICE
CAT4500(config-cmap)#
CAT4500(config-cmap)#class-map match-all VVLAN-CALL-SIGNALING
CAT4500(config-cmap)# match access-group name VVLAN-CALL-SIGNALING
CAT4500(config-cmap)#
CAT4500(config-cmap)#class-map match-all VVLAN-ANY
CAT4500(config-cmap)# match access-group name VVLAN-ANY
CAT4500(config-cmap)#
CAT4500(config-cmap)# policy-map IPPHONE-PC
CAT4500(config-pmap)#class VVLAN-VOICE
CAT4500(config-pmap-c)# set ip dscp 46
CAT4500(config-pmap-c)# police 128 kps 8000 byte exceed-action drop
CAT4500(config-pmap-c)# class VVLAN-CALL-SIGNALING
CAT4500(config-pmap-c)# set ip dscp 24
CAT4500(config-pmap-c)# police 32 kps 8000 byte exceed-action policed-dscp-transmit
CAT4500(config-pmap-c)# class VVLAN-ANY
CAT4500(config-pmap-c)# set ip dscp 0
CAT4500(config-pmap-c)# police 32 kps 8000 byte exceed-action policed-dscp-transmit
CAT4500(config-pmap-c)# class class-default
CAT4500(config-pmap-c)# set ip dscp 0
CAT4500(config-pmap-c)# police 5 mpbs 8000 byte exceed-action policed-dscp-transmit
CAT4500(config-pmap-c)# exit
CAT4500(config-pmap)# exit
CAT4500(config)#
CAT4500(config)#
CAT4500(config)#interface interface-id

```

```

CAT4500(config-if)# switchport voice vlan vvlan_id
CAT4500(config-if)# switchport access vlan dvlan_id
CAT4500(config-if)# qos trust device cisco-phone
CAT4500(config-if)# service-policy input IPPHONE-PC
CAT4500(config-if)# exit
CAT4500(config)#
CAT4500(config)#ip access list extended VVLAN-VOICE
CAT4500(config-ext-nacl)# permit udp Voice_IP_Subnet Subnet_Mask any range 16384 32767
dscp ef
CAT4500(config-ext-nacl)# exit
CAT4500(config)#ip access list extended VVLAN-CALL-SIGNALING
CAT4500(config-ext-nacl)# permit tcp Voice_IP_Subnet Subnet_Mask any range 2000 2002 dscp
cs3
CAT4500(config-ext-nacl)# permit tcp Voice_IP_Subnet Subnet_Mask any range 2000 2002 dscp
Af31
CAT4500(config-ext-nacl)# exit
CAT4500(config)#ip access list extended VVLAN-ANY
CAT4500(config-ext-nacl)# permit ip Voice_IP_Subnet Subnet_Mask any
CAT4500(config-ext-nacl)# end
CAT4500#

```

Cisco 6500

```

CAT6500> (enable) set qos cos-dscp-map 0 8 16 24 32 46 48 56
CAT6500> (enable) set qos policed-dscp-map 0, 24, 26, 46:8
CAT6500> (enable)
CAT6500> (enable) set qos policer aggregate VVLAN-VOICE rate 128 burst 8000 drop
CAT6500> (enable) set qos policer aggregate VVLAN-CALL-SIGNALING rate 32 burst 8000
policed-dscp
CAT6500> (enable) set qos policer aggregate VVLAN-ANY rate 5000 burst 8000 policed-dscp
CAT6500> (enable) set qos policer aggregate PC-DATA rate 5000 burst 8000 policed-dscp
CAT6500> (enable)
CAT6500> (enable) set qos acl ip IPPHONE-PC dscp 46 aggregate VVLAN-VOICE udp
Voice_IP_Subnet Subnet_Mask any range 16384 32767
CAT6500> (enable) set qos acl ip IPPHONE-PC dscp 24 aggregate VVLAN-CALL-SIGNALING tcp
Voice_IP_Subnet Subnet_Mask any range 2000 2002
CAT6500> (enable) set qos acl ip IPPHONE-PC dscp 0 aggregate VVLAN-ANY Voice_IP_Subnet
Subnet_Mask any
CAT6500> (enable) set qos acl ip IPPHONE-PC dscp 0 aggregate PC-DATA any
CAT6500> (enable) commit qos acl IPPHONE-PC
CAT6500> (enable) set vlan vvlan_id mod/port
CAT6500> (enable) set port qos mod/port trust-device ciscoipphone
CAT6500> (enable) set qos acl map IPPHONE-PC mod/port
CAT6500> (enable)

```

Software-Based Endpoints

Even though Cisco IP Communicator and IP SoftPhone mark their signaling packets and media packets respectively, Cisco recommends that you re-mark the DSCP value of the packets from the PC that is running Cisco IP Communicator or IP SoftPhone because the PC is not a trusted device in the network. Instead of using the full range of User Datagram Protocol (UDP) ports (16384 to 32767) for the media packets, Cisco IP Communicator and IP SoftPhone can be configured to use a specific UDP port.

Cisco Unified Video Advantage 2.0 brings video functionality to Cisco IP Communicator, therefore Cisco IP Communicator is able to make video calls. Cisco recommends that you configure an access control list (ACL) to match the video packets from Cisco Unified Video Advantage on UDP port 5445, and re-mark the traffic with DSCP value AF41. Unlike Cisco Unified IP Phones, Cisco IP Communicator

does not automatically mark its voice packets with DSCP value AF41 when a video call is made; instead, it marks its voice packets with DSCP value EF. Hence, the voice stream of a video call will be marked with DSCP value EF while the video stream of the call will be marked with DSCP value AF41.

For Cisco IP Communicator, you can specify the UDP port by using one of the following options:

- Specify the **RTP Port Range Start** and **RTP Port Range End** in the product-specific section of the IP Communicator configuration page.
- Select **Preferences > Audio Settings > Network > Port Range** and specify the port range.

If you use both options to set the UDP port and port range, settings from the second option will override the first option.

For Cisco IP SoftPhone, you can specify the UDP port and port range under **Network Audio Settings > Audio Output Port**.

The following sections list the QoS configuration commands for Cisco IP Communicator and IP SoftPhone on the most commonly deployed Cisco Catalyst switches.

Cisco 2950 with Enhanced Image

The Cisco Catalyst 2950 Series switches are not recommended for software-based endpoint QoS implementations due to two limitations:

- The Cisco 2950 does not support the **range** keyword to specify the UDP port range within an ACL configuration. The workaround to this limitation is to configure a single static UDP port for the Cisco IP SoftPhone to use, as described in preceding section.
- The Cisco 2950 supports only 1-Mbps increments on FastEthernet ports, which can create a fairly large hole to admit onto the network unauthorized traffic that might be mimicking call signaling or media.

Cisco 2970 or 3750

```

CAT2970 (config) #mls qos
CAT2970 (config) #mls qos map policed-dscp 0 24 26 46 to 8
CAT2970 (config) #
CAT2970 (config) #class-map match-all COMMUNICATOR-VOICE
CAT2970 (config-cmap) # match access-group name COMMUNICATOR-VOICE
CAT2970 (config-cmap) #class-map match-all COMMUNICATOR-VIDEO
CAT2970 (config-cmap) # match access-group name COMMUNICATOR-VIDEO
CAT2970 (config-cmap) #class-map match-all COMMUNICATOR-SIGNALING
CAT2970 (config-cmap) # match access-group name COMMUNICATOR-SIGNALING
CAT2970 (config-cmap) #exit
CAT2970 (config) #
CAT2970 (config) #policy-map COMMUNICATOR-PC
CAT2970 (config-pmap-c) #class COMMUNICATOR-VOICE
CAT2970 (config-pmap-c) # set ip dscp 46
CAT2970 (config-pmap-c) # police 128000 8000 exceed-action drop
CAT2970 (config-pmap-c) #class COMMUNICATOR-VIDEO
CAT2970 (config-pmap-c) # set ip dscp 34
CAT2970 (config-pmap-c) # police 5000000 8000 exceed-action policed-dscp-transmit
CAT2970 (config-pmap-c) #class COMMUNICATOR-SIGNALING
CAT2970 (config-pmap-c) # set ip dscp 24
CAT2970 (config-pmap-c) # police 32000 8000 exceed-action policed-dscp-transmit
CAT2970 (config-pmap-c) #class class-default
CAT2970 (config-pmap-c) # set ip dscp 0
CAT2970 (config-pmap-c) # police 5000000 8000 exceed-action policed-dscp transmit
CAT2970 (config-pmap-c) # exit
CAT2970 (config-pmap) #exit
CAT2970 (config) #
CAT2970 (config) #interface FastEthernet interface-id
CAT2970 (config-if) # switchport access vlan vvlan_id

```

```

CAT2970 (config-if) # switchport mode access
CAT2970 (config-if) # service-policy input COMMUNICATOR-PC
CAT2970 (config-if) # exit
CAT2970 (config) #ip access list extended COMMUNICATOR-VOICE
CAT2970 (config-ext-nacl) # permit udp host PC_IP_address eq fixed_port_number any
CAT2970 (config-ext-nacl) # exit
CAT2970 (config) #ip access list extended COMMUNICATOR-VIDEO
CAT2970 (config-ext-nacl) # permit udp host PC_IP_address any eq 5445
CAT2970 (config-ext-nacl) # exit
CAT2970 (config) #ip access-list extended COMMUNICATOR-SIGNALING
CAT2970 (config-ext-nacl) # permit tcp host PC_IP_address host CallManager_IP_address eq
2748 or 2000
CAT2970 (config-ext-nacl) # exit

```

Cisco 3550

```

CAT3550 (config) #mls qos
CAT3550 (config) #mls qos map cos-dscp 0 8 16 24 34 46 48 56
CAT3550 (config) #mls qos map policed-dscp 0 24 26 34 46 to 8
CAT3550 (config) #
CAT3550 (config) #class-map match-all COMMUNICATOR-VOICE
CAT3550 (config-cmap) # match access-group name COMMUNICATOR-VOICE
CAT3550 (config-cmap) #class-map match-all COMMUNICATOR-VIDEO
CAT3550 (config-cmap) # match access-group name COMMUNICATOR-VIDEO
CAT3550 (config-cmap) #class-map match-all COMMUNICATOR-SIGNALING
CAT3550 (config-cmap) # match access-group name COMMUNICATOR-SIGNALING
CAT3550 (config-cmap) #exit
CAT3550 (config) #
CAT3550 (config) #policy-map COMMUNICATOR -PC
CAT3550 (config-pmap-c) #class COMMUNICATOR -VOICE
CAT3550 (config-pmap-c) # set ip dscp 46
CAT3550 (config-pmap-c) # police 128000 8000 exceed-action drop
CAT3550 (config-pmap-c) #class COMMUNICATOR-VIDEO
CAT3550 (config-pmap-c) # set ip dscp 34
CAT3550 (config-pmap-c) # police 5000000 8000 exceed-action policed-dscp-transmit
CAT3550 (config-pmap-c) #class COMMUNICATOR-SIGNALING
CAT3550 (config-pmap-c) # set ip dscp 24
CAT3550 (config-pmap-c) # police 32000 8000 exceed-action policed-dscp-transmit
CAT3550 (config-pmap-c) #class class-default
CAT3550 (config-pmap-c) # set ip dscp 0
CAT3550 (config-pmap-c) # police 5000000 8000 exceed-action policed-dscp transmit
CAT3550 (config-pmap-c) # exit
CAT3550 (config-pmap) #exit
CAT3550 (config) #
CAT3550 (config) #interface FastEthernet interface-id
CAT3550 (config-if) # switchport access vlan vvlan_id
CAT3550 (config-if) # switchport mode access
CAT3550 (config-if) # service-policy input COMMUNICATOR-PC
CAT3550 (config-if) # exit
CAT3550 (config) #ip access list extended COMMUNICATOR-VOICE
CAT3550 (config-ext-nacl) # permit udp host PC_IP_address eq fixed_port_number any
CAT3550 (config-ext-nacl) # exit
CAT3550 (config) #ip access list extended COMMUNICATOR-VIDEO
CAT3550 (config-ext-nacl) # permit udp host PC_IP_address any eq 5445
CAT3550 (config-ext-nacl) # exit
CAT3550 (config) #ip access-list extended COMMUNICATOR-SIGNALING
CAT3550 (config-ext-nacl) # permit tcp host PC_IP_address host CallManager_IP_address eq
2748 or 2000
CAT3550 (config-ext-nacl) # exit

```

Cisco 4500 with SUP3, IV, or V

```

CAT4500(config) #qos
CAT4500(config) #qos map policed-dscp 0 24 26 34 46 to 8
CAT4500(config) #
CAT4500(config) #class-map match-all COMMUNICATOR-VOICE
CAT4500(config-cmap) # match access-group name COMMUNICATOR-VOICE
CAT4500(config-cmap) #class-map match-all COMMUNICATOR-VIDEO
CAT4500(config-cmap) # match access-group name COMMUNICATOR-VIDEO
CAT4500(config-cmap) #class-map match-all COMMUNICATOR-SIGNALING
CAT4500(config-cmap) # match access-group name COMMUNICATOR-SIGNALING
CAT4500(config-cmap) #exit
CAT4500(config) #
CAT4500(config) #policy-map COMMUNICATOR-PC
CAT4500(config-pmap-c) #class COMMUNICATOR-VOICE
CAT4500(config-pmap-c) # set ip dscp EF
CAT4500(config-pmap-c) # police 128 kps 8000 byte exceed-action drop
CAT4500(config-pmap-c) #class COMMUNICATOR-VIDEO
CAT4500(config-pmap-c) # set ip dscp AF41
CAT4500(config-pmap-c) # police 5 mbps 8000 byte exceed-action drop
CAT4500(config-pmap-c) #class COMMUNICATOR-SIGNALING
CAT4500(config-pmap-c) # set ip dscp CS3
CAT4500(config-pmap-c) # police 32000 kps 8000 byte exceed-action policed-dscp-transmit
CAT4500(config-pmap-c) #class class-default
CAT4500(config-pmap-c) # set ip dscp default
CAT4500(config-pmap-c) # police 5 mbps 8000 byte exceed-action policed-dscp transmit
CAT4500(config-pmap-c) # exit
CAT4500(config-pmap) #exit
CAT4500(config) #
CAT4500(config) #interface FastEthernet interface-id
CAT4500(config-if) # switchport access vlan vvlan_id
CAT4500(config-if) # switchport mode access
CAT4500(config-if) # service-policy input COMMUNICATOR-PC
CAT4500(config-if) # exit
CAT4500(config) #ip access list extended COMMUNICATOR-VOICE
CAT4500(config-ext-nacl) # permit udp host PC_IP_address eq fixed_port_number any
CAT4500(config-ext-nacl) # exit
CAT4500(config) #ip access list extended COMMUNICATOR-VIDEO
CAT4500(config-ext-nacl) # permit udp host PC_IP_address any eq 5445
CAT4500(config-ext-nacl) # exit
CAT4500(config) #ip access-list extended COMMUNICATOR-SIGNALING
CAT4500(config-ext-nacl) # permit tcp host PC_IP_address host CallManager_IP_address eq
2748 or 2000
CAT4500(config-ext-nacl) # exit

```

Cisco 6500

```

CAT6500> (enable) set qos enable
CAT6500> (enable) set qos policed-dscp-map 0, 24, 26, 34, 46:8
CAT6500> (enable)
CAT6500> (enable) set qos policer aggregate COMMUNICATOR-VOICE rate 128 burst 8000 drop
CAT6500> (enable) set qos policer aggregate COMMUNICATOR-VIDEO rate 5000 burst 8000
policed-dscp
CAT6500> (enable) set qos policer aggregate COMMUNICATOR-SIGNALING rate 32 burst 8000
policed-dscp
CAT6500> (enable) set qos policer aggregate PC-DATA rate 5000 burst 8000 policed-dscp
CAT6500> (enable)
CAT6500> (enable) set qos acl ip COMMUNICATOR-PC dscp 46 aggregate COMMUNICATOR-VOICE udp
host PC_IP_address eq fixed_port_number any
CAT6500> (enable) set qos acl ip COMMUNICATOR-PC dscp 34 aggregate COMMUNICATOR-VIDEO udp
host PC_IP_address any eq 5445
CAT6500> (enable) set qos acl ip COMMUNICATOR-PC dscp 24 aggregate COMMUNICATOR-SIGNALING
tcp host PC_IP_address host CallManager_IP_address eq 2748 or 2000

```

```

CAT6500> (enable) set qos acl ip COMMUNICATOR-PC dscp 0 aggregate PC-DATA any
CAT6500> (enable) commit qos acl COMMUNICATOR-PC
CAT6500> (enable) set vlan vvlan_id mod/port
CAT6500> (enable) set port qos mod/port trust untrusted
CAT6500> (enable) set qos acl map COMMUNICATOR-PC mod/port
CAT6500> (enable)

```

**Note**

The DSCP re-marking must be done by a Layer 3 capable switch. If the access layer switch (such as the Cisco Catalyst 2950 with Standard Image or the Cisco 3524XL) does not have this capability, then the DSCP re-marking must be done at the distribution layer switch.

Cisco Unified Wireless IP Phone 7920

By default, the Cisco Unified Wireless IP Phone 7920 marks its SCCP signaling messages using a Per-Hop Behavior (PHB) value of CS3 or a Differentiated Services Code Point (DSCP) value of 24 (this corresponds to a ToS value of 0x60), and it marks RTP voice packets using a PHB value of EF or a DSCP value of 46 (ToS of 0xB8). With proper queueing on the AP and configuration on the upstream first-hop switch to trust the AP's port, the wireless IP phone traffic will receive the same treatment as wired IP phone traffic. This practice allows the QoS settings to be consistent from LAN to WLAN environments.

In addition, the Cisco Unified Wireless IP Phone 7920 will automatically announce its presence to the AP using the Cisco Discovery Protocol (CDP). The CDP packets are sent from the wireless IP phone to the AP, and they identify the phone so that the AP can place all traffic to the phone in the high-priority queue.

While Ethernet switch ports can typically transmit and receive at 100 Mbps, 802.11b APs have a lower throughput rate that allows for a maximum data rate of 11 Mbps. Furthermore, wireless LANs are a shared medium and, due to contention for this medium, the actual throughput is substantially lower. This throughput mismatch means that, with a burst of traffic, the AP will drop packets, thus adding excessive processor burden to the unit and affecting performance.

By taking advantage of the policing and rate limiting capabilities of the Cisco Catalyst 3550 and 6500 Series switches, you can eliminate the need for the AP to drop excessive packets by configuring the upstream switch port to rate-limit or police traffic going to the AP. The switch port configurations in the following sections rate-limit the port(s) to a practical throughput of 7 Mbps for 802.11b and guarantee 1 Mbps for high-priority voice and control traffic. Furthermore, as indicated in the configuration examples, packets coming from the AP should be trusted and, based on the VLAN tag of each packet, the DSCP marking should either be maintained or marked down. Thus, packets sourced from the Cisco Unified Wireless IP Phone 7920 on the voice VLAN should maintain the appropriate DSCP marking, and packets source from data devices on the data VLAN should be remarked to a DSCP value of 0.

Cisco 3550

```

CAT3550(config)#mls qos
CAT3550(config)#mls qos map cos-dscp 0 8 16 24 32 46 48 56
CAT3550(config)#mls qos map policed-dscp 24 46 to 8
CAT3550(config)#mls qos aggregate-policer AGG-POL-1M-VOICE-OUT 1000000 8000 exceed-action
policed-dscp-transmit
CAT3550(config)#mls qos aggregate-policer AGG-POL-6M-DEFAULT-OUT 6000000 8000
exceed-action drop
CAT3550(config)#
CAT3550(config)#class-map match-all EGRESS-DSCP-0
CAT3550(config-cmap)#match ip dscp 0
CAT3550(config-cmap)#

```



