




Cisco Digital Media Suite 5.2 for Enterprise Medianet

DMS Overview

The Cisco Digital Media Suite (DMS) is a comprehensive suite comprised of Cisco Digital Signs, Cisco Cast, and Cisco Show and Share applications that allow companies to use digital media to increase sales, enhance customer experience, and facilitate learning. (See [Figure 1](#).)


Figure 1 *Cisco Medianet Application Components*



Cisco DMS is split into three distinct functional subsystems: Cisco Digital Signs, Cisco Cast, and Cisco Show and Share (see [Figure 2](#)).



Americas Headquarters:
Cisco Systems, Inc., 170 West Tasman Drive, San Jose, CA 95134-1706 USA

Figure 2 **DMS Functional Subsystems**

Cisco Digital Signs

The Cisco Digital Signs subsystem is centrally controlled. Media and messages are sent to the displays with minimal control at the endpoint, with optional touchscreens. Cisco Digital Media Manager (DMM) administrators control what is shown and when it is shown. Flash application content, messaging, and live or pre-recorded video can be displayed. Video can be displayed in a subset of the screen with Flash content around it, allowing for messaging and images to be placed around the borders of the screen.

Cisco Cast

With the Cisco Cast subsystem, the control shifts to the end user via a remote control; and with the latest release of IP phones and smartphones, touchscreens can be used to control content delivered. Cast has three sections: one to access VoDs, one to scroll through live channels, and a channel guide. The end user selects what is displayed, and the DMM administrators control what is available to be selected for viewing on the Cisco Cast displays. The displays are driven by a Cisco Digital Media Player (DMP).

Cisco Show and Share

The Cisco Show and Share subsystem is a social media system similar to YouTube, but for enterprise corporate use. Using a web interface, the end user can upload media, such as videos, and even generate content from a webcam directly into the Cisco Show and Share system. Show and Share can also be used to broadcast live content to any corporate user via the web interface. End users in any location, remote or local, are able to view live and pre-recorded content through a standard web browser.

Cisco Show and Share is a separate subsystem, but can be integrated with Cisco Cast and Cisco Digital Signs content by using the Cisco Media Experience Engine, discussed later in this document.

For more detailed information about specific Cisco DMS products, see the following URL:

<http://www.cisco.com/web/solutions/dms>.

Cisco DMS 5.1 to 5.2 Name Changes

Some Cisco DMS 5.1 names have been changed in DMS 5.2. [Table 1](#) shows both the old and updated names. This document uses DMS 5.2 names.

Table 1 *Cisco DMS 5.1 to 5.2 Name Changes*

Cisco DMS 5.1 Name	Cisco DMS 5.2 Name
Cisco Digital Media System	Cisco Digital Media Suite
Cisco Video Portal	Cisco Show and Share
Cisco Digital Signage	Cisco Digital Signs
Cisco Enterprise TV	Cisco Cast

Benefits of Using Cisco DMS

Cisco DMS is now the most compelling platform to effectively reach customers, employees, partners, and students with important information, news, training, and events.

Cisco Digital Signs

Cisco Digital Signs provides scalable, centralized management and publishing of high-quality content to networked, on-premise digital signs.

Increasingly, financial services organizations, retail stores, and educational institutions use Cisco DMS for digital signs. Additional examples of industry applications include the following:

- Sports and entertainment—High-definition event broadcasts, live streaming statistics, sales and marketing of products and services, and directional information are provided on digital signs and video walls throughout event venues, and in fan lounges and suites.
- Government—Digital signs provide useful information for people waiting in line at government offices to help speed transactions.
- Healthcare—Relevant healthcare information is provided through digital signs around hospitals, offering cost-effective training options for hospital personnel.

Cisco Cast

Cisco Cast allows organizations to deliver live and pre-recorded content, controlled by the end user. Content can include news, financial information, sales and marketing, learning/training, corporate communications, and entertainment videos.

On-screen menus and program guides give users access to Cast content, and organizations can customize lineups, as well as create their own content libraries. Users navigate through channel menus, selecting from live or on-demand content with a remote control, smart phone, or touchscreen.

Financial services organizations, retail stores, and educational institutions are using DMS for Cast, such as in the following settings:

- Executive offices—For delivery of live broadcast business channels over IP such as CNBC, MSNBC, and Bloomberg.

- Employee gathering areas—For delivery of executive broadcasts or employee communications videos, and infotainment such as the Weather Channel, CNN, ESPN, or internal TV channels.
- Customer-facing lobbies or waiting areas—For delivery of sales and marketing videos and infotainment such as the Weather Channel, CNN, and ESPN.

Cisco Show and Share

Cisco Show and Share gives customers, employees, partners, and students access to high-quality and compelling videos-on-demand (VoDs) and live webcasts at their desktops. Users can also upload video, edit video, and even record directly from a webcam into Cisco Show and Share. Digital media can be browsed, searched, and viewed over the network through a unique, easy-to-use Cisco Show and Share experience: anywhere, anytime.

Industries using Cisco Show and Share include the following:

- Government—Providing live and on-demand web-based access to city council meetings or events, and access to digital media-based information about relevant regulations and laws.
- Healthcare—Alleviating staff and resource shortages by providing patients, families, and friends with digital media-based “what to expect” patient material.
- Education—Providing students with a rich learning experience, with both live broadcasts and pre-recorded lectures, whether they are local or remote.

New Features in Cisco DMS 5.1 and 5.2

Moving from DMS Release 5.1 to 5.2 includes numerous new features as well as the introduction of an entirely new DMS subsystem, Cisco Show and Share, which replaces its predecessor, Cisco Video Portal. This section discusses the DMS features added in the 5.1 and 5.2 releases.

Cisco Digital Signs and Cisco Cast—New Features in 5.1

Cisco Digital Signs and Cisco Cast include the following new features in DMS Release 5.1:

- Support for Flash 10 applications with Cisco DMP 4400G
- Support for H.264 video with Cisco DMP 4400G
- Support for USB external storage devices attached to the Cisco DMP
- Addition of IP phone- and smartphone-based remote control for Cisco Cast
- Addition of scheduled content distribution to Cisco DMPs
- Common Internet File System (CIFS) support for integration with Cisco Wide Area Application Services (Cisco WAAS)

Cisco Digital Signs and Cisco Cast—New Features in 5.2

Cisco Digital Signs and Cisco Cast include the following new features in DMS Release 5.2:

- Touchscreen support for Cast and Digital Signs navigation
- Proof of play for Digital Signs
- Windows Media 9/VC-1 support for VoD content playback on the Cisco DMP 4400G
- MP3 audio support

- Real Time Streaming Protocol using TCP (RTSP-T) support for VoD content

Cisco Show and Share—New Features

Cisco Show and Share is a replacement for Cisco Video Portal. Video Portal was a structured one way media delivery mechanism, allowing users to select content but not contribute content. Show and Share is a social media system allowing uploading and editing of video content as well as direct webcam recording and editing.

Cisco DMS 5.2—Technical Details

This section provides more technical details of Cisco DMS 5.2, and includes the following topics:

- Cisco Digital Signs and Cisco Cast
- Cisco Show and Share


Cisco Digital Signs and Cisco Cast

Cisco Digital Signs and Cisco Cast use the same components with little exception and are grouped together throughout this document. However, actual implementation can have significant differences and requirements, from the network perspective.

Cisco Digital Signs and Cisco Cast Components

Figure 3 shows the core components of Digital Signs and Cast, and their normal locations within the enterprise network.

Figure 3 *Digital Signs and Cast Components*



The core components of Digital Signs and Cast are as follows:

- **Digital Media Manager (DMM)**
Controls and communicates with all critical Digital Signs/Cast components. All DMPs are managed directly from the DMM. Content distribution is mainly managed from the DMM, with some exceptions. For Digital Signs, scheduling and display of content through DMPs is managed through the DMM. For Cast, the program interface, live channel list, VoD list, and program guide data are managed through the DMM. The DMM is also capable of DMP maintenance, such as firmware upgrades.
- **Scientific Atlanta encoder**
Encodes live video input into an SD or HD, MPEG-2 or MPEG-4 multicast stream. The live input is normally a feed from a cable or satellite content provider, but may also be from other sources such as live cameras or DVD players. Scientific Atlanta encoders provide extremely high quality live stream encoding into multicast streams placed on the network. No recording or archive capabilities exist on the Scientific Atlanta encoders themselves.
- **Web server/content repository**
The web server holds all VoDs referenced by the DMM. All VoD streaming requests issued to the DMP are serviced from this server or servers. The web server may be a front-end for a content repository such as a SAN. The web server can be any standard web server such as Internet Information Services (IIS) on Windows-based servers and Apache on Linux-based servers. The content repository may also supply content through a CIFS share.
- **Digital Media Player (DMP)**
DMPs decode and display unicast VoDs and multicast live streamed video as well as Flash content. DMPs connect directly to large format displays through HDMI. Other output connections are available but normally not used. DMPs are controlled by the DMM. Content displayed with Digital Signs is controlled centrally through the DMM. Content displayed with Cast is controlled by end users through an infrared remote, IP phone interface, or web-enabled smartphone.
- **User control device for DMPs (not shown)**
Handheld infrared remote, IP phone-based control, or web interface for mobile devices may be used to control Cast content displayed in real-time.
- **Cisco Wide-Area Application Engines (WAEs)**
WAEs are not DMS core components but are used with Cisco Application and Content Networking System (ACNS) for content distribution and WAN optimization.



Note For more information about ACNS, see [Cisco Application and Content Networking System, page 33](#). For more information about Cisco Wide Area Application Services (WAAS), see [Cisco Wide Area Application Services, page 38](#).

Cisco Digital Signs and Cisco Cast Connections

This section discusses the connections between components, including protocols and content traversing those connections. Connections shown in [Figure 4](#) do not include content delivery network (CDN) components normally deployed with DMS. CDNs are covered later in this document.

Figure 4 Cisco Digital Signs and Enterprise Network Connections




Figure 4 shows significant connections and excludes some minor connections to reduce the complexity of the diagram. The three main connection/content types are as follows:

- **Control traffic**—Connections that contain control information and any connection that does not transport live or pre-recorded content.
- **Pre-recorded content**—Connections that transport pre-recorded content, including indirect content distribution and direct delivery to end devices.
- **Live content**—All live streaming content, including pre-recorded content streamed live through an encoder. This content uses multicast UDP for Digital Signs and Cast.

The following describes the numbered connections shown in Figure 4:

1. DMM to DMP

The connection between the DMM and the DMP is an HTTP connection used for delivering instructions to the DMP from the DMM. The DMP also pulls a Flash application for the Cast interface from the DMM.

