Stream Manager Video Surveillance Solutions Reference Network Design

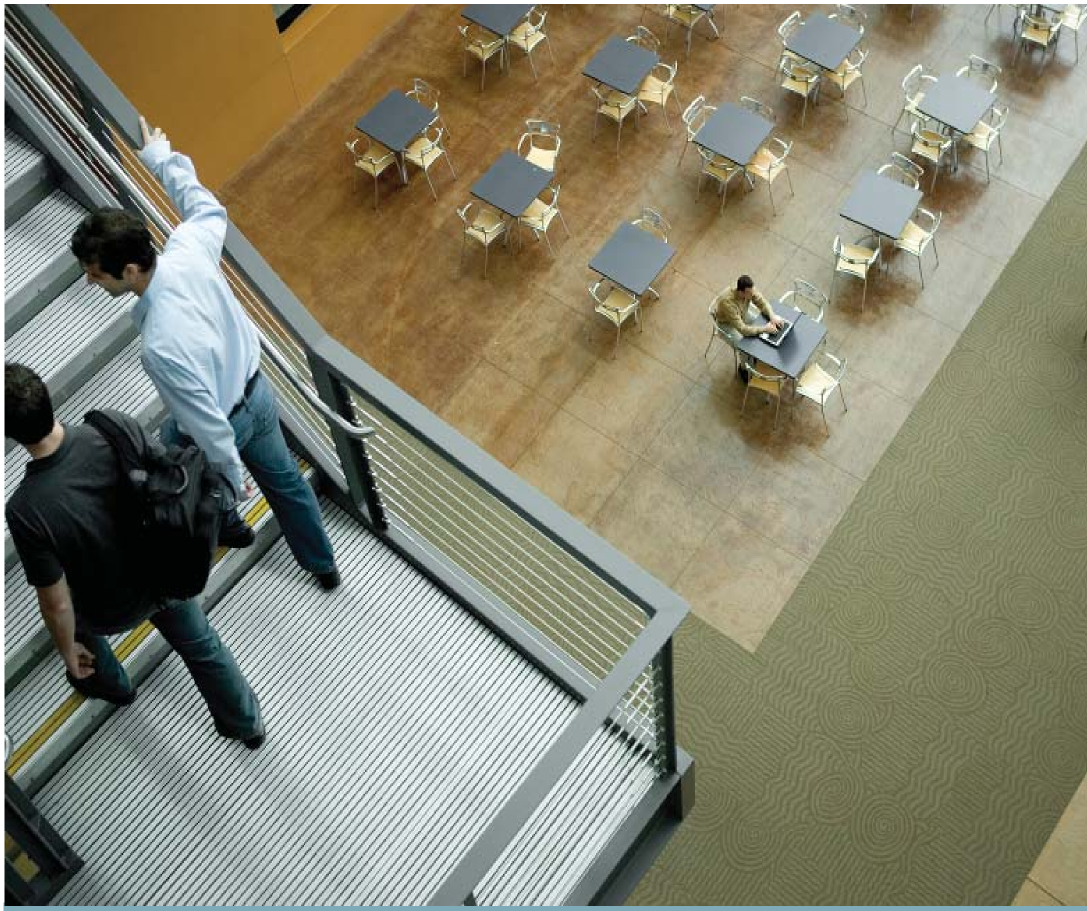

Table of Contents

Chapter 1: Video Surveillance Overview

This chapter provides an overview of the components that are used in a Cisco® Stream Manager Video Surveillance solution. These components are discussed in more detail throughout this Solutions Reference Network Design (SRND) Guide.

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 characteristics of live IP video surveillance traffic are similar to the traffic generated by IP video-conferencing systems. IP video surveillance traffic is bursty in nature, with many large packets and varying bandwidth consumption depending on motion. While video conferencing and IP phone calls tend to be short and intermittent in nature, video surveillance tends to have around-the-clock delivery requirements, and the availability of every frame can be critical during certain events.

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. Analog video can now be truly converged into a robust network environment with the intelligence and flexibility provided by the Cisco infrastructure.

Compared to the open, standards-based IT industry, the video surveillance market has seen limited innovation and still has many proprietary systems with unique features and management requirements.

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

Figure 1. Cisco Video Surveillance Solution

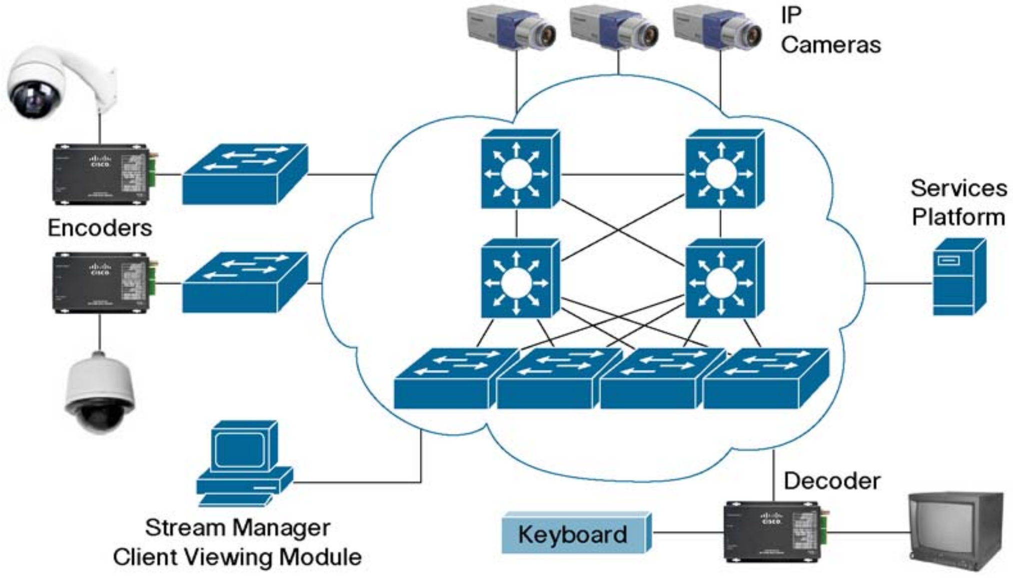

Solution Components

The following main components make up the Cisco Video Surveillance Solution:

- IP Gateway Encoders
- IP Gateway Decoders
- Convergence Chassis
- Recorders (Services Platforms)
- Cisco Stream Manager Software

These main components interact with a number of third-party analog video cameras, matrix switches and keyboards that are part of typical Closed Circuit TV (CCTV) environments.

Cisco IP Gateway Encoders and Decoders run the Cisco Stream Manager Gateway software, allowing them to become part of a “Virtual Matrix Switch (VMS)” that distributes video traffic over an IP network.

The solution can also be deployed as a Hybrid design that integrates with third-party matrix switches, or as a self-contained solution that provides video transport across an IP network infrastructure.

Cisco IP Gateway Encoders and Decoders Introduction

Cisco IP Gateway Encoders and Decoders are essential components of the video surveillance solution. By leveraging the built-in Cisco Stream Manager Gateway software, the Cisco IP Gateways operate over an intelligent IP network infrastructure and can become part of a highly available and fully distributed “Virtual Matrix Switch.”

IP Gateways contain embedded software libraries that allow them to translate transparently between different keyboard and Pan-Tilt-Zoom (PTZ) protocols. In this way, you can control analog camera functions using keyboards from a wide range of major manufacturers.

The Cisco IP Gateway Encoders and Decoders also provide integrated IP transport of audio channels and contact closure data.

Cisco IP Gateway Encoders

Analog cameras connect directly to the coaxial input on a Cisco IP Gateway Encoder. Encoders connect to the network using an Ethernet port.

Using MPEG-4 (advanced simple profile), the encoder digitizes and compresses the video stream, which can be displayed on a PC via the Cisco Stream Manager software or on a CCTV monitor using a Cisco IP Gateway Decoder (see Cisco IP Gateway Decoders below).

The Cisco IP Gateway Encoders use high-performance processors to provide broadcast-quality video at up to D1 resolution (720 x 480 NTSC or 720 x 576 PAL). 4CIF, 2CIF and CIF resolutions are also available and frame rates up to 30 frames per second are achievable in all cases.

Note: A single D1 video stream at 30 fps requires an average bandwidth of 3 Mbps, depending on the amount of motion in the camera's field of view (FOV). The greater the activity level in the FOV, the lower the compression will be relative to a static FOV. Chapter 6 provides more details about bandwidth requirement.

Cisco IP Gateway Encoders models CIVS-SG1AECMD-C16 and CIVS-SG1AECOD-FE can provide dual MPEG-4 video streams at different resolutions and frame rates. This functionality allows flexible configuration for live viewing at higher resolutions while recording at a lower resolution.

The Cisco IP Gateway Encoders also have the ability to switch dynamically to IP Multicast delivery when more than one destination/ receiver, either software or hardware based, or a Cisco Service Platform Video Recorder, requests a video stream.

Cisco IP Gateway Decoders

Similar to an IP Gateway Encoder, the IP Gateway Decoder connects to the network using its Ethernet port and to a traditional CCTV analog monitor using its coaxial output. IP video can be pulled from any camera on the network using the decoder and be displayed locally. You can also use a CCTV keyboard that is attached to the serial port of the decoder to control PTZ cameras that are attached to the network.

Cisco Convergence Chassis

The Cisco Convergence Chassis are 3RU, 16 slot modular devices that provide the flexibility of multiple system design options in a convenient and space saving package. There are currently three types of Cisco Convergence Chassis offered.

USB Convergence Chassis

The USB Convergence Chassis is used primarily when a design requires a high-density deployment. The chassis accommodates up to 64 input channels of analog video when an external power supply module is used. USB connections can then be made directly from the back of the chassis to the input of the video recorder.

10/100 FE Convergence Chassis

The 10/100 FE Convergence Chassis has 14 RJ-45 connectors associated with 14 card slots. Single channel Cisco IP Gateway Encoders and Decoders can be installed into the chassis and connected directly to the network. The chassis also provides the option of either internal or external power supplies.

Gigabit Ethernet Convergence Chassis

The Gigabit Ethernet Chassis has a built-in 20-port, fully managed Layer 2+ switch featuring 2 Gig-E copper or fiber ports, 14 10/100 ports accessible through the backplane, and 4 external 10/100 ports. This chassis minimizes or eliminates the need for external switches and minimizes cabling, while providing the design flexibility that additional fiber ports offer. An internal or external power supply is available.

Recorders

Recording is an essential component of a video surveillance solution and has traditionally been the weakest link in many environments that still use VCRs for recording.

The Cisco solution offers Service Platforms with different capacities to provide video recording and storage. The Services Platforms include the Cisco Video Surveillance Stream Manager software to provide continuous recording, scheduled recording or triggered recording.

Both live and recorded video can be viewed on a PC running the Cisco Stream Manager Client Viewing Module or from a traditional third-party analog console via the Cisco IP Gateway Decoders. The flexibility of an IP infrastructure makes possible the deployment of the Service Platforms locally or at remote network locations.

Video storage requirements are typically very large and measured in Terabytes (TB), with requirements to record continuously. Chapter 7 provides more details about Service Platforms and bandwidth storage requirements.

Cisco Stream Manager Software

The Cisco Video Stream Manager software is supported by Cisco IP Gateway Encoders, Cisco IP Gateway Decoders, and Cisco Service Platforms. The Cisco Stream Manager software provides virtual matrix switch functionality by leveraging the intelligent IP network infrastructure. It also provides multi-vendor interoperability supporting analog cameras, PCs, and specialized keyboard controllers and monitors designed for traditional CCTV systems.

Video Surveillance Designs

Historically, the traditional deployment model for video surveillance has been centralized and isolated from other systems, usually with separate management and cable infrastructures. This situation presents several challenges, including:

- High installation costs for coaxial and fiber cabling.
- Limited number of monitoring stations because cabling must be duplicated and typically is dedicated for this purpose.
- No integration with external sensory systems.

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.

Matrix Switch

In a traditional video surveillance environment, a matrix switch is the core element of the solution. Its basic function is to act as an array of video inputs and outputs, allowing users to control the display of different cameras and to switch control of PTZ functions. Figure 2 shows a traditional system with a matrix switch.

Figure 2. Traditional System with Matrix Switch

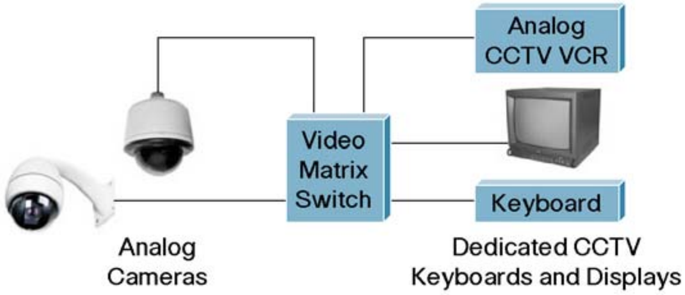

Because the matrix switch is the core of a traditional video surveillance environment, it exposes several challenges and limitations:

- The cable infrastructure is constricted by the physical location of the matrix switch, which tends to be in close proximity to the security operations center or to the main viewing location.
- The viewing of video surveillance streams is limited to the monitors that are connected to the matrix switch.
- Retrieving recorded video for analysis can be cumbersome.

Virtual Matrix Switch (VMS) Design

The Cisco Video Surveillance solution can operate without a matrix switch. This deployment allows video streams to be redirected to any monitor or PC via the Cisco IP Gateway Decoder or Stream Manager software. Once video is digitized and compressed by the Cisco IP Gateway Encoder, it can be directed to any viewing or recording location using the IP network infrastructure.

The Stream Manager Gateway software is supported on all Cisco IP Gateway Encoders and Decoders and provides a seamless integration with the IP network infrastructure. The software automatically discovers all Cisco Video Surveillance components and provides a single centralized management interface via the rich client based Cisco Stream Manager Configuration Module.

Figure 4 shows a system without a Matrix switch. This system provides the compression and transport of video traffic to any monitoring location in the following ways:

- Traditional analog cameras have their video encoded and placed on the network via IP Gateway Encoders.
- IP based cameras are connected directly to the network.

- Service Platform Recorders are placed wherever it is most convenient, either locally or remotely.
- Cisco Stream Manager Client Viewing Module, installed on a PC that is attached anywhere on the network, is used to manage, configure, view and monitor network based video.
- IP Gateway Decoders are used for retrieving video from the network for viewing on traditional analog monitors.

Figure 3. Virtual Matrix Switch

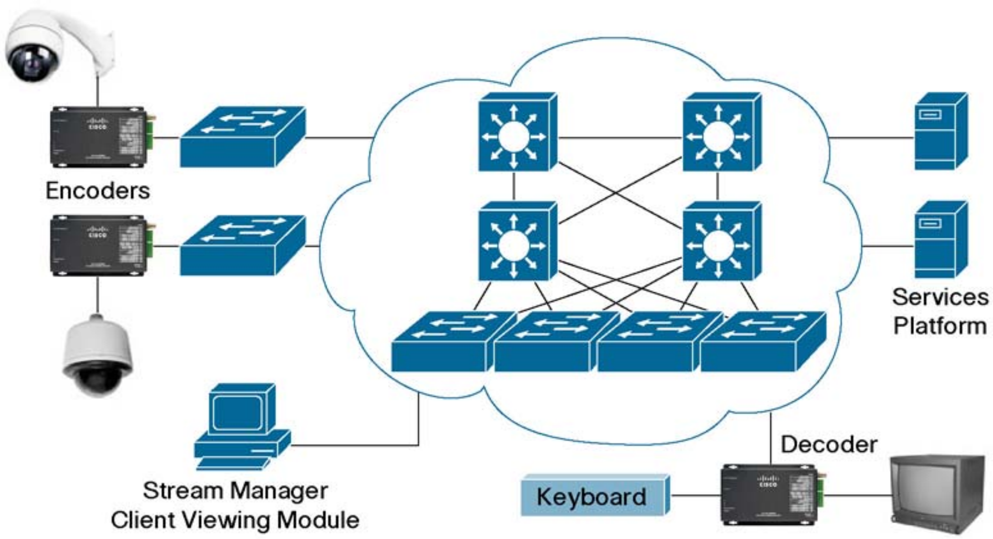

Hybrid Design

Traditional video surveillance deployments are purely analog and have not yet been able to benefit from a converged network approach. Rather than looking at a massive forklift upgrade, a hybrid system provides an easy migration to an IP-based solution.

Figure 4 shows a hybrid system in which a matrix switch seamlessly integrates with a Cisco Video Surveillance digital recording system. Many of the descriptions provided in the VMS design example are also relevant here.

Figure 4. Hybrid Design

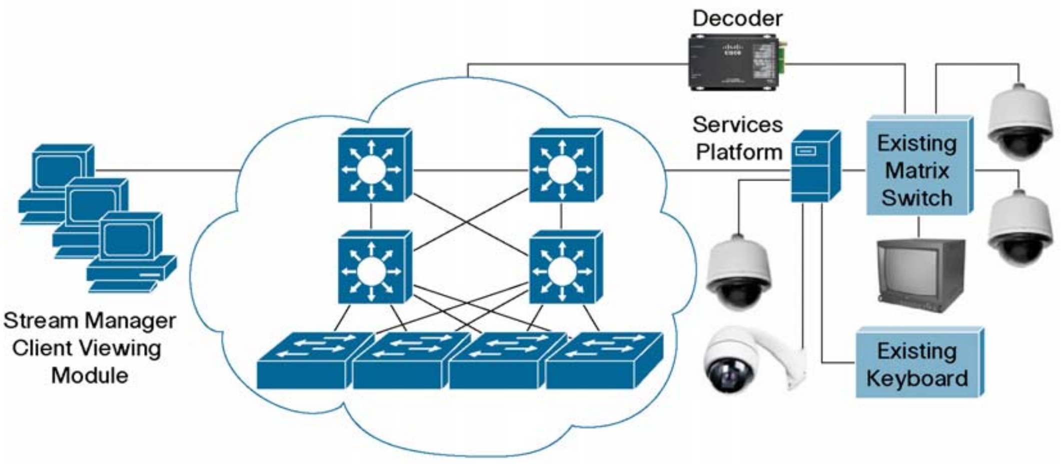

Video Resolutions and Frame Rates

Analog Video

Video surveillance solutions use a set of standard resolutions. Baseband analog video is received and transmitted from the coaxial BNC ports on Cisco Video Surveillance Gateways and Services Platforms. 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.

Digital Video

Table 1 shows resolutions for various video formats.

Table 1. Video Resolutions (in pixels)

Format	NTSC-Based	PAL-Based
QCIF	176 × 120	176 × 144
CIF	352 × 240	352 × 288
2CIF	702 x 240	704 x 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. The 4CIF and D1 resolutions are almost identical and sometimes these terms are used interchangeably.

Cisco Video Surveillance IP Gateways support different combinations of video resolutions and frame rates, based on the application requirements and bandwidth constraints. Each stream from a Cisco IP Gateway can be configured for any combination of the video resolutions and frame rates shown in Table 2.

Table 2. Video Resolutions, Frame Rates

Frame Rates (NTSC)	Frame Rates (PAL)	Video Resolutions
1.5 fps	1 fps	CIF
2 fps	2.5 fps	2CIF
3 fps	5 fps	4CIF
3.75 fps	6.25 fps	D1
5 fps	12.5 fps	
7.5 fps	25 fps	
10 fps		
15 fps		
29.97 fps		
30 fps		

The use of a D1 or 4CIF resolution in combination with a 30 or 25 frames per second rate results in video that is nearly indistinguishable from the source video when displayed on a monitor and viewed by the human eye. Lower resolutions and frame rates reduce the quality of the received image compared to the original, but can result in significantly less bandwidth consumption per stream on the IP network.

MPEG-4 and Video Codecs

Cisco Video Surveillance Gateways use MPEG-4 technology to convert the received analog NTSC or PAL video stream into a digital stream. A codec performs *encoding* and *decoding* on a digital video stream. MPEG-4 converts the incoming analog stream into a digital array of pixels and compresses the stream to reduce the amount of bandwidth consumed over the IP network. This compression is accomplished by reducing spatial redundancy within a given frame and the temporal redundancy between frames. Therefore, the amount of bandwidth consumed by a video stream can vary greatly depending on the complexity of the image and on the amount of motion or change over time within the stream.

In the world of IP networking, the term *frame* also 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. The human eye requires 25 images per second or more to perceive a sequence of images as smooth motion. In a video stream with fewer images per second, or a lower *frame rate*, motion usually is perceived as choppy or broken. At higher frame rates up to 30 frames per second, the video motion appears smooth. To achieve compression, video compression codecs such as MPEG-4 use these 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.

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. Therefore, for a given bit-rate, the quality of MPEG-4 codec is better than M-JPEG.

- **Wavelet** is a video compression approach that uses a discrete wavelet transform (DWT) on the entire image, unlike the discrete cosine transform (DCT) technology used in the MPEG codecs, which work on portions of the image. This compression results in image blurring under heavy compression, as opposed to the "blocky" (also known as pixelization) form of degradation common to overly compressed MPEG codecs. Wavelet technologies can produce high quality compressed video, but generally at higher bit-rates and with greater processor power requirements than MPEG-4.
- **MPEG-1 and MPEG-2** are predecessors to the MPEG-4 part 2 standard currently in use on Cisco Video Surveillance products. Both 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 part 2** is currently in use as the video compression technology within Cisco Video Surveillance Products. MPEG-4 offers improved quality over MPEG-2 compression. 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, and is widely deployed in the video surveillance industry.
- **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. H.264 is rapidly becoming the standard in the video conferencing world. At the time of this writing, H.264 is under active development for inclusion in a future Cisco Video Surveillance product offering.

Integration with IP Cameras

The Cisco video surveillance solution is focused primarily upon the ability to integrate with third-party analog cameras, which continue to account for 90 percent of professional grade video surveillance camera purchases. Many third-party vendors also offer IP Cameras. The IP camera market is growing primarily in market segments that do not require the highest level of quality or the ability to integrate with an installed base. Cisco currently has a development relationship with Panasonic to integrate its WV-NP244 IP camera with the Cisco Video Surveillance solution. The integration is through the interaction with the Panasonic camera as though it was a Cisco IP Gateway device.

The Panasonic configuration utility is used to initially assign an IP address, subnet mask, default gateway and necessary camera parameters to the WV-NP244. The user then adds up to 16 "Virtual Ports" on a Cisco Service Platform, where each port represents an individual Panasonic camera. The creation of these virtual ports allows the Cisco Service Platform to record video from these devices and to act as a gateway device for Cisco Decoders that request video from a WV-NP244 video source.

It is important to traffic engineer for no more than 16 cameras per Cisco Service Platform and to realize that by acting as a gateway device, the Cisco Service Platform becomes the source of any unicast/multicast streams that are requested by a decoder instead of by the camera itself. The reason for this traffic flow is that the Panasonic WV-NP244 does not have the ability to directly participate in the dynamic multicast discovery mechanism or video streaming mechanisms that the

Cisco video surveillance solution employs. Therefore, to enable interoperability, the Cisco Service Platform acts as a proxy for access on behalf of the Panasonic cameras.

Additional IP cameras will be integrated through the Cisco Technology Developers Program.

Chapter 2: Deployment Models

Campus

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, please review 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 5 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 5. Campus Network

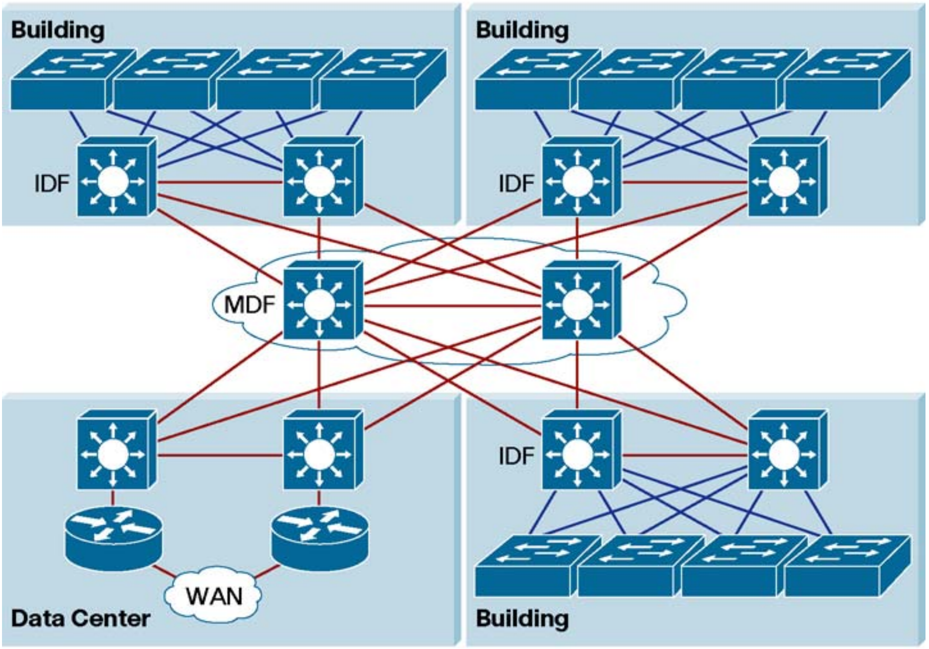

IP multicast must be enabled across the campus network before Cisco Video Surveillance devices can communicate. Multicast allows these devices to dynamically discover one another, and allows Cisco IP Gateway Encoders to transmit a single stream to be received by multiple endpoints. The Multicast section in Chapter 4 discusses deployment details for IP multicast and provides sample configurations for several scenarios.

Quality of service (QoS) is also 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. Chapter 6 provides more details about bandwidth requirements and how to protect video surveillance traffic from other network applications.

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 6 shows a typical Campus design with the three main layers.

Figure 6. Hierarchical Campus Design

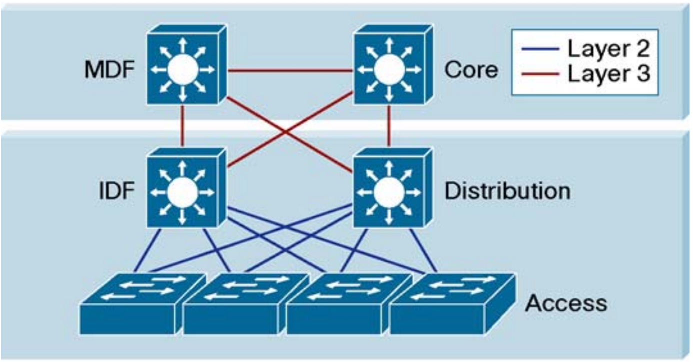

In smaller environments, it is typical to collapse the distribution and core layers into a single layer.

Default Gateway Redundancy

In a multilayer environment, each network host must be able to contact its default gateway at all times. The default gateway address is the IP address of the router that functions as the computer's gateway to other subnets. The default gateway is typically provided by DHCP, but it may be configured statically. In a network environment with redundant distribution layer routers, one of the routers is selected as the default gateway for a VLAN, and another router acts as a backup default gateway.

Figure 7. Gateway Redundancy

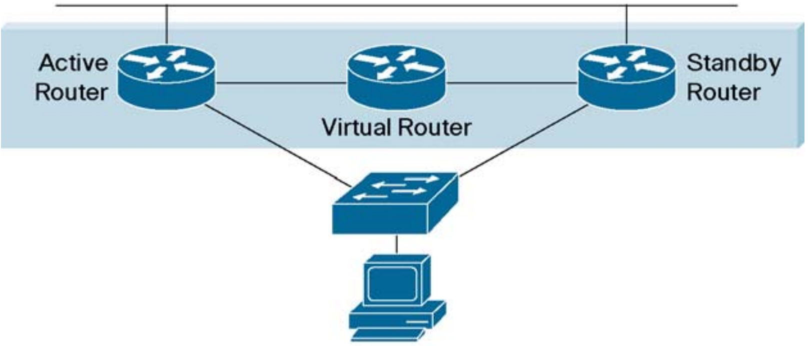

Cisco provides two main methods of gateway or first-hop redundancy:

Hot Standby Router Protocol (HSRP) – A Cisco proprietary protocol that provides path redundancy by sharing MAC addresses between gateways. A virtual MAC address and IP address are shared between two routers. HSRP routers send HSRP packets to the multicast group address 224.0.0.2, to which all routers listen and elect the active and standby routers.

In an HSRP group, there are typically:

- One active router – The active router forwards traffic.
- One standby router – If the active router fails, the standby router becomes active.

- One virtual router – While the virtual router is not an actual router, hosts use this address as their default gateway, and the current active router takes care of forwarding the traffic.

HSRP also is able to provide load balancing by allowing routers to be members of multiple HSRP standby groups. Typically, the active routers are divided among different subnets, making each router responsible for forwarding for some subnets, and for acting as a standby router for others.

Gateway Load Balancing Protocol (GLBP) – GLBP automatically selects simultaneous gateways, while HSRP only allows a single router to be active. By default, GLBP balances traffic on a per-host basis, using a round-robin scheme: the first device that comes up in a subnet uses the first router in the list, the second devices used the next, and so on. In this way, more than two devices can participate in the group.

VLANs

A Virtual LAN (VLAN) groups network devices in a single broadcast domain. This approach allows several devices to be logically grouped, regardless of their physical location. Such grouping keeps traffic local to the VLAN and provides several benefits:

- Contained broadcast and multicast traffic, which increases performance.
- Enhanced security. For security purposes, users can be grouped into a VLAN. Because passing traffic between different VLANs is achieved only through a router, the router provides a point where additional filtering and access lists may be implemented.
- Network administration is reduced when assigning hosts to specific VLANs on switch ports.

Chapter 5 provides more details about segmenting video surveillance traffic using VLANs.

Spanning Tree

A redundant network design provides several benefits, but it may also create potential loops, which can bring the network to a halt. The Spanning Tree Protocol (STP) prevents these loops from occurring and provides backup links through the redundant paths in the network.

Spanning Tree assigns roles to the switches and ports to make sure that only one path through the network is active at a time. The typical roles for a switch are root bridge, root ports, designated ports and non-designated ports.

Figure 8 shows how Spanning Tree eliminates the possibility of a bridging loop by blocking one of the redundant links. If the primary link to Switch 1 were to fail, STP would re-examine the network topology and enable the alternate path. Switch 1 has been elected as the *root bridge*, and all other switches calculate all the paths from themselves to the root. The next-hop bridge toward the root becomes the designated bridge as is responsible for sending bridge protocol data units (BPDUs) on that segment. BPDUs contain the appropriate information for STP configuration.

Figure 8. Spanning Tree

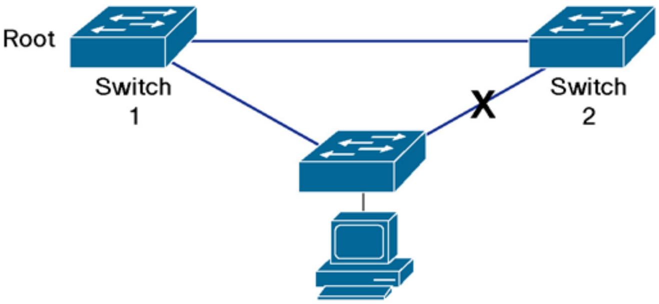

802.1D, the original STP protocol was designed with a conservative logic and a slow convergence time (up to 50 seconds). Current network environments demand quicker convergence times and new standards provide a quicker convergence during failures:

- 802.1w Rapid Spanning Tree (RSTP). Offers a fast convergence times because it does not rely on timers. Sub-second convergence times are possible in most cases. 8021w has become one of the most favored Spanning Tree protocols
- 802.1s Multiple Spanning Tree (MST). Traditionally, switches need to run one spanning tree process per VLAN. This approach can be CPU-intensive in an environment with a large number of VLANs. MST maps multiple VLANs with similar topologies to a single spanning tree instance.

To build and maintain a loop-free network, STP cycles the bridge ports through several states:

- Blocking – A port is in blocking mode upon startup and when the port is not the optimal path to the root bridge. Blocking ports do not forward traffic.
- Listening – A port is in listening mode when it is listening for other bridges. While listening, the port is not forwarding traffic.
- Learning – The bridge listens for other bridges and learns MAC addresses of network devices
- Forwarding – The port is operating normally, forwarding traffic.

The Cisco Spanning Tree toolkit offers several features to improve convergence time or add loop protection capabilities. These features include:

- Portfast – A port configured for portfast skips through the various STP stages and immediately starts forwarding traffic. This feature is intended for host devices and not for links that connect two switches
- Loop Guard – Protects against possible loops by detecting a unidirectional link. A unidirectional link is a port that is operationally in the up state but that is not receiving traffic. Loop Guard should not be enabled on ports that are enabled for portfast.
- Rootguard – Allows a device that is connected to a portfast-enabled port to be part of the STP, but it does not allow the device to become the root port. This feature is not compatible with Loop Guard.
- Portfast BPDU Guard – Prevents a switch port from receiving BPDU messages on a portfast-enabled port, preventing a loop if another switch is attached to that port. This feature disables the port completely when a BPDU is heard on a portfast port and requires manual intervention to bring the port out of error disabled (errDisabled) mode.

Routed Access Solution

Traditional hierarchical campus design models generally have used Layer 2 switching from the access layer, and Layer 3 switching or routing elsewhere. By using the Layer 3 capabilities available in most Cisco Catalyst® switches, a fully routed network design can be achieved that can reduce the network complexity and improve network availability. With a routing protocol such as EIGRP or OSPF at the access layer, the network can recover from failure in a deterministic way without requiring the fine-tuning of multiple protocols or devices.

Enabling a routing protocol at the access layer provides several benefits:

- Because Spanning Tree is not required for failover, the network is easier to deploy and troubleshoot. The risk for potential loops also goes away.
- The default gateway redundancy (HSRP, GLBP) requirement disappears because the default gateway is configured at the access layer.
- The need to match Spanning Tree and HSRP/GLBP priorities is no longer required.
- The connection between distribution and access layer is a routed interface; no VLAN trunks are required.
- Route summarization at the distribution layer.
- EIGRP can deliver sub-200 ms convergence times.
- Equal cost multi-path provides load balancing across the network

Figure 9 shows a design in which EIGRP provides a routing model to distribution and access switches. Each access layer provides its own unique subnets, which have to remain local within a wiring closet or access layer switch.

Figure 9. Routed Access Solution

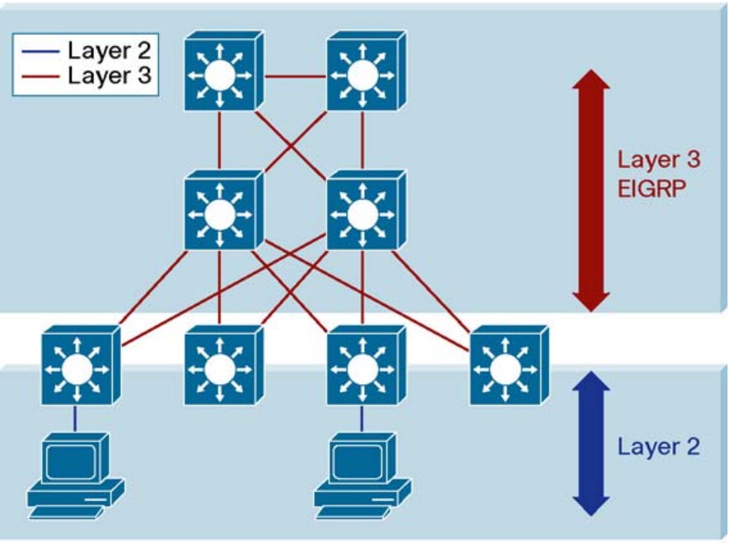

Wide-Area Networks

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.

Deploying a video surveillance solution through a WAN environment presents challenges that are not typically seen in a LAN. In a LAN environment, its 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 number of cameras, the number of locations and how many recorders can be supported at remote sites. The encoders may be configured with a lower resolution or frame rate to transport the video streams across a WAN connection. Chapter 6 includes a table with the expected bandwidth requirements for different traffic patterns.

Figure 10. Service Provider Network

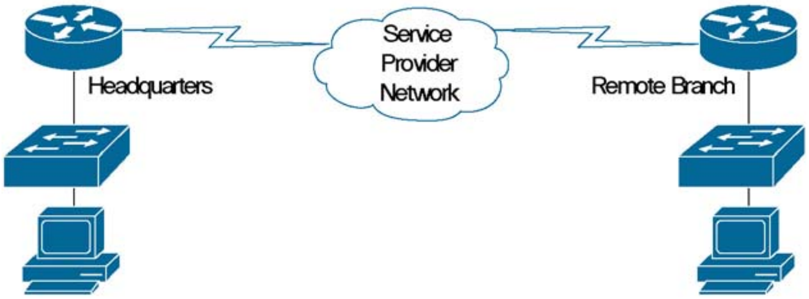

The placement of recording devices also becomes important. The video may be streamed to a central site using lower frame rates or resolution, but another attractive alternative is to deploy the recorders at the remote sites and stream the traffic using the LAN connectivity within the remote site.

Some Cisco IP Gateway Encoders can provide two simultaneous MPEG-4 video streams; one stream may be configured for storage (with lower resolution and frame rates) and a second stream may be used for live viewing with higher quality.

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

Table 3. 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 Gpbs	

Point-to-Point

A point-to-point or leased line is a link from one 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.

Figure 11 shows two locations connected via multiple T1 links using Multilink PPP.

Figure 11. Point-to-Point Links

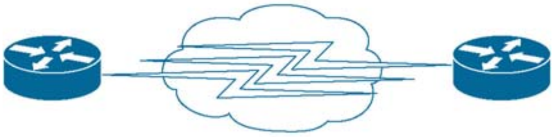

Hub and Spoke Topology

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. Frame Relay has been an efficient alternative to traditional leased-line circuits, allowing multiple users to share the network medium and the available bandwidth, but other technologies, such as MPLS, have begun to gradually displace Frame Relay.

Figure 12 shows a network with one central site and three remotes. The central site is the main connecting point.

Figure 12. Hub and Spoke (Star) Topology

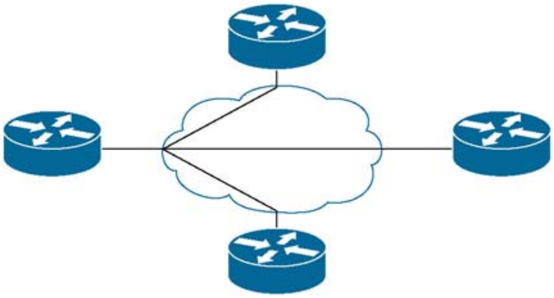

Wireless

Wireless LAN Infrastructure (WLAN)

Wireless LAN Infrastructure refers primarily to campus networks where each wireless access point (AP or WAP) is itself physically connected to a wired connection. Client devices associate to an AP to establish connection to services that are commonly available on the wired network. The IEEE 802.11 set of wireless networking standards has become commonplace in the world of IP networking. Cisco wireless products are standards-based, which provides for integration with other wireless solutions on the market.

The Cisco Unified Wireless Network architecture provides APs, client technologies, mobility services, and network unification for a tight integration to the wired network infrastructure. Video surveillance technologies have the opportunity to leverage the wireless network for support of remote cameras, or mobile clients receiving IP video streams, such as a laptops running Cisco Stream Manager software. However, certain caveats must be considered within the overall wireless network design to ensure success of the system:

- **Bandwidth Considerations:** The most widely deployed wireless LAN technology is 802.11b, which has a raw Radio Frequency (RF) data rate of 11Mbps. The maximum IP network throughput of 802.11b is around 6.5 Mbps, depending on the application. Careful consideration is required when planning to have a video surveillance endpoint transmit or receive streams over this limited bandwidth. 802.11g supports a theoretical raw RF data rate of 54 Mbps, but in practice speeds of 20 Mbps or less are more common. 802.11g is compatible with 802.11b, but when the technologies are intermingled, a least common denominator of throughput is delivered to the entire coverage area of the access point.

  The IEEE has formed a task group to develop a new amendment to the 802.11 standard known as 802.11n. This advancement proposes to increase bandwidth by as much as tenfold over 802.11g. Cisco is actively participating in the standards process. Expect this effort to offer significant benefits to video surveillance applications.
- **Shared Media:** 802.11 Wireless Access Point technology is by its nature a “shared media” network. Only one radio can transmit at a time on a given channel and access point. As more client devices attach to an AP, performance naturally degrades as each client gets a smaller share of the overall airtime.
- **Multicast Support:** On a wired network, IP multicast networking controls which ports receive packets destined to a certain multicast group. This function differentiates multicast from broadcast traffic, which is intentionally destined for every node on a subnet. Current 802.11b/g technology is not optimized for consistent multicast performance. Many corporate network administrators may choose not to support multicast applications on the wireless infrastructure for these reasons, and may block multicast traffic from wireless network segments.
- **Pitfalls of “Dual NICs”:** Many laptop computers today come with embedded network interface controllers (NICs) for both wireless and wired networks. Confusion can result when both of those NICs are active concurrently, because one of the two interfaces typically owns the default gateway path or route for all outbound traffic from the device. This situation can cause issues for many applications, particularly multicast applications such as the Cisco Stream Manager software. It is recommended that you activate only activate one NIC in such a device when it is being used for video surveillance functions. If both NICs are available, the wired NIC is most often the better choice. In situations where both NICs must be active, routing metrics for the NICs may be manipulated in the host operating system to control the proper path for video surveillance traffic to travel.
- **Service Interruption:** Wireless networking with 802.11b/g uses the 2.4G Hz portion of the RF spectrum. This space is unlicensed and is commonly used by devices such as cordless telephones. Interference is possible on the RF spectrum from competing 2.4 GHz network devices or from other sources, such as microwave ovens. Heavy use of the 2.4GHz portion of the spectrum by multiple devices can significantly impact multicast performance over wireless.

- **Quality of Service:** The Cisco Unified Wireless Network supports the 802.11e standard for wireless QoS, and is Wi-Fi Multimedia (WMM)-certified. These features enable prioritization of surveillance traffic over other traffic sources on the network. Linking to existing traffic classification and marking schemes from the wired network, 802.1p (CoS) and DSCP provides an end-to-end QoS architecture across the wired and wireless networks. However, 802.11 is still a shared medium, and the wireless prioritization does not provide an absolute bandwidth guarantee.

Wireless Mesh Networking

Wireless Mesh networking can use a combination of 2.4 GHz client access with 5.8 GHz backhaul links to create a “mesh” of access points that cover large areas without direct-wired access from all APs. It is possible to also use only the 5.8 GHz radios in mesh access points but connect a switch to the built-in Ethernet port and support wired client access. Wireless Mesh is a powerful technology for covering metropolitan areas or large campus environments with requirements for outdoor and mobile access. Cisco has successfully run IP video surveillance traffic over Cisco Wireless Mesh networking using 5.8 GHz radios and wired clients. When using 2.4GHz 802.11b/g, the caveats for WLAN also apply. Some additional items to consider for Mesh Networking include the following:

- **Maximum Transmission Unit (MTU):** Mesh networking uses AES encryption on the backhaul links between Mesh APs and Root APs (Root APs connect the mesh to the wired network). Cisco Video Surveillance IP Gateways produce streams of packets very near the MTU of Ethernet at 1514 Bytes. In order to transport these packets across the AES encrypted tunnel between APs, they must be further fragmented. In effect the total packet count approximately doubles, which results in additional bandwidth consumption by an already constrained network.
- **Backhaul Bandwidth:** The recommended data rate for backhaul links is 18 Mbps, which translates into a real IP networking throughput of approximately 9 Mbps. Depending on the size of the mesh, the number of video sources, and their resolutions and frame rates, careful traffic engineering must be performed to ensure that links are not overloaded with video traffic.
- **Multiple Hops:** Larger Mesh networks may have APs that are multiple hops from the Root AP at the edge of the wired network. Each hop effectively reduces the overall bandwidth that is available to a given endpoint device, because aggregation of multiple Mesh APs is the nature of mesh network design. This effect must be considered for traffic engineering within wireless networks for support of video surveillance.

Rings and Linear Topologies

The benefits of a hierarchical network design have been enjoyed by traditional enterprise networks, where a modular architecture allows easy addition of building blocks to support growth and maintain high availability.

The hierarchical design may not apply to every network environment, where an alternative design may be more appropriate due to physical constraints. In environments without a centralized wiring infrastructure or that cover wide geographic areas, such as parking garages, municipalities and railroads, an alternate design may be needed.

The concept of rings and linear topologies comes from daisy-chaining network devices to extend the physical reach of the cable infrastructure and to provide a basic level of redundancy.

Figure 13 shows devices connected in a linear and ring fashion, without a centralized cable infrastructure.

Figure 13. Rings and Linear

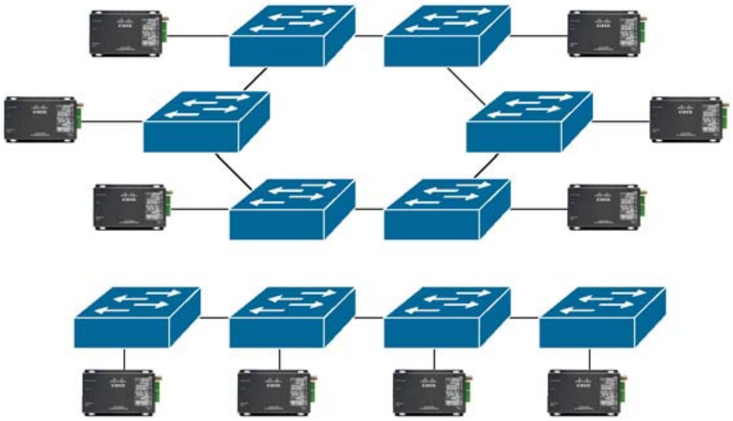

This design has several potential issues:

- The number of bridges is limited. The Spanning Tree 802.1d specification only allows for a diameter of seven bridge hops. Therefore, the number of bridge devices between any two devices in the network cannot be greater than seven.
- Fiber cable infrastructure needs to be carefully planned. When using Catalyst switches, the SFP/GBIC capabilities should be carefully addressed in terms of fiber type and distance limitations.
- At least one Layer 3 switch or router should be configured as a Multicast Rendezvous Point for redundancy. Using two switches with this capability enabled is desirable for redundancy. These topologies have unique design implications and are addressed in more detail in Chapter 4.

Chapter 3: Cisco Stream Manager Application Requirements

The Cisco Stream Manager switching and recording software provides virtual matrix switching to connect cameras, keyboards and monitors via the Cisco IP Gateways.

The Cisco Video Surveillance solution provides the Cisco Stream Manager Administration and Monitoring Module to manage all components.

