Video Surveillance Design Guide for UCS B- and C-Series Platforms

November 2012

This guide summarizes high-level design recommendations and best practices for implementing IP video surveillance on the Cisco® Unified Computing System™ (UCS) B- and C-Series platforms.

Contents

This document includes the following sections:

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

Introduction

This guide summarizes high-level design recommendations and best practices for implementing IP video surveillance on the Cisco UCS B- and C-Series platforms. In some instances, existing network equipment and topologies have the necessary configuration and performance characteristics to support high-quality IP video surveillance. In other instances, network hardware might require upgrading or reconfiguration to support increased bandwidth needed to support video.

Note This guide does not describe the configuration and operation of theVideo Surveillance Manager (VSM) products; however, for more detailed information, see the "More Information" section on page 21.

Note This guide assumes that the ESXi 5.0 Hypervisor is installed on the UCS B- and C–Series platforms.

Figure 1 represents a virtualized VSM application running on a UCS B-Series platform.

Figure 1 Cisco IP Video Surveillance on UCS B-Series Platform

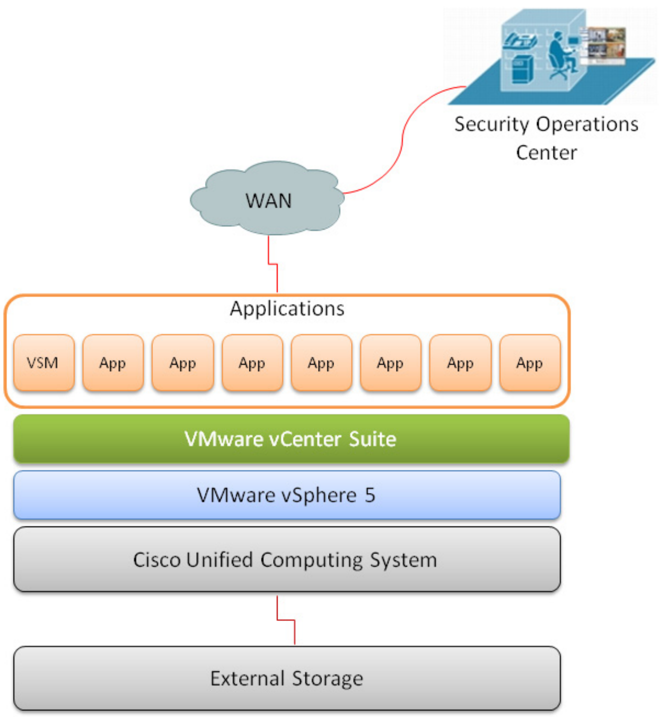

Solution Components

The required components for designing and deploying VSM on UCS B-and C–Series platforms include:

- **UCS B- and C-Series servers**—The Cisco UCS Servers can be deployed as rack-mount servers (C-Series) or blade servers (B-Series) running the ESXi 5.0 virtualization software.The B-Series servers deliver a scalable and flexible architecture to meet your data center needs while helping to reduce the total cost of ownership. The C-Series servers address fluctuating workload challenges through a varying balance of processing, memory, I/O, and internal storage resources.
- **Virtualized Video Surveillance Manager (VSM) software—**This software runs on UCS B-or C-Series server in a virtualized environment. The VSM software is available as an Open Virtual Appliance (OVA) file on Cisco.com. The OVA package is a tar file with the Open Virtualization Format (OVF) directory inside. Apart from the OVA file, the Video Surveillance Operations Manager (VSOM) virtual machine (VM) License and Video Surveillance Media Server (VSMS) VM License are also required.
- **Cameras—**The Cisco IP video surveillance camera and analog cameras are attached to encoders, analog gateway network modules for the integrated services router, or third-party IP surveillance cameras.
- **Network—**This component is comprised of the enterprise network—the Media Ready Network. The primary focus of this design guide is to reference the existing design baselines of branch office and wide-area network (WAN) while building on this base of knowledge with IP video Surveillance requirements, best practices, and design recommendations.

Logical Network Topology

Figure 2 illustrates the overall logical topology of the networking and video surveillance components, including a UCS B-Series containing the ESXi host running VSM and VSOM, various IP cameras, an external switch, and the operator workstations running the VSOM client.

