



Cisco Video Surveillance Manager Solutions Reference Guide



Contents

Chapter 1: Cisco Video Surveillance Manager Overview	1-1
Solution Benefits.....	1-1
Solution Components	1-2
IP Video Surveillance	1-2
Cisco Video Surveillance Media Server.....	1-4
Cisco Video Surveillance Operations Manager	1-4
Cisco Video Surveillance Virtual Matrix.....	1-6
Cisco Video Surveillance Encoding Server	1-6
Cisco Video Surveillance Storage System	1-7
Chapter 2: Video Traffic Flows.....	2-1
Video Surveillance Media Server.....	2-1
Video Surveillance Operations Manager	2-1
Operations Manager Views	2-2
Video Surveillance Virtual Matrix Switch	2-4
Proxy Processes.....	2-5
Direct Proxy	2-5
Parent-Child Proxies.....	2-6
Chapter 3: Protocols and Features.....	3-1
Video Resolutions and Codecs	3-1
Video Resolutions.....	3-1
Digital Video Codecs	3-1
MJPEG	3-2
MPEG-4.....	3-2
IP Unicast.....	3-3
IP Multicast.....	3-4
Multicast Addressing.....	3-5
Forwarding Multicast Traffic.....	3-5
Chapter 4: Application Requirements	4-1
Time Synchronization.....	4-1
TCP/UDP Transport	4-2
MJPEG	4-2
MPEG-4.....	4-2
Required TCP/UDP Ports	4-3
Securing Video Surveillance Traffic.....	4-3
Pan-Tilt-Zoom.....	4-6
Media Server System Requirements	4-7
Operations Manager System Requirements.....	4-7
Hardware Requirements	4-7
Software Requirements	4-7
Virtual Matrix System Requirements	4-7
Encoding Server System Requirements.....	4-8
Viewer Requirements	4-8
Chapter 5: Network Deployment Models.....	5-1
Campus Networks	5-1
Hierarchical Design.....	5-2
Wide Area Networks.....	5-3
Example 1: Network Bandwidth Usage.....	5-5
Example 2: Sites with Remote Storage	5-6
Example 3: Virtual Matrix Scenario.....	5-7
Example 4: Distributed Media Servers.....	5-8
Chapter 6: Video Storage	6-1
Archives.....	6-1
Extracting Video Clips.....	6-2
Calculating Storage Requirements.....	6-3
MJPEG	6-3
MPEG-4.....	6-4
IP Camera Video Settings	6-4
Storage Systems	6-6
Appendix A: Glossary	I
Appendix B: References	V

Chapter 1: Cisco Video Surveillance Manager Overview

This chapter provides an overview of the components that are used in a Cisco® Video Surveillance Manager solution. These components are discussed in more detail throughout this guide.

The Stream Manager product offering is covered in other design guides. These design guides may be found at <http://www.cisco.com/go/srmd>

Solution Benefits

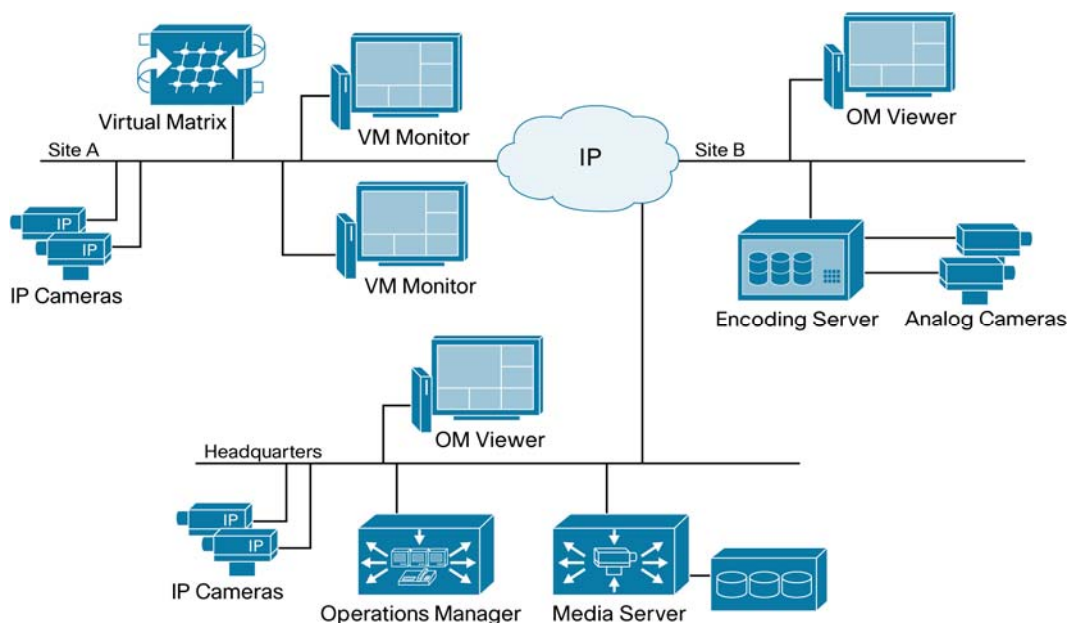
Video surveillance is a key component of the safety and security of many organizations, providing real-time monitoring of the environment, people and assets, and providing recording for investigation purposes. Benefits of Cisco's Video Surveillance solution include:

- Access to video at any time from any network location, enabling real-time incident response and investigation.
- Transfer of control and monitoring to any other point in the network in an emergency situation.
- Leverage existing investment in video surveillance and physical security equipment and technology.
- Ability for products from various vendors to interoperate in the same network.
- An open, standards-based infrastructure that enables the deployment and control of new security applications.

The Cisco Video Surveillance Solution relies on an IP network infrastructure to link all components. The designs of a highly available hierarchical network have been proven and tested for many years and allow applications to converge on an intelligent and resilient infrastructure.

Cisco offers a unique approach to moving different proprietary systems to a common IP backbone. This approach leverages other Cisco technologies, such as network security, routing, switching, network management and wireless. Video from IP cameras can now be truly converged into a robust network environment with the intelligence and flexibility provided by the Cisco infrastructure.

Figure 1 shows the Cisco Video Surveillance Manager solution using an Intelligent IP infrastructure as a transport.

Figure 1. Cisco Video Surveillance Solution

Solution Components

The following components make up the Cisco Video Surveillance Solution:

- **Cisco Video Surveillance Media Server:** The core component of the network-centric VSM. This software manages, stores, and delivers video for the network-centric video surveillance product portfolio.
- **Cisco Video Surveillance Operations Manager:** The Operations Manager authenticates and manages access to video feeds. It is a centralized administration tool for management of Media Servers, Virtual Matrixes, cameras, encoders, and viewers and for viewing network-based video.
- **Cisco Video Surveillance Virtual Matrix:** The Virtual Matrix monitors video feeds in command center and other 24-hour monitoring environments. It allows operators to control the video being displayed on multiple local and remote monitors.
- **Cisco Video Surveillance Encoding Server:** This single-box solution encodes, distributes, manages, and archives digital video feeds. Each server encodes up to 64 channels and provides up to 12 TB of storage.
- **Cisco Video Surveillance Storage System:** This complementary component allows the Media Server's internal storage to be combined with direct attached storage (DAS) and storage area networks (SANs). The Storage System allows video to be secured and accessed locally or remotely.

IP Video Surveillance

A video surveillance system that runs over an IP network infrastructure enables the video to be distributed to any number of sites, within the constraints of available bandwidth. The convergence of video surveillance into an existing IP network offers several benefits, including:

- Network-wide management. Devices are monitored over a single network for alarms or failures.

- Transfer of control and monitoring to any other point in the network in an emergency situation.
- Increased availability. IP networks offer a high level of redundancy that can extend to different physical locations.
- A system that can easily expand as business needs change.

Cisco's solution offers software and hardware to support video transmission, monitoring, recording, and management. Cisco video surveillance solutions work in unison with the advanced features and functions of the IP network infrastructure—switches, routers, and other network security devices—to enable secure, policy-based access to live or recorded video.

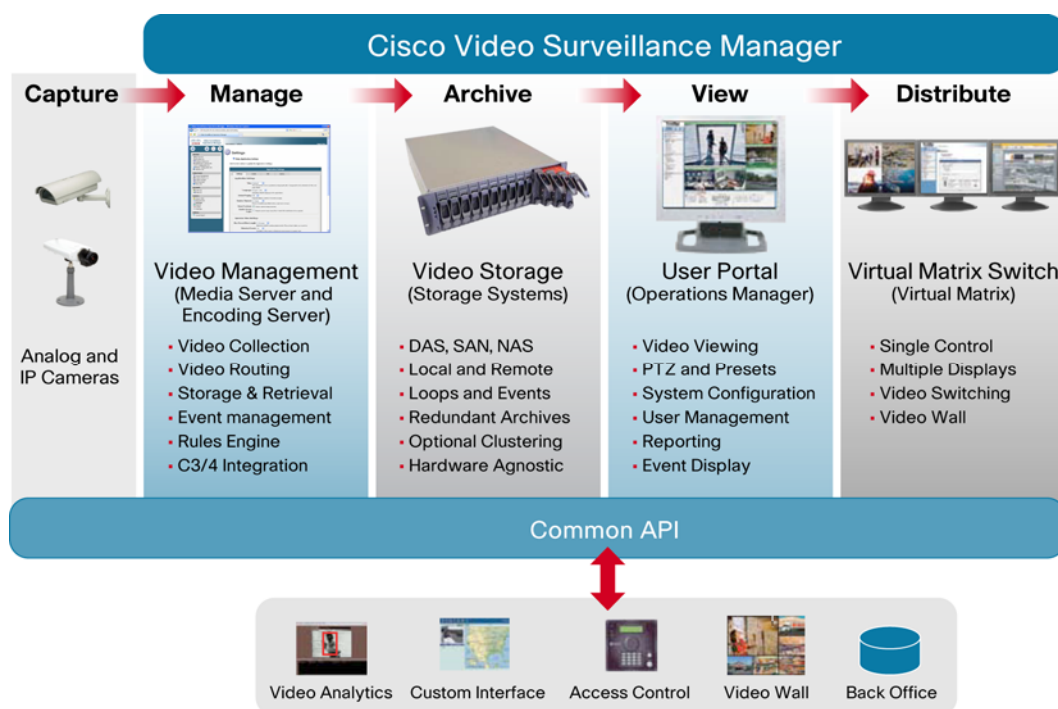
Cisco video surveillance products are deployed within the Cisco Intelligent Converged Environment architecture. Through this architecture, video can be accessed at any time from any place, enabling real-time incident response, investigation, and resolution. As an extension of the Cisco Self-Defending Network, the Cisco Intelligent Converged Environment enables customers to use existing investments in video surveillance and physical security while enhancing the safety of people and protection of assets.

The open, standards-based Cisco infrastructure enables the deployment and control of new security applications and maximizes the value of live and recorded video, interacting with multiple third-party video surveillance cameras, encoders and applications.

Unlike many other video surveillance offerings that use proprietary hardware, Cisco Video Surveillance software runs on Linux-based servers to allow for easy upgrades that support new features and support an evolving range of deployment scenarios.

Figure 2 shows the main components of the Cisco Video Surveillance Manager solution.

Figure 2. Cisco Video Surveillance Manager Solution




Cisco Video Surveillance Media Server

The Cisco Video Surveillance Media Server is the core component in the Cisco Video Surveillance Manager and performs the following networked video surveillance system functions:

- Collection and routing of video from a wide range of third-party cameras and video encoders over an IP network
- Event-tagging and recording of video for review and archival purposes
- Secure local, remote, and redundant video archive capabilities

In Figure 3 the Media Server is responsible for receiving video streams from different IP cameras and encoders and replicating them as necessary to different viewers.

Figure 3. Video Surveillance Media Server

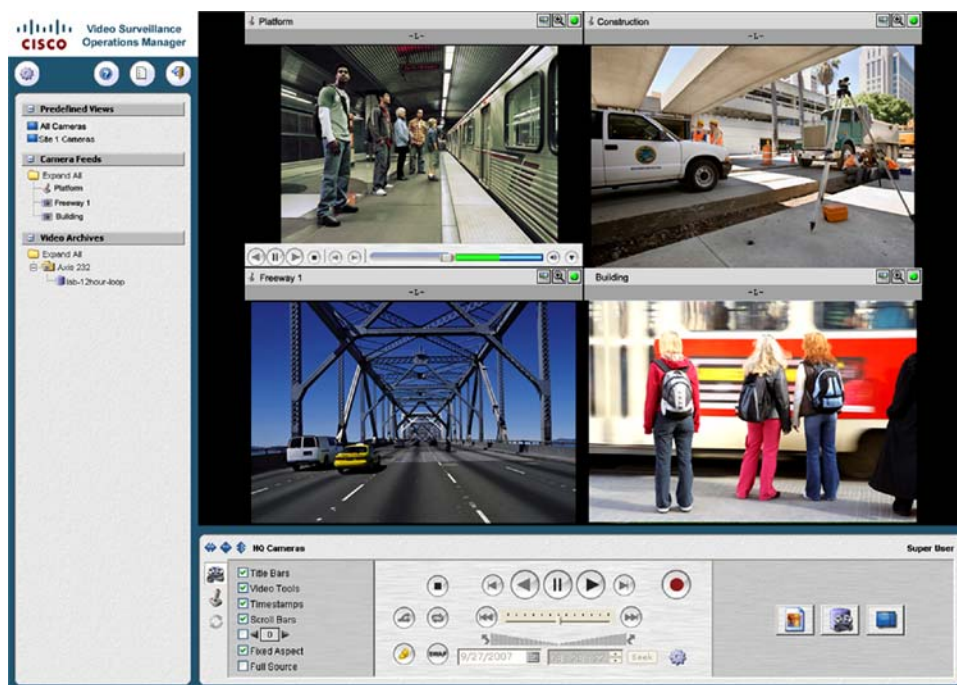


By using the power and advanced capabilities of today's IP networks, the Media Server software allows new third-party applications, additional users, cameras, and storage to be added over time. This system flexibility and scalability supports:

- Hundreds of simultaneous users viewing live or recorded video
- Standard video compression algorithms such as MJPEG, MPEG-2, and MPEG-4 simultaneously via a single Media Server or system
- Conservation of storage using events and loop-based archival options
- Integration with other security applications
- IT-caliber fault-tolerant storage for greater efficiency and easier maintenance

Cisco Video Surveillance Operations Manager

Working in conjunction with the Cisco Video Surveillance Media Server, the Cisco Video Surveillance Operations Manager enables organizations to quickly and effectively configure, manage and view video streams throughout the enterprise. Figure 4 shows the Operations Manager main screen, which is accessed via a web browser.