The following are the minimum system requirements to run the Cisco Stream Manager applications:

- PC with 2.8 GHz
- Windows XP PRO (SP2 + .Net Framework version 2.0)
- 1024 MB
- 10 GB free hard drive space
- 10/100 Mbps network card
- ATI video card

Cisco Stream Manager Configuration Module

The Cisco Stream Manager Configuration Module allows for complete configuration of all Cisco Video Surveillance products, including Cisco IP Gateway Encoders, Cisco IP Gateway Decoders, and Cisco Service Platforms.

The Cisco Video Surveillance products ship with a default IP address in the 192.168.0.x range. Before an administrator can initially view and configure the various Cisco Video Surveillance products, the server running the Cisco Stream Manager Configuration Module must be in this subnet or must have multicast connectivity to a subnet with this address scheme. The initial configuration allows the administrator to change the IP address of the Cisco Video Surveillance devices to match the production network environment.

The Cisco Stream Manager Configuration Module relies on the multicast group 235.1.1.1 for Stream Manager device discovery.

Configuration Software features include:

- Instant polling of connected devices – A single mouse-click allows administrators to poll all connected devices to gain immediate access to configuration setting.
- Intuitive data presentation – Devices can be sorted and displayed by name, IP address, device type or firmware revision, enabling a quick overview or detailed analysis of the system.
- Real-time feedback – Changes to brightness, contrast, color and hue are shown in the Cisco Stream Manager Configuration display window, improving the accuracy of image adjustments.

- Time synchronization – All connected equipment can be time synchronized. See Chapter 4 for more details.

Figure 14 shows the initial screen of the Cisco Stream Manager Configuration Module. This screen displays a list of devices that may be configured via the Stream Manager Configuration Module[1]. Clicking **Encoders** under **Network Devices** displays a list of the available IP Gateway Encoders. This list is dynamically discovered using the 235.1.1.1 multicast group and refreshed automatically every 60 seconds. Clicking **Refresh** in the lower section of the main area initiates a new multicast discovery. This refresh consists of the initiating station sending discovery messages to the 235.1.1.1 group address. The Cisco VS devices respond with their respective capabilities (for example, Cisco IP Gateway – Encoder, Firmware XYZ, Camera ID, IP Address).

Figure 14. Stream Manager Configuration Module

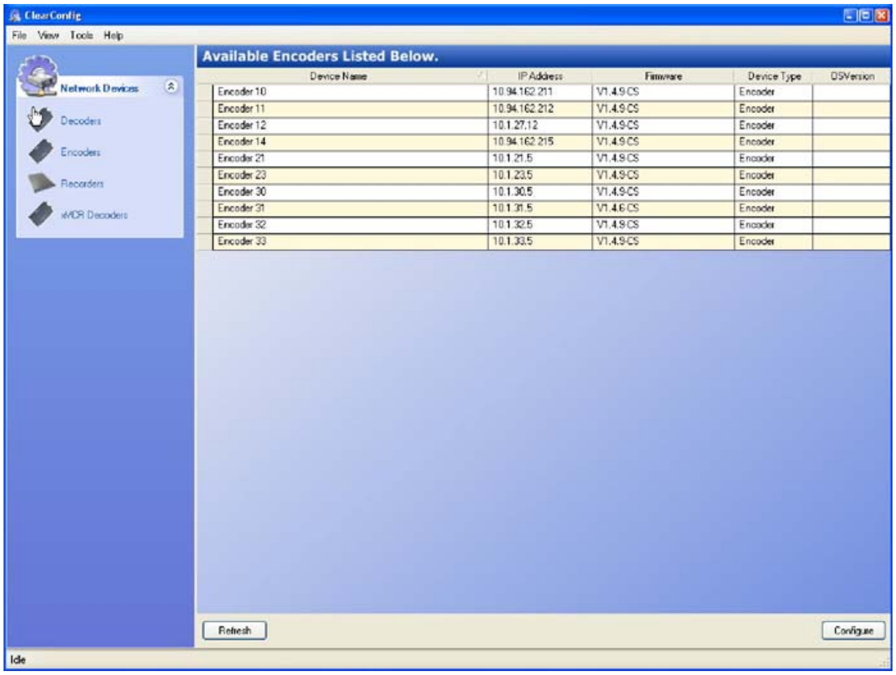

Available Encoders Listed Below.

Device Name	IP Address	Firmware	Device Type	OSVersion
Encoder 10	10.94.162.211	V1.4.9 CS	Encoder	
Encoder 11	10.94.162.212	V1.4.9 CS	Encoder	
Encoder 12	10.1.27.12	V1.4.9-CS	Encoder	
Encoder 14	10.94.162.215	V1.4.9 CS	Encoder	
Encoder 21	10.1.21.5	V1.4.9 CS	Encoder	
Encoder 23	10.1.23.5	V1.4.9-CS	Encoder	
Encoder 30	10.1.30.5	V1.4.9-CS	Encoder	
Encoder 31	10.1.31.5	V1.4.6 CS	Encoder	
Encoder 32	10.1.32.5	V1.4.9 CS	Encoder	
Encoder 33	10.1.33.5	V1.4.9-CS	Encoder	

By double-clicking a device, several parameters may be changed. To change the IP address of the device, click **Network Settings**:

Figure 15. Network Settings

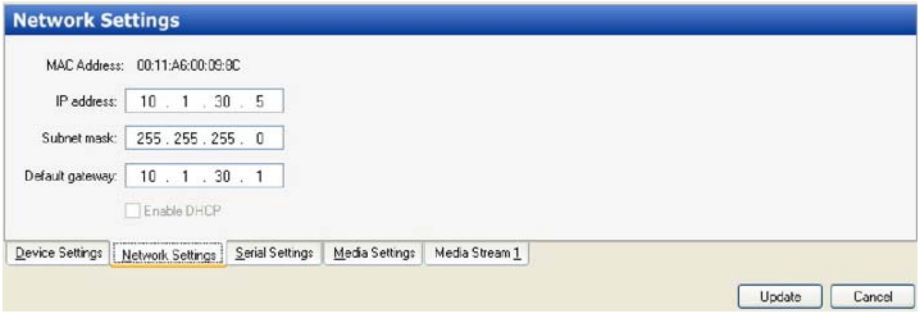

Network Settings

MAC Address: 00:11:A6:00:09:8C

IP address: 10 . 1 . 30 . 5

Subnet mask: 255 . 255 . 255 . 0

Default gateway: 10 . 1 . 30 . 1

□ Enable DHCP

[1] The screen captures have not been updated to reflect the updated application name.

The following screen shows the different stream settings for a Cisco IP Gateway Encoder:

Figure 16. Media Stream Settings

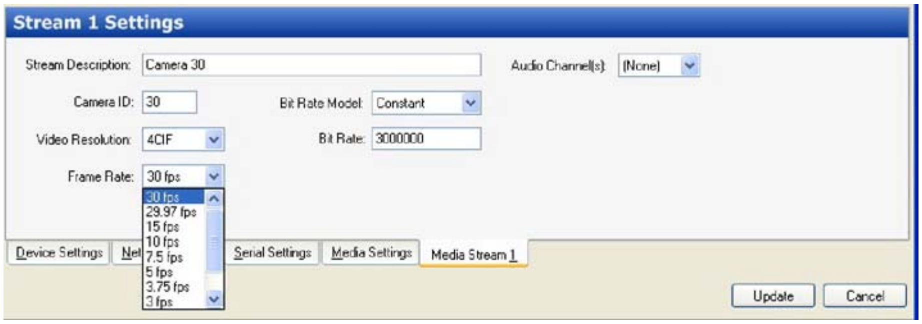

Stream 1 Settings

Stream Description: Camera 30

Audio Channel(s): (None)

Camera ID: 30

Bit Rate Model: Constant

Video Resolution: 4CIF

Bit Rate: 3000000

Frame Rate: 30 fps

To configure the Cisco Stream Manager Services Platform, click **Recorders** under **Network Devices.** The following screen shows how to configure a high-density Cisco Services Platform with 12 active cameras:

Figure 17. Configuring a Recorder

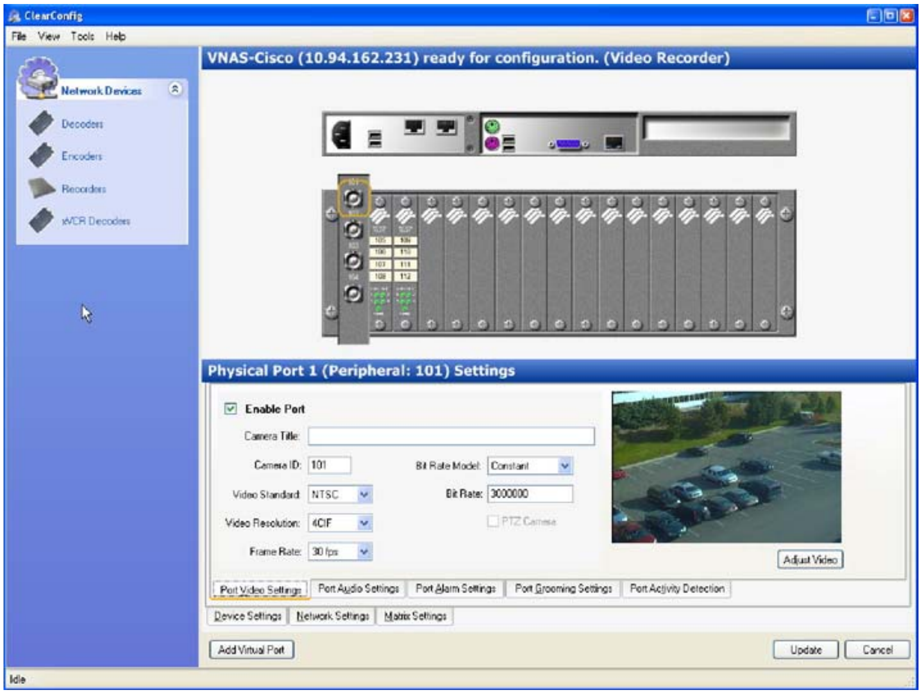

VNAS-Cisco (10.94.162.231) ready for configuration. (Video Recorder)

Physical Port 1 (Peripheral: 101) Settings

☑ Enable Port

Camera Title:

Camera ID: 101

Bit Rate Model: Constant

Video Standard: NTSC

Bit Rate: 3000000

Video Resolution: 4CIF

□ PTZ Camera

Frame Rate: 30 fps

Idle

Cisco Stream Manager Software

The Cisco Stream Manager software acts as a Cisco IP Gateway Decoder and allows the user to view and manage different video streams. This functionality is extended across the IP network and, given adequate bandwidth, is not constrained to a single geographic location It is important to note that the Cisco Stream Manager application currently is optimized to render up to ten simultaneous video streams (subject to platform performance capabilities and assuming the use of the recommended video card and driver software).

The Cisco Stream Manager software relies on the following multicast groups:

- 235.1.1.1 – Used for Stream Manager device discovery

- 239.1.1.2 – Used to synchronize clocks
- 239.1.1.3 – Used by the Cisco Stream Manager software to receive events and alarms
- 239.1.1.4 – Used to discover photo processors upon request

All Cisco IP Gateway Encoders, Decoders and Services Platforms rely on the 235.1.1.1 multicast group for device discovery and to communicate device capabilities, such as resolution and frame rates. The Cisco Stream Manager software relies on this multicast group to determine what devices are available and to display video streams. The Stream Manager software also relies on the multicast 239.1.1.2 and 239.1.1.3 to receive events and alarms and to synchronize its clock with the Service Platform acting as the time master. The bottom of Figure 18 shows the Events area, which is populated by these messages.

For authentication purposes, a checksum is calculated on each frame of video and embedded in the media frame header. The Stream Manager Player Module verifies each frame in real time to protect the video streams.

Cisco Stream Manager software features include:

- Display and playback at multiple resolutions and frame rates
- Instant replay of video streams
- Video data marked for follow-up or review is available in the Events area.
- Video playback by camera based on date and time
- Camera PTZ controls
- Built in matrix switch functionality

Figure 18. Stream Manager Client Viewing Module

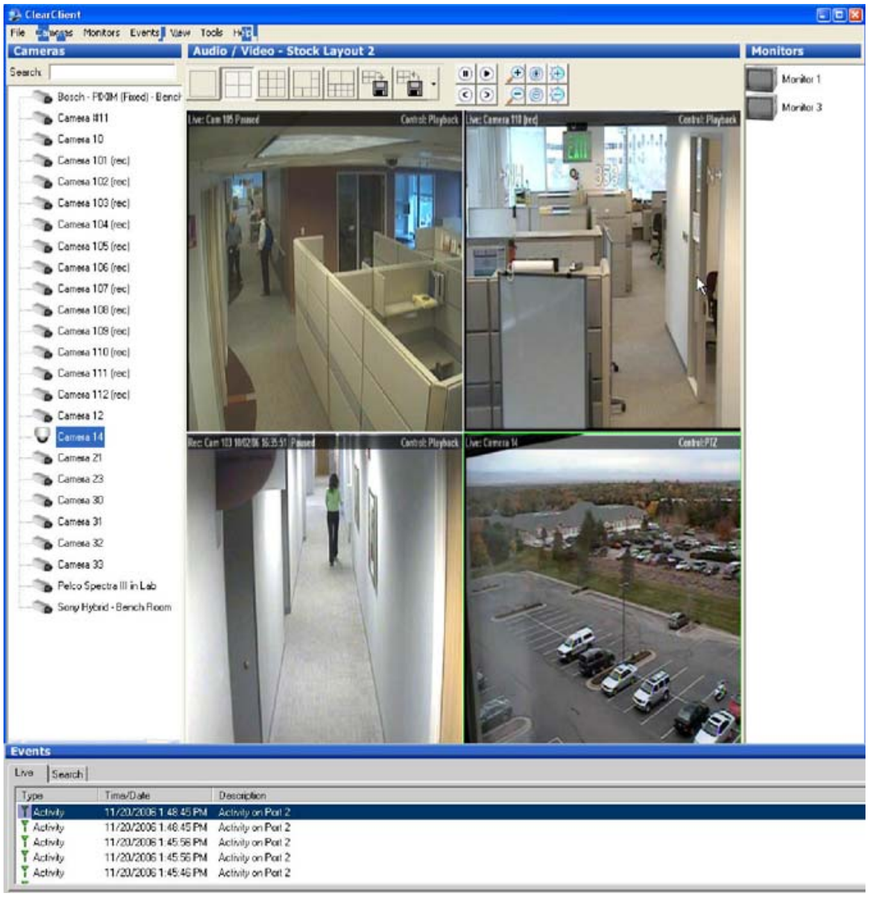

Type	Time/Date	Description
Activity	11/20/2006 1:48:45 PM	Activity on Port 2
Activity	11/20/2006 1:48:45 PM	Activity on Port 2
Activity	11/20/2006 1:45:56 PM	Activity on Port 2
Activity	11/20/2006 1:45:56 PM	Activity on Port 2
Activity	11/20/2006 1:45:46 PM	Activity on Port 2

Figure 18 shows the Cisco Stream Manager Client Viewing Module's main screen, which is divided in several sections:

- Camera Bank (left side of the screen) – Displays a list of cameras. The images from these cameras may be viewed by dragging them to the main screen.
- Monitor Bank (right side of screen) – Displays all monitors (Cisco IP Gateway Decoders). Images can be displayed on any monitor by dragging a camera from the camera bank to the monitor.
- Display Controls (top of screen, left) – These icons offer different layouts and the option to save and load custom layouts.
- Recorded Video Control (top of screen, center) – Playback functions for navigating stored video, such as pause, play frame forward, and play frame back.
- Events List (bottom of screen) – Includes events bookmarked for viewing or archiving. To review an event, drag it to one of the monitoring windows.

Cisco Stream Manager Administration and Monitoring Module

The Cisco Stream Manager Administration and Monitoring Module allows real-time monitoring and reporting of all Cisco Video Surveillance devices. It provides a general system overview and detailed statistics for all connected devices, and it provides health monitoring and storage utilization.

Cisco Stream Manager Administration and Monitoring Module features include:

- Intuitive data presentation – Shows all key device statistics, including storage capacity, recording rate, camera ID, configuration history and event log.
- Using a history log, the system provides a status report and history of all events and actions taken against the events.
- Equipment and peripheral search function – Allows administrator to conduct a search for any camera to display the associated encoding device.

Figure 19 shows the main window of the Cisco Stream Manager Administration and Monitoring Module and details about a storage device. It also shows the active ports on a Cisco Services Platform with several storage details, such as recording rate, average retention time and total storage capacity.

Figure 19. Cisco Stream Manager Administration and Monitoring Module

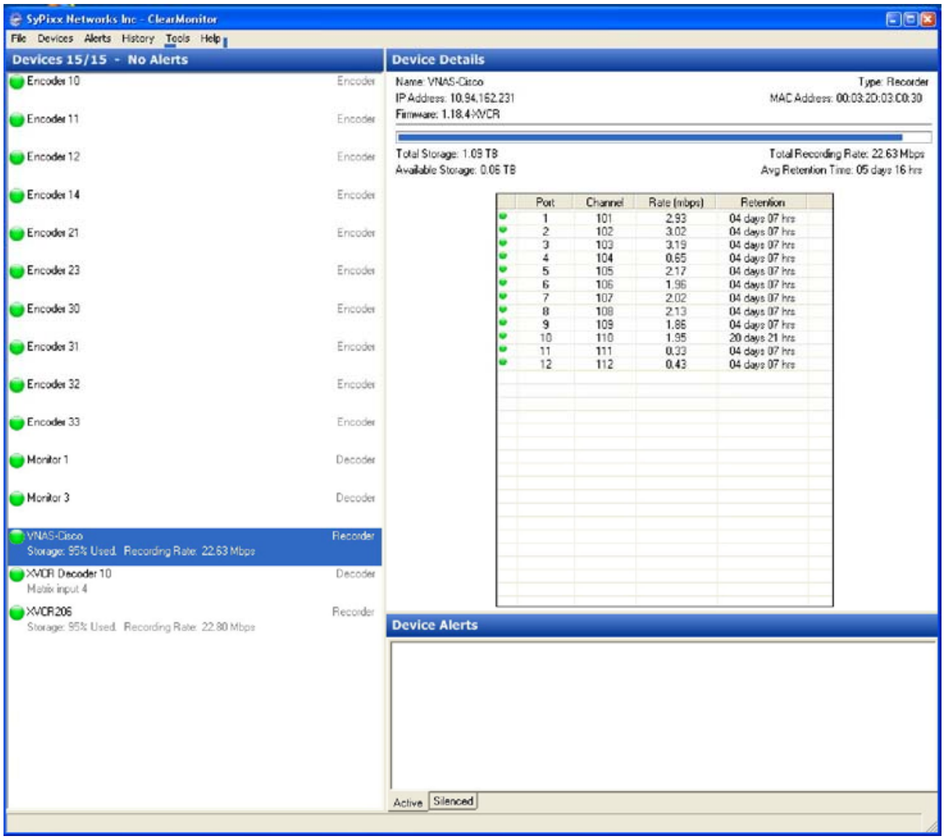

Encoder 10	Encoder
Encoder 11	Encoder
Encoder 12	Encoder
Encoder 14	Encoder
Encoder 21	Encoder
Encoder 23	Encoder
Encoder 30	Encoder
Encoder 31	Encoder
Encoder 32	Encoder
Encoder 33	Encoder
Monitor 1	Decoder
Monitor 3	Decoder
VNAS-Cisco Storage: 95% Used. Recording Rate: 22.63 Mbps	Recorder
XVCR Decoder 10 Matrix input 4	Decoder
XVCR206 Storage: 95% Used. Recording Rate: 22.80 Mbps	Recorder

Name: VNAS-Cisco
Type: Recorder
IP Address: 10.94.162.231
MAC Address: 00:03:2D:03:C0:30
Firmware: 1.18.4-XVCR

Total Storage: 1.09 TB
Total Recording Rate: 22.63 Mbps
Available Storage: 0.06 TB
Avg Retention Time: 05 days 16 hrs

	Port	Channel	Rate (mbps)	Retention
●	1	101	2.93	04 days 07 hrs
●	2	102	3.02	04 days 07 hrs
●	3	103	3.19	04 days 07 hrs
●	4	104	0.65	04 days 07 hrs
●	5	105	2.17	04 days 07 hrs
●	6	106	1.96	04 days 07 hrs
●	7	107	2.02	04 days 07 hrs
●	8	108	2.13	04 days 07 hrs
●	9	109	1.86	04 days 07 hrs
●	10	110	1.95	20 days 21 hrs
●	11	111	0.33	04 days 07 hrs
●	12	112	0.43	04 days 07 hrs

Chapter 4: Protocols and Features

IP Multicast Overview

Unlike most network applications that operate between one sender and one receiver, video and multimedia applications frequently require that one sender communicate with a group of receivers simultaneously. Using IP multicast to transmit this information reduces the overall network load and minimizes the affect on the source of the video from unnecessary replication of a common data stream.

The following basic types of IP traffic are used in a network environment:

- Unicast – Between one source address and one destination host address. This traffic can have a significant affect on the network if the number of hosts receiving this information is large.
- Broadcast – From a host address to a broadcast destination address, which typically includes all hosts on a segment.
- Multicast – 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.

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 the desired content. Since IP multicast 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.

Figure 20. Multicast Traffic

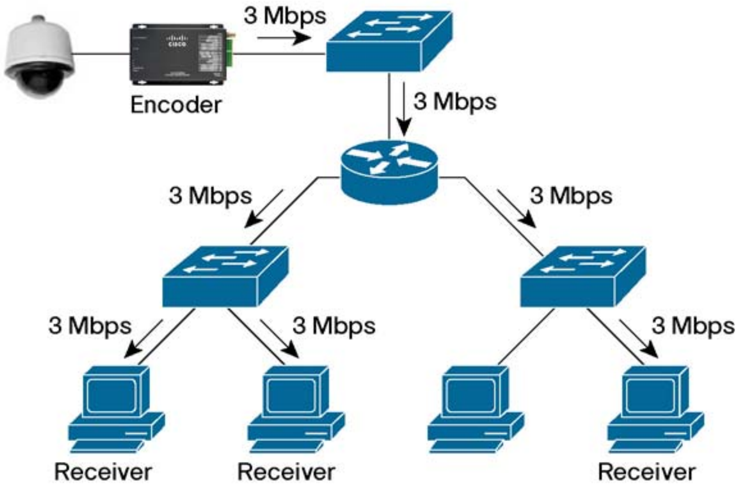

Figure 20 shows an example in which a Cisco IP Gateway Encoder is transmitting a video stream for a multicast group. Three Cisco Stream Manager Client Viewing Modules have requested to receive this video. Using multicast protocols, Cisco routers and switches replicate the video stream to only the segments and hosts that require it, using approximately 3 Mbps of encoder-to-network bandwidth.

Multicast Addresses

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.

The IANA maintains a list of Multicast addresses at:
http://www.iana.org/assignments/multicast-addresses.

The Cisco Video Surveillance solution uses the following multicast addresses:

- 235.1.1.1 – Used for device discovery
- 239.1.1.2 – Used for clock synchronization
- 239.1.1.3 – Used for alarms and events
- 239.1.1.4 – Used to discover photo processors upon request
- 239.255.x.x – Dynamically created by Cisco IP Gateway Encoders
- 239.x.x.x – Dynamically created by the Services Platform

Future software releases will allow users to change these addresses to minimize conflicts with existing network environments that may use these addresses for other multicast applications.

The address 239.255.x.x is formed by adding the last two octets of the encoder’s IP address to 239.255.x.x. For example, a Cisco IP Gateway Encoder with IP address 10.1.30.35 would use 239.255.30.35.

Cisco IP Gateway Encoders stream to this multicast address when a request for two or more streams is received. In this way, the same video stream can be seen simultaneously by different Cisco IP Gateway Decoders or Cisco Stream Manager Client Viewing Modules.

The address 239.x.x.x is formed by adding the last three octets of the Services Platform’s IP address to 239.x.x.x. This multicast group is used when a second request is received for a stream generated by the Services Platforms.

Note: In a future software release, IP Gateway Encoders will also use the 239.x.x.x to minimize possible address conflicts.

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 uses its own Layer 2, management and routing protocols. Figure 21 shows the interaction between these different protocols.

Figure 21. Interaction Between IGMP and PIM

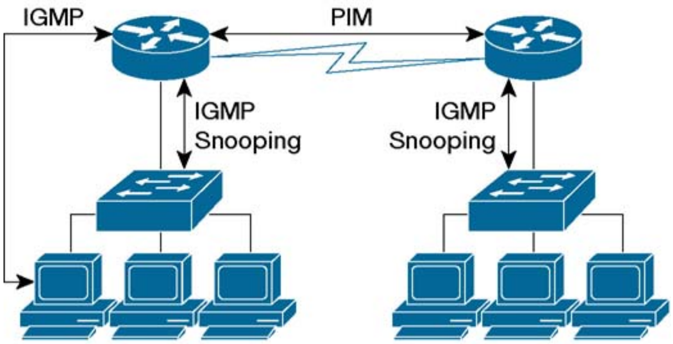

- PIM is the multicast routing protocol that is responsible for building multicast delivery trees and for enabling multicast packet forwarding.
- IGMP is used by hosts to dynamically register to multicast groups. The communication occurs between the router and the host.
- IGMP snooping is used to prevent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

Internet Group Management Protocol (IGMP)

To receive a video surveillance stream from a source, the host must support IGMP. Hosts register themselves in a multicast group by sending IGMP messages to their local multicast router. The routers and switches that are configured for IGMP periodically send out queries to discover which groups are still active on a particular subnet. Then, a multicast routing protocol (such as PIM) forms the multicast tree between the routers.

The original IGMP version 1 defined in RFC1112 has been extended to include new features. Table 4 describes the current IGMP versions.

Table 4. IGMP Versions

IGMP Version	Description
IGMPv1	Basic Query-Response mechanism that allows the multicast router to determine which multicast groups are active.
IGMPv2	Extends IGMP, allowing IGMP leave process, group-specific queries an explicit maximum query response time.
IGMPv3	Provides for source filtering, which enables a multicast receiver host to signal to a router which groups it wants to receive traffic from, and from which sources this traffic is expected.

By default, Cisco routers use IGMPv2, but IGMPv2 routers also work correctly in the presence of IGMPv1 hosts.

IGMP Snooping

Traditionally, Layer 2 switches treat multicast as they would a broadcast, and they forward multicast traffic to every port that belongs to the destination VLAN (broadcast domain) on the switch. This approach degrades the performance of the switch and the hosts, which must inadvertently process and dispose of the unwanted traffic.

IGMP snooping is one of the mechanisms that constrain IP multicast traffic in a Layer 2 switching environment by examining or "snooping" some Layer 3 information (IGMP Join/Leave messages) in the IGMP packets that are sent between the host and the router.

When the switch sees an IGMP host report from a particular multicast group, the switch adds the host's switch port to the multicast entry table. When the host sends a Leave group message, the switch removes the table entry of the host.

In most Catalyst switches, IGMP snooping is enabled by default, globally and per VLAN. After the VLAN is configured for multicast routing, no other configuration is necessary for the switch to dynamically access external multicast routers by using IGMP snooping.

Note: Most commodity LAN switches do not have Silicon enabled IGMP snooping, which results in either flooding or a situation in which the switch cannot adequately scale to support IP based video surveillance applications.

PIM-Protocol Independent Multicast

To forward multicast traffic through the network, IP Multicast relies on routers and switches that support PIM or other multicast protocols.

Instead of building a new routing table, PIM relies on the unicast routing table to perform multicast forwarding functions. The unicast routing table can be populated by any routing protocol, such as Enhanced Interior Gateway Routing Protocol (EIGRP), Open Shortest Path First (OSPF), Border Gateway Protocol (BGP) or static routes.

PIM Multicast Group Modes

With PIM, packet traffic for a multicast group is routed according to the rules of the mode configured for that specific multicast group. The group mode is independent of how any interface is configured on the router. The main modes for a multicast group are:

- PIM sparse mode (Classic PIM-SM).
- PIM bidirectional mode.
- PIM Source Specific Multicast (SSM) mode.

Additionally, two different versions of PIM can be used, PIM version 1 and PIM version 2.

PIM Multicast Interface Modes

PIM can operate in sparse mode or in dense mode. The interface mode determines how the router populates its multicast routing table and how it forwards multicast traffic that it receives from its directly connected LANs. The interface mode is independent from a multicast group's operating mode.

The following are some of the PIM interface modes supported on Cisco routers:

PIM Dense Mode

PIM Dense Mode (PIM-DM) uses a source distribution tree to forward multicast traffic; routers build this tree as soon as the multicast source becomes active. Routers build the tree assuming that there are multicast receivers at every network segment. After the tree is built, branches with no active receivers are pruned from the tree.

Due to its poor scalability and flooding properties, PIM-DM is not the multicast forwarding method that is preferred by enterprise and service provider customers. It is most useful in a lab environment.

PIM Sparse Mode

PIM Sparse Mode (PIM-SM) assumes that the members of the multicast group are distributed throughout the network and that some members may be sparsely populated while others may be densely populated. PIM-SM is more efficient than PIM-DM. It begins with an empty tree and assumes that there are no multicast receivers, unless the receivers join a group via IGMP. This approach eliminates the need to flood multicast router interfaces with unnecessary data.

PIM-SM is the preferred multicast routing protocol due to its scalability.

PIM Sparse-Dense

Cisco implemented an alternative to choosing just dense mode or sparse mode on an interface. PIM sparse-dense mode supports PIM-SM and PIM-DM simultaneously on an interface.

With PIM Sparse-Dense, individual groups can use either PIM-SM or PIM-DM, depending on whether an RP is available for that group. If a particular group is configured with an RP, the router treats the group as sparse mode; otherwise the group is treated as dense mode. An RP is required if the group is to be treated as a sparse group

Auto-RP provides a way to distribute the group-to-RP mappings. Sparse-dense mode allows the Auto-RP information to be distributed in dense mode while the multicast groups can be treated as sparse mode. This approach eliminates the requirement to configure a default RP at the leaf routers.

Rendezvous Point

In PIM sparse mode, all multicast routers must know how to reach the router acting as the RP. Multicast sources first send traffic to the RP. When a receiver wants to receive data, it also registers with the RP. In turn, traffic is forwarded to the receivers using a shared distribution tree.

Selecting a RP router is not a difficult decision. The RP is required only to start sessions with sources and receivers. It represents little overhead for the router acting as the RP.

In a small network, the RP address can be manually configured on each PIM router, but this manual approach can be hard to manage as the network grows or if the RP address were to change.

Instead of defining a static RP address manually, methods such as Auto-RP and Bootstrap Router can be employed to dynamically discover the RP address. These methods have several benefits:

- Manual configuration of RP addresses is not required, reducing overhead and configuration conflicts
- Multiple RPs can be configured to serve different multicast groups
- Load-splitting can be achieved by configuring RPs for different groups

Auto-RP

A main advantage of Auto-RP is that changes to the RP address take place only at the routers that are defined as RP, and not the leaf routers. Auto-RP provides a way to automate the distribution of group-to-RP mappings in a PIM network that supports PIM-SM. With Auto-RP, multiple routers can be configured as RP to act as backup to other RPs or to server different group ranges.

The IANA has assigned two group addresses for Auto-RP: 224.0.1.39 and 224.0.1.40. Auto-RP requires two basic components:

- Candidate RPs: Routers acting as RP advertise their willingness to be an RP via "RP-announcement" messages. These messages are sent to well-known group 224.0.1.39.
- RP Mapping Agents: This router sends the group-to-RP mappings to all other routers and resolves any possible conflicts. RP mapping agents join the 224.0.1.39 group to map the RPs to the associated groups, and join the group address 224.0.1.40 to advertise these mappings. All PIM routers join 224.0.1.40 to learn these RP-mappings.

Bootstrap Router (BSR)

Another way to distribute group-to-RP mappings is Bootstrap Router, but BSR only works with PIM version 2. BSR also uses candidate routers to relay RP information for a group, but BSR are carried within PIM messages, which travel from PIM router to PIM router.

With BSR, all routers interfaces can be configured as PIM sparse-mode, so BSR does not run the risk of reverting to dense mode operation. BSR also makes sense in environments where non-Cisco, PIM version 2 routers need to interoperate with Cisco routers.

Initial Multicast Configuration

When deploying multicast for the first time on a Catalyst switch, some steps must be followed to ensure correct multicast behavior.

In Figure 22, a Cisco IP Gateway Encoder is configured to provide video streams to decoders or Cisco Stream Manager clients. This simple environment needs to provide proper multicast connectivity and the option to expand to other switches or connect to an existing environment using PIM.

Figure 22. Multicast on a Single VLAN

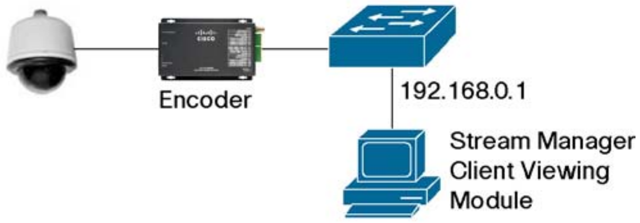

With a default configuration, all switch ports are on the same VLAN and IGMP snooping is enabled. While the switch is able to pass video traffic to be displayed on the Cisco Stream Manager Client Viewing Module, the switch is unable to constrain the multicast traffic to only the ports that need this multicast traffic.

Cisco IP Gateway Encoders and decoders support IGMP version 3, which is the default IGMP version for the Microsoft Windows XP operating system.

While video traffic flows to Cisco Stream Manager Client Viewing Modules, the traffic is flooded to all ports. In this case, IGMP snooping does not see an IGMP querier and is not aware of any multicast groups, as shown in the following example:

```
3750-1#show ip igmp snooping groups
3750-1#

3750-1#show ip igmp snooping  mrouter
Vlan    ports
----    -----

3750-1#show ip igmp snooping querier
Vlan      IP Address        IGMP Version   Port
```

Another simple way to verify that all ports in a given VLAN are not all being flooded is by visually inspecting the port LEDs in the front panel of the switch. If multicast is not enabled correctly, the switch will flood multicast traffic to all ports, affecting active ports that do not require the video stream traffic.

L2 Multicast Concepts

The IGMP protocol runs between a router and hosts, allowing the receivers to communicate with a multicast router (IGMP querier) to request multicast traffic.

IGMP snooping is designed to constrain multicast traffic to only the ports with active receivers, building a table that maps the multicast group to the ports that are requesting the video. Without IGMP snooping, the switch would flood packets to every port on the switch.

The Catalyst switch also maintains an mrouter port, which is the port that connects to the multicast router. The switch must have at least one mrouter port for IGMP snooping to work across switches.

There are several solutions to let the switch identify their mrouter port and make IGMP snooping work in this simple environment. This chapter presents two solutions and describes how to extend the IP PIM environment to a second switch.

Enable IGMP Querier Feature on an L2 Switch

If there are no multicast routers on a VLAN or there is no need to route multicast traffic to other subnets, a switch may be configured to act as the IGMP snooping querier.

This configuration is typical in a small environment, which may be as small as a single switch and no routers to provide multicast routing.

Figure 23. Initial Multicast Configuration

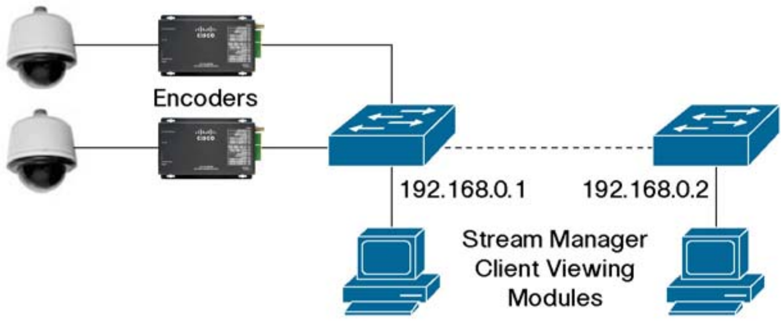

Enabling the **ip igmp snooping querier** command on the L2 switch constrains traffic to only the active multicast ports and makes the switch act as the querier for the VLAN.

An IP address must be configured on the VLAN interface to be used as the query source address. IP PIM is not required, making this configuration an attractive solution for L2 switches that do not support IP PIM or multicast-routing commands. If the switch is later connected to a multicast router, the IGMP querier functionality from this switch will be automatically disabled.

```
On 3750-1:
!
interface Vlan1
 ip address 192.168.0.1 255.255.255.0
!
ip igmp snooping querier

3750-1#show ip igmp snooping querier
Vlan      IP Address        IGMP Version   Port
-------------------------------------------------------------
1         192.168.0.1       v2             Switch

3750-1#show ip igmp snooping mrouter
Vlan    ports
----    -----
   1    Switch

3750-1#show ip igmp snooping groups
Vlan Group             Type    Version  Port List
----------------------------------------------------------------
1    235.1.1.1         igmp    v2       Gi1/0/10, Gi1/0/11
1    239.1.1.2         igmp    v2       Gi1/0/2, Gi1/0/11
1    239.1.1.3         igmp    v2       Gi1/0/2
1    239.255.0.10      igmp    v2       Gi1/0/1, Gi1/0/2
1    239.255.0.11      igmp    v2       Gi1/0/1, Gi1/0/2
1    239.255.255.250   igmp    v2       Gi1/0/1, Gi1/0/2
```

The following link provides a matrix with the IGMP support details for various Cisco Catalyst switches: http://www.cisco.com/en/US/customer/tech/tk828/technologies_tech_note09186a0080122a70.shtml#topic3.

Enable PIM on a Layer 3 or VLAN Interface

Catalyst switches are able to dynamically learn about the mrouter port by listening to either IGMP query messages or to PIM hellos from the multicast router.

Enabling PIM on an interface also enables IGMP on that interface, and the first switch on the network becomes the IGMP querier for the subnet. When additional PIM routers are added, the router with the lowest IP address on the subnet is elected as the IGMP querier.

The following configurations show a Catalyst 3750 switch acting as the IGMP querier by enabling the **ip pim sparse-dense-mode** command on the Vlan1 interface. The connection to a second Catalyst 3750 is done via an access or trunk port. This switch, 3750-2, has IGMP snooping enabled by default, so no other commands are necessary.

When enabling IP PIM on the interface for first time, the following message appears:

```
3750-1(config-if)#ip pim sparse-dense-mode
 WARNING: "ip multicast-routing distributed" is not configured.
           IP Multicast packets will not be forwarded
IP unicast CEF switching has to be enabled on the physical interface
```

Note: The ip multicast-routing command is not required if the multicast is to be constrained to a single VLAN and the multicast traffic is not to be forwarded to other interfaces or subnets.

```
on 3750-1:
!
interface Vlan1
 ip address 192.168.0.1 255.255.255.0
 ip pim sparse-dense-mode

show ip igmp groups
IGMP Connected Group Membership
Group Address   Interface      Uptime    Expires   Last Reporter
239.1.1.3       Vlan1          00:05:25  00:02:51  192.168.0.66
239.1.1.2       Vlan1          00:29:32  00:02:42  192.168.0.10
239.255.0.10    Vlan1          00:03:06  00:02:43  192.168.0.55
239.255.0.11    Vlan1          00:03:07  00:02:45  192.168.0.66
235.1.1.1       Vlan1          00:29:30  00:02:46  192.168.0.11
224.0.1.40      Vlan1          00:29:32  00:02:44  192.168.0.1

3750-1#show ip igmp snooping querier
Vlan      IP Address        IGMP Version   Port
---------------------------------------------------------------
1         192.168.0.1       v2             Router
```

```
3750-1#show ip igmp snooping groups
Vlan      Group         Type     Version  Port List
----------------------------------------------------------------
1         235.1.1.1     igmp     v2       Gi1/0/10, Gi1/0/15
1         239.1.1.2     igmp     v2       Gi1/0/10, Gi1/0/15
1         239.1.1.3     igmp     v2       Gi1/0/15
1         239.255.0.10  igmp     v2       Gi1/0/1, Gi1/0/15
1         239.255.0.11  igmp     v2       Gi1/0/1, Gi1/0/15


On 3750-2
!
interface Vlan1
 ip address 192.168.0.2 255.255.255.0


3750-2#show ip igmp snooping querier
Vlan      IP Address        IGMP Version   Port
---------------------------------------------------------------
1         192.168.0.1       v2             Gi1/0/15


3750-2#show ip igmp snooping groups
Vlan   Group           Type     Version     Port List
---------------------------------------------------------------
1      224.0.1.40    igmp      v2          Gi1/0/15
1      239.1.1.2     igmp      v2          Gi1/0/2, Gi1/0/15
1      239.1.1.3     igmp      v2          Gi1/0/2, Gi1/0/15
1      239.255.0.10  igmp      v2          Gi1/0/2, Gi1/0/15
1      239.255.0.11  igmp      v2          Gi1/0/2, Gi1/0/15


3750-2#show ip igmp snooping mrouter
Vlan    ports
----    -----
   1    Gi1/0/15(dynamic)
```

Enabling PIM on second switch:

A second switch to participate in the multicast routing can be added to the configuration that is shown in Figure 23. These two Layer 3 switches are now able to maintain IGMP and PIM tables and will be able to extend the multicast routing to other subnets. In this case, both switches will enable IP routing and multicast routing and will act as IP PIM neighbors.

Before enabling any routing commands, the 3750-2 switch is capable of IGPM snooping, but it is not aware of any IP PIM routing information:

```
3750-2#show ip pim neighbors
PIM Neighbor Table
Neighbor          Interface    Uptime/Expires    Ver   DR
Address                                                  Prio/Mode

#
```

To enable multicast routing on both switches, configure the **ip routing** and **ip multicast-routing distributed** commands on both switches. All interfaces also need the **ip pim sparse-dense-mode** command.

```
!
hostname 3750-1
!
ip multicast-routing distributed
ip routing
!
interface Vlan1
 ip address 192.168.0.1 255.255.255.0
 ip pim sparse-dense-mode

!
hostname 3750-2
ip routing
!
ip multicast-routing distributed
!
!
interface Vlan1
 ip address 192.168.0.2 255.255.255.0
 ip pim sparse-dense-mode

3750-1#show ip pim neighbor
PIM Neighbor Table
Neighbor      Interface     Uptime/Expires     Ver   DR
Address                                                Prio/Mode
192.168.0.2   Vlan1         00:08:35/00:01:32  v2    1 / DR S

3750-2#show ip pim neighbor
PIM Neighbor Table
Neighbor      Interface     Uptime/Expires     Ver   DR
Address                                                Prio/Mode
192.168.0.1   Vlan1         00:10:48/00:01:15   v2   1 / S
```

Both switches also are aware of the IGMP groups:

```
3750-1#show ip igmp groups
IGMP Connected Group Membership
Group Address    Interface    Uptime    Expires   Last Reporter
239.1.1.3        Vlan1        00:02:20  00:02:38  192.168.0.66
239.1.1.2        Vlan1        00:02:27  00:02:37  192.168.0.11
```

```
239.255.0.10     Vlan1        00:02:22  00:02:38  192.168.0.55
239.255.0.11     Vlan1        00:02:22  00:02:38  192.168.0.55
235.1.1.1        Vlan1        00:02:27  00:02:33  192.168.0.10
224.0.1.40       Vlan1        00:10:07  00:02:35  192.168.0.2

3750-2#show ip igmp groups
IGMP Connected Group Membership
Group Address    Interface    Uptime    Expires   Last Reporter
239.1.1.3        Vlan1        00:02:41  00:02:18  192.168.0.66
239.1.1.2        Vlan1        00:02:41  00:02:16  192.168.0.11
239.255.0.10     Vlan1        00:02:41  00:02:18  192.168.0.55
239.255.0.11     Vlan1        00:02:41  00:02:18  192.168.0.55
235.1.1.1        Vlan1        00:02:39  00:02:13  192.168.0.10
224.0.1.40       Vlan1        00:03:13  00:02:14  192.168.0.2
```

And both switches have the full IGMP snooping tables:

```
3750-1#show ip igmp snooping querier
Vlan      IP Address        IGMP Version   Port
---------------------------------------------------------------
1         192.168.0.1       v2             Router

3750-1#show ip igmp snooping groups
Vlan   Group          Type     Version     Port List
----------------------------------------------------------------
1      224.0.1.40     igmp                 Gi1/0/15
1      235.1.1.1      igmp     v2          Gi1/0/10, Gi1/0/15
1      239.1.1.2      igmp     v2          Gi1/0/10, Gi1/0/15
1      239.1.1.3      igmp     v2          Gi1/0/15
1      239.255.0.10   igmp     v2          Gi1/0/1, Gi1/0/15
1      239.255.0.11   igmp     v2          Gi1/0/1, Gi1/0/15

3750-2#show ip igmp snooping querier
Vlan      IP Address        IGMP Version   Port
---------------------------------------------------------------
1         192.168.0.1       v2             Gi1/0/15

3750-2#show ip igmp snooping groups
Vlan  Group         Type   Version Port List
---------------------------------------------------------------
1     224.0.1.40    igmp    v2     Gi1/0/15
1     235.1.1.1     igmp    v2     Gi1/0/11, Gi1/0/15
1     239.1.1.2     igmp    v2     Gi1/0/2, Gi1/0/11, Gi1/0/15
1     239.1.1.3     igmp    v2     Gi1/0/2, Gi1/0/15
1     239.255.0.10  igmp    v2     Gi1/0/2, Gi1/0/15
1     239.255.0.11  igmp    v2     Gi1/0/2, Gi1/0/15
```

While this simple configuration enables the two Layer 3 switches to fully participate in a multicast routing environment, other multicast features, such as the proper configuration of the RP, should be considered. The following section describes some methods for deploying the RP in a campus multicast environment.

Multicast Campus Deployment

Best-practices recommendations for a campus deployment can be followed when implementing the Cisco Video Surveillance solution in a network environment. These design principles have been tested extensively and have evolved to provide a solid foundation for existing and emerging applications. The hierarchical model has provided a design that is deterministic, highly available and easy to upgrade.

This section presents a design for a small campus network and a remote site that support IP multicast applications. Data and voice applications can also be supported in this design. The following design guides offer more information about how to deploy these technologies:

- Campus Design
- Unified Communications
- IP Multicast
- QoS
- High Availability

Appendix B provides links to these design guides.

The campus network example in this section consists of five different IDFs or closets, but these IDFs can also represent secondary buildings in a Campus network. The branch office consists of a single router and a Catalyst switch that supports one video stream. The WAN connectivity provides 8 MB of bandwidth via a multilink interface.

