Gateway Options

PSTN Gateway

In this type of deployment, the Ingress Voice Gateway is used as the PSTN Voice Gateway. The Ingress Voice Gateway is responsible for converting TDM speech to IP and for recognizing DTMF digits and converting them to RFC2833 events.

Note Unified CVP does not support passing SIP-Notify DTMF events.

In a centralized Unified CVP deployment, you can separate the VoiceXML functionality from the Ingress Voice Gateway to provide a separate PSTN ingress layer. The separate PSTN layer and VoiceXML enable the deployment to support many VoiceXML sessions and PSTN interfaces. For example, the Cisco AS5400XM can accept a DS3 connection, and support up to 648 DSOs. However, a gateway that is handling that many ingress calls cannot also support as many VoiceXML sessions. In such cases, the VoiceXML sessions are off-loaded to a separate farm of only VoiceXML Gateways, such as Cisco VVB.

Note You can use any TDM interface, supported by Cisco IOS gateway and Cisco IOS version, and compatible with Unified CVP.

VoiceXML Gateway with DTMF or ASR/TTS

The VoiceXML Gateway allows you to interact with the VoiceXML browser through DTMF tones or ASR/TTS. Because the gateway does not have PSTN interfaces, voice traffic is sent using Real-Time Transport Protocol (RTP) to the VoiceXML Gateway, and the RFC 2833 uses in-band signaling in RTP packets to transmit DTMF tones. An VoiceXML with DTMF or ASR and TTS allows you to increase the scale of the deployment and support hundreds of VoiceXML sessions.

In a centralized Unified CVP deployment, you could use a VoiceXML farm. For example, if you want to support 300 to 10,000 or more VoiceXML sessions, use the Cisco AS5350XM Gateway. The standalone AS5350XM can support many DTMF or ASR/TTS VoiceXML sessions per Voice Gateway. In addition, stack the AS5350XM Gateways to support large VoiceXML IVR farms. However, for better performance and higher capacity, and to avoid the need for stacking, you can use the 3945 or 3945-E Series Gateways.

In a distributed Unified CVP deployment, consider providing an extra layer of redundancy at the branch office. You can deploy a separate PSTN Gateway and a VoiceXML Gateway to provide an extra layer of redundancy. In addition, for a centralized Cisco Unified Communications Manager deployment, you must support Survivable Remote Site Telephony (SRST). The Cisco 2800 Series and 3800 Series Integrated Service routers and the newer 2900 Series and 3900 Series routers are the best choices for the Ingress Voice Gateway because they support SRST.

For a discussion of the advantages and disadvantages of each codec, see Voice Traffic.

VoiceXML and PSTN Gateway with DTMF or ASR/TTS

The most popular Ingress Voice Gateway is the combination VoiceXML Gateway and PSTN Interface Gateway. For a centralized Cisco Unified Communications Manager deployment, Survivable Remote Site Telephony (SRST) must be supported. The Cisco 2800 Series, 3800 Series Integrated Service routers, 2900 Series, and 3900 Series routers are the best choices for the Ingress Voice Gateway, because they support SRST.

TDM Interfaces

The Cisco AS5400XM Universal Gateway offers unparalleled capacity in only 2 rack units (2 RUs) and provides best-of-class voice, fax, and remote-access services.

The Cisco AS5350XM Universal Gateway is the one rack-unit (1 RU) gateway that supports 2-, 4-,8-, or 16-port T1/12-port E1 configurations and provides universal port data, voice, and fax services on any port at any time. The Cisco AS5350XM Universal Gateway offers high performance and high reliability in a compact, modular design. This cost-effective platform is ideally suited for internet service providers (ISPs) and enterprise companies that require innovative universal services.

