



Cisco Unified Contact Center

Revised: March 1, 2018

This chapter describes the Cisco Unified Contact Center solutions available with the Cisco Unified Communications System. It includes information on Cisco products such as Cisco Unified Contact Center Express, Cisco Unified Contact Center Enterprise, and Cisco Unified Customer Voice Portal. It also covers the design considerations for deploying these Cisco Unified Contact Center products with Cisco Unified Communications Manager and other Unified Communications components.

This chapter covers the following topics:

- [Cisco Contact Center Architecture, page 22-2](#)
- [Contact Center Deployment Models, page 22-12](#)
- [Design Considerations for Contact Center Deployments, page 22-17](#)
- [Capacity Planning for Contact Centers, page 22-21](#)
- [Video Customer Care, page 22-22](#)
- [Network Management Tools, page 22-23](#)

This chapter starts with a high-level overview of the main Cisco Unified Contact Center Portfolio. Then it covers the various Unified Communications deployment models for contact centers. Finally, it discusses design considerations on topics such as bandwidth, latency, Cisco Unified Communications Manager integration, and sizing.

The intent of this chapter is not to provide details on each contact center product and their various components but rather to discuss the design considerations for their integration with the Cisco Unified Communications System. Detailed design guidance for each Unified Contact Center product is covered in specific design guides for the Cisco Unified Contact Center Express, Cisco Unified Contact Center Enterprise, and Cisco Unified Customer Voice Portal products. Links to the product-specific design guides are listed at

<https://www.cisco.com/go/srnd>

What's New in This Chapter

Table 22-1 lists the topics that are new in this chapter or that have changed significantly from previous releases of this document.

Table 22-1 *New or Changed Information Since the Previous Release of This Document*

New or Revised Topic	Described in:	Revision Date
<p>The following products have reached end of sale (EoS) and have been removed from this document:</p> <ul style="list-style-type: none"> • Cisco MediaSense • Cisco Connected Analytics for Contact Center • Cisco Prime Collaboration Contact Center Assurance 	<p>For information on these products, refer to previous versions of the SRND, available at https://www.cisco.com/go/srnd.</p>	<p>March 1, 2018</p>

Cisco Contact Center Architecture

This chapter discusses the following main Cisco Contact Center products and related features:

- Cisco Unified Communications Manager (Unified CM) call queuing feature
- Cisco Unified Contact Center Enterprise (Unified CCE)
- Cisco Unified Customer Voice Portal (Unified CVP)
- Cisco Unified Contact Center Express (Unified CCX)

Cisco Unified CM Call Queuing

The Cisco Unified CM call queuing feature provides the capability for queuing the incoming callers to a hunt pilot number. With this option enabled, callers to the hunt pilot can be put in queue to wait for an available agent that is configured as a hunt member to answer the call. Callers receive an initial greeting announcement when they first enter the queue, and they hear periodic announcements while they are in queue. When an agent becomes available, the call is taken out of the queue and answered by the agent. For customers who need a basic contact center with very limited functionality, Cisco Unified CM call queuing can be an option. However, unlike the full-featured Cisco Contact Center products, the Unified CM call queuing option lacks much of the contact center functionality such as agent desktop, supervisor, and reporting capabilities. If customers require complete contact center functionality, Cisco Unified Contact Center Enterprise or Cisco Unified Contact Center Express should be used.

The hunt pilot line members can display the queue status about their associated hunt pilots from the phone screen, and the queue status provides the following information:

- Hunt pilot number
- Number of calls waiting in the queue
- Longest call waiting time

In addition, Unified CM call queuing provides statistics on the number of calls currently waiting in queue and the longest call waiting time, along with other statistics, through the serviceability counters based on the hunt pilot number. This allows the supervisor to monitor the queue status using the Real Time Monitoring Tool (RTMT).

For each hunt pilot, callers can be routed to an alternate configurable destination such as voicemail or another hunt pilot if any of the following situations occurs:

- The number of calls in the queue reaches the maximum that is set by the **Maximum Number of Callers Allowed in Queue** parameter.
- The wait time of a caller in queue exceeds the threshold that is configured by the **Maximum Wait Time in Queue** parameter.
- No hunt members are logged in or registered.

**Note**

For calls routed to the queue-enabled hunt pilot number through a SIP trunk, the SIP Rel1XX Options should be set to **Send PRACK if 1XX contains SDP** in the SIP profile associated with the SIP trunk.

For additional information on the Unified CM call queuing option, refer to the latest version of the *System Configuration Guide for Cisco Unified Communications Manager*, available at

<https://www.cisco.com/c/en/us/support/unified-communications/unified-communications-manager-callmanager/products-installation-and-configuration-guides-list.html>

Cisco Unified Contact Center Enterprise

Cisco Unified Contact Center Enterprise (Unified CCE) provides a contact center solution that enables you to integrate inbound and outbound voice applications with Internet applications, including real-time chat, Web collaboration, and email. This integration provides for unified capabilities, helping a single agent support multiple interactions simultaneously, regardless of the communications channel the customer has chosen. Because each interaction is unique and may require individualized service, Cisco provides contact center solutions to manage each interaction based on virtually any contact attribute. The Unified CCE deployments are typically used for large size contact centers and can support thousands of agents.

There is also a predesigned, bounded deployment model of Unified CCE: Cisco Packaged Contact Center Enterprise (Packaged CCE). Customers whose contact center requirements fit the boundaries of the solution can enjoy the advantages of the simplified management interface, smaller hardware footprint, and reduced time to install. Those customers can also benefit from the comprehensive feature set of Cisco Unified Contact Center Enterprise and Cisco Unified Customer Voice Portal. The solution comes packaged with Cisco Unified Intelligence Center for comprehensive reporting and Cisco Finesse desktop software for an enhanced, next-generation desktop experience. For more details on Packaged CCE, refer to the documentation at the following locations:

- <https://www.cisco.com/en/US/products/ps12586/index.html>
- https://www.cisco.com/en/US/products/ps12586/tsd_products_support_series_home.html
- https://docwiki.cisco.com/wiki/Packaged_CCE

Unified CCE employs the following major software components:

- Call Router

The Call Router makes all the decisions on how to route a call or customer contact.

- Logger

The Logger maintains the system database that stores contact center configurations and temporarily stores historical reporting data for distribution to the data servers. The combination of Call Router and Logger is called the *Central Controller*.

- Peripheral Gateway

The Peripheral Gateway (PG) interfaces to various "peripheral" devices, such as Cisco Unified CM, Cisco Unified IP Interactive Voice Response (Unified IP IVR), Cisco Unified CVP, or multichannel products such as Cisco Unified Web Interaction Manager (Unified WIM) and Cisco Unified E-Mail Interaction Manager (Unified EIM). A Peripheral Gateway that interfaces with Unified CM is also referred to as an *Agent PG*.

- CTI Server and CTI Object Server (CTI OS)

The CTI Server and CTI Object Server interface with the agent desktops. Agent desktops can be based on the Cisco Finesse agent and supervisor desktops, Finesse IP Phone Agent, or customer relationship management (CRM) connectors to third-party CRM applications.

- Administration & Data Server

The Administration & Data Server provides a configuration interface as well as real-time and historical data storage.

