



Non-Reference Designs

- [Configuration Limits and Scalability Constraints for Non-Reference Designs](#), on page 1
- [Noncompliant Topologies](#), on page 9
- [Non-Reference Design Core Components and Implementations](#), on page 29

Configuration Limits and Scalability Constraints for Non-Reference Designs

The following tables specify the configuration limits and scalability constraints for the Unified ICM/CCE products. These configuration limits are part of the Unified ICM/CCE product design constraints and were used for system sizing characteristics as tested by Cisco. Most of these system parameters (or combinations of these system parameters) form contribution factors that affect system capacity.

These limits apply to Non-Reference Designs, only. Do not use these limits to size a Reference Design solution. If you are not using the standard three coresident PG layout of one Agent PG, one VRU PG, and one MR PG, your design is a Non-Reference Design. The higher configuration limits that are listed for Reference Designs do not apply without the standard PG layout.

When you design your contact center, ensure that your design is deployed within these limits. (See the comments in the following table for more information.) Consult Cisco if you have special configuration requirements that might exceed specific parameters.



Important

This information serves as a quick reference. Check the *Virtualization for Unified Contact Center Enterprise* at http://www.cisco.com/c/dam/en/us/td/docs/voice_ip_comm/uc_system/virtualization/virtualization-unified-contact-center-enterprise.html and the *Contact Center Enterprise Compatibility Matrix* at <https://www.cisco.com/c/en/us/support/customer-collaboration/unified-contact-center-enterprise/products-device-support-tables-list.html> for more information on system constraints.

The compatibility matrix specifies all supported configurations and versions for Cisco Unified Contact Center Enterprise Release 10.0. The information in the compatibility matrix supersedes compatibility information in any other Cisco Unified Contact Center Enterprise documentation. If a configuration or version is not stated in the compatibility matrix, that configuration or version is not supported.

The check mark in the table indicates that a given parameter is applicable to the indicated Unified ICM/CCE product edition. See the notes at the end of this table.

Table 1: Configuration Limits and Scalability Constraints

Maximum Limit	Limit Value		Applies to		Comments
	<=450 agents (Unified CCE only)	>450 - <=12,000 agents	Unified CCE	Unified ICM	
Administrators (Users)	100	1000	✓	✓	Includes setup, config, and scripting users.
ECC (Extended Context Call) and User variables size (bytes)	2000	2000	✓	✓	Unified CVP and Outbound Option rely on a subset of this maximum limit for integration with Unified ICM. The maximum that is indicated is independent from the number of ECC and user variables used, with each representing approximately 50-bytes extra storage per record. The maximum includes both persistent and nonpersistent variables.
Number of Peripheral Variables (Call Variables)	10	10	✓	✓	
Peripheral Variable length (characters)	40	40	✓	✓	40 characters, excluding terminating NULL.
VRU PIMs on each VRU PG	N/A	10	✓	✓	Each CVP supports up to 3000 ports and 15 CPS, and each VRU PG can support a maximum of 12000 ports. Up to 10 VRU PIMs per VRU PG can be supported as long as the total ports on the VRU PG are less than 12000, and the total CPS across all the VRU PIMs is less than 60 CPS.
Maximum CPS per VRU PIM	N/A	15	✓	✓	Maximum CPS per VRU PG is 60.
VRU PIMs on each Generic PG	2	4	✓	✓	A Generic PG with Unified CM (CUCM) PIM only supports 1000 Ports total.
VRU PIMs on each System PG	N/A	5	✓	—	IP-IVR PIMs only.
TDM PIMs on each PG	N/A	5	—	✓	Multiple PIMs on a PG affect performance. Compared to a single PIM on each PG, multiple PIMs lower the total number of agents, VRU ports, and supported call volume. There is a maximum of one PIM on each TDM PG with CTI OS coresident.
MR PIMs on each MR PG	1	2	✓	✓	
UCM PIMs	1	12	✓	✓	

Maximum Limit	Limit Value		Applies to		Comments
	<=450 agents (Unified CCE only)	>450 - <=12,000 agents	Unified CCE	Unified ICM	
Maximum Number of PGs for each CUCM Cluster	4	4	✓	✓	
Duplex PGs on each ICM instance	1	150	✓	✓	For deployments with <= 450 agents and Cisco Outbound Option, there is another MR PG. See UCM limit above.
PIMs on each system (total)	4	150	✓	✓	One agent PIM, two VRU PIMs, and one MR PIM (applies to >= 450 agents only).
Configured agents on each system (total)	N/A	65,000	—	✓	
Configured agents on each system (total)	3000	76,000	✓	—	
Configured agents on each peripheral	3000	12,000	✓	✓	
Skill groups on each peripheral gateway	1500	4000	✓	✓	Configuration of precision queues creates a skill group per agent PG which counts toward the supported number of skill groups on each PG.
Skill groups on each system	1500	27,000	✓		Configuration of precision queues creates a skill group per agent PG which counts toward the supported number of skill groups on each system.
CTI OS on each Agent PIM	1	1	✓	✓	
Provisioning operations on each hour	30	120	✓	✓	For Configuration Manager, CCMP or AAS – maximum number of save operations across all ADSs in the solution in a 1-hour period. 200 changes per provisioning operation.
Maximum agents tracked in Agent State Trace		100	✓	✓	For more information on Agent State Trace, see the Configuration Guide for Cisco Unified ICM/Contact Center Enterprise at http://www.cisco.com/c/en/us/support/customer-collaboration/unified-contact-center-enterprise/products-installation-and-configuration-guides-list.html
SIP dialer ports on each dialer	N/A	1500	✓	—	This limit assumes that the model is distributed, numbers vary based on deployment.
SIP dialers on each PG pair (Side A + Side B)	N/A	1	✓	—	Only one dialer type can be installed per PG.

Configuration Limits and Scalability Constraints for Non-Reference Designs

Maximum Limit	Limit Value		Applies to		Comments
	<=450 agents (Unified CCE only)	>450 - <=12,000 agents	Unified CCE	Unified ICM	
SIP dialer Ports on each server (total)	N/A	1500	—	—	In multi-instance deployment.
Dialer ports on each system (total)	N/A	4000	✓	—	
Dialers on each system (total)	N/A	32	✓	—	
Campaigns on each system	N/A	600	✓	—	Sizing is required. Smaller deployments cannot support the full 600 campaigns.
Campaigns skill groups on each system	N/A	100	✓	—	Total skill groups from all campaigns.
Campaign skill groups on each campaign	N/A	20	✓	—	Limitation on skill groups for any given campaign (as long as the maximum 100 campaign skill groups per system not exceeded).
Dialed numbers on each system	1500	240,000	✓		
Labels configured on each system	500	100,000 for up to 4000 agents 160,000 for up to 12,000 agents	✓		
Call type skill groups on each interval	1000	30,000	✓	✓	Total call type skill group records.
Configured call types	500	10,000	✓	✓	Total call types configured.
Active call types	250	8000	✓	✓	
Precision Routing (PR) Attributes on each system	N/A	10,000	✓	—	
PR Attributes for each Agent	N/A	50	✓	—	
PR Precision Queues (PQ) on each system	N/A	4000	✓	—	
PR PQ Steps on each system	N/A	10,000	✓	—	
PR Terms for each PQ Step	N/A	10	✓	—	
PR Steps for each PQ	N/A	10	✓	—	

Maximum Limit	Limit Value		Applies to		Comments
	<=450 agents (Unified CCE only)	>450 - <=12,000 agents	Unified CCE	Unified ICM	
PR Unique attributes for each PQ	N/A	10	✓	—	
Unique skill groups and PQs in a supervisor team	50	50	✓	✓	



Note Deployments close to the maximum number of configured agents on a system can show performance degradation and failed call routing, especially if contending capacity limitations also approach maximum thresholds. Expert assistance from partner or professional services is necessary for capacity-related system planning. Parameters that most impact performance with large numbers of configured agents includes total number of system peripherals, routes, number of active agents, and overall call load. The points at highest risk for degradation are busy hours and the half-hour update period, during which the PG sends report data to the Central Controller. System administrators can lessen the impact of these issues by purging unused configured agents, retiring inactive peripherals, and maintaining systems at current maintenance release levels.

Unified CCE Solution

Additional Sizing Factors

Many variables in the Unified CCE configuration and deployment options can affect the server requirements and capacities. This section describes the major sizing variables and how they affect the capacity of the various Unified CCE components.

Busy Hour Call Attempts (BHCA)

The number of calls attempted during a busy hour is an important metric. As BHCA increases, there is an increase in the load on all Unified CCE components, most notably on Unified CM, Unified IP IVR, and the Unified CM PG. The capacity numbers for agents assume up to 30 calls per hour per agent. If a deployment requires more than 30 calls per hour per agent, it decreases the maximum number of supported agents for the agent PG. Handle such occurrences on a case-by-case basis.

Agents

The number of agents is another important metric that impacts the performance of most Unified CCE server components, including Unified CM clusters.

Average Skill Groups or Precision Queues Per Agent

The number of skill groups or precision queues per agent (which is independent of the total number of skills per system) significantly affects the following:

- CTI OS servers
- Finesse servers

- Agent PG
- Call Router
- Logger

Limit the number of skill groups and precision queues per agent to 5 or fewer, when possible. Periodically remove unused skill groups or precision queues so that they do not affect the system's performance. You can also manage the effects on the CTI OS Server by increasing the value for the frequency of statistical updates.

The Finesse server does not display statistics for unused skill groups. Therefore, the number of skill groups that are assigned to agents affects the performance of the Finesse server more than the total number of skill groups configured.

Queue (skill group) statistics are updated on the Finesse Desktop at 10-second intervals. The Finesse Desktop also supports a fixed number of queue statistics fields. These fields cannot be changed.

The first table shows examples of the number of skill groups or precision queues (PQ) per agent affecting the capacity of the Unified CCE system. The table shows the capacity for each CTI OS instance. The Finesse server supports the same number of agents and skill groups as CTI OS.

Unified CCE supports a maximum of 50 unique skill groups across all agents on a supervisor's team, including the supervisor's own skill groups. If this number is exceeded, all skill groups monitored by the supervisor still appear on the supervisor desktop. However, exceeding this number can cause performance issues and is not supported.



Note Each precision queue that you configure creates a skill group per Agent PG and counts toward the supported number of skill groups per PG. The skill groups are created in the same Media Routing Domain as the precision queue.

The numbers in this table are subject to specific hardware and software requirements.

Table 2: Sizing Effects of Skill Groups or Precision Queues for Each Agent (12,000 Agents)

Avg Configured PQ or SG for each Agent	System	Generic PG Limits			
	Max Concurrent Agent for each System	Max Concurrent Agent for each PG	Max Configured PQ or SG for each PG	Max Configured VRU Ports for each PG	Max Configured VRU PIMs for each PG
5	12000	2000	4000	1000	4
10	11038	1832	4000	1000	4
15	10078	1663	4000	1000	4
20	9116	1495	4000	1000	4
25	8156	1326	4000	1000	4
30	7194	1158	4000	1000	4
35	6234	989	4000	1000	4
40	5272	820	4000	1000	4
45	4312	652	4000	1000	4

Avg Configured PQ or SG for each Agent	System	Generic PG Limits			
	Max Concurrent Agent for each System	Max Concurrent Agent for each PG	Max Configured PQ or SG for each PG	Max Configured VRU Ports for each PG	Max Configured VRU PIMs for each PG
50	3350	484	4000	1000	4



Note CTI OS monitor mode applications are supported only at 20 or lower skill groups per agent.

Supervisors and Teams

The number of supervisors and team members can also be a factor impacting the CTI OS Server performance. Distribute your agents and supervisors across multiple teams and have each supervisor monitor only a few agents.



Note Supervisors can monitor only agents within their own team, and all of the agents must be configured on the same peripheral.



Note You can add a maximum of 50 agents per team. You can add a maximum of ten supervisors per team.

A Unified CCE system can support a maximum of 50 agents per supervisor with the assumptions below. If a particular environment requires more than 50 agents per supervisor, then use the following formula to ensure that there is no impact to the CTI OS Server and Supervisor desktop. The most important factor in this calculation is the number of updates per second.

$$X = (Y * (N + 1) / R) + ((Z * N * A) / 3600), \text{ rounded up to the next integer}$$

Where:

X = Number of updates per second received by the CTI OS Supervisor desktop.

Y = Weighted Average of Number of Skill Groups or Precision Queues per Agents. For example, if total of 10 agents have the following skill group distribution: 9 have 1 skill group and 1 agent has 12 Skill Groups. The number of skills per agent ('Y') is, $Y = 90\% * 1 + 10\% * 12 = 2.1$. (The number of configured statistics in the CTI OS server is 17.)

Z = Calls per hour per agent.

A = Number of agent states. (Varies based on call flow; average = 10.)

N = Number of agents per supervisor.

R = The skill group or precision queue refresh rate configured on the CTI OS Server. (Default = 10 seconds.)

$(Y * (N + 1) / R)$ = Number of updates per second, based on skill groups.

$(Z * N * A) / 3600$ = Number of updates per second, based on calls.

The CTI OS Supervisor desktop is not impacted as long as there are fewer than 31 updates per second. This threshold value is derived by using the above formula to calculate the update rate for 50 agents per supervisor (N = 50), as follows:

$$X = (5 * (50 + 1) / 10) + ((30 * 50 * 10) / 3600) = 25.5 + 5 = 31 \text{ updates per second}$$

The maximum number of agents per supervisor must not exceed 200 for any given configuration, still holding updates per second to a maximum of 31 with the above formula.

CTI OS Monitor Mode Applications

A CTI OS Monitor Mode application can affect the performance of the CTI OS Server. CTI OS supports only two such applications per server pair. Depending on the filter specified, the impact on the CPU utilization might degrade the performance of the Agent PG.

Unified CM Silent Monitor

Each silently monitored call adds more processing for the PG and Unified CM. Each silently monitored call is equivalent to two unmonitored calls to an agent. Make sure that the percentage of the monitored calls is within the capabilities of PG scalability.

CTI OS Skill Group Statistics Refresh Rate

The skill group statistics refresh rate can also affect the performance of CTI OS Server. Cisco requires that you do not lower the refresh rate below the default value of 10 seconds.

Call Types

The call type is also an important metric that affects the performance of most Unified CCE server components. An increase in the number of transfers and conferences increase the load on the system which decreases the total capacity.

Queuing

The Unified IP IVR and Unified Customer Voice Portal (CVP) place calls in a queue and play announcements until an agent answers the call. For sizing purposes, it is important to know whether:

- The VRU will handle all calls initially (call treatment), and direct the callers to agents after a short queuing period.
- The agents will handle calls immediately, and the VRU queues only unanswered calls when all agents are busy.

The answer to this question determines very different VRU sizing requirements and affects the performance of the Call Router/Logger and Voice Response Unit (VRU) PG.

Translation Route Pool

Sizing the translation route pool depends on the expected call arrival rate. Use the following formula to size the translation route pool:

$$\text{Translation route pool} = 20 * (\text{Calls per second})$$

This calculation is specific to Unified CCE. For more general Unified ICM deployments, consult your Cisco Account Team or Partner.

UCCE Script Complexity

As the complexity and/or number of UCCE scripts increase, the processor and memory overhead on the Call Router and VRU PG increases significantly. The delay time between replaying RunExternalScript also has an impact.

Reporting

Real-time reporting can have a significant effect on Logger and Rogger processing due to database access. A separate VM is required for an Administration & Data Server to off-load reporting overhead from the Logger and Rogger.

VRU Script Complexity

As VRU script complexity increases with features such as database queries, the load placed on the IP IVR Server and the Router also increases. There is no good rule or benchmark to characterize the Unified IP IVR performance when used for complex scripting, complex database queries, or transaction-based usage. Test complex VRU configurations in a lab or pilot deployment to determine the response time of database queries under various BHCA and how they affect the processor and memory for the VRU server, PG, and Router.

Unified IP IVR Self-Service Applications

In deployments where the Unified IP IVR is also used for self-service applications, the self-service applications are in addition to the Unified CCE load. Factor self-service applications into the sizing requirements as stated in the sizing tables above.

Third-Party Database and Cisco Resource Manager Connectivity

Carefully examine connectivity of any Unified CCE solution component to an external device and/or software to determine the overall effect on the solution. Cisco Unified CCE solutions are flexible and customizable, but they can also be complex. Contact centers are often mission-critical, revenue-generating, and customer-facing operations. Therefore, engage a Cisco Partner (or Cisco Advanced Services) with the appropriate experience and certifications to help you design your Unified CCE solution.

Expanded Call Context (ECC)

The ECC usage impacts PG, Router, Logger, and network bandwidth. There are many ways that ECC can be configured and used. The capacity impact varies based on ECC configuration, handled on a case-by-case basis.

Noncompliant Topologies

Parent/Child Architecture

The Unified CCE Gateway PG allows Unified CCE to appear as a traditional ACD connected to the Unified ICM system. The Unified CCE Gateway PG provides the Unified ICM system with a PG that communicates with the CTI interface of the Unified CCE System PG.

In the parent/child model, you configure the child Unified CCE to function on its own. Unified CCE does not need a connection to the parent to route calls to agents. This independence provides local survivability for mission-critical contact centers during connection failures between the child and parent.

The child system can automatically send configuration objects to the parent Unified ICM for insertion into the Unified ICM configuration. This process eliminates the need to configure objects twice (on the local ACD and again on the Unified ICM). You can also turn off this functionality for situations where the customer does not want automatic configuration updates. For example, you do not want automatic updates from an outsourcer child system that also supports agents for another customer.

The Unified CCE Gateway PG can connect to a Unified CCE child that is using the Unified CCE System PG. If the child has multiple Unified CCE System PGs and peripherals, install and configure a separate Unified CCE Gateway PG in the parent system for each child PG. When deployed on a separate VM, a Unified CCE Gateway PG can manage multiple child Unified CCE peripherals.

In the Unified CCE child, you can deploy Unified IP IVR or Unified CVP for call treatment and queuing. If you deploy Unified CVP, configure another VRU PG. This model does not follow the single peripheral model used when Unified IP IVR is deployed. For this reason, information about calls queued at the child (and queue time of a call) is not available on the parent. Any computation involving queue time by the parent is inaccurate, which can lead to adverse impacts to routing on the parent.

Parent/Child Components

The following sections describe the components used in Unified ICM (parent) and Unified CCE (child) deployments.

Unified ICM (Parent) Main Site

The Unified ICM main site contains the Unified ICM Central Controller. The main site has a redundant pair of Central Controllers. A Central Controller has Call Router and Logger servers. You can deploy the servers as individual Call Routers and Loggers and deploy the servers in two geographically distributed sites for extra fault tolerance.

The Unified ICM Central Controllers control Peripheral Gateways (PGs) at the main site. A redundant pair of VRU PGs controls Unified CVP across the architecture. You can insert more PGs at this layer to control TDM or legacy ACDs and VRUs. This approach can support a migration to Unified CCE or support outsource locations that use the TDM or legacy ACDs.

The Unified ICM parent does not support any directly controlled agents in this model. So, in parent/child deployments, Unified ICM does not support classic Unified CCE with a Unified Communications Manager PG installed on this Unified ICM parent. Control all agents externally to this Unified ICM parent system.

The Unified CVP VRU PG pair controls the Unified CVP Server. The CVP Server translates the VRU PG commands from Unified ICM into VXML and directs the VXML to the Voice Gateways (VGs) at the child sites. Calls from the main site can come into the remote call centers under control of the Unified CVP at the parent location. The parent then has control over the entire network queue of calls across all sites. The parent holds the calls in queue on the VGs at the sites until an agent becomes available.