2. Web server/content repository to the DMP

The connection between the web server and DMP is normally an HTTP connection to deliver pre-recorded content to the DMP. Other connections such as FTP and CIFS may be incorporated directly or indirectly through a Content Deliver Network and are described in [Application Network Services](#), page 33.

3. Scientific Atlanta encoder to DMP

The connection between the Scientific Atlanta encoder and the DMP is technically not a connection because it is connectionless UDP multicast traffic. Details of multicast operations with DMS is discussed in the [Multicast Design Considerations for Cisco DMS, page 24](#).

Cisco Digital Media Player Specifications

[Table 2](#) lists the capabilities and differences between the two Cisco DMP models: the 4305G and 4400G with Cisco DMS Release 5.2.

Table 2 *Cisco DMP Models—4305G and 4400G*

	4305G	4400G
Jitter Buffer	• 4 MB • 1500 ms for multicast	• 5.5 MB • 1500 ms for multicast
Multicast Support	IGMP v2	IGMP v2
Video Support	• MPEG-2 • MPEG-4 Part 2	• MPEG-2 • MPEG-4 Part 10 H.264 • WM9/VC-1 (VoD only)
Bandwidth Required	• MPEG-2 • SD—3 to 5 Mbps • HD—13 to 25 Mbps	• H.264/WM9 • SD—1.5 to 5 Mbps • HD—8 to 25 Mbps
Flash Application Support	Flash 7	Flash 10

Jitter Buffer

The jitter buffers in the DMPs are sufficient to deal with even extreme cases of jitter for live streams. The only reasonable scenario for failures resulting from exceeding the jitter buffer is when the jitter from streaming HD VoDs exceeds 1000 ms. A properly designed network should not allow this threshold to be exceeded.

Multicast Support

The DMPs join multicast MPEG-2 and H.264 streams as the only method of displaying live streaming video. DMPs support Internet Group Management Protocol (IGMP) v2, although a multicast source cannot be defined when defining multicast streams or *channels* within the Cast interface. This means that source-specific multicast cannot be fully implemented directly from the DMPs, and that all multicast join messages are sent as (*,g) messages. Source specific multicast may be implemented using methods described in the multicast section of this document.

Video Support

The following are common video formats:

- MPEG-4 Part 2

MPEG-4 Part 2 is a video compression technology developed by the Moving Picture Experts Group (MPEG). It belongs to the MPEG-4 ISO/IEC standard (ISO/IEC 14496-2). It is a discrete cosine transform compression standard, similar to previous standards such as MPEG-1 and MPEG-2. Several popular codecs, including DivX, Xvid, and Nero Digital, are implementations of this standard.

- MPEG-4 Part 10 (H.264)

H.264 is a standard for video compression, and is equivalent to MPEG-4 Part 10 or MPEG-4 for advanced video coding (AVC). As of 2008, this is the latest block-oriented motion-compensation-based codec standard developed by the ITU-T Video Coding Experts Group (VCEG) together with the ISO/IEC MPEG, and it was the product of a partnership effort known as the Joint Video Team (JVT). The ITU-T H.264 standard and the ISO/IEC MPEG-4 Part 10 standard (formally, ISO/IEC 14496-10) are jointly maintained so that they have identical technical content. The final drafting work on the first version of the standard was completed in May 2003.

- WMV9/VC-1

Windows Media Video 9 (WMV9) is a common Windows media format now supported for VoD playback only. WMV9 supports variable bit rate, average bit rate, and constant bit rate, as well as several important features including native support for interlaced video, non-square pixels, and frame interpolation. Unlike previous versions of WMV, WMV9 is more capable of efficient delivery of high-definition video content, at resolutions such as 720p and 1080p.

These formats are supported by the various Cisco DMPs as follows:

- Cisco DMP 4305G—Supports standard MPEG-2 streams (HD or SD) as well as the rarely used MPEG-4 Part 2. DMP 4305G does not support H.264.
- Cisco DMP 4400G—Supports standard MPEG-2 streams (HD or SD) as well as MPEG-4 Part 10, also known as H.264. WMV9 is also supported for VoD playback only.

The bandwidth requirements range between 1.5 and 25 Mbps, depending on several factors including whether the video is SD or HD and what codec is used. Bandwidth requirements are discussed in detail later in this document.

Flash Application Support

Flash applications are used with Cisco Digital Signs to display content and with Cisco Cast to display the navigation interface. The DMP 4400G introduces Flash 10 support, while the 4305G is limited to Flash 7.


Cisco Show and Share

Cisco Show and Share allows the delivery of live or pre-recorded content to desktop or laptop PCs. The following section details the Show and Share components and connections between them.

Cisco Show and Share Components

Figure 5 introduces the Cisco Show and Share components and their normal locations within the enterprise network.

Figure 5 Cisco Show and Share Components



The core Cisco Show and Share components are as follows:

- **Digital Media Manager (DMM)**
DMM controls and communicates with all critical Show and Share components. Content distribution is mainly controlled from the DMM, with some exceptions. The DMM is the central control and configuration point for the DMEs and Show and Share.
- **Digital Media Encoders (DME)**
The DME registers with the DMM and broadcasts live video to the streaming server using RTSP. DMEs are also capable of recording video that may be archived to the content repository as VoD content.
- **Web server/content repository**
The web server holds all VoDs referenced by the DMM. All VoD streaming requests issued to the DMP are serviced from this server. The web server may be a front-end for a content repository such as a SAN. The web server can be any standard web server, such as IIS on Windows-based servers and Apache on Linux-based servers. The content repository may also supply content through a CIFS share.
- **Show and Share server**
The Show and Share Server provides the web-based interface for clients. All navigation and authentication is completed through this server. Media is not normally stored on the Show and Share server.
- **Streaming server**
A streaming server is required for delivery of live content. Streaming servers provide stream splitting capabilities, allowing many clients to view a single live stream from a DME or pre-recorded source (live rebroadcast). Cisco ACNS has the option to enable stream splitting for live content and may be used as the streaming server. Windows Server 2003/2008 may be used with Windows Media Streaming Server enabled as the streaming server.

- Show and Share client

A Show and Share client is any desktop or laptop PC with a compatible browser that can access the Show and Share server.

Cisco Show and Share Connections

This section discusses the connections between components, including protocols and content traversing those connections. Figure 6 shows the significant connections, excluding some minor connections to reduce the complexity of the diagram.

Figure 6 Cisco Show and Share Connections




Figure 6 shows the following main connection/content types:

- Control traffic—Connections that contain control information and any connection that does not transport live or pre-recorded content.
- Pre-recorded content—Connections that transport pre-recorded content, including indirect content distribution and direct delivery to end devices.
- Live content—All live streaming content, including pre-recorded content streamed live through an encoder. This content is unicast for Show and Share with one exception. Cisco ACNS may be used to convert a unicast RTSP stream into multicast traffic.

The following describes the numbered connections shown in [Figure 6](#):

1. DME to DME
Between the DME and the DMM, HTTP is used to upload pre-recorded content from the DME.
2. DME to DMM
Between the DME and DMM, basic control traffic exists with the DME registering itself to the DMM.
3. DME to streaming server
DME sends all live streams to the streaming server through HTTP.
4. Web server/content repository to DMM
VoDs uploaded to the DMM are automatically pushed to the web server/content repository through FTP or SFTP.
5. DMM to Show and Share server
Configuration of the Show and Share server is pushed from the DMM. All configuration of the Show and Share server is managed through the DMM.
6. Web server/content repository to Show and Share client
VoD media is sent to the Show and Share client directly from the web server/content repository through HTTP.
7. Streaming server to Show and Share client
The streaming server sends live video streams directly to the Show and Share clients through RTSP-T, RTSP-U, or multicast UDP. Cisco ACNS or Cisco WAAS may be used in place of or in conjunction with the streaming server. Various stream splitting scenarios are discussed in [Application Network Services, page 33](#).
8. Show and Share server to Show and Share client
The web interface is provided by the Show and Share server to the Show and Share client through HTTP. No media is supplied by the Show and Share server.
9. Show and Share client to Show and Share server
The Show and Share client uploads media to the Show and Share server through HTTP upload. Both VoDs and live recording via the Show and Share client interface are transferred to the Show and Share server.
10. Show and Share server to web server/repository
The Show and Share server transfers media to the web server/content repository.

Network Requirements

This section contains the following topics:

- [Bandwidth Requirements](#)
- [Latency Requirements](#)
- [Jitter Requirements](#)
- [Packet Loss Requirements](#)
- [Tuning TCP Parameters](#)
- [Quality of Service](#)
- [Network Requirements Summary](#)

Bandwidth Requirements

Bandwidth requirements with DMS vary greatly, from 200 Kbps to 25 Mbps. The following sections discuss the bandwidth requirements for Cisco Digital Signs, Cisco Cast, and Cisco Show and Share.

Cisco Digital Signs and Cisco Cast

Cisco Digital Signs and Cisco Cast use the same hardware and therefore have the same bandwidth requirements for individual video streams. However, overall bandwidth usage differs because of different usage models for Digital Signs and Cast.

Digital Signs and Cast use the DMP to deliver live and pre-recorded streaming content to displays. Bandwidth used per stream is 1.5–5 Mbps for standard definition streaming video content and 8 to 25 Mbps for high definition streaming video content.

Digital Signs and Cast usage models differ in many ways, but the most important is in how the content is controlled. Digital Signs content is managed and controlled centrally, allowing a very predictable environment for content delivery. Cast content, while managed centrally, is controlled by the end user, using a remote, touchscreen, or smartphone to control the display of pre-recorded or live content in real time. With many Cast users independently controlling many individual endpoints, content delivery and bandwidth requirements may be significantly higher than requirements for Digital Signs.

With Cisco Digital Signs, streaming video content may be placed on a portion of the screen, with the remaining screen being used by Flash or media content such as information tickers, advertisements, images, or any other non-streaming content supported by the DMPs. Video resolution can be reduced for partial coverage of the screen. Reducing displayed video resolution allows the reduction of encoded stream resolution, lowering the bandwidth requirements.

With Cisco Cast, an additional function that may impact bandwidth requirements is the multicast delivery of live content, or TV channels, to the DMP. The end user has control of the delivery of these channels through the remote control. Switching between channels causes the DMP to start and stop, or join and leave, multicast streams. Rapid switching of channels may cause multiple streams to be started or joined at the same time. This significantly increases bandwidth requirements to the endpoint and may impact other parts of the network as well. Cisco DMS 5.2 has been improved to reduce the impact of channel switching. For more details, see [Multicast Design Considerations for Cisco DMS](#), page 24.