Figure 2 Logical Network Topology

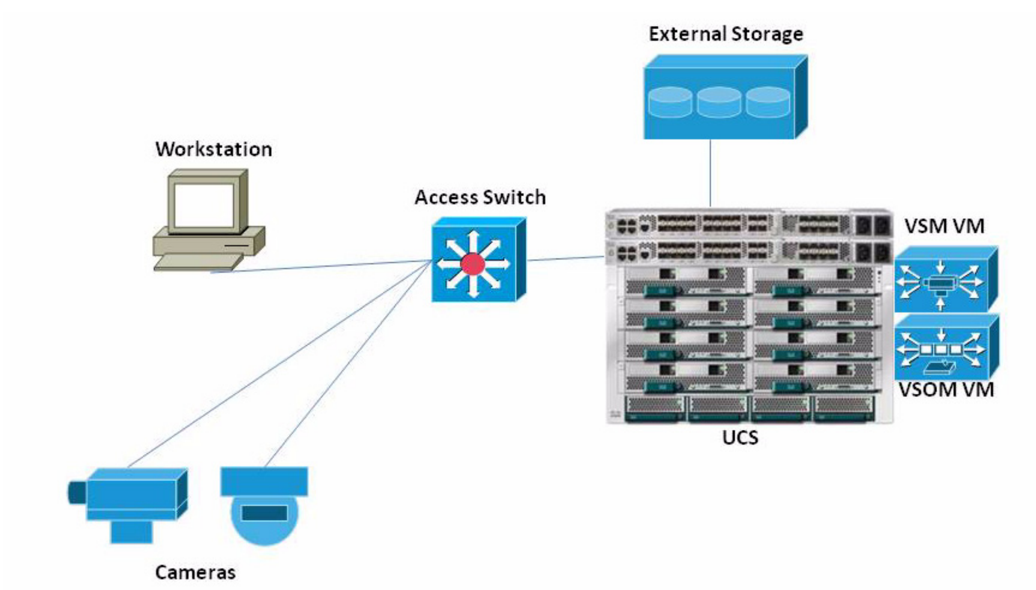

IP Network Infrastructure

Bandwidth

The bandwidth requirements for all video, but video surveillance in particular, is substantial compared to Voice over IP (VoIP). Common codecs used in VoIP deployments (G.711, G.729, G.726) use between 8 and 64 Kbps for voice encoding. A packet capture from a Cisco Video Surveillance 2600 Series IP Camera configured at a constant bit rate (CBR) target of 1 M for the MPEG-4 feed with audio-enabled on the camera. Control traffic (Hypertext Transfer Protocol [HTTP]/Real-Time Streaming Protocol [RTSP]) is also captured. Figure 3 illustrates the relationship between the amount of audio, video, and control plane.

Figure 3 Audio and Video Network Load

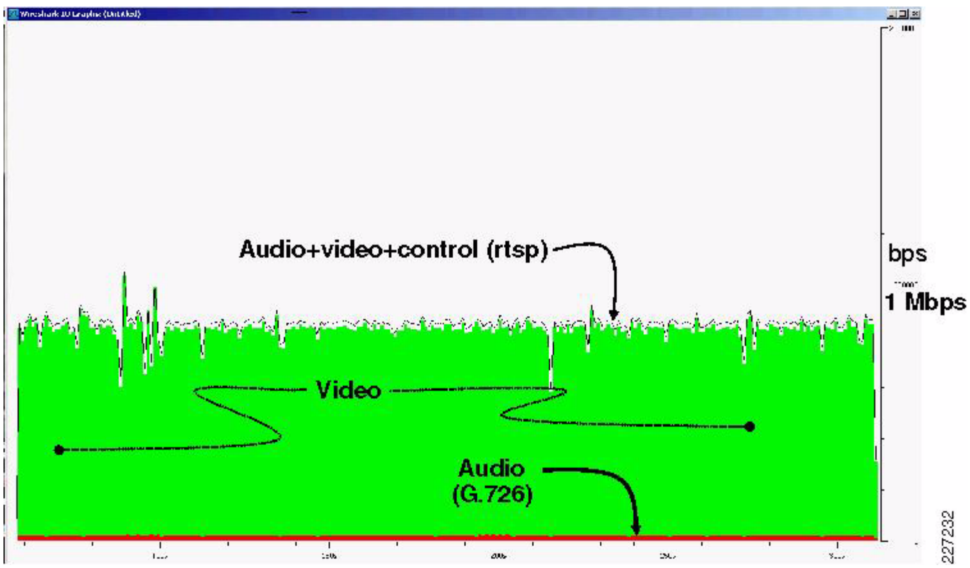

The bandwidth requirement for the video media stream is of a magnitude higher than audio (VoIP) and signaling (RTSP) for video in the enterprise network.