For the most current information about the various digital (T1/E1) and analog interfaces supported by the various voice gateways, see the latest product documentation available at the following sites:

- Cisco 2800 Series

  http://www.cisco.com/en/US/products/ps5854/tsd_products_support_series_home.html

- Cisco 3800 Series

http://www.cisco.com/en/US/products/ps5855/tsd_products_support_series_home.html

- Cisco AS5300

  http://www.cisco.com/en/US/products/hw/univgate/ps501/tsd_products_support_series_home.html

- Cisco 2900 Series
- http://www.cisco.com/en/US/products/ps10537/index.htm.
- Cisco 3900 Series
- http://www.cisco.com/en/US/products/ps10536/index.htm

Cisco Unified Border Element

The Cisco Unified Border Element (CUBE), formerly known as the Cisco Multiservice IP-to-IP Gateway is a session border controller (SBC) that provides connectivity between IP voice networks using SIP. The CUBE is supported in flow-through mode only so that all calls are routed through the CUBE.

Note Unlike flow-through mode, with flow-around mode you lose the ability to do DTMF interworking, transcoding, and other key functions, such as phone and media capabilities.

A Unified Border Element is needed when replacing a TDM voice circuit with an IP voice trunk, such as a SIP trunk, from a phone company. The CUBE serves as a feature demarcation point for connecting enterprises to service providers over IP voice trunks.

Note For outbound calls, CUBE supports Call Progress Analysis (CPA) on TDM circuits.

The CUBE has been tested with, and can be used in, any of the following scenarios:

- SIP-to-SIP connectivity between a third-party SIP device and Cisco Unified CVP over the SIP trunks certified by Cisco.
- SIP-to-SIP connectivity between Cisco Unified Communications Manager and Cisco Unified CVP.
- Coresidency of VoiceXML Gateway and CUBE for any of the above scenarios but with the limitation that the call flow does not work when the configurations listed here occur at the same time on the CUBE:
  - Survivability TCL script and incoming translation rules are configured under the same incoming dial-peer.
  - Transcoding between G.711 and G.729.
  - Header-passing between the call legs is enabled globally.

For CUBE session numbers, refer to:

http://www.cisco.com/en/US/prod/collateral/voicesw/ps6790/gatecont/ps5640/order_guide_c07_462222.html

For more information about using the CUBE with Unified CVP, including topologies and configurations, see *Cisco Unified Border Element for Contact Center Solutions* available at:

http://cisco.com/en/US/docs/voice_ip_comm/unified_communications/cubecc.html

Note Due to a limitation in Cisco IOS, the CUBE does not support midcall escalation or deescalation from audio to video, and conversely.

Using a SIP Trunk Without CUBE

When you connect to a third-party SIP device, including a SIP PSTN service provider, use a CUBE. If you do not place a CUBE between Unified CVP and the SIP device, ensure that both sides are compatible with thorough integration testing.

When connecting to a PSTN SIP Trunking service without a CUBE, carefully consider how to secure the connection between the contact center and the service provider. Also consider how to accomplish NAT and address hiding. Otherwise, the service-provider network can have full access to the contact center network. As the service-provider interconnect interface provided by Cisco, CUBE addresses both of these concerns.

Using Cisco ASR 1000 Series as Unified Border Element

Unified CVP supports Cisco IOS XE software Release 3.3.0S Enterprise with the following limitations:

- ASR 1000 Series do not support VXML. As a result, the VRU leg of the call must be routed to a separate VoiceXML Gateway. You must not use the **Send To Originator** setting on the CVP Call Server to route the IVR leg of the call back to the originating ASR CUBE Gateway, and standalone CVP calls must be routed to a separate VoiceXML Gateway.
- The global **Pass Thru SDP** setting on the ASR 1000 Series gateways is not supported with CVP deployments.
- ASR 1000 Series gateways do not support the TCP transport with SIP signaling when using the box-box hardware redundancy feature. The UDP transport is supported when failing the active ASR chassis to the standby chassis. It is important to note that the default TCP setting will not work with failover in this version of the ASR release. Therefore, UDP must be used on both the incoming and outgoing legs of the ASR CUBE for uninterrupted call control with CVP.
- Using the proxy servers to perform UDP to TCP Up-Conversion when receiving large size packet SIP messages, in a scenario where the proxy is in front of the ASR session border controller, the proxy servers should be turned off to ensure that UDP transport is used for the connection on the inbound call. Typically, a proxy server is positioned behind the session border controller in the deployment.
- A **sip-profile** configuration is needed on ASR 1000 Series for the courtesy callback feature only when deploying an IOS-XE version affected by CSCts00930. For more information on the defect, access the Bug Search Tool at https://sso.cisco.com/autho/forms/CDClogin.html. To configure the **sip-profile**, the following must be added:

