



# CHAPTER 1

## Unified CVP Architecture Overview

---

Last revised on: January 28, 2009

Over the past two decades, many customers have invested in TDM-based interactive voice response (IVR) applications to automate simple customer transactions such as checking account or 401K account inquires. In addition, many TDM-based IVR platforms were based on proprietary development environments and hardware platforms, which typically meant restricting the customer's integration options with automatic speech recognition (ASR) and text-to-speech (TTS) solutions. Over the past few years there has been a dramatic shift to using VoiceXML standards-based technology to support the next generation of IVR applications.

### What is VoiceXML?

Voice eXtensible Markup Language, or VoiceXML, is a markup language similar to HTML, that is used for developing IVR services and leverages the power of web development and content delivery. VoiceXML was designed for creating audio dialogs that feature synthesized speech, digitized audio, recognition of speech or dual-tone multifrequency (DTMF) key input, and recording of spoken input. It is a common language for content providers, tool providers, and platform providers, and it promotes service portability across implementation platforms. VoiceXML separates service logic from user interaction and presentation logic in VoiceXML voice web pages. It also shields application authors from low-level, platform-specific IVR and call control details. VoiceXML is easy to use for simple interactions, yet it provides language features to support complex IVR dialogs.

VoiceXML programs are rendered (or executed) by a VoiceXML browser, much like an HTML program is rendered via an internet browser (such as Internet Explorer). A Cisco Voice Gateway (or router) can provide the VoiceXML browser function. Specifically, the Cisco IOS VoiceXML Gateway provides the VoiceXML browser function. For small deployments, the Ingress Voice Gateway and VoiceXML Gateway are typically deployed in the same router.

In the most simple call processing scenario, upon arrival of a new call, the voice gateway dial peer matches the call to an available VoiceXML gateway port. The VoiceXML gateway port represents a Voice over IP (VoIP) endpoint and can be logically thought of as a voice response unit (VRU) port. Upon arrival of the new call, the VoiceXML gateway (that is, the VRU) sends an HTTP request to a Cisco Unified CVP VXML Server for instruction. The URL contained in the HTTP request correlates to a specific VoiceXML doc. In response to the HTTP request, the Unified CVP VXML Server sends the requested, dynamically generated VoiceXML doc to the VoiceXML gateway (that is, the voice browser) to be rendered. A typical VoiceXML doc is short and will prompt the caller for some input and then include the results in a new HTTP request that redirects the caller to another URL and VoiceXML doc.

Because a typical call will require numerous prompts and caller inputs, there will be numerous VoiceXML documents that need to be rendered and an infinite number of possible paths through these VoiceXML documents.

To logically link the many different VoiceXML documents that may need to be rendered and to greatly simplify the task of creating VoiceXML documents, a graphical scripting tool is often used to allow the IVR service developer to easily develop complete IVR services with conditional logic and customer relationship management (CRM) database integration. Cisco Unified Call Studio is one such scripting tool. The Cisco Unified CVP VXML Server is capable of executing scripts developed with Cisco Unified Call Studio, and both were designed to work with Cisco Unified CVP Server, Cisco Voice Gateways, Cisco VoiceXML Gateways, Cisco Unified Communications Manager, Cisco Unified Contact Center, and Cisco's VoIP-enabled LAN/WAN.

## What is the Cisco Unified Customer Voice Portal?

Unified CVP is both a product and a solution. There are specific software products that are delivered as part of the media kit for the Unified CVP product. However, the Unified CVP product relies upon other key components of the Unified CVP solution. From a solution perspective, Unified CVP is a VoiceXML-based solution that provides carrier-class IVR and IP switching services on Voice over IP (VoIP) networks. The Unified CVP feature set includes:

- IP-based IVR services

Unified CVP can perform basic prompt-and-collect or advanced self-service applications with CRM database integration as well as ASR and TTS integrated via Media Resource Control Protocol (MRCP). Unified CVP supports both audio and video IVR services.

- IP-based queuing treatment

Callers waiting for agents are presented personalized prompts and music and are prioritized based upon their CRM profile.

- Integration with Cisco Unified Contact Center

Unified CVP integrates with Cisco Unified Contact Center via a VRU Peripheral Gateway (PG). This integration allows Cisco Unified Contact Center Enterprise (Unified CCE) to control Unified CVP VoIP switching and IVR services. This integration also allows Unified CCE to control the agent selection application and to initiate the Real-Time Transport Protocol (RTP) stream transfer from the VoiceXML gateway to the selected agent. Unified CVP integration with Unified CCE requires that the traditional Cisco Unified Communications Manager PG be used for Unified CCE integration with Cisco Unified Communications Manager.

Unified CCE can be integrated with Unified CVP via the Cisco Unified Intelligent Contact Manager (ICM) System PG and the parent-child deployment model. This integration method provides callers with some simple up-front menus and prompts by the parent Unified ICM and Unified CVP, and it intelligently routes the calls via skill groups to the best Cisco Unified Contact Center Express or Enterprise child. Queuing control and agent selection are handled by the child contact center solution. In this model, it is also easy for a TDM automatic call distributor (ACD) to play the role of a child. All call transfers between Unified CVP and children will retain call data, and the ICM will provide enterprise-wide browser-based consolidated reporting.

