



Call Recording and Monitoring

Revised: May 24, 2013; OL-29367-04

Call monitoring and recording solutions provide a way to automatically monitor and record audio and video calls that traverse various components in a Unified Communications and Collaboration solution, such as Cisco IP Phones, Cisco Unified Border Element devices, or Cisco switches. These recordings can then be used by call centers and other enterprise for various purposes such as compliance purposes, transcription and speech analysis, or for podcasting and blogging purposes. This chapter provides an overview of various call recording solutions available for Cisco Unified Communications and Collaboration solutions for both audio and video calls. The chapter also outlines basic design considerations for call recording solutions embedded within a Cisco Unified Communications and Collaboration solution.

What's New in This Chapter

[Table 25-1](#) lists the topics that are new in this chapter or that have changed significantly from previous releases of this document.

Table 25-1 *New or Changed Information Since the Previous Release of This Document*

New or Revised Topic	Described in:	Revision Date
Cisco TelePresence Content Server	Cisco TelePresence Content Server, page 25-8	May 24, 2013
This chapter was first added to this document as new content	All sections of this chapter	April 2, 2013

Types of Monitoring and Recording Solutions

This section describes the following types of call recording and monitoring solutions:


- [SPAN-Based Solutions, page 25-2](#)
- [Silent Monitoring and Recording with Unified CM, page 25-3](#)
- [Cisco MediaSense, page 25-4](#)
- [Agent Desktop, page 25-8](#)
- [Cisco TelePresence Content Server, page 25-8](#)

SPAN-Based Solutions

Recording solutions based on a Switched Port Analyzer (SPAN) use the packet sniffing technology for recording calls. SPAN is a method of monitoring network traffic. When SPAN is enabled on a switch port or VLAN, the switch sends a copy of all network packets traversing that port or VLAN to another port where a recording or monitoring server (such as Cisco Unified Contact Center Express, Cisco Unified Workforce Optimization Quality Management, or a third-party recording server, for example) analyses those packets. It detects and decodes the VoIP packets embedded in the network traffic and stores them as audio on a storage device. SPAN can be enabled on the ports connected to Cisco Unified Communications Manager (Unified CM), a Cisco Voice Gateway, or Cisco IP Phones, as required. For example, for recording internal calls between IP phones, SPAN should be enabled on switch ports connected to the IP phones.

Figure 25-1 illustrates a SPAN-based recording solution deployment for recording internal calls. The ports marked as source ports connected to IP phones are mirrored to the destination port connected to the recording server.

Figure 25-1 SPAN-Based Recording Call Flow for Internal Calls



Several Cisco partners provide SPAN-based recording servers and applications for Cisco Unified Communications and Collaboration solutions. Refer to the specific partner product documentation for further technical details.

In addition, network traffic flow needs to be considered for appropriate bandwidth provisioning when port mirroring is enabled. Also consider the type of servers deployed for applications such as Cisco Unified CM and Cisco TelePresence Recording Server, especially when hosting these applications in a virtualized setup.

SPAN-Based Recording and Virtualization

This section reviews some common SPAN-based deployments with virtualization enabled and lists some of the limitations. Note that the SPAN feature requires the source and destination ports to be on the same switch. VMware provides support for the SPAN feature only on Cisco Video Distribution Suite (VDS) with VMware vSphere release 5.0.

In a virtualized setup, some of the Unified Communications and contact center applications and the port analyzer application may be deployed on virtual machines on the same host or on different hosts. There are some limitations to SPAN-based recording solutions in a virtualized setup. For example, the following features are not supported for deployments of Cisco Unified Contact Center Enterprise (Unified CCE) with virtualization:

- Remote silent monitoring
- SPAN-based silent monitoring and recording on Cisco Unified Computing System (UCS) B-Series chassis

**Note**

SPAN-based silent monitoring and recording is not supported on the UCS B-Series chassis. Cisco MediaSense does not support SPAN based recording or call monitoring.

Silent Monitoring and Recording with Unified CM

Cisco Unified CM silent monitoring and call recording solutions provide the ability to monitor and record customer conversations for compliance purpose. The Silent Call Monitoring feature allows a supervisor to listen to a conversation between an agent and a customer with neither the agent nor the customer aware of the supervisor's presence on the call. The Call Recording feature allows system administrators or authorized personnel to archive conversations between the agent and the customer.