  voice class sip-profiles 103

  request INVITE sip-header Call-Info add "X-Cisco-CCBProbe: <ccb param>"

ccb param is the **ccb** parameter defined in the survivability service. Add this sip-profile to the outgoing dial-peer to the CVP.

The following is a configuration example:

voice class sip-profiles 103

hoigogpoupcoioivc9iu i 8s66d8 8hxiciuvyd78zicvc8ayge

request INVITE sip-header Call-Info add "X-Cisco-CCBProbe:

id:192.168.1.50;loc:testbed04;trunks:10"

application

service survivability flash:survivability.tcl

param ccb id:192.168.1.52;loc:testbed04;trunks:10

dial-peer voice 700051 voip

description Comprehensive outbound route to CVP

destination-pattern 7000200T

session protocol sipv2

session target ipv4:192.168.1.20:5060

dtmf-relay rtp-nte

voice-class sip profiles 103

codec g711ulaw

no vad

- The following Survivability.tcl options are not applicable for use on the ASR because they are traditionally for POTS dial peers:
  - ani-dnis-split.
  - takeback-method.
  - -- *8.
  - -- hf.
  - icm-tbct.
  - digital-fxo.
- The following Survivability.tcl options are not supported: aa-name, standalone, and standalone-isntime.
  - The aa-name option is not supported because CME auto-attendant service is not supported on ASR.
  - The standalone and standalone-isntime options are not supported because there is no support for VXML on ASR.
- Due to ASR limitations, the following features are not supported:
  - Refer with Re-query
  - Legacy Transfer Connect using DTMF *8 label

- ASR 1000 does not terminate the TDM trunks. Therefore, the following TDM Gateway features do not apply to ASR 1000:
  - PSTN Gateway trunk and DS0 information for SIP calls to ICM
  - Resource Availability Indication (RAI) of DS0 trunk resources via SIP OPTIONS message to ICM

Note Because ASR 1000 represents the introduction of new equipment, to ensure success of ASR 1000 deployments, any UCCE/CVP contact center integration that utilizes the ASR 1000 requires an Assessment to Quality (A2Q) review. This review is required for new UCCE customers as well as existing UCCE customers who want to upgrade to the ASR 1000.

Note The Courtesy Call Back call flow does not work if ASR as CUBE is configured for the media flow-around instead of the media flow-through.

Using Cisco ISR as Unified Border Element

Unified CVP supports ISR with the following limitations:

- A **sip-profile** configuration is needed on ISR for the courtesy callback feature only when deploying an IOS-XE version affected by CSCts00930. For more information on the defect, access the Bug Search Tool at https://sso.cisco.com/autho/forms/CDClogin.html. To configure the **sip-profile**, the following must be added:

  voice class sip-profiles 103

  request INVITE sip-header Call-Info add "X-Cisco-CCBProbe: <ccb param>"

  ccb param is the “ccb” parameter defined in the survivability service. Add this **sip-profile** to the outgoing dial peer to the CVP.