While provisioning for this bandwidth requirement is a key element in planning for video in the enterprise network, there are other network requirements to be considered. For example, will the video traffic be segmented on both the local-area network (LAN) and WAN from other user traffic, either logically or physically? Is the video deployment an overlay on an existing network infrastructure or is it an entirely new deployment? Is IPsec encryption currently implemented? These factors must be considered, along with the bandwidth requirements.

QoS

Quality of Service (QoS) is a key element to managing network congestion during periods where bandwidth is constrained. QoS, however, does not eliminate bandwidth constraints; it manages the access to bandwidth by competing applications through prioritizing one application over another. QoS manages unfairness. Because the video quality for MPEG-4 and H.264 is highly dependent on little or no packet loss, IP video surveillance traffic must not be dropped by the enterprise QoS policy. Motion JPEG-based video does not suffer degradation in the image with packet loss due to lack of bandwidth, but the smoothness of motion is compromised. Several frames or even several seconds of video may be missing with no indication of loss. Because many video surveillance deployments are headless, the first time the video is viewed may be days or weeks after capture. If the quality is poor due to packet loss in the network, there is no recourse and the video data is worthless.

Security

Security focuses on controlling what users have access to a resource while in transit, at the originating node, or when it is processed or stored on a server. One aspect of IP networking is the any-to-any connectivity between networks and users. This strength is also a flaw. There is a certain population of users on the network that must have access to the video surveillance system, but many cannot be trusted to access this data. Video surveillance data is particularly sensitive because access to the system by unscrupulous individuals can expose the enterprise to financial loss and compromise personal safety. This design guide illustrates how to transport video traffic over LAN and WAN with IP security encryption, and also implement administrative controls on who has access to the network.

Network Services

One advantage of the any-to-any aspect of IP networks is resource and system access. The Network Time Protocol (NTP) and Syslog messages are examples of network services that IP cameras can request data from and send data to, which are either not available with analog-based systems or are more costly to implement. Additionally, local utilities like Power-over-Ethernet (PoE) and the Cisco Discovery Protocol (CDP) both lower the cost of installation and facilitate troubleshooting.

Network Management

Enterprise networks vary in degree of sophistication and maturity of network management. IP video surveillance, however, is one application that can greatly benefit from a proactive approach to the Fault, Configuration, Administration, Performance, and Security (FCAPS) model. For example, in headless deployments (video feeds that are not actively monitored by a person), the availability and network performance is critical to ensuring quality video recordings. The network management platforms and processes of the enterprise can help the physical security manager in detecting and reacting to an endpoint or network transport issues that could impact video quality.

Integration with Ancillary Subsystems

Physical security is one component of facilities management in many large organizations. Other components include door access control, which is often closely linked with video surveillance as a key component to the safety and security missions. To achieve the goal of a fully-converged network, the other Building Management Systems (BMSs), such as fire alarms, elevator control (to park elevators in the event of a fire), air quality monitoring (carbon monoxide and smoke detection), and lighting and heating/cooling must be able to communicate with the video surveillance systems.

The first step in achieving this goal is to IP-enable these devices and provide the network infrastructure to support their effective communication between systems. For example, if virtualization is enabled on the IP network to support video surveillance, a practical approach is to also include the BMS devices on the same address space, and in the same network segments, as the video surveillance devices. Typically, the bandwidth requirements of BMSs are very trivial to that of video surveillance, the end users of the data often report to the same organization heads and the likelihood of system integration (now or in the future) is high.

