

How Cisco IT Designed and Deployed Cisco Unity for Global Voicemail

Cisco Unity voicemail adoption saves equipment lease and management costs.

Cisco IT Case Study / Messaging / Unity: This case study describes the architecture and design solution behind Cisco IT's internal deployment of Cisco Unity within the Cisco global network, a leading-edge enterprise environment that is one of the largest and most complex in the world. Cisco customers can draw on Cisco IT's real-world experience in this area to help support similar enterprise deployments.

BACKGROUND

About Program Unity

Program Unity is a joint initiative between the Cisco® IT organization and Enterprise Communications Software Business Unit (ECSBU) to replace the existing Avaya Octel voice-messaging system with Cisco Unity™ technology.

This initiative implements a Cisco Unity voice-messaging desktop solution across the global Cisco network—the largest Cisco Unity deployment ever undertaken—to provide all users with a single, standard user interface and a range of voice-messaging function improvements. Three additional Cisco IT initiatives also fall under the domain of Program Unity:

- Adoption of a standard global dialing plan of 8+7 digits
- Consolidation of site-specific Cisco Call Manager servers to a centralized call processing (CCP) model in select data centers
- Migration of Cisco IP Contact Center (IPCC) agents and call routing applications to a limited number of CCP Cisco Call Manager clusters to use dedicated Cisco Unity voice-messaging servers for the IPCC population

The deployment of Cisco Unity Voice Messaging is the first step toward unified messaging. Program Unity is a critical element of the Cisco AVVID (Architecture for Voice, Video and Integrated Data) program. Cisco Unity represents a suite of IP-based messaging services that enable unified messaging capabilities with Microsoft Exchange and Lotus Domino. The implementation of Cisco Unity reflects the long-term commitment of Cisco to deliver an end-to-end architecture for the convergence of voice, video, and data across its own network. In addition, Program Unity is aligned with the IP telephony and centralized call processing initiatives and is consistent with the Cisco long-term goal of converging unified messaging with IP telephony to form an IP Communications solution.

Program Unity Team: Architecture and Design Engineering Tracks

Program Unity consists of a global cross-functional team with members chosen for their expertise in various disciplines and functions needed to further the initiative. The team was divided into eight tracks, each representing a major program deliverable: Architecture, Design, Support, Implementation, Automation, Communications, Training, and IPCC applications. In addition to the program manager, the team consisted of a lead for each track (with the exception of the Implementation track, which had a lead for each of the global regions), and representatives from the ECSBU.

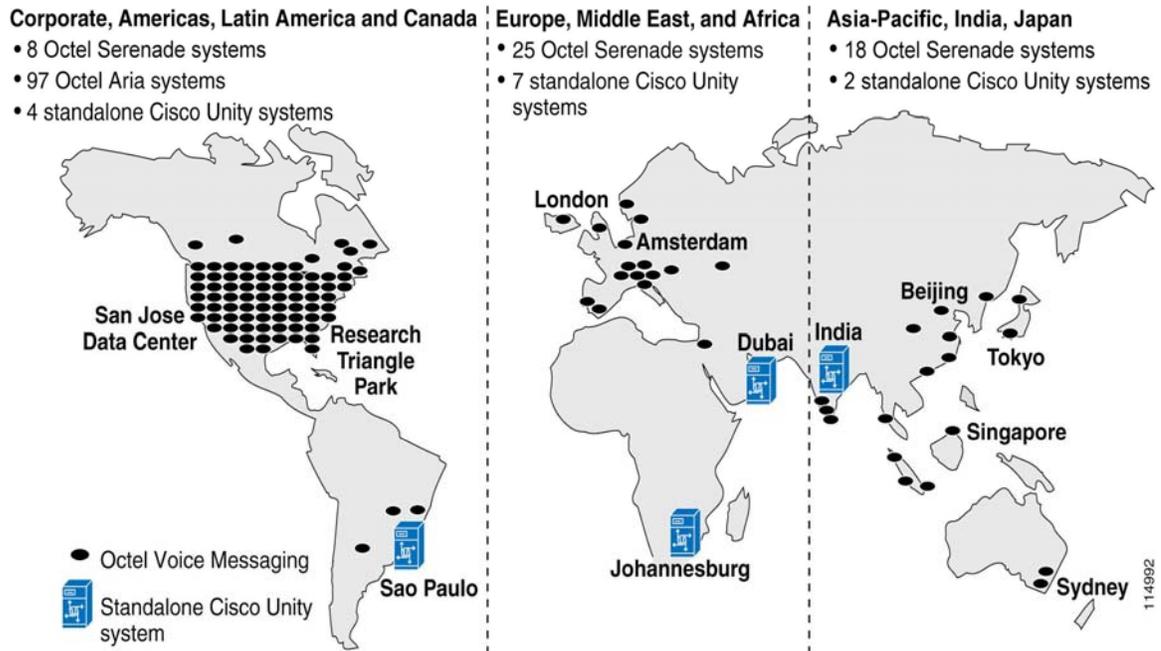
The Architecture track was assigned to develop the architectural solution for deployment of Cisco Unity Voice Messaging across the global enterprise, while the Design track defined the specific requirements for what functions

would (and would not) be available in the deployment.

Background: Cisco Voice-Messaging Environment

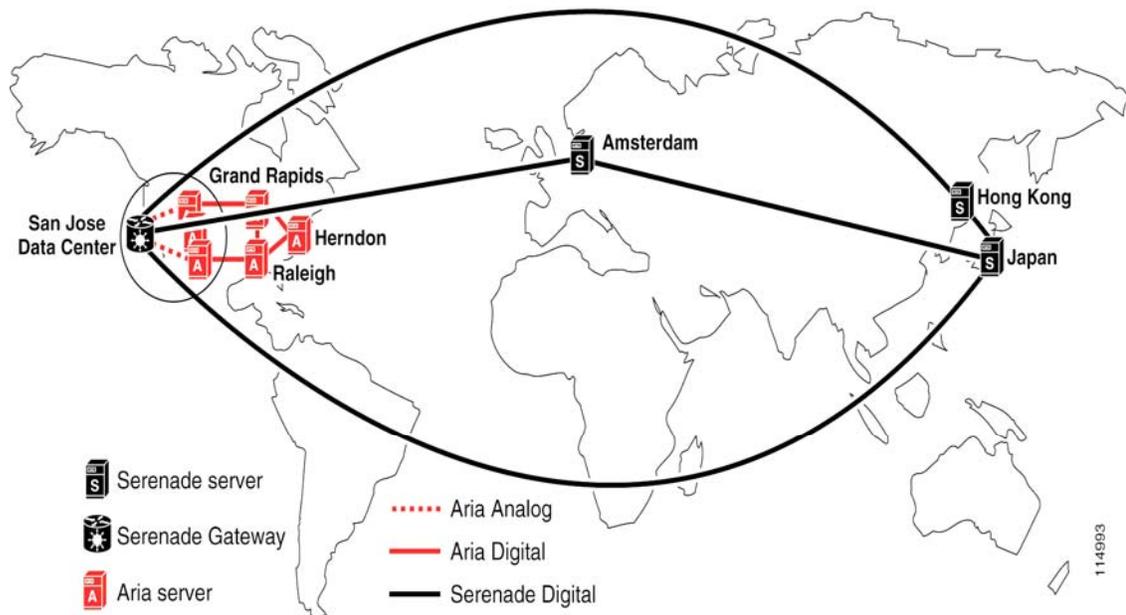
At Cisco, the voice messaging environment consisted of approximately 160 voice messaging systems distributed worldwide. These systems were a mix of two types of Avaya Octel systems—the Aria and the Serenade—as well as a few standalone Cisco Unity systems as shown in Figure 1.

Figure 1. Cisco Voice-Messaging Environment—160 Octel Systems (March 2003)



Background: Octel Message Routing

The Octel Aria servers were situated in the San Jose, California (corporate) and Americas regions, and the Octel Serenade servers primarily were in the Europe, Middle East, and Africa (EMEA), Americas International, and Asia-Pacific regions. The Aria and Serenade servers each use a different digital protocol to communicate that are not compatible with each other, and all communication between these two server types must use Octel analog networking. Because of poor line quality both inside and outside of North America, a Serenade server using a specialized gateway role was implemented in San Jose to route voice messages between critical Octel Aria servers in the United States and all Octel Serenade servers. For example, messages between select Aria servers in the United States and Serenade servers in the EMEA and Asia-Pacific regions are first routed using analog networking within the United States to the Serenade gateway server in San Jose, and are then routed digitally between the Serenade gateway and all other Serenade servers throughout the rest of the Cisco Octel voice network worldwide. See Figure 2.

Figure 2. Octel Voice Messaging Environment: Protocols and Message Flow

Typically, most companies with 160 Octel servers have an Avaya Interchange server to network the voice-messaging systems instead of using the Serenade gateway server. The presence of the Serenade gateway presented some voice-messaging numbering plan challenges for the migration process, but overall it had no effect on the Cisco Unity architecture.

