Cisco Hosted / Managed Unified Communications Services

Solution Reference Network Design (SRND)

Version 1.6(0)

EDCS 580462

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

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 following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with Cisco's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specifications in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

You can determine whether your equipment is causing interference by turning it off. If the interference stops, it was probably caused by the Cisco equipment or one of its peripheral devices. If the equipment causes interference to radio or television reception, try to correct the interference by using one or more of the following measures:

Turn the television or radio antenna until the interference stops.

Move the equipment to one side or the other of the television or radio.

Move the equipment farther away from the television or radio.

Plug the equipment into an outlet that is on a different circuit from the television or radio. (That is, make certain the equipment and the television or radio are on circuits controlled by different circuit breakers or fuses.)

Modifications to this product not authorized by Cisco Systems, Inc. could void the FCC approval and negate your authority to operate the product.

The following third-party software may be included with your product and will be subject to the software license agreement:

CiscoWorks software and documentation are based in part on HP OpenView under license from the Hewlett-Packard Company. HP OpenView is a trademark of the Hewlett-Packard Company. Copyright © 1992, 1993 Hewlett-Packard Company.

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.

Network Time Protocol (NTP). Copyright © 1992, David L. Mills. The University of Delaware makes no representations about the suitability of this software for any purpose.

Point-to-Point Protocol. Copyright © 1989, Carnegie-Mellon University. All rights reserved. The name of the University may not be used to endorse or promote products derived from this software without specific prior written permission.

The Cisco implementation of TN3270 is an adaptation of the TN3270, curses, and termcap programs developed by the University of California, Berkeley (UCB) as part of the UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981-1988, Regents of the University of California.

Cisco incorporates Fastmac and TrueView software and the RingRunner chip in some Token Ring products. Fastmac software is licensed to Cisco by Madge Networks Limited, and the RingRunner chip is licensed to Cisco by Madge NV. Fastmac, RingRunner, and TrueView are trademarks and in some jurisdictions registered trademarks of Madge Networks Limited. Copyright © 1995, Madge Networks Limited. All rights reserved.

Xremote is a trademark of Network Computing Devices, Inc. Copyright © 1989, Network Computing Devices, Inc., Mountain View, California. NCD makes no representations about the suitability of this software for any purpose.

The X Window System is a trademark of the X Consortium, Cambridge, Massachusetts. All rights reserved.

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 PRACTICAL 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.

AccessPath, AtmDirector, Browse with Me, CCIP, CCSI, CD-PAC, *CiscoLink*, the Cisco *Powered* Network logo, Cisco Systems Networking Academy, the Cisco Systems Networking Academy logo, Cisco Unity, Fast Step, Follow Me Browsing, FormShare, FrameShare, IGX, Internet Quotient, IP/VC, iQ Breakthrough, iQ Expertise, iQ FastTrack, the iQ logo, iQ Net Readiness Scorecard, MGX, the Networkers logo, ScriptBuilder, ScriptShare, SMARTnet, TransPath, Voice LAN, Wavelength Router, and WebViewer, Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, FastHub, FastSwitch, GigaStack, IOS, IP/TV, LightStream, MICA, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, SlideCast, StrataView Plus, Stratm, SwitchProbe, TeleRouter, and VCO are trademarks or registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site 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. (0110R).

Please refer to http://www.cisco.com/logo/ for the latest information on Cisco logos, branding and trademarks.

INTELLECTUAL PROPERTY RIGHTS:

THIS DOCUMENT CONTAINS VALUABLE TRADE SECRETS AND CONFIDENTIAL INFORMATION OF CISCO SYSTEMS, INC. AND IT'S SUPPLIERS, AND SHALL NOT BE DISCLOSED TO ANY PERSON, ORGANIZATION, OR ENTITY UNLESS SUCH DISCLOSURE IS SUBJECT TO THE PROVISIONS OF A WRITTEN NON-DISCLOSURE AND PROPRIETARY RIGHTS AGREEMENT OR INTELLECTUAL PROPERTY LICENSE AGREEMENT APPROVED BY CISCO SYSTEMS, INC. THE DISTRIBUTION OF THIS DOCUMENT DOES NOT GRANT ANY LICENSE IN OR RIGHTS, IN WHOLE OR IN PART, TO THE CONTENT, THE PRODUCT(S), TECHNOLOGY OF INTELLECTUAL PROPERTY DESCRIBED HEREIN.

H-UCS SRND Version 1.6(0) EDCS 580462

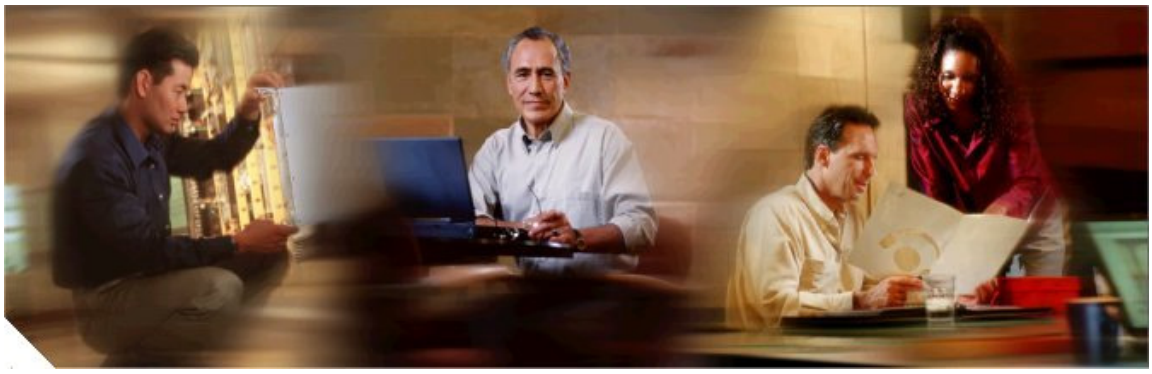

Contents

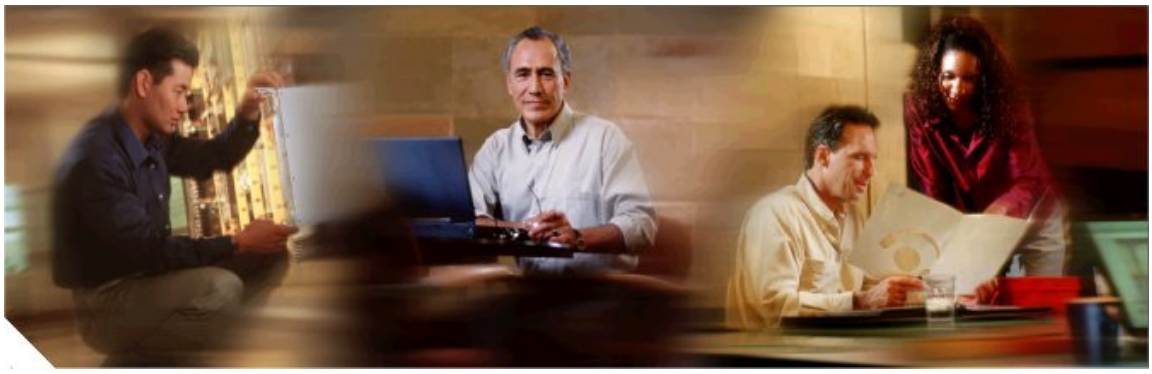

Introduction to the Hosted Managed Unified Communication Service Architecture

The Hosted Unified Communications Service (H/M-UCS) architecture is based on Cisco and some third party industry proven voice products, principally the Cisco Unified CallManager IP PBX platform, the Cisco PGW 2200 softswitch, and the Cisco voice gateway range of products.

Platform provisioning and service management is provided by a third-party operations support system (OSS) service management platform called Business Voice Services Manager (BVSM), which is developed by Vision OSS.

The Merion 3000 or 6000 voicemail systems from IP-Unity are used in the architecture to provide a multi tenant voicemail capability.

The Netwise CMG or ARC Console products are used where customers require an attendant console. The Netwise product is capable of serving multiple tenants from one hosted server platform where as separate ARC connect servers are required for each tenant in a multi tenant H/M-UCS platform.

The products and features that are used to build hosted-managed IP telephony systems are described here together with the requirements that should be considered when combining these products together to deliver a solution.

Cisco Partner Products in the H/M-UCS Architecture

By combining Cisco voice infrastructure products with products from Vision OSS, Netwise, and IP Unity, the H/M-UCS architecture provides service providers with a fully integrated, managed multi tenant voice platform that offers the key set of voice services typically required by business voice customers.

- Multi tenant Provisioning – Vision OSS BVSM
- Multi tenant Voice Mail – IP Unity Merion
- Multi tenant Attendant Console – Netwise
- Single Tenant Attendant Console - ARC

These third party products are not available directly from Cisco. System integrators (and/or the end customer) are therefore ultimately responsible for the purchase, performance and integration of these partner products and into the H/M-UCS architecture to meet the end customer requirements.

Business Requirements

Converged networking is rapidly becoming a key focus for large corporations and service providers who are striving to enhance organizational productivity and deliver significant return on investment. A key element of convergent network projects is the successful deployment and configuration of IP telephony networks that support the critical revenue-bearing voice, data, and video services that employees, customers, and partners expect.

H/M-UCS incorporates a streamlined service management framework that is essential to enable service providers and system integrators to offer a competitive IP based business voice service and services while significantly reducing integration and support costs.

By supporting virtualization and sharing of Cisco CallManager clusters between customers, H/M-UCS is able to preserve the advanced IP communication services offered by Cisco Unified Callmanager in a cost effective way to multiple customers.

Resource sharing in H/M-UCS is accomplished by the ability to virtualise most of the H/M-UCS components through partitioning of both dial plan and administration to ensure they can be shared with near feature parity for end users with Cisco enterprise unified communication deployments.

From a business perspective, the need to share one or more Cisco CallManager systems across different customers or tenants is also driven by the growing demand for IP telephony in the small-to-medium business (SMB) segment. With the advent of broadband VPN and the acceptance of convergence over IP as a way to control operating expenses and to quickly access new network services, managed IP telephony is quickly becoming part of the service bundle required by SMB customers.

The H/M-UCS architecture provides the framework and supporting applications that enable the execution of this efficient system design.

H/M-UCS Platform Partitioning

The key to addressing the requirement to provide managed or hosted IP voice services to their clients is the ability to partition (or virtualise) the service provider's or system integrator's infrastructure to support multiple clients (or tenants) and to devolve, in operational terms, as much of the administration to the clients themselves as possible.

Multi tenant voice services for H/M-UCS customers are based on capabilities provided by Cisco Unified CallManager IP PBX and Cisco PGW IP Class 4 transit switch platforms. The Cisco CallManager provides end user facing services to various tenants, and the Cisco PGW provides a routing function that mediates among all the tenants as well as to and from the PSTN and other zones in the overall architecture.

In the H/M-UCS solution, a Cisco CallManager is logically partitioned to host multiple customers, which provides a virtualization of call control resources. End customers receive broadly the same level and quality of service as if a separate Cisco CallManager system was dedicated to each of them.

Deployment Models

The H/M-UCS solution can be deployed using different models to adapt to different enterprise and service provider requirements for business voice services. The following section gives a brief outline of the various business voice services deployment models that can be considered within the H/M-UCS solution:

- Business Park Voice Services—Limited geographic area and multiple tenants. Examples include commercial business parks and airports, Internet cities, and city government initiatives.
- Large Multi-site Enterprise—Wide geographic area, such as a region, country, or multiple countries with a single tenant. Often a large enterprise may have a number of internal tenants or departments that need to be supported independently. This can be realised in either by the end users company hosting and or owning equipment or a Service provider (SP) hosts the equipment and manages the service.

- Service Provider Hosted—Multiple Small and Medium Enterprises—Wide geographic area and multiple tenants.

Cisco H/M-UCS Components

The following sections summarize the role of each element of the Cisco Hosted Unified Communications Service H/M-UCS solution architecture. Detailed software revisions for each component that forms part of the H/M-UCS 1.6.0 architecture are detailed in Appendix A – H/M-UCS Architecture Release 1.6 Software Versions.

Cisco Unified CallManager

A Cisco CallManager cluster or clusters are deployed within a network provider domain to provide service to IP phones located at a customer end-user facility. A CallManager cluster is a group of servers (typically up to 11 although smaller clusters may be built) running CallManager software that share the same database and therefore effectively behave as one logical IP PBX telephone switch. The Cisco CallManager clusters can be partitioned in a multi-tenant manner to provide segregated service to multiple enterprises deployed or in a dedicated managed manner to support a single large-scale enterprise.

CallManager versions 4.1 and 4.2 are currently supported in the 1.6.0 H/M-UCS architecture. Callmanager 5.0 support is scheduled for a later H/M-UCS release. The major feature that 5.0 brings to the architecture is the ability to provide call control to SIP endpoints as well as the existing SCCP endpoints that CallManager 4.1 and 4.2 can control but support for SIP endpoints within the H/M-UCS architecture is not anticipated in the initial CCM5.0 support phase. It is anticipated that SIP endpoint support will be added in a later release of H/M-UCS (probably 1.7) during 2007. The Cisco IP phone portfolio has just been enhanced with the 79X1 range of phones to provide a much richer feature set when being controlled by SIP than was previously possible with the 79X0 range of phones running SIP. When using an enhanced Cisco SIP phone, the feature set provided is very close to the feature set of a SCCP controlled phone. This enhanced feature set has required that Cisco extend the basic SIP protocol, which has been done in a standards supportable way, not by tunnelling a proprietary protocol over SIP which is an approach often seen in the marketplace. In doing this Cisco have been very careful not to compromise the support of generic (RFC3261) SIP devices.

For further information on the Cisco CallManager please see the datasheets:

CallManager 4.2 -

http://www.cisco.com/en/US/products/sw/voicesw/ps556/products_data_sheet0900aecd8042402c.html

CallManager 5.X -
http://www.cisco.com/en/US/products/sw/voicesw/ps556/products_data_sheet0900aecd8042403e.html

Gatekeeper

An H.323 gatekeeper is included in the H/M-UCS architecture to provide basic infrastructure capabilities. It provides registration capability for the Cisco PGW (via the HSI), local PSTN gateways and Cisco CallManager. The gatekeeper forces all routing to use the Cisco PGW, rather than routing directly between CallManager clusters.

Trunking Media Gateway (AS5400HPX/XM, AS5350, MGX8880 VXSM)

Cisco generally uses the Cisco AS5400HPX or AS5400XM as trunking media gateways in H/M-UCS deployments. Cisco also has the MGX8880 VXSM media gateway available, which provides a very high density gateway solution. The choice of gateway depends upon the total number of T1s or E1s required (in one physical location) to connect to the PSTN. Cisco have other products, such as the AS5350 or even the Cisco 38XX series of enterprise voice gateways, that may be used as the trunking gateway if the scale of

deployment is low or a widely distributed PSTN interconnect is required (such as when providing least cost routing to international destinations).

The Cisco AS5400 will connect to the PSTN through T1 or E1 trunks and will be controlled by the Cisco PGW2200 using MGCP. High-density (Up to 16 E1 or 1 CT3 with 648 simultaneous calls), low-power consumption and universal port digital signal processors (DSPs) make the Cisco AS5400 Series Universal Gateways ideal for many network deployment architectures especially TDM interconnect points of presence (POPs).

For further information on the AS5400 gateway please see the datasheet at:

http://www.cisco.com/en/US/products/hw/univgate/ps505/products_data_sheet0900aecd802efc92.html

Cisco PGW 2200

The Cisco PGW provides TDM PSTN interconnect , Customer TDM PBX interconnect and a Routing and Analysis Engine for the H/M-UCS architecture. It also provides carrier grade call detail records (CDRs) that may be used for billing by the network operator. The PGW also has a SIP stack that is only used to interconnect to the IP Unity Merion Voicemail system in the H/M-UCS 1.6 release but will also be used in future for interconnect to other SIP based service providers.

These are described below.

TDM PSTN Interconnect

The Cisco PGW provides connectivity for all services to the time-division multiplexing (TDM) based PSTN via Signalling System 7 (SS7) or ISDN Primary Rate Interface (PRI), depending on the nature of deployment. The Cisco PGW also incorporates some capabilities that can assist in meeting local regulatory requirements, such as support for the following services:

- Lawful intercept (LI) features, using Cisco Service-Independent Intercept (SII).
- Local number portability (LNP), using an onboard database or a Service Control Point (SCP) query.
- Malicious call identification (MCID), SS7 part only.
- Emergency services, depending on local requirements.

Lawful intercept (LI) features, using Cisco Service-Independent Intercept (SII).

H/M-UCS uses the Cisco Service Independent Intercept (Cisco SII) architecture to provide a network or service operator with the ability to intercept and duplicate voice conversations to or from individual E.164 numbers. The Cisco SII architecture provides only generic abilities to capture these conversations. In all cases a mediation partner is needed to provide the requisite interfaces to a law enforcement agencies.

In the H/M-UCS architecture, the Cisco PGW is used to provide information to the Lawful Intercept Mediation Device to ensure that calls can be appropriately monitored. During the latter stages of call setup the Cisco PGW sends the intercept related information (the RTP traffic send and receive IP addresses and ports) via a RADIUS interface to the Lawful Intercept Mediation Device to allow the device to configure media intercepts on the IP network itself and hence monitor the calls as required.

Business Voice Access (Customer Site TDM PBX Interconnect)

The Cisco PGW optionally provides the interconnect point for customer site based TDM PBXs. TDM PBX gateways are managed directly by the Cisco PGW (using MGCP with backhaul techniques). The Cisco PGW acts as the routing platform (effectively a transit PBX) for all traffic that is generated or

received in this environment. The DPNSS, Q.SIG and Q.931 protocols are supported on these PBX interfaces.

Routing and Analysis Engine

The Cisco PGW provides a routing engine for inter-domain routing. All platform components use the Cisco PGW to route calls that are not local, which ensures that the main dial plan and routing functions for the hosted platform are centrally located. The Cisco PGW includes A and B number analysis and modification functions, as well as regulatory capabilities that can be applied to satisfy local requirements. For SS7 connectivity, both the Cisco IP Transfer Point (ITP) and Cisco Signalling Link Terminal (SLT) signalling gateways are supported by the H/M-UCS baseline architecture.

For further information on the PGW2000 please see the datasheet at:

http://www.cisco.com/en/US/products/hw/vcallcon/ps2027/products_data_sheet09186a0080091b59.html

Voice Mail Platform

The H/M-HCS architecture allows integration with two different Voicemail and Unified Messaging platforms.

Cisco Unity can be used in the architecture to provide a feature rich Unified Messaging platform that is integrated tightly with Microsoft Exchange or Lotus Domino (depending on the version of Unity deployed). However, Cisco Unity is intrinsically a single tenant/customer product due to its tight integration with the customer email platform (it actually uses Exchange or Domino as the voicemail message store).

The IP Unity Mereon product has slightly less features than Cisco Unity, especially in the area of Unified Messaging integration (with Microsoft Exchange and Lotus Domino platforms). However it is inherently a multi-tenant product so provides the ability to have a single system with resources shared across multiple customers.

Other voicemail systems have successfully been integrated into the H/M-UCS architecture to meet specific customer requirements. An example of this is the integration of the VoiceRite unified messaging platform through a Q.SIG interface.

Cisco Unity Unified Messaging Option

Cisco Unity delivers unified messaging and intelligent voicemail capabilities to enterprise and mid-market customers with Microsoft Exchange and Lotus Domino environments. It delivers powerful unified messaging (e-mail, voice, and fax messages sent to one inbox) and intelligent voice messaging (full-featured voicemail with advanced functions).

In the past, e-mail, voice, and fax messages were delivered as separate media to different locations. The telephone provided the sole means for accessing voice messages and then could play messages back only in the order received. Faxes had to be manually retrieved from the nearest fax machine. Cisco Unity Unified Messaging integrates transparently with Microsoft Outlook or Lotus Domino e-mail clients to make handling all message types, e-mail, voice, and fax, easy and convenient, whether you are in the office or on the road. An intuitively designed interface makes it easy to access e-mail, voice, and fax messages from your desktop PC. Icons provide simple visual descriptions of each message type and because every message is delivered to one inbox, you can see the number, type, and status of all your communications at a single glance. You also can reply to, forward, and save your messages, regardless of media type, in public or personal Microsoft Exchange or Microsoft Outlook folders with just a click of the mouse.

With the text-to-speech (TTS) capability of Cisco Unity Unified Messaging, you get information about all your messages-and even hear the text portion of e-mail messages-over the telephone. You can then respond with a voice message and, depending on the capabilities of your fax server, print e-mail, attachments, and

incoming faxes on a nearby fax machine. Cisco Unity Unified Messaging also integrates with smart phones and other mobile devices to deliver all-in-one messaging anytime, anywhere. For example, mobile workers with Treo and BlackBerry devices can double-click to play a voice message within their PDA e-mail applications.

The Unity product itself is not multi-tenant capable so a separate Cisco Unity platform is required for each Customer on the H/M-UCS service

Integration of Cisco Unity into the H/M-UCS architecture is at the CallManager level (using the Cisco SCCP protocol in the H/M-UCS 1.6.0 design. Provisioning of the Cisco Unity environment within the H/M-UCS architecture through the BVSM application is not yet possible so Cisco Unity and the dial plan surrounding it has to be manually provisioned in this platform release. It is anticipated that BVSM provisioning of Cisco Unity will be added in a later platform release.

For further information on the Cisco Unity Messaging System please see the datasheets.

Unity 4.2 for Microsoft Exchange

http://www.cisco.com/en/US/products/sw/voicesw/ps2237/products_data_sheet0900aecd800fe148.html

Unity 4.2 for Lotus Domino

http://www.cisco.com/en/US/products/sw/voicesw/ps2237/products_data_sheet0900aecd800fe14d.html

IP Unity Option

The IP Unity Mereon Unified Messaging platform is used to provide voice mail services in a multi-tenant environment. Only the voice mail functions of the IP Unity platform are currently used in the H/M-UCS architecture.

The IP Unity voice mail system has been selected for its capability to support the multi-tenant H/M-UCS architecture. The BVSM platform is integrated with IP Unity to allow provisioning via the BVSM GUI of voice mail functionality on a per tenant basis. The interface between the PGW and the IP Unity voice mail system is SIP. The PGW forwards the incoming calls to the voice mail system. Once the caller leaves a message, the voice mail system notifies (using the SIP NOTIFY message) the PGW that a message was left for the user. The PGW supports only unsolicited subscription to the voice mail system. This means it does not need to send a SIP SUBSCRIBE message to the voice mail system for every user with voice mail service enabled.

PGW inter-works SIP and H.323 between IP Unity and CCM for Message Deposit, Retrieval and MWI (Message Waiting Indicator). The inter working of IP Unity via SIP and gateway fronted DPNSS/QSIG PBX is specifically not supported in the H/M-UCS 1.6 release so it is not possible to provide a hosted voicemail service for TDM PBX users.

The BVSM platform uses the IP Unity API (CORBA/XML) to define business groups, provision pilot numbers, add/delete mailboxes (assigned against a unique “internal” number and an “extension” number), and assign class of service.

For further information on the IP Unity Mereon please see the datasheet at:

http://www.ip-unity.com/solutions/media_server.asp?Section=solutions

Business CPE

Cisco IOS CPE, such as Cisco 2800 and Cisco 3800 series routers, are used to provide T1or E1 connections to customer site based PBX equipment if required. The Cisco PGW is used to directly control the gateways via D-channel backhaul and MGCP control. In this mode Q931, Q-SIG and DPNSS protocols are supported. The Cisco VG224 and VG248 and ATA analogue gateways can also be used in environments where traditional analogue telephony service is required.

The IOS based CPE may also be used to provide a local gateway capability to connect directly to the PSTN and optionally a backup capability, using Cisco SRST, which provides local call processing to the IP

phones should the WAN link to the hosted CallManagers become unavailable. In this application the CPE gateways are H.323 controlled from the PGW via a H.323 gatekeeper.

The BVSM can be used to automate the provisioning of these two applications.

Intelligent Provisioning Application (BVSM)

Intelligent provisioning and service management is a key feature of the H/M-UCS solution architecture. It provides a unified and integrated view of the Cisco components from a provisioning perspective. This functionality is provided by the BVSM product from Vision OSS.

BVSM is aware of the H/M-UCS architecture and can provision most of the major components, including Cisco CallManager, Cisco PGW, ISC DHCP server and IP Unity voice mail system. In addition to provisioning, the BVSM platform provides some of the non call related features of the service, including the following:

- Multilevel provisioning and tenant self-provisioning
- Devolved administration and user self-care
- Bulk migration of customers
- Number plan management
- Resource management
- Directory services
- Extension mobility
- Inventory Management
- IP Phone Services (XML phone browser) Subscription and Management

Currently BVSM does not provision the central gateway PSTN facing configuration that is required by the Cisco PGW or 100% of the data that is required for Business Voice Access (PBX inter-working). The Cisco Voice Services Provisioning Tool (VSPT) or Cisco PGW command line interface can be used to configure these aspects of the service.

BVSM version 3.1 6 also cannot provision the Netwise attendant console platform or associated dial plan on the Callmanager.

Billing and Measurement Server (BAMS)

In the H/M-UCS architecture call-detail records (CDRs) are generated by two components, the Cisco CallManager and the Cisco PGW. The Cisco CallManager generates CDRs records that allow a service provider to provide telephone usage reports. The CDRs generated by CallManager are not really suitable for billing but can be used to provide telephone use statistics or data about call quality.

Cisco PGW CDRs are aggregated and converted to common CDR formats by the Cisco PGW Billing and Measurements Server (BAMS). They observe carrier class accuracy and contain all necessary timestamps, as well as the called party information delivered to the Cisco PGW and the calling party information sent from the Cisco PGW. They can be used to generate bills for calls to the PSTN, between customers and between sites within a customer if desired.

For further information on the BAMS please see the documentation at:

http://www.cisco.com/en/US/products/sw/voicesw/ps522/products_user_guide_book09186a00800fd123.html

Firewall

The Cisco PIX 5XX series firewall, Cisco ASA or Cisco FWSM products are used in the H/M-UCS architecture to provide 3 functions :

- Securing the service provider owned infrastructure from attack or abuse from the customer networks.
- Securing the customers on the hosted voice service from each other whilst still allowing customers to call each other.
- Catering for customers with overlapping IP address spaces.

All of the firewall products mentioned are application aware in that they have special handing for SIP, SCCP, MGCP and H.323 as far as NAT and security are concerned.

The use of NAT requires that the firewall translates both the addresses in the IP headers and also any IP addresses embedded within the messages within the protocol.

All voice signalling protocols have difficult security requirements in that the signalling is used to negotiate UDP ports to carry the RTP voice streams themselves. Without special handing it is therefore necessary to leave open a large range of UDP ports so compromising the security that a firewall provides. A signalling protocol aware firewall can open UDP ports dynamically for the duration of the call to allow the voice streams to pass and then close off the port again at the end of the call so ensuring maximum security.

When performing this application aware firewall function the performance of the firewall is not usually determined by the pure ability to carry so many packets per second but often by the ability to process the voice signalling messages at the application layer. Cisco have tested the performance of our firewall products and optimized the protocol handling functions as far as is possible. As an example, a single PIX 535 firewall can (at 70% cpu load) handle around 3000 simultaneous calls with an average call hold time of 2 minutes.

One firewall (or an active / standby pair for resilience) is deployed between each customer and the service provider common network. The Catalyst 6000 FWSM also has the ability to provide up to 250 virtual firewall instances on one physical firewall blade, making it a good choice for service providers with large numbers of customers.

For further information on the Cisco Firewall products please see the datasheets at:

http://www.cisco.com/en/US/products/hw/vpndevc/ps2030/products_data_sheet0900aecd8040c5b5.html

http://www.cisco.com/en/US/products/hw/vpndevc/ps2030/products_data_sheet0900aecd8040c5b5.html

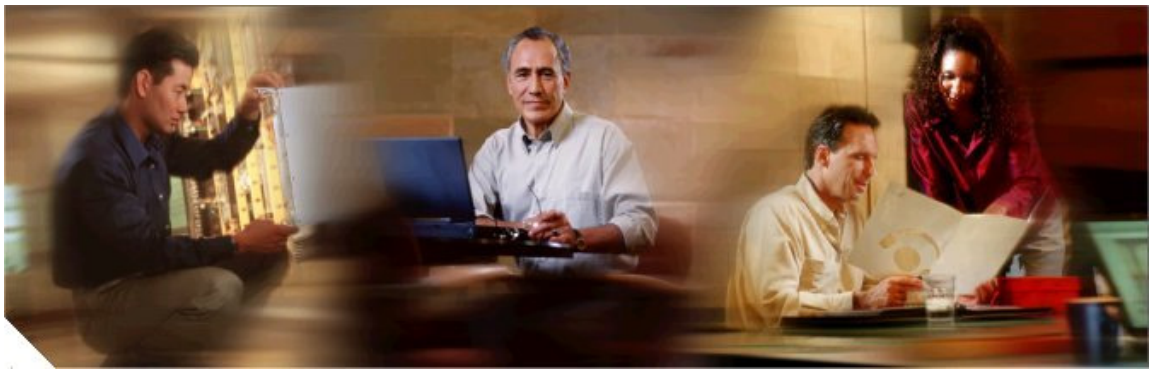

H/M-UCS Architecture Overview

Architecture Description

The baseline voice network architecture for the Hosted/Managed Unified Communications Services (H/M-UCS) architecture integrates the call control capability of a Cisco CallManager system and the routing and services function of a Cisco PGW softswitch to provide a set of capabilities that is useful for a broad range of H/M-UCS architecture deployment models. The resources of the Cisco CallManager system and the Cisco PGW can be shared among several tenants. A provisioning and service management facility, Business Voice Services Manager (BVSM), provides the ability to partition, or virtualise, those resources for individual tenants. In addition, the provisioning facility abstracts customers from the underlying data structures and routing schema and provides additional capabilities to directly support applications such as directory services and extension mobility for IP phones.

The H/M-UCS architecture is illustrated below.

Figure 1 H/M-UCS Network Architecture

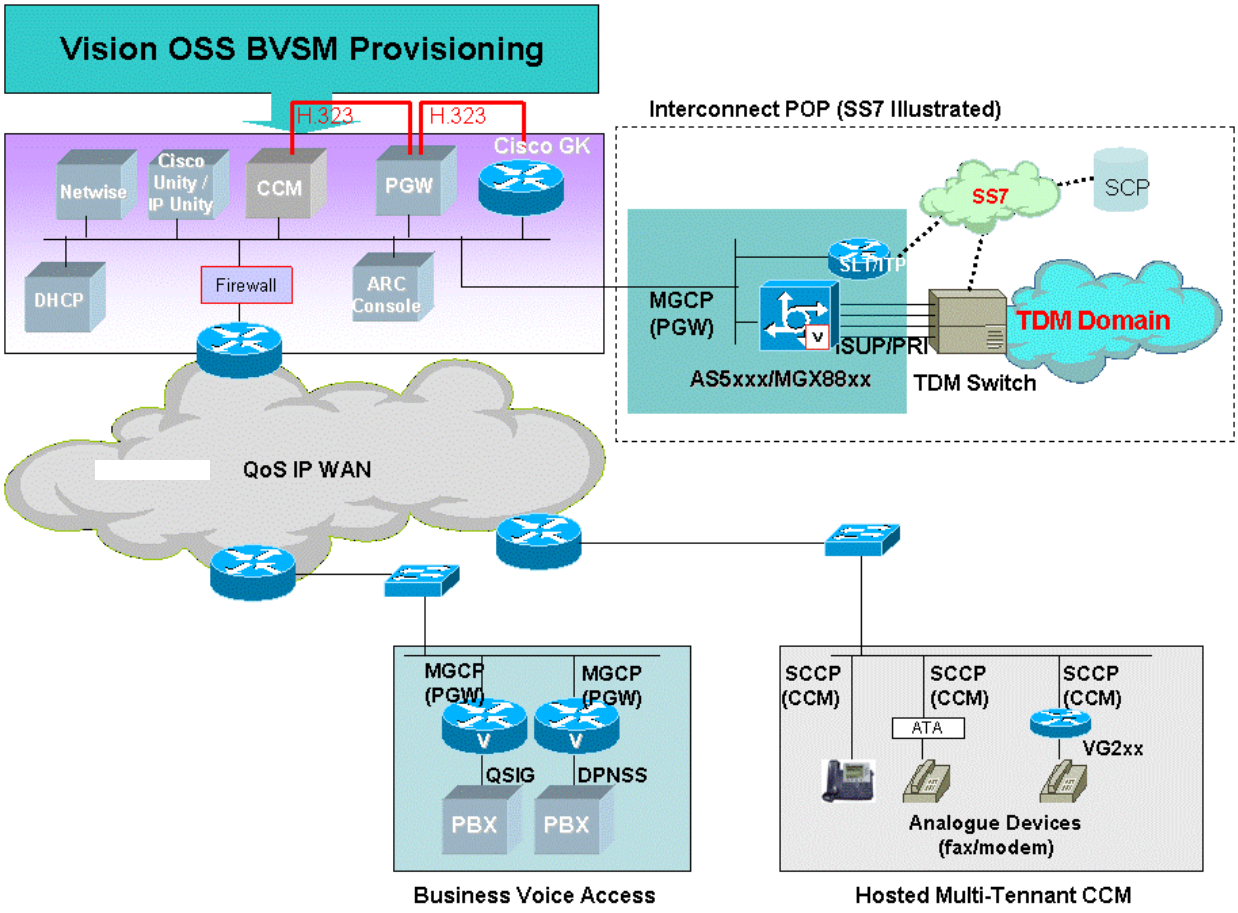

Multi-tenant voice services for H/M-UCS customers (tenants) are based on capabilities provided by Cisco CallManager and Cisco PGW platforms. The Cisco CallManager provides end user facing services to individual tenants, and the Cisco PGW provides a routing function that mediates among all the tenants, as well as to and from the PSTN and other zones in the overall architecture.

Call Scenarios

Different factors must be considered for calls that originate or terminate on Cisco CallManager or in the PSTN. Intra tenant calls within a single site are routed by Cisco CallManager (Case 1), while inter tenant calls, calls between sites and calls to and from the PSTN are routed through the Cisco PGW (Case 2, 3, 4). Figure 2 illustrates PSTN gateway routing using the Cisco PGW.

The following scenarios are described in this section.

- Calls Within a Single Site (Intra-tenant)
- Calls to Other Sites or Tenants (Inter-tenant)
- Calls from the PSTN
- Calls to the PSTN

Figure 2 PSTN Gateway Routing

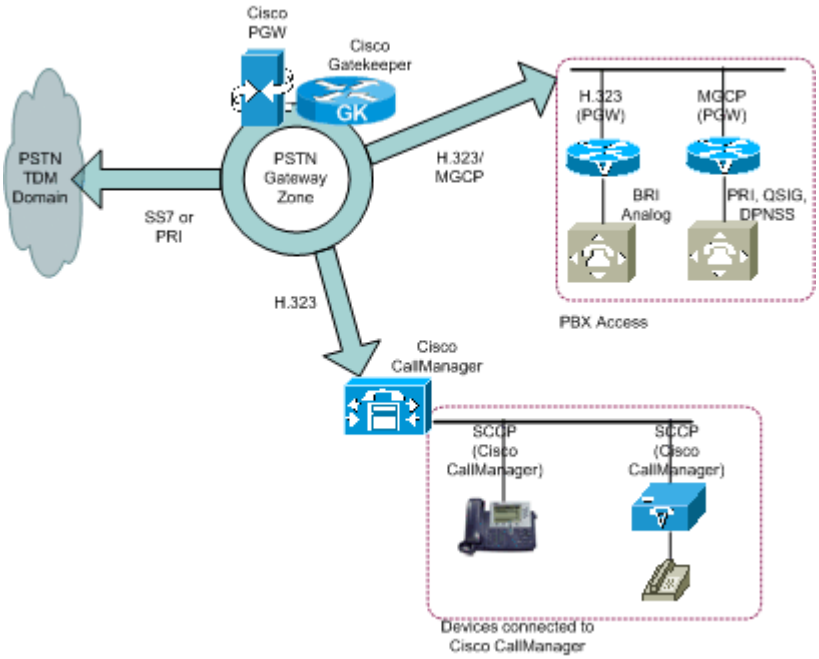

Calls Within a Single Site (Intra-tenant)

Calls that are made from one phone user to another phone user within the same site are also known as intra tenant calls. An intra tenant call is made when one phone user simply dials an extension number of another phone user in the same company. Figure 3 shows an example in which extension 501 in tenant B's company is calling extension 502, which is also in tenant B's company.

In an intra tenant call, the dialled extension number is matched by the internal calling search space (CSS) and the call is routed to the called party, which is presented with the extension number of the caller (three digits in this case). There is no Cisco PGW involvement in intra-tenant calls.

Figure 3 Intra Tenant Call

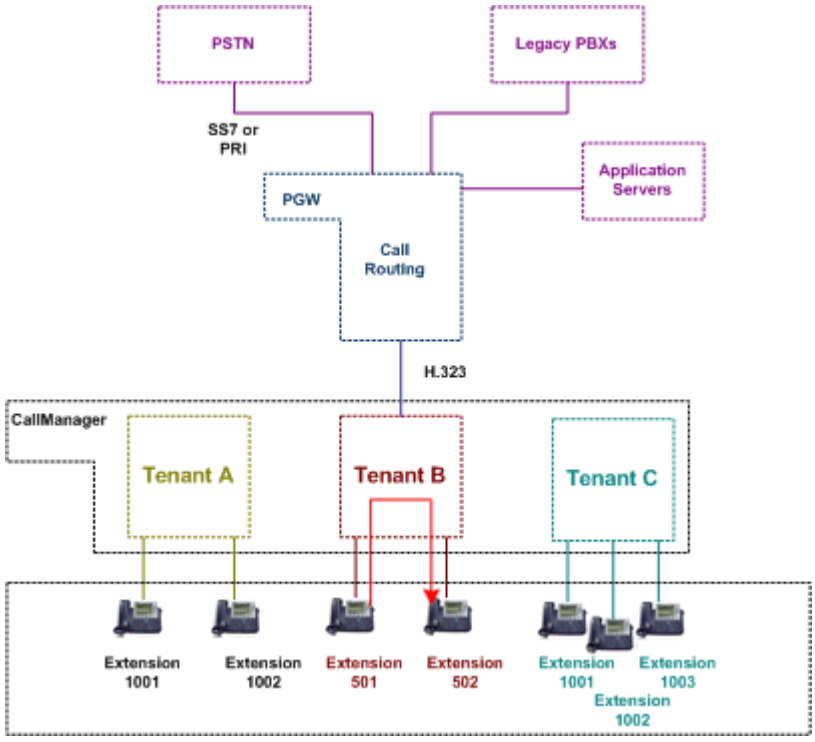

Calls to Other Sites or Tenants (Inter tenant)

Calls that are made from one tenant in an H/M-UCS to another tenant (or calls within a tenant but between two sites) in the same H/M-UCS are known as inter tenant calls. An inter tenant call is made when a phone user in one company dials an offnet prefix and the local number of the destination phone. This is also known as forced on net calling. Note that it is also possible to dial a national or international number for a destination phone and the call remains on-net. Figure 4 shows an example in which extension 501 in tenant B's company is dialling extension 1001 in tenant C's company.

The inter tenant call in Figure 4 is routed to the Cisco PGW and the off-net prefix is matched to a range of numbers assigned to tenant C. The call is terminated to the Cisco CallManager. The media is passed directly between two VoIP endpoints, where as from a signalling perspective the call effectively loops around the Cisco PGW. As the call arrives at the Cisco PGW, the calling-party number that is sent to the destination is changed to the public number of the device that initiated the call. The public number is the calling identity that is presented to the called user.

The Cisco PGW generates CDRs for inter-tenant calls.

Note that the inter tenant call flow logic also applies when a call terminates in another service zone within the VoIP network (such as a call between an IP phone and a directly connected TDM PBX).

Figure 4 Inter-tenant / Inter-site Call

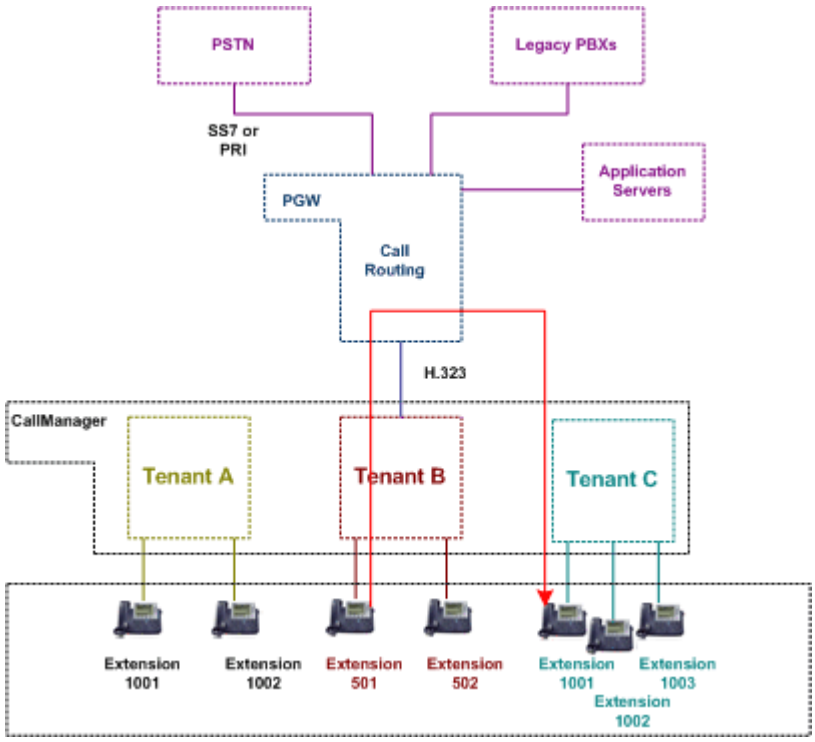

Calls from the PSTN

Calls from the PSTN are received at the H/M-UCS platform by the Cisco PGW at an E.164 number that has been assigned to one of the phones on a Cisco CallManager. Figure 5 shows a call from the PSTN that is routed to extension 1001 in tenant C's company. The equivalent E.164 number for extension 1001 is 498115551001.

When the call is received by the Cisco PGW, the E.164 prefix of 49811555 for the called party number (CdPn) is matched to tenant C and the appropriate dial plan is accessed. An H.323 SETUP message is sent to the Cisco CallManager. The Cisco CallManager matches the destination number and presents the call to the user.

Figure 5 Call from PSTN

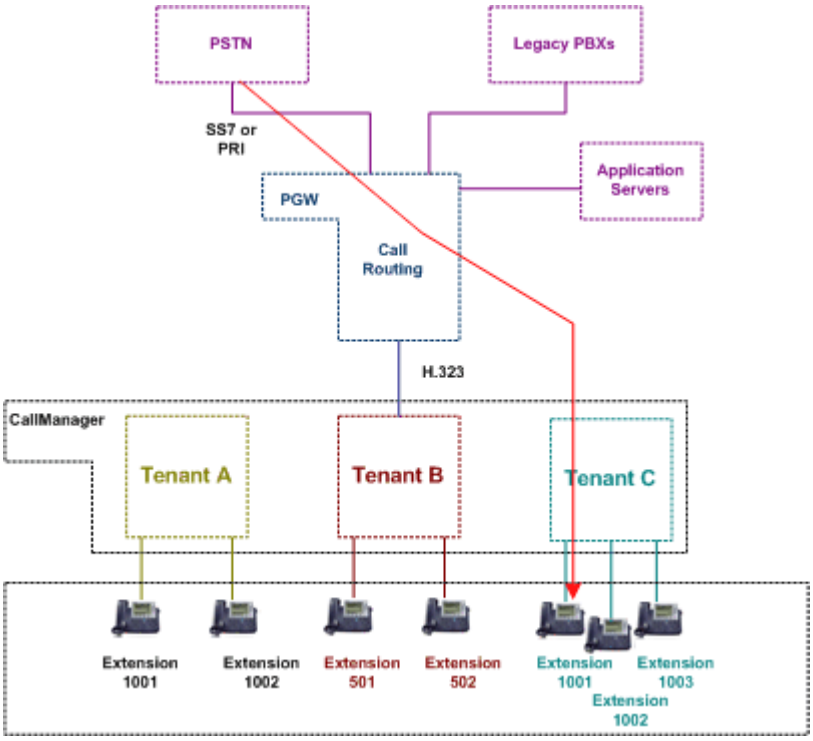

Calls to the PSTN

Calls to the PSTN are made when users from a particular H/M-UCS tenant make calls to external numbers by first dialling an off-net prefix, such as 9, followed by a local, national, or international number. The Figure below shows a call from extension 501 in tenant B's company to an international number, 00442088248777.