The VLAN numbers match the existing IP address range of the IDF and the Cisco IP Gateway Encoder number. For example, the IP address of IDF A is 10.1.30.x, on VLAN 30. HSRP has been configured between the two core/distribution switches and on the links connecting to the access switches.

For example, the address layout for IDF A or VLAN 30 is:

- 10.1.30.1 – Default gateway. Used by all host in IDF A
- 10.1.30.2 – IP address of active address (6504-1)
- 10.1.30.3 – IP address of standby router (6504-2)

6504-1 is the primary HSRP router for all VLANs. Figure 24 shows this topology in more detail.

Figure 24. Hierarchical Campus Design

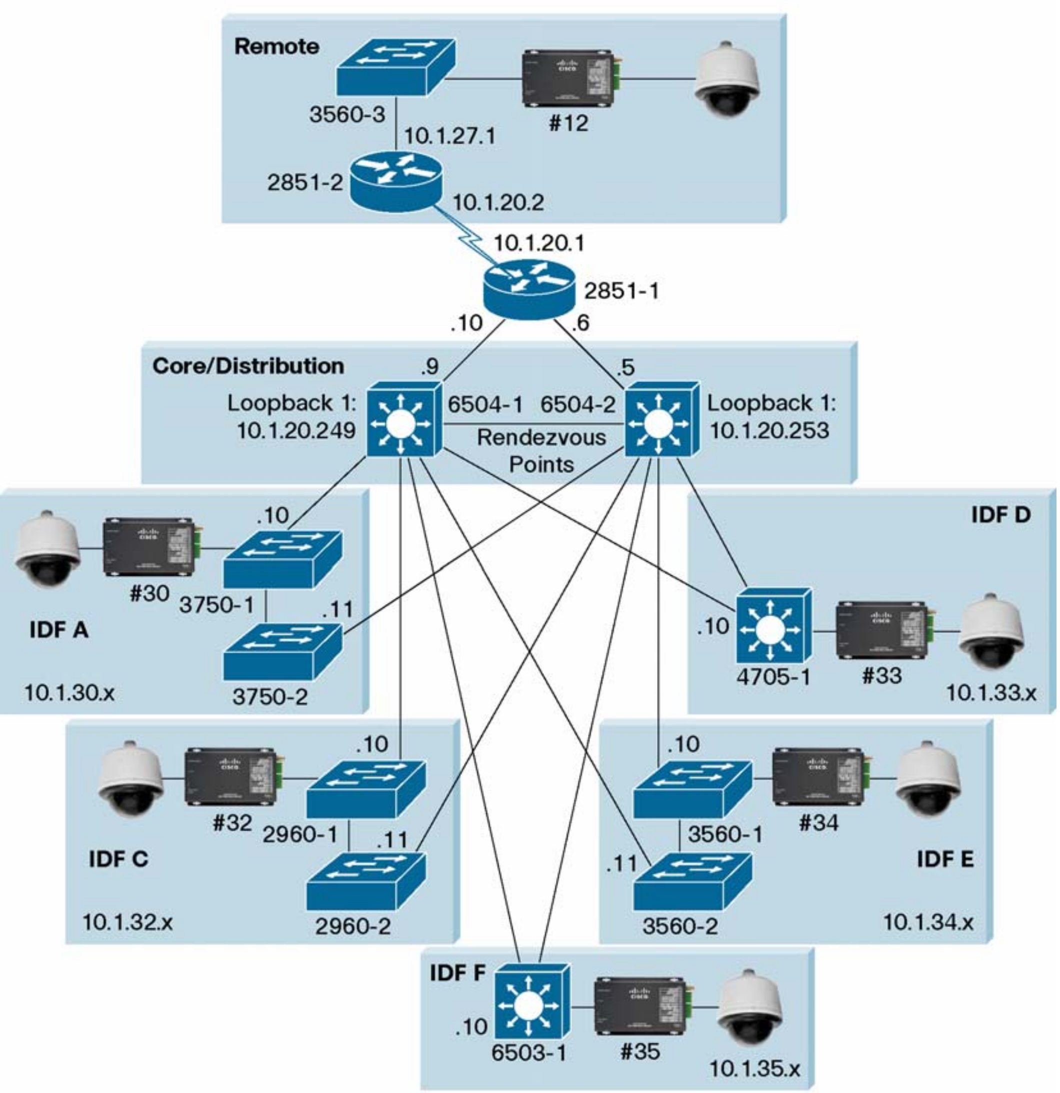

The connection between the core and the access switches is configured as a trunk, allowing only the required VLANs for that IDF to traverse the trunk. As an example, the following two interfaces connect between 6504-1 and IDFs A and B:

```
!
hostname 6504-1
!
interface GigabitEthernet4/1
 description 3750-1
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 30
 switchport trunk allowed vlan 30
 switchport mode trunk
 no ip address
!
interface GigabitEthernet4/4
 description 4507-1
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 33
 switchport trunk allowed vlan 33
 switchport mode trunk
 no ip address
```

The basic VLAN configurations for 6504-1 and 6504-2 are:

```
!
Hostname 6504-1
!
vtp domain ICE
vtp mode transparent
!
interface Vlan1
 no ip address
 shutdown
!
interface Vlan30
 description IDF A
 ip address 10.1.30.2 255.255.255.0
 standby 30 ip 10.1.30.1
 standby 30 priority 110
 standby 30 preempt
!
interface Vlan32
 description IDF C
 ip address 10.1.32.2 255.255.255.0
 standby 32 ip 10.1.32.1
 standby 32 priority 110
 standby 32 preempt
!
interface Vlan33
 description IDF D
 ip address 10.1.33.2 255.255.255.0
 standby 33 ip 10.1.33.1
 standby 33 priority 110
 standby 33 preempt
!
interface Vlan34
 description IDF E
 ip address 10.1.34.2 255.255.255.0
 standby 34 ip 10.1.34.1
 standby 34 priority 110
 standby 34 preempt
!
interface Vlan35
 description IDF F
 ip address 10.1.35.2 255.255.255.0
 standby 35 ip 10.1.35.1
 standby 35 priority 110
 standby 35 preempt
!
interface Vlan600
 ip address 10.94.162.194 255.255.255.192
```

```
!
hostname 6504-2
!
vtp domain ICE
vtp mode transparent

interface Vlan1
 no ip address
 shutdown
!
interface Vlan30
 description IDF A
 ip address 10.1.30.3 255.255.255.0
 standby 30 ip 10.1.30.1
 standby 30 priority 90
!
!
interface Vlan32
 ip address 10.1.32.3 255.255.255.0
 standby 32 ip 10.1.32.1
 standby 32 priority 90
!
!
!
interface Vlan33
 description IDF D
 ip address 10.1.33.3 255.255.255.0
 standby 33 ip 10.1.33.1
 standby 33 priority 90
!
!
interface Vlan34
 description IDF E
 ip address 10.1.34.3 255.255.255.0
 standby 34 ip 10.1.34.1
 standby 34 priority 90
!
!
interface Vlan35
 description IDF F
 ip address 10.1.35.3 255.255.255.0
 standby 35 ip 10.1.35.1
 standby 35 priority 90
!
interface Vlan600
ip address 10.94.162.250 255.255.255.192
!
```

The connection to the remote WAN location is created via two Cisco 2851 routers, using Multilink PPP to bundle several serial T1 interfaces. For redundancy, the Cisco 2851-1 also has two gigabit connections to the core switches.

```
!
hostname 2851-1
!
interface Multilink1
 ip address 10.1.20.1 255.255.255.252
 ppp multilink
 ppp multilink group 1
 max-reserved-bandwidth 100
!
interface GigabitEthernet0/0
 description to 6504-1 Port 2/33
 ip address 10.1.20.6 255.255.255.252
 duplex auto
 speed auto
!
interface GigabitEthernet0/1
 ip address 10.1.20.10 255.255.255.252
 duplex auto
 speed auto
!
interface Serial0/1/0
 no ip address
 encapsulation ppp
 clock rate 2016000
 clock rate 2016000
 dce-terminal-timing-enable
 ppp multilink group 1
 max-reserved-bandwidth 100
!
interface Serial0/1/1
 no ip address
 encapsulation ppp
 clock rate 2016000
 clock rate 2016000
 dce-terminal-timing-enable
 ppp multilink group 1
 max-reserved-bandwidth 100
!
. . .

(Output truncated)
```

Configuring a Rendezvous Point

While RPs can be deployed using static RP, that method does not provide redundancy or load balancing features and typically is used only in small deployments.

This chapter presents four methods for deploying RPs. Each method RP has its benefits and offers a different level of complexity.

- Static RP
- Auto-RP with sparse-dense mode
- Auto-RP with sparse-mode, using the Auto-RP listener feature
- Bootstrap router

The two core switches in Figure 24 (6504-1 and 6504-2) provide the RP functionality for this network.

All IDF switches are configured with IGMP snooping, which in most platforms is configured by default. The IDF switches also have a trunk connection to the core switches to provide VLAN connectivity for video and other applications.

Static RP

With Static RP, a router is selected as the RP for the network and all routers must be configured with this IP address manually.

In the following example, the loopback address of 6504-2 was selected as the RP and was manually configured in all routers.

```
!
ip multicast-routing
!
ip pim rp-address 10.1.20.253
!
```

In addition, Static RP supports only one RP. If the RP fails or changes, sparse mode stops operating and video surveillance traffic cannot traverse the network.

The following methods show alternative configuration methods, which overcome this limitation.

Auto-RP using sparse-dense-mode

IP multicast routing is enabled globally and all interfaces that are required to participate in the multicast domain have **ip pim sparse-dense-mode** enabled. The RPs are not manually configured on all routers. Instead, the two core switches are configured as the RP for the multicast domain using the IP PIM Auto-RP feature. This feature requires the configuration of an RP-mapping agent to determine which of the two core switches becomes the RP. The RP-mapping also provides group-to-RP mappings to all routers.

Appendix D provides full configurations for this section.

On both core switches, add the following commands:

```
!
ip multicast-routing
!
ip pim send-rp-announce loopback 1 scope 16
ip pim send-rp-discovery loopback 1 scope 16
!
```

With the **ip pim send-rp-announce** command, the core switch announces itself as the RP to the mapping agent, while the **ip pim send-rp-discovery** command sends discovery packets that tell other routers which group-to-RP mappings to use. The **scope** option determines how far (in hops) the message will be propagated.

On all interfaces that participate in the multicast domain, add the **ip pim sparse-dense-mode** interface command. The following example shows two VLAN interfaces and the connection to the Cisco 2851-1 enabled for ip pim sparse-dense-mode:

```
!
Hostname 6504-1
!
interface Vlan30
 description IDF A
 ip address 10.1.30.2 255.255.255.0
 ip pim sparse-dense-mode
 ip igmp version 3
 standby 30 ip 10.1.30.1
 standby 30 priority 110
 standby 30 preempt
!
interface Vlan32
 description IDF C
 ip address 10.1.32.2 255.255.255.0
 ip pim sparse-dense-mode
 ip igmp version 3
 standby 32 ip 10.1.32.1
 standby 32 priority 110
 standby 32 preempt
!

interface FastEthernet2/47
 description 2851-1
 ip address 10.1.20.9 255.255.255.252
 ip pim sparse-dense-mode
 ip igmp version 3
```

2851-1

```
!
ip multicast-routing
!
hostname 2851-1
!
interface Multilink1
 ip address 10.1.20.1 255.255.255.252
 ip pim sparse-dense-mode
 ppp multilink
 ppp multilink group 1
!
interface GigabitEthernet0/0
 description to 6504-1 Port 2/33
 ip address 10.1.20.6 255.255.255.252
 ip pim sparse-dense-mode
 ip igmp version 3
 duplex auto
 speed auto
!
interface GigabitEthernet0/1
 ip address 10.1.20.10 255.255.255.252
 ip pim sparse-dense-mode
 ip igmp version 3
 duplex auto
 speed auto
```

The **ip pim sparse-dense-mode** command allows the router to run in sparse mode if the RP is present. If RP is not available, the routers will switch to run in dense mode.

Every 60 seconds, a candidate RP sends an RP-Announcement message that details the group ranges for which it intends to serve as RP. This message is sent to 224.0.1.39. The PIM routers discover the RP information through the multicast address of 224.0.1.40.

The IDF switches are configured by default for IGMP snooping, so no additional configuration is required for them to join the multicast groups.

Verification commands

Basic show commands used to verify the multicast operations include the following:

Verify ip pim neighbors:

```
On the 6504-2:
6504-2#show ip pim neighbor
PIM Neighbor Table
Neighbor    Interface         Uptime/Expires    Ver   DR
Address                                                 Pri/Mode
10.1.20.6     FastEthernet2/33  00:27:34/00:01:16 v2  1 / DR S
10.1.20.18    FastEthernet2/38  00:56:01/00:01:25 v2  1 / DR S
10.1.31.2     Vlan31            00:56:33/00:01:24 v2  1 / S
10.1.32.2     Vlan32            00:56:33/00:01:20 v2  1 / S
10.1.33.2     Vlan33            00:56:33/00:01:16 v2  1 / S
10.1.34.2     Vlan34            00:56:33/00:01:23 v2  1 / S
10.1.42.2     Vlan42            00:56:33/00:01:21 v2  1 / S
10.1.43.2     Vlan43         00:56:33/00:01:15 v2  1 / S
10.94.165.2   Vlan500        00:56:27/00:01:27 v2  1 / S
10.1.37.2     Vlan600        00:56:18/00:01:28 v2  1 /
10.94.162.194 Vlan600        00:56:19/00:01:32 v2  1 / S
10.1.30.2     Vlan30         00:30:38/00:01:41 v2  1 / S

On the 2851-1:

2851-1#show ip pim neighbor
PIM Neighbor Table
Neighbor    Interface            Uptime/Expires     Ver   DR
Address                                                 Pri/Mode
10.1.20.2   Multilink1           4w6d/00:01:42      v2     1 / S
```

```
10.1.20.5   GigabitEthernet0/0   00:27:54/00:01:25 v2    1 / S
10.1.20.9   GigabitEthernet0/1   00:27:54/00:01:24 v2    1 / S
```

This command verifies that 6504-2 is currently the RP:

```
2851-1#show ip pim rp
Group: 235.1.1.1, RP: 10.1.20.253, v2, v1, uptime 00:18:35, expires
00:02:57
Group: 239.1.1.3, RP: 10.1.20.253, v2, v1, uptime 00:18:35, expires
00:02:57
Group: 239.255.27.12, RP: 10.1.20.253, v2, v1, uptime 00:18:35, expires
00:02:57
```

The 6504-2 is aware of the following IGMP groups:

```
6504-2#show ip igmp groups
IGMP Connected Group Membership
Group Address Interface        Uptime   Expires  Last       Reporter
235.1.1.1     Vlan30           00:30:38 00:02:53 10.1.30.5
235.1.1.1     Vlan32           00:56:30 00:01:56 10.1.32.5
235.1.1.1     Vlan34           00:56:31 00:02:59 10.1.34.5
235.1.1.1     Vlan31           00:56:32 00:02:59 10.1.31.5
235.1.1.1     Vlan33           00:56:32 00:01:57 0.0.0.0
235.1.1.1     Vlan600          00:56:37 00:01:25 10.94.162.217
239.1.1.2     Vlan600          00:56:37 00:01:25 10.94.162.217
239.1.1.3     Vlan600          00:56:37 00:01:25 10.94.162.217
239.255.30.5  Vlan600          00:56:37 00:01:27 10.94.162.206
239.255.27.12 Vlan600          00:16:41 00:01:26 10.94.162.206
224.0.1.39    FastEthernet2/38 00:17:31 00:02:12 10.1.20.17
224.0.1.39    FastEthernet2/33 00:17:31 00:02:30 10.1.20.5
224.0.1.39    Loopback1        00:17:31 00:02:23 10.1.20.253
224.0.1.39    Vlan30           00:18:25 00:02:59 10.1.30.2
224.0.1.39    Vlan600          00:18:25 00:01:22 10.94.162.250
224.0.1.39    Vlan500          00:18:25 00:02:02 10.94.165.2
224.0.1.39    Vlan43           00:18:25 00:02:00 10.1.43.2
224.0.1.39    Vlan42           00:18:25 00:02:02 10.1.42.2
224.0.1.39    Vlan34           00:18:25 00:02:59 10.1.34.2
224.0.1.39    Vlan33           00:18:25 00:02:59 10.1.33.3
224.0.1.39    Vlan32           00:18:25 00:02:58 10.1.32.2
224.0.1.39    Vlan31           00:18:25 00:02:54 10.1.31.3
224.0.1.40    Loopback1        00:57:33 00:02:24 10.1.20.253
239.255.162.222 Vlan600       00:56:05 00:01:23 10.94.162.222
```

The IDF switches have IGMP snooping enabled by default. This configuration can be verified with the following commands:

```
On 3750-1:
3750-1#show ip igmp snooping
Global IGMP Snooping configuration:
-----------------------------------
IGMP snooping              : Enabled
IGMPv3 snooping (minimal)  : Enabled
Report suppression         : Enabled
TCN solicit query          : Disabled
TCN flood query count      : 2
```

```
Last Member Query Interval : 1000


Vlan 30:
--------
IGMP snooping                       : Enabled
IGMPv2 immediate leave              : Disabled
Explicit host tracking              : Enabled
Multicast router learning mode      : pim-dvmrp
Last Member Query Interval          : 1000
CGMP interoperability mode          : IGMP_ONLY

3750-1#show ip igmp snooping groups
Vlan  Group         Type        Version     Port List
---------------------------------------------------------------
30    235.1.1.1     igmp        v3          Gi1/0/2, Gi1/0/5
```

Auto-RP using the Auto-RP Listener

Another way to propagate the RPs dynamically is to use Auto-RP with the autorp listener feature. The **ip pim autorp listener** global configuration command can enable Auto-RP to operate on interfaces that are configured for PIM sparse-mode. Before this command was available, Auto-RP required sparse-dense-mode. This command has been available since Cisco IOS® Software version 12.2(7).

The autorp listener feature maintains the availability of Auto-RP as the mechanism for distributing Group-to-RP mappings, which allows the two Auto-RP multicast groups (224.0.1.40 and 224.0.1.39) to function as dense-mode and the interfaces to function in sparse-mode, thereby eliminating the potential for dense mode flooding across a multicast domain.

To configure Auto-RP using the autorp listener feature:

On both core switches, add:

```
!
ip multicast-routing
!
ip pim autorp listener
ip pim send-rp-announce Loopback 1 scope 16
ip pim send-rp-discovery Loopback 1 scope 16
!
```

On all interfaces that participate in the multicast domain, add the **ip pim sparse-mode** interface command. The following example shows two VLAN interfaces and the connection to the Cisco 2851-1 enabled for ip pim sparse-mode:

```
!
Hostname 6504-1
interface Vlan30
 description IDF A
 ip address 10.1.30.2 255.255.255.0
 ip pim sparse-mode
 ip igmp version 3
 standby 30 ip 10.1.30.1
 standby 30 priority 110
 standby 30 preempt
!
interface Vlan32
 description IDF C
 ip address 10.1.32.2 255.255.255.0
 ip pim sparse-mode
 ip igmp version 3
 standby 32 ip 10.1.32.1
 standby 32 priority 110
 standby 32 preempt
!
interface FastEthernet2/47
 ip address 10.1.20.9 255.255.255.252
 ip pim sparse-mode
 ip igmp version 3

!
ip multicast-routing
!
hostname 2851-1
!
interface Multilink1
 ip address 10.1.20.1 255.255.255.252
 ip pim sparse-mode
 ppp multilink
 ppp multilink group 1
!
interface GigabitEthernet0/0
 description to 6504-1 Port 2/33
 ip address 10.1.20.6 255.255.255.252
 ip pim sparse-mode
 ip igmp version 3
 duplex auto
 speed auto
!
interface GigabitEthernet0/1
 ip address 10.1.20.10 255.255.255.252
 ip pim sparse-mode
 ip igmp version 3
 duplex auto
 speed auto
```

Verification Commands:

To verify the PIM neighbors:

```
On the 6504-2:

6504-2#show ip pim neighbor
PIM Neighbor Table
Neighbor      Interface          Uptime/Expires    Ver   DR
```

```
Address                                                Pri/Mode
10.1.20.6     FastEthernet2/33   00:05:37/00:01:31 v2  1 / DR S
10.1.20.18    FastEthernet2/38   00:05:36/00:01:35 v2  1 / DR S
10.1.30.2     Vlan30             00:06:07/00:01:32 v2  1 / S
10.1.31.2     Vlan31             00:06:07/00:01:31 v2  1 / S
10.1.32.2     Vlan32             00:06:07/00:01:30 v2  1 / S
10.1.33.2     Vlan33             00:06:07/00:01:31 v2  1 / S
10.1.34.2     Vlan34             00:06:07/00:01:32 v2  1 / S
10.1.42.2     Vlan42             00:06:07/00:01:31 v2  1 / S
10.1.43.2     Vlan43             00:06:07/00:01:29 v2  1 / S
10.94.165.2   Vlan500            00:06:02/00:01:36 v2  1 / S
10.1.37.2     Vlan600            00:05:53/00:01:17 v2  1 /
10.94.162.194 Vlan600            00:05:54/00:01:44 v2  1 / S
```

This command is used to verify that the 6504-2 is the RP:

```
6504-2#show ip pim rp
Group: 239.255.27.12, RP: 10.1.20.253, v2, v1, next RP-reachable in
00:01:27
Group: 239.255.30.5, RP: 10.1.20.253, v2, v1, next RP-reachable in
00:01:17
Group: 239.255.255.253, RP: 10.1.20.253, v2, v1, next RP-reachable in
00:00:20
Group: 239.255.255.250, RP: 10.1.20.253, v2, v1, next RP-reachable in
00:00:15
Group: 239.255.162.222, RP: 10.1.20.253, v2, v1, next RP-reachable in
00:00:19
Group: 235.1.1.1, RP: 10.1.20.253, v2, v1, next RP-reachable in 00:01:17
Group: 239.1.1.3, RP: 10.1.20.253, v2, v1, next RP-reachable in 00:01:17
Group: 239.1.1.2, RP: 10.1.20.253, v2, v1, next RP-reachable in 00:01:17
```

The 6504-2 is aware of the following IGMP groups:

```
6504-2#show ip igmp groups
IGMP Connected Group Membership
Group Address  Interface     Uptime    Expires   Last Reporter
235.1.1.1    Vlan32          00:06:05  00:02:46  10.1.32.5
235.1.1.1    Vlan34          00:06:06  00:02:40  10.1.34.5
235.1.1.1    Vlan31          00:06:07  00:02:42  10.1.31.5
235.1.1.1    Vlan30          00:06:07  00:02:09  10.1.30.5
235.1.1.1    Vlan33          00:06:07  00:02:10  0.0.0.0
235.1.1.1    Vlan600         00:06:12  00:01:44  10.94.162.219
239.1.1.2    Vlan600         00:06:12  00:01:46  10.94.162.219
239.1.1.3    Vlan600         00:06:12  00:01:51  10.94.162.233
239.255.30.5 Vlan600         00:06:12  00:01:45  10.94.162.206
. . .

(Output truncated)
```

The IDF switches have IGMP snooping enabled by default. This configuration can be verified with the following commands:

```
On 4507-1:
4507-1#show ip igmp snooping
Global IGMP Snooping configuration:
-----------------------------------
IGMP snooping             : Enabled
IGMPv3 snooping           : Enabled
Report suppression        : Enabled
TCN solicit query         : Disabled
TCN flood query count     : 2


Vlan 33:
--------
IGMP snooping                       : Enabled
IGMPv2 immediate leave              : Disabled
Explicit host tracking              : Enabled
Multicast router learning mode      : pim-dvmrp
CGMP interoperability mode          : IGMP_ONLY


4507-1#show ip igmp snooping groups
Vlan      Group          Version     Port List
---------------------------------------------------------
33        235.1.1.1      v3          Gi3/5
```

Bootstrap Router (BSR)

Bootstrap router (BSR), which was introduced with PIM version 2, provides another way to provide dynamic group-to-RP mappings. BSR is an open standard, while Auto-RP is proprietary to Cisco. The functionality provided by BSR is similar to Auto-RP and uses candidate routers for RP function and for relaying the RP information for a group. This information is distributed through BSR messages, which are carried within PIM messages.

Configuring BSR is similar to configuring Auto-RP, with the exception that the routers acting as the BSR. To configure BSR on both core switches:

```
!
ip multicast-routing
!
ip pim bsr-candidate Loopback1 0 192
ip pim rp-candidate Loopback1 priority 192
```

All interfaces participating in the multicast domain must have the command:

```
!
Interface x
  ip pim sparse-mode
```

The IDF switches also use IGMP snooping, which is enabled by default. The same verification commands used in the previous sections can be used to view the RP mapping and to verify the active IGMP groups.

Note: BSR lacks the ability to scope RP advertisements and cannot be deployed simultaneously with Auto-RP

Time Synchronization (NTP)

Network Time Protocol (NTP) is widely used to synchronize clocks of hosts, routers and other network elements in the Internet with a reliable time source. The Cisco Video Surveillance solution uses NTP to synchronize the time of its Cisco IP Gateway Decoders, recorder platforms and the Cisco Stream Manager software. Clock synchronization is critical when retrieving previously recorded video streams.

With the Cisco Video Surveillance solution, any of the recorder platforms, including the Service Platform and Integrated Service Platform can act as Time Master while the Cisco Stream Manager software running on them implements the NTP client functionality.

Figure 25. NTP Synchronization

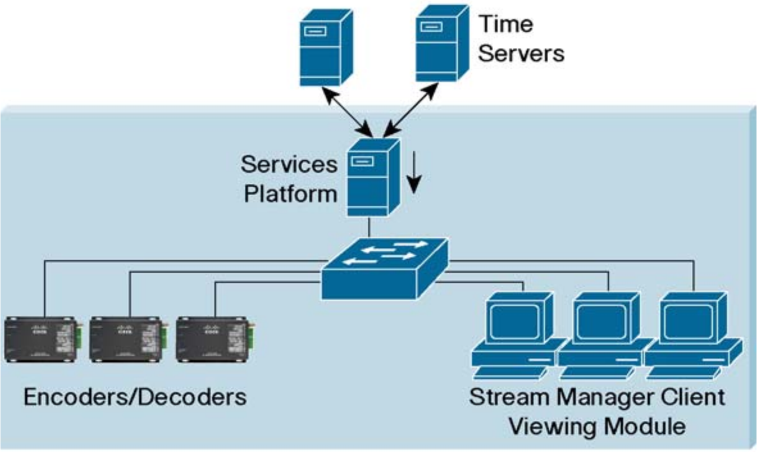

The Services Platforms relies on GMT time and must be configured with the correct time zone. PCs running the Cisco Stream Manager software obtain the time from the Time Master and must also be configured with the correct time zone.

When NTP is enabled on a Time Master platform, it makes provisions to configure two redundant external high stratum time sources (NTP servers) and estimates the accurate error of the local clock with reference to the time source to which it may be synchronized.

In a typical video surveillance solution, only one recorder platform is configured as Time Master. After this platform is time synchronized, it uses multicast to synchronize time on other devices, including Cisco IP Gateway Decoders, recorder platforms and Cisco Stream Manager Client Viewing Modules.

To configure the recorder platform as the Time Master, use the Cisco Stream Manager Configuration Module and select the appropriate recorder. Under the Device Settings tab, set the mode to ‘Time Master.’ Under the Time Master Setting’s pane, choose Sync Source to be NTP and set IP addresses of the primary and secondary servers. Click ‘update’ to save changes and reload the platform.

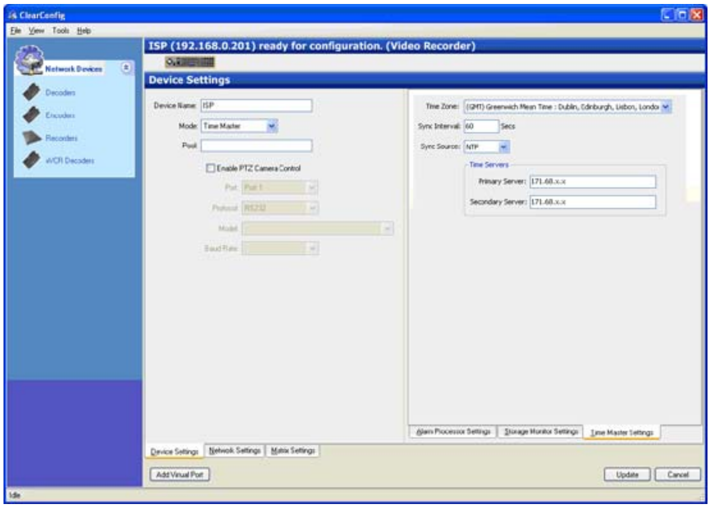

ISP (192.168.0.201) ready for configuration. (Video Recorder)

Device Settings

Device Name: ISP

Mode: Time Master

Pool:

☐ Enable PTZ Camera Control

Port: Port 1

Protocol: RS232

Model:

Baud Rate:

Time Zone: (GMT) Greenwich Mean Time : Dublin, Edinburgh, Lisbon, London

Sync Interval: 60 Secs

Sync Source: NTP

Time Servers

Primary Server: 171.68.x.x

Secondary Server: 171.68.x.x

To configure the Cisco Stream Manager software to synchronize the PC’s clock with the Time Master, select **Tools, Options** and select the General tab to enable **Synchronize clock with system.**

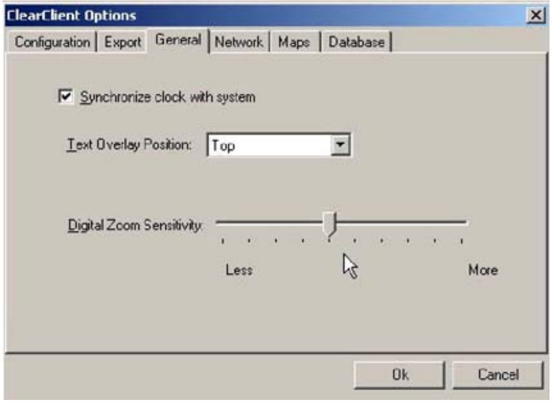

☑ Synchronize clock with system

Text Overlay Position: Top

Digital Zoom Sensitivity:

This procedure also synchronizes the PC’s clock with the NTP Time Master.

Host Addressing – DHCP

Dynamic Host Configuration Protocol (DHCP) allows a device to obtain IP address configuration from a server that is capable of maintaining and distributing unique address information. By using DHCP, a device connecting to a subnet dynamically learns its IP address and can be granted access to the network.

The current version of the Cisco Stream Manager Configuration Module (4.10.7) does not provide support for DHCP, but this feature will be added in a subsequent release. Currently, Cisco Video Surveillance products ship with an IP address that is in the 192.168.0.x subnet and that can easily be readdressed by using the Cisco Stream Manager Configuration Module.

Appendix C explains the steps that are used to configure a Cisco IP Gateway.

Chapter 5: Security

Protecting the Video Infrastructure

Security is one of the top concerns of enterprise networks, which utilize different layers of protection that are inherent in Cisco routers, switches and firewalls. Video surveillance is considered to be a mission-critical application that must be protected from unauthorized access or other forms of attack.

Cisco develops network-wide security products and technologies that enable the Cisco Self-Defending Network architecture. These products vary from security appliances to highly integrated hardware and software running on Cisco routers and Catalyst switches.

This chapter focuses on the following security features:

- Segmentation using Virtual LANs
- Network policies using access lists
- Protecting switch ports using Port Security
- Using stateful firewalls to protect video streams

Segmentation using Virtual LANs

Virtual LANs (VLANs) can be used to separate video surveillance traffic from other types of traffic. A VLAN allows a Catalyst switch to isolate traffic from different switch ports into different groups. These ports can be grouped by function, user community or application and can span across modules or switches.

The concept of VLANs is well understood and used by many enterprise networks. The following is a short review of the main concepts associated with VLANs:

- A broadcast domain is a set of devices for which a frame that is sent by one device is received by all devices in the same broadcast domain.
- A VLAN is fundamentally a broadcast domain
- By default, switch ports are configured in VLAN 1, but they can be assigned to participate in any VLAN
- Layer 2 switches only forward frames between devices in the same VLAN
- A Layer 3 switch or router is required to forward (route) packets between VLANs
- The devices in a VLAN typically also are in the same IP subnet
- Devices in different VLANs cannot communicate without a router

Figure 26 shows two VLANs configured on a Catalyst switch. VLAN 10 is dedicated to Cisco IP Gateway Encoders and Decoders, and VLAN 11 is reserved for the rest of the devices, such as PCs.

Devices that are assigned to VLAN 11 cannot view surveillance media streams, since they are only allowed in VLAN 10, but the monitoring station on VLAN 10 will have full access.

Figure 26. Virtual LANs

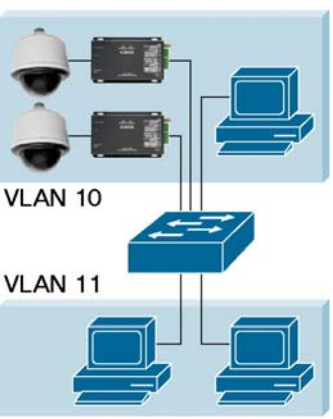

With VLAN trunking, multiple VLANs can be carried between switches across a single physical connection. These interconnected switches participate in the same Virtual Trunking Protocol (VTP) domain to provide access to the same VLANs.

Segmentation using Access Lists

In addition to segmenting the video surveillance environment with VLANs, Cisco routers and switches have a more granular filtering mechanism that uses Access Control Lists.

An Access Control List (ACL) provides a sequential list of **permit** and **deny** statements that are analyzed by a Cisco device to enforce different security policies. There are several general types of ACLs, such as numbered, named, time-based and rate-limiting. Each one is responsible for matching or filtering different protocol information.

ACLs also are used to block traffic from specified IP addresses or protocols and can used to filter and protect video surveillance media. High-level guidelines for ACLs include:

- An interface allows a maximum of one outbound ACL and one inbound ACL. These ACLs are specified with the **in** and **out** keywords.
- Since the Access Control Entries (ACEs) are analyzed from top to bottom, the order of the statements is critical. The router will stop processing an ACL as soon as it finds a match.
- Access lists always have an implicit *deny* statement at the end of the list. This statement denies or blocks any traffic that has not been explicitly permitted in the access list. The statement is never displayed and cannot be displayed with **show** commands.
- Access lists have a specific function that is defined by the **permit** and **deny** statements. Every ACL should be thoroughly tested to make sure that it is performing as expected.

The following two examples show ways to configure ACLS to do the following:

- Allow only specific Cisco IP Gateway Encoders and Decoders
- Allow a single video stream in dual-streaming configurations

These general examples focus on protecting the surveillance media stream, but they can be expanded to allow or deny other types of traffic or to match a specific network environment.

ACLs are simple to deploy and effective in limiting Cisco Video Surveillance streams. A larger deployment requires the use of administrative scoped zones to define multicast regions that are appropriate to a specific network environment. RFC 2365 provides best practices to deploy Administratively Scoped IP Multicast. Appendix B provides a link to this RFC.

Figure 27 shows a subnet with surveillance media traffic that needs to be protected to ensure that only authorized users are able to view surveillance video streams and that no other Cisco IP Gateway Encoders or Decoders are added to the subnet.

Assume that VLAN 600 contains recorders and Cisco IP Gateways that are dedicated to the security operation, and that IDF C is houses the Cisco IP Gateway Encoders and Decoders.

Figure 27. Access Lists

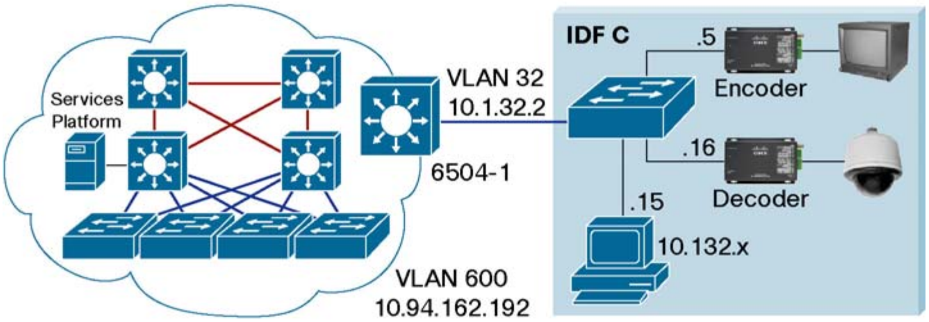

Example 1: Allow specific Cisco IP Gateways

One approach to selectively allow traffic from Cisco IP Gateway Encoders and Decoders is to specify what devices are allowed through by the access list, which relies on the implicit deny at the end of the access list to block all other traffic.

This approach requires more configuration commands and maintenance than other approaches, but it blocks unauthorized users from reaching surveillance media traffic.

The following two Access Control Entries (ACEs) allow the video surveillance control traffic. In this example, the **permit udp** commands allow traffic from any source, but they could be more restrictive and allow only traffic from specific subnets or devices:

```
!
hostname 6504-1
!
ip access-list extended Select_IP_Gateways
 permit udp any 239.0.0.0 0.255.255.255
 permit udp any host 235.1.1.1
```

To allow traffic from a specific Cisco IP Gateway to traverse the VLAN, the following commands may be used. These commands allow streams generated by the Cisco IP Gateway Encoder to be seen at other network locations. The 10.94.162.192 subnet houses all Service Platform recorders and other Cisco IP Gateway Encoders and Decoders and is permitted to reach the Cisco IP Gateway Encoder at 10.1.32.5.

```
!
hostname 6504-1
!
ip access-list extended Select_IP_Gateways
 permit ip 10.1.0.0 0.0.63.255 host 10.1.32.5
 permit ip 10.94.162.192 0.0.0.63 host 10.1.32.5
```

The following is a complete access list that allows traffic to a Cisco IP Gateway Encoder, to a Cisco IP Gateway Decoder and to the Cisco Stream Manager Client Viewing Module.

```
!
interface Vlan32
 ip address 10.1.32.2 255.255.255.0
 ip access-group Select_IP_Gateways out
!
ip access-list extended Select_IP_Gateways
 permit udp any host 235.1.1.1
 permit udp any 239.0.0.0 0.255.255.255
 remark **** Allow the Stream Manager Client ****
 permit ip 10.1.0.0 0.0.63.255 host 10.1.32.15
 permit ip 10.94.162.192 0.0.0.63 host 10.1.32.15
 remark **** Allow Cisco IP Gateway Decoder traffic ****
 permit ip 10.1.0.0 0.0.63.255 host 10.1.32.16
 permit ip 10.94.162.192 0.0.0.63 host 10.1.32.16
 remark **** Allow Cisco IP Gateway Encoder traffic ****
 permit ip 10.1.0.0 0.0.63.255 host 10.1.32.5
 permit ip 10.94.162.192 0.0.0.63 host 10.1.32.5
!
```

Note that the implicit deny at the end of the access list denies all other traffic, preventing unauthorized Cisco IP Gateway Encoders and Decoders from connecting to the network and preventing unauthorized Cisco Stream Manager devices from viewing any video streams.

This list can be expanded to deny traffic from specific devices. The following access list can be used to block traffic to the Cisco Stream Manager device with the IP address 10.1.32.15 and still allow traffic surveillance media traffic to other devices:

```
!
Hostname 6504-1
!
interface Vlan32
 ip address 10.1.32.2 255.255.255.0
 ip access-group Block_Client in
 ip access-group Select_IP_Gateways out
!
ip access-list extended Block_Client
 deny udp host 10.1.32.15 any
 permit ip any any
```

Example 2: Dual Streaming Encoders - Block a single stream

Single-port Cisco IP Gateway Encoders can provide two streams of the same video source, each with unique resolutions and frame rates. This feature is useful when the same surveillance media stream needs to be delivered to two different locations, each with different bandwidth capabilities.

Figure 28 shows the Cisco Stream Manager Configuration screen that enables the use of two different Media Streams:

Figure 28. Dual Stream Configuration

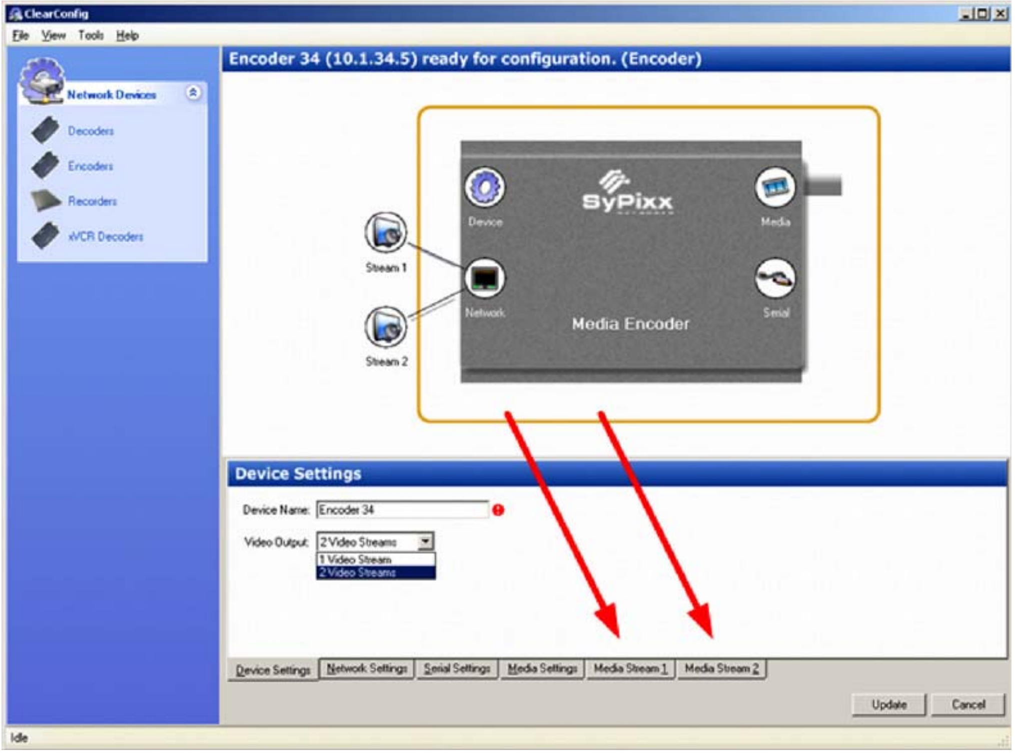

Encoder 34 (10.1.34.5) ready for configuration. (Encoder)

Device Settings

Device Name: Encoder 34

Video Output: 2 Video Streams

The UDP Ports that are used by the Cisco IP Gateways for multicast delivery are calculated as follows:

- For Video: 10,000 + (Camera ID * 2)
- For Audio: 10,000 + (Camera ID * 2) + 1

For example, a video stream from camera ID 34 would use port 10,068 and a video stream from camera ID 44 would use port 10,088.

While the example access control lists that are described in the previous sections can block all UDP traffic from going to specific devices, they cannot block just a single stream that comes from a Cisco IP Gateway Encoder. To block the second stream, defined as Camera 44, the ACL can be expanded to include UDP port number 10088:

```
!
hostname 6504-1
!
interface Vlan32
 ip address 10.1.32.2 255.255.255.0
 ip access-group Block_by_Stream out
!
ip access-list extended Block_by_Stream
 deny udp host 10.1.34.5 any eq 10088
 permit ip any any
```

This ACL allows the Cisco IP Gateway Encoder to still appear in the Cisco Stream Manager Configuration Module, and blocks only traffic from the second stream, or Camera 44 (UDP Port 10,088). This option is attractive in environments with limited bandwidth that need to allow only specific video streams.

Port Security

The Cisco port security feature provides an additional layer of protection to restrict unauthorized devices from connecting to switch ports. Port security identifies the MAC addresses of the hosts that can access a particular switch port and maintains a table with the MAC addresses that are allowed to connect.

Port Security is an attractive feature for protecting Cisco Video Surveillance components, including Cisco IP Gateway Encoders, Cisco IP Gateway Decoders and the Services Platforms. Without port security, an unauthorized device could be plugged into a switch port and gain full access to video surveillance video streams.

Two main modes can be configured on a switch port, which allows a system administrator to:

- Specify what MAC addresses can connect to a switch port
- Limit how many MAC addresses can be active on a single port

In port security terminology, a security violation occurs when either:

- The maximum number of MAC addresses has been added to the address table and a new host tries to connect to the port
- An address learned or configured on a secure port is seen on another interface in the same VLAN

If a violation occurs, the port can be configured to do one of the following:

- Shutdown – The port is error-disabled and is shut down immediately. SNMP trap and syslog messages are generated.
- Protect – The port continues to operate but drops frames from newer hosts when the maximum learned address has been exceeded.
- Restrict – Same as protect but also generates an SNMP trap and syslog message.

Figure 29. Port Security

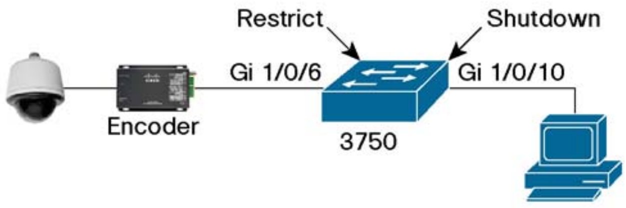

Figure 29 shows a Cisco IP Gateway Encoder connecting to port 1/0/6 on a Catalyst 3750. To prevent unauthorized devices from plugging into the port, the switch is configured with port security. The switch is not configured to shut down the port, but instead to restrict the port to only the encoder’s MAC address. In addition, a Cisco Stream Manager Client Viewing Module is configured on interface 1/0/10. Because this workstation is in a public area, port security is configured to shut down the port if an unauthorized device plugs into that port.

```
!
interface GigabitEthernet1/0/6
 description Encoder #30
 switchport access vlan 30
 switchport mode access
 switchport port-security
 switchport port-security violation restrict
 switchport port-security maximum 1
 switchport port-security mac-address 0011.a600.098c
 spanning-tree portfast
!
!
interface GigabitEthernet1/0/10
 description Stream Manager Client
 switchport access vlan 30
 switchport mode access
 switchport port-security
 switchport port-security violation shutdown
 switchport port-security maximum 1
 switchport port-security mac-address 000a.32af.4125
 spanning-tree portfast
!
```

If a port security violation occurs, the switch generates a syslog message and an SNMP trap. In the following message, the unauthorized device with MAC address 000a.412a.3af0 has tried to connect to interface gi1/0/10, which has generated a security violation and placed the interface into a disabled state.

```
21:48:39: %PM-4-ERR_DISABLE: psecure-violation error detected on
Gi1/0/10, putting Gi1/0/10 in err-disable state
21:48:39: %PORT_SECURITY-2-PSECURE_VIOLATION: Security violation
occurred, caused by MAC address 000a.412a.3af0 on port
GigabitEthernet1/0/10.
21:48:40: %LINEPROTO-5-UPDOWN: Line protocol on Interface
GigabitEthernet1/0/10, changed state to down
21:48:41: %LINK-3-UPDOWN: Interface GigabitEthernet1/0/10, changed state
to down
```

The **show port security** command can be used to verify this situation:

```
show port-security interface gigabitEthernet 1/0/10
Port Security              : Enabled
Port Status                : Secure-shutdown
Violation Mode             : Shutdown
Aging Time                 : 0 mins
Aging Type                 : Absolute
SecureStatic Address Aging : Disabled
Maximum MAC Addresses      : 1
Total MAC Addresses        : 1
Configured MAC Addresses   : 1
Sticky MAC Addresses       : 0
Last Source Address:Vlan   : 000a.412a.3af0:30
Security Violation Count   : 1
```

Firewalls

Network attacks can seriously affect network performance and business productivity. Network administrators should consider the security of their networks to be a key concern. Video surveillance is a critical business application and must be protected from attacks and unauthorized access. Firewalls provide an extra layer of protection against these attacks and can easily be used to segment a video surveillance environment.