Figure 4. Video Surveillance Operations Manager

The Operations Manager meets the diverse needs of administrators, systems integrators, and operators by providing:

- Multiple Web-based consoles to configure, manage, display, and control video throughout a customer's IP network
- The ability to manage a large number of Cisco Video Surveillance Media Servers, Cisco Video Surveillance Virtual Matrixes, cameras and users
- Customizable interface, ideal for branded application delivery
- Management of multiple Cisco Video Surveillance Media Servers
- Encoder and camera administration
- Scheduled and event-based video recording
- Interface to Media Server and Virtual Matrix software for pushing predefined views to multiple monitors
- User and role management
- Secure login
- Live and archived video views
- Friendly user interface for PTZ controls and presets, digital zoom, and instant replay
- Event setup and event notifications
- "Record Now" feature while viewing live video
- Archive review and clipping


Cisco Video Surveillance Virtual Matrix

The Cisco Video Surveillance Virtual Matrix software allows authorized security managers and operators to select and control video displayed on any number of digital monitors on a local and remote basis.

The software also permits integrated security applications to control digital video displayed on any number of digital monitors on a local or remote basis. The Virtual Matrix software uses the IP network to provide aggregation and transmission of video from cameras and recording platforms much like the function of a classic analog video matrix switch, offering capabilities that analog switches cannot deliver. The Virtual Matrix brings complete flexibility to the delivery of live and recorded video to demanding command centers, providing high availability access to network video for 24x7 monitoring applications.

Figure 5 demonstrates how operators can choose from any number of available cameras to be displayed on any system monitors within any custom video display patterns. The VM Monitors display the video streams defined on the Operations Manager on a single display or video wall displays. The Virtual Matrix also integrates with other systems to automatically display video in response to user-defined event triggers. These triggers can include access control and fire systems in buildings, outdoor motion sensors, or even radar systems for military applications.

Figure 5. Virtual Matrix Switch



The Virtual Matrix software provides:

- Ability to access and display video in remote command centers
- Easy integration with other intelligent systems
- Flexible delivery of both live and archived video
- Ability to control multiple video displays from a single station
- Display distribution to digital video walls

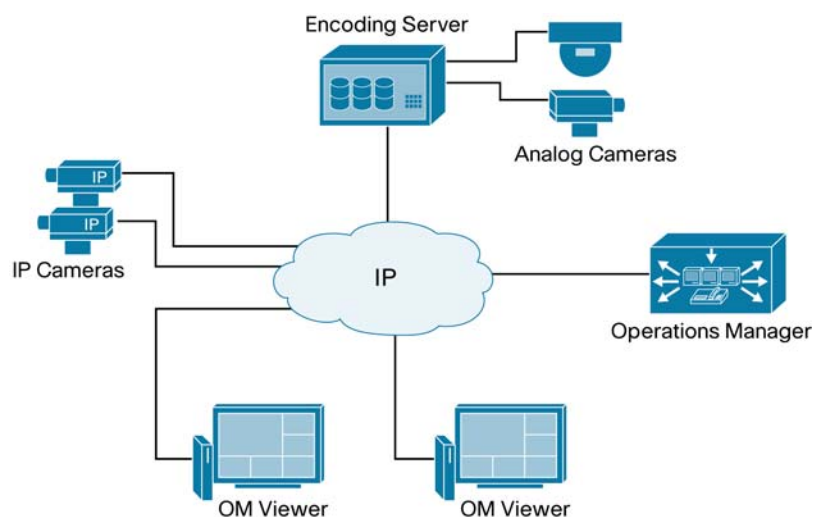
Cisco Video Surveillance Encoding Server

The Cisco Video Surveillance Encoding Server is an all-in-one appliance that encodes, distributes, manages, and archives digital video feeds. Each server encodes up to 64 channels and provides up to 12 TB of storage.

The Encoding Server can combine multiple video codecs in a single Encoding Server, including MJPEG and MPEG-4. With the addition of the Cisco Video Surveillance Operations Manager, the Encoding Server also provides administrators and operators with multiple Web-based consoles to configure, manage, display, and control video.

Figure 6 shows an Encoding Server receiving video streams directly from IP and analog cameras. The analog video streams are encoded into a video stream that can be archived and distributed to the different viewers. The Encoding Server acts as the Media Server and Encoding Server simultaneously.

Figure 6. Video Encoding Server



The Video Surveillance Encoding Server provides:

- Flexibility to use a broad array of analog cameras and IP cameras
- Simultaneous MJPEG and MPEG-4 encoding
- PTZ and alarm inputs
- Motion detection
- Scalable deployment with multiple sites, cameras, users, and storage
- Archives at various frame rate, duration, and location
- Efficient redundant multi-site archiving that conserves bandwidth
- Ability to connect to external storage
- The system is able to encoder up to:
 - 64 CIF channels
 - 32 2CIF or
 - 4 4CIF channels


Cisco Video Surveillance Storage System

As an integral component of the Cisco Video Surveillance Manager solution, the Cisco Video Surveillance Storage System provides flexible options for storing video and audio using cost-effective, IT-caliber storage devices.

The Storage System allows the Cisco Video Surveillance Media Server's internal storage to be combined with direct attached storage (DAS) and storage area networks (SANs). As a result, video can be efficiently secured and accessed wherever it is needed, locally or remotely.

Figure 7 shows a single server providing the functions of a Media Server, Operations Manager and Virtual Matrix that is also connected to an external storage system.

Figure 7. VS Storage System



The Storage System enables operations managers to institute video lifecycle management rules to ensure availability of the data. Video can be stored in redundant systems or in remote long-term archives to ensure the video is available when required.

Some of the Storage System features are:

- Media Server maintains internal storage up to 24 TB
- DAS arrays support up to 42 TB per array, 420 TB per rack
- Support for 3rd party SANS
- RAID 5 configuration available

Chapter 2: Video Traffic Flows

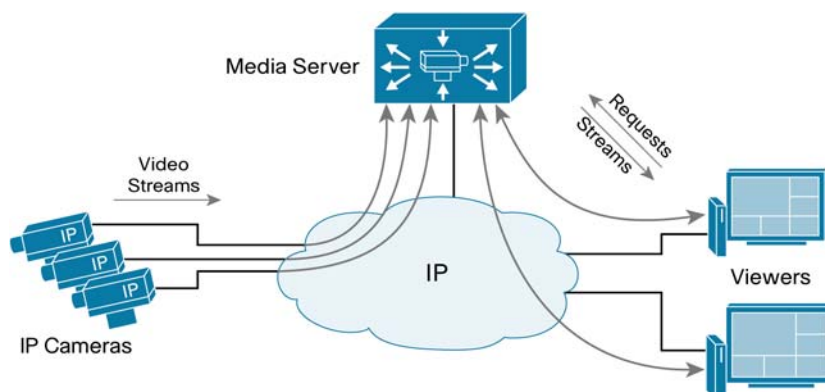
Each Video Surveillance Manager application plays a unique role in the deployment of a complete video streaming solution. When deploying and operating a Video Surveillance Manager environment, it is important to understand the video traffic flows of each application and how they interact with the system as a whole.

Video Surveillance Media Server

The Video Surveillance Media Server is the core component of the solution, providing for the collection and routing of video from IP cameras to viewers or other Media Servers. The system is capable of running on a single physical server or distributed across the network, scaling to handle thousands of cameras and users.

Figure 8 shows how IP cameras or encoders send a single video stream to the Media Server. The Media Server is responsible for distributing live and archived video streams to the viewers simultaneously over an IP network.

Figure 8. Media Server



For archive viewing, the Media Server receives video from the IP camera or encoder continuously (as configured per the archive settings) and only sends video streams to the viewer when requested.

In environments with remote branch locations, this becomes very efficient since traffic only needs to traverse the network when requested by remote viewers. Branch office traffic remains localized and does not have to traverse wide area connections unless is requested by users other users.

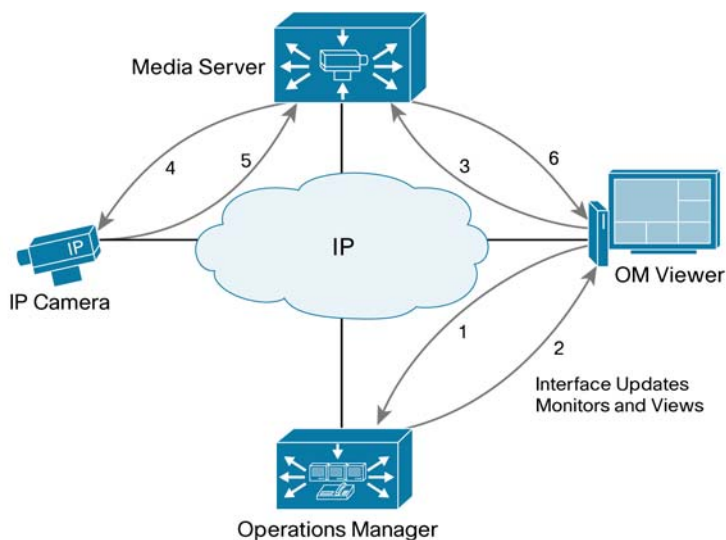
Video requests and video streams are delivered to the viewer using HTTP traffic (TCP port 80).

Video Surveillance Operations Manager

Viewers can access the Operations Manager via their Web browser. The Operations Manager is responsible for delivering a list of resource definitions, such as camera feeds, video archives and predefined views to the viewer. Once this information is provided to the viewer, the viewer communicates with the appropriate Media Server to request and receive video streams.

Figure 9 shows the traffic flow of video request from the viewer to the Operations Manager. The viewer is responsible for contacting the proper Media Server to receive video streams.

Figure 9. Operations Manager Traffic Flows



When the OM Viewer requests a video stream, the following steps occur as shown in Figure 9.

1. The user accesses the Operations Manager screen via an ActiveX web browser. This traffic can be over TCP port 80 (HTTP) or 43 (HTTPS).
2. The OM Viewer receives a complete list of resources, such as camera feeds, views and monitors. This information is sent each time the client starts or switches to the operator view. Since the OM Viewer has a complete list of resources, the operator may choose to view live or recorded video from any camera feed or predefined views.
3. The OM Viewers selects a video feed that is served by the Media Server and contacts the Media Server directly over TCP port 80.
4. The Media Server is the direct proxy for the IP camera and requests the video stream from the camera. This communication can be TCP, UDP, or multicast as configured by the Operations Manager.
5. The camera provides the video stream to the Media Server.
6. The Media Server replicates the requested video feed to the OM Viewer using IP unicast over TCP port 80. The connection remains active until the OM Viewer selects a different video feed.

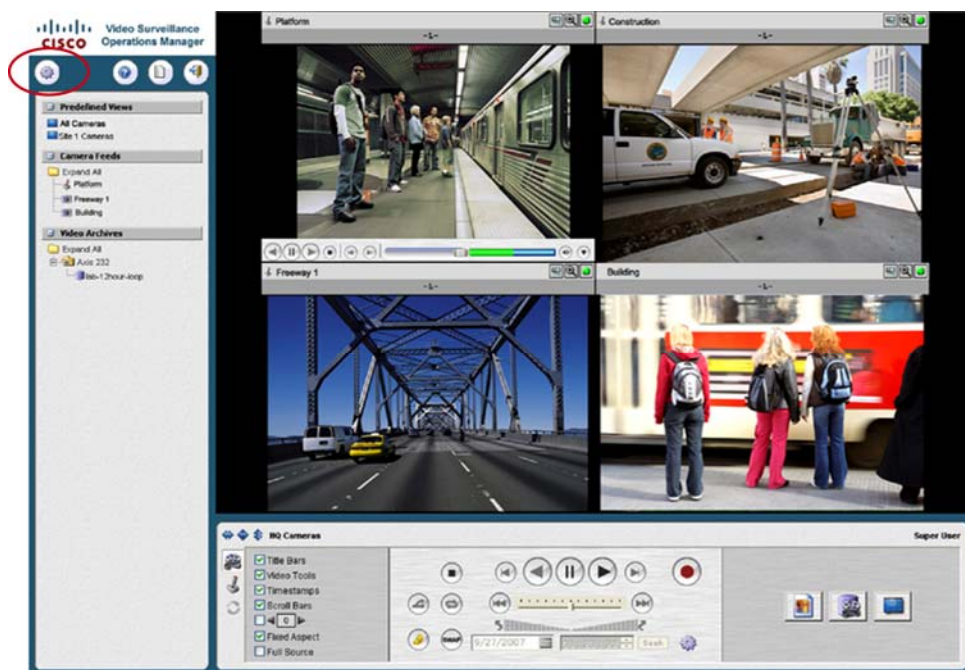
If another OM Viewer requests the video from the same IP Camera, the Media Server simply replicates the video stream as requested, and no additional requests are made to the camera.