Background: Why Cisco Unity at Cisco?

The primary benefit to Cisco in implementing Cisco Unity was to avoid paying millions of U.S. dollars each year to Avaya in hosting and service costs by replacing 160 Octel systems at 156 locations with 45 Cisco Unity voice messaging servers at ten data center locations. In addition to financial savings, Cisco wanted to create a large enterprise environment to run pre-released versions of Cisco Unity. Because the company was not using its own voice-messaging product, the sales force had a less compelling story to share with customers on the benefits of adopting Cisco Unity technology.

The company was motivated to make the change to Cisco Unity for several additional reasons:

- Reinforce commitment to be Cisco's "Own Best Customer."** Program Unity reflects the "Cisco on Cisco" vision. Replacing the Octel system with Cisco Unity on a global scale emphasizes the company commitment to be its own best customer and underscores the message of the company's quality commitment to the market. Capturing any lessons that the team learned while deploying Cisco Unity would be valuable to the Cisco channel partner helping enterprise customers migrate to Cisco Unity solutions. In addition, any product improvement ideas that were gained from this migration would be given to the engineers responsible for Cisco Unity products to improve future product releases.
- Reduce services costs.** In addition to the moves, adds, and changes service charges, Cisco was paying Avaya monthly high administrative costs, which could be attributed to the Octel system's highly distributed deployment model. By reducing the number of voice-messaging servers by more than 70 percent and consolidating the systems to 94 percent fewer locations, Cisco expects to save money on system administration and support.
- Improve availability in the event of WAN failure.** In the distributed Octel environment, Cisco CallManager servers situated at data centers are connected to Octel voice messaging servers located at remote sites. In the event of WAN circuit failure, the gateway connecting the two systems loses connectivity and the remote

sites lose access to voice mail. With a centralized Cisco Unity solution, should a WAN failure occur, voice-mail access would be routed using the public switched telephone network (PSTN) to the affected sites and preserve availability.

CHALLENGE

A Cisco Unity solution was needed quickly but the Program Unity team had to ensure sufficient program planning and preparation to minimize the effects on the end users migrating to Cisco Unity. Each month that Cisco continued to use the Octel system resulted in costs that the company paid Avaya for maintenance and support as well as lost business revenue because the sales force did not have a compelling migration story to share with customers. The Architecture track was assigned to define a scalable, reference architecture for the Cisco Unity IT global deployment. Other project deliverables included dial plan architecture, Cisco Unity architecture, site classification, proof of concept, and migration strategy.

The following sections describe the challenges, issues, and constraints that influenced the eventual Cisco Unity architecture solution.

Deploying Cisco Unity in a Changing Microsoft Exchange Environment

A global deployment of Cisco Unity has two technical prerequisites:

- The Microsoft Active Directory forest must be configured with all the required information to support Cisco Unity systems.
- Within this AD Forest, there must be a single Microsoft Exchange organization to handle all voice message storage and routing.

The Architecture team's initial objective was to develop an enterprise-wide unified messaging solution. However, unified messaging relies on a stable, global Microsoft Exchange environment. At the time the team was developing the Cisco Unity architecture, Cisco IT was in the process of globally deploying Microsoft Exchange 2000. The Exchange team encountered technical problems during this deployment that delayed the program rollout until upgrades to Microsoft Exchange 2003 and Outlook 2003 were available. Rather than also delay the replacement of the Avaya Octel systems, Cisco IT decided to deploy Cisco Unity as a voice-mail-only implementation in a separate Exchange 2000 organization and Active Directory forest environment so that both IT programs could proceed as quickly as possible.

The decision to create a separate Active Directory forest and Exchange 2000 organization for the Cisco Unity voice-mail deployment was conservative. As the largest known deployment of Cisco Unity, the project presented the team with challenges to provide a rapid deployment on a global scale. Creating this separate environment offered Cisco IT greater flexibility to make changes or modifications that might be problematic to test or diagnose in the shared production Active Directory environment.

The team also determined that creating this separate environment would significantly accelerate the schedule in two ways:

- By separating the Active Directory forest from the corporate production Active Directory environment, the team could eliminate any program schedule dependencies created by other IT projects.
- By creating the dedicated Exchange organization, the Cisco IT team was relieved from managing the user and technical issues that would inevitably arise from the transition to both a new e-mail system and a new voice-mail system at the same time.

While developing the Cisco Unity Voice Messaging architecture, the Architecture track considered the end goal of a global unified messaging deployment by following two guiding principles: First, the Cisco Unity solution servers would be placed as close as possible to the intended Exchange e-mail servers on the network and in the same data center location. Second, the Active Directory account aliases would be the same for both the production voice-mail Active

Directory forest and the production e-mail Active Directory forest. Both of these actions will improve the eventual migration to Unified Messaging after the Microsoft Exchange deployment is complete.

Integrating Voice and Data Technologies Across the Global Network

The Cisco IT Unity Voice Messaging solution has seven component technologies:

Web browser client application, as well as traditional telephone user interface

- Cisco Unity
- Cisco Unity Bridge
- Cisco CallManager
- Microsoft Exchange 2000 (E2K)
- Microsoft Active Directory
- Third-party fax application

Integrating these voice and data technologies on a global scale across the Cisco network remains one of the biggest challenges facing the IT infrastructure organization. A significant learning curve existed for the IT organization in understanding the integration issues between the voice technologies, Active Directory, and Microsoft Exchange. Most IT resources either had a strong background in the “data” or the “voice” technologies but not in both. Cisco IT addressed this issue by creating a strong cross-functional team whose members were familiar with both technologies and by working with the ECSBU to modify existing training or develop new training classes. For example, the Cisco Unity System Engineer course was updated with new material to address this challenge. Although this course was specifically customized for the Program Unity team, it is also available to Cisco customers via Cisco’s learning partners.

Maintaining Availability of Networked Voice Mail to All Users During Migration to a Cisco Unity System

The Architecture team decided that a global flash cut, where all sites worldwide would be transferred to the Cisco Unity solution, was not the best migration strategy because this approach could cause a Day 2 support spike from users unfamiliar with the new interface and would delay the start of the migration off the Octel systems. A primary program requirement was that at all times during the migration, all end users must be able to send and receive networked voice mail and be able to receive recorded voice name confirmation when addressing voice messages. Networked voice mail was used heavily by much of the Cisco user population and especially by the executive, sales, and administrative staffs, who must receive important voice messages. The flash cut approach could not guarantee that networked voice-mail service would not be disrupted. In addition, this approach would greatly impact the Cisco IT support staff, who would be unable to adequately provide training and service to so many end users at the same time without an unreasonable, although temporary, personnel increase. For all these reasons, the flash cut approach was considered an unfeasible migration strategy.

Developing an effective migration strategy to maintain both the Cisco Unity and Octel voice-messaging systems during the program build out, testing, and transition was paramount. The team determined that a realistic goal would be to attempt to convert 75 percent of the users as soon as possible and the remaining 25 percent when feasible, while still maintaining all networked voice-mail functions. This meant that the team would focus on migrating a small number of Octel servers that hosted large numbers of users. The best way to do this was to create a “hybrid” environment using the Cisco Unity Bridge, which provides Octel to Cisco Unity messaging interoperability and directory synchronization

Migration Strategy Overview

The migration strategy will be explained in greater detail in a future case study. Following is an overview of the

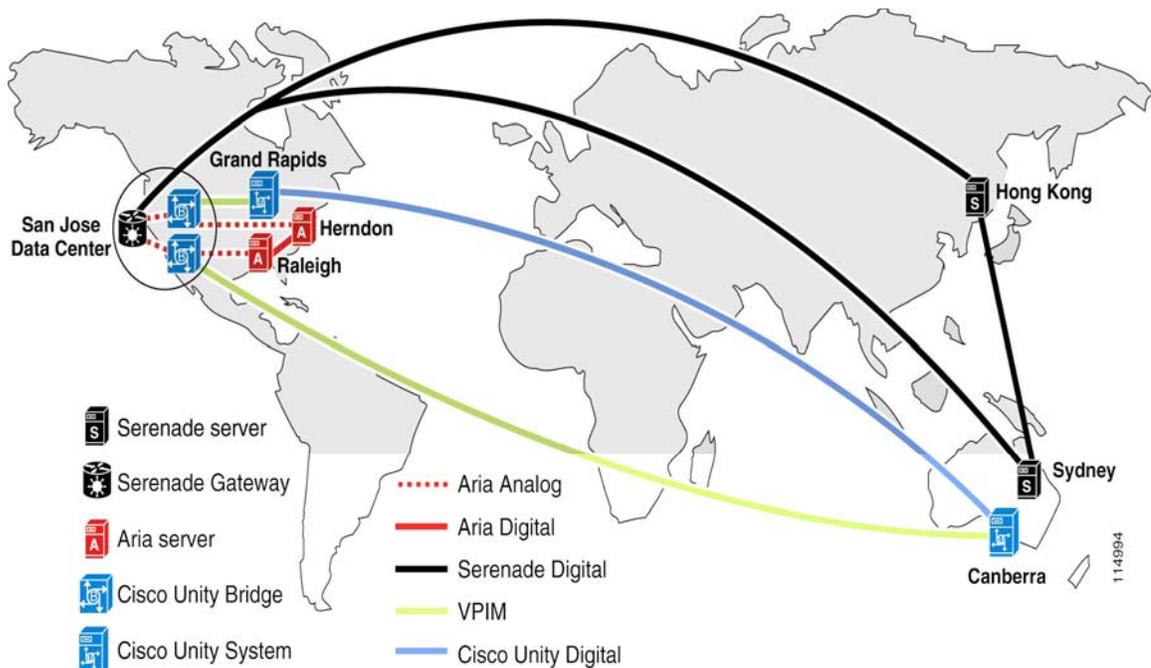
solution designed by the Architecture team:

The global implementation would be a two-phase strategy where 75 percent of the user community would be migrated in the first phase and the remainder Avaya Serenade systems and all the remaining U.S.-based Aria systems in the second phase.