When a user first dials an off-net prefix (for example, 9) and then dials a full E.164 number for an international destination, the call is forwarded to the Cisco PGW. The SETUP message received by the Cisco PGW contains the called-party number (CdPN) with the full number dialled (including the international prefix) and the calling-party number (CgPN) of the full internal number of the device that is making the call. The Cisco PGW maps the CgPN to its equivalent international E.164 number. The Cisco PGW analyzes the CdPN and determines that this call is an international call. The call is routed to its destination. Note that in the case of an international call, the country code is also added to the beginning of the CgPN.

Figure 6 Call to PSTN

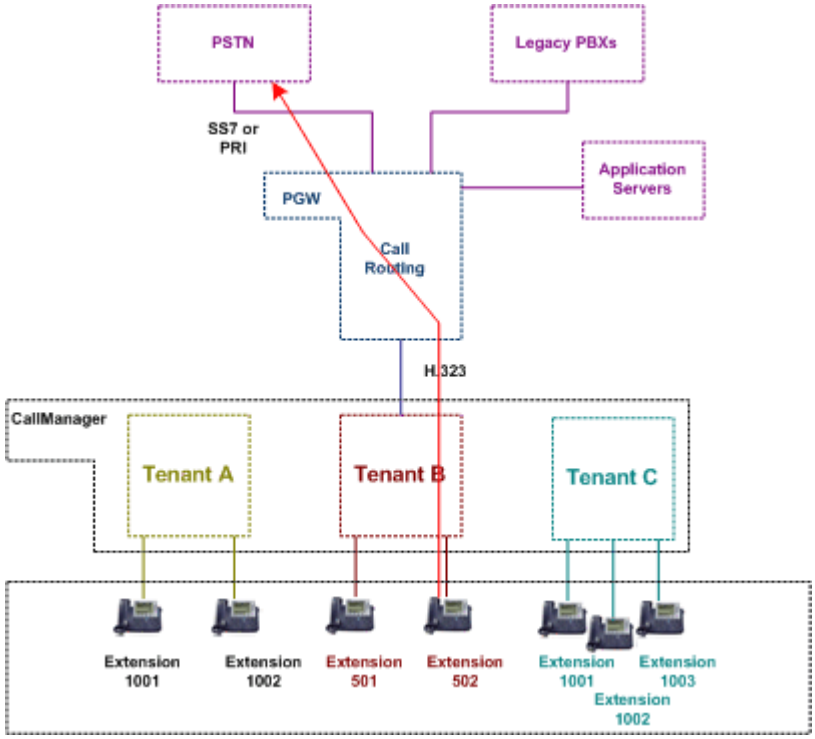

TDM PBX Integration

Customer site TDM PBXs can be connected into the H/M-UCS platform so that users can make and receive calls through a PBX to and from PSTN, between a PBX and H/M-UCS managed IP phones and between two TDM PBX.

This feature was supported in the large enterprise deployment model in the 1.5 H/M-UCS release and has now been extended in this 1.6 H/M-UCS release to include TDM PBX support in the hosted multi tenant deployment model.

PBX integration uses a Cisco PGW MGCP controlled VoIP gateway to provide connection to the TDM PBX using either the DPNSS, Q.SIG or Q.931 protocols. The Cisco PGW supports the necessary TDM signalling protocols and also provides the call routing function.

PBX devices can be connected into the H/M-UCS platform using the methods shown below. MGCP plus layer 3 protocol backhaul is used to supports E1 or T1 PRI based interfaces (including country variants as well as network and user mode) as well as Q.SIG and Digital Private Network Signalling System (DPNSS).


Note

CAS is not supported as a PBX interconnect method in the 1.6.0 H/M-UCS architecture.

Figure 7 PBX Services Integration

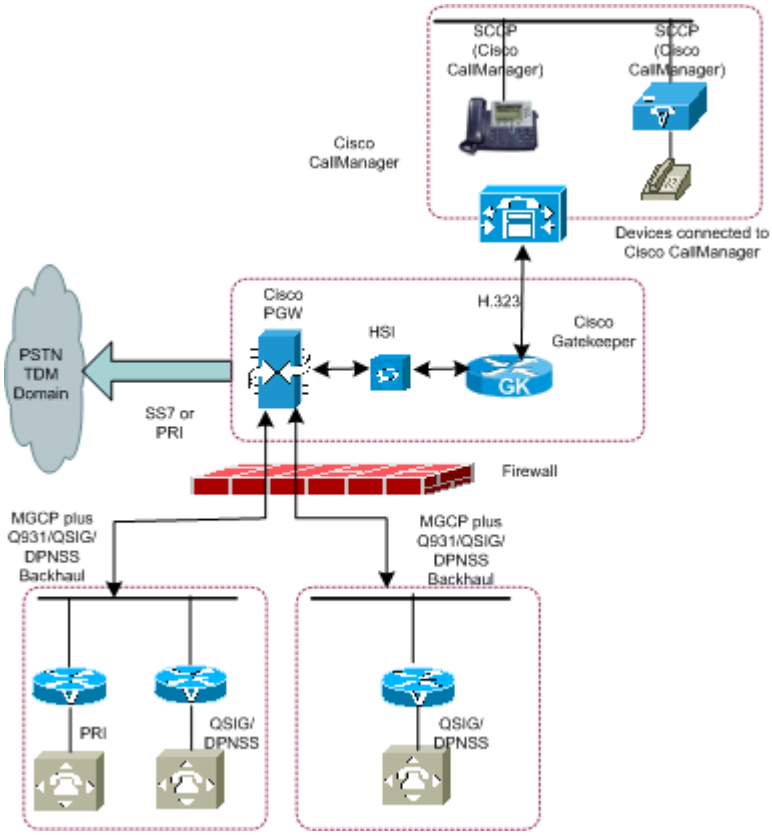

MGCP Connectivity

MGCP provides the simplest control mechanism for PBX gateways and, with backhaul of the layer 3 protocol, provides the most TDM like mechanism for connecting PRI, Q.SIG, and DPNSS PBXs via managed routers into the VoIP domain. MGCP does not work in isolation and requires a gateway to transport the TDM-based protocol over IP directly to the call control platform (in this case the Cisco PGW). In this model, the gateway is essentially passive, forwarding the signalling D channel to the Cisco PGW and acting totally under its control (via MGCP). Consequently there is very little configuration required in the gateway to support this model. All that needs to be done is to direct the D channel and the MGCP control interface to the Cisco PGW.

Table 1 MGCP standards that are supported by H/M-UCS

Reference	Protocol	Notes
RFC3435	MGCP1.0	
RFC2960	SCTP	Stream Control Transmission Protocol
RFC3057	IUA	ISDN Q.921-User Adaptation. Used for ISDN backhaul on some Cisco gateways.
	DUA	DPNSS User Adaptation Layer
N/A	PRI-backhaul	Pre Sigtrans backhaul technology. Used on gateways that have not yet implemented SCTP or IUA.

Reference	Protocol	Notes
Q.931	ISDN	European Telecommunication Standards Institute (ETSI) and country variants are also supported.
Q.699	ISDN ISUP	Mapping of ISDN to ISDN User Part (ISUP) and vice versa.
ECMA v1, ETSI	Q.SIG	In the H/M-UCS architecture QSIG connectivity is provided by the PGW rather than the CallManager.
BTNR188	DPNSS	In the H/M-UCS architecture DPNSS connectivity is provided by the PGW.

A Cisco PGW signalling interface assumes that a PBX is connected to it, so it provides most of the features that a TDM switch would provide across a similar interface.

Bearer Services Support

For TDM PBXs, the bearer services supported are the following:

- Speech (3.1KHz audio)
- Analogue Fax

Routing Cases

The following routing cases are supported by the H/M-UCS service:

- PBX to and from the PSTN (Multi tenant deployment model only)
- PBX to and from a PBX or IP phone belonging to a different customer on the H/M-UCS platform
- PBX to and from another PBX within same customer
- PBX to and from a H/M-UCS controlled IP phone within the same customer.

PBX Integration Billing

For PBX integration scenarios, the Cisco PGW forms the hub of the Business Voice Access service, regardless of the protocol that is used to communicate to endpoints. The Cisco PGW sees every call in this service domain, and the CDRs that are generated from the calls can be used to bill for this service.

TDM PBX Interconnect Number Formats

The table below identifies the format of Called and Calling numbers on the TDM trunk to a PBX connected to the H/M-UCS platform.

Table 2 Number Format on Trunks to TDM PBX

Call Description	Number format of CallED (B) Number	Number format of CallING (A) Number	Notes
Outgoing call from PBX to another PBX user within same tenant	Site code + Extn of destination	Site code +Extension of originator	Note : When using TDM PBX site codes cannot begin with the same digit as the PSTN access prefix (e.g. 9 in the UK)
Incoming call to PBX from another PBX within same customer	Site code + Extn of destination	Site code +Extension of originator	
Outgoing call from PBX to an IP phone within the same customer	Site code + Extn of destination	Site code +Extension of originator	
Incoming call to PBX from an IP phone within the same customer	Site code + Extn of destination	Site code +Extension of originator	
Outgoing call from PBX to PSTN (or to IP phone or PBX belonging to a different tenant on H/M-UCS platform) (Available in multi tenant H/M-UCS dial plans only)	PSTN access prefix (e.g. 9 in UK) followed by the called E164 number.	Site code + Extension of originator	The E164 number must be a national number prefixed with the national prefix for the country concerned (e.g.0 in UK) or an international number prefixed with the international prefix for the country concerned (e.g. 00 in the UK)
Incoming calls from PSTN to PBX PSTN (or from an IP phone or a PBX belonging to a different tenant on H/M-UCS platform) (Available in multi tenant H/M-UCS dial plans only)	Site code + Extn of destination	PSTN access prefix (e.g. 9) plus E164 number of originating party (note PSTN access prefix is not presented on calls that have been forwarded)	

Number Format Adjustment

Each PBX has two unique dial plans provisioned into the PGW. There is an ingress dial plan which receives calls from the PBX and there is an egress dial plan which passes calls to the PBX. BVSM handles the allocation of dial plan names automatically. The contents of these two dial plans may be manually provisioned to manipulate the PBX number formats between the PBX and the PGW. When manually provisioning digit manipulation it is necessary to understand which two dial plans are allocated to a particular PBX. To determine this, use BVSM and navigate to General Administration – Customers, select the Reseller/Customer where the PBX resides and then select the 'View PGW Config' button at the top of the 'Customer Management' screen. The correct PGW will then need to be selected on the next screen presented. The Customer Configuration screen is then shown as below.

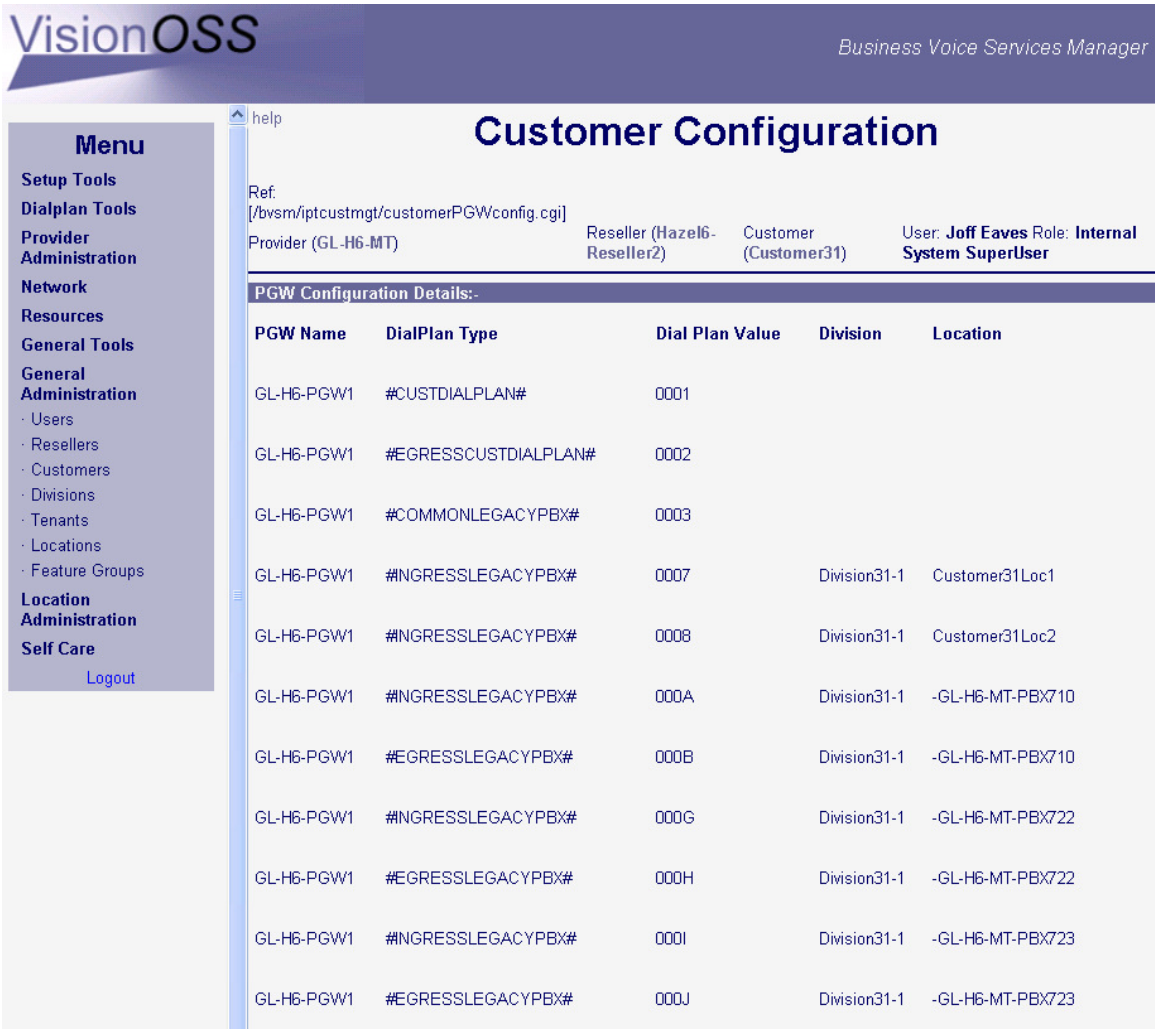

From this screen the user can determine which ingress and egress dial plans are allocated per PBX, e.g. PBX710 is allocated dial plans 000A and 000B as ingress and egress dial plans respectively. As a further note, the CUSTDIALPLAN is shown here as is the COMMONLEGACYPBX dial plan which are allocated per Customer and are also used in the TDM PBX call flows.

Caveats with H/M-UCS TDM PBX Implementation

BVSM does not provision the IOS gateways used for TDM PBX interconnect in this release, only the PGW is provisioned to support the gateway.

TDM PBX interconnect has so far only been designed with MGCP gateways connected to the PGW, H323 gateways and MGCP gateways to the CCM are not supported.

It is currently assumed (1.6.0 architecture release) that there is a one to one mapping between PBX internal numbers and E164 PSTN numbers i.e. a PBX site code and extension may be (123)789 and the E164 PSTN number would be <Area Code>123789. This was felt to be too restrictive so it is anticipated that this issue will be resolved in a near term maintenance release.

After a Gateway or E1 has been provisioned into the PGW, BVSM does not put the associated functions into service. Similarly with deleting, BVSM does not take the required functions out of service which will allow the E1 or Gateway to be removed from the PGW. Work on the BVSM is needed for this to be handled automatically and this will be the subject of a maintenance release. As a workaround, the service

state must be manipulated manually on the PGW. If the gateway name is c3745 for example, after adding the first E1, the following commands are required:

set-iplnk:iplnk1-c3745:IS
set-iplnk:iplnk2-c3745:IS

To delete the last E1, the following commands need to be manually entered in order to allow the PGW to remove the configuration:
set-iplnk:iplnk1-c3745:oos,confirm
set-iplnk:iplnk2-c3745:oos,confirm
set-dchan:dchan-c3745xy:oos – If protocol is PRI or QSIG (where xy is the gateway port assignment)
set-association:ass-c3745:oos,confirm – If protocol is DPNSS
set-iproute:ipr-c37451:oos,confirm
set-iproute:ipr-c37452:oos,confirm

If the PGW reports that a restart is required due to certain mml properties (e.g. MGCPHEARTBEATINTERVAL, OVERLAP, CUSTGRPID, MGCPDOMAINNAMEREMOTE) being set, BVSM does not restart the PGW. Administrators must schedule a restart and execute it manually on the PGW at a convenient time.

The DPNSS feature Callback and Q.SIG Call Completion are supported in the H/M-UCS multi tenant deployment model but only for calls between TDM PBX. Callback between DPNSS and CCM endpoints is not supported.

Geographic Redundancy

Geographic Redundancy is splitting the deployment of all the elements of the reference architecture between two geographically dispersed data centre sites (or Points of Presence (POPs) for disaster recovery purposes. L3 connectivity is required between the POPs.

CallManager Clustering Over the IP WAN to achieve Geographic Redundancy

A Cisco CallManager cluster can be deployed across multiple geographically distributed sites that are connected by an IP WAN with QoS features enabled. Each of the data centre sites will contain at least one Cisco CallManager subscriber in every cluster. The IP WAN between the data centre sites needs to be engineered to meet the following requirements.

- 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 Signalling (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 26 or PHB AF31). 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 SQL, LDAP, and other inter-server traffic.
- Signalling or call control traffic requires bandwidth when devices are registered across the WAN with a remote Cisco 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 section in the CallManager SRND.

- A maximum round-trip time (RTT) of 40 ms is allowed between any two servers in the Cisco 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.

Note

Under a site failure condition, such that the active site does not have the CCM cluster publisher some functions will not be available. Additions and modifications cannot be made to any part of the Callmanager configuration. Extension mobility users will not be able to log in or log out of the IP phones. Changes to phone call forward settings will not be possible.

In the H/M-UCS reference architecture two sites are used for geographic redundancy as shown in the figure below.

Figure 8 Callmanager Geographic Redundancy

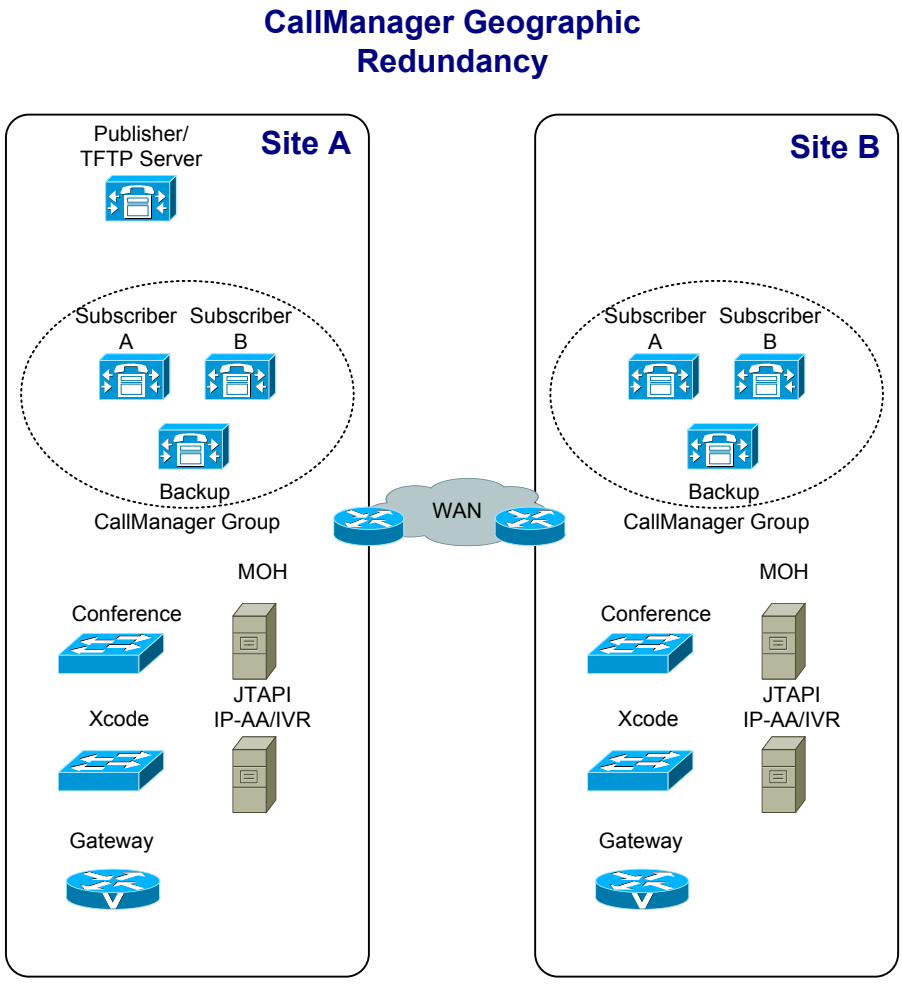

Separation of Cisco PGW Active and Standby Hosts over the IP WAN to achieve PGW Geographic Redundancy

The Cisco PGW 2200 product allows separation of Active and Standby hosts over a campus or a wide area network (WAN) from software releases 9.3(2) or later.

The table below outlines the criteria for geographic separation of the active and standby PGW 2200 hosts.

Table 3 PGW Geographic Separation Criteria

Condition	Requirement
Software release version	Cisco PGW 2200 software 9.3(2) or later (with associated operating system and hardware requirements).
Total end-to-end delay, one way	Must be less than 150 milliseconds.
Packet loss	Must not exceed 1% (preferably, less than 0.5%). Note : For packet loss rates below 0.5%, increase the RUDP receive window size (*. rudpWindowSz) to 64 for increased performance.
XECfgParm setting for replicator.reconnectInterval	This parameter defines the reconnect interval in number of seconds for the replicator during a switchover. Value: Any integer Default: 15 seconds Note – Set this value to 0 for a standalone Cisco PGW. For geographically separated Cisco PGW pairs, the following replicator timer values are recommended: On one PGW, replicator.reconnectInterval = 15 seconds On the other PGW, replicator.reconnectInterval = 20 seconds If the timer settings are the same on both Cisco PGWs, the additional latency between the separated Cisco PGWs may cause a problem in which the replicator links are continually reconnecting and then immediately disconnecting. The timer change prevents this problem.

The figure below shows the geographically redundant distribution of PGW components to be used in the H/M-UCS reference architecture.

Redundancy of the HSI component of the PGW architecture is achieved by having several HSIs active at the same time. For geographic redundancy this concept is extended to ensure that the HSIs present at a given data centre site can carry the total load required under data centre failure conditions. The dial plan on the PGW and the gatekeeper layer will ensure that calls are only presented to HSIs that are active and reachable from the active PGW.

Other components such as SLT for SS7 interconnect and the VoIP media gateways themselves can be replicated across data centres and scaled as required to carry all the calls under data centre failure conditions.

Figure 9 PGW Geographic Redundancy

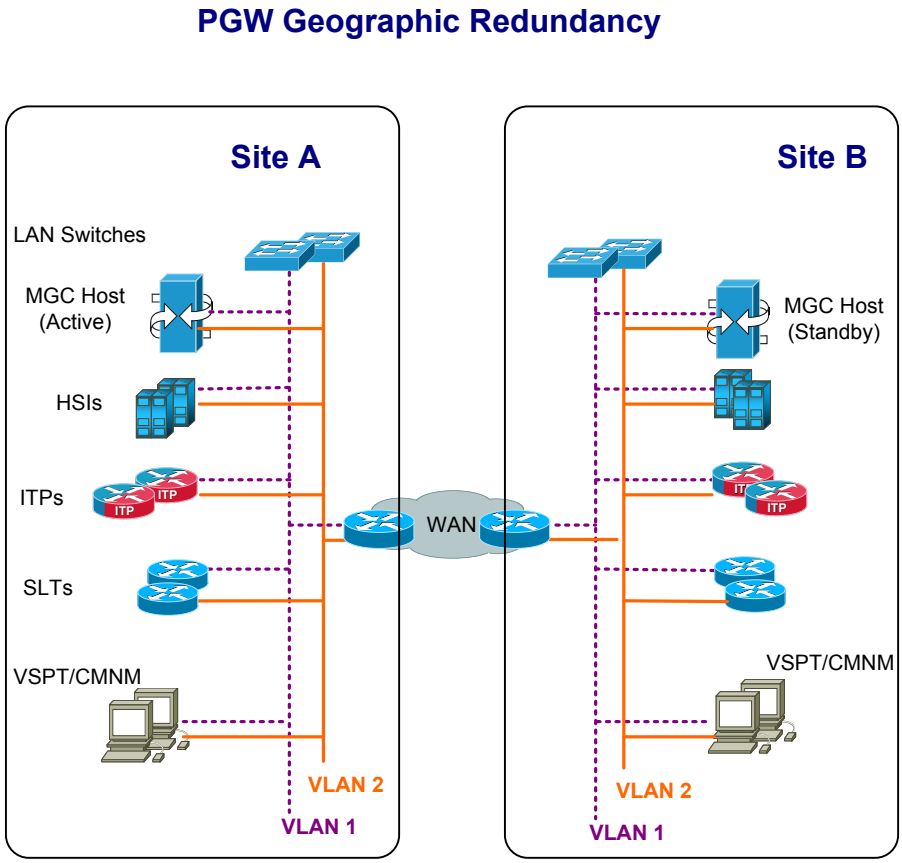

Gatekeeper Geographic Redundancy

Gatekeeper clustering is when two or more gatekeeper routers are grouped together and work in tandem to provide routing services to H.323 endpoints. One gatekeeper is designated as the primary and the others are designated as the alternates.

The redundancy design uses the configuration of gatekeeper clustering which allows for automatic registration, call routing and CAC redundancy between Gatekeeper routers within a Gatekeeper cluster. Gatekeeper clustering can provide for a maximum 5 Gatekeepers within a Gatekeeper cluster.

Clustered gatekeepers communicate via Gatekeeper Update Protocol (GUP). GUP is a TCP based Cisco proprietary protocol. GUP is used between Alternate Gatekeepers. This allows individual gatekeepers to collect group information such as registrations, bandwidth and individual load.

When endpoints first register, they are given a list of alternate gatekeepers in the RCF message in a priority order. The priorities are determined by the capacity available at each of the alternate gatekeepers, as reported to the primary gatekeeper, by the GUP announcement message. This list is updated with every RCF (for light weight RRQs) and the priorities are adjusted. In case of primary gatekeeper failure, the endpoints register with the gatekeeper highest in the priority order.

Gatekeeper Clustering using the GUP protocol has been designed to be used in remote redundant designs. The GUP protocol is not bandwidth intensive as it only sends two types of messages to each gatekeeper in the cluster. The first message is a registration announcement message which is sent by a gatekeeper whenever a new endpoint registers. This originates one or two IP packets and would be expected to be a very rare occurrence under normal circumstances.

The second type of message is a status update announcement message which occurs every 30 seconds which details the state of the gatekeeper and the zone status for all the zones managed by that gatekeeper. This however would still equate to a limited number of packets every 30 seconds.

The GUP protocol is TCP based and so should not be affected by packet loss. The protocol is based on 30 seconds announcements and is therefore also not affected by delay in any significant manner.

The figure below shows the design for geographically redundant gatekeeper clustering over the WAN to be used in the H/M-UCS reference architecture.

Figure 10 Gatekeeper Clustering Between Geographically Separated Sites

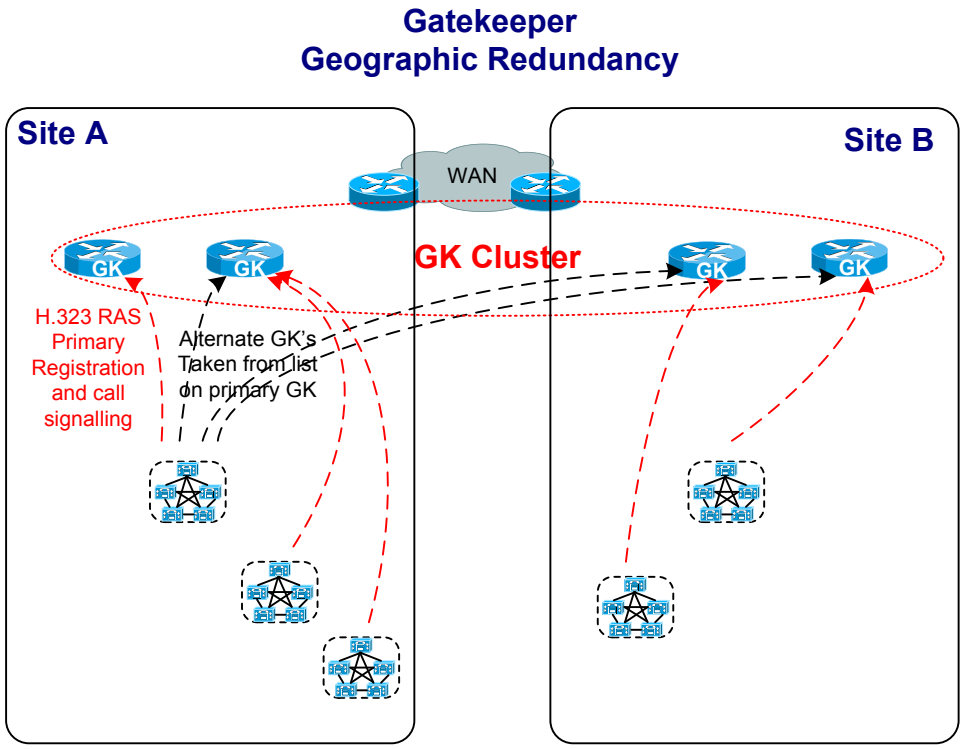


Note

With this gatekeeper architecture it is essential that the primary gatekeeper is available when new CallManagers are introduced to the system or CallManager subscribers are rebooted for any reason. Only the primary Gatekeeper is provisioned and known to the Callmanager cluster. The alternate gatekeepers are discovered through the initial registration with the primary gatekeeper, hence the need for the primary Gatekeeper to be available for initial registration. Once initial registration has sucessfully occurred the CallManager cluster is aware of, and uses, the alternate gatekeepers within the gatekeeper cluster.

BVSM Geographic Redundancy

There are key differences between Disaster Recovery and Automatic Failover systems.

Automatic Failover systems are designed to have multiple nodes where one or more nodes are 'active'. In the event of a node failing, the system is designed so that other nodes can automatically take over processing without system downtime or manual intervention.

Automatic Failover is used within a Vision OSS BVSM high availability cluster. Typically a 'VOSSDIR1' node and a 'VOSSDIR2' node operate in a Primary/Secondary configuration with automatic failover between them in the event of failure. Additional BVSM engines can be added to the cluster to provide

additional scalability. Clusters are designed such that the failure of a node does not effect the operation of the system and processing will continue without interruption.

It is not possible with BVSM version 3.1.6 to split a BVSM high availability cluster across two sites. Within the cluster, the provisioning database needs to be kept synchronized between the two VOSS Director appliances in real-time. Within a locally connected cluster this is achieved through the use of high speed disk replication. This is currently not supported over a wide area network link due to concerns such as with reliability of the cluster due to insufficient WAN bandwidth or too much latency preventing the replication from occurring in real time.

For disaster recovery purposes therefore a second BVSM high availability cluster is needed at the disaster recovery site. BVSM provides Disaster Recovery capabilities between "active" and "disaster recovery" clusters. This capability ensures that if an event completely takes down the 'active' cluster without warning, full service can restored via the 'disaster recovery' cluster without the loss of provisioning information.

The key to providing an effective Disaster Recovery capability for BVSM is to ensure that the underlying database of provisioning information is synchronized between the 'active' and 'Disaster Recovery' sites.

Within a BVSM cluster, the primary 'VOSSDIR' appliance stores all provisioning changes using a Postgres SQL database. This is continually replicated to the secondary 'VOSSDIR' appliance within the same cluster via high-speed disk replication.

The mechanism adopted by Vision OSS to replicate the database between "active" and "disaster recovery" sites is to use software provided by the SLONY-I project (www.slony.info) that has been specifically designed to provide master – slave database replication for the Postgres SQL database.

SLONY works by replicating *sets* of data between *nodes*. In the case of BVSM, the provisioning database can be considered as a single *set*. SLONY can then be configured to manage the replication of data from one master node to one (or more) slave nodes. Replication between nodes is managed on a database transaction basis to ensure that the database being replicated is always in a deterministic position.

The diagram below depicts the sequence of events that occur when a failure of the active site occurs. Note that initiation of these events is under manual control by the network operations staff.

Figure 11 BVSM Disaster Recovery Sequence of Events

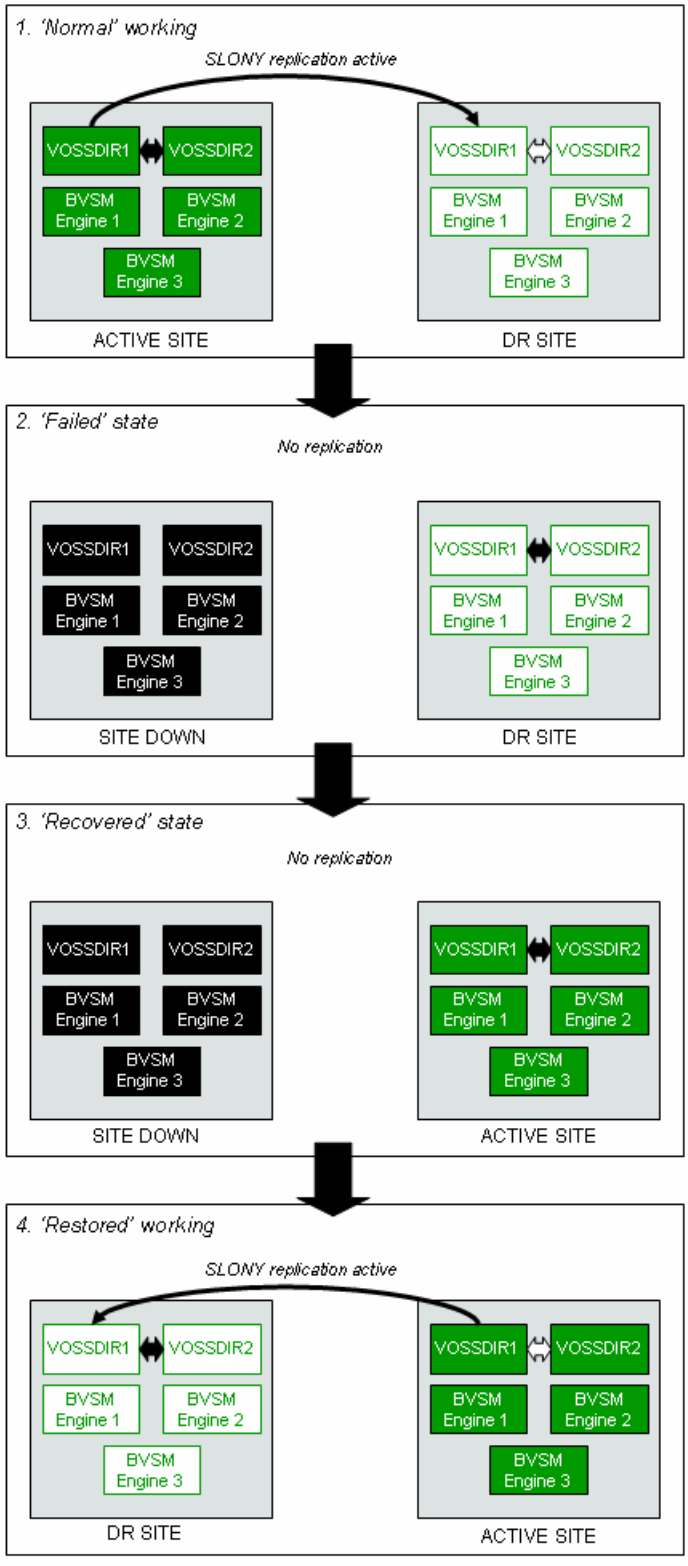

i. A 'disaster' has affected the Active site such that the whole site is 'down'.
ii. Provisioning is not possible immediately following the site failure.
iii. All transactions up to the point of the failure will have been replicated to the DR site.

i. The system operator decides to restore provisioning capability using the DR site.
ii. The DR site is made Active.
iii. Provisioning requests are pointed to the 'new' Active site (previously DR).
iv. Provisioning resumes.
v. There is no 'real-time' replication of transactions outside of the now Active cluster.

i. The failed site is restored.
ii. SLONY is configured to replicate from the 'Active' site to the restored site that is now in the role of DR.
iii. All transactions conducted since the site failure replicate across.
iv. There is no need to switch back the 'Active' and 'DR' roles to the original configuration unless designed by the system operator.
v. Provisioning continues throughout this stage.

Other design constraints when using BVSM in a disaster recovery situation need to be considered if there is a need to the need to maintain the phone directory service or extension mobility access. Extension mobility may not be available anyway in a DR situation due to the Callmanager Publisher being

inaccessible. Both of these services are hosted on the BVSM itself in the H/M-UCS architecture. If the "Directories URL" or EM service URL are configured on the phone using an explicit IP address, then in the event of a failure of the Active BVSM cluster, the phone will lose access to the phone directory. The preferred mechanism therefore is for the phone directory to be provisioned as a DNS name rather than an explicit IP address. This mechanism requires a DNS server to exist within the network and be managed correctly so that in the event of 'disaster' the phones will access directory and EM services via the DR cluster.

IP Unity Voicemail Geographic Redundancy

The Mereon Unified Messaging platform leverages its distributed IP-enabled clustering architecture at all layers, to enable the design of a wide range of deployment options that can support blended telephony and Internet traffic in a geographically distributed and redundant manner.

In the event of intra-site failure, the Mereon application server has built-in high availability for each of the subsystems, such as the mail server, web subsystem and call processing nodes.

To achieve geographic redundancy, the Mereon Unified Messaging platform integrates application server cluster(s) and media servers across multiple sites. Data synchronisation between active and DR sites is achieved using the synchronous data mirroring tools provided by the NAS vendor. Each site has a peer backup site where the NAS data is replicated. Real-time one-way mirroring of data is performed from the active site to the backup site to ensure that no messages or database updates are lost. This NAS replication often requires high bandwidth, low latency and even Layer 2 connectivity between sites.

In the event of whole site level failure, the call and web sessions are redirected to the peer backup site. The backup site is engineered to have the capability to handle the normal busy hour traffic during the site backup mode. The failover is accomplished in one of two ways:

1) remapping the IP addresses of the original DN to the surviving set of IP Address or

2) re-routing to an new DN in the PGW soft switch.

During a site failure, the failed site's voicemail calls as well as MWI Notification and web access are managed by the equipment at the peer backup site.

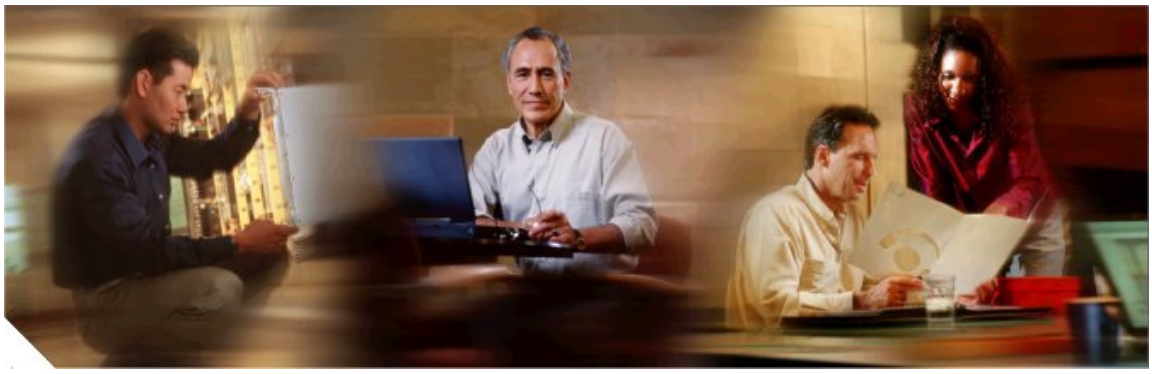

Deployment Models

The H/M-UCS solution can be deployed using different models to adapt to different enterprise and Service Provider requirements for business voice services. It was initially developed and tested for a geographically constrained business park or Internet city model consisting of a single infrastructure to support multiple single-site small businesses concentrated in a specific area. However, the H/M-UCS architecture has since been enhanced to offer additional deployment models in which either a centralized Unified Cisco CallManager can host multiple tenants/customers that are distributed over a wider region, or multiple, distributed Unified Cisco CallManager clusters can host a distributed single customer.

Business Voice Service Applications

The following sections give a brief outline of the various business voice services applications that can be considered within the H/M-UCS solution:

Business Park Voice Services - limited geographic area and multiple tenants. Examples include commercial business parks and airports, Internet cities, and city government initiatives.

Large Multi-site Enterprise – wide geographic area, such as a region, country, or multiple countries with a single tenant. Often a large enterprise may have a number of internal tenants that need to be supported independently

Service Provider Hosted - multiple Small and Medium Enterprises, wide geographic area and multiple tenants.

Business Park Voice Services

The initial focus of the H/M-UCS solution architecture and testing effort is a multi-tenant hosted IP telephony architecture that concentrates on business voice services. These services are most often offered by Service Providers to commercial customers in Business Park and Internet city developments, but are sometimes provided by commercial customers themselves in an enterprise deployment.

This type of service has the following characteristics:

- Managed by a Service Provider or by the enterprise itself.
- Targeted at small-to-medium businesses or local government departments, typically with only one physical site.
- Hosted IP telephony and TDM PBX connectivity is provided.
- Central or local connectivity to PSTN is provided.
- No multi-site or voice VPN capabilities are required.

Such architecture can be extended to support multiple business parks at a national or international level using shared infrastructure where possible.

Large Multi-site Enterprise Voice Services

In the single enterprise multi-site model, service is provided to a single organization that is spread across many locations. The H/M-UCS architecture can easily support large multi-site enterprises that want to consolidate their voice services infrastructure. This model can be managed by the enterprise itself or hosted by a service provider. Management of voice services can be handled centrally or devolved to branches that wish to manage their own services.

Figure 12 shows an example of a large multi-site enterprise model, in which a single large enterprise with multiple sites is provided with telephony service from equipment centrally located in a data centre. At each remote location, IP phones and PSTN gateways can be deployed, as well as gateways to TDM PBXs, which are not illustrated. The service can be controlled centrally by a single administrator for all sites. Alternatively, administrative domains can be created so that tasks are devolved to local site administration personnel.

In a multi-site deployment model for a single enterprise, off-net traffic can be breakout either centrally, locally, or a combination of both. Breakout points can be added wherever they are needed by simply adding small VoIP gateways.