Operations Manager Views

Based on user permissions, the user interface supports two main views:

Operator View: Allows operators to view and manage live and recorded video feeds based on user authorization. By default, the operator view is displayed during initial launch. When displayed, use to display the administrator view. Figure 10 shows the operator view.

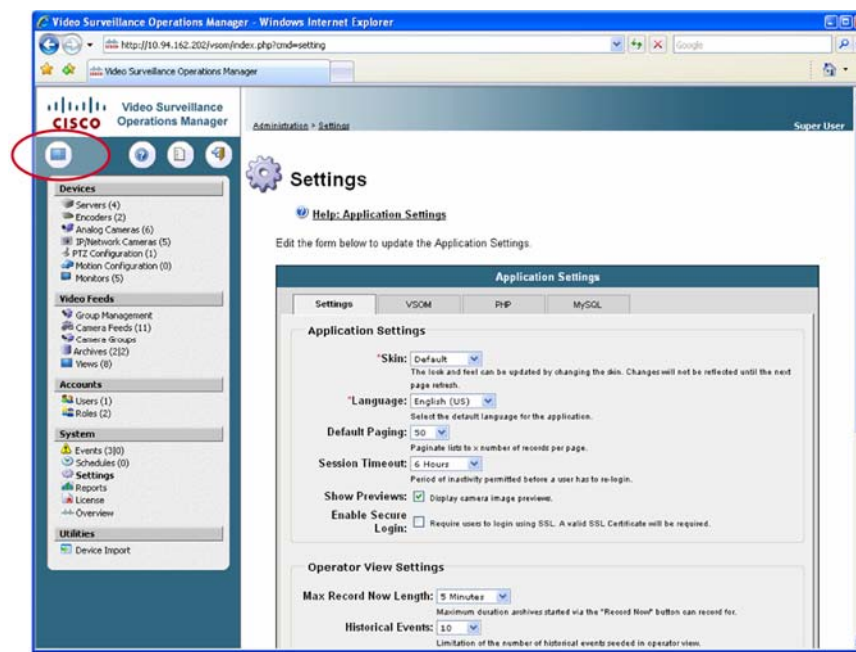
Figure 10. Operator View



Administrator view: Within this interface, administrators have access to configure the full solution, including servers, cameras, archives, user roles, etc. When displayed, use to switch back to the operator view.

Figure 11 shows the administrator view with the available resources displayed on the toolbar.

Figure 11. Administrator View



While the OM Viewer is in operator view mode, the Operations Manager sends a keep-alive message to the OM Viewer every two seconds. If changes are made to the list of resources (devices, monitors, feeds, etc.) the Operations Manager notifies the OM viewer during the normal updates.

Figure 12 shows the Refresh Signal button enabled to indicate that new resources are now available to the OM Viewer.

Figure 12. Refresh Signal Icon



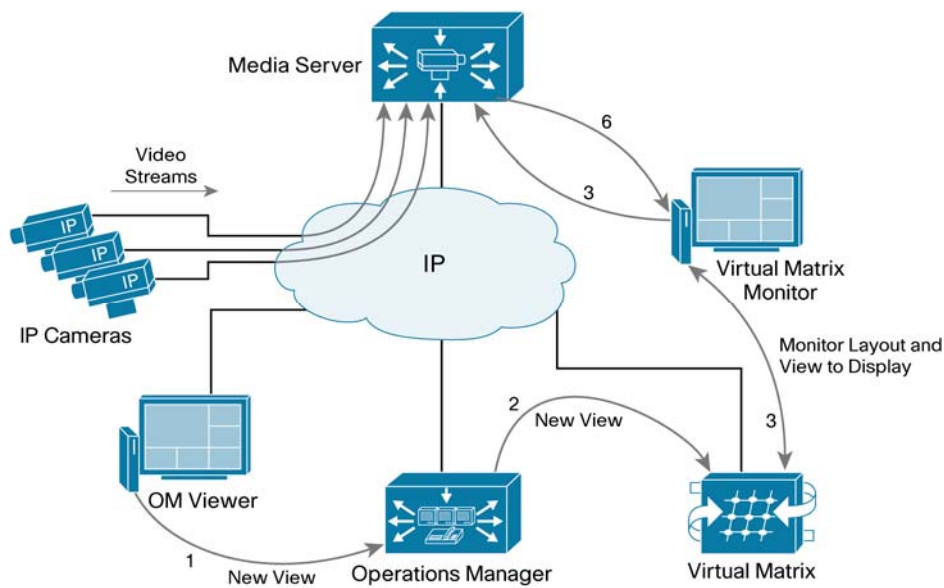
Keep-alive messages are not sent while the viewer is in the Administrator view.

Video Surveillance Virtual Matrix Switch

The Virtual Matrix server is responsible for providing detailed monitor layout to the Virtual Matrix Monitors and the position of each camera feed on the screen. A single Virtual Matrix server can be deployed to support a large number of Virtual Matrix monitors since the communication between the monitors and the server is required only during the initialization or when a new view is pushed to the monitor.

Once the monitor layout and views are pushed to the monitors, the monitors are responsible for contacting the appropriate Media Server(s) to request video streams.

Figure 13. Virtual Matrix



When requesting a new view for the Virtual Matrix Monitor, the following steps occur as shown in Figure 13.

1. The OM Viewer selects a new view to be displayed by the Virtual Matrix Monitor. The request is received by the Operations Manager
2. The Operations Manager sends the layout update to the Virtual Matrix server
3. The Virtual Matrix server pushes the new layout to the Virtual Matrix Monitor
4. Once the Virtual Matrix Monitor learns the new layout and the cameras to display, it contacts the appropriate Media Servers to request video streams
5. Video streams are sent from the Media Server directly to the Virtual Matrix Monitor.

The Virtual Matrix server sends a keep-alive message to the Virtual Matrix Monitor every three minutes to confirm that the display is still active.

Proxy Processes

Proxy processes allow for the replication of individual video feeds at different frame rates for multiple users or system processes. When a video feed is first registered with the Media Server, the server creates a proxy or process to manage connections and video streams from video sources into the Media Server.

The Media Server can support a large number of proxy processes on a single server or an architecture with distributed proxy processes on multiple Media Servers.

There are two types of proxy processes:

- Direct Proxy
- Parent-Child Proxy

Direct Proxy

A direct proxy is the process created on the Media Server to maintain connectivity with the edge device (IP camera or encoder). The proxy is capable of requesting video from the edge device with different video configurations such as frame rate and video resolution. One direct proxy exists for a given video stream.

In the example in Figure 14, the Media Server maintains connectivity and receives video from four different IP cameras. The Media server is responsible for replicating the video feeds to four different viewers.


Figure 14. Direct Proxy

Table 1 shows the active processes from Figure 14. The four OM Viewers are viewing live video from different cameras; each of the viewers is receiving the video feeds directly from the Media Server. The Media server is receiving four unique video streams, replicating them a total of 11 times.

Table 1. Active Processes

Video Source	Active Viewers	Number of Active Streams from the Media Server to Clients
Camera A	Viewer 1, Viewer 2, Viewer 3, Viewer 4	4
Camera B	Viewer 2, Viewer 2, Viewer 4	3 (two streams to Viewer 2)
Camera C	Viewer 1, Viewer 2, Viewer 4	3
Camera D	Viewer 1	1
Total Streams		11

Parent-Child Proxies

Video feeds can originate from the direct proxy or from a different Media Server. A proxy video feed can be the parent to another video feed served by a different Media Server. Parent proxies may be from remote or local hosts and may be nested in a hierarchy with inheritance rights.

A direct proxy becomes a parent when a child proxy is created. A child proxy receives its video directly from a parent proxy. A child proxy has the same resolution, quality, and media type of its parent, but in the case of MJPEG video streams, a lower frame rate may be configured for the child feed.

Parent-child proxies allow for more efficient network utilization by distributing video feeds closer to the viewers. This is very important in environments with remote branch offices or with limited bandwidth available for video delivery. By replicating a single video feed to a location with several viewers, the bandwidth requirements throughout the network are reduced.

In order to conserve bandwidth, the child process connects to the parent source only when video streaming is requested by a viewer.

In Figure 15, Media Server MS1 is acting as the parent for two feeds that are served by Media Server MS2. Video feeds from cameras A and B are replicated to Media Server MS2, which in turn can be served to a large number of users or other child feeds.

The environment in Figure 15 has generated a total of six proxy processes:

- Media Server MS1 is the direct proxy to four edge devices but also replicates eleven different video streams to other viewers or child feeds.
- Media Server MS2 has created two child proxy feeds, child A and child B. These feeds can be propagated to any viewers locally on Site B, reducing the bandwidth requirements across the wide area connections.

Figure 15. Parent-Child Proxy

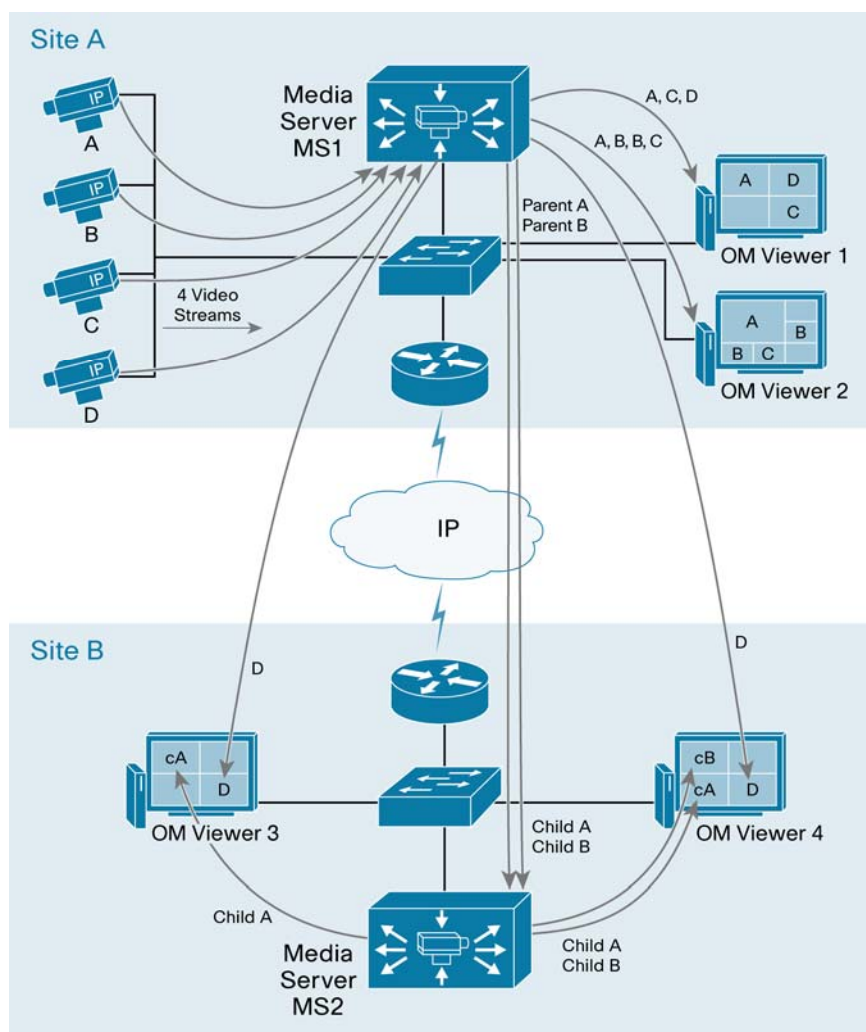


Table 2 shows the different streams required to distribute the video feeds from Figure 15.

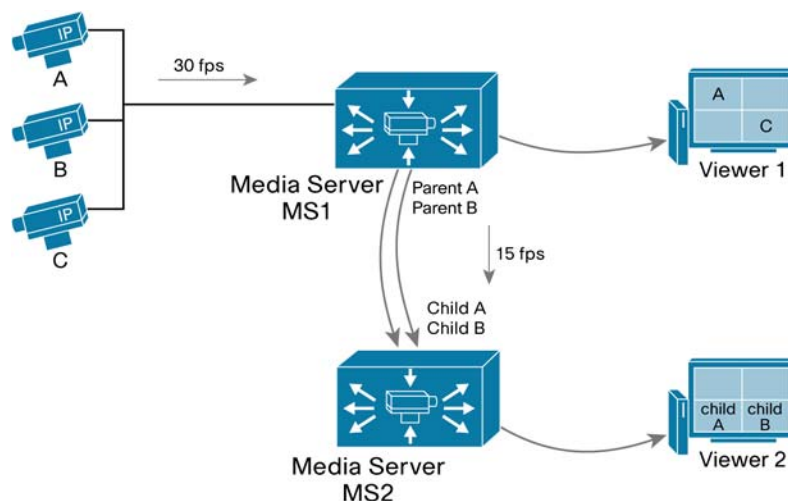
Table 2. Parent-Child Proxies

Video Source	Active Viewers	Number of Active Streams from the Media Server to Clients
Camera A	Viewer1, Viewer 2, MS2	3
Camera B	Viewer 2, Viewer 2, MS2	3 (two streams to OM Viewer 2)
Camera C	Viewer 1, Viewer 2	2
Camera D	Viewer 1, Viewer 3, Viewer 4	3
Parent A	MS2	1
Parent B	MS2	1
Child A	Viewer 3, Viewer 4	2
Child B	Viewer 4	1
Site B: Local Streams		3
Site B: Remote Streams		4

Since Media Servers do not provide transcoding features, the video quality and resolution remain the same for all child feeds. When using MJPEG streams, the frame rate can be lowered to reduce the bandwidth utilization by child feeds. Figure 16 shows an example of how frame rates can be lowered between parent and child feeds. The original video feed for all Cameras is 30 fps, but is reduced to 15 fps by child feeds A and B in order to conserve bandwidth.