The regional teams would build the infrastructure elements in the data centers (starting with the largest campus) in parallel with each other and test them in this order: Voice over IP (VoIP) networks for CCP, Active Directory Domain Controller/Global Catalog servers, Exchange routing servers, Cisco Unity Bridgehead servers, Cisco clusters, Exchange Message stores, Cisco Unity failover systems, IP interactive voice response (IP-IVR) servers, and fax servers.

Five Cisco Unity bridges would be centralized in San Jose, near the Serenade gateway server, to communicate to the remaining Octel sites. This approach would ensure digital message transmission for all messages into and out of the United States and analog transmission for all calls within the United States. See Figure 3.

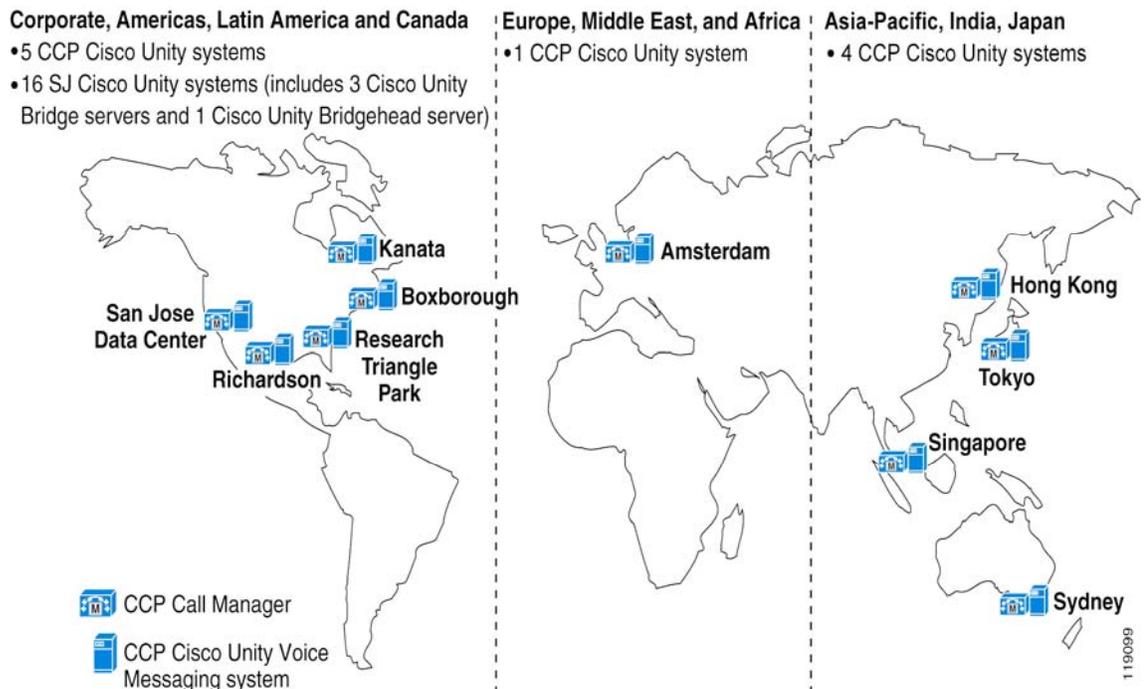
Figure 3. Cisco Voice-Messaging Environment: Protocols and Message Flow During Migration



SOLUTION

Program Unity will deploy a single Cisco Unity Voice Messaging solution across all Cisco geographic regions, replacing 160 Octel systems with a single Cisco Unity solution based on 45 Cisco Unity Voice Messaging servers. The solution developed by the Architecture and Design teams reflects the global nature of the program: All regional teams provided input into the solution architecture and design process, and the CCP solution and dial plan were designed from a global perspective to ensure that no conflicts or overlaps occur. See Figure 4

Figure 4. Cisco Voice-Messaging Environment—45 Cisco Unity Systems (Post Cisco Unity Deployment)



New Functions for Users

In addition to centralizing and consolidating the Cisco voice-messaging environment, the Program Unity solution offers the following new functions to subscribers:

- **Same Cisco Unity conversation for all subscribers worldwide.** The two Octel conversations (from the Aria and Serenade products) are replaced with a single global standard.
- **Web browser client for checking voice messages.** In addition to checking messages by phone, subscribers can use Internet Explorer to access the Cisco Unity Inbox from their desktops to review new and saved messages, listen to voice messages, and save or delete messages.
- **Cisco Unity Assistant.** This function is a Web-based client (Internet Explorer or Netscape) that subscribers can use to manage voice-messaging account settings. For example, subscribers can reset phone access passwords, customize the order in which messages are played, or configure their notification device schedules.
- **Distribution List Manager tool.** This function is a Web-based client that allows subscribers to share and manage their personal distribution lists.
- **Simple Mail Transfer Protocol (SMTP) notification.** Subscribers can use Cisco Unity to be notified of new messages by e-mail or text pager as well as traditional outcalling notification. The combination of using SMTP notification to an e-mail client as well as including a URL in the message to the Unity Inbox is referred to as “Desktop Messaging,” and illustrates how the Cisco IT Unity solution is more than a simple replacement of a legacy voice messaging system.

Cisco Unity Geographic Distribution

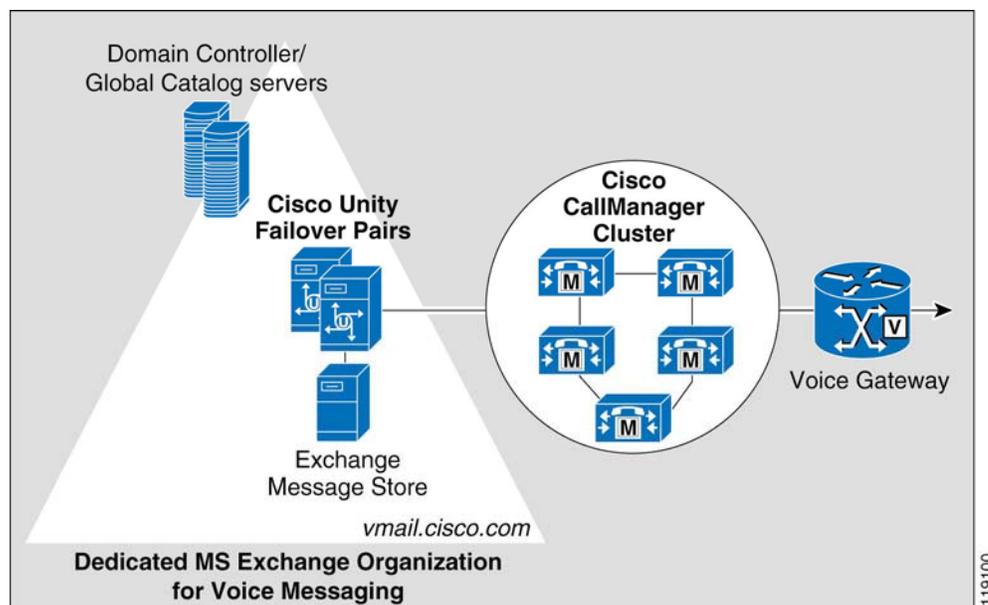
The Program Unity architecture solution components are located in any of ten data centers or more than 300 remote site locations.

Data Center Sites

Data center sites are major core sites with the necessary infrastructure to host a Cisco Technical Assistance Center, Cisco production and development servers, or serve as regional business centers. Ten data center sites worldwide host a complete Cisco Unity Voice Messaging system—Cisco Unity failover server pairs are co-located with at least one Cisco CallManager cluster (a single Cisco Unity server supports up to five Cisco CallManager clusters, if necessary), a pair of Active Directory Domain Controller/Global Catalog servers for redundancy, and an Exchange mailbox server for message storage. For message routing, a pair of Microsoft Exchange bridgehead servers are located in the San Jose, Research Triangle Park, Amsterdam, and Sydney data centers. All Cisco Unity voice messages are routed through these Exchange Bridgehead servers to data centers in the other regions.