The Cisco Unified CCE solution is based on the integration with Cisco Unified Communications Manager (Unified CM), which controls the agent phones. For deployments without Unified CM but with traditional ACD, use Cisco Unified Intelligent Contact Management Enterprise (Unified ICME) instead of Unified CCE.

The queuing and self-service functions are provided by Cisco Unified IP Interactive Voice Response (Unified IP IVR) or Cisco Unified Customer Voice Portal (Unified CVP) and are controlled by the Unified CCE Call Router.

Most of the Unified CCE components are required to be redundant, and these redundant instances are referred to as side A and side B instances. For example, Call Router A and Call Router B are redundant instances of the Call Router component running on two different virtual machines.

Agents can use a large variety of endpoints, including some video endpoints and some Cisco TelePresence endpoints such as the Cisco DX70 and DX80. For a list of supported endpoints, refer to the latest version of the *Unified CCE Solution Compatibility Matrix*, available at

<https://www.cisco.com/c/en/us/support/customer-collaboration/unified-contact-center-enterprise/products-device-support-tables-list.html>

Cisco Unified Customer Voice Portal

Cisco Unified Customer Voice Portal (Unified CVP) provides carrier-class voice and video IVR services on Voice over IP (VoIP) networks. It can perform basic prompt-and-collect or advanced self-service applications with CRM database integration and with automated speech recognition (ASR) and text-to-speech (TTS) integration. Unified CVP also provides IP-based call switching services by routing and transferring calls between voice gateways and IP endpoints.

Unified CVP is based on the Voice Extension Markup Language (VXML), which is an industry standard markup language similar to HTML and which is used to develop IVR services that leverage the power of web development and content delivery.

The Unified CVP solution employs the following main components:

- Unified CVP Call Server

The Unified CVP Call Server provides call control capabilities for SIP services. The Unified CVP Call Server can also integrate with the Unified CCE Call Router through the Intelligent Contact Management (ICM) service. The IVR service provides a platform to run VXML Micro applications and to create VoiceXML pages.

- Unified CVP VXML Server

This component executes complex IVR applications by exchanging VoiceXML pages with the VoiceXML gateway's built-in voice browser. Unified CVP VXML applications are written using Cisco Unified Call Studio and are deployed to the Unified CVP VXML Server for execution. Note that there is no RTP traffic going through the Unified CVP Call Server or the Unified CVP VXML Server.

- Cisco Voice Gateway

The Cisco Voice Gateway is the point at which a call enters or exits the Unified CVP system. The Cisco Voice Gateway could have a TDM interface to the PSTN. Alternatively, Cisco Unified Border Element could be used when the interface to the PSTN is an IP voice trunk.

- Cisco VoiceXML Gateway

The VoiceXML Gateway hosts the Cisco IOS Voice Browser. This component interprets VoiceXML pages from either the Unified CVP Server IVR Service or the Unified CVP VXML Server. The VoiceXML Gateway can play prompts based on .wav files to the caller and can accept input from the caller through DTMF input or speech (when integrated with Automatic Speech Recognition). It then returns the results to the controlling application and waits for further instructions.

The Cisco VoiceXML Gateway can be deployed on the same router as the Cisco Voice Gateway. This model is typically desirable in deployments with small branch offices. But the VoiceXML Gateway can also run on a separate router platform, and this model might be desirable in large or centralized deployments with multiple voice gateways.

- Video Media Server

A video media server in a Unified CVP comprehensive deployment enables video streaming for the Video in Queue feature. Cisco TelePresence Content Server can be used as a video media server.

Unified CVP can be deployed standalone or integrated with Unified CCE to offer voice and video self-service and queuing functions. The Unified CVP solution now supports the G.711 a-law codec end-to-end.

The Basic Video Service in Unified CVP is available when Unified CVP is deployed along with Cisco Contact Center Enterprise (Unified CCE) in a comprehensive deployment model. This service allows a video caller to interact with an audio-only IVR and subsequently connect with a video agent. It supports Cisco TelePresence endpoints such as the Cisco DX70 and DX80 as customer and agent endpoints. The video agents can also conference in a second audio-only agent by dialing a direct extension from their phone.

Video in Queue (VIQ) Basic Video is an optional feature in Unified CVP, and it can be enabled to play video to callers while they wait for a video-enabled agent or expert. Cisco TelePresence Content Server enables the video streaming. The caller can subsequently connect to a video agent.

For more information on Unified CVP system design and detailed call flows, refer to the latest version of the *Cisco Unified Customer Voice Portal Design Guide*, available at

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

Cisco Unified Contact Center Express

Cisco Unified Contact Center Express (Unified CCX) meets the needs of departmental, enterprise branch, or small to medium-sized companies that need easy-to-deploy, easy-to-use, highly available and sophisticated customer interaction management for up to 400 agents. It is designed to enhance the efficiency, availability, and security of customer contact interaction management by supporting a highly available virtual contact center with integrated self-service applications across multiple sites. To simplify the deployment, Unified CCX can be preloaded on a Cisco Business Edition 6000 or 7000 system.

Unified CCX integrates with Unified CM by means of JTAPI for call control. All the Unified CCX components, including the Unified CCX engine, Unified CCX database, Finesse Server, Unified CCX Outbound Dialer, Cisco Unified Intelligence Center, and Express E-mail Manager, are installed on a single virtual machine. For system redundancy, a second identical Unified CCX instances can be added to the deployment.

Unified CCX has built-in capabilities for inbound audio and video calls, silent monitoring, and Cisco Unified Intelligence Center reporting. Additional licensing and components can enhance the solution to support outbound dialing, call recording, email, chat, social network monitoring, and workforce optimization.

Unified CCX supports advanced features such as Automated Speech Recognition (ASR) and Text to Speech (TTS), HTTP, and VXML. It also supports products such as Cisco Unified Workforce Optimization to optimize performance and quality of the contact center. Agents can use a variety of video endpoints such as the Cisco Unified IP Phone 9900 Series with camera. For a list of supported endpoints, refer to the latest version of the *Unified CCX Software Compatibility Matrix*, available at

<https://www.cisco.com/c/en/us/support/customer-collaboration/unified-contact-center-express/products-device-support-tables-list.html>

Cisco Unified CCX includes IP IVR functionality for prompting, collecting, and queuing during customer inbound calls.

Cisco SocialMiner

Cisco SocialMiner is a social media customer-care solution that can help you proactively respond to customers and prospects by communicating through public social media networks such as Twitter, Facebook, or other public forums or blogging sites. By providing social media monitoring, queuing, and workflow to organize customer posts on social media networks and deliver them to your social media customer care team, your company can respond to customers in real time using the same social network the customers are using. For more information, refer to the documentation available at

<https://www.cisco.com/en/US/products/ps11349/index.html>

Universal Queue for Third-Party Multichannel Applications

Universal Queue describes the system's ability to route requests from various media channels to any agents in a contact center.