The example in Figure 16 also shows how video feeds can be replicated indefinitely between Media Servers. In this example, Media Server MS1 is the direct proxy to three IP camera feeds. In turn, two of the feeds are parents for feeds going into Media Server MS2.

Figure 16. MJPEG Frame Rate Reduction for Child Feeds



Note: The frame rate of a MJPEG child feed can only be equal to or lower than the parent feed

Chapter 3: Protocols and Features

The Cisco Video Surveillance Media Server supports IP endpoints that use Motion JPEG (MJPEG) or MPEG-4 codec technology. Both types of codecs have advantages and disadvantages when implemented in a video surveillance system. A system administrator may choose to use MJPEG on certain cameras and MPEG-4 on others depending on system goals and requirements.

Video Resolutions and Codecs

Video Resolutions

Video surveillance solutions use a set of standard resolutions. NTSC (National Television System Committee) and PAL (Phase Alternating Line) are the two prevalent analog video standards.

PAL is used mostly in Europe, China and Australia and specifies 625 lines per frame with a 50 Hz refresh rate. NTSC is used mostly in the United States, Canada, and portions of South America and specifies 525 lines per frame with a 59.94 Hz refresh rate.

These video standards are displayed in interlaced mode, which means that only half of the lines are refreshed in each cycle. Therefore, the refresh rate of PAL translates into 25 complete frames per second and NTSC translates into 30 (29.97) frames per second. Table 3 shows resolutions for common video formats.

Table 3. Video Resolutions (in pixels)

Format	NTSC-Based	PAL-Based
QCIF	176 × 120	176 × 144
CIF	352 × 240	352 × 288
2CIF	704 × 240	704 × 288
4CIF	704 × 480	704 × 576
D1	720 × 480	720 × 576

Note that the linear dimensions of 4CIF are twice as big as CIF. As a result, the screen area for 4CIF is four times that of CIF with higher bandwidth and storage requirements. The 4CIF and D1 resolutions are almost identical and sometimes these terms are used interchangeably.

Note: IP camera vendors may use different video resolutions. The Cisco Video Surveillance Manager solution supports the format delivered by the camera

Digital Video Codecs

A codec is a device or program that performs encoding and decoding on a digital video stream. In IP networking, the term frame refers to a single unit of traffic across an Ethernet or other Layer 2 network. In this document, frame primarily refers to one image within a video stream. A video frame can consist of multiple IP packets or Ethernet frames.

A video stream is fundamentally a sequence of still images. In a video stream with fewer images per second, or a lower frame rate, motion is normally perceived as choppy or broken. At higher frame rates up to 30 frames per second, the video motion appears smoother however, 15 frames per second video may be adequate for viewing and recording purposes.

Some of the most common digital video formats include:

- **Motion JPEG (M-JPEG)** is a format consisting of a sequence of compressed Joint Photographic Experts Group (JPEG) images. These images only benefit from spatial compression within the frame; there is no temporal compression leveraging change between frames. For this reason, the level of compression reached cannot compare to codecs that use a predictive frame approach.
- **MPEG-1 and MPEG-2** formats are Discrete Cosine Transform based with predictive frames and scalar quantization for additional compression. They are widely implemented, and MPEG-2 is still in common use on DVD and in most digital video broadcasting systems. Both formats consume a higher level of bandwidth for a comparable quality level than MPEG-4.
- **MPEG-4** introduced object-based encoding, which handles motion prediction by defining objects within the field of view. MPEG-4 offers an excellent quality level relative to network bandwidth and storage requirements.
- **H.264** is a technically equivalent standard to MPEG-4 part 10, and is also referred to as Advanced Video Codec or AVC. This emerging new standard offers the potential for greater compression and higher quality than existing compression technologies.

MJPEG

An MJPEG codec transmits video as a sequence of JPEG (Joint Photographic Experts Group) encoded images. Each image stands alone without the use of any predictive compression between frames. MJPEG is less computationally intensive than predictive codecs such as MPEG-4, so can be implemented with good performance on less expensive hardware. MJPEG can easily be recorded at a reduced frame rate by only sampling a subset of a live stream. For example, storing every third frame of a 30 frame per second video stream will result in a recorded archive at 10 frames per second.

MJPEG has a relatively high bandwidth requirement compared to MPEG-4. A 640x480 VGA resolution stream running at 30 frames per second can easily consume 5-10 Mbps. The bandwidth required is a function of the complexity of the image, in conjunction with tuning parameters that control the level of compression. Higher levels of compression reduce the bandwidth requirement but also reduce the quality of the decoded image. Since there is no predictive encoding between frames, the amount of motion or change in the image over time has no impact on bandwidth consumption.

MPEG-4

An MPEG-4 codec utilizes prediction algorithms to achieve higher levels of compression than MJPEG while preserving image quality. Periodic video frames called I-frames are transmitted as complete, standalone JPEG images similar to an MJPEG frame and are used as a reference point for the predictive frames. The remaining video frames (P-frames) contain only information that has changed since the previous frame.

To achieve compression, MPEG-4 relies on the following types of video frames:

- **I-frames** (intraframes, independently decodable): These frames are also referred to as key frames and contain all of the data that is required to display an image in a single frame.
- **P-frames** (predictive or predicted frames): This frame type contains only image data that

has changed from the previous frame.

- **B-frames** (bi-directional predictive frames): This frame type can reference data from both preceding frames and future frames. Referencing of future frames requires frame reordering within the codec.


The use of P-frames and B-frames within a video stream can drastically reduce the consumption of bandwidth compared to sending full image information in each frame. However, the resulting variance of the video frames' size contributes to the fluctuation in the bandwidth that a given stream uses. This is the nature of most codecs because the amount of compression that can be achieved varies greatly with the nature of the video source.

IP Unicast

Applications that rely on unicast transmissions send a copy of each packet between one source address and one destination host address. Unicast is simple to implement but hard to scale if the number of sessions is large. Since the same information has to be carried multiple times, the impact on network bandwidth requirements may be significant.

The communication between the Media Server and the viewers is always via IP Unicast, making the Media Server responsible for sending a single stream to each viewer. The example in Figure 17 shows five viewers requesting a single video stream from the Media Server. Assuming a single 1Mbps video feed, the bandwidth requirements are noted throughout each network link.

Figure 17. IP Unicast Traffic



Note: The Media Server only supports IP Unicast between the Media Server and the viewers

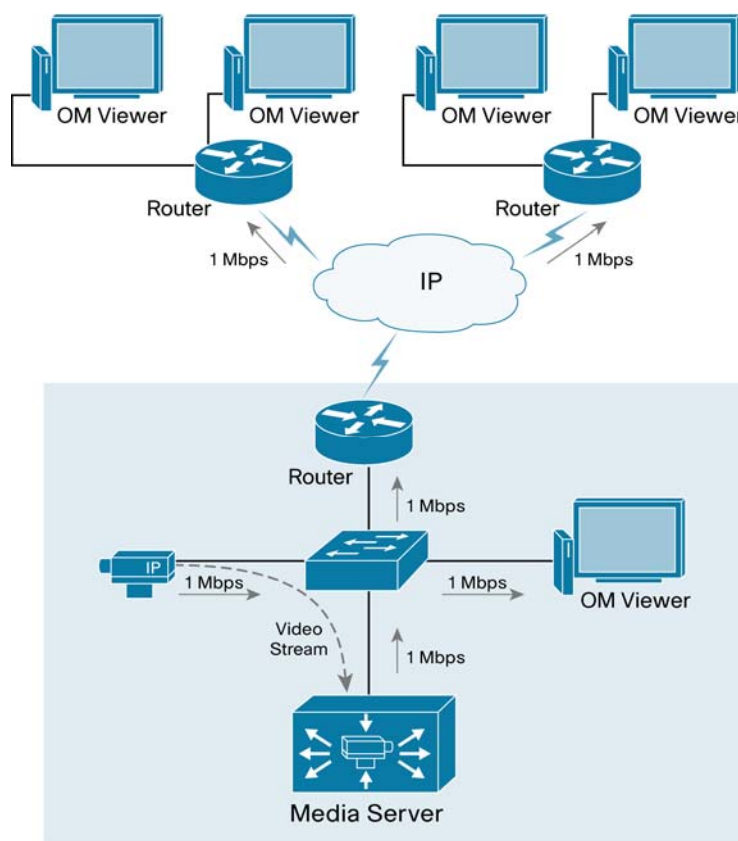
IP Multicast

In IP Multicast transmissions, a host sends one copy of each packet to a special address that can be used by several hosts interested in receiving the packets. Those hosts are members of a designated multicast group and can be located anywhere on the network. Using IP multicast to transmit video traffic reduces the overall network load and minimizes the impact on the source of the video from unnecessary replication of a common data stream.

By using multicast protocols, the hosts that want to receive traffic from a multicast group can join and leave the group dynamically. Hosts can be members of more than one group and must explicitly join a group before receiving content. Since IP multicast traffic relies on UDP, which, unlike TCP, has no built-in reliability mechanism such as flow control or error recovery mechanisms, tools such as QoS can improve the reliability of a multicast transmission.

Some edge devices may communicate with the Media Server using Unicast or Multicast communications. The use of IP Multicast offers some benefits when a video stream is to be archived by several Media Servers since only a single stream is required from the IP camera or encoder.

Figure 18 shows an example where a single multicast stream is generated by an IP camera and archived by two Media Servers. The Media Servers propagate the video streams to the viewers using IP Unicast transmission. Using multicast protocols, Cisco routers and switches replicate the video stream to only the segments and hosts that require it, using approximately 8 Mbps of bandwidth throughout the network.

Figure 18. IP Multicast

Note: The Media Server only supports IP Unicast between the Media Server and the viewers, but it can communicate via IP Multicast with edge devices that support IP Multicast

Multicast Addressing

IP multicast uses the Class D range of IP addresses, from 224.0.0.0 through 239.255.255.255. Within this range, several addresses are reserved by the Internet Assigned Numbers Authority (IANA):

- **224.0.0.0 through 224.0.0.255:** Link-Local addresses. Used by network protocols only in a local segment.
- **224.0.1.0 through 238.255.255.255:** Globally scoped addresses. Can be routed across the Internet or any organization. They are unique and globally significant.
- **239.0.0.0 through 239.255.255.255:** Used in private domains and not routed between domains. Similar to the IP address range from RFC1918.

Forwarding Multicast Traffic

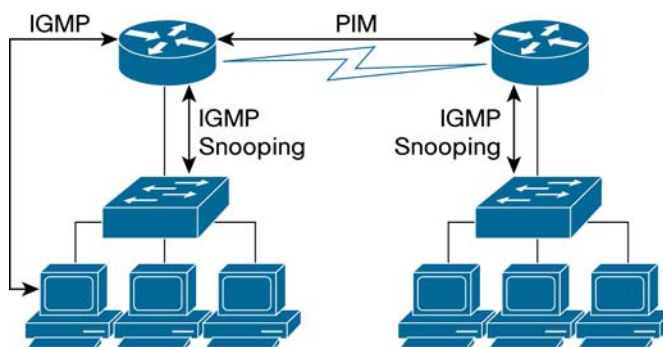
Forwarding multicast packets through a network is different than unicast routing. With unicast traffic, routers consider the destination address and how to find the single destination host. In multicast traffic, the source sends traffic to a multicast group address, which in turn can be reached by multiple hosts or receivers.

Routers rely on distribution trees to reach all multicast receivers. The two types of multicast trees are:

- **Source trees:** The root is located at the multicast source and a tree to all receivers is formed via the shortest path tree (SPT).
- **Shared trees:** The root is not necessarily the multicast source. The tree is shared by all sources relying on a defined common root. This shared root is the Rendezvous Point (RP).

Similar to IP unicast, IP multicast traffic uses its own Layer 2, management and routing protocols. Figure 19 shows the interaction between these different protocols.

Figure 19. Interaction Between IGMP and PIM



- PIM is the multicast routing protocol that is responsible for building multicast delivery trees and for enabling multicast packet forwarding.
- IGMP is used by hosts to dynamically register to multicast groups. The communication occurs between the router and the host.
- IGMP snooping is used to prevent multicast flows from flooding all ports on a VLAN. It does so by monitoring the Layer 3 IGMP packets.

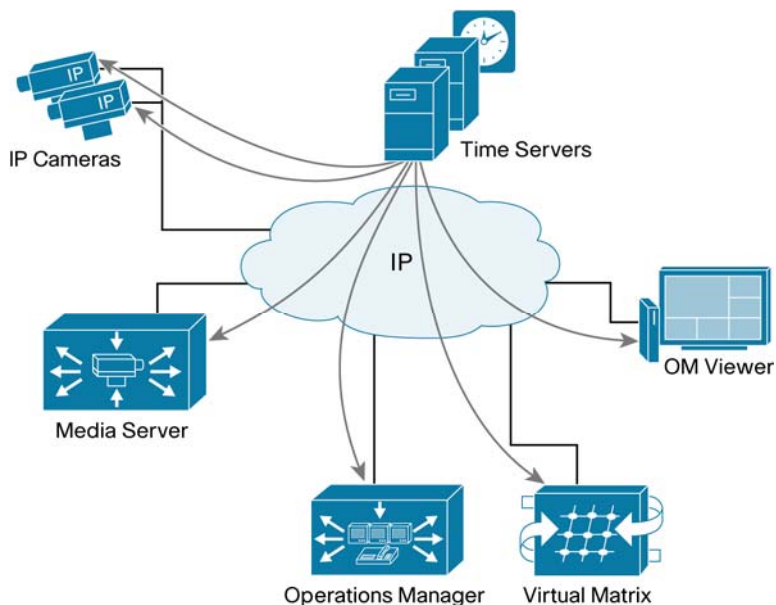
Chapter 4: Application Requirements

The Cisco Video Surveillance Media Manager application requires some network protocols and minimum hardware requirements to operate efficiently. This section focuses on some of the most important protocols and outlines the hardware requirements for the current software release.