```

CAT3550(config-cmap)#class-map match-all EGRESS-DSCP-8
CAT3550(config-cmap)#match ip dscp 8
CAT3550(config-cmap)#
CAT3550(config-cmap)#class-map match-all EGRESS-DSCP-16
CAT3550(config-cmap)#match ip dscp 16
CAT3550(config-cmap)#
CAT3550(config-cmap)#class-map match-all EGRESS-DSCP-32
CAT3550(config-cmap)#match ip dscp 32
CAT3550(config-cmap)#
CAT3550(config-cmap)#class-map match-all EGRESS-DSCP-48
CAT3550(config-cmap)#match ip dscp 48
CAT3550(config-cmap)#
CAT3550(config-cmap)#class-map match-all EGRESS-DSCP-56
CAT3550(config-cmap)#match ip dscp 56
CAT3550(config-cmap)#
CAT3550(config-cmap)#class-map match-all VOICE-SIGNALING
CAT3550(config-cmap)#match ip dscp 24
CAT3550(config-cmap)#
CAT3550(config-cmap)#class-map match-all VOICE
CAT3550(config-cmap)#match ip dscp 46
CAT3550(config-cmap)#
CAT3550(config-cmap)#class-map match-all INGRESS-DATA
CAT3550(config-cmap)#match any
CAT3550(config-cmap)#
CAT3550(config-cmap)#class-map match-all INGRESS-VVLAN-VOICE
CAT3550(config-cmap)#match vlan vvlan-id
CAT3550(config-cmap)#match class-map VOICE
CAT3550(config-cmap)#
CAT3550(config-cmap)#class-map match-all INGRESS-VVLAN-VOICE-SIGNALING
CAT3550(config-cmap)#match vlan vvlan-id
CAT3550(config-cmap)#match class-map VOICE-SIGNALING
CAT3550(config-cmap)#
CAT3550(config-cmap)#class-map match-all INGRESS-DVLAN
CAT3550(config-cmap)#match vlan dvlan-id
CAT3550(config-cmap)#match class-map INGRESS-DATA
CAT3550(config-cmap)#
CAT3550(config-cmap)#policy-map EGRESS-RATE-LIMITER
CAT3550(config-pmap)#class EGRESS-DSCP-0
CAT3550(config-pmap-c)#police aggregate AGG-POL-6M-DEFAULT-OUT
CAT3550(config-pmap-c)#
CAT3550(config-pmap-c)#class EGRESS-DSCP-8
CAT3550(config-pmap-c)#police aggregate AGG-POL-6M-DEFAULT-OUT
CAT3550(config-pmap-c)#
CAT3550(config-pmap-c)#class EGRESS-DSCP-16
CAT3550(config-pmap-c)#police aggregate AGG-POL-6M-DEFAULT-OUT
CAT3550(config-pmap-c)#
CAT3550(config-pmap-c)#class EGRESS-DSCP-32
CAT3550(config-pmap-c)#police aggregate AGG-POL-6M-DEFAULT-OUT
CAT3550(config-pmap-c)#
CAT3550(config-pmap-c)#class EGRESS-DSCP-48
CAT3550(config-pmap-c)#police aggregate AGG-POL-6M-DEFAULT-OUT
CAT3550(config-pmap-c)#class EGRESS-DSCP-56
CAT3550(config-pmap-c)#police aggregate AGG-POL-6M-DEFAULT-OUT
CAT3550(config-pmap-c)#
CAT3550(config-pmap-c)#class EGRESS-VOICE
CAT3550(config-pmap-c)#police aggregate AGG-POL-1M-VOICE-OUT
CAT3550(config-pmap-c)#
CAT3550(config-pmap-c)#class EGRESS-VOICE-SIGNALING
CAT3550(config-pmap-c)#police aggregate AGG-POL-1M-VOICE-OUT
CAT3550(config-pmap-c)#
CAT3550(config-pmap-c)#policy-map INGRESS-QOS
CAT3550(config-pmap)#class INGRESS-VVLAN-VOICE
CAT3550(config-pmap-c)#set ip dscp 46
CAT3550(config-pmap-c)#

```

```

CAT3550(config-pmap-c)#class INGRESS-VVLAN-CALL-SIGNALING
CAT3550(config-pmap-c)#set ip dscp 24
CAT3550(config-pmap-c)#
CAT3550(config-pmap-c)#class INGRESS-DVLAN
CAT3550(config-pmap-c)#set ip dscp 0
CAT3550(config-pmap-c)#
CAT3550(config-pmap-c)#class class-default
CAT3550(config-pmap-c)#set ip dscp 0
CAT3550(config-pmap-c)#
CAT3550(config-pmap-c)#interface interface id
CAT3550(config-if)#description 11Mb towards Wireless Access Point
CAT3550(config-if)#switchport access dvlan-id
CAT3550(config-if)#switchport voice vvlan-id
CAT3550(config-if)#mls qos trust dscp
CAT3550(config-if)#service-policy output EGRESS-RATE-LIMITER
CAT3550(config-if)#service-policy input INGRESS-QOS

```

Cisco 6500