Non-live content, known as VoD, has similar bandwidth requirements when streamed across the network. Remote locations may take advantage of a Content Delivery Network (CDN) for content pre-positioning to reduce network requirements for VoD. CDNs are discussed in [Application Network Services](#), page 33.

Cisco Show and Share

Cisco Show and Share generally requires 200 Kbps to 1.5 Mbps for live streaming video content. Requirements in the future will increase with the introduction of high-definition content.

VoDs have similar bandwidth requirements when streamed across the network. Remote locations may take advantage of a CDN for content pre-positioning to reduce bandwidth requirements for VoDs. CDNs are discussed in [Application Network Services, page 33](#).

Latency Requirements

The following sections discuss latency requirements for Cisco Digital Signs, Cisco Cast, and Cisco Show and Share.

Cisco Digital Signs and Cisco Cast

Cisco Digital Signs and Cisco Cast use the same hardware, and therefore have the same latency requirements.

Live Content

For live streaming content, moderate latency does not have a significant impact, with one exception. The Cisco Cast user may experience significant latency because of the delay when navigating the Cast interface, especially when changing channels of live content. Significant latency is rarely encountered with the large multicast streams sent to the DMPs because they are normally implemented in a campus environment.

Pre-Recorded Content

For pre-recorded video content, moderate latency does occur. Pre-recorded content is streamed through HTTP or RTSP-T, with large bandwidth demands because of the TCP mechanisms for transport. TCP is connection-oriented, unlike UDP, and requires acknowledgement of data sent. This process reduces the throughput maximum as latency increases, regardless of how much bandwidth is available. Adjustments to increase the efficiency of TCP are discussed in [Tuning TCP Parameters, page 20](#).

With TCP parameters set to optimal levels, tolerances for latency are still quite stringent because of the throughput needed. For SD video, latency must be less than 100 ms round trip. For HD video, latency must be less than 60 ms round trip. Delay beyond these thresholds causes the TCP data stream to slow because of the two-way acknowledgement-based communication.

Further TCP optimization can be achieved by using the Cisco WAAS, which implements TCP optimization transparently over WAN links, relaxing latency, loss, and jitter requirements. WAAS increases the latency tolerance for all HTTP and RTSP-T video content to 220 ms.

Cisco Show and Share

Live Content

Live streaming content may be delivered through one of the following three methods, each with different latency requirements:

- RTSP-T (TCP)

RTSP-T is TCP-based and has more stringent latency requirements than UDP-based streaming options. Adjustments to increase the efficiency of TCP are discussed in [Tuning TCP Parameters, page 20](#). With TCP parameters set to optimal levels, latency must be below 400 ms round trip. Use of Cisco WAAS doubles latency tolerances because of the built-in TCP optimization, bringing latency requirements to 800 ms round trip for RTSP-T content.

- RTSP-U (UDP)

RTSP-U is UDP-based and has much higher tolerances for latency than RTSP-T. Latency with RTSP-U may be as high as 2000 ms round trip.

- Multicast

Multicast is UDP-based and has the same latency requirements as RTSP-U, with a maximum of 2000 ms round trip. Though multicast has no feedback mechanism, and therefore has no theoretical limit to the amount of latency that may exist, exceeding 2000 ms is not recommended because of possible issues with join and leave latency for multicast streams.

Pre-Recorded Content

Pre-recorded content may be delivered through one of three methods. HTTP may be used to deliver pre-recorded content from any standard HTTP server. Windows Media Streaming Services may also be used to deliver pre-recorded content through RTSP-T and RTSP-U.

- HTTP and RTSP-T (TCP)

HTTP and RTSP-T are TCP-based and have more stringent latency requirements than UDP-based delivery options. Adjustments to increase the efficiency of TCP are discussed in [Tuning TCP Parameters, page 20](#). With TCP parameters set to optimal levels, latency must be below 400 ms round trip. Use of Cisco WAAS doubles latency tolerances because of the built-in TCP optimization, bringing latency requirements to 800 ms round trip for RTSP-T content.

- RTSP-U (UDP)

RTSP-U is UDP-based and has much higher tolerances for latency than RTSP-T. Latency with RTSP-U may be as high as 2000 ms round trip.

Jitter Requirements

The following sections discuss the jitter requirements for Cisco Digital Signs, Cisco Cast, and Cisco Show and Share.

Cisco Digital Signs and Cisco Cast

Cisco Digital Signs and Cisco Cast use the same hardware, and therefore have the same jitter requirements.

Live Content

For live streaming content, jitter is generally not an issue. Extreme jitter may have an impact if the jitter exceeds the buffer on the Cisco Digital Media Player. The Cisco DMP 4400G has a 1500 ms jitter buffer, which should account for all but the most extreme jitter scenarios, which are unlikely and reflect possible fundamental network problems. The impact of jitter exceeding the jitter buffer is similar to the impact on video quality from lost packets.

Jitter should not be confused with out-of-order packets. Out-of-order packets for live content are treated as lost packets by the DMP and are subject to the same requirements defined in [Packet Loss Requirements](#), page 16.

Pre-Recorded Content

For pre-recorded content, jitter requirements are quite stringent because HTTP is used as the delivery method. Pre-recorded content has strict latency requirements that limit the jitter tolerances to 30 ms.

Use of Cisco WAAS increases jitter tolerances to 60 ms.

Cisco Show and Share

Live Content

Live streaming content may be delivered through one of the following methods, each having different jitter requirements:

- RTSP-T (TCP)—TCP-based delivery has jitter limitations of 100 ms. Use of Cisco WAAS increases jitter tolerances to 200 ms.
- RTSP-U and multicast (UDP)—UDP-based delivery methods have much higher tolerances for jitter, which may reach as high as 400 ms.

Pre-Recorded Content

Pre-recorded content may be delivered through one of the following methods:

- HTTP and RTSP-T (TCP)—TCP-based delivery has jitter limitations of 100 ms. Use of Cisco WAAS increases jitter tolerances to 200 ms.
- RTSP-U (UDP)—UDP-based delivery methods have much higher tolerances for jitter that may reach as high as 400 ms.

Packet Loss Requirements

The following sections discuss the packet loss requirements for Cisco Digital Signs, Cisco Cast, and Cisco Show and Share.

Cisco Digital Signs and Cast

Live Streaming Multicast Content

For live streaming content, lost packets are not re-transmitted, and with the amount of compression used by the video codecs, even a single packet lost results in degraded the video quality. Avoiding any packet loss is the highest priority for live streaming video. With certain configurations, packet loss of 0.001 percent may be considered unacceptable over an extended period of time.

To emphasize the importance of this, the avoidance packet loss is the single most important factor when implementing live video with Cisco Digital Signs and Cisco Cast. Any packet loss may be visible and severely impact the video and audio quality of all DMPs experiencing that packet loss. Some codec and platform combinations are more susceptible to this loss than others. This section discusses three combinations for live video.

The Cisco DMP 4305G is capable of MPEG-2 and MPEG-4 Part 2. Because MPEG-4 Part 2 is rarely used, it was not examined for behavior with loss. The Cisco DMP 4400G is capable of MPEG-2 and MPEG-4 part 10 (H.264). Both combinations for live video are common and discussed below:

- Cisco DMP 4305G with MPEG-2

When MPEG-2 multicast streaming was examined with the DMP 4305G, even minimal loss caused significant issues with quality. Loss requirements for this combination, 4305G displaying MPEG-2, should be considered as 0 percent. This may be an unrealistic goal for some network implementations, but focusing on this goal keeps video quality at the highest level possible.

- Cisco DMP 4400G with MPEG-2

When MPEG-2 multicast streaming was examined with the DMP 4400G, results showed some tolerance for packet loss, unlike the DMP 4305G. Packet loss impacts can still be seen, but consequences are minimal for low packet loss. Packet loss levels of 0.05 percent did not cause any significant degradation of video, and losses of 0.1 percent, while more visible, did not disturb the viewing experience. A target of 0.05 percent packet loss or less is acceptable for this combination of codec and platform.

- Cisco DMP 4400G with MPEG-4 Part 10 (H.264)

When MPEG-4 Part 10 (H.264) multicast streaming was examined with the DMP 4400G, even minimal loss caused significant issues with quality. Loss requirements for this combination, DMP 4400G displaying MPEG-4 Part 10, should be looked at as 0 percent. This may be an unrealistic goal for some network implementations, but focusing on this goal keeps video quality at the highest level possible. Use of H.264 is highly desirable because of the reduced bandwidth requirements compared to MPEG-2, but caution should be taken when implementing on a network with inherent loss issues.

Pre-Recorded Content

For pre-recorded streaming content, packet loss requirements are still highly stringent because of the use of HTTP or RTSP-T to stream content. Packet loss of less than 0.05 percent is necessary for good quality video. This tolerance is increased to 0.1 percent with the use of Cisco WAAS for TCP optimization. Loss beyond these thresholds causes breaks in the audio, and the video periodically freezes. TCP retransmits slow the entire data stream, causing interruptions to video being displayed. At these thresholds, data streams are not delivered fast enough to the DMP, regardless of the speed of the network. Use of a CDN such as Cisco ACNS, Cisco Content Delivery System (CDS), Cisco WAAS, or the new DMS-Content Distribution (DMS-CD) functionality built into DMS 5.2 is highly recommended. These technologies are discussed in detail in [Application Network Services, page 33](#).

**Note**

Pre-recorded streaming content should not be confused with pre-recorded content that is pre-positioned through a CDN. Pre-positioned content, whether video or Flash applications, does not carry the same packet loss requirement because it uses background file transfer during pre-positioning. Excessive loss, however, can indicate a fundamental problem with the network and should be addressed.

Cisco Show and Share

Live Streaming Content

The predominant video type used with Cisco Show and Share is Windows Media Video (WMV), which uses the following three delivery types to deliver the live stream to the client, with each type having unique requirements and methods for handling packet loss:

- RTSP-T (TCP)

Packets lost at the network layer are recovered through resend by the TCP session, independent of the application. Loss should be kept under 0.5 percent for optimal performance. Use of Cisco WAAS increases loss tolerances to 1 percent.

- RTSP-U (UDP)

Packets lost at the network layer are recovered through UDP resend controlled at the application layer. The client application detects loss and requests the lost packet be resent. Loss should be kept under 0.5 percent for optimal performance.

- Multicast (UDP)

Packets lost at the network layer are repaired and recovered at the client layer through error-correcting code (ECC) error correction. Lost packets are not resent in a multicast implementation. ECC error correction on the client mitigates minimal packet loss. Loss above 0.1 percent begins to become noticeable; loss at 0.5 percent begins to seriously degrade video quality.