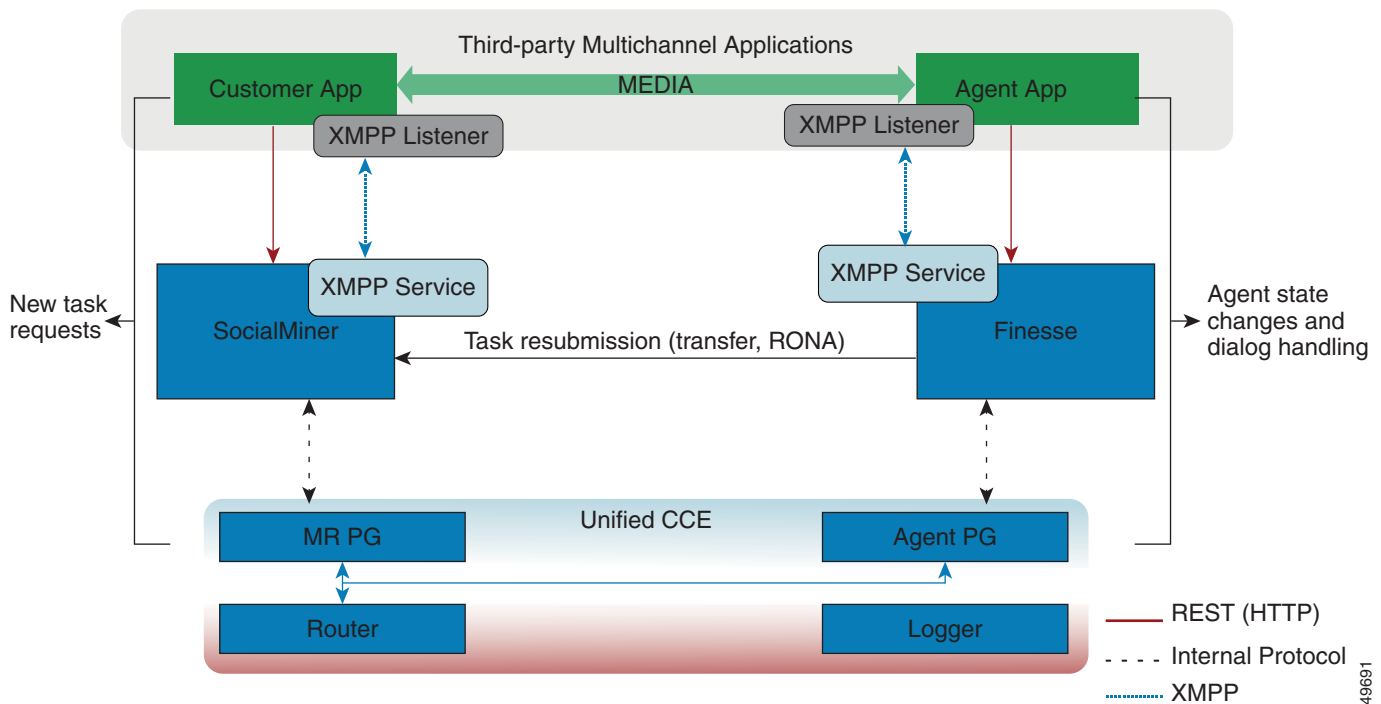
You can configure agents to handle a combination of voice calls, emails, chats, and so on. For example, you can configure an agent as a member of skill groups or precision queues in three different Media Routing Domains (MRDs) if the agent handles voice, email, and chat. You can design routing scripts to send requests to these agents based on business rules, regardless of the media. Agents signed into multiple MRDs may switch media on a task-by-task basis.

Universal Queuing APIs provide a standard way to request, queue, route, and handle third-party multichannel tasks in Unified CCE. (See [Figure 22-1](#).)

Contact Center customers or partners can develop applications using Cisco SocialMiner and Cisco Finesse APIs in order to use Universal Queue. The SocialMiner Task API enables applications to submit non-voice task requests to Unified CCE. The Finesse APIs enable agents to sign into various types of media and handle the tasks. Agents sign into and manage their state in each media independently.

Cisco partners can use the sample code available on Cisco DevNet (<https://developer.cisco.com/site/devnet/home/index.jsp>) as a guide for building these applications.

Figure 22-1 Universal Queue for Third-party Multichannel Applications Solution Components



SocialMiner and Universal Queue

Third-party multichannel applications use SocialMiner's Task API to submit non-voice tasks to Unified CCE. The API works in conjunction with SocialMiner task feeds, campaigns, and notifications to pass task requests to the contact center for routing.

The Task API supports the use of Call variables and Extended Call Context (ECC) variables for task requests. Use these variables to send customer-specific information with the request, including attributes of the media such as the chat room URL or the email handle.

Unified CCE and Universal Queue

Cisco Unified CCE provides the following functionality as part of Universal Queue:

- Processes the task request
- Provides estimated wait time for the task request
- Notifies SocialMiner when an agent has been selected
- Routes the task request to an agent, using either skill group or precision queue routing
- Reports on contact center activity across media

Finesse and Universal Queue

Cisco Finesse provides Universal Queue functionality via the Media API and Dialog API. With the Media API, agents using third-party multichannel applications can:

- Sign into multiple MRDs
- Change state in multiple MRDs

With the Dialog API, agents using third-party multichannel applications can handle tasks from multiple MRDs.

Administration and Management

Cisco Contact Center products have built-in administration and management capabilities. For example, Unified CCE can be administered using either the Configuration Manager tool that is installed with Unified CCE or the web-based administration tools for simplified execution of the most common administration and management tasks in the Contact Center Enterprise environment. In addition, REST API support allows third-party developers to create applications that can control many of the administration and support tasks.

Unified CVP can be administered with the Unified CVP Operations Console, also known as Operations, Administration, Maintenance, and Provisioning (OAMP).

In addition, Cisco Unified Contact Center Management Portal (Unified CCMP) can be deployed to simplify the operations and procedures for performing basic administrative functions such as managing agents and equipment. Unified CCMP is a browser-based management application designed for use by contact center system administrators, business users, and supervisors. It is a dense multi-tenant provisioning platform that overlays the Cisco Unified CCE, Unified CM, and Unified CVP equipment.

Reporting

Cisco Unified Intelligence Center is the main reporting tool for the Cisco Contact Center solutions. It is supported by Unified CCE, Unified CCX, and Unified CVP. This platform is a web-based application offering many Web 2.0 features, high scalability, performance, and advanced features such as the ability to integrate data from other Cisco Unified Communications products or third-party data sources.

Cisco Unified Intelligence Center gets source data from a database, such as an Unified CCE Administration & Data Server database or the Unified CVP Reporting Informix database. Reports are then generated and provided to a reporting client.

Multichannel Support

The Cisco Unified Enterprise solution supports web interaction and email interaction for multichannel support. Cisco Unified Web Interaction Manager (Unified WIM) technology helps ensure that communication can be established from nearly any web browser. Cisco Unified E-Mail Interaction Manager (Unified EIM) provides inbound email routing, automated or agent assisted email responses, real-time and historical reporting, and role-based hierarchical rights management for agents, supervisors, administrators, and knowledge base administrators.