```

CAT6500> (enable) set qos enable
CAT6500> (enable) set qos cos-dscp-map 0 8 16 24 32 46 48 56
CAT6500> (enable) set qos policed-dscp-map 24,46:8
CAT6500> (enable)
CAT6500> (enable) set qos policer microflow VOICE-OUT rate 1000 burst 32 policed-dscp
CAT6500> (enable) set qos policer microflow DATA-OUT rate 6000 burst 32 drop
CAT6500> (enable)
CAT6500> (enable) set qos acl ip AP-VOICE-EGRESS dscp 24 microflow VOICE-OUT ip any any
dscp-field 24
CAT6500> (enable) set qos acl ip AP-VOICE-EGRESS dscp 46 microflow VOICE-OUT ip any any
dscp-field 46
CAT6500> (enable) set qos acl ip AP-DATA-EGRESS dscp 0 microflow DATA-OUT ip any any
CAT6500> (enable)
CAT6500> (enable) set qos acl ip AP-VOICE-INGRESS trust-dscp ip any any
CAT6500> (enable) set qos acl ip AP-DATA-INGRESS dscp 0 ip any any
CAT6500> (enable)
CAT6500> (enable) set qos acl map AP-VOICE-EGRESS vvlan-id output
CAT6500> (enable) set qos acl map AP-DATA-EGRESS dvlan-id output
CAT6500> (enable) set qos acl map AP-VOICE-INGRESS vvlan-id input
CAT6500> (enable) set qos acl map AP-DATA-INGRESS dvlan-id input
CAT6500> (enable)
CAT6500> (enable) set port qos mod/port vlan-based
CAT6500> (enable)
CAT6500> (enable) set port qos mod/port trust trust-dscp
CAT6500> (enable)

```

Video Telephony Endpoints

This section discusses how the following types of endpoint devices classify traffic:

- [Cisco Unified Video Advantage with a Cisco Unified IP Phone](#), page 21-33
- [Cisco IP Video Phone 7985G](#), page 21-33
- [Sony and Tandberg SCCP Endpoints](#), page 21-34
- [H.323 Video Endpoints](#), page 21-34

Cisco Unified Video Advantage with a Cisco Unified IP Phone

The Cisco Unified Video Advantage application residing on the user's PC supports the classification of video packets using DSCP and, therefore, only at Layer 3. The current best practices for Cisco Unified Communications design recommend that the upstream Ethernet switch to which the phone is attached should be configured to trust the 802.1p CoS from the phone. Because the PC packets are unlikely to have an 802.1Q tag, they are unable to support 802.1p CoS bits. This lack of 802.1p support from the PC leaves the following possible options for providing QoS for Cisco Unified Video Advantage:

Option 1

If your current QoS model extends trust to the IP Phone, then the voice and signaling packets will be correctly marked as they ingress the network. With an additional ACL on the port to match UDP port 5445, the video media channel will also be classified to PHB AF41. Without this ACL, the video media would be classified Best Effort and would incur poor image quality and lip-sync issues. The same ACL could also be used to match the CAST connection between the Cisco Unified Video Advantage PC and the IP Phone, which uses TCP port 4224 (classifying it as CS3), although the benefit of doing so is minimal. The signaling packets from the PC, which is on the data VLAN, are returned over the same high-speed port onto the voice VLAN, therefore they are highly unlikely to encounter any congestion.

Option 2

The *Enterprise QoS Solution Reference Network Design Guide, Version 3.1* (available at <http://www.cisco.com/go/designzone>) presents another method. This alternative method recommends changing the port to trust the DSCP of incoming traffic instead of trusting CoS, and then running the incoming packets through a series of Per-Port/Per-VLAN Access Control Lists that match packets based on their TCP/UDP ports (along with other criteria) and police them to appropriate levels. For instance, Cisco Unified Video Advantage will mark its video packets with DSCP AF41, with the switch port set to trust DSCP. The packet will run through an ACL that matches it based on the fact that it is using UDP port 5445, is marked with DSCP AF41, and is coming in on the data VLAN. This ACL will then be used in a class map or policy map to trust the DSCP and police the traffic to N kbps (where N is the amount of video bandwidth you want to allow per port). Similar ACLs and policers will be present for the voice and signaling packets from the IP Phone in the voice VLAN.



Note

Option 2 should be used for Cisco Unified Video Advantage with Cisco IP Communicator because IP Communicator can mark only DSCP values and not CoS values.

Cisco IP Video Phone 7985G

Like many other Cisco Unified IP Phones, the Cisco IP Video Phone 7985G supports 802.1p/Q tagging for traffic originating from the phone and, because the Cisco IP Video Phone 7985G has a second Ethernet interface for PC access, traffic originating from attached devices as well. The current best practices for Cisco Unified Communications design recommend that the upstream Ethernet switch to which the phone is attached should be configured to trust the 802.1p CoS from the phone. Cisco recommends that trust not be extended to the PC port of the phone and, if the switch supports it, that you configure policers to limit the maximum amount of voice, video, and signaling traffic.

Sony and Tandberg SCCP Endpoints

Sony and Tandberg SCCP endpoints correctly mark their media and signaling packets at Layer 3 using DSCP. They do not, however, support 802.1Q and are therefore unable to classify using 802.1p CoS. If you use the UDP and TCP port-matching option, you would be able to classify the SCCP signaling correctly as CS3 and the video media as AF41; however, you would be unable to tell when a UDP port is being used in a voice-only call and should therefore be classified as EF. In such a case, the call admission control mechanisms would not be able to account for the bandwidth correctly. To avoid this situation, there is only one viable option for how to classify and trust traffic from a Sony or Tandberg endpoint:

Option 1

Trust DSCP on the port used by the Sony or Tandberg endpoint. If the switch allows it, configure policers to limit the maximum amount of EF, AF41, and CS3 traffic that can be received on that port. Any other device plugged into that port should not necessarily be trusted, even if its packets are classified using DSCP. This option may be acceptable if the Sony or Tandberg system is a permanent installation in an office or small conference room.

Because the Sony or Tandberg device does not support CDP, the VLAN placement of this endpoint requires manual modification if the requirement is to place it in the voice VLAN. The advantage of placing the endpoint directly in the voice VLAN is that it can be treated like any other IP Telephony endpoint in the system. The disadvantage is that the port might pose a security risk because it provides direct access to the voice VLAN. Alternatively, you can leave the Sony or Tandberg endpoint in the data VLAN, but you will have to provision access between the data and voice VLANs to permit SCCP signaling to Cisco Unified CallManager and to allow the UDP media streams to pass between the data and voice VLANs during voice or video calls.

H.323 Video Endpoints

This type of endpoint is potentially the most challenging from a QoS perspective due to the wide range of H.323 video endpoints, the variation in implementations, and the feature sets. There are two main QoS options for these endpoints; the first relies on the H.323 video endpoint to correctly mark all the traffic, and the second relies on detailed knowledge of the TCP and UDP ports used.

Option 1

If the endpoint correctly marks the media and signaling traffic (signaling should include H.225, H.245, and RAS), you could trust the classifications. Because it is unlikely that the endpoint supports 802.1Q (and therefore 802.1p CoS), you will probably have to use IP Precedence or DSCP in this case. The choice of classification type depends on the specific vendor, model, and software version.



Note

It is highly unlikely that an H.323 endpoint will mark its packets correctly.

Option 2

Using a combination of either source, destination, or both TCP and UDP port numbers (possibly including IP addresses as well), you could define an ACL that matches and classifies the traffic correctly. In addition, Cisco recommends that you also apply policers to limit the amount of each class of traffic that is admitted to the network. This option has the same potential as Option 1 for classifying voice-only calls incorrectly.

Endpoint Features Summary

The following tables summarize the features supported by the various endpoint devices discussed in this chapter:

- [Table 21-5](#) summarizes the Cisco Unified Communications features for Cisco analog gateways.
- [Table 21-6](#) summarizes the features for Cisco Basic IP Phones with SCCP protocol.
- [Table 21-7](#) summarizes the features for Cisco Business, Manager, and Executive IP Phones with SCCP protocol.
- [Table 21-8](#) summarizes the features for specialized endpoints including Cisco Unified IP Phones 7920, 7935, 7936G, and 7985G.
- [Table 21-9](#) summarizes the features for software-based devices including Cisco IP Communicator and Cisco IP SoftPhone.

Table 21-5 *Cisco Analog Gateway Features*

Feature	Analog Interface Cards	Ws-svc-cmm-24fxs	Ws-x6624-fxs	VG224	VG248	ATA 186 and 188
Ethernet Connection	N	N	N	Y ¹	Y ²	Y ³
Maximum number of Analog Ports	24 ⁴	72	24	24	48	2
Caller ID	Y	N	N	Y	Y	Y
Call Waiting	N	N	N	N	Y	Y
Caller ID on Call Waiting	N	N	N	N	Y	Y
Call Hold	N	N	N	Y ⁵	N	Y
Call Transfer	N	N	N	Y ⁵	Y	Y
Call Forward	N	N	N	N	Y ⁶	Y
Auto-Answer	N	N	N	N	N	N
Ad Hoc Conference	N	N	N	N	Y	Y
Meet-Me Conference	N	N	N	N	N	Y
Call Pickup	N	N	N	N	N	Y
Group Pickup	N	N	N	N	N	Y
Redial	N	N	N	N	Y ⁷	Y ⁷
Speed Dial	N	N	N	N	Y	Y
On-hook Dialing	N	N	N	N	N	N
Voice Mail Access	Y	Y	Y	Y	Y	Y ⁸
Message Waiting Indicator (MWI)	N	N	N	N	Y	Y ⁸
Survivable Remote Site Telephony (SRST) Support	N	N	N	Y	Y	Y
Music on Hold (MoH)	Y	Y	Y	N	Y	Y
Mute	N	N	N	N	N	N
Multilevel Precedence and Preemption (MLPP)	N	N	N	N	N	N

Table 21-5 Cisco Analog Gateway Features (continued)

Feature	Analog Interface Cards	Ws-svc-cmm-24fxs	Ws-x6624-fxs	VG224	VG248	ATA 186 and 188
Barge	N	N	N	N	N	N
cBarge	N	N	N	N	N	N
Call Preservation	N	N	N	N	Y ⁹	N
Call Admission Control	Y	N	N	N	N	N
Local Voice Busy-Out	Y	N	N	N	N	N
Private Line Automatic Ringdown (PLAR)	Y	N	N	N	N	Y
Hunt Group	Y	N	N	N	N	N
Dial Plan Mapping	Y	N	N	N	N	N
Supervisory Disconnect	Y	N	N	N	N	N
Signaling Packet ToS Value Marking	0x68	0x68 ¹⁰	0x68	0x68	0x68	0x68
Media Packet ToS Value Marking	0xB8	0xB8	0xB8	0xB8	0xB8	0xB8
Fax Pass-Through	Y ¹¹	Y	Y ¹²	Y	Y ¹¹	Y
Fax Relay	Y	Y	N	Y	Y	N
Skinny Client Control Protocol (SCCP)	N	N	N	N	Y	Y
Session Initiation Protocol (SIP)	N	N	N	Y	N	Y
H.323	Y	Y	N	Y	N	Y
Media Gateway Control Protocol (MGCP)	Y	Y	Y	Y	N	Y ¹³
G.711	Y	Y	Y	Y	Y	Y
G.722	N	N	N	N	N	N
G.723	Y	Y	N	N	N	Y
G.726	Y	N	N	N	N	N
G.729	Y	Y	Y	Y	Y	Y
Voice Activity Detection (VAD)	Y	Y	N	Y	N	Y
Comfort Noise Generation (CNG)	Y	Y	N	Y	N	Y

1. Two 10/100 Base-T.
2. One 10/100 Base-T.
3. Two 10/100 Base-T for ATA 188; one 10 Base-T for ATA 186.
4. The EVM-HD-8FXS/DID provides eight ports on the baseboard and can be configured for FXS or DID signaling. In addition, it has room for two EM-HDA-8FXS as extension modules.
5. H.323 and SIP call control.
6. Call Forward All.
7. Last Number Redial.
8. Only on SCCP and SIP version.
9. Supported on VG248 version 1.2 or later.
10. It marks MGCP signaling on UDP port 2427, but it marks the MGCP keep-alive packets as best-effort on TCP port 2428.
11. Fax pass-through and fax relay.

12. Fax pass-through.
13. Cisco Unified CallManager does not support MGCP with the ATA.

Table 21-6 Cisco Basic IP Phones with SCCP

Features	7902G	7905G	7906G	7910G	7910 +SW	7911G	7912G/G-A
Ethernet Connection	Y ¹	Y ¹	Y ²	Y ¹	Y ³	Y ³	Y ³
Ethernet Switch (PC port)	N	N	Y	N	Y	Y	Y ⁴
Cisco Power-Over-Ethernet (PoE)	Y	Y	Y	Y	Y	Y	Y
IEEE 802.3af Power-Over-Ethernet (PoE)	N	N	Y	N	N	Y	N
Localization	N	Y	Y	N	N	Y	Y
Directory Number	1	1	1	1	1	1	1
Maximum number of calls per line	200	200	200	200	200	200	200
Liquid Crystal Display	N	Y	Y	Y	Y	Y	Y
Caller ID	N	Y	Y	Y	Y	Y	Y
Call Waiting	N	Y	Y	Y	Y	Y	Y
Caller ID on Call Waiting	N	Y	Y	Y	Y	Y	Y
Call Hold	Y	Y	Y	Y	Y	Y	Y
Blind Transfer	N	N	N	N	N	N	N
Early-attended Transfer	Y	Y	Y	Y	Y	Y	Y
Consultative Transfer	Y	Y	Y	Y	Y	Y	Y
Call Forward	Y	Y	Y	Y	Y	Y	Y
Auto-Answer	N	Y ⁵	Y ⁵	N	N	Y ⁵	Y ⁵
Ad Hoc Conference	Y	Y	Y	Y	Y	Y	Y
Meet-Me Conference	N	Y	Y	Y	Y	Y	Y
Call Pickup	N	Y	Y	Y	Y	Y	Y
Group Pickup	N	Y	Y	Y	Y	Y	Y
Call Pickup Notification	N	Y	Y	N	N	Y	Y
Directed Call Park	N	N	N	N	N	N	N
Logging out of Hunt Groups	N	Y	Y	N	N	Y	Y
Redial	Y ⁶	Y ⁶	Y ⁶	Y ⁶	Y ⁶	Y ⁶	Y ⁶
Speed Dial	Y	Y	Y	Y	Y	Y	Y
On-hook Dialing	N	Y	Y	Y	Y	Y	Y
Voice Mail Access	Y	Y	Y	Y	Y	Y	Y
Message Waiting Indicator (MWI)	Y	Y	Y	Y	Y	Y	Y
Video call	N	N	N	N	N	N	N
Survivable Remote Site Telephony (SRST) Support	Y	Y	Y	Y	Y	Y	Y
Unicast MoH	Y	Y	Y	Y	Y	Y	Y

Table 21-6 Cisco Basic IP Phones with SCCP (continued)

Features	7902G	7905G	7906G	7910G	7910 +SW	7911G	7912G/G-A
Multicast MoH	Y	Y	Y	Y	Y	Y	Y
Tone on Hold	Y	Y	Y	Y	Y	Y	Y
Speaker	N	Y ⁵	Y ⁵	Y ⁵	Y ⁵	Y ⁵	Y ⁵
Headset Jack	N	N	N	N	N	N	N
Mute	N	N	N	Y	Y	N	N
Multilevel Precedence and Preemption (MLPP)	Y	Y	Y	Y	Y	Y	Y
Barge	N	N	N	N	N	N	Y
cBarge	N	Y	Y	N	N	Y	Y
Disable General Attribute Registration Protocol (GARP)	Y	Y	Y	Y	Y	Y	Y
Signaling and Media Encryption	N	N	N	N	N	N	N
Signaling Integrity	N	N	N	N	N	N	N
Manufacturing-Installed Certificate (X.509v3)	N	N	N	N	N	N	N
Field-Installed Certificate	N	N	N	N	N	N	N
Third-Party XML Service	N	Y	Y	N	N	Y	Y
External Microphone and Speaker	N	N	N	N	N	N	N
Dial plan	N	N	N	N	N	N	N
Signaling Packet ToS Value Marking	0x60	0x60	0x60	0x60	0x60	0x60	0x60
Media Packet ToS Value Marking	0xB8	0xB8	0xB8	0xB8	0xB8	0xB8	0xB8
G.711	Y	Y	Y	Y	Y	Y	Y
G.722	N	N	N	N	N	N	N
G.723	N	Y	N	N	N	N	N
G.726	N	Y	N	N	N	N	N
G.729	Y	Y	Y	Y	Y	Y	Y
Wideband Audio	N	N	N	N	N	N	N
Wideband Video	N	N	N	N	N	N	N
Voice Activity Detection (VAD)	Y	Y	Y	Y	Y	Y	Y
Comfort Noise Generation (CNG)	Y	Y	Y	Y	Y	Y	Y
Voice Quality Metrics	N	N	Y	N	N	Y	N
DTMF - SCCP	Y	Y	Y	Y	Y	Y	Y
DTMF - RFC2833	Y	Y	Y	Y	Y	Y	Y

1. One 10 Base-T.
2. One 10/100 Base-T.
3. Two 10/100 Base-T.
4. The Cisco Unified IP Phone 7912GA has an enhanced version of Ethernet switch.
5. One-way audio monitor mode.

6. Last Number Redial.

Table 21-7 Cisco Business, Manager, and Executive IP Phones with SCCP

Feature	7940G	7941G/G-GE	7960G	7961G/G-GE	7970G	7971G-GE
Ethernet Connection	Y ¹	Y ²	Y ¹	Y ³	Y ¹	Y ⁴
Ethernet Switch (PC port)	Y	Y	Y	Y	Y	Y
Cisco Power-Over-Ethernet (PoE)	Y	Y	Y	Y	Y ⁵	Y
IEEE 802.3af Power-Over-Ethernet (PoE)	N	Y	N	Y	Y ⁵	Y
Localization	Y	Y	Y	Y	Y	Y
Directory Number	2	2	6	6	8	8
Maximum number of calls per line	200	200	200	200	200	200
Liquid Crystal Display	Y	Y	Y	Y	Y	Y
Caller ID	Y	Y	Y	Y	Y	Y
Call Waiting	Y	Y	Y	Y	Y	Y
Caller ID on Call Waiting	Y	Y	Y	Y	Y	Y
Call Hold	Y	Y	Y	Y	Y	Y
Blind Transfer	N	N	N	N	N	N
Early-Attended Transfer	Y	Y	Y	Y	Y	Y
Consultative Transfer	Y	Y	Y	Y	Y	Y
Call Forward	Y	Y	Y	Y	Y	Y
Auto-Answer	Y	Y	Y	Y	Y	Y
Ad Hoc Conference	Y	Y	Y	Y	Y	Y
Meet-Me Conference	Y	Y	Y	Y	Y	Y
Call Pickup	Y	Y	Y	Y	Y	Y
Group Pickup	Y	Y	Y	Y	Y	Y
Call Pickup Notification	Y	Y	Y	Y	Y	Y
Directed Call Park	Y	Y	Y	Y	Y	Y
Logging out of Hunt Groups	Y	Y	Y	Y	Y	Y
Redial	Y ⁶	Y ⁶	Y ⁶	Y ⁶	Y ⁶	Y ⁶
Speed Dial	Y	Y	Y	Y	Y	Y
On-hook Dialing	Y	Y	Y	Y	Y	Y
Voice Mail Access	Y	Y	Y	Y	Y	Y
Message Waiting Indicator (MWI)	Y	Y	Y	Y	Y	Y
Video call	Y	Y	Y	Y	Y	Y
Survivable Remote Site Telephony (SRST) Support	Y	Y	Y	Y	Y	Y
Unicast MoH	Y	Y	Y	Y	Y	Y

Table 21-7 Cisco Business, Manager, and Executive IP Phones with SCCP (continued)

Feature	7940G	7941G/G-GE	7960G	7961G/G-GE	7970G	7971G-GE
Multicast MoH	Y	Y	Y	Y	Y	Y
Tone on Hold	Y	Y	Y	Y	Y	Y
Speaker	Y	Y	Y	Y	Y	Y
Headset Jack	Y	Y	Y	Y	Y	Y
Mute	Y	Y	Y	Y	Y	Y
Multilevel Precedence and Preemption (MLPP)	Y	Y	Y	Y	Y	Y
Barge	Y	Y	Y	Y	Y	Y
cBarge	Y	Y	Y	Y	Y	Y
Disable General Attribute Registration Protocol (GARP)	Y	Y	Y	Y	Y	Y
Signaling and Media Encryption	Y	Y	Y	Y	Y	Y
Signaling Integrity	Y	Y	Y	Y	Y	Y
Manufacturing-Installed Certificate (X.509v3)	N	Y	N	Y	Y	Y
Field-Installed Certificate	Y	N	Y	N	N	N
Third-Party XML Service	Y	Y	Y	Y	Y	Y
External Microphone and Speaker	Y	Y	Y	Y	Y	Y
Dial Plan	N	N	N	N	N	N
Signaling Packet ToS Value Marking	0x60	0x60	0x60	0x60	0x60	0x60
Media Packet ToS Value Marking	0xB8	0xB8	0xB8	0xB8	0xB8	0xB8
G.711	Y	Y	Y	Y	Y	Y
G.722	N	N	N	N	N	N
G.723	N	N	N	N	N	N
G.726	N	N	N	N	N	N
G.729	Y	Y	Y	Y	Y	Y
Wideband Audio	N	N	N	N	N	N
Wideband Video	N	N	N	N	N	N
Voice Activity Detection (VAD)	Y	Y	Y	Y	Y	Y
Comfort Noise Generation (CNG)	Y	Y	Y	Y	Y	Y
Voice Quality Metrics	Y	Y	Y	Y	Y	Y
DTMF - SCCP	Y	Y	Y	Y	Y	Y
DTMF - RFC2833	Y	Y	Y	Y	Y	Y

1. Two 10/100 Base-T.
2. The Cisco Unified IP Phone 7941G has two 10/100 Mbps Ethernet switches, and the Cisco Unified IP Phone 7941GE has two 10/100/1000 Mbps Ethernet switches.
3. The Cisco Unified IP Phone 7961G has two 10/100 Mbps Ethernet switches, and the Cisco Unified IP Phone 7961GE has two 10/100/1000 Mbps Ethernet switches.

4. The Cisco Unified IP Phone 7971G has two 10/100 Mbps Ethernet switches, and the Cisco Unified IP Phone 7971GE has two 10/100/1000 Mbps Ethernet switches.
5. The external power adaptor (CP-PWR-CUBE-3) must be used in order to have full display brightness.
6. Last Number Redial.

Table 21-8 Specialized Endpoints

Feature	7920	7936	7985G
Ethernet Connection	N	Y ¹	Y ²
Ethernet Switch (PC port)	N	N	Y
Cisco Power-Over-Ethernet (PoE)	N	N	N
IEEE 802.3af Power-Over-Ethernet (PoE)	N	N	Y
Localization	Y	N	Y
Directory Number	12	1	2
Max number of calls per line	2	2	100
Liquid Crystal Display	Y	Y	Y
Caller ID	Y	Y	Y
Call Waiting	Y	Y	Y
Caller ID on Call Waiting	Y	Y	Y
Call Hold	Y	Y	Y
Blind Transfer	N	N	N
Early-Attended Transfer	Y	Y	Y
Consultative Transfer	Y	Y	Y
Call Forward	Y	Y	Y
Auto-Answer	N	N	Y
Ad Hoc Conference	Y	Y	Y
Meet-Me Conference	Y	Y	Y
Call Pickup	Y	Y	Y
Group Pickup	Y	Y	Y
Call Pickup Notification	N	N	N
Directed Call Park	N	N	N
Logging out of Hunt Groups	N	N	N
Redial	Y ³	Y	Y
Speed Dial	Y	N	Y
On-hook Dialing	Y	Y	Y
Voice Mail Access	Y	N	Y
Message Waiting Indicator (MWI)	Y	N	Y
Video call	N	N	Y
Survivable Remote Site Telephony (SRST) Support	Y	Y	Y ⁴

Table 21-8 Specialized Endpoints (continued)

Feature	7920	7936	7985G
Unicast MoH	Y	Y	Y
Multicast MoH	Y	Y	N
Tone on Hold	Y	Y	Y
Speaker	N	Y	Y
Headset Jack	Y	N	Y
Mute	Y	Y	Y
Multilevel Precedence and Preemption (MLPP)	N	N	Y
Barge	N	N	Y
cBarge	N	N	Y
Disable General Attribute Registration Protocol (GARP)	N	N	N
Signaling and Media Encryption	Y	N	N
Signaling Integrity	N	N	N
Manufacturing-Installed Certificate (X.509v3)	N	N	N
Field-Installed Certificate	N	N	N
Third-Party XML Service	Y	N	N
External Microphone and Speaker	N	N	N
Signaling Packet ToS Value Marking	0x60	0x60	0x60
Media Packet ToS Value Marking	0xB8	0xB8	0x88
G.711	Y	Y	Y
G.722	N	N	Y
G.723	N	N	N
G.726	N	N	N
G.729	Y	Y	Y
Wideband Audio	N	N	N
Wideband Video	N	N	N
H.261	N	N	Y
H.263	N	N	Y
H.263+	N	N	Y
H.264	N	N	Y
Voice Activity Detection (VAD)	Y	Y	Y
Comfort Noise Generation (CNG)	Y	Y	Y
Voice Quality Metrics	N	N	N
DTMF - H.245	N	N	N
DTMF - SCCP	Y	Y	Y
DTMF - RFC2833	N	N	N

1. One 10/100 Base-T.
2. Two 10/100 Base-T.
3. Last Number Redial.
4. Only audio supported on SRST.

Table 21-9 Software Device Features

Feature	IP Communicator	IP SoftPhone
Directory Number	8	6
Caller ID	Y	Y
Call Waiting	Y	Y
Caller ID on Call Waiting	Y	Y
Call Hold	Y	Y
Call Transfer	Y	Y
Call Forward	Y	Y
Auto-Answer	Y	Y
Ad Hoc Conference	Y	Y
Meet-Me Conference	Y	N
Call Pickup	Y	N
Group Pickup	Y	N
Call Pickup Notification	Y	N
Directed Call Park	Y	N
Logging out of Hunt Groups	Y	N
Redial	Y	Y
Speed Dial	Y	N
On-hook Dialing	Y	Y
Voice Mail Access	Y	Y
Message Waiting Indicator (MWI)	Y	Y
Video call	N	N
Survivable Remote Site Telephony (SRST) Support	Y	N
Music on Hold (MoH)	Y	Y
Speaker	Y	Y
Mute	Y	Y
Multilevel Precedence and Preemption (MLPP)	Y	Y
Barge	Y	N
cBarge	N	N
Disable General Attribute Registration Protocol (GARP)	Y	N
Signaling and Media Encryption	N	N

Table 21-9 Software Device Features (continued)

Feature	IP Communicator	IP SoftPhone
Signaling Integrity	N	N
Manufacturing-Installed Certificate (X.509v3)	N	N
Field-Installed Certificate	N	N
Third-Party XML Service	Y	N
Signaling Packet ToS Value Marking	0x60	0x60
Media Packet ToS Value Marking	0xB8	0xB8
Skinny Client Control Protocol (SCCP)	Y	N