  The following is a configuration example:

  voice class sip-profiles 103

  request INVITE sip-header Call-Info add "X-Cisco-CCBProbe:

  id:192.168.1.50;loc:testbed04;trunks:10"

  application

  service survivability flash:survivability.tcl

  param ccb id:192.168.1.52;loc:testbed04;trunks:10

  dial-peer voice 700051 voip

  description Comprehensive outbound route to CVP

  destination-pattern 7000200T

  session protocol sipv2

  session target ipv4:192.168.1.20:5060

dtmf-relay rtp-nte

voice-class sip profiles 103

codec g711ulaw

no vad

Note

- For ISR versions, see the *Hardware and System Software Specification for Cisco Unified Customer Voice Portal* at http://www.cisco.com/c/en/us/support/customer-collaboration/unified-customer-voice-portal/products-technical-reference-list.html.
- The Courtesy Call Back call flow does not work if ISR as CUBE is configured for the media flow-around instead of the media flow-through.

Mixed G.729 and G.711 Codec Support

Transcoders (DSPs) are required if the two endpoints participating in the call cannot negotiate a common codec. Therefore, midcall codec negotiation greatly reduces the need for transcoders.

CVP supports mixed G.711 and G.729 codecs in Standalone and Comprehensive SIP deployments with Cisco Unified Border Element Enterprise Edition (CUBE) and Cisco Unified Communications Manager (Unified CM). Calls that are ingressed through a SIP trunk from the carrier to a CUBE require Cisco IOS 15.1(2)T or later for mixed codec support. You can use any combination of codecs on the legs of a call. For example, a caller can place a call using the G.729 codec, hear an IVR prompt played using the G.711 codec, be transferred to the first Agent using the G.729 codec, and then transferred to the second agent using the G.711 codec.

A typical use case where transcoders may be required is when phones in a WAN connected location only support the G729 codec, and CVP is set up for G711 support. In this case, when these phones call into CVP, Unified Communications Manager engages transcoders. For inbound calls that arrive from a gateway or CUBE can start with G711 at CVP, then later renegotiate to G729 with the agents without the need for transcoders.

Transcoders (DSPs) are controlled by CUBE and Unified Communications Manager depending on the call flow. Because most of the service providers support midcall codec negotiation, transcoders in CUBE are not necessary. You commonly need transcoders controlled by Unified Communications Manager to support call flows, in which the phone supporting G729 is calling into CVP supporting G711.

Generate G729 Prompts for Unified CVP

To generate the G.729 prompts for Unified CVP, perform the following procedure:

- Convert the audio files from G.711 to G.729 format using the Music on Hold (MOH) audio translator.
- Change the G.729 compression identifier in the file header.

Convert the Audio Files from G.711 to G.729 Format

Procedure

Step 1 Log in to the Cisco Unified CM Administration portal and select **Media Resources** > **MOH Audio File Management**.

Step 2 Click **Upload File** and select the G.711 audio files individually.

Step 3 Click **Media Resources** > **MOH Audio File Management** and check whether the audio files have been converted to G.729 format. If the conversion was successful, the recording length of audio files has a nonzero value.

Step 4 Copy the converted audio files to your Windows server using the Secure File Transfer Protocol (SFTP) Server.

Note Do not add spaces when you rename the audio files.

Step 5 Use putty to sign in to the Unified Communications Manager Server as an administrator.

Step 6 From the command prompt, run **file get activelog mohprep/*g729.wav** and provide the SFTP prompts.

Change the G.729 Compression Identifier in the File Header

The G.729 files that the Unified Communications Manager generates have a non-standard compression codec tag in the file header. The VXML Gateway cannot play these audio files, as it does not recognize the codec type. Change the compression codec type value to convert the audio files into the standard G729r8 format.

Use the following procedure to change the compression codec type number in the file header from 0x0133 to the standard 0x14db, G729r8 format.

Procedure

Step 1 Create a folder in the Unified CVP directory. Copy the G.729 audio files that have a nonstandard compression codec tag in the file header into the new folder location.

Step 2 From the command prompt, navigate to the `C:\Cisco\CVP\bin` folder.