Unified CCE Call Center (Child) Site

The Unified CCE call center contains a Unified Communications Manager cluster for local IP-PBX functionality and call control for the IP phones. A local Unified IP IVR also provides local call queuing for the Unified CCE site. A redundant pair of Unified CCE Gateway PGs connects this site over the WAN to the Central Controller at the Unified ICM parent. You can deploy the Unified CCE Gateway PGs on separate VMs or co-incident with the CCE System PG with the following caveats:

- If the Unified CCE Gateway PG and Unified CCE System PG Instance Numbers are the same, then use different PG numbers for those PGs.

- If the Unified CCE Gateway PG and Unified CCE System PG Instance Numbers are different, then you can use the same PG number for those PGs.
- Do not add other PGs (such as VRU PG or MR PG) to this VM.
- Follow the scalability limits of the coresident Unified CCE Gateway PG and Unified CCE System PG.

The Unified CCE Gateway PGs provide real-time event data and agent states to the parent from the Unified CCE child. The Unified CCE Gateway PGs also capture some, but not all, configuration data and send it to the parent Unified ICM configuration database.

You can replace the Unified IP IVR at the child site with a local Unified CVP instance. Unified CVP is not integrated as part of the System PG for the Agent Controller. The installation for Unified CCE with Unified CVP defines a separate VRU PG specifically for Unified CVP. Because Unified CVP is not part of the System PG, the Unified CCE Gateway PG does not report calls in queue or treatment to the parent Unified ICM. If your parent routing requires queuing or treatment information, deploying with Unified CVP might not meet your needs.

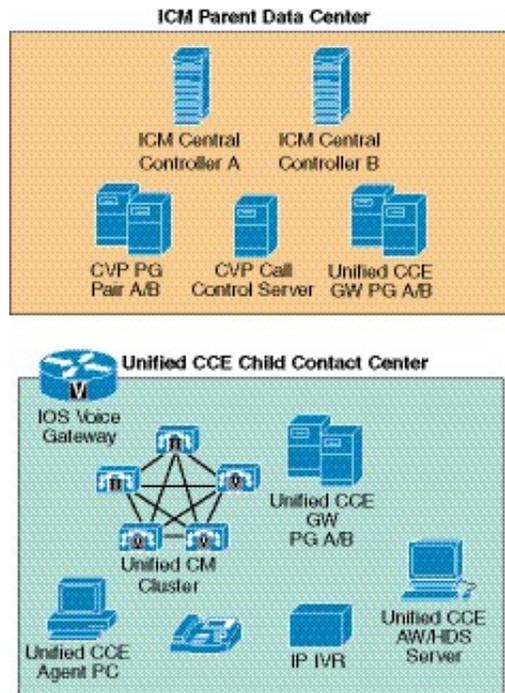
A local Unified CCE child system provides ACD functionality. You can size the Unified CCE child system as a Rogger with a separate Unified CCE Agent PG server. The Rogger contains a Call Router and Logger. The set of redundant Agent PG servers contain the System PG for Unified Communications Manager and Unified IP IVR, CTI Server and CTI OS Server, and the optional VRU PG for Unified CVP.

In either configuration, you need a separate Administration & Data Server to host the configuration and scripting tools for the system, as well as the optional Historical Database Server, and the web-based Unified Intelligence Center reporting tool.

Unified CCE Gateway PGs at Main Site

You can deploy the Unified CCE Gateway PG at the Unified ICM main site as shown in the following figure. This deployment model allows you to manage and control the Unified CCE Gateway PGs centrally. When different companies own and manage the child and parent sites, that condition can force you to deploy the Unified CCE Gateway PG at the Unified ICM main site. For example, an Outsourcer/Service Bureau manages child site and connects to the Unified CCE Gateway PGs at the parent site.

Figure 1: Parent/Child Deployment with Unified CCE Gateway PGs at Main Site



There are several drawbacks with moving the Unified CCE Gateway PGs to the main site. One drawback is recovering reporting data after a network failure. If the network connection between the parent site and the Unified CCE System PGs at the child drops, all reporting at the parent site is lost for that period.



Note You can deploy the Unified CCE Gateway PG locally to the Unified CCE System PG. Then, if the connection between the parent and child sites drops, the historical data in the parent site updates when the network connection comes back.

Another drawback with centralizing the Unified CCE Gateway PGs is higher network bandwidth requirements for the connections between the PGs.

Bandwidth for Unified CCE Gateway PG to Central Controller

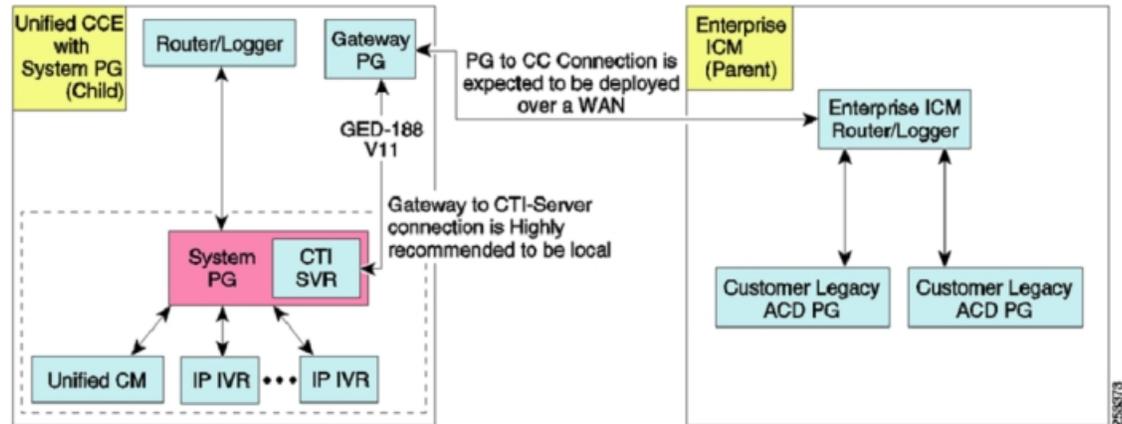
No special considerations are necessary for the PG-to-CC connection over other TDM PGs.

If you do not use agent reporting, then leave that setting disabled to avoid sending unnecessary data over the link.

Bandwidth for Unified CCE Gateway PG to System PG

The following figure shows the connection between the parent PG/PIM and the child system PG.

Figure 2: Connection Between Gateway PG and System PG



Note In general, you deploy the Gateway PG on the same VM with the System PG that it is monitoring. For an outsource model, you can deploy the Gateway PG remote from the System PG.

The following factors affect the amount of data coming over the link once it is initialized:

- Message sizes can vary depending on their content (such as the size of extensions, agent IDs, and call data). For example, a Route Request with no data is a small message. If all call variables and ECC variables are populated with large values, the data drastically affects the size of the message.
- Call scenarios can cause great variation in the number of messages per call that are sent over the line. A simple call scenario can send 21 messages over the line. More complex call scenarios involving queuing, hold retrieves, conferences, or transfers send even more messages over the line for each call.
- The more skill groups to which an agent belongs, the more messages are sent over the line. In a simple call scenario, each additional skill group adds two messages per call. These messages are approximately 110 bytes each, depending on field sizes.

Bandwidth Calculation for Basic Call Flow

A basic call flow (simple ACD call without other steps) with a single skill group typically generates 21 messages. Plan for a minimum required bandwidth of 2700 bytes/second.

In a basic call flow, there are four places where call variables and ECC data can be sent. If you use call data and ECC variables, they are sent four times during the call flow. Using many variables could easily cause the 2700 bytes/second of estimated bandwidth per call to more than double.



Note Call variables used on the child PG are sent to the parent PG regardless of their use or the setting of the MAPVAR parameter. For example, assume that the child PG uses call variables 1 through 8 but the parent PG never uses those variables. If MAPVAR = EEEEEEEEE, meaning Export all but Import nothing, the variables are sent to the PG where the filtering takes place. The bandwidth is still required. But, if the map setting is MAPVAR = IIIIIIIII, Import all but Export nothing, then bandwidth is spared. Call variable data is not sent to the child PG on a ROUTE_SELECT response.

Basic Call Flow Example

Assume a call rate of 300 simple calls per minute (five calls per second). The agents are all in a single skill group with no passing of call variables or ECC data. The required bandwidth in this case is:

$$5 * 2700 \text{ bps} = 13,500 \text{ bps} = 108 \text{ kbps of required bandwidth}$$



Note A more complex call flow or a call flow involving call data could easily increase this bandwidth requirement.

Unified CCE System Peripheral

The Unified CCE System Peripheral acts as a single logical Unified ICM peripheral that combines the functionality the VRU peripherals and a Unified Communications Manager peripheral. Unified CCE treats the Unified IP IVR and Unified Communications Manager peripherals as a single peripheral eliminating the need to translation-route calls to the Unified IP IVR for treatment and queuing. If multiple Unified IP IVRs are configured, the Unified CCE System peripheral automatically load-balances calls between the Unified IP IVRs that have available capacity.

As a single peripheral, Termination Call Detail (TCD) records and other reporting data include the information for the call during the entire time the call is on the peripheral. Instead of getting up to three TCDs for each call (one for the original route, one for the VRU, and one for the agent handle time), the Unified CCE System PG generates only a single record.

The Unified CCE System PG does not support Unified CVP. All queuing and treatment in the Unified CCE System PG uses Unified IP IVR. You can use a separate Unified CVP on its own PG with the Unified CCE System Peripheral.

Parent/Child Limitations

Precision Routing

Precision Routing is not supported in a parent/child deployment. Precision Routing does not support the Unified CCE System Peripheral.

Multichannel Routing

In parent/child configurations, there is no multichannel routing and integration through the parent Unified ICM. Media Routing PGs must connect to the child Unified CCE. A separate Cisco Interaction Manager or partition is required for each child.

Enterprise Unified CCE Peripheral deployments

You cannot use Enterprise Unified CCE (CallManager and VRU deployed as separate peripherals) deployments where Unified CCE is a child to a Unified ICM. Use a Unified CCE System Peripheral deployment for that solution.

Whisper Announcement

Whisper Announcement only supports a specific Parent/Child configuration. That configuration queues calls and sources whisper announcements with an Unified IP IVR on the child system PG. If you also use agent greetings, the configuration also requires a dedicated CVP at the child on a dedicated VRU PG to provide agent greetings. Cisco must approve any use of Whisper Announcement in Parent/Child configurations. Cisco must analyze and approve any such designs.

Whisper Announcement with Unified IP IVR in Parent/Child has no impact on agent sizing on the Child System PG, but incurs great impact on the Unified IP IVR.

Agent Greeting

Agent Greeting only supports a specific Parent/Child configuration. That configuration queues calls at an Unified IP IVR on the child system PG. The configuration requires a dedicated CVP at the child on a dedicated VRU PG to provide the agent greetings. Cisco must approve any use of Agent Greeting in Parent/Child configurations. Cisco must analyze and approve any such designs.

Your configuration must meet the following conditions:

- Use Unified IP IVR on the child for call treatment, queuing, and Whisper Announcement.
- Use CVP on the child only for Agent Greeting. Do not use CVP for call queuing. Use a separate VRU PG for the CVP.

Network Consultative Transfer

In a parent/child deployment, you cannot transfer calls terminating on child systems by network consultative transfer (NCT) through the routing clients on the parent. Although NCT works for TDM ACDs, and the parent/child deployment architecture is similar, the parent/child deployment does not work the same.

For a TDM PG, the CTI Server is connected to the ACD PG, which is part of the parent system. This arrangement is the same as a CTI-Server connected to the Gateway PG. In parent/child deployments, CTI connects to the child PG. Having CTI connected to the child PG does not provide the network call ID and other information that are needed for allow network consultative transfer.



Note When a post route is initiated to the parent system from the child, network blind transfer is possible using any client (for example, Unified CVP) on the parent system.

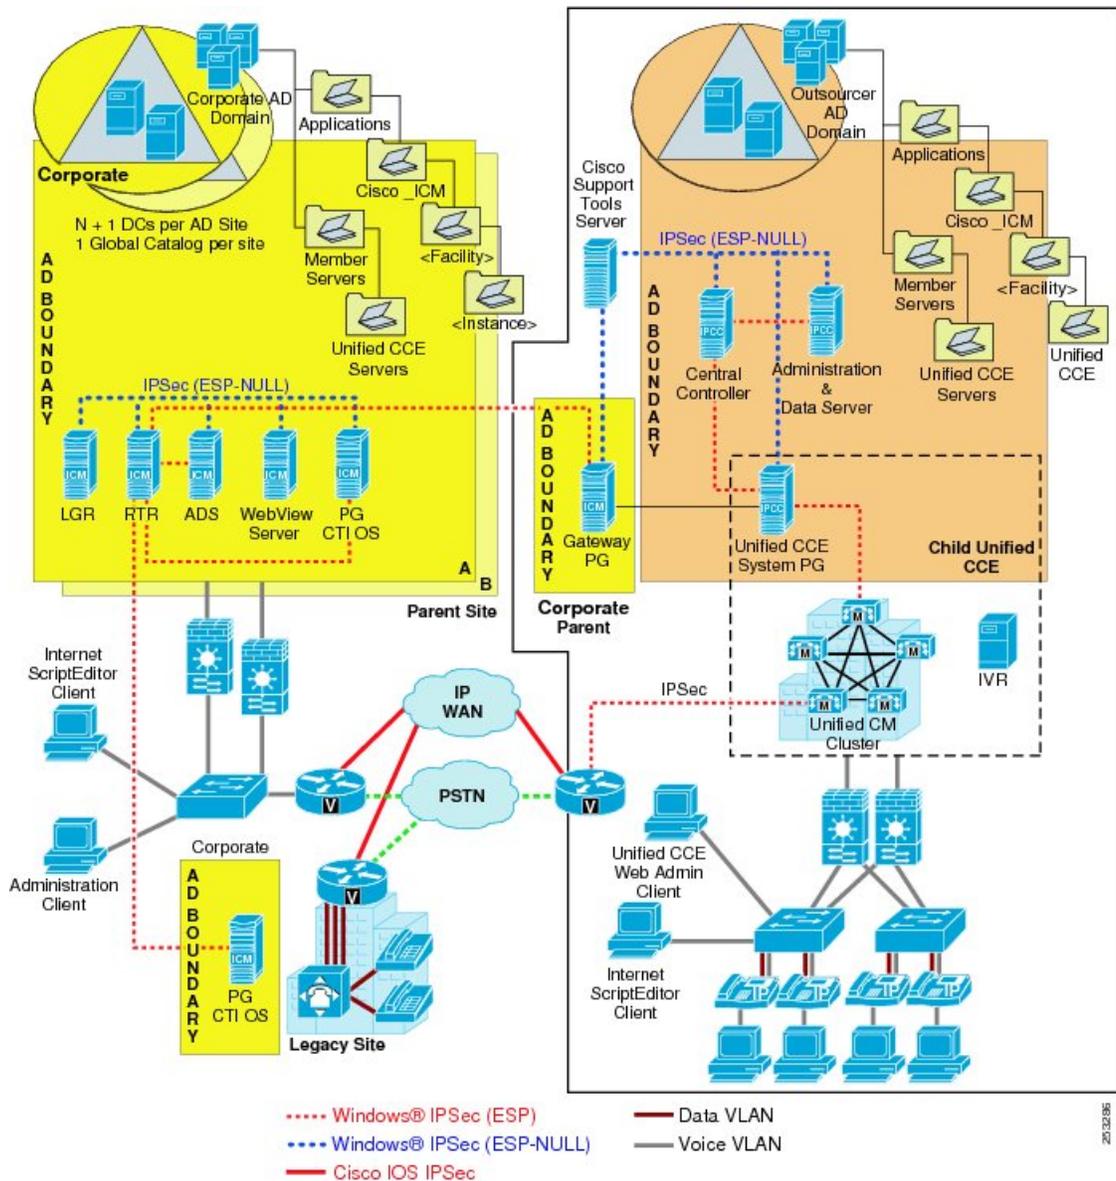
Active Directory Deployments for Parent/Child

You can deploy parent/child systems on the same AD Domain or Forest, and in disparate AD environments. The common scenario for this deployment occurs when an outsourced contact center site houses the child Unified CCE system. In this case, the parent AD domain contains the Gateway PG that is a parent node. (Do *not* use Workgroup membership because of the administration limitations.) This deployment supports remote branch offices with PGs that are members of the central site domain which contains the Routers, Loggers, and Distributors.

The topology shown in the following figure represents the AD Boundaries for the two AD domains in this deployment. The figure also shows to which domain the application servers are joined. The parent AD Domain Boundary extends beyond the main site. The parent AD Boundary includes the Unified ICM Central Controllers and accompanying servers with the ACD PG (at the legacy site) and Gateway PG at the child Unified CCE site. The child Unified CCE site and its AD Boundary have the Unified CCE servers as members. The AD

Boundary can be part of an outsourcer corporate AD environment, or the AD Boundary can be a dedicated AD domain for Unified CCE.

Figure 3: Active Directory and Firewall Deployment



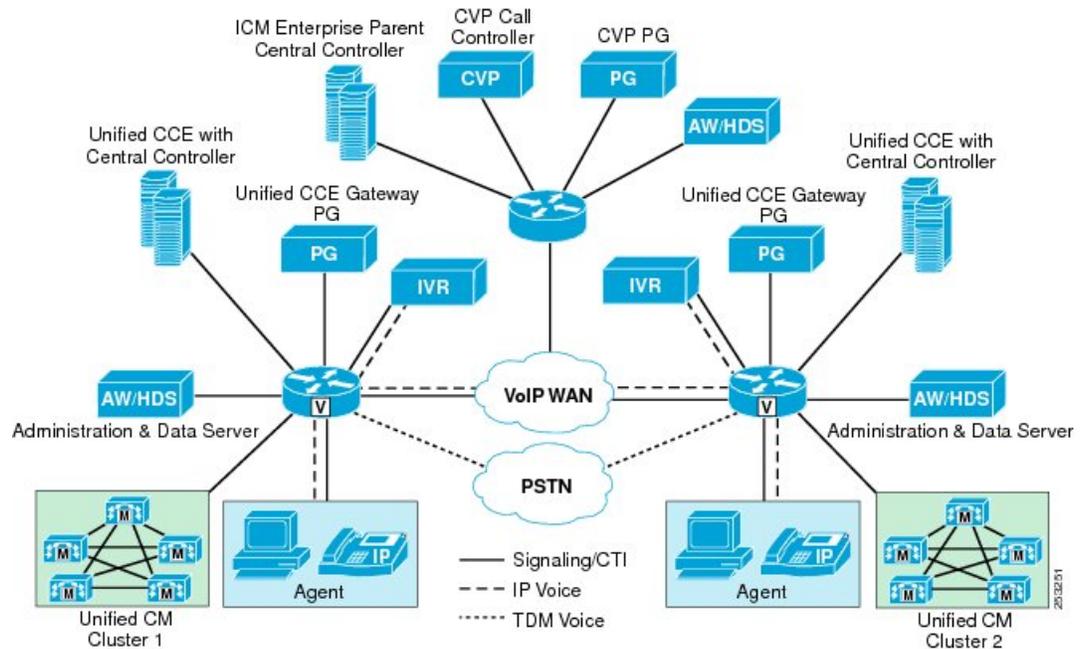
Unified CCMP for Parent/Child Deployments

In parent/child deployments, a single Unified CCMP instance connects to each of the child Unified CCE Administration & Data servers. Configure these servers as physically separate Primary Administration & Data Servers. Each child instance appears as a *tenant* within Unified CCMP. Resources added through Unified CCMP are linked to a tenant. The standard process replicates the added resource from the Unified CCE child to its parent.