Session Initiation Protocol (SIP)	N	N
H.323	N	Y
Media Gateway Control Protocol (MGCP)	N	N
Telephony Application Programming Interface (TAPI)	N	Y
G.711	Y	Y
G.722	N	N
G.723	N	Y
G.726	N	N
G.729	Y	Y
Wideband Audio	Y	N
Wideband Video	N	N
Voice Activity Detection (VAD)	Y	N
Comfort Noise Generation (CNG)	Y	N
Voice Quality Metrics	Y	N



CHAPTER 22

Device Mobility

Last revised on: February 13, 2008

In Cisco Unified CallManager, a site or a physical location is identified using various settings such as locations, regions, calling search spaces, and media resources. Cisco Unified IP Phones residing in a particular site are statically configured with these settings. Cisco Unified CallManager uses these settings for proper call establishment, call routing, media resource selection, and so forth. However, when mobile phones such as Cisco IP Communicator or Cisco Unified Wireless IP Phones are moved from their home site to a remote site, they retain the home settings that are statically configured on the phones. Cisco Unified CallManager then uses these home settings on the phones in the remote site. This situation is undesirable because it can cause problems with call routing, codec selection, media resource selection, and other call processing functions.

Cisco Unified Call Manager Release 4.2 introduces a new feature called Device Mobility, which enables Cisco Unified CallManager to determine if the IP phone is at its home location or at a roaming location. Cisco Unified CallManager uses the device's IP subnets to determine the exact location of the IP phone. By enabling device mobility within a cluster, mobile users can roam from one site to another, thus acquiring the site-specific settings. Cisco Unified CallManager then uses these dynamically allocated settings for call routing, codec section, media resource selection, and so forth.


This chapter explains the design consideration for implementing device mobility within a Cisco Unified CallManager cluster.

Need for Device Mobility

This section explains the need for device mobility when there are many mobile users in a Cisco Unified CallManager cluster.

Figure 22-1 illustrates a hypothetical network with a Cisco Unified CallManager cluster using Release 4.1 or earlier, located at the headquarter site (HQ). The cluster has two remote sites, Branch1 and Branch2. All intra-site calls use G.711 voice codecs, while all inter-site calls (calls across the IP WAN) use G.729 voice codecs. Each site has a PSTN gateway for external calls.

Figure 22-1 Example Network with Two Remote Sites



When a user in Branch1 moves to Branch2 and calls a PSTN user in Denver, the following behavior occurs:

- Cisco Unified CallManager is not aware that the user has moved from Branch1 to Branch2. An external call to the PSTN is sent over the WAN to the Branch1 gateway and then out to the PSTN. Thus, the mobile user continues to use its home gateway for all PSTN calls.
- The mobile user and Branch1 gateway are in the same Cisco Unified CallManager region and location. Location-based call admission control is applicable only for devices in different locations, and an intra-region call uses the G.711 voice codec. Thus, the call over the IP WAN to the Branch1 gateway uses the G.711 codec and is not tracked by Cisco Unified CallManager for purposes of call admission control. This behavior can result in over-subscription of the IP WAN bandwidth if all the remote links are low-speed links.
- The mobile user creates a conference by adding multiple Branch2 users to the existing call with the PSTN user in Denver. The mobile user uses the conferencing resource that is on the Branch1 gateway, therefore all conference streams flow over the IP WAN.

Device Mobility Feature


Cisco Unified CallManager 4.2 introduces the Device Mobility feature, which helps solve the problems mentioned above. This section briefly explains how the feature works. However, for a detailed explanation of this feature, refer to the product documentation available on <http://www.cisco.com>.

Some of the new device mobility terms introduced in Cisco Unified CallManager 4.2 include:

- **Device Mobility Info** — Configures IP subnets and associates device pools to the IP subnets.
- **Device Mobility Group** — Defines a logical group of sites with similar dialing patterns (for example, US_dmg and EUR_dmg in Figure 22-2).
- **Physical Location** — Defines the physical location of a device pool. In other words, this element defines the geographic location of IP phones and other devices associated with the device pool. (For example, all San Jose IP phones in Figure 22-2 are defined by physical location SJ_phyloc.)


Figure 22-2 illustrates the relationship between all these terms.

Figure 22-2 Relationship of Device Mobility Components



Cisco Unified CallManager 4.2 assigns a device pool to an IP phone based on the device's IP subnet. The following steps, illustrated in Figure 22-3, describe the behavior:

1. The IP phone tries to register to Cisco Unified CallManager by sending its IP address in the Skinny Client Control Protocol (SCCP) registration message.
2. Cisco Unified CallManager derives the device's IP subnet and matches it with the subnet configured in the Device Mobility Info.
3. If the subnet matches, Cisco Unified CallManager provides the device with a new configuration based on the device pool configuration.

Figure 22-3 Phone Registration Process

Cisco Unified CallManager 4.2 adds new parameters to the device pool configuration to accommodate Device Mobility. These parameters are of the following two main types:

- **Roaming Sensitive Settings**

The parameters under these settings will override the device-level settings when the device is roaming within or outside a Device Mobility Group. The parameters included in these settings are:

- Date/time Group
- Region
- Media Resource Group List
- Location
- Network Locale
- SRST Reference
- Physical Location
- Device Mobility Group

The roaming sensitive settings primarily help in achieving proper call admission control and voice codec selection because the location and region configurations are used based on the device's roaming device pool.

For more details on various call admission control techniques, see the chapter on [Call Admission Control, page 9-1](#).

The roaming sensitive settings also update the media resource group list (MRGL) so that appropriate remote media resources are used for music on hold, conferencing, transcoding, and so forth, thus utilizing the network efficiently.

The roaming sensitive settings also update the Survivable Remote Site Telephony (SRST) gateway. Mobile users register to a different SRST gateway while roaming. This registration can affect the dialing behavior when the roaming phones are in SRST mode. For more information on dial plan design considerations, see [Device Mobility, page 10-26](#).

- **Device Mobility Related Settings**

The parameters under these settings will override the device-level settings only when the device is roaming within a Device Mobility Group. The parameters included in these settings are:


- Device Mobility Calling Search Space
- AAR Calling Search Space
- AAR Group

The device mobility related settings affect the dial plan because the calling search space dictates the patterns that can be dialed or the devices that can be reached.

Device Mobility Group, as explained earlier, defines a logical group of sites with similar dialing patterns (for example, sites having the same PSTN access codes and so forth). With this guideline, all sites have similar dialing patterns in the site-specific calling search spaces. Sites having different dialing behavior are in a different Device Mobility Group. A user roaming within a Device Mobility Group may preserve his dialing behavior at the remote location even after receiving a new calling search space. A user roaming outside the Device Mobility Group may still preserve his dialing behavior at the remote location because he uses his home calling search space.

However, if a Device Mobility Group is defined with sites having different dialing patterns (for example, sites having different PSTN access codes), then a user roaming within that Device Mobility Group might not preserve his same dialing behavior at all locations. A user might have to dial digits differently at different locations after receiving a new calling search space at each location. This behavior can be confusing for users.

The flowchart in [Figure 22-4](#) represents the operation of the Device Mobility feature.

Figure 22-4 Operation of the Device Mobility Feature

The Cisco Unified CallManager publisher server must be available for Device Mobility feature operation. If the publisher is not available, mobile users will register with the call-processing server when they roam to other sites and will continue to use their home device pool configurations. When the

publisher becomes available after the user has moved from its home location to a remote location, the phone will continue to use its home device pool configurations until the phone is reset. After the reset, the phone will update its configurations with roaming device pool settings.

The following guidelines apply to the Device Mobility feature:

- If the overlapping parameters listed in [Figure 22-4](#) have the same configurations on the device as well as the device pool, then these parameters may be set to NONE on the device. These parameters must then be configured on the device pool. This practice can greatly reduce the amount of configuration because the devices do not have to be configured individually with all the parameters.
- Define one physical location per site. A site may have more than one device pool.
- Define sites with similar dialing patterns for PSTN or external/off-net access with the same Device Mobility Group. Sites with different dialing patterns for PSTN or external/off-net access may be defined with the same Device Mobility Group. However, the users must be educated properly about the various dialing patterns or access codes that are used.
- A "catch-all" Device Mobility Info with IP subnet 0.0.0.0 may be defined for all non-defined subnets, depending on the company policy. This Device Mobility Info may be used to assign a device pool that can restrict access or usage of the network resources. (For example, the device pool may be configured with a calling search space NONE that will block any calls from the device associated with this device pool while roaming.) However, by doing so, administrators must be aware of the fact that this will block all calls, even 911 or other emergency calls. The calling search space may be configured with partitions that will give access only to 911 or other emergency calls.

Dial Plan Design Considerations

When using the Device Mobility feature, the dialing behavior of the phone depends on the roaming (or home) location of a phone. As discussed earlier, the device mobility related settings within the device pool affect the call flow behavior because the calling search space dictates the reachability of destination patterns within Cisco Unified CallManager. This section discusses several dial plan approaches for Device Mobility.

For detailed explanations of various dial plan approaches, see the chapter on [Dial Plan, page 10-1](#).


Device Mobility Considerations for Building Classes of Service

Typically, a mobile user should have the same calling privileges while he is roaming as he would have at his home location. The chapter on [Dial Plan, page 10-1](#), discusses two approaches for building classes of service: [Traditional Approach, page 22-8](#), and [Line/Device Approach, page 22-9](#).

Traditional Approach

Figure 22-5 shows an example of building classes of service with the traditional approach when using Device Mobility in a cluster.

Figure 22-5 Traditional Approach to Building Classes of Service



The following guidelines apply to this approach:

- Use the calling search space on the phone device for path selection to route calls to internal or external destinations. Do not set the calling search space to NONE.
- Use the calling search space on the phone device to assign calling privileges.

- Create a device pool per site and assign a calling search space to it based on required class of service. In the example above, the international calling search space (Site1_INTL_CSS) is assigned to the device pool. Because the configuration on the device takes precedence over the device pool configuration in the home location, the users continue to use the calling search space on the phone device.
- Device Mobility and Extension Mobility may be enabled for an IP phone; however, mobile users must be aware that the line calling search space of the IP phones may be used to define classes of service. Follow the extension mobility considerations with the traditional approach discussed in the chapter on [Dial Plan, page 10-1](#). When you enable both Extension Mobility and Device Mobility for a mobile user, Cisco recommends using the line/device approach.
- Assign the same class of service or calling privileges to all mobile users in the cluster. For the example in [Figure 22-5](#), a mobile user at Site 1 has Site1_National_CSS as the home device calling search space, and a mobile user at Site 2 has Site1_INTL_CSS as the home device calling search space. If both users roam to Site 2, then both users are assigned the Site2_INTL_CSS calling search space by the Site2_INTL_DP device pool, thus giving the mobile user from Site 1 a different class of service. Even if multiple device pools are created, each with different calling search spaces assigned, it is not possible to assign a correct roaming device pool with the same class of service as the home location because device pools assignments are based on IP subnet and not based on the user's capabilities.
- If not all mobile users have the same classes of service, then consider defining each site as a Device Mobility Group. This method ensures that mobile users preserve their calling privileges when they are roaming. However, this method will route all external PSTN calls using the home gateway.
- Consider using the line/device approach if customer policies require mobile users with different classes of service and the requirements cannot be satisfied by defining each site as a Device Mobility Group.


Line/Device Approach

Cisco Unified CallManager concatenates the line and device calling search spaces for a given IP phone. The following key concepts apply to the line/device approach:

- The device calling search space provides call routing information.
- The line calling search space provides class-of-service information.

With the Device Mobility feature, the device calling search space is dynamically associated to the phone based on its location. The key concept of the line/device remains the same when using Device Mobility. The line calling search space provides the class-of-service information, while the roaming or home device calling search space that is selected provides the call routing information.

[Figure 22-6](#) shows an example of building classes of service with the line/device approach when using Device Mobility in a cluster.

Figure 22-6 Line/Device Approach to Building Classes of Service

Cisco recommends using the line/device approach for building classes of service. This model has significant advantages when using Device Mobility because it greatly reduces the number of device pools needed, as indicated by the following formula:

$$\text{Total device pools} = (\text{Number of sites})$$

The following design considerations apply to this approach:

- The calling search space on the phone device can be configured to NONE. The calling search space configuration on the device pool will be assigned to the phone device. This method can greatly reduce the amount of configuration because you do not have to configure the phones individually with a device calling search space.
- There is no restriction on having the same classes of service or calling privileges for all mobile users. Because the classes of service are defined using the line calling search space, a mobile user keeps the same class of service while roaming.
- A mobile user may have both Device Mobility and Extension Mobility enabled in his profile.

Choosing a Dial Plan Model

As discussed in the chapter on [Dial Plan](#), page 10-1, there are three main approaches for the dial plan model:


- Uniform on-net dialing
- Variable length on-net dialing with partitioned addressing
- Variable length on-net dialing with flat addressing

The following sections present various dial plan models combined with an approach for building classes of service.

Uniform On-Net Dialing Using the Line/Device Approach

Figure 22-7 shows a uniform on-net dial plan for Device Mobility.

Figure 22-7 Uniform On-Net Dial Plan for Device Mobility



190169


This is the most basic dial plan model, and it has the following characteristics:

- Mobile users can use abbreviated dialing (four digits for the example in [Figure 22-7](#)) from any location.
- Abbreviated speed dialing for internal extensions continues to work on the mobile user's phone in roaming locations.
- Mobile users use a "roaming" calling search space when they are at a remote location. Cisco recommends having the same access codes for PSTN calls at all sites. If the PSTN access codes are not the same, users must learn the different access codes.

Variable Length On-Net Dialing with Partitioned Addressing Using the Line/Device Approach

Figure 22-8 shows a variable-length on-net dial plan with partitioned addressing for Device Mobility.

Figure 22-8 Variable-Length On-Net Dial Plan with Partitioned Addressing for Device Mobility



190170


The following design considerations apply to the dial plan model in [Figure 22-8](#):

- Calls might be routed to the wrong destination when mobile users use abbreviated dialing from a roaming location. For the example in [Figure 22-8](#), assume that mobile user 1 in Site 1 with extension 1000 moves to Site 2. If user 1 dials 1001 with the intention of calling a person in Site 1, the call will be routed to extension 1001 in Site 2 instead. If this behavior is not desired, consider defining each site as a Device Mobility Group. However, users must be aware that, for any external PSTN calls, the mobile phone continues to use the home gateway and therefore consumes WAN bandwidth.
- Additional device calling search spaces may be configured for roaming users with access only to the PSTN and translation partitions. This configuration will need at least one additional device pool and calling search space per site. Thus, N sites will need N device pools and N calling search spaces. However, this configuration will not require defining each site as a Device Mobility Group.
- Abbreviated speed dials should not be used. Cisco recommends configuring speed dials in a universal way that enables the users to use speed dials at any location. For example, users may configure speed dials using E.164 numbers or using site codes and access codes.
- Overlapping extensions at multiple sites might cause problems when a roaming user registers to a remote SRST gateway. For the example in [Figure 22-8](#), assume that mobile user A in Site 1 with extension 1000 moves to Site 2. Also assume that the WAN link at Site 2 goes down, causing the phones to register to an SRST gateway at Site 2. An incoming call on the SRST gateway for extension 1000 will be routed to actual extension 1000 in Site 2 as well as to the mobile user with extension 1000. Thus, calls might not be routed properly. This issue can be avoided by using unique extensions throughout the network.

Variable Length On-Net Dialing with Flat Addressing Using the Line/Device Approach

Figure 22-9 shows a variable-length on-net dial plan with flat addressing for Device Mobility.