Figure 12 Single-Enterprise, Multi-site Deployment

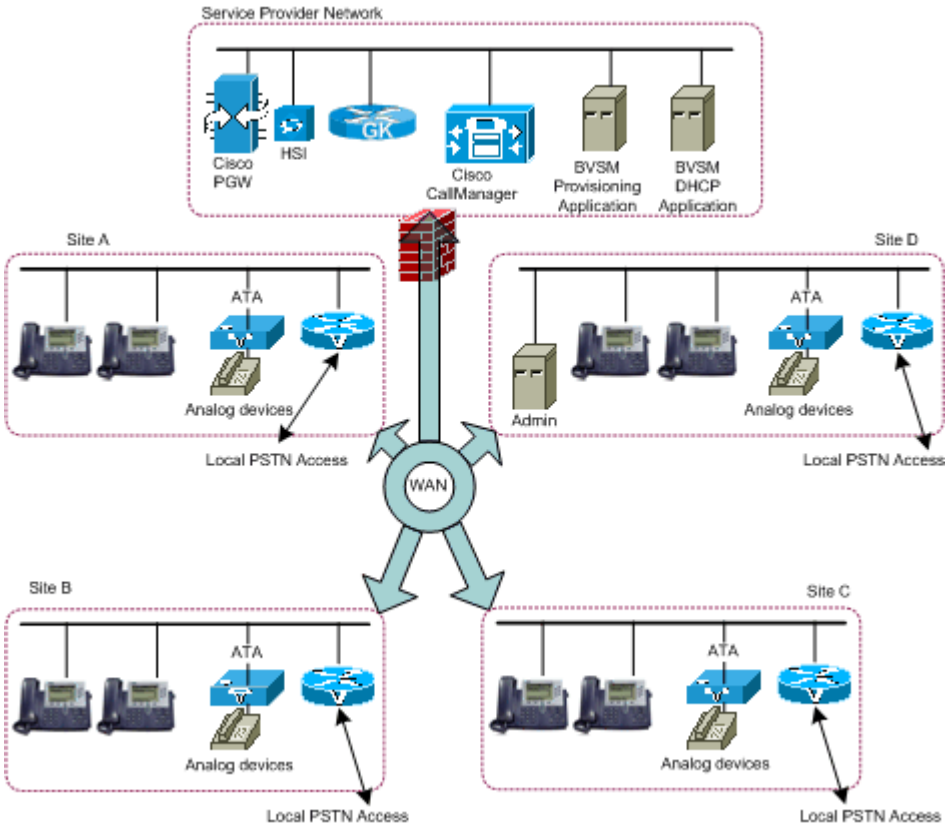

Service Provider Hosted, Multiple Enterprise Voice Services

The hosted multiple enterprise model enables business voice services to a number of small and medium enterprises, each with a single site or multiple sites distributed over a wide geographic area. Managed and hosted IP telephony service is provided by a Service Provider to multiple tenants over multiple sites, using one or more Unified Cisco CallManager clusters. This model represents the most complex deployment of the H/M-UCS architecture, and it is shown in Figure 13.

Service providers host all of the service related platforms in their domain in a secure environment behind firewalls, so that only relevant types of traffic can access the shared platform. The service provider could also host all of the PSTN gateways, which can be centralized or distributed, and are usually shared across

all customers. This service can be combined with the more traditional voice VPN service to support enterprises with TDM PBXs and also enterprises with hybrid environments in which some sites are served with hosted IP telephony and other sites remain on TDM PBX technology. TDM VPN service is provided by the Cisco PGW 2200, and is part of the Cisco Service Provider Business Voice Services portfolio.

In this scenario it is very important to consider the underlying network infrastructure over which the service will operate (such as MPLS or other IP network), along with the effect that this kind of service will have on the enterprise LAN environment.

Figure 13 Multiple Enterprise Deployment

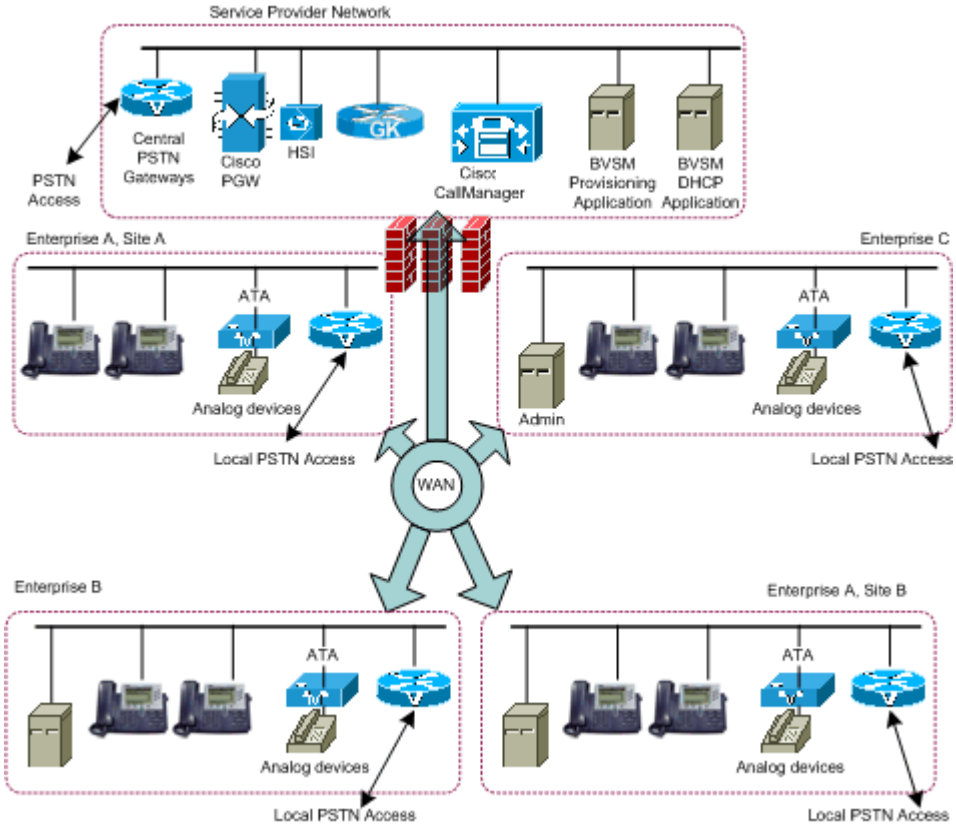

Cisco H/M-UCS Deployment Models

Following are the five Cisco H/M-UCS deployment models with their key characteristics.

Hosted

Hosted Multi-Tenant

Both PGW and Unified Cisco CallManager are located in the Service Provider network and shared by multiple Enterprises

Partitioned dial plan shared by PGW and Unified Cisco CallManager

Hosted Dedicated

PGW located in the Service Provider network and can be shared by multiple enterprises.

Unified Cisco CallManager dedicated to a single enterprise and can be located in either the Service Provider or the Enterprise network.

PGW partitioned dial plan

Unified Cisco CallManager dial plan can be partitioned to serve Enterprise divisions/departments.

Note

Although hosted dedicated deployment is possible with the current H/M-UCS dial plan, the dedicated Callmanager is still running a H/M-UCS multi tenant dial plan, with only one customer defined in the associated BVSM hardware group.

Figure 14 Hosted-UCS deployment models

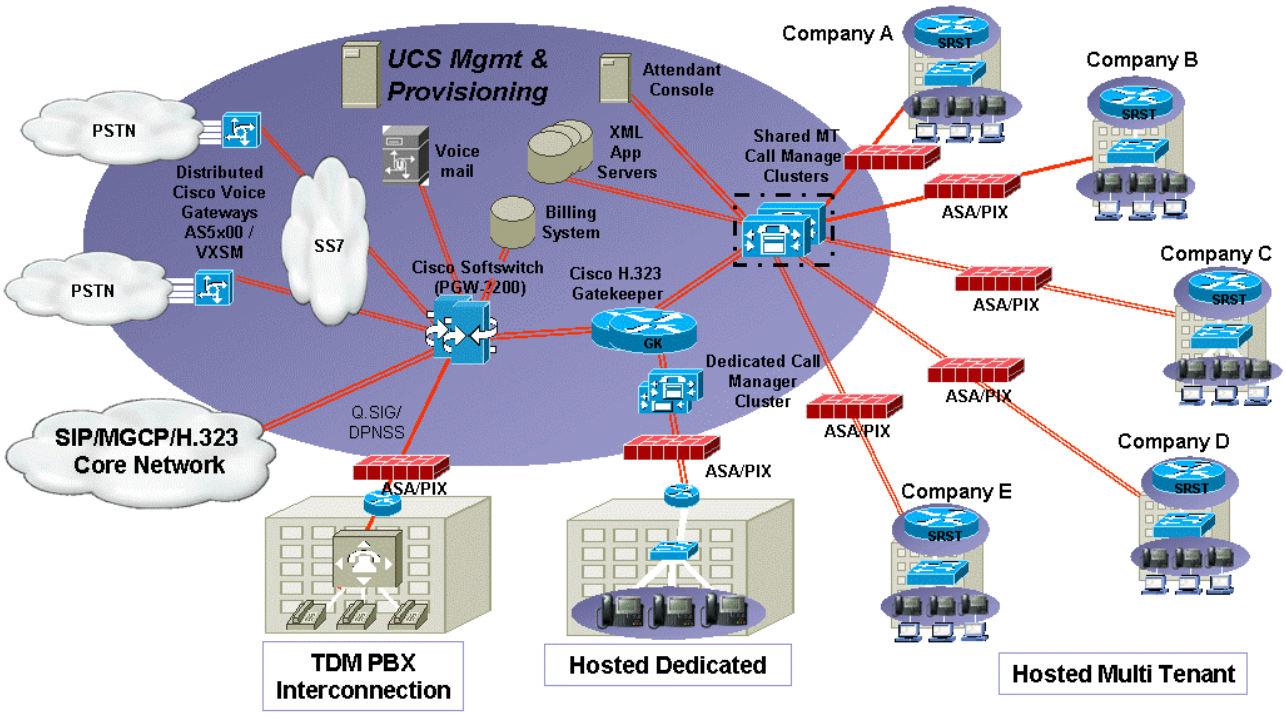

Managed

Note

Although three managed deployment models have been defined, only the Managed Dedicated (managed Multi Cluster) model is currently supported by the H/M-UCS 1.6 architecture. The others managed models will be delivered in a later release.

Managed Dedicated (also known as Managed Multi-Cluster)

Both PGW and Unified Cisco CallManager are located in the Enterprise network and dedicated to that Enterprise

Dial plan can be partitioned to serve Enterprise divisions/departments

Managed with PSTN Trunking

This deployment is the same as a typical enterprise based Unified Cisco CallManager deployment today except for the ability to manage and provision the Callmanager remotely from a provisioning application.

PGW is not present in this architecture. Unified Cisco CallManager is located in the Enterprise network and dedicated to that Enterprise.

Managed with On-net Trunking

SIP or H.323 trunk from an enterprise based Callmanager to the Service Provider's network to provide access to the PSTN.

Unified Cisco CallManager is located in the Enterprise network and dedicated to that Enterprise.

Figure 15 Managed-UCS deployment models.

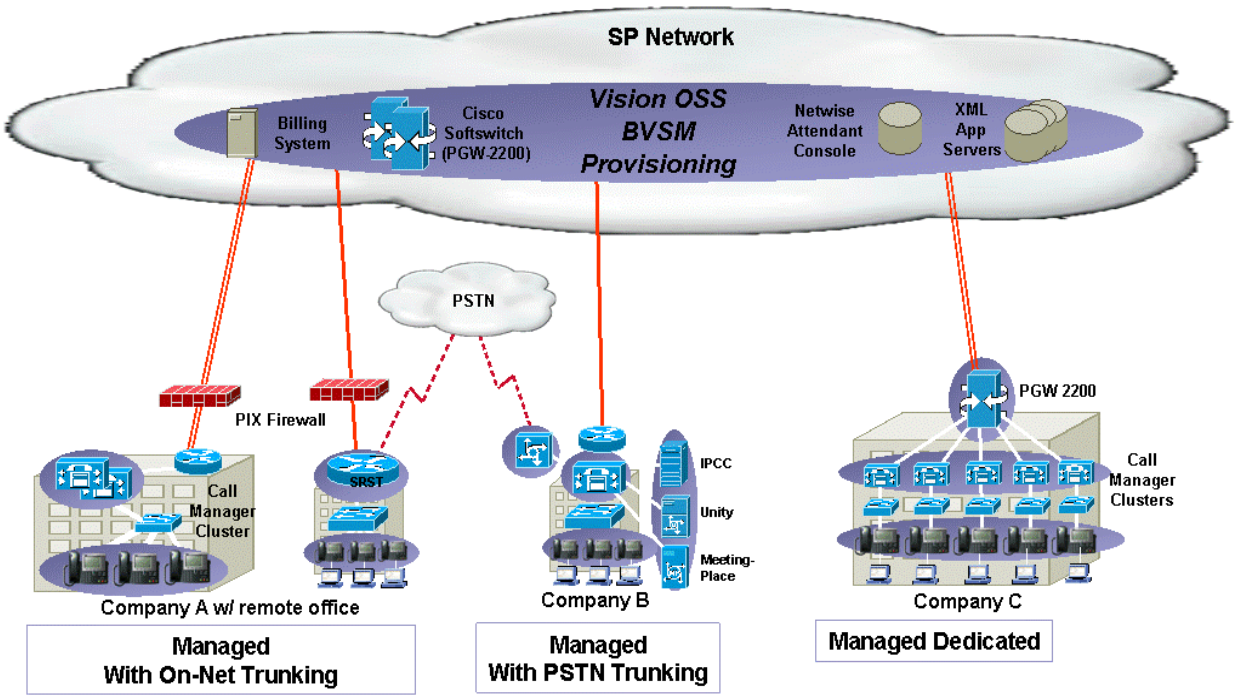

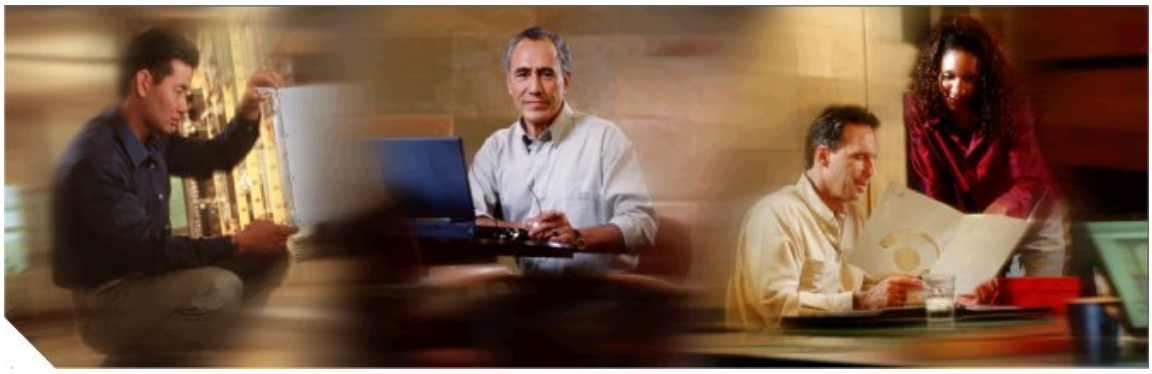

Network Infrastructure

This chapter provides some guidelines for the IP infrastructure that needs to be in place to support a Cisco H/M-UCS deployment. The service provider or systems integrator needs to ensure that the IP network infrastructure is designed appropriately so as to provide acceptable voice quality end to end between VoIP endpoints.

Voice over IP traffic, unlike most other IP applications, is not at all tolerant to IP packet loss or variations in the end to end delay (jitter) that IP networks often experience. The only way to provide acceptable voice quality over an IP network therefore is to place strict design requirements on IP packet loss, packet delay, and jitter. It is often necessary to use several of the Quality of Service (QoS) mechanisms (e.g. queuing, traffic shaping and packet fragmentation features) that are available on Cisco switches and routers throughout the network.

Telephony is often regarded as a "mission critical" service and for that reason redundant devices and network links that provide quick convergence after network failures or topology changes are also important to ensure a highly available infrastructure.

In Multi tenant H/M-UCS deployments some consideration needs to be made to accommodate the security and traffic isolation requirements between customers on the service.

The BVSM provisioning platform configures only the Cisco PGW, Cisco Unified CallManager, Cisco HSIs, Cisco Gatekeepers and Cisco Voice Gateways. It is not normally used to configure LAN or WAN switches and routers, or security related elements of the network infrastructure although it can be used to configure CPE routers and switches in some applications. The design and configuration of the IP network infrastructure components is therefore largely outside the scope of the H/M-UCS base design and can change considerably depending on particular customer or service provider application for the H/M-UCS architecture.

The chapter has sections on each of the following topics.

- Quality of Service in the campus
- Quality of service on the WAN
- Call Admission Control
- IP Addressing, Security and NAT
- WAN Resilience and SRST

Further information on IP network design as it relates to IP telephony can be found on CCO. Specific documents that are of interest are :

Enterprise QoS Solution Reference Network Design Guide

http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/ccmigration_09186a008049b062.pdf

The Network Infrastructure chapter in the Callmanager 4.X SRND

http://www.cisco.com/en/US/products/sw/voicesw/ps556/products_implementation_design_guide_chapter09186a00806e8c42.html

Quality of service in the Campus

Until fairly 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. The figure below illustrates the typical oversubscription that occurs in LAN infrastructures.

Figure 16 Data Traffic Oversubscription in the LAN

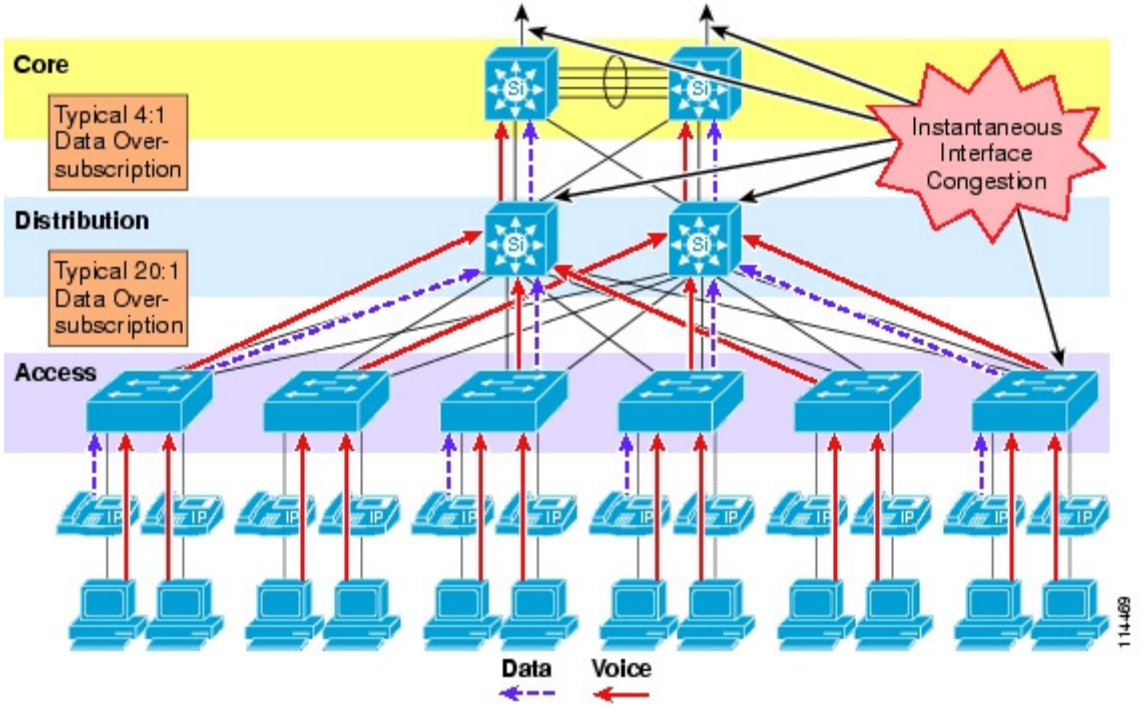

This oversubscription, coupled with individual traffic volumes and the cumulative effects of multiple independent traffic sources, can result in the egress interface buffers becoming full instantaneously, thus causing additional packets to drop when they attempt to enter the egress buffer. The fact that 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 re-established 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

- Queuing or scheduling
- Bandwidth provisioning
- IP Phone QoS Behaviour
- Separation of Voice and Data VLANS

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

Traffic Classification

Traffic 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).

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 signalling and voice RTP streams at the source, and it adheres to the values presented in Table 3-2. As such, the IP phone can and should classify traffic flows.

The table below lists the traffic classification requirements for the LAN infrastructure.

Table 4 Traffic Classification Guidelines for Various Types of Network Traffic

	Layer-3 Classification			**Layer-2 Classification**
Application	**IP Precedence (IPP)**	**Per-Hop Behaviour (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
Transactional data	2	AF21	18	2
Network management	2	CS2	16	2
Scavenger	1	CS1	8	1
Best effort	0	0	0	0

Note

[1] The recommended DSCP/PHB marking for call control signalling 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 signalling traffic as 26/AF31. Therefore, in the interim, Cisco recommends that both AF31 and CS3 be reserved for call signalling.

Table 5 Traffic Classification for Video Telephony

Traffic Type	Classification
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 signalling	Call signalling 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 signalling class is applicable to all voice signalling protocols (such as SCCP, MGCP, H323 and CTIQBE) as well as video signalling protocols (such as SCCP, H.323 and RAS).

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

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

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

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

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 signalling or media flows. Contrasting the bandwidth requirements of a single G.711 voice call (approximately 86 kbps) to the raw bandwidth of a Fast Ethernet 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.

QoS Behaviour of IP Phones

At the heart of a Cisco IP Phone is a 3 port 10/100 switch. One port, P0, is an internal port used for connecting the actual voice electronics in the phone. Port P1 is used to connect a daisy chained PC and Port P2 is used to uplink to the wiring closet Ethernet switch. Each port has 4 queues with a single threshold (4Q1T) configuration. One of these queues, Queue 0, is a high priority queue for all BPDU and CoS=5 traffic. These queues are all serviced in a round-robin fashion with a timer used on the high priority queue. If this timer expires while the queue scheduler is servicing the other queues, the scheduler will automatically move back to the high priority queue and empty it's buffer, ensuring voice quality. The queuing scheme on an IP Phone is shown in Figure 17 below. Because the IP Phone's high priority queue is accessible to any Layer 2 CoS=5 traffic, it's critical to make sure the PC connected to the IP Phone's access port is not classifying traffic, as well. The recommended method for doing this is to extend the Ethernet switch's trust boundary to the IP Phone and not beyond. On Catalyst 6000s, this "trust extension" is done using set port trust-ext command in CatOS 5.5 and above. On the Catalyst 3534-PWR, use the priority extend CoS 0 command. These commands instruct the IP Phone to mark all data traffic from the attached PC as CoS=0.

Figure 17 Traffic Remarking Within the IP Phone

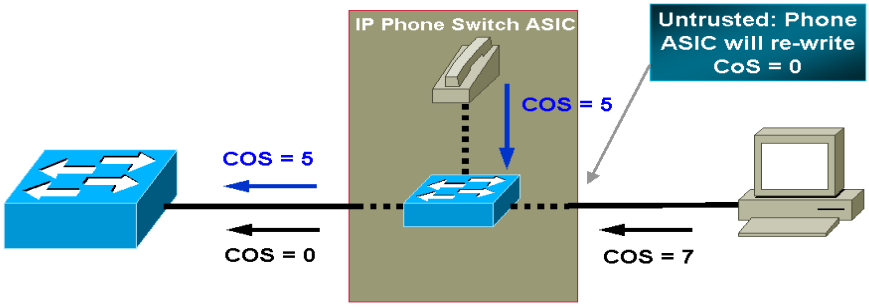

Separation of Voice and Data VLANS

When deploying voice on an IP network, Cisco recommends that two VLANs are used at the access layer: a native VLAN for data traffic (VLANs 10, 11, 30, 31, and 32 in Figure 18) 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 18).

Figure 18 Separation of Voice and Data VLANS

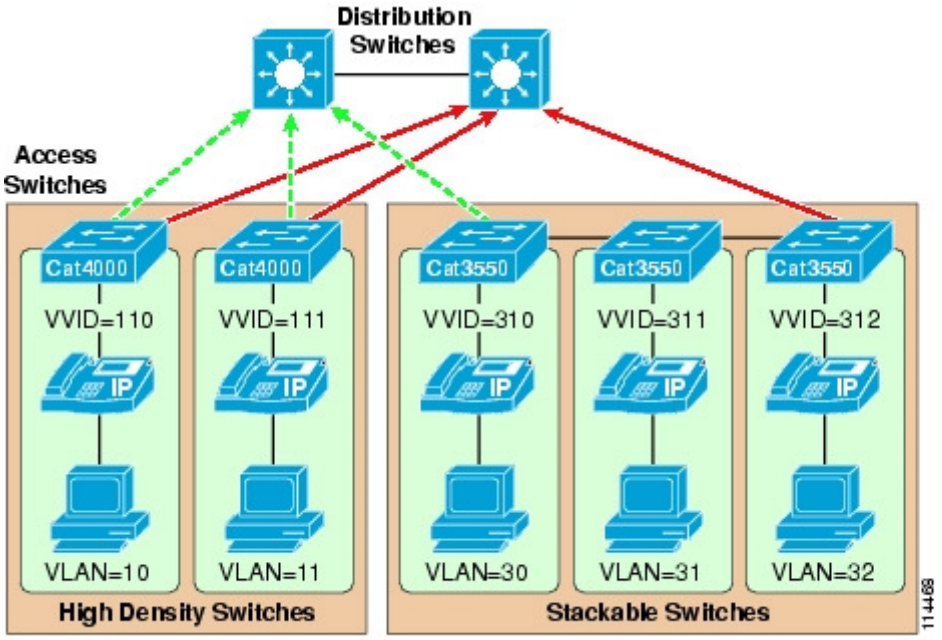

Separate voice and data VLANs are recommended for the following reasons:

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.

Quality of Service (QoS) in the WAN

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 6 lists the QoS features and tools required for the WAN infrastructure to achieve this goal.

Table 6 QoS Features and Tools Required to Support IPT by 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
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
- Link Efficiency Techniques
- Traffic Shaping

Traffic Prioritisation

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 behaviour 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.

The figure below shows an example prioritization scheme.

Figure 19 Optimized Queuing for VoIP over the WAN

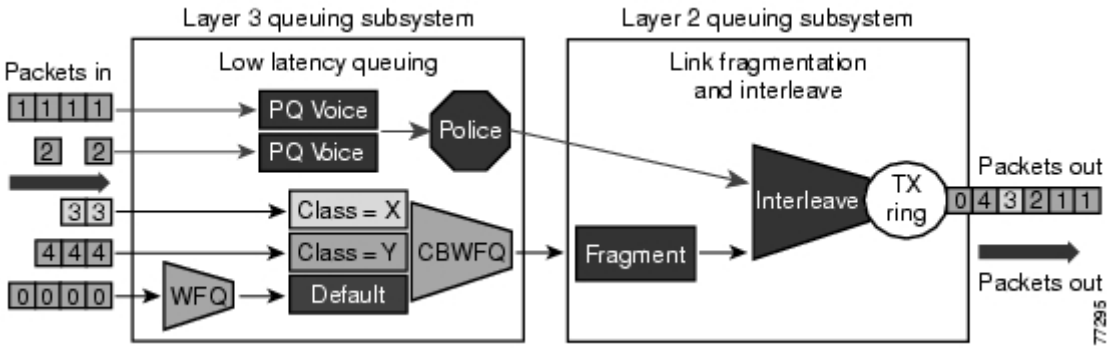

Cisco recommends the following prioritization criteria for LLQ:

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 behaviour (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).

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

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 queuing has occurred. Beginning in Cisco IOS Release 12.(2)2T and later, cRTP provides a feedback mechanism to the LLQ class-based queuing 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 7 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 7 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 7 Voice Class Bandwidth 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 – Note 1
12.2(2)T or later	240 kbps	100 kbps

1) In this case 140 kbps of extra 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.

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 the figure below. 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 20 Link Fragmentation and Interleaving (LFI)

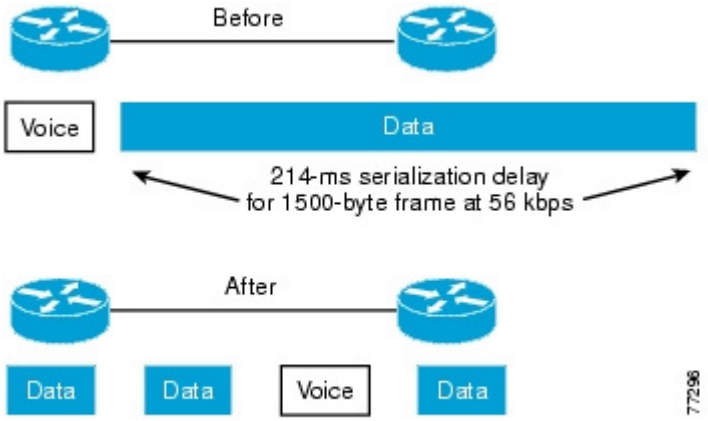

Voice-Adaptive Fragmentation (VAF)

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 signalling 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)). 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.

The figure below illustrates the main reasons why traffic shaping is needed when transporting voice and data on the same IP WAN.

Figure 21 Traffic Shaping with Frame Relay and ATM

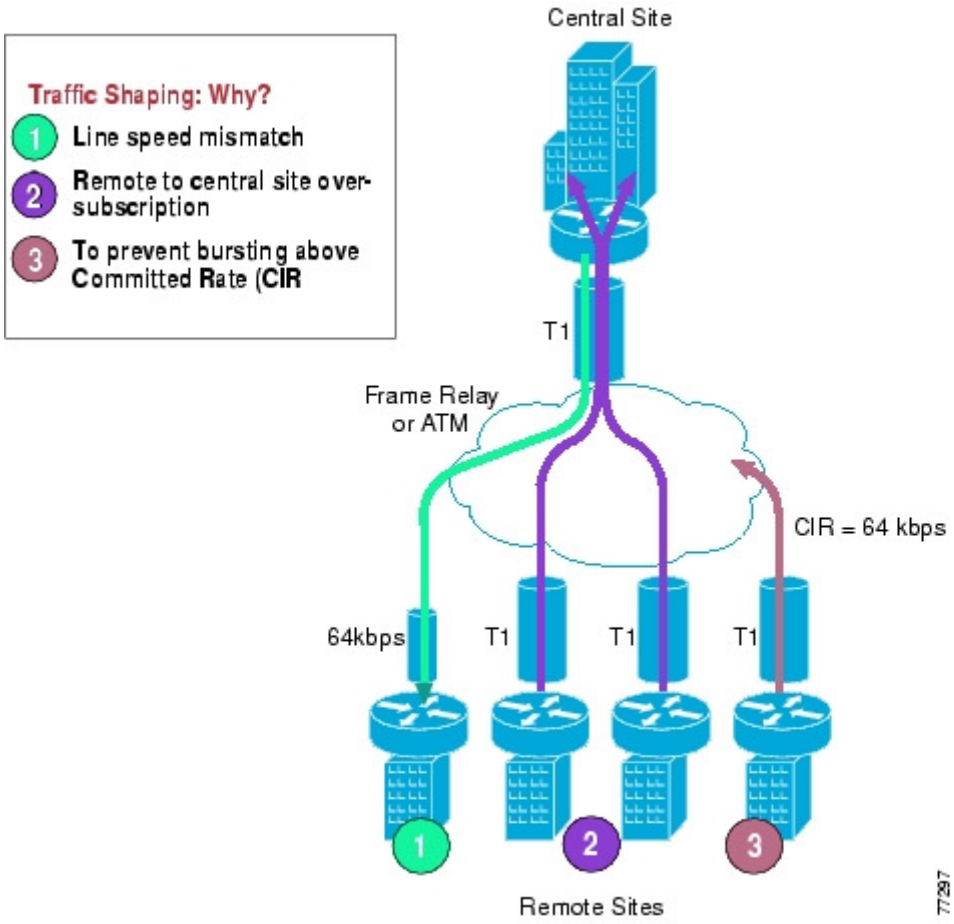

As shown above, there are three scenarios where traffic shaping needs to be used:

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.

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.

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. The figure below illustrates this mechanism with a generic example, where R is the rate with traffic shaping applied.

Figure 22 Traffic Shaping Mechanism

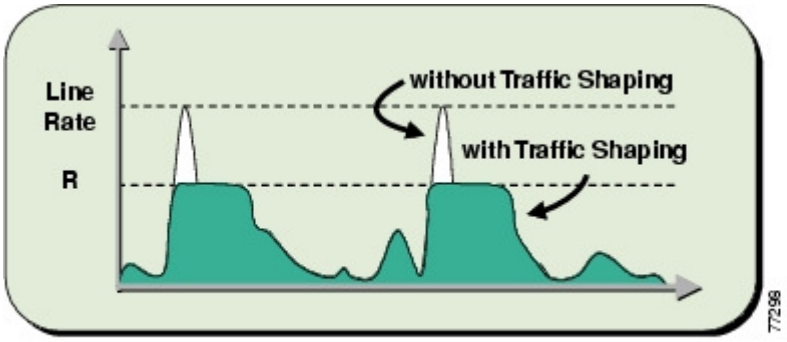

Voice-Adaptive Traffic Shaping (VATS)

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 signalling 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)), 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 behaviour 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/univercd/cc/td/doc/product/software/ios122/122newft/122t/122t15/ft_vats.htm

Call Admission Control

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 shown below.

Figure 23 Why Call Admission Control is Needed

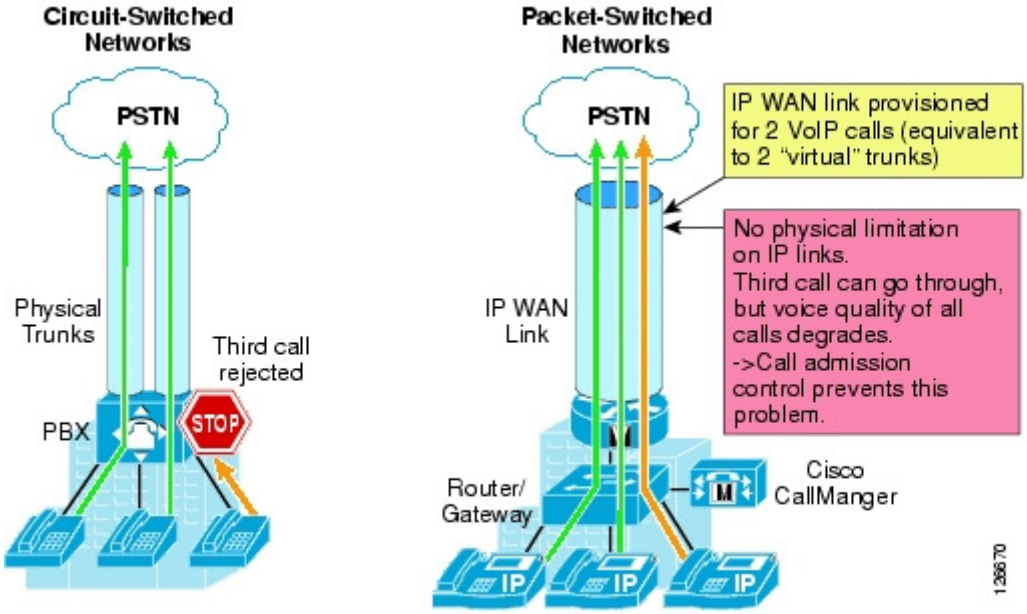

As shown on the left side of Figure 23, traditional TDM based PBXs operate within circuit-switched networks, where a circuit is established each time a call is set up. As a consequence, when a TDM PBX is connected to the PSTN or to another PBX, a certain number of physical trunks must be provisioned. 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 23. 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).

Call Admission Control 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 section explains these best practices in more detail.

The following recommendations apply to H/M-UCS deployments:

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 MPLS network itself is the "hub" of the network and effectively in location <None>.

For MPLS topologies with multiple CallManager clusters, use Cisco Unified CallManager static locations, with every site in a location. Leave intercluster trunks (if any) in the <None> location. Generally it is necessary to design such that all the devices in a given site (spoke) are controlled by one Callmanager cluster so allowing the locations admission control mechanism on that cluster can be aware of all the calls going in and out of the site. If this is not possible, for example at a gateway POP site where PGW controlled gateways are shared between several CallManager clusters, then it is necessary to statically allocate a certain percentage of the available bandwidth into that site to each cluster for control by the locations mechanism.

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 analyses the principles of topology unaware call admission control (an example of which is the CallManager locations mechanism used in H/M-UCS) and its limitations, then it presents the principles of topology aware call admission control (which we are developing in the form of RSVP for future versions of H/M-UCS).

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 (e.g. CallManager or PGW), aimed at limiting the number of simultaneous calls to or from a remote site connected via the IP WAN.

As shown in Figure 24, 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 or send it to voicemail.

Figure 24 Principles of Topology Unaware Call Admission Control

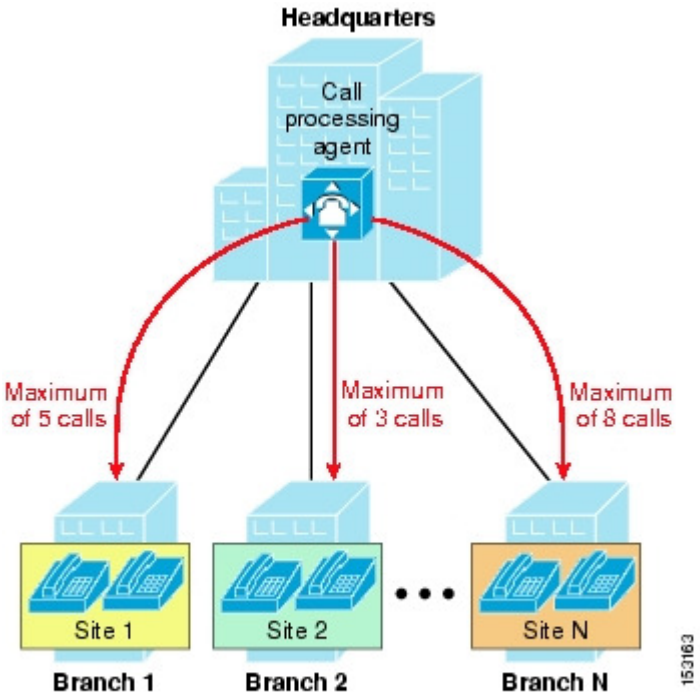

Because of their reliance on static configurations, topology unaware call admission control mechanisms can generally be deployed only in networks with a relatively simple IP WAN topology. In fact, most of these mechanisms mandate a simple hub and spoke topology or a simple MPLS based topology, as shown below.

Figure 25 Domain of Applicability of Topology-Unaware Call Admission Control

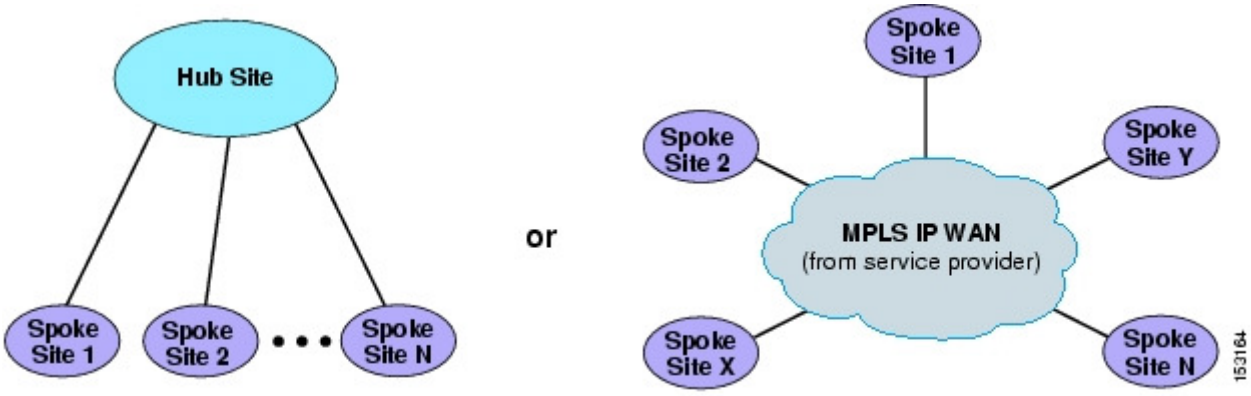

In a hub and spoke network or MPLS based network, each spoke site is assigned to a "site" within the call processing agent, and the number of calls or amount of bandwidth for that "site" is configured to match the bandwidth available for voice (and/or video) on the IP WAN link that connects the spoke 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. The figure below 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 26 Topology Characteristics of Typical Enterprise Networks

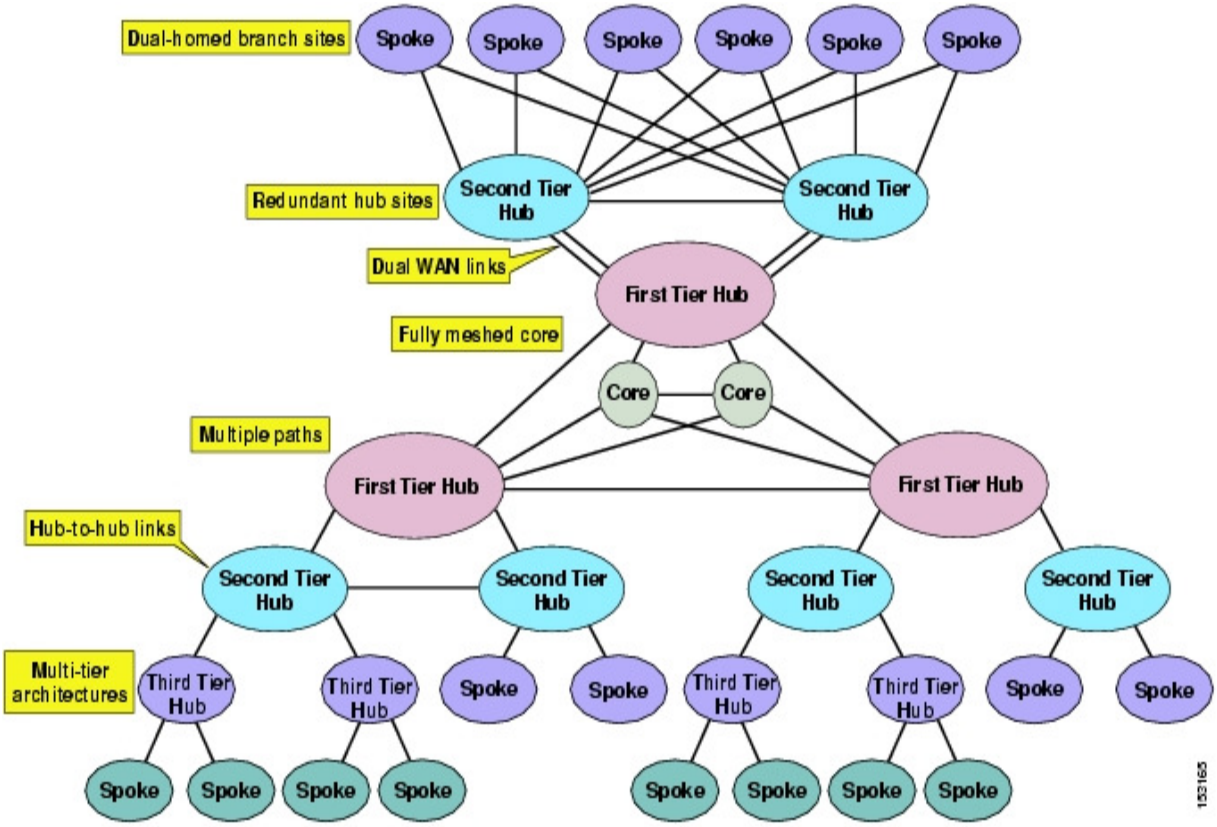

It is sometimes possible to adapt a topology unaware call admission control mechanism to a complex network topology, but there are limitations in terms of when this approach can be used and what behaviour can be achieved. For example, consider the simple case of a branch site connected to a hub site 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 below 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 27 Topology-Unaware Call Admission Control in Presence of Dual Links

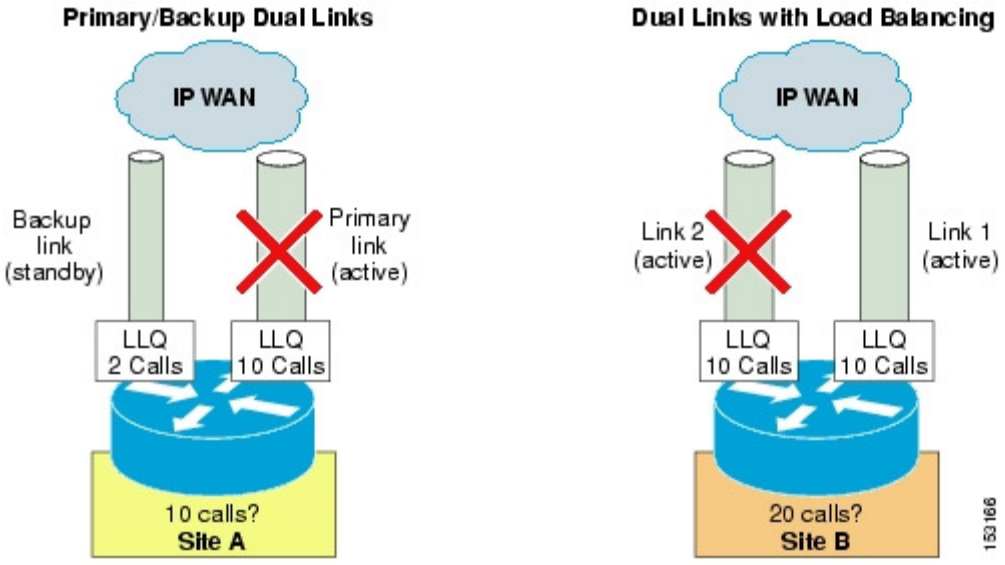

For the first example shown on left above, branch office A is normally connected to the IP WAN via a primary link, whose Low Latency Queuing (LLQ) bandwidth is provisioned to allow a maximum of 10 simultaneous calls. When this primary link fails, a smaller backup link becomes active and preserves the 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 suboptimal 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.