Most sites are hosted on a single Cisco Unity server (see Figure 5). The largest campus sites, San Jose and Research Triangle Park, are hosted on multiple Cisco Unity servers.

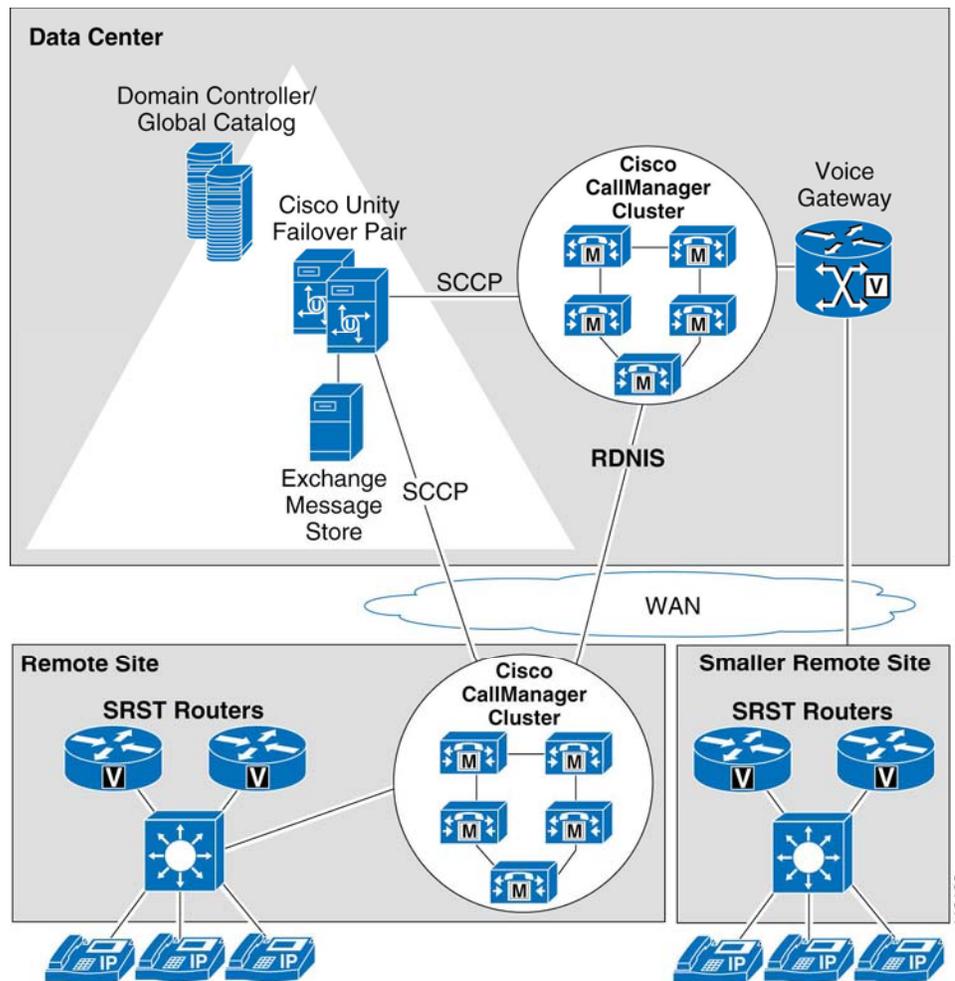
Figure 5. Data Center Site Architecture



Remote Sites

Remote sites are sites directly connected to the data center for their voice-messaging services (for example, New York, Herndon, Paris, or Munich). These sites can be large, often with more than 250 users. Remote sites contain a Cisco CallManager cluster at the site location but do not contain Cisco Unity, Domain Controller/Global Catalog, or Exchange servers; instead, these sites rely on the data center site for these services.

Smaller sales offices do not contain a Cisco CallManager cluster and instead rely on the WAN connection to the data center site for all voice call and messaging functions. The only function locally present at these sites is Survivable Remote Site Telephony (SRST), which maintains system availability if the WAN goes down and the phones transfer to the local router. SRST also maintains access to messaging if the WAN connection fails, by routing voice mail using the PSTN until the WAN link is restored (see Figure 6).

Figure 6. Remote Site Architecture

Standardized Global Voice-Messaging Numbering Plan

Cisco Unity allows subscribers to send voice messages in two ways: By leaving a message when the caller does not answer the phone or by logging into their Cisco Unity system using either the Unity Inbox or the phone to address a message to one or more recipients. Subscribers can address messages to other subscribers by spelling the recipient's name or by entering a number that has been assigned to the recipient using the organization's network voice-messaging numbering plan.

Ideally, network voice-messaging numbering plans are set up to be consistent with phone system dial plans. This allows voice-mail subscribers to send network voice messages by using the same sequence of numbers they use when placing a phone call. Within the Cisco Octel voice-messaging environment, the Octel servers used either a four- or five-digit mailbox address for sending messages to users serviced by the same Octel server (that is, at the same site). If users wanted to send a voice message to recipients on another Octel server, a seven-digit address was needed (this was a location identifier prefix plus the recipient's extension). The exception to this was the San Jose corporate site, which was large enough—an Aria domain of five Octel Aria servers—that all users within this domain could send messages to all other San Jose users using a five-digit mailbox address. However, these users needed to use the seven-digit number to send messages to users outside the San Jose Aria domain.

An important goal of Program Unity was to create a standardized global voice-messaging numbering plan. With the move to VoIP, Cisco employees could make IP telephony phone calls within their site by dialing the recipient's four-

or five-digit extension. Calls to different sites required first pressing 8 (the access code) and then dialing the full seven-digit phone number. One of the goals of the new voice-messaging numbering plan was to ensure global alignment of the phone dial plan and the voice-messaging numbering plan as much as possible. That means that Cisco users could use one number either to call a user at another site on a Cisco VoIP network or to send a voice message to another subscriber.

In the new voice-messaging numbering plan, telephony and voice mail are considered part of the same dial plan. As with VOIP, Unity subscribers press 8 plus seven unique digits to send a voice message. All mailbox addresses are seven digits long and correspond to the telephony directory number. The primary impact to subscribers is that when sending voice messages to someone at the same site, they are now required to dial seven digits instead of the previous four or five. However, Cisco Unity enables subscribers to address messages by spelling the recipient's name, so this may be a faster and preferred method.

The San Jose site was large enough to continue allowing subscribers to call and send messages within the Cisco Unity dialing domain by pressing the five-digit extension. San Jose employees will continue to use the seven-digit address for making phone calls or sending messages outside of this domain. In addition, the five-digit address also applies to subscribers outside the San Jose domain who may want to send messages to recipients inside this domain: subscribers will need only to dial the five-digit address rather than seven digits to send a San Jose employee a message.

Centralized Voice-Messaging Environment

For a previous "Cisco on Cisco" initiative, Cisco IT had centralized Cisco CallManager in data center locations to great success. The next goal was to centralize Cisco Unity servers. With the CCP architecture model, Cisco CallManager clusters at regional data centers provide call processing services to remote sites. By centralizing and reducing the number of voice-messaging servers, it was believed that Cisco IT could save significant cost in system maintenance and administration, reduce operational burden, and provide the same standard voice-messaging features and functions for all employees.

Using Cisco CallManager with SRST Improves Availability

Centralizing the Cisco Unity servers offered improved availability. By using the Cisco CallManager IP telephony solution combined with the SRST feature in Cisco IOS® Software, voice mail would remain available to subscribers in case of WAN failure. If the network has problems, the telephones fail over to the router that is running SRST. The router then sends all calls out of the PSTN. Even if the WAN link to the hub site goes down, Cisco remote office employees can use their IP phones to dial out to the PSTN, receive inbound calls over the PSTN, and make station-to-station calls.

With the current distributed Octel environment, voice mail is not available if the WAN fails. In this situation, the gateway connecting the two systems loses connectivity and the remote sites lose access to voice mail. However, by implementing a centralized Cisco Unity solution with SRST, voice mail would be routed over the PSTN to the affected sites and availability is preserved.

CCP and SRST are not necessary for the Cisco Unity voice-messaging desktop implementation, but they are important building blocks toward the eventual goal of a centralized Unified Messaging solution. If the Cisco CallManager system is in failback mode, voice mail has limited support. It is possible to forward all incoming calls to one number if the line is busy, or if there is no answer, or to forward calls to the main voice-mail number so that callers can reenter the extension of the person they were trying to call. In failback mode, it is not possible to have the call go straight to the recipient's mailbox.