Step 3 Perform one of these steps:

- To convert audio files individually, from the command prompt, run **<UCMHeaderFixer.exe Audio file Name>*.***.
- To perform bulk conversion of audio files, from the command prompt, run **UCMHeaderFixer.exe Folder Path**.

The script runs and the audio file is converted from name.g729.wav file into name.wav format.

Step 4 Use the Operations Console to upload the converted audio files to the IOS Gateway.

ISO Gateway Choices

Unified CVP uses IOS Gateways for two purposes: TDM ingress and VoiceXML rendering. Any Cisco gateway supported by Unified CVP can be used for either purpose or both. However, depending on your deployment model, you can use one of the following functions:

- Model #1: Standalone Self-Service

  All calls use both ingress and VoiceXML.

- Model #2: Call Director

  All calls use ingress only.

- Model #3a: Comprehensive Using Unified ICM Micro-Apps

  All calls use ingress, and some calls use VoiceXML.

- Model #3b: Comprehensive Using Unified CVP VXML Server

  All calls use ingress, and some calls use VoiceXML.

- Model #4: VRU Only with NIC Controlled Routing

  All calls use both ingress and VoiceXML.

In cases where both Ingress and VoiceXML are required, you can choose to run both functions on the same gateways or you can choose to designate some gateways for ingress and others for VoiceXML. Use the following guidelines to determine whether the functions should be combined or split:

- In classical branch office deployments, where the call is queued at the branch where it arrived, ingress and VoiceXML functions must always be combined.
- In cases where many non-CVP PSTN connections share the gateways, it is submitted to dedicated Ingress for that purpose, and use separate VoiceXML Gateways.
- VoiceXML-only gateways are less costly because they do not require DSP farms or TDM cards. Use a spreadsheet to determine which way you obtain the best price.
- With relatively low call volume, it is usually better to combine the functions for redundancy purposes. Two combined gateways are better than one of each because the loss of one gateway still allows calls to be processed, though at a lower capacity.

The next decision is whether to use Cisco Integrated Service Router (ISR) Gateways (Cisco 2800 or 3800 series routers), ISR-G2 (2900 or 3900 Series routers), or the Cisco AS5x00 Series Gateways. ISR Gateways are used only in branch office sites and AS5x00 Series Gateways are used in centralized data center sites.

You might sometimes have difficulty determining what constitutes a branch office, and which gateway is used. The following guidelines can help with that determination:

- The classical definition of branch offices, for which you must use ISR Gateways, includes:
  - Multiple sites where TDM calls arrive from the PSTN.
  - Those sites are separated from the data centers where most of the Unified CVP equipment resides.
  - One gateway is used at each site.

- If you have sites where you are stacking multiple gateways for any reason, then those sites are data center sites and should use Cisco AS5x00 Series Gateways.

For more information on the Cisco AS5x00 Series Gateways, refer to the technical specifications available at http://www.cisco.com/en/US/products/hw/univgate/ps501/index.html.

For more information on the Cisco Integrated Service Routers (ISRs), refer to the documentation available at http://www.cisco.com/en/US/products/hw/routers/index.html.

IOS Gateway Sizing

Individual Cisco gateways can handle various call capacities depending on whether they are doing ingress only, VoiceXML only, or a combination of the two. IOS Voice Gateways doing VoiceXML activities also have different call capacities depending on whether they are supporting ASR or TTS activities, and on the type of VoiceXML application being executed. For instance, an intensive JavaScript application reduces call capacity. Gateways using HTTPS, have lower call capacity as compared to HTTP.

In general, gateways that perform ingress only can be sized according to the number of TDM cables that can be connected to them. For gateways that are combined or VoiceXML-only, it is important to ensure that the overall CPU usage is less than 75 percent on average. The numbers in the Maximum Number of VoiceXML Session tables are based on Unified CVP VoiceXML documents; other applications that generate more complex VoiceXML documents have a higher impact on performance. The following factors affect CPU usage:

- Calls per second (CPS)
- Maximum concurrent calls
- Maximum concurrent VoiceXML sessions

Before sizing the IOS Voice Gateways, use the Unified CCE Resource Calculator to determine the maximum number of trunks (DS0s) and VoiceXML IVR ports required to support the entire solution.