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. This type of admission control is not yet available in the H/M-UCS 1.6 architecture.

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 signalling protocol.

The Resource Reservation Protocol (RSVP) is the first significant industry standard signalling 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.

IP Addressing, Security and NAT

The deployment of a multi tenant hosted or managed VoIP service provides several challenges in the IP network design in order to provide sufficient security between customers and accommodate overlapping customer IP address space.

The key concept is splitting the complete platform into several logical security and IP address domains. There is one of these domains for each customer and one common one which contains the service providers call control and gateways and also, very importantly, provides routing between each customer domain for voice traffic, such that customers on the platform can call each other. The concept is illustrated in the diagram below.

Figure 28 Logical Security and IP Address Domains

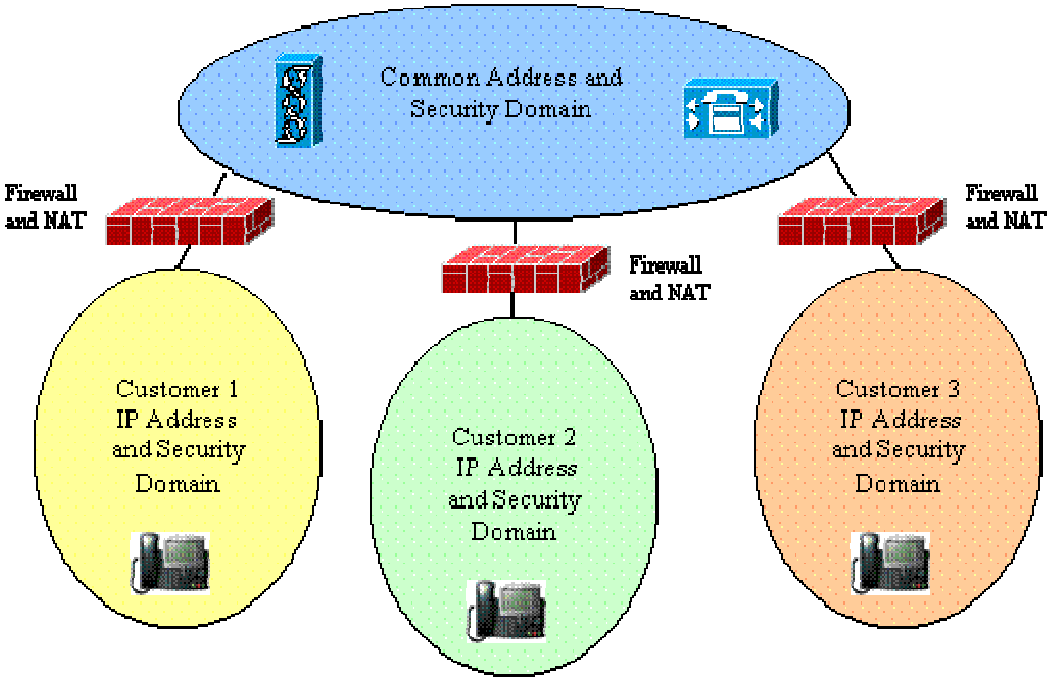

One key point that needs to be addressed when formulating a design is how far out from the core the common address and security domain should come. There is a trade off between having the common area just in the core, which has the negative that all voice traffic between customers has to route through the core and fully distributing the common address and security domain to the edge of the network (as a separate VPN) which has the negative that the amount of equipment needed to provide the firewall function can be quite large and complex to administer. The choice of positioning of the firewalls, and

hence how far out from the core the common address and security domain extends, is down to individual service provider or system integrator system requirements.

The key element in this design is the firewall between the common address and security domain and the customer individual domains. This provides both security and network address translation between the individual customer domains and the common service provider domain.

Security

Each firewall depicted in Figure 28 is configured to only allow a very limited set of traffic types from the customer domain into the common domain. This protects the infrastructure equipment hosted in the service provider common domain from attack from the customer domains. It also has the effect of protecting customers from each other as the only routing between customers is through the common domain.

Figure 29 NAT and Security at the PIX Firewall

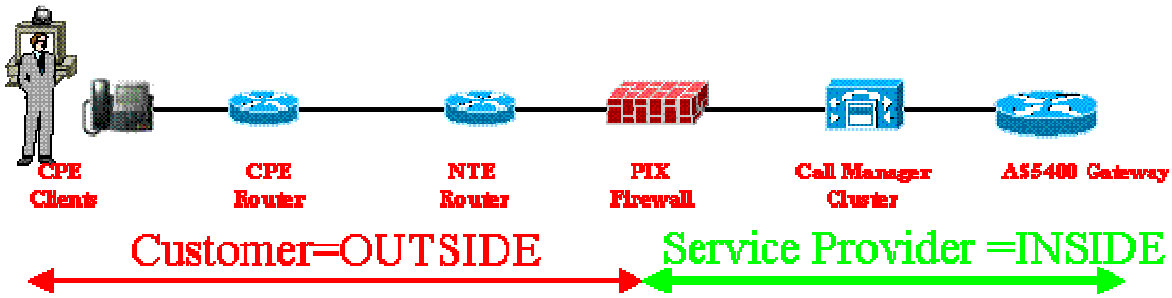

The PIX firewall is application aware as far as SCCP, SIP, H323 and MGCP are concerned so UDP ports for the media streams are opened dynamically based on the signalling protocol call setup. An example of the rules that can be implemented on the firewall to protect both the service provider from the customer and the customers from each other are show in the table below.

Table 8 Firewall Rules

Rule	Function
SCCP (TCP port 2000) to the Call Managers only.	This is to allow the Cisco Unified CallManagers to control the phones in the customer domain using the SCCP protocol
TAPI (CTIQBE) to the Call Managers running CTI manager only	Optional, used for third party phone call control or for call control to TAPI based soft phones or software applications.
HTTP (TCP port 80) to the Publisher Call Managers and BVSM only	Required for access to phone XML services hosted on the Cisco Unified CallManager and BVSM (e.g. directory) and also for customer self provisioning of BVSM using a web browser.
TFTP (UDP port 69) to the TFTP server only	Required to allow phones to download their configuration files and software updates.
H.323 (and H.245), to the Call Manager and maybe the HSI and gatekeeper if customer site applications that use H323 are required (e.g. a customer site located PSTN gateway using H.323).	Only required to support H.323 endpoints in the customer address space. Applications for this include H.323 video terminals and site located PSTN gateways.

Rule	Function
RTP Traffic UDP ports are opened dynamically by the ALG function within the firewall by MGCP, H.323, TAPI, SIP and SCCP Call Control	To allow voice to flow between customers and to PSTN gateways and conference bridges hosted in the common domain.
MGCP (UDP 2427/2428) to the PGW. Various backhaul protocols also need to be allowed to the PGW depending upon the L3 protocol at the gateway (e.g. Sigtrans)	To allow the PGW to control customer site located PSTN and PBX gateways.

NAT

Customers sharing a single H/M-UCS based service provider hosted service may have internal IP address space ranges that overlap. NAT can be configured on the firewalls to translate the private (non unique) addresses actually on the IP phones into addresses that are unique in the service provider (shared) domain. Not only must these addresses be unique, they must be reachable from the individual customer networks to allow voice calls to flow between customers through the common address and security domain. To achieve this routes must be injected to each customers domain that represent the address scheme used in the common address and security domain. It is therefore important that no customer subscribing to the service is actually using the address space chosen for the common service provider's IP address and security domain. The NAT pools that are configured on each customer firewall must be large enough to supply addresses to all the IP phones that are deployed in the customer concerned.

NAT normally prevents VoIP services from operating correctly due to embedded IP addresses at the application layer within virtually all VoIP signalling protocols. However, Cisco has implemented an Application Layer Gateway (ALG) function in the PIX, ASA and FWSM firewalls that intelligently manipulates the source and destination addresses of the IP packets and also can manipulate associated addresses embedded in the application layer. In PIX Version 7.X the ALG function is VoIP protocol aware for the SIP, H.323, MGCP, SCCP (SCCP) and CTIQBE (CallManager CTI) protocols. The ALG function or "fixup" must be enabled for the firewall to operate correctly with these VoIP protocols.

An example of a H/M-UCS implementation with NAT is shown in the figure below.

Figure 30 ALG Operation for a Call Within One Tenant

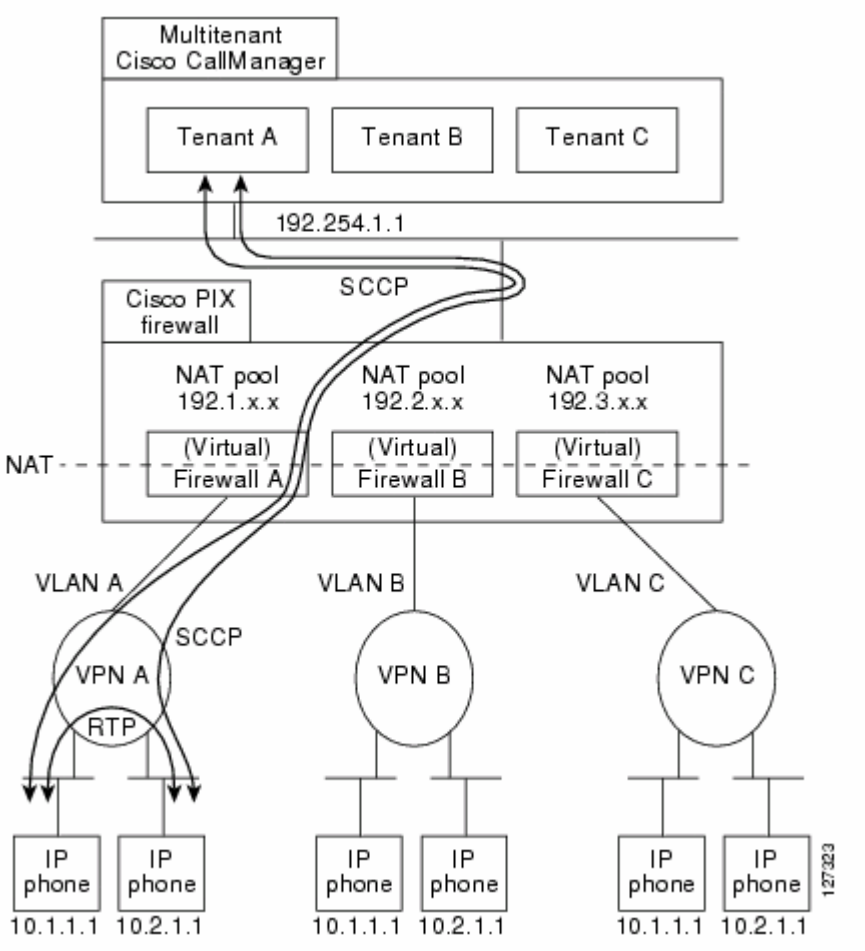

- Because both phones in this example are from the same tenant, they are on the same MPLS VPN. The tenant DHCP server has assigned 10.1.1.1 and 10.2.1.1 as the addresses for the two IP phones.
- The phones register with the SP's Cisco Unified CallManager at address 192.254.1.1. The phone registration flows through the firewall, and due to the NAT translations for the internal addresses, Phone 1 (internal address 10.1.1.1) is registered with the CallManager with external address 192.1.1.1 from the pool of address space that was allotted for this purpose and Phone 2 (internal address 10.2.1.1) is assigned the external address 192.1.1.2. The Cisco Unified CallManager sees IP phones registered with the addresses 192.1.1.1 and 192.1.1.2 and is not aware of the internal addresses that actually have been assigned to the phones.
- When a call is setup between the two phones the Callmanager instructs each phone to send its media to the external address (i.e. 192.1.X.X) of the other phone, but because these SCCP signalling messages are sent through the SCCP aware ALG function on the firewall, the media addresses within them get translated to the internal (actual) addresses of the IP phones (i.e. 10.1.X.X. As a result the RTP stream containing the voice flows directly between the two phones across the customers MPLS VPN.

In the second example, calls are sent between different customers in different IP address and security domains hence the media streams are routed via the SP common address and security domain as this is the only possible IP route between the two customers networks. Figure 31 below shows this scenario.

Figure 31 ALG Operation for a Call Between Tenants

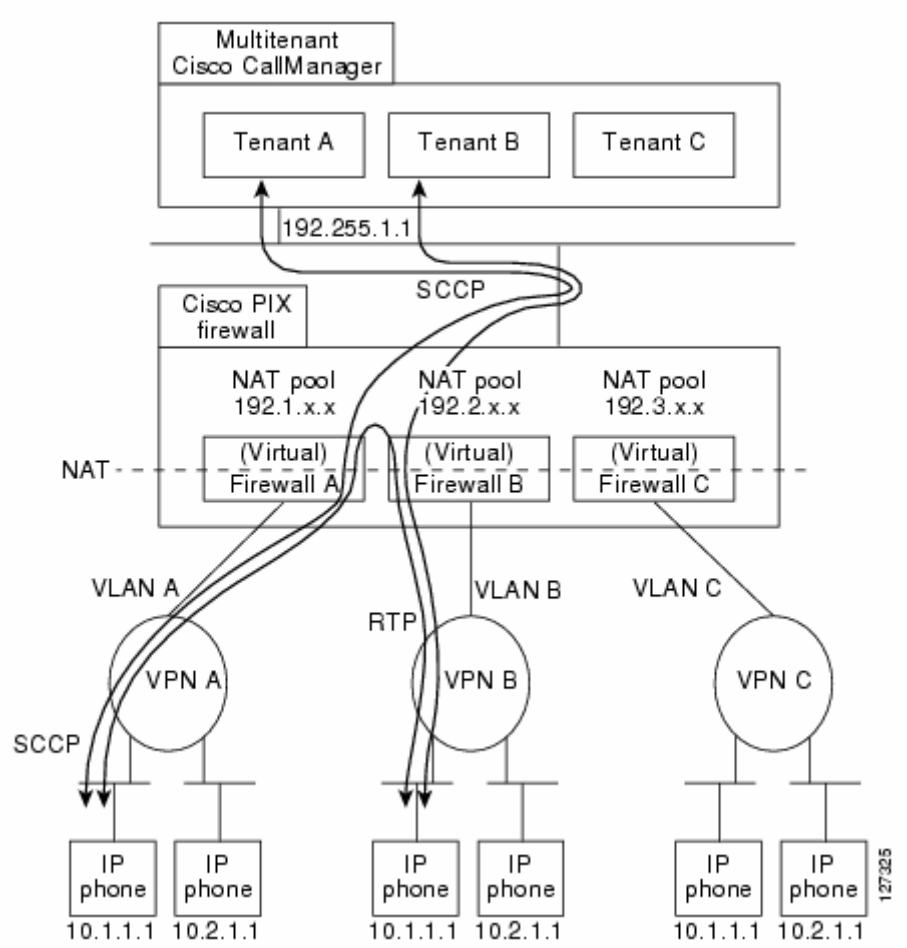

- The phone registration process is the same as explained in the previous example. All endpoints register with the Cisco Unified CallManager thorough the ALG function on the firewall, hence by the time the registration is seen by the Callmanager it is seen as from the external (192.1.X.X for customer A phones and 192.2.X.X for customer B phones. This means that Callmanager sees all phones as on unique addresses even although the real addresses of the phones are overlapping between tenants A and B.
- When a call is initiated to a number in another tenant, the destination media address delivered to the phones is the NAT translated address of the other phone. Consider the following examples:

Tenant A phone: Actual phone address = 10.1.1.1, NAT address =192.1.1.1 Phone in Tenant B receives 192.1.1.1 as the destination media address for the call.

Tenant B phone: Actual phone address = 10.1.1.1, NAT address =192.2.1.1 Phone in Tenant B receives 192.2.1.1 as the destination media address for the call.

- The destination IP addresses are not local to the customer so the media packets need to be routed towards the service provider. The service provider or tenant themselves needs to inject a route for all address space in the common domain (e.g. 192.0.0.0/8 in this example) to route packets to this address space towards the service providers network.
- For both sides of the media, the SP network routes to the destinations, which are addresses under its control (actually NAT pools on each of the firewalls). Both media streams arrive at the required firewall interface, and both destination addresses are then translated via the firewall NAT and ALG functions to the correct internal enterprise representation.

DHCP and IP Address Allocation

The following recommendations for address allocation issues should be considered when designing an IP telephony service that is based H/M-UCS:

- IP phones generally should obtain their IP addresses via DHCP. IP phones can be manually configured with network parameters, but for an installation of any meaningful size the configuration will need to be done automatically using DHCP servers.
- The DHCP server must return DHCP Option 150 to direct tenant IP phones to the Cisco Unified CallManager cluster TFTP server, from which they will receive their software image and configuration data.

Each enterprise has the option of controlling its own DHCP servers and specific address range or outsourcing the management of these tasks to its SP. The H/M-UCS solution is therefore able to support both a centralized and devolved DHCP server models.

It is possible to use the BVSM as the DHCP server for voice VLANS. To do this it is necessary to ensure that phones are on separate VLANS to data devices. DHCP requests are directed to the BVSM for those VLANS by configuring a "IP Helper Address" in the router that is acting as the default gateway for the VLAN. BVSM can manages DHCP servers which run either on the same servers as the BVSM itself or on other servers within the BVSM cluster. When using BVSM DHCP for IP address allocation several function are available from BVSM that are not possible when the enterprise (or another independent DHCP server within the SP) are used. These are :

- An "auto move" function for phones such that the location of a phone is automatically identified by the source subnet of the DHCP request and can then optionally be moved automatically by BVSM into that location to allow administration and further configuration by a location or customer administrators. This feature makes the initial setup of phones much simpler as there is no need to ensure a particular phone with a given MAC address ends up on a particular site.
- Rogue phone movement prevention between location is also prevented by the same mechanism. Once a role out of phones has been completed a phone that is allocated to a customer location will not be allowed to receive an IP address from the BVSM controlled DHCP server if it is plugged in another location (subnet). It is essential that CallManager locations admission control is configured to correctly restrict calls should the WAN bandwidth to a location be exhausted so preventing unauthorised phone movements between location is essential in maintaining correct admission control operation and hence voice quality.


Note

BVSM managed DHCP servers cannot currently be used in situations where the IP addresses allocated to sets of phone overlap. This is typically when NAT is being used on the firewall between the service provider and the customer. It is therefore necessary to avoid the use of BVSM managed DHCP servers in deployments that use NAT and have overlapping IP address space. Local or unmanaged central DHCP servers must be used for these deployments. When using unmanaged DHCP servers within a H/M-UCS environment it is not possible to use the "automove" BVSM feature or prevent rogue phone movements as described above. Other BVSM features may be dependant on IP source address such as directory and extension mobility. Please contact Vision OSS for details.

For deployments where the enterprise controls its IP space, the enterprise must be notified of the DHCP option parameters (specifically option 150) that are necessary to support voice service. A DHCP response to an IP phone must contain the responses below.

Mandatory responses :

- Option 003: Default gateway
- Option 150: TFTP server IP address

Optional responses:

- Option 006: DNS servers
- Option 015: Domain name
- Option 066: Boot server hostname (TFTP server)


Note

When using NAT with overlapping IP address space or port address translation (PAT) some features managed by BVSM need careful configuration to be useable. These are multi tenant extension mobility and directory. Both of these features currently use the source IP address to identify the phone. This limitation is due to be fixed in a BVSM release in 2007.

WAN Resilience

It is often desirable to design the IP WAN with resilient links, both in the core and often in the "last mile" connection to the customer site where the phones are located. This can be achieved in two ways.

- Providing IP WAN network resilience using techniques such as backup links or ISDN to ensure that the phones remain connected and registered at all times to the hosted CallManagers.
- By using local call processing that takes over in the event of a WAN failure and allows the phones to carry on making calls and provide a local connection to the PSTN. The Cisco technique for achieving this is called SRST (survivable remote site telephony).

Both of these techniques are discussed below.

IP WAN Network Resilience

The IP WAN to an end customer site may be made resilient by providing backup links. For large sites these may be the same WAN technology as the main links to the site (e.g. MPLS) but for small sites it may be necessary, for cost reasons, to use alternative WAN technologies for the backup links such as ISDN. In both cases the key consideration when deploying a backup link to a site is admission control. As discussed earlier in this chapter, the CallManager "locations" mechanism used currently in the H/M-UCS architecture for call admission control is topology unaware. What this means in practice is that the number of calls to a given location or site cannot change just because a WAN link in the network has failed and hence a backup link is in use. Until we have practical topology aware admission control mechanism incorporated into the H/M-UCS architecture, such as RSVP, the backup links to sites must be designed to carry the same number of **calls** to the site as the main link to the site. Note that does not necessarily mean that the backup link needs to be the same bandwidth as the main link as a useful technique is to deploy cRTP on a lower bandwidth backup link to achieve the same call capacity as the main link to a site.

A very cost effective backup WAN technology in many countries is ISDN. Using cRTP allows 4 G.729 voice calls to be carried reliably over a single ISDN 64Kbps channel or 8 calls over a BRI (2 B channels). This is often enough call capacity for a small branch within a banking or retail network. The main link to a branch to carry 8 G.729 calls without cRTP would have to be around 256Kbps. Using cRTP just over backup links and not in normal operation reduces the processing required at the head end routers in a hub and spoke network. It is also not possible to run cRTP over some WAN technologies such as public MPLS services.

SRST

Survivable Remote Site Telephony (SRST) is a software feature that enables Cisco IOS routers, which are usually used for the data access, to function as the call control for IP phones when primary call control is down or when WAN link is out of service.

SRST is a capability embedded in Cisco IOS that runs on the local branch office IP Telephony router. In the event of a WAN link failure, SRST allows a Cisco router to perform backup call processing for Cisco IP Phones. SRST automatically detects a failure in the network, and using Cisco's SNAP (Simple Network Automated Provisioning) capabilities, initiates a process to auto configure the router in "fallback" mode to provide call processing backup redundancy for the IP phones in that office. Calls in progress are sustained

for the duration of the call. However, calls in transition and calls that have not yet connected are dropped and must be reinitiated once Cisco IP phones re-establish connection to their local SRST router. SRST does not route the calls to the PGW. Any calls to remote destinations are sent directly to the PSTN through local ISDN interfaces. While in fallback mode to an SRST, Cisco IP phones periodically attempt to re-establish a connection with the primary Cisco Unified CallManager. When a connection is re-established with the primary Cisco Unified CallManager, Cisco IP phones automatically cancel their registration with the SRST router.

One major consideration when using SRST in a H/M-UCS environment is how the incoming calls are routed to a site. If the site normally uses a local gateway for receiving incoming calls from the PSTN then this gateway can continue to operate normally when the site WAN link fails and the phones are being controlled by the SRST router. If however, the incoming calls to a site are normally routed from the PSTN into a set of shared centrally located gateways then, when in a WAN failure situation, it is necessary to provide extra dial plan to route these calls back out to the PSTN and into the SRST local isdn gateway router. This routing configuration is not automated by BVSM and can be quite complex. Cisco therefore recommends that SRST is not generally used as the WAN resilience mechanism in situations where central gateways are being used to route incoming PSTN calls to a site.

The main roles of the SRST router are:

- Provide call control to endpoints in remote sites when primary and backup Cisco Unified CallManager are down or when the WAN link to Cisco Unified CallManager is down
- Provide a localized PSTN breakout trunk or trunks to allow users dialling out to PSTN and making emergency calls when operating in fallback mode

Within the H/M-UCS architecture, the following are provisioned automatically by BVSM when configuring SRST functionality :

- Enable the Cisco Unified CallManager “fallback” mode in the SRST router
- Provision the same time and date format as the one used in the primary Cisco Unified CallManager
- Enable the router to receive messages from the Cisco IP phones through specified IP addresses and ports
- Configure the maximum number of Cisco IP phones that can be supported by the router. The maximum number is platform dependent as shown below.

Table 9 SRST Phone Support by Platform

Platform	Number of Phones Supported
Cisco 1751-V, Cisco 1760, Cisco 1760-V, and Cisco 2801	Up to 24 phones
Cisco 2600XM, and Cisco 2811	Up to36 phones
Cisco 2650, Cisco 2650XM, and Cisco 2821 routers	Up to 48 phones
Cisco 2851 router	Up to 96 phones
Cisco 3725 router	Up to 144 phones
Cisco 3825 router	Up to 336 phones
Cisco 3745	Up to 480 phones
Cisco Catalyst® 6500 Communication Media Module (CMM)	Up to 480 phones
Cisco 3845 router	Up to 720 phones

- Create translation rules in the SRST router to allow extension dialling between IP phones, allow off-net PSTN calls and define the off-net prefix. (e.g. ‘9’)

- Create one or more dial-peers, depending on the voice traffic, to allow off-net prefixed calls to egress via these breakout interfaces (PRI, BRI, etc.)
- Replace the A number's CPID + RID + SLC with a valid E.164 DDI prefix
- Allow outbound emergency calls:
- Create one or more dial-peers, depending on the voice traffic, to allow emergency calls to egress via these breakout interfaces (PRI, BRI, etc.)
- Allow inbound calls from the PSTN
- Translate the B number's E.164 to CPID + RID + SLC + Extension number


Note

Some Callmanager configuration (e.g. provisioning SRST reference) is not automated by BVSM in the 1.6 H/M-UCS release due to limitations with the AXL provisioning interface on CCM. These items need to be manually configured.

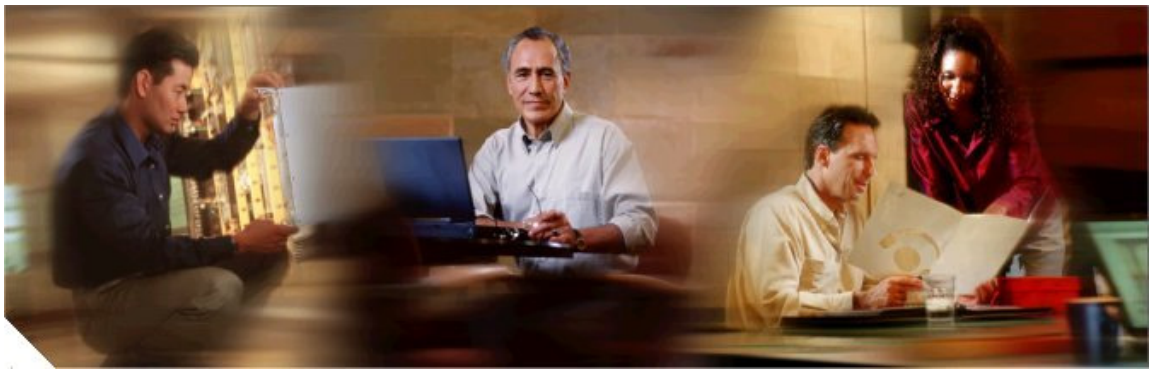

Service Description

This chapter describes the set of services that are provided by Cisco H/M-UCS. It includes the following:

- Bearer Services (including speech, video, fax, modem and DTMF)
- Subscriber Voice Services
- Subscriber Video Services
- Voice Mail Services
- Advanced XML Services
- Attendant Console Services

Bearer Services

Cisco H/M-UCS supports bearer services that include voice, video, analogue fax, modem or data using analogue telephone adaptors (ATAs).

The specifics of each of the bearer services are:

Speech

- G.711 - Bandwidth: 64 kbps
- G.729 - Bandwidth: 8 kbps
- Cisco Wideband Audio Codec - Bandwidth 256 kbps (although supported by the architecture this codec cannot be selected via the BVSM provisioning)


Note

BVSM version 3.1.6 provides limited control over region configuration. Only one codec may be selected platform wide for use within all regions and one codec may be selected for use platform wide between all regions.

If a single codec is selected platform wide (i.e. the same codec is selected both with and between regions) then BVSM does not provision separate regions per site, but a single region platform wide. This was a deliberate design decision to make a platform with many sites quicker to provision and more scalable.

Video

- H.263

– Bandwidth: 128 kbps to 1.5 Mbps

– Resolution: Common Intermediate Format (CIF) and Quarter Common Intermediate Format (QCIF)

– Frame Rate: up to 30 frames per second (fps)

- Cisco VT Camera Wideband Video Codec

– Bandwidth: 7 Mbps

– Resolution: 320x240

– Frame Rate: up to 30 fps


Note

BVSM version 3.1.6 provides no control over video region configuration. The video region is left at the default Callmanager setting (normally 384Kbps) unless manually adjusted.

Fax

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 reconverting the digital data to analog for the receiving fax machine.

Most Cisco voice gateways currently support three methods to transmit fax traffic across the IP network:

- Cisco Fax Relay - In fax relay mode, the gateways terminate the T.30 fax signaling. This protocol is not supported by the AS5400 gateway.
- Fax Pass Through - In fax pass through mode, the gateways pass the fax call as normal through the speech channel. The gateways have to change the voice codec to G.711 and disable echo cancellation to allow fax modulation to pass successfully through the speech channel. The Cisco ATA supports fax pass through but does not support either of the two fax relay methods.
- T.38 Fax Relay – T.38 is the standardized fax relay method. If differs from Cisco FAX relay in that it is negotiated through end to end through the signaling rather than in the RTP stream. Callmanager only supports T.38 negotiation when using H.323 to control a gateway and H.323 analogue gateways are only supported in the large enterprise single tenant H/M-UCS deployment model. There is therefore not currently a T.38 solution currently for the multi tenant H/M-UCS deployment model.

Modem

In general, there are two mechanisms for supporting modem sessions over an IP network using voice gateways - modem pass-through and modem relay

Currently, modem pass-through is the only mechanism supported on Cisco voice gateways. Modem pass-through is the transport of modem signals through a packet network using the 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.

Dual Tone Multi Frequency (DTMF) Relay

DTMF relay capability, specifically out-of-band DTMF, separates DTMF digits from the voice stream and sends them as signalling indications through the gateway protocol (H.323, SCCP, or MGCP) signalling 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.

Unsupported Bearer Service

The following bearer services are specifically not currently supported by the H/M-UCS solution.:

- ISDN data
- Group 4 fax..

Subscriber Voice Services

Cisco H/M-UCS provides support for a large subset of Cisco Unified CallManager voice services. The matrix below lists all the services with any caveats.

Table 10 H/M-UCS and CCM Feature Comparison Matrix - Key

Symbol	Meaning
√M	Available but Manual Configuration Required
N/A	Not Applicable
√B	A similar feature is supported by BVSM directly
√	Available
N	Not Available
√*	Available but with caveats, see notes

Table 11 H/M-UCS and CCM Feature Comparison Matrix

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
CallManager 4.1.3	N/A	√	
CallManager 4.2.3	N/A	√	
Abbreviated Dialing	√	√	

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
Ad-hoc Conferencing Enhancements ▼			
1. Allows any participant in an ad-hoc conference to add a party into or remove any participant from the conference	√	√	Supported in H/M-UCS1.6 without need for manual CCM config (supported via soft key). The default soft key template includes this soft key
2. Conference list and drop any party	√	√	
3. Drop last conference party	√	√M	Service parameter and soft key template needs to be configured manually allows this to be supported in H/M-UCS1.6
4. Conference Chaining	√	√	
Annunciator	√	√	
Answer/Release	√	√	
Application Programmable Interfaces ▼			
1. Call Detail Records (CDR) and Call Management Records (CMR) API	√	√	
2. Computer Telephony Integration (CTI) API	√	√	
3. AXL SOAP Database API	√	√	
4. LDAP Directory Integration API	√	√	
5. IP Phone XML Services API	√	√	
CallManager Attendant Console	√	√	Supported in non multitenant environments. This product is very difficult to use through firewalls without leaving the firewall open on many ports.
Audible and Visual Indication of Ringing Line ▼	√	√	
1. Disable Audible/Visual Indication of Ringing Line	√	N	
Distinctive Ring (Internal vs. External Call) ▼			
1. External vs. Internal Trunk Designation	√	N	The H/M-UCS architecture does not currently support this feature.
2. Distinctive Ring Per Line	√	√	
User Configurable Ring Settings via User Options Web Page	√	N	The CallManager Web Interface is not supported in the H/M-UCS architecture In H/M-UCS, the Ringing tone can be selected via the phone only.
Audio and Video Codec Support ▼			Currently only the G.711 and G.729 audio codecs are supported by the H/M-UCS architecture.
3. G.711a	√	√	

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
4. G.711u	√	√	
5. G.722	√	N	
6. G.723	√	N	
7. G.723.1	√	N	
8. G.728	√	N	
9. G.729a	√	√	
10. G.729b	√	√	
11. G.729ab	√	√M	Silence suppression can be turned on cluster wide via manual configuration of a service parameter on CCM.
12. GSM-EFR	√	N	
13. GSM-FR	√	N	
14. Cisco Wideband Audio Codec (16-bit/16Khz)	√	N	Supported only on 7910, 7940, 7960, 7970, 7971
15. H.261	√	√	Supported only on H.323 and SCCP devices
16. H.263	√	√	Supported only on H.323 and SCCP devices
17. H.264	√	√	Supported only on SCCP devices
18. Cisco Wideband Video Codec (7Mbps)	√	√	Supported only on Cisco VT Advantage endpoints
Authentication/Encryption ▼			
CallManager Administration Web Page Authentication	√	√B	BVSM equivalent. Access to CCM is via BVSM in the H/M-UCS architecture except for initial setup.
Multi-Level Administration	√	√B	BVSM equivalent although MLA is also still available for restricting direct access to CCM (e.g. for troubleshooting in H/M-UCS)
CallManager User Options Web Page Authentication	√	√B	BVSM equivalent
CallManager Web Page Encryption (Https)	√	√B	BVSM equivalent
LDAP over SSL	√	√	The Callmanager LDAP database is not used extensively for user data in the H/M-UCS environment. Its use is mainly for CTI interface logon. BVSM provides the user directory in a H/M-UCS environment.
Extension Mobility Username/ PIN Authentication	√	√B	BVSM takes over authentication of extension mobility logon against its own database in order to provide multitenant secure operation.
SQL Database Access Authentication ▼			

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
1. Windows Integrated Authentication for SQL Database Access	√	√	The H/M-UCS architecture has no effect on this feature, it is still used within the CCM cluster.
2. TFTP Directory Access Restrictions	√	√	The H/M-UCS architecture has no effect on this feature, it is still used within the CCM cluster.
Phone File Authentication ▼			
1. Signed Firmware Loads	√	√	The H/M-UCS architecture has no effect on this feature, it is still used within the CCM cluster.
2. Signed Configuration Files	√	√	The H/M-UCS architecture has no effect on this feature, it is still used within the CCM cluster.
Device Authentication/Encryption ▼			Device authentication and encryption is not tested as part of the H/M-UCS architecture. Configuration of these features is not available through BVSM. It is not possible to use signaling authentication or encryption through firewalls hence these features can generally not be used in a hosted environment.
1. Signaling Authentication	√	N	Available only on 7940, 7960, 7970 and 7971 IP Phone models
2. Signaling Encryption	√	N	Available only on 7970 and 7971 in 4.0, support added for 7940 and 7960 in 4.1(2)
3. Media Encryption	√	N	Available only on 7970 and 7971 in 4.0, support added for 7940 and 7960 in 4.1(2)
4. Visual Indication of Device Authentication/ Encryption	√	N	Available only on 7940, 7960, 7970 and 7971 IP Phone models
Auto-Answer/Intercom ▼			
1. Auto-Answer to Headset (with Zip Tone)	√	√	
2. Auto-Answer to Speakerphone (hands-free Intercom)	√	√	Supported only on 7940, 7960, 7970 and 7971 IP Phone models
Automated Change Notification/ Database Replication	√	√	The H/M-UCS architecture has no effect on this feature, it is still used within the CCM cluster.
Automated Installation and Recovery ▼			
1. Cisco IP Telephony Backup and Restore System (BARS)	√	√	The H/M-UCS architecture has no effect on this feature, it is still used within the CCM cluster.
2. Automated System-Wide Software and Feature Upgrades	√	√	The H/M-UCS architecture has no effect on this feature, it is still used within the CCM cluster.
Automatic Alternate Routing (AAR) ▼			

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
1. AAR for on-net (site-to-site) calls	√	N	AAR supported to enable forwarding calls on no bandwidth only
2. AAR support for forwarding calls on no bandwidth	√	√	AAR supported to enable forwarding calls on no bandwidth only
PSTN failover on route being unavailable	√	√	The H/M-UCS architecture supports gateway failover between central gateways or between local gateways on the same customer site only.
Automatic Attenuation/Gain Adjustment	√	√	
Automatic Bandwidth Selection	√	√	
Automatic Number Identification (ANI)	√	√	
Auto-Registration	√	√B	BVSM has an equivalent feature called Automove that can be used in the H/M-UCS architecture to automate phone provisioning.
Barge	√	√*	The IP network a in hosted environment may cause it to be available intra-site only due to the mixing of streams being performed on the phone itself. This feature also has the restriction that only the G.711 codec is supported due to the mixer on the phone not being able to support G.729. The feature is superseded by Conference Barge which is supported in H/M-UCS and overcomes both these limitations.
Call Admission Control (CAC) ▼			
1. Call Admission Control (CAC) - intercluster and intracluster	√	√*	The H/M-UCS solution provides CAC via the CallManager Region/Location configuration which works intra cluster only
2. Multisite (cross-WAN) capability with intersite CAC	√	√	
Call Back ▼			
1. Call Back over Q.SIG trunks (i.e. Call Completion)	√	√*	The H/M-UCS architecture supports QSIG trunks from the PGW to TDM PBX. Call completion is supported on these trunks between TDM PBX only. Future versions of H/M-UCS with PGW 9.7.3 should support interoperation of some QSIG features to Callmanager Phones via H.323 annex M1.
2. Call Back over DPNSS trunks	N	√*	The H/M-UCS architecture supports DPNSS trunks from the PGW to TDM PBX. Call back is supported on these trunks between TDM PBX only. Future versions of H/M-UCS with PGW 9.7.3 should support interoperation of some DPNSS features to Callmanager Phones via H.323 annex M1 with the PGW providing DPNSS to Q.SIG (via annex M1) interworking.

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
3. Call Back Station to Station	√	√*	Supported fully when using the large enterprise (Single Tenant) H/M/UCS dial plan. When using the Hosted Multi Tenant H/M-UCS dial plan, callback is supported for intra-site calls only (so is not especially useful) due to intersite calls traversing the PGW. This should be available in H/M-UCS 1.7 where PGW 9.7(3) will allow full support by running H.323 annex M1 between the CCM and the PGW.
Call Coverage ▼			
1. Using Call Forward Busy (CFA), Call Forward No Answer (CFNA)	√	√	
2. Using TCD Hunt Groups	√	√*	Cisco attendant console use only
3. Using Native Hunt Groups	√	√	
4. Per-User Enhanced Call Coverage Paths	√	√*	Not all Cisco Callmanager Hunt Group features can be configured through BVSM
5. Forwarding based on internal/external calls	√	√*	Only available when using the large enterprise (single tenant) H/M-UCS dial plan due to the fact that all internal calls in this dial plan stay on the Callmanager cluster or traverse clusters via CCM intercluster trunks which can be marked as internal.
6. Forwarding out of a coverage path	√	N	
7. Timer for maximum time in coverage path	√	N	
8. Time of day	√	N	
Call Forwarding ▼			
1. Call Forward All (CFA)	√	√	
2. Per-Line Configurable Call Forward Busy Trigger	√	√	
3. Call Forward No Answer (CFNA)	√	√	
4. Per-Line Configurable Ring No Answer Timeout	√	√	
5. Call Forward Number Expansion to Voicemail	√	√	
6. Call Forward Reason Codes to Voicemail	√	√	Supported in H/M-UCS but voicemail dependent (e.g. Cisco Unity does not use this)
7. Call Forward on Failure (CFOF)	√	√M	This feature is used on Callmanager with CTI route points and ports only to allow call handling when CTI applications fail. It may be manually configured as part of a CTI application integration in the same way as it would way using Callmanager directly.

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
8. Call Forward on Unregistered	√	√	
Call Hold and Retrieve	√	√	
Call Join	√	√	
Call Park	√	√	
Call Pickup (Pickup) and Group Call Pickup (GPickUp)	√	√	
Call Pickup Enhancements ▼			
Call Pickup Notification	√	√M	Configuration of this feature is not available through BVSM
One touch call pickup	√	√M	
One touch group pickup	√	√M	
Call Preservation	√	√	
Call preservation-redundancy and automated failover-on call-processing failure	√	√	
Call Status per Line	√	√	Supported only on 7905, 7910, 7912, 7920, 7935, 7936, 7940, 7960, 7970 and 7971 IP Phone models
Call Waiting/Call Retrieve per Line	√	√	
Consecutive Call Waiting/Alerting per Line	√	√	
Calling Line Identification (CLID)	√	√	
Call-by-Call Calling Line ID Restriction (CLIR)	√	√	
Per Phone Call Display Restrictions	√	√M	
Calling Name Identification (CNID)	√	√	
CNID over Q.931 Facility Information Element	√	√	CNID is supported from the H/M-UCS architecture on all trunk types (Q.SIG, DPNSS, SS7 and Q.931) via the PGW. The H/M-UCS solution does have a problem with connected name update on CCM originated calls when call forwarded from DPNSS PBX as connected name is not updated through the HSI due to CSCsh53557.
Centralized System Administration, Monitoring and Reporting ▼			The H/M-UCS solution includes BVSM for centralized Provisioning and Service Management.
1. Cisco CallManager Serviceability Tools	√	√	
2. Bulk Administration Tool (BAT)	√	√	The H/M-UCS architecture includes BVSM to provide the provisioning functionality, including bulk loading functionality

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
3. Bulk Export Utility	√	N	Callmanager based bulk export is still available to use but there is no bulk export capability on BVSM.
4. Syslog Support for Debugging Output	√	√	
Device Pools	√	√	BVSM manages all CallManager configuration so direct management of specific entities (e.g. Device Pools) is usually unnecessary.
Device Search Filtering Criteria	√	√B	BVSM has equivalent functionality
External Route Plan Wizard	√	√	BVSM manages all CallManager configuration so direct management of specific entities is usually unnecessary.
Integrated HTTPD Server on IP Phones	√	√	
Performance Monitoring and Alarms	√	√	
Quality Reporting Tool (QRT)	√	√	
Call Detail Records (CDR) and Call Management Records (CMR) ▼	√	√	
1. Single CDR Repository per CallManager Cluster	√	√	
2. CDR Analysis and Reporting Tool (CAR)	√	√	ART was enhanced and renamed to CAR in release 3.2
System Event Reporting	√	√	
Zero Cost Automated Phone Moves	√	√	The BVSM provisioning system has a variety of phone registration and move capabilities
Zero Cost Automated Phone Adds	√	√	BVSM allows a variety of phone registration capabilities
Cisco ATA-186 2-Port Analog Gateway Support	√	√	
Cisco Attendant Console/WebAttendant ▼	√	√	WebAttendant renamed to Cisco Attendant Console. Note Cisco Attendant console can only be used in single tenant deployments. The protocols running between the attendant console and the TDC on Callmanager does not lend itself to being firewalled.
1. Longest Idle Support	√	√	
2. Pop-to-Top on New Call	√	√	The H/M-UCS architecture does not effect the availability of this feature
3. Accessibility Enhancements	√	√	The H/M-UCS architecture does not effect the availability of this feature
Cisco Conference Connection (CCC) Support	√	√	The H/M-UCS architecture does not effect the availability of this feature

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
Cisco Discovery Protocol (CDP) Support	√	√	The H/M-UCS architecture does not effect the availability of this feature
Cisco Emergency Responder (CER) Support	√	N	Cisco Emergency Responder is not supported in the H/M-UCS 1.6 architecture.
Cisco IP Automated Attendant (IP AA) Support	√	N	The H/M-UCS architecture does not currently support this product (it was not tested as part of the 1.6 release).
Cisco IP Integrated Contact Distribution (IP ICD) Support	√	N	The H/M-UCS architecture does not currently support this product.
Cisco IP Phone 79xx Models ▼			
1. Cisco IP Phone 7902	√	√	SCCP
2. Cisco IP Phone 7905	√	√	SCCP
3. Cisco IP Phone 7906	√	N	Not tested to work with H/M-UCS yet although phone types can be added through BVSM GUI so 7906 could be added without BVSM upgrade. Only supported on CCM 4.2(3)via a dev pack install.
4. Cisco IP Phone 7905G	√	√	SCCP
5. Cisco IP Phone 7910 and 7910+SW	√	√	SCCP
6. Cisco IP Phone 7911	√	√	SCCP
7. Cisco IP Phone 7912	√	√	SCCP
8. Cisco IP Phone 7912G	√	√	SCCP
9. Cisco IP Phone 7914 Expansion Module	√	√	
10. Cisco IP Phone 7920	√	√	SCCP
11. Cisco IP Phone 7921	√	N	Not tested to work with H/M-UCS yet although phone types can be added through BVSM GUI so 7921 could be added without BVSM upgrade. Only supported on CCM 4.2(3)via a dev pack install.
12. Cisco IP Phone 7940	√	√	SCCP
13. Cisco IP Phone 7940G	√	√	SCCP
14. Cisco IP Phone 7941G	√	√	SCCP
15. Cisco IP Phone 7941G-GE	√	√	SCCP
16. Cisco IP Phone 7960	√	√	SCCP
17. Cisco IP Phone 7960G	√	√	SCCP
18. Cisco IP Phone 7961G	√	√	SCCP
19. Cisco IP Phone 7961G-GE	√	√	SCCP

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
20. Cisco IP Phone 7970	√	√	SCCP
21. Cisco IP Phone 7970G	√	√	SCCP
22. Cisco IP Phone 7971	√	√	SCCP
23. Cisco IP Phone 7985G	√	√	SCCP Video
24. Cisco IP Conference Station 7935	√	√	
25. Cisco IP Conference Station 7936	√	√	
26. Cisco IP SoftPhone	√	N	JTAPI/CTI - H/M-UCS does not support this phone type which is obsolete.
27. Cisco IP Communicator	√	√	SCCP
Cisco IP Phone Productivity Applications (IPPA)	√	N	The H/M-UCS architecture does not support this application
Cisco Personal Address Book	√	√B	BVSM provides the Personal Address book functionality in the H/M-UCS architecture
Cisco Personal Assistant (PA) Support	√	N	The H/M-UCS architecture does not support this application
Cisco Security Agent Support	√	√	The H.M-UCS architecture does not effect the availability of this feature.
Cisco VG248 48-Port Analog Gateway Support	√	N	
Cisco VG248 48-Port Analog Gateway Support	√	√*	The VG224 is only supported under H.323 control in the large enterprise deployment model. It is not supported in the multi tenant deployment model.
SMDI Voicemail Integration through VG248 analog ports	√	N	The H/M-UCS architecture does not support SMDI Voicemail integration
Click-to-Dial/ Click-to-Call	√	√	The H/M-UCS architecture supports a dial from Outlook feature based on TAPI.
Cisco WebDialer Click-to-Dial Service	√	N	
Client Matter Codes (CMC)	√	√*	CMC implementation has some restrictions in the H/M-UCS architecture in that CMC codes cannot overlap between tenants.
Closest-Match Routing	√	√	
Clustering ▼	√	√	
1. 10,000 Devices Per Cluster	√	√	These are maximum numbers. Some features may reduce the achievable maximum in some circumstances.
2. 30,000 Devices Per Cluster on 7845 class servers	√	√	These are maximum numbers. Some features may reduce the achievable maximum in some circumstances.

