



Cisco Catalyst DSP Resources for Transcoding and Conferencing

This chapter describes Cisco Catalyst digital signal processor (DSP) resources and how they are used for transcoding and conferencing. The modules include the WS-X4604-GWY for the Cisco Catalyst 4000 and the WS-X6608-T1 (WS-X6608-E1 for countries outside the USA) for the Cisco Catalyst 6000. The modules, available for use with Cisco CallManager, can perform conferencing and Media Termination Point (MTP) transcoding services in addition to serving as a PSTN gateway.

This chapter covers the following topics:

- [Understanding Cisco Catalyst DSP Resources, page 21-2](#)
- [DSP Resource Manager, page 21-4](#)
- [Cisco Catalyst Conferencing Services, page 21-11](#)
- [Cisco Catalyst 4000 Voice Services, page 21-13](#)
- [Cisco Catalyst 6000 Voice Services, page 21-15](#)
- [Requirements and System Limits, page 21-17](#)
- [Where to Find More Information, page 21-18](#)

Understanding Cisco Catalyst DSP Resources

The DSP resources on the Cisco Catalyst 4000 and 6000 gateway modules essentially provide hardware support for IP telephony features offered by Cisco CallManager. These features include hardware-enabled voice conferencing, hardware-based MTP support for supplementary services, and transcoding services.

Catalyst-enabled conferencing designates the ability to support voice conferences in hardware. DSPs convert voice over IP (VoIP) sessions into time-division-multiplexing (TDM) streams, which can then be mixed into a multiparty conference call.

The Cisco Catalyst 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 using gateways and clients that do not support the H.323v2 feature of OpenLogicalChannel and CloseLogicalChannel with the EmptyCapabilitiesSet. MTP, available as a software feature, can run on Cisco CallManager or a separate Windows NT server. When MTP is running in software on Cisco CallManager, the resource supports 24 MTP sessions. When MTP is running on a separate Windows NT server, the resource supports up to 48 MTP sessions. The Cisco Catalyst gateway modules can support this same functionality, but they provide the service in the hardware.

Transcoding in effect provides an IP-to-IP voice gateway service. A transcoding node can convert a G.711 voice stream into a low-bit-rate (LBR), compressed voice stream, such as G.729a. This process proves critical for enabling applications such as integrated voice response (IVR), voice messaging, and conference calls over low-speed IP WANs. Cisco CallManager currently supports MTP transcoding only on the Cisco Catalyst voice gateways.

[Table 21-1](#) shows DSP resources that can be configured on the Cisco Catalyst voice services modules.

Table 21-1 Cisco Catalyst DSP Resource Matrix

Cisco Catalyst Voice Modules	PSTN Gateway Sessions	Conferencing Sessions	MTP Transcoding Sessions
Catalyst 4000 WS-X4604-GWY	G.711 only <ul style="list-style-type: none"> • 96 calls 	G.711 only <ul style="list-style-type: none"> • 24 conference participants • Maximum of 4 conferences of 6 participants each 	To G.711 <ul style="list-style-type: none"> • 16 MTP transcoding sessions
Catalyst 6000 WS-6608-T1 or WS-6608-E1	WS-6608-T1 <ul style="list-style-type: none"> • 24 calls per physical DS1 port • 192 calls per module WS-6608-E1 <ul style="list-style-type: none"> • 30 calls per physical DS1 port • 240 calls per module 	G.711 or G.723 <ul style="list-style-type: none"> • 32 conferencing participants per physical port • Maximum conference size of 6 participants • 256 conference participants per module G.729 <ul style="list-style-type: none"> • 24 conferencing participants per physical port • Maximum conference size of 6 participants • 192 conference participants per module 	The following capacities also apply to simultaneous transcoding and conferencing: G.723 to G.711 <ul style="list-style-type: none"> • 31 MTP transcoding sessions per physical port • 248 sessions per module G.729 to G.711 <ul style="list-style-type: none"> • 24 MTP transcoding sessions per physical port • 192 sessions per module

DSP Resource Manager

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 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.

Each Cisco CallManager must have its own conference and MTP transcoder resources, because DSP resources can register with only one Cisco CallManager at a time. When a request is received from the signaling layers for a session, the first available DSP from the respective pool (transcoding or conferencing) is assigned along with the first available channel. DSPRM maintains a set of MAX limits (such as maximum conference sessions per DSP or maximum transcoding session per DSP) for each DSP.