Parent/Child Deployments Across Sites

The Parent/Child deployment provides local, distributed call processing with a local Unified Communications Manager and Unified CCE at each site (child). A centralized Unified ICM Enterprise parent controls the child site for enterprise-wide routing, reporting, and call control. This deployment is more tolerant of WAN outages, with each site continuing to operate during an outage. The following figure shows this deployment.

Figure 4: Multisite Deployment with Distributed Call Processing and Parent/Child



In this design, there is a parent Unified ICM Enterprise system deployed with Unified CVP and its own Administration & Data server. Each distributed child site is a complete Unified CCE deployment consisting of Central Controller on one or more VMs. A local Administration & Data Server for Unified CCE performs configuration, scripting, and reporting tasks for that specific site. A Unified CCE Gateway PG connects Unified CCE to the Unified ICM parent and it is part of the Peripheral Gateways deployed on the parent Unified ICM.

An optional deployment for the Unified CCE Gateway PG is to colocate it with the Unified CCE System PG, under the following guidelines:

- If the Unified CCE Gateway PG and Unified CCE System PG Instance Numbers are the same, then use different PG numbers for the PGs.
- If the Unified CCE Gateway PG and Unified CCE System PG Instance Numbers are different, then you can use the same PG number for the PGs.
- Do not add any other PGs (such as a VRU PG or MR PG) to this VM.

In this design, the local Unified CCE deployments act as their own local IP ACDs with no visibility to any of the other sites in the system. Agents at Site 1 cannot see any of the calls or reports from Site 2 in this deployment. Only the Unified ICM Enterprise parent system has visibility to all activity at all sites connected to the Unified ICM Enterprise system.

The Unified CVP at the Unified ICM parent site controls the calls coming into the distributed sites. Unified CVP provides call queuing and treatment in the VXML Browser in the Voice Gateway (VG). Configure the Unified CVP on the parent to use Unified CVP Router Requery to take control of the call during a failure or answer timeout. The child Unified CCE cannot terminate the ingress call to a child Unified CVP or a child Unified IP IVR. The local Unified IP IVR servers only provide a local backup for the connection from these VGs to the parent Unified CVP Call Control server. The local Unified IP IVR also provides local queue treatment for calls that the local agents do not answer (RONA) rather than sending the call to the Unified CVP to be queued.

The child Unified CCE deployments can also transfer calls across the system between the sites using Unified ICM post-routing by the Unified CCE Gateway PG. The Unified CCE Gateway PG allows the child Unified CCE to ask the Unified ICM to transfer a call to the best agent at another site or to queue it centrally for the next available agent.

Unlike traditional Unified CCE deployments with distributed Unified Communications Manager Peripheral Gateways, the parent/child deployment provides for complete local redundancy at the contact center site. The local Unified CCE takes over call processing for inbound calls from the Unified CVP gateways and provides local call queuing and treatment in the local Unified IP IVR. This configuration provides complete redundancy and 100% up-time for contact centers that cannot be down because of a WAN failure.

For customers who have Unified ICM already installed with their TDM ACD platforms, this approach can be useful when they want to do the following:

- Add new sites with Unified CCE
- Convert an existing site to Unified CCE

The approach allows the Unified ICM to continue performing enterprise-wide routing and reporting across all the sites while inserting new Unified CCE technology on a site-by-site basis.



Note Unified CVP can be at both the parent and child. The call flows are similar for Unified CVP at the parent and IP IVR at the child. One key difference is that information about queued calls at the child Unified CVP are not available at the parent (through the Unified CCE Gateway PG). This difference means that you cannot use routing elements like the minimum expected delay (MED) over services or CallsQNow in the parent.

Advantages

- Unified CVP provides a virtual network queue across all the distributed sites controlled by the parent Unified ICM. The parent Unified ICM has visibility into all the distributed sites and sends the call to the next available agent from the virtual queue.
- Each distributed site can scale up to the maximum number of supported agents on a single Unified CCE deployment. Multiple Unified CCE Central Controllers can be connected to a single cluster to scale up to the maximum number of supported agents per cluster. The Unified CCE Gateway PG on the parent connects the Unified CCE systems to the parent Unified ICM. The Unified CCE Gateway PG can scale up to the maximum number of supported agents per parent Unified ICM Enterprise system.
- All or most VoIP traffic can be contained within the LAN of each site, if desired. The QoS WAN is required for voice calls to be transferred across sites. Use of a PSTN transfer service (for example, Take Back and Transfer or Transfer Connect) eliminates that need. If desired, a small portion of calls arriving at a particular site can be queued for agent resources at other sites to improve customer service levels.

- Unified ICM pre-routing can load-balance calls based on agent or Unified CVP session availability. Unified ICM can also route calls to the best site to reduce WAN usage for VoIP traffic.
- Failure at any one site has no impact on operations at another site.
- Each site can be sized according to the requirements for that site.
- The parent Unified ICM Central Controller provides centralized management for configuration of routing for all calls within the enterprise.
- The parent Unified ICM Central Controller provides the capability to create a single enterprise-wide queue.
- The parent Unified ICM Central Controller provides consolidated reporting for all sites.

Disadvantages

The parent/child deployment usually requires a higher number of VMs. The extra VMs are needed for the increased number of software components (additional Unified CCE Gateway PGs required if colocating with Unified CCE System PG is not an option, additional Central Controller for each child, and so forth).

Requirements

- Colocate the Unified CCE Gateway PG, Unified Communications Manager cluster, Unified IP IVR, and Unified CCE (if possible) at the contact center site.
- The communication link from the parent Unified ICM Central Controller to the Unified CCE Gateway PG must be sized properly and provisioned for bandwidth and QoS.
- Gatekeeper-based or RSVP agent-based call admission control is used to reroute calls between sites over the PSTN when WAN bandwidth is not available. It is best to ensure that adequate WAN bandwidth exists between sites for the maximum amount of calling that can occur.
- If the communication link between the Unified CCE Gateway PG and the parent Unified ICM Central Controller is lost, then all contact center routing for calls at that site is put under control of the local Unified CCE. Unified CVP-controlled ingress Voice Gateways have survivability TCL scripts to redirect inbound calls to local Unified Communications Manager CTI route points and the local Unified IP IVR are used to handle local queuing and treatment during the WAN outage. This feature of the parent/child deployment provides complete local survivability for the call center.
- While two intercluster call legs for the same call do not cause unnecessary RTP streams, two separate call signaling control paths remain intact between the two clusters (producing logical hair-pinning and reducing the number of intercluster trunks by two). Consider the percentage of intersite transfers when sizing intercluster trunks capacities.
- Latency between parent Unified ICM Central Controllers and remote Unified CCE Gateway PGs must not exceed 200 ms one way (400-ms round trip).

Geographically Redundant Child Sites (Using Unified IP IVR)

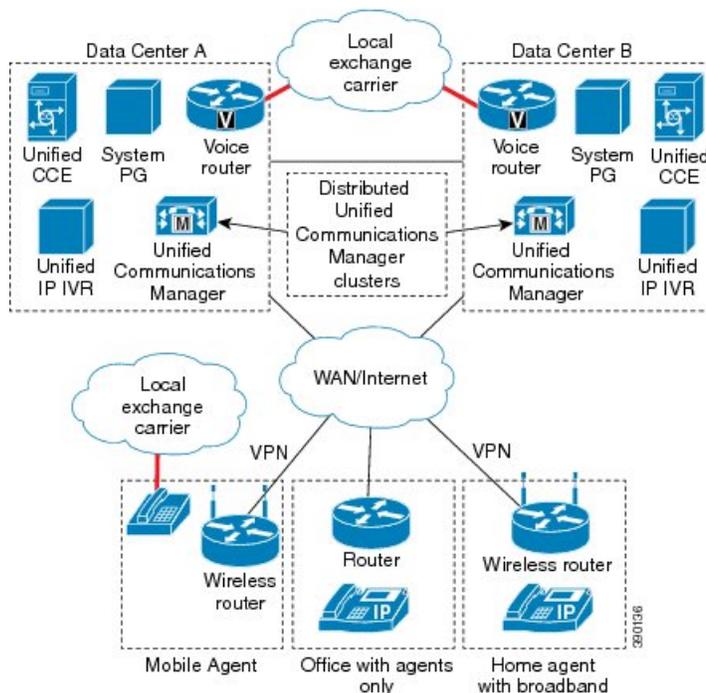
Globalization, security, and disaster recovery considerations are driving business to diversity locations across multiple regions. In addition, organizations want to distribute workloads between computers, share network resources effectively, and increase the availability of critical applications. Geographically redundant sites split critical applications across two sites. Enterprises deploy geographically redundant sites to minimize planned or unplanned downtime and share data across regions.

Geographically redundant sites have a minimum of two load balancers, one in each site. You can use two load balancers for each site for local redundancy.

Geographically Redundant Child Sites with CoW

The following figure shows geographically redundant Child sites with clustering over the WAN (CoW) using Unified IP IVR.

Figure 5: Geographically Redundant Child Sites with Clustering over WAN (Using Unified IP IVR)



Geographically redundant sites use clustering over the WAN, Unified Communications Manager clusters, and 1:1 redundancy for IP IVR, SIP proxy, voice gateways, and Cisco Unified Intelligence Center for example.

Latency requirements across the high-availability (HA) WAN must meet the current Cisco Unified Communications requirements for clustering over the WAN. Unified Communications Manager allows a maximum latency of 40 ms one way (80 ms for a round trip).

Certain fault tolerant networks can carry all your traffic on a single network, for example, Multiprotocol Label Switching (MPLS) or SONET. For such networks, keep the public and private traffic on separate routes within the network and respect standard latency and bandwidth.

Provision WAN connections to the agent sites with bandwidth for voice, control, and CTI. You can have a local voice gateway at remote sites for local and 911 calls. For more information, see Cisco Unified Communications and Collaboration Solutions Design Guidance at http://www.cisco.com/en/US/docs/voice_ip_comm/uc_system/design/guides/UCgoList.html

In a balanced deployment, central site outages include loss of half of the ingress gateways. Scale gateways to handle the full load in both sites if one site fails.

Carrier call routing must be able to route calls to the alternate site during failures.

A single Unified CCE system peripheral controls all the Unified IP IVRs and the Unified Communications Manager. Unified IP IVR distributes the calls to the least loaded Unified IP IVR. A Unified IP IVR in Site B

can treat calls coming into Site A. Both A- and B-sides of Unified CCE can identify all the Unified IP IVRs. PIM activation logic determines if the A- or B- side PIM connects to each of the Unified IP IVRs. This process means that the PG at Site A can connect to the Unified IP IVR at Site B. Make sure that you size the WAN appropriately.

To avoid the bandwidth overhead, consider using Unified CVP for clustering over the WAN deployments. Unified CVP allows higher scalability per cluster than Unified IP IVR.

Unified IP IVR-Based Child Sites with Distributed Unified Communications Manager

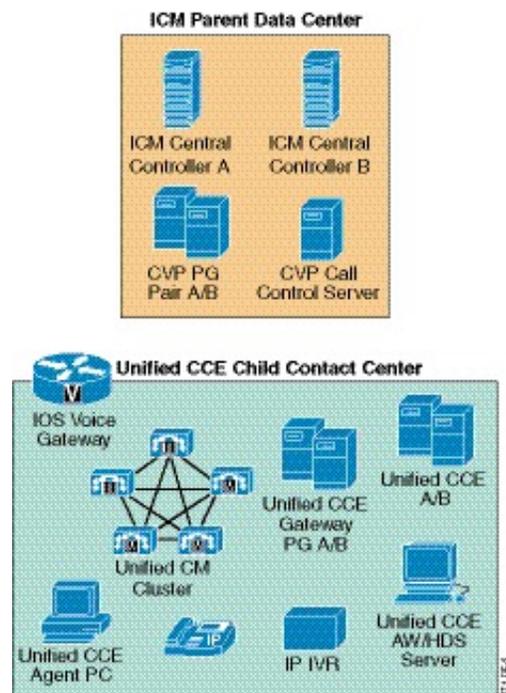
If you have remote offices with agents, gateways, and Unified Communications Manager clusters, the clusters at the sites are typically independent.

The remote office has a WAN connection back to the main site. Each cluster is independent, with its own agents and PG pairs. Each site uses subscribers that are local to the site because JTAPI is not supported over the WAN. For example, **Site A** cannot use the subscribers in **Site B**. The Unified CCE central controller, Unified Intelligence Center, load balancer, SIP proxy server, and Unified IP IVR are located in the data centers. TDM and VXML voice gateways are located at the remote office with local PSTN trunks.

Unified CCE High Availability with Unified ICM

In a parent/child deployment, the Unified ICM acts as the parent controlling one or more Unified CCE child ACDs. The Unified ICM system is the network call routing engine for the contact centers. The network queuing uses the Unified CVP and Unified CCE Gateway PGs to connect child Unified CCE systems. The child Unified CCE systems are individual ACD systems fully functional with local call processing in case they lose their WAN connection to the parent system. This configuration provides a high level of redundancy and availability to the Unified CCE solution. These qualities allow sites to remain functional as Unified CCE sites even if they are cut off from centralized call processing resources.

Figure 6: Parent/Child Deployment



Parent/Child Call Flows

The following sections describe the call flows between the parent and child.

Typical Inbound PSTN Call Flow

In a typical inbound call flow from the PSTN, the carrier network directs calls to the contact center sites using a predefined percent allocation or automatic routing method. These calls terminate in the Unified CVP VGs (Voice Gateways) at the contact center locations under control of Unified CVP on the Unified ICM parent .

The inbound call flow is as follows:

1. The call arrives on the Unified CVP VG at the Parent ICM Site.
2. The Unified CVP VG maps the call by dialed number to a particular Unified CVP Server at the parent site. The VG then sends a new call event to the Unified CVP Server.
3. The Unified CVP Server sends the new call event message to the Unified CVP VRU PG at the Unified ICM parent site.
4. The Unified CVP PG sends the new call message to the Unified ICM parent. The Unified ICM uses the inbound dialed number to qualify a routing script to determine the proper call treatment (messaging) or agent groups to consider for the call.
5. Unified ICM instructs Unified CVP to hold the call in the VG at the site. Unified CVP waits for an available agent while directing specific instructions to play .wav files for hold music to the caller.
6. When an agent becomes available, the Unified ICM instructs Unified CVP to transfer the call to the available agent by using a translation route. (The agent is not at the same physical site, but located across the WAN.) Any data collected about the call in the Unified ICM parent Unified CVP is transferred to the remote PG (either a TDM, legacy PG, or one of the Unified CCE Gateway PGs for Unified CCE).
7. The call arrives at the targeted site on a specific translation route DNIS that the Unified ICM parent selected. The PG at the child site is expecting a call to arrive on this DNIS to match up with any pre-call CTI data associated with the call. The local ACD or Unified CCE performs a post-route request to the PG (either TDM PG or Gateway PG depending on the target ACD) to request the CTI data as well as the final destination for the call (typically the lead number for the skill group of the available agent).
8. If the agent is no longer available, Unified CVP at the parent site uses the Router Requery function in the ICM Call Routing Script to select another target automatically.

Post-route Call Flow

You use post-routing for calls already at a peripheral ACD or VRU that you want to route to another agent or location intelligently. If an agent gets a call in the ACD or Unified CCE that must go to a different skill group or location, the agent can use the post-route functionality to reroute the call.

The post-route call flow is as follows:

1. The agent transfers the call to the local CTI route point for reroute treatment using the CTI agent desktop.
2. The reroute application or script makes a post-route request to the Unified ICM parent by using the local Unified CCE Gateway PG connection.
3. The Unified ICM parent maps the CTI route point from Unified CCE as the dialed number and uses that number to select a routing script. This script returns a label or routing instruction that can move the call to another site, into a different skill group on the same site, or to a Unified CVP node for queuing.
4. Unified CCE receives the post-route response from the Unified ICM parent system. Unified CCE then uses the returned routing label as a transfer number to send the call to the next destination.

Parent/Child Fault Tolerance

The parent/child model provides for fault tolerance to maintain a complete ACD with Unified CCE deployed at the site, with local IP-PBX and call treatment and queuing functionality.

Unified CCE Child Loses Connection to Unified ICM

A WAN failure between the child and the parent isolates the local Unified CCE system from the parent and the Unified CVP VG. Calls coming into the site no longer get treatment from the Unified CVP under control of the Unified ICM parent. Replicate the following functionality locally, depending on the configuration at the child site:

- For Unified CCE child configurations using local IP IVR resources for queue and treatment:
 - The local VG must have dial peer statements to pass control of the calls to the local Unified CM cluster if the parent Unified CVP Server cannot be reached. The local Unified CM cluster must have CTI route points mapped to the inbound DNIS or dialed numbers that the local VG presents if the parent Unified CVP Server is not reached.
 - Configure the local IP IVR with appropriate audio files and applications that the child system can call locally to provide basic call treatment.
 - The child CCE Routing Script must handle queuing of calls for agents in local skill groups. The script must instruct the IP IVR to play the treatment in-queue while waiting for an agent.
 - To allow the agents full access to customer data for routing and screen pops, provision locally any data lookup or external CTI access that parent system normally provides.
 - Any post-routing transfer scripts fail during this outage, so configure Unified CCE to handle this outage or prevent the post-route scripts from being accessed.
- For Unified CCE child configurations using local Unified CVP resources for queue and treatment:
 - The local VG must have dial peer statements to pass control of the calls to the local Unified CVP Server at the child site. To process these calls locally at the child site, configure in the child Unified CCE the inbound DNIS or dialed numbers that the local VG presents to the child Unified CVP.
 - Configure the local VXML Gateways and Unified CVP Servers with appropriate .wav files and applications that the child system can call locally to provide basic call treatment.
 - Replicate self-service or Unified CVP Studio VXML applications that the parent Unified ICM normally provides. Use the Unified CVP Server (web application server) at the child site to generate the dynamic VXML for these applications.
 - The child Unified CCE Routing Script must handle queuing of calls for agents in local skill groups. The script must instruct the local Unified CVP at the child site to play the treatment in-queue while waiting for an agent.
 - To allow the agents full access to customer data for routing and screen pops, provision locally any data lookup or external CTI access that parent system normally provides.
 - Any post-routing transfer scripts fail during this outage, so configure Unified CCE to handle this outage or prevent the post-route scripts from being accessed.

Unified CCE Gateway PG Cannot Connect to Unified ICM

If the Unified CCE gateway PG fails or cannot communicate with the Unified ICM parent, the Unified ICM parent cannot detect the state of the agents at the child. But, in some cases, the Unified ICM parent Unified CVP can still control the inbound calls. In this case, the Unified ICM parent does not know if the remote Unified CCE gateway PG failed or if the actual Unified CCE ACD failed locally.

The Unified ICM at the parent location can automatically route around this site, considering it down until the PG comes back online and reports agent states again. Alternatively, the Unified ICM can also direct a percentage

of calls as blind transfers to the site Unified CCE using the local inbound CTI route points on Unified CM. This method would present calls with no CTI data from Unified CVP, but it would allow the agents at the site to continue to get calls locally with their Unified CCE system.

If the local Unified CCE child system fails, the Unified CCE gateway PG cannot connect to it. The Unified ICM parent then considers all agents to be off-line and not available. If local cluster receives calls while the child Unified CCE system is down, the call-forward-on-failure processing takes over the call for the CTI route point. This method redirects the call to another site or an answering resource to play a message telling the caller there is a problem and to call again later.

Reporting and Configuration Impacts

If the Unified CCE child disconnects from the Unified ICM parent, the local ACD still collects reporting data and allows local administrators to change the child routing scripts and configuration. The Unified CCE gateway PG at the child site caches these objects and stores them to be sent when the Unified ICM parent is available. This functionality is available only if the Unified CCE gateway PG is colocated at the child Unified CCE site.