Figure 22-9 Variable-Length On-Net Dial Plan with Flat Addressing for Device Mobility



190171

The following design considerations apply to the dial plan model in Figure 22-9:

- Mobile users cannot use abbreviated dialing after roaming to another site because calls might be routed to the wrong destination. If this behavior is not desired, consider defining each site as a Device Mobility Group. However, users must be aware that, for any external PSTN calls, the mobile phone continues to use the home gateway and therefore consumes WAN bandwidth.
- Additional device calling search spaces may be configured for roaming users with access only to the PSTN and internal phones partitions. This configuration will need at least one additional device pool and calling search space per site. Thus, N sites will need N device pools and N calling search spaces. However, this configuration will not require defining each site as a Device Mobility Group.
- Mobile users registered with a remote SRST gateway have unique extensions. However, mobile users must be aware that no PSTN user can call them when they are registered to a remote SRST gateway.

Design Guidelines for Using a VPN


This section briefly explains the design guidelines for enabling the Device Mobility feature for IP software phones or Cisco IP Communicator. Most users may have software phones that connect to the Cisco Unified CallManager cluster using a virtual private network (VPN) connection over the internet.

For information on deploying a VPN, refer to the various VPN design guides available at

<http://www.cisco.com/go/designzone>

Figure 22-10 shows an example of an IP software phone user connected over a VPN to the Cisco Unified CallManager cluster.

Figure 22-10 VPN Connection to Cisco IP Communicator



VPN users must adhere to the following guidelines:

- Configure Device Mobility Info (DMI) with the IP subnets distributed or owned by the VPN concentrators.
- Associate the DMI with the same device pool that is used for devices co-located with the VPN concentrators. However, parameters such as calling privileges, network locale, and so forth, must be taken into consideration.
- Educate the users to use the nearest VPN concentrator.

These guidelines ensure that call admission control is correctly applied on the enterprise WAN.



APPENDIX **A**

Recommended Hardware and Software Combinations

Last revised on: February 13, 2008

For the most recent information on recommended hardware platforms, software releases, and firmware versions for Cisco Unified Communications System deployments based on Cisco Unified CallManager, frequently refer to the System Test Release Matrix, latest Release Notes, and other documentation available at the following location:

http://www.cisco.com/en/US/products/sw/voicesw/ps556/tsd_products_support_integrated_systems_documentation_list.html



Note

The platforms and software versions recommended in the System Release Matrix and Release Notes are not the only supportable deployment options. They represent the combinations of hardware and software that Cisco has subjected to the most extensive system-level testing. This ongoing testing is conducted using a variety of deployment models, several end-station size categories, and realistic call flows, traffic patterns, and use cases. For information on other possible hardware and software options for Cisco Unified Communications deployments, contact your Cisco account representative.



GLOSSARY

Last revised on: February 13, 2008

A

AA	Automated attendant
AAD	Alerts and Activities Display
AAR	Automated Alternate Routing
AC	Cisco Attendant Console
ACD	Automatic call distribution
ACF	Admission Confirm
ACL	access control list
ACS	Access Control Server
AD	Microsoft Active Directory
ADUC	Active Directory Users and Computers
AFT	ALI Formatting Tool
AGM	Cisco Access Gateway Module
ALG	Application Layer Gateway
ALI	Automatic Location Identification
AMI	Alternate mark inversion
AMIS	Audio Messaging Interchange Specification
ANI	Automatic Number Identification
AP	Access point
API	Application Programming Interface
ARJ	Admission Reject
ARP	Address Resolution Protocol
ARQ	Admission Request

ASA	Adaptive Security Appliance
ASP	Active server page
ASR	Automatic speech recognition
ATA	Cisco Analog Telephone Adapter
ATM	Asynchronous Transfer Mode

B

BAT	Cisco Bulk Administration Tool
BBWC	Battery-backed write cache
BGP	Border Gateway Protocol
BHCA	Busy hour call attempts
BHCC	Busy hour call completions
BPDU	Bridge protocol data unit
bps	Bits per second
BRI	Basic Rate Interface
BTN	Bill-to number

C

CA	Certificate Authority
CAC	Call admission control
CAM	Content-addressable memory
CAMA	Centralized Automatic Message Accounting
CAPF	Certificate Authority Proxy Function
CAR	Cisco CDR Analysis and Reporting
CAS	Channel Associated Signaling
CBWFQ	Class-Based Weighted Fair Queuing
CCA	Clear channel assessment
CCE	Cisco Unified Contact Center Enterprise

CCS	Common channel signaling
CDP	Cisco Discovery Protocol
CDR	Call detail record
CFUR	Call Forward Unregistered
CGI	Common Gateway Interface
CIR	Committed information rate
CKM	Cisco Centralized Key Management
CLEC	Competitive local exchange carrier
CLID	Calling line identifier
CMC	Client Matter Code
CME	Cisco Unified CallManager Express
CMI	Cisco Messaging Interface
CMM	Cisco Communication Media Module
CNG	Comfort noise generation
CO	Central office
COM	Component Object Model
COR	Class of restriction
CoS	Class of service
CPCA	Cisco Unity Personal Assistant
CPI	Cisco Product Identification tool
CPN	Calling party number
CRS	Cisco Customer Response Solution
cRTP	Compressed Real-Time Transport Protocol
CSUF	Cross-Stack UplinkFast
CTI	Computer telephony integration
CUE	Cisco Unity Express

D

DC	Domain controller
DDNS	Dynamic Domain Name Server
DHCP	Dynamic Host Configuration Protocol
DID	Direct inward dial
DIT	Directory Information Tree
DMZ	Demilitarized zone
DN	Directory number
DNIS	Dialed number identification service
DNS	Domain Name System
DoS	Denial of service
DPA	Digital PBX Adapter
DSCP	Differentiated Services Code Point
DSE	Digital set emulation
DSP	Digital signal processor
DTMF	Dual tone multifrequency
DTPC	Dynamic Transmit Power Control
DUC	Domino Unified Communications Services

E

E&M	Receive and transmit, or ear and mouth
EAP	Extensible Authentication Protocol
EC	Echo cancellation
ECM	Error correction mode
ECS	Empty Capabilities Set
EI	Enhanced Image
EIGRP	Enhanced Interior Gateway Routing Protocol
ELIN	Emergency location identification number

EM	Extension Mobility
ER	Cisco Emergency Responder
ERL	Emergency response location
ESF	Extended Super Frame

F

FAC	Forced Account Code
FCC	Federal Communications Commission
FIFO	First-in, first-out
FR	Frame Relay
FWSM	Firewall Services Module
FXO	Foreign Exchange Office
FXS	Foreign Exchange Station

G

GARP	General Attribute Registration Protocol
GC	Global catalog
GKTMP	Gatekeeper Transaction Message Protocol
GMS	Greeting management system
GPO	Group Policy Object
GUI	Graphical user interface
GUP	Gatekeeper Update Protocol

H

H.225D	H.225 daemon
HP	Hewlett-Packard
HSRP	Hot Standby Router Protocol

HTTP Hyper-Text Transfer Protocol

Hz Hertz

I

IANA Internet Assigned Numbers Authority

IAPP Inter-Access Point Protocol

ICCS Intra-Cluster Communication Signaling

ICMP Internet Control Message Protocol

ICS IBM Cabling System

ICT Intercluster trunk

IETF Internet Engineering Task Force

IGMP Internet Group Management Protocol

IIS Microsoft Internet Information Server

IntServ Integrated Services

IntServ/DiffServ Integrated Services/Differentiated Services

IP Internet Protocol

IPCC Cisco IP Contact Center

IPIPGW IP-to-IP gateway

IPSec IP Security

ISO International Standards Organization

ITEM CiscoWorks IP Telephony Environment Monitor

ITU International Telecommunication Union

IVR Interactive voice response

J

JTAPI Java Telephony Application Programming Interface

K

kbps	Kilobits per second
KPML	Key Press Markup Language

L

LAN	Local area network
LBR	Low bit-rate
LCD	Liquid crystal display
LCF	Location Confirm
LDAP	Lightweight Directory Access Protocol
LDAPS	LDAP over SSL
LDIF	LDAP Data Interchange Format
LDN	Listed directory number
LEAP	Lightweight Extensible Authentication Protocol
LEC	Local Exchange Carrier
LFI	Link fragmentation and interleaving
LLQ	Low-latency queuing
LRJ	Location Reject
LRQ	Location Request
LSC	Label Switch Controller

M

MAC	Media Access Control
MAN	Metropolitan area network
Mbps	Megabits per second
MCM	Multimedia Conference Manager
MCS	Media Convergence Server
MCU	Multipoint Control Unit

MFT	Multiflex trunk
MGCP	Media Gateway Control Protocol
MIC	Manufacturing installed certificate
MIPS	Millions of instructions per second
MISTP	Multiple Instance Spanning Tree Protocol
MITM	Man-in-the-middle
MLA	Cisco Multi-Level Administration
MLP	Multilink Point-to-Point Protocol
MLPP	Multilevel Precedence and Preemption
MLTS	Multi-line telephone system
MoH	Music on hold
MPLS	Multiprotocol Label Switching
MRG	Media resource group
MRGL	Media resource group list
ms	Millisecond
MTP	Media termination point
mW	Milli-Watt
MWI	Message Waiting Indicator

N

NAT	Network Address Translation
NENA	National Emergency Number Association
NFAS	Non-Facility Associated Signaling
NIC	Network interface card
NPA	Numbering Plan Area
NSE	Named Service Event
NSF	Network Specific Facilities

NTE Named Telephony Event

NTP Network Time Protocol

O

OSPF Open Shortest Path First

OU Organizational unit

P

PAC Protected Access Credential

PBX Private branch exchange

PC Personal computer

PCI Peripheral Component Interconnect

PCM Pulse code modulation

PD Powered device

PHB Per-hop behavior

PIN Personal identification number

PINX Private integrated services network exchange

PIX Private Internet Exchange

PLAR Private Line Automatic Ringdown

PoE Power over Ethernet

POTS Plain old telephone service

pps Packets per second

PQ Priority Queue

PRI Primary Rate Interface

PSAP Public safety answering point

PSE Power source equipment

PSTN Public switched telephone network

PVC Permanent virtual circuit

Q

QBE	Quick Buffer Encoding
QBSS	QoS Basic Service Set
QoS	Quality of Service
QSIG	Q signaling

R

RADIUS	Remote Authentication Dial-In User Service
RAS	Registration Admission Status
RCP	Remote Copy Protocol
RDNIS	Redirected Dialed Number Information Service
RF	Radio frequency
RFC	Request for Comments
RIP	Routing Information Protocol
RSP	Route/Switch Processor
RSSI	Relative Signal Strength Indicator
RSTP	Rapid Spanning Tree Protocol
RSVP	Resource Reservation Protocol
RTMT	Cisco Real-Time Monitoring Tool
RTP	Real-Time Transport Protocol
RTT	Round-trip time

S

S1, S2, S3, and S4	Severity levels for service requests
SCCP	Skinny Client Control Protocol
SCSI	Small Computer System Interface
SDK	Software Development Kit

SDL	Signal Distribution Layer
SDP	Session Description Protocol
SE	Cisco Systems Engineer
SF	Super Frame
SI	Standard Image
SIP	Session Initiation Protocol
SIW	Service Inter-Working
SLB	Server load balancing
SLDAP	Secure LDAP
SMDI	Simplified Message Desk Interface
SNMP	Simple Network Management Protocol
SQL	Structured Query Language
SRND	Solution Reference Network Design
SRST	Survivable Remote Site Telephony
SRTP	Secure Real-Time Transport Protocol
SRV	Server
SS7	Signaling System 7
SSID	Service set identifier
SSL	Secure Socket Layer
STP	Spanning Tree Protocol
SUP1	Cisco Supervisor Engine 1
SUP2	Cisco Supervisor Engine 2
SUP2+	Cisco Supervisor Engine 2+
SUP3	Cisco Supervisor Engine 3

T

TAC	Cisco Technical Assistance Center
TAPI	Telephony Application Programming Interface
TCDD	Telephony Call Dispatcher
TCER	Total Character Error Rate
TCL	Tool Command Language
TCP	Transmission Control Protocol
TCS	Terminal Capabilities Set
TDD	Telephone Device for the Deaf
TDM	Time-division multiplexing
TEHO	Tail-end hop-off
TFTP	Trivial File Transfer Protocol
TLS	Transport layer security
ToD	Time of day
ToS	Type of service
TRaP	Telephone record and playback
TSP	Telephony Service Provider
TTL	Time to live
TTS	Text-to-speech
TTY	Terminal teletype
TUI	Telephony user interface

U

UAC	User agent client
UAS	User agent server
UDC	Universal data connector
UDLD	UniDirectional Link Detection
UDP	User Datagram Protocol

UN	Unsolicited SIP Notify
UNC	Universal Naming Convention
UPS	Uninterrupted power supply
URI	Uniform resource identifier
USB	Universal Serial Bus
UTIM	Cisco Unity Telephony Integration Manager
UTP	Unshielded twisted pair
UUIE	User-to-User Information Element

V

V3PN	Cisco Voice and Video Enabled Virtual Private Network
VAD	Voice activity detection
VAF	Voice-Adaptive Fragmentation
VATS	Voice-Adaptive Traffic Shaping
VIC	Voice interface card
VLAN	Virtual local area network
VMO	ViewMail for Outlook
VoIP	Voice over IP
VoPSTN	Voice over the PSTN
VPN	Virtual private network
VWIC	Voice/WAN interface card

W

WAN	Wide area network
WEP	Wired Equivalent Privacy
WFQ	Weighted fair queuing
WINS	Windows Internet Naming Service

WLAN	Wireless local area network
WLSM	Cisco Wireless LAN Services Module

X

XML	Extensible Markup Language



INDEX

Symbols

! in route patterns [10-12](#)
<None> location [9-12](#)
@ in route patterns [10-11](#)

Numerics

1700 Series Routers [6-9, 6-11](#)
1A and 2A cabling [3-20](#)
2800 Series Routers [6-8, 6-11, 6-17, 6-21](#)
3500 Series Video Gateways [4-31](#)
3800 Series Routers [6-8, 6-11, 6-17, 6-21](#)
4ESS [4-16, 4-17](#)
508 conformance [2-27](#)
5ESS [4-16, 4-17](#)
7914 Expansion Module [21-7](#)
7920 Wireless IP Phone [17-46, 21-12, 21-30](#)
7935 IP Conference Station [21-16](#)
7936 IP Conference Station [21-16](#)
7985G IP Video Phone [21-19, 21-20, 21-33](#)
802.1s [3-4](#)
802.1w [3-4, 3-6](#)
802.3af PoE [3-19](#)
9.@ route pattern [10-11](#)
911 calls [10-57, 11-1](#)

A

AA [14-1](#)
AAR
 dial plan considerations [10-22](#)
 for video calls [4-35, 17-7](#)

 for Voice over PSTN [2-10, 2-12](#)
 voicemail [10-24](#)
 with Cisco Unity [13-7](#)
 with hunt pilot [10-88](#)
abbreviated dialing [10-3](#)
abbreviations [GL-1](#)
access codes [10-7, 10-23](#)
access control list (ACL) [20-24, 20-26, 21-33](#)
Access Control Server (ACS) [3-63](#)
accessibility of IP Telephony features [2-27](#)
Access Layer [3-4](#)
access point (AP) [3-58, 3-61, 21-12](#)
access ports [15-25](#)
ACF [10-40](#)
ACL [20-24, 20-26, 21-33](#)
acronyms [GL-1](#)
ACS [3-63](#)
Active Directory (AD) [3-63, 18-7, 18-8, 18-11, 18-13](#)
AD [3-63, 18-7, 18-8, 18-11, 18-13](#)
Adaptive Security Appliance (ASA) [20-29, 20-31](#)
addresses
 Admission Request (ARQ) [10-40](#)
 flat [10-54, 10-63, 22-15](#)
 H.323 clients [15-37](#)
 MAC [20-13](#)
 partitioned [10-54, 10-58, 22-13](#)
 resolution [10-40, 10-41](#)
 security [20-4](#)
Address Resolution Protocol (ARP) [3-61, 20-20](#)
Admission Confirm (ACF) [10-40](#)
admission control (*see* call admission control)
Admission Reject (ARJ) [10-40](#)
Admission Request (ARQ) [10-40](#)

advanced formulas for bandwidth calculations [3-55](#)

AFT [11-19](#)

agents for call processing [1-4, 2-17](#)

Aironet [17-45](#)

algorithms for video bandwidth selection [15-13](#)

ALI [11-4, 11-19](#)

ALI Formatting Tool (AFT) [11-19](#)

all trunks busy [11-11](#)

alternate

- endpoints [5-11](#)
- gatekeeper [5-11, 8-21](#)
- TFTP file locations [3-15](#)

analog

- gateways [4-2, 4-13, 4-22, 21-2](#)
- interface modules [21-2, 21-4](#)

Analog Telephone Adaptor (ATA) [21-6, 21-22](#)

ANI [4-13, 11-4, 11-6, 11-7](#)

Annex M1 [5-11](#)

annunciator [6-17](#)

answer supervision [11-12](#)

antivirus [20-40](#)

AP [3-58, 3-61, 21-12](#)

Application ID for RSVP [3-42, 3-51](#)

applications

- for video telephony [17-42](#)
- security [20-39](#)
- third-party [1-2](#)

architecture

- for directories [18-6](#)
- for IP Telephony [1-2](#)

area code [10-23](#)

ARJ [10-40](#)

ARP [3-61, 20-20](#)

ARQ [10-40](#)

ASA [20-29, 20-31](#)

asynchronous H.323 client [17-24, 17-28](#)

Asynchronous Transfer Mode (ATM) [2-6, 2-16, 3-28](#)

ATA [21-6, 21-22](#)

ATM [2-6, 2-16, 3-28](#)

Attendant Console (AC) [17-44](#)

attending a video conference [15-32](#)

audio conferences [15-8, 15-23, 15-26](#)

audio-only calls [17-7](#)

Audio Server [15-5, 15-39](#)

audio sources [7-3, 7-9](#)

authentication

- of phones [3-62, 20-10, 21-12](#)
- open [21-13](#)
- shared key [21-13](#)

auto-detection [8-27](#)

automated alternate routing (AAR)

- dial plan considerations [10-22](#)
- for video calls [4-35, 17-7](#)
- for Voice over PSTN [2-10, 2-12](#)
- voicemail [10-24](#)
- with Cisco Unity [13-7](#)
- with hunt pilot [10-88](#)

automated attendant (AA) [14-1](#)

Automatic Location Identification (ALI) [11-4, 11-19](#)

Automatic Number Identification (ANI) [4-13, 11-4, 11-6, 11-7](#)

AUTO negotiate [3-20](#)

average conferencing minutes per month [15-7](#)

B

BackboneFast [3-6](#)

bandwidth

- advanced formulas [3-55](#)
- algorithm for selecting video bandwidth [15-13](#)
- best-effort [3-27](#)
- call control traffic [3-53, 3-54, 3-57](#)
- codec selection [15-12](#)
- consumption [3-44, 3-47](#)
- for Cisco Unified MeetingPlace Express [16-7](#)
- for Cisco Unity [13-6](#)
- for RSVP [3-49](#)
- for screen sharing [16-7](#)

- for shared line appearances [3-55](#)
- for virtual tie lines [3-57](#)
- for web applications [16-7](#)
- for wireless networks [3-65](#)
- general rule [2-19](#)
- guaranteed [3-27](#)
- management of [9-15](#)
- provisioning [3-24, 3-27, 3-44, 15-12](#)
- regions [15-12](#)
- request for [5-11](#)
- requirements for call admission control [9-13](#)
- requirements for gatekeepers [9-15](#)
- voice class requirements [3-31](#)

basic IP phones [21-6](#)

BAT [18-10](#)

B-Channel [4-38](#)

beacons [3-62](#)

Bearer Capabilities Information Element (bearer-caps) [4-41](#)

bearer-caps command [4-41](#)

bearer traffic [3-46, 3-49](#)

benefits of

- distributed call processing [2-16](#)
- single-site deployment [2-3](#)

best-effort bandwidth [3-27](#)

best practices for

- call admission control [9-2](#)
- centralized call processing [2-7](#)
- Cisco Unity Express (CUE) [14-5](#)
- directory integration [18-13](#)
- distributed call processing [2-16](#)
- fax support [4-20](#)
- IP-to-IP gateway [8-29, 9-19](#)
- line/device approach to building classes of service [10-80](#)
- modem support [4-22](#)
- music on hold [7-8](#)
- RSVP [3-44](#)
- single-site deployment [2-4](#)

- WAN design [3-25](#)

BHCA [2-24, 8-15, 8-16, 10-90](#)

BHCC [10-90](#)

bill-to number (BTN) [11-5](#)

binding of channels [4-38](#)

BPDU [3-6](#)

branch office router [7-17](#)

bridge protocol data unit (BPDU) [3-6](#)

BTN [11-5](#)

buffer size for jitter [15-14](#)

built-in conferencing [6-10](#)

Bulk Administration Tool (BAT) [18-10](#)

bump in the wire [20-34](#)

bursting [3-33](#)

business IP phones [21-6](#)

busy hour call attempts (BHCA) [2-24, 8-15, 8-16, 10-90](#)

busy hour call completions (BHCC) [10-90](#)

busy-out channels [4-38](#)

C

C5421 chipset [6-5](#)

C542 chipset [6-6](#)

C549 chipset [6-5](#)

C5510 chipset [6-4](#)

cabling

- Category 3 [3-20](#)
- IBM type 1A and 2A [3-20](#)

CAC (*see* call admission control)

calculations for server capacities [8-13](#)

call admission control

- bandwidth management [9-15](#)
- bandwidth requirements [9-13](#)
- bandwidth settings [15-12](#)
- best practices [9-2](#)
- case study [9-42](#)
- centralized call processing [9-26, 9-30, 9-36](#)
- components [9-11](#)
- described [9-1, 15-12](#)

- design considerations [9-25](#)
- distributed call processing [9-27, 9-33, 9-39](#)
- elements [9-11](#)
- for Cisco IP Communicator [21-8](#)
- for Cisco IP SoftPhone [21-11](#)
- for Cisco Unified MeetingPlace Express [16-7](#)
- for music on hold [7-16](#)
- for wireless access points [21-15](#)
- gatekeeper [8-18, 9-15, 10-38](#)
- locations [17-6](#)
- moving devices to a new location [11-13, 22-2](#)
- MPLS [9-11](#)
- regions [15-12, 17-3](#)
- RSVP [3-43](#)
- static locations [9-12](#)
- topologies [9-25](#)
- topology-aware [9-7](#)
- topology-unaware [9-3](#)