Call Preservation

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 switchback to the higher order Cisco CallManager.

When a switchover or switchback registration occurs, the DSPRM reports the number of active connections along with the number of RTP streams. Once the call disconnects, the DSPRM detects the media streaming failure by reacting to the Internet Control Message Protocol (ICMP) “port not reachable” indications received on the RTP port and reports this information to the currently active Cisco CallManager using SCCPs “StationMediaStreamingFailure” message. The Cisco CallManager can update the status of the specified resource that becomes available for use in a new connection establishment. Disconnect supervision describes this method.

Catalyst MTP Transcoding Services

Introducing the WAN into an IP telephony implementation forces the issue of voice compression. Once a WAN-enabled network is implemented, voice compression between sites represents the recommended design choice in order 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.

MTP Transcoding Design Details

The following points summarize the design capabilities and requirements of the MTP transcoding:

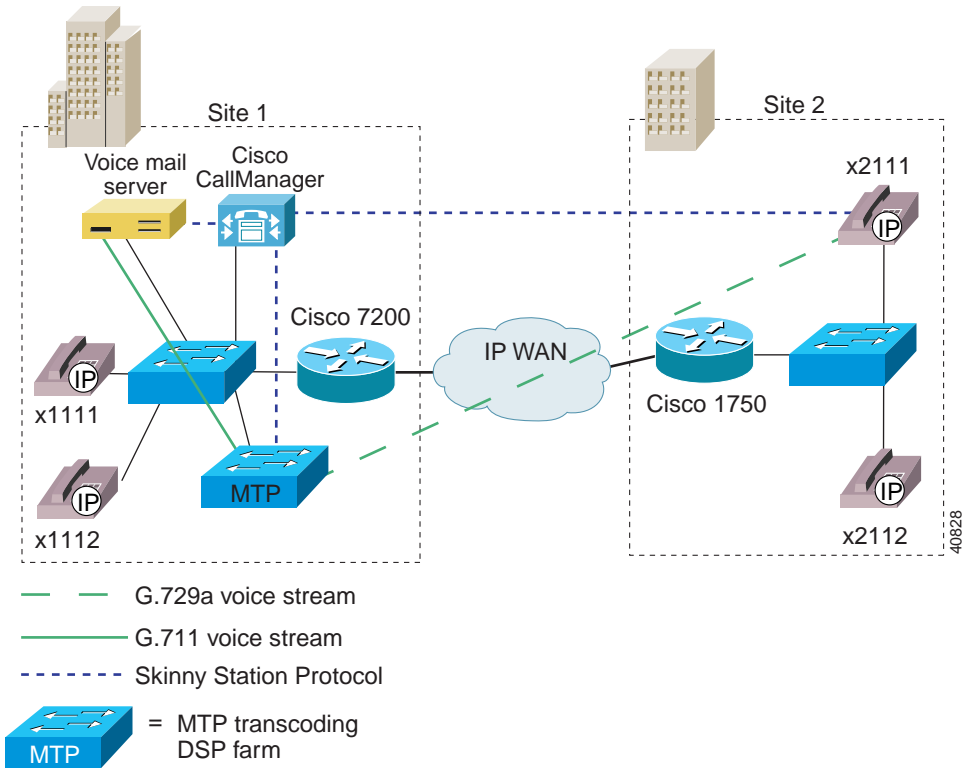
- Provision MTP transcoding resources appropriately for the number of IP WAN callers to G.711 endpoints.
- The Cisco Catalyst 4000 WS-X4604-GWY module supports 16 transcoding sessions per module.
- The Cisco Catalyst 6000 WS-X6608-T1 or WS-X6608-E1 modules support 31 G.723 or G.711 transcoding sessions per physical port (248 per module) or 24 G.729 to G.711 transcoding sessions per physical port (192 per module).
- Support transcoding only in low bit rate to high bit rate (G.729a or G.723.1 to G.711), or vice versa, configurations.

- Each Cisco CallManager must have its own MTP transcoding resources.
- Each transcode has its own jitter buffer of 20-40 ms.