Other High Availability Considerations

You can install multichannel components, such as Cisco Unified Web and E-mail Interaction Manager and Cisco Outbound Option, only at the child Unified CCE level, not at the parent. They are treated as nodal implementations on a site-by-site basis.

Traditional ACD Integration

Enterprises that want to integrate traditional ACDs with their Unified CCE use a parent/child deployment. In these deployments, the Unified ICM and Unified CCE each have a Central Controller or a hybrid deployment where Unified Communications Manager PGs with TDM ACD PGs, Gateway PGs, or both use the same Central Controller. Several options exist within those categories depending on how the calls are routed within the deployment.

Hybrid Deployment with Fixed PSTN Delivery

As an alternative to pre-routing calls from the PSTN, the PSTN can deliver calls to just one site or to split the calls across the two sites according to a set of static rules provisioned in the PSTN. When the call arrives at either site, either the traditional ACD or the Unified Communications Manager generates a route request to the hybrid Unified ICM/CCE to determine which site is best for this call. If the call is for an agent at the opposite site from where the call was originally routed, you require TDM circuits between sites. The determination of where calls are routed (and if and when they are transferred between sites) depends on the enterprise business environment, objectives, and cost components.

Hybrid Deployment with Unified CVP

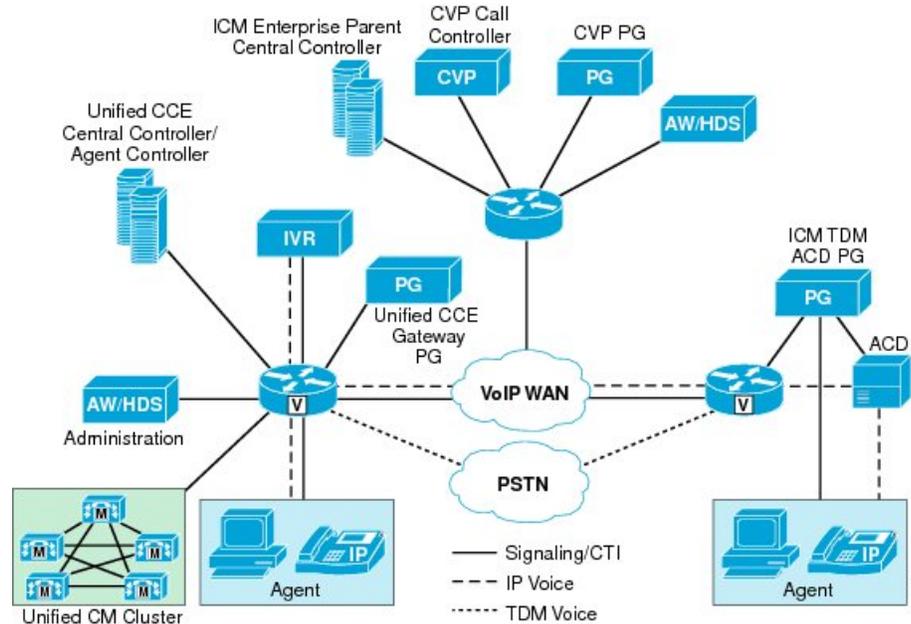
Customers can choose to front end all calls with Unified CVP to provide initial call treatment and queuing across both the TDM ACD and Unified CCE agents.

In this design, all calls first come to the Voice Gateway (VG) controlled by Unified CVP. The Unified ICM/CCE Call Router then directs the calls. Unified ICM/CCE uses the PG connections to the TDM ACD and Unified CCE PG to monitor for available agents. Calls are queued in Unified CVP until an agent becomes available in either environment. When a call transfers to the TDM ACD, the call comes into the VG on a T1 interface from the PSTN carrier network and goes out on a second physical T1 interface to appear as a trunk on the TDM ACD (known as "hairpinning the call"). Most TDM ACDs cannot accept inbound calls in IP from the VG and require this physical T1 interface connection. Unified CCE agents receive their calls directly over the IP voice network.

ACD Integration and Parent/Child Deployments

The following figure illustrates the parent/child.

Figure 7: Parent/Child Integration of Traditional ACD with Unified CCE



In this model, the Unified ICM Enterprise parent has PGs connected to a Unified CCE System PG at one site and a Unified ICM TDM ACD PG at a second site. In this model, Unified ICM still provides virtual enterprise-wide routing, call treatment, and queuing with the distributed Unified CVP Voice Gateways at the sites. Unified ICM also has full visibility to all the sites for agents and calls in progress. The difference in this model is that Unified CCE provides local survivability. If it loses connection to the Unified ICM parent, the calls are still treated locally just as they are at the TDM ACD site.

Traditional VRU Integration

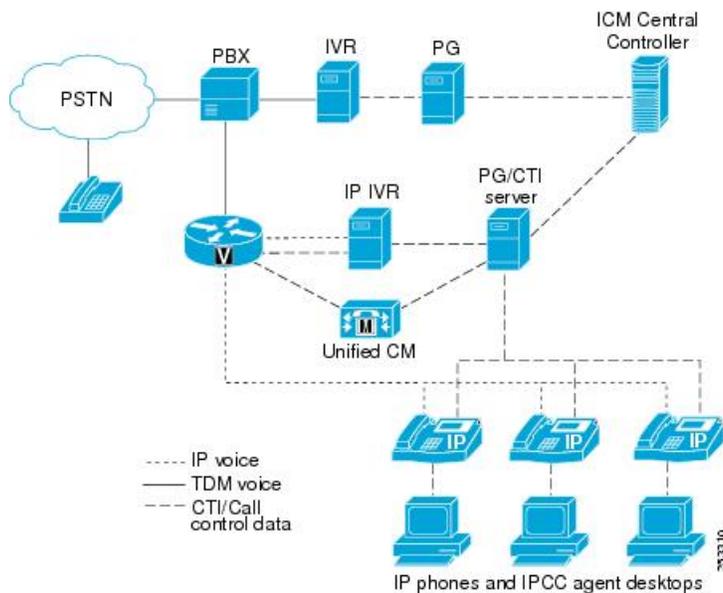
There are numerous ways to integrate traditional VRUs into a Unified CCE deployment. The following sections discuss the factors that can determine the best way. The primary consideration is determining how to eliminate or reduce VRU double trunking when transferring the call from the VRU.

Integration Through PBX Transfer

Many call centers have existing traditional VRU applications that they are not prepared to rewrite. To preserve these VRU applications and integrate them into a Unified CCE environment, the VRU must have an interface to Unified CCE. That VRU interface to Unified CCE is the Service Control Interface (SCI). The SCI enables the VRU to receive queuing instructions from Unified CCE. In the PBX model, the SCI is not required.

Even if the VRU has the SCI interface, deploy Unified CVP or Unified IP IVR for all call queuing. This method prevents any extra use of the traditional VRU ports. In addition, use of the Unified IP IVR for queuing provides a way to re-queue calls on subsequent transfers or RONA treatment.

Figure 8: Traditional VRU Integration Using PBX Transfer



In this design, calls come first to the PBX from the PSTN carrier network on a standard T1 trunk interface. The PBX typically uses a hunt group to transfer the call to the VRU, putting all VRU ports into the hunt group as agents in auto available mode. The PBX looks like the PSTN to Unified CCE because it does not have a PG connected to the PBX. Unified CCE cannot track the call from the original delivery to the VRU. Unified CCE only has call data from the time the call arrived at the VRU and the VRU informed Unified CCE of the call.

When the caller opts out of the VRU application, the VRU sends a Post-Route to Unified CCE. Unified CCE looks at the agent states across the system. Unified CCE then selects the agent to send the call to or translation-routes the call to the Unified IP IVR for queuing.

When the call is sent to an agent or into the queue, the call comes into the PBX from the PSTN on a T1 trunk port and then goes out to a VG on a second T1 trunk port in the PBX. This connection is used for the life of the call.

If you want to track the call from its entry at the PBX or if you want to capture the caller ANI or original dialed number, you can install a PG on the PBX. The PBX can request (through a Post-Route to Unified CCE) which VRU port to send the call to behind the PBX. The PBX cannot use a hunt group to deliver the call from the PBX to the VRU. Unified CCE requires direct DNIS termination to ensure that the translation route maintains the call data collected in the PBX and makes it available to the VRU.

Integration Through PSTN Transfer

This model is similar to the PBX Transfer model. In this model, the VRU invokes a PSTN transfer (instead of a PBX transfer) to release the traditional VRU port. Again, the Unified IP IVR does all queuing to avoid any additional occupancy of the traditional VRU ports and any double trunking in the VRU. Unified CCE passes any call data collected by the traditional VRU application to the agent desktop or Unified IP IVR.

In this model, the TDM VRU is set up as a farm of VRU platforms that have direct PSTN connections for inbound calls. The VRU has a PG connection to Unified CCE which tracks all calls in the system. When a caller opts out of the VRU treatment, the VRU sends a post-route request to Unified CCE. Unified CCE returns a label that directs the call either to an agent or to the Unified IP IVR for queuing.

The label that is returned to the TDM VRU instructs it to send an in-band transfer command using transfer tones (*8 with a destination label in the carrier network). The VRU out-pulses these tones to the service provider with tone generation or plays the tones by using a recorded file.

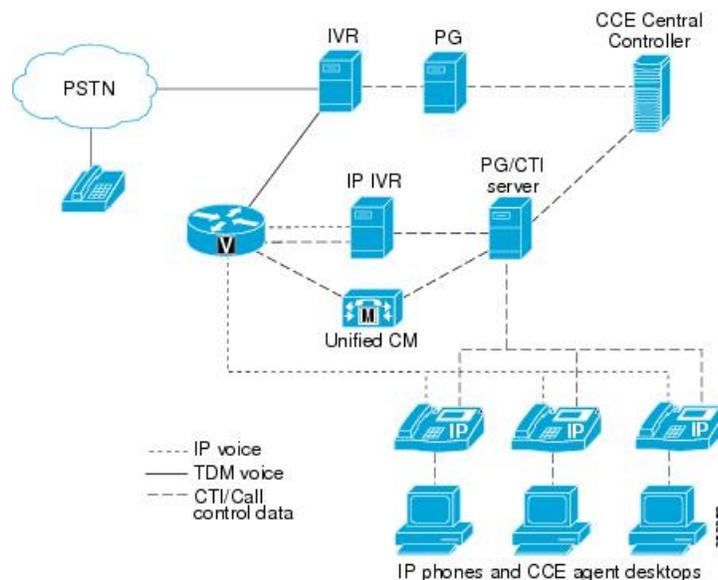
Integration Through VRU Double Trunking

Some traditional VRU applications have a high success rate where most callers are completely self-served in the traditional VRU and only a small percentage of callers are transferred to an agent. For those cases, you can double-trunk the calls in the traditional VRU for that small percentage of calls.

Unlike the previous model, if the traditional VRU has a Service Control Interface (SCI), then the initial call queuing is done on the traditional VRU. In this case, VRU double trunking does not use a second traditional VRU port to transfer the call to the Unified IP IVR. When the traditional VRU does the initial queuing, only one traditional VRU port is used for the call. However, any subsequent queuing because of transfers or RONA treatment must be done on the Unified IP IVR to avoid any double trunking.

If the traditional VRU does not have an SCI interface, then the VRU generates a post-route request to Unified CCE to determine where the call is transferred. All queuing in that scenario happens on the Unified IP IVR.

Figure 9: Traditional VRU Integration Using VRU Double Trunking



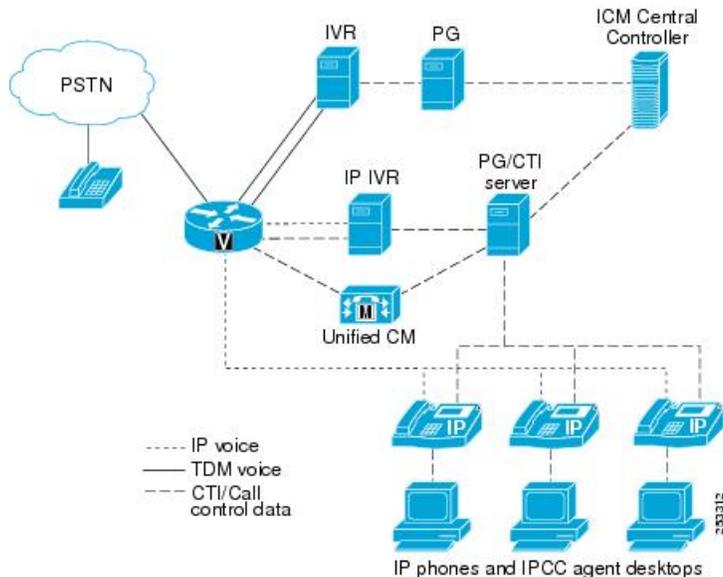
In this model, the TDM VRU is set up as a farm of VRU platforms that have direct PSTN connections for inbound calls. The VRU has a PG connection to Unified CCE which tracks all calls in the system. When a caller opts out of the VRU treatment, the VRU sends a post-route request to Unified CCE. Unified CCE returns a label that either directs the call to an agent or queues the call locally on the TDM VRU using the Service Control Interface (SCI). The TDM VRU does the transfer to the agent by selecting a second port to hairpin the call to the Voice Gateway and to the Unified CCE agent. The hairpin connection takes up two ports for the time the call is at the agent.

Unified Communications Manager Transfer and VRU Double Trunking

Over time, you can migrate the traditional VRU applications to Unified CVP or Unified IP IVR. If a few traditional VRU applications still exist for specific scenarios, then the VRU can connect to a second Voice Gateway (VG). The Unified Communications Manager routes calls arriving at the VG from the PSTN. Unified Communications Manager can route specific DNs to the traditional VRU or let Unified CCE, Unified CVP,

or Unified IP IVR determine when to transfer calls to the traditional VRU. To transfer calls in the traditional VRU to a Unified CCE agent, use a second VRU port, trunk, and VG port during the call. Ensure that transfer scenarios cannot create multiple loops or voice quality suffers.

Figure 10: Unified Communications Manager Transfer and VRU Double Trunking



In this model, Unified CVP can front end the TDM VRU using the VG to determine where to provide call treatment. Alternately, you can use Unified IP IVR and Unified Communications Manager with Unified CCE for this purpose.

With Unified CVP, calls coming into the VG immediately start a routing dialog with Unified CCE using the Service Control Interface (SCI). Based on the initial dialed number or prompting in Unified CVP, Unified CCE decides to send the call to the TDM VRU or to Unified CVP for a specific self-service application. If the call was sent to the TDM VRU, the TDM VRU sends a route request to Unified CCE when the caller opts out. The reply is not sent back to the TDM VRU but back to Unified CVP as the original routing client. Unified CVP then takes the call leg away from the TDM VRU. Unified CVP then transfers the call to the Unified CCE agent over the VoIP network or holds the call in a queue locally in the VG.

With Unified Communications Manager, calls in the VG use a subscriber CTI route point to send a route request to Unified CCE for the proper call treatment device. If the CTI route point indicates an application that still is on the TDM VRU, Unified CCE instructs the subscriber to transfer the call to the TDM VRU by hairpinning the call using a second T1 port on the VG. Unified CCE can also instruct the subscriber to translation-route the call to the Unified IP IVR for call processing or prompting. Unified CCE can then make a subsequent transfer to the TDM VRU for further processing. When the caller opts out of the TDM VRU, it sends a post-route request to Unified CCE for a label to the TDM VRU. This label instructs the TDM VRU to transfer the call using a second T1 port on the VRU back to the VG. The VG transfers the call over to the Unified CCE agent under the Unified Communications Manager dial plan.

In the model controlled by Unified Communications Manager, the VG initially receives calls and sends them to the TDM VRU on a second T1 port. When the VRU returns the call to the Unified CCE agent, it uses a second TDM VRU port and a third port on the VG. All three VG ports are in use as long as the agent is talking with the caller. Both of the TDM VRU ports are used for the remainder of the call.

Non-Reference Design Core Components and Implementations

Using MGCP Gateways

Cisco Unified CVP requires the deployment of a SIP Gateway. However, you require the use of MGCP0.1 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 MGCP0.1 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 UnifiedCVP.

When sending calls through Unified CM to UnifiedCVP, 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 UnifiedCVP, 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 UnifiedCVP to ensure that the calls are queued at the appropriate site or Voice Browser.

**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

Network VRU Types

This section discusses the Network VRU types for Unified ICM, and how Unified ICM relates to Unified CVP deployments. It includes the following topics:

- Unified ICM Network VRUs
- Unified CVP Type 10 VRU

- Unified CVP Type 7 VRU (Correlation ID Mechanism)
- Unified CVP Type 8 VRU (Translation Route ID Mechanism)



Note The terms voice response unit (VRU) and interactive voice response (IVR) are used interchangeably throughout this document.

Unified ICM Network VRU

Unified ICM perceives calls that need IVR treatment as having two portions: the Switch leg and the VRU leg. The Switch is the entity that first receives the call from the network or caller. The VRU is the entity that plays audio and preforms prompt-and-collect functions. If you use Unified ICM, Unified CVP can participate in the Switch role or the VRU role, or both. In a network deployment, multiple Unified CVP devices provide the Switch and VRU portions independently.

The call delivery to VRU can be based on either a Correlation ID or a translation route ID function, depending on the network capability to pass the call reference identification to the VRU. Call reference identification is needed because Unified ICM has to correlate the two legs of the same call in order to provide instructions for completing the call. In the Unified ICM application, the VRU supplies this call reference ID to Unified ICM when the VRU asks for instructions on how to process the incoming call that it receives from the switch. This method enables Unified ICM to retrieve the appropriate call context for this same call, which at this stage is to proceed to the IVR portion of the call.

- **Correlation ID**—This method is used if the network can pass the call reference ID to the VRU when the VRU is located in the network with the switch and the call signaling can carry this information (for example, the Correlation ID information is appended to the dialed digits when Unified ICM is used). This function usually applies to calls being transferred within the VoIP network.
- **Translation Route ID**—This method is used when the VRU is reachable across the PSTN (for example, the VRU is at the customer premise) and the network cannot carry the call reference ID information in delivering the call to the VRU. You must configure a temporary directory number (known as a translation route label) in Unified ICM to reach the VRU, and the network routes the call normally to the VRU as with other directory number routing in the PSTN. When the VRU requests instructions from Unified ICM, the VRU supplies this label (which can be a subset of the received digits), and Unified ICM can correlate the two portions of the same call. Generally, the PSTN carrier contains a set of translation route labels to be used for this purpose.



Note The deployed VRU can be located in the network (Network VRU) or at the customer premises. At the customer premises, a Network Applications Manager (NAM) is deployed in the network and a Customer ICM (CICM) is deployed at the customer premises. The corresponding Correlation ID or Translation Route ID is used, depending on the location of the VRU.

Unified CVP Type 10 VRU

Unified CVP Type10 VRU simplifies the configuration requirements in Unified CVP Comprehensive Model deployments. Use the Type 10 VRU for new installations except for the VRU-only deployments. In deployments that need to use ICM Customers, you cannot initiate a two-step transfer from the UnifiedCVP VRU switch leg to a completely separate UnifiedCVP (for example, a two-step CVP-to-CVP transfer using SendToVRU). You are require to use a translation route for these two-step transfers to work.

Type 10 Network VRU operates as follows:

- Transferred routing client responsibilities are handed off to the Unified CVP switch leg.
- An automatic transfer to the Unified CVP VRU leg occurs resulting in a second transfer when calls are originated by the VRU, ACD, or Cisco Unified Communications Manager (UnifiedCM).
- For calls originated by Unified CM, the Correlation ID transfer mechanism is used. The Correlation ID is automatically added to the end of the transfer label defined in the Type 10 Network VRU configuration.
- The final transfer to the Unified CVP VRU leg is similar to a Type 7 transfer, which includes a RELEASE message to be sent to the VRU prior to any transfer.