Time Synchronization

The Network Time Protocol (NTP) is widely used to synchronize clocks of viewers, application servers, routers and other network elements with a reliable time source. The Cisco Video Surveillance Manager solution relies on NTP to synchronize the time of all its applications and viewers. Clock synchronization is critical when retrieving previously recorded video streams. Figure 20 shows how the NTP servers propagate the current time to IP cameras, viewers and application servers.

Figure 20. Network Time Protocol (NTP)



The application servers should be configured to receive time from an NTP server. If an external NTP server is not available, one of the application servers may be configured to act as the NTP server for all devices in the environment, as shown in Figure 21. It is a good practice to have more than one server providing the current time.

Figure 21. Network Time Protocol

The viewer's workstations should also be synchronized with a reliable time source.

NTP servers keep time in Universal Time (UTC), and each device on the network is configured for the proper geographical time zone. The conversion to the proper local time is handled by the operation system on each device.

TCP/UDP Transport

IP cameras and encoders communicate with the Media Server in different ways, depending on the manufacturer. Some edge devices may support only MJPEG over TCP (Transmission Control Protocol), while others may also support MPEG-4 over UDP (User Datagram Protocol).

MJPEG

MJPEG is typically transported via the TCP protocol. TCP provides guaranteed delivery of packets by requiring acknowledgement by the receiver. Packets that are not acknowledged will be retransmitted. The retransmission of TCP can be beneficial for slightly congested networks or networks with some level of inherent packet loss such as a wireless transport. Live video rendering at the receiving end may appear to stall or be choppy when packets are retransmitted, but with the use of MJPEG each image stands alone so the images that are displayed are typically of good quality.

MPEG-4

MPEG-4 video is typically transmitted over UDP or RTP (Real-time Transport Protocol). UDP does not guarantee delivery and provides no facility for retransmission of lost packets. RTP/UDP transport is most suitable for networks with very little packet loss and bandwidth that is guaranteed through Quality of Service (QoS) mechanisms. MPEG-4 over RTP/UDP is relatively intolerant to packet loss, if there is loss in the stream there will typically be visible artifacts and degradation of quality in the decoded images. UDP transport does provide the option of IP Multicast delivery, where a single stream may be received by multiple endpoints. In an IP Multicast configuration, the internetworking devices handle replication of packets for multiple recipients. This reduces the processing load on the video encoder or IP camera and can also reduce bandwidth consumption


on the network.

Some IP cameras and encoders also provide for TCP transport of MPEG-4. TCP encapsulation can be beneficial for networks with inherent packet loss. TCP may be useful especially for fixed cameras and streams that are only being recorded and not typically viewed live. TCP transport induces a little more latency in the transport due to the required packet acknowledgements, so may not be a desirable configuration for use with a PTZ controlled camera.

Required TCP/UDP Ports

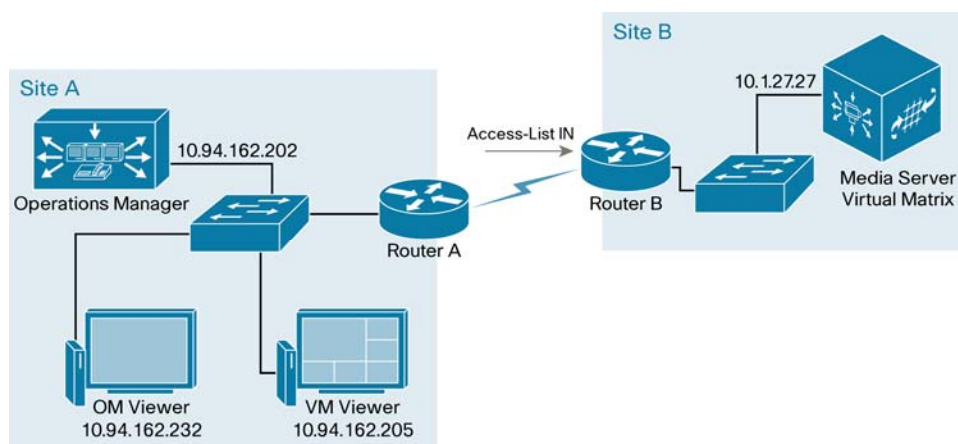
The example in Figure 22 shows that the communication between the Media Server and viewers relies on TCP port 80 (HTTP) while the communication between edge devices and the Media Server may vary. The communication between the Virtual Matrix Server and the VM Monitor is typically over TCP port 1066 while the communication between the Virtual Matrix Server and the Operations Manager is typically over TCP port 8086.

Figure 22. TCP/UDP Ports



Securing Video Surveillance Traffic

Figure 23 shows an example where the viewers and the Operations Manager are separated from the Media Server and Virtual Matrix via Cisco IOS routers. The Media Server and Virtual Matrix applications are installed on the same server. In this example IOS access lists are utilized but other security devices, such as the Cisco ASA 5500 Adaptive Security Appliance may be used.

Figure 23. Traffic Filtering

The following access lists shows simple ways to block traffic to these resources and control what devices can receive video streams. The same examples can be used if a firewall is in place to protect video streams.

Note: The syntax may vary when using different IOS or firewall devices.

The following access list may be applied to Router B to allow traffic destined for the servers on Site B.

```
interface Multilink1
 ip address 10.1.20.2 255.255.255.252
 ip access-group ALLOW_VMS in
 ppp multilink
 ppp multilink group 1
!
!
ip access-list extended ALLOW_VSM_TRAFFIC
 permit tcp any host 10.1.27.27 eq www
 permit tcp any host 10.1.27.27 eq 1066
 permit tcp any host 10.1.27.27 eq 8086
 deny ip any any
```

The access-list is applied to the Multilink1 interface on the incoming direction and specifies what traffic can reach the server Site B. This access list allows any hosts to reach the 10.1.27.27 server and blocks all other types of traffic. Access lists have an implicit deny statement at the end of the list in order to block traffic types that were not explicitly permitted with the access list.

The access-list only allows the following traffic types to reach the server with IP address 10.1.27.27:

- HTTP traffic. This traffic is required for all viewers to reach the Media Server and receive video streams
- TCP port 1066, required by the VM monitor client to reach the Virtual Matrix server
- TCP port 8086, required by the Operations Manager to reach the Virtual Matrix server

The following example shows an access-list with more granular statements to allow traffic only from specific hosts and block any other hosts from access video streams.

```

interface Multilink1
 ip address 10.1.20.2 255.255.255.252
 ip access-group ALLOW_VSM_HOSTS in
 ppp multilink
 ppp multilink group 1
!
!
ip access-list extended ALLOW_VSM_HOSTS
 permit tcp host 10.94.162.202 host 10.1.27.27 eq 8086
 permit tcp host 10.94.162.205 host 10.1.27.27 eq 1066
 permit tcp host 10.94.162.232 host 10.1.27.27 eq www
 permit tcp host 10.94.162.205 host 10.1.27.27 eq www
 deny ip any any

```

This access list is also applied to the incoming traffic of Router B and only allows traffic from the hosts on Site A to reach the server resources at 10.1.27.27. This example allows the network administrator to ensure that video streams reach only the intended recipients.

The diagram in Figure 23 does not show IP Cameras or encoders but the traffic from those devices can also be blocked or configured to reach only the intended Media Server acting as the direct proxy.

For MJPEG transmission, Media Servers communicate with edge devices using TCP port 80 (HTTP) but in some cases a different transmission and protocol may be selected.

When using MPEG-4 video transmission, the Media Server communicates with cameras using unique UDP port numbers. The ports listed in Table 4 show the UDP ports used by different manufacturers.

Table 4. UDP Ports Used for MPEG-4

Edge Device Model	UDP Ports
Axis	16400
Bosch	6001–60001
Cisco	65000
Cornet	16400
Mango	2000
Panasonic	1024
Smartsigth	19000
Sony	1024
Teleste	
MPEG2	16400
MPEG4	16100–65534
Vbrick	18000

The network path must allow for the appropriate TCP and UDP ports to travel freely between edge devices, application servers, and viewing stations. If Access Control Lists (ACL) or firewalls are deployed between the devices, they should be configuration to allow traffic between all video surveillance devices.

Pan-Tilt-Zoom

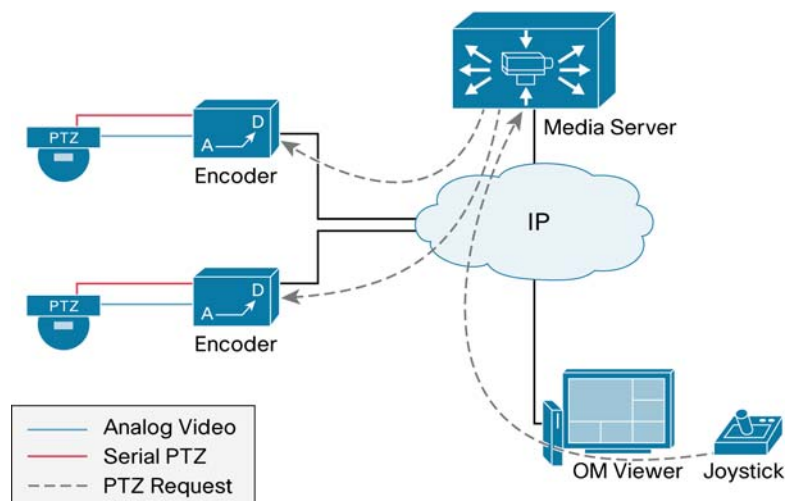
The Cisco Video Surveillance Manager solution supports the configuration of PTZ cameras connected to encoders or as IP cameras. In order to support Pan-Tilt-Zoom (PTZ) connectivity, the encoder should be able to connect to the camera via a serial interface. Appendix A provides a matrix of third-party systems supported.

The Video Surveillance Manager solution supports the following PTZ protocols:

- Bosch
- Cohu
- J2 Vision
- Pelco D
- Pelco P

Figure 24 shows how an analog camera can be connected to an IP encoder to convert its video feed to an IP video format. The encoder also connects via a serial cable to the analog camera. When the OM viewer requests PTZ control via the joystick, the Media Server intercepts the request and communicates the request to the encoder. Once the request is received by the encoder, a serial communication takes place between the encoder and the analog camera.

Figure 24. Pan-Tilt-Zoom Via Encoders



Media Server System Requirements

The Cisco Video Surveillance Media Server runs on standard commercial, off-the-shelf computer server equipment based on Intel processors.

The minimum requirements for the Media Servers are:

- 3.0 GHz Intel Pentium 4 Processor
- 1 GB RAM
- Capability for up to 24 TB Internal Storage or Fiber Channel or SCSI external storage options
- Gigabit Ethernet Interface
- Fiber Channel or SCSI External Storage Options
- SUSE Linux Enterprise Server (SLES) version 9 Service pack 3 or SUSE Linux Enterprise Server (SLES) version 10 Service pack 1 (32 bit) or SUSE Linux Enterprise Server (SLES) version 10 Service pack 1 (64 bit) or Red Hat Linux (64 bit) Server Enterprise 4 Operating System

Operations Manager System Requirements

The Cisco Video Surveillance Operations Manager may run on the same server as the Video Surveillance Media Server. A single Operations Manager server is designed to control all Media Servers and Virtual Matrix servers in the environment.

Hardware Requirements

The following are the minimum hardware requirements for the Video Surveillance Operations Manager:

- 3.0 GHz Intel Pentium 4 Processor
- 1 GB RAM
- 200 GB HDD
- Gigabit Ethernet Interface

Software Requirements

The Video Surveillance Operations Manager module shall include the following software:

- Video Surveillance Operations Manager Application
- SUSE Linux Enterprise Server (SLES) version 9 Service pack 3 or SUSE Linux Enterprise Server (SLES) version 10 Service pack 1 (32 bit) or SUSE Linux Enterprise Server (SLES) version 10 Service pack 1 (64 bit) or Red Hat Linux (64 bit) Server Enterprise 4 Operating System

Virtual Matrix System Requirements

A single Virtual Matrix server can be deployed to support a large number of Virtual Matrix monitors since the communication between the monitor and the server is required only during the client initialization or when a new view is pushed to the viewer.

Encoding Server System Requirements

The Encoding Server specifications are:

- Intel Dual-Core Pentium 4
- 3U rack unit
- 1 GB RAM
- Up to 12 TB internal storage
- 10/100 or dual Gigabit Ethernet Interface
- SUSE Linux Enterprise Server version 9 with Service Pack 3 (32 bit) or SUSE Linux Enterprise Server version 10 SP1 (32 bit)

Viewer Requirements

The Viewing Software allows an individual operator's PC to access and view video streams. The Viewing Software may also play video archive files without a browser or connection to the video surveillance system host.

The viewing PC shall meet the following requirements:

- 3.0 GHz Intel Pentium 4
- 1GB RAM
- 10/100 Ethernet adapter
- Windows XP or Windows Vista (32 bit) and Internet Explorer 6 or 7
- DirectX 9.0c
- NVIDIA or ATI AGP Graphics adaptor with 128 MB RAM

Chapter 5: Network Deployment Models

This chapter provides a high-level overview of different deployment models and highlights the typical requirements of campus and wide area networks. Cisco's Enterprise Systems Engineering team offers detailed network designs that have been deployed by enterprise customers to provide enhanced availability and performance. These designs may be found at the Cisco Validated Design Program site at: <http://www.cisco.com/go/cvd>

Campus Networks


An infrastructure that supports physical security applications requires several features from a traditional campus design. A hierarchical campus design approach has been widely tested, deployed, and documented. This section provides a high-level overview and highlights some of the design requirements that may apply to a video surveillance solution. For a more detailed review of Campus designs refer to the Campus Design documents in the Reference section.