- callback
 - for emergency services [11-8, 11-14](#)
 - from the PSAP [11-8, 11-14](#)
- call control traffic [3-53, 3-57](#)
- call detail record (CDR) [2-20](#)
- call flows
 - MeetingPlace audio calls [15-23](#)
 - MeetingPlace video calls [15-34](#)
 - music on hold [7-5, 7-20](#)
- Call Forward Unregistered (CFUR) [10-19](#)
- calling line ID (CLID) [4-13, 10-13](#)
- calling party number (CPN) [11-5](#)
- calling privileges [10-15, 10-47](#)
- calling restrictions [10-15, 10-47](#)
- calling search spaces [10-15, 10-17, 10-58, 10-75, 10-79](#)
- CallManager (*see* Cisco Unified CallManager)
- CallManager Capacity Tool [8-12, 8-13, 8-14](#)
- CallManager Express (CME) [2-17, 8-27](#)
- call processing
 - agents [1-4, 2-17](#)
 - centralized [2-4, 9-26, 9-30, 9-36, 13-12, 13-13](#)
 - distributed [2-14, 9-27, 9-39](#)
 - distributed deployments [9-33](#)
 - guidelines [8-1](#)
 - hardware platforms [8-2](#)
 - redundancy [4-2, 8-6](#)
 - subscriber server [8-6](#)
 - with gatekeeper [8-18](#)
- call-related traffic [3-57](#)
- call routing for emergency calls [11-18](#)
- calls
 - 911 [11-1](#)
 - audio-only [17-7](#)
 - classification of [10-13](#)
 - coverage of [10-87](#)
 - emergency [10-57](#)
 - flow between clusters [17-8](#)
 - forwarding [10-19, 10-83](#)
 - H.323 [5-10](#)
 - hairpinning [13-13](#)
 - hold [7-6](#)
 - inbound [4-33, 4-38, 10-57, 10-63, 10-70](#)
 - music on hold [7-1](#)
 - number of calls per DSP resource [6-4, 6-5, 6-6](#)
 - outbound [4-34, 4-39, 10-57, 10-61, 10-67](#)
 - privileges [10-15](#)
 - restrictions [10-47](#)
 - routing [4-33, 4-34, 10-9, 10-35, 10-38, 11-18](#)
 - scenarios [17-8](#)
 - signaling [4-40, 4-41](#)
 - speed of [3-49](#)
 - survivability [4-12](#)
 - tromboning [13-13](#)
 - types supported [17-3](#)
 - within a cluster [10-57, 10-61, 10-66](#)
- CAM [20-13](#)
- CAMA [11-5](#)
- campus
 - access switch [3-3](#)
 - infrastructure requirements [3-1](#)

- cancellation of echo [4-21](#)
- CanMapAlias [5-11](#)
- capability exchange for T.38 fax relay [4-29](#)
- capacity planning for
 - Cisco Unified CallManager servers [8-12, 8-13, 8-14](#)
 - CTI route points and ports [8-17](#)
 - gateways [8-16](#)
 - music on hold [7-14](#)
 - phones [8-15](#)
 - third-party controlled lines [8-17](#)
 - trunks [8-17](#)
 - wireless networks [21-13](#)
- Capacity Tool [8-12, 8-13, 8-14](#)
- CAR [18-10](#)
- cascaded conferences [15-25, 15-28, 15-37](#)
- case study of call admission control [9-42](#)
- Category 3 cabling [3-20](#)
- CCA [3-62](#)
- CCE [17-44](#)
- CDP [20-11, 21-16](#)
- CDR [2-20](#)
- CDR Analysis and Reporting (CAR) [18-10](#)
- Centralized Automatic Message Accounting (CAMA) [11-5](#)
- centralized call processing
 - call admission control [9-26, 9-30, 9-36](#)
 - call coverage [10-88](#)
 - centralized messaging [13-12](#)
 - deployment model [2-4](#)
 - distributed messaging [13-13](#)
 - hunt lists [10-88](#)
 - Voice over the PSTN [2-10](#)
- centralized gatekeeper deployment [10-42](#)
- centralized messaging [13-2, 13-12, 13-16, 13-20](#)
- CFUR [10-19](#)
- channels
 - binding [4-38](#)
 - for video calls [4-38](#)
 - for wireless devices [3-59](#)
 - rollover [4-38](#)
- chipsets
 - C542 [6-6](#)
 - C5421 [6-5](#)
 - C549 [6-5](#)
 - C5510 [6-4](#)
- CIF [21-21](#)
- CIR [3-33](#)
- Cisco Aironet [17-45](#)
- ciscoatGUID [18-6](#)
- ciscoatUserProfile [18-6](#)
- ciscoatUserProfileString [18-6](#)
- Cisco CallManager
 - upgrades [18-23](#)
- Cisco Centralized Key Management (Cisco CKM) [21-13](#)
- Cisco Discovery Protocol (CDP) [20-11, 21-16](#)
- Cisco Emergency Responder (ER) [11-9, 11-13, 17-43](#)
- Cisco IOS
 - calling privileges [10-47](#)
 - call routing [10-35, 10-38](#)
 - classes of service [10-84](#)
 - digit manipulation [10-49](#)
 - DSP resources supported [6-4, 6-5, 6-6](#)
 - gatekeeper [17-19](#)
 - gateways [4-26, 4-27](#)
 - minimum release required [21-4](#)
 - software MTP [6-16](#)
- Cisco IP Communicator [21-26, 21-35](#)
- Cisco IP Conference Station [21-16, 21-22](#)
- Cisco IP SoftPhone [11-14, 17-45, 21-9, 21-26, 21-35](#)
- Cisco IP Voice Media Streaming Application [6-17](#)
- Cisco LEAP [3-63, 21-12, 21-13](#)
- Cisco Messaging Interface (CMI) [12-2](#)
- Cisco Multimedia Conference Manager (MCM) [5-11, 17-33](#)
- Cisco Security Agent [20-39](#)
- Cisco Unified CallManager
 - Capacity Tool [8-12, 8-13, 8-14](#)

- configuration of Cisco IOS gateways for fax/modem support [4-27](#)
- described [1-4](#)
- groups [2-23, 2-26](#)
- H.323 [5-9](#)
- integration with MeetingPlace [15-1](#)
- integration with MeetingPlace Express [16-1](#)
- load balancing [15-42](#)
- migration from Release 4.0 [17-42](#)
- redundancy [15-41](#)
- redundancy groups [15-38](#)
- regions [15-12](#)
- Release 3.3 [10-29](#)
- Release 4.0 [10-29](#)
- Cisco Unified CallManager Assistant (Unified CM Assistant) [17-43](#)
- Cisco Unified CallManager Express (CME) [2-17, 8-27](#)
- Cisco Unified Contact Center Enterprise (Unified CCE) [17-44](#)
- Cisco Unified IP-IVR [17-18, 17-44](#)
- Cisco Unified MeetingPlace [17-45](#)
- Cisco Unified Video Advantage
 - classification of traffic [21-33](#)
 - described [21-16](#)
- Cisco Unified Wireless IP Phone 7920 [17-46](#)
- Cisco Unity [13-1](#)
- Cisco Unity Express (CUE) [14-1](#)
- Cisco Unity Personal Assistant (CPCA) [13-2](#)
- Cisco Unity Telephony Integration Manager (UTIM) [13-9, 13-11](#)
- CKM [21-13](#)
- classes of service for users [10-72, 10-84, 22-7](#)
- classification of
 - calls [10-13](#)
 - traffic [3-4, 3-22, 3-64, 15-11, 21-21, 21-32](#)
- class of restriction (COR) [10-47, 10-84](#)
- Class of Service (CoS) [3-4, 15-11, 21-22](#)
- clear channel assessment (CCA) [3-62](#)
- CLEC [11-4](#)
- CLID [4-13, 10-13](#)
- Client Matter Code (CMC) [10-14](#)
- clients
 - H.323 [17-24](#)
 - zones [17-32](#)
- clipping [2-7](#)
- clocking source for fax/modem support [4-27](#)
- clustering over the WAN
 - Cisco Unity [13-16, 13-18](#)
 - described [2-17](#)
 - failover with Cisco Unity [13-19](#)
 - local failover [2-22](#)
 - MeetingPlace [15-4](#)
 - MeetingPlace Express [16-5](#)
 - music on hold [7-20](#)
 - remote failover [2-26](#)
 - troubleshooting [2-21](#)
 - WAN considerations [2-18](#)
- clusters
 - design guidelines [8-2](#)
 - Emergency Responder (ER) [11-18, 11-19](#)
 - multiple, for Cisco Unity [13-5](#)
 - redundancy [8-7](#)
 - services [8-3](#)
- CMC [10-14](#)
- CME [2-17, 8-27](#)
- CMI [12-2](#)
- CMM [7-3, 21-5](#)
- codecs
 - complexity modes [6-2](#)
 - flex mode [6-2](#)
 - for MeetingPlace [15-12, 15-17](#)
 - for music on hold [7-8](#)
 - for video telephony [21-20](#)
 - low bit-rate (LBR) [6-24](#)
 - selection of [21-8, 21-10](#)
 - supported by complexity mode [6-3](#)
 - supported by endpoint devices [17-4, 21-21](#)
 - types [7-3, 21-8, 21-10](#)
- collaboration

- capabilities [1-6](#)
- solutions [17-45](#)
- COM [18-3](#)
- combined deployment models for messaging [13-15](#)
- Committed Information Rate (CIR) [3-33](#)
- Common Intermediate Format (CIF) [21-21](#)
- Communication Media Module (CMM) [7-3, 21-5](#)
- Communicator [21-7, 21-26, 21-35](#)
- competitive local exchange carrier (CLEC) [11-4](#)
- complexity modes for codecs [6-2](#)
- Component Object Model (COM) [18-3](#)
- components of
 - Device Mobility [22-3](#)
 - IP Video Telephony [17-1](#)
 - MeetingPlace [15-5](#)
 - messaging system [13-5](#)
- compressed Real-Time Transport Protocol (cRTP) [3-28, 3-31](#)
- Computer Telephony Integration (CTI) [8-10, 14-2, 17-2, 17-43](#)
- Conference Station [21-16, 21-22](#)
- conferencing
 - attending a video conference [15-32](#)
 - audio [15-23](#)
 - audio-only using video endpoints [15-26](#)
 - built-in resource [6-10](#)
 - capabilities [1-6](#)
 - cascading [15-25, 15-28, 15-37](#)
 - described [6-7, 15-23](#)
 - hardware resources [6-8, 6-9](#)
 - MeetingPlace [1-6](#)
 - network utilization [15-13](#)
 - ports [15-25, 15-28, 15-31](#)
 - resources [6-7, 17-11, 17-18](#)
 - rich media [1-1](#)
 - scheduling [15-25, 15-32](#)
 - sizing the system [15-7](#)
 - software resources [6-8](#)
 - usage calculations [15-8](#)
 - video [15-29](#)
 - voice [15-23](#)
 - web [15-13, 15-26](#)
- configuration examples for
 - access control list (ACL) [20-25, 20-27](#)
 - ATA 188 and IP phones [21-22](#)
 - Cisco Unified CallManager Express [8-27](#)
 - DHCP snooping [20-18](#)
 - Dynamic ARP Inspection [20-21](#)
 - endpoint gatekeeper [17-39](#)
 - fax/modem support [4-26, 4-27](#)
 - firewalls [20-35, 20-37](#)
 - gatekeeper [8-18](#)
 - IP Source Guard [20-23](#)
 - IP-to-IP gateways [9-21](#)
 - lobby phone security [20-42](#)
 - QoS [21-21](#)
 - software-based endpoints [21-26](#)
 - switch port security [20-15](#)
 - VG224 gateways [21-21](#)
 - VG248 gateways [21-21](#)
 - via-zone gatekeepers [9-21](#)
 - Wireless IP Phone 7920 [21-30](#)
 - zones [17-31](#)
- conformance with Section 508 [2-27](#)
- connectivity
 - between MeetingPlace and IP Telephony components [15-11](#)
 - options for the WAN [2-6, 2-16](#)
- console
 - for attendants [17-44](#)
- Contact Center [1-1, 17-44](#)
- content-addressable memory (CAM) [20-13](#)
- continuous meetings [15-28, 15-33](#)
- continuous meeting server [15-39](#)
- continuous-presence conference view [17-13](#)
- control signaling [3-53, 3-57](#)
- COR [10-47, 10-84](#)
- Core Layer [3-9](#)

co-resident servers [3-14, 7-3](#)
 core switch [3-3](#)
 CoS [3-4, 15-11, 21-22](#)
 coverage of calls [10-87](#)
 CPCA [13-2](#)
 CPN [11-5](#)
 cRTP [3-28, 3-31](#)
 CTI [8-10, 14-2, 17-2, 17-43](#)
 CTI Manager [8-3, 8-10](#)
 CTI-QBE [14-2](#)
 CTI route points [6-16, 8-17](#)
 CUE [14-1](#)
 customer contact [1-1](#)
 Customer Directory Configuration Plugin [18-9](#)
 cutover [19-1, 19-2](#)

D

DAI [20-19, 20-20](#)
 database
 on Web Server [15-27](#)
 replication [8-4](#)
 SQL [18-6](#)
 data center [3-10, 20-39](#)
 data sent per web session [15-13](#)
 DC Directory [18-13](#)
 delay
 of packets [2-18, 2-20, 4-20, 4-22](#)
 variation (jitter) [4-20, 4-22](#)
 Demilitarized Zone (DMZ) [15-14, 15-28, 16-5, 20-44](#)
 deployment models
 clustering over the WAN [2-17, 7-20, 15-4, 16-5](#)
 combined for messaging [13-15](#)
 described [2-1](#)
 DHCP [3-13](#)
 for Cisco Unity [13-2](#)
 for Cisco Unity Express [14-2](#)
 for Unified MeetingPlace [15-2](#)
 for Unified MeetingPlace Express [16-1](#)

 multisite dial plan [10-52](#)
 multisite WAN with centralized call processing [2-4, 6-24, 7-15, 10-88, 15-4, 16-3](#)
 multisite WAN with distributed call processing [2-14, 6-25, 7-19, 10-53, 10-89, 15-4, 16-4](#)
 music on hold [7-15](#)
 single site [2-2, 6-23, 7-15, 15-3, 16-2](#)
 voice over the PSTN [2-10](#)
 desktop phones [21-6](#)
 destination of a call [10-22](#)
 device mobility
 described [22-1](#)
 dial plan [10-26, 22-7, 22-11](#)
 feature components and operation [22-3](#)
 Group [22-3](#)
 Info [22-3](#)
 operation flowchart [22-6](#)
 parameter settings [22-4](#)
 Physical Location [22-3](#)
 using a VPN [22-16](#)
 devices
 hunt list [10-90](#)
 limits per server [8-14](#)
 line group [10-33](#)
 mobility of [11-13, 22-2](#)
 pools [2-23, 2-26](#)
 route group [10-15](#)
 DHCP
 binding information [20-19](#)
 deployment options [3-13](#)
 described [3-11](#)
 lease times [3-12](#)
 Option 150 [3-12](#)
 servers [3-14](#)
 Snooping [20-15, 20-19](#)
 starvation attack [20-17](#)
 dialed pattern recognition [10-2](#)
 dial-in conferences [17-18](#)
 dial peers [10-35, 10-47, 10-49](#)

- dial plan
 - 911 calls [11-1](#)
 - abbreviated dialing [10-3](#)
 - access codes [10-7](#)
 - approaches to [10-54](#)
 - calling privileges [10-15, 10-47](#)
 - calling search space [10-75, 10-79](#)
 - call routing [10-9](#)
 - classes of service [10-72, 10-84, 22-7](#)
 - design considerations [10-51, 22-7](#)
 - device mobility [22-7](#)
 - dial peers [10-35, 10-47, 10-49](#)
 - distribution of digits [10-5](#)
 - elements [10-7](#)
 - emergency call string [11-10](#)
 - Extension Mobility [10-27, 10-75, 10-81](#)
 - for Cisco Unified MeetingPlace [15-37](#)
 - for Device Mobility [10-26, 22-7, 22-11](#)
 - for distributed call processing [10-53](#)
 - for multisite deployments [10-52](#)
 - for Voice over PSTN [2-13](#)
 - functions [10-1](#)
 - hunt lists [10-29, 10-32](#)
 - international calls [10-12](#)
 - line groups [10-29, 10-32](#)
 - number of digits [10-4](#)
 - on-net vs. off-net [10-3](#)
 - overlapping extensions [10-4, 10-58](#)
 - partitions [10-75, 10-79](#)
 - planning considerations [10-2, 10-7](#)
 - shared line appearance [11-14](#)
 - site codes [10-6](#)
 - string length [10-4](#)
 - uniform on-net dialing [10-4, 10-56, 22-11](#)
 - variable length on-net dialing [10-6, 10-58, 10-63, 22-13, 22-15](#)
 - voicemail [10-58, 10-63, 10-70](#)
- DID [4-13, 11-5](#)
- differential threshold [21-14](#)
- differentiated services code point (DSCP) [3-4, 3-29, 15-11](#)
- DiffServ [15-11](#)
- digital gateways [4-2, 4-15, 4-22](#)
- Digital PBX Adapter (DPA) [12-4, 12-6](#)
- digital set emulation (DSE) [12-4](#)
- digital signal processor (*see* DSP resources)
- digital trunks [15-21](#)
- digit manipulation [4-33, 10-12, 10-21, 10-49](#)
- Direct Inward Dial (DID) [4-13, 11-5](#)
- directories
 - access [18-3](#)
 - architecture [18-6](#)
 - Domain Controller (DC) [18-13](#)
 - integration with
 - Cisco Unified CallManager 4.x [18-12](#)
 - integration with Cisco Unified CallManager 5.0 [18-6](#)
 - integration with Cisco Unified MeetingPlace Express [16-9](#)
 - integration with IP telephony system [18-1, 18-2, 18-13](#)
 - LDAP [18-1](#)
 - maintaining [18-22](#)
 - permissions [18-17](#)
 - schema [18-1](#)
 - security [18-10](#)
- directory gatekeeper [8-24, 10-45](#)
- Directory Information Tree (DIT) [18-6](#)
- directory number (DN) [10-90](#)
- disaster recovery [15-39](#)
- distributed call processing [2-14, 9-27, 9-33, 9-39, 10-89](#)
- distributed gatekeeper deployment [10-44](#)
- distributed messaging [13-3, 13-13, 13-18](#)
- Distribution Layer [3-7](#)
- distribution of digits in a dial plan [10-5](#)
- DIT [18-6](#)
- DMZ [15-14, 15-28, 16-5, 20-44](#)
- DN [10-90](#)
- DNS [3-10, 15-14, 15-28, 18-13](#)
- Domain Controller (DC) [18-13](#)
- Domain Name System (DNS) [3-10, 15-14, 15-28, 18-13](#)

domains [18-11](#)

Domino Unified Communications Services (DUC) [13-2](#)

DPA [12-4, 12-6](#)

DS0 [8-16](#)

DSCP [3-4, 3-29, 15-11](#)

DSE [12-4](#)

DSP resources

- C5421 chipset [6-5](#)
- C542 chipset [6-6](#)
- C549 chipset [6-5](#)
- C5510 chipset [6-4](#)
- calculations [6-22](#)
- described [6-2](#)
- for voice termination [6-3](#)
- in multisite deployment model [2-5, 2-15](#)
- in single-site deployment model [2-2](#)
- number of calls [6-4, 6-5, 6-6](#)
- PVDM [6-20](#)

DTMF [4-2, 4-6, 5-11, 6-14, 6-15](#)

dual PBX integration [12-5, 12-6](#)

dual tone multifrequency (DTMF) [4-2, 4-6, 5-11, 6-14, 6-15](#)

DUC [13-2](#)

dynamic ANI interface [11-7](#)

Dynamic ARP Inspection (DAI) [20-19, 20-20](#)

dynamic H.323 addressing [15-37](#)

Dynamic Host Configuration Protocol (DHCP) [3-11, 20-15, 20-17, 20-19](#)

E

E.164 address [10-67, 10-68, 11-4, 11-5, 11-7, 15-37](#)

E1 trunks [15-22](#)

E911 [11-1, 11-3](#)

EAP [3-62, 3-63](#)

EAP-FAST [3-62, 21-12](#)

echo cancellation [4-21](#)

ECM [4-21](#)

ECS [17-2](#)

efficiency of links [3-31](#)

elements of a dial plan [10-7](#)

ELIN [11-6, 11-7](#)

emergency calls [10-57](#)

emergency call string [11-10](#)

emergency location identification number (ELIN) [11-6, 11-7](#)

Emergency Responder (ER) [10-57, 11-9, 11-13, 17-43](#)

emergency response location (ERL) [11-6, 11-7, 11-13](#)

emergency services [11-1](#)

EMP [15-31, 17-11](#)

Empty Capabilities Set (ECS) [17-2](#)

encryption

- for phones [20-10](#)
- for signaling [3-54, 3-55](#)

endpoints

- alternate [5-11](#)
- analog gateways [21-2](#)
- codecs supported [17-4](#)
- defined [1-5](#)
- directory access [18-3](#)
- features [21-35](#)
- gatekeeper [17-21, 17-23](#)
- gatekeeper output [8-24](#)
- gatekeeper registration [8-24](#)
- H.323 [21-34](#)
- H.323 clients [17-24](#)
- hiding IP addresses [6-26](#)
- line group devices [10-33](#)
- SCCP [21-16](#)
- software-based [1-5, 21-7, 21-26](#)
- Sony [21-20](#)
- supplementary services [6-12](#)
- Tandberg [21-20, 21-34](#)
- time to live [17-38](#)
- types of [21-1](#)
- video [1-5, 15-26, 17-1, 21-16, 21-32](#)
- wireless [1-6, 21-12](#)

Enhanced Media Processor (EMP) [15-31, 17-11](#)

Enterprise MCM [8-18](#)

- equations for calculating
 - bandwidth [3-54, 3-55](#)
 - calling search spaces [10-75, 10-79](#)
 - partitions [10-75, 10-79](#)
- ER [10-57, 11-13, 17-43](#)
- ERL [11-6, 11-7, 11-13](#)
- Error Correction Mode (ECM) [4-21](#)
- error rate [2-21](#)
- ettercap virus [20-20](#)
- example configurations [17-31, 17-39](#)
- executive IP phones [21-7](#)
- Expansion Module 7914 [21-7](#)
- Extensible Authentication Protocol (EAP) [3-62, 3-63, 21-12](#)
- Extensible Authentication Protocol-Flexible Authentication via Secure Tunneling (EAP-FAST) [3-62, 21-12](#)
- Extension Mobility (EM)
 - dial plan [10-27, 10-75, 10-81](#)
 - profiles [8-16](#)
 - server capacity and performance [8-11, 8-16](#)
- extensions, overlapping [10-4](#)

F

- FAC [10-13](#)
- failover
 - between subscriber servers [2-20](#)
 - Cisco Unity [13-4, 13-19](#)
 - clustering over the WAN [2-22, 2-26](#)
 - to PSTN [10-67, 10-68](#)
- Fast Start [6-12](#)
- fax
 - clocking source [4-27](#)
 - Error Correction Mode [4-21](#)
 - features supported [4-25](#)
 - gateway support for [4-2, 4-20](#)
 - interface modules [21-2, 21-3](#)
 - interoperability of features [4-24](#)
 - pass-through mode [4-20](#)

- protocols supported [4-23](#)
- relay mode [4-20](#)
- supported platforms and features [4-22](#)
- T.38 [4-28](#)
- features of endpoints [21-35](#)
- firewalls
 - around gateways [20-29](#)
 - bump in the road [20-34](#)
 - centralized deployment [20-44](#)
 - configuration example [20-35, 20-37](#)
 - described [20-31](#)
 - routed mode [20-33, 20-36](#)
 - stealth mode [20-34](#)
 - transparent mode [20-34, 20-37](#)
- Firewall Services Module (FWSM) [20-29, 20-31, 20-36, 20-37](#)
- flash used for music on hold [7-17](#)
- flat addressing [10-54, 10-63, 22-15](#)
- flex mode for codecs [6-2](#)
- floating ports [15-31](#)
- flows for
 - calls between clusters [17-8](#)
 - call signaling [15-23](#)
 - video conference calls [15-34](#)
- Forced Account Codes (FAC) [10-13](#)
- Foreign Exchange Office (FXO) [11-6](#)
- Foreign Exchange Station (FXS) [12-3](#)
- forwarding calls [10-19, 10-83](#)
- Frame Relay [2-6, 2-16, 3-28](#)
- French national numbering plan [10-79](#)
- full-duplex [3-20](#)
- FWSM [20-29, 20-31, 20-36, 20-37](#)
- FXO [11-6](#)
- FXS [12-3](#)

G

- GARP [20-6, 20-20](#)
- gatekeeper

- alternate [5-11, 8-21](#)
- call admission control [2-16, 9-15](#)
- call routing [10-38](#)
- centralized deployment [10-42](#)
- clustering [8-21](#)
- configuration examples [8-18](#)
- described [15-37, 17-19](#)
- design considerations [8-18](#)
- directory [8-24, 10-45](#)
- distributed deployment [10-44](#)
- for Cisco Unified MeetingPlace Express [16-11](#)
- for endpoints [8-24, 17-21, 17-23](#)
- geographical resiliency [17-21](#)
- H.225 trunks [5-2, 5-10](#)
- incompatibilities [17-21](#)
- intercluster trunks [5-2](#)
- IOS [17-19](#)
- legacy [9-22](#)
- MCU registration [15-38](#)
- MeetingPlace registration [15-38](#)
- output example [8-24](#)
- proxy [17-33, 17-35, 17-37](#)
- redundancy [8-18, 8-24](#)
- roles [17-21](#)
- scalability [17-21](#)
- summary [17-39](#)
- supported platforms [17-22](#)
- trunk redundancy [5-4](#)
- via-zone [9-18, 9-23, 10-41](#)
- zones [9-15, 17-31](#)
- gatekeeper-controlled
 - H.225 trunks [5-2, 5-10](#)
 - H.323 client [17-24, 17-28](#)
 - intercluster trunks [5-2](#)
- Gatekeeper Transaction Message Protocol (GKTMP) [5-11](#)
- Gatekeeper Update Protocol (GUP) [5-4, 8-21](#)