You need to define a single Type 10 Network VRU in Unified CVP implementations without the ICM Customers feature (that is, in UnifiedCVP implementations with a single Network VRU), and associate all UnifiedICM VRU scripts. One label for the UnifiedCVP Switch leg routing client, transfers the call to the Unified CVP VRU leg. If calls are transferred to Unified CVP from Unified CM, another label for the UnifiedCM routing client, and this label should be different from the label used for the CVP Routing Client. This label transfers the call to the Unified CVP Switch leg. The Unified ICM Router sends this label to UnifiedCM with a Correlation ID concatenated to it. You must configure Unified CM to handle these arbitrary extra digits.

Configure the Unified CVP Switch leg peripheral to point to the same Type 10 Network VRU. Also, associate all incoming dialed numbers for calls that are to be transferred to Unified CVP with a Customer Instance that points to the same Type 10 Network VRU.

For calls that originate at a Call Routing Interface VRU or at a TDM ACD, a TranslationRouteToVRU node is required to transfer the call to a Unified CVP's Switch leg peripheral. For all other calls, use either a SendToVRU node, a node that contains automatic SendToVRU function (such as the queuing nodes), or a RunExternalScript.



Note Type 5 and Type 2 VRU types are not supported. Instead of these VRU types, use Type 10 VRU.

Correlation ID

Configure the appropriate route patterns for the Translation Route DNIS or VRU Label with Correlation ID appended. You use the Correlation ID method with Type 10 VRUs. Configure the route pattern in Unified CM to allow appending extra digits, such as an exclamation point (!), to the end of the pattern.

Unified CVP Type 7 VRU (Correlation ID Function)

When the VRU functions as an IVR with the Correlation ID function, Unified ICM uses Type3 and Type7 to designate suboperations of the VRU with the Peripheral Gateway in the Correlation ID scheme. Both Type3 and Type7 VRUs can be reached with the Correlation ID function, and a Peripheral Gateway is needed to control the VRU. However, the difference between these two types is in how they release the VRU leg and how they connect the call to the final destination.

In Type 3, the switch that delivers the call to the VRU can take the call from the VRU and connect it to a destination (or agent).

In Type 7, the switch cannot take the call away from the VRU. When the IVR treatment is complete, Unified ICM must disconnect or release the VRU leg before the final connect message can be sent to the Switch leg to instruct the switch to connect the call to the destination.

When used as an Intelligent Peripheral IVR, Unified CVP supports only Type 7 because it gets a positive indication from Unified ICM when its VRU leg is no longer needed (as opposed to waiting for the Voice Browser to inform it that the call has been pulled away). Type 3 is deprecated.

A call has two legs: the Switch leg and the VRU leg. Different Unified CVP hardware can be used for each leg. A service node along with a Unified CVP for VRU leg with Peripheral Gateway acting as VRU Type 7 can be used to complete the IVR application (for example, self service and queuing).



Note Use Type 10 VRU for all new implementations of Unified CVP using Unified ICM 7.1 or greater, except as VRU Only.

For configuration examples of the Unified CVP application with VRU Type 7, see the *Configuration Guide for Cisco Unified Customer Voice Portal* at http://www.cisco.com/en/US/products/sw/custcosw/ps1006/products_installation_and_configuration_guides_list.html.

Unified CVP Type 8 VRU (Translation Route ID Function)

When the VRU functions as an IVR with the Translation Route ID function, Unified ICM uses Type8 or Type10 to designate suboperations of the VRU through the Peripheral Gateway in the translation route scheme. Both Type8 and Type10 VRUs can be reached through the Translation Route ID mechanism, and Peripheral Gateway is needed to control the VRU. However, they differ in how they connect the call to the final destination.

In Type 8, the switch that delivers the call to the VRU can take the call from the VRU and connect it to a destination or agent.

When the switch cannot take the call away from the VRU to deliver it to an agent, use Type 10. In that case, when the IVR treatment is complete, Unified ICM sends the final connect message to the VRU (rather than to the original switch) to connect the call to the destination. The VRU assumes control of the switching responsibilities when it receives the call. This process is known as handoff.

Similar to the Correlation ID, there are two legs of the call: the Switch leg and the VRU leg. Use Unified CVP for either the Switch leg or the VRU leg. For example, when Network Interface Controller (NIC), NAM, or CICM is taken, configure Unified CVP as Type8 or Type10 in the VRU leg.



Note Use Type 10 VRU for new implementations of Unified CVP using Unified ICM 7.1 or greater, except as VRU Only.

For configuration examples of the Unified CVP application with VRU Type 8 or Type 10, see *Configuration Guide for Cisco Unified Customer Voice Portal* at http://www.cisco.com/en/US/products/sw/custcosw/ps1006/products_installation_and_configuration_guides_list.html.

Network VRU Types and Unified CVP Deployment

In Unified ICM, a Network VRU is a configuration database entity that you can access using Network VRU Explorer. A Network VRU entry contains the following information:

- Type—A number from 7, 8, and 10, which corresponds to one of the VRU types.

- **Labels**—A list of labels that you use in Unified ICM to transfer a call to the particular Network VRU. These labels are relevant only for Network VRUs of Type 7 or 10 (that is, those VRU types that use the Correlation ID Mechanism to transfer calls). Each label consists of two parts:
 - A digit string, which becomes a Dialed Number Identification Service (DNIS). A SIP Proxy Server or a static route table (when SIP is used), or gateway dial peers understand DNIS.
 - A routing client, or switch leg peripheral. Each peripheral device that acts as a Switch leg must have its label, although the digit strings are the same in all cases.

Network VRU configuration entries have no value until they are associated with active calls. Unified ICM association is made at the following locations:

- In the **Advanced** tab for a given peripheral in the PG Explorer tool
- In the Customer Instance configuration in the Unified ICM Instance Explorer tool
- In every VRU Script configuration in the VRU Script List tool

Depending on the protocol-level call flow, the currently used Unified ICM Enterprise looks at either the peripheral or the Customer Instance to determine how to transfer a call to a VRU. The Unified ICM Enterprise examines the Network VRU associated with the Switch leg peripheral when the call first arrives on a Switch leg, and examines the Network VRU that is associated with the VRU leg peripheral when the call is being transferred to the VRU using the Translation Route Mechanism. The Unified ICM Enterprise examines the Network VRU that is associated with the Customer Instance when the call is being transferred to the VRU using the Correlation ID Mechanism.

Unified ICM Enterprise also checks the Network VRU that is associated with the VRU Script every time it encounters a RunExternalScript node in its routing script. If Unified ICM determines that the call is currently not connected to the designated Network VRU, the VRU Script is not executed.

Unified ICM Enterprise Release 7.1 introduced Network VRU Type 10, which simplifies the configuration of Network VRUs for Unified CVP. For most call flow models, a single Type 10 Network VRU replaces the place of the Type 2, 3, 7, or 8 Network VRUs that are associated with the Customer Instance and the switch, and VRU leg peripherals. VRU Only (Model #4a) still requires Type 7 or 8.



Note For existing deployments, the previously suggested VRU types work in the similar way, new installations are required to use Type 10. Existing deployments should switch to Type 10 on upgrade.

Standalone Self-Service

Standalone Self-Service usually does not communicate with Unified ICM VRU scripts, so a Network VRU setting is not required. Standalone Self-Service with Unified ICM Label Lookup does not use the VRU scripts in Unified ICM. It issues a Route Request to the VRU Peripheral Gateway (PG) Routing Client, which does not require this Network VRU.

Standalone Self-Service Deployments ASR-TTS

If the ASR/TTS MRCP server fails, the following conditions apply to the call disposition:

- Calls in progress in standalone deployments are disconnected. Calls in progress in ICM-integrated deployments can be recovered using scripting techniques to work around the error. For example, retry the request, transfer to an agent or label, switch to the prerecorded prompts and DTMF-only input for

the rest of the call, an error will occur with an END script node, to invoke survivability on the originating gateway.



Note Cisco VVB has a built-in load-balancing mechanism that uses a round-robin technique. If the present ASR/TTS MRCP server fails, then the next request for MRCP resource will get to the next server in the server group.

In a call, if the selected ASR/TTS MRCP server responds with a failure to the setup request, then the VVB retries only once to set up with another server. If the VXML application has defined a preferred server for ASR dialog or TTS, then retry is not attempted.

For configuration steps, see the section “Configure Speech Server” in *CVP Configuration Guide*.

- New calls in ICM-integrated deployments are directed transparently to an alternate ASR/TTS Server if a backup ASR/TTS Server is configured on the gateway.

Call Director

Unified CCE and Unified CVP are responsible for call switching only. It does not provide queuing or self-service, so there is no VRU leg. A Network VRU setting is not required.

VRU-Only

In VRU-Only, the call arrives at Unified ICM through an ICM-NIC interface. Initially, Unified CVP is not responsible for the Switch leg; its only purpose is as a VRU. However, depending on the type of NIC used, it may be required to take over the Switch leg after it receives the call.

The following sections describe the VRU-Only types.

VRU-Only with NIC Controlled Routing



-
- Note** VRU-Only with NIC Controlled Routing has the following assumptions:
- A fully functional NIC can deliver the call temporarily to a Network VRU (that is, to Unified CVP's VRU leg) and then retrieve the call and deliver it to an agent when that agent is available.
 - If the agent is capable of requesting that the call is retransferred to another agent or back into queue or self-service, the NIC can retrieve the call from the agent and redelivering it as requested.
-

There are two variants of VRU-Only with NIC Controlled Routing, depending on whether the Correlation ID or the Translation Route function is used to transfer calls to the VRU. Most NICs (most PSTN networks) cannot transfer a call to a particular destination directory number and carry an arbitrary Correlation ID along with it. The destination device can pass back to Unified ICM to make the Correlation ID transfer mechanism function properly. For most NICs, you must use the Translation Route function.

However, a few exceptions to this rule exist, in which case the Correlation ID function can be used. The NICs transmits a Correlation ID include Call Routing Service Protocol (CRSP), and Telecom Italia Mobile (TIM).

However, because this capability also depends on the PSTN devices that connect behind the NIC, check with your PSTN carrier to determine whether the Correlation ID can be passed through to the destination.

If the NIC is capable of sending the Correlation ID, the incoming dialed numbers must all be associated with a Customer Instance that is associated with a Type 7 Network VRU. The Type7 Network VRU must contain labels that are associated to the NIC routing client, and all the VRU Scripts must also be associated with that same Type 7 Network VRU. The peripherals do not need to be associated with any Network VRU. We recommend that to execute the Unified ICM routing script SendToVRU node prior to the first RunExternalScript node.

If the NIC cannot send a Correlation ID, the incoming dialed numbers must all be associated with a Customer Instance that is not associated with any Network VRU. However, the Unified CVP peripherals must be associated with a Network VRU of Type8, and all the VRU Scripts must also be associated with that same Type8 Network VRU. In this case, it is necessary to insert a TranslationRouteToVRU node in the routing script prior to the first RunExternalScript node. If the call is going to the VRU leg because it is being queued, generally the TranslationRouteToVRU node appears after the Queue node. In that way, you can avoid an unnecessary delivery and removal from Unified CVP when the requested agent is already available.

VRU-Only with NIC Controlled Prerouting



Note VRU-Only with NIC Controlled Prerouting assumes a less capable NIC that can deliver the call only once, either to a VRU or to an agent. After the call is delivered to retrieve a call and then it is redelivered somewhere else, Unified CVP takes control of the switching responsibilities for the call. Unified ICM considers this process as a handoff.

Calls that fit VRU-Only with NIC Controlled Prerouting must use the Translation Route function to transfer calls to the VRU. A handoff cannot be implemented using the Correlation ID function.

To implement VRU-Only with NIC Controlled Prerouting with Unified ICM 7.1 and later, the incoming dialed numbers must all be associated with a Customer Instance that is associated with a Type 10 Network VRU. The VRU labels are associated with the Unified CVP routing client, not the NIC. The Unified CVP peripherals and VRU Scripts must be associated with the Type10 Network VRU. You need to insert a TranslationRouteToVRU node in the routing script, followed by a SendToVRU node, before the first RunExternalScript node. If the call is going to the VRU leg because it is being queued, these two nodes should appear after the Queue node. An unnecessary delivery and removal from Unified CVP can be avoided if the requested agent is already available.



Note Two different VRU transfer nodes are required. The first node transfers the call away from the NIC with a handoff, and it establishes Unified CVP as a Switch leg device for this call. Physically the call is delivered to an Ingress Gateway. The second transfer delivers the call to the Voice Browser and also establishes Unified CVP as the call's VRU device.

Unified ICM Hosted Deployments

Unified ICM Hosted implementations incorporate a two-level hierarchy of Unified ICM systems. The Network Application Manager (NAM) is at the top level, and one or more Customer ICMs (CICMs) are below it. Both the NAM and CICM are complete ICMs, with a communication link between them known as Intelligent Network Call Routing Protocol (INCRP). Each CICM functions in an isolated way; it is unaware of the other

CICMs and that the NAM is another ICM. A CICM has no connection to the other CICMs, but its connection to the NAM is through the INCRP NIC.

Unified ICM Hosted is for service providers who provide contact center services to multiple customers. The NAM routes calls to the appropriate CICM for each customer. Your customers run their own contact centers with their own Automatic Call Distributors (ACDs) connected to PGs at their own premises. The PGs connect to their assigned CICMs at your main site. You, as the service provider, own and host the NAM and all the CICMs, but your customers own and host all the ACDs. You own the PGs for those ACDs, but the PGs are located at the customer's premises, next to the ACDs. Those PGs connect either to an assigned CICM or to the NAM.

For ICM scripting, all incoming calls initially invoke an appropriate NAM routing script. That script has the primary responsibility of identifying the appropriate target customer. After the NAM, these actions identify the appropriate target customer:

- The script delegates control to a routing script that runs on that customer's CICM.
- The CICM-based routing script selects the appropriate ACD to receive the call. That script can return the necessary translation route label to the NAM.
- The NAM instructs its routing client to deliver the call to the designated target ACD. The NAM places the call into queue at a Service Control VRU, if either of the following happens:
 - The CICM routing script determines that no ACD can currently take the call.
 - The CICM routing script cannot yet identify which ACD should take the call.
- The CICM routing script issues Network VRU Script requests through the NAM to that VRU until a routing decision is made.

Some hosted customers use this topology to get more calls or more PGs through their ICM setup. Other customers use CICMs, not for customer contact centers, but for external customers. In these situations, the NAM handles the same number of calls as the CICM, and the CICM machines is located far away from the NAM.



Note The NAM and CICM architecture is an old design that assumes all contact centers run on TDM-based ACDs. The addition of VoIP routing and Unified CCE with direct agent routing made this architecture more complicated.

CVP in Hosted Environments

Unified CVP in a Unified ICM Hosted environment is a self-service or queuing platform on the NAM in the service provider's main site. With CVP, a service provider can:

- Route calls to the appropriate customer-owned ACDs.
- Retain control of calls that are queued for those ACDs.
- Provide either basic prompt-and-collect capability or full-featured self-service applications to their customers.

The last case typically incorporates VXML Servers into the network. You can host the VXML Servers or your customer can host them. Either you or your customers can write and own the self-service applications.

Allowing your customer to own or host the VXML Servers is a convenient solution for self-service applications that reference back-end services. This solution allows your customer to retain control of that interaction within its own enterprise network. Your customer only transmits VXML over HTTP to your VXML Gateway.

In many ICM Hosted environments, when the service provider is a PSTN carrier, all of the actual call routing occurs through an ICM NIC. The same situation applies when a PG routes calls using (typically) the ICM NIC. However, the service provider might not have a NIC interface at all. Then, all calls arrive through TDM interfaces, such as T3 or E3. In those cases, CVP handles the Switch and the VRU legs.

CVP Placement and Call Routing in Hosted Environments

In the Unified ICM Hosted environment, CVP connects at the NAM in a valid Network VRU. Some deployments might also require a CVP at the CICM level, instead of or in addition to the NAM level. Consider these guidelines on where to place Unified CVP components:

- When a CVP at the NAM handles both the Switch and VRU legs, use the Correlation ID transfer function. Either the NAM or the CICM routing script can execute the SendToVRU node. (Place the RunExternalScript nodes in the same script that executes the SendToVRU.)
- If a CVP at the NAM handles the VRU leg and a NIC handles the Switch leg, you can use either the Correlation ID or the Translation Route transfer functions. This depends on the capabilities of the NIC. These guidelines also apply:
 - If you use Correlation ID transfer, place the SendToVRU node in either the NAM or the CICM routing script. (Place the RunExternalScript nodes in the same script that executes the SendToVRU.)
 - If you use Translation Route transfer, place the TranslationRouteToVRU node and all RunExternalScript nodes in the NAM routing script. The call is queued (or treated with prompt-and-collect) before selecting a particular CICM. This configuration does not facilitate queuing. However, this configuration enables service providers to offer initial prompt-and-collect before delegating control to the CICM.
- If a CVP at the CICM handles the VRU leg and a NIC handles the Switch leg, you can only use the Translation Route transfer method. Place the TranslationRouteToVRU node and all RunExternalScript nodes in the CICM routing script.

Adding calls initiated by Unified CM or an ACD creates additional constraints. Both of these devices are considered ACDs from the ICM perspective and they connect at the CICM level. For new calls, rather than transfers, the route request comes from the ACD and the resulting label returns to the ACD. Unified CM or an ACD cannot send a Correlation ID upon transfer. You can only use the Translation Route transfer method. This limitation means that the transfer destination connects at the CICM level, and not the NAM level.

For transfers, your customers might want either blind network transfers or warm consultative transfers. The following guidelines apply to these transfers:

- **Blind network transfers**—If the original call comes to the NAM through a NIC or a CVP Switch leg, the CICM passes the transfer label to the NAM to the original Switch leg device. Blind network transfers have two subcases:
 - If the Switch leg device is CVP or a NIC that can handle Correlation ID, use the Correlation ID transfer function. Place the SendToVRU node and all RunExternalScript nodes in the CICM routing script. Connect the CVP VRU leg to the NAM. This combination of blind transfer with Correlation ID transfer is suitable for CVP.

- If the Switch leg device is a NIC that cannot handle Correlation ID, use the Translation Route transfer method. Connect the CVP VRU leg device to the CICM.



Note In this situation, you need extra dedicated CVP Call Servers at the CICM level because you cannot use the NAM-level CVP Call Servers.

- **Warm consultative transfers**—CVP supports warm consultative transfers only for Unified CCE agents transferring calls to other Unified CCE agents. For such calls, CVP owns the initial Switch leg for the inbound call. Use the Translation Route transfer method because Unified CM cannot handle Correlation ID transfers. Connect the CVP VRU leg device to the CICM.

For TDM agents, you use the ACD functions instead of CVP. When the incoming calls to Unified CCE agents arrive through a NIC, use the Unified ICM Network Consultative Transfer facility instead of CVP.



Note In this situation, you need extra dedicated CVP Call Servers at the CICM level because you cannot use the NAM-level CVP Call Servers.

Network VRU Type in Hosted Environments

In an ICM Hosted environment, there are two types of ICM systems, the NAM and the CICM. You configure Network VRU types differently in the NAM and CICM.

The NAM gets incoming calls either from the NIC or from Unified CVP, and is aware of the Unified CVP VRU leg device. Configure the NAM Network VRU types exactly as an independent ICM, operating with those devices. You can ignore that the transfer labels sometimes come from CICM when configuring Network VRU types. The CICM sees incoming calls that arrive from the Intelligent Network Call Routing Protocol (INCRP) NIC.