A traditional campus design should provide:

- **High availability:** Avoid single points of failure and provide fast and predictable convergence times.
- **Scalability:** Support the addition of new services without major infrastructure changes.
- **Simplicity:** Ease of management with predictable failover and traffic paths.

A highly available network is a network that provides connectivity at all times. As applications have become more critical, the network has become more and more important to businesses. A network design should provide a level of redundancy where no points of failure exist in critical hardware components. This design can be achieved by deploying redundant hardware (processors, line cards and links) and by allowing hardware to be swapped without interrupting the operation of devices.

The enterprise campus network shown in Figure 25 is a typical campus network. It provides connectivity to several environments such as IDFs, secondary buildings, data centers and wide area sites. An Intermediate Distribution Frame (IDF) is the cable infrastructure used for interconnecting end user devices to the Main Distribution Frame (MDF) or other buildings and is typically located at a building wiring closet.

Figure 25. Campus Network

Quality of service (QoS) is critical in a converged environment where voice, video, and data traverse the same network infrastructure. Video surveillance traffic is sensitive to packet loss, delay and delay variation (jitter) in the network. Cisco switches and routers provide the QoS features that are required to protect critical network applications from these effects.

Hierarchical Design

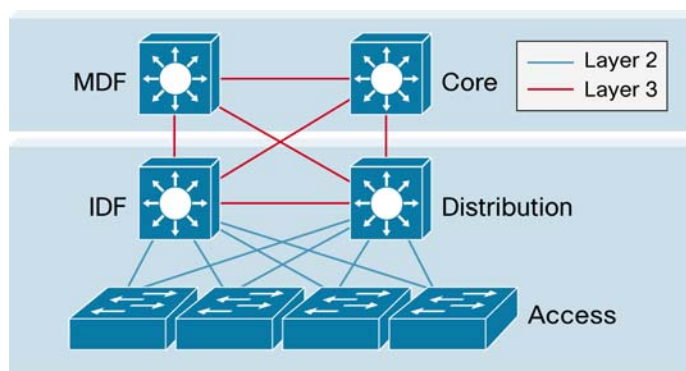
The goal of a campus design is to provide highly available and modular connectivity by separating buildings, floors, and servers into smaller groups. This multilayer approach combines Layer 2 switching (based on MAC addresses) and Layer 3 switching or routing (based on IP address) capabilities to achieve a robust, highly available campus network. This design helps reduce failure domains by providing appropriate redundancy and reducing possible loops or broadcast storms.

With its modular approach, the hierarchical design has proven to be the most effective in a campus environment. The following are the primary layers of a hierarchical campus design:

- **Core layer:** Provides high-speed transport between distribution-layer devices and core resources. The network's backbone.
- **Distribution layer:** Implements policies and provides connectivity to wiring closets. This layer provides first-hop redundancy such as Hot Standby Router Protocol (HSRP) and Gateway Load Balancing Protocol (GLBP).
- **Access layer:** User and workgroup access to the network. Security and QoS can be defined at this layer and propagated to the higher layers.

Figure 26 shows a typical Campus design with the three main layers.

Figure 26. Hierarchical Campus Design




In smaller environments, it is typical to collapse the distribution and core layers into a single layer.

Wide Area Networks

A wide-area network (WAN) is used to connect different local-area networks (LANs) and typically covers a broad geographic area. WAN services are leased from service providers who provide different speeds and connectivity options.

Figure 27 shows how a remote branch office relies on the connectivity provided by a WAN Service Provider.

Figure 27. Service Provider Network



Deploying a video surveillance solution through a WAN environment presents challenges that are not typically seen in a LAN. In a LAN environment it is common to see 1 Gbps and 10 Gbps of bandwidth, while in a WAN environment, most connections are less than 10 Mbps; many remote connections operate on a single T1 (1.544 Mbps) or less.

These inherent bandwidth constraints require careful evaluation of the placement of cameras and Media Servers and how many viewers can be supported at remote sites simultaneously. By using Child proxies, bandwidth requirements can be reduced to transport video streams across WAN connections.

The placement of recording devices also becomes important. The video may be streamed to a central site using lower frame rates or resolution, but another attractive alternative is to deploy Media Servers at the remote sites and stream the traffic using the LAN connectivity within the remote site.

The following tables show typical links that are offered by service providers:

Table 5. Service Provider Links

Digital Signal Level	Speed	"T"	Channels or DS0s
DS0	64 kbps	–	1
DS1	1.544 Mbps	T1	24
DS3	44.736 Mbps	T3	672

SONET Signal Level	Speed	SDH Equivalent
STS-OC-1	51.84 Mbps	STM-0
STS-OC-3	155.52 Mbps	STM-1
STS-OC-12	622.08Mbps	STM-4
STS-OC-48	2488.32 Mbps	STM-16
STS-OC-192	9.952 Gbps	

A point-to-point or leased line is a link from a primary site to a remote site using a connection through a carrier network. The link is considered private and is used exclusively by the customer. The circuit usually is priced based on the distance and bandwidth requirements of the connected sites.

Technologies such as Multilink PPP allow several links to be bundled to appear as a single link to upper routing protocols. In this configuration, several links can aggregate their bandwidth and be managed with only one network address. Because video surveillance traffic requirements tend to be larger than other IP voice and data applications, this feature is attractive for video surveillance applications.


Hub and spoke, also known as star topology, relies on a central site router that acts as the connection for other remote sites. Frame Relay uses hub and spoke topology predominantly due to its cost benefits, but other technologies, such as MPLS, have mostly displaced Frame Relay.

Example 1: Network Bandwidth Usage

Figure 28 shows a simple scenario with two sites. Each site has a Media Server and each Media Server is the direct proxy for an IP camera. Three OM Viewers are active in Site A and each IP camera is generating 1Mbps of network traffic. For simplicity the Operations Manager has been removed from this graphic.

Two OM Viewers are displaying video streams from Camera 1 and Camera 2 while one OM Viewer is displaying three video streams: two streams from Camera 1 and one stream from Camera 2. The network bandwidth required to display video streams for Camera 2 in Site A are relatively small for a LAN environment, but the traffic from Camera 1 can be significant for WAN environments since four different 1Mbps streams have to traverse the WAN locations.

Figure 28. Network Bandwidth Requirements




Example 2: Sites with Remote Storage

Figure 29 shows how Media Servers can be deployed at different WAN locations in order to minimize the bandwidth requirements. By deploying the Media Servers close to viewers and edge devices, the network traffic remains local to each site. Archiving video streams at each location is also an attractive solution to minimize the network traffic between sites.

In this example Site A and Site C have Media Servers acting as direct proxies and archives for the IP cameras. Since both sites are archiving and distributing video to the OM Viewers locally, the network traffic remains local to each site.

Site B can function without a local Media Server, but all video streams have to traverse the WAN connections. Since Media Server A is the direct proxy for Camera B, the 1Mbps stream has to reach Media Server A before reaching any OM Viewers. A total of 3Mbps would be required in order for both OM Viewers in Site B to receive video from Camera B.

Figure 29. Sites with Remote Storage



Example 3: Virtual Matrix Scenario

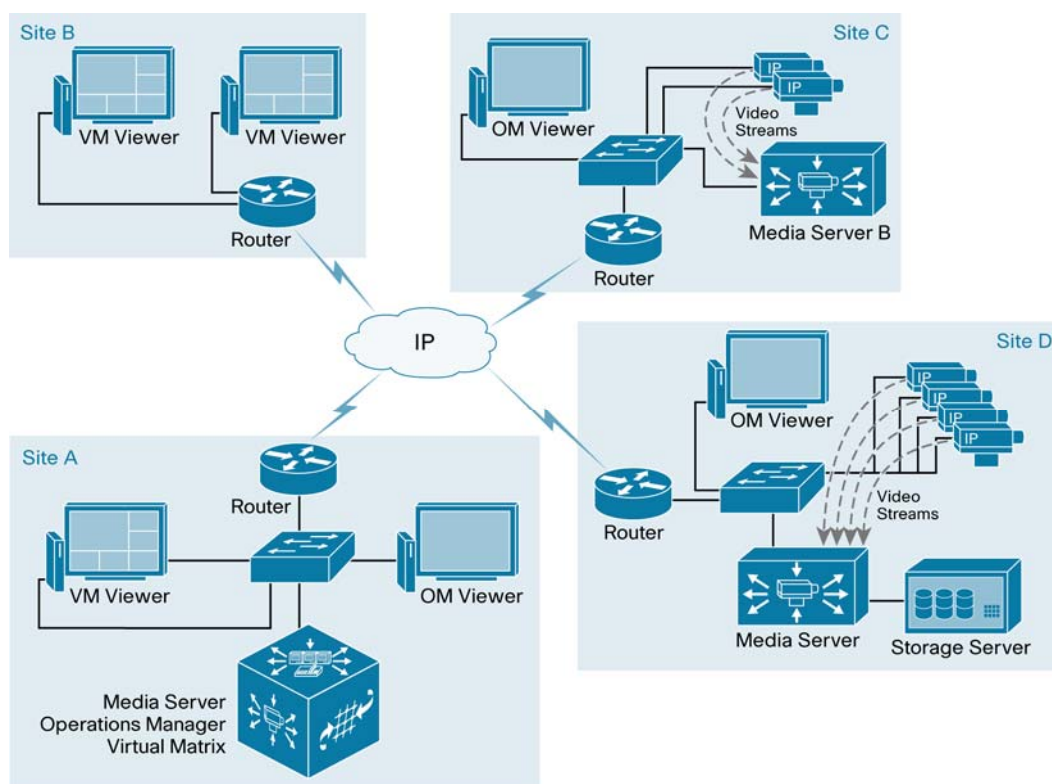
Figure 30 shows an example that includes a Virtual Matrix Server and VM Monitors located at two different sites. The Server on Site A is acting as the Media Server, Operations Manager, and Virtual Matrix for the environment. In order to reduce bandwidth traffic, Media Servers are also installed on Site C and Site D.

A single Operations Manager and a single Virtual Matrix are adequate to support this scenario.

Since the cameras are located on Site C and Site D, they are able to serve the local OM Viewers at those sites.

The Media Server on Site A can also be configured with child feeds that come from the remote Media Servers and provide those feeds locally to viewers and monitors on Site A.

Figure 30. Virtual Matrix Scenario




Example 4: Distributed Media Servers

Figure 31 shows a deployment with several remote sites, each with a local Media Server acting as the direct proxy and archive for local IP cameras.

In this scenario, all recording occurs at the remote sites and live video streams are viewed by OM Viewers and VM Monitors (video walls) at the headquarters.

The Media Server at the headquarters could also have Parent-Child proxies to each remote Media Server and request the remote streams only when required at the headquarters. This would have less bandwidth impact when the same stream is requested by more than one viewer since the traffic would be contained locally in the headquarters LAN.

Figure 31. Distributed Media Servers



Chapter 6: Video Storage

The video surveillance storage system provides multiple options to store video and audio files. The internal storage of the Media Server may be augmented by using direct attached or SAN storage. The video surveillance storage system can store video in loops, one-time archives, or event clips triggered by alarm systems providing for redundant and remote long-term archival.

Archives

An archive is a collection of video data from any given video source. This enables a feed from a camera or encoder to be stored in multiple locations and formats to be viewed at a later time.

There are two main types of archives:

- **One-Time:** where the archive recording terminates at the pre-set specified date and time.
- **Continuous Loop:** where the archive continuously records until the archive is stopped. Loop archives reuse the space (first-in-first-out) allocated after every completion of the specified loop time.

The archives may also be scheduled to begin at a certain date and time and may be configured to run using a recurring schedule.

Figure 32 shows the available options when defining a new archive using the Operations Manager. In this example a one-time archive is scheduled to start on 12/23/2007 and stop on 12/28/2007. If the Continuous Loop option was selected, the archive would record seven days before reusing the archive space.

Figure 32. Archive Types

Archive Type:

One-Time:

Start Date/Time: 12/23/2007 @ 22 : 00

Stop Date/Time: 12/28/2007 @ 22 : 00

Continuous Loop:

Loop Duration: 7 Days

Figure 33 shows an archive scheduled to occur every weekend during the month of December of 2007. The archive will last 24 hours, beginning at 10:00pm every Saturday and Sunday.

Figure 33. Scheduled Archives

Recurring Schedule:

Scheduled archives will not be a one-time or continuous loop.

Start Date/Time: 12/1/2007 @ 22 : 00

Duration: 24 Hours

Range End: No End Date
 End After 1 Occurrences
 End By: 12/31/2007

Recurs Every: 1 Hours Days Weeks Months Years

On: Sunday Monday Tuesday Wednesday
 Thursday Friday Saturday

Extracting Video Clips

Video clips may be generated from live video or from existing archives.

To generate a clip while viewing live video, click on the Record Now icon from the operator view of the Operations Manager, as shown in Figure 34. By default, the Record Now feature creates a 5-minute video clip, but the length may be changed under the Settings screen of the administrator view.

Figure 34. Extracting Video Clips



An archive clip may be also extracted from a One-Time clip or from a Continuous Loop clip using the Operations Manager. Select the archive of interest from the Video Archives section and click the Archive Clip icon (see Figure 34), the Archive Clip Form appears, as shown in Figure 35.