IP-to-IP Packet Transcoding and Voice Compression

You can easily 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 two of the new 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 using different compression algorithms. That is, when 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. However, if the called party at the central site is unavailable, the call potentially rolls to an application that supports G.711 only. In this case, a packet-to-packet gateway transcodes the G.729a voice stream to G.711 to leave a message with the voice-messaging server. See [Figure 21-1](#).

Figure 21-1 IP-to-IP Packet Gateway Transcoding for the WAN with Centralized Call Processing

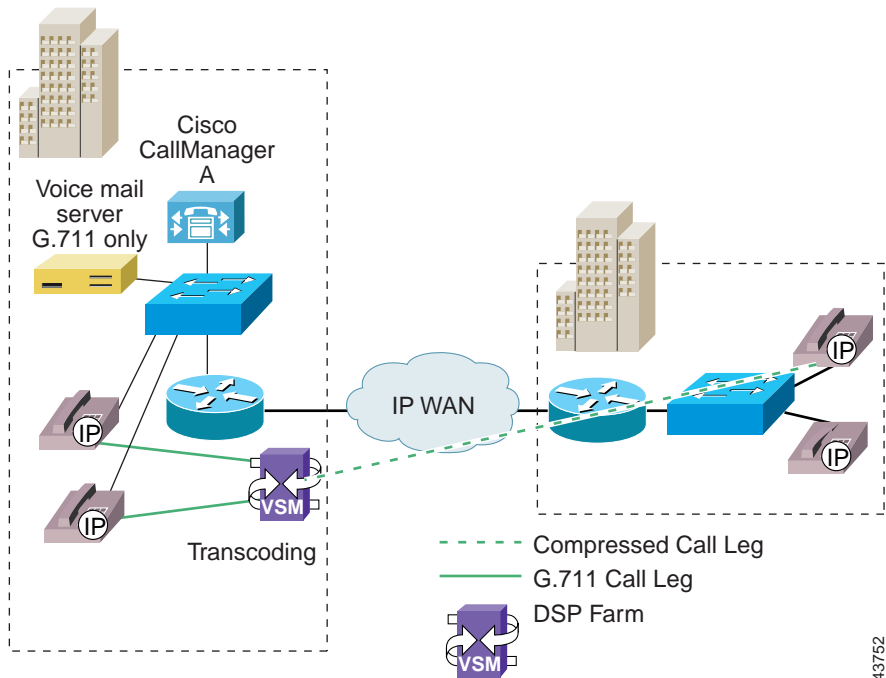


Voice Compression, IP-to-IP Packet Transcoding, and Conferencing

Connecting sites across an IP WAN for conference calls presents a complex scenario. In this scenario, the Cisco Catalyst modules must perform the conferencing service as well as the IP-to-IP transcoding service to uncompress the WAN IP voice connection. In [Figure 21-2](#), 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 are sent to single logical conferencing port where all transcoding and summing takes place.

Figure 21-2 Multisite WAN Using Centralized MTP Transcoding and Conferencing Services




43752

IP-to-IP Packet Transcoding Across Intercluster Trunks

H.323v2 intercluster trunks connect Cisco CallManager clusters. If transcoding services are needed between clusters, configure the intercluster trunks with MTP. In this case, all calls between clusters route through the MTP/transcoding devices in each cluster. The Cisco Catalyst 6000 module uses the MTP service regardless whether transcoding is needed for that particular intercluster call.

Cisco CallManager supports compressed voice call connection through the MTP service if a hardware MTP is used. [Figure 21-3](#) shows an intercluster call flow.

Figure 21-3 Intercluster Call Flow with Transcoding



43753

The following list gives intercluster MTP/transcoding details:

- If transcoding is required between Cisco CallManager clusters, you must configure the H.323 intercluster trunk with an MTP resource.
- All calls between Cisco CallManager clusters go through MTPs.
- 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 H.323 intercluster trunks

Cisco Catalyst Conferencing Services

To scale IP telephony systems in large enterprise environments, you must use hardware-based conferencing. The hardware for the Cisco Catalyst 4000 and 6000 switch families keeps this requirement in mind. The Cisco Catalyst voice modules can handle conferencing in hardware, eliminating the requirement of running a software conferencing service on a Windows NT server in the IP telephony network.