Traditionally, GRE tunneling was used to allow multicast traffic to traverse firewalls because firewalls were not capable of participating in multicast routing. The most effective way to pass multicast through a firewall was to embed it within a logical point-to-point tunnel between router pairs. The challenges with this approach are that it involves a lot of manual configuration and maintenance and that that it bypasses the firewall's inspection engine. A simple misconfiguration can allow traffic that is not within the security policy to slip past the firewall via the tunnel. Cisco security appliances can fully participate in a PIM/IGMP multicast environment.

This section provides the basic steps to configure a Cisco ASA Security Appliance or a PIX firewall to protect a video surveillance environment. This environment may be a second campus location or a separate network subnet. Firewalls can be deployed to protect any subnet or network location.

One scenario is to deploy a firewall to protect Cisco IP Gateway Encoders and Decoders at the primary location (where the Services Platforms may be located) and to deploy additional firewalls to protect Cisco IP Gateway Encoders deployed at other locations.

Figure 30. Firewall Deployments

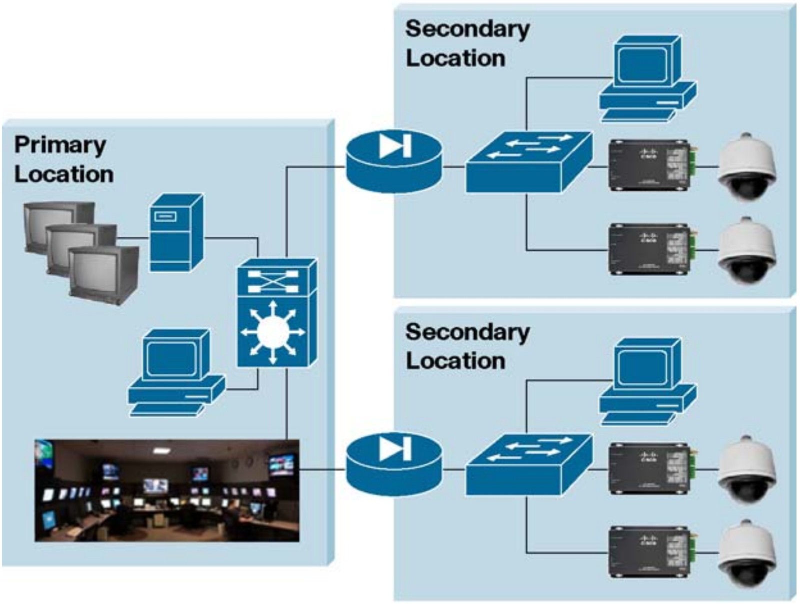

Detailed Topology

Figure 31 shows a detailed topology protecting a secondary location. This topology includes IP addresses and the location of the Rendezvous Point, which is located in the Primary Campus.

Figure 31. Topology

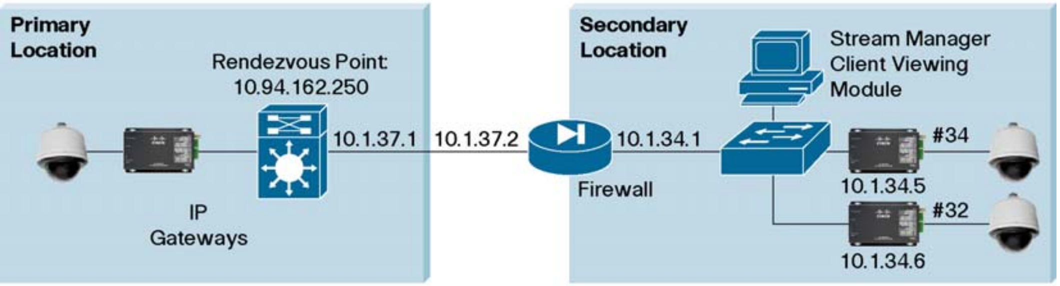

This topology shows a second location that is protected by a firewall. The Cisco Stream Manager Client Viewing Module should be able to view all video streams from both locations and devices in the primary location should be able to view and manage Encoders #32 and #34, which are located at the secondary location. These concepts apply to an environment with a separate building in a campus or simply a separate wiring closet.

Multicast Support

While multicast support was traditionally supported via GRE Tunnels, the Cisco ASA Security Appliance can participate in a multicast environment with routers running Protocol-Independent Multicast (PIM) Sparse Mode, eliminating the need for the GRE tunnel approach.

Version 6.x of the PIX software had limited multicast support. It only supported *IGMP proxy agent*, or *stub router*, and was only able to forward IGMP messages between hosts and multicast routers.

Beginning with Version 7, both the PIX and Cisco ASA appliance can fully participate in multicast routing and support PIM-SM and bi-directional PIM. This capability allows the firewall to distribute multicast traffic dynamically with other PIM routers.

Basic Firewall Configuration

By default, the firewall is configured to block traffic that originated from outside interfaces. The configurations that are used in the following sections were collected using a Cisco ASA 5500 Adaptive Security Appliance (v.7.2.1), but the same commands apply to a PIX Security Appliance.

The following basic steps configure the host name, and two interfaces: Outside and Inside. The security levels on these interfaces are typically assigned as a default, where a default security level of 0 is assigned to the outside interface and 100, the highest security level is assigned to the inside interface. Other networks, such as DMZ or server networks, can be configured with a security level somewhere between 0 and 100.

The following commands were applied using the console port:

```
#configure terminal

hostname ASA-1
names
!
interface GigabitEthernet0/0
 description Outside
 nameif Outside
 security-level 0
 ip address 10.1.37.2 255.255.255.0
 no shutdown
!
interface GigabitEthernet0/1
 description Inside
 nameif inside
 security-level 100
 ip address 10.1.34.1 255.255.255.0
 no shutdown
!
```

User Authentication

The following commands change the default password and enable local authentication, which provides more restrictive administrative access control:

```
!
passwd cisco123
enable password cisco123
!
```

The **passwd** command is used for Telnet and SSH connections. By default, this password is set to “cisco”. The **enable password** sets the password for the privileged EXEC mode, which is the highest privilege level.

IP Routing

For routing purposes, the ASA supports static routes and routing protocols such as RIP and OSPF. To communicate with other networks, the ASA requires a routing table.

A default route is simply a static route to 0.0.0.0/0 as the destination IP address and is used by the appliance to send packets to IP addresses that are not learned by a routing protocol or a static route. To define a default route that points to the interface named **outside,** enter the following command:

```
ASA-1# configure terminal
ASA-1(config)# route outside 0.0.0.0 0.0.0.0 10.1.37.1 1
!
```

Logging

Firewalls provide a logging feature that is useful for monitoring and troubleshooting. The logs provided should be monitor to determine if there are any attacks against the site or if there are any traffic conditions that may affect the video surveillance streams. The following basic logging commands enable logging, using the internal buffer and severity level of 4, for warning conditions. These messages will also be displayed by the Adaptive Security Device Manager.

```
!
logging enable
logging timestamp
logging buffered warnings
logging asdm warnings
!
```

Adaptive Security Device Manager (ASDM)

The ASDM provides a browser-based, Java applet to configure and monitor the PIX and ASA security appliances.

The following commands enable ASDM access to the security appliance:

```
!
asdm image disk0:/asdm521.bin
!
http server enable
http 10.1.34.0 255.255.255.0 inside
!
```

These commands specify the location of the ASDM image and enable only devices that are on the inside (protected) network and that have IP addresses in the 10.1.34.1 – 254 range to use ASDM and access the HTTP server..

To start ASDM from a web browser, simply use the IP address of the target appliance with an HTTPS connection:

https://10.1.34.1

Note: Make sure to enter **https**, not **http**

At this point you may choose to install the ASDM application or run it as a Java applet. Figure 32 shows the initial ASDM page:

Figure 32. ADSM

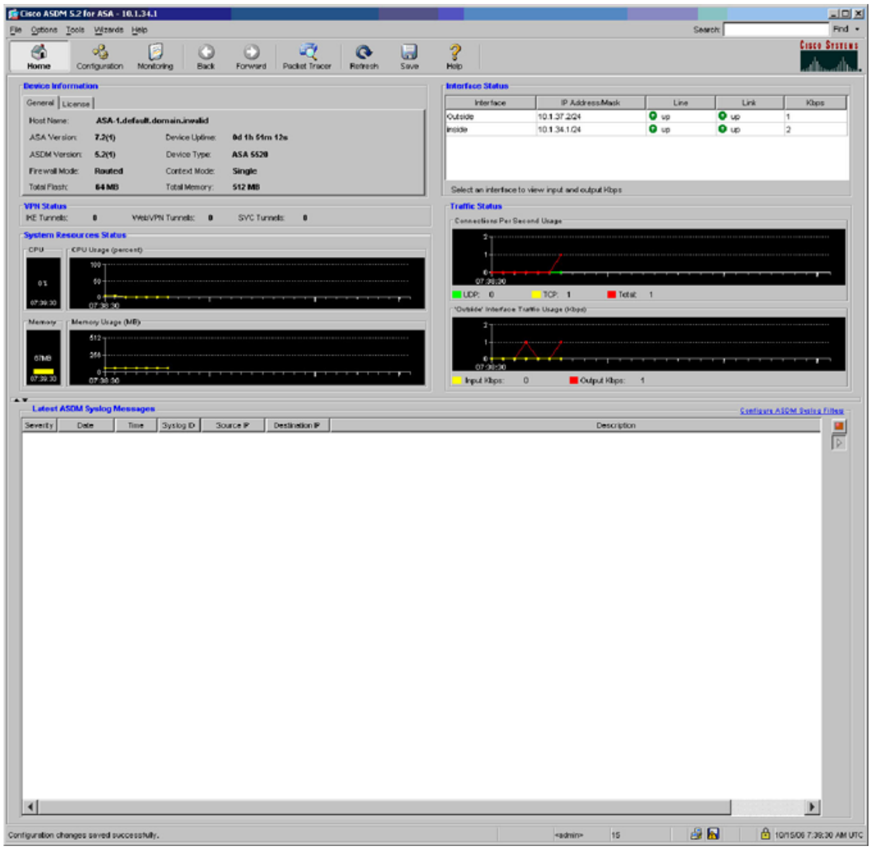

Cisco ASDM 5.2 for ASA - 10.1.34.1

File Options Tools Wizards Help
Search: Find
Home Configuration Monitoring Back Forward Packet Tracer Refresh Save Help

Device Information

General | License

Host Name: ASA-1.default.domain.invalid
ASA Version: 7.2(1) Device Uptime: 0d 1h 51m 12s
ASDM Version: 5.2(1) Device Type: ASA 5520
Firewall Mode: Routed Context Mode: Single
Total Flash: 64 MB Total Memory: 512 MB

Interface Status

Interface	IP Address/Mask	Line	Link	Kbps
Outside	10.1.37.2/24	● up	● up	1
Inside	10.1.34.1/24	● up	● up	2

Select an interface to view input and output Kbps

VPN Status
IKE Tunnels: 0 WebVPN Tunnels: 0 SVC Tunnels: 0

Traffic Status

Latest ASDM Syslog Messages

Configure ASDM Syslog Filters

Severity	Date	Time	Syslog ID	Source IP	Destination IP	Description

Configuration changes saved successfully. <admin> 15 10/15/06 7:38:30 AM UTC

From this point, the configuration can be fully accomplished using ASDM or CLI. In this document, most commands are entered in CLI and the ASDM will be used as a graphical monitoring tool.

Configuring Multicast Routing

To configure PIM routing on the security appliance, follow these steps:

1. Enable Multicast routing. The following command enables PIM and IGMP on every interface:

```
ASA-1# configure terminal
ASA-1(config)#multicast-routing
```

2. Define the Rendezvous Point.

```
ASA-1(config)#pim rp-address 10.94.162.250
```

Note: Because the ASA cannot participate in dynamic RP discovery mechanisms such as Auto-RP or BSR, the RP must be defined manually. Multiple RPs can be defined by using the acl argument of the **pim rp-address** command.

Multicast and basic device verification

Before applying access lists to restrict and allow multicast traffic between Cisco IP Gateway Encoders and Decoders, a review of the current state is recommended:

PIM and IGMP Verification:

To verify the current PIM and IGMP status:

```
ASA-1# show pim interface

Address          Interface PIM  Nbr   Hello  DR     DR
                                Count Intvl  Prior
10.1.37.2        Outside   on   1     30     1      this system
10.1.34.1        inside    on   0     30     1      this system
```

The following igmp groups have been formed on the inside interface:

```
ASA-1# show igmp groups

IGMP Connected Group MembershipGroup Address    Interface
Uptime    Expires   Last Reporter
235.1.1.1        inside         00:17:37  00:03:46  10.1.34.6
239.1.1.2        inside         00:17:36  00:03:48  10.1.34.55
239.1.1.3        inside         00:17:32  00:03:50  10.1.34.55
```

The ASA has successfully become a PIM neighbor with 10.1.37.1:

```
ASA-1# show pim neighbor

Neighbor Address  Interface    Uptime    Expires DR pri Bidir

10.1.37.1         Outside      00:20:13  00:01:44 1
```

This information verifies that the firewall has become a PIM neighbor with the 6504-2 switch, which is also acting as the RP.

To display the current multicast routing table:

```
ASA-1# show mroute

Multicast Routing Table
Flags: D - Dense, S - Sparse, B - Bidir Group, s - SSM Group,
       C - Connected, L - Local, I - Received Source Specific Host
Report,
       P - Pruned, R - RP-bit set, F - Register flag, T - SPT-bit set,
       J - Join SPT
Timers: Uptime/Expires
Interface state: Interface, State

(*, 235.1.1.1), 00:58:56/never, RP 10.94.162.250, flags: SCJ
  Incoming interface: Outside
  RPF nbr: 10.1.37.1
  Outgoing interface list:
    inside, Forward, 00:58:56/never

(10.1.34.55, 235.1.1.1), 00:00:48/00:02:41, flags: SFJT
  Incoming interface: inside
  RPF nbr: 10.1.34.55
  Outgoing interface list:
    Outside, Forward, 00:00:48/00:02:46
```

```
(10.94.162.206, 235.1.1.1), 00:00:52/00:03:07, flags: SJT
  Incoming interface: Outside
  RPF nbr: 10.1.37.1
  Immediate Outgoing interface list: Null

(10.94.162.217, 235.1.1.1), 00:00:35/00:02:54, flags: SJT
  Incoming interface: Outside
  RPF nbr: 10.1.37.1
  Immediate Outgoing interface list: Null

(10.94.162.219, 235.1.1.1), 00:00:39/00:02:50, flags: SJT
  Incoming interface: Outside
  RPF nbr: 10.1.37.1
  Immediate Outgoing interface list: Null

(10.94.162.232, 235.1.1.1), 00:00:32/00:03:27, flags: SJT
  Incoming interface: Outside
  RPF nbr: 10.1.37.1
  Immediate Outgoing interface list: Null

(10.94.162.233, 235.1.1.1), 00:00:30/00:02:59, flags: SJT
  Incoming interface: Outside
  RPF nbr: 10.1.37.1
  Immediate Outgoing interface list: Null

. . . (Output truncated)
```

Appendix D contains the full firewall configuration.

Configuring Firewall Access Rules

By default, for ASA and PIX appliances, traffic from a higher security level (for example, inside) can access an interface with a lower security level (for example, outside), but outside traffic destined for an inside network is denied. Other firewalls may have a different behavior.

Object Grouping has been available since PIX version 6.2 and has been widely used to group objects such as hosts (servers, clients), services and networks to apply policies and rules to the whole group. This capability has reduced the configuration complexity and has helped to reduce the number of access rules required.

The following Object Groups can be used to group the Cisco IP Gateways and the services that are required to allow video surveillance traffic through the ASA:

```
!
object-group network IP_GATEWAYS
 description Encoders and Decoders
 network-object host 10.1.34.5
 network-object host 10.1.34.6
 network-object host 10.1.34.55
object-group network Multicast_Control
 network-object host 235.1.1.1
 network-object host 239.1.1.2
 network-object host 239.1.1.3
```

```
object-group icmp-type Standard_ICMP
 icmp-object echo
 icmp-object echo-reply
 icmp-object time-exceeded
 icmp-object unreachable
```

To allow inbound traffic that is normally blocked by the ASA, an access list must be configured. The following commands allow traffic between IP Gateway Encoders and Decoders to flow between the inside and outside networks.

```
!
access-list Outside_ACL extended permit udp any object-group IP_GATEWAYS
access-list Outside_ACL extended permit udp any object-group
Multicast_Control
access-list Outside_ACL extended permit icmp any object-group IP_GATEWAYS
object-group Standard_ICMP
access-list Outside_ACL extended permit udp any 239.0.0.0 255.0.0.0
```

While this configuration accomplishes the goal of allowing specific traffic, a more restrictive access rule could offer additional protection:

- Instead of using **any**, a more restrictive access rule could specify only the addresses that have access to the video streams.
- Instead of using **udp any**, an access rule could only allow the specific UDP ports required. The “Segmentation using Access Lists” section in this chapter has more details on these port numbers.

To apply the Outside_ACL to the Outside interface add the following command:

```
!
access-group Outside_ACL in interface Outside
```

Verification

At this point, the Cisco Stream Manager Configuration Module should be able to display all encoders and decoders in the network.

Figure 33. Stream Manager Configuration

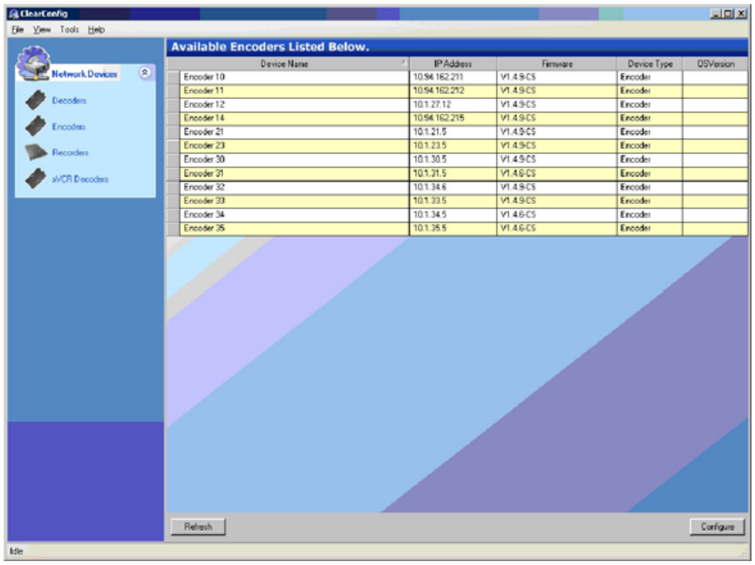

Available Encoders Listed Below.

Device Name	IP Address	Firmware	Device Type	OSVersion
Encoder 10	10.94.162.211	V1.4.9-CS	Encoder	
Encoder 11	10.94.162.212	V1.4.9-CS	Encoder	
Encoder 12	10.1.27.12	V1.4.9-CS	Encoder	
Encoder 14	10.94.162.215	V1.4.9-CS	Encoder	
Encoder 21	10.1.21.5	V1.4.9-CS	Encoder	
Encoder 23	10.1.23.5	V1.4.9-CS	Encoder	
Encoder 30	10.1.30.5	V1.4.9-CS	Encoder	
Encoder 31	10.1.31.5	V1.4.6-CS	Encoder	
Encoder 32	10.1.34.6	V1.4.9-CS	Encoder	
Encoder 33	10.1.33.5	V1.4.9-CS	Encoder	
Encoder 34	10.1.34.5	V1.4.6-CS	Encoder	
Encoder 35	10.1.35.5	V1.4.6-CS	Encoder	

The Cisco Stream Manager software can display the available cameras and their associated video streams. Video streams from cameras #32 and #34 should also be available at other locations.

IP Fragmentation

When an IP packet is too large to be transmitted by an interface, the oversized packet is split into two or more IP fragments, each small enough to be transmitted on the selected network. If a packet is fragmented, each fragment becomes its own packet, which contains its own IP header and is routed independently of other packets.

The destination host is responsible for reassembling the packets, even if the packets arrive out of order, which makes the host for the most part transparent to transport layer protocols such as TCP and UDP.

Frames of live surveillance video are typically too large for the MTU size of Ethernet and are broken up into many fragment packets before being transmitted to the network.

Some firewalls may have trouble processing IP fragments correctly. If the firewall is configured to allow non-initial fragments with insufficient information to properly match the filter, a firewall attack could occur. To properly enforce a security policy, the firewall must track all the fragments of a datagram to evaluate the validity of the traffic.

Figure 34 shows quality issues with the current Cisco Stream Manager camera streams.

Figure 34. Fragmentation Issues

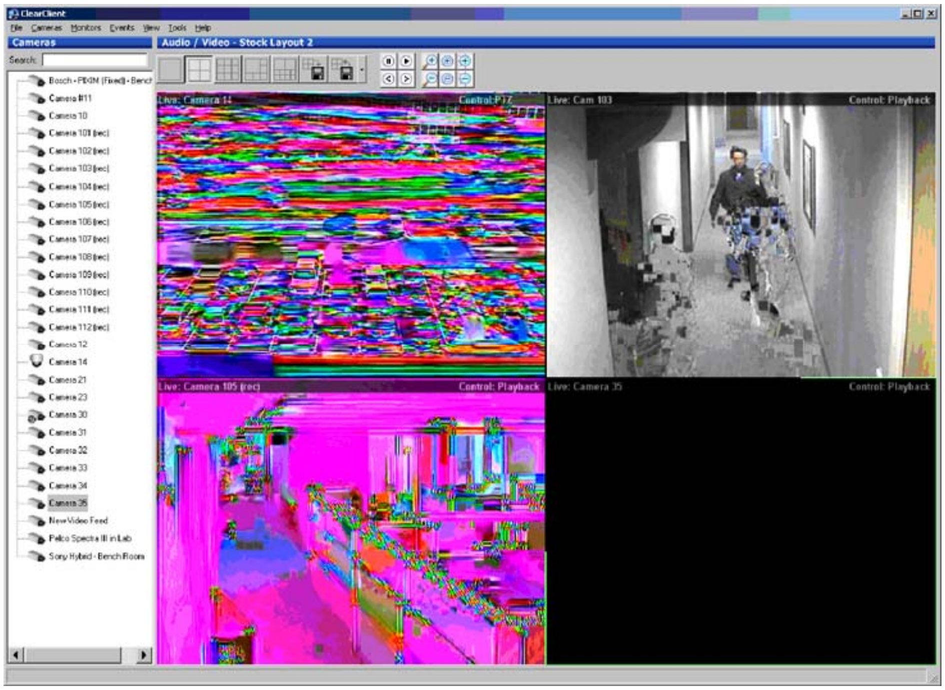

In Figure 34 the Cisco Stream Manager Client Viewing Module was not able to display one of the video streams due to fragmentation. Three other video streams are severely affected by packet fragmentation, making them unusable.

A simple way to check for packet fragmentation is to review the ASDM Syslog Messages, which appear on the Home Screen.

Figure 35. ASDM Messages

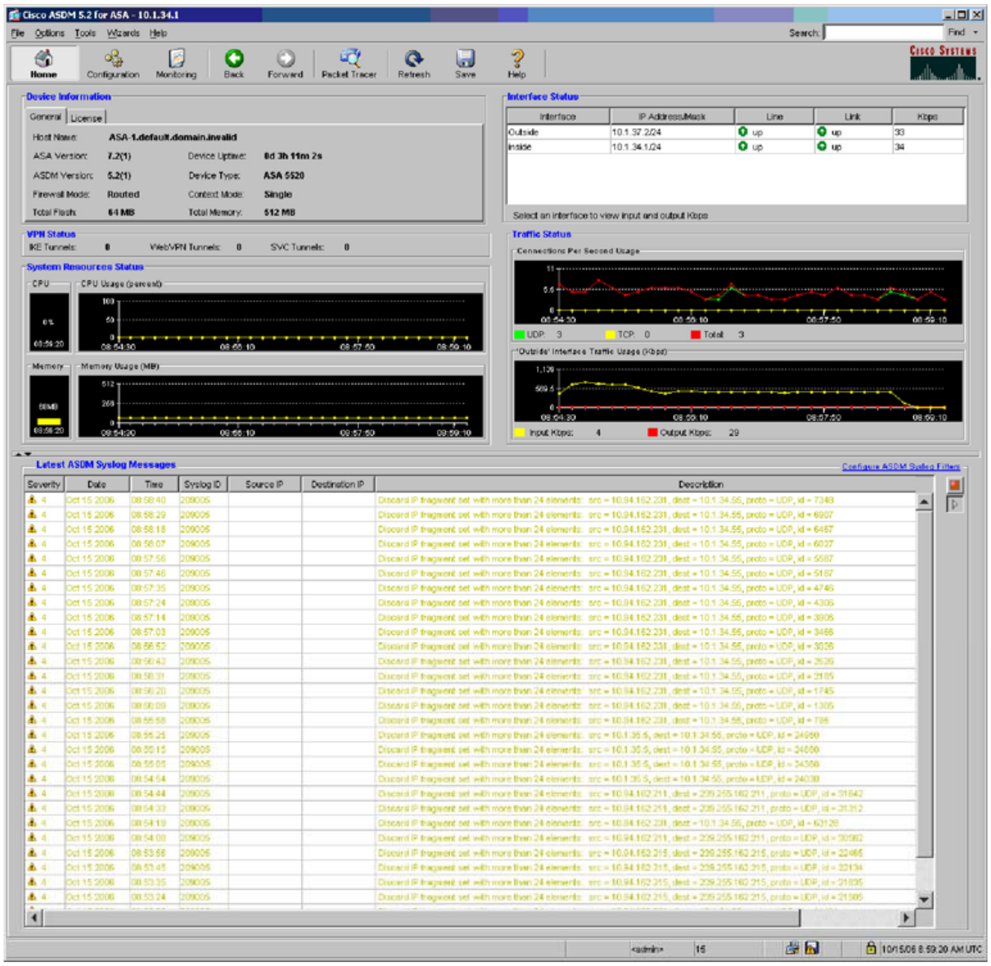
Cisco ASDM 5.2 for ASA - 10.1.34.1
File Options Tools Wizards Help
Search: Find
Home Configuration Monitoring Back Forward Packet Tracer Refresh Save Help
Cisco Systems
General License
CPU Usage (percent)
100
50
0
08:54:30 08:55:10 08:57:50 08:59:10
Memory Usage (MB)
512
256
0
08:54:20 08:55:10 08:57:50 08:59:10
Connections Per Second Usage
11
5.5
0
'Outside' Interface Traffic Usage (Kbps)
1,129
589.5
0

Device Information

Host Name: ASA-1.default.domain.invalid

ASA Version: 7.2(1) Device Uptime: 0d 3h 11m 2s

ASDM Version: 5.2(1) Device Type: ASA 5520

Firewall Mode: Routed Context Mode: Single

Total Flash: 64 MB Total Memory: 512 MB

VPN Status

IKE Tunnels: 0 WebVPN Tunnels: 0 SVC Tunnels: 0

System Resources Status

CPU: 0% 08:59:20

Memory: 60MB 08:59:20

Interface Status

Interface	IP Address/Mask	Line	Link	Kbps
Outside	10.1.37.2/24	up	up	33
inside	10.1.34.1/24	up	up	34

Select an interface to view input and output Kbps

Traffic Status

🟩 UDP: 3 🟨 TCP: 0 🟥 Total: 3

🟨 Input Kbps: 4 🟥 Output Kbps: 29

Latest ASDM Syslog Messages

Configure ASDM Syslog Filters

Severity	Date	Time	Syslog ID	Source IP	Destination IP	Description
⚠ 4	Oct 15 2006	08:58:40	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 7348
⚠ 4	Oct 15 2006	08:58:29	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 6907
⚠ 4	Oct 15 2006	08:58:18	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 6467
⚠ 4	Oct 15 2006	08:58:07	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 6027
⚠ 4	Oct 15 2006	08:57:56	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 5587
⚠ 4	Oct 15 2006	08:57:46	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 5187
⚠ 4	Oct 15 2006	08:57:35	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 4746
⚠ 4	Oct 15 2006	08:57:24	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 4306
⚠ 4	Oct 15 2006	08:57:14	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 3806
⚠ 4	Oct 15 2006	08:57:03	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 3466
⚠ 4	Oct 15 2006	08:56:52	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 3026
⚠ 4	Oct 15 2006	08:56:42	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 2626
⚠ 4	Oct 15 2006	08:56:31	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 2185
⚠ 4	Oct 15 2006	08:56:20	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 1745
⚠ 4	Oct 15 2006	08:56:09	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 1305
⚠ 4	Oct 15 2006	08:55:58	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 786
⚠ 4	Oct 15 2006	08:55:25	209005			Discard IP fragment set with more than 24 elements: src = 10.1.35.5, dest = 10.1.34.55, proto = UDP, id = 24960
⚠ 4	Oct 15 2006	08:55:15	209005			Discard IP fragment set with more than 24 elements: src = 10.1.35.5, dest = 10.1.34.55, proto = UDP, id = 24600
⚠ 4	Oct 15 2006	08:55:05	209005			Discard IP fragment set with more than 24 elements: src = 10.1.35.5, dest = 10.1.34.55, proto = UDP, id = 24360
⚠ 4	Oct 15 2006	08:54:54	209005			Discard IP fragment set with more than 24 elements: src = 10.1.35.5, dest = 10.1.34.55, proto = UDP, id = 24030
⚠ 4	Oct 15 2006	08:54:44	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.211, dest = 239.255.162.211, proto = UDP, id = 31642
⚠ 4	Oct 15 2006	08:54:33	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.211, dest = 239.255.162.211, proto = UDP, id = 31312
⚠ 4	Oct 15 2006	08:54:19	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 63128
⚠ 4	Oct 15 2006	08:54:08	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.211, dest = 239.255.162.211, proto = UDP, id = 30582
⚠ 4	Oct 15 2006	08:53:58	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.215, dest = 239.255.162.215, proto = UDP, id = 22466
⚠ 4	Oct 15 2006	08:53:45	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.215, dest = 239.255.162.215, proto = UDP, id = 22134
⚠ 4	Oct 15 2006	08:53:35	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.215, dest = 239.255.162.215, proto = UDP, id = 21835
⚠ 4	Oct 15 2006	08:53:24	209005			Discard IP fragment set with more than 24 elements: src = 10.94.162.215, dest = 239.255.162.215, proto = UDP, id = 21505

<admin> 15 10/15/06 8:59:20 AM UTC

These messages show repeated warnings that the ASA is dropping fragmented packets. These messages can be seen at the CLI console and can be sent to an external syslog server:

```
ASA-1# show logging

Syslog ID: 209005
Description: Discard IP fragment set with more than 24 elements: src=
10.94.162.231, dest = 10.1.34.55, proto = UDP, id = 7348
```

With the default configuration, the ASA considers this fragmentation a possible intrusion and discards the packet. By default, the ASA accepts up to 24 fragments to reconstruct a full IP packet.

To specify a different number of packets that may be fragmented, use the **fragment** command. The syntax of this command is:

fragment {**size** | **chain** | **timeout** *limit*} [*interface*]

The chain size number depends on the media stream settings of the Cisco IP Gateway Typically, higher resolution and frame rates translate into a higher number of fragmented packets.

By using the **fragment** command with a **chain** limit of 45, the ASA allows all video streams to pass without being blocked. This parameter should be adjusted and monitored to match the proper environment requirements.

```
!
fragment chain 45 Outside
fragment chain 45 inside
```

Figure 36. No Packet Fragmentation

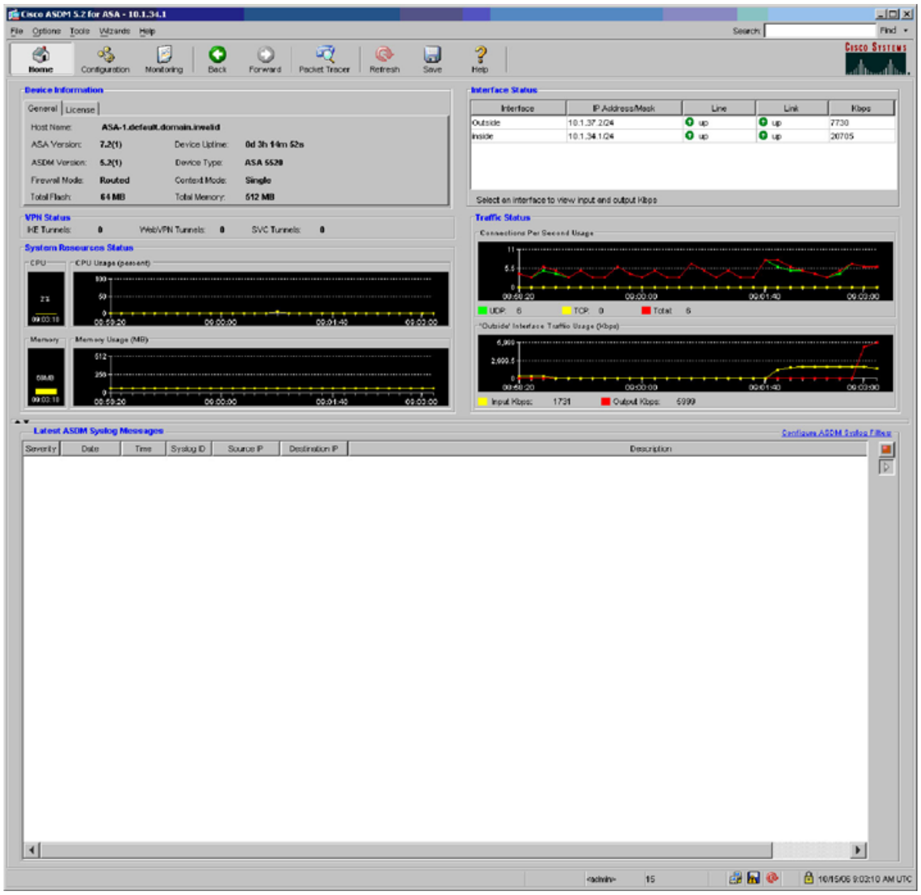

Device Information

Host Name: ASA-1.default.domain.invalid

ASA Version: 7.2(1) Device Uptime: 0d 3h 14m 52s

ASDM Version: 5.2(1) Device Type: ASA 5520

Firewall Mode: Routed Context Mode: Single

Total Flash: 64 MB Total Memory: 512 MB

Interface Status

Interface	IP Address/Mask	Line	Link	Kbps
Outside	10.1.37.2/24	● up	● up	7730
inside	10.1.34.1/24	● up	● up	20705

Select an interface to view input and output Kbps

VPN Status

IKE Tunnels: 0 WebVPN Tunnels: 0 SVC Tunnels: 0

System Resources Status

CPU

Memory

Traffic Status

🟩 UDP: 6 🟨 TCP: 0 🟥 Total: 6

🟨 Input Kbps: 1731 🟥 Output Kbps: 5999

Latest ASDM Syslog Messages

Severity	Date	Time	Syslog ID	Source IP	Destination IP	Description

Figure 37. Stream Manager Client Viewing Module

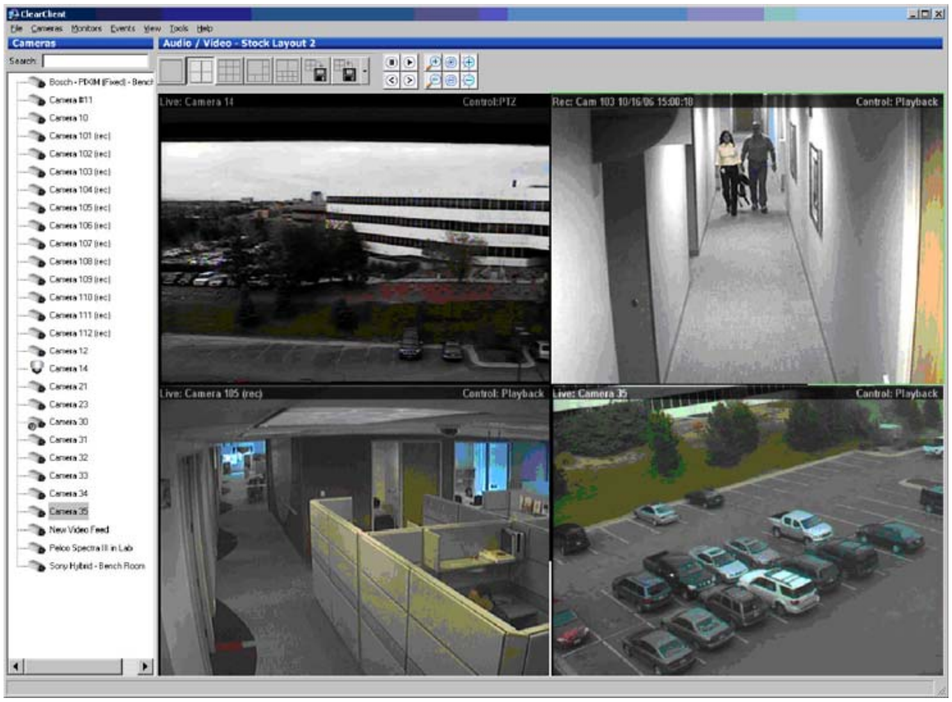

Figure 37 shows the video streams after raising the fragment parameters.

When the same video stream is displayed by more than one Cisco IP Gateway Decoder, the stream switches automatically to multicast. In our example, the IP addresses of the encoders are 10.1.34.5 and 10.1.34.6, and the following multicast groups have been created for these encoders:

- 239.255.34.5
- 239.255.34.6

These multicast groups are formed by adding the last two octets of the encoder's IP address to 239.255.x.x.

```
ASA-1# show igmp groups
IGMP Connected Group MembershipGroup Address    Interface
Uptime    Expires   Last Reporter
235.1.1.1     inside         01:18:14  00:03:41  10.1.34.5
239.1.1.2     inside         01:18:13  00:03:38  10.1.34.55
239.1.1.3     inside         01:18:09  00:03:40  10.1.34.55
239.255.34.5  inside         00:00:19  00:04:00  10.1.34.55
239.255.34.6  inside         00:00:19  00:04:00  10.1.34.55
```

More Restrictive Firewall Access List

Access lists can be used in different ways, with different restriction levels. A more restrictive Access list can be configured that only allows traffic between specific Cisco IP Gateway Encoders and Decoders.

The following access list provides a more granular control of the traffic parameters, but it can be harder to manage in a large environment. The following object groups can be configured to specify the Cisco IP Gateways and recorders on the network. While these object groups need to be maintained as the number of devices grows, this configuration adds some level of protection by allowing video traffic to be viewed only by the approved devices.

```
!
object-group network Multicast_Control
 network-object host 235.1.1.1
 network-object host 239.1.1.2
 network-object host 239.1.1.3
object-group icmp-type Standard_ICMP
 icmp-object echo
 icmp-object echo-reply
 icmp-object time-exceeded
 icmp-object unreachable
object-group network Inside_IPGateways
 network-object host 10.1.34.5
 network-object host 10.1.34.6
 network-object host 10.1.34.55
object-group network Outside_Recorders
 network-object host 10.94.162.231
object-group network Outside_IPGateways
 network-object host 10.94.162.211
 network-object host 10.94.162.212
 network-object host 10.94.162.215
 network-object host 10.1.27.12
 network-object host 10.1.21.5
 network-object host 10.1.23.5
 network-object host 10.1.30.5
 network-object host 10.94.162.206
 network-object host 10.94.162.217
 network-object host 10.94.162.232
 network-object host 10.94.162.233
 network-object host 10.94.162.219
 network-object host 10.1.33.5
 network-object host 10.1.31.5
 network-object host 10.1.35.5
 network-object host 10.94.162.207
 network-object host 10.94.162.234
 group-object Outside_Recorders
```

Note that the Outside_Recorders group is nested within the Outside_IPGateways object group. This configuration allows all members of the Outside_Recorders group to be included in the Outside_IPGateways group.

The access list to be applied to the Outside interface could include:

```
!
access-list OUTSIDE extended permit udp object-group Outside_IPGateways
object-group Inside_IPGateways
access-list OUTSIDE extended permit icmp object-group Outside_IPGateways
object-group Inside_IPGateways object-group Standard_ICMP
access-list OUTSIDE extended permit udp object-group Outside_IPGateways
object-group Multicast_Control
```

```
access-list OUTSIDE extended permit udp object-group Outside_Recorders
239.0.0.0 255.0.0.0
access-list OUTSIDE extended permit udp object-group Outside_IPGateways
239.255.0.0 255.255.0.0
!
!
access-group OUTSIDE in interface Outside
```

With this configuration, only authorized devices from the outside can view video from the inside. We can also add an ACL to the inside interface so that only authorized devices from the inside can view video from the outside.

Chapter 6: Traffic Engineering

Video Surveillance

Cisco Video Surveillance products use the MPEG-4 video compression codec to convert analog video streams to digital format for live viewing and disk storage. Additional details about the MPEG-4 codec may be found in the MPEG4 and Video Codecs section of this document.

When live Cisco Video Surveillance streams are transported across the IP network, they have specific characteristics. Each frame of live analog video that is processed by a Cisco Video Surveillance Gateway can result in multiple IP packets across the network. A single video I-frame may be larger than 60 KB. Such frames are too large for the Maximum Transmission Unit (MTU) size of the Ethernet interface on the encoder, which is 1518 Bytes. These large video frames are fragmented into smaller Ethernet-framed IP packets to stay smaller than the Ethernet MTU. This process results in potentially dozens of packets representing a single frame of video. The packet fragments of each frame are followed up by a User Datagram Protocol (UDP) header packet, which contains the Real-Time Transport Protocol (RTP) sequencing information for the entire video frame. Depending on the frame rate, this process occurs from 1 to 30 times per second during the video stream.

Because of the varying types of images, the amount of motion, and the natural processing of a video compression codec, live video streams are considered “bursty” in nature by IP networking standards. The streams are also sensitive to delay, loss, and jitter in the network. For these reasons, proper bandwidth provisioning and priority queuing in the Quality of Service (QoS) configuration are critical to a successful implementation.

Recorded Video Surveillance

When replaying a recorded video stream across the IP network, the traffic characteristics are significantly different from those of a live stream. Recorded video traffic is transferred to the hardware or software decoder using the TCP protocol, which guarantees delivery. Latency or delay in the network is less critical than a live stream because the decoder builds a buffer of several seconds of video for replay. This buffer is continuously filled with video information that is carried via TCP packets from the recorder.

Because of the nature of this stream, a different QoS approach is recommended for recorded streams. This approach uses Class Based Weighted Fair Queuing and Traffic Shaping features to protect and smooth the traffic across a Wide Area Network link. The average bandwidth consumed by live and recorded streams remains similar over time.

Audio Surveillance

Cisco Video Surveillance IP Gateways can transmit unidirectional audio from an encoder to a decoder in the same direction as the video surveillance stream. The input and output on the analog audio side is standard line level stereo, such as that on PC sound cards and home audio equipment. The audio is transmitted across the IP network in raw Pulse Code Modulation (PCM,) using 16-bit samples and 20 milliseconds (ms) of audio per RTP packet. The sampling rate used for the encoding is configurable from 8 kHz to 48 kHz. Higher sampling rates result in both higher quality audio and higher bandwidth consumption.

Because there is no audio compression algorithm and no silence suppression in use, the audio streams consume a steady amount of bandwidth. The 20 ms sample size produces a steady rate of 50 samples per second. At lower sample rates, the audio data may fit into a single IP packet for each sample. At higher sample rates, fragmentation is used to divide the audio sample across multiple IP packets.

Audio traffic over IP is sensitive to packet loss, delay and jitter on the network. Proper bandwidth provisioning and QoS are critical to a successful implementation. Care should be taken not to overlook audio bandwidth and storage requirements during planning for video surveillance needs.

Recorded Audio Surveillance

Replay of recorded audio surveillance is delivered across the network by using the TCP protocol. This traffic should be placed into the same class-based queue as recorded video surveillance. The bandwidth requirements for the recorded audio are similar to those of the live stream that is being replayed. Note that recorded audio is only available when using direct analog audio feeds into the audio-enabled Services Platform. Audio recording is not yet implemented for streams from the IP network.

Planning Traffic Flows in a Surveillance Network

Traffic engineering for surveillance media IP traffic requires insight into the physical and logical topology of the underlying network. You must take into consideration what paths media streams will take with a fully operational network and what alternate paths Spanning Tree and routing protocols will cause traffic to take in the case of various link outages. Networking features such as equal cost multi-path routing may allow traffic to load balance on a fully operational network, but you should always consider the potential load on a single link that can be caused by something as simple as an interface failure or a bad cable.

Cisco Video Surveillance provides a powerful distributed architecture to deliver media streams anywhere across an IP network. However, this architecture is not topology-aware, and it contains no connection admission control (CAC) mechanisms between endpoints. Thoughtful placement of sources and destinations of the media flows is critical to ensuring the overall performance of the system. Queue capacity that has been set aside to carry a specific number of live or recorded media streams can be overrun if additional, unplanned streams are passed over the same path. The result does not only affect the unplanned streams. It can degrade performance for all streams in the queue.

Such performance degradation must be considered in a WAN environment that uses a priority queue for live video, which polices (discards) nonconforming (excess) traffic. Policing occurs across all flows in the queue, not just across the unplanned flows. Therefore, the design of the surveillance network must use application layer and network layer controls to ensure that only the intended number of flows is presented to the queue.