Pre-recorded Content

Pre-recorded content that is streamed live to Cisco Show and Share clients has the same requirements as live video, allowing up to 0.5 percent loss before risking quality issues. Content may be delivered through HTTP, RTSP-T, or RTSP-U, all of which have the same loss requirement of less than 0.5 percent. For HTTP and RTSP-T, use of Cisco WAAS increases loss tolerances to 1 percent.

Cisco DMS Network Requirements

Table 3 shows the Cisco DMS network requirements.

Table 3 Cisco DMS Network Requirements

	Type	Transport	Codec	Bandwidth (Mbps)	Latency (ms RTT)	Jitter (ms)	Loss (% packets)
Cisco Show and Share							
	VoD	HTTP	WMV	0.2–1.5	<400 ¹	<100 ¹	< 0.5
		RTSP			<2000	<500	< 0.5
	Live	RTSP	WMV	0.2–1.5	<2000	<500	< 0.5
		Multicast UDP			<2000	<500	< 0.1
Cisco Digital Signs/Cisco Cast							
	VoDs SD	HTTP	MPEG-2	3–5	<100 ¹	<30 ¹	< 0.05
			H.264	1.5–2.5	<100 ¹	<30 ¹	
			WMV				
	VoDs HD	HTTP	MPEG-2	13–15	<60 ¹	<30 ¹	< 0.05
			H.264	8–12	<60 ¹	<30 ^v	
	Live SD	Multicast UDP	MPEG-2	3–5	<2000	<1000	DMP 4400G—<0.05
							DMP 4305G—0
	Live HD	Multicast UDP	MPEG-2	13–15	<2000	<1000	DMP 4400G—<0.05
							DMP 4305G—0
			H.264	8–12	<2000	<1000	0

1. Latency and Jitter tolerances significantly increased by optimizing TCP through the use of the Cisco WAAS.

Cisco DMS Tolerance Increases with Cisco WAAS

Table 4 shows how the Cisco DMS tolerance increases with the use of Cisco WAAS.

Table 4 Cisco DMS with Cisco WAAS

Application	Aspect	Maximum Tolerance without WAAS	Maximum Tolerance with WAAS
Cisco Show and Share (RTSP-T/HTTP)	Latency (ms RTT)	400	800
	Jitter (ms)	100	200
	Loss	0.5%	1%
Cisco Digital Signs and Cisco Cast (RTSP-T/HTTP)	Latency (ms RTT)	60	220
	Jitter (ms)	30	60
	Loss	0.05%	0.1%

Implementing Cisco WAAS increases the efficiency of WAN links significantly. [Table 4](#) shows the increase in tolerances for DMS based on the Transport Flow Optimization (TFO) feature, which optimizes TCP-based traffic. TFO is one of several benefits provided by WAAS. For more information about Cisco WAAS, see [Application Network Services, page 33](#).

Cisco DMS Generalized Network Requirements

[Table 3](#) details specific requirements for each aspect of Cisco DMS. Taking all the information into account, a basic list of network requirement characteristics can be created. This list is highly generalized and should be used to express only the general requirements and behavior of DMS on the network. This list also assumes content pre-positioning for VoDs accessed through TCP delivery mechanisms and optimization of TCP delivery traversing WAN links with Cisco WAAS.

[Figure 7](#) shows the latency, jitter, and packet loss settings for DMS.

Figure 7 Latency, Jitter, and Packet Loss Requirements

Latency	• High Tolerance
Jitter	• High Tolerance
Loss	• Critical - Low to No Tolerance

226873

Tuning TCP Parameters

Latency requirements can be reduced by tuning certain TCP parameters. Increasing TCP window size may have a significant impact on throughput with moderate latency. The following equation and example show how to calculate effective values for the TCP window size.

The basic equation is as follows:

$$TcpWindowSize(bytes) = Total\ Available\ Bandwidth\ (KBps) \times Latency\ Round\ Trip\ (ms)$$

Examining a 100 Mbps connection with 200 ms of latency round trip, the equation is built and calculated as follows:

- 100 Mbps = 12,500 Kbps
- $TcpWindowSize(bytes) = 12,500\ Kbps \times 200\ ms$
- $TcpWindowSize(bytes) = 2,500,000\ bytes$

The conclusion from this calculation is to set the TCP window size to 2,500,000 bytes.

Each operating system has a different process for adjusting TCP settings. Details for adjusting TCP parameters in Windows 2003 Server R2 are provided in the next section.

Windows 2003 Server R2 TCP Tuning

In Windows 2003 Server R2, TCP settings are adjusted in the registry at the following location:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Tcpip\Parameters

For optimal performance, keys should be edited, or added if they do not exist, as shown in [Table 5](#).

Table 5 *TPC Parameters*

DefaultTTL	64
EnableRSS	1
EnableTCPA	1
EnableTCPChimney	1
SackOpts	1
Tcp1323Opts	1
TcpMaxDupAcks	2
GlobalMaxTcpWindowSize	2500000
TcpWindowSize	2500000

All values in [Table 5](#) are in decimal and are type REG_DWORD. The value of 2500000 represents 2,500,000 bytes in the example above and would be adjusted with different network implementations.

Cisco WAAS TCP Tuning

When using Cisco WAAS, tuning TCP parameters is critical. The same formula may be used as above to tune the Cisco WAAS TCP parameters, as shown in [Figure 8](#).

Figure 8 *Tuning Cisco WAAS TCP Parameters*

Quality of Service

Figure 9 lists the QoS settings for DMS.

Figure 9 QoS Settings


Cisco Video Application	RFC4594 Application Class	DSCP	Reference RFC
	Network Control	CS6	RFC 2474
	VoIP Telephony	EF	RFC 3246
IP Video Surveillance – DMS Multicast	Broadcast Video	CS5	RFC 2474
CUPC	Multimedia Conferencing	AF41	RFC 2597
TelePresence	Real-Time Interactive	CS4	RFC 2474
Digital Media Systems	Multimedia Streaming	AF31	RFC 2597
	Call-Signaling	CS3	RFC 2474
	Transactional Data	AF21	RFC 2597
	OAM	CS2	RFC 2474
	Bulk Data	AF11	RFC 2597
	Best Effort	DF	RFC 2474
YouTube/BitTorrent/etc.	Scavenger	CS1	RFC 3662

226786

Cisco Digital Signs and Cast

Figure 10 shows an example of QoS classification for Cisco Digital Signs and Cisco Cast.

Figure 10 QoS Traffic Classification for Cisco Digital Signs and Cisco Cast



253733

For live streaming video, a Scientific Atlanta encoder is recommended. Scientific Atlanta encoders allow Differentiated Services Code Point (DSCP) to be set for multicast streams. This value needs to be set to CS5 through the web-based management interface running on the encoder. With the proper DSCP value set on the encoder, remarking on the switch is not needed and the port may be set to trust DSCP.

Live multicast streaming with Cisco Cast may require special consideration from a network implementation standpoint. Policing is generally not recommended for traffic of this nature because of its extreme sensitivity to loss. Cisco Cast may require the use of policing to prevent significant queue overrun situations caused by multiple multicast streams being sent to the DMPs simultaneously. Other bandwidth control or call admission control (CAC) techniques exist that may be implemented instead of queue policing. These techniques are discussed in [Multicast Design Considerations for Cisco DMS, page 24](#).


Pre-recorded content and applications are delivered through HTTP from any third-party web server or CIFS share. If the web server or CIFS share has the option to set the DSCP value of exiting packets, the port(s) for the web server may be trusted. Remarking is also an option for content coming from the web server(s) or CIFS share(s). If the servers are dedicated for use with DMS, marking of all traffic may be implemented.

Pre-positioning traffic should not be marked because it involves scheduled file transfers across the network.

Cisco Show and Share

Figure 11 shows an example of QoS classification for Cisco Show and Share.

Figure 11 QoS Traffic Classification for Cisco Show and Share



Cisco Show and Share real-time streaming may be implemented in multiple ways. Prioritization of real-time RTSP-T or RTSP-U media streams is suggested with marking of DSCP value of AF31. The application class of AF31 is for multimedia streaming, which best matches the type and nature of the Show and Share streaming traffic.

VoD traffic streamed across the network may also be marked AF31. VoD traffic that has been pre-positioned generally does not need prioritization because it is available on the local network with no WAN traversal necessary.

Pre-positioning traffic should not be marked because it involves scheduled file transfers across the network.

Network Requirements Summary

Cisco Digital Signs and Cisco Cast have a high tolerance for latency and jitter but little to no tolerance for loss when displaying live content. Table 3 shows that loss tolerances are between 0.05 and 0 percent for live streaming video, depending on the platform and codec used. Latency and jitter requirements become more stringent for pre-recorded content because the delivery is TCP-based. Content pre-positioning for VoDs via a content delivery network, and optimization of TCP delivery traversing WAN links with Cisco WAAS may be necessary.

Cisco Show and Share has high tolerances for latency and jitter as well as minimal tolerances for loss when displaying live content. Latency and jitter requirements become more stringent for pre-recorded content. Content pre-positioning for VoDs via a content delivery network, and optimization of TCP delivery traversing WAN links with Cisco WAAS may be necessary.

Multicast Design Considerations for Cisco DMS

Multicast is the only transport option for delivering live video to the Cisco DMP. Multicast is an optional method for delivering live streaming video to Cisco Show and Share clients. Multicast differs greatly from other transport mechanisms in its implementation, impact on the network, distribution ability, security, and reliability. This section describes multicast design considerations for both Cisco Digital Signs/Cisco Cast and Cisco Show and Share, focusing primarily on the delivery of content to Cisco DMPs.

Basics of Multicast

This section assumes some knowledge of multicast. Multicast technology primers and other technical content are available online. Only the bare minimum of information is presented here that describes basic multicast operation.

Multicast has the following five main implementation protocol options, referred to as Protocol Independent Multicast (PIM) types:

- *Dense mode*—Dense mode uses the “push” model, which floods traffic throughout the network and is pruned back where it is unwanted. Flood and prune typically occurs every three minutes.
- *Sparse mode*—Sparse mode uses the “pull” model, which sends traffic only upon request. Explicit joins are issued by the multicast clients to request a multicast stream. A rendezvous point (RP) is used on a designated router (DR) to connect senders and receivers.
- *Sparse-dense mode*—A combination of the previous two modes that allows for both push and pull behavior.
- *Source-specific multicast (SSM)*—SSM requires the multicast clients to specify the multicast source explicitly. In this mode, an RP is not needed.
- *Bi-directional mode*—Bi-directional mode is used for implementations requiring all multicast clients to be multicast sources simultaneously.


With the DMS, sparse mode or SSM are recommended and assumed for the rest of this section.