G711 Codec Selection Assures Best Possible Voice-Message Quality

To ensure a smooth functional integration of the Cisco Unity Voice Messaging system, all Cisco Unity servers will be configured using a single design standard as defined by the Design team. An important part of this standard is the

selection and use of the voice codec. All Cisco Unity servers will be configured to create and store messages using the G711 codec (using 64 Kbps for voice), and the G729a codec (which uses only 8 Kbps for voice) will be used to accommodate smaller locations with greater bandwidth constraints. No transcoding on the Cisco Unity servers will be allowed. This standard allows the highest voice quality at all the major locations with the highest voice-messaging traffic, allowing Cisco employees and external callers to experience the best possible service. It also allows low-bandwidth locations to transcode to G729a if necessary across the WAN.

Cisco Unity Solution Components

The Cisco IT Unity Voice Messaging solution comprises the following components that are installed on a Cisco VoIP network to support 35,000 end users at more than 250 locations worldwide:

- Cisco Unity failover servers
- Cisco Unity Bridgehead server
- Cisco Unity Bridge servers
- Cisco CallManager cluster
- Cisco IP-IVR servers
- Microsoft Active Directory Domain Controller/Global Catalog (DC/GC) servers
- Microsoft Exchange 2000 Message Store
- Microsoft Exchange Message Routing servers

Cisco Unity System

Each Cisco Unity system consists of two Cisco Unity servers: A primary server and a secondary failover server for redundancy. Each server has 72 voice ports. A single server model is used interchangeably for either Cisco Unity or Microsoft Exchange. Microsoft Windows 2000 Advanced Server software is installed on both the Cisco Unity and Exchange servers.

Because the majority of the users at remote sites are sales personnel who use voice messaging more than the engineering community (who usually are located at the campus sites), each Cisco Unity system is sized based on a 1:20 port-to-user ratio to support a maximum of 1500 users. The San Jose and Research Triangle Park campuses are unique because of their concentration of engineering staff who use voice mail less frequently than other types of Cisco employees; each Cisco Unity system at these sites will be configured to support up to 3000 users or 40 users per port, depending on the maximum number of ports available. Technically, the Cisco Unity servers can support up to 7500 users if the message server traffic warrants additional users.

Subscribers are mapped to their home Cisco Unity server based on phone number. One or two 1500 direct inward dialing (DID) phone number blocks are allocated to a Cisco Unity server depending on whether the server will be used at a campus or to host remote offices.

Cisco Unity Bridge

The Cisco Unity Bridge acts as a gateway between the Cisco Unity and Octel systems. It uses analog connectivity to transport messages to and from the Octel systems and SMTP to deliver messages or directory updates to Exchange. Each Cisco Unity Bridge server has 24 ports where all voice cards are homed in the associated expansion chassis, but each server is capable of being expanded to up to 48 ports if necessary. All Cisco Unity Bridge servers run on Windows 2000 Standard Edition. The Serenade gateway server has 32 analog ports that interact with a single Cisco Unity Bridge server.

The existing Cisco maintenance contract with Avaya did not include the kind of Call Detail reports that are typically available to customers, so the message traffic between Octel systems and the peak and average utilization of the voice ports was not specifically known. However, based on user patterns and the migration plan, the team decided to have five Cisco Unity Bridge servers (one server is used as a spare). This allocation probably is more than sufficient

to meet company needs.

The Cisco Unity Bridge servers will be located in San Jose and configured to communicate with the Octel servers. Messages to or from any Cisco Unity subscriber, regardless of where they reside on the Cisco Unity network, always will be routed using the specific Bridge configured to communicate with the Octel subscriber's Octel home server. The assignment of Octel nodes to Cisco Unity Bridge servers initially will be based on region and geographic location but also will be influenced by anticipated traffic loads. As traffic patterns change during the course of the migration, the distribution of Octel nodes with which each Cisco Unity Bridge communicates can be modified if necessary to even out the message load distribution.

In addition, a single Cisco Unity Bridgehead server will be configured to centrally manage all Cisco Unity Bridge subscribers during the Octel migration phase, and then to manage the Voice Profile for Internet Mail (VPIM) users during the migration to Unified Messaging. The Cisco Unity Bridgehead server will also be located in San Jose to administer the subscriber and delivery location information associated with the four Cisco Unity Bridge servers.

The four Cisco Unity Bridge Servers are allocated as follows:

- Bridge 1—San Jose, Research Triangle Park, one-third of Americas region locations
- Bridge 2—San Jose, one-third of Americas region locations, Partner locations
- Bridge 3—one-third of Americas region locations
- Bridge 4—San Jose-based Serenade International Gateway server traffic (EMEA and Asia-Pacific region locations)

Cisco CallManager Cluster

Cisco Unity may integrate with up to five Cisco CallManager clusters using the TSP (TAPI Service Provider). At data centers, each Cisco CallManager cluster contains five servers; at remote sites, the Cisco CallManager clusters contain three servers. Each Cisco CallManager cluster must be configured with 144 voice-mail ports per Cisco Unity pair (72 ports for the primary server and 72 ports for the failover server). The Cisco Unity systems are usually co-located in the same data center as the Cisco CallManager cluster that supports them. Exceptions to this are sites in the United States that are considered too large for the CCP SRST solution. For these sites, the Cisco Unity server integrates with a Cisco CallManager server at the remote site as well as with the CallManager server in the data center.

Each set of 144 voice-mail ports is matched to a voice-mail pilot number. This pilot number is set to a voice-mail profile that is assigned to the subscriber phones to use when the "Messages" button is pressed. All Cisco Unity servers in the data centers will integrate with Cisco CallManager clusters that have been networked together using Remote Network Device Interface Specification (RNDIS). In effect, all Cisco Unity systems in the data centers view the Cisco CallManager clusters as one globally networked switch. This IP telephony foundation for the Cisco Unity Voice Messaging solution allows Cisco IT to use voice-messaging integration features in a way that does not occur with traditional switch integrations.

Microsoft Active Directory

The Program Unity Voice Messaging solution includes a single flat Microsoft Active Directory forest that contains all the Domain Controller/Global Catalog, Cisco Unity, and Exchange servers. It is separate from the Cisco production Active Directory forest that users log in to from a personal computer. This Active Directory implementation will be provisioned and maintained with a full mirror of employee accounts for all active Cisco employees. Integrated within this Active Directory environment will be a single Exchange 2000 organization that serves as the basis for all voice-message storage and routing. No mail clients are allowed to connect to these message servers. In addition to their phone, subscribers can access and manage voice messages from their desktops using Cisco Unity Inbox, a Web-based client application. The Cisco Unity servers can support up to 300 simultaneous Web-browser connections, which exceeds the 72 ports of simultaneous voice traffic that could ever occur. The Cisco Unity servers will join this

Active Directory forest, but the Cisco Unity Bridge servers will be in their own workgroup. The Cisco CallManager servers do not use Active Directory for any directory services.

A flat standalone Active Directory forest (vmail.cisco.com) will be used, without any child domains. Each data center site will include at least one pair of domain controller servers for redundancy. These domain controllers also will act as Global Catalog servers for the Exchange 2000 organization. Region object segmentation will be achieved through the use of organizational units within Active Directory. This voice-mail forest will be provisioned using the Cisco IT Enterprise Management organization's existing hub-and-spoke architecture provisioning model, which allows Cisco IT to replicate content simultaneously to both the production voice-mail forest and the production e-mail Active Directory forest. Because the San Jose data center has a sufficient number of Exchange servers, this data center will have at least one Domain Controller/Global Catalog server for every four Exchange message store servers.

Microsoft Exchange 2000 Message Store

Each Cisco Unity system is associated with one or two Microsoft Exchange 2000 servers that act as its primary message store, with a capacity of up to 2500 Cisco Unity subscriber mailboxes. Larger sites may include two Microsoft Exchange 2000 servers. Depending on their entitlement, subscribers are allocated either 60 or 180 minutes of storage. For subscribers with 60 minutes of storage, the estimated mailbox size is approximately 28.8 MB. This number is then tripled to accommodate any additional log or message store requirements. Using these calculations, 168 GBs should adequately support 3000 subscribers (however, no more than 2500 subscribers are expected to be contained on any single Exchange server). Microsoft Exchange Enterprise Edition is installed on all Exchange servers.

If a subscriber reaches 100 percent of the allocated storage, the Cisco Unity system sends a warning message. If the subscriber does not delete messages and reaches 135 percent of the storage quota, a second warning message is sent and Microsoft Exchange prohibits the subscriber to send voice messages. If the subscriber still does nothing and reaches 150 percent of the allocated mailbox storage, a final warning is sent and Exchange will not allow the subscriber to send or receive further messages.