Figure 35. Archive Clip Form

Archive Clip Form

Start: [Set]
12/27/2007 15 : 40 : 11

Stop: [Set]
12/27/2007 15 : 40 : 11

Save as: This Computer

Save Clip

The operator may specify the start and stop times for the archive clip and where to save the video clip. By selecting This Computer, the operator may save the video clip to the local desktop in one of the two common video formats: AVI and WMV.

The Video Surveillance Manager also supports two server-side video formats. Video formats played using the Cisco Review Player include date and time information that is typically unavailable in more portable or less proprietary formats. This time and date information is a common requirement in surveillance environments when reviewing recorded video.

The operator also has the option to save the clip on the server in one of the following formats:

- **BWM:** The BWM format does not offer a digital signature to verify against tampering of the original clip.
- **BWX:** BMX files are more secure and contain a digital signature that verifies video contents. With this format, the digital signature pass phrase must be remembered to view video clips.

Note: Operators do not have permission to view BWM or BMX clips

Calculating Storage Requirements

MJPEG

When using MJPEG streams, the frame size of each image plays a key role in estimating the storage and transmission requirements. Since each frame is unique and varies according to the image complexity, it is difficult to provide a guide that provides fixed frame sizes. An IP camera that provides images with low complexity will generate smaller frame sizes. Smaller frames will require less bandwidth and storage capacity.

Table 6 shows typical frame sizes for different image resolution and image quality. As an example, each frame on a 4CIF stream with 100% quality will require 320 kilobytes (KB).

Note: These numbers are simply estimates. Each camera will provide different frame sizes based on the image complexity.

Table 6. Sample MJPEG Frame Sizes

Resolution	Frame Size		
	Quality 50%	Quality 75%	Quality 100%
QCIF	4	12	20
CIF	12	24	72
4CIF	36	72	320

* All frame sizes in kilobytes

The following formula is used to calculate the bandwidth requirements for MJPEG streams:

$$\text{MJPEG storage} = \text{Average Frame size} \times \text{Frame rate} \times \text{duration}$$

Example 1: For an 8-hour archive of a QCIF video stream with 50% quality and 15 frames per second:

$$\begin{aligned} 4 \text{ KB} \times 15\text{fps} \times 3600\text{s} &= 216,000 \text{ KB/ hour} \\ &= 216\text{MB /hour} \times 8 \text{ hours} \\ &= 1.728 \text{ GB} \end{aligned}$$

Example 2: For a 24-hour archive of a 4CIF video stream with 100% quality and 5 frames per second:

$$\begin{aligned}
 320 \text{ KB} \times 5 \text{ fps} \times 3600 \text{ s} &= 5,760,000 \text{ KB /hour} \\
 &= 5,760 \text{ MB /hour} = 5.76 \text{ GB /hour} \times 24 \text{ hours} \\
 &= 138.24 \text{ GB}
 \end{aligned}$$

MPEG-4

Rather than standalone images, MPEG-4 streams take into account video frames and the size of a given video frame varies widely between I-frames and predictive frames. Typically MPEG-4 requires less bandwidth and storage capacity when using higher frame rates.

The following formula is used to calculate the bandwidth requirements for MPEG-4 streams:

$$\text{MPEG4 storage} = \text{Bit rate (kbps)} \times \text{duration}$$

The target bit rate is configured on the camera and is already expressed in bits per second.

Example 1: For an 8-hour video stream with target bit rate of 768kbps:

$$\begin{aligned}
 768 \text{ kbps} / 8 \text{ bits/s} &= 96 \text{ KB /second} \times 3600 \text{ s} \\
 &= 345,600 \text{ KB/hour} / 1000 \\
 &= 345.6 \text{ MB/hour} \times 8 \text{ hours} \\
 &= 2.764 \text{ GB}
 \end{aligned}$$

IP Camera Video Settings

When creating a new IP camera, several settings play a role in providing the appropriate image quality and frame rate.

When using MJPEG video streams, the following image settings may be configured: Resolution, Frame Rate and Image Quality, as shown in Figure 36.

The frame rate setting determines the amount of video data generated at a given amount of time. In the example of Figure 36, the Media Server requests 30 frames per second from the edge device.

Figure 36. MJPEG Video Settings

The screenshot shows a configuration window titled "Camera Feed" with the following settings:

- *Server: -- Choose One --
- *Media Type: JPEG
- *Format: NTSC
- *Resolution: 4CIF (640 x 480)
- *Transport: TCP, UDP
- Multicast Address: [] (Leave blank for unicast)
- *Frame Rate: 30
- *Quality: 100

A red oval highlights the *Frame Rate and *Quality settings.

For MPEG-4 videos streams, the Resolution, Bit rate and Quality may be configured, as shown in Figure 37.

The bit rate setting specifies the amount of bandwidth required for the MPEG-4 video stream. Figure 37 shows an example for a 4CIF MPEG-4 stream generating 2Mbps of traffic. Higher values generate more video data every second, translating into smoother video and a more accurate representation of the field of view. A higher value also translates into larger archive file sizes.

Figure 37. MPEG-4 Video Settings

The screenshot shows the 'Camera Feed' configuration window. It includes several dropdown menus and radio buttons. The 'Server' dropdown is set to '-- Choose One --'. The 'Media Type' dropdown is set to 'MPEG-4'. The 'Format' dropdown is set to 'NTSC'. The 'Resolution' dropdown is set to '4CIF (640 x 480)'. The 'Transport' section has radio buttons for 'TCP' and 'UDP'. There is a 'Multicast Address' text field with a note '(Leave blank for unicast)'. The 'Bitrate' dropdown is set to '2000'. The 'Quality' slider is positioned towards the right end of its range. A red oval highlights the 'Bitrate' and 'Quality' settings.

The Quality scroll bar only sets two priority settings:

- A number between 50 and 99 indicates the same priority for the image quality.
- A number between 1 and 49 indicates the same priority for the frame rate

If the Quality scroll bar is between 50 and 99, the image quality has priority while trying to maintain the bit rate. Typically lower frame rates are returned when using this setting.

If the Quality scroll bar is set between 1 and 49, the frame rate has priority. The requested frame rate is as specified in the proxy_mpeg4_settings.xml file, located in the /usr/BWhttpd/drivers directory of the Media Server. Table 7 shows these settings.

Table 7. Frames Per Second Values When Quality is Set to Less Than 50

Resolution	Bit Rate (kbps)	Frames per Second
QCIF	385–1024	30
	129–384	10
	65–128	7
	0–64	5
CIF	769–1024	30
	385–768	15
	129–384	10
	0–128	5
2CIF	1500–5000	30
	769–1499	15
	385–768	10
	129–385	5
	0–128	2

Resolution	Bit Rate (kbps)	Frames per Second
4CIF	1500–5000	30
	769–1499	15
	385–768	10
	129–385	5
	0–128	2

Note: The settings of the proxy_axis_mpeg4.xml or other proxy_mpeg4 files may be changed in future releases without notice.

Storage Systems

The following tables show the current Cisco storage products and their available storage capacity:

Table 8. Storage Systems

	Media Server Internal Storage	Storage System SS-3U Series	Storage System SS-4U Series
Connectivity	Internal	Fibre Channel	Fibre Channel
Dimensions	1, 2, 3, or 5 U rack units	3 U rack units	4 U rack units
Gross Capacity	.5 to 24 TB	4 to 10.5 TB	14 to 42TB
RAID 5 Capacity	1 to 20TB	3 to 9 TB	12 to 36 TB
Cache	–	512 MB	512 MB
Host Interfaces	IDE	2 Fibre Channel	2 Fibre Channel
Data Transfer Rates	100 Mbps	2 Gbps	2 Gbps
RAID Level	–	0, 1, 5	0, 1, 5
Power	–	Redundant power supply	Dual redundant power supply

Table 9. Internal Storage Servers

Part Number	Description	Hard Drives					Storage Available		
		Size	Qty	Capacity	Parity/RAID Sets	Spare	For OS	Usable	For Archives
CIVS-MS1R-500	Cisco VS Media Server, 1RU, 1x500GB	500	1	500	0	0		500	455
CIVS-MS1R-1500	Cisco VS Media Server, 1RU, 3x500GB	500	3	1500	1	0	15	985	896
CIVS-MS1R-2250	Cisco VS Media Server, 1RU, 3x750GB	750	3	2250	1	0	15	1485	1351
CIVS-MS1R-3000	Cisco VS Media Server, 1RU, 3x1000GB	1000	3	3000	1	0	15	1985	1806
CIVS-MS2R-3000	Cisco VS Media Server, 2RU, 6x500GB	500	6	3000	1	1	15	1985	1806
CIVS-MS2R-4500	Cisco VS Media Server, 2RU, 6x750GB	750	6	4500	1	1	15	2985	2716
CIVS-MS2R-6000	Cisco VS Media Server, 2RU, 6x1000GB	1000	6	6000	1	1	15	3985	3626
CIVS-MS3R-6000	Cisco VS Media Server, 3RU, 12x500GB	500	12	6000	1	1	0	5000	4550
CIVS-MS3R-9000	Cisco VS Media Server, 3RU, 12x750GB	750	12	9000	1	1	0	7500	6825
CIVS-MS3R-12000	Cisco VS Media Server, 3RU, 12x1000GB	1000	12	12000	1	1	0	10000	9100

Part Number	Description	Hard Drives					Storage Available		
		Size	Qty	Capacity	Parity/RAID Sets	Spare	For OS	Usable	For Archives
CIVS-MS5R-18000	Cisco VS Media Server, 5RU, 24x750GB	750	24	18000	2	2	0	15000	13650
CIVS-MS5R-24000	Cisco VS Media Server, 5RU, 24x1000GB	1000	24	24000	2	2	0	20000	18200

Table 10. Storage Arrays

Part Number	Description	Hard Drives					Storage Available	
		Size	Qty	Capacity	Parity/RAID Sets	Spare	Usable	For Archives
CIVS-SS-3U-4000	Cisco VS 3U Storage System With 8x500GB Drives	500	8	4000	1	1	3000	2730
CIVS-SS-3U-6000	Cisco VS 3U Storage System With 8x750GB Drives	750	8	6000	1	1	4500	4095
CIVS-SS-3U-7000	Cisco VS 3U Storage System With 14x500GB Drives	500	14	7000	1	1	6000	5460
CIVS-SS-3U-10500	Cisco VS 3U Storage System With 14x750GB Drives	750	14	10500	1	1	9000	8190
CIVS-SS-4U-14000	Cisco VS 4U Storage System With 28x500GB Drives	500	28	14000	2	2	12000	10920
CIVS-SS-4U-21000	Cisco VS 4U Storage System With 42x500GB Drives	500	42	21000	3	3	18000	16380
CIVS-SS-4U-28000	Cisco VS 4U Storage System With 28x1000GB Drives	1000	28	28000	2	2	24000	21840
CIVS-SS-4U-31500	Cisco VS 4U Storage System With 42x750GB Drives	750	42	31500	3	3	27000	24570
CIVS-SS-4U-42000	Cisco VS 4U Storage System With 42x1000GB Drives	1000	42	42000	3	3	36000	32760

Table 11. Encoding Servers

Part Number	Description	Hard Drives					Storage Available	
		Size	Qty	Capacity	Parity/RAID Sets	Spare	Usable	For Archives
CIVS-ES-16-1500	Cisco VS Encoding Server: 3U, 16 Channel, 3x500GB	500	3	1500	1	0	1000	910
CIVS-ES-16-3500	Cisco VS Encoding Server: 3U, 16 Channel, 7x500GB	500	7	3500	1	1	2500	2275
CIVS-ES-xx-6000	Cisco VS Encoding Server: 3U, xx Channel, 12x500GB	500	12	6000	1	1	5000	4550
CIVS-ES-xx-9000	Cisco VS Encoding Server: 3U, xx Channel, 12x750GB	750	12	9000	1	1	7500	6825
CIVS-ES-xx-12000	Cisco VS Encoding Server: 3U, xx Channel, 12x1000GB	1000	12	12000	1	1	10000	9100

Appendix A: Glossary

A	
Alarm	The action or event that triggers an alarm for which an event profile is logged. Events can be caused by an encoder with serial contact closures, a motion detected above defined thresholds, or another application using the soft-trigger command API.
Alarm Trigger	The action or event that triggers an alarm for which an event profile is logged. Events can be caused by an encoder with serial contact closures, a motion detected above defined thresholds, another application using the soft-trigger command API, or a window or door opening/closing.
Alert	The action or event that triggers an alarm for which an event profile is logged. Events can be caused by an encoder with serial contact closures, a motion detected above defined thresholds, or another application using the soft-trigger command API.
API	Application Programming Interface
Archive	A place in which records or historical documents are stored and/or preserved. An archive is a collection of video data from any given proxy source. This enables a feed from a camera-encoder to be stored in multiple locations and formats to be viewed at a later time. There are three types of archives: Regular, where the archive recording terminates after a pre-set time duration lapses and is stored for the duration of its Days-to-Live. Loop, where the archive continuously records until the archive is stopped. Loop archives reuse the space (first-in-first-out) allocated after every completion of the specified loop time. Clip, the source of the archive is extracted from one of the previous two types and is stored for the duration of its Days-to-Live.
Archive Clip	The source of the archive that is extracted from one of the other two types and stored for the duration of its Days-to-Live.
Archive Server	Programs which receive incoming video streams or loops, interprets them, and takes the applicable action.
Archiver	An application that manages off-line storage of video/audio onto back-up tapes, floppy disks, optical disks, etc.
C	
Camera Controls	Permits users to change the camera lens direction and field view depth. Panning a camera moves its field of view back and forth along a horizontal axis. Tilting commands move it up and down the vertical axis. Zooming a camera moves objects closer to or further from the field of view. Many of these cameras also include focus and iris control. A camera may have a subset of these features such as zoom, pan, or tilt only.
Camera Drivers	Responsible for converting standardized URL commands supported by the module into binary control protocols read by a specific camera model.
Child Proxy	An agent, process, or function that acts as a substitute or stand-in for another. A proxy is a process that is started on a host acting as a source for a camera and encoder. This enables a single camera-encoder source to be viewed and recorded by hundreds of clients. There are three types of proxies: A "direct" proxy is the initial or direct connection between the edge camera-encoder source. By definition at least one direct proxy exists for a given video source. A "parent" proxy is the source of a nested or child proxy. Parent proxies may be from remote or local hosts. Proxies are nested in a hierarchy with inheritance rights. A "child" proxy is the result of a nested or parent proxy. Child proxies run on the local host. Proxies are nested in a hierarchy with inheritance rights. A child proxy has the same resolution, quality, and media type of its parent, but can have a lower framerate for motion JPEG.
Clip	A place in which records or historical documents are stored and/or preserved. An archive is a collection of video data from any given proxy source. This enables a feed from a camera-encoder to be stored in multiple locations and formats to be viewed at a later time. There are three types of archives: Regular: where the archive recording terminates after a pre-set time duration lapses and is stored for the duration of its Days-to-Live. Loop: where the archive continuously records until the archive is stopped. Loop archives reuse the space (first-in-first-out) allocated after every completion of the specified loop time. Clip: the source of the archive is extracted from one of the previous two types and is stored for the duration of its Days-to-Live.