Unified CVP integration is not directly supported with the Unified CCE System PG (which is also used by System Unified CCE). The Unified CCE System PG supports only the Cisco Unified IP IVR. Unified CVP works only with System PG children via the parent-child deployment model. Unified CVP can also provide IVR services for Unified CCE outbound IVR campaigns and post-call customer surveys.

- IP-based call switching

Unified CVP can route and transfer calls between voice gateways and IP endpoints. Voice gateways provide natural integration of TDM ACDs and PBXs with the PSTN. After completing the routing or transfer of a call, Unified CVP maintains H.323 or SIP call control to provide for subsequent call control. This call control enables Unified CVP to provide switching services similar to takeback-and-transfer (TNT) between IP endpoints via the Unified ICM Enterprise (ICME) interface. Integration with Cisco Unified Presence Server and a gatekeeper helps to provide easily managed dial plans.
- Unified CVP Operations Console

Unified CVP has a single management portal for administrative and configuration tasks for both Unified CVP product components and portions of some Unified CVP solution components (such as gateways). Unified CVP also integrates with Cisco Contact Center Support Tools, but Support Tools must run on a separate physical server.
- VRU reporting

Unified CVP provides a centralized reporting database for access to real-time and historical reporting.

## Unified CVP Components

As mentioned previously, Unified CVP is both a product and a solution. The Cisco Unified Customer Voice Portal (CVP) product consists of the following components:

- [Unified CVP Server, page 1-4](#)
- [Unified CVP VXML Server, page 1-4](#)
- [Cisco Unified Call Studio, page 1-5](#)
- [Unified CVP Reporting Server, page 1-5](#)
- [Unified CVP Operations Console Server, page 1-6](#)
- [Unified CVP Video Media Server, page 1-6](#)

The following components of the Unified CVP solution are not part of the Unified CVP product but are still required for a complete solution:

- [Cisco Ingress Voice Gateway, page 1-6](#)
- [Cisco VoiceXML Gateway, page 1-7](#)
- [Cisco Egress Gateway, page 1-7](#)
- [Multipoint Conference Unit, page 1-7](#)
- [Cisco Unified Communications Manager, page 1-8](#)
- [Cisco Unified Contact Center, page 1-9](#)
- [Cisco Gatekeeper, page 1-9](#)
- [SIP Proxy Server, page 1-9](#)
- [DNS Server, page 1-10](#)
- [Content Services Switch, page 1-11](#)
- [Third-Party Media Server, page 1-12](#)
- [Third-Party Automatic Speech Recognition \(ASR\) and Text-to-Speech \(TTS\) Servers, page 1-12](#)

The following sections discuss each of these components in more detail. Depending on the particular deployment model you choose, some of the above components might not be required.

## Unified CVP Server

The Unified CVP Server component provides the following independent services, which all run on the same Windows 2003 server:

- Unified CVP Server: SIP Service

This service implements a SIP Back-to-Back User Agent (B2BUA). This B2BUA accepts SIP invites from ingress voice gateways and typically directs those new calls to an available VoiceXML gateway port. After completing call setup, the Unified CVP B2BUA acts as an active intermediary for any subsequent call control. While the Unified CVP SIP signaling is hairpinned through this service, this service does not touch the RTP traffic.

Integrated into this B2BUA is the ability to interact with the Cisco Unified ICM via the ICM Service. This integration provides the ability for the SIP Service to query the Unified ICM for routing instruction and service control. This integration also allows Unified ICM to initiate subsequent call control to do things such as requesting that a caller be transferred from queue to an agent or transferred from one agent to another agent.

- Unified CVP Server: ICM Service

This service is responsible for all communication between Unified CVP components and Unified ICM. It sends and receives messages on behalf of the SIP Service, the IVR Service, and the H.323 Service.

- Unified CVP Server: IVR Service

This service creates the VoiceXML pages that implement the Unified CVP Microapplications based on Run VRU Script instructions received from Unified ICM. The IVR Service functions as the VRU leg (in Unified ICM Enterprise parlance), and calls must be transferred to it from the SIP Service in order to execute microapplications. The VoiceXML pages created by this module are sent to the VoiceXML gateway to be executed.

- Unified CVP Server: H.323 Service (Formerly known as the Unified CVP Voice Browser)

This service interacts with the IVR Service to relay call arrival, release, and transfer call control between it and the other H.323 components. This service is needed only for deployments using H.323.

A Unified CVP Call Server can be deployed co-resident with the Unified CVP VXML Server or Media Server.

For hardware details, refer to the latest version of the *Hardware and System Software Specification for Cisco Unified CVP* (formerly called the *Bill of Materials*), available at

[http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod\\_technical\\_reference\\_list.html](http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod_technical_reference_list.html)

## Unified CVP VXML Server

This component executes complex IVR applications by exchanging VoiceXML pages with the VoiceXML gateway's built-in voice browser. Like almost all other Unified CVP product components, it runs within a Java 2 Enterprise Edition (J2EE) application server environment such as Tomcat or WebSphere, and many customers add their own custom-built or off-the-shelf J2EE components to interact with back-end hosts and services. Unified CVP VXML Server applications are written using

Cisco Unified Call Studio and are deployed to the Unified CVP VXML Server for execution. The applications are invoked on an as-needed basis by a special microapplication which must be executed from within the Unified ICME routing script.

The Unified CVP VXML Server can also be deployed in a Standalone configuration that does not include any Unified ICME components. In this model, applications are invoked as a direct result of calls arriving in the VoiceXML gateway, and a single post-application transfer is allowed.