When planning a surveillance network, consider both the static media streams and the dynamic media streams. Many streams will be static. For example, 20 cameras in a facility that all stream to a local recorder or 4 critical cameras that are always watched by the security staff at an operations center. These streams are commonly present 24 hour a day, 7 days a week.

The dynamic streams are more challenging to plan for, and create the most risk of saturating limited bandwidth. Dynamic streams can be created in a number of ways. For example, an operator with a CCTV monitor and keyboard that are attached to a decoder can cause dynamic streams by frequently switching the monitor to various cameras across the network, some of which could be located over low-bandwidth WAN connections. Another example are the dynamic streams caused by an operator with the Cisco Stream Manager Client Viewing Module, which can be configured to view up to ten concurrent live or recorded video streams.

One approach to controlling the amount of streams that traverse a given path is to configure limitations on the flow of multicast traffic within the network where applicable.

Sources of Media Flows

Sources of media traffic to consider on a surveillance network include more than just the UDP streams of live video encoders. Recorders also source streams of TCP traffic when video is being reviewed. Streams of recorded video may be created by a planned session when a team is investigating a past incident or performing a routine review of captured video. It also could also be an ad-hoc session such as when a decoder or the Cisco Stream Manager Client Viewing Module requests video for an “instant replay” and scrolls back in time to pull video from a recorder. Decoders also produce PTZ control traffic that is considered part of the media stream because it is key to the interactive experience of the operator who is controlling the PTZ camera. Decoders also produce traffic which controls rewind, fast forward, play and pause on video content that is requested from a recorder. Any third-party products that are integrated with the overall physical security system must also be considered, such as third-party IP-based video surveillance cameras.

Destinations of Media Flows

Destinations of media traffic to consider on a surveillance network include more than just hardware-based decoders. The recording systems are typically the destination for the largest number of flows in the network, as it is common to record all surveillance streams for future review. The Cisco Stream Manager application is another potential destination. Any third-party products that are integrated with the overall physical security system, such as video analytics applications, must also be considered.

The Cisco Stream Manager Client Viewing Module provides a software-based decoder, which offers highly flexible configuration options and which can be placed anywhere on the network. This flexibility results in additional considerations for bandwidth planning. The most challenging scenario for which to engineer is the Cisco Stream Manager Client Viewing Module installed on a laptop computer that can rove anywhere in a WAN environment with a wide-open multicast policy. This device could call up to 10 concurrent streams to any location in the network using a standard screen layout. You must either control the availability and use of such a device administratively through company policy or limit its ability to join certain multicast groups through router and switch configuration. See the “Segmentation – VLANs and ACLs” section in the “Security” chapter for more detailed information.

Unicast and Multicast Bandwidth Considerations

Cisco Video Surveillance gateways can dynamically transition between unicast and multicast streaming modes for live streams, based on the number of endpoints that request a given stream. Live streams being viewed are often multicast even when viewed by a single source, because a recorder frequently is subscribing to the same stream. Single-port Cisco Video Surveillance gateways also have the capability to dual-stream the media at two different resolutions and frame rates. This capability allows customers to use a high-quality stream for live viewing while using a lower bit-rate stream for recording to conserve disk space.

A decoder requests a live media stream from an encoder when the camera ID (also known as peripheral ID) is entered on a traditional CCTV keyboard. When a PC that is running the Cisco Stream Manager Client Viewing Module requests a media stream from an encoder, the Stream Description field that is configured on the encoder is displayed. When using dual-streaming, it is important to ensure that only the streams intended for live viewing are chosen from CCTV keyboards and Cisco Stream Manager PCs. If both streams from an encoder are triggered into multicast mode by multiple requests, the link capacity in a WAN environment can be flooded.

Figure 38. Camera Selection in Stream Manager Client Viewing Module

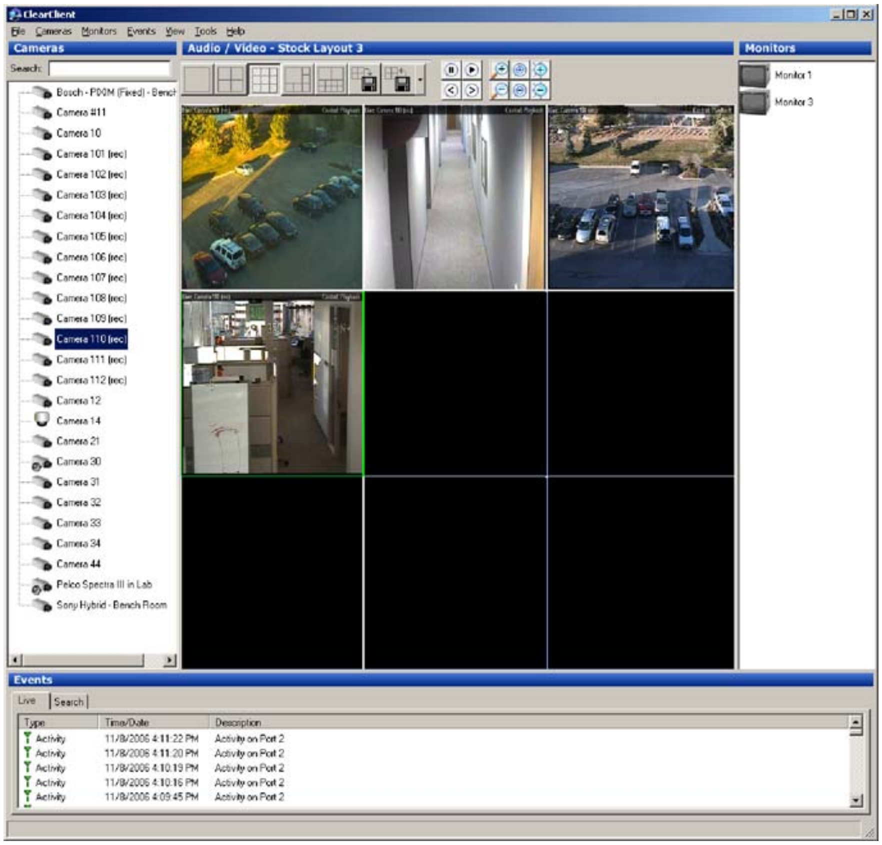
ClearClient
File Cameras Monitors Events View Tools Help
Cameras
Search:
Audio / Video - Stock Layout 3
Monitors
Monitor 1
Monitor 3
Events
Live
Search

- Bosch - P00M (Fixed) - Bench
- Camera #11
- Camera 10
- Camera 101 (rec)
- Camera 102 (rec)
- Camera 103 (rec)
- Camera 104 (rec)
- Camera 105 (rec)
- Camera 106 (rec)
- Camera 107 (rec)
- Camera 108 (rec)
- Camera 109 (rec)
- Camera 110 (rec)
- Camera 111 (rec)
- Camera 112 (rec)
- Camera 12
- Camera 14
- Camera 21
- Camera 30
- Camera 31
- Camera 32
- Camera 33
- Camera 34
- Camera 44
- Pelco Spectra III in Lab
- Sony Hybrid - Bench Room

Type	Time/Date	Description
Activity	11/8/2006 4:11:22 PM	Activity on Port 2
Activity	11/8/2006 4:11:20 PM	Activity on Port 2
Activity	11/8/2006 4:10:19 PM	Activity on Port 2
Activity	11/8/2006 4:10:16 PM	Activity on Port 2
Activity	11/8/2006 4:09:45 PM	Activity on Port 2

Another multicast-related bandwidth consideration is the use of live audio streaming from encoders. When audio streaming is configured, it is streamed to the same unicast or multicast address as the video media. Differentiation is achieved between streams by different User Datagram Protocol (UDP) source and destination port numbers. A hardware decoder or Cisco Stream Manager Client Viewing Module that requests the video stream also receives the audio stream, regardless of whether the destination is equipped to support replay of the audio stream. Always ensure that audio bandwidth is included in bandwidth estimates for network circuit provisioning.

Bandwidth Provisioning for Surveillance

Video Bandwidth Consumption Guidelines

The Cisco Stream Manager Configuration Module allows the configuration of the Bit Rate Model on a Cisco Video Surveillance Gateway encoder to Variable or Constant. These settings directly control the MPEG-4 CODEC behavior on the encoder. To achieve video compression, one of the MPEG-4 CODEC processes uses a mathematical process called *quantization* to translate arrays of video information into mathematical approximations that require less bandwidth to transmit. These approximations allow the arrays of data to be reproduced on the decoder side, with some level of loss. The more aggressive the quantization scale, the more loss is induced, which can affect video quality.

The “Constant” Bit Rate Model (CBR) allows you to set a target bit rate for the encoder. The result over time is a consistent level of bandwidth consumption for the provisioning of network bandwidth and storage capacity. For this reason, the Constant Bit Rate model is the preferred approach, especially when implementing video surveillance on a Wide Area Network. The CBR model is used as a reference for all traffic engineering and storage capacity planning examples in this document.

> Note: The CBR bit rate is a target and does not result in an absolutely flat level of bandwidth consumption. Also, “CBR” in this context is not the same as it is in the context of other networking technologies such as Asynchronous Transfer Mode (ATM.)

When using CBR, the MPEG-4 CODEC manipulates the quantization scale on a frame-by-frame basis to attempt to meet the target bit rate. The result is a significant leveling of the bursty nature of video bandwidth consumption, although the peak values during times of heavy motion in the video stream may still exceed the target bit rate. The downside of CBR is also that when there is little or no motion or change in the source video stream, the CODEC still produces IP packets that consume the target bit rate. Figure 39 shows Constant Bit Rate Model configuration in Cisco Stream Manager Configuration Module.

Figure 39. Constant Bit Rate Model in Cisco Stream Manager

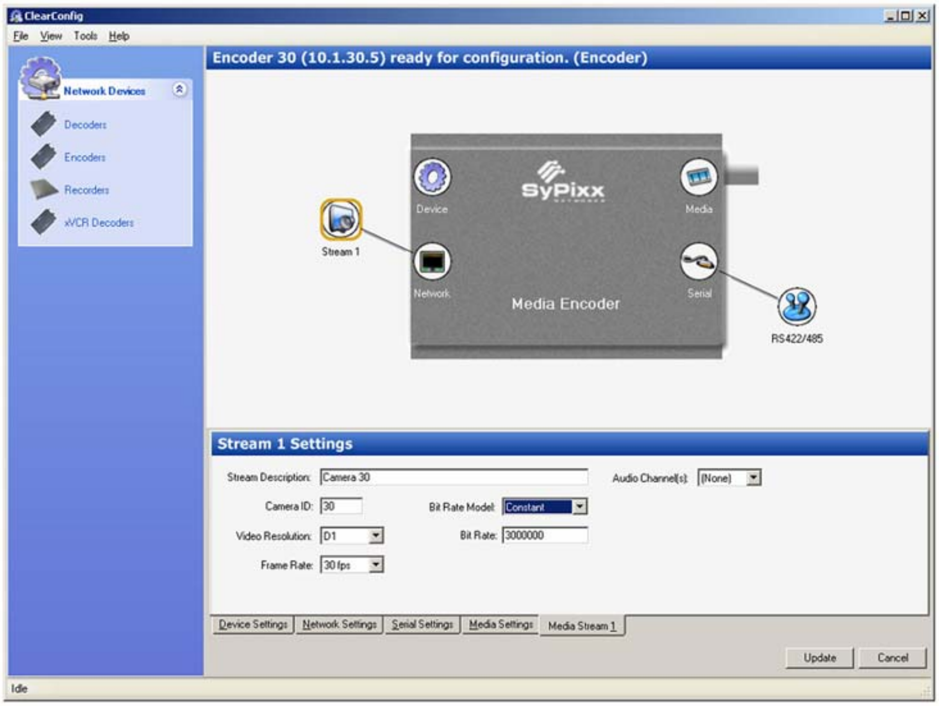

Encoder 30 (10.1.30.5) ready for configuration. (Encoder)

Stream 1 Settings

Stream Description: Camera 30

Audio Channel(s): (None)

Camera ID: 30

Bit Rate Model: Constant

Video Resolution: D1

Bit Rate: 3000000

Frame Rate: 30 fps

Idle

When choosing a bit rate for CBR, enter the number of bits per second using the Cisco Stream Manager Configuration Module. As frame rates and resolutions decrease, so does the amount of bandwidth required to produce good quality video output. Table 5 provides values that are good starting points for setting the target rate based on frame rate and resolution. The target rate may be set as desired during configuration, but be aware that too low a target rate setting may adversely affect quality.

Table 5. CBR Rate Guidelines

CBR Rate Guidelines by Resolution and Frame Rate			
NTSC	**Resolution**		
Frame Rate	**CIF**	**2CIF**	**4CIF or D1**
1.5	155000	230000	450000
2	200000	315000	600000
3	260000	410000	770000
3.75	300000	475000	935000
5	330000	525000	1050000
7.5	400000	750000	1400000
10	530000	900000	1700000
15	600000	1100000	2200000
30	850000	1600000	3000000
PAL	**Resolution**		
Frame Rate	**CIF**	**2CIF**	**4CIF or D1**
1	180000	270000	540000
2.5	330000	525000	1050000
5	425000	700000	1400000
6.25	475000	800000	1600000
12.5	700000	1200000	2400000
25	1100000	1800000	3400000
Values in Bits per Second			

When using CBR, the MPEG-4 CODEC outputs video information at an average of approximately the target rate over time. To estimate IP network bandwidth consumption for network provisioning, adding 10% to the target rate provides a reasonable basis for the stream’s real consumption.

The Variable Bit Rate Model (VBR) allows you to designate a single level of quantization on the encoder. The default level of 85% represents a moderate level of quantization that gains some valuable compression, but in most cases it produces video in that looks like the original to the human eye. As the quality percentage is set lower on the encoder, bandwidth savings becomes greater, but there is more of a chance that video quality will be affected on the decoder side. When using Variable Bit Rate, the default Quality Percentage of 85% provides a good starting point for a compromise between bandwidth and quality.

On an IP network with limited bandwidth, as in most WAN links, the use of Variable Bit Rate Model would make accurate bandwidth provisioning more difficult. Certain video sources may cause significant variability in the IP network bandwidth consumption. For this reason, VBR also makes planning of disk capacity for storage more challenging. VBR does, however, ensure a more consistent quality level from frame to frame because the quantization level is held constant.

Figure 40 shows the configuration of Variable Bit Rate Model in Cisco Stream Manager Configuration Module.

Figure 40. Variable Bit Rate Model in Cisco Stream Manager

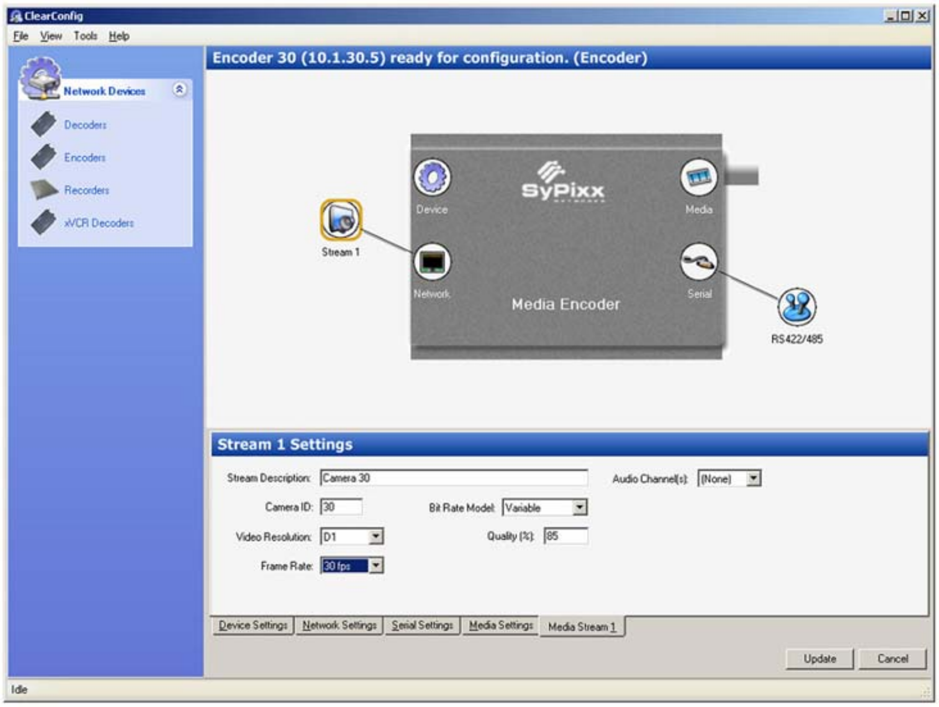

Encoder 30 (10.1.30.5) ready for configuration. (Encoder)

Stream 1 Settings

Stream Description: Camera 30

Audio Channel(s): (None)

Camera ID: 30

Bit Rate Model: Variable

Video Resolution: D1

Quality (%): 85

Frame Rate: 30 fps

Decoder Video Buffering

When configuring a decoder using the Cisco Stream Manager Configuration Module, under the Media Settings tab there is a sliding control for Video Buffering. This option controls the buffer space that is available on the decoder for storage of video frames when some source of jitter affects the pace of incoming frames. This buffer is not the network-oriented packet buffer, but a frame buffer of the sequence of video frames.

The Low Latency setting keeps the buffer size small to ensure that response time is crisp for PTZ controlled cameras. The Smooth Video setting does the opposite, maximizing the size of the video frame buffer to 15 frames. This buffer is not automatically forcefully filled by the decoder; the control only defines how many frames of video the decoder can buffer before drops occur. If more frames reside in the buffer than the level set with this option, the decoder systematically drops frames to reduce the frames stored in the buffer. For PTZ environments it is recommended that the default buffering be kept to the value of two frames to ensure that latency does not affect usability of PTZ-controlled cameras. For fixed camera environments, this value can be increased, but note that the buffer fills only when forced to do so by variability in network conditions. Figure 41 illustrates the Video Buffering configuration in Cisco Stream Manager Configuration Module.

Figure 41. Decoder Video Buffering in Cisco Stream Manager

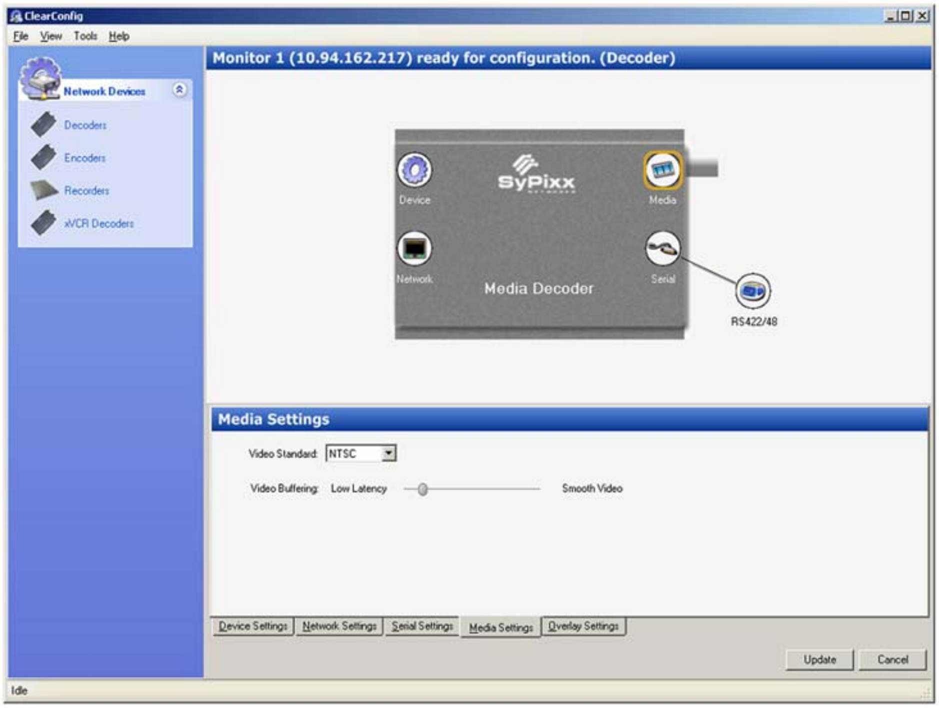

Monitor 1 (10.94.162.217) ready for configuration. (Decoder)

Media Settings

Video Standard: NTSC

Video Buffering: Low Latency — Smooth Video

Idle

Audio Bandwidth Consumption Guidelines

Live audio surveillance flows that are generated by Cisco Video Surveillance IP Gateways are composed of uncompressed raw PCM samples, and as such have smooth and consistent bandwidth consumption. Equal sized individual packets or groups of packets (depending on the sampling rate and mono or stereo configuration) are transmitted from the encoder at 50 samples per second, based on the 20 ms sample size. Table 6 shows the bandwidth consumption data in Kbps for audio streams that are generated by Cisco Video Surveillance gateways.

Table 6. Audio Bandwidth Consumption

Audio Bandwidth Consumption, Steady per Stream		
	Channels	
Sample Rate kHz	**Single**	**Stereo**
8	157.6	285.6
16	285.6	541.6
32	541.6	1067.2
48	811.2	1592.8
Cisco IP VS Gateway encoder raw PCM 16-bit Values in Kbps		

Bandwidth Provisioning Requirements

As described in the “Characteristics of Surveillance Traffic” section, video streams that are generated across the IP network tend to be bursty in nature. Even when using CBR for video surveillance, the actual bit rate of live streams on the IP network can fluctuate above the target rate. The peaks in bandwidth typically are generated by a field of view with a complex image that is not easily compressed, or by a great deal of motion or change. Table 9 provides an estimate of peak bandwidth consumption for various frame rate and resolution combinations. These numbers are accurate only when encoders are configured using the CBR target rate guidelines in Table 5.

When provisioning Wide Area Network links for video surveillance traffic, adequate bandwidth must be available to forward traffic from all video streams at their concurrent peak potential bandwidth consumption. While it may seem unlikely that there will be intensive motion in the field of view of all cameras in a facility at the same time, provisioning for that possibility is the only way to guarantee the integrity of all video streams across the network. The nature of an IP network that runs video surveillance traffic may be *converged*, which indicates a shared network running video, voice and data applications, or *dedicated*, in which video surveillance traffic and its associated control traffic are the only application in use.

Converged Network

The Cisco best practice for managing latency sensitive traffic from multimedia applications, such as video and voice over IP, is that such traffic should consume only 33% of the total bandwidth of a given link. This recommendation has been established through testing by Cisco Enterprise Systems Engineering (ESE) and is documented in detail in the *Enterprise QoS Solution Reference Network Design Guide,* which is available at http://www.cisco.com/go/srnd. One of the foundations of this recommendation states that the variability of the latency sensitive traffic coupled with its prioritization could cause the available bandwidth for other applications to fluctuate, which causes adverse affects in perceived end-user application performance.

Based on the 33% guideline, one must consider what real-time multimedia traffic types must be placed in one or more priority queues for a converged WAN. This traffic could include live video surveillance, IP telephony, video conferencing, or other video applications. Add up the peak bandwidth requirements for all of these traffic types to determine the total size required for the priority queues. Then consider data applications that will use the remaining portion of the link that is not consumed by traffic in the priority queues. The total bandwidth that is provisioned for the shared link should be at least three times the size of the priority queue itself. In other words, the total of the priority queued traffic should not exceed 33% of the total link bandwidth.

Another guideline is that voice, video, and data requirements added together should be targeting at most 75% of the physical link that is provisioned. This configuration provides headroom for *control plane* traffic on the network, such as IP routing protocols, and provides room for bursty data applications. For additional details about bandwidth provisioning, refer to the *Quality of Service Solutions Reference Network Design Guide.* As noted in that guide, the 33% guideline is not an inflexible rule for network design, but it is a conservative baseline starting point that has been shown to be successful on many converged Cisco customer networks.

Dedicated Network

In some cases, Cisco customers may choose to implement Cisco Video Surveillance solutions on an IP network that has been provisioned solely for that purpose. Such a network may be considered dedicated to this purpose if no other data, voice, or video applications are contending for bandwidth on lower speed links. For such a network, more flexibility can be exercised regarding the percentage of the link that is allocated to real-time multimedia traffic.

In this case, the challenge is not contention from voice or data applications, but video surveillance streams contending with other video surveillance streams for the shared bandwidth. Clearly, we cannot expect to provision a link to carry 100% bandwidth utilization filled by IP-based video surveillance traffic. We need to provide headroom on the link for control plane traffic, and for handling contention between several video streams, all of which produce large IP packets. We can expect to exceed the 33% guideline significantly, and even approach 75% of the link bandwidth with success. In planning such a network, it is critical to use the peak expected bandwidth for each aggregated stream instead of the average bandwidth when calculating the total bandwidth requirement. This approach allows for additional headroom on the link because odds are usually low that all of the video streams reach their peaks concurrently.

Quality of Service

The Cisco Video Surveillance IP Gateway Encoders, Decoders, and Cisco Video Surveillance Services Platform place streams of IP packet traffic onto the network infrastructure. In IP networks where multiple types of traffic such as data, voice, or video conferencing traffic may be sharing the same network, there is a potential for network congestion to affect the delivery of video surveillance packets to their destination. Quality of Service (QoS) within the IP network allows a system administrator to protect video surveillance traffic from other types of traffic on the network.

Goals of QoS for Surveillance Traffic

Congestion can occur on an IP network anywhere that there is more traffic attempting to traverse a given link than the speed of that link can support. In addition, QoS is required to protect critical traffic anywhere that there is a link speed transition within a network, such as Gigabit Ethernet (1000 Mbps) to 100 Mbps Ethernet, or a transition to a low-speed WAN link. Otherwise, packet loss, delay, and jitter can occur in the IP path, and result in poor quality video or a complete loss of video at the decoder end.

> Note: Other IP applications commonly in use on enterprise networks, such as e-mail and web browsing, are more tolerant of these effects in the network; the user response time expectations are not as tight and the applications can retransmit lost packets using the TCP protocol

Loss: Loss is a relative measure of the number of packets lost compared to the number of packets transmitted, typically expressed as a percentage. Packet Loss can occur if there more packets trying to traverse a given link than the link speed supports. An outbound network interface can buffer a certain amount of traffic, but under heavy load, some packets will ultimately be discarded. QoS tools allow the prioritization of certain types of IP packets, controlling which packets can be discarded, and which cannot. Video decoders can tolerate some packet loss and still produce a stream of live video output, but the quality of the displayed images degrades rapidly as more of the required packets are dropped. With the current Cisco Video Surveillance CODEC implementation (Cisco Video Surveillance IP Gateway Encoder software 1.7.1 and previous), loss of a single packet in a live stream will result in the decoder abandoning all packets that belong to the same video frame. Packet loss of 1% could translate to a much higher percentage of video information lost to the decoder. This issue is especially a concern if the dropped frame is an I-frame, because it affects the display of all subsequent frames that reference it until the next I-frame is received. Because surveillance video may required to reconstruct critical information in a physical security application, there is no amount of packet loss that can arbitrarily be deemed acceptable for this application. Loss of a single packet that results in loss of an I-frame at a critical time could mean that authorities are unable to see something like a license plate number or other critical information at a key split-second window of time.

Delay: Network propagation delay is the amount of time that it takes for IP packets to traverse the network from the sender to the receiver. In the case of video surveillance traffic, the estimate of total end-to-end delay should take into account the time that is required to encode, transmit, buffer, and decode the video to be displayed at the receiving end. This delay is especially critical in a live Pan-Tilt-Zoom (PTZ) environment where the commands from the operator joystick or keyboard must result in a rapid corresponding change in the displayed image, so that interaction with the system is acceptable. As a guideline, the network propagation delay target for a PTZ environment should be less than 50 milliseconds. The total end-to-end delay including codec processing in a PTZ environment should be less than 500 ms. Network delay is less critical in a fixed-camera environment. If the delay is constant, the gateways can establish sessions and stream video over networks with several hundred milliseconds of pure network delay (not including codec and receive buffer delay).

Jitter: Also commonly referred to as delay variation, jitter measures the difference in end-to-end delay for IP packets within a given video surveillance stream. For example, if one packet takes 50 ms to traverse the network and the next packet takes 80 ms, the jitter value is 30 ms between them. Decoders have some limited de-jitter buffering capability designed to smooth the incoming IP packet traffic for decoding and playback. In networks with excessive jitter, these packets can overrun or under run the available buffer space, which affects video playback quality. Jitter can cause additional frames to be stored in the video buffer on the decoder, which injects additional end-to-end delay into the operator's response time in PTZ camera environments. Cisco recommends that the maximum jitter in a network carrying surveillance media be less than 10 ms.

> Note: One additional effect that networking can have on IP traffic is packets arriving at their destination out-of-order. This can be generated by networks with multiple paths to the same destination. Cisco IP Surveillance Gateways can perform some RTP packet reordering for live video streams at the decoder. However, the overall jitter between the delayed packets must remain within the jitter guidelines of 10 ms. For live audio streams, packet-reordering is not supported. When audio packets are received out of sequence, quality is affected because samples in the late packets are dropped.

Classification and Marking of Surveillance Traffic

A QoS implementation for surveillance traffic requires classification, marking, queuing, and scheduling capabilities. Classification entails identifying packets that belong to an IP stream, ideally as close to the network edge as possible. Marking these packets allows them to be easily identified as they pass through additional internetworking devices (routers and switches). Given the different characteristics of live and recorded surveillance streams, the two traffic types should carry different markings to ensure appropriate treatment across the network. The preferred methods for marking packets for QoS treatment are Class of Service (CoS) marking at Layer 2 of the OSI model (see Figure 42) and DiffServ Code Point (DSCP) marking at layer 3 of the OSI model (see Figure 43.)

> Note: CoS uses the 802.1p User Priority bits within the 802.1Q header. 802.1Q headers are typically found only on Ethernet frames that are traversing a VLAN trunk. However, CoS may also be assigned to packets at the ingress port before they traverse the backplane of a Cisco Ethernet switch. If CoS is to be used as the trusted marking for traffic flowing between switches, the link must be configured as an 802.1Q VLAN trunk.

Figure 42. Class of Service in the 802.1Q/p Header

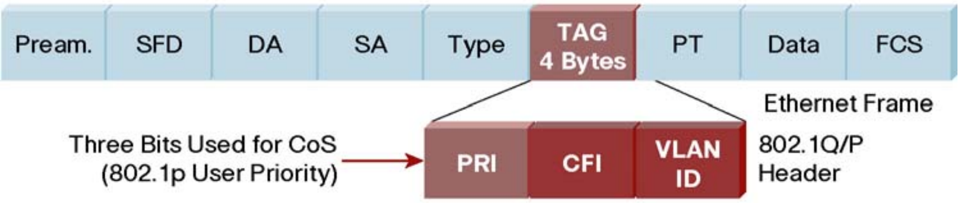

Figure 43. DSCP in the IP Header

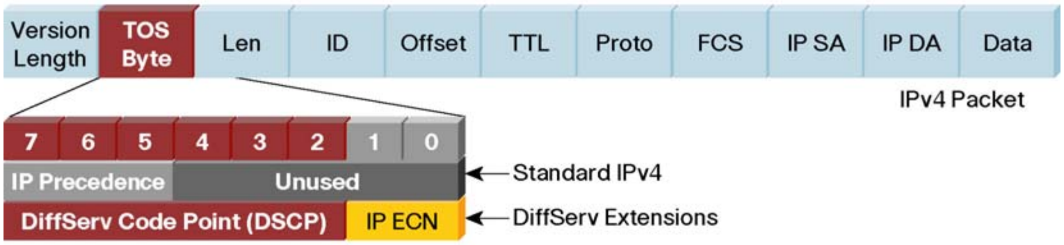

Ideally, Cisco Video Surveillance products will mark their packets with DSCP identifiers at the edge of the network. Then the access layer switches only need to be configured to trust these DSCP markings, and they can carry forward across the entire network. However, the current Cisco Video Surveillance products do not offer this capability (configurable DSCP marking capability is planned for a future release). In the interim, the access layer switches can be configured to mark traffic on behalf of the endpoints. Example configurations for this approach are provided in the “Campus QoS Tools” section of this document.

More on DiffServ Code Point

DSCP was initially defined in IETF RFC 2474 as a way to use the first six bits of the Type of Service (TOS) byte of the IP packet header to identify traffic for differentiated treatment in the network. The TOS byte previously had a definition using the first three bits to provide a similar marking for differentiated services called “IP Precedence.” However, the limitation of eight values from 0 – 7 did not provide enough flexibility to handle the various network traffic types that require identification. So the IETF further defined the DSCP values as the first six bits of the TOS byte, providing 64 potential values for the marking or tag.

The values of 0 – 63 are not recommended for use as random definitions of different traffic types, however. DSCP has been designed to be somewhat backwards compatible with the older IP Precedence definition. In this way, if the first three bits of a DSCP resolve to what would have been a 7 in IP Precedence, that is considered a higher class of traffic than if they resolved to a 6, which is higher than 5, and so on. These main definitions using the first three bits are now referred to as Class Selectors. The remaining three bits are used for more granular definition within the Class Selectors for attributes such as Drop Preference. The IETF has developed RFC 2597 and RFC 3246, which define recommendations for use of the additional bits within DSCP.

Default DSCP markings of Class Selector 4 (also expressed as DSCP 32 decimal, CS4 class identifier, or DSCP 100000 binary) are recommended for live video and live audio media traffic originating from Cisco Video Surveillance IP Gateways and Services Platforms. This marking should also be used for decoder-originated PTZ traffic, because it is conceptually part of the interactive media for the operator viewing the PTZ controlled video, and as such is sensitive to latency. This recommendation concurs with the most recent Cisco QoS Baseline guidelines for default values, as shown in Table 7.

Table 7. QoS Baseline 2.0 Default Markings

Application	DSCP Name	DSCP Value	Reference
Interactive Voice Media	EF	101110	RFC 3246
Interactive Video & Associated Voice Media	AF41, AF42, AF43	100010 100100 100110	RFC 2597
Streaming Video	Class 4	100000	RFC 2474 section 4.2.2
IP Routing	Class 6	110000	RFC 2474 section 4.2.2
Telephony Signaling (voice & video)	Class 3	011000	RFC 2474 section 4.2.2
Other	Default or class 0	000000	RFC 2474 section 4.1

In addition, the most recent informational IETF RFC 4594 defines a Real-Time Interactive Service Class using the suggested marking of CS4. The interactive nature of PTZ camera control, coupled with the fact that the MPEG-4 video streams are not rate adaptive to network conditions, fits the RFC guidelines for the class.

Table 8. DSCP to Service Class Mapping in RFC 4594

Service Class Name	DSCP Name	DSCP Value	Application Examples
Network Control	CS6	110000	Network routing
Telephony	EF	101110	IP Telephony bearer
Signaling	CS5	101000	IP Telephony signaling
Multimedia Conferencing	AF41, AF42, AF43	100010, 100100, 100110	H.323/V2 video conferencing (adaptive)
Real-Time Interactive	CS4	100000	Video conferencing and Interactive gaming
Multimedia Streaming	AF31, AF32, AF33	011010, 011100, 011110	Streaming video and audio on demand
Broadcast Video	CS3	011000	Broadcast TV & live events
Low-Latency Data	AF21, AF22, AF23	010010, 010100, 010110	Client/server transactions Web-based ordering
OAM	CS2	010000	OAM&P
High-Throughput Data	AF11, AF12, AF13	001010, 001100, 001110	Store and forward applications
Standard	DF (CS0)	000000	Undifferentiated applications
Low-Priority Data	CS1	001000	Any flow that has no BW assurance

DSCP default or recommended markings are only guidelines. The markings themselves have meaning only in the context of the Per Hop Behavior (PHB) definitions on the routers. PHB definitions consist of the queuing and scheduling behavior.

For the replay of recorded video streams, the traffic does not closely match the characteristics of an interactive or streaming class. The TCP protocol is used, and the traffic patterns are similar to a file-transfer within a data application. The large replay buffer at the decoder that is used for recorded streams can compensate for delay and jitter in the network. For this reason, a different DSCP marking is recommended for the replay of recorded streams. This DSCP marking should still provide preferential treatment, without the strict priority queuing that is often used for CS4 or EF DSCP traffic. The AF31 DSCP value provides a default marking to provide identification of this traffic, and corresponds to the Multimedia Streaming service class as defined in RFC 4594.

Cooperating with Cisco TelePresence

Cisco TelePresence is a new category of high-definition virtual presence solutions that utilize advanced visual, audio, and interactive technologies to create an in person experience over the network. TelePresence video network traffic requirements also correspond to the RFC 4594 Real-Time Interactive service class and carry a default DSCP marking of CS4. In some cases, it may be beneficial for live Cisco Video Surveillance and Cisco TelePresence traffic to carry different DSCP markings to easily distinguish the traffic types on the network. For example, multiple priority queues may be used to independently police the two traffic types, and DSCP may be used as the entrance criteria for the queues.

To achieve this differentiation, an alternative marking approach for live video surveillance streams is a CS5 identifier. While CS3 is more in line with the Broadcast Video service class definition of RFC 4594, the CS3 marking is used by Cisco IP Telephony products for call setup traffic. The use of the CS5 identifier as an alternative ensures that the live video surveillance traffic is distinguished from IP telephony call control. CS5 would be associated with a CoS marking of 5 through the DSCP-CoS map tables on Catalyst switches. Carrying these markings, it would still be placed into the hardware priority queue with the default CoS to queue mappings on Catalyst switches.

Queuing and Scheduling of Surveillance Traffic

Queues consist of data structures in memory that hold packets waiting to exit an interface. The mechanism for selecting the next packet to exit the interface is a scheduler. Queuing algorithms, such as weighted fair queuing and low latency queuing, consist of multiple queues and a scheduler that empties the queues in a specific sequence.

Cisco IOS Routers use software based queuing algorithms, which can be applied to interfaces.

- Weighted fair queuing (WFQ or fair-queue) is a flow-based queuing mechanism that ensures that all flows exiting an interface receive a fair share of bandwidth, while respecting a priority scheme (based on IP Precedence marking in some implementations).
- Class Based Weighted Fair Queuing (CBWFQ) is a sophisticated algorithm that allows a system administrator to allocate traffic to multiple queues based on characteristics of the traffic as defined within an Access Control List (ACL.)
- Low Latency Queuing (LLQ) is a special form of CBWFQ with a Priority Queue (PQ) layered on top of the Class Based queues. The PQ is unique in that it is exhaustive. That is, any traffic in the PQ is prioritized and scheduled over all other traffic on the interface. For this reason, it is important to properly provision and then police (limit) the PQ to ensure that it does not starve out other traffic that needs to exit an interface.

Cisco Catalyst Switches have queuing structures built into hardware. Common notation for the queuing capabilities of an interface or line card on the Catalyst 6500 platform uses the following abbreviations:

- p – Priority Queue,
- q – Standard Queue,
- t – Drop Threshold.

For example, the notation 1p2q2t implies one priority queue, two standard queues, and two drop thresholds within each queue, which defines how packets may be discarded when the queue fills and runs out of memory buffer space. Other switches have specific queue numbers set aside in hardware to be the priority queue. DSCP at Layer 3, and or Class of Service (CoS) markings at Layer 2 of packet headers are used to control entry to the hardware based queues.

QoS is not a substitute for proper network provisioning of link speeds to handle traffic loads. Also, provisioning a high level of bandwidth on a link is not a substitute or replacement for configuring QoS on the network. Network statistics can provide the utilization of an interface over time, such as 15% average utilization over 5 minutes. However, a low 5-minute average utilization statistic is not an indication that the network is never congested. IP data applications tend to be bursty in nature, and networks with reasonably low average utilization could actually be intermittently at 100% utilization for several hundred milliseconds at a time. Also, anywhere that there is a speed transition in a network device, for example from Gigabit or 10/100 Ethernet down to T-1 or other low-speed WAN transport, QoS is required to properly manage how the traffic is queued and scheduled. If there is 200 ms of additional network delay induced to a packet flow due to congestion, and that flow happens to be a PTZ controlled surveillance camera where the total delay budget is approximately 400 – 500 ms, the camera operator may experience jerky or slow camera reaction. QoS and proper bandwidth provisioning are part of the overall traffic engineering that is required to successfully deploy video surveillance and its associated applications on an IP network.

Campus QoS Tools

Converged campus networks generally consist of a minimum of 10/100 ports servicing end nodes, and Gigabit or 10 Gigabit links between switches that form the network. In such an environment, it would take 100 encoders running high motion traffic at D1/4CIF 30 frames per second to consume one-third of the total bandwidth on a Gigabit link.

- 1/3 of Gigabit (1000 Mbps) link approximately 330 Mbps
- 3,296 Kbps * 100 streams = 329.6 Mbps

In networks with many cameras, consider this requirement when provisioning campus links. Also consider the paths that video streams take in a fully operational network, and in a network with various link failures.

The general recommendation for high bandwidth environments is to place live surveillance traffic into the hardware priority queue on switching platforms. This configuration is accomplished through the inherent queuing mechanism on the switches with respect to CoS and DSCP markings. When taking this approach, ensure in the system design that the amount of traffic presented to a given link not exceed the priority queue size. Excessive traffic entering a strict priority queue can starve other traffic on the link of bandwidth.

On a high-bandwidth campus network, recorded surveillance streams can be transmitted within a standard data queue, as opposed to a strict hardware priority queue. The TCP transport protocol and large decoder buffer will compensate for minor variability in any but the most heavily saturated of campus networks. Bandwidth allocation for replay streams is discussed in the “WAN QoS” section of this document. Because recorded streams are generated only by the Cisco Video Surveillance Services Platform and Integrated Services Platform, a specific QoS approach is recommended to differentiate streams generated by these devices.

CoS and DSCP Mapping

CoS and DSCP are mapped to each other within Cisco Catalyst multilayer switches. In this way, the switch hardware can fill in the corresponding value when one of the two markings is trusted on an ingress interface. The default values for surveillance media include CoS 4, and DSCP 32. DSCP 32 is the decimal expression equivalent to CS4 or DSCP 100000 (binary). The Command Line

Interface (CLI) output below is from a Catalyst 3750 switch. It illustrates the DSCP-CoS map and the CoS-DSCP map.

```
3750-1#show mls qos maps dscp-cos
   Dscp-cos map:
     d1 :  d2 0  1  2  3  4  5  6  7  8  9
     ---------------------------------------
      0 :    00 00 00 00 00 00 00 00 01 01
      1 :    01 01 01 01 01 01 02 02 02 02
      2 :    02 02 02 02 03 03 03 03 03 03
      3 :    03 03 04 04 04 04 04 04 04 04
      4 :    05 05 05 05 05 05 05 05 06 06
      5 :    06 06 06 06 06 06 07 07 07 07
      6 :    07 07 07 07

   3750-1#show mls qos maps cos-dscp
   Cos-dscp map:
        cos:   0  1  2  3  4  5  6  7
     --------------------------------
       dscp:   0  8 16 24 32 40 48 56
```

The DSCP-CoS map is used to synchronize the CoS marking as necessary when DSCP is the known or trusted marking. The d1 value in the table above represents the high-order digit in a decimal expression of DSCP and the d2 value represents the low-order digit. For example, for DSCP 32, d1 = 3 and d2 = 2. The values in the body of the table are the resulting CoS values. For example, live video surveillance traffic set at DSCP 32 is mapped to CoS 4. Recorded video surveillance traffic set at DSCP 26 (equivalent to AF31 or 011010) is mapped to CoS 3.

The CoS-DSCP map is used to synchronize the DSCP marking as necessary when CoS is the known or trusted marking. The CoS-DSCP map in the example above is a little more intuitive, because there are less original values to map when starting with CoS. CoS 4 used for live video surveillance traffic maps to a DSCP value of 32. For recorded video surveillance traffic, the DSCP value is always the one trusted or set by the switch with a QoS service policy, so the CoS-DSCP table is not used to update a DSCP marking on these packets.

Additional tables exist in the switching platforms to control which DSCP and CoS values are placed into which hardware queues. These tables can vary by switching platform, so additional details about these tables and their configuration is provided in the platform-specific configuration example sections of this document.

Access Layer Requirements

As described in the “Deployment Models” section of this document, a Cisco best-practices hierarchical campus network design has multiple layers of switches, and each switch has specific functions. The Access Layer switches are where end node devices such as Cisco Video Surveillance IP Gateways and Service Platforms plug into the network. The Access Layer defines the edge of the network, which is the ingress point where traffic must be classified and marked for treatment further in the network.

Future implementations of Cisco Video Surveillance will rely on configurable DSCP markings originating from the end node devices themselves. The Access Layer switches will need to be configured to trust these markings and preserve them on ingress.

Cisco Video Surveillance IP Gateways running version 1.6.7 or earlier need an alternate method to appropriately classify and mark traffic for DSCP. One alternative is to create an IOS service policy using Access Control Lists (ACLs) to identify specific traffic and to force DSCP markings onto the packets at the ingress port. This approach allows the most granularity for different traffic types, but it adds the most complexity. This approach is required for the Cisco Video Surveillance Services Platform and Integrated Services Platform (SP/ISP) devices, because they have the capability to produce both live and recorded streams.

A simpler alternative that can be used for Cisco Video Surveillance IP Gateways is to use ingress packet marking to force a CoS or DSCP tag onto all incoming traffic from directly attached encoders and decoders. The downside of this approach is that all traffic, not just the live surveillance media packets, will then carry the high-priority tags. However, the non-media traffic issued from these devices is minimal in comparison to the media streams themselves, so in most cases the small amount of additional traffic being prioritized is not a problem for a properly provisioned network. **These recommendations are considered an interim solution until endpoint DSCP marking is available in Cisco Video Surveillance products.**

Access layer switches also receive live surveillance media traffic from distribution layer switches. As an example is an access layer that hosts decoders or recorders as directly attached devices. Another example is an access layer switch that hosts only encoders that are attached to PTZ dome cameras. The incoming PTZ control traffic is latency sensitive, so it is considered part of the interactive media stream. The simplest QoS approach is to configure all ports that connect to other trusted switches in the topology to trust incoming DSCP markings.

Distribution and Core Layer Requirements

The remaining layers of the network must also be configured for QoS. QoS configuration is required end-to-end, from the port where the IP traffic enters the network to the port where it exits. If not properly configured, a distribution or core layer switch could re-mark DSCP on all incoming packets down to zero, and live video surveillance traffic would be reduced to best-effort services from the network. When QoS is enabled on distribution and core layer switches, the switches must be configured to trust DSCP markings so that ingress traffic preserves these markings.