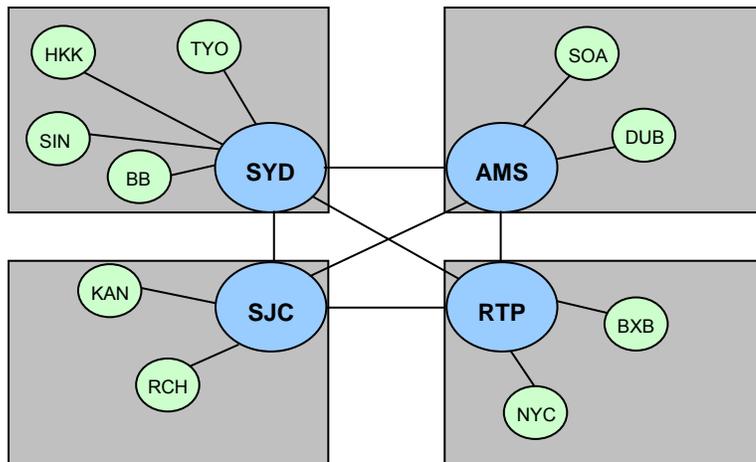
Other Microsoft Exchange 2000 message store details include:

- Use of the G711 codec provides an average message size of 480 KB per minute.
- Deleted messages are retained in a subscriber's deleted items folder for two days before being automatically and permanently deleted, to allow the subscriber time to recover any accidentally deleted voice messages.

Exchange 2000 Message Routing Servers

Message routing is handled through dedicated Microsoft Exchange bridgehead servers. Route groups are created for the Sydney, San Jose, Research Triangle Park, and Amsterdam data centers. Two Exchange bridgehead servers will be assigned to each of these route groups. All Cisco Unity voice messages are routed through these bridgehead servers to the other regional data centers. In addition, because all Bridge servers are located in San Jose, the Cisco Unity voice connector will be installed on both of the San Jose Exchange bridgehead servers to handle voice-mail networking between the Cisco Unity and Octel servers, as well as the eventual use of VPIM networking between Cisco Unity servers within different Exchange organizations.

Within Cisco, alpha development environments also use Cisco Unity. Cisco IT initially plans to use the Cisco Unity Bridge as a gateway between these alpha environments and the production Cisco Unity voice-messaging environment to gain more real-world experience with the Cisco Unity Bridge. However, after all of the Octel systems are removed, Cisco Unity VPIM networking will be used to route messages between the Cisco Unity servers in the production voice messaging Active Directory forest and those in the alpha development environments (see Figure 7).

Figure 7. Cisco Unity IT Deployment Voice-Messaging Routing Plan

In each regional IT center, two Exchange bridgehead servers act as the hub for routing messages within the region as well as for routing messages to the regional data centers. Route groups are connected with Exchange route group connectors, which are configured to route between each site's bridgehead servers. They are organized in a hub-and-spoke configuration with San Jose as the hub. For redundancy, there is full mesh routing between the Exchange bridgehead servers.

Architecture Summary

- Consolidate from current 160 Octel systems to 45 centralized systems using Cisco Unity Voice Messaging
- Deploy new global voice calling and messaging dial plan along with the CCP solution to ensure co-located centralized call and message processing
- Use Cisco Unity Bridge to preserve voice-messaging networking to legacy Octel systems during the migration period
- Consider the eventual migration plan to Unified Messaging:
 - Align network placement of Cisco Unity solution servers with intended Exchange e-mail servers
 - Create similar Active Directory account aliases to ease subsequent migration to Unified Messaging
- Use separate Active Directory forest and Exchange infrastructure from that being used to provide e-mail service to speed up the deployment process

SOLUTION DESIGN EXAMPLES BY REGION

The following sections provide specific data center and region-specific examples to illustrate aspects of the Cisco Unity Voice Messaging solution architecture.

San Jose Region: San Jose Data Center

The San Jose data center is a good example of the redundancy elements that are in place to ensure a highly reliable solution for a large campus site. This site supports approximately 20,000 Cisco Unity subscribers. Like the other data centers, the San Jose site has spare Exchange, DC/DC, and Cisco Unity servers. However, because the Cisco Unity Bridge servers are also located in San Jose, this data center also includes a spare Cisco Unity Bridge server for redundancy (see figure 8).

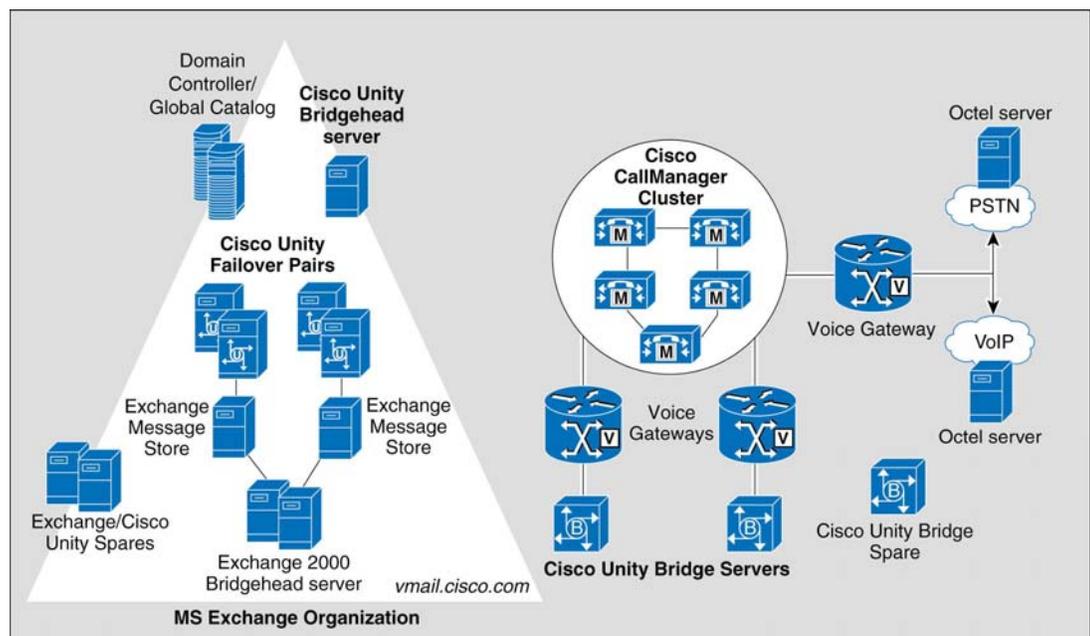
The decisions about where to place potential single points of failure in the Cisco Unity solution strategically were based on cost-benefit analysis. For example, there is only one Cisco Unity Bridgehead server, which serves as a

centralized point of administration for the Cisco Unity Bridge subscribers rather than for hosting any Cisco Unity subscribers. In the event that the Bridgehead server became unavailable, any changes to the Cisco Unity Bridge subscribers would be held until the server was restored.

The team also decided against redundant Exchange Message stores or a storage area network (SAN) backend for storage because Cisco Unity has a message repository that enables subscribers to receive and listen to outside caller messages if the Exchange Message store is unavailable.

The Architecture team based these decisions on the presumption that Cisco IT would migrate to a full Unified Messaging solution after completion of the migration off the Octel systems. Because the Cisco IT deployment of Microsoft Exchange 2000 had a high-availability solution with SAN, the team determined that it was not cost-effective to purchase additional servers for Exchange clustering or SAN support for a limited time. Regarding the Cisco Unity Bridge servers, the team designed the Bridge server solution so that it could be reconfigured quickly to redistribute traffic if a Cisco Unity Bridge server fails. Finally, all servers are backed up over the network with a recovery policy greater than what is in place for the Octel systems today.