Dense mode may cause significant issues, especially with Cisco Cast, because it can limit functionality and overload network links. Dense mode is not recommended for implementation of the Cisco DMS. Bi-directional mode is not needed because Cisco DMPs receive only multicasts.

Multicast Use in Cisco DMS

Figure 12 shows an example of multicast use with Cisco Digital Signs and Cisco Cast.

Figure 12 Multicast Use with Cisco Digital Signs and Cisco Cast



With Digital Signs and Cast, multicast-delivered UDP is used for all live video content. Live content must be delivered by multicast; no unicast option exists at this time. The Cisco DMP sends join and leave IGMP messages to control delivery of multicast content originating from a Scientific Atlanta encoder, or other compatible source.




Note

With the introduction of DMS 5.2, the Cisco DMPs now support RTSP-T for streaming pre-recorded files, but not for streaming live content.

With Cisco Show and Share, multicast may optionally be used to delivery live content to client stations, as shown in Figure 13.

Figure 13 Multicast Use with Cisco Show and Share



Multicast may originate from Windows Media Streaming Services running Windows 2003 Server, Windows 2008 Server, or Cisco ACNS. Windows Media Streaming Services allows for fallback to RTSP-T and RTSP-U unicast, allowing multicast to be implemented with non-multicast capable clients still able to receive the stream.

Implementing Multicast for Cisco DMS

The following sections summarize the implementation of multicast for the Cisco DMS:

- PIM Sparse Mode
- Source Specific Multicast

PIM Sparse Mode

Implementation of the Cisco DMS with PIM sparse mode has no specific caveats, other than those explicitly discussed later in this section, and are not discussed in detail. Implementation with PIM sparse mode is a recommended strategy.

For more information, see the following URL: <http://www.cisco.com/go/multicast>.

Source Specific Multicast

Source-specific multicast (SSM) is a recommended strategy that avoids the need for a rendezvous points (RP) and a shared tree for multicast flows. SSM is more efficient, secure, and scalable than implementing PIM sparse mode alone. SSM may be used exclusively or co-exist with PIM sparse mode.

SSM does require some additional configuration on the last hop routers when compared to PIM sparse mode configuration. Implementation and configuration of SSM for DMS is discussed in this section.

Multicast JOIN messages have basically two critical pieces of information in them to identify the desired multicast stream:

- Group (G)—Multicast IP address of the multicast stream
- Source (S)—IP address of the device sending the multicast stream

Most multicast clients will JOIN a group but not know the source, sending a (*,G) JOIN message. If a multicast client knows the source, they can send a (S,G) JOIN message. This is known as a *source specific JOIN*.

SSM requires the client to explicitly define the source, and that all components support Internet Group Management Protocol (IGMP) v2. Although the Cisco DMPs do support IGMP v2, they do not at this time have any way of identifying or specifying a source. The router that receives the IGMP JOIN from the DMP receives a (*,G) JOIN. This router is referred to as the *last-hop* router for the multicast stream because it is the last hop before the stream reaches the Cisco DMP, or the closest router to the Cisco DMP.

The last-hop router must *translate* the (*,G) to (S,G) if SSM is to be used with the Cisco DMPs. This is accomplished by defining static IGMP mappings in the Cisco IOS configuration, as follows:

```
ip igmp ssm-map enable
ip igmp ssm-map static <acl-1> <source-1 IP address>
ip igmp ssm-map static <acl-2> <source-2 IP address>
```

With large, complex networks, static mapping on all last-hop routers might not be a scalable or easily maintainable solution. Domain Name System (DNS)-based SSM mapping can be used to maintain all mappings with centrally managed DNS. DNS-based SSM mapping enables you to configure the last-hop router to perform a reverse-DNS lookup to determine sources sending to groups. When DNS-based SSM mapping is configured, the router constructs a domain name that includes the group address *G* and performs a reverse lookup into DNS. The router looks up the IP address resource records to be returned for this constructed domain name and uses the returned IP addresses as the source addresses associated

with this group. SSM mapping supports up to 20 sources for each group. The router joins all sources configured for a group; although, with the Cisco DMPs, a single source would normally be used. The following Cisco IOS global configuration commands are used:

```
ip domain-server <IP address>
ip domain-name <domain name>
ip igmp ssm-map enable
ip igmp ssm-map query dns
```

To prevent (*,G) entries from being forwarded when SSM is desired, you can use the **ip pim ssm range** command. This Cisco IOS global command prevents requests from being forwarded for specific groups if the source is not present. The following Cisco IOS global configuration command is used:

```
ip pim ssm range <acl-1>
```

Figure 14 and Figure 15 illustrate examples of SSM implementations for DMS, both with and without DNS.

Figure 14 SSM Without DNS





Figure 14 shows the following process:

1. The Cisco DMP learns the multicast group address and port from the Cisco DMM.
2. The Cisco DMP sends an IGMP (*,G) join to the campus router.
3. The campus router matches the multicast address (G) to a multicast source (S).
4. The campus router sends a PIM (S,G) to the data center router, the next hop toward the source.
5. Multicast data flows to client across the shortest path.

Figure 15 SSM with DNS



226887

Figure 15 shows the following process:

1. The Cisco DMP learns the multicast group address and port from the Cisco DMM.
2. The Cisco DMP sends an IGMP (*,G) join to the campus router.
3. The campus router performs a reverse-DNS lookup of the group (G) and receives the source (S).
4. The campus router sends a PIM (S,G) to the data center router, the next hop toward the source.
5. Multicast data flows to client across shortest path.

Multicast Considerations with Cisco Technologies

The following subsections summarize multicast considerations with Cisco technologies:

- Multicast with VSS
- Multicast with Cisco VRF-Lite

Multicast with VSS

The Cisco Virtual Switching System (VSS) distribution block design is a significant change from either the routed access or multi-tier designs. The introduction of the Cisco Catalyst 6500 VSS and Stackwise/Stackwise-Plus in the Cisco Catalyst 3750/3750E provides the opportunity to make a significant change to the way switch and link redundancy is implemented. In the past, multiple Layer 3/Layer 2 access switches were connected to two redundant distribution switches, and the configuration of the network control protocols, such as the Hot Standby Routing Protocol (HSRP), 802.1D spanning tree, and Enhanced Interior Gateway Routing Protocol (Enhanced IGRP), determined the way in which the switches forwarded traffic over each of the uplinks and the network recovered in the event of a switch or link failure. With the introduction of the virtual switch concept, the distribution switch pair can now be configured to run as a single logical switch. By converting the redundant physical

distribution switches into a single logical switch, a significant change is made to the topology of the network. Rather than an access switch configured with two uplinks to two distribution switches, and needing a control protocol to determine which of the uplinks to use, now the access switch has a single multichassis EtherChannel (MEC) upstream link connected to a single distribution switch.

Implementation of multicast in a VSS environment is no different than implementation in a non-VSS environment.

For more information about Cisco VSS, see the following URL: <http://www.cisco.com/go/vss>.

Multicast with Cisco VRF-Lite

The use of a switched VLAN-based design has provided a number of advantages, including increased capacity, better isolation, and improved manageability. However, the flexibility that VLANs offer has had the greatest impact on campus designs. The ability to dynamically reconfigure the network and add new subnets or business groups (without having to physically replace the network) has provided huge cost and operational benefits. The modern campus networking environment exists largely because of the capabilities that VLAN virtualization provides. Although VLANs provide flexibility in dynamically segmenting groups of devices, VLANs do have some limitations. As a Layer 2 virtualization technique, VLANs are bound by the rules of Layer 2 network design. In the structured hierarchical campus design VLANs do not have the flexibility to span large domains.

The use of Virtualized Routing and Forwarding (VRF) with Generic Routing Encapsulation (GRE), 802.1q, and Multiprotocol Label Switching (MPLS) tagging to create virtual private networks (VPNs) in the campus combine to provide one approach for extending the configuration flexibility offered by VLANs across the entire campus, and, if required, through the entire network. VRFs provide the ability to have separate routing and forwarding instances inside one physical switch. Each VRF has its own Layer 3 forwarding table. Any device in a specific VRF can be Layer 3 switched directly (in other words, routed) to another device in the same VRF, but cannot directly reach one in another VRF. This is similar to the way in which each VLAN in each switch has its own Layer 2 forwarding and flooding domain. Any device in a VLAN can directly reach another device at Layer 2 in the same VLAN, but not a device in another VLAN, unless traffic is forwarded by a Layer 3 router.

Taking into account all the details presented in the previous sections about the characteristics of a VRF-Lite end-to-end deployment, you can additionally enable the multicast functionality inside each defined virtual network by simply mirroring the multicast configuration already in place in the underlying infrastructure (the global table) on platforms that support multicast functionality within a VRF. All required multicast components can be virtualized. A dedicated multicast routing table is available in the context of each defined VRF, while a separate instance of the required multicast protocols, such as PIM, IGMP, or Multicast Source Discovery Protocol (MSDP), can be enabled inside each VRF. Even the RP discovery mechanism and placement can be inherited from the global table multicast deployment.

For more information about Cisco VRF-Lite, see the network virtualization design guides at the following URL: <http://www.cisco.com/go/cvd>.

Other Multicast Design Considerations

This section includes the following topics:

- Multicast Across a WAN
- Protecting Multicast-Enabled WAN Links from Rogue Clients
- Unique Cisco Cast and Cisco DMP Considerations for Multicast
- Multicast Bandwidth Control and Call Admission Control

Multicast Across a WAN

Multicast streams with the Cisco DMS are generally in the range of 3–15 Mbps but can be as high as 25 Mbps. Because of the heavy bandwidth requirements of the Cisco DMS for live multicast feeds, only higher speed WAN links should be considered for remote implementation of the Cisco DMPs. Cisco ACNS may not be used to transport multicast across a WAN as a unicast stream because of compatibility limitations.

Multicast streams with Cisco Show and Share are relatively small in comparison, ranging from 200 Kbps to 1.5 Mbps. These multicast streams are transportable across WAN links with multicast enabled and appropriate bandwidth available. Cisco ACNS may also be used to transport unicast RTSP streams across a WAN, converting them into multicast at the remote site. More information regarding the use of ACNS for unicast to multicast conversion is available in [Application Network Services](#), page 33.

Protecting Multicast-Enabled WAN Links from Rogue Clients

Be careful with slower WAN links that have multicast enabled for other applications. Multicast applications, such as music-on-hold (MoH), with Cisco Unified Communications Manager for IP telephony may be implemented across slower WAN links. Client PCs would be able to send multicast join messages to the network, joining and subsequently receiving these DMS multicast streams, which can easily exceed the capacity of the WAN link.