The first task for configuring QoS on distribution and core switches is to enable QoS globally. Then, ports that connect to other trusted switches (which should be all ports on distribution or core-specific switches) should be configured to trust the DSCP markings for incoming traffic. This configuration informs the switch receiving the packets that the source port is known and trusted and that the DSCP markings on these packets can reliably be used to make queuing decisions on the switch. Each switch in the path across the network must be properly configured for QoS so that DSCP and CoS markings are maintained end-to-end. This approach provides an easy way for core and distribution switches to identify traffic that belongs to surveillance streams without requiring complex ACL configuration at every hop across the network.

Enabling DSCP trust on switch ports requires that all ingress ports on the network be secured and properly configured for QoS policy. This configuration ensures that only the traffic that is intended for a specific prioritized queue is directed to that queue. If ports are improperly configured in the access layer, it could be possible for other applications to inject into the network traffic into that receives priority DSCP and CoS markings, which can hijack the bandwidth and possibly overrun the priority queue. QoS is fundamentally a system of “managed unfairness” where certain types of traffic receive preferential treatment at the expense of other traffic types. The only way to ensure that DSCP markings can be trusted at all points across the network is to carefully configure and manage all ingress points for network traffic according to a documented QoS policy that has been approved by the management of the organization operating the network. This approach is referred to as creating a trust boundary in the network to ensure that all CoS and DSCP markings on packets can be relied upon for queuing and scheduling.

Access Layer Service Policy Configuration Example for Services Platforms

Cisco Video Surveillance SP/ISP devices differ from the IP Gateways because they have the capability to transmit both live and recorded streams to the network. Live analog streams entering a 4-port Cisco Video Surveillance IP Gateway card in a USB Convergence Chassis or an ISP rely on the SP or ISP to transmit the live stream on the network. And the recorded streams on disk are played back across the IP network from the same devices.

Because multiple traffic types are transmitted by the same device, a more sophisticated classification and marking approach is required until the Cisco Video Surveillance endpoints have the ability to mark DSCP on their own. This approach uses the IOS Modular QoS Command Line Interface, or MQC. This interim approach is recommended until endpoint DSCP marking becomes available. The MQC approach for ingress classification and marking is platform independent and can be applied across various models of switches in the access layer. Additional hardware-specific QoS configuration requirements must be implemented to ensure that the markings generated by the service policy have the desired effect on queuing and scheduling of traffic. These requirements are described in each corresponding hardware-specific section of this document.

The Modular QoS CLI allows users to create traffic policies and attach these policies to interfaces. A traffic policy contains a traffic class and one or more QoS features. A traffic class classifies traffic, and the QoS features in the traffic policy determine how to treat the classified traffic. Modular QoS CLI configuration includes contains the following steps, which are described more thoroughly in the *Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.4*.

Step 1 Define traffic classes with the class-map command.

Step 2 Create a traffic policy by associating the traffic classes with one or more QoS features (using the policy-map command).

Step 3 Attach the traffic policy to the interface with the service-policy command.

The definition of a traffic class requires the use of a match statement. In this example configuration, Access Control Lists (ACLs) are used to match traffic belonging to live and recorded surveillance media streams. These ACLs can be general because the service policy to classify and mark ingress traffic is applied only to the ports where SP/ISP devices are attached.

```
6504-1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
6504-1(config)#ip access-list extended VS-LIVE-ACL
6504-1(config-ext-nacl)# permit udp any any
6504-1(config-ext-nacl)#ip access-list extended VS-REPLAY-ACL
6504-1(config-ext-nacl)# permit tcp any any
```

The next step (Step 1 from the MQC description above) is to define traffic classes with the class-map command.

```
6504-1(config)#class-map match-all VS-REPLAY-CLASS
6504-1(config-cmap)#  match access-group name VS-REPLAY-ACL
6504-1(config-cmap)#class-map match-all VS-LIVE-CLASS
6504-1(config-cmap)#  match access-group name VS-LIVE-ACL
```

This step uses the ACL that we created to define class names that are associated with live and recorded surveillance traffic. This approach may seem like an extra step in our simplified example, but the MQC is a robust structure that has been optimized for flexibility in aggregating QoS configuration to reduce redundancy in the application of QoS configuration to devices. The class-map structure provides a great number of match condition options that can increase the granularity of definition of the traffic in question

Use the policy-map command to associate the traffic classes with one or more QoS features. In this case, the only feature that we associate with the traffic classes is to mark the packets with DSCP for further queuing and scheduling across the network.

```
6504-1(config)#policy-map VS-NVR-INGRESS
6504-1(config-pmap)#  class VS-LIVE-CLASS
6504-1(config-pmap-c)#   set dscp cs4
6504-1(config-pmap-c)#  class VS-REPLAY-CLASS
6504-1(config-pmap-c)#   set dscp af31
```

After the policy map is created, we can attach the policy to an interface with the service-policy command. The service-policy must be applied in a specific direction, to inbound or outbound traffic. In this case, we want to apply the traffic in the inbound direction to the switch ports where SP/ISP devices are attached.

```
6504-1(config)#interface FastEthernet2/30
6504-1(config-if)#service-policy input VS-NVR-INGRESS
```

The resulting service policy status can be viewed on the interface using the show policy-map interface <interface id> command. This command is important for viewing the status of the service policy. It displays the rate of traffic offered to the classes so that you can ensure that traffic is meeting the defined criteria and being marked as expected.

```
6504-1#show policy-map interface fastethernet2/10
 FastEthernet2/10

  Service-policy input: VS-NVR-INGRESS

    class-map: VS-LIVE-CLASS (match-all)
      Match: access-group name VS-LIVE-ACL
      set dscp 32:
      Earl in slot 1 :
        154275104 bytes
        5 minute offered rate 2301432 bps
        aggregate-forwarded 154275104 bytes
    class-map: VS-REPLAY-CLASS (match-all)
      Match: access-group name VS-REPLAY-ACL
      set dscp 26:
      Earl in slot 1 :
        0 bytes
        5 minute offered rate 0 bps
        aggregate-forwarded 0 bytes

    Class-map: class-default (match-any)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: any
```

Catalyst 2960/3560/3750 Access Configuration Example

Enable QoS Globally

The Cisco Catalyst 2960, 3560, and 3750 series of switches are fixed-configuration switching solutions that offer sophisticated intelligence and QoS at the network edge. The Catalyst 3750 features StackWise technology, which joins multiple switches across a 32-Gbps stack interconnect. The QoS architecture and configuration command set is similar for all of these switches. The following example uses a Catalyst 3750 stack. To enable QoS for surveillance traffic, the first step is to globally enable QoS.

```
3750-1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
3750-1(config)#mls qos
3750-1(config)#
```

Assign VS Ports for CoS or DSCP Marking

The next step is to determine which ports have encoders attached. Until DSCP marking is available from the edge, these ports should be configured on Catalyst 2960/3560/3750 switches to mark CoS for live streams at the ingress port. First identify the CoS value to be used (the recommended value is 4). Then specify that the port should override any existing CoS value with the configured value. The incoming packets should not be carrying CoS values at ingress, because the access switch port will not be a trunk. However, this configuration is required.

```
3750-1(config-if)#interface gigabit 1/0/5
3750-1(config-if)#mls qos cos 4
3750-1(config-if)#mls qos cos override
```

When DSCP marking is supported directly by the Cisco Video Surveillance end nodes, the preferred configuration will be to set DSCP within the end nodes and then to configure the switch ports to trust DSCP. The Cisco Catalyst switch will also assign CoS to the packets at ingress based on the CoS-DSCP mapping table. End node DSCP marking will provide greater granularity of traffic control, with independent configuration on the surveillance end nodes for DSCP markings of different traffic types.

```
3750-1(config)#interface gigabit 1/0/6
3750-1(config-if)#mls qos trust dscp
```

When taking either of these approaches to force or honor QoS marking onto traffic entering a switch port, it is critical to ensure that these ports are physically secure and are not used for other applications. It is anticipated that most encoder and recorder ports are in data center environments, or in a secured area where analog coaxial cable connections are terminated. When a standalone encoder is used near the camera to convert to Ethernet, ensure the encoder ports are not readily accessible for someone to “borrow” the cable and connect a laptop or other endpoint. With the CoS override approach, any traffic that enters the network from such a point will carry CoS and ultimately DSCP marking, and enter the priority queue. With the trust DSCP approach, the end node can manipulate its DSCP values to anything desired and have them honored by the network. Also ensure that the CoS or DSCP marking configuration is removed when a port is reallocated so that it is is no longer used for surveillance media.

The default DSCP-CoS and CoS-DSCP mappings on the Catalyst 2960/3560/3750 switches result in live video surveillance traffic marked with CoS 4 to be mapped or re-marked to DSCP 32 (also called CS4 or DSCP 100000). This mapping provisions the packets to be properly queued at further hops across the network.

Enable Output Priority (Expedite) Queues

By default, the Catalyst 3750 platform uses a Shaped Round Robin (SRR) model for output queuing. A priority queue (also referred to as an *Expedite Queue*) also may be enabled. When this expedite queue is enabled, it is identified as queue number 1, and there are four total queues, numbered 1 – 4, on the platform. Identify the ports that will be egress ports for surveillance media traffic, and use the “priority-queue out” command to enable this capability on a per-port basis. Apply this configuration to ports which uplink to the distribution layer and to any end node port that is attached to a decoder or a recorder.

```
3750-1(config)#interface gigabit 1/0/1
3750-1(config-if)#priority-queue out
```

The “show mls qos interface <interface identifier> queueing” command allows verification that the egress priority queue is enabled

```
3750-1#show mls qos interface gigabit 1/0/1 q
GigabitEthernet1/0/1
Egress Priority Queue : enabled
Shaped queue weights (absolute) :  25 0 0 0
Shared queue weights  :  25 25 25 25
The port bandwidth limit : 100  (Operational Bandwidth:100.0)
The port is mapped to qset : 1
```

Configure CoS and DSCP Queue Maps

After the priority queue is enabled, ensure that live surveillance media traffic, marked with CoS 4 and DSCP 32, is placed in the appropriate queue. The Catalyst 3750 platform by default places this traffic into queue 4. The “show mls qos maps dscp-output-q” command displays the mapping of DSCP to queues. As in the DSCP-CoS map table, the d1 value corresponds to the high-order value of the decimal DSCP and the d2 value corresponds to the low-order value. For example, for DSCP 32, d1=3 and d2=2. The body of the table consists of two values per cell: the first is the queue number and the second is the drop threshold.

```
3750-1#show mls qos map dscp-out
   Dscp-outputq-threshold map:
     d1 :d2    0     1     2     3     4     5     6     7     8     9
     ------------------------------------------------------------
      0 :    02-01 02-01 02-01 02-01 02-01 02-01 02-01 02-01 02-01 02-01
      1 :    02-01 02-01 02-01 02-01 02-01 02-01 03-01 03-01 03-01 03-01
      2 :    03-01 03-01 03-01 03-01 03-01 03-01 03-01 03-01 03-01 03-01
      3 :    03-01 03-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01
      4 :    01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 04-01 04-01
      5 :    04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01
      6 :    04-01 04-01 04-01 04-01
```

The “show mls qos maps cos-output-q” command shows the mapping of CoS values to queues.

```
3750-1#show mls qos map cos-output-q
   Cos-outputq-threshold map:
              cos:  0   1   2   3   4   5   6   7
              ------------------------------------
  queue-threshold: 2-1 2-1 3-1 3-1 4-1 1-1 4-1 4-1
```

Because both of these markings default to placement of traffic in queue 4, they need to be updated to ensure that surveillance media traffic is placed in queue 1, the priority queue. In config mode, the **mls qos srr-queue** command allows you to manipulate the maps to place the traffic into the correct queue.

```
3750-1(config)#mls qos srr-queue output dscp-map queue 1 threshold 1 32

3750-1#show mls qos map dscp-out
   Dscp-outputq-threshold map:
     d1 :d2    0     1     2     3     4     5     6     7     8     9
     ------------------------------------------------------------
      0 :    02-01 02-01 02-01 02-01 02-01 02-01 02-01 02-01 02-01 02-01
      1 :    02-01 02-01 02-01 02-01 02-01 02-01 03-01 03-01 03-01 03-01
      2 :    03-01 03-01 03-01 03-01 03-01 03-01 03-01 03-01 03-01 03-01
      3 :    03-01 03-01 01-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01
      4 :    01-01 01-01 01-01 01-01 01-01 01-01 01-01 01-01 04-01 04-01
      5 :    04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01 04-01
      6 :    04-01 04-01 04-01 04-01

3750-1(config)#mls qos srr-queue output cos-map queue 1 threshold 1 4

3750-1#show mls qos map cos-output-q
   Cos-outputq-threshold map:
              cos:  0   1   2   3   4   5   6   7
              ------------------------------------
  queue-threshold: 2-1 2-1 3-1 3-1 1-1 1-1 4-1 4-1
```

With these steps completed, we have ensured that live surveillance media traffic is marked with CoS 4 and DSCP 32, enabled a priority queue on the egress ports facing the distribution layer, and altered the queue mapping to ensure that traffic carrying these markings is placed in the priority queue. This configuration ensures that live surveillance media packets that are sent across the uplinks have priority over applications that use the other queues on the switch.

Provisioning Access Layer Switches for Ingress Traffic

When configuring QoS for surveillance media traffic, we are primarily concerned with egress queuing in the direction of the video stream, and ingress DSCP trust in the same direction. Several design scenarios result in traffic flowing from the Distribution Layer toward the Access Layer that requires prioritization. For example, access switches that have decoders or recorders attached or have Cisco Stream Manager Client Viewing Module installed receive surveillance media streams on uplink ports. Even encoders that are configured for PTZ control receive traffic that comes in the opposite direction from the voice and audio media flows.

Best practices dictate preserving DSCP markings end-to end across the life of the marked flow. To do so, the uplink ports that face the distribution layer should be configured to trust incoming DSCP. Inbound DSCP trust configuration syntax is described in the “Distribution/Core and Linear Configuration Example” section.

Catalyst 2960/3560/3750 Distribution/Core and Linear Configuration Example

To cause QoS processing to occur on the switch, the first step is to globally enable QoS processing:

```
3750-1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
3750-1(config)#mls qos
3750-1(config)#
```

Then, determine which ports will carry inbound surveillance traffic (inbound from the perspective of the switch itself,) or other inbound traffic types that have valid DSCP markings. Configure these ports to trust the DSCP markings on packets that enter the interface.

```
3750-1(config-if)#interface gigabit 1/0/1
3750-1(config-if)#mls qos trust DSCP
```

As with the Access Layer configuration example, the priority queue (also called *Expedite Queue*) must be enabled on the ports of the switch that surveillance media traffic egresses. Also, queue mappings must be altered to ensure that surveillance media traffic enters the priority queues. Refer to the configuration examples in the “Catalyst 3750 Access” section for the appropriate syntax to complete these tasks.

StackWise vs. Daisy-Chain or Linear Topologies

The Catalyst 3750 StackWise technology allows all switches within the stack to appear as a single entity. The switches are managed through a single interface, and there is no requirement to configure additional QoS statements for traffic between the switches.

Cisco 2960 and 3560 switches are often used in daisy-chain configurations. This approachs is sometimes referred to as *stacking* in an enterprise access layer. The Catalyst 3750 can also be used in this way, in the absence of a StackWise interconnect cable. Many custom topologies also are used for dedicated surveillance networks. These topologies, which do not fall into the traditional Hierarchical Network Design model, are described in the “Rings and Linear Topologies” section of this document. Any links between switches in such a topology should have ports configured to trust DSCP on ingress. Physical access must be controlled and all ingress ports must be properly configured to ensure that DSCP markings on all packets can truly be trusted within the QoS policy of the organization.

Reference Documents:

Catalyst 3750 Switch Software Configuration Guide, 12.2(25)SEE, Configuring QoS
www.cisco.com/en/US/partner/products/hw/switches/ps5023/products_configuration_guide_chapter09186a00805a6504.html

Catalyst 4500 Access Configuration Example

Enable QoS Globally

The Catalyst 4500 series of switches are high-performance, chassis-based systems with advanced QoS capabilities. The QoS configuration syntax is similar to that found on the Catalyst 2960/3560/3750 series switches, with the exception that the “mls” keyword is omitted in the syntax of most commands. The platform <the Catalyst 4500?> ships with QoS disabled globally, so the first step is to access the system and enable QoS.

```
4507-1#show qos
QoS is disabled globally
IP header DSCP rewrite is enabled

4507-1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
4507-1(config)#qos
4507-1(config)#^Z

4507-1#show qos
QoS is enabled globally
IP header DSCP rewrite is enabled
```

Assign VS Ports for DSCP Marking

Next, determine which ports have attached encoders or other devices that produce live surveillance media. Until DSCP marking is available from the edge, these ports should be configured on a Catalyst 4500 platform to mark DSCP at the ingress port. To do so, identify the DSCP value to be used (the recommended value is 32). Then, apply the value as the default ingress DSCP for the port.

```
4507-1(config)#interface gigabit 3/5
4507-1(config-if)#qos dscp 32
```

When DSCP marking is supported directly by Cisco Video Surveillance end nodes, the preferred configuration will be to set DSCP within the end nodes and configure the switch ports to trust DSCP. The DSCP-CoS map in the switch will cause the switch to also assign the associated CoS value to packets at ingress. This approach will improve granularity of traffic control. It also will provide for independently configuring DSCP for different traffic types on surveillance end nodes.

```
4507-1(config)#interface gigabit 3/6
4507-1(config-if)#qos trust dscp
```

When using either of these approaches to force or honor QoS marking onto traffic entering a switch port, ensure that these ports are physically secure and that they are not used for other applications.

Typically, many encoder and recorder ports are in data center environments or in secured areas where analog coaxial cable connections terminate. When a standalone encoder is physically deployed near a camera, ensure that ports are not readily accessible for someone to “borrow” the cable and connect a laptop or other endpoint. With the forced DSCP approach, any traffic that enters the network from such a point carries DSCP and CoS marking, and enters the priority queue. With the trust DSCP approach, the end node could manipulate its DSCP value to anything desired, and have it honored by the network. Also ensure that the DSCP marking configuration is removed when a port is no longer used for surveillance media.

The default DSCP-CoS and CoS-DSCP mappings on the Catalyst 4500 platform result in traffic that is marked with DSCP 32 (also called CS4 or DSCP 100000) to be mapped or re-marked for CoS 4. In this way, packets are properly queued at later hops across the network.

Enable Transmit Strict-Priority Queue

The Catalyst 4500 supports four egress queues, which may be configured in either 4Q1T or 1P3Q1T modes. The strict-priority queue is transmit-queue 3. To enable the strict-priority transmit queue on an interface, use the tx-queue 3 interface command followed by the priority high sub-command. It is a best practice to place a peak bandwidth limitation on the link, which protects unintended traffic that might enter the priority queue from saturating the link and starving out all other traffic. Identify the ports that will be egress ports for live surveillance media traffic and configure them for a strict priority queue. In an access layer switch, these ports typically are the ones that connect to the distribution layer.

```
4507-1(config)#interface gigabit 1/3
4507-1(config-if)#tx-queue 3
4507-1(config-if-tx-queue)#priority high
4507-1(config-if-tx-queue)#bandwidth percent 30
```

The “show qos interface <interface identifier>” command allows verification that the egress priority queue is enabled.

```
4507-1#show qos interface gigabit 1/3
QoS is enabled globally
Port QoS is enabled
Administrative Port Trust State: 'untrusted'
Operational Port Trust State: 'untrusted'
Trust device: none
Default DSCP: 0 Default CoS: 0
Appliance trust: none
Tx-Queue   Bandwidth   ShapeRate   Priority   QueueSize
             (bps)       (bps)                (packets)
  1        250000000   disabled    N/A        2080
  2        250000000   disabled    N/A        2080
  3        300000000   disabled    high       2080
  4        250000000   disabled    N/A        2080
```

Verify DSCP Queue Maps

After the priority queue has been enabled, ensure that live surveillance media traffic, marked with DSCP 32, is placed in the appropriate queue. By default, the Catalyst 4500 platform places this traffic into queue 3, which has been enabled as the priority queue. The “show qos map dscp tx-queue” command displays the mapping of DSCP to queues. As in the DSCP-CoS mapping tables, the d1 value corresponds to the high-order value of the decimal DSCP, and the d2 value corresponds to the low-order value. For example, for DSCP 32, d1 = 3 and d2 = 2. The body of the table consists of the queue numbers that are used for transmit. The default configuration is shown below.

```
4507-1#show qos map dscp tx-queue
DSCP-TxQueue Mapping Table (dscp = d1d2)
d1 : d2  0  1  2  3  4  5  6  7  8  9
-------------------------------------
 0 :    01 01 01 01 01 01 01 01 01 01
 1 :    01 01 01 01 01 01 02 02 02 02
 2 :    02 02 02 02 02 02 02 02 02 02
 3 :    02 02 03 03 03 03 03 03 03 03
 4 :    03 03 03 03 03 03 03 03 04 04
 5 :    04 04 04 04 04 04 04 04 04 04
 6 :    04 04 04 04
```

These steps ensure that live surveillance media traffic is marked with CoS 4 and DSCP 32, enable a priority queue on the egress ports that face the distribution layer, and alter the queue mapping to ensure that traffic carrying these markings is placed into the priority queue. This configuration ensures that live surveillance media packets receive priority across the uplinks over applications using the remaining queues on the switch.

Provisioning Access Layer Switches for Ingress Traffic

When configuring QoS for surveillance media traffic, we are primarily concerned with egress queuing in the direction of the video stream and with ingress DSCP trust in the same direction. Several design scenarios result in traffic that flows from the Distribution Layer toward the Access Layer and that requires prioritization. For example, access switches that have decoders, the Cisco Stream Manager software, or recorders attached receive surveillance media streams on uplink ports. Encoders that are configured for PTZ control receive that traffic coming in the opposite direction of the voice and audio media flows.

Best practices dictate preserving DSCP markings end-to-end across the life of the marked flow. To do so, the uplink ports facing the distribution layer should be configured to trust incoming DSCP. Inbound DSCP trust configuration syntax is described in the “Distribution/Core and Linear Configuration Example” section.

Catalyst 4500 Distribution/Core and Linear Configuration Example

To enable any QoS processing to occur on the switch, the first step is to globally enable QoS processing:

```
4507-1#show qos
QoS is disabled globally
IP header DSCP rewrite is enabled

4507-1#configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
4507-1(config)#qos
4507-1(config)#^Z

4507-1#show qos
QoS is enabled globally
IP header DSCP rewrite is enabled
```

Then, determine which ports will carry inbound live surveillance traffic (inbound from the perspective of the switch) or other traffic types that have valid DSCP markings. Configure these ports to trust the DSCP markings on packets that enter the interface.

```
4507-1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
4507-1(config)#int gig 1/3
4507-1(config-if)#qos trust dscp
4507-1(config-if)#^Z
4507-1#show qos interface gigabit 1/3
QoS is enabled globally
Port QoS is enabled
Administrative Port Trust State: 'dscp'
Operational Port Trust State: 'dscp'
Trust device: none
Default DSCP: 0 Default CoS: 0
Appliance trust: none
Tx-Queue   Bandwidth   ShapeRate   Priority   QueueSize
             (bps)       (bps)                (packets)
  1        250000000   disabled    N/A        2080
  2        250000000   disabled    N/A        2080
  3        300000000   disabled    high       2080
  4        250000000   disabled    N/A        2080
```

As with the Access Layer configuration example, the priority queue must be enabled on the switch ports that surveillance media traffic egresses. Also, queue mappings must be altered to ensure that live surveillance media traffic enters the priority queues. Refer to the configuration examples in the Catalyst 4500 Access section above for the syntax to perform this configuration.

Catalyst 6500 Access Configuration Example

Enable QoS Globally

The Catalyst 6500 series is the Cisco flagship campus switching product, with the most flexible suite of line cards and the availability of powerful services modules. There are two operating system loads for the Catalyst 6500: Catalyst OS and 6500 Supervisor IOS. Supervisor IOS provides command line syntax similar to that of the Catalyst 4500 and the 3750 platforms. Over time, it will become the single OS available for the Catalyst 6500, so examples in this document focus on Supervisor IOS syntax.

The Catalyst 6500 default configuration has QoS disabled, so the first step in configuring QoS on the Catalyst 6500 is to enable QoS globally.

```
6503-1#show mls qos
  QoS is disabled globally
6503-1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
6503-1(config)#mls qos
6503-1(config)#^Z
6503-1#show mls qos
  QoS is enabled globally
  Policy marking depends on port_trust
  QoS ip packet dscp rewrite enabled globally
  Input mode for GRE Tunnel is Pipe mode
  Input mode for MPLS is Pipe mode
  Vlan or Portchannel(Multi-Earl) policies supported: Yes
  Egress policies supported: Yes

 ----- Module [1] -----
  QoS global counters:
    Total packets: 6
    IP shortcut packets: 0
    Packets dropped by policing: 0
    IP packets with TOS changed by policing: 6
    IP packets with COS changed by policing: 0
    Non-IP packets with COS changed by policing: 0
    MPLS packets with EXP changed by policing: 0
```

Assign VS ports for DSCP Marking

Next, determine which ports have attached encoders or other devices that produce live surveillance media. Until DSCP marking is available from the edge, these ports should be configured on a Catalyst 6500 platform to mark CoS at the ingress port. To do so, identify the CoS value to be used (the recommended value is 4). Then, apply the value as the default ingress CoS for the port. The Catalyst 6500 platform requires that **mls qos trust cos** be enabled on the port if the switch is to alter the CoS markings of packets at ingress.

```
6503-1(config)#interface FastEthernet 2/5
6503-1(config)# mls qos trust cos
6503-1(config-if)#mls qos cos 4
```

When DSCP marking is supported by Cisco Video Surveillance end nodes, the preferred configuration will be to set DSCP within the end nodes and then configure the switch ports to trust DSCP. With DSCP trust enabled, the switch will assign the associated CoS value from the DSCP-CoS map to packets at ingress. This approach will improve granularity of traffic control. It also will provide for independently configuring DSCP of different traffic types on surveillance end nodes.

```
6503-1(config)#interface FastEthernet 2/6
6503-1(config-if)#mls qos trust dscp
```

When using either of these approaches to force or honor QoS marking onto traffic entering a switch port, ensure that these ports are physically secure and that they are not used for other applications. Typically, most encoder and recorder ports are in data center environments or in secured areas where analog coaxial cable connections terminate. When a standalone encoder is deployed physically near the camera, ensure those ports are not readily accessible for someone to “borrow” the cable and connect a laptop or other endpoint. With the forced CoS approach, any traffic that enters the network from such a point carries CoS and DSCP marking, and enters the priority queue. With the trust DSCP approach, the end node could manipulate its DSCP value to anything desired, and have it honored by the network. Also ensure that the DSCP marking configuration is removed when a port is no longer used for surveillance media.

The default DSCP-CoS and CoS-DSCP mappings on the Catalyst 6500 platform result in traffic that is marked with DSCP 32 (also called CS4 or DSCP 100000) to also be mapped or re-marked for CoS 4. In this way, packets are properly queued at later hops across the network.

```
6503-1#show mls qos maps dscp-cos
   Dscp-cos map:                                  (dscp= d1d2)
     d1 :  d2 0  1  2  3  4  5  6  7  8  9
     -------------------------------------
      0 :    00 00 00 00 00 00 00 00 01 01
      1 :    01 01 01 01 01 01 02 02 02 02
      2 :    02 02 02 02 03 03 03 03 03 03
      3 :    03 03 04 04 04 04 04 04 04 04
      4 :    05 05 05 05 05 05 05 05 06 06
      5 :    06 06 06 06 06 06 07 07 07 07
      6 :    07 07 07 07

6503-1#show mls qos maps cos-dscp
   Cos-dscp map:
         cos:   0  1  2  3  4  5  6  7
     ------------------------------------
        dscp:   0  8 16 24 32 40 48 56
```

Queuing Surveillance Media on the Catalyst 6500

The Catalyst 6500 has a distributed QoS architecture. The classification and marking functions are controlled by the central processor, while queuing implementations are hardware-specific on individual line cards. For a detailed explanation of the QoS capabilities of the Catalyst 6500 and of individual line cards, refer to *Quality of Service Solutions Reference Network Design Guide.*

For live surveillance media in the access layer of a hierarchical campus network design, we are primarily concerned with the egress queuing on the ports that face the distribution layer. The objective is to ensure that live surveillance media traffic is placed in the priority queue on uplink ports that face the distribution switches. To view the default queuing configuration of the port in question, use the **show queueing interface *<interface id>*** command.

> Note: The **show queueing interface** command displays configuration of both transmit and receive queuing configurations, each of which include extensive detail. The receive queuing command output has been truncated from the example below Ensure you are viewing the transmit portion of the command output when verifying configuration.

```
6503-1#show queueing interface GigabitEthernet 1/1
Interface GigabitEthernet1/1 queueing strategy:  Weighted Round-Robin
  Port QoS is enabled
  Port is untrusted
  Extend trust state: not trusted [COS = 0]
  Default COS is 0
    Queueing Mode In Tx direction: mode-cos
    Transmit queues [type = 1p3q8t]:
    Queue Id    Scheduling  Num of thresholds
    -----------------------------------------
       01         WRR                 08
       02         WRR                 08
       03         WRR                 08
       04         Priority            01
    WRR bandwidth ratios:  100[queue 1] 150[queue 2] 200[queue 3]
    queue-limit ratios:     50[queue 1]  20[queue 2]  15[queue 3]  15[Pri
Queue]

    queue tail-drop-thresholds
    --------------------------
    1     70[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
    2     70[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
    3     100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]

    queue random-detect-min-thresholds
    ----------------------------------
      1    40[1] 70[2] 70[3] 70[4] 70[5] 70[6] 70[7] 70[8]
      2    40[1] 70[2] 70[3] 70[4] 70[5] 70[6] 70[7] 70[8]
      3    70[1] 70[2] 70[3] 70[4] 70[5] 70[6] 70[7] 70[8]

    queue random-detect-max-thresholds
    ----------------------------------
      1    70[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
      2    70[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]
      3    100[1] 100[2] 100[3] 100[4] 100[5] 100[6] 100[7] 100[8]

    WRED disabled queues:

    queue thresh cos-map
    ---------------------------------------
    1     1      0
    1     2      1
    1     3
    1     4
    1     5
    1     6
    1     7
    1     8
    2     1      2
    2     2      3 4
    2     3
    2     4
    2     5
    2     6
    2     7
    2     8
    3     1      6 7
    3     2
    3     3
    3     4
    3     5
    3     6
    3     7
    3     8
    4     1      5
```

In the example above, 1p3q8t is the hardware queuing configuration of port GigabitEthernet 1/1. This designation is shorthand for one priority queue, three hardware based queues, and eight drop thresholds per queue. The queuing mode for transmit traffic queue distribution is CoS based, which means that CoS markings decide which queue traffic is placed in for egress scheduling. The queue thresh(old) cos-map table shows that the default configuration places CoS 4 traffic into queue 2 in the second drop threshold.

To alter the cos-map table for the priority queue, use the priority-queue cos-map command to place CoS 4 traffic into the priority queue. Notice that this queuing configuration is propagated to all of the gigabit ports on the module.

```
6503-1(config-if)#priority-queue cos-map 1 4
Propagating cos-map configuration to:  Gi1/1 Gi1/2 Gi1/3 Gi1/4 Gi1/5
Gi1/6 Gi1/7 Gi1/8 Gi1/9
```

These steps ensure that live surveillance media traffic is marked with CoS 4 and DSCP 32, enable a priority queue on the egress ports that face the distribution layer, and alter the queue mapping to ensure that traffic carrying these markings is placed into the priority queue. This configuration ensures that live surveillance media packets receive priority across the uplinks over applications using the remaining queues on the switch.

Provisioning Access Layer Switches for Ingress Traffic

When configuring QoS for surveillance media traffic, we are primarily concerned with egress queuing in the direction of the video stream and with ingress DSCP trust in the same direction. Several design scenarios result in traffic that flows from the Distribution Layer out toward the Access Layer and that requires prioritization. For example, access switches that have decoders, the Cisco Stream Manager software, or recorders attached receive surveillance media streams on uplink ports. Encoders that are configured for PTZ control receive that traffic coming in the opposite direction of the voice and audio media flows.

Best practices dictate preserving DSCP markings end-to end across the life of the marked flow. To do so, the uplink ports facing the distribution layer should be configured to trust incoming DSCP. Inbound DSCP trust configuration syntax is described in the “Distribution/Core and Linear Configuration Example” section.

Catalyst 6500 Distribution/Core Configuration Example

To enable any QoS processing to occur on the switch, the first step is to globally enable QoS processing:

```
6503-1#show mls qos
  QoS is disabled globally
6503-1#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
6503-1(config)#mls qos
6503-1(config)#^Z
6503-1#show mls qos
  QoS is enabled globally
  Policy marking depends on port_trust
  QoS ip packet dscp rewrite enabled globally
  Input mode for GRE Tunnel is Pipe mode
  Input mode for MPLS is Pipe mode
  Vlan or Portchannel(Multi-Earl) policies supported: Yes
  Egress policies supported: Yes

 ----- Module [1] -----
  QoS global counters:
    Total packets: 6
    IP shortcut packets: 0
    Packets dropped by policing: 0
    IP packets with TOS changed by policing: 6
    IP packets with COS changed by policing: 0
    Non-IP packets with COS changed by policing: 0
    MPLS packets with EXP changed by policing: 0
```

Then, determine which ports will carry inbound live surveillance traffic (inbound from the perspective of the switch) or other traffic types that have valid DSCP markings. Configure these ports to trust the DSCP markings on packets that enter the interface.

```
6503-1(config)#interface GigabitEthernet 1/1
6503-1(config-if)#mls qos trust dscp
6503-1#show queueing interface gigabit 1/1
Interface GigabitEthernet1/1 queueing strategy:  Weighted Round-Robin
  Port QoS is enabled
  Trust state: trust DSCP
       ...
```

As with the Access Layer configuration example, the queue mappings must be altered to ensure that live surveillance media traffic enters the priority queues. Refer to the configuration examples in the Catalyst 6500 Access section above for the syntax to perform this configuration.

WAN QoS Tools

Cisco IOS routers perform queuing and scheduling processes in software instead of hardware. Because the bandwidth in use on the WAN is generally much lower than the gigabit or 10 gigabit speeds of the campus, the software model allows flexibility in configuration and features while staying within the bounds of the processing capability of the platform. The primary QoS tool for live surveillance media traffic on a converged network is Low Latency Queuing (LLQ.) LLQ is an enhanced implementation of Class Based Weighted Fair Queuing (CBWFQ) with the addition of an exhaustive priority queue for latency- sensitive traffic. LLQ also is the QoS tool of choice for managing IP telephony and video conferencing traffic on a Cisco network, due to the similar sensitivity to loss, delay, and jitter that characterizes IP multimedia applications. When LLQ is used to manage video traffic and voice on a WAN, dual priority queues may be configured to ensure that video traffic is policed independently from voice traffic. This configuration allows applications to share the priority bandwidth without impacting each other’s performance. For additional information about LLQ, CBWFQ, and priority queue configuration for Cisco IP telephony and video conferencing, refer to *Quality of Service Solutions Reference Network Design Guide.*

LLQ allows the configuration of a software-based priority queue structure that can be assigned to a router interface. The priority queue may be configured as a percentage of the available bandwidth or as a specific bit-rate.

On a WAN, where bandwidth is typically constrained, specific steps must be taken to ensure the delivery of recorded surveillance streams, where applicable. The Transport Control Protocol (TCP) used for recorded streams and the large decoder buffer preclude the requirement for a strict LLQ priority queue. However, it is common for WAN links to congest, so a Class Based Weighted Fair Queue with an assigned bandwidth is recommended for recorded traffic. This approach ensures that a portion of the link bandwidth is available to carry recorded surveillance traffic when required. Traffic shaping also is recommended to ensure that the stream of recorded traffic is smoothed as it is transmitted across the WAN, leaving bandwidth available for other applications that share the link. The recommended IOS software based queuing model for traffic associated with Cisco Video Surveillance is shown in Figure 44.

Figure 44. Video Surveillance Egress Queuing Model

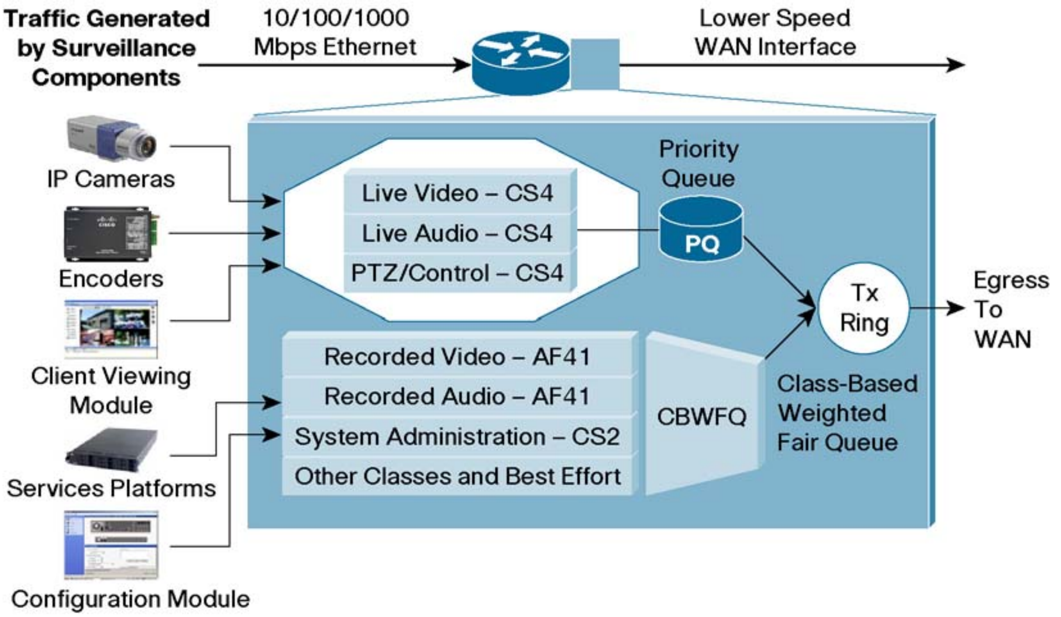

Priority Queuing Estimates for MPEG-4 Traffic

Video traffic in general is bursty by the standards of IP networking. The MPEG-4 compression codec technology uses different classes of video frames to accomplish compression of an overall stream. The bandwidth requirements of a given stream also are affected by the complexity of the image being displayed and by the amount of change or motion in the source analog NTSC or PAL signal.

IOS LLQ accepts two configuration parameters for a priority queue: the rate of the queue expressed in KB per second, and an optional burst size configured in bytes. If the burst size is not configured, IOS generates a default burst size equal to one-fifth of the rate per second. This default value works well for traffic that is smooth in its network use, such as live audio surveillance or Cisco IP Telephony. Live surveillance video traffic is burstier. Increasing the burst size above the default-generated value can allow traffic to stay within the limits of the queue without needing to increase the overall rate of the PQ.

In a WAN environment with constrained bandwidth, there is a tradeoff between the required quality level of a video stream and the bandwidth that must be provisioned to carry the stream. The greater the resolution of the compressed video, the more bandwidth is required to carry the stream. The greater the frame rate of a video stream, the more bandwidth required. When using the Constant Bit Rate model (CBR) on the encoder, as recommended for a WAN environment, average bandwidth consumption can be estimated as the target CBR bit rate plus ten percent. This calculation applies to live and to recorded video streams. When sizing a priority queue that carries live surveillance video traffic, peak video bandwidth consumption, not averages, must be used.

Table 9 provides guideline values for PQ and burst sizing. These values are based on testing that was conducted using Cisco Video Surveillance Gateways with Cisco IOS 2851 ISR routers configured for LLQ. Priority queue and burst size requirements were determined to support live video surveillance traffic at various resolutions and frame rates. These tests were conducted with the encoders configured using the target values in CBR Rate Guidelines Because WAN QoS and LLQ are software based, this configuration is valid across the range of Cisco enterprise routing platforms. During LLQ testing, the tested link was fully saturated with background traffic flowing in the same direction as the video stream. <Edits to the table: Suggest removing the "***" in both cases; Change "Guideline CBR Rates" to "guideline CBR rates"; change "Kbps" to "kbps"; change "Bytes" to "bytes"

Table 9. IOS LLQ PQ and Burst Guidelines

IOS LLQ Priority Queue and Burst Sizes when using CBR Rate Guidelines			
NTSC	**Resolution**		
Frame Rate	**CIF**	**2CIF**	**4CIF or D1**
1.5	255/25575	379/37950	742/74250
2	330/33000	519/51975	990/99000
3	429/42900	676/67650	1270/127050
3.75	495/49500	783/78375	1542/154275
5	544/40837	866/64968	1732/129937
7.5	660/33000	1237/61875	2310/115500
10	768/38425	1305/65250	2465/123250
15	870/43500	1595/79750	3190/159500
30	1232/61625	2320/116000	4350/217500
PAL	**Resolution**		
Frame Rate	**CIF**	**2CIF**	**4CIF or D1**
1	305/30525	445/44550	891/89100
2.5	545/54445	866/86625	1733/173250
5	701/52594	1155/86625	2310/173250
6.25	784/58781	1320/99000	2640/198000
12.5	1015/50750	1740/87000	3480/174000
25	1595/79750	2610/130500	4930/246500
Values are only for use when using SRND Guideline CBR Rates Each cell contains PQ size in Kbps / Burst size in Bytes			

Note: The PQ and Burst sizing guidelines in Table 9 are valid only when used in context with the CBR Rate guidelines that are provided in Table 5.

The numbers in Table 9 are for a single video stream. When several streams are configured across a WAN link, obtain provisioning guidance by multiplying the values for the given resolution and frame rate by the number of streams. A single stream evaluation represents a worst-case scenario The real-world use of multiple video streams is generally lower than the aggregate number, because statistically the streams tend to not all have high motion levels concurrently. In addition, the codecs are be on different framing cycles, so the more streams that are on the link, the less likely that they will all transmit high-bandwidth I-frames at the same instant within each second. However, in order to guarantee there will be no drops in the Priority Queue, it must be sized to handle the possibility of all streams hitting their peak bandwidth consumption concurrently.

Estimating Priority Queue and Burst Sizes using Formulas

The IOS LLQ priority queue and burst sizes in this document were validated in lab testing, but were originally calculated using formulas. The formulas take the CBR target setting as an input to derive the recommended queue and burst sizes. The formula used changes according to the frame rate of the source video stream. This was found to be necessary in lab testing, and corresponds to a video stream with fewer frames per second being burstier by nature than a stream with a greater number of frames per second. The greater number of frames results in a smoothing of traffic on the network, even though the overall bit rate is generally higher. The following formulas were used to derive the PQ and burst settings in IOS LLQ PQ and Burst Guidelines:

30 – 10 frames per second

Priority Queue = CBR Bit Rate * 1.45
Burst bytes = Priority Queue bit rate * 1Byte/8bits * .4

7.5 frames per second

Priority Queue = CBR Bit Rate * 1.65
Burst bytes = Priority Queue bit rate * 1Byte/8bits * .4

5 frames per second

Priority Queue = CBR Bit Rate * 1.65
Burst bytes = Priority Queue bit rate * 1Byte/8bits * .6

3.75 – 1.5 frames per second

Priority Queue = CBR Bit Rate * 1.65
Burst bytes = Priority Queue bit rate * 1Byte/8bits * .8

These formulas may also be used to create a starting point for CBR rates that are set differently from the guidelines in Table 5. This provides additional flexibility in how the CBR rate is set while still providing guidance in terms of priority queue and burst configuration. The optimal priority queue and burst settings are as low as possible on the network without having packets dropped (policed) from the queue, the formulas provide a conservative approximation with some headroom.

LLQ and CBWFQ Configuration Example

LLQ allows the configuration of a software-based priority queue structure that can be assigned to a router interface. The priority queue may be configured as a percentage of the available bandwidth, or as a specific bit-rate. Specific bit-rates were used in testing and the examples in this document, as they are more easily portable to transport links of various bandwidths.

An LLQ priority queue bandwidth may be configured as a bit-rate only, or as a combination of a bit-rate and a burst value. The bit-rate is expressed in Kilobits per second, while the burst value is expressed in Bytes. These two parameters combined control the queuing, scheduling, and policing of the priority queue. Increasing the default burst rate for the priority queue to a higher number than the size of the default burst generated by IOS allows a much lower overall bit rate to be allocated to the queue itself.

The LLQ and CBWFQ configuration process benefits from the use of use of the IOS Modular QoS Command Line Interface, or MQC. The Modular QoS CLI is a structure that allows users to create traffic policies and attach these policies to interfaces. A traffic policy contains a traffic class and one or more QoS features. A traffic class is used to classify traffic, while the QoS features in the traffic policy determine how to treat the classified traffic. Modular QoS CLI configuration includes contains the following three steps, which are detailed more thoroughly in the *Cisco IOS Quality of Service Solutions Configuration Guide, Release 12.4*.

Step 1 Define traffic classes with the class-map command.

Step 2 Create a traffic policy by associating the traffic classes with one or more QoS features (using the policy-map command).

Step 3 Attach the traffic policy to the interface with the service-policy command.

The definition of a traffic class requires the use of a match statement. In our example configuration, we will use Access Control Lists (ACLs) to match traffic belonging to live and recorded surveillance media streams. Building the Access Control lists will be simplified by the DSCP marking process described in the Campus portion of this document. Build named ACLs to permit traffic that is marked with DSCP CS4 for live video, and AF31 for recorded.

```
2851-1#configure terminal
Enter configuration commands, one per line.  End with CNTL/Z.
2851-1(config)#ip access-list extended VS-LIVE-ACL
2851-1(config-ext-nacl)# permit ip any any dscp cs4
2851-1(config-ext-nacl)#ip access-list extended VS-REPLAY-ACL
2851-1(config-ext-nacl)# permit ip any any dscp af31
```

The ACLs defined above are relying on the validity of the DSCP markings on live and recorded surveillance media traffic for identification. If more granularity is required, the IOS ACL configuration is very robust in terms of source/destination IP addresses, port ranges, and such to filter down to specific traffic. However, part of the purpose of DSCP marking is to avoid the use of complex ACL's for classification at every hop within the network. If additional granularity is required, it may be a better practice to instead differentiate the DSCP values used for traffic flows from different applications.