Video Data-Mining and Analytics

The end goal of migrating from analog-based systems to IP-enabled video surveillance is to move the application from targeting loss prevention, compliance, safety, and security to obtain a greater business value by increasing sales and reducing expenses and exposure to liability. Data mining is the process of detecting some pattern in data. One video surveillance application can be to analyze video feeds to detect certain colors or articles of clothing to identify groups of gang members among patrons at a shopping mall. Video analytics use data mining techniques to detect patterns in data. Video analytics may be performed at the endpoint (IP camera) on specialized digital signal processors (DSPs) by a third-party analytics vendor, or by servers within the enterprise data center. One application of video analytics is to detect the queue length of checkout lines and inform management to increase or decrease the staffing at cash registers to more fully use staff.

In the future, the analysis output of video data can be more economically valuable than the loss prevention role of video surveillance to many retail organizations.

NTP Considerations

NTP must be configured on the VM and ESXi host. Time sync on VMware tools should be disabled on the VM. More detailed instructions are available in the *Video Surveillance Deployment Guide for UCS B- and C-Series Platforms*.

Performance and Scalability

Table 1, Table 2, and Table 3 list and define the performance and scalability settings.

Table 1 B-Series Blade Servers with External FC SAN Storage

Video Stream Bit Rate	Number of Camera Streams per VM	Number of VMs per Blade
1 Mbps	250	4
2 Mbps	130	4
4 Mbps	65	4

Table 2 C-Series Rack-Mount Servers with External FC SAN Storage

Video Stream Bit Rate	Number of Camera Streams per VM	Number of VMs per Server
1 Mbps	250	3
2 Mbps	130	3
4 Mbps	65	3

Table 3 C-Series Rack-Mount Servers with Internal RAID-based Storage (Total I/O Limited to 250 Mbps)

Video Stream Bit Rate	Number of Camera Streams per VM	Number of VMs per Server
1 Mbps	250	1
2 Mbps	130	1
4 Mbps	65	1

Storage Considerations

Apart from the onboard storage, external storage can also be used by the Media Server. Usable storage is based on the redundant array of independent disks (redundant array of independent disks [RAID]) level used. For information about adding a media partition to a VSM VM, see the *Video Surveillance Deployment Guide for UCS B- and C-Series Platforms*.

Note VSM camera recording cannot span across media partitions. This limitation should be considered when performing storage calculation.

Fiber Channel-based SAN Storage

The guest operating system (OS) is SLES 10 SP1 32-bit, hence, the maximum size of each disk that can be added is limited to 16 TB.

To overcome the VMware limitation of 2 TB per virtual machine disk (vmdk), the storage area network (SAN) logical unit numbers (LUNs) are attached to the VM as Raw Device Mappings (RDMs).

Internal Storage on C-Series Servers

The onboard C-Series disks are subject to the VMware vmdk limitation of 2 TB each. Depending on the available storage (number of hard disks and RAID configuration), multiple disks of 2 TB (or less) must be added to the VMs on the UCS C-Series servers.

For more information on how to add the media hard disk, see the *Video Surveillance Deployment Guide for UCS B- and C-Series Platforms*.

External Storage Considerations

The Cisco Physical Security 4 Rack Unit (4RU) Storage system can be used for external storage. The Cisco Physical Security Storage System 4RU (CPS-SS: 4-RU) and the Cisco Physical Security Storage System 4RU-EX represent the next generation of high-performance, high-density, and energy-efficient storage for the Cisco Video Surveillance System.

Up to 60 disks can be installed in only 4RUs, and the Automated Massive Array of Idle Disk (AutoMAID) technology can reduce energy costs by up to 85%. Innovative active drawer technology enables one person to service the system while it remains in full production. This is a significant improvement over competitive systems that require several people to replace a single disk or fan.

The 4RU storage system delivers an enterprise-class solution with high-density, high-performance storage that is ideal for backup to disk or bulk data storage. It includes:

- Two Fiber Channel ports
- Two 6-Gigabit serial attached small computer system interface (SCSI) (SAS) ports, as well as an Ethernet port for web-based administration
- Up to four 8-Gigabit Fiber Channel ports and four 6-Gigabit SAS ports
- Ability to add more I/O ports to help deliver even more bandwidth and connectivity

The 4RU-EX high-capacity expansion chassis offers 60 drive bays in only 4RUs to provide the most storage capacity in the smallest footprint, with enterprise-class reliability and ease of use. Attaching the expansion chassis to the 4RU storage system can help your organization increase efficiency and capacity while reducing power, space, and cost.