The Unified CVP VXML Server can be installed co-resident with a Unified CVP Call Server or the Media Server.

The Unified CVP VXML Server can execute on Windows 2003 or AIX 5.3 servers. For hardware requirements and details, refer to the latest version of the *Hardware and System Software Specification for Cisco Unified CVP* (formerly called the *Bill of Materials*), available at

[http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod\\_technical\\_reference\\_list.html](http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod_technical_reference_list.html)

## Cisco Unified Call Studio

This component is the service creation environment (script editor) for Unified CVP VXML Server applications. It is based on the open source Eclipse framework, and it provides advanced drag-and-drop graphical editing as well as the ability to insert vendor-supplied and custom-developed plug-ins that enable applications to interact with other services in the network. The Cisco Unified Call Studio is essentially an offline tool whose only interaction with the Unified CVP VXML Server is to deliver compiled applications and plugged-in components for execution.

The Cisco Unified Call Studio can execute on Windows XP or Windows Vista workstations or servers. Because the license is associated with the IP address of the machine on which it is running, customers typically designate one or more data center servers for that purpose. Cisco Unified Call Studio cannot run on machines also running a headless version of the Cisco Security Agent.

For additional hardware details, refer to the latest version of the *Hardware and System Software Specification for Cisco Unified CVP* (formerly called the *Bill of Materials*), available at

[http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod\\_technical\\_reference\\_list.html](http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod_technical_reference_list.html)

## Unified CVP Reporting Server

The Unified CVP Reporting Server is a Windows 2003 server that hosts an IBM Informix Dynamic Server (IDS) database management system. The Reporting Server provides consolidated historical reporting for a distributed self-service deployment. The database schema is prescribed by the Unified CVP product, but the schema is fully published so that customers can develop custom reports based on it. The Reporting Server receives reporting data from the IVR Service, the SIP Service (if used), and the Unified CVP VXML Server. The Reporting Server depends on the Call Server to receive call records. For Standalone Unified CVP VXML Server deployments, one Call Server is needed per Reporting Server. The Reporting Server must be local to the Call Server(s) and Unified CVP VXML Server(s) that it is servicing; deploying the Reporting Server at a remote location across the WAN is not supported. Multiple Reporting Servers should be used and placed at each site when Call Server(s) and Unified CVP VXML Server(s) exist at multiple locations.

The Reporting Server does not itself perform database administrative and maintenance activities such as backups or purging. However, Unified CVP provides access to such maintenance tasks through the Operations Console.

## Unified CVP Operations Console Server

The Operations Console Server is a Windows 2003 server that provides an Operations Console for the browser-based administration and configuration for all Unified CVP product components, and it offers shortcuts into the administration and configuration interfaces of other Unified CVP solution components. This is a required component in all Unified CVP deployments and must run on a separate machine, except in lab deployments.

The Operations Console also offers a direct link to Support Tools, which can collect trace logs and perform other diagnostic and instrumentation functions on many solution components. The Operations Console is, in effect, the dashboard from which an entire Unified CVP deployment can be managed.

The Operations Console must itself be configured with a map of the deployed solution network. It can then collect and maintain configuration information from each deployed component. Both the network map and the configuration information are stored locally on the server, where it can be backed up by off-the-shelf backup tools. A web browser-based user interface, the Operations Console, provides the ability to both display and modify the network map and the stored configuration data and to distribute such modifications to the affected solution components.

The Operations Console can display two views of configuration parameters for managed components. The runtime view shows the status of all configuration parameters as those components are currently using them. The configured or offline view shows the status of all configuration parameters that are stored in the Operations Server database and will be deployed to the device the next time a Save and Deploy option is executed.

The Operations Console allows configuration parameters to be updated or preconfigured even when the target component is not online or running. If the target server (without its services) comes online, the user can apply the configured settings to that server. These settings will become active when that server's services also come online. Only then will they be reflected in the runtime view.

The Operations Console Server is not a redundant component. As such, the Operations Console Server cannot be duplicated within a deployment. Backups of the configuration database should be taken regularly or whenever changes are made.

## Unified CVP Video Media Server

The Video Media Server (VMS) is a Unified CVP component that provides live and recorded video content to callers with video-enabled devices. The VMS stores, manages, and streams 3GP videos and is accessible via two web interfaces: the Cisco Unified Customer Voice Portal Video Media Server Administration and the Cisco Unified Customer Voice Portal Video Desktop.

This component is required for Full Video Service deployments only.

## Cisco Ingress Voice Gateway

The Cisco Ingress Voice Gateway is the point at which an incoming call enters the Unified CVP system. It terminates TDM calls on one side and implements VoIP on the other side. It serves as a pivot point for extension of calls from the TDM environment to VoIP endpoints. Therefore, WAN bandwidth is conserved because no hairpinning of the media stream occurs. Unified CVP Ingress Voice Gateways support both SIP and H.323. Media Gateway Control Protocol (MGCP) voice gateways are supported if they are registered with Cisco Unified Communications Manager.

Supported gateways include the Cisco 2800 Series, 3800 Series, 5350XM, and 5400 XM. The Cisco AS5850 ERSC is also supported as an ingress voice gateway. For the most current list of supported gateways, refer to the latest version of the *Hardware and System Software Specification for Cisco Unified CVP* (formerly called the *Bill of Materials*), available at

[http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod\\_technical\\_reference\\_list.html](http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod_technical_reference_list.html)

## Cisco VoiceXML Gateway

The VoiceXML Gateway hosts the Cisco IOS Voice Browser. This component interprets VoiceXML pages from either the Unified CVP Server IVR Service or the Unified CVP VXML Server. The VoiceXML Gateway encodes .wav files and accepts DTMF input. It then returns the results to the controlling application and waits for further instructions.