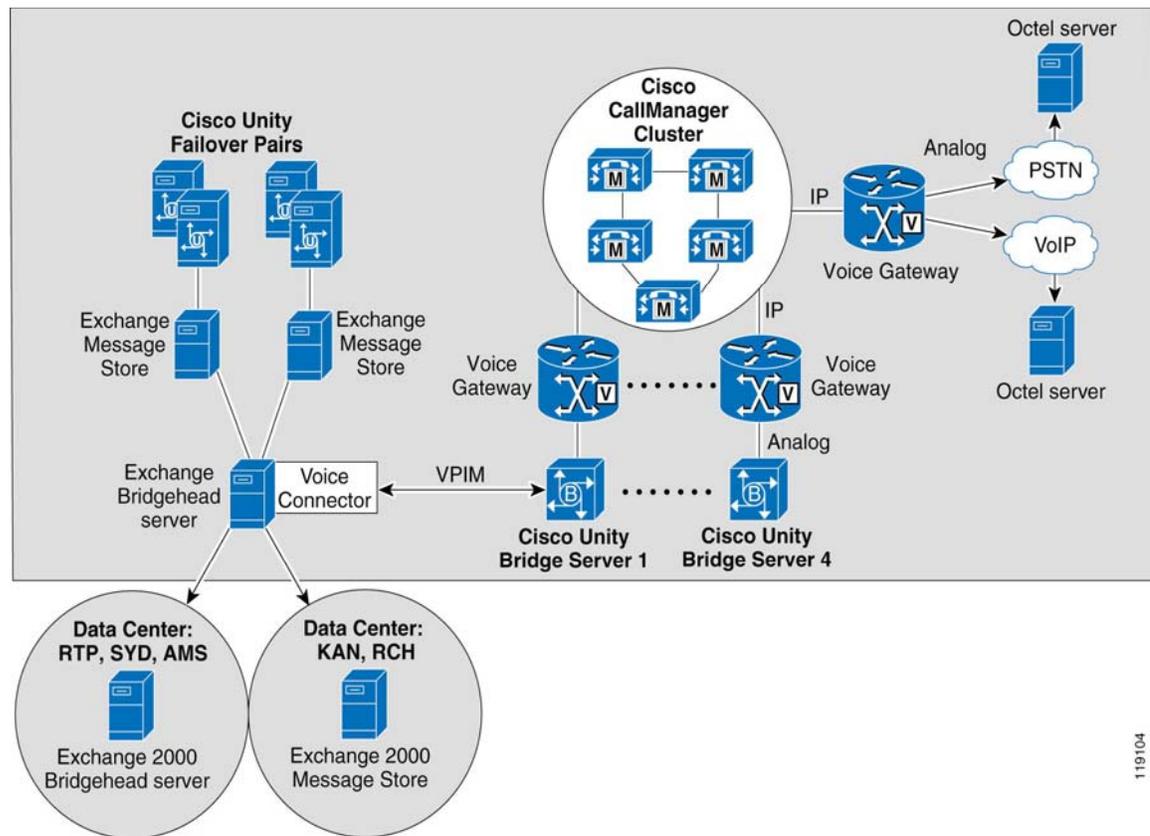
Figure 8. Cisco Data Center Architecture: Cisco Unity Solution Reliability (San Jose Data Center)



Message Routing

Another view of the San Jose data center illustrates how voice messages flow between the Cisco Unity systems and the Octel systems during the migration period. Figure 9 shows the following potential message routing patterns:

- Messages can be routed between Exchange servers for Cisco Unity to Cisco Unity subscriber messaging. If the Cisco Unity recipient is in a different data center, the message is routed by the Exchange Bridgehead server to the other data center.
- If the message is going from a San Jose Cisco Unity subscriber to an Octel user, the message is routed by the Cisco Unity voice connector on the Exchange Bridgehead server in San Jose. The message goes from the Exchange Bridgehead server to the Cisco Unity Bridge server and then out of the data center to the Octel server using Cisco VoIP network or the PSTN.
- In the same way, a message from a Cisco Unity server in a different data center is routed using Exchange to the Cisco Unity voice connector in San Jose.

Figure 9. Cisco Unity Solution: Message Routing During Migration (San Jose Data Center Example)

119104

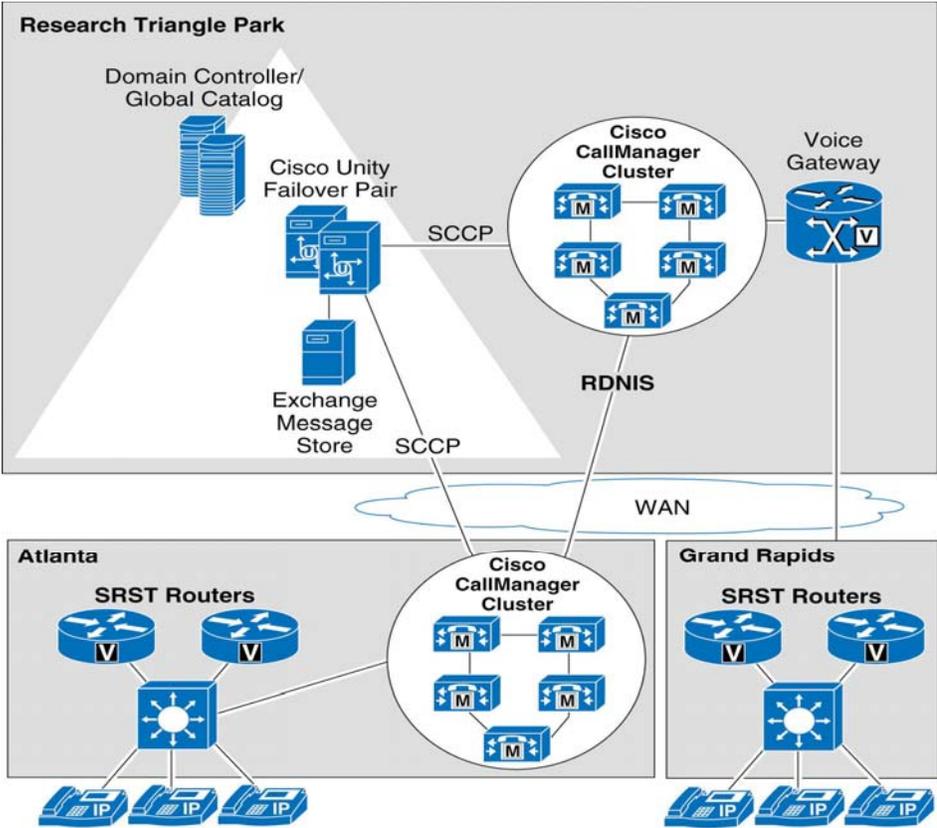
Americas Region: Research Triangle Park Data Center

The Americas region consists of centralized Cisco Unity Voice Messaging systems and Cisco CallManager clusters located in five data centers (San Jose, CA; Boxborough, MA; Richardson, TX; Research Triangle Park, NC; and Kanata, Canada). The Research Triangle Park data center is a good example of both a campus solution and a centralized site where remote sites connect to the data center for voice-messaging services. Each Cisco Unity system can be integrated with up to five Cisco CallManager clusters. For the centralized site solution at Research Triangle Park, it is expected that the data center there will service both the small remote sites (such as Grand Rapids) that have only SRST functions and the larger remote sites like Atlanta that have enough users to require a Cisco CallManager server (see figure 10).

Centralizing the voice-messaging services into a few data centers is a significant factor in reducing Day 2 operational costs associated with Cisco Unity servers. Cisco IT, like many customers, had a number of voice-mail systems and private branch exchanges located at the physical site of the users that were being serviced by those systems. By using a complete end-to-end IP Communications solution, Cisco IT replaced many of the site-based voice-messaging systems with fewer data center-based Cisco Unity servers that will service subscribers on multiple sites. This ultimately reduces the total number of servers required to service Cisco employees.

The Cisco CallManager servers are expected to be networked together using RNDIS so that to the Cisco Unity servers in all data centers they appear to be connected to one giant network switch. This will provide Cisco employees worldwide with voice-messaging features that transcend physical site boundaries. For example, voice-messaging features that are based on the switch integration such as Identified Subscriber Messaging, Reply, or Live Reply will be available in a more economical manner and on a global scale with this IP Communications solution.

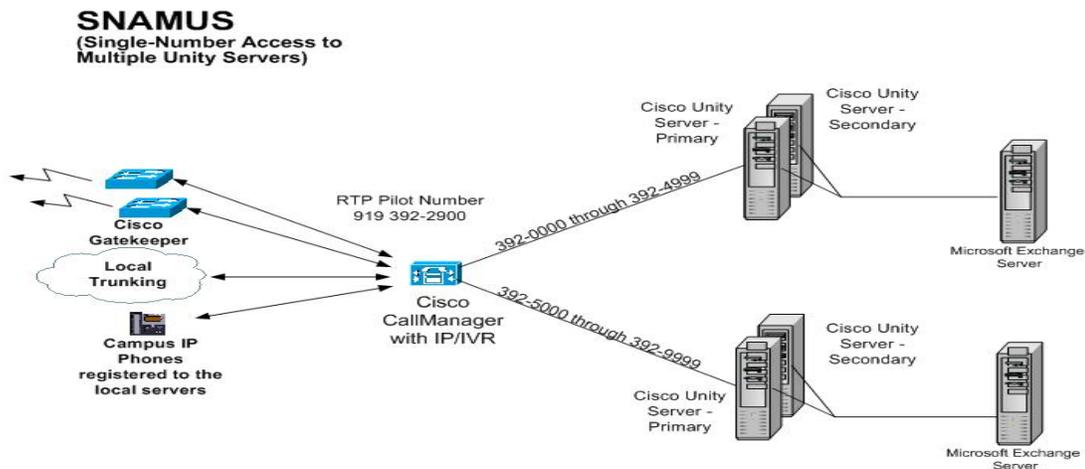
Figure 10. Cisco Unity Remote Site: Voice-Messaging Integration (Research Triangle Park Data Center)



Single Number Access to Multiple Cisco Unity Servers

The Research Triangle Park campus solution uses a Cisco IP IVR solution to provide single number access to multiple Cisco Unity servers (see figure 11). Research Triangle Park subscribers can dial a single number to access messages without having to know or look up the number for their specific Cisco Unity server.