The next step is "step 1" from the MQC description above, define traffic classes with the class-map command.

```
2851-1(config)#class-map match-all VS-REPLAY-CLASS
2851-1(config-cmap)# match access-group name VS-REPLAY-ACL
2851-1(config-cmap)#class-map match-all VS-LIVE-CLASS
2851-1(config-cmap)# match access-group name VS-LIVE-ACL
```

This step essentially uses the ACL we created to define a class name associated to surveillance traffic. This may seem like an extra step in our very simplified example, however, the MQC is a very robust structure that has been optimized for flexibility in aggregating QoS configuration to reduce redundancy in the application of Qos configuration to devices. The class-map structure provides a great number of "match" condition options that can also increase the granularity of definition of the traffic in question. Instead of using an access list, you could also use a direct match statement such as "match ip dscp cs4" if you do not require additional granularity to identify the traffic.

```
2851-1(config)#class-map match-all VS-REPLAY-CLASS
2851-1(config-cmap)# match dscp af31
2851-1(config-cmap)#class-map match-all VS-LIVE-CLASS
2851-1(config-cmap)# match dscp cs4
```

Use the policy-map command to associate the traffic classes with one or more QoS features. In this case, we want to associate the class named “VS-LIVE-CLASS” with the QoS feature of a priority queue, and the “VS-REPLAY-CLASS with the features of bandwidth assignment and traffic shaping. The PQ and Burst values for a single 4CIF 30 frames per second stream were found in Table 9. The bandwidth reservation value is derived from 110% of the CBR target, which is 3000000 bps, found in Table 5. The traffic shaping value is set to be slightly higher than the bandwidth setting, at 115% of the CBR target.

```
2851-1(config)# policy-map HQ-WAN
2851-1(config-pmap)# class VS-LIVE-CLASS
2851-1(config-pmap-c)#  priority 4350 217500
2851-1(config-pmap-c)# class VS-REPLAY-CLASS
2851-1(config-pmap-c)#  bandwidth 3300
2851-1(config-pmap-c)#  shape average 3450000
```

Once the policy map has been created, we can attach the policy to an interface with the service-policy command. The service-policy must be applied in a specific direction, applied to inbound or outbound traffic. In this case, we want to apply the traffic to a WAN egress port in the outbound direction. This will prioritize the live surveillance media streams over other traffic utilizing the WAN link, while reserving bandwidth for replay of recorded streams on the same link.

```
2851-2(config)#interface multilink 1
2851-2(config-if)# service-policy output HQ-WAN
```

The resulting service policy status can be viewed on the interface using the show policy-map interface <interface id> command. This is a very important command for viewing the status of the service policy, showing the rate of traffic offered to the priority queue, and whether or not drops are occurring in the queue. If drops are occurring in the priority queue, it is an indication that either the queue is under-provisioned, or unintended traffic is hijacking a portion of the queue bandwidth. Statistics are also displayed for the class-based queue for recorded traffic, as well as the traffic shaping features which are enabled.

```
2851-1#show policy-map interface multilink 1
 Multilink1

  Service-policy output: HQ-WAN

    Class-map: VS-LIVE-CLASS (match-all)
      7135 packets, 10045771 bytes
      5 minute offered rate 3054000 bps, drop rate 0 bps
      Match: access-group name VS-LIVE-ACL
      Queueing
        Strict Priority
        Output Queue: Conversation 264
        Bandwidth 4350 (kbps) Burst 217500 (Bytes)
        (pkts matched/bytes matched) 4866/6706446
        (total drops/bytes drops) 0/0
    Class-map: VS-REPLAY-CLASS (match-all)
      0 packets, 0 bytes
      5 minute offered rate 0 bps, drop rate 0 bps
      Match: access-group name VS-REPLAY-ACL
      Queueing
        Output Queue: Conversation 265
        Bandwidth 3300 (kbps) Max Threshold 64 (packets)
        (pkts matched/bytes matched) 0/0
        (depth/total drops/no-buffer drops) 0/0/0
      Traffic Shaping
          Target/Average   Byte   Sustain   Excess    Interval  Increment
             Rate           Limit  bits/int  bits/int  (ms)      (bytes)
          3450000/3450000   20700  82800     82800     24        10350

        Adapt  Queue     Packets   Bytes     Packets   Bytes     Shaping
        Active Depth                         Delayed   Delayed   Active
        -      0         0         0         0         0         no

    Class-map: class-default (match-any)
      34 packets, 6074 bytes
      5 minute offered rate 1000 bps, drop rate 0 bps
      Match: any
2851-1#
```

The command output above illustrates the VS-LIVE-CLASS setup for a Strict Priority queue of 4350Kbps, and a Burst of 217500 Bytes. The VS-REPLAY-CLASS is configured with an assigned bandwidth of 3300Kbps, and traffic shaping with a target rate of 3450000 bps.

Alternate QoS Configuration for Dedicated Networks

When dedicated networks are implemented for surveillance traffic, a different approach is required to provide adequate throughput on lower speed WAN connections. It would not be the best use of resources to provision a 45 Mbps DS-3, or a 10 Mbps Metro Ethernet service, and then only use 33% of the link for multimedia traffic. However, a balance must be achieved between over-provisioning and under-provisioning WAN links for the required services they are to support. Cisco best practices recommend only 75% of link bandwidth be consumed by the requirements of Data, Voice, and Video applications on a WAN link. This recommendation still holds true, so when Data and Voice requirements are zero, Video should consume no more than 75% of the link on its own.

When configuring a dedicated WAN link for surveillance media, there is no less important traffic to de-prioritize, so it does not make sense to implement the “managed unfairness” of an LLQ/PQ configuration. Instead, a simple Weighted Fair Queuing configuration has proven to help allocate bandwidth between several streams running on a single shared link. Weighted Fair Queuing is aware of individual flows or “conversations” on a given link, and performs a scheduling approach designed to ensure all flows get a fair share of the network bandwidth.

Configuration Example for Weighted Fair Queuing

Weighted fair queuing (WFQ) can also be applied to the WAN egress interface using an MQC configuration as shown in the LLQ example. The difference is that instead of creating a named class, you use the built-in class “default-class” to apply the settings to all traffic in the policy. (All traffic not covered by other classes, which is all if no other classes are configured.) Following is an example of how to convert the LLQ Policy Map from the previous configuration, to a simple WFQ default configuration, and also an example of the resulting “show queuing interface” command output.

```
2851-2(config)#policy-map BRANCH-WAN-EDGE
2851-2(config-pmap)# class class-default
2851-2(config-pmap-c)#  fair-queue
2851-2#show queueing interface multilink 1
Interface Multilink1 queueing strategy: fair
  Input queue: 0/75/0/0 (size/max/drops/flushes); Total output drops: 848
  Queueing strategy: Class-based queueing
  Output queue: 1/1000/64/0 (size/max total/threshold/drops)
     Conversations  1/2/256 (active/max active/max total)
     Reserved Conversations 0/0 (allocated/max allocated)
     Available Bandwidth 8192 kilobits/sec

  (depth/weight/total drops/no-buffer drops/interleaves) 1/6476/0/0/0
  Conversation 88, linktype: ip, length: 1089
  source: 10.1.27.12, destination: 239.255.27.12, id: 0x5A46, ttl: 4,
  prot: 17
```

Weighted Fair Queuing may also be applied to an interface by simply using the “fair-queue” interface command as opposed to applying a service-policy using the MQC. However, MQC creates a more modular configuration, and also exposes additional management capability such as the Class Based QoS SNMP MIB. Use of the MQC is considered the best practice for application of QoS policies on IOS router platforms with software-based QoS configurations.

Chapter 7: Storage and Retrieval of Surveillance Media

Cisco Video Surveillance solutions include the capability to store video and audio surveillance media on disk drives for future review. This capability is provided through Cisco Video Surveillance Services Platforms and Integrated Services Platforms, which provide Network Digital Video Recorder (NVR) features to the overall solution. These devices participate in the same multicast discovery architecture as Cisco Video Surveillance IP Gateways, are configured through Cisco Stream Manager Configuration Module, and are managed through Cisco Stream Manager Administration and Monitoring Module.

Storage Functional Model

Many products and options make up the overall Cisco Video Surveillance solution. This approach provides flexibility in running over various network topologies and in integrating with existing matrix switches and third-party systems. To understand the functions of the various storage-related products in the solution, it is helpful to understand the conceptual storage model.

The primary functional areas of this process are:

Encoding

Even when the source video does not need to traverse the IP network to get to a recorder, it still needs to be encoded in MPEG-4 to be stored on disk. This process converts the analog NTSC or PAL stream to frames, each of which contains a digital array of pixels. The codec further compresses the digital information, which reduces the overall disk consumption of the streams.

Network Video Recorder (NVR)

The Network Video Recorder (NVR) function is the set of software processes that receives the encoded streams and writes them to disk. Review of stored video is provided over the IP network through the NVR. The NVR also manages the available disk resources and grooms from disk video samples that are outside of the retention period. The Cisco Video Surveillance solution uses various models of Cisco Services Platforms as the NVR.

Disk Storage

The digitally encoded video samples are written to disk storage by the NVR. Disk drives are located in the Services Platform (SP) and Integrated Services Platform (ISP). The ISP may be ordered in various configurations, including a JBOD and RAID5 array. Cisco has tested and deployed configurations using external disk arrays from Nexsan Technologies, such as the ATABoy and SATABeast products. Currently, only disk drives located in the SP and ISP chassis are offered directly from Cisco Systems, Inc.

Note: The acronym JBOD originated from the phrase “just a bunch of disks” and refers to a group of hard drives concatenated into a single logical volume with no redundancy. The term RAID5 refers to a level of redundancy within the “redundant array of inexpensive disks” hierarchy. RAID5 involves striping data across the disk array while devoting one hard disk worth of space for storing redundancy information. This approach provides fault tolerance for an array so that it can lose any single hard drive and continue to operate. When the failed drive is replaced, data can be reconstructed from the stored redundancy information to recreate the failed drive's data.

Retrieval

Surveillance media that has been stored to disk may be retrieved and viewed through a hardware or software decoder. A simple example is using a CCTV keyboard and monitor that are connected to a Cisco Video Surveillance IP Gateway, pressing the instant-replay key, and using the joystick to view recently stored video. Alternately, one can use the Cisco Stream Manager Client Viewing Module. This module provides the Browse feature to locate stored video by date and time and the Activity Search feature, which can be based on a user-defined portion of the display. Export of stored video to a portable file for convenient transport is also supported.

Connectivity Functions

Each of the conceptual blocks of functionality that are described below must communicate to another block to form a complete solution. The transport media involved varies depending on the products used. IP networking is the most pervasive transport and will provide an option to cover all of these connectivity requirements. USB connectivity also is an option from the encoding function to the NVR. Disk drive connectivity is offered using SCSI or fiber channel interfaces, and iSCSI connectivity is a future possibility. In a platform such as the ISP, a single chassis integrates much of this connectivity, which offers a simple and efficient solution.

Figure 45 illustrates the functional blocks of the Cisco Video Surveillance storage solution. The roles of various Cisco Video Surveillance products are shown as they relate to the model.

Figure 45. Storage Functional Model

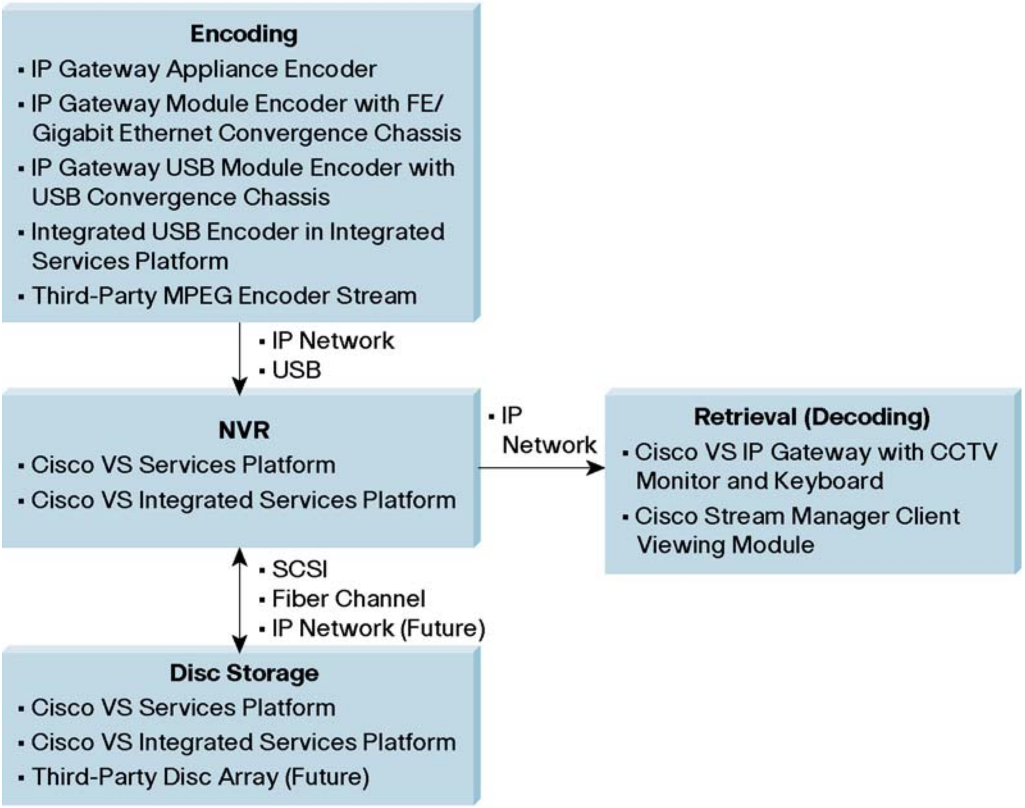

Encoding
- IP Gateway Appliance Encoder
- IP Gateway Module Encoder with FE/ Gigabit Ethernet Convergence Chassis
- IP Gateway USB Module Encoder with USB Convergence Chassis
- Integrated USB Encoder in Integrated Services Platform
- Third-Party MPEG Encoder Stream

- IP Network
- USB

NVR
- Cisco VS Services Platform
- Cisco VS Integrated Services Platform

- IP Network

Retrieval (Decoding)
- Cisco VS IP Gateway with CCTV Monitor and Keyboard
- Cisco Stream Manager Client Viewing Module

- SCSI
- Fiber Channel
- IP Network (Future)

Disc Storage
- Cisco VS Services Platform
- Cisco VS Integrated Services Platform
- Third-Party Disc Array (Future)

Using USB Connectivity from Encoding to NVR

As shown in Figure 45, several Cisco Video Surveillance products can perform the MPEG-4 encoding function. The USB Convergence Chassis, populated with four-port IP Gateway modules, provides a high density (up to 64 analog ports) encoding solution that can connect directly to an SP or ISP with USB cables. The video streams must be digitally encoded before the NVR processes in the SP or ISP can store them to disk. The ISP is differentiated by the ability to house up to three of the four-port USB module encoders within the chassis, which can terminate up to twelve direct analog ports.

Using IP Connectivity from Encoding to NVR

The ISP also provides storage for live surveillance media streams that are received across the IP network. Fast Ethernet or Gigabit Ethernet based Convergence Chassis house virtual ports that may be created on the system through the Cisco Stream Manager Configuration Module. A virtual port provides the definitions necessary for the SP or ISP to receive the stream across the IP network. Figure 46 shows an example of virtual port configuration on the ISP. The device pictured at the top of the screen represents the actual back of the ISP and the device that resembles card slots with RJ-45 ports is a representation of the virtual ports that have been added to the system.

Figure 46. Virtual Port Configuration on ISP

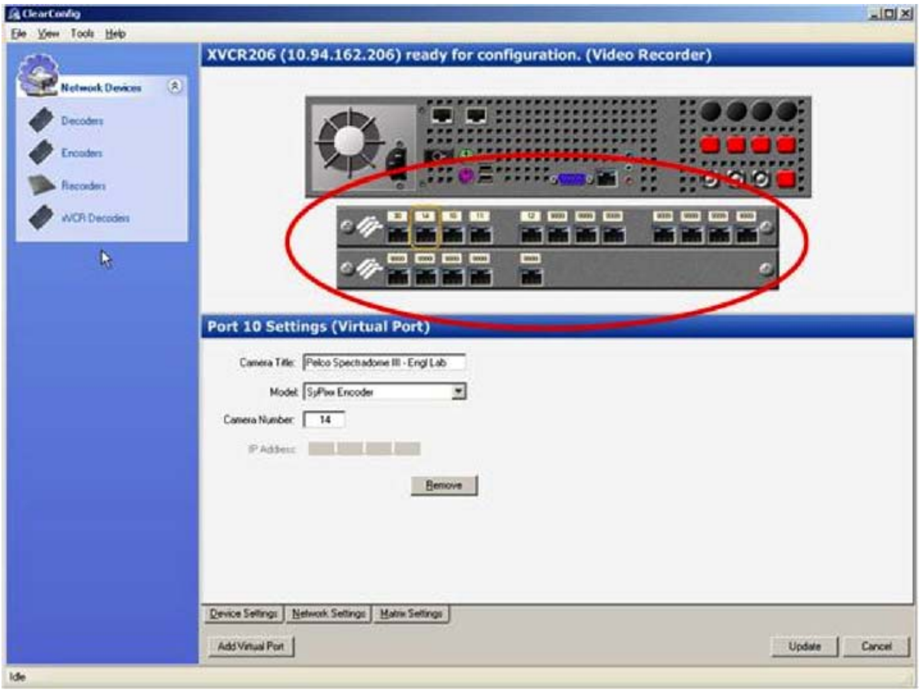

XVCR206 (10.94.162.206) ready for configuration. (Video Recorder)

Port 10 Settings (Virtual Port)

Camera Title: Pelco Spectradome III - Engl Lab

Model: SyPixx Encoder

Camera Number: 14

IP Address:

Planning Storage in a WAN Deployment

The low speed links in a typical WAN environment can significantly limit flexibility in the placement of the NVR function. As with other IP applications, when designing an IP Video Surveillance system to operate over a WAN, there is a financial trade-off between a centralized and a distributed approach. Centralization reduces equipment cost and operational maintenance costs, but a distributed approach reduces the recurring cost of carrier bandwidth required for WAN services. Cisco Video Surveillance products are flexible and can work with either approach, within the constraints of adequate bandwidth provisioning and proper QoS.

Provisioning Guidelines for WAN Storage Deployments

When designing the placement of the NVR function in a WAN video surveillance deployment, follow these guidelines:

- Consider the NVR as a potential video source and as a video destination. Provisioning a WAN topology to support occasional video review of a single stream can be more cost-effective than provisioning bandwidth for concurrent recording of all streams.
- Consider the tradeoff between higher and lower resolution and frame rates, and their associated costs. Consider the affect on WAN bandwidth requirements and storage requirements.
- Consider using dual-streaming, which is available on single-port Cisco Video Surveillance IP Gateways, to locally view live video at a high resolution and frame rate, while recording over the network at a lower resolution and/or frame rate. (Although dual-streaming is a powerful feature, single-port encoders do not provide the density or price point of the USB Convergence Chassis solution.)
- Use the bandwidth provisioning and QoS recommendations in the “Traffic Engineering” section to ensure adequate network performance for all surveillance media streams in the system.

Centralized Storage Deployment

There are many advantages of a centralized storage deployment. Equipment can be located in a single, environmentally controlled facility where it is in close proximity to the technical staff that provides support. Fewer, larger capacity systems can be used, saving the costs of managing multiple pieces of equipment. Video surveillance review can be performed centrally for any location across the WAN. A disadvantage of a centralized storage deployment is the cost required to provision bandwidth on the WAN links to support one or more video streams. As the cost per MB of WAN services drops and the efficiency of video codec technology improves, the WAN circuit costs of a centralized approach may become within reach for more organizations. Organizations such as municipalities or transportation organizations may have physical topologies that are geographically dispersed but have dedicated fiber or other high-bandwidth transport. These organizations may be excellent candidates for a more centralized approach.

Distributed Storage Deployment

Distributed deployment of storage equipment significantly reduces the bandwidth requirements on a WAN. When the security staff also is distributed, this approach may improve the process of video surveillance review. Multiple smaller systems allow a deployment to scale, and this approach creates a resilient network without a large central point of failure. A disadvantage of the distributed storage approach is that the cost per hour of storage for several small systems may be more than the cost for fewer large, high-capacity systems. The cost of maintaining more devices in the overall network also has an affect.

A storage deployment does not have to take only a centralized or a distributed approach. For example, some consolidation of sites where adequate bandwidth is readily available may be a good approach, without requiring overall centralization. Ultimately, considerations of who needs to access the video from where on the network must be balanced with the financial constraints of carrier-provided bandwidth and with the actual user requirements of the video. For example, if the video is to be used in legal proceedings, it may be necessary to record at a high video resolution for the video evidence to be admissible in court.

Storage Capacity Planning

Disk space consumption for storage is closely related to network bandwidth consumption. Both are measures of the amount of data that is required to represent video streams in a digital format over time. The CBR target setting defined on the encoder controls the amount of bits per second used by the codec to represent the stream. For an estimate of the number of bytes that storing one second of that stream will consume on disk, convert the CBR target into bytes by dividing by 8.

Calculating Storage Capacity Requirements

Calculating storage capacity is a basic math problem. The primary factors are the duration of the video to be stored multiplied by the rate at which the disk space is consumed. This result is calculated for each stream and then is summed for all streams to be stored on a common disk array.

This calculation results in an overall capacity requirement for the desired capability of the system. In some cases, this calculation may indicate more disk storage is required than is appropriate for the budget of an installation. In that case, consider the following options, which can reduce disk consumption.

- Resolution and/or frame rate of the stored stream can be reduced, which reduces the storage capacity required.
- CBR target bit rate can be reduced, which reduces the storage capacity required. The target bit rate guidelines in Table 5 are optimized for each combination of resolution and frame rate. In many cases, a lower target rate can be configured and reasonable quality achieved, but aggressive compression can compromise quality when there is high complexity or changes in the source video stream.
- The retention time for a given stream on an ISP or USB Convergence Chassis can be reduced, which frees disk capacity for storing new video streams.
- Dual-streaming (available from single-port Cisco VS IP Gateways only) allows live viewing of a higher quality stream, while a lower resolution and frame rate may be chosen for the stream directed to the NVR.
- Port Activity Detection can be used on an ISP or USB Convergence Chassis to flag for retention only video samples that contain motion in a specified area of the field of view.

Port Activity Detection

Port Activity Detection is a powerful tool to optimize limited disk space resources. While continuous video is streamed to the NVR from a given camera, Port Activity Detection allows retention of stored video to be governed by the presence of activity or change in the display received by a port. This setting can be tuned to a specific region of the display, such as a doorway or a critical item. It can also be tuned for sensitivity and should be tested upon configuration to ensure that the desired result is achieved. The NVR can also be configured to save video for a time after detecting motion. This video is then protected as the NVR processes groom the disk, while static video with no activity may be discarded and the space reclaimed.

Note: Port Activity Detection and Retention settings are only available for cameras that terminate directly on a Cisco Integrated Services Platform or Cisco USB Convergence Chassis. On these platforms, the encoding process is occurring locally to the SP or ISP. IP based streams which terminate on Virtual Ports do not support Port Activity Detection or configurable Retention.

Figure 47 shows an example of the Port Activity setup from the Cisco Stream Manager Configuration Module.

Figure 47. Port Activity Detection Configuration

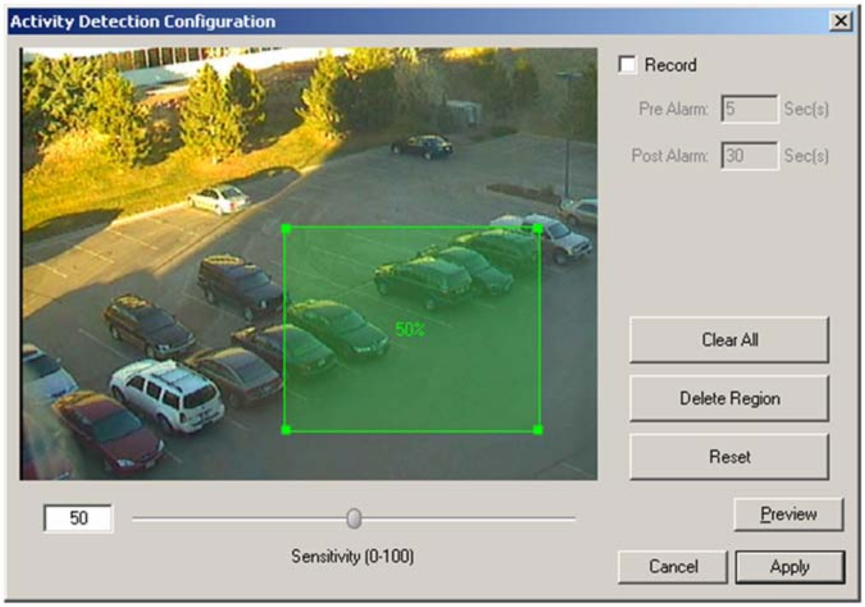

Activity Detection Configuration

☐ Record

Pre Alarm: 5 Sec(s)

Post Alarm: 30 Sec(s)

Sensitivity (0-100): 50

The Port Activity Detection Configuration dialog is launched from the Port Activity Detection tab on the Port Setting screen in the Cisco Stream Manager Configuration Module.

Chapter 8: Traffic and Storage Engineering Case Study

Designing an IP video surveillance network requires careful capacity planning for bandwidth use and storage disk space. The following case study of traffic engineering and storage capacity planning for a hypothetical system illustrates the application of information provided in this document.

Overview

Consider the IP video surveillance requirements of Company X. This company has existing analog cameras that are widely deployed throughout its facilities. The company has decided to move its surveillance networking to a converged IP network, using a virtual matrix switch model, and implement new storage equipment providing instant replay capabilities directly from their existing CCTV monitors and keyboards.

Company X's requirements include:

- 48 Cameras at the main campus, currently wired with coaxial cabling terminating at a centralized matrix switch.
- Two remote sites connected to the main campus via IP WAN links, each with 16 surveillance cameras.
- Required resolution of 2CIF, and frame rate of 7.5 frames per second for all cameras.
- Video surveillance recording equipment will be deployed in each of the three locations (main campus and two remote sites).
- A security operations center at the main campus must be able to view any main campus camera live or recorded.
- The main campus security operations center must be able to review at least one recorded stream from each remote site.

Figure 48. Company X Network Topology

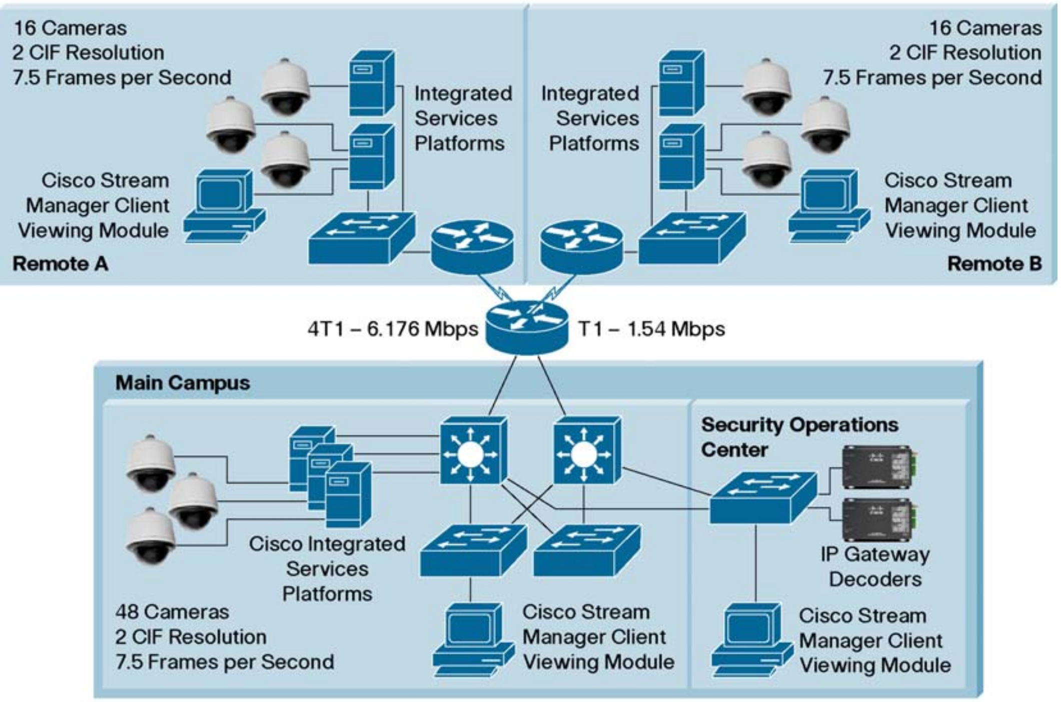
16 Cameras
2 CIF Resolution
7.5 Frames per Second
Integrated Services Platforms
Cisco Stream Manager Client Viewing Module
Remote A
16 Cameras
2 CIF Resolution
7.5 Frames per Second
Integrated Services Platforms
Cisco Stream Manager Client Viewing Module
Remote B
4T1 – 6.176 Mbps
T1 – 1.54 Mbps
Main Campus
Cisco Integrated Services Platforms
48 Cameras
2 CIF Resolution
7.5 Frames per Second
Cisco Stream Manager Client Viewing Module
Security Operations Center
IP Gateway Decoders
Cisco Stream Manager Client Viewing Module

IP Networking Requirements

For the campus environment, the first step is to calculate the bandwidth consumption for an individual camera at each level of resolution. The standard formula for this calculation is:

Bandwidth = Image Size * Rate * Normalization

Using the MPEG-4 codec, we are dealing with video frames instead of standalone images, and the size of a given video frame varies widely between I-frames and predictive frames. Therefore, the image size is not a consistent value and will vary widely during the stream. To come up with bandwidth consumption estimates for MPEG-4, we can use the CBR rate set on the encoder, which is already expressed in bits per second.

We begin with the values from Table 5. To convert these numbers from raw codec targets to anticipated IP network bandwidth, multiply the values by 1.1 (110%) to account for IP network packet header overhead.

2CIF resolution, 7.5 frames per second ~ 750 kbps * 1.1 = 825 kbps

Main Campus Bandwidth Requirements

The next step is to determine the total bandwidth that all main campus cameras will generate on the network. We will use a normalization ratio in the formula of 1 Mbps / 1000 kbps to convert the resulting number into Mbps.

48 2CIF streams * 825 kbps per camera * 1Mbps/1000 kbps= 39.6 Mbps

Gigabit Ethernet runs at 1000 Mbps on all of the connections between switches in the Company X network. At a total video bandwidth load of 39.6 Mbps, we should not be concerned with saturating the network with the required video traffic. QoS configuration best practices should be followed to prioritize live video surveillance traffic on the campus network.

Remote Site Bandwidth Requirements

Next, consider the two remote sites. Company X has requested the capability to view one recorded stream from each remote site over a WAN link from storage at 2CIF/7.5 fps. Remote site A is connected to the main campus with four T-1 circuits multiplexed with Multilink Point to Point Protocol (MLPPP), for an aggregate bandwidth of 6.176 Mbps. Remote site B is connected to the main campus with a single T-1 circuit running at 1.544 Mbps. Since we are dealing with low speed links compared to the Gigabit Ethernet of the campus, we need to carefully consider bandwidth consumption and QoS configuration.

Cisco design guidelines for QoS on a converged WAN link using Low Latency Queuing recommend that no more than 33% of the link be allocated to a priority queue. In this case, we will be using a Class-Based Queue with a bandwidth assignment for recorded traffic, but we still need to be concerned with the effect of the intermittent recorded flow on other applications that share the link. When viewing a recorded stream, a portion of the WAN link will be consumed completely by forwarding the surveillance traffic. Since the recorded stream will not always be present on the link, we need to consider the varying effect of the presence of the recorded video stream on the performance of other applications sharing the link, similar to the provisioning of a priority queue.

Remote Site A: 33% of 6.176 Mbps = 2.038 Mbps (or 2,038 kbps)

Remote Site B: 33% of 1.544 Mbps = .509 Mbps (or 509 kbps)

Because we are dealing with a recorded stream, our QoS configuration on this link will consist of a CBWFQ that is traffic- shaped to 115% of the CBR target rate. This configuration provides space for the IP and TCP packet header overhead that is required to transport the video across the network. For a single 2CIF 7.5 frames per second stream with a configured CBR rate of 750:

750000 bits per second CBR * 1.15 * 1 KB/1000 bits = 862.5 kbps

We compare this estimated bandwidth requirement of a single recorded stream to the available bandwidth to determine if more than 33% of the capacity of existing WAN links is consumed:

Remote Site A: 862.5 kbps is less than 2038 kbps, which is OK.

Remote Site B: 862.5 kbps is more than 509 kbps, which indicates a provisioning problem.

The bandwidth requirement of remote site A fits well within the 33% guideline for their converged WAN link. The bandwidth requirements of remote site B are much greater than the 33% guideline, which creates a bandwidth-provisioning problem. This situation can be resolved in these ways:

- **Decrease the resolution and/or frames per second.** We could choose to throttle the resolution or frame rate of the recorded streams down at remote site B, changing the requirements to fit within the network bandwidth.
- **Decrease the configured CBR target rate.** It may be possible to achieve the required quality level at 2CIF 7.5 fps by reducing the CBR target. Video samples of the required field of view could be assessed at a lower setting to determine if the desired quality is achieved.
- **Increase the bandwidth on the WAN link to remote site B.** To determine a total provisioned link size that is three times the size of the required bandwidth, use this calculation: 862.5 kbps * 3 = 2587.5 kbps. We could surpass this number by increasing the bandwidth to remote site B with one additional T-1, for a total of two T-1s multiplexed with MLPPP and 3088 kbps of bandwidth.

When implementing IP video surveillance on a converged network with other applications running, WAN link provisioning often must be increased to meet the application requirements. In the simplified example of Company X, we did not address other real-time applications, such as live video surveillance streams, video conferencing, and IP telephony, which may require priority queuing over the WAN. Network provisioning must always take into account the requirements of all applications that need to share the network.

Storage Disk Capacity Requirements

Company X has decided to place Cisco Video Surveillance Integrated Services Platforms as Network Video Recorders (NVRs) at the main campus and at each of the remote sites. Based on the bandwidth analysis in the previous section this approach is a wise choice; fully centralized storage would not be an option considering the bandwidth limitations of the existing WAN infrastructure. We can run a standard formula to calculate the amount of disk space that is required for each site. This calculation is similar to the bandwidth requirement, with these differences.

- The numbers for bandwidth should take into account the overhead of the IP network layers, IP, User Datagram Protocol, and RTP headers. The numbers for storage do not, because these layers are stripped by the NVR as the payload of the packets is reassembled into video frames and stored to disk. The basic CBR target numbers should be used to estimate storage capacity requirements.
- Retention time should be considered in the storage formula. How much video needs to be kept in case it is required for review? Increasing retention affects the total amount of storage space that is required.
- Consider the activity level if Port Activity Detection (or motion detection) is configured for storage. This level may be expressed in terms of a percentage of overall time, based on the duration that the NVR keeps specific event, and the predicted number of events in a given period.

Note: Port Activity Detection and Retention settings are only available for cameras that terminate directly on a Cisco Integrated Services Platform or Cisco USB Convergence Chassis. On these platforms, the encoding process is occurring on modular encoders which are either directly installed in the ISP, or in a chassis which is USB attached to the SP. IP based streams which terminate on Virtual Ports do not support Port Activity Detection or configurable Retention.

The overall standard storage calculation formula is the following:

Storage = Image Size * Rate * Retention * Activity * Normalization

For purposes of our case study, let’s break this formula into steps. First, calculate storage requirements where Port Activity Detection is not configured. Using the CBR target rate, we save a step in the formula, and the result is same number as the product of image size times rate in the formula above. Because disk capacity is commonly expressed in bytes, convert this number from bits per second to bytes per second using a normalization factor of dividing by 8. Company X has a policy of retaining surveillance video for 10 days.

2CIF resolution, 7.5 fps ~ 750 kbps

750 kbps * (1 Byte / 8 bits) = 93.75 Kilobytes per second

93.75 KBps * 10 days * 86400 seconds/day = 81,000,000 KB, or 81 GB of storage per stream

To calculate the total storage requirement for each site based on the number of streams per location:

Main campus, 48 streams * 81 GB = 3888 GB (3.9 Terabytes)

Remote site A, 16 streams * 81 GB = 1296 GB (1.3 Terabytes)

Remote site B, 16 streams * 81 GB = 1296 GB (1.3 Terabytes)

From this step of the calculation, it is clear that Port Activity Detection is an attractive option. The Cisco Video Surveillance Service Platform and Integrated Service Platform provide configuration of Port Activity Detection with a pre-alarm and post-alarm window, which allows retention of video only when motion is detected. Using a pre-alarm setting of 3 seconds, and a post-alarm setting of 5 seconds, we have an event time window of 8 seconds. In a real-world example, the estimated number of events could be set differently for each stream, but in our case study, we will assume that all cameras have a consistent expected event level of 3500 events per day. So to calculate an activity percentage:

Event duration seconds * events per day = recorded seconds per day

Recorded duration seconds / total time in one day = Activity Percentage

8 seconds * 3500 events per day = 28000 seconds

28000 seconds / 86400 seconds per day = .324 = 32.4% Activity

Now we can apply the activity percentage to our original storage requirements per site to calculate actual storage requirements per site with expected activity levels:

Main campus, 3888 GB * .324 = 1259.71 GB (1.3 Terabytes)

Remote site A, 1296 GB * .324 = 419.9 GB

Remote site B, 1296 GB * .324 = 419.9 GB

Company X now has a reasonable baseline number to predict storage requirements for their new installation. Monitoring capacity and verifying assumptions is an important part of operating the system. Most systems tend to grow, rather than shrink, and an under-provisioned system does not meet requirements, so being conservative when estimating disk capacity can pay off in the long run.

Appendix A: Glossary

A

Alarm Panel

A device that aggregates inputs from multiple sensors, monitors their signals, and translates them into actionable events. Such events may include audible or visual alerts in a facility as well as sending an alert to a monitoring center indicating the event’s status.

Alert

A message sent to security personnel indicating the location and nature of an emergency or threat.

Aperture

The opening of a lens that controls the amount of light reaching the surface of the pickup device. The size of the aperture is controlled by the iris adjustment.

Attenuation

A decrease or loss of signal. Within a fiber or coaxial-cabled surveillance system, this causes degradation in the video image (e.g. jitter, noise, loss of signal).

Auto White Balance

A feature on color cameras that constantly monitors the light and adjusts its color to maintain white areas.

B

B Frames

B Frames can be found in MPEG video streams. They are not full frames and predicted by both the previous frame(s) and the next frames. Since they are not full frames, they use the least amount of space to store the frame data.

Back Light Compensation (BLC)

A feature on cameras that electronically compensates for high background lighting to give detail which would normally be silhouetted.

Background Noise

The total system noise independent of the presence of absence of a signal. The signal itself is not included in the measure of a system’s noise.

Biometrics

The identification of a user based on a physical characteristic, such as a fingerprint, iris, face, voice or handwriting.

Brightness

(aka, Luminance).The visual perception of an area reflecting or emitting light.

BPDU

Bridge Protocol Data Unit. Spanning-Tree Protocol hello packet that is sent out at configurable intervals to exchange information among bridges in the network.

C

Camera

An optical device capable of viewing a given area and translating that view into an electronic signal.

Camera Housing (aka, Camera Enclosure)

The hardware mount encasing a camera generally used to protect it from harsh environmental conditions and vandalism.

Central Station

A remote location that is designed to monitor signals from physical security systems.

Channel

A single video signal.

Closed-Circuit Television (CCTV)

A television system in which signals are distributed via cables to a closed network of monitors. This system is most often used for security surveillance in small, closed areas like buildings or parking garages.

C-mount/CS-mount

Lenses are available in two different mounts: C-mount and CS-mount. C-mount lenses have a flange back distance of 17.5mm and CS-mount lenses have 12.5mm. Many cameras can accept either type of lens, but it is important to make sure that camera and lens are compatible and set up properly. C-mount lenses can be used on CS-mount cameras by utilizing a 5mm adapter or adjusting the camera for C-mount lenses. Because of the shorter back focal distance, CS-mount lenses can only be used on CS-mount cameras.

Coaxial Cable

(aka, Coax). A type of cable that is capable of passing a range of frequencies with low loss. It consists of a hollow metallic shield in which one or more center conductors are put in place and isolated from one another and from the shield.

Codec

Coder-decoder. Refers to the technique used to translate an analog signal into a digital format, and vice-versa.

Common Intermediate Format (CIF)

The term CIF is used to mean specific video resolution: 352x288 in PAL 352x240 in NTSC. CIF is 1/4th of "full resolution" TV, also called D1

Composite Video

A type of cable capable of passing a range of frequencies with low loss. It consists of a hollow metallic shield in which one or more center conductors are put in place and isolated from one another and from the shield.

Compression

Compressing digital data means removing redundancies so that it takes up less space. There are two main forms of compression, lossy and lossless. Lossless compression only takes away a certain amount of data so that it can be returned to its original complete state. Lossy compression however, will sacrifice more data to produce better compression. MPEG-4 is a lossy compression that can keep high quality but can reduce the amount of space a video file needs tremendously.

Console (CCTV)

The part of a monitoring station an operator uses to control surveillance cameras. Usually consists of a joystick for PTZ control and a set of numbered buttons allowing the operator to switch cameras displayed on an attached monitor. It may also refer to the entire structure at a monitoring station that houses the keyboards, joysticks, monitors, phones, etc. for controlling the physical security system.

Contrast

The ratio of light to dark portions of a video image.

Crosstalk

The transfer of signals between adjacent systems. Often causes interference that degrades the signal's output making it difficult to read or jamming the signal all together.

D

Day and Night

Refers to a video camera's ability to provide images in both lighted and dark conditions by changing the imaging format from color to black-and-white, respectively.

Decoder

A hardware or software device that employs a codec to translate a signal from its digital form into an analog output for display on a monitor.

Depth of Field

The distance between two objects, front to back, which is in focus in a televised scene. With a greater depth of field, more of the scene, near to far, is in focus.

Digital PTZ

(aka, ePTZ). The capability to virtually pan-tilt-zoom within a digital image. The feature does not require the ability to mechanically move a camera or its focus. Currently an emerging feature of megapixel cameras.

Digital Video Recorder (DVR)

Digital Video Recorder is the industry standard term applied to PC-based or embedded systems that encode and record video images to a computer hard drive. DVRs provide a quicker method of retrieving the recorded information unlike media such as VHS tapes and other equipment that stores information in a sequential manner. DVRs are often integrated into enterprise networks through a single Ethernet interface yet they terminate multiple analog cameras, typically four, eight or sixteen. (See also **Network Video Recorder**.)

Digital Watermark A digital watermark is an embedded identifying mark that cannot be removed from a digital document. It contains hidden identification data and is often used to provide surveillance video image integrity for forensic purposes.

Dome Camera A video imaging device contained within a demisphere. Generally supports the ability to change its focus (i.e. camera PTZ inside the dome) within the field-of-view allowable by the dome itself.

Dwell Time The length of time a video switcher holds on a camera before moving on to the next in sequence.

E

Encoder A hardware or software device that employs a codec to translate an analog video signal into a digital form.

EIGRP Enhanced Interior Gateway Routing Protocol. Advanced version of IGRP developed by Cisco. Provides superior convergence properties and operating efficiency, and combines the advantages of link state protocols with those of distance vector protocols.

F

Field The image produced by a single vertical sweep of a camera or monitor. One field contains 262.5 lines in NTSC systems; consequently, two fields form a single image or frame. (See also **Interlacing**.)

Field of View (FOV) A camera's area of focus (i.e. what it can see).

Fixed Iris Lens A lens in which the aperture is manually adjusted to maintain proper light levels on the faceplate of the camera pickup device.

Format (Analog CCTV) An image's size produced by a camera or displayed on a monitor.

Format (Digital CCTV or IPVS) The resolution and codec used by a camera or PC monitor to display an image.

Frame The total area of the picture that is scanned. With interlaced video, the frame is comprised of two fields.

Frame Rate See **Frames Per Second**.

Frames Per Second (FPS) A measure of a camera's rate of output of single snapshots. Also known as images per second and frame rate.

H

Horizontal Resolution The maximum number of individual picture elements that can be distinguished in a single scanning line.

I

I Frame Unlike B Frames, I Frames do not depend on the previous or next frames to predict a full image. An I Frame is basically a full image. In compressed video there are only a limited number of I Frames to keep the size of the video files down. Since it is a Full Frame, an I Frame uses more space than a B Frame.

Image Size (Lenses) Reference to the size of an image formed by the lens onto the camera pickup device. The current standards are: 1", 2/3”, 1/2”, 1/3” and 1/4” measured diagonally.

Infrared Camera Refers to a camera’s ability to provide enhanced visual images during extremely low-light conditions (e.g. at night) using infrared technology.

Infrared Illuminator A device that emits infrared light in order to enhance an IR camera’s ability to detect objects under low light conditions.

Interlacing The process of joining two separate video fields to create a single frame.

Interleaving A method used in alarms or activity detection that allows extra frames of video from alarmed cameras to be added to a time multiplexed sequence while a state of alarm exists. This is specific to analog CCTV.

IP or Network Camera A video imaging device that natively attaches to an Ethernet network and delivers its images in IP packets. It differs from its analog equivalents in that it does not require an external encoder to translate the video into a digital signal nor to attach to the IP network.

IP Video Surveillance (IPVS) Refers to the system or process of monitoring an area by using an IP network as the transport for remote video signals. The components of an IPVS system include edge devices such as IP cameras, IP encoders, or DVRs; an IP network for transport; recording devices such as NVRs; monitoring stations including legacy monitors and consoles served through decoders or PCs running monitoring software; and management software for configuration and maintenance.

Iris A camera’s eye. An adjustable opening that controls the amount of light entering a camera from its lens projected onto the camera’s imager.

J

Jitter Jitter is a measure of the variability over time of the latency across a network. A very low amount of jitter is important for real-time applications using voice and video.

Joystick The part of a surveillance system console that allows an operator to steer a camera into different positions.

K

Keypad A device that provides a user interface to control a security system or subsystem. Typically includes a numerical 10-key touchpad to allow entering of passcodes and commands. See also **Console**.

L

Lag The image retention of an object after the object has been scanned. Sometimes, it causes image smearing.

Lens Front most portion of a camera’s optics that allows it to focus on an area of interest. In many applications, a camera’s lens is interchangeable (See **C-Mount Lens, Auto Iris Lens,** and **Fixed Iris Lens**.)

Level Control Main iris control. Used to set the auto-iris circuit to a video level desired by the user. After setup, the circuit will adjust the iris to maintain this video level in changing lighting conditions. Turning the control toward High will open the iris; toward Low will close the iris.