**Note**

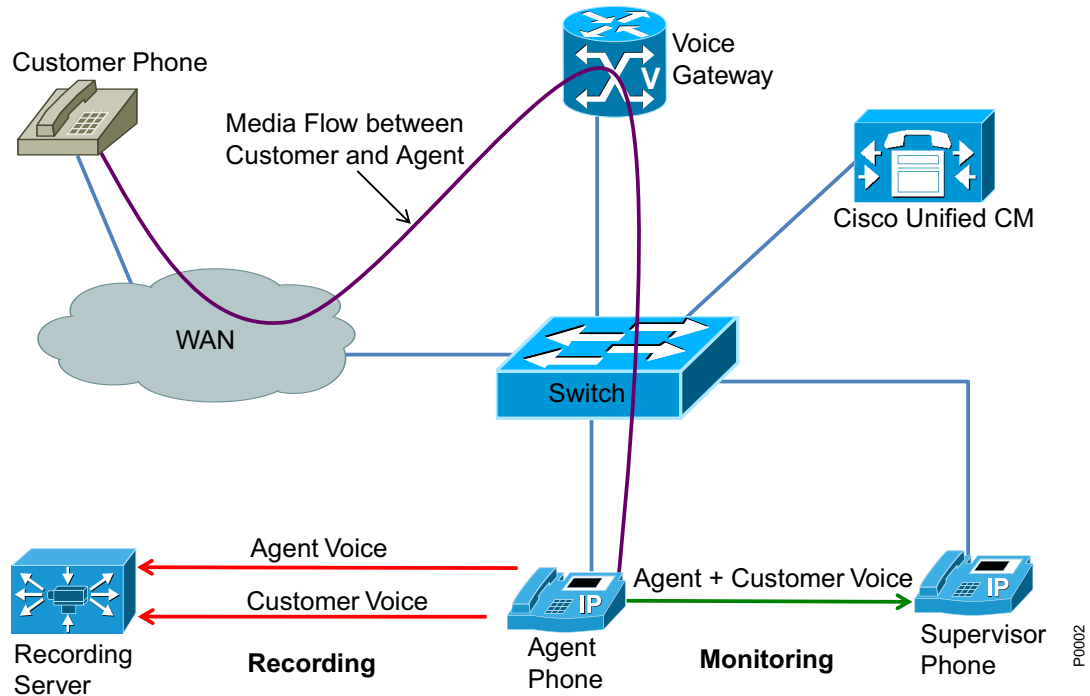
Cisco Unified Communications Manager supports the Silent Call Monitoring and Call Recording features only within a single cluster.

Cisco Unified CM uses an architecture based on the IP phones to provide call monitoring and recording for audio calls. It uses the IP phone built-in bridge (BIB) to enable call monitoring and recording.

For call monitoring, the agent phone combines the two voice RTP streams on the agent phone and sends the resulting stream to the supervisor phone. Call monitoring can be invoked by call center applications through the JTAPI or TAPI interfaces of Unified CM.

The call recording device is configured on Unified CM as a SIP trunk device. For recording, the agent phone forks the two streams to the recording server. The two streams, one for the agent voice and one for the customer voice, get recorded separately. The recorded streams can be mixed to produce the conversation using third-party applications.

[Figure 25-2](#) illustrates the basic call recording setup using Unified CM Silent Call Monitoring and Call Recording features. Agent and supervisor phones are CTI controlled.

Figure 25-2 Call Monitoring and Recording with Unified CM

For details on the Unified CM Silent Call Monitoring and Call Recording features, refer to the latest version of the *Cisco Unified Communications Manager Features and Services Guide*, available at

http://www.cisco.com/en/US/products/sw/voicesw/ps556/prod_maintenance_guides_list.html

Cisco MediaSense

Cisco MediaSense is a SIP-based, network-level service that provides voice and video media recording capabilities for network devices. It is fully integrated into the Unified CM architecture and can capture and store VoIP conversations that traverse appropriately configured Unified CM IP phones or Cisco Unified Border Element devices by invoking media forking capabilities on the IP phones and Cisco Unified Border Element devices. In addition, an IP phone user or SIP endpoint device may call the Cisco MediaSense system directly in order to leave a recording that consists of only media generated by that user. Such recordings may include video as well as audio, thus offering a simple and easy method for recording video blogs and podcasts. While the recording is in progress, it can also be streamed live using a media player such as VLC or Apple QuickTime. Cisco MediaSense uses an HTTP interface to access and play back recordings.