Table 4 Example External Storage Partitioning

Array 1 (10 Disk Array with 2 TB Drives, RAID Level 5)
RAW Storage: 20.0 TB
Array Size: 18 TB
Usable: 16.7 TB

For more information about Cisco Physical 4RU Security Storage, see http://www.cisco.com/en/US/products/ps12501/index.html.

Because the usable disk size is greater than the 16 TB LUN addressable by the 32-bit guest OS, partition the RAID array into three LUNs:

- **LUN0**: 100.0 GB
- **LUN1**: 8.35 TB
- **LUN2**: 8.35 TB

Note **LUN0** can be used to host the VM root partition. **LUN1** and **LUN2** can be used as media partitions in the VSM to store video.

Storage Networking

Storage costs are growing faster than server costs. Businesses need more efficient, cost-effective, and responsive SANs, without increasing overall power consumption. In today's environments that must comply with government regulations for data recovery, 24-hour access to critical information is imperative. Cisco SANs provide:

- **Multiprotocol storage networking—**Achieves lower total cost of ownership (TCO) and greater agility by integrating disparate protocols, such as Fiber Channel (FC), Fiber Connectivity (FICON), Fiber Channel over Ethernet (FCoE), Internet Small Computer System Interface (iSCSI), and Fiber Channel over IP (FCIP).
- **Unified OS and management tools—**Reduces operating expenses and provides operational simplicity, seamless interoperability, and consistent features.
- **Enterprise-class storage connectivity—**Supports significantly higher virtualized workloads, to increase availability, scalability, and performance.
- **Services-oriented SANs**—Extends any network service to any device, regardless of protocol, speed, vendor, or location.

Cisco MDS 9000 Switches: Flexible, Agile, Cost-Effective

With Cisco MDS 9000 switches, you can optimize physical resources, reduce operating costs such as energy use, and decrease capital expenditures to achieve lower TCO.

With industry-leading features and advanced capabilities, Cisco MDS 9000 switches continue to be the leading platform for hosting intelligent multilayer services that fill business needs, such as:

- SAN consolidation
- Storage encryption
- Disaster recovery
- Data mobility
- SAN security
- Replication

The MDS 9000 product family also facilitates server storage and fabric virtualization to provide an end-to-end virtualization solution for the data center. SANs are central to the Cisco Data Center Business Advantage architecture. They provide a networking platform that helps IT departments achieve lower total cost of ownership, enhanced resiliency, and greater agility.

Note Cisco recommends using MDS switches for your SANs. For more information about MDS switches, see http://www.cisco.com/go/mds.

Backup Servers

The VSMS storage backup allows you to configure a secondary server to use if a VSMS becomes unavailable. You can specify any VSMS as a backup server. For more details about adding a backup server, see http://www.cisco.com/en/US/docs/security/physical_security/video_surveillance/network/vsm/6_3_2/user_guide/vsmug.pdf.

Recovery Considerations

To support the recovery of VSM 6.3.2, set aside 65 GB of disk space, 32 GB for the actual VSM VM, and 32 GB for the recovery.

To back up video archives, the system must have at least one VSMS backup server. The VSMS backup server must have a repository for archive backups from other servers. For more details, see http://www.cisco.com/en/US/docs/security/physical_security/video_surveillance/network/vsm/6_3_2/user_guide/vsmug.pdf.

Design Recommendations for Deployment Models

A typical IP video surveillance deployment in an enterprise network consists of one or more campus locations running a Cisco Video Surveillance Media Server, Video Surveillance Operations Manager, and Video Surveillance Virtual Matrix on an Intel-based Linux Enterprise Server OS (Cisco Physical Security Multiservices Platform or Cisco UCS Platform). Deployment on a standalone hardware or UCS C-Series server is targeted at locations with more than 32 video surveillance cameras. Branches that require 1-to-32 video surveillance streams can incorporate the Cisco Services Ready Engine (SRE) with ESXi running a video surveillance virtual image to run the Cisco Video Surveillance Media Server and Video Surveillance Operations Manager software. Branch offices and teleworker locations may view and administer the video surveillance system—as may external organizations connected either through an Extranet or the public Internet through a global IP connectivity and a web browser. Figure 4 illustrates the topology and application services deployed in an enterprise-wide implementation of IP-based video surveillance.