M

Manual Iris Lens A lens with a manual adjustment to set the iris opening (F stop) in a fixed position. Generally used for fixed lighting applications. (See also **Fixed Iris Lens**.)

Matrix Switch

A video signal device able to route any of its inputs (i.e. cameras) to any of its outputs (i.e. Monitors and recorders). Through a matrix switch, the relation of inputs to outputs is a one-to-one connection unless a looping device is introduced. The actual number of inputs to outputs is generally not one-to-one. Inputs usually exceed the number of outputs available. Matrix switches are usually located at a security operations center, where all video concentrates and displays on multiple monitors. Users control the matrix via a joystick and keyboard that allows switching and the remote control of pan-tilt-zoom cameras.

Mega-Pixel Camera

An IP camera capable of providing extremely detailed image resolution (on the order of HDTV quality). Mega-pixel loosely refers to a single image as containing multi-million pixels.

MJPEG

Motion JPEG. A video codec standard that describes a technique whose frames are digitized independently using the JPEG format. Each frame is replayed in sequence to produce the movement. The encoding provides bandwidth efficiency (as best as possible) through JPEG's compression of each image.

Monitor

A CRT used to display live and recorded analog video.

Monitoring

The sending of alarm, trouble, and other signals to a remote location such as a security operations center.

Motion Detection (Video)

The process of analyzing a camera's video signal to determine if there is any movement (pixel changes) in the picture and then subsequently trigger an alarm.

Motion Detector

A device designed to detect movement within the premises. There are multiple types: active and passive.

Moving Picture Experts Group (MPEG)

A set of international standards (ISO) describing a video encoding method that utilizes mathematical prediction techniques to represent video as a series of a reference frame (full images) and set of subsequent changes to the image (predictive frames). Bandwidth efficiency is achieved through the transmission of the predictive frames. The most common MPEG standards used in video surveillance include MPEG-2 and MPEG-4.

Multilink PPP

Method of splitting, recombining, and sequencing datagrams across multiple logical data links.

N

Network Video Recorder (NVR)

A PC or network appliance running special software used to capture and store images emanating from IP cameras and encoders. An NVR differs from a DVR in that it provides no encoding of analog video signals; in other words, it has no video inputs. Typically the NVR acquires video by attaching to the source over an IP network. (See also **Digital Video Recorder**.)

NTSC (National Television Systems Committee)

A committee that worked with the FCC in formulating the standards for the United States color television system. NTSC specifies a resolution of 480 lines at 30 frames per second. (See also **PAL**.)

O

OSPF

Open Shortest Path First. Link-state, hierarchical IGP routing algorithm proposed as a successor to RIP in the Internet community. OSPF features include least-cost routing, multipath routing, and load balancing.

P

Passive Sensor

An intrusion detection device that does not use a transmitter to produce a signal but instead simply detects energy emitted near the sensor.

Phase Alternating Line (PAL)

The European standard for video resolution that specifies 625 lines/frame at 25 frames/second. (See also **NTSC**.)

Physical Security

The use of personnel, equipment, and procedures to control the access to a facility and its assets.

Pinhole Lens

A lens used for applications where the camera must be hidden. The front of the lens has a small opening to allow the lens to view an area through a small hole in a wall or ceiling.

PTZ (Pan-tilt-zoom)

Describes the capability to change a camera’s field of view through three planes of reference. Panning refers to physically sweeping a camera from side-to-side (xy-plane) whereas tilting is the ability to move it up-and-down (azimuth). Zooming changes a camera’s lens magnification giving the visual effect that the point-of-focus is closer or further away.

R

Resolution

A measure of the ability of a camera, encoder or video system to reproduce detail. In analog systems, resolution usually refers to the number of lines that make up an image. Whereas with digital systems, resolution gives a measure of the number of pixels used to generate the image.

S

Saturation

The vividness of color.

Security Operations Center (SOC)

The command center where security personnel monitor and respond to security and safety related incidents.

Signal to Noise Ratio (S/N)

The ratio between a useful video signal and unwanted noise.

U

UTP

Unshielded twisted pair. A cable medium with one or more pairs of twisted insulated copper conductors bound in a single

Z

Zoom (Digital)

The process of magnifying a video image by using computational algorithms on the digital signal.

Zoom (Optical)

The process of magnifying a video image by changing a lens’ focal length.

Zoom Lens

A lens that may be effectively used as a standard or telephoto lens by varying its focal length.

Zoom Ratio

The ratio of the starting focal length (wide position) to the ending focal length (telephoto position) of a zoom lens. A lens with a 10X zoom ratio will magnify the image at the wide-angle end by 10 times.

Appendix B: References

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

Cisco Video Surveillance Stream Manager Software: Installation and Upgrade Guides:
http://www.cisco.com/en/US/products/ps6940/prod_installation_guides_list.html

Design Guides:

- *Designing a Campus Network for High Availability* - http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/cdccont_0900aecd801a8a2d.pdf
- *Cisco AVVID Network Infrastructure IP Multicast Design (SRND)*- http://www.cisco.com/application/pdf/en/us/guest/tech/tk363/c1501/ccmigration_09186a008015e7cc.pdf
- *Enterprise QoS Solution Reference Network Design Guide* - http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/ccmigration_09186a008049b062.pdf
- *Cisco Unified Communications SRND Based on Cisco Unified CallManager 5.0* - http://www.cisco.com/en/US/products/sw/voicesw/ps556/products_implementation_design_guide_book09186a00806492bb.html
- *HA Campus Recovery Analysis* - http://www.cisco.com/application/pdf/en/us/guest/netsol/ns432/c649/cdccont_0900aecd801a89fc.pdf

List of multicast addresses, maintained by IANA:
http://www.iana.org/assignments/multicast-addresses

Cisco IOS IP Multicast Configuration Guide:
http://www.cisco.com/en/US/partner/products/ps6350/products_configuration_guide_book09186a0080435b9f.html

Catalyst 3750 Switch Software Configuration Guide, 12.2(25)SEE, “Configuring QoS”:
http://www.cisco.com/en/US/partner/products/hw/switches/ps5023/products_configuration_guide_chapter09186a00805a6504.html

RFC 2365, "Administratively Scoped IP Multicast," Best Current Practice:
http://www.ietf.org/rfc/rfc2365.txt

Appendix C: Initial Cisco IP Gateway Encoder/Decoder Configuration

Cisco Video Surveillance components ship with a default IP address that may need to be configured to operate in the proper network environment.

The following IP addresses are used by new devices or by devices that have been set to factory defaults:

- 192.168.0.100 – Cisco IP Gateway Encoders and Decoders
- 192.168.0.200 – Services Platforms

Steps to Change the IP Address of a Cisco IP Gateway Encoder

To configure a new Cisco IP Gateway Encoder to join an existing network, follow these steps:

1. Change the IP address of the PC that is running the Cisco Stream Manager Configuration Module. Choose an IP address in the 192.168.0.0 subnet. For example:

   IP Address: 192.168.0.25

   Subnet Mask: 255.255.255.0

2. Launch the Cisco Stream Manager Configuration Module and click Encoders.

The new device appears as Encoder #100, with its default IP address of 192.68.0.100:

Figure 49. New Encoder Configuration

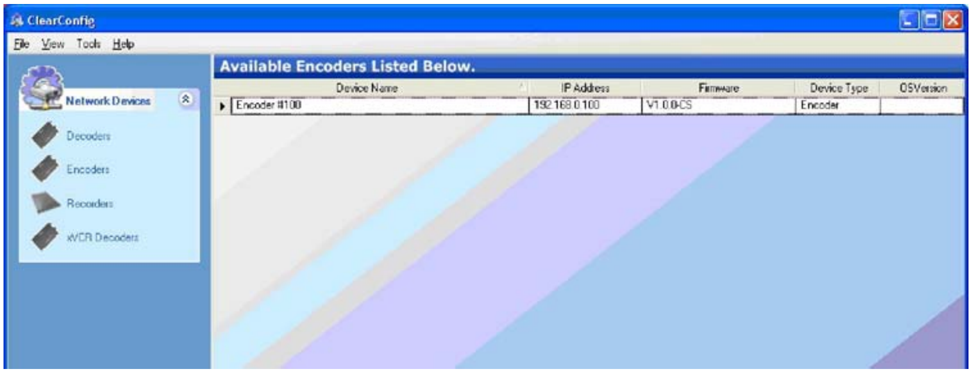

Available Encoders Listed Below.

Device Name	IP Address	Firmware	Device Type	OSVersion
Encoder #100	192.168.0.100	V1.0.0-CS	Encoder	

Note that the firmware version is also set to its factory defaults. In this case, it is version 1.0.0CS.

Optional: Before changing the IP address of this encoder to join the proper network subnet, upgrade the firmware using the Stream Manager Upgrade module.

3. Launch the Cisco Stream Manager Administration and Monitoring Module, and specify the location of the new firmware version. Highlight Encoder #100 and click Next. The upgrade process begins:

Figure 50. Firmware Upgrade

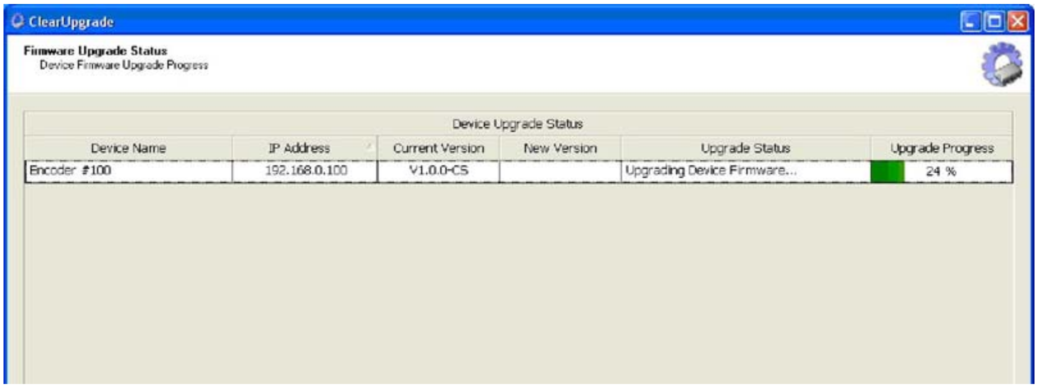

Firmware Upgrade Status
Device Firmware Upgrade Progress

Device Upgrade Status

Device Name	IP Address	Current Version	New Version	Upgrade Status	Upgrade Progress
Encoder #100	192.168.0.100	V1.0.0-CS		Upgrading Device Firmware...	24 %

When the process completes, the Cisco IP Gateway Encoder automatically reboots with the new firmware image.

4. Launch the Cisco Stream Manager Configuration Module and select the new encoder. Double-click Encoder#100 and click Network Settings.
5. Change the Network Settings to reflect the encoder's new address
6. Click Update. This action causes the Cisco IP Gateway Encoder to reboot and join the new subnet.

Note: Do not connect more than one new Cisco IP Gateway Encoder to a network segment at the same time. Each encoder comes from the factory with the same IP address. Make sure to assign a unique IP address to each device.

Figure 51. Network Settings

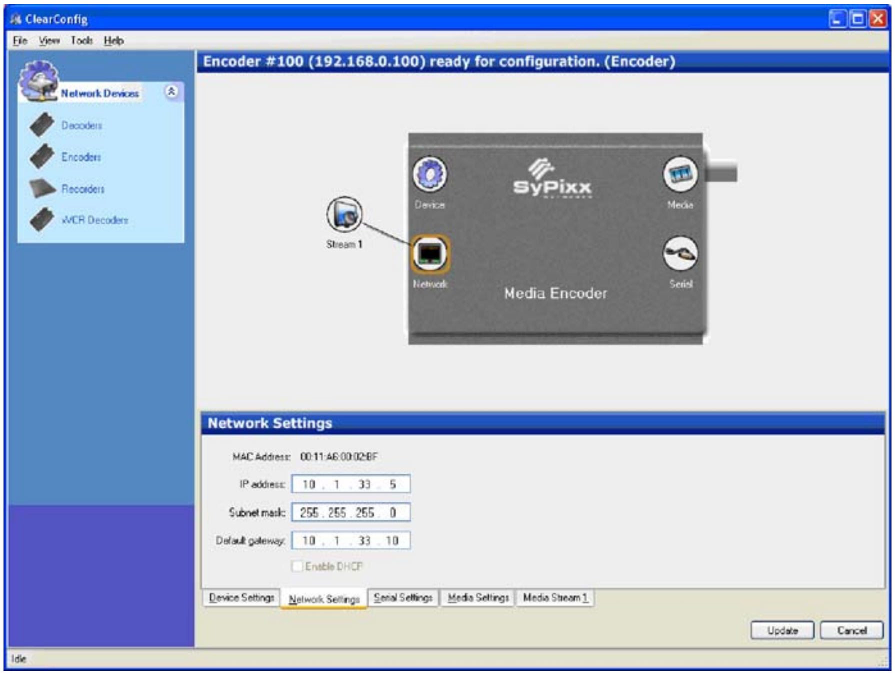

Encoder #100 (192.168.0.100) ready for configuration. (Encoder)

Network Settings

MAC Address: 00:11:A6:00:02:BF
IP address: 10 . 1 . 33 . 5
Subnet mask: 255 . 255 . 255 . 0
Default gateway: 10 . 1 . 33 . 10
□ Enable DHCP

Because the IP address of the encoder has changed, the IP address of the PC running the Cisco Stream Manager Configuration Module also needs to change to the new subnet (10.1.33.x in this case).

7. If necessary, adjust other settings, such as Media and Serial parameters.

Reset Modules to Factory Default

To restore the factory default configuration to a module, follow these steps:

1. Disconnect any inputs (Ethernet and video) to the module. If the module is mounted within a Convergence Chassis, loosen the module retainer screws and unseat the module. If the module is surface mounted, unplug the 2-pin power plug momentarily and then reseat it.
2. Use a pencil or small tool to press and hold the recessed FACTORY RESET button, which is located at the top of the module edge, and while pressing the reset button, reseat the module.
3. Continue pressing the reset button until all the module LEDs illuminate solid green, and then release the button. (This activity can take up to 30 seconds.) The module LEDs remain illuminated for several minutes.
4. When the LEDs return to normal states (all off except Power), tighten the retainer screws (for Convergence Chassis modules) or reseat the surface mount module. The module is now reset to factory default levels.

Appendix D: Device Configurations

The following devices were used while testing the scenarios in this design guide:

Hardware	Software
Catalyst 6504-E	12.2(18)SXE5
Catalyst 6503	12.2(18)SXF5
Catalyst 4507	12.2(25)EWA1
Catalyst 3750	12.2(25)SEE1
Catalyst 2960	12.2(25)SEE1
Catalyst 3560	12.2(25)SEE1
Cisco 2851 Routers	12.4(3e)
ASA 5500 Series Appliance	7.2(1)

IP Gateway Encoders and IP Gateway Decoders: Firmware V.1.4.9-CS

Stream Manager Configuration Module: Version 4.10.7.0

Stream Manager Client Viewing Module: Version 3.1.9.0

Stream Manager Administration and Monitoring Module: Version 2.4.1.0

Multicast Campus Deployment

Auto-RP using sparse-dense-mode, with QoS enabled

```
!
hostname 6504-1
!
!
ip multicast-routing
vtp domain ICE
vtp mode transparent
mls qos
!
spanning-tree mode pvst
spanning-tree loopguard default
spanning-tree portfast default
spanning-tree portfast bpduguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
spanning-tree vlan 30,32,34,40 priority 24576
spanning-tree vlan 31,33,35 priority 28672
!
class-map match-all VS-REPLAY-CLASS
  match access-group name VS-REPLAY-ACL
```

```
class-map match-all VS-LIVE-CLASS
  match access-group name VS-LIVE-ACL
!
!
policy-map VS-NVR-INGRESS
  class VS-LIVE-CLASS
   set dscp cs4
  class VS-REPLAY-CLASS
   set dscp af31
!
!
interface Loopback1
 ip address 10.1.20.249 255.255.255.252
 ip pim sparse-dense-mode
 ip igmp version 3
!
interface GigabitEthernet1/1
 description Connection to  6504-2
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 600
 switchport mode trunk
 no ip address
 wrr-queue cos-map 2 1 6
 priority-queue cos-map 1 4 5
 rcv-queue cos-map 1 3 6
 mls qos trust dscp
!
interface GigabitEthernet1/2
 description Uplink to 3750
 switchport
 switchport access vlan 600
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 600
 switchport trunk allowed vlan 500,600
 switchport mode trunk
 no ip address
 wrr-queue cos-map 2 1 6
 priority-queue cos-map 1 4 5
 rcv-queue cos-map 1 3 6
 spanning-tree bpdufilter enable
!
interface FastEthernet2/47
 description 2851-1
 ip address 10.1.20.9 255.255.255.252
 ip pim sparse-dense-mode
 ip igmp version 3
 wrr-queue cos-map 2 1 2 3
 priority-queue cos-map 1 4 5
 rcv-queue cos-map 1 1 0 1 2 3 6 7
 spanning-tree portfast
!
```

```
interface GigabitEthernet4/1
 description 3750-1
 switchport
 switchport access vlan 30
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 30
 switchport trunk allowed vlan 30
 switchport mode trunk
 no ip address
 wrr-queue cos-map 2 2 3
 priority-queue cos-map 1 4 5
 mls qos trust dscp
 spanning-tree guard loop
!
interface GigabitEthernet4/3
 description 2960-1
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 32
 switchport trunk allowed vlan 1,32
 switchport mode trunk
 no ip address
 wrr-queue cos-map 2 2 3
 priority-queue cos-map 1 4 5
 mls qos trust dscp
 spanning-tree guard loop
!
interface GigabitEthernet4/4
 description 4507-1
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 33
 switchport trunk allowed vlan 33
 switchport mode trunk
 no ip address
 wrr-queue cos-map 2 2 3
 priority-queue cos-map 1 4 5
 mls qos trust dscp
 spanning-tree guard loop
!
interface GigabitEthernet4/5
 description 3560-1
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 34
 switchport trunk allowed vlan 34
 switchport mode trunk
 no ip address
 wrr-queue cos-map 2 2 3
 priority-queue cos-map 1 4 5
 mls qos trust dscp
 spanning-tree guard loop
```

```
!
interface GigabitEthernet4/6
 description 6503-1
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 35
 switchport trunk allowed vlan 35
 switchport mode trunk
 no ip address
 wrr-queue cos-map 2 2 3
 priority-queue cos-map 1 4 5
 mls qos trust dscp
 spanning-tree guard loop
!
interface Vlan1
 no ip address
 shutdown
!
interface Vlan30
 description IDF A
 ip address 10.1.30.2 255.255.255.0
 no ip redirects
 no ip unreachables
 no ip proxy-arp
 ip pim sparse-dense-mode
 ip igmp version 3
 standby 30 ip 10.1.30.1
 standby 30 priority 110
 standby 30 preempt
!
interface Vlan32
 description IDF C
 ip address 10.1.32.2 255.255.255.0
 no ip redirects
 no ip unreachables
 no ip proxy-arp
 ip pim sparse-dense-mode
 ip igmp version 3
 standby 32 ip 10.1.32.1
 standby 32 priority 110
 standby 32 preempt
!
interface Vlan33
 description IDF D
 ip address 10.1.33.2 255.255.255.0
 no ip redirects
 no ip unreachables
 no ip proxy-arp
 ip pim sparse-dense-mode
 ip igmp version 3
 standby 33 ip 10.1.33.1
 standby 33 priority 110
 standby 33 preempt
!
interface Vlan34
 description IDF E
 ip address 10.1.34.2 255.255.255.0
 no ip redirects
 no ip unreachables
 no ip proxy-arp
 ip pim sparse-dense-mode
 ip igmp version 3
 standby 34 ip 10.1.34.1
 standby 34 priority 110
 standby 34 preempt
!
interface Vlan35
 description IDF F
 ip address 10.1.35.2 255.255.255.0
 no ip redirects
 no ip unreachables
 ip pim sparse-dense-mode
 ip igmp version 3
 standby 35 ip 10.1.35.1
 standby 35 priority 90
 standby 35 preempt
!
interface Vlan500
 ip address 10.94.165.2 255.255.255.192
 ip pim sparse-dense-mode
 ip igmp version 3
!
interface Vlan600
 ip address 10.94.162.194 255.255.255.192
 ip pim sparse-dense-mode
 ip igmp version 3
!
router eigrp 999
 passive-interface Vlan30
 passive-interface Vlan31
 passive-interface Vlan32
 passive-interface Vlan33
 passive-interface Vlan34
 passive-interface Vlan35
 passive-interface Vlan42
 passive-interface Vlan43
 network 10.0.0.0
 no auto-summary
!
ip classless
!
ip pim send-rp-announce Loopback1 scope 16
ip pim send-rp-discovery Loopback1 scope 16
!
```

```
ip access-list extended VS-LIVE-ACL
 permit udp any any
ip access-list extended VS-REPLAY-ACL
 permit tcp any any
!
!
!
line con 0
 exec-timeout 45 0
line vty 0 4
 exec-timeout 35 0
 login
!
end

!
hostname 6504-2
!
ip multicast-routing
vtp domain ICE
vtp mode transparent
mls qos
!
spanning-tree mode pvst
spanning-tree loopguard default
spanning-tree portfast default
spanning-tree portfast bpduguard default
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
spanning-tree vlan 30,32,34 priority 28672
spanning-tree vlan 31,33,35 priority 24576
!
!
interface Loopback1
 description Loopback1
 ip address 10.1.20.253 255.255.255.252
 ip pim sparse-dense-mode
 ip igmp version 3
!
interface GigabitEthernet1/1
 description Connection to  6504-1
 switchport
 switchport access vlan 600
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 600
 no ip address
 wrr-queue cos-map 2 1 6
 priority-queue cos-map 1 4 5
 rcv-queue cos-map 1 3 6
 mls qos trust dscp
!
interface FastEthernet2/33
```

```
 description 2851-1
 ip address 10.1.20.5 255.255.255.252
 ip pim sparse-dense-mode
 ip igmp version 3
!
interface GigabitEthernet4/1
 description 3750-1
 switchport
 switchport access vlan 30
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 30
 switchport trunk allowed vlan 30
 switchport mode trunk
 no ip address
 wrr-queue cos-map 2 2 3
 priority-queue cos-map 1 4 5
 mls qos trust dscp
 spanning-tree guard loop
!
interface GigabitEthernet4/3
 description 2960-2
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 32
 switchport trunk allowed vlan 1,32
 switchport mode trunk
 no ip address
 wrr-queue cos-map 2 2 3
 priority-queue cos-map 1 4 5
 mls qos trust dscp
 spanning-tree guard loop
!
interface GigabitEthernet4/4
 description 4507-1
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 33
 switchport trunk allowed vlan 33
 switchport mode trunk
 no ip address
 wrr-queue cos-map 2 2 3
 priority-queue cos-map 1 4 5
 mls qos trust dscp
 spanning-tree guard loop
!
interface GigabitEthernet4/5
 description 3560-2
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 34
 switchport trunk allowed vlan 34
 switchport mode trunk
 no ip address
 wrr-queue cos-map 2 2 3
 priority-queue cos-map 1 4 5
 mls qos trust dscp
 spanning-tree guard loop
!
interface GigabitEthernet4/6
 description 6503-1
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 35
 switchport trunk allowed vlan 35
 switchport mode trunk
 no ip address
 wrr-queue cos-map 2 2 3
 priority-queue cos-map 1 4 5
 mls qos trust dscp
 spanning-tree guard loop
!
interface Vlan1
 no ip address
 shutdown
!
interface Vlan30
 description IDF A
 ip address 10.1.30.3 255.255.255.0
 no ip redirects
 no ip unreachables
 no ip proxy-arp
 ip pim sparse-dense-mode
 ip igmp version 3
 standby 30 ip 10.1.30.1
 standby 30 priority 90
!
interface Vlan32
 ip address 10.1.32.3 255.255.255.0
 no ip redirects
 no ip unreachables
 no ip proxy-arp
 ip pim sparse-dense-mode
 ip igmp version 3
 standby 32 ip 10.1.32.1
 standby 32 priority 90
!
interface Vlan33
 description IDF D
 ip address 10.1.33.3 255.255.255.0
 no ip redirects
 no ip unreachables
 no ip proxy-arp
 ip pim sparse-dense-mode
 ip igmp version 3
 standby 33 ip 10.1.33.1
 standby 33 priority 90
!
interface Vlan34
 description IDF E
 ip address 10.1.34.3 255.255.255.0
 no ip redirects
 no ip unreachables
 no ip proxy-arp
 ip pim sparse-dense-mode
 ip igmp version 3
 standby 34 ip 10.1.34.1
 standby 34 priority 90
!
interface Vlan35
 description IDF F
 ip address 10.1.35.3 255.255.255.0
 no ip redirects
 no ip unreachables
 ip pim sparse-dense-mode
 ip igmp version 3
 standby 35 ip 10.1.35.1
 standby 35 priority 110
!
interface Vlan42
 ip address 10.1.42.3 255.255.255.0
 no ip redirects
 no ip unreachables
 no ip proxy-arp
 ip pim sparse-dense-mode
 ip igmp version 3
 standby 42 ip 10.1.42.1
 standby 42 priority 90
!
interface Vlan43
 description IDF D
 ip address 10.1.43.3 255.255.255.0
 no ip redirects
 no ip unreachables
 no ip proxy-arp
 ip pim sparse-dense-mode
 ip igmp version 3
 standby 43 ip 10.1.43.1
 standby 43 priority 90
!
interface Vlan500
 ip address 10.94.165.3 255.255.255.192
 ip pim sparse-dense-mode
 ip igmp version 3
!
interface Vlan600
 ip address 10.94.162.250 255.255.255.192
 ip pim sparse-dense-mode
 ip igmp version 3
!
router eigrp 999
 passive-interface Vlan30
 passive-interface Vlan31
 passive-interface Vlan32
 passive-interface Vlan33
 passive-interface Vlan34
 passive-interface Vlan35
 passive-interface Vlan42
 passive-interface Vlan43
 network 10.0.0.0
 no auto-summary
!
ip pim send-rp-announce Loopback1 scope 16
ip pim send-rp-discovery Loopback1 scope 16
!
line con 0
line vty 0 4
 login
!
!
end

!
hostname 3750-1
!
udld enable
!
mls qos srr-queue output cos-map queue 1 threshold 1  4
mls qos srr-queue output dscp-map queue 1 threshold 1  32
mls qos
!
spanning-tree mode pvst
spanning-tree extend system-id
spanning-tree uplinkfast
!
interface GigabitEthernet1/0/1
 description to 6504_2
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 30
 switchport trunk allowed vlan 30
 switchport mode dynamic desirable
 priority-queue out
 mls qos trust dscp
!
interface GigabitEthernet1/0/2
 description to 6504_1
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 30
 switchport trunk allowed vlan 30
 switchport mode dynamic desirable
 priority-queue out
 mls qos trust dscp
!
interface GigabitEthernet1/0/5
 switchport access vlan 30
 switchport mode access
 mls qos cos 4
 mls qos cos override
 spanning-tree portfast
!
interface Vlan1
 no ip address
 no ip route-cache
 no ip mroute-cache
!
interface Vlan30
 ip address 10.1.30.10 255.255.255.0
 no ip route-cache
 no ip mroute-cache
!
ip default-gateway 10.1.30.1
ip classless
ip http server
!
!
ip access-list extended Block_Client
!
access-list 1 permit any
!
!
line con 0
line vty 0 4
 login
!
end

!
hostname 2960-1
!
!
vtp domain ICE
vtp mode transparent
!
mls qos srr-queue output cos-map queue 1 threshold 1  4
mls qos srr-queue output dscp-map queue 1 threshold 1  32
mls qos
!
!
spanning-tree mode pvst
spanning-tree extend system-id
!
```

```
vlan 32
 name IDF_C_VS
!
!
interface FastEthernet0/5
 description Encoder 32
 switchport access vlan 32
 switchport mode access
 mls qos cos 4
 mls qos cos override
 spanning-tree portfast
 spanning-tree bpduguard enable
!
interface GigabitEthernet0/1
 description to 6504-1
 switchport trunk native vlan 32
 switchport trunk allowed vlan 1,32
 switchport mode trunk
 priority-queue out
 mls qos trust dscp
!
interface GigabitEthernet0/2
 description 2960-2
 switchport trunk native vlan 32
 switchport trunk allowed vlan 32
 switchport mode trunk
 priority-queue out
 mls qos trust dscp
!
interface Vlan1
 no ip address
 no ip route-cache
 shutdown
!
interface Vlan32
 ip address 10.1.32.10 255.255.255.0
 no ip route-cache
!
ip default-gateway 10.1.32.1
!
line con 0
 exec-timeout 45 0
line vty 0 4
 login
line vty 5 15
 no login
!
!
end

!
```

```
hostname 4507-1
!
vtp domain ICE
vtp mode transparent
udld enable
!
spanning-tree mode pvst
spanning-tree extend system-id
spanning-tree uplinkfast
!
vlan 33
 name IDF_D_VS
!
!
interface GigabitEthernet1/3
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 33
 switchport trunk allowed vlan 33,43
 switchport mode trunk
 qos trust dscp
 tx-queue 3
   bandwidth percent 30
   priority high
 spanning-tree guard loop
!
interface GigabitEthernet2/3
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 33
 switchport trunk allowed vlan 33,43
 switchport mode trunk
 qos trust dscp
 tx-queue 3
   bandwidth percent 30
   priority high
 spanning-tree guard loop
!
interface GigabitEthernet3/5
 description Encoder #33
 switchport access vlan 33
 qos dscp 32
 spanning-tree portfast
 spanning-tree bpduguard enable
!
interface GigabitEthernet3/6
 switchport access vlan 33
 qos trust dscp
!
interface Vlan1
 no ip address
!
interface Vlan33
 ip address 10.1.33.10 255.255.255.0
```

```
!
ip route 0.0.0.0 0.0.0.0 10.1.33.1
!
line con 0
 stopbits 1
line vty 0 4
 login
!
end

!
hostname 3560-1
!
vtp domain ICE
vtp mode transparent
!
mls qos srr-queue output cos-map queue 1 threshold 1  4
mls qos srr-queue output dscp-map queue 1 threshold 1  32
mls qos
!
!
spanning-tree mode pvst
!

vlan 34
 name IDF_E_VS
!
interface FastEthernet0/5
 description Encoder 34
 switchport access vlan 34
 switchport mode access
 mls qos cos 4
 mls qos cos override
 spanning-tree portfast
!
interface GigabitEthernet0/1
 description connection to 6504-1
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 34
 switchport trunk allowed vlan 1,34
 switchport mode trunk
 priority-queue out
!
interface GigabitEthernet0/2
 description to 3560-2
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 34
 switchport trunk allowed vlan 1,34
 switchport mode trunk
 priority-queue out
!
interface Vlan1
```

```
 no ip address
 shutdown
!
interface Vlan34
 ip address 10.1.34.10 255.255.255.0
 no ip mroute-cache
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.1.34.1
!
line con 0
line vty 0 4
 login
!
end

!
hostname 6503-1
!
ip multicast-routing
udld enable

vtp domain ICE
vtp mode transparent
mls ip multicast flow-stat-timer 9
mls qos
!
spanning-tree mode pvst
spanning-tree extend system-id
spanning-tree uplinkfast
!
vlan 35
 name IDF_F_VS
!
interface GigabitEthernet1/1
 description to 6504-1
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 35
 switchport trunk allowed vlan 35
 switchport mode trunk
 no ip address
 wrr-queue cos-map 2 2 3
 priority-queue cos-map 1 4 5
 mls qos trust dscp
!
interface GigabitEthernet1/2
 description to 6504-2
 switchport
 switchport trunk encapsulation dot1q
 switchport trunk native vlan 35
 switchport trunk allowed vlan 35
 switchport mode trunk
 no ip address
 wrr-queue cos-map 2 2 3
 priority-queue cos-map 1 4 5
!
interface FastEthernet2/5
 description Encoder 35
 switchport
 switchport access vlan 35
 no ip address
 mls qos trust cos
 mls qos cos 4
!
interface Vlan1
 no ip address
 shutdown
!
interface Vlan35
 ip address 10.1.35.10 255.255.255.0
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.1.35.1
!
!
line con 0
line vty 0 4
 login
!
!
end

!
!
hostname 2851-1
!
ip cef
!
ip multicast-routing
!
controller T1 0/0/0
 framing esf
 clock source internal
 linecode b8zs
 cablelength short 133
 channel-group 1 timeslots 1-24
!
controller T1 0/0/1
 framing esf
 clock source internal
 linecode b8zs
 cablelength short 133
 channel-group 1 timeslots 1-24
```

```
!
class-map match-all VS-REPLAY-CLASS
 match access-group name VS-REPLAY-ACL
class-map match-all VS-LIVE-CLASS
 match access-group name VS-LIVE-ACL
!
!
policy-map HQ-WAN
 class VS-LIVE-CLASS
  priority 4350 217500
 class VS-REPLAY-CLASS
  bandwidth 3300
  shape average 3450000
!
!
!
!
!
interface Multilink1
 ip address 10.1.20.1 255.255.255.252
 ip pim sparse-dense-mode
 ppp multilink
 ppp multilink group 1
 max-reserved-bandwidth 100
 service-policy output HQ-WAN
!
interface GigabitEthernet0/0
 description to 6504-1 Port 2/33
 ip address 10.1.20.6 255.255.255.252
 ip pim sparse-dense-mode
 ip igmp version 3
 duplex auto
 speed auto
!
interface GigabitEthernet0/1
 ip address 10.1.20.10 255.255.255.252
 ip pim sparse-dense-mode
 ip igmp version 3
 duplex auto
 speed auto
!
interface Serial0/0/0:1
 no ip address
 encapsulation ppp
 ppp multilink group 1
 max-reserved-bandwidth 100
!
interface Serial0/0/1:1
 no ip address
 encapsulation ppp
 ppp multilink group 1
 max-reserved-bandwidth 100
```

```
!
interface Serial0/1/0
 no ip address
 encapsulation ppp
 clock rate 2016000
 clock rate 2016000
 dce-terminal-timing-enable
 ppp multilink group 1
 max-reserved-bandwidth 100
!
interface Serial0/1/1
 no ip address
 encapsulation ppp
 clock rate 2016000
 clock rate 2016000
 dce-terminal-timing-enable
 ppp multilink group 1
 max-reserved-bandwidth 100
!
interface Serial0/1/2
 no ip address
 encapsulation ppp
 clock rate 2016000
 clock rate 2016000
 dce-terminal-timing-enable
 ppp multilink group 1
 max-reserved-bandwidth 100
!
interface Serial0/1/3
 no ip address
 encapsulation ppp
 clock rate 2016000
 clock rate 2016000
 dce-terminal-timing-enable
 ppp multilink group 1
 max-reserved-bandwidth 100
!
interface Vlan1
 no ip address
!
interface Vlan600
 no ip address
!
router eigrp 999
 network 10.0.0.0
 no auto-summary
!
ip classless
!
ip access-list extended VS-LIVE-ACL
 permit ip any any dscp cs4
ip access-list extended VS-REPLAY-ACL
```

```
 permit ip any any dscp af31
!
!
line con 0
line aux 0
line vty 0 4
 login
line vty 5 15
 privilege level 15
 login local
 transport input telnet
!

end


!
!
hostname 2851-2
!
ip cef
!
ip multicast-routing
!
!
controller T1 0/0/1
 shutdown
 framing esf
 linecode b8zs
 cablelength short 133
 channel-group 1 timeslots 1-24
!
class-map match-all VIDEO-SURV-CLASS
 match access-group name VIDEO-SURV-ACL
!
!
policy-map BRANCH-WAN-EDGE
 class VIDEO-SURV-CLASS
  priority 1015 50000
 class class-default
  fair-queue
!
interface Multilink1
 ip address 10.1.20.2 255.255.255.252
 ip pim sparse-dense-mode
 ppp multilink
 ppp multilink group 1
 max-reserved-bandwidth 100
 service-policy output BRANCH-WAN-EDGE
!
interface GigabitEthernet0/0
 ip address 10.1.27.1 255.255.255.0
 ip pim sparse-dense-mode
 ip igmp version 3
 duplex auto
 speed auto
 arp timeout 10000
!
interface Serial0/0/0:1
 no ip address
 encapsulation ppp
 ppp multilink group 1
 max-reserved-bandwidth 100
!
interface Serial0/0/1:1
 no ip address
 encapsulation ppp
 ppp multilink group 1
 max-reserved-bandwidth 100
!
interface Serial0/1/0
 no ip address
 encapsulation ppp
 ppp multilink group 1
 max-reserved-bandwidth 100
!
interface Serial0/1/1
 no ip address
 encapsulation ppp
 ppp multilink group 1
 max-reserved-bandwidth 100
!
interface Serial0/1/2
 no ip address
 encapsulation ppp
 ppp multilink group 1
 max-reserved-bandwidth 100
!
interface Serial0/1/3
 no ip address
 encapsulation ppp
 ppp multilink group 1
 max-reserved-bandwidth 100
!
interface Vlan1
 no ip address
!
router eigrp 999
 network 10.0.0.0
 no auto-summary
!
ip access-list extended VIDEO-SURV-ACL
 permit ip any any dscp cs4
!
```

```
line con 0
line aux 0
line vty 0 4
 privilege level 15
 login
 transport input telnet
!
!
end
```

Firewall Configurations

With simple ACL

```
ASA Version 7.2(1)
!
hostname ASA-1
domain-name default.domain.invalid
enable password 9jNfZuG3TC5tCVH0 encrypted
multicast-routing
names
!
interface GigabitEthernet0/0
 description Outside
 nameif Outside
 security-level 0
 ip address 10.1.37.2 255.255.255.0
!
interface GigabitEthernet0/1
 description Inside
 nameif inside
 security-level 100
 ip address 10.1.34.1 255.255.255.0
!
interface GigabitEthernet0/2
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/3
 shutdown
 no nameif
 no security-level
 no ip address
!
interface Management0/0
 shutdown
 no nameif
 no security-level
 no ip address
!
```

```
passwd 9jNfZuG3TC5tCVH0 encrypted
pim rp-address 10.94.162.250
ftp mode passive
dns server-group DefaultDNS
 domain-name default.domain.invalid
object-group network IP_GATEWAYS
 description Encoders and Decoders
 network-object host 10.1.34.5
 network-object host 10.1.34.6
 network-object host 10.1.34.55
object-group network Multicast_Control
 network-object host 235.1.1.1
 network-object host 239.1.1.2
 network-object host 239.1.1.3
object-group icmp-type Standard_ICMP
 icmp-object echo
 icmp-object echo-reply
 icmp-object time-exceeded
 icmp-object unreachable
access-list Outside_ACL extended permit udp any object-group IP_GATEWAYS
access-list Outside_ACL extended permit udp any object-group
Multicast_Control
access-list Outside_ACL extended permit icmp any object-group IP_GATEWAYS
object-group Standard_ICMP
access-list Outside_ACL extended permit udp any 239.0.0.0 255.0.0.0
pager lines 24
logging enable
logging timestamp
logging buffered warnings
logging asdm warnings
mtu Outside 1500
mtu inside 1500
no failover
asdm image disk0:/asdm521.bin
no asdm history enable
arp timeout 14400
access-group Outside_ACL in interface Outside
route Outside 0.0.0.0 0.0.0.0 10.1.37.1 1
timeout xlate 3:00:00
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00 h225 1:00:00 mgcp 0:05:00 mgcp-pat
0:05:00
timeout sip 0:30:00 sip_media 0:02:00 sip-invite 0:03:00 sip-disconnect
0:02:00
timeout uauth 0:05:00 absolute
http server enable
http 10.1.34.0 255.255.255.0 inside
no snmp-server location
no snmp-server contact
snmp-server enable traps snmp authentication linkup linkdown coldstart
fragment chain 45 Outside
fragment chain 45 inside
```

```
telnet timeout 5
ssh timeout 5
console timeout 0
!
class-map inspection_default
 match default-inspection-traffic
!
!
policy-map type inspect dns preset_dns_map
 parameters
  message-length maximum 512
policy-map global_policy
 class inspection_default
  inspect dns preset_dns_map
  inspect ftp
  inspect h323 h225
  inspect h323 ras
  inspect netbios
  inspect rsh
  inspect rtsp
  inspect skinny
  inspect esmtp
  inspect sqlnet
  inspect sunrpc
  inspect tftp
  inspect sip
  inspect xdmcp
!
service-policy global_policy global
prompt hostname context
Cryptochecksum:5d588df5283032ad40cbf202d75b8509
: end
```

With more restrictive ACL

```
: Saved
:
ASA Version 7.2(1)
!
hostname ASA-1
domain-name default.domain.invalid
enable password 9jNfZuG3TC5tCVH0 encrypted
multicast-routing
names
dns-guard
!
interface GigabitEthernet0/0
 description Outside
 nameif Outside
 security-level 0
 ip address 10.1.37.2 255.255.255.0
!
interface GigabitEthernet0/1
```

```
 description Inside
 nameif inside
 security-level 100
 ip address 10.1.34.1 255.255.255.0
!
interface GigabitEthernet0/2
 shutdown
 no nameif
 no security-level
 no ip address
!
interface GigabitEthernet0/3
 shutdown
 no nameif
 no security-level
 no ip address
!
interface Management0/0
 shutdown
 no nameif
 no security-level
 no ip address
!
passwd 9jNfZuG3TC5tCVH0 encrypted
pim rp-address 10.94.162.250
ftp mode passive
dns server-group DefaultDNS
 domain-name default.domain.invalid
object-group network IP_GATEWAYS
 description Encoders and Decoders
 network-object host 10.1.34.5
 network-object host 10.1.34.6
 network-object host 10.1.34.55
object-group network Multicast_Control
 network-object host 235.1.1.1
 network-object host 239.1.1.2
 network-object host 239.1.1.3
object-group icmp-type Standard_ICMP
 icmp-object echo
 icmp-object echo-reply
 icmp-object time-exceeded
 icmp-object unreachable
object-group network Inside_IPGateways
 network-object host 10.1.34.5
 network-object host 10.1.34.6
 network-object host 10.1.34.55
object-group network Outside_Recorders
 network-object host 10.94.162.231
object-group network Outside_IPGateways
 network-object host 10.94.162.211
 network-object host 10.94.162.212
 network-object host 10.94.162.215
 network-object host 10.1.27.12
 network-object host 10.1.21.5
 network-object host 10.1.23.5
 network-object host 10.1.30.5
 network-object host 10.94.162.206
 network-object host 10.94.162.217
 network-object host 10.94.162.232
 network-object host 10.94.162.233
 network-object host 10.94.162.219
 network-object host 10.1.33.5
 network-object host 10.1.31.5
 network-object host 10.1.35.5
 network-object host 10.94.162.207
 network-object host 10.94.162.234
 group-object Outside_Recorders
access-list OUTSIDE extended permit udp object-group Outside_IPGateways
object-group Inside_IPGateways
access-list OUTSIDE extended permit icmp object-group Outside_IPGateways
object-group Inside_IPGateways object-group Standard_ICMP
access-list OUTSIDE extended permit udp object-group Outside_IPGateways
object-group Multicast_Control
access-list OUTSIDE extended permit udp object-group Outside_Recorders
239.0.0.0 255.0.0.0
access-list OUTSIDE extended permit udp object-group Outside_IPGateways
239.255.0.0 255.255.0.0
pager lines 24
logging enable
logging timestamp
logging buffered warnings
logging asdm warnings
mtu Outside 1500
mtu inside 1500
no failover
asdm image disk0:/asdm521.bin
no asdm history enable
arp timeout 14400
access-group OUTSIDE in interface Outside
route Outside 0.0.0.0 0.0.0.0 10.1.37.1 1
timeout xlate 3:00:00
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00 h225 1:00:00 mgcp 0:05:00 mgcp-pat
0:05:00
timeout sip 0:30:00 sip_media 0:02:00 sip-invite 0:03:00 sip-disconnect
0:02:00
timeout uauth 0:05:00 absolute
http server enable
http 10.1.34.0 255.255.255.0 inside
no snmp-server location
no snmp-server contact
snmp-server enable traps snmp authentication linkup linkdown coldstart
fragment chain 45 Outside
fragment chain 45 inside
```

```
telnet timeout 5
ssh timeout 5
console timeout 0
!
class-map inspection_default
 match default-inspection-traffic
!
!
policy-map type inspect dns migrated_dns_map_1
 parameters
  message-length maximum 512
policy-map global_policy
 class inspection_default
  inspect dns migrated_dns_map_1
  inspect ftp
  inspect h323 h225
  inspect h323 ras
  inspect netbios
  inspect rsh
  inspect rtsp
  inspect skinny
  inspect esmtp
  inspect sqlnet
  inspect sunrpc
  inspect tftp
  inspect sip
  inspect xdmcp
!
service-policy global_policy global
prompt hostname context
Cryptochecksum:358189cdb1f0bac013568ce1240dab3f
: endno asdm history enable
arp timeout 14400
access-group OUTSIDE in interface Outside
route Outside 0.0.0.0 0.0.0.0 10.1.37.1 1
timeout xlate 3:00:00
timeout conn 1:00:00 half-closed 0:10:00 udp 0:02:00 icmp 0:00:02
timeout sunrpc 0:10:00 h323 0:05:00 h225 1:00:00 mgcp 0:05:00 mgcp-pat
0:05:00
timeout sip 0:30:00 sip_media 0:02:00 sip-invite 0:03:00 sip-disconnect
0:02:00
timeout uauth 0:05:00 absolute
http server enable
http 10.1.34.0 255.255.255.0 inside
no snmp-server location
no snmp-server contact
snmp-server enable traps snmp authentication linkup linkdown coldstart
fragment chain 45 Outside
fragment chain 45 inside
telnet timeout 5
ssh timeout 5
console timeout 0
```

```
!
class-map inspection_default
 match default-inspection-traffic
!
!
policy-map type inspect dns migrated_dns_map_1
 parameters
  message-length maximum 512
policy-map global_policy
 class inspection_default
  inspect dns migrated_dns_map_1
  inspect ftp
  inspect h323 h225
  inspect h323 ras
  inspect netbios
  inspect rsh
  inspect rtsp
  inspect skinny
  inspect esmtp
  inspect sqlnet
  inspect sunrpc
  inspect tftp
  inspect sip
  inspect xdmcp
!
service-policy global_policy global
prompt hostname context
Cryptochecksum:8a2ab2239d533fd94a2f397c416d4545
end
```

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Europe Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: +31 0 800 020 0791
Fax: +31 0 20 357 1100

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.**

Printed in USA