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
Conference Barge	√	√M	Supersedes Barge where mixing is done on the phone. This feature uses the conference bridges for mixing and hence can support G.729 as well as G.711. Achieved by setting service parameter and defining soft key template manually in H/M-UCS 1.6
Computer Telephony Integration (CTI) Support ▼	√	√	
1. CTI Manager Redundancy for TAPI/JTAPI applications	√	√M	The H/M-UCS architecture does not effect the availability of this feature.
2. JTAPI Control of Analog (FXS) Gateway Ports	√	√M	The H/M-UCS architecture does not effect the availability of this feature.
Context-Sensitive Help	√	√B	BVSM has an equivalent feature
Contrast	√	√	
Date/Time Zone Display Format configurable per phone	√	√M	BVSM does not provision time zones but they could be provisioned manually on the CallManager and associated with the location based device pools
Device Mobility	√	N	Device mobility is a new feature in Callmanager 4.2 that has yet to be picked up by the H/M-UCS architecture or the BVSM provisioning system.
Dial Plan Partitioning – Multi location	√	√	The H/M-UCS architecture includes a partitioned dial plan with full class of restriction support.
Dial Plan Partitioning and Class of Restrictions	√	√	The H/M-UCS architecture includes a partitioned dial plan with full class of restriction support.
Dialed Number Analyzer (DNA)	√	√	
Dialed Number Identification Service (DNIS)	√	√	
Redirected Number Identification Service (RDNIS)	√	√	
Outbound RDNIS to H.323 Gateways	√	√	
Digit Analysis and Translation	√	√	
Digit analysis and call treatment (digit string insertion, deletion, stripping, dial access codes, digit string translation)	√	√	Handled by combination of CallManager and PGW in the H/M-UCS design.
DSP resource management including distributed and topologically aware Resource Sharing	√	√	BVSM does not effect the availability of this feature on CCM. BVSM 3.1.6 now provisions CCM media resource groups in a more flexible way.
Direct Inward Dial (DID/DDI) Support	√	√	
Direct Outward Dial (DOD) Support	√	√	

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
Direct Transfer	√	√	
Directories ▼			
1. Corporate Directory	√	√B	BVSM provides this feature
2. Personal Directory	√	√B	BVSM provides this feature
3. Missed Calls	√	√	
4. Placed Calls	√	√	
5. Received Calls	√	√	
Distributed Call Processing▼			
1. Deployment of devices and applications across an IP network	√	√	
2. Virtual clusters of up to eight Cisco Unified CallManager subscriber servers for scalability, redundancy, and load balancing	√	√	
3. Maximum of 7500 IP phones per Cisco Unified CallManager server and 30,000 per server cluster (configuration-dependent)	√	√	Cluster support for 30K phones/cluster (config-dependent, e.g. <500 Locations/cluster limitation also applies) available H/M-UCS 1.6
4. Maximum of 100,000 busy-hour call completions (BHCCs) per Cisco Unified CallManager server and 250,000 per server cluster (configuration-dependent)	√	√	
5. Intercluster scalability to more than 100 sites or clusters through H.323 gatekeeper	√	√*	Clusters connected together by both Gatekeeper and PGW in H/M-UCS architecture.
6. Intracluster feature and management transparency	√	√	
Dual-Tone Multi-Frequency (DTMF) Support ▼			
1. SCCP Out of Band	√	√	
2. MGCP Out of Band	√	√	
3. TAPI/JTAPI Out of Band	√	√	
4. H.323 H.245 Alpha-Numeric Out of Band	√	√	
5. SIP RFC2833 Inband	√	N	The H/M-UCS 1.6 architecture does not support a SIP trunk to a third party device. RFC2833 support on Callmanager requires a Media Termination Point (MTP) for every SIP call.

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
Dynamic Host Configuration Protocol (DHCP) Support	√	√	BVSM provides a DHCP server with enhanced functionality including the ability to detect a phones location based on IP subnet and then automatically move it to that location within the configuration database.
Emergency 911 (E911) Support	√	N	Cisco Emergency Responder is not supported as part of the H/M-UCS 1.6 architecture.
End Call	√	√	
eXtensible Markup Language (XML) Support ▼	√	√	
1. XML Tag Enhancements	√	√	
2. Assigning an XML Service to a Line Button	√	N	
Extension Mobility ▼	√	√	
1. Native Extension Mobility	√	√	BVSM enhances the extension mobility operation in a H.M-UCS architecture to provide security between tenants in multi tenant deployments.
2. Extension Mobility Support of 7905, 7912, 7970 and CIPC	√	√	
FAX/Modem over IP Support ▼			
1. Fax Pass-Through	√	√	
2. Cisco Fax-Relay	√	√	This can be supported in a H/M-UCS environment provided the gateways chosen support the protocol. Cisco AS5400 does not support Cisco Fax Relay.
3. T.38 Fax-Relay for H323	√	√*	Supported only on H.323 gateways in the single tenant large enterprise H/M-UCS deployment model.
4. Modem Pass-Through	√	√	
5. Cisco Modem-Relay	√	√	
Forced Authentication Codes (FAC)	√	N	
Foreign eXchange Office (FXO)/Foreign eXchange Station (FXS)	√	√	
H.323 Support ▼	√	√	
1. H.323 Gatekeeper Support	√	√	
2. Multiple Gatekeepers per CallManager Cluster	√	√	
3. 1,000 H.323 Calls per Node	√	√	
4. H.323 Trunks and Scalability Improvements	√	√	

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
5. H.323 Videoconferencing Support	√	√	
6. H.323 FastStart Signaling Support ▼	√	√	
7. Inbound H323 FastStart	√	N	H.323 Fast Start is not used in the H/M-UCS architecture
8. Outbound H323 FastStart	√	N	H.323 Fast Start is not used in the H/M-UCS architecture
9. H.323 Overlap Sending and Receiving	√	N	H.323 overlap sending is not used in the H/M-UCS architecture
10. H.323 Annex M.1 support for H.323 gateways and H.225 trunks	√	√	This is supported in the H/M-UCS architecture on intercluster trunks used within the large enterprise (single tenant) deployment model.
Hands-Free Speakerphone (SPKR) Support	√	√	
Hold/Resume	√	√	
Music on Hold	√	√	
Tone on Hold	√	√	
Hookflash / Hookflash Transfer ▼	√	√	
1. Inbound Hookflash/Hookflash Transfer on Analog Ports	√	√	
2. Outbound Hookflash/Hookflash Transfer on Analog Trunks	√	N	H/M-UCS does not support Analogue Trunks
3. Inbound Hookflash Transfer on Digital Trunks	√	N	Supported only on T1-CAS Trunks and T1-CAS trunks are not supported within the H/M-UCS architecture.
HTML Help Access From Phone	√	√	
Hunt Groups ▼			
1. TCD Longest-Idle Hunting	√	√M	Telephone call dispatcher is used for the Cisco attendant console only
2. TCD Broadcast Hunting	√	√M	Telephone call dispatcher is used for the Cisco attendant console only
3. TCD Queuing	√	√M	Telephone call dispatcher is used for the Cisco attendant console only
4. Native Hunt Groups (Circular, Longest-Idle, Linear)	√	√	
5. Log in or log out of Hunt Group	√	√	This is a new feature in BVSM 3.1.6 which is part of the H/M-UCS 1.6 release.
6. Hunt Group Enhancements - Broadcast	√	√	
7. Hunt Group Enhancements - Circular	√	√	
8. Hunt Group Enhancements - Longest Idle	√	√	

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
9. Hunt Group Enhancements - Linear	√	√	
IP Manager Assistant - Shared Line and Proxy Line	√	N	The H/M-UCS architecture does not support IPMA
Inline Power ▼	√	√	
1. IEEE 802.3af Power over Ethernet (PoE)	√	√	Support for 802.3af PoE on the phones does not depend on a particular release of CallManager
Internationalization/Localization ▼			
1. Globalization of Cisco IP Phone 79xx Series	√	√	
2. Configurable User and Network Locales	√	√	
3. Downloadable User and Network Locales	√	√	
4. International Dial Plan	√	√	H/M-UCS does not use the Callmanager international dial plan but implements its own multi country capable dial plan.
ISDN Basic Rate Interface (BRI) Support ▼			
1. H.323 Controlled ISDN BRI	√	√M	The H/M-UCS solution can support BRI interfaces as local PSTN gateways H323 controlled from the PGW. Some manual configuration is required as the architecture only supports configuration of E1 or T1 physical interfaces from BVSM.
2. MGCP-Controlled ISDN BRI	√	N	
3. SIP-Controlled ISDN BRI	√	N	
Join	√	√	
Last Number Redial (On-net and Off-net)	√	√	
Least-Cost Routing (LCR)	√	N	
Light-Weight Directory Access Protocol (LDAP) Support	√	√	
External LDAP Directory Integration Support	√	N	
Line	√	√	
Multiple Line Appearances per Phone	√	√	
Multiple Calls per Line	√	√	
Select Specified Line Appearance	√	√	
Shared/Bridged Line Appearances	√	√	
Multiple Calls per Shared Line	√	√	

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
Privacy	√	√	
Configurable Call Forward Busy Trigger per Line Appearance	√	√	
IPMA Support for Shared Line Appearances	√	N	IPMA is not supported in the H/M-UCS solution
Unassigned Directory Numbers	√	√	
Manager-Assistant Services ▼	√	N	IPMA is not supported in the H/M-UCS solution
1. IPMA Support on Cisco IP Phone 7940	√	N	IPMA is not supported in the H/M-UCS solution
2. IPMA Support for Shared Line Appearances	√	N	IPMA is not supported in the H/M-UCS solution
Malicious Call ID	√	√	
Malicious Call ID and Trace	√	N	
Mappable Soft keys	√	√M	Soft key templates cannot be created from BVSM but, once created manually on the Callmanager BVSM can assign them to phones.
Media Gateway Control Protocol (MGCP) Support ▼	√	√	The H/M-UCS architecture defines the MGCP gateways for use as Central PSTN Gateways and for supporting TDM PBX's. The PGW is used to control all MGCP gateways, not the CallManager.
1. MGCP Gateway Fallback to H.323	√	N	The H/M-UCS architecture does not include MGCP fallback to H.323 due to the MGCP version that PGW uses (1.0) not being supported in this mode by IOS gateways.
2. MGCP ISDN T1/E1 PRI and T1-CAS with Q.931 Backhaul	√	√*	T1-CAS is not supported in the H/M-UCS architecture
3. Network-Specific Facilities (NSF)	√	√	
Messages Button	√	√	
Message Waiting Indicator (MWI)	√	√	
Multi-Level Precedence and Preemption (MLPP) ▼	√	N	The H/M-UCS solution does not support MLPP
1. Routine, Priority, Immediate, Flash and Flash Override	√	N	The H/M-UCS solution does not support MLPP
2. Executive Override	√	N	The H/M-UCS solution does not support MLPP
3. MLPP over MGCP-Controlled T1/E1 PRI and T1-CAS Trunks	√	N	The H/M-UCS solution does not support MLPP
4. MLPP over SCCP-Controlled ISDN BRI Ports	√	N	The H/M-UCS solution does not support MLPP

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
5. MLPP over H.323 Inter-Cluster Trunks	√	N	The H/M-UCS solution does not support MLPP
6. MLPP Based on Locations-Based Call Admission Control (CAC)	√	N	The H/M-UCS solution does not support MLPP
Multi-Party Conferencing ▼			
1. Ad Hoc Conferencing (up to 4 participants)	√	√	
2. Drop Last Conference Party	√	√	
3. List All/Drop Any Conference Party	√	√	
4. Drop Conference When Initiator Leaves	√	√	
5. Drop Conference When No OnNet Parties Remain	√	N	The H/M-UCS architecture does not provide sufficient information to CallManager regarding internal calls vs. external calls
6. Release Conference Bridge When Only Two Parties Remain	√	√	
7. Multiparty conference-ad-hoc with add-on, meet-me features	√	√	
8. Meet-Me Conferencing (up to 10 participants)	√	√	
Multiple Calls per Line Appearance	√	√	
Multiple Line Appearances per Phone	√	√	
Mute	√	√	
Mute Capability from Speakerphone and Handset	√	√	
New Call	√	√	
North American Numbering Plan (NANP)	√	N/A	H/M-UCS design loads country specific number plans into CallManager via BVSM
Non-NANP Support	√	N/A	H/M-UCS design loads country specific number plans into CallManager via BVSM
Off-Premise Extension (OPX) Support	√	√	
On-Hook Dialing	√	√	
Outbound Call Blocking	√	√	H/M-UCS supports outbound call restrictions
Overlap Sending/Receiving	√	N	
Paperless Phone Labels	√	√	
Privacy	√	√	
Privacy Button	√	√M	
Private Line Auto Ringdown (PLAR) Support	√	√	

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
Q.SIG Support ▼			
1. Basic Call	√	√	PBX to PBX (DPNSS or QSIG) and PBX to CCM phones
2. Call Back ISO/IEC 13870: 2nd Ed, 2001-07 (CCBS, CCNR)	√	√*	PBX to PBX only (DPNSS or QSIG)
3. Calling Line Identification Presentation (CLIP)	√	√	PBX to PBX (DPNSS or QSIG) and PBX to CCM phones
4. Calling Name Identification Presentation (CNIP)	√	√	PBX to PBX (DPNSS or QSIG) and PBX to CCM phones
5. Connected Name Identification Presentation (CONP)	√	√*	PBX to PBX (DPNSS or QSIG) and PBX to CCM (except call transfer cases- see CSCsh53557)
6. Calling/Connected Line Identification Restriction (CLIR)	√	√	
7. Calling/Connected Name Identification Restriction (CNIR)	√	√	
8. Alerting Name specified in ISO 13868 as part of the SS-CONP feature	√	N	
9. Message Waiting Indicator (MWI)	√	√*	This has only been tested to interoperate with the Voicerite Voicemail system
10. Support for QSIG and Non-QSIG Devices in Route Lists	√	N/A	
11. Call Diversion (SS-CFB(Busy))	√	√	
12. Call Diversion (SS-CFNR(No Answer))	√	√	
13. Call Diversion (SS-CFU(Unconditional))	√	√	
14. Call Diversion by Reroute	√	√	
15. Call Completion (a.k.a. Call Back)	√	√*	PBX to PBX only (DPNSS or QSIG)
16. Call Transfer by Join	√	√	
17. Loop Prevention	√	N	PBX to PBX only
18. Loop Detection	√	√	PBX to PBX only
19. Diversion Counter and Reason	√	√	
20. Diversion To Number	√	√	
21. Diverting Number	√	√	
22. Original Called Name & Number	√	√	
23. Original Diversion Reason	√	√	
24. Redirecting Name	√	√	

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
25. Path Replacement ISO/IEC 13863: 2nd Ed. 1998 and ISO/IEC 13974: 2nd Ed. 1999.	√	√*	PBX to PBX only (DPNSS or QSIG)
26. Q.SIG over H.323 (Annex M.1)	√	N	Requires PGW 9.7(3) for interoperation with PGW, This is planned in the H/M-UCS 1.7 release
Quality of Service (QoS) ▼	√	√	
Differentiated Services (DiffServ) and IP Precedence (ToS)	√	√	
802.1p Class of Service (CoS)	√	√	
QoS Statistics	√	√	
Recent Dial List ▼			
1. Calls To Phone	√	√	
2. Calls From Phone	√	√	
3. Auto Dial	√	√	
4. Edit Dial	√	√	
Redundancy/Fail Over for Call Processing Servers	√	√	
CTI Manager Redundancy for TAPI/JTAPI applications	√	√	
Security ▼			
1. Configurable operation modes: non-secure or secure	√	N	
2. Device authentication: Embedded X.509v3 certificate in new model phones. CAPF used to install locally significant certificate in phones.	√	N	
3. Data Integrity: TLS cipher "NULL-SHA" supported. Messages appended with SHA1 hash of the message to ensure that the message is not altered on the wire and can be trusted.	√	N	
4. Secure HTTP (HTTPS) support for Callmanager Administration access	√	√	BVSM also has HTTPS access to its web interface.
5. Privacy: CallManager supports encryption of signaling and media.	√	N	
6. Secure Sockets Layer (SSL) for directory	√	√B	Callmanager directory is not used in the H/M-UCS environment. BVSM provides the directory service.

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
7. USB eToken containing a Cisco rooted X.509v3 Certificate is used to generate a Certificate Trust List (CTL) file for the phones as well as configuring the security mode of the cluster.	√	N	
8. Phone Security: TFTP files (configuration and firmware loads) are signed with the self-signed certificate of the TFTP server. The CallManager system admin will be able to disable http and telnet on the IP phones	√	√	
Service URL-Single-Button Access to IP Phone Service	√	√	
Settings	√	√	
Contrast/LCD Contrast	√	√	
Ring Type (on a per line basis)	√	√	
Network Configuration/Network Settings	√	√	
Status/Phone Info	√	√	
Single Directory Number, Multiple Phones-Bridged Line Appearances	√	√	
Simple Network Management Protocol (SNMP) Support ▼	√	√	
1. CISCO-CCM-MIB SNMP MIB	√	√	
SRST	√	√	
Speed Dials ▼	√	√	
2. System Speed Dials	√	√	
3. User-Programmable Speed Dials	√	√	
4. Multiple Speed Dials per Phone	√	√	
5. Configuration of speed dials from Phone	N	√	BVSM provides this via an IP Phone XML service.
Station Volume Controls (Audio, Ringer)	√	√	
Supplementary Services	√	√	
T1/E1 PRI Support	√	√	
T1/E1-CAS Support	√	N	CAS gateways are not currently supported as within the H/M-UCS architecture for interconnection to TDM PBX
Third Party Applications Support ▼			
1. Broadcast paging—through foreign exchange station (FXS)	√	√	Requires third party paging system.

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
2. Simple Messaging Desktop Interface (SMDI) for message waiting indication	√	N	SMDI integration to Voicemail systems is not supported as part of the H/M-UCS architecture
3. Hook-flash feature support on selected FXS gateways	√	√	
4. TAPI 2.1 service provider (TSP) interface	√	√	
5. JTAPI 2.0 service provider interface	√	√	
6. Billing and call statistics	√	√	
7. Configuration database API (Cisco AVVID XML Layer)	√	√	
Time of day, day of week, day of year routing/restrictions	√	√	
Toll Restriction/Toll Fraud Support ▼	√	√	
Forced authorization codes	√	N	
Conferencing Enhancements (for Toll Fraud) ▼	√	√M	manually provisioned service parameters and must have same behavior for whole cluster
1. Drop Conference When Initiator Leaves	√	√M	manually provisioned service parameters and must have same behavior for whole cluster
2. Drop Conference When No OnNet Parties Remain	√	N	On net vs. off net information not available in H/M-UCS design.
3. Drop conference call when originator hangs up	√	√M	manually provisioned service parameters and must have same behavior for whole cluster
External Transfer Restrictions (for Toll Fraud) ▼	√	√	
External vs. Internal Trunk Designation	√	N	The H/M-UCS architecture does not currently support this feature
Prevent trunk-to-trunk transfer	√	N	The H/M-UCS architecture does not currently support this feature
Transcoding and Media Termination Point (MTP) Support	√	√	
Transfer ▼	√	√	
1. Blind Transfer	√	√	
2. Consultative Transfer	√	√	
3. Consultative Transfer Support in Cisco Attendant Console	√	√	
4. Direct Transfer	√	√	
5. Direct Transfer of Two Parties on a Line	√	√	
Trivial File Transfer Protocol (TFTP) Support	√	√	

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
User-Programmable Speed Dials, Call Forwarding and Services	√	√	
Video Telephony Support ▼	√	√	
1. Cisco VT Advantage Support	√	√	
2. Tandberg SCCP Video Endpoints Support	√	√M	These third party video endpoints have not been tested as part of the H/M-UCS architecture but there is nothing in the architecture that should prevent them from working. The phone type would have to be added manually to BVSM 3.1.6.
3. SCCP Support on Cisco IP/VC 3500 Series Multipoint Conference Units (MCUs)	√	√*	
4. H.264 Video Codec Support on SCCP Endpoints	√	√	
5. Mid-Call Video	√	√	
6. User-Selectable Video Display Mode	√	√	
7. Participant Information in Multipoint Video Conferences	√	√	
8. Dynamic H.323 Client Addressing Support	√	√	
Virus Protection Certification	√	√	
Voice Activity Detection (VAD)/Silence Suppression Support	√	√	
Comfort Noise Generation	√	√	
Voice Mail Support ▼			
1. SCCP Voice Mail Integration with Cisco Unity	√	√M	Cisco Unity (dedicated per tenant) is supported via manual config and integration.
SMDI Voicemail Integration	√	N	
2. Cisco Digital Port Adapter (DPA) 7600 Series Support	√	N	
3. Q.SIG Voice Mail Integration with MGCP-controlled Trunks	√	N	H/M-UCS supports Q.SIG integration with Voicerite
4. Q.SIG Voice Mail Integration with Annex M.1 H.323 Trunks	√	N	
5. Per-Line Configurable MWI	√	√	
6. Voice Mail Profiles	√	√	
7. Immediate Divert to Voice Mail	√	√	
Voice Quality Statistics (call-by-call) basis ▼	√	√	

Feature	CCM 4.2.3	H/M-UCS 1.6	Comments
Cumulative conceal ratio	√	√	
Interval conceal ratio	√	√	
Interval conceal ratio maximum	√	√	
Conceal seconds	√	√	
Severely concealed seconds	√	√	
MOS listening quality k-factor	√	√	
MOS listening quality k-factor minimum	√	√	
MOS listening quality k-factor maximum	√	√	
MOS listening quality k-factor average	√	√	
Volume Control	√	√	
Web Services Access from Phone	√	√	
Web Dialer-Click to Dial	√	√M	Web Dialer has not been tested as part of the H/M-UCS 1.6 architecture but it should be possibly to manually integrated it into the solution.
Year 2000 Compliance	√	√	

Client Matter Codes (CMC)

BVSM version 3.1.6 includes a new feature for provisioning CallManager client matter codes. In BVSM these are configured as “Billing Codes”. Client matter codes on the CallManager are enabled at the route pattern level hence the feature also requires changes to the CallManager model loader to add the necessary patterns and associated classes of service. Once enabled through the model, the feature is used by adding a BVSM feature group calling up the class of service that requires a CMC and this feature group is then associated with phones.

Billing codes themselves can be managed from the BVSM which allows the to be added in ranges and then associated with resellers, customers, divisions, locations, tenants, areas or users. It should be noted that BVSM assumes billing codes are unique across the platform it controls so different resellers, customers etc cannot use the same billing code.

Call Quality Reporting Using the Callmanager 4.2, PGW and Gateway “K Factor” Features

This feature allows H/M-UCS platform operators to measure voice quality end to end without use of separate probes and hence commit to a voice quality SLA in contracts with their customers.

There are three sources for voice quality metrics – IP Phones (with CCM 4.2), voice gateways and trunking gateways.

The voice quality metrics for IP Phones get written to the CallManager CMR records on the Callmanager cluster publisher and those from the gateways controlled by the PGW are reported in PGW/BAMs CDRs. These CMRs and CDRs are retrieved by the service provider by existing procedures. It is beyond the scope of this document to describe the post-processing that may be needed to extract useful information from the CMRs/CDRs related to voice quality metrics.

K-Factor Information from SCCP IP Phones

CallManager 4.2 supports voice quality statistics that originate from the following devices SCCP IP Phones:

- 7911G
- 7940G
- 7941G, 7941G-GE
- 7960G
- 7961G, 7961G-GE
- 7970G, 7971G-GE

Voice quality data stored in CallManager CMR records include K-Factor, Concealed Seconds, Severely Concealed Seconds & Concealed Ratio VQ metrics. The table below shows the CMR field name and equivalent phone display name viewed by pressing the “Help / ?” button twice during a call.

Table 12 Voice Quality Metrics

Description	CMR Field Name	Phone Display Name
Cumulative Conceal Ratio	CCR	Cum Conceal Ratio
Interval Conceal Ratio	ICR	Interval Conceal Ratio
Interval Conceal Ratio Max	ICRmx	Max Conceal Ratio
Conceal Secs	CS	Conceal Secs
Severely Conceal Secs	SCS	Severely Conceal Secs
MOS Listening Quality K-factor	MLQK	MOS LQK
MOS Listening Quality K-factor Min	MLQKmn	Min MOS LQK
MOS Listening Quality K-factor Max	MLQKmx	Max MOS LQK
MOS Listening Quality K-factor Avg	MLQKav	Avg MOS LQK

In addition to the existing voice quality data supported by the SCCP phone, new measurements are added in the new varVQMetrics field. This string field contains voice quality metrics separated by semicolons.

The format of the string is either fieldName=value or fieldName=value/precision.

Below is an example of CMR record with voice quality statistics.

“MLQK=3.5146;MLQKav=3.3999;MLQKmn=3.0;MLQKmx=3.5146;ICR=0.0;CCR=0.0156;ICRmx=0.0500;CS=15;SCS=2”

See “Cisco CallManager 4.2(1) Call Detail Record Definition” for full CDR/CMR definition at

http://www.cisco.com/en/US/products/sw/voicesw/ps556/products_programming_usage_guide09186a00805538a5.html

K-Factor Information from PGW Controlled Gateways

The PGW supports the collection of digital signal processor (DSP) statistics from voice gateways in the PGW call detail records (CDRs). These DSP statistics are provided in the Media Gateway Control Protocol (MGCP) Delete Connection (DLCX) message.

An IOS software feature, DSP Voice Quality in DLCX Messages, has modified the DSP statistics gathering function on Cisco media gateways. This feature provides a way to trace a MGCP call between a PGW and the Cisco IOS gateway by including the MGCP call ID and the DS0 and DSP channel ID in call-active and call history records. These DSP statistics are sent as part of the MGCP Delete Connection (DLCX) message. By correlating an MGCP call on the PGW with the call record on the gateway, additional statistics from the DSP can be understood and debugged for problems related to voice quality.

It is recommend that the DSP Voice Quality Statistics in DLCX Messages feature is not turned on for all of the media gateways associated with a single PGW 2200 unless they are needed. Doing so can impact the call processing performance of the PGW unnecessarily.

The DSP voice quality statistics are available in Cisco IOS release 12.4(4)T and can be used only with the following gateways which use the c5510 DSP:

- Cisco 1760
- Cisco 1751
- Cisco 26xx (must use the NM-HDV2, NM-HD, or EVM-HDM)
- Cisco 36xx (must use the NM-HDV2, NM-HD, or EVM-HDM)
- Cisco 37xx (must use the NM-HDV2, NM-HD, or EVM-HDM)
- Cisco 28xx
- Cisco 38xx
- Cisco AS5350XM
- Cisco AS5351XM
- Cisco AS5400XM

The following voice quality parameters are available:

- DSP/Endpoint Statistics
- DSP/MOS/K-factor Statistics
- DSP/Concealment Statistics
- DSP/R-factor Statistics
- DSP/User Concealment Statistics
- DSP/Delay Statistics

When the PGW software patch CSCOnn026 is installed, the functionality for this feature is automatically activated. No configuration or provisioning changes are required to activate this feature.

The BAMs generates output files that include voice-gateway Quality-of-Service (QoS) statistics produced by the voice gateways, collected by the Cisco PGW 2200.

The BAMS 3.20 User Guide Appendix D: Quality of Service Statistics Output details the changes to accommodate the new QoS statistics.

The functionality to generate QoS statistics output files is included in a patch of BAMS named CSCOBAMS320QOS.pkg

The PGW call data elements (CDEs) that were modified for this feature are described in:

http://www.cisco.com/univercd/cc/td/doc/product/access/sc/rel9/mgcfm/96/fmdspen.htm#wp1018968

TDM PBX Services

Cisco H/M-UCS provides TDM PBX integration to DPNSS, PRI and QSIG PBXs. The architecture and signalling is described else where in this SRND. Below are the features supported for each type of PBX.

DPNSS PBX

DPNSS PBX to DPNSS PBX Supplementary Service Interworking

All DPNSS telephony services inter-worked transparently. Please refer to http://www.cisco.com/univercd/cc/td/doc/product/access/sc/rel9/mgcfm/941fm/fmdpnss.htm

DPNSS PBX to PSTN Supplementary Service Interworking

- Call Hold
- Caller Identity
- Caller Identity Restriction
- 3 Party Calls
- Call Waiting

DPNSS PBX to Cisco CallManager Supplementary Service Interworking

- Call Hold
- Caller Identity
- Caller Name
- Caller Identity Restriction
- Message Waiting Indication (Although PGW does support this feature it is not supported in the H/M-UCS 1.6 architecture)
- 3 Party Calls
- Call Waiting
- Call Back When Free (Although PGW does support this feature it is not supported in the H/M-UCS architecture)

For full description of DPNSS Supplementary Services Interworking with Cisco CallManager please refer to http://www.cisco.com/univercd/cc/td/doc/product/access/sc/rel9/mgcfm/96/fmdpsup2.htm and http://www.cisco.com/univercd/cc/td/doc/product/access/sc/rel9/mgcfm/96/index.htm

QSIG PBX

QSIG PBX to QSIG PBX Supplementary Service Interworking

All QSIG telephony services inter-worked transparently. Please refer to http://www.cisco.com/univercd/cc/td/doc/product/access/sc/rel9/mgcfm/941fm/fm_qsig.htm

QSIG PBX to PSTN Supplementary Service Interworking

- Caller Identity
- 3 Party Calls
- Call Forwarding
- Call Waiting

QSIG PBX to Cisco CallManager Supplementary Service Interworking

- Caller Identity
- Message Waiting Indication
- 3 Party Calls
- Call Forwarding
- Call Waiting

Advanced XML Services

Simple XML services, such as employee directories and information services can be made easily available to all staff via their IP phones, allowing them to become more efficient and improving employee satisfaction.

XML control of the phone display is available on Cisco IP Phone models 797X, 796X and 794X to offer customized services through the phone display and keypad. Cisco IP Phones are capable of launching XML applications that enable the display of interactive content with text and graphics on the phone's LCD display.

Typical information services that might be supplied to a phone include company information, weather information, stock quotes, and news quotes.

More advanced, problem-solving services can also be created and deployed to increase productivity, gain a competitive advantage, and even generate revenue. Potential examples include organizational or corporate announcements, weather information (e.g. road ice warning), stock quotes and other news quotes (e.g. competitor news), inventory check, flight schedule, time card, shipment tracking, instant messaging, Outlook Integration (Calendar & Email control), Lotus Integration, facility management and conference room booking.

Cisco IP Phone models 797X, 796X and 794X provide a Services button. When the user presses this button, the phone uses its HTTP client to load a specific URL that contains a menu of services which the user has subscribed to for their phone. The user then chooses a service from the listing. XML primitives can be used to override soft key set definitions and to add new keys with associated URLs for application control.

Cisco H/M-UCS allows a system administrator to setup URL links to services and then create packages of services via the GUI. Customer administrators and end users can then subscribe to these services on a per user basis, accessible via the services button on Cisco IP phones. On subscribing to these services, Cisco H/M-UCS creates event data records to allow service provider billing (if appropriate).

Voice Mail Services

Multitenant Voice Mail

For the Cisco H/M-UCS reference architecture, the IP Unity Mereon Unified Messaging platform is used to provide voice mail services. The IP Unity voice mail system has been selected for its capability to support the multi-tenant voicemail service in the Cisco H/M-UCS reference architecture. The BVSM platform is expected to be integrated with IP Unity to allow provisioning via the BVSM GUI of voice mail functionality on a per-tenant-site basis.

The following services are provided to subscribers on the IP Unity voice mail system.

- Multi-Tenant Voicemail
- Voicemail retrieval via a telephone
- Voicemail recording
- Message Waiting Indication using IP Phone Lamp
- English Language support
- Customisable prompts
- G.729 and G.711 codec support
- Personalised greetings

Single Tenant Voice Mail

For Cisco H/M-UCS deployments that require large single tenant hosted voice mail services the Cisco Unity voice mail platform is recommended. The Cisco Unity product itself does not operate in a multi tenant environment in that no segmentation is possible within the Cisco Unity data store, and everything from administration to management has been designed on the assumption that a server will deployed into a single company. The Cisco Unity is manually provisioned and uses SCCP interface to Cisco Unified CallManager. When deployed in a multi tenant H/M-UCS environment separate Unity instances will be required per tenant. It is not possible to provision the Cisco Unity product from the BVSM provisioning system in the 1.6 architecture release.

The following Cisco Unity features are available for single tenant deployments:

- Saving of voice mail.
- Listening to voice mail.
- Deleting of voice mail.
- Replying to voice mail.
- Personal directory listings.
- System directory listings that are centrally administered.
- End-user administration via phone.
- End-user administration via web-based GUI.
- Transfer to external number.
- Out dial message waiting indication
- MWI via lamp on phone
- Unified Messaging with Microsoft Exchange or Lotus Domino Integration

Attendant Console Services (Netwise CTC/CMG/NOW)

The Netwise attendant console is a third party component that was chosen due to its ability to support a multi tenant attendant console environment from a single set of shared hardware and Callmanager resources (CTI route points and CTI ports). Major services and functionality includes:

- Line status information in list view and on call
- "Camp on" functionality
- Pop-up on incoming call with greeting phrase in highlighted text
- Application on top on incoming call
- Queue list with extensive information
- Support for entering, changing and deleting of activity and contact profile for a selected group of extensions
- Support for sending of messages to a selected group, e.g. via e-mail
- Integrated web browser with access to Quick Info directory
- Visitor management (optional)
- Customizable workspace
- Support for visually disabled
- Activity database
- User directory
- Quick Info directory
- Interfaces to email system or calendar system
- Configuration and administration tools
- Support for database synchronization
- Support for Microsoft Windows Clustering solution

Cisco IP Communicator

The Cisco IP Communicator is a Microsoft Windows-based application that delivers enhanced telephony support through personal computers. This application allows computers with the functionality of IP Phones, providing voice calls on the road, in the office, or from wherever users may have access to the corporate network.

A number of advanced call features are available depending on the Cisco Unified Cisco Unified CallManager system including:

- Support of multiple lines or directory numbers
- Configurable speed dials
- Calling name and number display
- Call waiting
- Call forward
- Call transfer
- Three-way calling (conference)
- Park
- Pick-up
- Redial

- Call hold
- Barge
- Call back (where architecture allows)
- Extension mobility

Cisco Unified Video Advantage

Video support in Cisco H/M-UCS is based on Cisco Unified Video Advantage and 7985 Video phone. SCCP based Video conferencing is support for intra-site mode with ad-hoc functionality using IP/VC 3511, 3540 Multimedia Control Unit with EMP. Note that video calls are supported within sites, between sites and even between tenants on a multi tenant H/M-UCS platform.

The Cisco Unified Video Advantage brings video telephony functionality to Cisco Unified IP Phones (7900 Series and Cisco IP Communicator soft phone application). It is a video telephony solution consisting of Cisco Unified Video Advantage software and Cisco VT Camera II, a video telephony 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-pushing or mouse-clicking.

When registered to Cisco Unified CallManager, the Cisco Unified Video Advantage-enabled phone has the features and functionality of a full-featured IP videophone.

Cisco Unified Video Advantage supports the following features:

- Phone association choice-Users can choose to place Cisco Unified Video Advantage video calls with either a Cisco Unified IP Phone or Cisco IP Communicator.
- Camera on/off-Users can choose to view incoming video only by turning off their camera.
- Video check-Users can check their video before calls are placed or received.
- Mute video on audio mute option-When users mute the audio on the phone, video is automatically muted until the audio is resumed.
- Easy access to video controls-Controls for showing the console, video window options, and video window position are conveniently available from the video windows.
- Video signal indicators-Quality of incoming and outgoing video signals are graphically displayed.
- Connectivity and status indicators-Graphics indicate the state and availability of the connections to the associated phone device and camera, including muted calls and "no available video."
- Supported USB webcams: Cisco VT Camera and Cisco VT Camera II (Cisco Unified Video Advantage 2.0 release and later)
- H.263, H.264, Cisco VT Camera Wideband video codecs - 50 kbps to 1.5 Mbps bit rates
- Video formats (up to 30 frames per second): 352 x 288, 320 x 240, 176 x 144, and 160 x 120

The following device and features will not be supported in current H/M-UCS release:

- 3rd Party SCCP Phones such as Tandberg (1000, T-550), Sony (PCS-1,PCS-TL50) etc
- 3rd Party H.323 Phones such as Tandberg, Sony etc
- IP/VC (3526, 3540) H.320 ISDN Video Gateways
- IP/VC(3511 BRI Video gateway)
- IP/VC (3521/3526 PRI video gateway)

Cisco Unified Personal Communicator

The Cisco Unified Personal Communicator is planned for a future release of the Cisco H/M-UCS reference architecture.

Cisco Unified Video Conferencing

Video conferencing is support for intra-site mode with ad-hoc functionality using IP/VC 3511, 3540 Multimedia Control Unit (MCU) with EMP. See Desktop client services and Cisco Unified Video Advantage.

Wireless Phone Services

The Cisco 7920 IP Phone extends the capability of IP Phones from 10/100 Ethernet to 802.11b WLANs. The 7920 provides a multi-line appearance with functionality similar to Cisco's existing 79xx IP Phones. In addition, the 7920 provides enhanced WLAN security and QoS for operation in 802.11b networks.

Cisco Unified Mobility Manager

The Cisco Unified Mobility Manager and Cisco Mobile Connect are planned for a future release of the Cisco H/M-UCS reference architecture.

Cisco Unified Meeting Place

The Cisco Meeting Place conferencing is planned for a future release of the Cisco H/M-UCS reference architecture. The introduction will be phased, integration with manual provisioning being available first followed by full provisioning by BVSM.

Cisco Unified Presence Server

The Cisco Unified Presence Server is planned for a future release of the Cisco H/M-UCS reference architecture.

Cisco Fax Server Services

The Cisco Fax Server is planned for a future release of the Cisco H/M-UCS reference architecture. Note Cisco Fax Relay and pass-through are supported, see bearer services above.

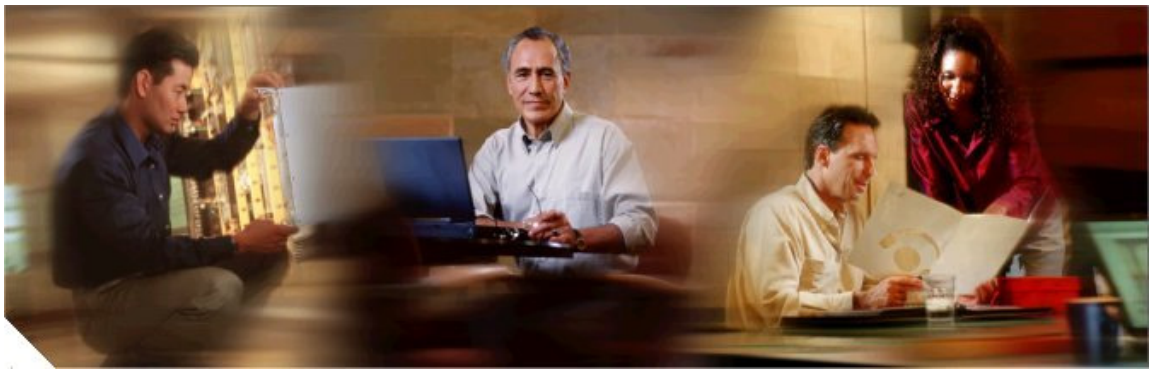

PSTN Interconnection

In the H/M-UCS reference architecture PSTN access is controlled by the PGW. Depending on where the PSTN gateways are located access to the PSTN can be classified as follows:

- Centralized Access
- Local Access
- Mixed (Centralized + Local) Access

Centralized Access

In the centralized access mode PSTN interconnection is achieved via the PGW interconnecting to the PSTN via SS7 or ISDN and inter-working with the H.323 VoIP network. The PSTN gateways are located in one or more centralized points of interconnection and are shared by all the customers serviced by the H/M-UCS deployment.

All PSTN calls (in-bound and out-bound) for all customers are handled by the same set of PGW controlled gateways. Out-bound gateway selection can be made intelligent my manual provisioning on the PGW to provide routing to the gateways based on criteria such as the called number (B number) (so as to provide least cost routing), or some other criteria applicable to all customers on the platform.

Forced On Net

The normal operation of a H/M-UCS platform with centralised gateways is to “force on net” that is calls that are to a destination that is served by the H/M-UCS platform itself, either within the same tenant or a different tenant, are not routed to the PSTN but remain within the platform itself. It is possible to configure a H/M-UCS platform to always pass E164 dialled calls to the centralised gateway devices, irrespective of whether the call is destined for a subscriber on the H/M-UCS platform or not. This may be useful in situations where inter tenant calls need to be routed to an upstream switch for billing purposes for example, or in situations where legal requirements dictate that the platform operator may not carry calls between customers (e.g. the platform operator does not have a carrier licence).

The determination of this behaviour is a platform wide decision made at install time by loading different PGW config models into the BVSM.

Figure 32 Centralized PSTN Access model.