Figure 4 VSM Deployment Models

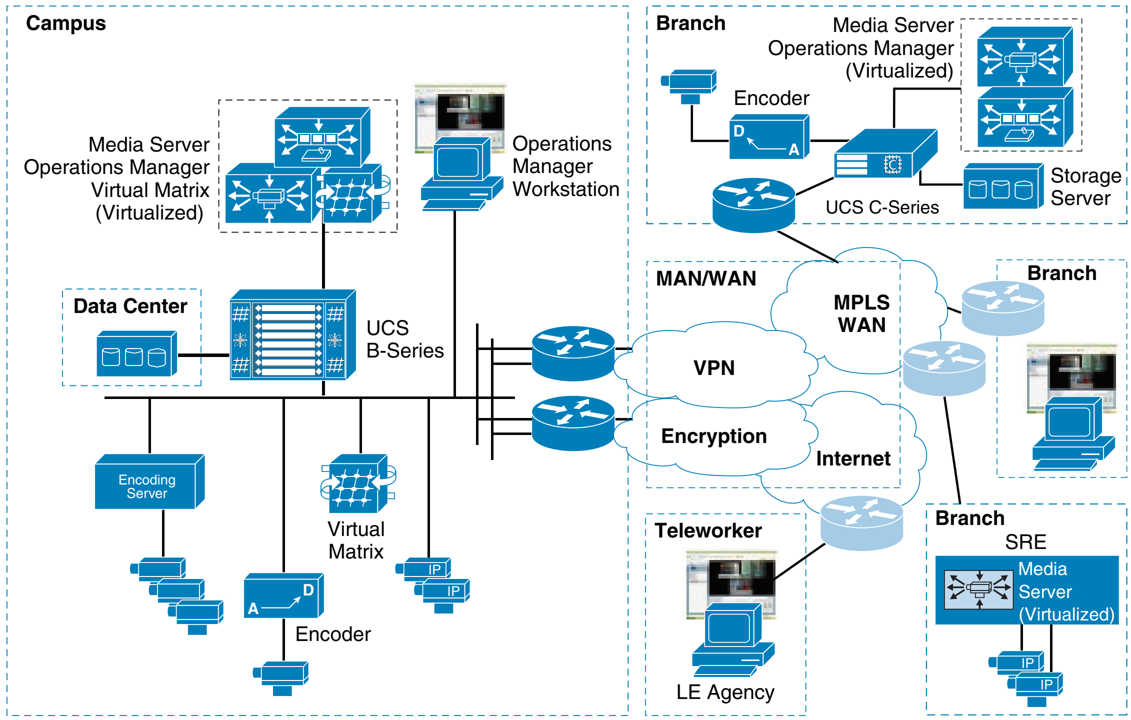

The branch locations are connected to the enterprise campus by WAN technologies, including Metro Ethernet, private line, the public Internet, or a Layer-2 or Layer-3 Multiprotocol Label Switching (MPLS) Virtual Private Network (VPN) deployment. With a Layer-2 MPLS deployments (Pseudowire), IP cameras can be Ethernet-attached to a remote switch and can transport images through the carrier network, provisioned and managed by the Video Surveillance Operations Manager, either at a branch location or a central location. Branches attached through a Layer-3 MPLS network, leased line, or over the Internet can support viewing stations and IP cameras that can be managed by either the campus or branch deployment.

Cisco technologies, such as Dynamic Multipoint Virtual Private Network (DMVPN), can be overlaid on to the WAN transport to provide data privacy and authentication by way of IP security (IPsec) encryption. To ensure the prioritization of voice, video, and mission-critical applications over the WAN, QoS is deployed on the WAN. Where multiple WAN links exist, Performance Routing (PfR) can be enabled to provide intelligent path selection and the ability to route around brownouts and transient failures, thereby enhancing what can be provided by traditional routing protocols, such as the Enhanced Internal Gateway Routing Protocol (EIGRP).

The decision as to whether a specific environment should implement the Cisco Video Surveillance on UCS Express at a branch location and archive data at the branch—or provision cameras off the campus implementation of the Cisco Video Surveillance Manager—depends on the number of cameras, the resolution, frame or bit rate of the camera, quality factors of the cameras, and the bandwidth cost and availability at the remote locations.