Note

Most but not all Cisco Unified IP Phones support media forking. Those that do not support media forking cannot be used for phone-based recording. For a list of IP phones that support phone-based media recording with Cisco MediaSense, refer to the latest version of the *Solution Reference Network Design for Cisco MediaSense*, available at

http://www.cisco.com/en/US/products/ps11389/products_implementation_design_guides_list.html.

Cisco MediaSense also provides an administration and reporting interface to configure the cluster and manage recordings. It provides secure media storage on an encrypted storage area network (SAN). It does not currently support secure media relay using sRTP or other means.

Deployment

Cisco MediaSense can be deployed only on top of a VMware hypervisor running on a Cisco C-Series or B-Series platform. It is deployed as a single server or as a cluster of up to a five nodes, depending upon the required capacity of the system. It can also be deployed on a Cisco Services-Ready Engine (SRE) platform with up to two SRE modules. In a multi-node setup, there are three types of servers:

- Primary — Provides both database and media operations.
- Secondary — Provides high availability for the database as well as both database and media operations.
- Expansion — Provides additional capacity for media operations but no data operations. It is not supported in Cisco SRE deployments.

When deploying multiple Cisco MediaSense clusters, Cisco recommends partitioning the IP phones carefully among the various clusters so that each IP phone gets recorded by only one cluster.




Note

SIP proxy servers are not supported between Cisco MediaSense and Unified CM or Cisco Unified Border Element.

Cisco Unified CM Deployments

Figure 25-3 illustrates a basic Unified CM deployment for call recording. Once a call is established from a signaling perspective, the media flows directly between the external phone and the internal IP phone. The IP phone is configured to fork media to Cisco MediaSense for call recording. If the call gets transferred to another IP phone, call recording ends unless the phone that accepts the transferred call is also configured for recording.

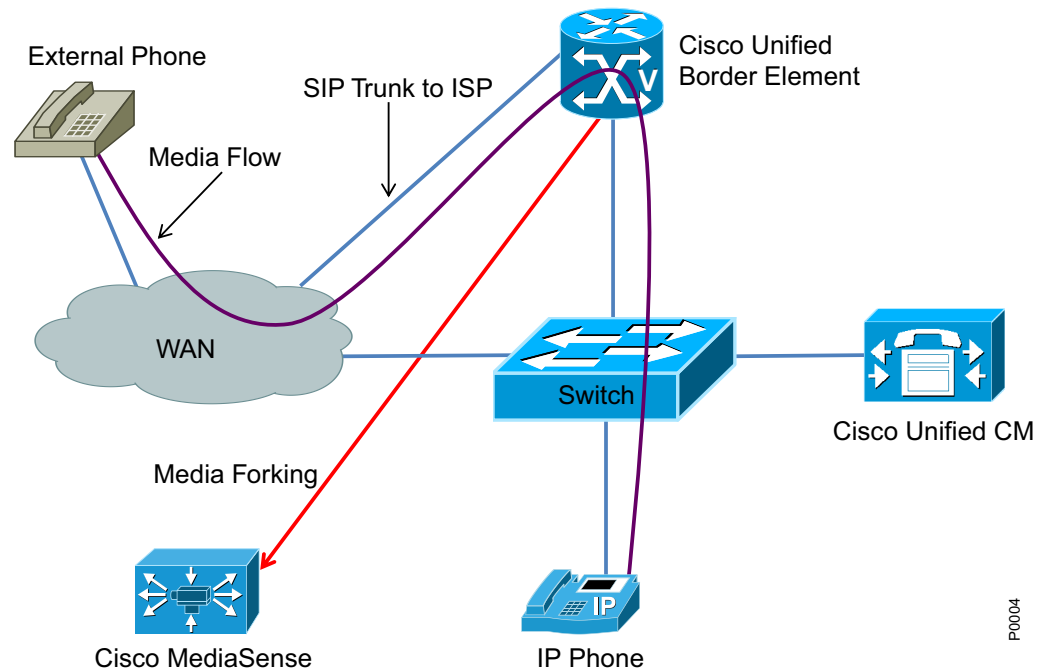
Figure 25-3 Call Recording with Cisco Unified CM and MediaSense



Cisco Unified Border Element Deployments

Cisco Unified Border Element media forking provides the ability to capture the end-to-end conversation from a caller's perspective, no matter how the call traverses through the enterprise. Figure 25-4 illustrates a basic Cisco Unified Border Element deployment. The Cisco Unified Border Element device does media forking by means of a recorder profile configuration attached to one or more dial peers. Cisco recommends attaching the recording profile to the outbound dial peer.