For more design information on these products, refer to the *Cisco Unified E-Mail and Web Interaction Manager Solution Reference Network Design Guide*, available at

https://www.cisco.com/en/US/products/ps7236/products_implementation_design_guides_list.html

Recording and Silent Monitoring

Cisco Unified Contact Center solutions provide recording and silent monitoring capabilities based on the following mechanisms:

- The SPAN feature in Cisco switches
This feature replicates the network traffic to a destination port to which a Cisco contact center server is connected.
- Unified CM and media replication by the built-in-bridge (BIB) in Cisco IP Phones
With this option, Unified CM is involved in setting up the recording flows and can perform call admission control for those flows.
- Media forking by Cisco Unified Border Element gateway

For more details on call recording and monitoring, see the chapter on [Call Recording and Monitoring](#), page 23-1.

Contact Sharing

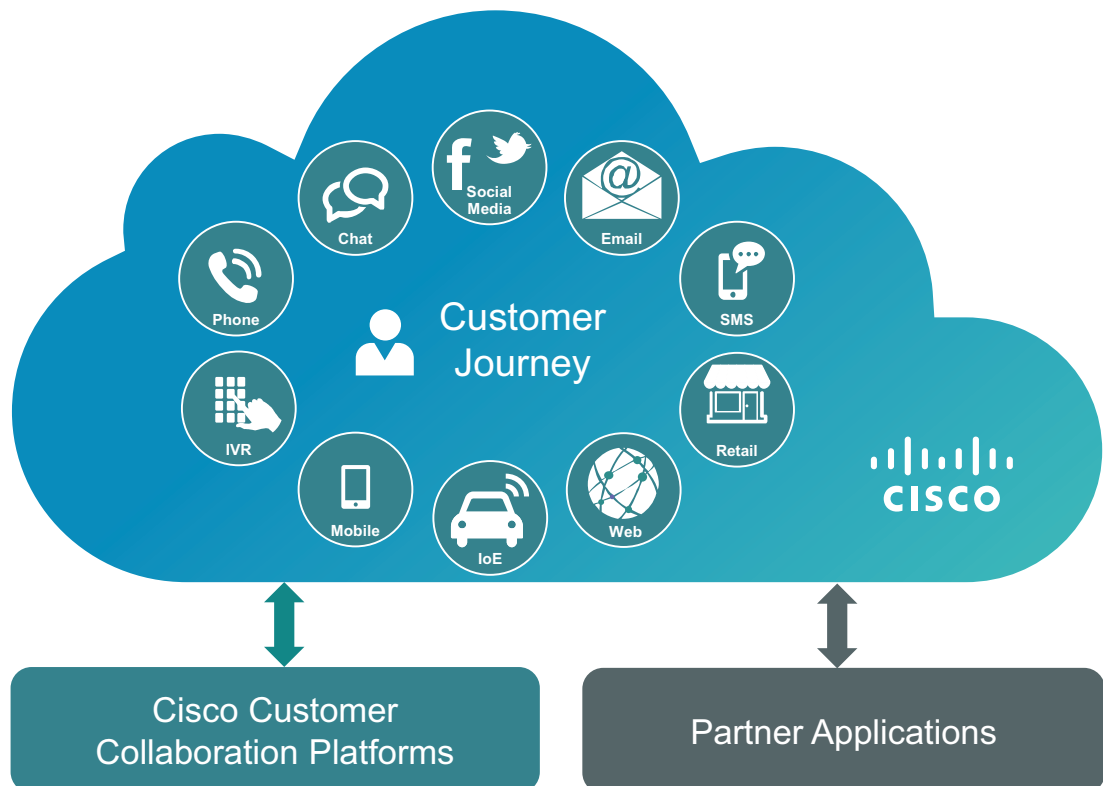
Contact Sharing enables large contact centers to grow larger. Centralized self-service (IVR ICM deployment model) uses a Contact Sharing routing node to distribute calls to two Unified CCE instances for horizontal scaling. Live Data is a prerequisite for Contact Sharing and must be installed and configured prior to use of Contact Sharing. Contact Sharing also requires the IVR Cisco Intelligent Contact Management (ICM) deployment model to be enabled in the deployment. For details about Contact Sharing, refer to the latest version of the *Cisco Unified Contact Center Enterprise Features Guide*, available at

<https://www.cisco.com/c/en/us/support/customer-collaboration/unified-contact-center-enterprise/products-feature-guides-list.html>

Context Service

Context Service is a cloud-based storage service that provides a repository for customer journey data. It enables Cisco Contact Center customers to deliver a seamless omnichannel experience through integration with other Cisco Customer Collaboration products as well as APIs for third-party integration, as depicted in [Figure 22-2](#).

Figure 22-2 Context Service Integration

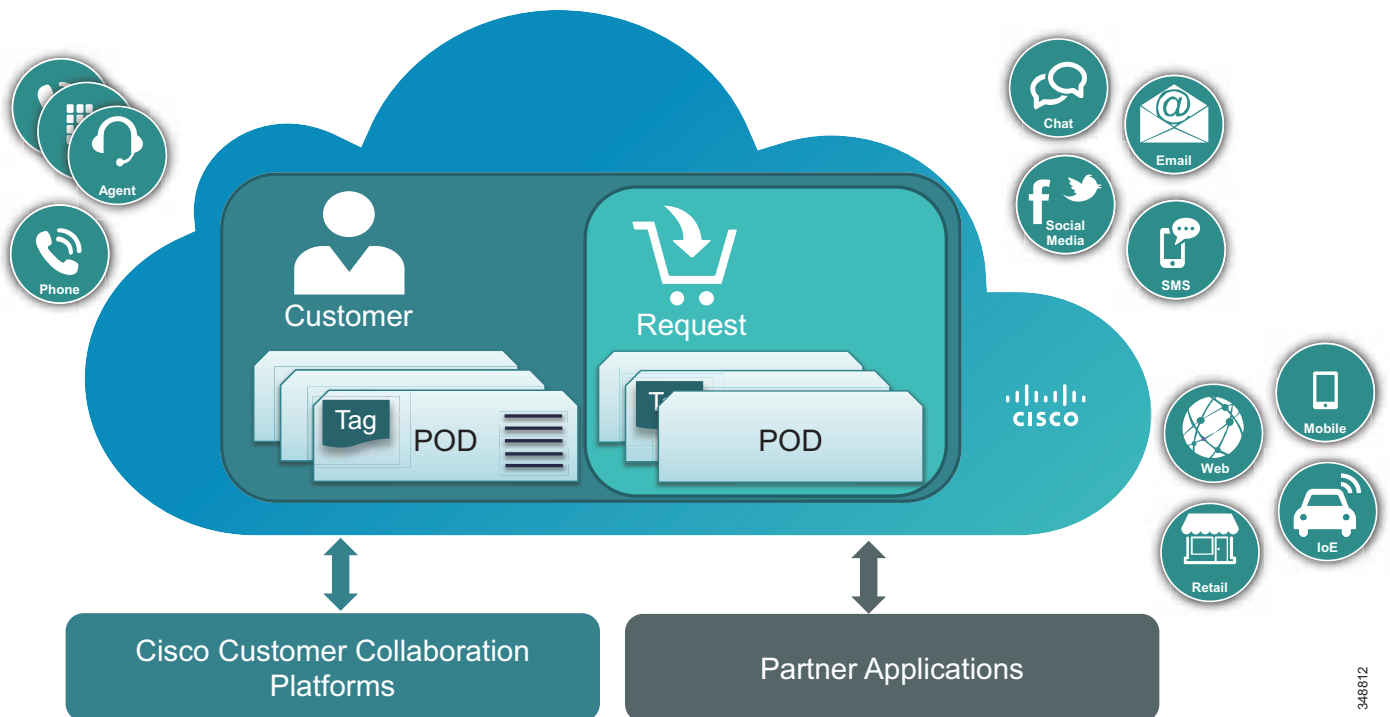


348811

Context Service allows any application to write and read customer journey activity. Cisco Contact Center customers, referred to in this section as the *business*, have access to Context Service from within their Cisco Contact Center platforms. The Cisco Contact Center platforms are enabled and optionally configured to post context data about contact center interactions.