Conferencing Design Details

The following points summarize the design capabilities and requirements of the Cisco Catalyst voice modules:

- Support exists for a maximum of six participants per conference call.
- The Cisco Catalyst 4000 WS-X4604-GWY module supports 24 conference participants per module.
- The Cisco Catalyst 4000 WS-X4604-GWY module supports conferencing for G.711 voice streams only. Transcoding can convert G.729a or G.723.1 to G.711 for conference calls.
- The Cisco Catalyst 6000 WS-X6608-T1 or WS-X6608-E1 modules support 32 G.711 or G.723 conference participants per physical port (256 per module) or 24 G.729 conference participants per physical port (192 per module).

- The Cisco Catalyst 6000 WS-X6608-T1 or WS-X6608-E1 modules can support both uncompressed and compressed VoIP conference calls.
- Make sure each Cisco CallManager has its own conference and MTP transcoding resources because the DSP resources can register with only one Cisco CallManager at a time. Cisco CallManagers cannot share DSP resources.

The Cisco Catalyst 4000 module, the WS-X4604-GWY, can support up to four simultaneous conference calls of six callers each. The Cisco Catalyst 6000 T1 or E1 PSTN gateway module, the WS-X6608, also can support conferencing. After the WS-X6608 is added as a T1 or E1 Cisco AVVID gateway, you can configure it, on a per-port basis, for conferencing services. The Cisco Catalyst 6000 conferencing module supports up to six callers per conference call with a maximum of 32 simultaneous G.711 or G.723 conference callers per configured logical port. This configuration results in a maximum of 256 conference participants per module with G.711 or G.723 calls.

See [Table 21-1](#) for a summary of conference call densities for each module.

Both the WS-X4604-GWY and WS-X6608-T1 (or WS-X6608-E1) modules use Skinny Station Protocol to communicate with Cisco CallManager when providing conferencing or transcoding services. The Cisco Catalyst 6000 voice-conferencing solution can support both compressed and uncompressed conference attendees.

On the Cisco Catalyst 4000, only G.711, or uncompressed, calls can join to a conference call. When the conferencing service registers with Cisco CallManager, using Skinny Station Protocol, it announces that only G.711 voice calls can connect. If any compressed calls request to be joined to a conference call, Cisco CallManager connects them to a transcoding port first, to convert the compressed voice call to G.711. Once the G.711 connections are associated with a particular conferencing session (maximum of six participants per conference call), the call converts to a TDM stream and passes to the summing logic, which combines the streams.


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.

You should also observe the following recommendations when configuring conferencing services:

- When provisioning an enterprise with conference ports, first determine how many callers will attempt to join the conference calls from a compressed Cisco CallManager region. Once 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).

Figure 21-4 illustrates the components used in Cisco Catalyst conferencing services.

Figure 21-4 Cisco Catalyst Conferencing Services



Cisco Catalyst 4000 Voice Services

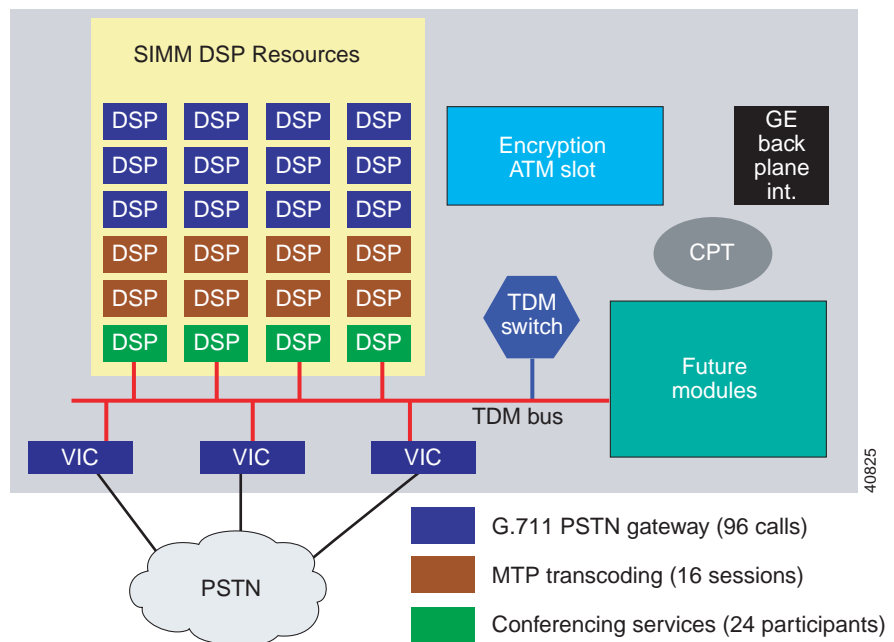
The PSTN gateway and voice services module for the Cisco Catalyst 4003 and 4006 switches supports three analog voice interface cards (VICs) with two ports each or one T1/E1 card with two ports and two analog VICs. Provisioning choices for the VIC interfaces any combination of Foreign Exchange Office (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. Once the gateway mode is enabled, the 24 DSPs (4 SIMMs with 6 DSPs each) for the module occur as follows:

- PSTN gateway: 96 channels of G.711 voice and
- Conferencing: 24 channels of G.711 conferencing and
- MTP transcoding: 16 channels of LBR-G.711 transcoding

Figure 21-5 shows a physical representation of the Cisco Catalyst 4000 voice gateway module in gateway mode.

Figure 21-5 Cisco Catalyst Voice Gateway Module in Gateway Mode



Gateway mode designates the default configuration. You can change the conferencing-to-transcoding ratios from the CLI.

The following configuration notes apply 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. For all PSTN gateway functions, the Cisco Catalyst 4000 module uses H.323v2 and is configured identically to a Cisco IOS gateway. From the Cisco CallManager configuration window, add the Cisco Catalyst 4000 gateway as an H.323 gateway.
- 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 (SCCP).
- You can define a primary, secondary, and tertiary Cisco CallManager for both the conferencing services and MTP transcoding services.

Cisco Catalyst 6000 Voice Services

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 PRI interfaces, each of which has its own CPU and DSPs. Once the card is added from Cisco CallManager as a voice gateway, you configure it as a conferencing or MTP transcoding node. 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 node for each conferencing node.


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 provided by DHCP. 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. Once configured through the Cisco CallManager interface, each port can support one of the following configurations:

- PSTN gateway mode: 24 sessions on the WS-6608-T1 module; 30 sessions on the WS-6608-E1
- Conferencing mode: 32 conferencing sessions for G.711 or G.723; 24 conferencing sessions for G.729
- MTP mode: 31 MTP transcoding sessions for G.723 to G.711; 24 MTP transcoding sessions for G.729 to G.711

Figure 21-6 shows one possible configuration of the Cisco Catalyst 6000 voice gateway module. This diagram shows two of the module's eight ports configured in PSTN gateway mode, three ports in conferencing mode, and three ports in MTP transcoding mode.

Figure 21-6 Cisco Catalyst 6000 Voice Gateway Module



Requirements and System Limits

The following sections describe the Cisco Catalyst DSP resources requirements and system limits for transcoding and conferencing.

MTP Transcoding Caveats

The following summary gives caveats that apply to Cisco Catalyst MTP transcoding:

- Cisco Catalyst MTP transcoding service only supports LBR codec-to-G.711 conversion and vice versa. No support exists for LBR-to-LBR codec conversion.
- On the Cisco Catalyst 6000, transcoding services cannot cross port boundaries.
- Make sure each Cisco CallManager has its own MTP transcoding resource configured.
- If transcoding is required between Cisco CallManager clusters, make sure the H.323 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 heard.

See [Table 21-1](#) for a list of transcoding capabilities for each module.

Conferencing Caveats

The following caveats apply to Cisco Catalyst conferencing services:

- The Cisco Catalyst 4000 conferencing services support G.711 connections only, unless an MTP transcoding service is used.
- On the Cisco Catalyst 6000, conferencing services cannot cross port boundaries.
- Each Cisco CallManager must have its own conferencing resource configured.

The “[DSP Resource Manager](#)” section on page 21-4 discusses conference calls across an IP WAN.

Where to Find More Information

Related Topics

- [Transcoders, page 18-1](#)
- [Conference Bridges, page 17-1](#)

Additional Cisco Documentation

- *Cisco CallManager Administration Guide*
- *Cisco IP Phone Administration Guide for Cisco CallManager*
- Cisco IP Phone user documentation and release notes (all models)