Associate the dialed numbers that arrive from the NAM with a Customer Instance for the corresponding Network VRU on a CICM. Associate that Network VRU with all VRU scripts. Provide the same label that you need in the NAM Network VRU definition, but with the INCRP NIC as its routing client. No peripherals have Network VRUs configured.

For more information on Network VRU Type support, see the *Configuration Guide for Cisco Unified Customer Voice Portal* at <http://www.cisco.com/c/en/us/support/customer-collaboration/unified-customer-voice-portal/products-installation-and-configuration-guides-list.html>.

Unified CM, ACD Call Deployments, and Sizing Implications

The information in this section applies to ACDs and Cisco Unified Communications Manager (UnifiedCM) integrations that use Unified CVP instead of Cisco IPIVR for queuing. If Unified CVP is considered, these devices share the following characteristics:

- Manage agents and can be destinations for transfers.
- Can issue Route Requests and be Switch leg devices.
- Although they can be Switch leg devices, they cannot handle more than one transfer and they might not be able to handle the Correlation ID.

- Manage agents and transfer calls to the destination.
- Route requests and be switch leg devices. However, the device cannot handle Correlation ID and more than one transfer.

A UnifiedCM or ACD user issues a Route Request for one of the following reasons:

- Connect to another agent in a particular skill group
- Reach a self-service application
- Blind-transfer a previously received call to one of the above entities

A UnifiedCM user might also issue a Route Request for one of the following reasons:

- Deliver a successful outbound call from the Unified ICM Outbound dialer to a self-service application based on Unified CVP
- Warm-transfer a call that the user had previously received to either a particular skill group or a self-service application

Each of the above calls invokes an Unified ICM routing script. The script searches for an available destination agent or service and if an appropriate destination is found, it sends the corresponding label either back to the ACD or, if blind-transferring an existing call, to the original caller's Switch leg device. If it needs to queue the call or if the ultimate destination is intended to be a self-service application rather than an agent or service, the script sends a VRU translation route label either back to the ACD or, if transferring an existing call through blind-transfer, to the original caller's Switch leg device.

If the above sequence results in transferring the call to Unified CVP's VRU leg device, a second transfer is done to deliver it to a Voice Browser. To ensure that these events take place, the following Unified ICM configuration elements are required:

- For new calls from the ACD or warm transfers of existing calls:
 - Configure the Unified CVP peripheral to be associated with a Type10 Network VRU.
 - Associate the dialed number that the ACD dialed with a Customer Instance that is associated with a Type 10 Network VRU.
 - When an ACD is not configured UnifiedCM, the routing script that is invoked by the ACD dialed number must contain a TranslationRouteToVRU node to get the call to Unified CVP's Switch leg, followed by a SendToVRU node to get the call to the Voice Browser and Unified CVP's VRU leg.
 - The routing script that is invoked by UnifiedCM should use a SendToVRU node to send the call to Unified CVP using the Correlation ID. The Type10 VRU performs an automatic second transfer to the Voice Browser VRU leg.
 - Associate all the VRU scripts that are executed by that routing script with the Type10 Network VRU.
- For blind transfers of existing calls:
 - The Unified CVP peripheral can be associated with any Network VRU.
 - The dialed number that the ACD dialed must be associated with a Customer Instance that is associated with a Type 10 Network VRU.

- The routing script that is invoked by the ACD dialed number must contain a SendToVRU node to send the call to the Voice Browser and Unified CVP's VRU leg.
- All the VRU scripts that are executed by that routing script must be associated with the Type10 Network VRU.

When Unified ICM chooses an agent or ACD destination label for a call, it tries to find one that lists a routing client that can accept that label. For calls originated by an ACD or UnifiedCM that are not blind transfers of existing calls, the only routing client is the ACD or UnifiedCM, after the call is transferred to Unified CVP, because of the handoff operation, the only routing client is the Unified CVP Switch leg. However, in the case of blind transfers of existing calls, two routing clients are possible:

- The Call Server switch leg that delivered the original call.
- The ACD or UnifiedCM. For calls that originate through Unified CVP, you can prioritize Unified CVP labels above ACD or UnifiedCM labels by checking the **Network Transfer Preferred** check box in the **Unified ICM Setup** screen for the Unified CVP peripheral.

In addition to the CTI Agent Desktop, the system recognizes a configured routing Dialed Number to intercept the transfer and run the Unified CCE routing script without sending the transfer commands to Unified CM through JTAPI.

Third-Party VRU

A third-party TDM VRU can be used in any of the following ways:

- As the initial routing client (using the GED-125 Call Routing Interface)
- As a VRU (using the GED-125 Call Routing Interface)
- As a Service Control VRU (using the GED-125 Service Control Interface)

In the first and second operations, the VRU works as an ACD. Similar to ACD, the VRU can be a destination for calls that arrive from another source. Calls can even be translation-routed to such devices to carry call context information. (This operation is known as a traditional translation route, not a TranslationRouteToVRU). Also like an ACD, the VRU can issue its own Route Requests and invoke routing scripts to transfer the call to subsequent destinations or even to Unified CVP for self-service operations. These transfers almost always use the Translation Route transfer function.

In the third operation, the VRU replaces either Unified CVP's Switch leg or Unified CVP VRU leg, or it can also replace Unified CVP.

Release Trunk Transfer

Release Trunk Transfer releases the ingress trunk and removes Unified CVP and the gateway from the call control loop. These transfers have the following characteristics:

- They can be invoked by VXMLServer (Standalone Call Flow Model) or using Unified ICM.
- Unified ICM Network Transfer using Unified CVP as the routing client does not work because Unified CVP can no longer control the call.
- They are blind, that is, if the transfer fails for any reason then Unified ICM does not recover control of the call. Router Requery is not supported.

- They cause the switch leg to terminate resulting in a Telephony Call Dispatcher (TCD) record being written to the database for the call even though the caller is still potentially talking to an agent. This behavior differs from other types of transfers in which the TCD record is not finalized until the caller hangs up.
- As the ingress trunk is released, you do not have to size gateways to include calls that have been transferred using Release Trunk Transfer. This behavior differs from other types of transfers in which gateway resources continue to be occupied until the caller hangs up.
- Because Unified CVP is no longer monitoring the call, you do not have to size Call Servers to include calls that have been transferred using Release Trunk Transfer. Additionally, Unified CVP Call Director port licenses are not required.

The signaling methods to trigger a release trunk transfer are:

- Takeback-and-Transfer
- Hookflash and Wink
- Two B Channel Transfer

Takeback-and-Transfer

Takeback-and-Transfer (TNT), also known as Transfer Connect, is a transfer method where dual tone multifrequency (DTMF) tones are outpulsed to the PSTN by Unified CVP. TNT outpulses DTMF tones to the PSTN. A typical DTMF sequence is *8xxxx, where xxxx represents a new routing label for the PSTN. Upon detection of a TNT DTMF sequence, the PSTN drops the call leg to the Ingress Gateway port, and then reroutes the caller to a new PSTN location, such as a TDM ACD location. This method is offered by a few PSTN service providers.

Customers can use TNT, if they have an existing ACD site but no IVR and want to use Unified CVP as an IVR. Over time, customers may need to transition agents from the TDM ACD to Unified CCE and use Unified CVP as an IVR, queueing point, and transfer pivot point. Using Unified CVP as more than just an IVR eliminates the need for TNT services.

In Unified CVP deployments with Unified ICM, the DTMF routing label that is outpulsed can be a Unified ICM translation routing label to enable passing of call data to another Unified ICM peripheral, such as a TDM ACD. In this scenario, Unified CVP views the call as completed, and Unified CVP call control is ended. With TNT, if the transfer to the termination point fails, Unified CVP cannot reroute the call. Using some TNT services, you can reroute the callback to Unified CVP. However, Unified CVP treats this call as a new call.

Two B Channel Transfer

Two B Channel Transfer (TBCT) is an Integrated Services Digital Network (ISDN)-based release trunk signaling function that is offered by some public switched telephone network (PSTN) service providers. When a TBCT is invoked, the Ingress Gateway places the initial inbound call on hold briefly while a second call leg (ISDN B Channel) is used to call the termination point. When the termination point answers the call, the gateway sends ISDN signaling to the PSTN switch to request to complete the transfer, bridge the call through the PSTN switch, and remove the call from the Ingress Gateway. As with a TNT transfer, the termination point might be a TDM PBX or ACD connected to the PSTN.

This process may be necessary for a customer with an existing ACD site but no IVR, who wants to use Unified CVP initially as just an IVR. Over time, the customer might want to transition agents from the TDM ACD to Cisco Unified CCE and use Unified CVP as an IVR, queueing point, and transfer pivot point (which eliminates the need for TBCT services and using Unified CVP to perform reroute on transfer failure).

VXML Transfer

VXML call control is supported only in Standalone deployments in which call control is provided by the VXMLServer. In Unified ICM integrated deployments, ICM controls all calls.

The VXMLServer supports the following types of transfers:

Table 4: Types of VXML Transfers

Release Trunk Transfer	VXML Blind Transfer	VXML Bridged Transfer
Incoming call is released from the Ingress Voice Gateway.	Call is bridged to an Egress Voice Gateway or a VoIP endpoint. However, the VXMLServer releases all subsequent call control.	Call is bridged to an Egress Voice Gateway or a VoIP endpoint. However, the VXMLServer retains call control so that it can return a caller to an IVR application or transfer the caller to another termination point.
Invoked using the subdialog_return element. VXMLServer can invoke a TNT transfer, Two B Channel transfer, HookFlash/Wink transfers, and SIP Refer Transfers. For TDM Release Trunk Transfers (TNT, TBCT and Hookflash/Wink), the VXML Gateway must be combined with the Ingress Gateway for the Release Trunk Transfer to work.	Invoked using the Transfer element in Cisco Unified Call Studio. These transfers transfer the call to any dial peer that is configured in the gateway.	Bridged transfers do not terminate the script. The VXMLServer waits until either the ingress or the destination call ends. The script ends only if the ingress call leg hangs up. If the destination call leg hangs up first, the script recovers control and continues with additional self-service activity. Note that the VXMLServer port license remains in use for the duration of a bridged transfer, even though the script is not actually performing any processing.

VXML blind transfers differ from VXML bridged transfers in the following ways:

- VXML blind transfers do not support call progress supervision; bridged transfers support it. This means that if a blind transfer fails, the VXML Server script does not recover control and cannot attempt a different destination or take remedial action.
- VXML blind transfers cause the VXMLServer script to end. Always connect the “done exit” branch from a blind transfer node to a subdialog_return and a hang up node.



Note Cisco VVB supports Blind Transfer only under VXML Transfer.

Non-Reference Call Flows

Unified CVP offers several call flows to support different needs. The appropriate deployment for your contact center depends on the call flow, geographic distribution, and server configurations that best meet your needs.

The Comprehensive call flow is the only one that you can use in a Reference Design. You can use the following alternate call flows in a Non-Reference Design:

- UnifiedCVP VXMLServer (standalone)—Provides a standalone VRU with no integration to Unified CCE for queuing control or subsequent call control. Used to deploy self-service VXML applications.
- Call Director—Provides IP switching services only. Use this call flow if you want to:
 - Only use Unified CVP to provide Unified CCE with VoIP call switching.
 - Prompt and collect data using third-party VRUs and ACDs.
 - Avoid using the Unified CVP VXML Server.
- VRU Only—Provides VRU services, queuing treatment, and switching for PSTN endpoints. This model relies on the PSTN to transfer calls between call termination endpoints. Use this call flow if you want to:
 - Use Unified CVP to provide Unified CCE with VRU services, including integrated self-service applications and initial prompt and collect.
 - Avoid using Unified CVP for switching calls.
 - Use an optional Unified CVP VXML Server.
 - Prompt or collect data using optional ASR and TTS services.

Standalone Self-Service

Standalone Self-Service does not involve Unified ICM. It is implemented when a UnifiedCM user dials a directory number that connects to a VXML Gateway and invokes a Unified CVP VXMLServer application. The VXML Gateway is configured in UnifiedCM as a SIP trunk. The call flow is as follows:

1. A caller dials a route pattern.
2. UnifiedCM directs the call to the VXML Gateway.
3. The VXML Gateway invokes a voice browser session based on the configured Unified CVP self-service application.
4. The Unified CVP self-service application makes an HTTP request to the UnifiedCVP VXMLServer.
5. The UnifiedCVP VXMLServer starts a self-service application.
6. The UnifiedCVP VXMLServer and VXML Gateway exchange HTTP requests and VXML responses.
7. The caller hangs up.

Call Director

Call Director has only switching and has no VRU leg. Calls originated by UnifiedCM are always delivered directly to their targets or rejected. No queuing or self-service is involved.

Call Director assumes that the call is truly originating from UnifiedCM. It excludes calls that originally arrived through a Ingress VXML Gateway and were transferred to UnifiedCM, and are now transferred again. These situations are rare because UnifiedCM can usually handle those transfers itself. There are exceptions, however, such as when the target is an ACD other than UnifiedCM.

Call Director requires the following items are configured:

- UnifiedCM route point that invokes a Unified CCE script
- Unified CVP configured as a Type 10 NetworkVRU
- VRU translation routes to Unified CVP
- Translation route Dialed Number Identification Service (DNIS) numbers configured in the Unified CVP Call Server
- UnifiedCM configured with a SIP trunk
- UnifiedCM route patterns for Translation Route DNIS

The call flow is as follows:

1. A caller dials a route point.
2. Unified ICM invokes a routing script.
3. The routing script encounters a TranslationRouteToVRU node to transfer the call to Unified CVP. (Unified CVP is configured as a Type 10 NetworkVRU.)
4. Unified ICM returns the translation route label to UnifiedCM.
5. UnifiedCM consults the SIP Proxy to locate the Unified CVP Call Server.
6. UnifiedCM connects the call to the Unified CVP Call Server.
7. The routing script encounters a Select or Label node, and it selects a target label.
8. Unified ICM returns the target label to the Unified CVP Call Server (not to the device that issued the route request).
9. The Unified CVP Call Server consults the SIP Proxy to locate the destination device.
10. The Unified CVP Call Server communicates through SIP with the target device and instructs UnifiedCM to establish a media stream to it.

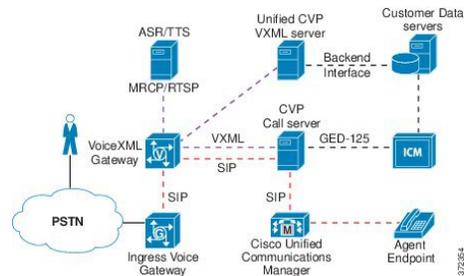
If the target device issues another route request to Unified ICM. This part of the call flow is not possible without the initial TranslationRouteToVRU mentioned in Step3.
11. Unified ICM invokes a new routing script.
12. The routing script encounters a Select or Label node, and it selects a target label.
13. Unified ICM returns the target label to the Unified CVP Call Server (not to the device that issued the route request).
14. The Unified CVP Call Server consults the SIP Proxy to locate the destination device.
15. The Unified CVP Call Server communicates using SIP with the target device and instructs UnifiedCM to establish a media stream to the device.

VXML Server (Standalone)

The VXML Server (standalone) functional deployment provides organizations with a standalone IVR solution for automated self-service. Callers can access Unified CVP from a local, long-distance, or toll-free numbers that terminate at Unified CVP Ingress Voice Gateways. Callers can also access Unified CVP from VoIP

endpoints number that terminates at Unified CVP Ingress Voice Gateways. Callers can also access Unified CVP from VoIP endpoints. The following figure illustrates this deployment.

Figure 11: VXML Server (Standalone) Functional Deployment



This deployment includes the following components:

- Ingress Voice Gateway
- Virtualized Voice Browser/VXML Gateway
- VXML Server
- Unified Call Studio
- Operations Console Server

Following are the optional components of this deployment:

- Automatic Speech Recognition/Text-to-Speech (ASR/TTS) Server
- Third-Party Media Server
- Egress Voice Gateway
- Reporting Server
- CVP Call Server
- ICM
- Cisco Unified Communications Manager

Protocol-Level Call Flow

1. A call arrives at the Ingress VXML Gateway through TDM or SIP. The gateway performs inbound Plain Old Telephone Service (POTS) or VoIP dial-peer matching.



Note Ingress Voice Gateway and VXML Gateway deployment can either be separate entities or co-located (using IOS VXML Gateway). The self-service application is invoked on VXML Gateway.

2. The selected VXML Gateway port invokes the Unified CVP self-service script.
3. The self-service script invokes the Unified CVP standalone bootstrap VXML document, which sends an HTTP request to the configured IP address of the VXMLServer.

4. The VXML service function, which resides in the CVP Server, runs the application that is specified in the HTTP URL and returns a dynamically generated VXML document to the VXML Gateway. The Unified CVP VXML Service may access back-end systems to incorporate personalized data into the VXML document that is sent to the VXML gateway.
5. The VXML Gateway parses and renders the VXML document. For spoken output, the VXML Gateway either retrieves and plays back prerecorded audio files that are referenced in the VXML document, or it streams media from a text-to-speech (TTS) server. DTMF is detected on VXML Gateway.
6. The VXML Gateway submits an HTTP request containing the results of the caller input to the VXML Server. This server runs the application that is specified in the HTTP URL again and returns a dynamically generated VXML document to the VXML Gateway. The dialog continues through repetition of Steps 5 and 6.
7. The VRU dialog ends when the caller hangs up, the application releases, or the application initiates a transfer.

Transfers and Subsequent Call Control

You can use the VXML Server (Standalone) functional deployment to transfer callers to another endpoint. Transfers to either VoIP (for example, Cisco Unified Communications Manager) or TDM (for example, Egress Voice Gateway to PSTN or TDM ACD). IVR application data cannot be passed to the new endpoint with this deployment. If the endpoint is a TDM ACD, the agent screen popup window does not appear.

This deployment supports the following call transfers:

- VXML Bridged Transfer
- VXML Blind Transfer
- Release Trunk Transfer (TNT, hookflash, TBCT, and SIP Refer)



Note Cisco VVB supports the Blind Transfer feature only.

Call Director

The Call Director deployment provides an organization with a method to route and transfer calls across a VoIP network. For example, organizations can use this deployment with multiple TDM ACD and TDM IVR locations that are integrated with Unified ICM through an ACD or IVR Peripheral Gateway. The organization wants to use the Cisco Unified Intelligent Contact Management (Unified ICM) to route and transfer calls intelligently across these locations without using PSTN prerouting or Release Trunk Transfer services. In this deployment, Unified CVP and Unified ICM can also pass call data between these ACD and IVR locations. Unified ICM can also provide beginning-to-end reporting for all calls.

The Call Director is often the first step in the migration from a TDM-based contact center to a VoIP-based contact center. When an organization is ready to implement CVP-based IVR services and Cisco Unified Contact Center Enterprise (Unified CCE), the organization can migrate their Unified CVP deployment to the Comprehensive functional deployment.

Callers can access Unified CVP from local, long-distance, or toll-free numbers that terminate at Unified CVP Ingress Voice Gateways. Callers can also access Unified CVP from VoIP endpoints.

Following are the required components of this Call Director deployment:

- Ingress Voice Gateway
- Egress Voice Gateway
- CVP Server
- Operations Console Server
- Cisco Unified ICM Enterprise
- SIP Proxy Server (for SIP deployments)

The reporting server is an optional component of this model because there is very little call information stored in the UnifiedCVP reporting database.