Figure 25-4 Call Recording with Cisco Unified Border Element and MediaSense



Note

Network-based recording is supported only for SIP-to-SIP call flows. Cisco Unified Border Element software with media forking runs only on Cisco Integrated Services Routers Generation 2 (ISR G2). Media forking is not supported on Cisco Aggregation Services Routers (ASR).

Any requirements around call recording need to be considered when doing capacity planning for Cisco Unified Border Element devices because they require additional DSP resources and memory resources. For the memory requirements, Cisco recommends provisioning the Cisco Unified Border Element devices with the maximum amount of memory when enabling call recording. Also, media forking increases bandwidth usage on the link between the Cisco Unified Border Element and the Cisco MediaSense server. The percentage of calls getting recorded needs to be factored in when calculating bandwidth requirement.

For details on configuring Cisco Unified Border Element devices to enable network-based recording, refer to the section on *Network-Based Recording Using Cisco UBE* in the *Cisco Unified Border Element Protocol-Independent Features and Setup Configuration Guide*, available at

http://www.cisco.com/en/US/docs/ios-xml/ios/voice/cube_proto/configuration/15-mt/cube-proto-15-mt-book.html

Cisco SME Deployments

In a deployment of Cisco Unified CM Session Management Edition (SME), Cisco MediaSense is supported only in the leaf clusters. The phones that need to be recorded and the Cisco MediaSense cluster must be part of the same SME leaf cluster. Separate MediaSense clusters need to be deployed for different SME leaf clusters. Cisco MediaSense deployed in an SME leaf cluster can record calls only for that leaf cluster.

For more details on other supported deployments available with Cisco MediaSense, refer to the section on *Solution-Level Deployment Models* in the latest version of the *Solution Reference Network Design for Cisco MediaSense*, available at

http://www.cisco.com/en/US/products/ps11389/products_implementation_design_guides_list.html

Agent Desktop

Agent desktop monitoring and recording solutions are specific to contact center deployments that enable supervisors to do silent monitoring and initiate call recording when needed. Several agent desktop monitoring and recording solutions are available, such as:

- Cisco Agent Desktop (CAD) Silent Monitoring and Recording
- Cisco Remote Silent Monitoring (RSM)

These solutions are described in details in the following documents:


- *Cisco Unified Contact Center Enterprise Solution Reference Network Design*, available at http://www.cisco.com/en/US/products/sw/custcosw/ps1844/products_implementation_design_guides_list.html
- *Cisco Unified Contact Center Express Solution Reference Network Design*, available at http://www.cisco.com/en/US/products/sw/custcosw/ps1846/products_implementation_design_guides_list.html

Cisco TelePresence Content Server

The Cisco TelePresence Content Server is a network appliance that provides the ability to record and stream Cisco TelePresence and third-party video conferences and multimedia presentations that can be distributed to media devices or shared through applications such as Cisco Show and Share.

The Cisco TelePresence Content Server can be used to record content and create media from any H.323 or SIP videoconference endpoint. Cisco TelePresence Content Server release 5.3 supports up to 10 recording ports. The TelePresence Content Server solution can be deployed as a single Content Server or as a cluster with up to 10 servers in a Cisco TelePresence Video Communication Server (VCS) cluster. Clustering several servers together increases the total recording and playback capacity. A cluster can have a mix of 5-port and 10-port servers. The cluster uses a network load balancing (NLB) solution that distributes incoming requests across the cluster. Each Cisco TelePresence Content Server cluster is registered to a Cisco VCS cluster that load-balances the inbound calls. Note that all servers in a cluster must be located at the same physical site, within a network round-trip time (RTT) to the Network Attached Storage (NAS) and Structured Query Language (SQL) servers not exceeding 10 ms.

Figure 25-5 illustrates TelePresence Content Server clustering.