For almost all Unified CVP deployment models, sizing is based on the maximum number of concurrent VoiceXML sessions and VoIP calls. The following tables list this information for different versions of Cisco IOS.

Note The performance numbers listed in the Table 9, Table 10, and Table 11 are equivalent for MRCPv1 and MRCPv2.

Table 1: Maximum Number of VoiceXML Sessions Supported by Cisco Voice Gateways (Cisco IOS Release 15.1.4.M7 and Later)

VoiceXML Gateway CPU Capacity for Cisco IOS Release 15.1.4.M7 or Later					
Platform	**VXML Only**		**VXML + PSTN**		**Memory**
	DTMF	**ASR**	**DTMF**	**ASR**	**Recommended**
5000XM	200	135	155	104	512 MB
3825	130	85	102	68	512 MB

VoiceXML Gateway CPU Capacity for Cisco IOS Release 15.1.4.M7 or Later					
Platform	**VXML Only**		**VXML + PSTN**		**Memory**
	DTMF	**ASR**	**DTMF**	**ASR**	**Recommended**
3845	160	105	125	83	512 MB
2901	12	8	9	6	2 GB
2911	60	40	47	31	2 GB
2921	90	60	71	48	2 GB
2951	120	80	95	64	2 GB
3925	240	160	190	127	2 GB
3945	340	228	270	180	2 GB
3925E	475	450	380	375	2 GB
3945E	580	550	460	450	2 GB
Based on ISO 15.1.4.M7, G.711, basic calls, Ethernet egress, CPU NTE 75% (5000XM 80%)					

Table 2: Maximum Number of VoiceXML Sessions Supported by Cisco Voice Gateways Executing Intensive JavaScript Applications (Cisco IOS Release 15.1.4.M7 and Later)

Cisco Voice Gateway Platform	Dedicated VoiceXML Gateway		Voice Gateway and VoiceXML		
	VoiceXML and DTMF	**VoiceXML and ASR/TTS**	**VoiceXML and DTMF**	**VoiceXML and ASR/TTS**	**Memory Recommended**
AS5350XM	105	85	110	70	512 MB (default)
AS5400XM	105	85	110	70	512 MB (default)

Table 3: Maximum Number of VoiceXML Sessions Supported by Cisco Voice Gateways Using HTTPS (Cisco IOS Release 15.1.4.M7 and Later)

Cisco Voice Gateway Platform	Dedicated VoiceXML Gateway		Voice Gateway and VoiceXML		
	VoiceXML and DTMF	**VoiceXML and ASR/TTS**	**VoiceXML and DTMF**	**VoiceXML and ASR/TTS**	**Memory Recommended**
3945E	510	342	408	270	2 GB
AS5350XM	155	120	138	95	512 MB (default)
AS5400XM	155	120	138	95	512 MB (default)

Note The performance numbers listed in Table 11 are only for selected models of Cisco Voice Gateways using HTTPS. Use the HTTPS performance numbers of the 3945E router, to estimate the performance numbers for router models that are not listed in Table 11.

Note Performance numbers for the Cisco 3825 Series and 3845 Series Integrated Services Routers (ISRs) are higher when the Ingress Voice Gateway and the VoiceXML Gateway functions reside on the same router (coresident deployment). When the call is connected to the VoiceXML Gateway from the Ingress Voice Gateway, the media flows directly between the two. In a coresident deployment, the gateway does not have to spend CPU cycles to packetize and de-packetize the RTP packets. Hence, by saving these CPU cycles, the gateway can support increased VoiceXML sessions.

This note does **not** apply to Cisco IOS Release 15.0.1M and Cisco IOS 15.1.4.M7.