- gateways
 - 911 services [11-11](#)
 - all trunks busy [11-11](#)
 - analog [4-2, 4-13, 4-22, 21-2](#)
 - automated alternative routing [4-35](#)
 - blocking [11-11](#)
 - capabilities [4-41](#)
 - capacity calculations [8-16](#)
 - Cisco IOS [4-26, 4-27](#)
 - Cisco Unified Videoconferencing 3500 Series Video Gateways [4-31](#)
 - configuration examples for fax/modem support [4-26](#)
 - configuration in Cisco Unified CallManager [4-40](#)
 - controlled with Named Service Event (NSE) [4-28](#)
 - core feature requirements [4-3](#)
 - digital [4-2, 4-15, 4-22](#)
 - digit manipulation [4-33](#)
 - DS0 [8-16](#)
 - fax support [4-20](#)
 - features [21-35](#)
 - firewalls [20-29](#)
 - for local failover [2-25](#)
 - for music on hold [7-3](#)
 - for video telephony [4-31](#)
 - H.320 [17-30, 17-36](#)
 - H.323/SIP [15-5, 15-40](#)
 - IP [15-5, 15-40](#)
 - IP-to-IP [9-17, 9-23, 10-41](#)
 - modem support [4-21](#)
 - placement [11-11](#)
 - protocols [4-3](#)
 - QoS configuration examples [21-21](#)
 - QSIG support [4-19](#)
 - redundancy [4-10](#)
 - security [20-28](#)
 - selection of [4-2](#)
 - service prefixes [4-34](#)
 - SIP [4-7, 4-11](#)
 - site-specific requirements [4-12](#)
 - V.34 modem support [4-22](#)
 - V.90 modem support [4-22](#)

VG224 [4-13, 12-2, 21-5](#)
 VG248 [4-26, 12-2, 12-3, 21-5](#)
 voice applications [4-1, 21-2, 21-5](#)
 WS-X6624 [12-2, 12-3](#)
 zone prefixes [17-37](#)
 Gateway System Integrity Manager (GWSIM) [15-11, 15-17](#)
 general security [20-1](#)
 generic topologies [9-41](#)
 geographical resiliency [17-21](#)
 GKTMP [5-11](#)
 glossary [GL-1](#)
 GPO [18-19](#)
 Gratuitous Address Resolution Protocol (GARP) [20-6, 20-20](#)
 Group Policy Object (GPO) [18-19](#)
 groups for
 call routing [10-14](#)
 Cisco Unified CallManager redundancy [5-2, 8-6](#)
 Emergency Responder (ER) [11-15, 11-16](#)
 line numbers (hunting) [10-29](#)
 media resources [6-1](#)
 ports [13-5](#)
 guaranteed bandwidth [3-27](#)
 GUP [5-4, 8-21](#)
 GWSIM [15-11, 15-17](#)

H

H.225 [5-2](#)
 H.225 trunks [5-2, 5-10](#)
 H.245 [4-29](#)
 H.320 [17-30, 17-36](#)
 H.323
 analog gateways [4-13](#)
 Annex M1 [5-11](#)
 call hairpinning [8-27](#)
 calls [5-10](#)
 classes of service [10-84](#)
 clients [15-38, 17-24, 17-33](#)

 dial peers for call routing [10-35](#)
 digital gateways [4-15, 4-16, 4-17](#)
 dynamic addressing [15-37](#)
 Fast Start [6-12](#)
 fax and modem support [4-23](#)
 gateways [4-3, 16-9](#)
 in Cisco Unified CallManager [5-9](#)
 in single-site deployment model [2-4](#)
 MCU resources [15-9, 17-16](#)
 SIP IP Gateway [15-5, 15-19, 15-40](#)
 supplementary services [6-12](#)
 support on Cisco Unified MeetingPlace [15-18](#)
 T.38 fax relay [4-30](#)
 trunks [5-2, 5-8](#)
 video endpoints [17-2, 21-34](#)
 zones prefixes [17-33](#)
 hairpinning [8-27, 13-13](#)
 half-duplex [3-20](#)
 hardware
 analog interface modules [21-4](#)
 audio conferencing bridge [6-8, 6-9](#)
 DSP resources [6-4, 6-5, 6-6](#)
 gatekeepers [8-18](#)
 media resource capacities [6-20](#)
 MTP resources [6-17](#)
 music on hold [7-12](#)
 recommendations [2-2, A-1](#)
 transcoder [6-11, 6-12](#)
 types of platforms [8-2](#)
 headers for voice packets [3-46](#)
 hiding IP addresses of endpoints [6-26](#)
 high availability of
 network services [3-4](#)
 voice services [2-8](#)
 high-availability servers [8-2](#)
 high-density analog interface modules [21-3](#)
 high-performance servers [8-2](#)
 HLog softkey [10-33](#)
 hold [7-1, 7-6](#)

holdee [7-5](#)
 holder [7-5](#)
 Hot Standby Router Protocol (HSRP) [2-16, 3-7, 8-18, 8-19](#)
 HSRP [2-16, 3-7, 8-18, 8-19](#)
 hub-and-spoke topology [3-3, 3-25, 9-15, 9-25, 10-38](#)
 hunt
 groups [10-29, 10-33](#)
 lists [10-29, 10-32, 10-90](#)
 logging out of hunt group [10-33](#)
 pilot [10-29, 10-31, 10-90](#)

IBM Cabling System (ICS) [3-20](#)
 ICCS [2-19, 2-24, 8-4](#)
 ICMP [4-11](#)
 ICS [3-20](#)
 IDS [2-19, 20-29](#)
 IM gateway [15-2](#)
 immediate meetings [15-28, 15-33](#)
 impairments without QoS [3-24](#)
 inbound calls [4-33, 4-38, 10-57, 10-63, 10-70](#)
 incompatibilities [17-21](#)
 Informix Dynamic Server (IDS) [2-19](#)
 infrastructure (*see* network infrastructure)
 infrastructure gatekeeper [17-21](#)
 inline power [3-19](#)
 Instant Messaging (IM) Gateway [15-2](#)
 Integrated Services (IntServ) model [3-39, 3-44](#)
 Integrated Services/Differentiated Services (IntServ/DiffServ) model [3-41, 3-44](#)
 integrating MeetingPlace with IP Telephony [15-1, 16-1](#)
 integrations with Cisco Unity [13-5](#)
 interactive voice response (IVR) [2-4, 17-18, 17-44](#)
 intercluster trunks
 gatekeeper controlled [5-2](#)
 non-gatekeeper controlled [5-3](#)
 interface modules [21-2](#)
 interface types for 911 calls [11-4](#)

interference to wireless communications [3-60](#)
 international calls [10-12](#)
 Internet Control Message Protocol (ICMP) [4-11](#)
 Internet Protocol (IP) [15-17](#)
 interoperability
 of fax and modem features [4-24](#)
 protocols [15-17](#)
 Intra-Cluster Communication Signaling (ICCS) [2-19, 2-24, 8-4](#)
 Intrusion Detection System (IDS) [20-29](#)
 IntServ/DiffServ model [3-41, 3-44](#)
 IntServ model [3-39, 3-44](#)
 invia [9-18, 10-41, 17-31](#)
 IOS
 calling privileges [10-47](#)
 call routing [10-35, 10-38](#)
 classes of service [10-84](#)
 digit manipulation [10-49](#)
 DSP resources supported [6-4, 6-5, 6-6](#)
 Gatekeeper [17-19](#)
 minimum release required [21-4](#)
 software MTP [6-16](#)
 IP [15-17](#)
 IP/H323 feature set [8-18](#)
 IP/VC 3500 Series Video Gateways [4-31](#)
 IP/VC Enhanced Media Processor (EMP) [15-31](#)
 IP/VC MCU [15-30, 15-38, 15-42](#)
 IP addresses
 hiding [6-26](#)
 security [20-4](#)
 IP Communicator [1-5, 21-7, 21-26, 21-35](#)
 IP Conference Station [21-16, 21-22](#)
 IP Gateway [15-5, 15-40](#)
 IP-IP gateway (IPIPGW) [9-17, 9-23, 10-41](#)
 IPIPGW [9-17, 9-23, 10-41](#)
 IPMA [18-10](#)
 IP Manager Assistant (IPMA) [18-10](#)
 IP phones [21-6](#)
 IP ports [15-22](#)

IP Precedence [3-4, 3-29, 15-11](#)
 IP PSTN [6-26](#)
 IPSec [2-6, 2-16](#)
 IP Security Protocol (IPSec) [2-6, 2-16](#)
 IPSG [20-22](#)
 IP SoftPhone [17-45](#)
 IP Source Guard (IPSG) [20-22](#)
 IP Telephony [1-1, 1-2, 15-1](#)
 IP-to-IP gateway (IPIPGW) [9-17, 9-23, 10-41](#)
 IP videoconferencing (IP/VC) [15-30, 15-38, 15-42](#)
 IP Video Telephony
 components [17-1](#)
 described [1-1, 1-6, 17-1](#)
 security [20-9](#)
 IP VOICE feature set [8-27](#)
 IP Voice Media Streaming Application [6-8, 6-16, 6-17, 6-19, 8-11](#)
 ISDN [2-8, 2-10, 4-38](#)
 IVR [2-4, 17-18, 17-44](#)

J

jitter [2-18, 4-20, 4-22, 15-14](#)
 JTAPI [8-10, 17-2](#)

K

keep-alive message [15-34](#)
 Keypad Markup Language (KPML) [10-2](#)
 KPML [10-2](#)

L

LAN infrastructure [3-4](#)
 Layer 2 [2-16, 3-4](#)
 Layer 3 [3-4](#)
 layers of security [20-2](#)
 LBR [6-24](#)
 LCF [8-24, 10-41](#)

LCR [4-37](#)
 LDAP [8-4, 18-1](#)
 LDAP Data Interchange Format (LDIF) [18-9, 18-15](#)
 LDAP over Secure Socket Layer (LDAPS) [18-10](#)
 LDAPS [18-10](#)
 LDIF [18-9, 18-15](#)
 LDN [11-5](#)
 LEAP [3-62, 3-63, 21-12, 21-13](#)
 leased lines [2-6, 2-16, 3-28](#)
 lease times for DHCP [3-12](#)
 least-cost routing (LCR) [4-37](#)
 LEC [11-2, 11-11](#)
 lecture-style meetings [15-28](#)
 legacy gatekeeper [9-22](#)
 LFI [3-28, 3-31, 3-32](#)
 licenses [15-7](#)
 Lightweight Directory Access Protocol (LDAP) [8-4, 18-1](#)
 line/device approach to classes of service [10-76, 22-9](#)
 line appearances [3-55, 8-15](#)
 line group devices [10-33](#)
 line groups [10-29, 10-32, 10-90](#)
 line speed mismatch [3-33](#)
 link efficiency [3-31](#)
 link fragmentation and interleaving (LFI) [3-28, 3-31, 3-32](#)
 listed directory number (LDN) [11-5](#)
 LLQ [3-28, 3-29](#)
 LMHOSTS file [3-10](#)
 load balancing [3-17, 5-4, 5-7, 8-9, 15-39](#)
 lobby phone security [20-42](#)
 local dialing area [10-25](#)
 local exchange carrier (LEC) [11-2, 11-11](#)
 local failover deployment model [2-22](#)
 Location Confirm (LCF) [8-24, 10-41](#)
 Location Reject (LRJ) [10-41](#)
 Location Request (LRQ) [8-24, 10-41](#)
 locations
 <None> [9-12](#)
 static [9-12, 17-6](#)
 logout from hunt groups [10-33](#)

loose gateway [4-28](#)
 loss of packets [4-20, 4-22](#)
 low bit-rate (LBR) codecs [6-24](#)
 low-density analog interface modules [21-2](#)
 low-latency queuing (LLQ) [3-28, 3-29](#)
 LRJ [10-41](#)
 LRQ [8-24, 10-41](#)
 LRQ blast [8-24](#)

M

MAC address [20-13](#)
 MACs [11-9, 18-22](#)
 Manager Assistant [17-43](#)
 manager IP phones [21-6](#)
 manipulation of digits [10-21, 10-49](#)
 marking traffic [15-11](#)
 masking IP addresses of endpoints [6-26](#)
 MC [17-11](#)
 MCM [5-11, 8-18, 17-19, 17-33](#)
 MCS [15-26](#)
 MCU
 capacity and sizing [15-9, 17-18](#)
 configuration [15-30, 17-29](#)
 for video telephony [17-1, 17-11](#)
 redundancy [15-42](#)
 registration with gatekeeper [15-38](#)
 with H.323 [17-16](#)
 with Skinny Client Control Protocol (SCCP) [17-14](#)
 zone prefixes [17-35](#)
 zones [17-34](#)
 Media Convergence Server (MCS) [15-26](#)
 Media Gateway Control Protocol (MGCP) [2-4, 4-3, 4-14, 4-18, 4-23, 17-2](#)
 media resource group (MRG) [6-22, 17-15](#)
 media resource group list (MRGL) [6-22, 17-15](#)
 media resources
 described [6-1](#)
 design guidelines [6-22](#)
 for local failover [2-25](#)
 hardware and software capacities [6-20](#)
 PVDM [6-20](#)
 security [20-28](#)
 Media Streaming Application [6-8, 6-16, 6-17, 6-19, 8-11](#)
 media termination point (MTP)
 described [6-12](#)
 for PSTN calls [6-26](#)
 hiding IP addresses of endpoints [6-26](#)
 in multisite deployment model [2-5, 2-15](#)
 in single-site deployment model [2-2](#)
 Named Telephony Events [6-13](#)
 requirements for trunks [8-17](#)
 with H.323 trunk [5-8](#)
 with SIP trunk [5-11](#)
 MeetingPlace
 Audio Server [15-5, 15-39](#)
 components [15-5](#)
 connectivity with IP Telephony network [15-11](#)
 described [1-6](#)
 Express [16-1](#)
 gatekeeper registration [15-38](#)
 H.323/SIP IP Gateway [15-5, 15-19, 15-40](#)
 integration with IP Telephony [15-1](#)
 ports used by components [15-14, 15-18, 15-27, 15-28](#)
 protocols supported [15-17](#)
 server recommendations [15-2](#)
 sizing the system [15-7](#)
 Video application [15-29, 15-42](#)
 video conferences [17-45](#)
 Web Server [15-26, 15-41](#)
 meetings
 (*see also* conferencing)
 segmented [15-28](#)
 types [15-24, 15-28, 15-33](#)
 Message Waiting Indicator (MWI) [12-6, 14-2](#)
 messaging
 bandwidth management [13-6](#)
 capabilities [1-6](#)

- centralized [13-2, 13-12, 13-16, 13-20](#)
- Cisco Unity [13-1](#)
- combined deployment models [13-15](#)
- deployment models [13-2](#)
- distributed [13-3, 13-13, 13-18](#)
- failover [13-4, 13-19](#)
- system components [13-5](#)
- MGCP [2-4, 4-3, 4-14, 4-18, 4-23, 17-2](#)
- Microsoft Active Directory (AD) [18-7, 18-8, 18-11, 18-13](#)
- Microsoft ViewMail for Outlook (VMO) [13-2](#)
- migration
 - from Cisco Unified CallManger 4.0 [17-42](#)
 - parallel cutover [19-2](#)
 - phased method [19-1](#)
 - to IP Telephony [19-1](#)
- MISTP [3-4](#)
- mixed system port capacities [15-22](#)
- MLA [18-10](#)
- MLP [3-28](#)
- MLPP [6-17](#)
- MLTS [11-1](#)
- models for deployments (*see* deployment models)
- modem
 - clocking source [4-27](#)
 - features supported [4-22, 4-25](#)
 - gateway support for [4-2, 4-21](#)
 - interoperability of features [4-24](#)
 - pass-through mode [4-21](#)
 - platforms supported [4-22](#)
 - protocols supported [4-23](#)
 - relay mode [4-21](#)
 - upspeed [4-21](#)
 - V.34 [4-22](#)
 - V.90 [4-22](#)
- MoH [2-25, 7-1](#)
- moves, adds, and changes (MACs) [11-9, 18-22](#)
- MP [17-11, 17-13](#)
- MPLS [2-6, 2-16, 3-25, 3-28, 9-11, 9-34](#)
- MRG [6-22, 17-15](#)
- MRGL [6-22, 17-15](#)
- MTP
 - audio conferencing bridge [6-17](#)
 - described [6-12](#)
 - for PSTN calls [6-26](#)
 - hardware resources [6-17](#)
 - hiding IP addresses of endpoints [6-26](#)
 - in multisite deployment model [2-5, 2-15](#)
 - in single-site deployment model [2-2](#)
 - Named Telephony Events [6-13](#)
 - requirements for trunks [8-17](#)
 - software resources [6-16](#)
 - with H.323 trunk [5-8](#)
 - with SIP trunk [5-11](#)
- multicast music on hold [7-2, 7-8, 7-10, 7-17, 7-20](#)
- multicast traffic on WLAN [3-60](#)
- Multi-Level Administration (MLA) [18-10](#)
- Multilevel Precedence Preemption (MLPP) [6-17](#)
- multi-line telephone system (MLTS) [11-1](#)
- Multilink Point-to-Point Protocol (MLP) [3-28](#)
- multimedia collaboration [1-6](#)
- Multimedia Conference Manager (MCM) [5-11, 8-18, 17-19](#)
- multiple Cisco Unified CallManager servers [13-20](#)
- multiple clusters for Cisco Unity [13-5](#)
- Multiple Instance Spanning Tree Protocol (MISTP) [3-4](#)
- multipoint conferencing [17-11](#)
- Multipoint Controller (MC) [17-11](#)
- Multipoint Control Unit (MCU)
 - capacity and sizing [15-9, 17-18](#)
 - configuration [15-30, 17-29](#)
 - for video telephony [17-1, 17-11](#)
 - redundancy [15-42](#)
 - registration with gatekeeper [15-38](#)
 - with H.323 [17-16](#)
 - with Skinny Client Control Protocol (SCCP) [17-14](#)
- Multipoint Processor (MP) [17-11, 17-13](#)
- Multiprotocol Label Switching (MPLS) [2-6, 2-16, 3-25, 3-28, 9-11, 9-34](#)
- multiserver meetings [15-25, 15-28, 15-37](#)

multisite dial plan [10-52](#)
 multisite WAN deployment model
 with centralized call processing [2-4, 6-24, 7-15, 10-88, 15-4, 16-3](#)
 with distributed call processing [2-14, 6-25, 7-19, 10-89, 15-4, 16-4](#)
 music on hold (MoH) [2-25, 7-1](#)
 MWI [12-6, 14-2](#)

N

Named Service Event (NSE) [4-23, 4-28](#)
 Named Telephony Event (NTE) [4-7, 6-13](#)
 National Emergency Number Association (NENA) [11-6, 11-19](#)
 native transcoding with Cisco Unity [13-8](#)
 NENA [11-6, 11-19](#)
 Netscape Directory Server [18-11](#)
 network hold [7-6](#)
 network infrastructure
 access layer [3-4](#)
 core layer [3-9](#)
 distribution layer [3-7](#)
 for Cisco Unified MeetingPlace [15-10](#)
 high availability [3-4](#)
 LAN [3-4](#)
 overview [1-3](#)
 requirements [3-1](#)
 roles [3-3](#)
 security [20-3](#)
 WAN [3-25](#)
 WLAN [3-58](#)
 network modules [6-21](#)
 network services [3-10](#)
 Network Specific Facilities (NSF) [4-17](#)
 Network Time Protocol (NTP) [3-18, 15-14](#)
 network utilization for web conferencing [15-13](#)
 NFAS [2-4, 4-17](#)
 NM-HD-1V/2V/2VE module [6-8, 6-11, 6-17](#)

NM-HDV2 module [6-8, 6-11, 6-17](#)
 NM-HDV module [6-9, 6-11](#)
 nomadic phones [11-9](#)
 None location [9-12](#)
 Non-Facility Associated Signaling (NFAS) [2-4, 4-17](#)
 non-gatekeeper controlled H.323 client [17-24, 17-28](#)
 non-gatekeeper controlled intercluster trunks [5-3](#)
 non-IOS hardware platforms [6-6](#)
 NPA [10-23](#)
 NSE [4-23, 4-28](#)
 NSF [4-17](#)
 NTE [4-7, 6-13](#)
 NTP [3-18, 15-14](#)
 Numbering Plan Area (NPA) [10-23](#)
 number of digits dialed [10-4](#)

O

off-net dialing [10-3](#)
 on-net dialing [10-3, 10-4, 10-6, 10-56, 10-58, 10-63](#)
 open authentication [3-62, 21-12, 21-13](#)
 open forum meetings [15-28](#)
 Open Shortest Path First (OSPF) [20-33](#)
 Option 150 [3-11, 3-12](#)
 organizational unit (OU) [18-6, 18-15, 18-19](#)
 OSPF [20-33](#)
 OU [18-6, 18-15, 18-19](#)
 outbound calls [4-34, 4-39, 10-57, 10-61, 10-67](#)
 outvia [9-18, 10-41, 17-31](#)
 overbook ports [15-25, 15-31](#)
 overlap
 receiving [10-12](#)
 sending [10-12](#)
 overlapping
 channels [3-59](#)
 dial plans [10-58](#)
 extensions [10-4](#)
 oversubscription of a link [3-33](#)

P

- PA [17-44](#)
- PAC [3-62](#), [21-12](#)
- packets
 - delay [2-18](#), [2-20](#), [4-22](#), [15-14](#)
 - headers [3-46](#)
 - jitter [2-18](#)
 - loss of [2-19](#), [4-20](#)
- parallel cutover [19-2](#)
- parameters
 - for Device Mobility [22-4](#)
- partitioned addressing [10-54](#), [10-58](#), [22-13](#)
- partitions [10-15](#), [10-16](#), [10-58](#), [10-75](#), [10-79](#)
- passive-interface** command [3-9](#)
- pattern recognition in dialing [10-2](#)
- PC
 - Access to Voice VLAN [21-16](#)
 - port on IP phone [20-5](#), [21-16](#)
- PCS-1 video endpoint [21-20](#)
- PCS-TL50 video endpoint [21-20](#)
- performance
 - call rate [8-1](#)
 - of servers [8-12](#)
- permissions for directory access [18-17](#)
- Per-Port/Per-VLAN ACLs [21-34](#)
- Personal Assistant (PA) [17-44](#)
- Personal Communicator [1-5](#)
- phased migration [19-1](#)
- phones
 - 7914 Expansion Module [21-7](#)
 - 7920 Wireless IP Phone [21-12](#), [21-30](#)
 - 7985G IP Video Phone [21-19](#), [21-20](#), [21-33](#)
 - authentication and encryption [20-10](#)
 - basic models [21-6](#)
 - built-in conferencing [6-10](#)
 - business models [21-6](#)
 - capacity calculations [8-15](#)
 - configuration [21-14](#)
 - desktop IP models [21-6](#)
 - executive models [21-7](#)
 - features [21-35](#)
 - line appearances [8-15](#)
 - location for 911 purposes [11-9](#)
 - manager models [21-6](#)
 - nomadic [11-9](#)
 - PC port [20-5](#)
 - QoS [21-22](#)
 - roaming [3-59](#), [21-14](#)
 - SCCP [10-8](#)
 - security [20-5](#), [20-42](#)
 - settings [20-9](#)
 - software-based [1-5](#), [21-7](#), [21-26](#)
 - user input [10-8](#)
 - video telephony [21-32](#)
 - web access [20-8](#)
 - wireless [1-6](#), [21-12](#), [21-30](#)
 - with Cisco Unified Video Advantage [21-16](#)
 - with VT Advantage [17-1](#)
- physical security [20-4](#)
- pilot number for hunt lists [10-29](#), [10-31](#), [10-90](#)
- ping utility [2-20](#)
- PINX [12-6](#)
- PIX [20-29](#), [20-31](#)
- plain old telephone service (POTS) [11-6](#)
- platforms [2-2](#), [8-2](#), [8-18](#), [17-22](#)
- plugins [18-9](#)
- PoE [3-19](#)
- policy
 - for network security [20-1](#)
 - for RSVP [3-51](#)
- PortFast [3-6](#)
- port groups for Cisco Unity Connection [13-5](#)
- ports
 - access [20-14](#)
 - allocation on MCU [15-9](#)
 - capacities in mixed systems [15-22](#)
 - CTI [8-17](#)

- E1 [15-22](#)
 - enable/disable [21-16](#)
 - for call signaling [4-40](#)
 - for Cisco Unified Video Advantage [21-33](#)
 - for integration of Cisco Unity with Cisco Unified CallManager [13-9, 13-11](#)
 - groups [13-5](#)
 - IP [15-22](#)
 - managing [15-25, 15-31](#)
 - on the IP phone [20-5](#)
 - PC connection [21-16](#)
 - reserving all [15-28](#)
 - security [20-12](#)
 - T1 [15-22](#)
 - TCP [15-14, 15-18, 15-27, 15-28](#)
 - UDP [15-14, 15-18](#)
- positive disconnect supervision [12-11](#)
- POTS [11-6](#)
- Power over Ethernet (PoE) [3-19](#)
- precedence settings for network traffic [3-4, 3-29](#)
- prefixes
 - for access code [10-23](#)
 - for gatekeeper zones [9-23](#)
 - gateway [17-31](#)
 - MCU [17-30](#)
 - service [4-34, 17-17](#)
 - zones [17-33, 17-35, 17-37](#)
- PRI [11-5](#)
- Primary Rate Interface (PRI) [11-5](#)
- prioritization of traffic [3-29](#)
- priority, urgent [10-13](#)
- Priority Queue [3-51](#)
- Private Integrated Services Network Exchange (PINX) [12-6](#)
- Private Internet Exchange (PIX) [20-29, 20-31](#)
- privileges for making calls [10-15, 10-47](#)
- profiles for Extension Mobility [8-16](#)
- progress_ind alert enable 8** command [11-12](#)
- propagation of database [8-4](#)
- Protected Access Credential (PAC) [3-62, 21-12](#)
- Protocol Auto Detect [5-10](#)
- protocols
 - ARP [3-61, 20-20](#)
 - CDP [20-11, 21-16](#)
 - cRTP [3-28, 3-31](#)
 - DHCP [3-11, 20-15, 20-17, 20-19](#)
 - features supported [17-3](#)
 - GARP [20-6, 20-20](#)
 - GKTMP [5-11](#)
 - GUP [5-4, 8-21](#)
 - GWSIM [15-17](#)
 - H.225 [5-2, 5-10](#)
 - H.245 [4-29](#)
 - H.320 [17-30, 17-36](#)
 - H.323 [2-4, 4-3, 4-13, 4-15, 4-16, 4-17, 4-23, 4-30, 5-2, 5-8, 8-27, 10-35, 10-84, 15-9, 15-18, 15-37, 15-38, 17-2, 17-16, 17-24, 21-34](#)
 - HSRP [2-16, 3-7, 8-18, 8-19](#)
 - interoperability [15-17](#)
 - IP [15-17](#)
 - IPSec [2-6, 2-16](#)
 - JTAPI [17-2](#)
 - LDAP [8-4, 18-1](#)
 - MGCP [2-4, 4-3, 4-14, 4-18, 4-23, 17-2](#)
 - MISTP [3-4](#)
 - MLP [3-28](#)
 - MPLS [9-11](#)
 - NTP [3-18, 15-14](#)
 - RAS [10-38, 17-19](#)
 - RCP [20-20](#)
 - RIP [20-33](#)