While this might seem an unlikely issue, clients at remote locations might attempt to join and view the video streams if they gain knowledge of specific stream addresses. Common, free programs such as Video LAN Client (VLC) can be used to join, decode, and view a Cisco DMS multicast stream.

Several options are available to protect against this issue, including the following:

- Implementation of VLAN segmentation—Any VLAN containing Cisco DMS multicast traffic can be restricted from traversing the WAN.
- Implementation of access control lists (ACLs)—All multicast traffic may be blocked, or specific Cisco DMS multicast traffic may be blocked, using source address or destination multicast address.
- Implementation of virtualization technologies, such as VRF-Lite—VRF-Lite can be used to restrict access to specific devices and ports by placing Cisco DMS components into a separate virtual network.


Unique Cisco Cast and Cisco DMP Considerations for Multicast

With DMS 5.1 and earlier, the interface for Cast displays a preview of each live channel as the channels are scrolled through. With DMS 5.2, the preview is not shown for each channel unless manually selected; however, selecting a channel to display full screen and then scrolling through other channels rapidly causes the same behavior discussed below.

To display the channel, a multicast stream must be sent to the Cisco DMP. The Cisco DMP issues a multicast join to begin receiving the desired multicast stream. As the user scrolls through channels, the Cisco DMP sends a multicast leave for the previous stream and a join for the next, showing the next channel.

Multicast streams, however, do not stop the moment the multicast leave message is sent. Several concurrent streams might end up being active at the same time if scrolling very quickly through channels. As many as ten streams can be concurrently active to a Cisco DMP. This can be a problem, especially when the streams are in HD format. Ten concurrent 15-Mbps HD streams create a total of 150 Mbps of traffic to the Cisco DMP. This is shown in Figure 16, which shows actual traffic to a Cisco DMP during channel hopping.

Figure 16 Cisco DMP Channel Hopping Behavior



Multicast Bandwidth Control and Call Admission Control

To address the issue of overrunning queues and links with excessive multicast traffic, the following options exist for bandwidth control and CAC:

- Policing multicast traffic—Policing the queue used for Cisco DMP multicast traffic prevents link saturation from excessive multicast traffic, but has the drawback of impacting all other traffic (including other multicast flows) in the same queue.
- IGMP state limit feature—This feature introduces the capability to limit the number of *mroute* states resulting from IGMP membership states per interface, per subinterface, or globally.
- Per interface *mroute* state limit feature—This feature allows limiting the number of *mroute* states on an interface for various ACL-classified sets of multicast traffic.

Policing Multicast Traffic

To help prevent overrunning links with multicast traffic, policing can be put in place on queues supporting the traffic. Doing this can affect the quality of video displayed by other Cisco DMPs serviced by the same link, and is generally discouraged on queues containing multiple streams. The impact should be momentary if related to channel hopping, and is more desirable than the alternative of overrunning the link. Links supporting a single Cisco DMP can implement this technique if no other flows exist in the same queue used for delivering multicast traffic to the Cisco DMPs.

Cisco Digital Signs/Cisco Cast multicast traffic is normally marked with a DSCP value of CS5. The queue assigned to handle Digital Signs/Cast multicast streams should be implemented as a bandwidth queue with policing enabled. Policing limits the amount of bandwidth available to the multicast streams, dropping any traffic in excess of the defined allowable bandwidth and preventing the entire link from being overrun.

IGMP State Limit Feature

The IGMP State Limit feature introduces the capability to limit the number of mroute states resulting from IGMP membership states per interface, per subinterface, or globally. Membership reports exceeding the configured limits are not entered into the IGMP cache, and traffic for the excess membership reports is not forwarded.

Per-interface and global IGMP limits operate independently of each other. Both per-interface and global IGMP limits can be configured on the same router. A membership report that exceeds either the per-interface or the global state limit is ignored.

To prevent multicast flooding because of channel hopping with the DMP, every interface with a DMP connected can be configured to limit the IGMP membership state to one. This strategy allows only one multicast flow, or channel, at a time.


Per Interface Mroute State Limit Feature

The Per Interface Mroute State Limit feature allows you to limit the number of mroute states on an interface for different ACL-classified sets of multicast traffic. This feature can be used to prevent denial-of-service (DoS) attacks, or to provide a multicast CAC mechanism when all the multicast flows generally use the same amount of bandwidth.

To prevent multicast flooding because of channel hopping with the Cisco DMP, upstream interfaces can be configured to limit the number of mroute states on that particular interface. This strategy allows you to define stream limits that prevent overruns because of channel hopping behavior, and to set overall limits on the number of streams for a particular connection. This is a form of CAC for Cast live streams.

The IGMP State Limit feature and Per Interface Mroute State Limit feature can be combined to block the channel hopping issue at the source while also controlling the overall utilization of particular links by limiting the number of multicast flows across them. (See [Figure 17](#).)

Figure 17 Using IGMP State Limit and Per Interface Mroute State Limit



Application Network Services

This section focuses on application network services (ANS) that provide optimized content distribution. Cisco provides the following offerings for content distribution:

- Cisco Application and Content Networking System (ACNS)
- Cisco Wide Area Application Services (WAAS)
- Cisco Content Delivery System (CDS)

In addition to these offerings, Cisco DMS 5.2 now includes basic content distribution capabilities for Cisco Digital Signs and Cisco Cast: Cisco Digital Media Suites Content Delivery.

All of these technologies are described in some detail, with the exception of CDS, which is summarized. In-depth analysis and testing of CDS is outside the scope of this version of the document.

In addition to content distribution, the optimization technology performance routing (PfR) is also described.

Cisco Application and Content Networking System

This section provides the following descriptions for Cisco Application and Content Networking System (Cisco ACNS):

- Cisco ACNS Overview
- Cisco ACNS Components
- Cisco ACNS Deployment
- Cisco ACNS with Cisco Digital Signs/Cisco Cast
- Cisco ACNS with Cisco Show and Share

Cisco ACNS Overview

Cisco ACNS is a powerful digital media delivery solution that enables enterprises to deploy a highly scalable, reliable, and manageable network infrastructure for video business applications across a WAN. Cisco ACNS can be used to pre-position VoD content for all three Cisco DMS subsystems. Benefits provided by Cisco ACNS include the following:

- Reducing demand on the WAN during peak hours
- Allowing multiple branch clients to access content transferred across the WAN only once
- Ensuring quality of content by accessing locally

Cisco ACNS can be used with Show and Share live streaming content, performing stream transport and splitting. This reduces bandwidth requirements on the WAN by sending a single stream to a remote location, regardless of the number of clients at that location.

Cisco ACNS can use one of the following three methods for providing content to remote users:

- Proxy

In proxy mode, the web browser of the end user must be explicitly configured to use the IP address or host name of the Cisco Content Engine (CE). There is no need for additional hardware, such as Layer 4 switches, content routers, or WCCP-enabled routers to intercept user requests. If the CE has the requested content, the request is serviced from this cached storage. If the content is not in CE storage, the content is requested from the origin web server and served to the requesting client by way of the CE.

- Web Cache Communications Protocol (WCCP) v2

When configured for WCCP routing, the CE receives requests from its assigned router and compares the content request against the content currently stored in its cache. If the CE has the requested content (cache hit), the request is served from this cached storage. If the CE does not have the content (cache miss), the content is requested from the origin web server and served to the requesting client through the CE. No proxy configuration is required at the end user browser.

- Content routing

Content routing uses HTTP or RTSP redirection through a content router. The Cisco Content Router (CR) determines which content engine is best-suited to deliver the desired content to the client by comparing the source IP address of the client end system against a table of address ranges assigned to the CEs, known as the *coverage zone*. The coverage zone provides information on the proximity of client end systems to the CEs based on the IP address of each client. The CR can then choose the closest, best-suited CE to serve the website request to the client.

Cisco ACNS Components

All Cisco ACNS components use the Cisco Wide Area Application Engine (Cisco WAE) as the operating platform. Key Cisco ACNS components include the following: Cisco Content Distribution Manager (CDM), Cisco Content Engine (CE), and Cisco Content Router (CR).

Cisco Content Distribution Manager

The primary role of a CDM is to perform centralized content and device management. In the Cisco ACNS 5.5 network, the CDM manages both content acquisition and distribution and also manages policy settings and configurations on individual CE. Through the CDM GUI, the network administrator can specify what content is to be distributed and to whom. The CDM also allows the administrator to monitor network nodes and apply changes, such as software upgrades, to groupings of nodes from a central location.

Cisco Content Engine

The primary role of a CE in the Cisco ACNS network is to serve client requests for content. CEs also play a major role in content request routing and in channel distribution of content. The root CE acquires the content from the origin servers and distributes that content to other CEs.

Cisco Content Router

The primary role of a CR in the Cisco ACNS network is to redirect client requests for content to the closest CE containing that content. In the Cisco ACNS 5.5 network, CRs use DNS to ensure that the CR receives all requests for content. When the CR receives these requests, it can redirect the end user to the content.

Cisco ACNS Deployment


Cisco ACNS deployment involves implementing a CDM, optional CR, and root CE in the data center or campus where the original content resides; and CEs in branch locations as well as other campus locations. The CDM controls the configuration of the CR and CEs as well as holding the schedules of content positioning. Source, destination, time, and bandwidth use are all maintained and controlled from the CDM.

For more information about Cisco ACNS, see the following URL: <http://www.cisco.com/go/acns>.

Cisco ACNS with Cisco Digital Signs/Cisco Cast

Cisco ACNS is a very powerful content distribution system that, when used with Cisco Digital Signs and Cisco Cast, provides highly configurable content pre-positioning and caching. (See [Figure 18](#).)

Figure 18 Cisco Digital Signs/Cisco Cast with Cisco ACNS



Pre-recorded Content

Cisco ACNS provides significant benefits for VoDs, applications, and other Digital Signs and Cast content:

- Pre-positions content to edge CEs or to the storage available on a Cisco DMP itself
- Performs scheduled transfers
- Controls bandwidth used for transfers
- No duplication of transfers across the WAN; content is transferred once for all remote Cisco DMPs
- Fully integrated into Digital Media Manager for Digital Signs and Cast modules

After creating channels using the Cisco ACNS CDM and assigning devices to those channels, the DMM can be used to configure scheduled transfers. Scheduled transfers are then pushed to the Cisco ACNS CDM in the form of manifest files. Cisco DMS administrators can schedule off-hours transfers, defining both when to transfer content and how much bandwidth to use across specific links.


Live Content

Multicast is required as the transport for Digital Signs and Cast live content. Cisco ACNS does not support unicast to multicast conversion for the delivery method used by Digital Signs and Cast. For this reason, Cisco ACNS provides no benefit for live content.