Context Service stores this data in an element called a Piece of Data (POD). A POD can store any metadata about the consumer interaction, except for the media (such as audio recording). Businesses choose which fields (metadata) to store in the POD and the level of data privacy for each field. PODs can be organized by customer and also grouped together as part of a collection of interactions called a Request (see [Figure 22-3](#)). Context Service also provides tagging capability to group PODs for correlation, trending, and analytics.

Figure 22-3 Pieces of Data (PODs) and Requests



Context Service is hosted on Cisco Intercloud, which is an ecosystem of Cisco and partner data centers that is managed and operated by the Cisco data center team across the globe. Context Service follows a data privacy model very similar to Cisco Spark, in which each business controls access to its data. The data is encrypted/decrypted on-premises at the client and stored as an encrypted blob in Cisco data centers. Businesses can choose to host the encryption keys (Keystore) on their premises. This is analogous to valuables stored in a safety deposit box (locker) at a bank; even though the valuables are in the bank, the customer has the key to the deposit box and controls access to it. This is a newer approach to data privacy, and it puts the customer in control without the overhead of hosting a private cloud. Context Service provides a level of data privacy classification so that businesses can store their customers' Personally Identifiable Information (PII) separately from other encrypted data, and thereby allowing businesses to provide third-party analytics vendors with controlled access to their encrypted data without giving access to their customers' PII data.

348812

Context Service can store data for Universal Queue task contacts. When Context Service is enabled, SocialMiner selects pieces of data from an incoming task request and saves it in a POD in the cloud. You can specify the media type of the POD in the task request. If you don't specify the media type, then the media type **event** is set to the POD.

Context Service is managed by Cisco Collaboration Management (Atlas), which is the management portal for all Cisco cloud collaboration offerings, including Cisco Spark. Cisco partners and businesses use Collaboration Management to connect on-premises clients, manage the POD data model (fields), monitor POD usage, and so forth.

Context Service provides an open API and Java/JS SDK to make it easy for technology partners to integrate their applications with Context Service.

For more details on Cisco Context Service, refer to the Cisco Unified Contact Center design guides available at <https://www.cisco.com/go/srnd>.

Cisco Virtualized Voice Browser

Cisco Virtualized Voice Browser (Cisco VVB) provides a platform for interpreting VoiceXML documents. When a new call arrives at the contact center, the VVB allocates a VXML port that represents the VoIP endpoint. Cisco VVB sends HTTP requests to the Cisco Unified Customer Voice Portal (Unified CVP) VXML server. In response to the HTTP request, the Unified CVP VXML server executes the request and sends a dynamically generated VXML document. For more information about Cisco VVB, refer to the Cisco Virtualized Voice Browser design considerations and installation and configuration options documented in the latest version of the *Installing and Configuring Guide for Cisco HCS*, available at

<https://www.cisco.com/c/en/us/support/unified-communications/hosted-collaboration-solution-contact-center/products-installation-guides-list.html>

Contact Center Deployment Models

This section describes the various design models used for deploying Cisco Unified Contact Center solutions. For more details on these deployment models, refer to the Cisco Unified Contact Center design guides available at <https://www.cisco.com/go/srnd>.

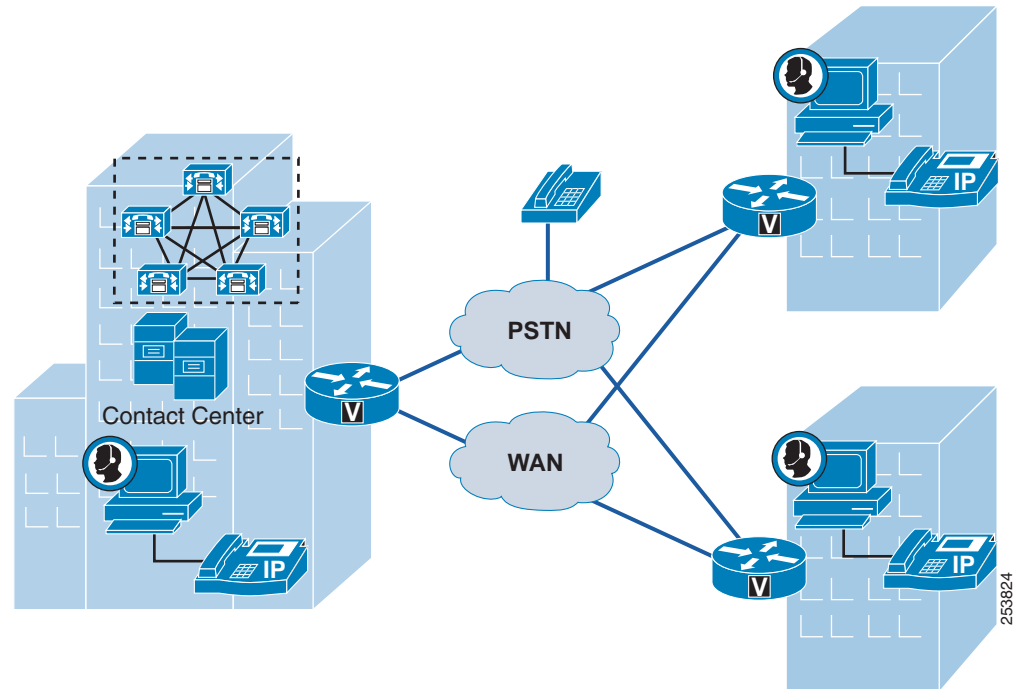
Single-Site Contact Center

In this deployment, all the components such as call processing agents, voice gateways, and contact center applications are in the same site. Agents and supervisors are also located at that site. The main benefit of the single-site deployment model is that there is no WAN connectivity required and, therefore, no need to use a low-bandwidth codec such as G.729, transcoders, compressed Real-Time Transport Protocol (cRTP), or call admission control.

Multisite Contact Center with Centralized Call Processing

A multisite deployment with centralized call processing consists of a single call processing cluster that provides services for many remote sites and uses the IP WAN. Cisco Contact Center applications (Unified CCE, Unified CCX, and Unified CVP) are also typically centralized to reduce the overall costs of management and administration. [Figure 22-4](#) illustrates this type of deployment.

Figure 22-4 Multisite Contact Center with Centralized Call Processing



Because the agents or the voice gateways in this type of deployment are located in remote sites, it is important to consider the bandwidth requirements between the sites. It is also important to carefully configure call admission control, Quality of Service (QoS), codecs, and so forth. For more information on the general design considerations for Unified Communications solutions, refer to the chapter on [Collaboration Deployment Models](#), page 10-1.