 - routing [3-9](#)
 - RSTP [3-4, 3-6](#)
 - RSVP [3-25, 3-35, 9-7, 9-17](#)
 - RTP [2-16, 15-11, 15-17, 17-2](#)
 - SCCP [4-3, 4-23, 10-2, 10-8, 15-9, 15-12, 17-2, 17-14, 21-16, 21-20](#)
 - SDP [4-29, 6-13](#)

SIP [2-16, 4-7, 4-11, 4-13, 4-15, 4-16, 4-17, 5-11, 6-17, 10-2, 15-19, 17-2](#)
 SMDI [12-1](#)
 SNMP [11-9](#)
 SRTP [3-46](#)
 STP [3-6](#)
 TAPI [17-2](#)
 TCP [15-14, 15-18, 15-27, 15-28](#)
 TFTP [3-12, 3-14, 8-3, 8-10, 21-16](#)
 UDP [2-16, 5-4, 15-14, 15-17, 15-18](#)
 provisioning
 H.320 gateways [17-30](#)
 H.323 clients [17-24](#)
 MCUs [17-29](#)
 servers [8-12, 8-13, 8-14](#)
 proxy
 for gatekeeper [8-18, 17-33, 17-35, 17-37](#)
 server for SIP [15-19](#)
 PSAP [11-2, 11-8, 11-14](#)
 PSTN [2-2, 2-6, 2-10, 2-15, 6-26, 10-22, 11-1, 15-21](#)
 public safety answering point (PSAP) [11-2, 11-8, 11-14](#)
 Public Switched Telephone Network (PSTN) [2-2, 2-6, 2-15, 10-22, 11-1, 15-21](#)
 publisher server [2-20, 8-5, 18-6](#)
 PVDM [6-20](#)

Q

Q.SIG [5-11](#)
 QBE [14-2](#)
 QBSS [3-62, 21-14](#)
 QBSS-Differential Threshold [21-14](#)
 QCIF [21-21](#)
 QoS
 Cisco Unified MeetingPlace Express [16-8](#)
 configuration examples [21-21](#)
 for Cisco Unified MeetingPlace [15-11](#)
 for LAN [3-21](#)
 for music on hold [7-11](#)

 for security [20-24](#)
 for WAN [3-25, 3-28](#)
 for wireless LAN [3-64](#)
 general [1-4](#)
 RSVP [3-38](#)
 QoS Basic Service Set (QBSS) [3-62, 21-14](#)
 QSIG [4-15, 4-19, 12-6, 15-22, 19-3](#)
 Quality of Service (QoS)
 configuration examples [21-21](#)
 for Cisco Unified MeetingPlace [15-11](#)
 for Cisco Unified MeetingPlace Express [16-8](#)
 for LAN [3-21](#)
 for music on hold [7-11](#)
 for security [20-24](#)
 for WAN [3-25, 3-28](#)
 for wireless LAN [3-64](#)
 general [1-4](#)
 RSVP [3-38](#)
 quantity of
 calls [8-17](#)
 gateways [8-16](#)
 phones [8-15](#)
 trunks [8-17](#)
 Quarter Common Intermediate Format (QCIF) [21-21](#)
 queue depth [3-56](#)
 queuing of voice traffic [3-24, 3-65](#)
 Quick Buffer Encoding (QBE) [14-2](#)
 quiescent traffic [3-57](#)

R

radio frequency (RF) [21-12](#)
 RADIUS [3-63](#)
 Rapid Spanning Tree Protocol (RSTP) [3-4, 3-6](#)
 RAS [5-4, 9-15, 10-38, 15-18, 15-37, 17-19](#)
 RASAggregator trunk [17-23, 17-28](#)
 Rate Matching (RM) module [17-11, 17-14](#)
 rate of error [2-21](#)
 RBOC [11-2](#)

- RCF [17-38](#)
- RCP [20-20](#)
- RDNIS [13-7](#)
- Real-Time Monitoring Tool (RTMT) [18-10](#)
- Real-Time Transport Protocol (RTP) [2-16, 15-11, 15-17, 17-2](#)
- recommended hardware and software versions [2-2, A-1](#)
- Redirected Dialed Number Information Service (RDNIS) [13-7](#)
- redundancy
 - call processing [8-6](#)
 - Cisco Unified MeetingPlace Express [16-12](#)
 - cluster configurations [8-7](#)
 - for Cisco Unified MeetingPlace [15-39](#)
 - for music on hold [7-11](#)
 - for remote sites [2-8](#)
 - for SIP trunks [5-12](#)
 - for trunks [5-4](#)
 - gatekeeper [8-18](#)
 - gateway support for [4-2, 4-10](#)
 - groups [15-38](#)
 - IP-to-IP gateways [9-20](#)
 - load balancing [8-9](#)
 - TFTP services [3-16](#)
- Regional Bell Operating Company (RBOC) [11-2](#)
- regions [15-12, 17-3, 17-5](#)
- Registration Admission Status (RAS) [5-4, 9-15, 10-38, 15-18, 15-37, 17-19](#)
- Registration Confirm (RCF) [17-38](#)
- Registration Request (RRQ) [17-38](#)
- Relative Signal Strength Indicator (RSSI) [21-14](#)
- Remote Authentication Dial-In User Service (RADIUS) [3-63](#)
- Remote Copy Protocol (RCP) [20-20](#)
- remote failover deployment model [2-26](#)
- remote site survivability [2-8](#)
- re-packetization of a stream [6-12](#)
- replication of database [8-4](#)
- request for
 - bandwidth [5-11](#)
- reservationless meetings [15-24, 15-28, 15-33](#)
- Reservationless Single Number Access (RSNA) [15-38](#)
- reserve all ports [15-28](#)
- resilience [5-4, 8-1](#)
- resolution of addresses [10-40, 10-41](#)
- Resource Reservation Protocol (RSVP) [3-25, 3-35, 9-7, 9-17](#)
- Retry Video Call as Audio [17-7](#)
- RF [21-12](#)
- RFC 2833 [4-7, 6-13](#)
- rich-media conferencing [1-1](#)
- RIP [20-33](#)
- RJ-45 [3-20](#)
- RM [17-11, 17-14](#)
- roaming [3-59, 11-9, 21-14](#)
- Roaming Sensitive Settings [22-4](#)
- rogue
 - DHCP server [20-15](#)
 - network extensions [20-14](#)
- roles
 - in the network infrastructure [3-3](#)
 - of a gatekeeper [17-21](#)
- rollover of channels [4-38](#)
- root guard [3-6](#)
- round-trip time (RTT) [2-20, 2-24](#)
- Route/Switch Processor (RSP) [4-21](#)
- routed firewall
 - ASA or PIX [20-33](#)
 - FWSM [20-36](#)
- routers
 - access control list (ACL) [20-26](#)
 - branch office [7-17](#)
 - flash [7-17](#)
 - roles and features [3-3](#)
 - RSVP [3-38](#)
 - selective for E911 [11-3](#)
- routes
 - filters [10-11](#)
 - group devices [10-15](#)
 - groups [10-12, 10-14](#)
 - lists [10-14](#)

- patterns [10-9, 10-11](#)
 - selection of [10-24](#)
 - routing
 - calling line ID [10-13](#)
 - calls [10-9, 10-35, 10-38](#)
 - digit manipulation [10-12](#)
 - inbound calls [4-33](#)
 - least-cost [4-37](#)
 - outbound calls [4-34](#)
 - protocols [3-9](#)
 - time-of-day (ToD) [10-34](#)
 - Routing Information Protocol (RIP) [20-33](#)
 - RRQ [17-38](#)
 - RSNA [15-38](#)
 - RSP [4-21](#)
 - RSSI [21-14](#)
 - RSSI-Differential Threshold [21-14](#)
 - RSTP [3-4, 3-6](#)
 - RSVP
 - call admission control [9-7](#)
 - described [3-35](#)
 - IP-to-IP Gateway [9-17](#)
 - WAN infrastructure [3-25](#)
 - RTMT [18-10](#)
 - RTP [2-16, 15-11, 15-17, 17-2](#)
 - RTT [2-20, 2-24](#)
-
- S**
- scalability of
 - Cisco Unified CallManager [8-1](#)
 - gatekeepers [17-21](#)
 - SCCP
 - dialed pattern recognition [10-2](#)
 - fax and modem support [4-23](#)
 - gateway support for [4-3](#)
 - MCU resources [17-14](#)
 - phones [10-8](#)
 - ports on MCU [15-9](#)
 - user input on phones [10-8](#)
 - video endpoints [15-12, 17-2, 21-16, 21-20](#)
 - scheduling conferences [15-25, 15-32](#)
 - schema [18-1](#)
 - Schema Master [18-9, 18-15](#)
 - screen sharing [16-7](#)
 - SDK [18-3](#)
 - SDP [4-29, 6-13](#)
 - Section 255 [2-27](#)
 - Section 508 [2-27](#)
 - Secure Socket Layer (SSL) [18-10](#)
 - security
 - access control list (ACL) [20-24, 20-26](#)
 - antivirus [20-40](#)
 - Cisco Security Agent [20-39](#)
 - configuration example [20-15, 20-18, 20-21, 20-23, 20-25, 20-27, 20-35, 20-37, 20-42](#)
 - data center [20-39](#)
 - DHCP Snooping [20-15](#)
 - DHCP starvation attack [20-17](#)
 - directories [18-10](#)
 - firewalls [20-31, 20-44](#)
 - gateways [20-28](#)
 - infrastructure [20-3](#)
 - in general [1-6, 20-1](#)
 - layers [20-2](#)
 - lobby phone example [20-42](#)
 - MAC CAM flooding [20-13](#)
 - media resources [20-28](#)
 - MeetingPlace Express [16-5](#)
 - PC port on the phone [20-5](#)
 - phones [20-5](#)
 - phone settings [20-9](#)
 - physical access [20-4](#)
 - policy [20-1](#)
 - QoS [20-24](#)
 - rogue network extensions [20-14](#)
 - servers [20-39, 20-41](#)
 - switch port [20-12](#)

- Video Capabilities [20-9](#)
- voice VLAN [20-7](#)
- web access [20-8](#)
- wireless network [3-62](#)
- segmented meetings [15-28](#)
- selecting the proper route [10-24](#)
- selective router [11-3](#)
- separate integrations for Cisco Unity [13-5](#)
- Sequenced Routing Update Protocol (SRTP) [3-46](#)
- sequential LRQs [8-24](#)
- Server Load Balancing (SLB) [18-13](#)
- servers
 - capacity planning [8-12, 8-13, 8-14](#)
 - co-resident [3-14, 7-3](#)
 - CTI Manager [8-10](#)
 - data center [3-10](#)
 - domain [18-11](#)
 - farm [3-10](#)
 - for Cisco Unified CallManager [8-2](#)
 - for DHCP [3-14](#)
 - for media resources [6-1](#)
 - for music on hold [7-3, 7-4, 7-12](#)
 - high-availability [8-2](#)
 - high-performance [8-2](#)
 - maximum number of devices [8-14](#)
 - multiple Cisco Unified CallManager servers [13-20](#)
 - performance [8-12](#)
 - publisher [2-20, 8-5](#)
 - recommended deployments [15-2](#)
 - redundancy [15-39](#)
 - security [20-39, 20-41](#)
 - shadow server [15-39](#)
 - standalone [3-14, 7-3](#)
 - subscriber [2-20, 8-6](#)
 - TFTP [8-10](#)
 - types [8-2](#)
- Service Inter-Working (SIW) [2-6, 2-16, 3-28](#)
- services
 - prefix [4-34, 17-17, 17-30, 17-31](#)
 - supplementary [4-3](#)
 - template [17-17](#)
 - within a cluster [8-3](#)
- service set identifier (SSID) [3-58, 3-61](#)
- Session Description Protocol (SDP) [4-29, 6-13](#)
- Session Initiation Protocol (SIP)
 - analog gateways [4-13](#)
 - annunciator [6-17](#)
 - dialed pattern recognition [10-2](#)
 - digital gateways [4-15, 4-16, 4-17](#)
 - for distributed call processing [2-16](#)
 - gateways [4-11, 15-19](#)
 - gateway support for [4-7](#)
 - redundancy [5-12](#)
 - trunks [5-11, 16-10](#)
 - video endpoints [17-2](#)
- settings for IP phones [20-9](#)
- shadow server [15-39](#)
- shaping traffic [3-32](#)
- shared
 - key authentication [21-13](#)
 - line appearances [3-55, 11-14](#)
 - T.120 applications [17-45](#)
- shielded twisted-pair (STP) [3-20](#)
- signaling encryption [3-54, 3-55](#)
- Signaling System 7 [2-4](#)
- Simple Network Management Protocol (SNMP) [11-9](#)
- Simplified Message Desk Interface (SMDI) [12-1](#)
- single site
 - deployment model [2-2, 6-23, 7-15, 15-3, 16-2](#)
 - messaging model [13-2](#)
- SIP
 - analog gateways [4-13](#)
 - annunciator [6-17](#)
 - dialed pattern recognition [10-2](#)
 - digital gateways [4-15, 4-16, 4-17](#)
 - for distributed call processing [2-16](#)
 - gateways [4-11, 15-19](#)
 - gateway support for [4-7](#)

- redundancy [5-12](#)
- trunks [5-11, 16-10](#)
- video endpoints [17-2](#)
- site
 - dialing codes [10-6, 10-70](#)
 - survey for wireless network [21-12](#)
- SIW [2-6, 2-16, 3-28](#)
- sizing
 - CallManager servers [8-12, 8-13, 8-14](#)
 - Cisco Unified MeetingPlace [15-7](#)
 - MCUs [17-18](#)
- Skinnny Client Control Protocol (SCCP)
 - diald pattern recognition [10-2](#)
 - fax and modem support [4-23](#)
 - gateway support for [4-3](#)
 - MCU resources [17-14](#)
 - phones [10-8](#)
 - ports on MCU [15-9](#)
 - user input on phones [10-8](#)
 - video endpoints [15-12, 17-2, 21-16, 21-20](#)
- SLB [18-13](#)
- SMDI [12-1](#)
- SNMP [11-9](#)
- snooping [20-15](#)
- soft clients [11-14](#)
- SoftPhone [11-14, 17-45, 21-9, 21-26, 21-35](#)
- software
 - audio conferencing bridge [6-8](#)
 - endpoints [21-7](#)
 - media resource capacities [6-20](#)
 - MTP resources [6-16](#)
 - phones [21-35](#)
 - recommendations [2-2](#)
 - versions [2-2, 21-4, 21-5, A-1](#)
- Software Development Kit (SDK) [18-3](#)
- Sony endpoints [21-20](#)
- Source Guard [20-22](#)
- Spanning Tree Protocol (STP) [3-6](#)
- speed of calls [3-49](#)
- SQL [18-6](#)
- SQL database [15-27](#)
- SRST [2-6, 2-8, 7-17, 8-3, 10-89, 11-3](#)
- SRTP [3-46](#)
- SS7 [2-4](#)
- SSID [3-58, 3-61](#)
- SSL [18-10](#)
- standalone server [3-14, 7-3](#)
- standard meetings [15-24, 15-33](#)
- standard server [8-2](#)
- standby preempt** command [3-7](#)
- standby track** command [3-7](#)
- star topology [9-25](#)
- static ANI interface [11-7](#)
- static locations [9-12](#)
- Static Wired Equivalent Privacy (WEP) [3-62, 3-63](#)
- stealth firewall [20-34](#)
- STP [3-6, 3-20](#)
- string length [10-4](#)
- subnets [17-37](#)
- subscriber server [2-20, 8-6, 18-6](#)
- summary of endpoint gatekeepers [17-39](#)
- supplementary services
 - for H.323 endpoints [6-12](#)
 - on gateways [4-3, 4-7](#)
- supported
 - call types [17-3](#)
 - codecs [17-4, 21-21](#)
 - platforms for gatekeepers [17-22](#)
 - protocol features [17-3](#)
 - protocols [17-2](#)
- survey of wireless network [21-12](#)
- survivability of calls [4-12](#)
- Survivable Remote Site Telephony (SRST) [2-6, 2-8, 7-17, 8-3, 10-89, 11-3](#)
- switches
 - port security [20-12](#)
 - roles and features [3-3](#)
- synchronous H.323 client [17-24](#)

T

- T.120 application sharing [17-45](#)
- T.38 fax relay [4-28](#)
- T-1000 video endpoint [21-20](#)
- T1 trunks [15-21](#)
- T-550 video endpoint [21-20](#)
- Tail-End Hop-Off (TEHO) [10-52](#)
- Tandberg endpoints
 - classification of traffic [21-34](#)
 - described [17-1, 21-20](#)
- TAPI [8-10, 17-2](#)
- TCP [15-14, 15-18, 15-27, 15-28](#)
- TCP/UDP ports [21-33](#)
- TCS [17-9](#)
- TEHO [10-52](#)
- Telecommunications Act [2-27](#)
- telephone record and playback (TRaP) [13-2](#)
- telephone user interface (TUI) [13-2](#)
- templates to define service settings [17-17](#)
- Terminal Capabilities Set (TCS) [17-9](#)
- termination of calls [6-2](#)
- test calls for 911 [11-14](#)
- TFTP [3-12, 3-14, 8-3, 8-10, 21-16](#)
- third-party
 - controlled lines [8-17](#)
 - software applications [1-2](#)
 - video endpoints [21-20](#)
 - voicemail systems [12-1, 12-12](#)
- threshold, differential [21-14](#)
- time-of-day (ToD) routing [10-34](#)
- timers for call signaling [4-41](#)
- time synchronization [3-18, 3-19](#)
- Time to Live (TTL) [17-38](#)
- ToD [10-34](#)
- Token Ring [3-20](#)
- topology
 - case study [9-42](#)
 - for call admission control [9-25](#)
 - generic [9-41](#)
 - hub-and-spoke [9-15, 9-25, 10-38](#)
 - MPLS-based [9-34](#)
 - star [9-25](#)
 - two-tier hub-and-spoke [9-29](#)
- topology-aware call admission control [9-7](#)
- topology-unaware call admission control [9-3](#)
- ToS [15-11](#)
- tracking domain [11-18](#)
- traditional approach to classes of service [10-72, 22-8](#)
- traffic
 - bearer traffic [3-46, 3-49](#)
 - call control [3-53, 3-57](#)
 - call-related [3-57](#)
 - classification [3-4, 3-22, 3-64, 15-11, 21-21, 21-32](#)
 - prioritization [3-29](#)
 - provisioning for [3-46](#)
 - queuing [3-24, 3-65](#)
 - quiescent [3-57](#)
 - shaping [3-32](#)
 - video bearer traffic [3-48, 3-50](#)
 - voice bearer traffic [3-46, 3-50](#)
- transcoding
 - Cisco Unity [13-8](#)
 - described [6-10](#)
 - hardware resources [6-11, 6-12](#)
 - IP PSTN [6-26](#)
 - resources [6-11](#)
- translation of digits
 - patterns [10-21](#)
 - voice translation profiles [10-49](#)
- Transmission Control Protocol (TCP) [15-14, 15-18, 15-27, 15-28](#)
- transparent firewall
 - ASA or PIX [20-34](#)
 - FWSM [20-37](#)
- TRaP [13-2](#)
- Trivial File Transfer Protocol (TFTP) [3-12, 3-14, 8-3, 8-10, 21-16](#)

tromboning [13-13](#)

troubleshooting for clustering over the WAN [2-21](#)

trunks

- capacity calculations [8-17](#)
- described [5-1](#)
- digital [15-21](#)
- E1 [15-22](#)
- gatekeeper controlled [5-2](#)
- H.225 [5-2, 5-10](#)
- H.323 [5-2, 5-8](#)
- intercluster, gatekeeper controlled [5-2](#)
- intercluster, non-gatekeeper controlled [5-3](#)
- load balancing [5-4, 5-7](#)
- MTP requirements [8-17](#)
- RASAggregator [17-23, 17-28](#)
- redundancy [5-4](#)
- SIP [5-11, 6-17](#)
- T1 [15-21](#)

trust [21-21](#)

TTL [17-38](#)

TUI [13-2](#)

Tunneled Q.SIG [5-11](#)

two-tier hub-and-spoke topology [9-29](#)

Type of Service (ToS) [15-11](#)

U

UAC [21-6](#)

UAS [21-6](#)

UDC [3-20](#)

UDLD [3-6](#)

UDP [2-16, 3-31, 5-4, 15-14, 15-17, 15-18](#)

UMDirectoryConfiguration.ini file [18-22](#)

UN [4-7](#)

unicast music on hold [7-2, 7-8, 7-10, 7-20](#)

UniDirectional Link Detection (UDLD) [3-6](#)

Unified CCE [17-44](#)

Unified CM Assistant [17-43](#)

Unified Communications [1-1](#)

unified messaging (*see also* messaging) [13-1](#)

Unified Personal Communicator [1-5](#)

Unified Video Advantage

- classification of traffic [21-33](#)
- described [21-16](#)

uniform on-net dial plan [10-4, 10-56, 22-11](#)

uninterrupted power supplies (UPS) [3-19](#)

Unity [13-1](#)

Unity Express [14-1](#)

Unity Telephony Integration Manager (UTIM) [13-5, 13-9, 13-11](#)

universal data connector (UDC) [3-20](#)

Unsolicited SIP Notify (UN) [4-7](#)

upgrades to Cisco CallManager [18-23](#)

UplinkFast [3-6](#)

UPS [3-19](#)

upspeed [4-21](#)

Urgent Priority [10-13](#)

user agent client (UAC) [21-6](#)

user agent server (UAS) [21-6](#)

User Creation Base [18-19](#)

User Datagram Protocol (UDP) [2-16, 3-31, 5-4, 15-14, 15-17, 15-18](#)

user hold [7-6](#)

users

- classes of service [10-72, 10-84](#)
- input on phones [10-8](#)
- licenses [15-7](#)

User Search Base [18-15, 18-19](#)

User-to-User Information Element (UUIE) [5-10](#)

utilization of

- DSOs [8-16](#)
- phones [8-15](#)
- trunks [8-17](#)

UTIM [13-5, 13-9, 13-11](#)

UUIE [5-10](#)

V

- V.34 modems [4-22](#)
- V.90 modems [4-22](#)
- V3PN [2-6, 2-16](#)
- VAD [4-21, 8-11, 17-12](#)
- VAF [3-32](#)
- variable length on-net dial plan [10-6, 10-58, 10-63, 22-13, 22-15](#)
- variation in packet delay [15-14](#)
- VATS [3-34](#)
- VG224 Voice Gateway [4-13, 12-2, 21-5, 21-21](#)
- VG248 Analog Phone Gateway [4-26, 12-2, 12-3, 21-5, 21-21](#)
- via-zone gatekeeper [9-18, 9-23, 10-41](#)
- VIC [21-2, 21-3](#)
- video
 - bearer traffic [3-48, 3-50](#)
 - capabilities [20-9](#)
 - conferences [15-9, 15-29](#)
 - conferencing ports [15-31](#)
 - described [17-1](#)
 - enable/disable [21-16](#)
 - endpoints [1-5, 15-26, 17-1, 21-16, 21-32](#)
 - features [1-1, 1-6](#)
 - gateways [4-31](#)
 - MeetingPlace Video application [15-29, 15-42](#)
 - traffic classification [3-23, 21-32](#)
 - VLAN [20-11](#)
- Video Capabilities [20-9](#)
- video telephony (*see* IP Video Telephony)
- ViewMail for Outlook (VMO) [13-2](#)
- virtual LAN (VLAN) [3-4, 3-58, 21-21](#)
- Virtual Private Network (VPN) [2-6, 2-16, 22-16](#)
- virtual tie lines [3-57](#)
- VLAN
 - access control list (ACL) [20-24](#)
 - number of devices per VLAN [3-4](#)
 - separate VLANs for voice and data [3-58](#)
 - video [20-11](#)
 - VLAN ID [21-21](#)
 - voice [20-7, 20-11](#)
- VMO [13-2](#)
- voice
 - bandwidth requirements [3-31](#)
 - bearer traffic [3-46, 3-50](#)
 - conferences [15-8, 15-23, 15-26](#)
 - gateways [4-1, 21-2, 21-5](#)
 - link [15-29](#)
 - port integration [13-9, 13-11](#)
 - termination [6-2](#)
 - translation profiles [10-49](#)
 - VLAN [20-7, 20-11](#)
- voice/WAN interface card (VWIC) [21-2](#)
- voice-activated conference view [17-12](#)
- voice activity detection (VAD) [4-21, 8-11, 17-12](#)
- Voice-Adaptive Fragmentation (VAF) [3-32](#)
- Voice-Adaptive Traffic Shaping (VATS) [3-34](#)
- Voice and Video Enabled IPsec VPN (V3PN) [2-6, 2-16](#)
- voice interface card (VIC) [21-2, 21-3](#)
- voicemail
 - automated alternate routing (AAR) [10-24](#)
 - centralized [12-6](#)
 - Cisco Unity [13-1](#)
 - Cisco Unity Express [14-1](#)
 - dial plan [10-58, 10-63, 10-70](#)
 - dual PBX integration [12-5, 12-6](#)
 - for local failover [2-25](#)
 - integration with IP telephony system [12-1](#)
 - positive disconnect supervision [12-11](#)
 - third-party systems [12-1, 12-10, 12-12](#)
 - unified messaging [13-1](#)
- voice over IP (VoIP) [3-46](#)
- voice over the PSTN (VoPSTN) [2-10](#)
- voice rtp send-recv** command [11-12](#)
- VoIP [3-46](#)
- VoPSTN [2-10](#)
- VPN [2-6, 2-16, 22-16](#)
- VWIC [21-2](#)

W

Wait for Far-End to Send TCS [17-9](#)

WAN

aggregation router [3-3](#)

infrastructure [3-25](#)

web

access from IP phone [20-8](#)

applications [16-7](#)

conferences [15-8, 15-13, 15-26](#)

server [15-26, 15-41](#)

weighted fair queuing [3-29](#)

WEP [3-62, 3-63, 21-12](#)

Wi-Fi Protected Access (WPA) [3-63, 21-12](#)

Wi-Fi Protected Access Pre-Shared Key
(WPA-PSK) [21-13](#)

wildcard route pattern [10-11, 10-12](#)

Windows Internet Naming Service (WINS) [3-14](#)

WINS [3-14](#)

Wired Equivalent Privacy (WEP) [3-63, 21-12](#)

wireless

endpoints [21-12](#)

IP Phone 7920 [17-46, 21-12, 21-30](#)

IP phones [1-6, 21-12, 21-30](#)

LAN [3-58](#)

networking solutions [17-45](#)

wireless LAN (WLAN) [3-58](#)

WLAN infrastructure [3-58](#)

WPA [3-63, 21-12](#)

WPA-PSK [21-13](#)

WS-SVC-CMM-ACT module [6-9, 6-11, 6-17](#)

WS-X6608-E1 module [6-9, 6-12, 6-17](#)

WS-X6608-T1 module [6-9, 6-12, 6-17](#)

WS-X6624-FXS analog interface module [21-5](#)

WS-X6624 module [12-2, 12-3](#)

XML services [17-46](#)

Z

zones

clients [17-32](#)

configuration on gatekeeper [17-31](#)

for gatekeepers [9-15](#)

H.320 gateways [17-36](#)

MCU [17-34](#)

prefixes [9-23, 17-33, 17-35, 17-37](#)

subnets [17-37](#)

X

XML [15-20](#)