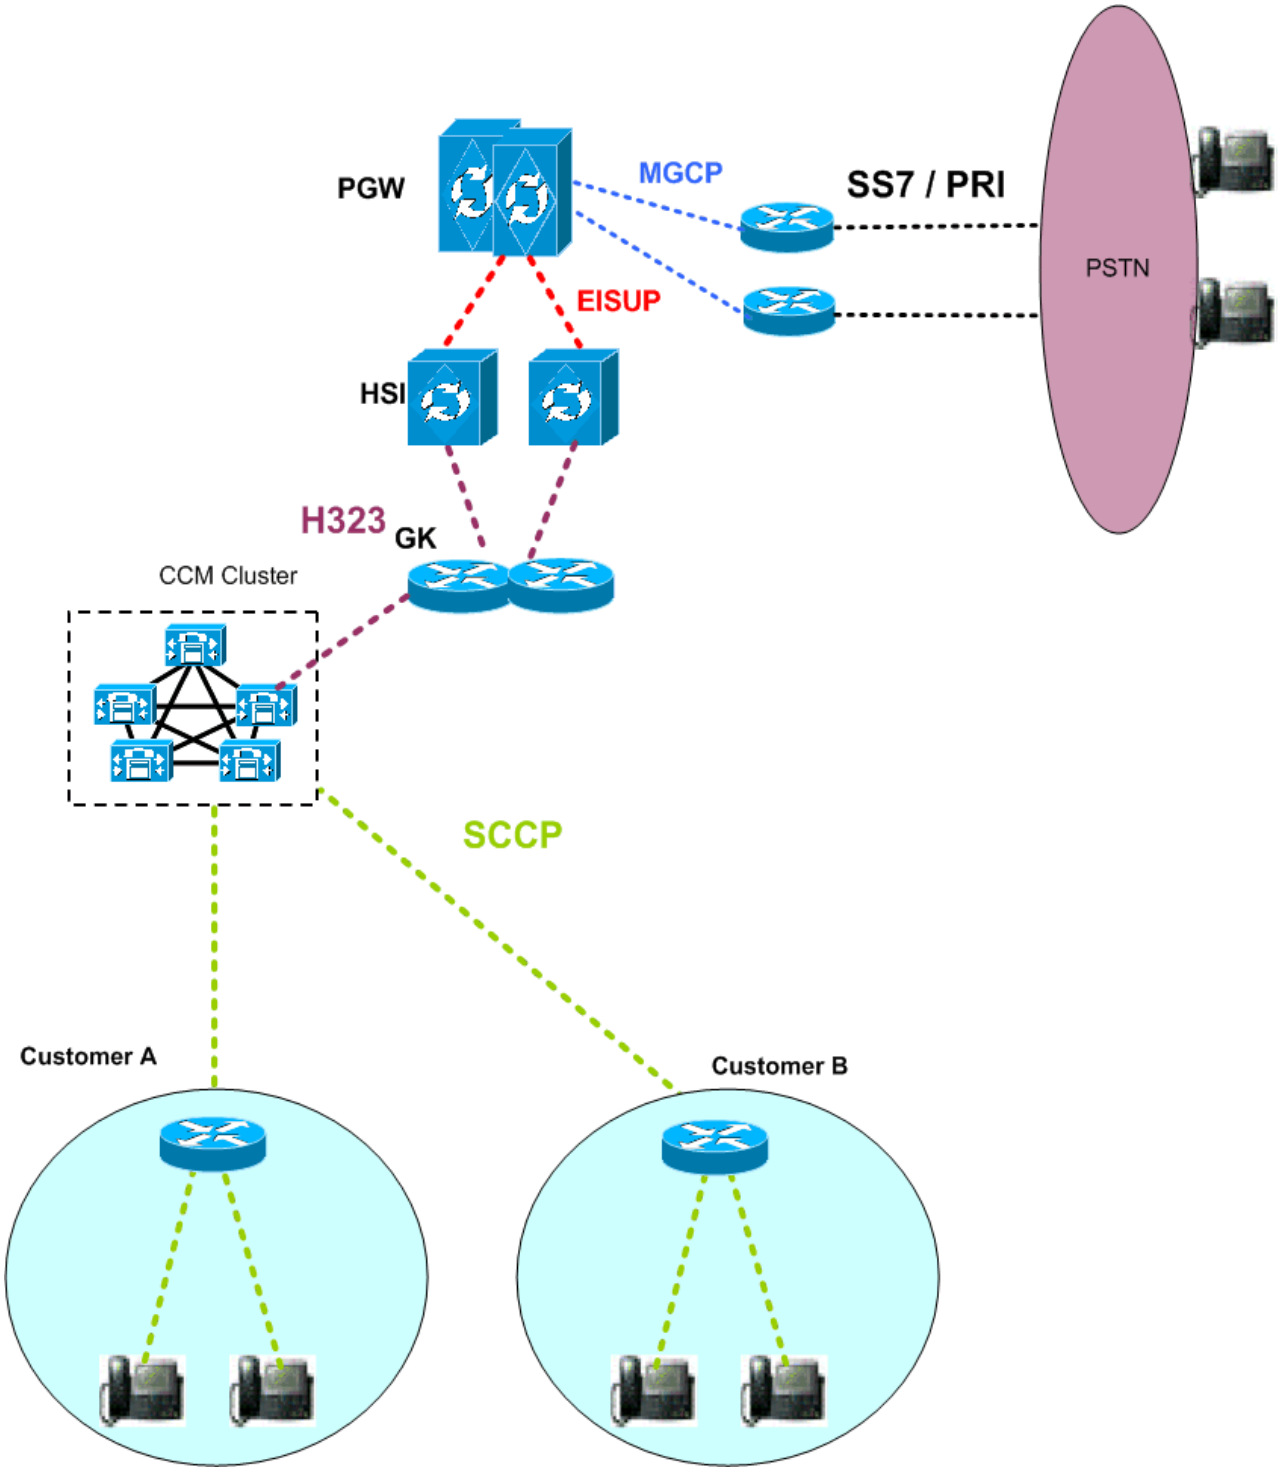

Local Access

In the local access mode PSTN interconnection is achieved via the PGW interconnecting to the PSTN via ISDN PRI though an H.323 controlled VoIP gateway. The PSTN gateways are located at the customer premises or in the geographic vicinity of the customer.

In this model one or more customers (provided they are co located) can share a local gateway for PSTN access. This means that inbound and outbound gateway selection is done on a per location basis. The PSTN gateway may be optionally configured to support SRST in the event that the CallManager connection fails.

Caveats

Following are the caveats for the local access mode:

- Maximum of 13 locations per gateway (due to the number of translation rules allowed on IOS – 15 allowed total but 2 are used for PSTN access and emergency calls).
- When in SRST fallback mode, IP phones must use E164 numbers to call one another.
- While in SRST fallback mode, the DN displayed in the IP phone is the phone's BVSM internally assigned number.
- A maximum of 10,000 local PSTN gateways can be supported.
- The only Local Gateway selection algorithm supported in H/M-UCS 1.6 is "near end hopoff" i.e. If a site has a local gateway calls originated from phones on that site are the only calls that use that gateway. No overflow mechanisms are supported to allow sites with local gateways to use a central pool of gateways or another site's local gateways.

Figure 33 Local PSTN Access model.

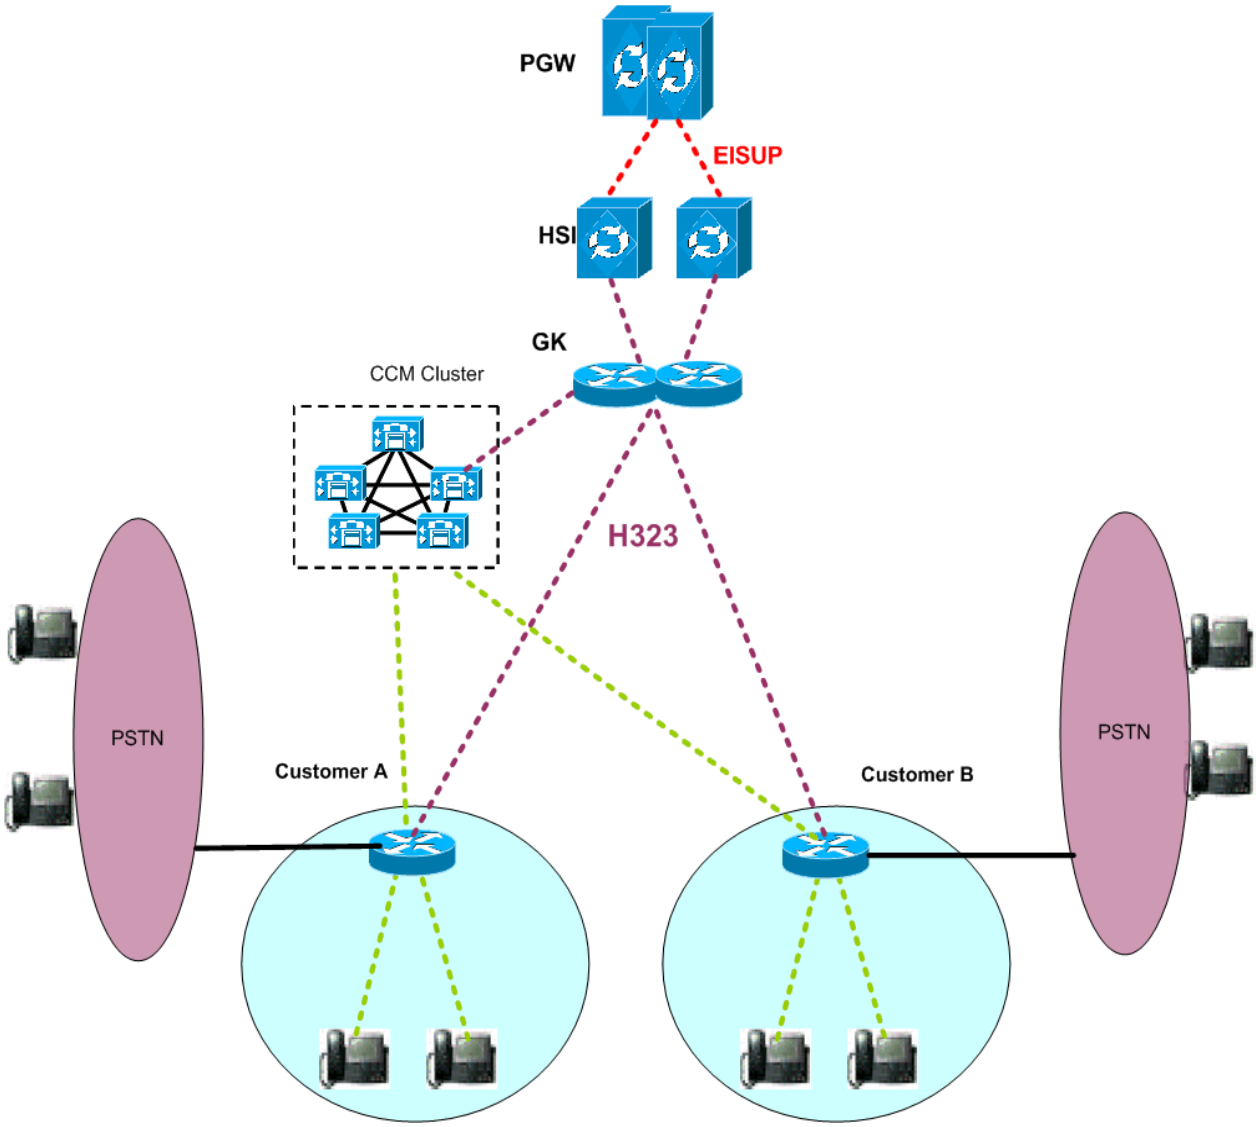

Mixed (Centralized + Local) Access

The mixed access mode is a combination of the two previous models. PSTN interconnection is achieved via the PGW configured in signalling mode interconnecting to the PSTN via SS7 or ISDN and inter-

working with the H.323 VoIP network. Some PSTN gateways are located in one or more centralized points of interconnection and can be shared by all the customers serviced by the H/M-UCS deployment, and some other PSTN gateways are located at the customer premises or in the geographic vicinity of the customer. The gateway selection rules supported in the H/M-UCS 1.6 architecture are limited to the following.

- If a give site has a local gateway then outgoing PSTN calls from phones on that site are always directed to use it.
- Sites without a local gateway always use the central pool of gateways for PSTN access. They cannot use local gateways at another site.
- No overflow or failure routing mechanisms are supported between central gateways and local gateways or between local gateways at different sites.

Figure 34 Mixed PSTN Access model.

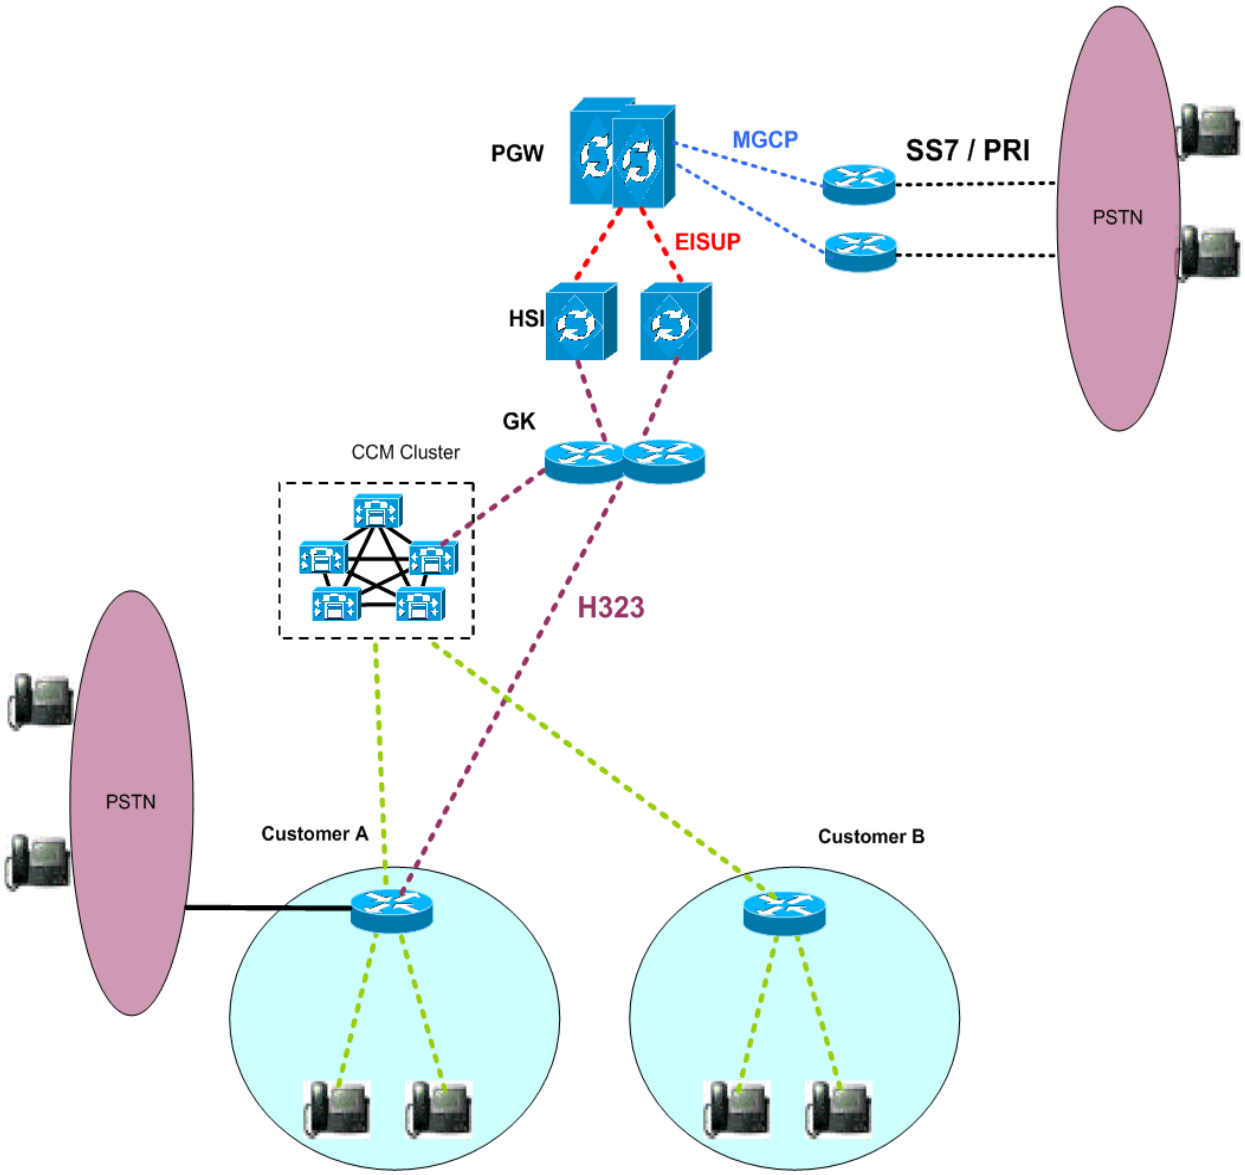

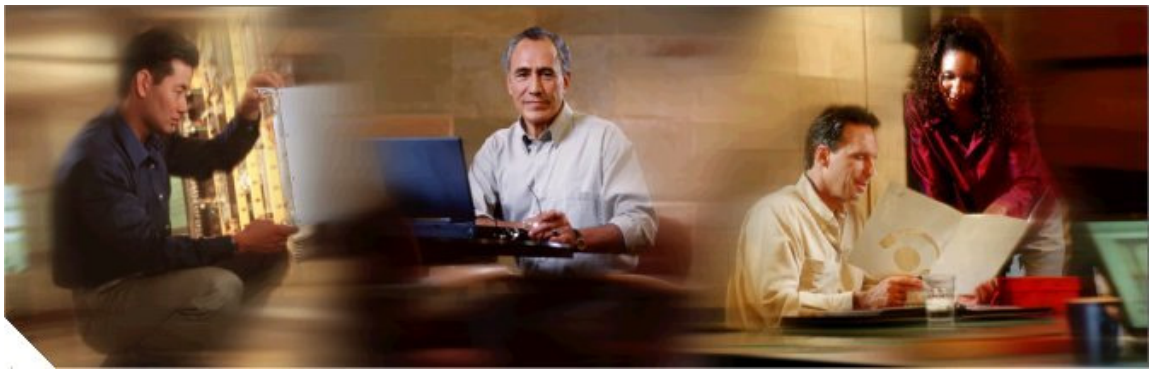

Voicemail

This chapter discusses the considerations for deploying voice messaging system in an H/M-UCS environment

For the multi tenant H/M-UCS program the Mereon Unified Messaging platform from IP Unity is used to provide voice mail services. Only the voice mail functions of the IP Unity platform are currently supported.

In both single and multi tenant deployments Cisco Unity may be used to provide more advanced capabilities such as unified messaging and fax. In multitenant deployments, separate Cisco Unity systems are required for each tenant as the Cisco Unity product itself is only single tenant capable.

Integration and provisioning of Cisco Unity into the H/M-UCS 1.6 architecture has to be performed manually as BVSM cannot currently provision Cisco Unity or the associated dial plan.

Some of the relevant considerations when deploying Cisco Unity in H/M-UCS environment are documented below.

IP Unity Mereon Voicemail

The Mereon voice mail system from IP Unity has been selected for its capability to support multi-tenant operation within the architecture. It should be noted that specific details of the design of the IP Unity components themselves are not covered in this SRND as the design required should be arrived at after discussion of requirements between the systems integrator and IP Unity system engineers. The overall responsibility of the IP Unity design lies with the system integrator and not with Cisco.

The Vision OSS BVSM platform is integrated with IP Unity to allow provisioning via the BVSM GUI of voice mail functionality on a per-tenant basis. The BVSM platform uses the IP Unity API (CORBA/XML) to define business groups, provision pilot numbers, add/delete mailboxes (assigned against a unique “internal” number and an “extension” number), and assign class of service.

The interface between the H/M-UCS platform and the IP Unity voice mail system is SIP. The CallManager forwards the incoming calls (on busy, no answer or all calls if configured) via the PGW to the voice mail system over a SIP trunk. Once the caller leaves a message, the voice mail system notifies the PGW that a message was left for the user using a SIP NOTIFY message. PGW inter-works SIP and H.323 between IP Unity and CCM for Message Deposit, Retrieval and MWI (Message Waiting Indicator).

The inter working of IP Unity via SIP and gateway fronted DPNSS/QSIG PBX is specifically not supported in the H/M-UCS 1.6 release so it is not possible to provide a hosted voicemail service for TDM PBX users.

IP Unity Platform Description

The platforms consist of the Mereon 6000 and Mereon 3000 Media Servers as well as Mereon Application Servers.

The IP Unity Media Servers consist of the Mereon 6000 for carriers and the Mereon 3000 for enterprises and carriers with smaller subscriber bases. Both servers are based on a modular design that allows you to

add additional media or speech processing cards without any service interruption to the subscribers. The Mereon 6000 platform scales from as few as 100 ports to more than 10,000 ports in a single 14RU shelf. The Mereon 3000 scales from 100 ports to 400 ports.

The Mereon Application Server offers service providers a robust service execution environment that delivers scalability and reliability to applications deployed on it. The standard Java-based runtime engine is exclusively designed for event-driven applications that require extremely high transaction throughput, while also supporting a distributed and replicated event processing environment.

Figure 35 H/M-UCS with IP Unity Platform Overview

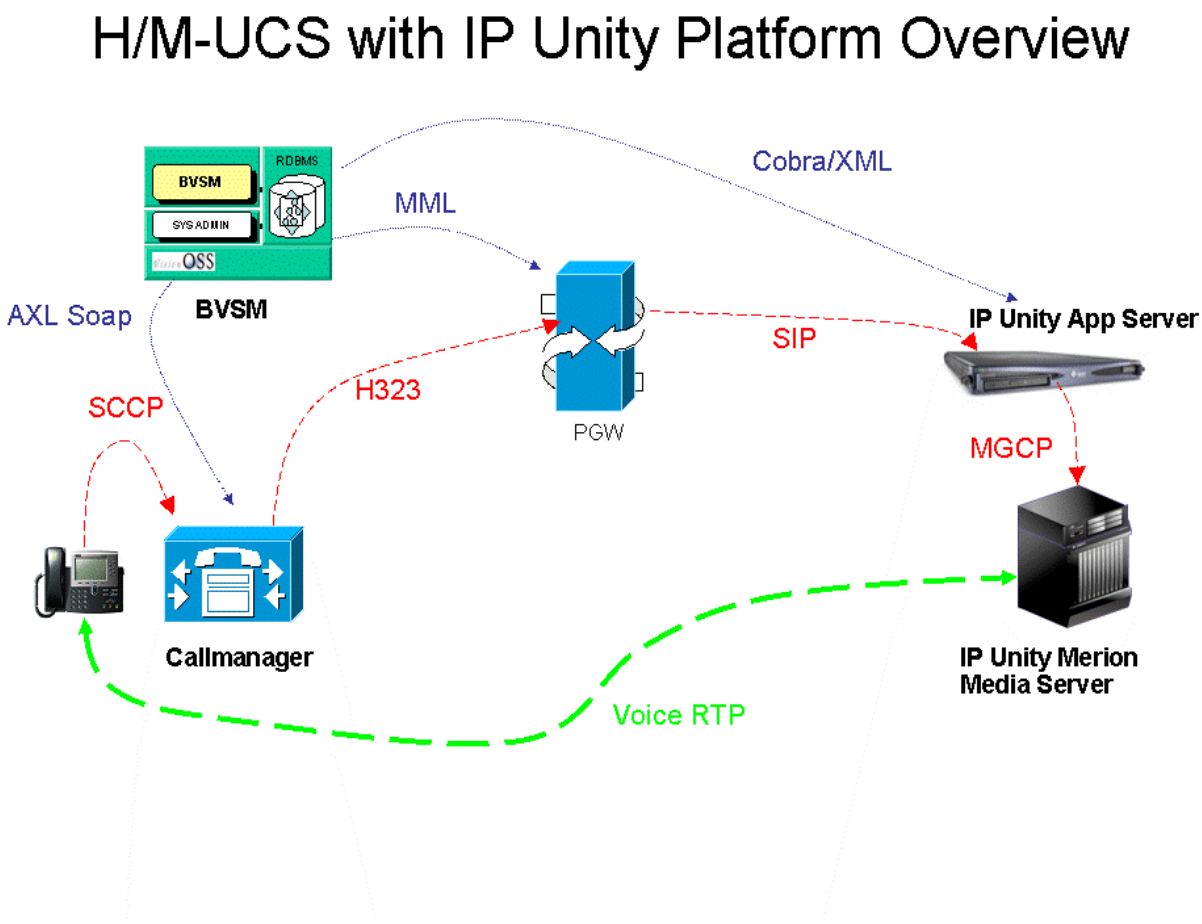

H/M-UCS IP Unity Integration Call Flows

The main IP Unity voicemail call flows supported are shown described here :

Voicemail Retrieval

A user dialling into IP unity to retrieve their voicemail. Remote access from PSTN is also similar to this flow. CCM is configured by BVSM with per-location voice mail Pilot numbers and profiles in order to initiate this call when an IP Phone messages button is pressed. PGW is configured by BVSM with per-location dial plan to determine the customer (based on CPID+RID information in the A-number) and to route the call to the customer's specific voicemail pilot on the IP Unity system.

Figure 36 Voicemail Retrieval

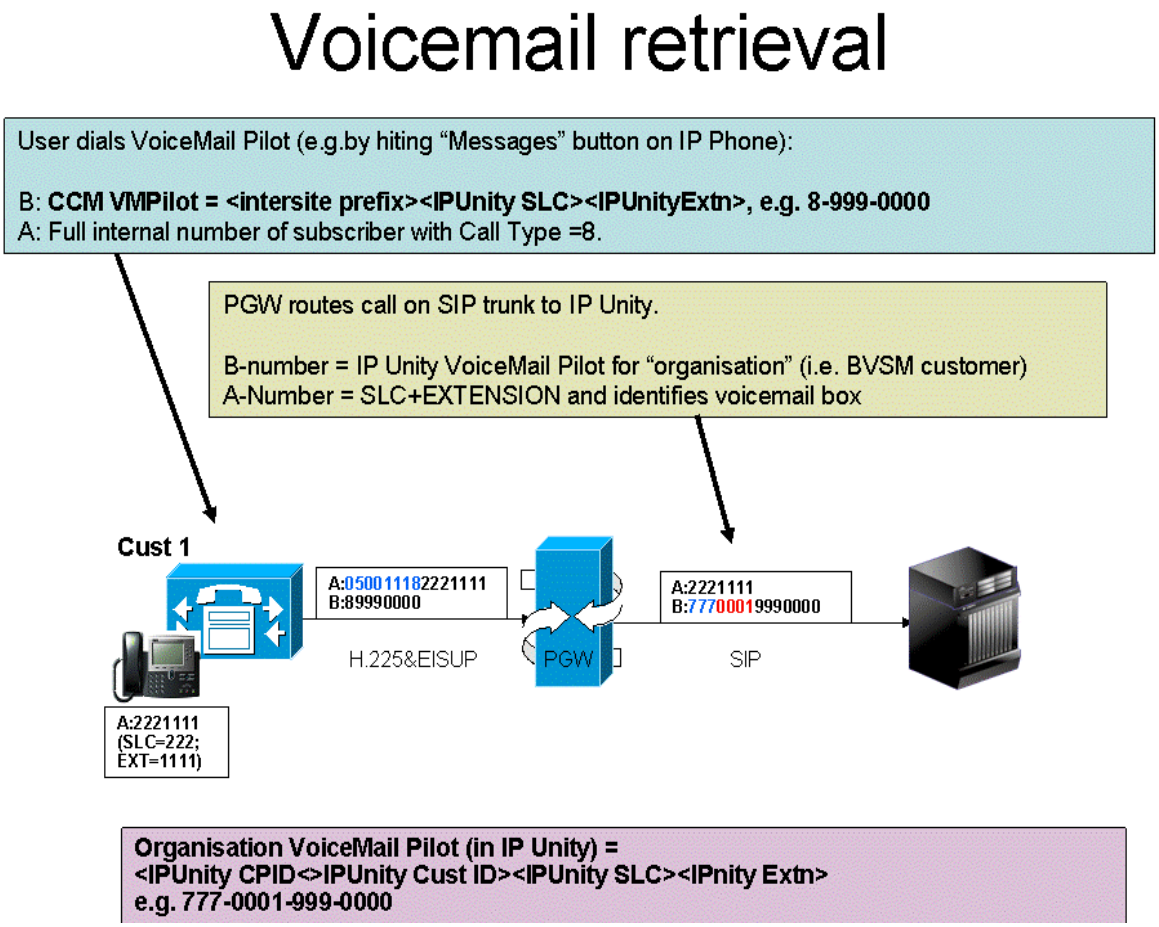

Leaving a Voicemail

A Call Forward to voicemail to leave a message for a user. This flow is similar to the previous call flow, but includes support for setting of redirected number such that voicemail is left in the correct voicemail box.

Figure 37 Call Forward to Voicemail

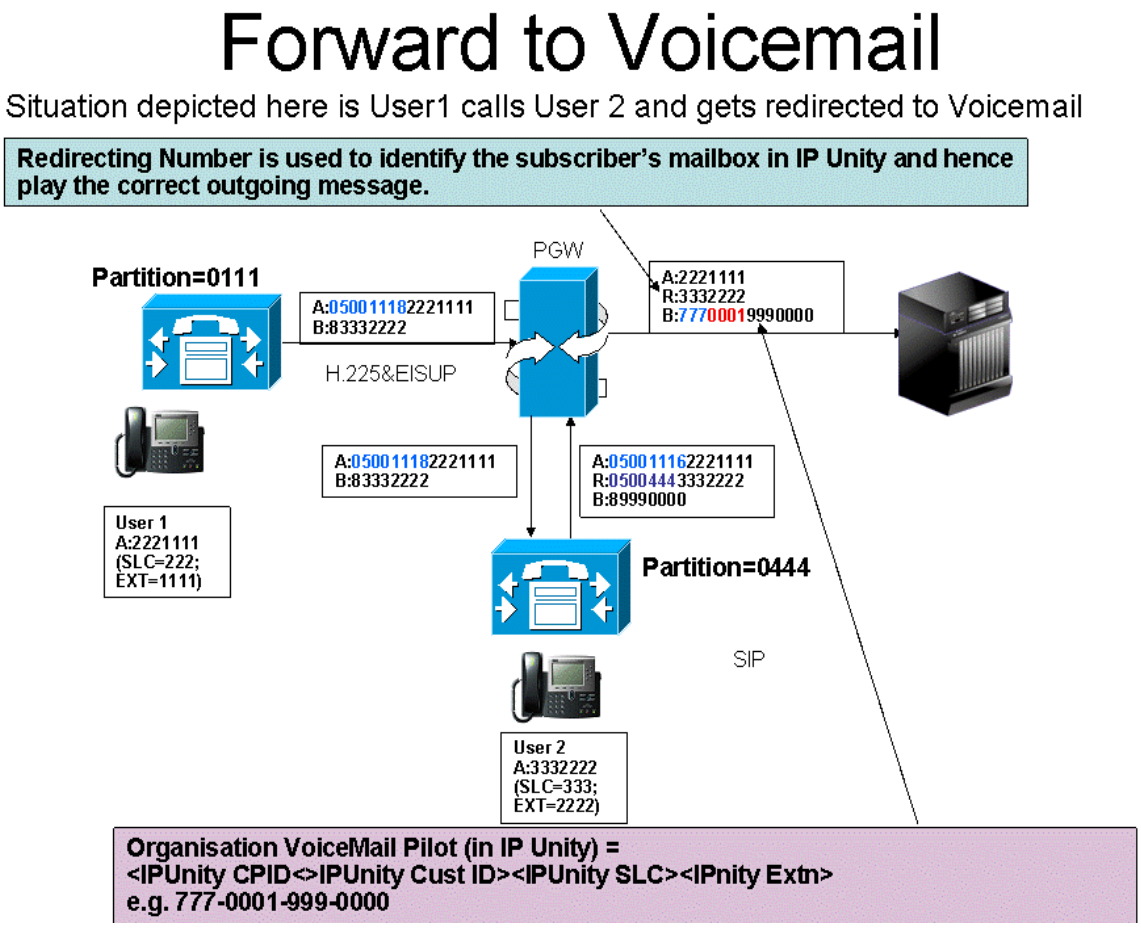

Changing state of Message Waiting Indicator (MWI)

MWI setting by IP Unity. This call flow requires PGW logic to analyze the A-number (which will be a per-customer pilot number) and B-number from IP unity and to manipulate numbers such that the call is routed to the MWI On or Off device on the CCM Cluster where the corresponding IP Phone resides.

Figure 38 MWI Call flow from IP Unity to CCM

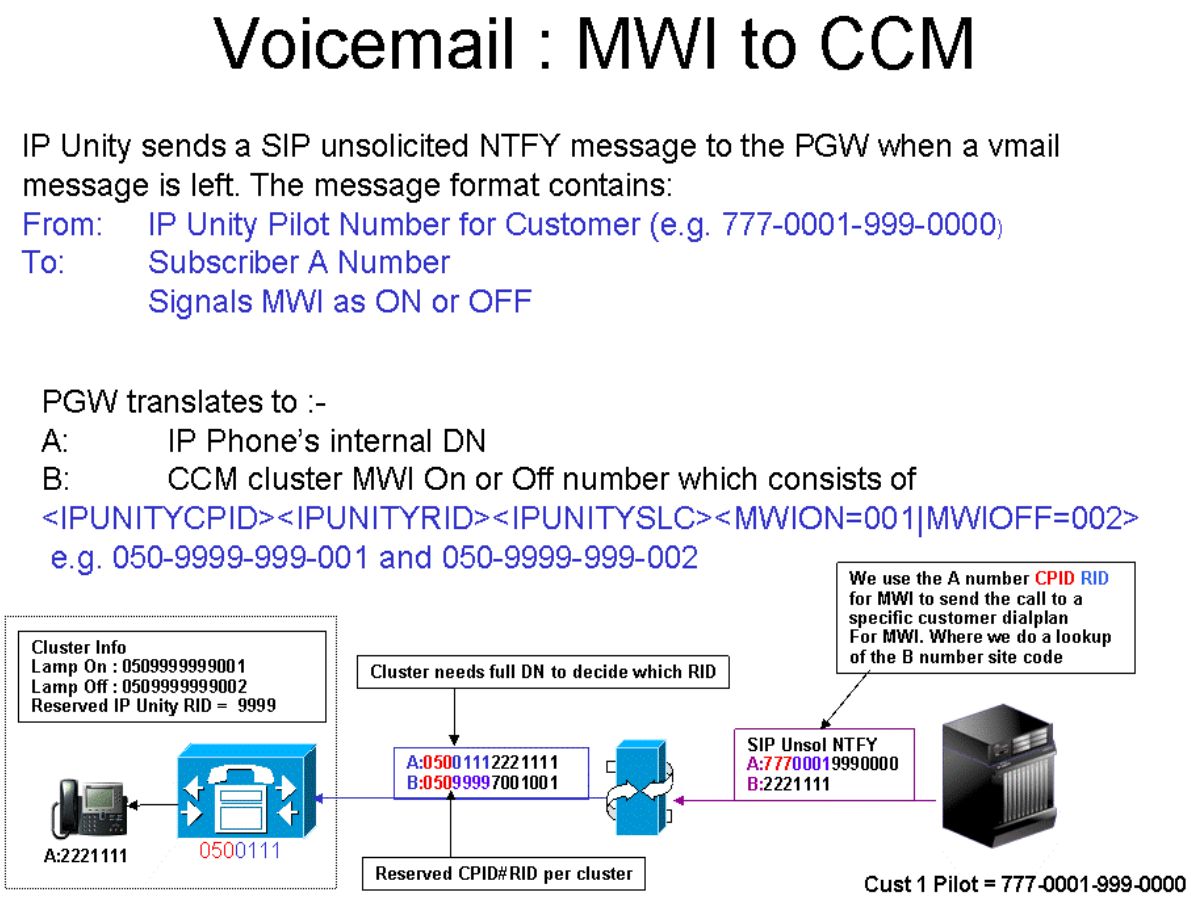

Cisco Unity

Integration of Cisco Unity into the H/M-UCS architecture has been tested as part of the 1.6 H/M-UCS release but provisioning of the Cisco Unity platform itself by BVSM is a roadmap item for a later H/M-UCS release. Cisco Unity is integrated with H/M-UCS in the 1.6 release by direct integration of Unity with CallManager using SCCP pretty much as it is in a pure CallManager environment.

Cisco Unity provides full unified messaging capabilities that allow both fax and e-mail to be integrated with voice messaging into a single inbox hosted either on Microsoft Exchange or Lotus Domino. Unified messaging brings major benefits to organizations, such as reduced administration costs, increased productivity, and increased efficiency once users are familiar with the new features.

The key considerations when using Cisco Unity in a H/M-UCS environment are listed below.

- Cisco Unity is not a multi tenant product. This means that the directory, administration and architecture of the product does not have the partitioning necessary to provide security between tenants and the features necessary for multi tenant deployments. As a result Cisco Unity may only be deployed in single tenant H/M-UCS deployments or in situations where it is acceptable to use one Cisco Unity system per tenant deployed on a multi tenant H/M-UCS architecture.
- Unity should be integrated into the architecture directly with CallManager using SCCP as described in the CallManager integration guide available at :

http://www.cisco.com/univercd/cc/td/doc/product/voice/c_unity/integuid/callma41/index.htm

- BVSM currently does not provision the Cisco Unity platform itself or the integration of the platform into the H/M-UCS dial plan. It is therefore necessary to manually provision both the

Cisco Unity platform itself and the CallManager to provide the integration required. These two tasks have to be performed manually.

- In a H-UCS integration the subscriber's voicemail box is configured on Unity as the site location code plus extension of the subscriber. A voicemail profile is configured per location on the CCM to strip the full internal number down to just the SLC plus Extension for the unity system.
- The Voicemail ports and pilot number sit, as far as the dial plan is concerned, within one location that a customer has, the "unity site". This means the voicemail system is callable throughout a customer's locations by calling the site location code of the "unity site" and the extension number allocated to the voicemail pilot. This technique also means that some extension DNs on the "Unity site" cannot be used for real phones as they are already used by the voicemail system.
- The subscriber configuration in the Cisco Unity system needs to include the full internal number or FINT of the subscriber in the !alternate MWI extension" field to ensure correct MWI operation.
- A feature guide for the manual Cisco Unity integration into H/M-UCS is under development and will be made available to system integrators.

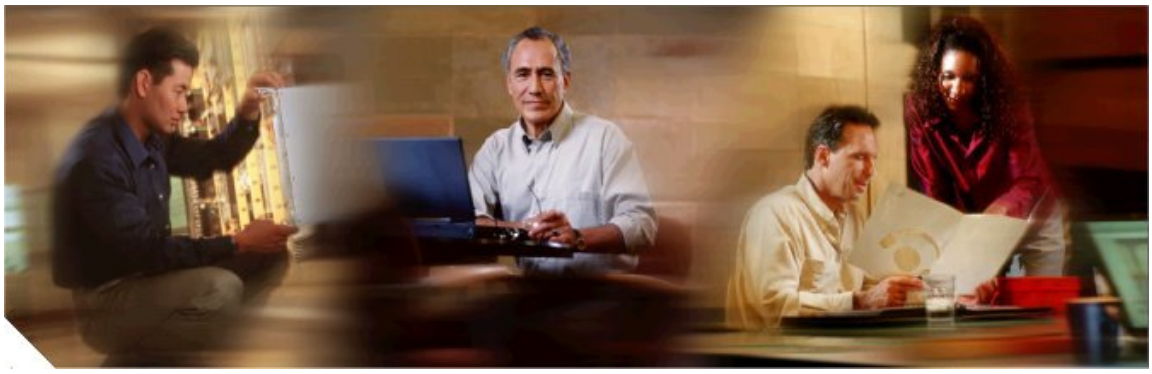

Conferencing, Tones and Announcements

Conferencing

This section provides conferencing guidelines for the H/M-UCS reference architecture.

Cisco recommends that you provide the following minimum conferencing resources for your users:

- Ad Hoc (up to 4 people) conferencing resources for at least 5% of the user base
- Meet-Me (up to 10 people) conferencing resources for at least 5% of the user base

With CallManager you can also use MRGs and MRGLs to separate resources based on geographical location, thereby conserving WAN bandwidth whenever possible. This configuration is not currently available through BVSM.

Each type of conferencing resource has different limits on the maximum number of participants in a single bridge. If multiple resources are defined in an MRGL, the first available resource is used. Two Cisco Unified CallManager service parameters control the maximum number of participants: Maximum Ad Hoc Conferences, which can be 3 to 64, and Maximum Meet Me Conference Unicast, which can be 1 to 128. The default setting for both is 4. The maximum number of participants for the conference resource overrides the service parameter settings.

To maximize the conference bridge size, set the service parameters to match the smallest bridge size of the resources in the MRGL. Otherwise, maximize bridge size by using only resources of homogenous characteristics, and set the service parameters to match.

CTI applications and the Drop Any Party feature do not support more than 16 participants. Although some Ad Hoc conference resources can support more than 16 participants, these applications display only the 16 most recent participants.

Tones and Announcements

In the H/M-UCS reference architecture there are two methods of providing tones and announcements.

- With the Cisco Unified CallManager Voice Media Streaming Application that implements the Annunciator Server. (This is the same application that provides MTPs, MoH and conference bridges)
- AS5400 Media Gateway

Cisco Unified CallManager Annunciator Server

The 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 IOS gateways and intercluster trunks. These devices require support for call progress tone (ringback tone).
- System messages.

During call failure conditions, the system plays a streaming message to the end user. These conditions are, a dialled number that the system cannot recognize, a call that is not routed due to a service disruption and 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. It also plays ringback tone to an Ad Hoc conference when a participant is added to the conference until the participant answers the call.

- Play a ringback tone to SIP client when the SIP client is being transferred from an IP Phone. Does not apply to H/M-UCS
- Play Vacant Code Announcement (VCA) “Your call cannot be completed as dialled…..”
- Play Isolated Code Announcement (VCA) “A Service disruption has prevented the completion of your call…”
- Play ringback and other tones for transfer on H323 intercluster trunks
- Play busy tone based on error generated by CTI for H323 related calls
- Play audible tones at the beginning and end of recording sessions

The following features are not part of H/M-UCS therefore are not used:

- Cisco Multilevel Precedence Preemption (MLPP) *This will not apply to H/M-UCS*
- Integration via SIP trunk *This will not apply to H/M-UCS*
- 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. Both call progress tones (busy, alerting, ringback) and DTMF tones are supported.

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; 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.

Cisco Unified CallManager 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 Unified 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. Multiple standalone servers can be added to support the required number of streams.

AS5400 based Tones and Announcements

There are a number of scenarios where an Announcement Server (AS) controlled by the PGW is required to play tones and announcements for inter-tenant calls, called to ported in-numbers, incoming call from PSTN etc. This may not be a requirement for all deployments and customers. In cases where this is required a centralized AS5400 can be used to perform this function. It should be noted that the tones and announcements will need to be customized to the PSTN and is therefore the responsibility of the system integrator or service provider to configure this manually.

The following describes how the PGW (PSTN Interconnect) and AS interact with one another in order to play an appropriate announcement to the calling party.

When the PGW receives a release message for dial plans that route out to the PSTN or loopback routes, it examines the cause and location value within that message. If the cause and location value received match a specific set of values configured into the PGW, then the PGW will alter the CdPN and re-route the call to an announcement server, rather than propagating the release message back to the calling party. In order that the correct announcement is played related to the cause / location value, each announcement stored on the AS has a unique called party number. The PGW is configured to change the CdPN, so that it maps to the correct announcement when the call is re-routed to the AS, based upon the cause / location value received.

The AS is configured with multiple dial peers, on its PRI interface. Each announcement or tone shall have its own dial peer. Whenever a dial peer is matched, the GW will invoke the TCL script associated with that dial peer, which causes the appropriate tone and / or announcement to be played.

When the call reaches the alerting state the AS establishes the RTP/RTCP bearers directly between itself and the calling endpoint. This requires the PGW to overwrite the PSTN GW IP address with that of the AS. No ring tone is generated. Instead the AS plays the announcement relating to the received CdPN towards the calling endpoint. The announcement and / or tones are repeated as specified in the figure above. Once the announcement has been played out the AS will, after a defined time period, clear the call back towards the calling party if the calling user has not yet initiated call clearing procedures.