SIP-Level Call Flow

VoIP-Based Prerouting

1. A call arrives at the Ingress Voice Gateway and sends a SIP INVITE/SIP/SIP message to the SIP Proxy Server, which forwards the request to the CVP Server SIP Service.
2. The SIP Service sends a route request to Unified ICM using the CVP Server ICM Service and the VRU Peripheral Gateway. This route request invokes Unified ICM to run a routing script based on the dialed number and other criteria.
3. The Unified ICM routing script selects a target and returns a translation route label to the CVP Server SIP Service. The Server SIP Service then signals through the SIP Proxy Server to the Egress Voice Gateway (which connects to the TDM termination point) and the Ingress Voice Gateway to enable the call to be set up between the Ingress and Egress Voice Gateways. While the RTP stream flows directly between the Ingress and Egress Voice Gateways, call control signaling flows through Unified CVP to allow subsequent call control.
4. When the call arrives at the selected termination, the termination equipment sends a request to its Peripheral Gateway for routing instructions. This step resolves the translation route and allows any call data from the previous Unified ICM script to be passed to the selected termination. If the selected termination is a TDM IVR platform, self-service is provided and the caller can either release or request to be transferred to a live agent. If the selected termination is a TDM ACD platform, then the caller is queued until an available agent is selected by the TDM ACD. Call data can then be displayed on the agent screen. After receiving live assistance, the caller can either release or request to be transferred to another agent.

VoIP-Based Transfer

1. Regardless of whether the call is initially routed to a TDM IVR or ACD location, a caller can request a call to be transferred to another location. When the transfer occurs, the TDM IVR or ACD sends a postroute request with call data (with its Peripheral Gateway) to Unified ICM.
2. When Unified ICM receives this postroute request, it runs a routing script based on the transfer dialed number and other criteria. The Unified ICM routing script selects a new target for the call and then signals to the CVP Server SIP Service to release the call leg to the originally selected termination and to extend the call to a new termination.
3. When the call arrives at the new termination, the termination equipment sends a request to its PG for routing instructions. This step resolves a translation route that is allocated for this call to this new termination location, and it allows any call data from the previous location (IVR port or agent) to be passed to the new termination. Calls can continue to be transferred between locations using this same VoIP-based transfer call flow.

Transfers and Subsequent Call Control

In addition to the transfers managed by Unified ICM, the Call Director deployment can transfer calls to non-ICM terminations or invoke a Release Trunk Transfer in the PSTN. If a call is transferred to a non-ICM termination, then no call data can be passed to the termination; no further call control is possible for that call; and the cradle-to-grave call reporting, that Unified ICM captures, is completed. In the case of a Release Trunk Transfer, the Ingress Voice Gateway port is released; no call data can be passed to the termination; and no further call control is possible for that call. If the Release Trunk Transfer is translation-routed to another ICM peripheral, call data and cradle-to-grave reporting can be maintained.

Call Server cancels the transfer request and sends a transfer failure indication to Unified ICM. Following are the reasons for it:

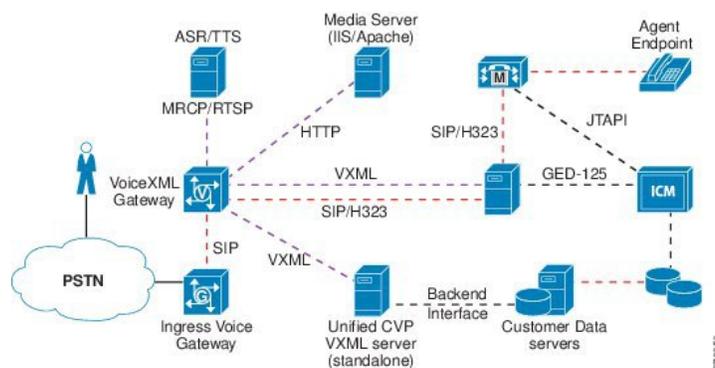
- A selected termination (for either a new or transferred call) returns a connection failure or busy status.
- The destination phone rings until it exceeds the ring-no-answer (RNA) timeout setting of Call Server.

This scenario invokes the Router Requery operation. The Unified ICM routing script then recovers control and selects a different target or takes a remedial action.

VRU-Only

The VRU-Only functional deployment provides self-service applications and queuing treatment for organizations that use advanced PSTN switching services that are controlled using a Cisco Unified ICM PSTN Network Interface Controller (NIC). Two Unified ICM PSTN NICs allow subsequent call control of calls in the PSTN: the NIC and the Carrier Routing Service Protocol (CRSP) NIC. These NICs allow Unified ICM to route calls intelligently to Unified ICM peripherals, such as ACDs and IVRs. They also allow Unified ICM to invoke mid-call transfers in the PSTN. The following figure illustrates this model.

Figure 12: VRU-Only Functional Deployment



In the VRU-Only functional deployment:

- Unified ICM routes a call before it is routed to another calls to Ingress Voice Gateway for call treatment and queuing. When an agent becomes available, Unified ICM instructs the PSTN to transfer the call to that agent. The agents can be Cisco Unified Contact Center Enterprise agents, Cisco Unified Contact Center Express agents, or ACD agents. If necessary, Unified ICM can request the PSTN (using the NIC) to transfer the call, just as Unified ICM can request Unified CVP to transfer the call.
- Ingress Voice Gateway is a Unified ICM-managed PSTN termination point that provides VRU services using a VXML Gateway, the VXML server, the ICM Service, and Unified ICM.
- The SIP Service is not used for call control. All call control and switching is controlled by Unified ICM and the PSTN.

- Unified ICM can pass call data between these termination points (for a pop-up window or other intelligent treatment) and provide reporting for all calls.

Following are the required components of this deployment:

- Ingress Voice Gateway
- VXML Gateway
- CVP Server
- Unified Call Studio
- Cisco Unified ICM Enterprise and NIC (CRSP)

Following are the optional components of this deployment:

- ASR/TTS Server
- Third-Party Media Server
- SIP Proxy Server (for SIP deployments)
- Reporting Server

Non-Reference Design Deployments of the Unified CM

A single UnifiedCM can originate and receive calls from SIP. PSTN calls that arrived on MGCP Voice Gateways registered with Unified CM can be routed or transferred to Unified CVP only through SIP (and not through Cisco Unified Border Element).



Note SCCP-based line side protocol is not supported in newer phones. SCCP-based line side protocol is supported in older phones in earlier releases. However, fresh installation of or enabling the deprecated phones is not supported.

For information on phones supported in the Cisco Unified Call Manager (CUCM) releases, see the *Deprecated Phone Models* section of the CUCM compatibility information for the relevant release at the following URL: <https://www.cisco.com/c/en/us/support/unified-communications/unified-communications-manager-callmanager/products-device-support-tables-list.html>

UnifiedCM is an optional component in the Unified CVP solution. Its use in the solution depends on the type of call center being deployed. TDM-based call centers using ACDs, for example, typically do not use UnifiedCM (except when they are migrated to Cisco Unified CCE). The same can be said for self-service applications that use the Unified CVP Standalone self-service deployment model. UnifiedCM is used as part of the Unified CCE solution. Call center agents are part of an IP solution that uses Cisco IP Phones, or when migrated from TDM ACDs.

Cisco Unified Communications Manager is a software application that controls the Voice Gateways and IP phones, providing the foundation for a VoIP solution. Unified Communications Manager runs on Cisco Unified Computing System (UCS) hardware or a specification-based equivalent. The software running on a VM is referred to as a Unified Communications Manager server. Multiple Unified Communications Manager servers can be grouped into a cluster to provide for scalability and fault tolerance. Unified Communications Manager communicates with the gateways using standard protocols such as Media Gateway Control Protocol (MGCP)

and Session Initiation Protocol (SIP). Unified Communications Manager communicates with the IP phones using SIP or Skinny Call Control Protocol (SCCP). For details on Unified Communications Manager call processing capabilities and clustering options, see the latest version of the *Cisco Collaboration System Solution Reference Network Designs* at <http://www.cisco.com/go/ucsrnd>.

Unified CCE Agent Desktop Options

Cisco offers the following interfaces for Unified CCE agents:

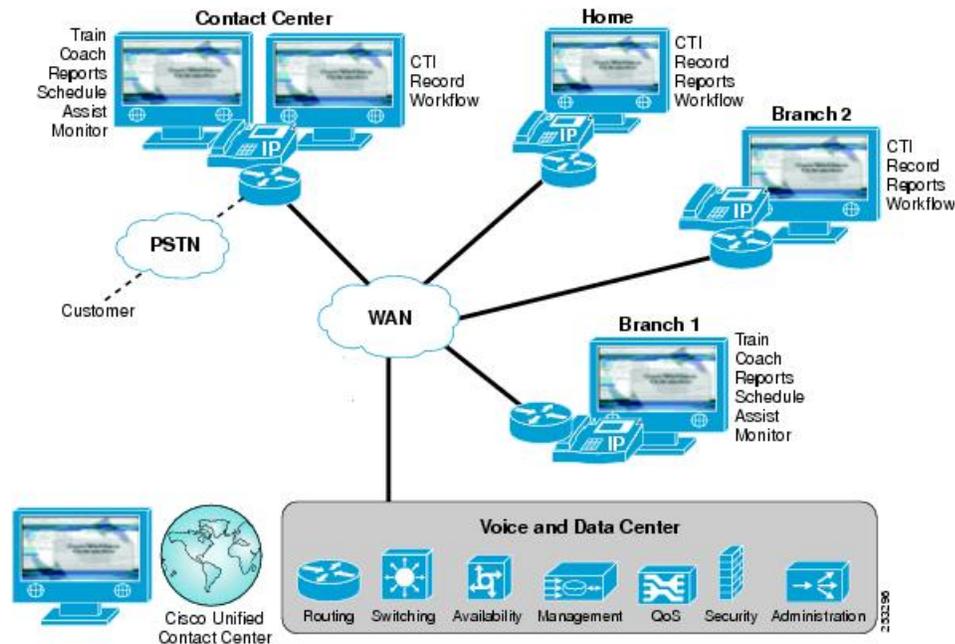
- **Cisco Finesse Desktop**—Cisco Finesse is a web-based desktop solution that allows for the extension of the desktop through standardized web components. Cisco Finesse offers:
 - a browser-based solution
 - a limited-feature Cisco IP Phone-based solution
 - an extensible desktop interface using standard OpenSocial gadgets
 - server features available to applications using documented REST APIs
- **Cisco CTI OS Desktop Toolkit**—The CTI OS Desktop Toolkit provides a software toolkit for building custom desktops, desktop integrations into third-party applications, or server-to-server integrations to third-party applications. you can use CTI OS in a secure mode, but it reduces the agent capacity of each PG by 25%.



Note Only Non-Reference Designs that use the Avaya PG or the Parent/Child topology can use CTI OS desktops. Cisco Finesse is the required dekstop for all other contact center enterprise solutions.

These integrated solutions enable call control (Answer, Drop, Hold, Un-Hold, Blind or Warm Transfers, and Conferences), outbound calls, consultative calls, and delivery and manipulation of Call Context Data (CTI screen pop).

Figure 13: Variety of Agent Interfaces for Unified CCE



Agents who use a third-party CRM user interface that is connected through a CRM Connector can be supervised using a CTI OS Desktop Toolkit-based supervisor desktop.

Desktop Architecture

Desktop applications typically run on agent or supervisor desktops, Administration & Data servers, or administration clients. Services that support the desktop applications can run on the Unified CCE Peripheral Gateway (PG) server or on their own server. For each PG, there is one set of active desktop services.

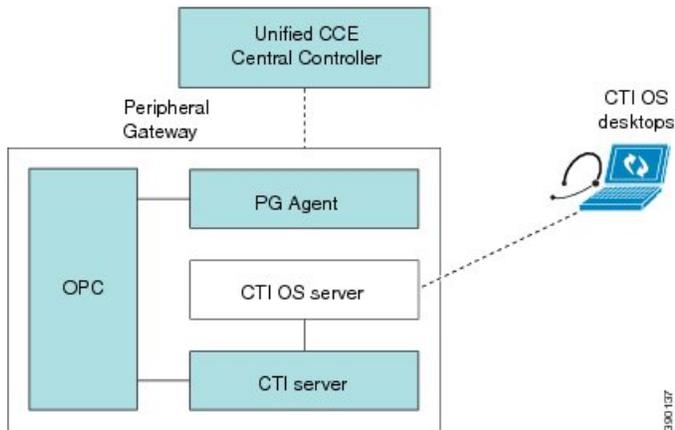
CTI Object Server

The CTI Object server (CTI OS) is a high-performance, scalable, fault-tolerant, server-based solution for deploying CTI applications. CTI OS is a required component for the CTI Toolkit Desktop. The CTI OS server runs as a redundant pair, one server on each VM that hosts an Agent PG.

Desktop applications pass communications, such as agent state change requests and call control, to the CTI OS server. CTI OS is a single point of integration for CTI Toolkit Desktops and third-party applications, such as Customer Relationship Management (CRM) systems, data mining, and workflow solutions.

The CTI Object server connects to the CTI server over TCP/IP and forwards call control and agent requests to the CTI server.

Figure 14: CTI OS Desktop Architecture



The CTI OS server also manages CTI Toolkit desktop configuration and behavior information, simplifying customization, updates, and maintenance, and supporting remote management.

CTI Object Server Services

- Desktop security: Supports secure socket connections between the CTI Object server on the PG and the agent, supervisor, or administrator desktop PC. Any CTI application built using the CTI OS Desktop Toolkit (CTI Toolkit) C++/COM CIL SDK can use the desktop security feature.



Note Desktop Security is not currently available in the .NET and Java CILs.

- Quality of Service (QoS): Supports packet prioritization with the network for desktop call control messages.



Note QoS is not currently available in the .NET and Java CILs.

- Failover recovery: Supports automatic agent sign-in upon failover.
- Chat: Supports message passing and the text chat feature between agents and supervisors.
- Silent Monitoring: Supports VoIP monitoring of active calls. The CTI Object server communicates with the Silent Monitor Service (SMS) to start or stop the VoIP packet stream forwarding.

You deploy the CTI Object server in redundant pairs, one on Agent PG A and one on Agent PG B . Both CTI OS servers are active simultaneously. The CTI Toolkit desktop applications randomly connect to one of the two servers. If the connection to the original server fails, the desktops automatically fail over to the alternate server.



Note The CTI OS server interfaces to any desktop application built using the CTI Toolkit SDK.

Agent Desktops

A Unified CCE deployment requires an agent desktop application. The agent uses this desktop for agent state control and call control. In addition to these required features, the desktop can provide other useful features.

Cisco offers the following primary types of Unified CCE agent desktop applications:

- Cisco Finesse: A browser-based agent desktop solution that provides a gadget-based architecture for extending base agent functionality.
- CTI Toolkit Desktop: An agent desktop application built with the CTI Toolkit. The desktop supports full customization and integration with other applications, customer databases, and Customer Relationship Management (CRM) applications.



Note The CTI Toolkit Desktop was deprecated in Unified CCE Release 11.0(1). Do not include the CTI Toolkit Desktop in new deployments. Support for the CTI Toolkit Desktop will be removed in a future release.

- Cisco Unified CRM Connector for Siebel: A CTI driver for the Siebel Communication Server.

Cisco partners offer the following types of agent desktop applications:

- Partner agent desktops: Custom agent desktop applications are available through Cisco Technology Partners. These applications are based on the CTI Toolkit and are not discussed individually in this document. The Finesse REST API also enables partner desktop integration.
- Prepackaged CRM integrations: CRM integrations are available through Cisco Unified CRM Technology Partners. These integrations are based on the CTI Toolkit and are not discussed individually in this document.

Supervisor Desktops

In addition to the agent desktop application, Cisco offers a supervisor desktop application. The contact center supervisor uses this application to monitor the agent state for members of their team. The supervisor desktop also allows the silent monitoring of agents during calls.

Cisco offers the following types of Unified CCE supervisor desktop applications:

- Cisco Finesse: A fully browser-based supervisor application that extends the base Finesse agent desktop with supervisor capabilities.
- CTI Toolkit supervisor desktop: A supervisor desktop application built with the CTI Toolkit. The desktop supports customization and integration with other applications, customer databases, and CRM applications.
- Supervisor desktop applications offered through Cisco partners.
- Prepackaged CRM integrations: CRM integrations are available through Cisco Unified CRM Technology Partners. These integrations are based on the CTI Toolkit and are not discussed individually in this document.



Note CAD also provides a supervisor desktop. For more information, see the CAD documentation.

Cisco Unified CRM Connector for Siebel Solution

The Cisco Unified CRM Connector for Siebel is a component that you install to enable integration of Unified CCE with the Siebel CRM environment. In this solution, the Siebel agent desktop provides the agent state and call control interface. The connector is built on top of the CTI OS Desktop Toolkit C++ CIL. The Siebel desktop uses the connector to communicate with the CTI OS server.

The Cisco Unified CRM Connector for Siebel does not support the following:

- Agent Greeting Enable/Disable
- Recording new agent greetings

For more information about the Cisco Unified CRM Connector for Siebel, see product documentation at http://www.cisco.com/en/US/products/ps9117/tsd_products_support_series_home.html. For more information about the Siebel eBusiness solution, see the [Siebel website](#).

Cisco Unified IP IVR

The Unified IP IVR provides prompting, collecting, and queuing capability for the Unified CCE solution. Unified IP IVR does not provide call control as Unified CVP does because it is behind Unified Communications Manager and under the control of the Unified CCE software by way of the Service Control Interface (SCI). When an agent becomes available, the Unified CCE software instructs the Unified IP IVR to transfer the call to the selected agent phone. The Unified IP IVR then requests Unified Communications Manager to transfer the call to the selected agent phone.

Unified IP IVR is a software application that runs on Cisco Unified Computing System (UCS) hardware or a specification-based equivalent. You can deploy multiple Unified IP IVR servers with a single Unified Communications Manager cluster under control of Unified CCE.

Unified IP IVR has no physical telephony trunks or interfaces like a traditional VRU. The telephony trunks are terminated at the Voice Gateway. Unified Communications Manager provides the call processing and switching to set up a G.711 or G.729 Real-Time Transport Protocol (RTP) stream from the Voice Gateway to the Unified IP IVR. The Unified IP IVR communicates with Unified Communications Manager through the Java Telephony Application Programming Interface (JTAPI), and the Unified IP IVR communicates with Unified CCE through the Service Control Interface (SCI) with a VRU Peripheral Gateway or System Peripheral Gateway.

For deployments requiring complete fault tolerance, a minimum of two Unified IP IVRs are required.

In Unified CCE deployments with only Unified IP IVR, a Unified CM cluster can support about 2000 Unified CCE agents. These limits assume that the BHCA call load and all configured devices are spread equally among the eight call processing subscribers with 1:1 redundancy. These capacities can vary, depending on your specific deployment. Size your solution with the *Cisco Unified Communications Manager Capacity Tool*.

A subscriber can support a maximum of 500 agents. In a fail-over scenario, the primary subscriber supports a maximum of 1000 agents.

For deployments requiring large numbers of VRU ports, use Unified CVP instead of Unified IP IVR. Unified IP IVR ports place a significant call processing burden on Unified CM, while Unified CVP does not. Thus, Unified CCE deployments with Unified CVP allow more agents and higher BHCA rates per cluster.

JTAPI Communications for Unified IP IVR

The Unified IP IVR sign-in process establishes JTAPI communications between the Unified CM cluster and the application. The CTI Manager communicates through JTAPI to Unified IP IVR. Every subscriber within a cluster runs a CTI Manager instance. But, the Unified CM PIM on the PG communicates with only one CTI Manager (and thus one node) in the cluster. That connected CTI Manager passes CTI messages for the other nodes within the cluster. Each redundant pair of PGs shares a unique JTAPI user ID.

Each Unified IP IVR communicates with only one CTI Manager within the cluster. The user ID is how the CTI Manager tracks the different processes.

Unified IP IVR is not deployed in redundant pairs. But, if its primary CTI Manager is out of service, Unified IP IVR can fail over to another CTI Manager within the cluster.