The numbers in Table 9, Table 10, and Table 11 assume that the only activities running on the gateway are VXML with basic routing and IP connectivity. If you intend to run extra applications such as fax, security, and normal business calls, then the capacity numbers presented here should be prorated accordingly. The numbers mentioned in the Voice Gateway and VoiceXML column mean that the indicated number of VoiceXML sessions and voice calls can be supported simultaneous on the same gateway. For example, in Table 9 the 500XM can terminate a maximum of 200 PSTN calls, and those 200 PSTN calls could have 200 corresponding VoiceXML sessions at the same time.

The numbers represent performance with scripts generated by Unified CVP Studio running on the Unified CVP VXML Server. Other VoiceXML applications might perform differently. These figures apply if the CPU utilization does not exceed more than 75 percent Voice Activity Detection (VAD) is turned off, and your system is running VoiceXML v2.0 and MRCP v2 with Cisco IOS Release 15.1.4.M7 and later.

Note These performance numbers are accurate when used with either the Cisco Call Server or Cisco Unified CVP VXML Server. Performance can, and often does, vary with different applications. Performance from external VoiceXML applications (such as Nuance OSDMs) might not be representative of the performance when interoperating with non-Cisco applications. Ensure that the CPU usage is less than 75 percent on average and that adequate memory is available on Cisco gateways at full load when running external VoiceXML applications. Contact the application provider of the desired VoiceXML application for performance and availability information. External VoiceXML applications are not provided by Cisco, and Cisco makes no claims or warranties regarding the performance, stability, or feature capabilities of the application when interoperating in a Cisco environment.

Note Cisco does not specifically test or qualify mixes of traffic because there are infinite combinations. All numbers should be seen as guidelines only and varies from one implementation to the next based on configurations and traffic patterns. The systems are required to be engineered for worst-case traffic (all ASR) if the types of calls that are offered to the VoiceXML Gateway are not known or cannot be predicted.

If you run VoiceXML on one of the Cisco 2900 and 3900 Series gateways, more licenses (FL-VXML-1 or FL-VXML-12) are required.

Consult the following links to ensure that the concurrent call load and call arrival rates do not exceed the listed capacities:

- Model comparison:

  http://www.cisco.com/c/en/us/support/customer-collaboration/unified-customer-voice-portal/products-implementation-design-guides-list.html.

- Gateway sizing for Contact Center traffic:

  http://www.cisco.com/c/en/us/support/unified-communications/unified-communications-manager-callmanager/products-implementation-design-guides-list.html.

Also consider how much DRAM and flash memory to order. The capacity that comes with the machine by default is usually sufficient for most purposes. However, if your application requires large numbers of distinct .wav files (as with complex self-service applications) or if your application has unusually large .wav files (as with extended voice messages or music files), you might want to increase the amount of DRAM in order to expand your flash memory. The use of DRAM for prompt caching is discussed in detail in the chapter on Media File Options.

Note HTTP cache can only be extended to 100 MB in the current Cisco IOS releases.

Cisco VVB Sizing

Cisco VVB has call capacities is defined based on the call support for ASR or TTS activities and on the type of VoiceXML application being executed. For instance, an intensive JavaScript application reduces call capacity and VVB using HTTPS has lower call capacity as compared to HTTP.

It is important to ensure that the overall CPU usage is less than 75 percent on average. The numbers in the Maximum Number of VoiceXML Sessions tables is based on VoiceXML documents; other applications that generate more complex VoiceXML documents have a higher impact on performance. The following factors affect CPU usage:

- Calls per second (cps)
- Maximum concurrent VoiceXML sessions
- Complexity of VoiceXML applications

Before sizing the Cisco VVB, use the Unified CCE Resource Calculator to determine the maximum number of trunks (DS0s) and VoiceXML IVR ports, which is required to support the entire solution.