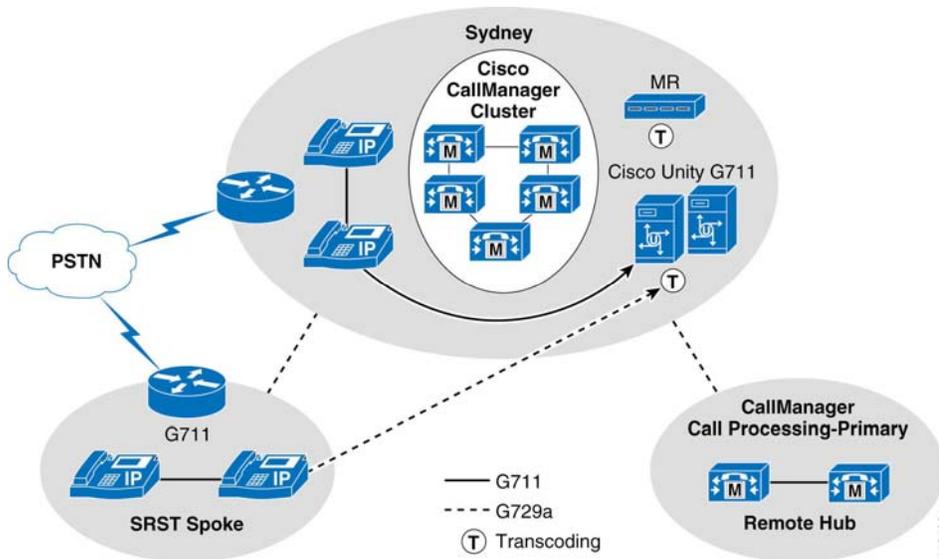
Figure 11. Systems Network Architecture Message Units Solution for Research Triangle Park



Asia-Pacific Region: Transcoding G711 Codec

The G711 codec standard provides the highest voice quality at the major locations with the highest voice-messaging traffic, allowing Cisco employees and external callers to experience the best possible service. The standard also allows low-bandwidth locations to transcode to G729a if necessary across the WAN. In Figure 12, the hardware resources on the router do the necessary transcoding for this remote site solution so that minimal bandwidth is used over the WAN.

Figure 12. Asia Pacific Region: Transcoding G711 Codec



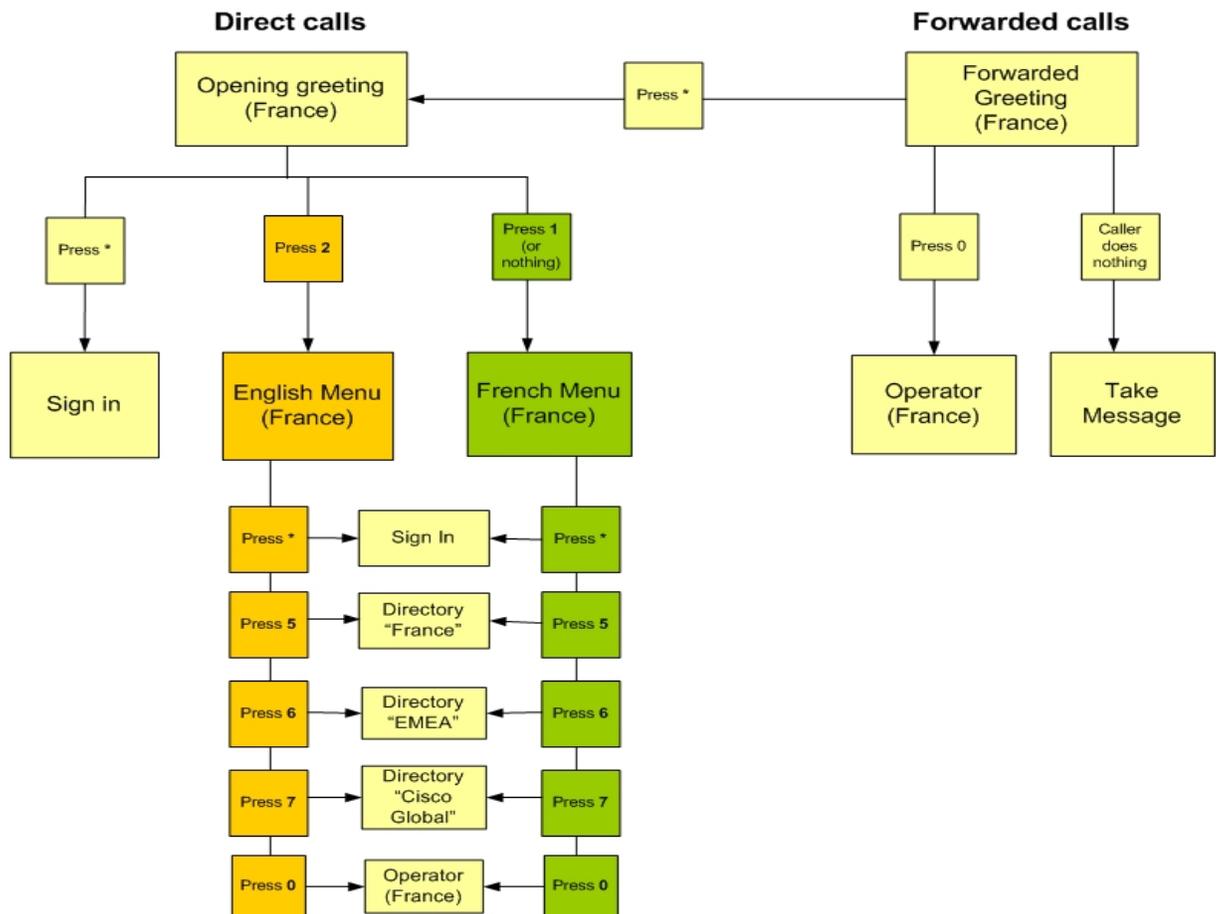
EMEA Region: Automated Attendant Feature

The Cisco Unity Automated Attendant feature automatically answers and directs calls without using a human operator. For example, the Automated Attendant handles the initial greeting and menu that a caller hears when calling Cisco Unity voice mail. This feature can be expanded to suit an organization's needs—for example, to reduce the amount of call traffic handled by human operators.

For the EMEA region with multiple sites using multiple languages, the Automated Attendant feature gives callers the choice of hearing English or the local language. Various factors, such as differing time zones and the need for local

operators at respective EMEA sites, were addressed by the design standards for configuring Cisco Unity Automated Attendants. The example in Figure 14 is a small portion of the automated attendant design. In addition, the team configured multiple alpha directory handlers to accommodate having different sites on the same Cisco Unity system so that callers will be presented with a directory of the users who are local to that particular site.

Figure 13. EMEA Automated Attendant for France



NEXT STEPS AND ANTICIPATED RESULTS

Beginning in June 2004, the Program Unity architectural design will be validated at several small sites with a centralized call and message processing environment to ensure that sufficient network bandwidth is available. During the first implementation phase of the project, several smaller production sites will be used to validate the reliability of the voice-messaging interoperability solution with the Octel systems before the larger San Jose and Research Triangle Park site migration.

Anticipated results from testing at these sites include validating the team's global design solution for replacing a complex legacy voice-messaging environment, as well as gaining additional feedback for integration into the Cisco Unity design solution process to guide future Cisco Unity deployments by enterprise customers.

LESSONS LEARNED

In developing the Program Unity solution, lessons learned so far include:

- **Know your environment.** During the planning, architecture, and design phases of a large Cisco Unity deployment, it is essential to know as much about your existing environment as possible. For example, voice-

mail traffic patterns, number of users, port utilization, distribution list information, etc.

- **Get user feedback early.** If possible, conduct focus groups or user surveys early in the planning process and before the architecture and design phase, to minimize the number of changes to your eventual architecture solution. User feedback also can highlight issues and needs in your training, communications, and Day 2 support planning. For the Program Unity team, providing more emphasis on end user segmentation and gaining a greater understanding of user needs earlier in the process would help ensure that the Cisco Unity architecture solution meets subscriber service-level expectation
- **Define requirements as thoroughly as possible to prevent unanticipated tasks.** Establishing the requirements baseline as early as possible allows the program scope to be clearly understood by team members and better facilitates the development of a comprehensive architecture and design solution.

FOR MORE INFORMATION

To read the entire case study or for additional Cisco IT case studies on a variety of business solutions, visit Cisco on Cisco: Inside Cisco IT www.cisco.com/go/ciscoit

NOTE

This publication describes how Cisco has benefited from the deployment of its own products. Many factors may have contributed to the results and benefits described; Cisco does not guarantee comparable results elsewhere.

CISCO PROVIDES THIS PUBLICATION AS IS WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some jurisdictions do not allow disclaimer of express or implied warranties, therefore this disclaimer may not apply to you.



Americas 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: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

©2007 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0705R)