Unified IP IVR uses the same types of JTAPI messages as an Agent PG. Unlike the Agent PG, Unified IP IVR provides both the application and the devices that are monitored and controlled.

The devices that Unified CCE monitors and controls are the physical phones. The Unified IP IVR does not have physical ports like a traditional VRU. Unified IP IVR ports are logical ports called CTI Ports. For each CTI Port on Unified IP IVR, there needs to be a CTI Port device defined in Unified Communications Manager.

Unlike a traditional PBX or telephony switch, Unified Communications Manager does not select the Unified IP IVR port to which it sends the call. When a call is made to a DN that is associated through a CTI Route Point with a Unified IP IVR JTAPI user, the subscriber asks the Unified IP IVR which CTI Port handles the call. If Unified IP IVR has an available CTI Port, Unified IP IVR responds to the routing control request with the device identifier of the CTI Port to handle that call.

SIP sends Dual Tone Multi-Frequency (DTMF) digits, however Unified IP IVR and Unified Communications Manager only support out-of-band DTMF digits. JTAPI messages from the cluster notify Unified IP IVR of caller-entered DTMF digits. The cluster uses an MTP resource to convert in-band signaling to out-of-band signaling. CTI ports only support out-of-band DTMF digits. If your deployment includes SIP phones or gateways, provision sufficient MTP resources to support the conversion. The Mobile Agent feature also requires extra MTP resources for this conversion.

The following scenarios are examples of Unified IP IVR device and call control. When an available CTI Port is allocated to the call, a Unified IP IVR workflow starts within Unified IP IVR. When the workflow executes the accept step, a JTAPI message is sent to the subscriber to answer the call for that CTI Port. When the Unified IP IVR workflow wants the call transferred or released, the workflow again instructs the subscriber on what to do with that call.

When a caller releases the call while interacting with the Unified IP IVR, the VG detects the caller release. The VG notifies the subscriber with the Media Gateway Control Protocol (MGCP), which then notifies the Unified IP IVR with JTAPI. When the VG detects DTMF tones, the VG notifies the subscriber through H.245 or MGCP, which then notifies the Unified IP IVR through JTAPI.

In order for the CTI Port device control and monitoring to occur, associate the CTI Port devices on Unified Communications Manager with the appropriate Unified IP IVR JTAPI user ID. If you have two 150-port Unified IP IVRs, you have 300 CTI ports. Associate half of the CTI ports with JTAPI user Unified IP IVR 1, and associate the other half of the CTI ports with JTAPI user Unified IP IVR 2.

While you can configure Unified Communications Manager to route calls to Unified IP IVRs on its own, Unified CCE routes calls to the Unified IP IVRs in a Unified CCE environment (even if you have only one Unified IP IVR and all calls require an initial VRU treatment). Doing so ensures proper Unified CCE reporting. For deployments with multiple Unified IP IVRs, this routing practice also allows Unified CCE to load-balance calls across the multiple Unified IP IVRs.

CTI Manager Failover with Unified IP IVR

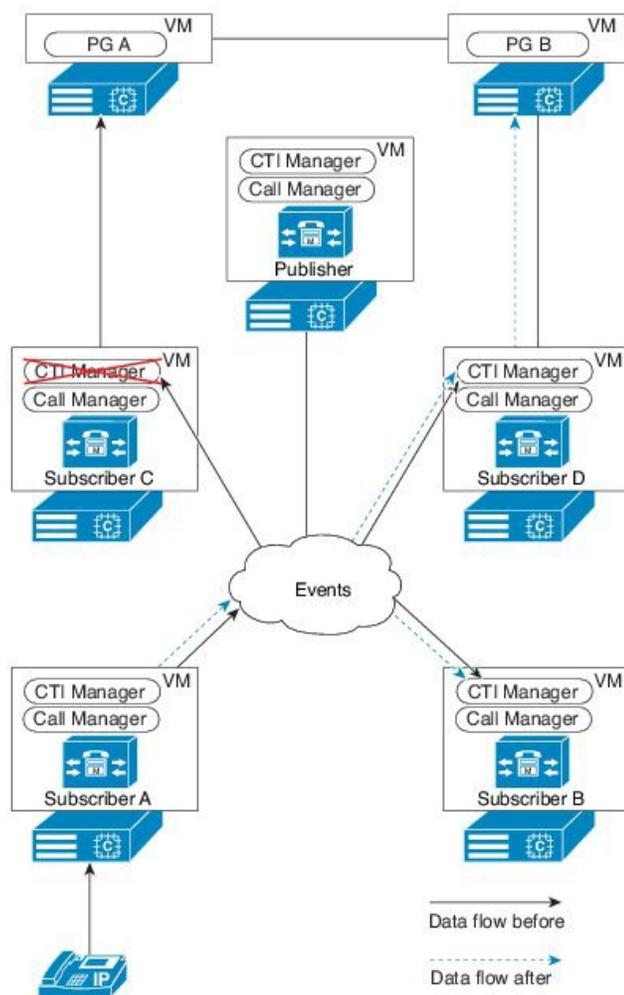
Each Agent PG can support only one CTI Manager connection. While each subscriber has a CTI Manager, only two subscribers normally connect to the Agent PGs. You would have to add another pair of Agent PGs to enable all subscribers in a four-subscriber cluster to connect directly to an Agent PG.

The following figure shows the failure of a CTI Manager with a connection to the Agent PG. Only subscribers C and D are configured to connect to the Agent PGs.

The following conditions apply to this scenario:

- For redundancy, all phones and gateways that are registered with subscriber A use subscriber B as their backup server.
- The CTI Managers on subscribers C and D provide JTAPI services for the Agent PGs.

Figure 15: CTI Manager with Agent PG Connection Fails



Unlike Unified CVP, Unified IP IVR depends on the CTI Manager for call control. In Unified IP IVR deployments, failure of a CTI Manager with an Agent PG connection causes the Unified IP IVR JTAPI

subsystem to shut down. This shutdown causes the Unified IP IVR server to drop all voice calls that the server is processing.

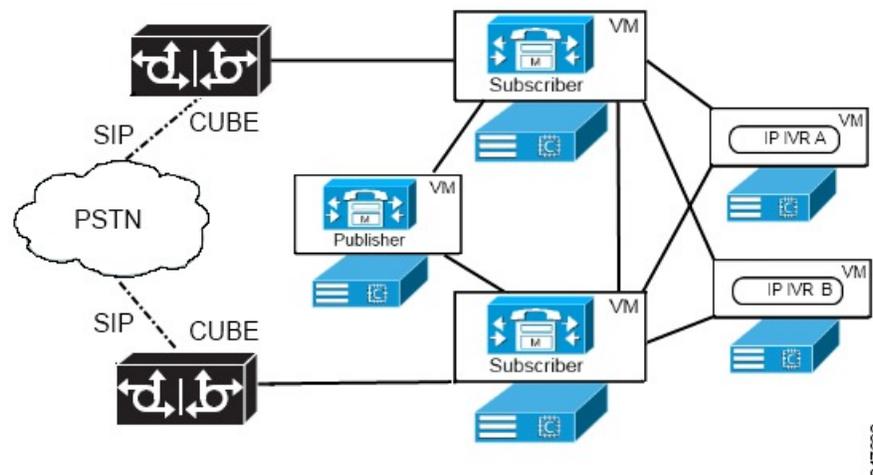
Then, the JTAPI subsystem restarts and connects to the CTI Manager on the backup subscriber. The Unified IP IVR reregisters all the CTI ports that are associated with the Unified IP IVR JTAPI user. After all the Unified CM devices are successfully registered, the server resumes its VRU functions and handles new calls.

Unified IP IVR Design Considerations

Cisco Unified IP IVR can establish JTAPI connections with two CTI Managers on different subscribers in the Unified Communications Manager cluster. This feature enables Unified IP IVR redundancy at the CTI Manager level. You can gain more redundancy by deploying multiple Unified IP IVR servers. Multiple Unified IP IVR servers enable call routing scripts to load balance calls between the available IP IVR resources.

The following figure shows two Unified IP IVR servers set up redundantly with a cluster. Set up the Unified IP IVR group with each server connected to the CTI Manager on a different Unified Communications Manager subscriber. Then, add a second CTI Manager as a backup to each Unified IP IVR server. If the primary CTI Manager fails, the Unified IP IVR server fails over to the backup CTI Manager.

Figure 16: IP IVR High Availability Deployment



347693

High Availability Through Call Forwarding

You can use the following call forwarding features in Unified Communications Manager to manage Unified IP IVR port usage:

- Forward Busy—Forwards calls to another port or route point when Unified Communications Manager detects that the port is busy.
- Forward No Answer—Forwards calls to another port or route point when Unified Communications Manager detects that a port did not pick up a call within the timeout period.
- Forward on Failure—Forwards calls to another port or route point when Unified Communications Manager detects a port failure caused by an application error.

When using the call forwarding features, do not establish a path back to the CTI port that initiated the call forwarding. Such paths create loops when all the Unified IP IVR servers are unavailable.

High Availability Through Call Flow Routing Scripts

You can use Unified CCE call flow routing scripts to support high availability. Check the Unified IP IVR Peripheral Status with a call flow routing script before sending calls to Unified IP IVR. This check prevents calls from queuing to an inactive Unified IP IVR. For example, create a Translation Route to the Voice Response Unit (VRU) to select the Unified IP IVR with the most idle ports. This method distributes the calls evenly on a call-by-call basis. You can modify this method to load balance ports across multiple Unified IP IVRs. This method can address all the Unified IP IVRs on the cluster in the same Translation Route or Send to VRU node.



Note If the Unified IP IVR server fails, all calls at the Unified IP IVR are dropped. Minimize the impact of such failures by distributing calls across multiple Unified IP IVR servers. In Unified IP IVR, a default script prevents loss of calls if the Unified IP IVR loses the link to the VRU PG.

Outbound Option with Unified IP IVR

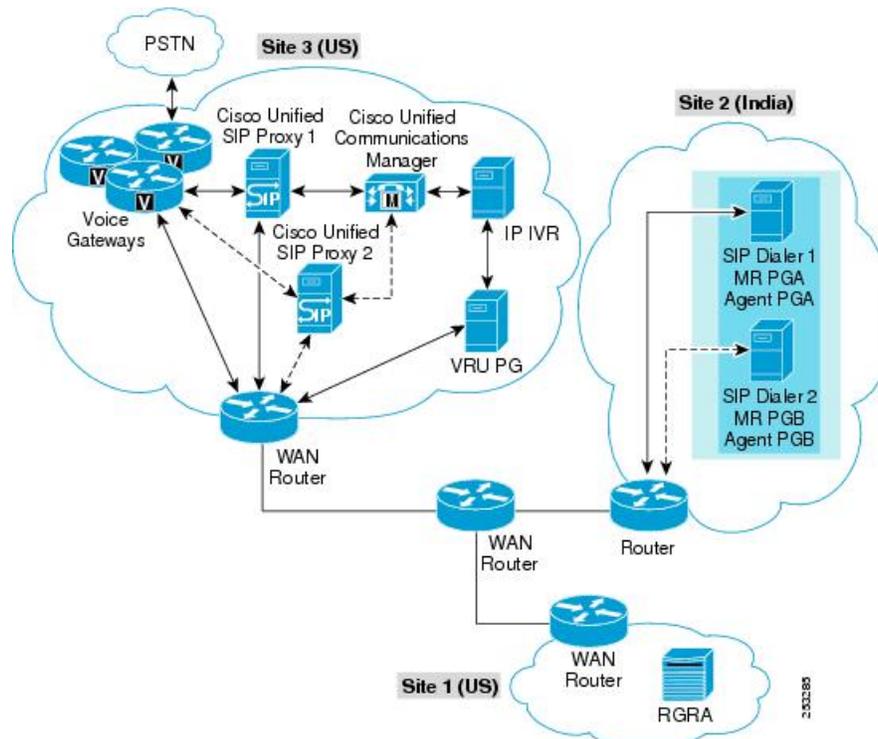
When you use Unified IP IVR with Outbound Option, front-end all calls through Unified CM. This deployment results in a higher call load on the subscribers. Because the subscriber supports only five calls per second, distribute calls that are transferred to agents and the VRU across multiple subscribers with a Unified SIP Proxy server.

For some deployments, if the outbound call center is in one country (for example, India) and the customers are in another country (for example, US), then consider the WAN network structure. When using Unified IP IVR deployments for transfer to a VRU campaign, and the Unified IP IVR is separated from the Outbound Option Dialer servers by a WAN. Provide Unified IP IVR with its own Unified CM cluster to reduce the WAN traffic.

Distributed Deployment of Outbound Option for Transfer to VRU Campaign

This figure shows a distributed deployment for a transfer-to-VRU campaign with Unified IP IVR.

Figure 17: Distributed Deployment Example for Transfer-to-VRU Campaign with Unified IP IVR



In the solution in this figure, note these points:

- The Voice Gateway and Router/Logger A servers are distributed between two sites (Site 1 and 3) in the United States.
- The Unified CM cluster is located at Site 3 (United States) along with the VRU PG.
- The redundant VRU PGs are at Site 3 (United States).
- IP IVR is included at Site 3 (United States).
- The redundant MRPG/Dialer and redundant Agent PGs are installed on the same VM at Site 2.
- The SIP Dialer uses the Voice Gateways located at Site 3 (United States).
- The Voice Gateways are included in the diagram with CT3 interface at Site 3 (United States). These routers will provide 1:1 redundancy for Dialer calls.
- The redundant Unified SIP Proxy servers are at Site 2 to avoid the WAN SIP signaling traffic to transfer live outbound calls.
- Each SIP Dialer connects to its own Unified SIP Proxy server at Site 2. Each Unified SIP Proxy server controls the set of Voice Gateways at Site 3 (United States).
- The Unified SIP Proxy servers provide (N+1) redundancy.

If you enable recording at the SIP Dialer, the WAN bandwidth requirements are:

- Bandwidth for each answered outbound call:
 - G.711 Codec calls require a WAN bandwidth of 80 kbps for the Call Progress Analysis time period.

- G.729 Codec calls require a WAN bandwidth of 26 kbps for the Call Progress Analysis time period.
- Bandwidth for each alerting outbound calls:
 - G.711 Codec calls require a WAN bandwidth of 80 kbps
 - G.729 Codec calls require a WAN bandwidth of 26 kbps
- Outbound calls being queued for self-serviced at Unified IP IVR do not require WAN bandwidth.

Unified Intelligent Contact Manager

Unified CCE may be deployed with Unified ICM to form a complete enterprise call management system. Unified ICM interfaces with ACDs from various vendors (including Cisco Unified CCE), VRUs (including Cisco Unified IP IVR and Unified CVP), and telephony network systems to efficiently and effectively treat customer contacts such as calls and e-mail regardless of where the resources are located in the enterprise.

Unified CCE may be deployed in a *hybrid* model with Unified ICM where the CallRouter, Logger, Administrative & Data Servers, and other components are shared between Unified ICM and Unified CCE.

Alternatively, Unified CCE may be deployed in a parent/child model where Unified ICM is the parent and Unified CCE is the child. This closely resembles the deployment model of Unified ICM with ACDs from other vendors. It is used for a highly scalable deployment because it provides CallRouters, data servers, and so forth for each product; although there are more components to manage and maintain. It is also used for a distributed model where isolation is needed between Unified ICM and Unified CCE, such as in an outsourced operation.

Unified ICM Configuration

- With Cisco Unified ICM7.0, to perform subsequent call control through Unified CVP, always use translation route to route the call to Unified CVP as a Type 2 NetworkVRU before delivering the call to its next destination. This practice passes the control to Unified CVP as an charge of subsequent call transfers because UnifiedCM cannot receive any further labels.
- To perform any queuing treatment, prompt and collect, or self-service applications, always follow translation route with a SendToVRU node. SendToVRU can sometimes be invoked implicitly by a Queue node or a RunExternalScript node, but you should not rely on that method. Always include an actual SendToVRU node.
- With Cisco Unified ICM7.1, to perform subsequent call control through Unified CVP, a translation route is not necessary if you use a Type 10 NetworkVRU. The Type 10 VRU uses the Correlation ID method to perform a transfer from UnifiedCM to Unified CVP using a SendToVRU node. When the SendToVRU node is used with a Type 10 VRU, an initial transfer to Unified CVP hands off call control to Unified CVP, and then an automatic second transfer to the VRU leg is performed to deliver the call to a Voice Browser for VRU treatment.



Note This call flow and all others in this document assume that you are using Cisco Unified ICM7.0(0) or later.

- When the SendToVRU node is used with a Type 10 VRU, an initial transfer to Unified CVP hands off call control to Unified CVP, and then an automatic second transfer to the VRU leg is performed to deliver the call to a Voice Browser for VRU treatment.

Generic Peripheral Gateway

Generic PG—Combines an Agent PG and a VRU PG

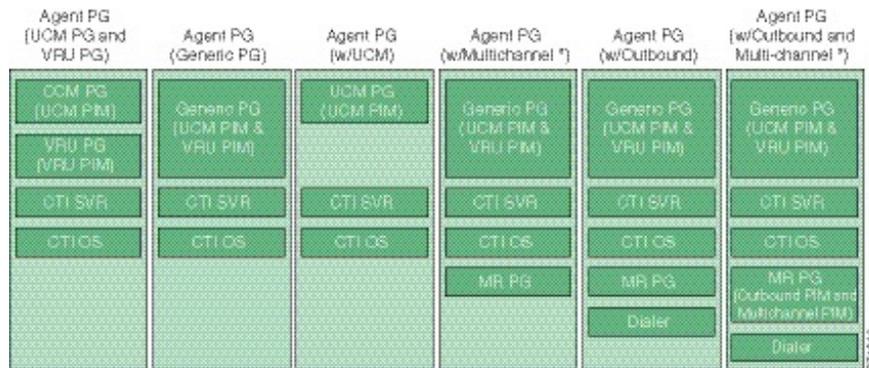
The last class of PG is the Generic PG. This PG can combine PIMs of different types. So, you can deploy the Agent PG and VRU PG independently. You can also combine those functions in a single Generic PG.

Peripheral Gateway and Server Options

A Unified CCE Peripheral Gateway (PG) translates messages coming from the Unified Communications Manager servers, the Unified IP IVR, Unified CVP, or voice response units (VRUs) into common internally formatted messages that are then sent to and understood by Unified CCE. In the reverse, it also translates Unified CCE messages so that they can be sent to and understood by the peripheral devices.

The figures below illustrate various configuration options for the Agent PG with CTI OS.

Figure 18: Agent PG Configuration Options with CTI OS



* with CCE, the MR PG for outbound and multi-channel is co-resident with the Agent PG. With System CCE, the MR PG for outbound (outbound controller) is co-resident with the Agent PG, but not the MR PG for multi-channel (multi-channel controller).

The table below gives sizing guidelines for PGs and PIMs.

Table 5: PG and PIM Sizing Guidelines

Sizing variable	Guidelines based on Unified CCE Release 10
Maximum number of PGs per Unified CCE	150
Maximum number of PG types per VM	Up to two PG types are permitted per VM, but each VM must meet the maximum agent and VRU port limitations.
Maximum number of Unified Communications Manager PGs per VM	Only one Unified Communications Manager PG, Generic PG, or System PG is allowed per VM.

Sizing variable	Guidelines based on Unified CCE Release 10
Maximum number of Unified Communications Manager PIMs per PG	1
Can PGs be remote from Unified CCE?	Yes
Can PGs be remote from Unified Communications Manager?	No
Maximum number of VRUs controlled by one Unified Communications Manager	See the <i>Cisco Collaboration System Solution Reference Network Designs</i> at http://www.cisco.com/go/ucsrnd .
Maximum number of CTI servers per PG	1
Can PG be co-resident with Cisco Unified Communications Manager?	No