Cisco ACNS with Cisco Show and Share

Cisco ACNS is a very powerful content distribution system that, when used with Cisco Show and Share, provides highly configurable content pre-positioning and caching as well as live stream splitting and conversion. (See [Figure 19](#).)

Figure 19 Cisco Show and Share with Cisco ACNS



Pre-recorded Content

Cisco ACNS provides significant benefits for VoDs, applications, and other Digital Signs and Cast content:

- Content may be pre-positioned to edge CE
- Scheduled transfers
- Controls bandwidth used for transfers
- No duplication of transfers across the WAN; content is transferred once

Cisco ACNS, when used with Show and Share, has similar benefits as Digital Signs and Cast for VoDs, with the exception of integration to the Show and Share module. Administration of content distribution is handled through the Cisco ACNS CDM for Show and Share content.

Live Content

Cisco ACNS provides several benefits for optimizing live Show and Share content:

- Primary streaming server, converting the HTTP stream from the Digital Media Encoder to RTSP or multicast streams for clients
- RTSP-T and RTSP-U stream splitting at the edge
- RSTP-T and RTSP-U unicast to multicast conversion at the edge

Primary Streaming Server

For Show and Share, the Digital Media Encoders do not send live streams. Cisco ACNS may be used in place of Microsoft Windows 2003 Server or any other source as the primary streaming server, receiving the live HTTP stream directly from the Digital Media Encoder. As the primary streaming server, the root CE runs Windows Media Streaming Services, allowing clients to connect directly to the root CE or serve edge Cisco ACNS CEs for remote site stream splitting.

Stream Splitting

Central and remote CEs are capable of receiving a single Windows Media stream via either RTSP or HTTP and serving that stream to multiple client machines, reducing the load across lower speed links.

Stream Conversion

Cisco ACNS is capable of converting Windows Media stream transport from HTTP, RTSP-T, or RTSP-U to RTSP-T, RTSP-U, or Multicast UDP.

Conversion of unicast to multicast allows an organization to do the following:

- Send a unicast stream across WAN links to remote sites while still taking advantage of the scaling capabilities of multicast at those remote sites
- Send a multicast stream out to the campus network, converting the unicast source directly from the Digital Media Encoder to multicast

One caveat with unicast to multicast conversion is the way in which multicast streams are enabled from a unicast source. The unicast stream must be live before the unicast-to-multicast conversion can be enabled. This is a manual, two-step process for each stream to be enabled: first starting the unicast stream, and then enabling the conversion. Unicast-to-multicast conversion works best for “always-on” live sources such as TV channels because the enablement process is executed only once.

Cisco Wide Area Application Services

Cisco Wide Area Application Services (WAAS) can be implemented to increase the efficiency of video delivery for all three Cisco DMS subsystems.

This section provides the following descriptions:

- [Cisco WAAS Overview](#)
- [Cisco WAAS Benefits for Video](#)
- [Cisco WAAS Components](#)
- [Cisco WAAS Deployment](#)
- [Cisco WAAS with Digital Signs/Cast](#)
- [Cisco WAAS with Cisco Show and Share](#)

Cisco WAAS Overview

Cisco WAAS is a comprehensive WAN optimization solution that accelerates applications over the WAN, delivers video to the branch office, and provides local hosting of branch office IT services. Cisco WAAS allows IT departments to centralize applications and storage in the data center while maintaining LAN-like application performance, and provides locally hosted IT services while reducing the branch office device footprint.

Cisco WAAS helps organizations to accomplish these primary IT objectives:

- Application acceleration—Improves productivity of remote employees
- IT consolidation and WAN optimization—Minimizes branch IT costs
- Branch IT agility—Responds rapidly to changing business needs
- Simplified data protection—Eases compliance and business continuity

Unlike other WAN optimization products, the Cisco WAAS solution offers the following unique advantages:

- Validated by application vendors with proven designs that lower risks
- Integrates into your network to ease network operations and management
- Minimizes your ownership costs by reducing device footprints and complexity

Cisco WAAS Benefits for Video

Cisco WAAS provides the following four main benefits for Cisco DMS video delivery:

- Windows Media stream splitting—Starting with Cisco WAAS 4.1, the Cisco WAAS Video Optimizer feature allows Windows Media stream splitting, requiring only one live stream to traverse WAN links regardless of how many clients are located at the remote site.
- CIFS share—With the release of Cisco DMS 5.2, all Cisco DMS subsystems now support CIFS. CIFS allows Cisco DMPs and Show and Share clients to easily access content that has been pre-positioned or cached at remote locations by Cisco WAAS.
- TCP optimization—Cisco WAAS optimizes TCP sessions across WAN links with its TFO feature, allowing both live and pre-recorded video content using TCP transfer mechanisms to be delivered more efficiently. TCP optimization relaxes the latency, jitter, and loss restrictions that limit throughput.
- Data redundancy elimination (DRE)—The DRE feature of Cisco WAAS caches data at the edge, preventing the same data from traversing the WAN more than once. Pre-recorded video, or VoDs, are sent from the local WAE for requests subsequent to the initial request.
- Reduction of non-video traffic—Though Cisco WAAS compression of video traffic is minimal because of the high compression already implemented with the video codecs, compression and optimization of non-video traffic creates more available bandwidth for video.

Cisco WAAS Components

Key Cisco WAAS components are as follows:

- Cisco WAAS Central Manager—Every Cisco WAAS network must have one primary Cisco WAAS Central Manager device that is responsible for managing the other Cisco WAAS devices in your network. The Cisco WAAS Central Manager hosts the Cisco WAAS Central Manager GUI, a

web-based interface that allows configuration, management, and monitoring of the Cisco WAAS devices in your network. The Cisco WAAS Central Manager resides on a dedicated Cisco WAE appliance. The Cisco WAAS Central Manager does not run on a Cisco WAE network module.

- Cisco WAAS Application Accelerator—Application accelerators run on dedicated Cisco WAE appliances or network modules. Transparent traffic interception and optimization occurs between application accelerators. Application accelerators are managed and controlled by the Cisco WAAS Central Manager.

Cisco WAAS Deployment


Cisco WAAS differs greatly from Cisco ACNS in its purpose and implementation, although it uses the same primary hardware, Cisco WAEs, and is deployed in a similar manner. Cisco WAAS deployment involves implementing a Cisco WAAS Central Manager and one or more Cisco WAAS Application Accelerators in the data center or campus where the original content resides, and Cisco WAAS Application Accelerators in branch locations as well as other campus locations. The Cisco WAAS Central Manager controls the configuration of the application accelerators.

For more information about Cisco WAAS, see the following URL: <http://www.cisco.com/go/waas>.

Cisco WAAS with Digital Signs/Cast

Figure 20 shows the Digital Signs/Cast deployed with WAAS.

Figure 20 Digital Signs/Cast with WAAS



226784

Pre-recorded Content

With the release of Cisco DMS 5.2, Cisco DMPs can now mount CIFS shares, enabling Digital Signs and Cast to take advantage of the pre-positioning abilities of Cisco WAAS. Cisco DMPs mount a CIFS share on their local Cisco WAE running Cisco WAAS Application Accelerator. Cisco WAAS pre-positions content to remote Cisco WAEs to be available upon request from the Cisco DMPs via the CIFS mount.

CIFS support in Digital Signs and Cast is all or nothing in the Cisco DMS Release 5.2. When CIFS is enabled for Cisco DMPs, all Cisco DMPs use the CIFS share. In the Cisco DMS Release 5.2, CIFS support cannot be specified on a per-DMP or per-DMP group level. If no local Cisco WAE exists, the Cisco DMP is forced to mount a remote CIFS share, which might not be as efficient across WAN links as the traditional HTTP-based transfer. Future releases will have additional configuration options for CIFS.

DRE support in WAAS caches content at the edge as it is being accessed for the first time. Although this is not pre-positioning, it is a very effective way to distribute content locally with no additional configuration in DMS. DRE works transparently, allowing locally cached content to be distributed to the DMP as if it is coming from the remote HTTP or RTSP server.


In addition to pre-positioning and caching, Cisco WAAS TCP optimizations help deliver pre-recorded content across WAN links directly to Cisco DMPs, especially links with significant latency. Cisco DMPs use HTTP as the transport for pre-recorded content, which has some limitations when significant latency, jitter, or loss is present in the path, because of the TCP mechanisms for acknowledgement and recovery. Cisco WAAS TCP optimizations provide local acknowledgements, as well as other mechanisms to increase the efficiency of TCP connections, while remaining completely transparent to the endpoints using the TCP connection. Efficiency gains are detailed in [Network Requirements, page 13](#).

Live Content

Cisco WAAS provides little benefit for live content with Digital Signs and Cast because that content is multicast only and highly compressed. Cisco WAAS increases bandwidth available to live streams traversing multicast-enabled WAN links by optimizing and compressing other WAN traffic.

Cisco WAAS with Cisco Show and Share

[Figure 21](#) shows Cisco Show and Share deployed with Cisco WAAS.

Figure 21 Cisco Show and Share with WAAS

Pre-recorded Content

Cisco WAAS can pre-position content accessible through CIFS shares. Show and Share clients are commonly running operating systems capable of mounting a CIFS share. The CIFS share on the local Cisco WAE running the Cisco WAAS Application Accelerator is mounted via the operating system running on the Show and Share client. Requests for content may then be directed to the CIFS share. Consistent authentication and mounting of the CIFS share might become a problem when implemented in client operating systems. Proper IT control and mount scripting on client machines is necessary to effectively implement the content pre-positioning and CIFS sharing capability of Cisco WAAS for Show and Share.

In addition to pre-positioning, Cisco WAAS TCP optimizations help deliver HTTP and RTSP-T based pre-recorded content across WAN links directly to Show and Share clients, especially links with significant latency. Show and Share uses HTTP, RTSP-T, and RTSP-U as the transport methods for delivering pre-recorded content. HTTP and RTSP-T are TCP-based and have some limitations when significant latency, jitter, or loss is present in the path, because of the TCP mechanisms for acknowledgement and recovery. Cisco WAAS TCP optimizations provide local acknowledgements, as well as other mechanisms, to increase the efficiency of TCP connections while remaining completely transparent to the endpoints using the TCP connection.

Live Content

Windows Media stream-splitting is available with the Cisco WAAS Video Optimizer, available in the 4.1 and later releases of Cisco WAAS. The Cisco WAAS Windows Media stream-splitter is capable of taking a source of Live Windows Media stream and delivering it via RTSP-T to many client machines. The WAAS Video Optimizer feature reduces the load across WAN links by requiring only one stream to traverse the WAN to remote branch or campus locations, regardless of the number of clients viewing that stream.