D	
Direct Proxy	An agent, process, or function that acts as a substitute or stand-in for another. A proxy is a process that is started on a host acting as a source for a camera and encoder. This enables a single camera-encoder source to be viewed and recorded by hundreds of clients. There are three types of proxies: A "direct" proxy is the initial or direct connection between the edge camera-encoder source. By definition at least one direct proxy exists for a given video source. A "parent" proxy is the source of a nested or child proxy. Parent proxies may be from remote or local hosts. Proxies are nested in a hierarchy with inheritance rights. A "child" proxy is the result of a nested or parent proxy. Child proxies run on the local host. Proxies are nested in a hierarchy with inheritance rights. A child proxy has the same resolution, quality, and media type of its parent, but can have a lower frame rate for motion JPEG.
DVR	Digital Video Recorder/Recording: broadcasts on a hard disk drive which can then be played back at a later time
E	
Encoder Driver	Sends the output of a camera driver to the encoder to which the camera is attached (via the network protocol supported by a particular type of encoder).
ES	Cisco Video Surveillance Encoding Server
Event	When an incident or event occurs, it is captured by a device or application and is tagged. An event is a collection of information about an incident, including name, associated video sources, and a timestamp. If the event setup includes triggered clips, an event will have trigger tracking or video data associated directly with it. Users will need to use the event log to refer to times within a referenced archive, typically a master loop. By using the API to seek to a specific UTC timestamp, events can be used to look up occurrences in an archive that were not necessarily associated with the original event.
Event Setup	A collection of processes and configurations designed to track and notify when alarms or alerts are triggered. Types of event profiles includes event trigger tracking only, event triggers with archive clips, and motion detection. When an event profile includes a trigger from an encoder, part of the profile includes scripts copied to the encoder which release an event notification. When an event profile includes event triggered clips, a pre-post buffer archive is started from the proxies associated with the event profile. Once a trigger occurs, a clip is extracted from the pre-post buffer.
F	
Feed	The transmission of a video signal from point to point.
FPS	Frames Per Second
Frame Rate	The rate at which the source is being recorded. For motion JPEG sources, the play rate is the number of frames-per-second or fps. For MPEG sources, the play rate is the number of megabits-per-second or Mbps and kilobits per second or Kbps.
H	
HTTP	Hypertext Transfer Protocol
J	
J2EE	Java 2 Enterprise Edition
JPEG	JPEG (pronounced "jay-peg") stands for Joint Photographic Experts Group, the original name of the committee that wrote the standard. JPEG is designed for compressing full color or gray-scale images of natural, real-world scenes. JPEG is "lossy," meaning that the decompressed image is not exactly the same as the original. A useful property of JPEG is that the degree of lossiness can be varied by adjusting compression parameters. This means that the image maker can trade off file size against output image quality. The play rate is the number of frames-per-second or fps.
K	
Kbps	The rate at which the source is being recorded. For motion JPEG sources, the play rate is the number of frames-per-second or fps. For MPEG sources, the play rate is the number of megabits-per-second or Mbps and kilobits per second or Kbps.
L	
Layout	The geometric description of one or more video panes.
LDAP	Lightweight Directory Access Protocol
Loop	A loop is a hardware or software device which feeds the incoming signal or data back to the sender. It is used to aid in debugging physical connection problems.
M	
Mbps	The rate at which the source is being recorded. For motion JPEG sources, the play rate is the number of frames-per-second or fps. For MPEG sources, the play rate is the number of megabits-per-second or Mbps and kilobits per second or Kbps.
Media Server	A device that processes multimedia applications.

MPEG	MPEG (pronounced “em-peg”) stands for Moving Picture Experts Group and is the name of family of standards used for the compression of digital video and audio sequences. MPEG files are smaller for and use very sophisticated compression techniques. The play rate is the number of megabits-per-second or Mbps and kilobits per second or Kbps.
N	
NTSC	National Television System Committee
P	
Pan-Tilt-Zoom Controls	Permits users to change the camera lens direction and field view depth. Panning a camera moves its field of view back and forth along a horizontal axis. Tilting commands move it up and down the vertical axis. Zooming a camera moves objects closer to or further from the field of view. Many of these cameras also include focus and iris control. A camera may have a subset of these features such as zoom, pan, or tilt only.
Parent proxy	An agent, process, or function that acts as a substitute or stand-in for another. A proxy is a process that is started on a host acting as a source for a camera and encoder. This enables a single camera-encoder source to be viewed and recorded by hundreds of clients. There are three types of proxies: A “direct” proxy is the initial or direct connection between the edge camera-encoder source. By definition at least one direct proxy exists for a given video source. A “parent” proxy is the source of a nested or child proxy. Parent proxies may be from remote or local hosts. Proxies are nested in a hierarchy with inheritance rights. A “child” proxy is the result of a nested or parent proxy. Child proxies run on the local host. Proxies are nested in a hierarchy with inheritance rights. A child proxy has the same resolution, quality, and media type of its parent, but can have a lower frame rate for motion JPEG.
Proxy	An agent, process, or function that acts as a substitute or stand-in for another. A proxy is a process that is started on a host acting as a source for a camera and encoder. This enables a single camera-encoder source to be viewed and recorded by hundreds of clients. There are three types of proxies: A “direct” proxy is the initial or direct connection between the edge camera-encoder source. By definition at least one direct proxy exists for a given video source. A “parent” proxy is the source of a nested or child proxy. Parent proxies may be from remote or local hosts. Proxies are nested in a hierarchy with inheritance rights. A “child” proxy is the result of a nested or parent proxy. Child proxies run on the local host. Proxies are nested in a hierarchy with inheritance rights. A child proxy has the same resolution, quality, and media type of its parent, but can have a lower frame rate for motion JPEG.
Proxy Command	A URL-based API that is neither application-platform nor programming language specific. Commands are sent to dynamically loaded modules (e.g. info.bwt, command.bwt, event.bwt, &c.) using arguments in the form of name-value pairs.
Proxy Server	An agent, process, or function that acts as a substitute or stand-in for another. A proxy is a process that is started on a host acting as a source for a camera and encoder. This enables a single camera-encoder source to be viewed and recorded by hundreds of clients. There are three types of proxies: A “direct” proxy is the initial or direct connection between the edge camera-encoder source. By definition at least one direct proxy exists for a given video source. A “parent” proxy is the source of a nested or child proxy. Parent proxies may be from remote or local hosts. Proxies are nested in a hierarchy with inheritance rights. A “child” proxy is the result of a nested or parent proxy. Child proxies run on the local host. Proxies are nested in a hierarchy with inheritance rights. A child proxy has the same resolution, quality, and media type of its parent, but can have a lower frame rate for motion JPEG.
Proxy Source	An agent, process, or function that acts as a substitute or stand-in for another. A proxy is a process that is started on a host acting as a source for a camera and encoder. This enables a single camera-encoder source to be viewed and recorded by hundreds of clients. There are three types of proxies: A “direct” proxy is the initial or direct connection between the edge camera-encoder source. By definition at least one direct proxy exists for a given video source. A “parent” proxy is the source of a nested or child proxy. Parent proxies may be from remote or local hosts. Proxies are nested in a hierarchy with inheritance rights. A “child” proxy is the result of a nested or parent proxy. Child proxies run on the local host. Proxies are nested in a hierarchy with inheritance rights. A child proxy has the same resolution, quality, and media type of its parent, but can have a lower frame rate for motion JPEG.
PTZ: Pan Tilt Zoom	Permits users to change the camera lens direction and field view depth. Panning a camera moves its field of view back and forth along a horizontal axis. Tilting commands move it up and down the vertical axis. Zooming a camera moves objects closer to or further from the field of view. Many of these cameras also include focus and iris control. A camera may have a subset of these features such as zoom, pan, or tilt only.
R	
Rate	The rate at which the source is being recorded. For motion JPEG sources, the play rate is the number of frames-per-second or fps. For MPEG sources, the play rate is the number of megabits-per-second or Mbps and kilobits per second or Kbps.
Record Rate	The rate at which the source is being recorded. For motion JPEG sources, the play rate is the number of frames-per-second or fps. For MPEG sources, the play rate is the number of megabits-per-second or Mbps and kilobits per second or Kbps.

Recording	A place in which records or historical documents are stored and/or preserved. An archive is a collection of video data from any given proxy source. This enables a feed from a camera-encoder to be stored in multiple locations and formats to be viewed at a later time. There are three types of archives: Regular, where the archive recording terminates after a pre-set time duration lapses and is stored for the duration of its Days-to-Live. Loop, where the archive continuously records until the archive is stopped. Loop archives reuse the space (first-in-first-out) allocated after every completion of the specified loop time. Clip, the source of the archive is extracted from one of the previous two types and is stored for the duration of its Days-to-Live.
Recording Archive	An archive whose state is running/recording. A running regular archive gathers additional data and increases in size. A running loop archive gathers more data and reuses its allocated space. Regular archives that have not reached their duration and loops that are still recording are running. Running archives have a Days-to-Live value of v^{-1} which does not update until they have stopped.
Repository	A central place where data is stored and maintained. A repository can be a place where multiple databases or files are located for distribution over a network, or a repository can be a location that is directly accessible to the user without having to travel across a network.
S	
Stopped Archive	An archive whose state is stopped. A shelved archive does not gather additional data or increase in size. Regular archives, clips, recordings, and loops that have reached their duration are considered shelved. Shelved archives are stored for the duration of their Days-to-Live.
Stored Archive	An archive whose state is stopped. A shelved archive does not gather additional data or increase in size. Regular archives, clips, recordings, and loops that have reached their duration are considered shelved. Shelved archives are stored for the duration of their Days-to-Live.
Stream	Any data transmission that occurs in a continuous flow.
T	
Tagged Event	When an incident or event occurs, it is captured by a device or application and is tagged. An event is a collection of information about an incident, including name, associated video sources, and a timestamp. If the event setup includes triggered clips, an event will have trigger tracking or video data associated directly with it. Users will need to use the event log to refer to times within a referenced archive, typically a master loop. By using the API to seek to a specific timestamp, events can be used to look up occurrences in an archive that were not necessarily associated with the original event.
Time stamp	An international and universal time system. Representation of time used by computers and many programming languages are most often accurate down to the millisecond. UTC values are used to track archive date/time values and records when events are triggered.
Trap	Used to report alerts or other asynchronous events pertaining to a managed subsystem.
Trigger	The action or event that triggers an alarm for which an event profile is logged. Events can be caused by an encoder with serial contact closures, a motion detected above defined thresholds, or another application using the soft-trigger command API.
U	
UI	User Interface
Update Proxy	Changes the registered information for a proxy source so that the proxy process will serve multiple videos as required. Once a proxy has been updated, all requests for that proxy will be served via the new feed. All clients requesting the feeds will be switched. Proxies are not trans-coded meaning some attributes may not be changed once registered.
V	
Video Feed	The transmission of a video signal from point to point. View: A layout, dwell time, and media sources. VM: Cisco Video Surveillance Virtual Matrix Client VMR: Video Mixing Renderer
W	
Window	All or a portion of the camera view. The display can contain multiple windows either by stacking (only the top one is entirely visible) or tiling (all are visible) or a combination of both.
WMV	Windows Media Video

Appendix B: References

Physical Security Products:

http://www.cisco.com/en/US/products/ps6918/Products_Sub_Category_Home.html

Design Guides:

- Cisco Validated Designs
<http://www.cisco.com/go/cvd>
- Cisco Solutions Reference Network Designs
<http://www.cisco.com/go/srmd>



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV
Amsterdam, The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

CCDE, CCENT, Cisco Eos, Cisco StadiumVision, the Cisco logo, DCE, and Welcome to the Human Network are trademarks. Changing the Way We Work, Live, Play, and Learn is a service mark, and Access Registrar, Aironet, AsyncOS, Bringing the Meeting To You, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, CCVP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Collaboration Without Limitation, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Event Center, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, iQuick Study, IronPort, the IronPort logo, LightStream, Linksys, MediaTone, MeetingPlace, MGX, Networkers, Networking Academy, Network Registrar, PCNow, PIX, PowerPanels, ProConnect, ScriptShare, SenderBase, SMARTnet, Spectrum Expert, StackWise, The Fastest Way to Increase Your Internet Quotient, TransPath, WebEx, and the WebEx logo are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0803R)