The figure below shows the high-level the call flow sequence for a call being re-routed to an AS. This shows the AS being used for an IP Phone call, but the same principle applies to a PBX call.

Figure 39 Announcement Server Call Flows

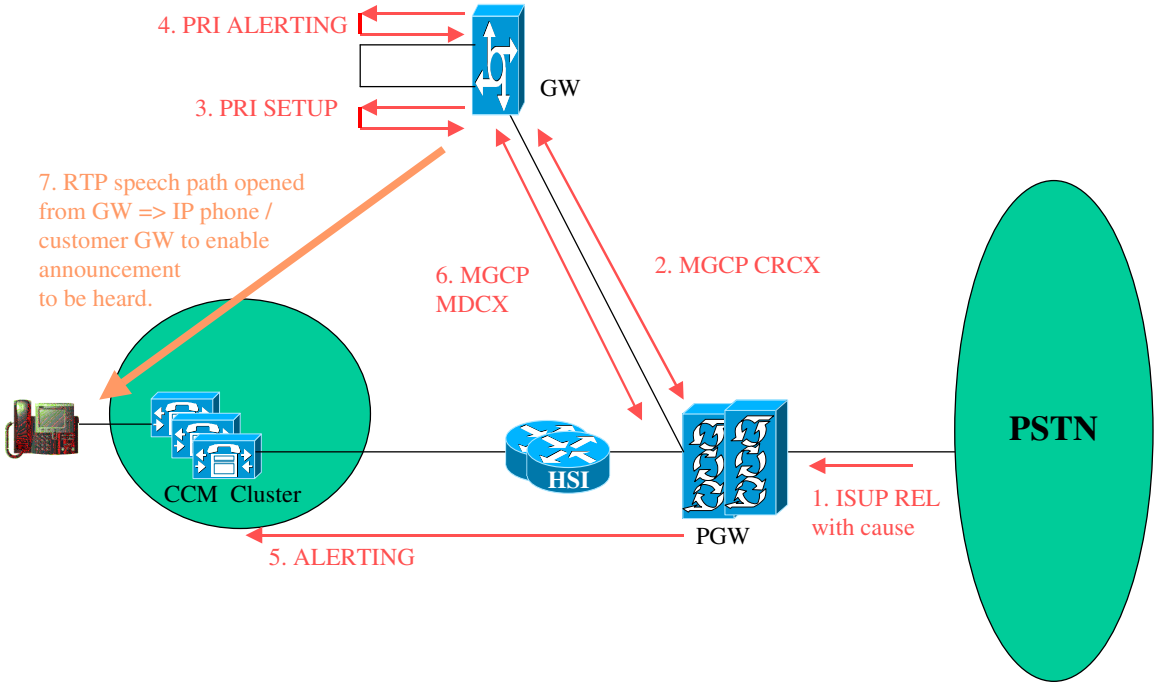

PGW Re-routing

The PGW is also configured to re-route calls to an AS based on the cause values received. This enables announcements / tones to be played for PBXs or IP Phones destined calls that did not complete. The PGW modifies the CdPN accordingly to include the AS prefix and BTNR number based upon the received cause value. The PGW (PSTN Interconnect) routes this call to it's local AS, or to the remote AS as a secondary choice, based upon the CdPN, as described above.

AS Addressing Format

The addressing format used by the PGW to re-route to the ASs is as follows:

- The number will be 6 digits in length. From left to right, the first 3 digits will act as the Announcement Server ID prefix. Digits 4-6 define the announcement/Tone ID, describing which announcement/tones should be generated.
- Announcement Server ID prefix 399 - is used to identify the AS.

Figure 40 Announcement Server Address Format

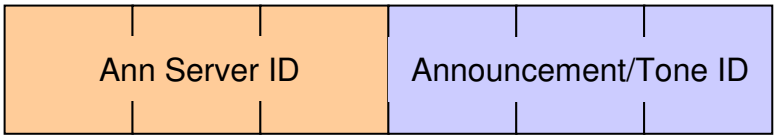

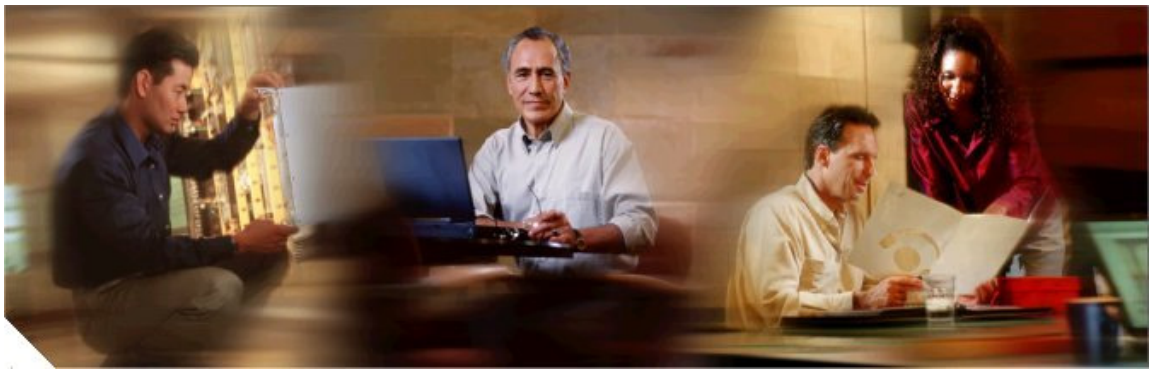

Music on Hold

Multi tenant music on hold (MOH) is the ability to provide an audio stream to callers while they are on hold with one of the tenants of the service provider.

MOH Audio Source and Server Selection

The H/M-UCS reference architecture can deliver multiple audio streams for use in MOH. When considering MOH deployments, it is important to understand the factors that determine which MOH stream a device being placed on hold will receive. In a multi tenant environment, this becomes especially important as groups of phones and PSTN callers are expected to receive tenant-specific MOH content.

The MOH stream that a device receives is determined by the following two factors:

- The user- or network-hold audio source that has been configured for the phone or device that places a call on hold.
- The media resource group list (MRGL) of the phone or device that has been placed on hold. An MRGL is a method that allows an administrator to allocate media resources to particular devices by providing a prioritized grouping of media resource groups (MRGs), which are logical groupings of media resources.

Figure 41 shows a typical hold scenario and illustrates how the MOH stream is determined. The following events occur in this example:

1. Phone A and phone B are connected in a call, and phone B places the call on hold.
2. Cisco CallManager examines the available audio sources and determines that there are four audio sources shared by the MOH servers within the cluster.
3. Cisco CallManager determines the audio source and the MOH server for the audio source that is being played to the call on hold. In this example, Cisco CallManager looks at the user-hold audio source of the device that placed the call on hold (phone B) and determines that audio-source 2 is the audio source that will be streamed to the phone being placed on hold (phone A).
4. Next, Cisco CallManager looks at the MRGL of the phone that has been placed on hold (phone A) and determines that MRGL A is the media resource group list that the stream will come from. This MRGL has MRG A assigned to it. MOH server A is the first resource in MRG A. With this information, Cisco CallManager now knows that audio-source 2 should be sent from MOH server A to the phone that has been placed on hold (phone A).

For more information about MOH, refer to the Cisco CallManager documentation.

A new feature in the H/M-UCS 1.6 architecture is the ability to create and manipulate Callmanager media groups in a more selective way via the BVSM provisioning application (feature introduced in BVSM version 3.1.6). BVSM now supports the creation of a number of CallManager media services such as Music on Hold servers, Conference bridges and Transcoders from both the GUI and the bulk loader spreadsheets. BVSM can also create CCM Media resource groups and place these media services into them. CCM Media Resource Lists can also be provisioned via BVSM and then be available to be selected as part of the manage location screen. The result is that that specific media resources may be allocated for

use on a per location basis via BVSM in much the same manner as they can when provisioning Cisco CallManager directly.

Figure 41 MOH Stream Selection: Audio Hold Source and MRGL

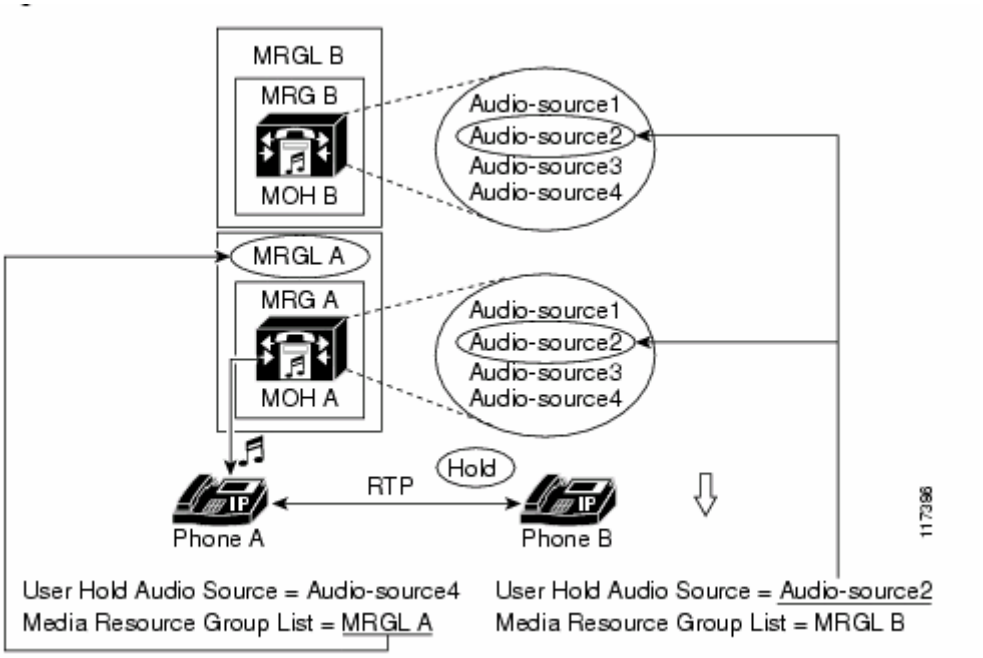

MOH Deployment Models

The following deployment models are discussed in this section:

- Tenant-Shared MOH Deployment
- Tenant-Dedicated External Media Server Deployment

Tenant-Shared MOH Deployment

Once the audio source and server selection mechanisms are understood, it is easy to see that, in a multi tenant environment, MOH servers can deliver MOH streams to groups of tenant phones or devices based on the MRGLs that have been configured. In the example in Figure 42, Tenants A and B are assigned to MRGL A. MRGL A contains MRG A, which contains MOH server A as a resource. Tenants C, D, E, and F are assigned to MRGL B, MRG B, and MOH server B resources. All tenants use the same 51 audio sources, which are configured and shared among all the MOH servers. Each tenant's phones can be grouped and pointed to an appropriate generic shared audio source. In this example, all of tenant A's phones have been configured with user/network-hold audio source Generic1, while tenant B's phones have been configured in two groups. One group is configured with hold audio source Generic 51 and one group is configured for Generic 2 as the hold audio source. Tenants C, D, E, and F are configured similarly. Meanwhile, all gateways are assigned to MRGL GW, which in turn contains MRG GW. MRG GW contains one or more of the configured MOH servers (in this case MOH server B).

Figure 42 Tenant-Shared MOH Server Deployment

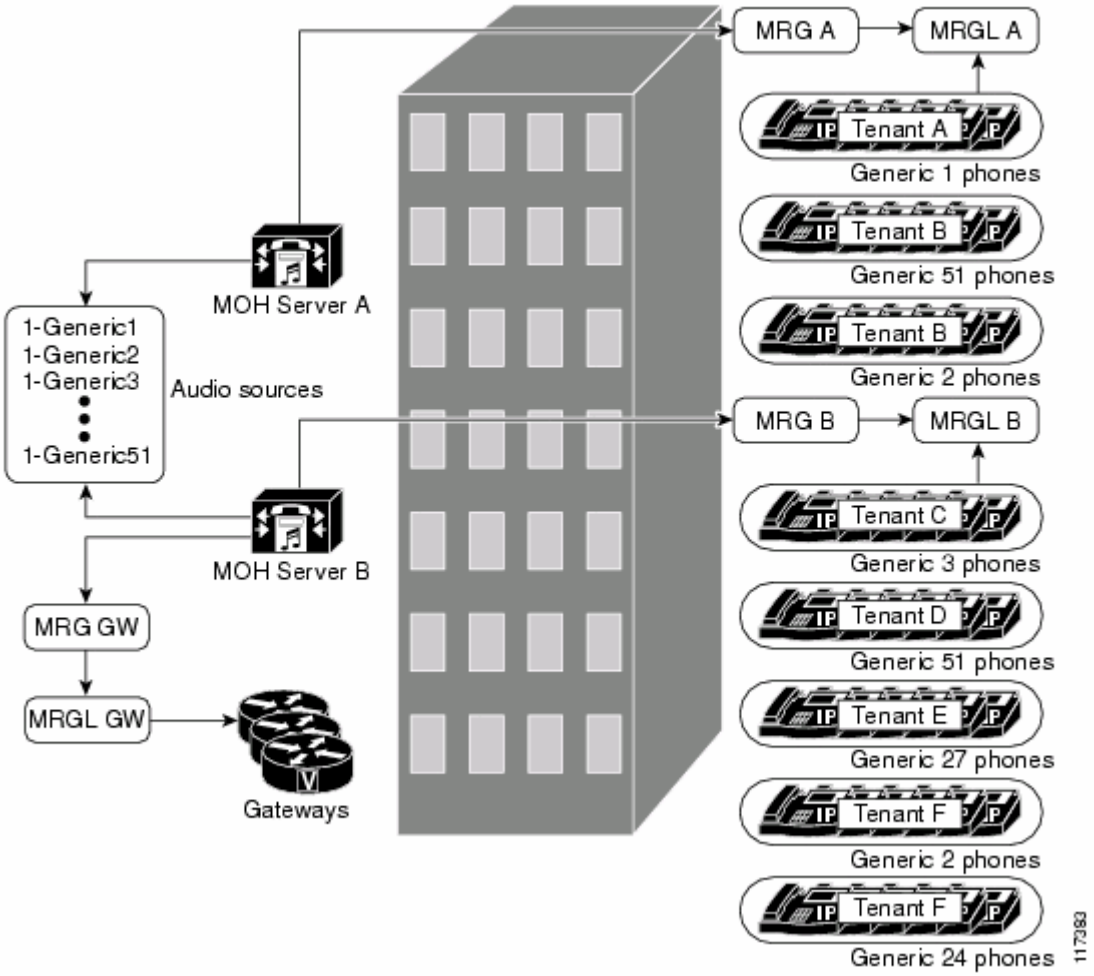

One important thing to consider with the tenant-shared MOH model is that the audio sources are shared among all the MOH servers within the cluster, so irrespective of the number of MOH servers in the cluster, there will only be a maximum of 51 unique MOH streams among all the servers and tenants in the cluster. Because this is the case, providing dedicated or customized MOH streams for individual tenants can be problematic. If a single dedicated stream is needed for each tenant, then up to 50 tenants can be handled with this deployment model. If more than 50 tenants are present, a dedicated MOH stream per tenant is not possible. Some tenants will have to share MOH audio streams.

As noted, up to 20 MOH servers can be configured for this deployment model, so a large number of tenants can be handled if dedicated MOH audio source are not required. However, MOH server capacity must be considered when determining the number of tenants that can be supported. MOH server capacity depends on the following factors:

- Type of server deployment: co resident, standalone, or both.
- MOH transport mechanism: unicast, multicast, or both.
- % of calls on hold at a given time.

For example, for 12 co resident MOH servers and 8 standalone MOH servers on high-end platforms, there is a limit of 2240 streams ([12 co resident servers * 20 MOH sessions] + [8 standalone servers * 250 MOH sessions] = 2240). Depending on the transport mechanisms used and the desired Erlang ratio, this MOH stream capacity will help determine the number of tenants or number of phones supported. Regardless of the number of tenants, care should be taken to evenly distribute tenant load among all configured MOH servers using MRGLs.

Tenant-Dedicated External Media Server Deployment

Another way to deploy MOH in a multi tenant environment is to use an external media server for each tenant in the network. As illustrated in Figure 43, a single MOH server is all that is required for this model. However, each tenant must have an external media server of some sort that can stream multicast audio to the network.

In the example in Figure 43, each tenant has a dedicated media server. Each tenant's phones can be grouped and configured to select the appropriate audio source, which will come from the local media server rather than from the MOH server. This audio source will use identical multicast group addresses. Thus, the phones that are configured with a user/network-hold audio source that is addressed with multicast group 239.1.1.1 on the MOH server will receive the audio source that is locally streamed to that same group address by the external media server. In this example, tenant C's phones have been grouped so that some of the phones are configured with an MOH audio source that multicasts to 239.1.1.37 and some of the phones are configured with an MOH audio source that multicasts to 239.1.1.201. In both cases, these audio streams come from tenant C's local media server. Note that this type of deployment also requires tenant-dedicated gateways. As shown in Figure 43, each tenant also requires a local gateway for PSTN access. Dedicated gateways are required for each tenant because the gateway must reside on the subnet with the local media server in order for the gateway to have access to the locally significant customized audio streams coming from that media server.

Figure 43 Tenant-Dedicated External Media Server Deployment

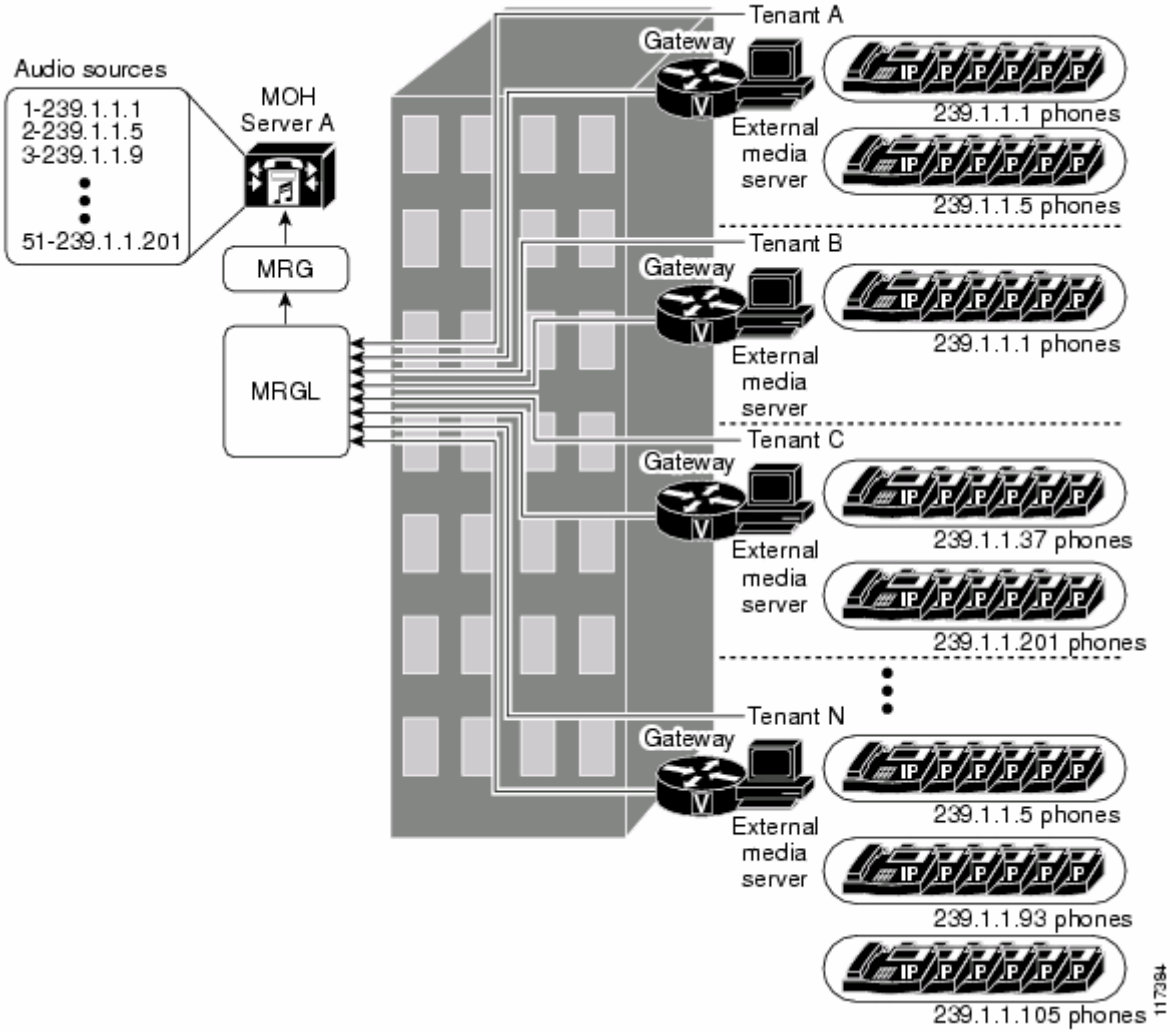

To make sure that MOH audio sources do not travel across tenant subnets, potentially causing a tenant's phone or gateway to receive a generic cluster-wide audio stream rather than a locally significant stream, tenant-dedicated media server deployments also require one of the following settings:

- A setting of 0 in the Max Hops field in the MOH Server Configuration window for each audio source on the MOH servers to keep the audio source in its own subnet.
- An access list (ACL) on the local MOH server subnet to disallow multicast streams from travelling off the local subnet.

A tenant-dedicated media server deployment requires significantly more resources than the tenant-shared deployment model because a media server and gateway are required for each tenant. For this reason, management and maintenance of these resources becomes decentralized. However, this model provides the largest number of dedicated, tenant-specific audio streams per tenant, as each tenant can have up to 51 unique dedicated MOH audio streams. Furthermore, calls coming from the PSTN receive tenant-specific MOH audio and only a single co resident MOH server is required for the entire cluster, as stream capacity is now dictated by the local media servers. This deployment model supports only multicast MOH. Unicast-only phones (for example, Cisco Wireless IP Phone 7920) can receive unicast streams coming from the MOH server, but these streams are generic cluster-wide audio streams. Note that if there are many unicast-only devices on the network, more than one MOH server may be required to meet audio stream capacity needs.

One major caveat with this deployment model is that any inter tenant calls that are placed on hold will receive local MOH rather than the MOH of the tenant that places the call on hold. This is problematic if the local MOH streams contain identifying information for a particular tenant. For example, if a tenant A phone calls a tenant B phone and the tenant B phone places the call on hold, the tenant A phone will receive the MOH stream specified by the user-hold audio source that was configured for the tenant B phone, but tenant A will receive this MOH stream locally. So instead of hearing a MOH stream like: "Thank you for calling Company B. Please hold...", tenant A's phone will receive a stream like: "Thank you for calling Company A. Please hold...". There is no workaround for this issue unless this inter tenant call is routed through the PSTN via tenant-dedicated gateways.

Multi-tenant MOH Deployment Comparisons

The table below compares the two deployment models outlined in the previous sections. Note that in some cases these two deployment models may be combined in order to handle specific multi tenant environments.

Table 13 MOH Deployment Model Comparison

	Tenant-Shared MOH Server	Tenant-Dedicated External Media Server
Number of Cisco CallManager MOH Servers	Up to 20	1
Audio Stream Selection	51 shared streams[1] (One of these streams must be dedicated for PSTN callers.)	51 unique streams for each tenant
Unicast and Multicast Support	Unicast and multicast	Multicast only
Gateways	Shared gateways	Dedicated gateways required
PSTN Caller Experience	PSTN caller receives generic, or shared, MOH	PSTN caller receives tenant-specific MOH

Note

[1] Streams can be dedicated, up to one for each of 50 tenants.

MOH Design and Configuration Best Practices

Proper MOH design practices should be followed when deploying MOH in multi tenant environments. Redundancy, server capacity, resource provisioning, transport mechanisms, multicast addressing, and codec selection should all be considered when deploying MOH.

For more information about MOH, MOH configuration, and design best practices refer to the MOH chapter in the latest Cisco Unified Communications SRND, located at the Solution Reference Network Design web site at http://www.cisco.com/go/srnd.

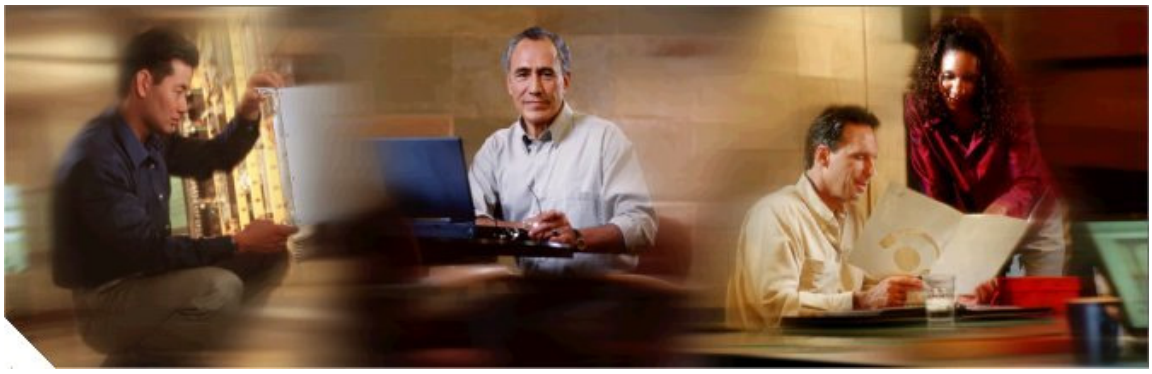

Phone Services and Directory Integration

In a standard Callmanager deployment, phone services and directory are managed and hosted by the Cisco CallManager.

This chapter describes the differences in the phone services and directory integration that are available in a multi tenant H/M-UCS deployment where the BVSM from Vision OSS takes over hosting and managing these functions and provides the separation necessary for multi tenant operation.

Multi-tenant Phone Services

To allow service providers to offer in a multi-tenant environment the same type of value and phone applications that are available with standard Cisco CallManager, the Business Voice Services Manager (BVSM) from Vision OSS has been developed to host some of the XML phone services and directory functions that would normally be hosted on CallManager itself.

In a multi-tenant environment, all phones are usually configured to get the XML service information screen and the directory screen from the BVSM rather than from the Cisco CallManager system.

The following phone services are managed by the BVSM:

Directories Key

In addition to missed, received, and placed calls, the Directories URL can be set up to retrieve a user (corporate) directory and Personal Directory that are supported by a web application running on the BVSM platform. The BVSM-based directory service can be segmented by tenant to comply with the multi-tenant nature of the services and can search for users on multiple Cisco CallManager clusters.

Services Key

The Services URL retrieves a list of user-subscribed phone services that has been configured in the BVSM. The services are validated and are specific to the user and tenant who subscribe to a particular phone.

Extension Mobility

BVSM also implements a version of the extension mobility feature for environments with multiple tenants. This feature ensures that a user from one tenant cannot roam to another tenant's phone.

Figure 44 Phone Services Provisioning

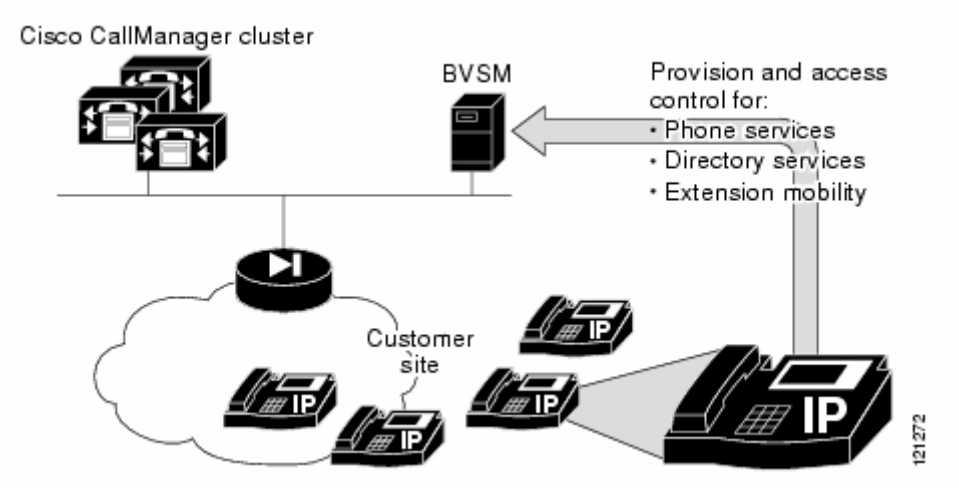

BVSM also supports the addition of XML services for which BVSM acts as the access portal, much like the Cisco CallManager support in an enterprise environment. BVSM displays the menu of services available configured for a particular phone.

Multi-tenant Directory Integration

Multi-tenant directory integration is facilitated by the BVSM platform, which includes a Phone-based XML directory service (the same functionality as provided by CallManager on single enterprise Cisco CallManager-based systems) to provide corporate directory on the phone from the BVSM platform LDAP directory. This directory capability replaces the use of internal Cisco CallManager LDAP directories and allows directory services to be provided in multi CCM cluster and multi tenant environments.

In a multi tenant environment, whether multiple internal tenants in a large enterprise or multiple separate enterprises in a shared, hosted environment BVSM enforces a level of security that ensures that uses see only their own partitioned area of the telephone directory.

Some operators may also wish to offer value added directory services such as offering companies access to an external directory (e.g. published telephone numbers for companies within an Internet city or business park or access to white and yellow page service). These services, which fall into the value-added category, may be offered by a subscription or on demand. In these scenarios, BVSM can provide the subscription access mechanism by passing event data records (EDRs) to a billing system.

BVSM stores telephone directory information within the database on the resilient multi server VOSS Manager platform architecture, which fails over automatically in the event of a single server failure.

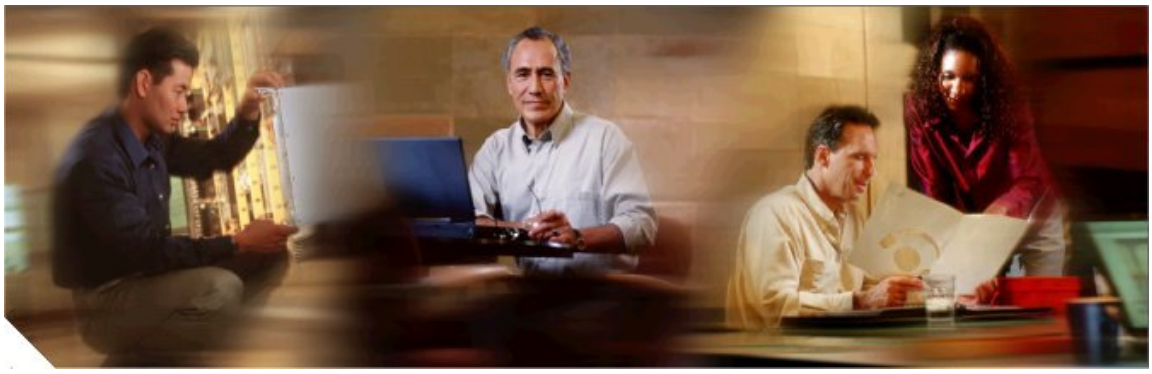

Wireless IP Phones

This chapter contrasts the use of the Cisco Wireless IP Phone 7920 in an H/M-UCS environment with its use in a standard enterprise deployment. For more information about wireless IP telephony, refer to the latest 7920 design guide, the Wireless LAN design guide and the CallManager 4.X SRND available on CCO at the following links :

http://www.cisco.com/en/US/products/hw/phones/ps379/products_implementation_design_guide_book09186a00802a029a.html

http://www.cisco.com/application/pdf/en/us/guest/netsol/ns178/c649/ccmigration_09186a00800d67eb.pdf

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

Wireless telephony presents some limitations with regards to roaming, radio design, and maximum number of concurrent calls and users. Deployment of wireless IP phones should only be considered if these limitations are understood.

This chapter contains the following sections:

- RF Design
- QOS
- Security
- Network Sizing
- Roaming
- VLAN Usage
- BVSM DHCP Considerations
- Admission Control Considerations

RF design

One of the most important design challenges when deploying wireless IP phones is to create a proper radio frequency (RF) design. Important issues to consider are coverage, signal strength, radio overlap, and possible conflicts between 802.11b and 802.11g deployments. Each of these is addressed briefly in the sections that follow.

Coverage

A major difference between wireless voice networks and wireless data-only networks is the type of radio coverage that must be designed. Voice clients are likely to be used in different types of locations than data clients. Voice clients are frequently found in elevators, parking lots, restrooms, and restaurants. They are

more likely to cluster together in waiting areas, stores, restaurants, and cafes. And, finally, voice clients are more likely to roam while they are in use.

For a H/M-UCS environment, two coverage models can be considered: full coverage, in which a voice client operates throughout a facility; and hot-spot coverage, in which a voice client operates only in clearly indicated areas throughout a facility.

Minimum RSSI level

To ensure proper roaming and voice quality, a minimum radio Received Signal Strength Indication (RSSI) of 20 should be available throughout the coverage area.

Radio Overlap

Radio coverage should provide sufficient overlap to allow roaming of the voice client. Roaming is needed when a client is moving, but also when the access point (AP) with which the client device is associated is too busy to support another voice call, as indicated by a high QOS Basis Service Set (QBSS) level that forces the client to re-associate to a less-busy AP.

802.11b Impact on 802.11g

With the introduction of 802.11g, many organizations are deploying 802.11g as their wireless access technology. Be aware that the presence of active 802.11b clients (such as the Cisco Wireless IP Phone 7920) forces an 802.11g network to use a protection mechanism which greatly reduces the throughput of the wireless network. Customers that require both high-speed wireless access and wireless phones should consider using separate access points for 802.11g and for 802.11b clients, although this design is nearly impossible to achieve because 802.11 has just three non-overlapping radio channels. This limitation makes it difficult to create two separate radio networks with sufficient overlap to allow roaming. Another solution is to use 802.11a for high-speed data access. Although 802.11a has limitations with regards to range, combined 802.11a/b/g adapters (like the Cisco Aironet AIR-CB21AG) will allow data clients to access 802.11a networks when they are available and fall back to 802.11g/b when they are not.

QOS

QOS requirements for wireless voice clients in a H/M-UCS environment are comparable to those of a normal enterprise environment. These are documented in the 7920 deployment guide at the link listed above.

Security

It is unlikely that a H/M-UCS based wireless voice network will be deployed without a form of security, encryption, or both. The current Cisco Wireless IP Phone 7920 firmware supports Wired Equivalent Privacy (WEP) and Cisco Light Extensible Authentication Protocol (LEAP).

When Cisco LEAP is used and a phone is roaming, the phone must reissue its Cisco LEAP authorization request. If the authenticating RADIUS server—such as a Cisco Access Control Server (ACS)—is using an external database to process the authorization requests, the delay of the Cisco LEAP authorization can become unpredictable and can interfere with seamless roaming. It is recommended that you use an account that can be authorized locally for voice clients.

Network Sizing

The following set of design rules is recommended:

- The maximum number of clients per access point is 15 to 25.
- The maximum number of simultaneous voice calls per access point is 7 when using G.711 or 8 when using G.729.
- The maximum number of access points per Layer 2 roaming domain is 30.
- The maximum number of voice clients per Layer 2 roaming domain is 450 to 600.

Roaming

For wireless voice clients it is important to distinguish between different types of roaming:

- Layer 2 roaming, in which a client roams to an access point in the same VLAN without changing its IP address.
- Layer 3 roaming, in which a client roams to an access point in a different VLAN.

The Cisco Wireless IP Phone 7920 supports roaming without interruption of an active call at Layer 2 only unless the Cisco Catalyst 6000 WLSM is deployed.

Without WLSM, should a wireless phone perform a Layer 3 roam, it must also perform a warm start to renew its IP address. This limitation means that if a user is on a call, a Layer 3 roam results in the call being dropped and the phone must reregister with CallManager with its new IP address before service is resumed.

When devices roam at Layer 3, they move from one AP to another AP and cross a subnet boundary. With the release of the new Cisco Catalyst 6500 Series Wireless LAN Services Module (WLSM), the Cisco 7920 Wireless IP Phone now supports call survivable Layer 3 roaming while using Static WEP. This is achieved by allowing the phone to maintain its IP address even when roaming between access points in different subnets.

The addition of Cisco Centralized Key Management (Cisco CKM) enables the Cisco 7920 phone to achieve full Layer 3 mobility while using LEAP.

VLAN Usage

Cisco wireless APs support multiple VLANs on the same radio interface. Each VLAN is identified by its own service set identifier (SSID) and encryption scheme. In implementing a wireless network for voice clients there are two approaches to assigning SSIDs:

- A single SSID can be used for voice traffic throughout the entire campus.
- An SSID per customer can be used for that customer's voice traffic (and can optionally be extended with a public SSID when roaming to shared spaces).

AP firmware supports up to 16 VLANs per AP; the Cisco Wireless IP Phone 7920 supports configuration of up to four SSIDs (enabling roaming in up to four separate wireless VLANs).

BVSM DHCP Considerations

When the Vision OSS BVSM is being used to host the DHCP server(s) for the phone VLANs there are several additional considerations when deploying 7920 wireless IP phones.

The first consideration is that, unlike a typical desk phone voice VLAN, it is more unusual for a wireless VLAN to be dedicated to only phones. If the wireless VLAN is being used for devices other than phones then it is not possible to use BVSM as the DHCP server. The best design technique to deploy to avoid this issue is to use separate wireless SSIDs and VLANs for phones and other wireless data devices.

Another consideration is a problem caused by the behaviour of wireless IP phone users. Wireless phones are much more likely to legitimately move between locations than fixed desk phones. Unfortunately the BVSM DHCP server has a function designed to specifically prevent phones from moving between locations without intervention from the administrator of the system. This function is designed to avoid the problems associated with moving devices between locations such as incorrect call admission control and incorrect local gateway selection. Until special handling of DHCP for wireless VLANS is implemented within the BVSM it is therefore not possible to use the DHCP functions within the BVSM for the wireless VLANs and still allow the device mobility between sites that 7920 deployments typically need.

Admission Control Considerations

The H/M-UCS architecture makes use of the "admission control by locations" CallManager feature to provide call admission control between the locations configured on the system. This Cisco CallManager feature does not allow for the concept of device mobility between sites. The location of a device has to be manually configured (or configured via BVSM in the case of H/M-UCS) on CallManager.

Wireless IP phone users typically want the ability of moving devices between company locations freely in much the same way that a mobile phone user would expect their phone to work if it was within coverage. Unfortunately this need for device mobility therefore breaks the admission control by location mechanisms.

CallManager 4.2 (which is supported in the H/M-UCS version 1.6 described by this SRND) introduces the capability of dynamically allocating the location according to the IP subnet that the phone is found in. This feature is not supported currently in the H/M-UCS 1.6 architecture, but when it is will allow wireless IP phones to move between locations without breaking the admission control mechanisms. Until then it is effectively not possible to do admission control for wireless IP phones if there is a requirement for the phones to move freely between locations.

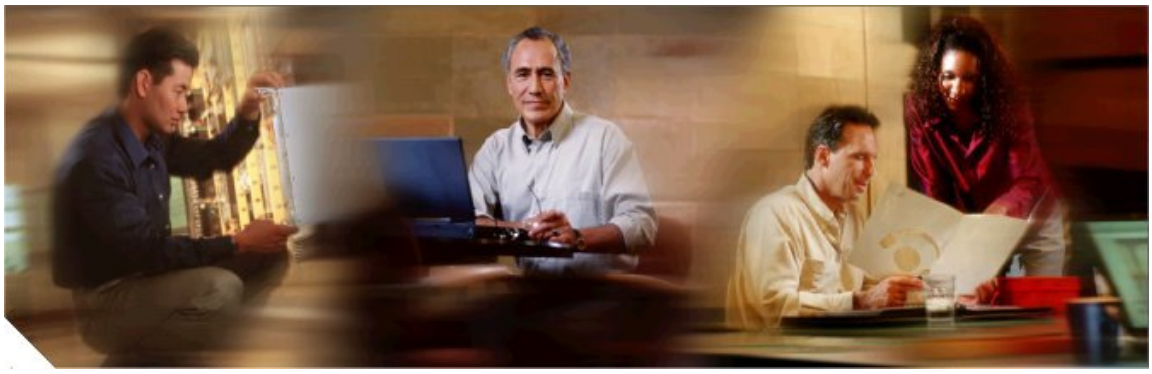

Attendant Console Management

Part of the overall Cisco H/M-UCS solution offering is the ability to integrate with attendant console systems in both large single enterprise and multi tenant environments.

This chapter describes the design considerations for incorporating attendant consoles within a Cisco H/M-UCS environment.

Cisco H/M-UCS Attendant Console Support

The H/M-UCS architecture has been tested to allow incorporation of three different types of attendant console. Each attendant console type has advantages and disadvantages and it is certainly not appropriate to use some of the alternatives in all H/M-UCS deployment situations. The three types supported are :

- Cisco Attendant Console
- ARC Console from ARC Solutions
- NOW Attendant Console from Netwise

Each attendant Console type has limitations and can only be deployed in certain operating environments when used in a H/M-UCS architecture. These are covered in the table below.

Table 14 Attendant Console Types and Limitations

Feature	Cisco Attendant Console	ARC Connect	Netwise NOW
Can be used through Firewalls.	No	Yes	Yes
Suitable for use when a Callmanager cluster is being used for multiple customers.	No	Yes	Yes
Multitenant Capabilities	Not suitable for use in H/M-UCS multitenant environments	Requires a separate instance of the product per customer in H/M-UCS multitennant environments	A single server instance can service multiple tenants.

Feature	Cisco Attendant Console	ARC Connect	Netwise NOW
Multicluster Capabilities	No capability to use across Callmanager clusters	The H/M-UCS 1.6 integration assumed single cluster operation for a given ARC connect server hence a single customer needs to be constrained to a single cluster	Multicluster operation is supported.
Provisioned Through BVSM	Yes, in large enterprise deployment model only	No	Planned early 2007 but not part of 1.6 H/M-UCS release

Integration with Cisco H/M-UCS Architecture

All the attendant console products mentioned integrate directly with Callmanager using the CTI interface (either JTAPI or TAPI) available via the CTI manager service. The dial plan needs to be modified to allow CTI route points and ports to be incorporated to connect to the Attendant console servers. Detailed integration documents are either available or under preparation.

Scalability of a CallManager Cluster when connected to an Attendant Console

The H/M-UCS solution provides the following attendant console features for deployments with a single Cisco Unified CallManager cluster:

- The ARC and Netwise attendant consoles use CTI for line state monitoring. The Cisco attendant console uses an internal interface for monitoring line state so is potentially more scalable.
- CTI Line-state monitoring is a heavy overhead on the Cisco Callmanager servers so must be factored into the design calculations when determining how many phones can be supported on a single cluster. The Callmanager capacity tool can be used to determine the impact and is available online at http://www.cisco.com/cgi-bin/CT/CCMCT/ct.cgi.
- The absolute maximum number of CTI monitored phones on a cluster formed with lower performance servers (7835 and below) is 3200 and on high performance servers (7845) it is 10000. To achieve these numbers four CTI managers are required in the cluster and the load must be evenly balanced between them.

Netwise Multi Tenant Attendant Console Architecture

In contrast to single-enterprise deployments of Cisco Unified CallManager with associated attendant console capabilities, the challenge here is to provide an attendant console platform that can simultaneously support multiple consoles, with each console managing operator calls into separate tenant sites without conflict.

The Netwise attendant console platform is capable of integrating into the H/M-UCS reference solution and has been selected for its capability to support the Cisco multi-tenant H/M-UCS solution as well as installations of major TDM PBX vendors. This platform also offers a PBX-neutral attendant console capability for customers who have a mixed TDM- and Cisco-IP-PBX infrastructure.

Netwise System Architecture

The system solution consists of resilient pairs of Netwise CTC and CMG servers working in conjunction with client attendant console applications running Netwise attendant console services (NOW) on locally hosted desktop PCs.

The CTC servers manage the task of queuing and routing inbound operator calls, integrating with Cisco Unified CallManager via TAPI. The CTC servers also use TAPI to perform line state monitoring.

The CMG servers meanwhile maintain databases of phones and end users. The CMG database can be partitioned on a per tenant basis so that directory services for each tenant customer can be managed separately.

Provisioning of Attendant Console Services

In H/M-UCS Version 1.6 (BVSM version 3.1.6) the configuration of the Netwise attendant console and associated integration into Callmanager still has to be performed manually. Integration documentation for to help with this task is under preparation.