The stream interception and splitting process is accomplished transparently, with no configuration needed in Windows Media Streaming Services, Cisco DMS, or the client station.



Tip

When implementing the Cisco WAAS Video Optimizer feature, all RTSP streams traversing a WAAS optimized WAN link must be TCP-based. Windows Media Streaming Services has the option to disable TCP and force UDP. Such a configuration renders previously working video streams non-functional. [Figure 22](#) shows the Windows Media Streaming Server interface with both TCP and UDP enabled, which is the recommended configuration. Also note that the “Show all plug-in categories” check box must be checked to access the protocol settings.

Figure 22 Windows Media Streaming Server Interface



Cisco Content Delivery System

The Cisco Content Delivery System (CDS) comprises one or more networked Cisco Content Delivery Engines, each optimized for one or more tasks such as content ingest, storage, caching, or streaming. CDS is not discussed in detail in this document.

Cisco Content Delivery Engines

The Cisco CDS consists of networked Cisco Content Delivery Engines (CDEs), which you can group into arrays to work as a single logical system. You can easily attach additional engines for almost unlimited video storage and streaming capacity.

Cisco Content Delivery Applications

Cisco Content Delivery Applications are the software elements of the Cisco CDS and implement content processes on top of Cisco CDEs, providing functions that include ingest, storage, caching, personalization, and streaming.

For more information about Cisco CDS, see the following URL: <http://www.cisco.com/go/cds>.

Cisco Digital Media Suites Content Delivery

Although Cisco DMS Content Delivery (DMS-CD) is not a separate content delivery and optimization offering, it is a basic option for content delivery. Cisco DMS-CD is a new capability introduced in Cisco DMS 5.2 allowing basic distribution of content for Digital Signs and Cast.

Cisco DMS-CD allows scheduled delivery of media and application content to remote Cisco DMPs directly. Content is stored either on the internal Cisco DMP Flash or external USB storage media connected to the Cisco DMP. Cisco DMS-CD is similar to other content distribution technologies, such as Cisco ACNS, in its basic function of scheduled content distribution. It differs significantly in efficiency in certain deployment models.

Cisco DMS-CD sends the desired content to each Cisco DMP, but does not send content that exists on the Cisco DMP storage media, preventing unnecessary duplication across the WAN. Because Cisco DMS-CD delivers directly to the Cisco DMPs without any remote caching engine involved, sites with many remote Cisco DMPs might require other means for efficient content distribution.

Content is sent over the WAN to each Cisco DMP, resulting in content being sent multiple times. WAN utilization can be represented by multiplying the size of the content by the number of Cisco DMPs at a location. Using Cisco ACNS or Cisco WAAS allows content to traverse the WAN once, regardless of how many Cisco DMPs are at a particular location.

Performance Routing

Performance Routing (PfR), an optimization technology, can be used to ensure that Cisco DMS video content takes the optimal path through the WAN. With dual-tier branch deployments, PfR helps ensure video quality by picking the less congested path for video streams.

PfR with Digital Signs/Cast

PfR may be used to ensure optimal delivery of pre-recorded content to Cisco DMPs. This might be beneficial for the delivery of VoDs in real-time to Cisco DMPs with Cast. Using PfR with live multicast streams is not a supported option at this time.

PfR with Show and Share

PfR may be used with Show and Share for both live and pre-recorded content to ensure that the optimal path is chosen for delivery of video to the endpoints.

For more information about Cisco Performance Routing, see the following URL:

<http://www.cisco.com/go/pfr>.

Cisco Media Experience Engine

This section discusses the Cisco Media Experience Engine (MXE) and includes the following topics:

- [Cisco MXE Overview](#)
- [Cisco MXE 3000 Integration with Cisco DMS](#)
- [Cisco MXE 3500 Integration with Cisco DMS](#)

Cisco MXE Overview

The Cisco MXE is a video converter. Video of different resolutions and formats using different codecs can be converted into any other form, allowing video to be converted between different systems and devices.

Cisco MXE 3000 Integration with Cisco DMS

The Cisco MXE 3000 provides the ability to convert from one video format to another and bridges one of the gaps between Cisco Show and Share and Cisco Digital Signs/Cast. Video formats for VoDs supported by Show and Share are not compatible with the Digital Media Players used by Cast and Digital Signs. The same is true for the opposite. With MXE 3000, these formats can be converted from one to another, making VoDs compatible with both systems. [Table 6](#) illustrates the video format support with DMS and the need for MXE to convert from one format to another.

Table 6 *Cisco DMS Video Format Support*

Video Format	Digital Signs/Cast	Show and Share
MPEG-2	X	
MPEG-4 Part 2	X	
MPEG-4 Part 10 H.264	X	
WMV	X ¹	X
FLV		X
Quicktime		X

1. Supported on Cisco DMP 4400G for VoD playback only.

In addition, VoDs created for Cisco Show and Share specifically may be supplied in different formats. The Cisco MXE allows conversion of these different formats to a common single format if desired.

Figure 23 displays the Cisco MXE 3000 placed between the content repository for Cisco Show and Share and the content repository for Cisco Digital Signs and Cisco Cast. These can be separate repositories or a single repository.

Figure 23 Cisco MXE 3000 with DMS




Network requirements of the Cisco MXE 3000 are limited at this time because it is a post-production video conversion system. The network is used for file transfers of pre- and post-production video content using basic transfer methods such as FTP. No special mechanisms need to be applied to this traffic.

Cisco MXE 3500 Integration with Cisco DMS

The Cisco MXE 3500 adds to the capabilities of the MXE 3000 with the ability to transcode live multicast streams to Cisco Show and Share, providing the bridge for live content to be shown on all DMS endpoints from one feed. (See Figure 24.)


Figure 24 Cisco MXE 3500 with DMS



The following sequence takes place:

1. The Scientific Atlanta encoder sends a multicast video stream to the network compatible with the Cisco DMPs.
2. The Cisco MXE 3500 joins the multicast stream to ingest the video stream.
3. The Cisco MXE 3500 converts the video stream to a unicast, Windows Media HTTP-based video stream.
4. The Windows Media Streaming server receives the HTTP-based stream and distributes it as RTSP or multicast.

Another way to look at how the Cisco MXE 3500 fits into DMS is to think of it as a Cisco DMP and Cisco DME connected back-to-back, providing the bridge for live content to be shown on all DMS endpoints from one feed. Figure 25 shows a DMP and DME in place of the Cisco MXE 3500.

Figure 25 Cisco DMP and Cisco DME with DMS


The following sequence takes place:

1. The Scientific Atlanta encoder sends a multicast video stream to the network compatible with the Cisco DMPs.
2. The DMP joins the multicast stream to ingest the video stream.
3. The DMP sends the decoded video stream to the DME as raw video.
4. The DME ingests the video and converts it to a unicast, Windows Media HTTP-based video stream.
5. The Windows Media Streaming server receives the HTTP-based stream and distributes it as RTSP or multicast.

Places in the Network Architecture Design Considerations

Figure 26 shows the places in the network (PIN) architecture.

Figure 26 *Places in the Network Architecture Design*



253787

Data Center PIN Design Considerations

The data center is home to the computational power, storage, and applications necessary to support an enterprise business. The data center infrastructure is central to the IT architecture, from which all content is sourced or passes through. The data center holds the core DMS components and is the central point of content distribution.

Storage of VoD content can reside in the data center and be distributed through a content delivery network. Storage requirements can be quite substantial as video content moves more toward high definition. Direct attached storage or SANs can be used to provide capacity and redundancy benefits.

For more information about the data center design, see the Design Zone for the Data Center at the following URL: <http://www.cisco.com/go/designzone>.

Campus PIN Design Considerations

Within the campus network, one of the primary design considerations is multicast enablement. Live video content delivery with Digital Signs and Cast requires multicast. The Digital Media Player supports IGMPv2, and the use of SSM is recommended, if possible.

Within the campus network, the use of technologies to segment or virtualize the network can be implemented, such as VRF-Lite to segment some or all video content, if desired.


For more information about the campus design, see the Design Zone for Campus at the following URL: <http://www.cisco.com/go/designzone>.

Branch PIN Design Considerations

Branch locations can be designed with a single WAN connection or multiple WAN connections. Multiple connections combined with redundant Cisco Integrated Services Routers (ISRs) offer the highest availability. Dual links and ISRs also enable the use of additional optimization services such as PfR. Cisco WAAS and Cisco ACNS can be implemented with Cisco WAE network modules, providing an integrated approach to content distribution and optimization.

Single-Tier Branch Architecture


This profile is recommended for smaller enterprise branches that do not require platform redundancy and a large user base. This profile consists of an ISR as the access router, with an Integrated EtherSwitch network module for LAN and WAN connectivity. High availability is achieved through a T1 link with an ADSL backup. This profile is intended for branch networks that want to incorporate as many services as possible into a single platform solution. This profile is also highly cost-effective and contains the least number of devices to manage at the branch. The drawback to this profile is network resiliency and capacity planning. By having a single platform solution, there is a common point of failure. There is no platform redundancy, so a network can affect users. User capacity is also limited in this design to the number of LAN ports that the ISR platforms can support. For future growth, either an external desktop switch must be used, or another router platform is needed for additional slot capacity. [Figure 27](#) shows a single tier branch.

Figure 27 **Single Tier Branch Architecture**

Dual-Tier Branch Architecture

This profile is based on legacy branch networks that exist today. The intent of this profile is to illustrate how to apply advanced services within a branch network without requiring a forklift upgrade or the redesign of an existing network. This profile consists of two ISR access routers connected to an external switch. Dual WAN links and hardware redundancy provide a greater level of high availability compared to the single tier branch profile, at the expense of additional equipment costs and more components to manage at the branch. This branch is typical of most branches in traditional enterprise branch networks. [Figure 28](#) shows a dual tier branch.

Figure 28 **Dual Tier Branch Architecture**



For more information about the branch design, see the Design Zone for Branch at the following URL:
<http://www.cisco.com/go/designzone>.

WAN/MAN PIN Design Considerations

Cisco DMS services available to branch locations depend on the WAN type. A high speed multicast-enabled WAN is required for live video delivery to Cisco Digital Signs and Cisco Cast. Live video deliver to Cisco Show and Share clients does not have this requirement and can be implemented over a lower speed non-multicast WAN.

Enabling QoS on the WAN is essential for any critical content being delivered. Within DMS, live streaming content relies on a QoS-enabled WAN for guaranteed delivery meeting specific requirements for latency, jitter, and especially loss.

For more information about the WAN/MAN design, see the Design Zone for MAN/WAN at the following URL: <http://www.cisco.com/go/designzone>.