For almost all Unified CVP deployment models, sizing is based on the maximum number of concurrent VoiceXML sessions and VoIP calls.

Note The following performance numbers listed in ASR and TTS columns are applicable only for MRCP v1.

Table 4: Maximum Number of VoiceXML Sessions Supported by Cisco VVB

CPU Capacity for Cisco VVB					
System Specification	**DTMF**	**ASR**	**TTS**	**HTTPS**	**Standalone java script**
4 CPU, 8-GB RAM	600	400	400	480	200

Note

- Table 11 displays the numbers that represent the performance when the activities running on the gateway are only VXML with basic routing and IP connectivity. If you intend to run more activities such as security, call tracing, monitoring, then the capacity numbers presented here should be prorated accordingly.
- The numbers also represent performance with VoiceXML pages generated by Unified CVP Call Studio applications running on the Unified CVP VXML Server. Other VoiceXML applications might perform differently. These figures apply if the CPU utilization does not exceed more than 75 percent and your system is running VoiceXML v2.0 and MRCP v1.

  These performance numbers are accurate when used with Cisco Unified CVP VXML Server and moderately complex VoiceXML applications. Performance can, and often does, vary with different applications. Performance from external VoiceXML applications (such as Nuance OSDMs) is not representative of the performance when interoperating with non-Cisco applications. Ensure that the CPU usage is less than 75 percent on average and that adequate memory is available on Cisco VVB at full load when running external VoiceXML applications. Contact the application provider of the desired VoiceXML application for performance and availability information.
- External VoiceXML applications are not provided by Cisco, and Cisco makes no claims or warranties regarding the performance, stability, or feature capabilities of the application when interoperating in a Cisco environment.
- HTTP cache can be extended to 2 GB in the current Cisco VVB releases.

Using MGCP Gateways

Cisco Unified CVP requires the deployment of a SIP Gateway. However, you require the use of MGCP 0.1 voice gateways with Unified CM deployments for purposes of overlap sending, NSF, and Q.SIG support. The following design considerations apply to deploying Cisco Unified CVP in this environment:

- Design and plan a phased migration of each MGCP voice gateway to SIP.
- Implement both MGCP 0.1 and SIP.

  Because of the way MGCP works, a PSTN interface using MGCP can be used for MGCP only. If you want to use MGCP for regular Unified CM calls and SIP for Unified CVP calls, you need two PSTN circuits.
- Deploy a second SIP voice gateway at each Unified CVP location.
- Send calls through the Unified CM to Unified CVP.

When sending calls through Unified CM to Unified CVP, the following guidelines apply:

- The Unified CVP survivability.tcl script cannot be used in this solution. If the remote site is disconnected from the central site, the call is dropped.
- There is an additional hit on the performance of Unified CM. This is because, in a normal Unified CVP deployment, Unified CM resources are not used until the call is sent to the agent. In this model, Unified CM resources are used for all calls to Unified CVP, even if they terminate in self-service. This is in addition to the calls that are extended to agents. If all calls are eventually extended to an agent, the performance impact on Unified CM is approximately double that of a normal Unified CVP deployment. This factor alone typically limits this scenario to small call centers only.
- In order to queue calls at the edge, use the **sigdigits** feature in Unified CVP to ensure that the calls are queued at the appropriate site or VoiceXML Gateway. For more information on how the **sigdigits** feature works, see the chapters on Distributed Deployment and Unified CVP Design for High Availability.

Note The Cisco Unified CVP provides the flexibility to add, modify, remove, or deploy Unified CVP in many scenarios to facilitate interoperability with third-party devices. Not all SIP service providers support advanced features such as REFER, 302 Redirect Messages, DTMF-based take-back-and-transfer, or data transport (UUI, GTD, NSS, and so forth). Refer to the interoperability note available at the following location for information on the interoperability support for SBC when deployed in place of Cisco CUBE, http://www.cisco.com/en/US/solutions/ns340/ns414/ns728/voice_portal.html