The Cisco VoiceXML Gateway can be deployed on the same router as the Unified CVP Ingress Voice Gateway. This model is typically desirable in deployments with small branch offices. But the VoiceXML Gateway can also run on a separate router platform, and this model is typically desirable in deployments with large or multiple voice gateways, where only a small percentage of the traffic is for Unified CVP. This model allows an organization to share PSTN trunks between normal office users and contact center agents and to route calls based upon the dialed number.

The Cisco VoiceXML Gateway can encode .wav files stored in flash memory or on a third-party media server. Prompts retrieved from a third-party media server can be cached in the router to reduce WAN bandwidth and prevent poor voice quality. The VoiceXML doc will provide a pointer to the location of the .wav file to be played or it will provide the address of a TTS server to generate a .wav file. The VoiceXML Gateway interacts with ASR and TTS servers via MRCP.

Supported VoiceXML Gateways include the Cisco 2800 Series, 3800 Series, 5350XM, and 5400 XM. For the most current list of supported VoiceXML Gateways, refer to the latest version of the *Hardware and System Software Specification for Cisco Unified CVP* (formerly called the *Bill of Materials*), available at

[http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod\\_technical\\_reference\\_list.html](http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod_technical_reference_list.html)

## Cisco Egress Gateway

The Egress Voice Gateway is used only when calls need to be extended to TDM networks or equipment such as the PSTN or a TDM ACD. While the RTP stream goes between the ingress and egress voice gateway ports, the signaling stream logically goes through the Unified CVP Server and ICM in order to allow subsequent call control (such as transfers).

## Multipoint Conference Unit

The Multipoint Conference Unit (MCU) consists of three components, the MCU Chassis, the MCU card, and the EMP card (sometimes referred to as the MVP card). The Cisco MCU Chassis, MCU card, and EMP/MVP card are available under the IPVC-3545-xxx Cisco part numbers.



### Note

A firmware conversion is required to make the Cisco MCU and EMP/MVP cards compatible with the software components from Radvision, Inc., that are used in the Unified CVP Full Video Service. A firmware conversion utility for the MCU and EMP cards is provided by Radvision on the installation CD for their software components.

The MCU and EMP register to and communicate with the Radvision IVP Server. Multiple MCUs and EMPs can be used to provide scalability and redundancy and are pooled together for the IVP Server to utilize dynamically when needed throughout a call.

The MCU and EMP are required for Full Video Service deployments only.

## Video Endpoints

When using the Unified CVP Basic Video Service, the following video endpoints are supported:

- Cisco Unified IP Phone 7985G
- Cisco Unified Video Advantage
- Cisco TelePresence

When using the Unified CVP Full Video Service, the following video endpoints are supported:

- Cisco Unified IP Phone 7985G
- Cisco Unified Video Advantage
- Radvision 3G Video Gateway and 3G Mobile Video Phone

## Cisco Unified Communications Manager

Cisco Unified Communications Manager (Unified CM) is the main call processing component of a Cisco Unified Communications system. It manages and switches VoIP calls among IP phones. Unified CM combines with Cisco Unified Intelligent Contact Manager Enterprise (Unified ICME) to form Cisco Unified Contact Center Enterprise (Unified CCE). Unified CVP interacts with Unified CM primarily as a means for sending PSTN-originated calls to Unified CCE agents. SIP gateway calls are routed to an available Unified CM SIP trunk, and H.323 gateway calls are routed to an available Unified CM H.323 trunk.

The following common scenarios require calls to Unified CVP to originate from Unified CM endpoints:

- A typical office worker (not an agent) on an IP phone dials an internal help desk number.
- An agent initiates a consultative transfer that gets routed to a Unified CVP queue point.
- A Cisco Unified Outbound Dialer port transfers a live call to a Unified CVP port for an IVR campaign.

A single Unified CM can originate and receive calls from both SIP and H.323 devices. PSTN calls that arrived on MGCP voice gateways registered with Unified CM can also be routed or transferred to Unified CVP via either SIP or H.323, depending upon the deployment model chosen.

Unified CM is an optional component in the Unified CVP solution. Its use in the solution depends on the type of call center being deployed. Pure TDM-based call centers using ACDs, for example, typically do not use Unified CM (except when migrating to Cisco Unified CCE), nor do strictly self-service applications using the Unified CVP Standalone Self-Service deployment model. Unified CM generally is used as part of the Cisco Unified CCE solution, in which call center agents are part of an IP solution using Cisco IP Phones, or when migrating from TDM ACDs.

Only specific versions of Unified CM are compatible with Unified CVP solutions. Unified CVP is supported with SIP only if Cisco Unified CM 5.0 or later release is used. Unified CVP is supported with H.323 for Cisco Unified CM 4.x or later releases. For full details on version compatibility, consult the latest version of the *Hardware and System Software Specification for Cisco Unified CVP* (formerly called the *Bill of Materials*), available at

[http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod\\_technical\\_reference\\_list.html](http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod_technical_reference_list.html)

## Cisco Unified Contact Center

Either Cisco Unified CCE or Cisco Unified Intelligent Contact Management (ICM) is a required component when advanced call control (IP switching, transfers to agents, and so forth) is required in Unified CVP. The Hosted versions of these products might also be used for this purpose. Unified ICM provides call-center agent management capabilities and call scripting capabilities. Variable storage capability and database access through the Unified CCE or Unified ICM application gateways are also powerful tools. A Unified CVP application can take advantage of these capabilities because Unified CVP applications can be called from within a Unified CCE or Unified ICM script in a non-standalone Unified CVP deployment model.

The Unified CVP Call Server maintains a GED-125 Service Control Interface connection to Unified CCE or Unified ICM. GED-125 is a third-party-control protocol in which a single socket connection is used to control many telephone calls. From the point of view of Unified CCE or Unified ICM, the Unified CVP Call Server is a voice response unit (VRU) connected to Unified CCE or Unified ICM, just as all other GED-125 VRUs are connected. Unified CVP is simply a VRU peripheral to Unified CCE or Unified ICM.

## Cisco Gatekeeper

The gatekeeper is a network element used by H.323 gateways for call routing. It is an optional component. However, most H.323 installations incorporate an H.323 gatekeeper for dial plan configuration and bandwidth management.

One scenario for gatekeeper usage with Unified CVP is to map specific dialed numbers to specific Unified CVP Servers or VoiceXML gateways. Another scenario for gatekeeper usage with Unified CVP is to load-balance new calls to a set of Unified CVP Servers or VoiceXML gateways. A third scenario for gatekeeper usage with Unified CVP is to route the transfer of callers from a VoiceXML gateway port to a Cisco IP Phone.

Two gatekeeper failover mechanisms are supported: Hot Standby Router Protocol (HSRP) and Alternate Gatekeepers.

For more information on H.323 gatekeepers, refer to the Cisco gatekeeper product documentation available online at Cisco.com.

## SIP Proxy Server

The SIP Proxy Server is the component that routes individual SIP messages among SIP endpoints. It plays a key role in Unified CVP's high-availability architecture for call switching. It is designed to support multiple SIP endpoints of various types and to implement load balancing and failover among these endpoints. Currently the only supported SIP Proxy Server is the Cisco Unified Presence Server, which has a built-in SIP Proxy function. Any other SIP Proxy Server must be tested by the customer with Unified CVP to ensure compatibility.

The Cisco Unified Presence Server SIP Proxy cannot use DNS SRV for outbound calls; it must be configured with multiple static routes in order to do load balancing and failover. (The Cisco Unified Presence Server does support the DNS SRV feature, but it is not supported with Unified CVP deployments.) The static routes can point to an IP address or a regular DNS A host record.

Unified CVP may also be deployed without a SIP Proxy Server. In such cases, some of the same functions can be provided by the Unified CVP Server SIP Service. If a SIP Proxy Server is not used, then Ingress Gateways and Unified CMs must point directly to Unified CVP. In such a deployment, load balancing is done via DNS SRV lookups from the gateway to the DNS Server. Load balancing of calls outbound from Unified CVP (outbound call leg) can be done in a similar fashion.

The following guidelines apply to the use of a Cisco Unified Presence server as a SIP Proxy:

- A Unified CM publisher is required in order to install Cisco Unified Presence. Therefore, you need at least one Unified CM publisher if you plan on using the Cisco Unified Presence server as a SIP Proxy (even for a TDM-only deployment with no Unified CM or Unified CCE agents). Unified CM does not need any Device License Units to perform this function.
- A single Cisco Unified CM 6.x publisher can support two Cisco Unified Presence servers; therefore, a single Unified CM 6.x cluster can have two Cisco Unified Presence servers. If you need more than two Cisco Unified Presence servers, additional Unified CM 6.x publishers are required. However, a Cisco Unified CM 7.x publisher can support six Cisco Unified Presence nodes per cluster, contained in three dual-node sub-clusters.
- The Cisco Unified Presence servers can be clustered over the WAN, provided the following conditions are met:
  - Clustering over the WAN with Cisco Unified Presence servers requires either Cisco Unified CM 6.x or 7.x.
  - The Cisco Unified Presence servers must not be used for presence or for registrations.
  - The Cisco Unified Presence servers can be used in SIP Proxy mode only.
  - The WAN round-trip delay between the Cisco Unified Presence server and Unified CM must not exceed 40 msec.

In situations where redundancy across two or more sites is required but clustering over the WAN with Cisco Unified Presence servers is not needed or cannot be supported, you will need at least one Unified CM publisher and one Cisco Unified Presence server at each site. Cisco Unified Presence configuration data is not shared between clusters, so you must configure each Cisco Unified Presence server with dial plan information.

- If you have multiple Cisco Unified Presence servers, in order for them to provide redundancy to Unified CVP, you must configure a DNS SRV record that provides load balancing and/or failover pointing at both servers. You then configure Unified CVP to use this single DNS SRV record as the SIP Proxy Server.
- If you have multiple Cisco Unified Communications Manager clusters, you do not need a Cisco Unified Presence server attached to each cluster for SIP Proxy functionality. It is possible for a single Cisco Unified Presence server to provide SIP Proxy services for multiple clusters. However, depending on where the clusters are located, you might want to have multiple Cisco Unified Presence servers for redundancy (such as with clustering over the WAN).

## DNS Server

This optional component may be installed anywhere in the network. Its purpose in general is to resolve hostnames to IP addresses. However, it can also be configured with multiple IP addresses for the same hostname. Successive requests for a given hostname can be configured to return IP addresses in a

round-robin fashion. Unified CVP's SIP design can make use of this capability in order to implement a form of load balancing among multiple like components. Keep in mind that a DNS SRV lookup is performed for each and every call. If the DNS server is slow to respond, is unavailable, is across the WAN, or so forth, this will affect performance.

The DNS Server comes into play during SIP interactions in the following situations:

- When a call arrives at an Ingress Gateway, the dial peer can use DNS to alternate calls between the two SIP Proxy Servers. The SIP Proxy Servers can also use DNS to distribute incoming calls among multiple SIP Services. If SIP Proxy Servers are not being used, then the Ingress Gateway can use DNS directly to distribute inbound calls among multiple SIP Services.
- When the SIP Service is instructed by Unified ICME to transfer the call to the VRU leg, it can use DNS to alternate such requests between two SIP Proxy Servers. If SIP Proxy Servers are not being used, the SIP Service can use DNS directly to distribute VRU legs among multiple VoiceXML Gateways.
- When transferring a call to an agent using a SIP Proxy Server, the Cisco Unified Presence Server SIP Proxy cannot use DNS SRV for outbound calls; it must be configured with multiple static routes in order to do load balancing and failover. (The Cisco Unified Presence Server itself supports DNS SRV, but it has not been tested in Unified CVP deployments.) The static routes can point to an IP address or a regular DNS A host record. If SIP Proxy Servers are not being used, then the SIP Service can use DNS to locate the target agent's IP address.

Note that use of the DNS Server for SIP routing is entirely optional in Unified CVP. It is not required to have a dedicated DNS Server, but the existing DNS server needs to handle the additional load of Unified CVP. For every call destined for Unified CVP that comes into the network, there will be approximately 3 to 4 DNS lookups. You can determine the number of DNS queries per second by determining the number of calls per second for the solution, and multiplying that number by 4.

## Radvision IVP Server

The Radvision Interactive Video Platform (IVP) is a media server platform for the development and delivery of enhanced interactive video communication services and applications. An IVP Server (IVP Controller) is a component of the Radvision IVP platform that communicates and controls the various components of the IVP, including the MCU and EMP. The IVP Server also communicates with the Unified CVP Call Server via an XML interface in order to notify Unified CVP of new calls and to receive instructions from Unified CVP on how to handle them, such as playing Video IVR prompts, transferring calls to agents, or streaming video files to callers as directed by an agent.

This component is used for Full Video Service deployments only.

## Content Services Switch

The Content Services Switch (CSS) is a load-balancing device designed to provide robust, highly available and scalable network services for data centers. The CSS can be logically placed between one or more VoiceXML Gateways and one or more Unified CVP VXML Servers, Media Servers, and ASR/TTS Servers. Various mechanisms allow the CSS to implement transparent load balancing and failover across these servers. One of these mechanisms is the stateful redundancy mechanism, which is called Adaptive Session Redundancy in CSS terms. Adaptive Session Redundancy (ASR) provides session-level redundancy for applications where active flows (including TCP and UDP) must continue without interruption, even if the master CSS fails-over to the backup CSS.

CSS is an optional device, but it is highly recommended. Without it, the IVR Service implements a "poor man's failover" mechanism, but it is not load-balanced and various retries and delays are part of the algorithm, all of which can be avoided if CSS is used.

The CSS is normally deployed as a Virtual Router Redundancy Protocol (VRRP) pair. VRRP provides box-to-box redundancy for CSS pairs. For session-level redundancy (stateful failover), CSS pairs could use the Adaptive Session Redundancy option to minimize VXML Server license port usage during a CSS failover. VRRP is useful in all deployment models except for Call Director call flows, which do not require use of Unified CVP VXML Servers, Media Servers, or ASR/TTS servers. If SSL is used in the solution, you will need an SSL module for the CSS 11503 or 11506 chassis.

## Third-Party Media Server

The media server component is a simple web server, such as Microsoft IIS or Apache, and is an optional component that can provide prerecorded audio files, external VoiceXML documents, or external ASR grammars to the gateway. Because some of these files can be stored in local flash memory on the gateways, the media server can be an optional component. However, in practice, most installations use a centralized media server to simplify distribution of prerecorded customer prompt updates. Media server functionality can also include a caching engine. The gateways themselves, however, can also do prompt caching when configured for caching. Typical media servers used are Microsoft IIS and Apache, both of which are HTTP-based.

As with ASR/TTS Servers, Media Servers may be deployed simplex, as a redundant pair, or with CSS in a farm. Note that the VoiceXML Gateway caches .wav files it retrieves from the Media Server. In most deployments, the Media Server encounters extremely low traffic from Unified CVP.

The Media Server can be installed co-resident with the Unified CVP Call Server or Unified CVP VXML Server.

For the most current information on media servers, refer to the latest version of the *Hardware and System Software Specification for Cisco Unified CVP* (formerly called the *Bill of Materials*), available at

[http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod\\_technical\\_reference\\_list.html](http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod_technical_reference_list.html)

## Third-Party Automatic Speech Recognition (ASR) and Text-to-Speech (TTS) Servers

This component provides speech recognition services and text-to-speech services for the VoiceXML Gateway. Communication between the ASR/TTS server(s) and the VoiceXML Gateway uses Media Resource Control Protocol (MRCP). MRCP v1 can be used on the VoiceXML Gateway when the application is based on either Micro-Apps or the Unified CVP VXML Server. MRCP v2 can be used on the VoiceXML Gateway only with applications that are based on the Unified CVP VXML Server.

For capacity and redundancy reasons, a CSS is usually used to mediate between a VoiceXML Gateway and a farm of ASR/TTS servers. If CSS is not used, then a VoiceXML Gateway can support a maximum of two ASR/TTS Servers.

Cisco does not sell, OEM, or support any ASR/TTS software or servers. Cisco does, however, test Unified CVP with ScanSoft, Nuance, and IBM offerings. A certification process is currently being developed to allow additional vendors to qualify their products against Unified CVP VoiceXML, and the World Wide Web Consortium (W3C) provides a rich feature set to support the ASR grammars. The simplest to implement and support is inline grammars, by which the set of acceptable customer responses is passed to the gateway. Another form is external grammars, by which Unified ICM passes a pointer to an external grammar source. The Cisco Unified CVP VXML Server adds this pointer to the VoiceXML

document that it sends to the VoiceXML Gateway, which then loads the grammar and uses it to check ASR input from the caller. In this case, the customer is responsible for creating the grammar file. A third type of grammar is the built-in grammar. For a complete explanation of grammar formats, consult the W3C website at

<http://www.w3.org/TR/speech-grammar/>

The text for TTS is passed directly from the Unified CVP VXML Server to the gateway. This action is referred to as *inline TTS* in this document.

The actual speech recognition and speech synthesis are performed by a separate server that interfaces directly to the VoiceXML gateway via Media Resource Control Protocol (MRCP). Currently, ScanSoft, Nuance, and IBM are the supported ASR/TTS engines. These ASR/TTS engines also support (with limitations) voice recognition and synthesis for multiple languages. For the latest information on supported languages and limitations of these ASR/TTS engines, refer to the following websites:

- Nuance and ScanSoft

<http://www.nuance.com>

- IBM

<http://www-306.ibm.com/software/voice/>

Note that these are third-party products, which the customer or partner must purchase directly from the vendor. The customer also receives technical support directly from the vendor. That does not, however, mean that the vendor's latest software version can be used. Unified CVP is carefully tested with specific versions of each vendor's product, and Cisco Technical Assistance Center (TAC) will not support Unified CVP customers who use different ASR/TTS versions than those which have been tested with Cisco Unified CVP. For additional details on supported ASR and TTS products, consult the latest version of the *Hardware and System Software Specification for Cisco Unified CVP* (formerly called the *Bill of Materials*), available at

[http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod\\_technical\\_reference\\_list.html](http://www.cisco.com/en/US/products/sw/custcosw/ps1006/prod_technical_reference_list.html)

## Call Flows

This section describes the Unified CVP call flows for both SIP and H.323 calls.

### Typical SIP Unified CVP Call Flow

A typical SIP Unified CVP call will arrive at an ingress voice gateway and send an invite message to the SIP Proxy Server, which will determine the IP address of the Unified CVP Server for the dialed number. The SIP Proxy Server will then forward the invite to the selected Unified CVP Server SIP Service. The SIP Service consults Unified ICME via the Unified CVP Server ICM Service. This consultation causes Unified ICME to run a routing script. The routing script will typically initiate a transfer of the call to a VoiceXML Gateway port via the SIP Service. The VoiceXML Gateway then sends a message to the IVR Service, which then requests scripted instructions from Unified ICME. Unified ICME then exchanges VRU instructions with the VoiceXML Gateway via the IVR Service. Among these VRU instructions can be requests to invoke more sophisticated applications on the Unified CVP VXML Server. Such requests will result in multiple exchanges between the Unified CVP VXML Server and the VoiceXML Gateway in order to provide self-service.

If the customer wants to transfer to a live agent, the Unified ICME routing script queues for an available agent. While waiting for an available agent, the Unified ICME provides additional instruction to the VoiceXML Gateway to provide queueing treatment to the caller. When an agent becomes available,

Unified ICME sends a message to the Unified CVP Server SIP Service, which forwards a message via the SIP Proxy Server to the Ingress Gateway and to Unified CM to transfer the call away from the VoiceXML Gateway port and deliver it to the Unified CM agent IP phone.

## Typical H.323 Unified CVP Call Flow

A typical H.323 Unified CVP call will arrive at an ingress voice gateway and send a route request to the H.323 Gatekeeper to determine to which Unified CVP Server to send this new call. The ingress gateway then interacts with the Unified CVP Server H.323 Service, which then consults Unified ICME via the Unified CVP Server IVR Service and Unified CVP Server ICM Service. This consultation causes Unified ICME to run a routing script. The routing script will typically initiate a transfer of the call to a VoiceXML Gateway port via the ICM, IVR, and H.323 Services. The VoiceXML Gateway then sends a message to the H.323 Service, which then requests scripted instructions from Unified ICME via the IVR and ICM Services. Unified ICME then exchanges VRU instructions with the VoiceXML Gateway via the ICM, IVR, and H.323 IVR Services. Among these VRU instructions can be requests to invoke more sophisticated applications on the Unified CVP VXML Server. Such requests will result in multiple exchanges between the Unified CVP VXML Server and the VoiceXML Gateway in order to provide self-service.