Figure 25-5 Cisco TelePresence Content Server Clustering

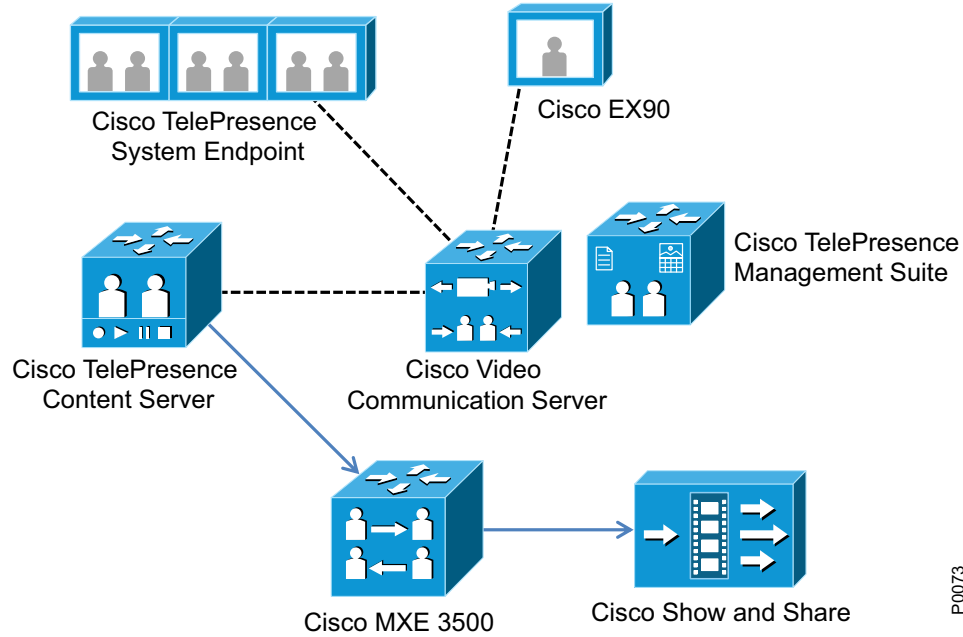
P0072

Cisco TelePresence Content Server Deployments

The Cisco TelePresence Content Server is not supported with Unified CM-only deployments. The following deployments are supported for the Cisco TelePresence Content Server recording solution:

- Cisco TelePresence Content Server registered directly to Cisco Video Communication Server (VCS)
Cisco recommends deploying the TelePresence Content Server solution with Cisco VCS to enable all features.
- Cisco TelePresence Content Server registered to Cisco IOS Gatekeeper
This deployment does not support all the features (for example, SIP functionality) that are available in a TelePresence Content Server deployment with Cisco VCS.
- Standalone Cisco TelePresence Content Server
This is the most simplistic deployment but is not recommended. This deployment does not provide any call control and has other limitations such as support for only a single recording alias.

Figure 25-6 illustrates a sample deployment where the Cisco TelePresence Content Server is registered to Cisco VCS. The Cisco TelePresence System endpoint and Cisco EX90 are also registered to Cisco VCS. A Cisco Media Experience Engine (MXE) 3500 is also used to transcode the video recordings and publish to Cisco Show and Share. The TelePresence Content Server records the call between the two endpoints by joining the TelePresence bridge. It then sends the recorded video to the MXE 3500 by means of File Transfer Protocol (FTP). The MXE 3500 transcodes the video and publishes it to the Cisco Show and Share application.

Figure 25-6 Deployment of Cisco TelePresence Content Server with Cisco VCS

P0073

For details on the Cisco Telepresence Content Server, refer to the latest version of the *Cisco TelePresence Content Server Administration and User Guide*, available at

http://www.cisco.com/en/US/products/ps11347/prod_maintenance_guides_list.html

Capacity Planning for Monitoring and Recording

Enabling any type of monitoring and/or call recording impacts the overall Unified Communications system capacity. 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. Consider the following points when doing capacity planning for Unified Communications systems with call recording enabled:

- With Unified CM call recording, each recorded call counts as two calls towards the BHCA cluster capacity.
- Bandwidth requirements increase when media forking is enabled on IP phones or Cisco Unified Border Element devices to send forked media to the recording server. In case of agent desktop monitoring and recording, the bandwidth utilization can be bursty, depending on how many calls are being monitored or recorded at a given time.
- DSP resource utilization is impacted on Cisco Unified Border Element when doing media forking to Cisco MediaSense.
- Memory utilization on the router increases for each call that is recorded.
- In cases where CTI applications interact with Cisco Unified CM to invoke recording and monitoring, you should consider the Unified CM cluster deployment model and load-balance the CTI applications across the cluster.

Due to the complexity associated with sizing, all deployments must be sized with the Cisco Unified Communications Sizing Tool, available to Cisco employees and partners only (with proper login authentication) at

<http://tools.cisco.com/cucst>