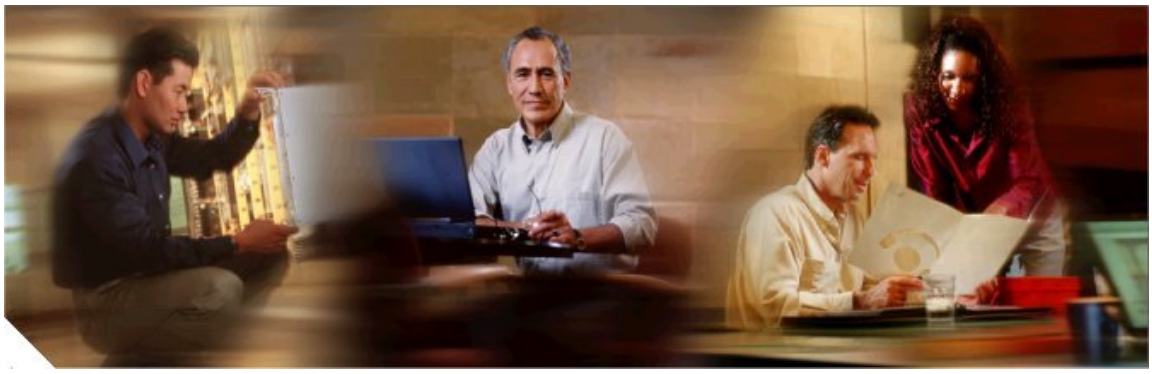

Regulatory Requirements

The H/M-UCS reference architecture must fulfil specific service provider (SP) regulatory requirements for the geographic area in which it is deployed. For most deployment models the H/M-UCS reference architecture uses the Cisco PSTN Gateway (Cisco PGW) as the call routing engine and takes advantage of capabilities within the Cisco PGW to satisfy those local regulatory requirements. For the deployment models where the PGW does not act as the main call routing engine the reference architecture uses Cisco Unified CallManager capabilities or capabilities from other network elements to fulfil these requirements.

The capabilities listed in the following sections apply to all services in the H/M-UCS architecture.

- SS7 Interconnection
- CLIP/CLIR and CLI-Related Requirements
- Number Portability
- Lawful Intercept
- Emergency Services
- Malicious Call Identification
- Billing Accuracy

SS7 Interconnection

The Cisco PGW provides an SS7 interface for Cisco voice and data products. It is suited to providing a carrier-class interconnect capability for H/M-UCS voice services. The Cisco PGW has a large library of locale-specific ISDN User Part (ISUP) variants that can be leveraged, and the Cisco PGW is the most widely deployed and homologated platform of its kind.

The Cisco PGW is capable of internetworking any supported protocol to SS7 in as standard a manner as possible. Supported protocols include TDM standards, such as PRI, DPNSS and Q.SIG; and IP protocols, such as H.323. Some key standards that are supported are listed in the table below.

Table 15 Supported SS7 Related Standards

Reference	Protocol
H.246 Annex C	H.323/ISUP internetworking
Q.699	ISDN/ISUP internetworking

CLIP/CLIR and CLI-Related Requirements

Many requirements that relate to calling-line identification (CLI) must be satisfied in a regulated environment. Some requirements involve the integrity of the information that is being transported and other requirements involve the privacy rights of the concerned end user.

Most networks base their support of the CLI feature on the capabilities described by ITU-T Recommendation Q.951 and the appropriate ISUP capabilities and interoperability described by ITU-T Recommendations Q.76x and Q.699. This support is generally found both in networks with ISDN user access (as described by the Q.95x documents) and in networks with analogue user access.

CLI support in the H/M-UCS solution architecture provides the following key capabilities:

- CLIP/CLIR
- CLI Validation
- Multiple CLIs and Presentation Numbers

CLIP/CLIR

The most basic CLI-handling services are CLI Presentation (CLIP) and CLI Restriction (CLIR), which must be provided in any H/M-UCS network. It must be possible to present and restrict the presentation of CLIs based on a combination of received user preference (configurable) and sent user preference, which is indicated by the CLIP indicator sub parameter in the CLI. Many countries require networks to provide the ability for each user to set the CLIR indicator on a call-by-call basis by dialling a prefix before the destination number.

The Cisco PGW facilitates compliance to the legal aspects of CLIP/CLIR in the following ways:

- The Cisco PGW can obey the presentation indication of a given CLI as delivered in a call establishment message. If the call is to be forwarded to a terminating service that does not have the ability to obey the CLI presentation information, or if the call is to be passed forward on a protocol that does not have the ability to carry this information, the Cisco PGW can act upon the restriction indication by withholding the CLI appropriately.
- For calls that are delivered to the Cisco PGW with a CLI but without a presentation or restriction indication, it is possible to define a default, customer-specific setting to be applied to all calls from that particular origin (origin is defined as a trunk group or IP address). The end user may override the default setting on a call-by-call basis by dialling a customer-specific prefix, which is acted upon and stripped before the call is forwarded.

CLI Validation

For digital and IP based interfaces, the value of the CLI is generally sent to the network in some form of service establishment message (such as a Q.931 setup). Because the device that is sending this request is generally outside the control of the SP, it must be possible for the SP to review, verify, and screen the identity of the sender to ensure that the sender is using appropriate values and is not trying to masquerade as someone else (for instance, by spoofing the CLI of a different user).

The Cisco PGW has capabilities that enable a CLI to be verified against a physical route into the platform (for PBX connections) or a source IP address (for native IP protocols). The Cisco PGW analyzes the received identity information and verifies that it is the expected value or is within the range of expected values. If the CLI verification passes, the identity of the source is marked as "user-provided, screened, and passed." If the CLI verification fails, it is possible to replace the delivered identity value with one that represents the true origin of the call. The identity is then marked as network provided.

Multiple CLIs and Presentation Numbers

Some countries require the support of multiple CLIs per call. These CLIs generally include a true network CLI that is used for identification purposes and a presentation number. Multiple CLIs can be generated in either of the following ways:

- Multiple CLIs might be generated by a customer when more than one CLI parameter is delivered in a call establishment message (possible with Q.931). In this case, a network operator will at least screen and verify the true network CLI. Additional CLIs may also be screened, but these are usually only presentation numbers that have business significance to the end user. In any case, multiple CLIs are passed forward as the call progresses.
- Multiple CLIs can also be generated if a user delivers a single presentation CLI and the network operator (as part of the screening function) assigns a true network CLI that can be used to identify the origin of the call. Both values are typically forwarded in TDM environments.

The capabilities of the Cisco PGW have been extended to allow the handling of multiple CLIs. The Cisco PGW can pass forward multiple CLIs and also can generate an additional CLI to be passed forward.


Note

The provisioning of multiple CLIs or separate presentation numbers is not possible via BVSM.

Number Portability

Number portability is the ability for subscribers to take their telephone number with them when they change providers. This service is regulated in many countries. SPs who directly host customer numbers and number ranges are required to provide the ability to port-in and port-out these numbers and number ranges.

The H/M-UCS solution uses the following mechanisms for number portability:

- Donor-onward-routing method—A donor network maintains the portability information for ported-out numbers and reroutes the incoming calls for ported-out numbers onward towards the recipient network.
- All-call-query method—All SPs can access portability information for every ported number in the country. The portability information database can reside inside local SP switches, inside a central repository for each SP, or inside a central repository that is shared by multiple SPs. Before a call is routed to another network, a check of the directory number is performed in the originating network to determine whether the number has been ported. Through access to a local number portability (LNP) database, new routing information is retrieved.

Compared to the all-call-query method, the donor-onward-routing method leads to less efficient routing of calls, because the call is looped in and out of the donor network in most cases. The donor-onward-routing method is, however, more generic, because it assumes fewer capabilities for network platforms.

The typical format for ported numbers in the ISUP Initial Address Message (IAM) message uses one of the following options:

- Standard Nature Of Address N(S)N=3 and an Address parameter with the following parameters concatenated: Ported Prefix (PP); Routing Number (RN), which identifies the recipient network; and Directory Number, which contains the E.164 number.
- Specific Nature Of Address N(S)N=8 and an Address parameter with the following parameters concatenated: Routing Number (RN), which identifies the recipient network, and Directory Number, which contains the E.164 number.


Note

Because number portability is a regulated service, it is normally provided over an SS7 interface to a PSTN or other operator. Number portability is not typically required for PRI interfaces to PSTNs.

The Cisco PGW supports either of the local number portability mechanisms described previously, as well as the less common query-on-release method. The Cisco PGW provides the following two options for locating the portability information database:

- Internal number portability database inside the Cisco PGW—The Cisco PGW has an onboard local number portability database that is accessed as a result of number analysis. This database matches a number with a carrier prefix that is used to route the call to the operator who now provides service to the destination number.
- External Intelligent Network platform—The Cisco PGW also supports the European Telecommunications Standards Institute (ETSI) Intelligent Network Application Protocol (INAP), either via traditional SS7 or Stream Control Transmission Protocol (SCTP)/SCCP-User Adaptation Layer (SUA), as an interface to an external Service Control Point (SCP). The interface can be configured to initiate a dialog to this SCP on the detection of a configurable trigger. The Cisco PGW supports a subset of ETSI INAP Capability Set 1 (CS-1).

For business customers hosted by the H/M-UCS architecture it is not expected that individual numbers will be ported. If a customer tenant moves, its entire number range will be moved, so the amount of data in the portability database is always relatively trivial, whether the database is found on the Cisco PGW or on another network platform.

LNP can be applied because the Cisco PGW is involved as the routing platform in every inter-business call, whether the businesses concerned have TDM PBXs or are hosted directly via the multi tenant business voice service.

The next two illustrations show examples to demonstrate number portability in an H/M-UCS business voice service with three tenants: A, B, and C. User 1 in company B wants to call user 2 in company C. The call is routed through the Cisco PGW.

Table 16 Internal and external numbers for the users in the example.

User	Company	Extension	Internal Directory Number	E.164 (National) Number
1	B	501	1.002.00501	08115595501
2	C	1001	1.003.01001	08115501001

Figure 45 Inter-tenant Call Without Number Portability

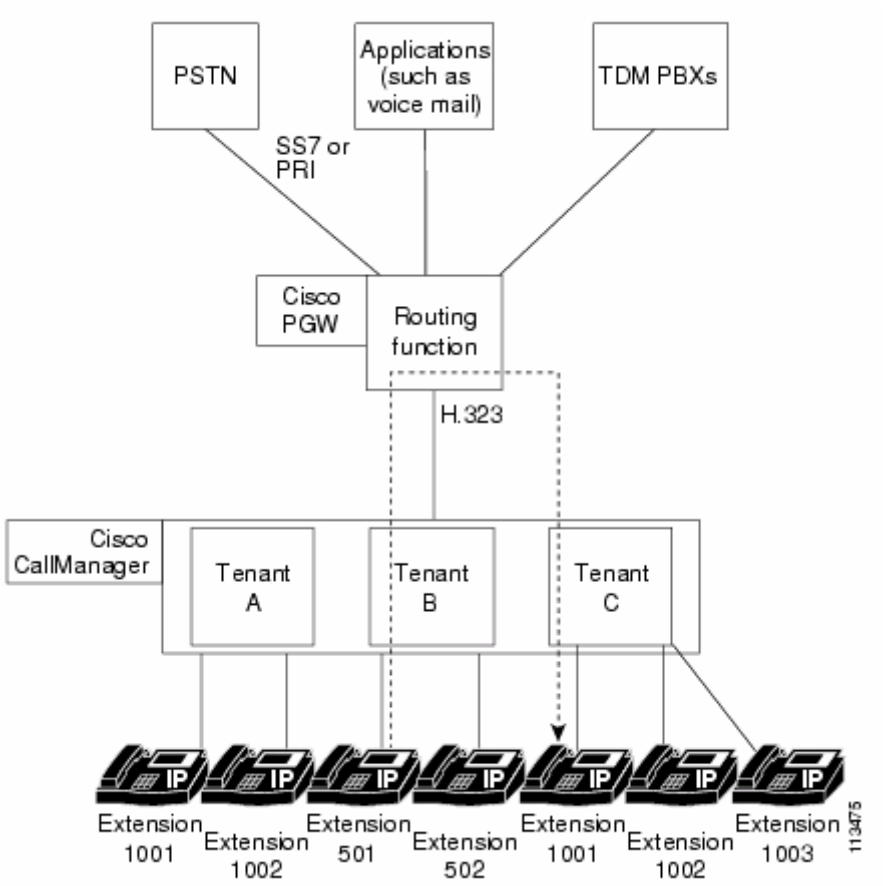

The diagram above 42 shows an H/M-UCS network and a call being routed from Extension 501 in tenant B to extension 1001 in tenant C.

Tenant C then changes SPs and gets its new service from a PSTN operator. Tenant C moves its E.164 number range along with its service. The E.164 number range previously used by tenant C is no longer hosted in the H/M-UCS network.

An LNP table is populated with the carrier routing information for the new SP in order to provide a donor-onward-routing capability for networks that still deliver traffic to the VoIP platform for this number range or to provide an all-call-query service for calls to these numbers that originate from a VoIP platform. This number portability table must be queried whenever a setup request is received for the ported-out number range. This situation is illustrated in the diagram below.


Note

Because company C is no longer being served by the Cisco Cisco Unified CallManager, its data can be removed and internal number range reused. The external E.164 number ranges can be marked by BVSM to prevent their reuse.

Figure 46 Local Number Portability Example

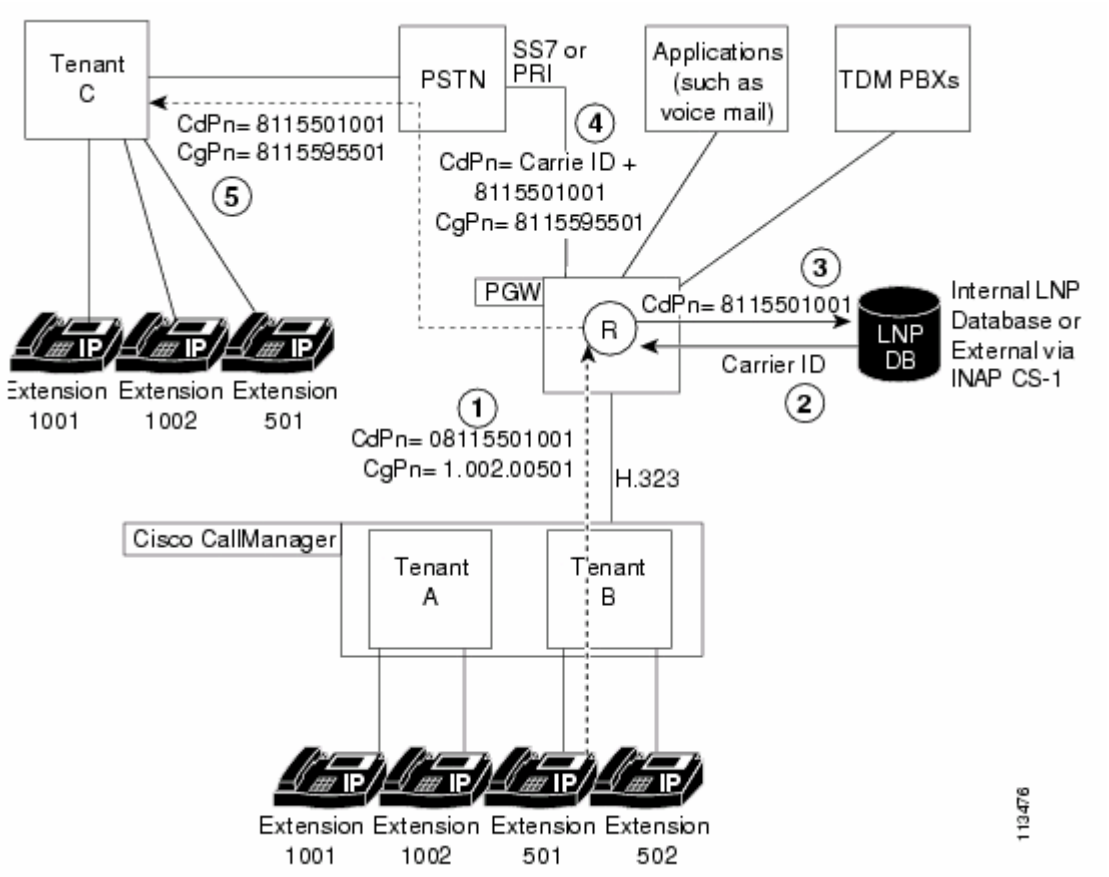

The figure above illustrates that tenant C is now hosted by a PSTN provider. A call from tenant B to tenant C uses the following call flow (the numbers refer to the numbers in the figure):

1. A user at extension 501 in tenant B picks up the handset and dials an E.164 number, 08115501001. The call is routed to the Cisco PGW in the same manner as all external calls. The Cisco PGW manipulates the CLI of the call to generate the public E.164 number of the calling-party number (CgPN) (extension 501). When the Cisco PGW analyzes the called-party number (CdPN), a trigger is set off to query the LNP database. This method assumes that triggers are set against individual dialled numbers or dialled number ranges. It is also possible to configure an all-call-query model, but this is less common in the H/M-UCS environment.
2. The Cisco PGW queries the LNP database using the CdPN (08115501001) as the key. Note that the database can be located on the Cisco PGW or on a different network platform that is reached via an ETSI INAP interface or possibly through a SIP interface.
3. The local number portability database returns the carrier ID of the carrier that now owns this number range. The carrier ID is inserted in front of the CdPN and any modifications that are required to the Nature-of-Address are made. The call is analyzed for routing information, and is routed on the basis of the LNP carrier ID rather than on the basis of the originally dialled E.164 number.
4. The call is routed to the PSTN operator that now hosts company C. The PSTN operator recognizes the call as an LNP call, and identifies its own carrier ID. The carrier ID is stripped from the number and the originally dialled E.164 number is analyzed for termination.
5. The call is routed to the new company C environment with a CdPN of 08115501001 and a CgPN of 08115595501.


Note

BVSM does not provision any of the number portability features of the PGW. They must be configured manually via the PGW command line. VSPT could also be used to input some configuration.

Lawful Intercept

The H/M-UCS reference architecture uses the Cisco Service Independent Intercept (Cisco SII) architecture to provide a network or service operator with the ability to intercept and duplicate voice conversations to or from individual E.164 numbers. The Cisco SII architecture provides only generic abilities to capture these conversations; in all cases, a mediation partner is needed to provide the requisite interfaces to a law enforcement agency (LEA).

The intercept capability described in this section provides the capability to intercept all calls that originate or terminate to a specific number within the VoIP environment, with the single exception of calls that remain within a single enterprise (calls between two extensions belonging to the same business tenant).

The general architecture for Lawful Intercept is shown in Figure 47. A third-party mediation device is employed to provide all the interfaces to the LEA. The mediation device also interfaces to the relevant Cisco call-processing platforms so that intercepts can be enabled and disabled against specific E.164 numbers to VoIP trunking gateways. Calls that can be intercepted are calls made within the VoIP platform, to or from the PSTN and Layer 2/Layer 3 aggregation devices. The Cisco call-processing platforms are instructed to duplicate data streams (RTP) and send them to the mediation device.

The following key interfaces are illustrated in the figure below :

- HI1—Used by a legal authority to send an HI1 (Handover Interface 1) message (usually a fax) to the SP Lawful Intercept administration system with the identity of the target subject.
- HI2—Used to transfer call-related information (such as call-setup information) regarding a call made to or from the target identity.
- HI3— Used to transfer data, such as duplicated RTP packets, from the trunking gateways, the Layer 2/Layer 3 aggregation devices, or both.
- Target identity provisioning—Used by the operator of the Lawful Intercept mediation device to provision a target identity subscriber number in the Lawful Intercept mediation device, which then provisions a number in the Cisco PGW.
- RADIUS—Used by the Cisco PGW to send intercept-related information to the Lawful Intercept mediation device. The RADIUS server provides information about a call-establishment attempt in progress either from or to the target subject. This information is sufficient for the Lawful Intercept device to identify the network devices (gateways or Layer 2/Layer 3 aggregation devices) that must be instructed to duplicate the RTP packets that form the call.
- Simple Network Management Protocol, Version 3 (SNMP V3)—Used by the Lawful Intercept mediation device to send intercept requests to the gateway or Layer 2/Layer 3 aggregation device to duplicate and forward RTP packets (one intercept request is sent for each RTP direction).
- Media/RTP—Used to forward RTP packets to the Lawful Intercept mediation device.

Figure 47 Generic H/M-UCS Lawful Intercept Architecture

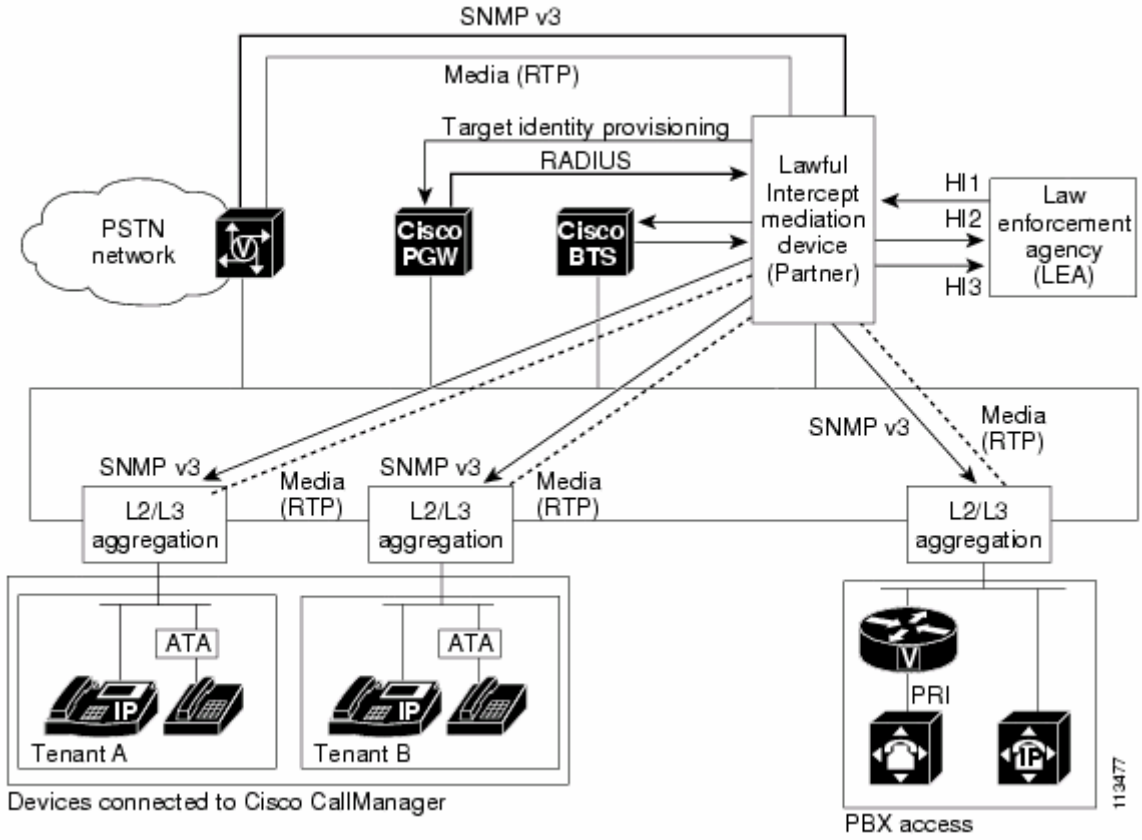

In the H/M-UCS architecture, the Cisco PGW must be able to interface to the Lawful Intercept mediation device to ensure that all calls can be appropriately monitored. The calls that must be monitored include all calls made from that extension to external destinations and all calls made from external originators to that extension. Calls from the target extension to other extensions within the same tenant and calls from other extensions within the same tenant to the target extension are not monitored.

The following example involves the case of a target identity being assigned to a particular extension within a tenant (tenant A). This example (Figure 48) demonstrates an Lawful Intercept for a call that is made from the target extension to an extension at a different tenant (tenant B) in the same H/M-UCS system.

Figure 48 Example of Intercepted Call Between Different Tenants

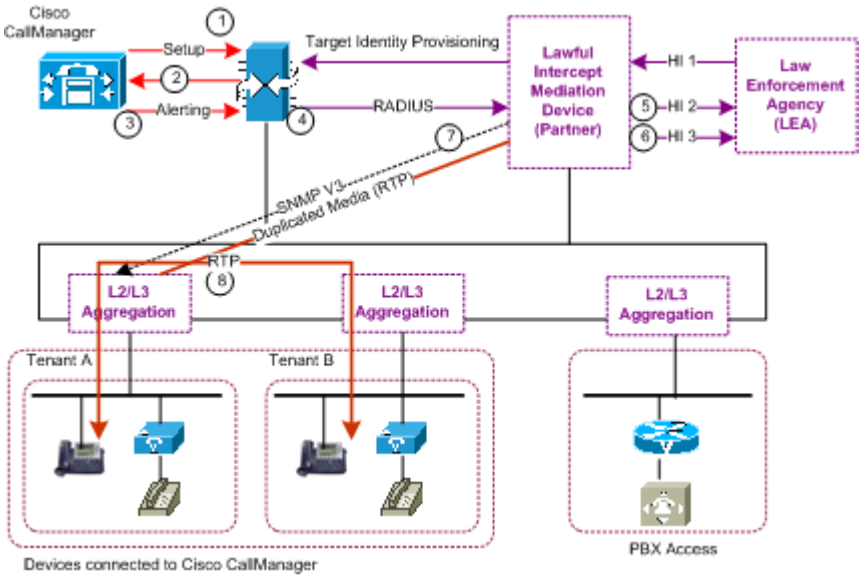

The target extension was provisioned prior to the initiation of the call. The call and its subsequent interception proceed in the following sequence. The numbers refer to those in the figure.

1. The target extension initiates a call and Cisco Unified CallManager sends a SETUP message to Cisco PGW.
2. The Cisco PGW performs the normal call processing. During the process of determining the call routing, the Cisco PGW identifies the call as originating from a target extension, and indicates that the call should be intercepted. The Cisco PGW continues normal call processing and sends a SETUP message back to the Cisco Unified CallManager to initiate a normal inter-tenant call.
3. The Cisco Unified CallManager sends an ALERTING message to the Cisco PGW.
4. The Cisco PGW sends the intercept-related information via a RADIUS interface to the Lawful Intercept mediation device. At this time, the Cisco PGW must know the specifications of the RTP traffic in both directions (the send and receive IP addresses and ports) of the call to be intercepted.
5. The Lawful Intercept mediation device sends Handover Interface 2 (HI2, intercept-related information) to the LEA, which provides details exchanged in Setup signalling for the call (such as who called whom).
6. The Lawful Intercept mediation device initiates a call to establish Handover Interface 3 (Call Content) with the LEA.
7. The Lawful Intercept mediation device sends two Intercept Requests to the Layer 2/Layer 3 aggregation device using two SNMP V3 interfaces (one for each RTP stream). The Lawful Intercept mediation device is responsible for maintaining a network map that allows it to determine, based on the information forwarded from the call-control platform via the RADIUS interface, which network elements to instruct to duplicate and forward the RTP packets.
8. The aggregation device duplicates the IP RTP packets that were specified in the Intercept Request filter specification and forwards them in a tunnel to the Lawful Intercept mediation device.

Responsibility for Compliance

- In many, if not most instances, Cisco's products will have to be combined with products and tools from other vendors in order to provide a complete LI network solution.
- Cisco will provide products and features that we believe (but do not guarantee) will enable our Service Provider customers to implement networks that comply with applicable Lawful Intercept laws.
- It is the responsibility of the Service Provider to ensure that their network is compliant with applicable Lawful Intercept laws.

Emergency Services

Calls to emergency services must be routed to the service bureau that serves the originating subscriber. The caller location function that is required for emergency services calls in the United States does not apply outside the United States.

Emergency service calls have two key requirements:

- The origin of the call must be easily identifiable. Identification of call origin is generally achieved by the normal CLI verification process and, in some cases, by adding a community identity to the CLI.
- The call must be routed to the nearest service centre .

The exact requirements for emergency calls vary on a country-by-country basis. You should verify local requirements before stating any degree of compliance for a particular H/M-UCS solution architecture

There requirements can be addressed in two different and complementary ways:

- PGW Emergency Services Processing
- Cisco Emergency Responder (not supported in H/M-UCS version 1.6)

PGW Emergency Services Processing

The Cisco PGW can easily meet the two key emergency services requirements listed above. The normal CLI screening and verification functions are applied first to ensure that a CLI is valid. Then, as the call passes into the CdPN analysis function, the call is identified as one with an emergency services destination. Once the emergency services destination is detected, the Cisco PGW inserts an optional community identification element before the CLI if required to do so and can also be configured to change the CLI to a site specific CLI upon which a call can be returned (such as to a building security desk or reception). After the community identification element has been added, the destination can be modified, based on the call origination, to the address of the appropriate emergency services centre and the call can be routed. In some countries, it is necessary to reserve small numbers of circuits in trunk groups to be used only for emergency calls, which is achieved by using the Cisco PGW's flexibility when defining trunk groups.

Cisco Emergency Responder

The 1.6 release of the H/M-UCS architecture does not support the Cisco Emergency responder due to difficulties experienced with the integration (it was supported in H/M-UCS version 1.5). It is anticipated that support will be restored in a maintenance release of 1.6 during 2007.

Cisco Emergency Responder (Cisco ER) helps manage emergency calls in the telephony network so that the SP can comply with local ordinances concerning the handling of emergency calls. In North America, these local ordinances are called "enhanced 911," or E911. Other countries and locales might have similar ordinances.

Because emergency call ordinances can differ from location to location within a country, region, state, or even metropolitan area, Cisco ER includes the flexibility needed to fit the emergency call configuration to specific local requirements. However, because ordinances do differ from location to location, and because security requirements differ from company to company, it is required to do extensive planning and research before you deploying Cisco ER in a manner that fits the local legal and security needs.

Enhanced 911 (E911) extends the basic 911 emergency call standard to make it more reliable.

In basic 911 in North America, if a caller dials 911, the call is routed to a public safety answering point (PSAP), also called the 911 operator. The PSAP is responsible for talking to the caller and arranging the appropriate emergency response, such as sending police, fire, or ambulance teams.

E911 extends this standard with these requirements:

- The emergency call must be routed to the local PSAP based on the location of the caller. (In basic 911, the call simply needs to be routed to some PSAP, not necessarily the local one.)
- The caller's location information must be displayed at the emergency operator's terminal. This information is obtained by querying an automatic location information (ALI) database.

In E911, the location of the caller is determined by the emergency location identification number (ELIN), which is a phone number the PSAP can dial to reconnect to the emergency caller if the emergency call is cut off for any reason, or if the PSAP simply needs to talk to the caller again. The emergency call is routed to the PSAP based on the location information associated with this number. For multi-line phone systems, such as an office system, the ELIN can be associated with more than one telephone by grouping the phones in an emergency response location (ERL). In this case, the location the PSAP receives would be the address of an office building. For large buildings, the location would include additional information such as floor or region on a floor. Each ERL requires a unique ELIN.

In addition to these general E911 requirements, each locality can further extend or limit these requirements. For example, a city ordinance might include specific limitations on the size of an ERL (such as, no larger than 7,000 square feet), or on the number of phones that can be included in an ERL (such as, no more than

48 phones). You must work with your service provider and local government to determine the exact E911 requirements in your area.

Malicious Call Identification

Malicious Call Identification (MCID) is another regulatory capability that must be provided in most service provider environments. For a situation in which an end user makes a complaint to regulatory authorities about the receipt of malicious calls, this service is a tool that can be employed to readily identify the originator of such calls. Note that this service is made available only on a temporary basis to victims of malicious calls to aid in detection activity, and is not always available to end users.

For the H/M-UCS solution architecture, the enabling of this feature depends on whether the user who requires the MCID service is located in a foreign network such as the PSTN or is located in the VoIP platform.

For example, a user initiates MCID service from a foreign network and the user who is making the malicious call is located in the H/M-UCS VoIP platform. From an SP perspective, the foreign network connectivity is provided by the SS7 protocol. Either of the following scenarios is possible:

- A network-verified CLI was passed forward when the malicious call was initiated; this is the most common scenario. In this case, when the end user in the foreign network initiates an MCID request, the service is wholly handled by the local exchange to which the user is connected. Because that local exchange is a foreign network, there is no impact on the VoIP platform.
- No network-verified CLI was passed or an option to prevent the call being released forward is required. In this case, the foreign network will send a request message (Identification Request or IDR) and will expect a response that provides some indication of the user's identity or the network's identity (Identification Response or IRS). Optionally, indication can also be received backwards to prevent the originator from releasing the call forward.

The H/M-UCS architecture forwards a network-verified CLI for all calls leaving the VoIP platform. For the exceptional case in which a CLI might not be forwarded, the Cisco PGW supports the backward CLI request mechanism. Note that the ability to prevent a call from being released forward will depend on the capabilities of the VoIP equipment that initiates the call.

In a second example, a user initiates MCID service from the H/M-UCS VoIP platform. The receiver of the malicious call is located within the VoIP platform and the originator of the call is located in another part of the VoIP platform or in a foreign network.

If a user faces malicious calls from the outside they can use a soft key for “Malicious Call Identification” (MCID) if that feature is enabled. Once they hit the soft key during the call the Cisco Unified CallManager CDR entry in the database will include a flag for MCID. The system can also alert the Cisco Unified CallManager administrator by either Simple Network Management Protocol (SNMP) or Real-time Monitoring Tool (RTMT).

If the Public Switched Telephone Network (PSTN) and the H.323 or Session Initiation Protocol (SIP) gateway support it, a malicious call trace can be requested to allow the off-network system to take actions such as tracing the call and/or notifying legal authorities. Cisco H.323 and SIP gateways provide the option to forward MCID to the law enforcement over external ISDN lines using Interactive Voice Response (IVR) scripts.

There are several standards that have to be supported by the PSTN to interact with Cisco Unified CallManager for MCID:

- Q.932—Digital Subscriber Signalling System No. 1
- DSS 1—Generic Procedures for the Control of ISDN Supplementary Services
- Q.951.7—Stage 3 description for number identification supplementary services using DSS 1: MCID
- EN 300 130-1—ISDN MCID supplementary service

Note

Currently the H/M-UCS reference architecture does not support MCID requests to off-network systems.

Billing Accuracy

Some countries require that the timestamps placed on CDRs must be accurate to within 10 ms. This mandatory legal requirement can be satisfied by the Cisco PGW for all business traffic and all traffic to and from the PSTN.

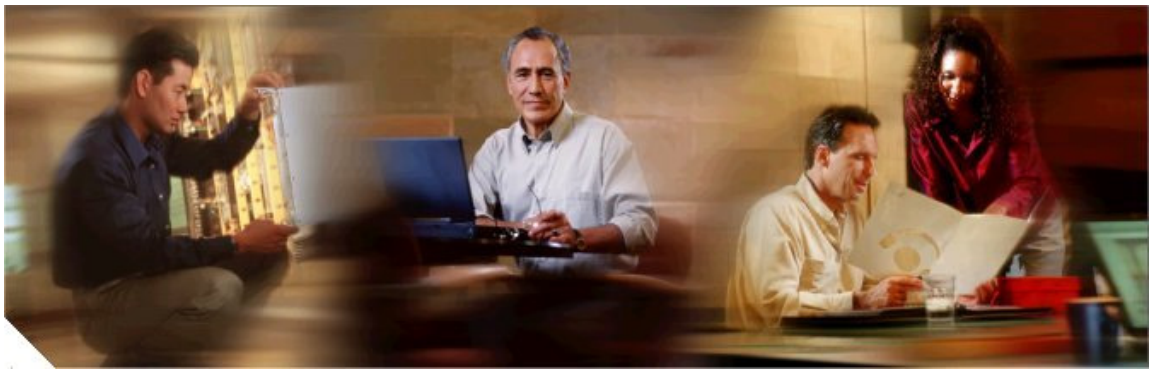

Appendix A – H/M-UCS Architecture Release 1.6 Software Versions

Components	Release	Comments
VoSSManager		
BVSM for Hosted Applications - Eclipse Dual Xeon PC	3.1.6 RC3	
Call Control - Enterprise		
CCM - MCS-7835, 7845	4.2(3)	Based on CCM Magical release and Golden Bridge 4.4 release set
CCM OS	2000.4.3a SR4	See http://www.cisco.com/univercd/cc/td/doc/product/voice/c_callmg/osbios.htm for updates
Call Control, Routing, PSTN Interconnect		
PGW - Netra 440, 1405, 1125, T4, V210	9.6(1) P39 S38, Solaris 8/04	
PGW HSI - T1, V120, V210	4.2 Patch 10, Solaris 8/04	
PGW BAMs	3.20 Patch 6 + QOS Patch	
PGW VSPT	2.6.1 Patch 5	
Gatekeepers		
Gatekeepers - 2821	12.4(7c)	
Gatekeepers - 2851	12.4(7c)	
Gatekeepers - 2891	12.4(7c)	
Gatekeepers - 3640	12.4(7c)	
Gatekeepers - 3660	12.4(7c)	
Gatekeepers - 3745	12.4(7c)	
Gatekeepers - 3825	12.4(7c)	

Gatekeepers - 3845	12.4(7c)	
Messaging - Cisco		
Cisco Unity	4.2(1)	
Cisco Unity TSP	8.1(3)	
Cisco Unity - MS Exchange	Exch 2003 SP2	
Messaging – Third Party		
IP Unity - Mereon 3000, 6000, V120, Sun 100	3.1 SP2	
VoiceRite	4.2.0.x + 8.1.1.40	
Voice Mail DPNSS Gateways	See DPNSS GWs list below	
Voice Mail QSIG Gateways - 1760	See QSIG GWs list below	
Conferencing		
Cat 65xx CMM Module, Cat OS and (IOS, Sup720,Sup32,Sup2)	CMM - 12.4(7c) CatOS 8.5.4 Cat Sup IOS - 12.2(18)SXF4	
DPNSS/QSIG/PRI PBX Voice GWs - PGW Controlled		
2620XM	12.4(7c)	
2650XM	12.4(7c)	
2651XM	12.4(7c)	
2691XM	12.4(7c)	
2801	12.4(7c)	
2811 - DPNSS	12.4(7c)	
3660	12.4(7c)	
3725 - PRI	12.4(7c)	
3745 - DPNSS, PRI	12.4(7c)	
3825 - DPNSS, PRI	12.4(7c)	
3845 - PRI	12.4(7c)	
SRST Gateways - H.323 GK HSI PGW Controlled		
1760-V	3.3 / 12.4(7c)	

2801	3.3 / 12.4(7c)	
2600XM	3.3 / 12.4(7c)	
2811	3.3 / 12.4(7c)	
2650XM	3.3 / 12.4(7c)	
2821	3.3 / 12.4(7c)	
2851	3.3 / 12.4(7c)	
3725	3.3 / 12.4(7c)	
3745	3.3 / 12.4(7c)	
3825	3.3 / 12.4(7c)	
3845	3.3 / 12.4(7c)	
Trunking Gateways		
AS5350 (EoS 22 Dec 2006, last ship 22 Mar 2007)	12.4(7c)	
AS5400 (EoS 22 Dec 2006, last ship 22 Mar 2007)	12.4(7c)	
AS5400-HPX (EoS 22 Dec 2006, last ship 22 Mar 2007)	12.4(7c)	
AS5350-XM	12.4(7c)	
AS5400-XM	12.4(7c)	
CPEs (FXS)		
ATA 186	3.2(3)	
ATA 188	3.2(3)	EoS - http://www.cisco.com/en/US/products/hw/gatecont/ps514/prod_eol_notice0900aecd804047cb.html
VG224 (IOS)	12.4(7c)	
VG248 (non-IOS)	1.3(1)ES8.2	
Endpoints		

7902G, 7905G, 7911G, 7912G, 7920, 7935, 7936, 7940G, 7941G 7960G, 7961G, 7970G, 7971G	Bundled	Refer to http://www.cisco.com/univercd/cc/td/doc/systems/unified/uc511/relnotes/rnipt511.htm
7985 (SCCP Video)	Bundled	Refer to http://www.cisco.com/univercd/cc/td/doc/systems/unified/uc511/relnotes/rnipt511.htm
IP Phone Expansion Module 7914	Bundled	For 7960G (CP-7960G), 7961G (CP-7961G) , 7970G (CP-7970G) and 7971G-GE (CP-7971G-GE)
CIPC (IP Communicator)	2.0(1a)	with CCM 4.1, 4.2 and 5.0
Firewall/NAT		
ASA 5510	7.2(1)	
ASA 5520	7.2(1)	
ASA 5540	7.2(1)	
PIX 501	7.2(1)	
PIX 506	7.2(1)	
PIX 515E	7.2(1)	
PIX 525	7.2(1)	
PIX 535	7.2(1)	
FWSM Firewall Blade for Cat 6500	3.1(3)	Supports SCCP Video
Security Applications		
Cisco Security Agent - Management Centre	4.5	
Cisco Security Agent (Engine/Policy) - CallManager	4.5.1.655 / 2.0(5)	Hot fixes for CSA http://www.cisco.com/en/US/products/sw/secursw/ps5057/index.html http://www.cisco.com/cgi-bin/tablebuild.pl/csahf-crypto
Cisco Security Agent (Engine/Policy) - Unity	4.5.1.639 / 2.0(3)	Hot fixes for CSA http://www.cisco.com/en/US/products/sw/secursw/ps5057/index.html http://www.cisco.com/cgi-bin/tablebuild.pl/csahf-crypto
Cisco Security Agent (Engine/Policy) - ICM	4.5.1.639 / 2.0(2)	Hot fixes for CSA http://www.cisco.com/en/US/products/sw/secursw/ps5057/index.html http://www.cisco.com/cgi-bin/tablebuild.pl/csahf-crypto
Cisco Security Agent (Engine/Policy) - CER	4.5.1.655 / 2.0(5)	Hot fixes for CSA http://www.cisco.com/en/US/products/sw/secursw/ps5057/index.html http://www.cisco.com/cgi-bin/tablebuild.pl/csahf-crypto
Cisco Security Agent (Engine/Policy) - CVP	4.5.1.639 / 2.0(0)	Hot fixes for CSA http://www.cisco.com/en/US/products/sw/secursw/ps5057/index.html http://www.cisco.com/cgi-bin/tablebuild.pl/csahf-crypto
McAfee VirusScan Enterprise, Symantec AntiVirus Corporate Edition, Trend Micro -	See URL	http://www.cisco.com/en/US/products/sw/voicesw/ps556/prod_bulletin0900aecd800f8572.html

http://www.cisco.com/en/US/products/sw/voicesw/ps556/prod_bulletin0900aecd800f8572.html		
Video		
CUVA (VT Advantage) with IP Phone	2.0(2)	SCCP only
CUVA & CIPC Integration, no IP Phone	2.0(1) / 2.0(1a)	SCCP only
Network and Element Management		
Cisco Unified Operations Manager - CUOM	2.0	(Version 1.1 EoS/EoL Last order date 17 May 2007)
Cisco Unified Service Monitor - CUSM	2.0	(Version 1.1 EoS/EoL Last order date 17 May 2007)
Cisco MGC Node Manager - CMNM	2.6.1	Includes CEMF 3.2 SP7 and VSPT
Billing And Measurement Server - BAMS	3.20	For PGW
Cisco Adaptive Security Device Manager - CASDM	5.2	ASA/PIX Element Manager
Cisco Info Centre - CIC	7.0.3	OEMed from IBM. Fault Mgmt, receives SNMP traps
Cisco Network Registrar - CNR	6.1	
CiscoWorks LAN Mgmt Solution - LMS	2.5.1	

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

European Headquarters
Cisco Systems Europe
11 Rue Camille Desmoulins
92782 Issy-Les-Moulineaux
Cedex 9
France
www-europe.cisco.com
Tel: 33 1 58 04 60 00
Fax: 33 1 58 04 61 00

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems Australia, Pty., Ltd
Level 9, 80 Pacific Highway
P.O. Box 469
North Sydney
NSW 2060 Australia
www.cisco.com
Tel: +61 2 8448 7100
Fax: +61 2 9957 4350

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the

Cisco Web site at www.cisco.com/go/offices.

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China • Colombia • Costa Rica • Croatia • Czech Republic Denmark • Dubai, UAE
Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico
The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Singapore • Slovakia • Slovenia
South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe