



Cisco DSP Resources for Transcoding, Conferencing, and MTP

This chapter describes how Cisco digital signal processor (DSP) resources are used for transcoding and conferencing. The modules, which are available for use with Cisco CallManager, can perform conferencing, Media Termination Point (MTP), and transcoding services in addition to serving as a PSTN gateway.

This chapter covers the following topics:

- [Understanding Cisco DSP Resources, page 28-1](#)
- [Hardware-Based MTP/ Transcoding Services, page 28-2](#)
- [Hardware-Based Conferencing Services, page 28-4](#)
- [Supported Cisco Catalyst Gateways and Cisco Access Routers, page 28-5](#)
- [Where to Find More Information, page 28-10](#)

Understanding Cisco DSP Resources

DSP resources on the Cisco gateway, for example, Catalyst 4000 (WS-X4604-GWY), Catalyst 6000 (WS-6608-T1 or WS-6608-E1), Cisco 2600, Cisco 2600XM, Cisco 2800, Cisco 3600, Cisco 3700, Cisco 3800, or Cisco VG200, provide hardware support for IP telephony features that are offered by Cisco CallManager. These features include hardware-enabled voice conferencing, hardware-based MTP support for supplementary services, and transcoding services.

**Note**

Verify with your Cisco account manager which devices support conferencing, media termination points, and transcoding services.

The DSP resource management (DSPRM) maintains the state for each DSP channel and the DSP. DSPRM maintains a resource table for each DSP. The following responsibilities belong to DSPRM:

- Discover the on-board DSP SIMM modules and, based on the user configuration, determine the type of application image that a DSP uses.
- Reset DSPs, bring up DSPs, and download application images to DSP.
- Maintain the DSP initialization states and the resource states and manage the DSP resources (allocation, deallocation, and error handling of all DSP channels for transcoding and conferencing).

- Interface with the backplane Protocol Control Information (PCI) driver for sending and receiving DSP control messages.
- Handle failure cases, such as DSP crashes and session terminations.
- Provide a keepalive mechanism between the DSPs and the primary and backup Cisco CallManagers. The primary Cisco CallManager can use this keepalive to determine when DSPs are no longer available.
- Perform periodic DSP resource checks.

When a request is received from the signaling layers for a session, the system assigns the first available DSP from the respective pool (transcoding or conferencing), as determined by media resource groups and media resource group lists, along with the first available channel. DS prm maintains a set of MAX limits (such as maximum conference sessions per DSP or maximum transcoding session per DSP) for each DSP.

A switchover occurs when a higher order Cisco CallManager becomes inactive or when the communication link between the DSPs and the higher order Cisco CallManager disconnects. A switchback occurs when the higher order Cisco CallManager becomes active again and DSPs can switch back to the higher order Cisco CallManager. During a switchover and switchback, the gateway preserves active calls. When the call ends, the gateway detects RTP inactivity, DSP resources release, and updates occur on the Cisco CallManager.

Hardware-Based MTP/ Transcoding Services

Introducing the WAN into an IP telephony implementation forces the issue of voice compression. After a WAN-enabled network is implemented, voice compression between sites represents the recommended design choice to save WAN bandwidth. This choice presents the question of how WAN users use the conferencing services or IP-enabled applications, which support only G.711 voice connections. Using hardware-based Media Termination Point (MTP)/transcoding services to convert the compressed voice streams into G.711 provides the solution.

The MTP service can act either like the original software MTP resource or as a transcoding MTP resource. An MTP service can provide supplementary services such as hold, transfer, and conferencing when the service is using gateways and clients that do not support the H.323v2 feature of EmptyCapabilitiesSet. The MTP, provided by the Cisco IP Voice Media Streaming Application service, can be activated as co-resident with Cisco CallManager or activated separately without Cisco CallManager. Both of these services operate on the Cisco CallManager appliance (server). The Cisco IP Voice Media Streaming Application service installs as a component with Cisco CallManager; however, for a dedicated MTP server, the Cisco CallManager service would not be activated (only the Cisco Voice IP Voice Media Streaming Application service).

When MTP is running in software on Cisco CallManager, the resource supports 48 MTP sessions. When MTP is running on a separate Cisco CallManager appliance (server), the resource supports up to 128 MTP sessions. In addition, Cisco Voice Gateway Routers also have the ability to provide MTP services.

Observe the following design capabilities and requirements for MTP transcoding:

- Provision MTP transcoding resources appropriately for the number of IP WAN callers to G.711 endpoints.
- Each transcoder has its own jitter buffer of 20-40 ms.

The following summary gives caveats that apply to MTP transcoding:

- Make sure that each Cisco CallManager has its own configured MTP transcoding resource.
- If transcoding is required between Cisco CallManager clusters, make sure that the intercluster trunk is configured with an MTP resource. All calls between Cisco CallManager clusters will go through the MTPs.
- If all n MTP transcoding sessions are utilized, and an $n + 1$ connection is attempted, the next call will complete without using the MTP transcoding resource. If this call attempted to use the software MTP function to provide supplementary services, the call would connect, but any attempt to use supplementary services would fail and could result in call disconnection. If the call attempted to use the transcoding features, the call would connect directly, but no audio would be received. If a transcoder is required but not available, the call would not connect.

For specific information on the number of sessions that are supported, see the “[Supported Cisco Catalyst Gateways and Cisco Access Routers](#)” section on page 28-5.


IP-to-IP Packet Transcoding and Voice Compression

You can configure voice compression between IP phones through the use of regions and locations in Cisco CallManager. However, the Cisco Catalyst conferencing services and some applications currently support only G.711, or uncompressed, connections. For these situations, MTP transcoding or packet-to-packet gateway functionality provides modules for the Cisco Catalyst 4000 and Cisco Catalyst 6000. A packet-to-packet gateway designates a device with DSPs that has the job of transcoding between voice streams by using different compression algorithms. For example, a user on an IP phone at a remote location calls a user at the central location. Cisco CallManager instructs the remote IP phone to use compressed voice, or G.729a, only for the WAN call. If the called party at the central site is unavailable, the call may roll to an application that supports G.711 only. In this case, a packet-to-packet gateway encodes the G.729a voice stream to G.711 to leave a message with the voice-messaging server.

Voice Compression, IP-to-IP Packet Transcoding, and Conferencing

Connecting sites across an IP WAN for conference calls presents a complex scenario. In this scenario, the modules must perform the conferencing service as well as the IP-to-IP transcoding service to uncompress the WAN IP voice connection. In [Figure 28-1](#), a remote user joins a conference call at the central location. This three-participant conference call uses seven DSP channels on the Catalyst 4000 module and three DSP channels on the Cisco Catalyst 6000. The following list gives the channel usage:

- Cisco Catalyst 4000
 - One DSP channel to convert the IP WAN G.729a voice call into G.711
 - Three conferencing DSP channels to convert the G.711 streams into TDM for the summing DSP
 - Three channels from the summing DSP to mix the three callers together
- Cisco Catalyst 6000
 - Three conferencing DSP channels. On the Cisco Catalyst 6000, all voice streams get sent to single logical conferencing port where all transcoding and summing takes place.

Figure 28-1 Multisite WAN Using Centralized MTP Transcoding and Conferencing Services

43752

IP-to-IP Packet Transcoding Across Intercluster Trunks

Intercluster trunks connect Cisco CallManager clusters. Intercluster trunks allocate a transcoder on a dynamic basis.

The Cisco Catalyst 6000 module uses the MTP service regardless of whether transcoding is needed for a particular intercluster call. Cisco CallManager supports compressed voice call connection through the MTP service if a hardware MTP is used.

The following list gives intercluster MTP/transcoding details:

- Outbound intercluster calls will use an MTP/transcoding resource from the Cisco CallManager from which the call originates.
- Inbound intercluster call will use the MTP/resource from the Cisco CallManager that terminates the inbound intercluster trunk.
- Allocate additional DSP MTP/transcoding resources to Cisco CallManagers terminating intercluster trunks.
- For compressed callers, you can accurately provision the MTP transcoding resources.

Hardware-Based Conferencing Services

Hardware-enabled conferencing designates the ability to support voice conferences by using DSPs to perform the mixing of voice streams to create multiparty conference sessions. The voice streams connect to conferences through packet or time-division-multiplexing (TDM) interfaces.

The network modules, depending on the type, support both uncompressed and compressed VOIP conference calls. The modules use Skinny Client Control Protocol to communicate with Cisco CallManager to provide conferencing services. When the conferencing service registers with

Cisco CallManager, it announces that only G.711 calls can connect to the conference. If any compressed calls request to join a conference, Cisco CallManager connects them to a transcoding port first to convert the compressed call to G.711.

Observe the following recommendations when you are configuring conferencing services:

- When you are provisioning an enterprise with conference ports, first determine how many callers will attempt to join the conference calls from a compressed Cisco CallManager region. After you know the number of compressed callers, you can accurately provision the MTP transcoding resources.
- Conference bridges can register with more than one Cisco CallManager at a time, and Cisco CallManagers can share DSP resources through the Media Resource Manager (MRM).

For specific information on the number of sessions that are supported, see the “[Supported Cisco Catalyst Gateways and Cisco Access Routers](#)” section on page 28-5.

Supported Cisco Catalyst Gateways and Cisco Access Routers

For specific information on the number of supported conferencing, transcoding, and MTP sessions for Cisco Catalyst Gateways and Cisco Access Routers, see the following sections:

- [Cisco Catalyst 4000 WS-X4604-GWY, page 28-5](#)
- [Cisco Catalyst 6000 WS-6608-T1 or WS-6608-E1, page 28-6](#)
- [Cisco 2600, Cisco 2600XM, Cisco 2800, Cisco 3600, Cisco 3700, Cisco 3800, and Cisco VG200 for NM-HDV, page 28-8](#)
- [Cisco 2600XM, Cisco 2691, Cisco 2800, Cisco 3600, Cisco 3700, and Cisco 3800 for NM-HD and NM-HDV2, page 28-8](#)

Cisco Catalyst 4000 WS-X4604-GWY

The PSTN gateway and voice services module for the Cisco Catalyst 4003 and 4006 switches supports three analog voice interface cards (VICs) with two ports each or one T1/E1 card with two ports and two analog VICs. Provisioning choices for the VIC interfaces include any combination of Foreign Exchange Office (FXO), Foreign Exchange Station (FXS), or Ear & Mouth (E&M). Additionally, when configured as an IP telephony gateway from the command-line interface (CLI), this module can support conferencing and transcoding services.

You can configure the Cisco Catalyst 4000 voice gateway module in either toll bypass mode or gateway mode; however, you can configure the module conferencing and transcoding resources only in gateway mode. Gateway mode designates the default configuration. From the CLI, you can change the conferencing-to-transcoding ratios. After the gateway mode is enabled, the 24 DSPs (4 SIMMs with 6 DSPs each) for the module occur as described in the following bullets:

- Over the PSTN gateway using G.711 only—96 calls
- In a G.711 conference only—24 conference participants; maximum of 4 conferences of 6 participants each

Unlike the WS-X6608-x1, which can mix all conference call participants, the Cisco Catalyst 4000 WS-X4604-GWY module sums only the three dominant speakers. The WS-X4604-GWY dynamically adjusts for the dominant speakers and determines dominance primarily by voice volume, not including any background noise.

**Caution**

The Cisco Catalyst 4000 conferencing services support G.711 connections only, unless an MTP transcoding service is used.

- Transcoding to G.711—16 MTP transcoding sessions

The following information applies to the Cisco Catalyst 4000 module:

- The WS-X4604-GWY uses a Cisco IOS interface for initial device configuration. All additional configuration for voice features takes place in Cisco CallManager.
- The WS-X4604-GWY can operate as a PSTN gateway (toll bypass mode) as well as a hardware-based transcoder or conference bridge (gateway mode). To configure this module as a DSP farm (gateway mode), enter one or both of the following CLI commands:

```
voicecard conference
voicecard transcode
```

- The WS-X4604-GWY requires its own local IP address in addition to the IP address for Cisco CallManager. Specify a loopback IP address for the local Signaling Connection Control Part.
- You can define a primary, secondary, and tertiary Cisco CallManager for both the conferencing and MTP transcoding services.


Cisco Catalyst 6000 WS-6608-T1 or WS-6608-E1

The WS-6608-T1 (or WS-6608-E1 for European countries) designates the same module that provides T1 or E1 PSTN gateway support for the Cisco Catalyst 6000. This module comprises eight channel-associated-signaling (CAS) or primary rate interface (PRI) interfaces, each of which has its own CPU and DSPs. After the card is added from Cisco CallManager as a voice gateway, you configure it as a conferencing or MTP transcoding resource. Each port acts independently of the other ports on the module. Specifically, you can configure each port only as a PSTN gateway interface, a conferencing node, or an MTP transcoding node. In most configurations, configure a transcoding resource for each conferencing resource.

Whether acting as a PSTN gateway, a conferencing resource, or an MTP transcoding resource, each port on the module requires its own IP address. Configure the port to have either a static IP address or an IP address that the DHCP provides. If a static IP is entered, you must also add a TFTP server address because the ports actually get all configuration information from the downloaded TFTP configuration file.

[Figure 28-2](#) shows one possible configuration of the Cisco Catalyst 6000 voice gateway module. This diagram shows two of the eight ports of the module as configured in PSTN gateway mode, three ports in conferencing mode, and three ports in MTP transcoding mode.

Figure 28-2 Cisco Catalyst 6000 Voice Gateway Module



After a port is configured through the Cisco CallManager interface, each port can support one of the following configurations:

- WS-6608-T1 over the PSTN gateway —24 calls per physical DS1 port; 192 calls per module
- WS-6608-E1 over the PSTN gateway—30 calls per physical DS1 port; 240 calls per module
- For a G.711 or G.723 conference—32 conferencing participants per physical port; maximum conference size of 16 participants
- For a G.729 conference—24 conferencing participants per physical port; maximum conference size of 16 participants



After the WS-X6608 is added as a T1 or E1 Cisco gateway, you can configure it, on a per-port basis, for conferencing services.

On the Cisco Catalyst 6000, conferencing services cannot cross port boundaries.

The following capacities apply to simultaneous transcoding and conferencing:

- For transcoding from G.723 to G.711—32 MTP transcoding sessions per physical port; 256 sessions per module
- For transcoding from G.729 to G.711—24 MTP transcoding sessions per physical port; 192 sessions per module

Cisco 2600, Cisco 2600XM, Cisco 2800, Cisco 3600, Cisco 3700, Cisco 3800, and Cisco VG200 for NM-HDV

NM-HDV supports the previous Cisco gateways.

The following list represents the maximum number of sessions:

- G.711, G.729, GSM FR, and GSM EFR conference sessions—Per network module, 15



Tip Maximum participants per conference session equals 6.

- Transcoding from G.711 to G.729—Per network module, 60
- Transcoding from G.711 to GSM FR/GSM EFR—Per network module, 45



Caution On these gateways, transcoding services cannot cross port boundaries.

Cisco MTP transcoding service only supports HBR codec to G.711 conversion and vice versa. No support exists for LBR-to-LBR codec conversion.

Cisco 2600XM, Cisco 2691, Cisco 2800, Cisco 3600, Cisco 3700, and Cisco 3800 for NM-HD and NM-HDV2



Tip The NM-HDV2 does not support the Cisco 3660.

The following list represents the maximum number of sessions that are available for conferences, transcoding, and MTP for NM-HD and NM-HDV2:

Per NM-HD-1V/2V

- G.711 only conference—8 sessions
- G.729, G.729a, G.729ab, and G.729b conference—2 sessions
- GSM FR conference—Not applicable
- GSM EFR conference—Not applicable



Tip Maximum number of participants per conference equals 8.

- Transcoding for G.711 to G.729a/G.729ab/GSMFR—8 sessions
- Transcoding for G.711 to G.729/G.729b/GSM EFR—6 sessions

Per NM-HDV2

- G.711 only conference—50 sessions
- G.729, G.729a, G.729ab, G.729b conference—32 sessions
- GSM FR conference—14 sessions

- GSM EFR conference—10 sessions
- Transcoding for G.711 to G.729a/G.729ab/GSMFR—128 sessions
- Transcoding for G.711 to G.729/G.729b/GSM EFR—96 sessions

**Tip**

For a software MTP (DSP-less with same packetization period for both devices supporting G.711 to G.711 or G.729 to G.729 codecs), 500 sessions can occur per gateway; for a hardware MTP (with DSP, using G.711 codec only), 200 sessions can occur per NM-HDV2 and 48 per NM-HD.

Per 2801/2811 (2 PVDM2-64)

- G.711 only conference—50 sessions
- G.729, G.729a, G.729ab, G.729b conference—16 sessions
- GSM FR conference—7 sessions
- GSM EFR conference—5 sessions
- Transcoding for G.711 to G.729a/G.729ab/GSMFR—64 sessions
- Transcoding for G.711 to G.729/G.729b/GSM EFR—48 sessions

Per 2821/2851 (3 PVDM2-64)

- G.711 only conference—50 sessions
- G.729, G.729a, G.729ab, G.729b conference—24 sessions
- GSM FR conference—10 sessions
- GSM EFR conference—8 sessions
- Transcoding for G.711 to G.729a/G.729ab/GSMFR—96 sessions
- Transcoding for G.711 to G.729/G.729b/GSM EFR—72 sessions

Per 3825/3845 (4 PVDM2-64)

- G.711 only conference—50 sessions
- G.729, G.729a, G.729ab, G.729b conference—32 sessions
- GSM FR conference—14 sessions
- GSM EFR conference—10 sessions
- Transcoding for G.711 to G.729a/G.729ab/GSMFR—128 sessions
- Transcoding for G.711 to G.729/G.729b/GSM EFR—96 sessions

**Tip**

Maximum number of participants per conference equals 8.

Where to Find More Information**Related Topics**

- [Transcoders, page 25-1](#)
- [Conference Bridges, page 24-1](#)
- [Media Termination Points, page 27-1](#)

Additional Cisco Documentation

- *Cisco CallManager Administration Guide*