Contact center deployments in a Unified Communications system typically have the following additional bandwidth requirements:

- The traffic volume handled by the agents is higher than that of typical users, and therefore voice and signaling traffic is also higher for agents.
- Agents and supervisors use desktops with screen popup, reports and statistics, and so forth. This causes data traffic between the agent or supervisor desktops and the contact center servers. In addition, bandwidth calculations must account for reporting information if, for example, an agent or supervisor is remote and pulls data from a server in a central location. For more information and guidance, refer to the design guides for the individual Cisco Contact Center products, available at <https://www.cisco.com/go/srnd>.
- Depending on type of IVR solution, there could be traffic between the voice gateway and the IVR system. For example, if the voice gateways are distributed and calls arrive at a voice gateway located in a remote site with Unified IP IVR, there would be voice traffic across the WAN between the voice gateway and Unified IP IVR. With Unified CVP, the call could be queued at the remote site, with the VXML Gateway providing call treatment and queuing and therefore avoiding voice traffic across the WAN for IVR and reducing overall WAN bandwidth requirements.

Remote agents (for example, agents working from home) are also supported with Cisco Unified Contact Center. There are mainly two solutions. The first one requires the agent to use an IP phone that is connected to the central site by a broadband internet connection. In this solution, the phone is CTI


controlled by the Cisco Unified Contact Center application. The second solution is based on Cisco Unified Mobile Agent, which enables an agent to participate in a call center with any PSTN phone such as cell phone.

Multisite Contact Center with Distributed Call Processing

The model for a multisite deployment with distributed call processing consists of multiple sites, each with its own call processing cluster connected to an IP WAN. This section assumes that each Unified CM cluster has agents registered to it.


A Unified CCX deployment cannot be shared across multiple Unified CM clusters. Each Unified CM cluster requires its own Unified CCX deployment, as illustrated in [Figure 22-5](#).

Figure 22-5 Multisite Unified CCX Deployment with Distributed Call Processing



Requirements for Unified CCE differ from Unified CCX. A single Unified CCE system can span across multiple Unified CM clusters distributed across multiple geographic locations. A Unified CCE Agent PGs must be installed in each Unified CM cluster location and could be physically remote from the Unified CCE Central Controller (Call Router + Logger). [Figure 22-6](#) illustrates this type of deployment and highlights the placement of the Agent PG.

Figure 22-6 Multisite Unified CCE Deployment with Distributed Call Processing



If you require multiple contact center deployments, you could connect those deployments through Unified ICM by using the parent/child deployment model to form a single virtual contact center. The parent/child model provides several benefits, such as enterprise queuing and enterprise reporting across all the contact center deployments. It also provides complete site redundancy and higher scalability. For more details on the parent/child model, refer to the following documents:

- *Cisco Unified Contact Center Enterprise Design Guide*, available at <https://www.cisco.com/c/en/us/support/customer-collaboration/unified-contact-center-enterprise/products-implementation-design-guides-list.html>
- *Cisco Contact Center Gateway Deployment Guide for Cisco Unified ICME/CCE*, available at <https://www.cisco.com/c/en/us/support/customer-collaboration/unified-contact-center-enterprise/products-installation-guides-list.html>

Similarly to the multisite model with centralized call processing, multisite deployments with distributed call processing require careful configuration of QoS, call admission control, codecs, and so forth.

Clustering Over the IP WAN

In this deployment model, a single Unified CM cluster is deployed across multiple sites that are connected by an IP WAN with QoS features enabled. Cisco Unified Contact Center solutions can be deployed with this model. In fact, the Cisco Unified Contact Center components themselves can also be clustered over the WAN.


For example, with Unified CCE, the side A components could be remote from the Unified CCE side B components and separated from them by an IP WAN connection. (For more details on Unified CCE high availability, see [High Availability for Contact Centers, page 22-17](#).) The following design considerations apply to this type of deployment:

- The IP WAN between the two sites must be highly available, with no single point of failure. For example, the IP WAN links, routers, and switches must be redundant. WAN link redundancy could be achieved with multiple WAN links or with a SONET ring, which is highly resilient and has built-in redundancy.
- The Agent Peripheral Gateway (PG) and the CTI Manager to which it is connected must be located in the same data center. Because of the large amount of redirect and transfer traffic and additional CTI traffic, the Intra-Cluster Communication Signaling (ICCS) bandwidth requirements between the Unified CM nodes are higher when deploying Unified CCE.
- If the primary Unified CCE and Unified CM nodes are located in one site and the secondary Unified CCE and Unified CM nodes are in another site, the maximum latency between the two sites is dictated by the Unified CM latency requirement of 80 ms round trip time (RTT). However, if the Unified CCE nodes are in different locations than the Unified CM nodes, it is possible to have a higher latency between the redundant Unified CCE nodes.

Figure 22-7 illustrates a deployment of Unified CCE using clustering over the WAN. For more details, refer to the *Cisco Unified Contact Center Enterprise Design Guide*, available at

<https://www.cisco.com/c/en/us/support/customer-collaboration/unified-contact-center-enterprise/products-implementation-design-guides-list.html>


Figure 22-7 Unified CCE Deployment with Clustering Over the WAN



With Unified CCX and Unified IP IVR solutions, the primary Unified CCX or Unified IP IVR node could also be remote from the backup node. The requirements for Unified CCX deployments are different than the ones for Unified CCE deployments. For example, redundant WAN links are not required with Unified CCX. Also, the maximum latency between the primary and backup Unified CCX nodes is 80 ms RTT. Figure 22-8 illustrates this type of deployment. For more details, refer to the *Cisco Unified Contact Center Express Design Guide*, available at

<https://www.cisco.com/c/en/us/support/customer-collaboration/unified-contact-center-express/products-implementation-design-guides-list.html>

Figure 22-8 Unified CCX Deployment with Clustering Over the WAN



Design Considerations for Contact Center Deployments

This section summarizes the following major design considerations for contact center deployments:

- [High Availability for Contact Centers](#), page 22-17
- [Bandwidth, Latency, and QoS Considerations](#), page 22-18
- [Call Admission Control](#), page 22-19
- [Integration with Unified CM](#), page 22-20
- [Other Design Considerations for Contact Centers](#), page 22-20

High Availability for Contact Centers

All Cisco Unified Contact Center products provide high availability. For example, with Unified CCX or Unified IP IVR, you could add a second identical Unified CCX or Unified IP IVR node to provide high availability. The second node can reside in the same data center as the primary node or, if geographic redundancy is required, the second node can reside in a different data center across the WAN from the primary node (see [Clustering Over the IP WAN](#), page 22-15). One of the nodes would be the active node and would handle all the call processing. The other node would be in standby mode and become active only if the primary node fails. Unified CVP also supports high available deployments with multiple Unified CVP nodes, voice gateways, VXML gateways, SIP proxies, and so forth.

With Unified CCE, most of the components are required to be redundant, and the redundant instances are referred to as side A and side B instances. For example, Call Router A and Call Router B are redundant instances of the Call Router module (process) running on two different virtual machines. This redundant configuration is also referred to as *duplex mode*. The Call Routers run in synchronized execution across the two instances, which means both sides of the duplex instances process every call. Other components, such as the Peripheral Gateways, run in hot-standby mode, meaning that only one of the Peripheral Gateways is actually active at any given time.

In addition to the redundancy of the Unified Contact Center components themselves, their integration with Unified CM can also be redundant. For example, each Unified CCX or Unified IP IVR node can connect to a primary CTI Manager and also to a backup CTI Manager in case the primary CTI Manager fails. With Unified CCE, a PG side A would connect to a primary CTI Manager, while the redundant PG side B connects to the secondary CTI Manager, thus providing high availability if one CTI Manager fails.

For more details, refer to the Cisco Unified Contact Center design guides available at <https://www.cisco.com/go/srnd>.

Bandwidth, Latency, and QoS Considerations

This section describes how to provision WAN bandwidth in a multisite contact center deployment, taking into account various types of call control traffic and real-time voice traffic. It is important to understand the latency and QoS parameters because adequate bandwidth provisioning and implementation of QoS are critical components in the success of contact center deployments.

Bandwidth Provisioning

Contact center solutions require sufficient WAN bandwidth to accommodate the following main types of traffic:

- Voice traffic between the ingress gateway and the IVR system. With Unified IP IVR, if the Unified IP IVR cluster is in a central location and PSTN gateways are in remote locations, there will be voice traffic over the WAN. With Unified CVP, it is possible to queue the call at the edge and therefore keep the voice traffic local to the remote site to avoid voice traffic across a WAN link. Video queuing is also supported with the Unified CVP Video in Queue (ViQ) feature, so also consider the video traffic between the caller and the video media server.
- Voice traffic between the ingress gateway and the agent, or voice traffic between the caller and agent for internal calls. There could also be video traffic between the caller and the agent if the contact center deployment supports video.
- Voice and video signaling traffic. This is typically for the signaling traffic between the ingress gateway or caller endpoint and Unified CM, and between the agent phone and Unified CM.
- VXML Gateway traffic if Unified CVP is deployed. The traffic includes media file retrieval from the media server and VXML documents exchanged with the VXML server.
- Data traffic between the Finesse agent or supervisor desktop and the application server(s) hosting Finesse gadgets.
- Reporting traffic between the reporting user and the Unified Contact Center Reporting server.
- Traffic between Unified Contact Center servers if they are remote from each other. For example, this type of traffic occurs with clustering over the IP WAN or with multisite and distributed call processing with PGs remote from the Unified CCE Central Controller.
- Additional Intra-Cluster Communication Signaling (ICCS) traffic between the Unified CM subscribers due to the large amount of redirect and transfer traffic and additional CTI traffic.
- Voice traffic due to recording and silent monitoring. Depending on the solution, one or two RTP streams could be sent in order to silently monitor or record the conversation with an agent.

Bandwidth calculations and guidelines are provided in the Cisco Unified Contact Center design guides available at <https://www.cisco.com/go/srnd>.

Latency

Agents and supervisors can be located remotely from the call processing components and the contact center. Technically, the delay between the Finesse Server and the Finesse desktop could be very high because of high time-out values. Long latency will affect the user experience and might cause confusion or become unacceptable from the user perspective. For example, the phone could start ringing but the desktop might not be updated until later.

Latency requirements between the contact center and the call processing components, and between the contact center components themselves, depend on the contact center solutions. For example, the Unified CCX redundant nodes can be located remotely from each other, with a maximum latency of 80 ms RTT. With Unified CCE, the maximum latency between the Unified CCE components and Unified CM, or between the Unified CCE components themselves, is greater than 80 ms RTT.

For more details, refer to the Cisco Unified Contact Center design guides available at <https://www.cisco.com/go/srnd>.

QoS

Similar to deployments with other Unified Communications components, contact center deployments require the configuration of Quality of Service (QoS) to prioritize time-sensitive or critical traffic. QoS marking for voice and voice signaling in a contact center environment is the same as with other Unified Communications deployments. Traffic specific to the contact center must be marked with specific QoS markings. For example, some of the traffic for the Unified CCE private network must be marked as AF31, while other traffic must be marked as AF11. The QoS marking recommendations and QoS design guidance are documented for each Unified Contact Center solution in their respective Cisco Unified Contact Center design guides available at <https://www.cisco.com/go/srnd>.

Call Admission Control

Similar to deployments with other Unified Communications components, contact center deployments require careful provisioning of call admission control. The same mechanisms described in the chapter on [Bandwidth Management, page 13-1](#), also apply to contact center environments.

Voice traffic associated with silent monitoring and recording might not be accounted for in the call admission control calculation. For example, voice traffic from silent monitoring and recording by Unified CM (voice traffic forked at the phone) is properly accounted for, but voice traffic from desktop-based silent monitoring (desktop connected to the back of the agent IP phone) is not counted in call admission control calculations.

Call admission control for Mobile Agent and Unified CVP involves special considerations. For more details, refer to the Cisco Unified Contact Center design guides available at <https://www.cisco.com/go/srnd>.


Integration with Unified CM

The following design considerations apply when integrating Cisco Unified Contact Center components with Unified CM:

- For administration and upgrade purposes, Cisco recommends separate Unified CM clusters for contact center and non-contact center deployments. If separate clusters are not possible, then Cisco recommends separate Unified CM subscriber nodes for contact center and non-contact center applications.
- With contact center deployments, Cisco recommends that you do not use a 2:1 redundancy scheme for Unified CM. Use 1:1 redundancy to provide higher resiliency and faster upgrades.
- The integration between Unified CM and Unified CCX, Unified IP IVR, or Unified CCE is done through JTAPI. The Unified CCX cluster connects to a primary CTI Manager and also has a backup connection to a secondary CTI Manager. With Unified CCE, the Agent PG connects to only one CTI Manager. The redundant Agent PG connects to the backup CTI Manager only. If the primary CTI Manager fails, the primary Agent PG will also fail and trigger the failover.
- A single PG can control and monitor agent phones on all Unified CM subscriber pairs in a centralized deployment, as illustrated in [Figure 22-9](#).
- It is possible to integrate multiple Unified CCX deployments with a single Unified CM cluster.

For more details on Unified CM integration, refer to the Cisco Unified Contact Center design guides available at <https://www.cisco.com/go/srnd>.

Figure 22-9 Deployment with One Agent PG and Four Unified CM Subscriber Pairs



Other Design Considerations for Contact Centers

The following additional design considerations apply in the situations indicated:

- Because Unified CVP allows queuing at the edge, deploying Unified CVP instead of Unified IP IVR could lower the bandwidth requirements for multisite deployments.
- Most of the Cisco Unified Contact Center products and components can be installed in a virtualized environment based on VMware.
- Media termination point (MTP) resources might be required in some scenarios. For example, with Mobile Agents and inbound calls through SIP trunks, MTPs are required for the associated CTI ports when RFC 2833 is negotiated. MTPs are also required in some scenarios with Unified CVP. With Unified CCX Extend and Connect, MTPs are required for the associated CTI Remote Device when RFC 2833 is negotiated.

- Transcoders might be required. For example, if phones in a WAN- connected location support only the G.729 codec but Unified CVP is configured for G.711 support, then Unified CM will engage transcoders. However, an inbound call that arrives from a gateway or Cisco Unified Border Element can start with G.711 at Unified CVP then later renegotiate to G.729 with the agents without the need for transcoders.
- Some third-party contact center products are also supported with Unified CM. The integration with Unified CM could be based on JTAPI and could use CTI ports for call treatment and queuing and CTI route points. To size Unified CM correctly, it is important to have a good understanding of the call flows and their impact on Unified CM. It is also important to understand how the redundancy is implemented and whether or not it impacts Unified CM or CTI scalability.

For more detailed design considerations, refer to the Cisco Unified Contact Center design guides available at <https://www.cisco.com/go/srnd>.

Capacity Planning for Contact Centers

All deployments must be sized with the Cisco Collaboration Sizing Tool. This tool performs sizing of the contact center products such as Unified CCE, Unified IP IVR, Unified CVP, and Unified CCX. It determines the contact center resources required for your deployment, such as number of agents, number of IVR ports, and number of gateway ports. In addition to performing sizing for the contact center components themselves, the tool also sizes the rest of the Unified Communications solution, including Unified CM and voice gateways. This tool is available to Cisco employees and partners only (with proper login authentication) at <https://cucst.cloudapps.cisco.com/landing>.

In general, sizing of the contact center depends heavily on the busy hour call attempts (BHCA) for calls coming into the contact center. It also depends on other parameters such as the Service Level Goal and Target Answer Time. For example, a deployment where 90% of the calls must be answered within 30 seconds will require more contact center resources than a deployment where 80% of the calls must be answered within 2 minutes. Another parameter that impacts the sizing is whether Finesse or Finesse IP Phone Agent is used. Use the Unified CST for sizing, and consult the respective Cisco Unified Contact Center design guides, available at <https://www.cisco.com/go/srnd>, for more details.

The contact center design also impacts Unified CM sizing. The following considerations apply to sizing Unified CM when it is deployed in contact center solutions:

- The maximum number of Unified CCE agents in a single Unified CM cluster depends on the IVR solution. With Unified IP IVR, CTI route points and CTI ports are used during the call treatment queuing, which consume Unified CM resources. With Unified CVP, the call treatment and queuing are typically handled by the VXML Gateway, Unified CVP VXML server, and Unified CVP call server, with no impact on Unified CM. Therefore, a single Unified CM cluster can support more agents with Unified CVP than with Unified IP IVR.
- The Unified CCE Mobile Agent feature relies on CTI ports and therefore needs additional resources from Unified CM subscribers. Therefore, Unified CM scalability is reduced when Mobile Agents are deployed.
- With Unified CCE deployments, SIP dialing is supported. With the SIP dialer, each outbound call is placed directly from the SIP dialer port to the egress voice gateway. The call reaches Unified CM only when the call is transferred to an agent. Therefore, Unified CM capacity is much higher when the SIP dialer is used.

- When sizing Unified CM, it is also important to account for any additional CTI applications. For example, some PC clients can control a phone remotely through CTI. Some call recording applications can also integrate directly with Unified CM through the CTI Manager and can monitor agent phones, which could require additional resources from Unified CM. For more details, refer to [Computer Telephony Integration \(CTI\), page 9-28](#), and to the Cisco Unified Contact Center design guides available at <https://www.cisco.com/go/srnd>.
- Some silent monitoring and recording solutions (such as the silent monitoring and recording feature based on Unified CM) consume resources from Unified CM, whereas other solutions such as SPAN or desktop silent monitoring and recording do not.
- Again, due to the complexity associated with sizing, all deployments must be sized with the Cisco Collaboration Sizing Tool, available to Cisco employees and partners only (with proper login authentication) at <https://cucst.cloudapps.cisco.com/landing>.

For more details, refer to the Cisco Unified Contact Center design guides available at <https://www.cisco.com/go/srnd>.

Video Customer Care

For high-touch customer engagements, enhancing the customer care options to include video-enabled customer experiences can greatly improve the interaction for the both the customer and the agent.

Cisco Remote Expert Solution

The Cisco Remote Expert Solution enables customers as well as internal employees to connect with experts across multiple channels. It also delivers a consistent, interactive experience that helps optimize revenue, improve expert productivity, and build customer loyalty. Cisco Remote Expert creates a virtual pool of specialists, manages their availability, and quickly connects customers with experts across multiple channels and devices, using high-quality audio and video.

Designed to deliver a consistent customer and employee experience across multiple touch points and devices, the Cisco Remote Expert Solution is an end-to-end, multichannel collaboration platform that establishes a new industry benchmark for customer care and provides the following benefits:

- Improved response time
Customers can reach your experts over video with the touch of a button from personal devices, from kiosks, or from customer workstations within your store, branch, or clinic.
- Increased sales close ratios
Cisco Remote Expert can intelligently route customers to the right resource required to satisfy product and service inquiries.
- Improved cross-sell and up-sell opportunities
Customers can engage with highly trained experts who can address the customers' needs and suggest adjacent products and services.
- Increased productivity
Subject-matter experts can use a single platform to reach customers, regardless of device or location.

The Cisco Remote Expert Solution employs industry-leading, high-quality collaboration products and services supported by a Cisco Validated Design reference architecture and partner ecosystem.

For more details on Remote Expert, refer to the *Cisco Remote Expert Solution Design Guide*, available at https://www.cisco.com/c/en/us/solutions/enterprise/design-zone/remote_expert.html

Network Management Tools

Unified CCE is managed with the Simple Network Management Protocol (SNMP). Unified CCE devices have a built-in SNMP agent infrastructure that supports SNMP v1, v2c, and v3, and it exposes instrumentation defined by the CISCO-CONTACT-CENTER-APPS-MIB. This MIB provides configuration, discovery, and health instrumentation that can be monitored by standard SNMP management stations. Moreover, Unified CCE provides a rich set of SNMP notifications that alert administrators of any faults in the system. Unified CCE also provides a standard syslog event feed (conforming to RFC 3164) for those administrators who want to take advantage of a more verbose set of events.

For more information about configuring the Unified CCE SNMP agent infrastructure and the syslog feed, refer to the *SNMP Guide for Cisco Unified ICM/Contact Center Enterprise & Hosted*, available at

<https://www.cisco.com/c/en/us/support/customer-collaboration/unified-contact-center-enterprise/products-installation-and-configuration-guides-list.html>

Unified CVP health monitoring can be performed by using any SNMP standard monitoring tool to get a detailed visual and tabular representation of the health of the solution network. All Unified CVP product components and most Unified CVP solution components also issue SNMP traps and statistics that can be delivered to any standard SNMP management station or monitoring tool.

Unified CCX can also be managed with SNMP and a syslog interface.

Cisco Prime Collaboration can also help manage a Contact Center deployment. For example, Cisco Prime Collaboration Assurance can be used to monitor the number of active calls, number of inbound calls per second, or number of agents logged on.

Also, the Prime Contact Center Assurance Module can be added if Prime Collaboration Assurance Advanced has previously been implemented. The Prime Contact Center Assurance Module diagrams the topology of your customer care environment and the relationship between components. It provides event correlation to speed up error root cause analysis, provides a performance dashboard to help detect and fix performance issue, and provides call trace analysis to help identify devices that break a call flow.