If the customer wants to transfer to a live agent, the Unified ICME routing script queues for an available agent. While waiting for an available agent, the Unified ICME provides additional instruction to the VoiceXML Gateway to provide queueing treatment to the caller. When an agent becomes available, Unified ICME sends a message to the Unified CVP Server H.323 Service, which then queries the H.323 gatekeeper for routing instructions to the appropriate Unified CM H.323 trunk. Then the H.323 Service signals to the Ingress Gateway and Unified CM to transfer the call away from the VoiceXML Gateway port and deliver it to the Unified CM agent IP phone.

## Design Process

Unified CVP has an infinite number of possible deployment models due to its modularity, flexibility, resiliency, and scalability.

When designing a Unified CVP deployment, it is most important first to choose a functional deployment model. After choosing a functional deployment model, Unified CVP solution designers must determine where the Unified CVP components are going to be deployed (in the data center or at a branch). Then Unified CVP solution designers must choose the amount of availability and resiliency that is justifiable or required. Finally, Unified CVP solution designers must size the deployment to provide the justifiable or required grade of service for the initial deployment and near-term growth.

## Choosing SIP or H.323

When choosing the call control protocol for your Unified CVP implementation, the following factors can have a direct influence on your decision:

- Characteristics of using SIP with Unified CVP
  - SIP is the preferred call control protocol for Unified CVP. New features are not planned for H.323.
  - SIP supports a higher port capacity per Unified CVP call server.
  - SIP can use a static dial plan on the Call Server (suitable for simple deployments).

- Transferred legs use a Ring-No-Answer timeout based on the destination pattern (rather than the single server-wide setting used with H.323)
- Third-party SIP endpoints are supported (subject to testing).
- Video calls are supported.
- Failover and load balancing are performed with the SRV mechanism on a per-call basis, but no heartbeats are maintained and endpoints are not registered.
- Characteristics of using H.323 with Unified CVP:
  - H.323 provides stability and product maturity.
  - H.323 supports post-call survey.
  - H.323 provides load balancing and failover using a gatekeeper, with registered endpoints and a heartbeat mechanism (thus providing faster failover time).
  - GKTMP is supported.
  - Hookflash transfers are supported.

## Functional Deployment Models

As mentioned previously, the first step in the design process is typically to determine what functionality you need. Unified CVP offers the following functional deployment models:

- Unified CVP VXML Server (standalone) — Provides a standalone VRU with no integration to Unified ICM for queuing control or subsequent call control.
- Call Director — Provides IP switching services only.
- Comprehensive — Provides IVR services, queue treatment, and IP switching services. The previously described typical call flows use this functional deployment model.
- VRU Only — Provides IVR services, queuing treatment, and switching for SS7/IN PSTN endpoints. This model relies upon the PSTN to transfer calls between call termination endpoints.

For more details and design considerations for each of these functional deployment models, see the chapter on [Functional Deployment Models, page 2-1](#).

## Distributed Network Options

After choosing a functional deployment model, Unified CVP solution designers must determine where the Unified CVP components will be deployed. Unified CVP deployment can use one of the following primary distributed network options:

- Combined Branch Gateways — Allows for call treatment at the edge and integration of locally dialed numbers into the enterprise virtual contact center. This option can be either a combined Ingress and VoiceXML gateway or separate gateways, but typically they are combined when deployed in a branch.
- Branch Ingress Gateways with Centralized VoiceXML Gateways — Allows for integration of locally dialed numbers and resource grouping of VoiceXML gateways. This option might be desirable for organizations with many medium to large branches but with few contact center calls in those branches. However, VRU announcements would have to traverse the WAN from the VoiceXML Gateway to the Ingress Gateway.

- Branch Egress Gateways — Allows for calls to be transferred across the WAN to remote TDM terminations.
- Branch Agents — Allows for a virtual contact center where agents can be located anywhere on the IP network.

It is also possible to use a combination of these distributed options. For more details and design considerations for each of these distributed network options, see the chapter on [Distributed Deployments, page 3-1](#).

## High Availability Options

After choosing a functional deployment model and distributed deployment options, Unified CVP solution designers must choose the amount of availability required. Unified CVP solution designers can increase solution availability in the following areas:

- Multiple gateways, Unified CVP Servers, Unified CVP VXML Servers, VRU PGs, Cisco Unified Presence Servers, and gatekeepers — Allow for inbound and outbound call processing and IVR services to continue upon component failure.
- Multiple call processing locations — Allow for call processing to continue in the event of a loss of another call processing location.
- Redundant WAN links — Allow Unified CVP call processing to occur upon failure of individual WAN links.
- Content Services Switches — Provide an efficient means to remove failed Unified CVP Servers, Unified CVP VXML Servers, and Media Servers from the load-balancing algorithms being used for those components.

It is also possible to use a combination of these high availability options to be utilized. For more details and design considerations for each of these high-availability network options, see the chapter on [Designing Unified CVP for High Availability, page 4-1](#).

## Scalability Options

After choosing the functional model and the distributed and high-availability deployment options, Unified CVP solution designers must then size their solution and select appropriate hardware. To make Unified CVP deployments larger, Unified CVP supports multiple gateways, Unified CVP Servers, and Unified CVP VXML Servers. To load-balance HTTP requests efficiently to multiple Unified CVP Servers, Unified CVP VXML Servers, and media stores, you can use the Content Services Switch (CSS). For more details on choosing appropriate hardware for your deployment, see the chapter on [Sizing, page 14-1](#).