Note Cisco recommendeds to deploy VSM on UCS Express only if you plan to record up to 32 streams 1 M or 15 streams @ 2 M or 7 streams @ 4 M. Consider dedicating the entire SRE 9xx for video surveillance. In cases where implementing cameras is the only requirement, it may be practical to transport the camera feeds across the WAN for archiving. However, in most deployments, local storage is necessary due to the bandwidth required and the bandwidth costs. If the branch requires more flexibility with respect to the number of cameras, then we recommend to deploy video surveillance on UCS C-Series servers. A typical enterprise deployment consists of one or more campus locations running the Cisco Video Surveillance Media Server. For enterprise-level deployments, we recommend to use the UCS B-Series servers in the data center.

Solution Characteristics

The following list defines some general solution characteristics for an IP video surveillance deployment:

- An IP network infrastructure is required to link all components.
- IP cameras are under the control of and feed Media Servers. The VSOM interface is the viewing station portal into the video archives and live feeds.
- The amount of disk storage for archiving camera feeds depends on factors that include the retention period requirements, image resolution, image quality, format, and encoding. Storage requirements might be difficult to plan and predict.
- Encryption through IPsec may be implemented between video endpoints to ensure data privacy, integrity, and authentication.
- Virtual Routing and Forwarding table (VRF-Lite), virtual LANs (VLANs), and other network virtualization techniques may be used to segment the video endpoints and servers.
- Viewing stations are PCs running Internet Explorer (IE) with Active-X controls. The PC must have a sufficient central processing unit (CPU) clock rate to decode the video feeds.
- Camera feeds traverse the IP network from the camera source to the Media Server either as Motion JPEG (MJPEG) or MPEG-4.
- MJPEG is typically transported via the Transmission Control Protocol (TCP). TCP provides guaranteed delivery of packets by requiring receiver acknowledgement. Packets that are not acknowledged are retransmitted.
- With MJPEG, each image stands alone; therefore, the images displayed are of good quality.
- MPEG-4 video is typically transmitted over the User Datagram Protocol (UDP), Real-Time Transport Protocol (RTP), or Real-Time Streaming Protocol (RTSP). UDP does not guarantee delivery and provides no facility for retransmission of lost packets.
- UDP transport provides the option of IP multicast (IPMC) delivery; however, it is not universally supported.
- Deploying a video surveillance solution through a WAN environment presents challenges that are not typically seen in a LAN.
- WAN bandwidth is most costly and the available transport types are dependent on the service provider offering available in the geographic area.

WAN Considerations

WAN is used to connect different LANs and typically includes a broad geographic area. WAN services are leased from service providers who provide different speeds and connectivity options.

Figure 5 illustrates how a remote branch office relies on the connectivity provided by a WAN service provider.

Figure 5 Service Provider Network

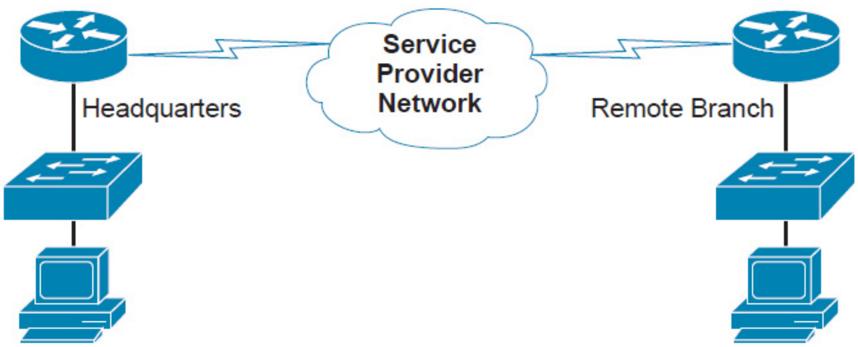

Deploying a video surveillance solution through a WAN environment presents challenges that are not typically seen in a LAN. In a LAN environment, it is common to see 1 Gbps and 10 Gbps of bandwidth, while in a WAN environment, most connections are less than 10 Mbps; many remote connections operate on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

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

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

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

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

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

Example 1—Network Bandwidth Usage

Figure 6 illustrates a simple scenario with two sites. Each site has a Media Server and each is the direct proxy for an IP camera. Three Operations Manager (OM) viewers are active in Site A and each IP cameras generates 1 Mbps of network traffic.

Figure 6 Network Bandwidth Requirements

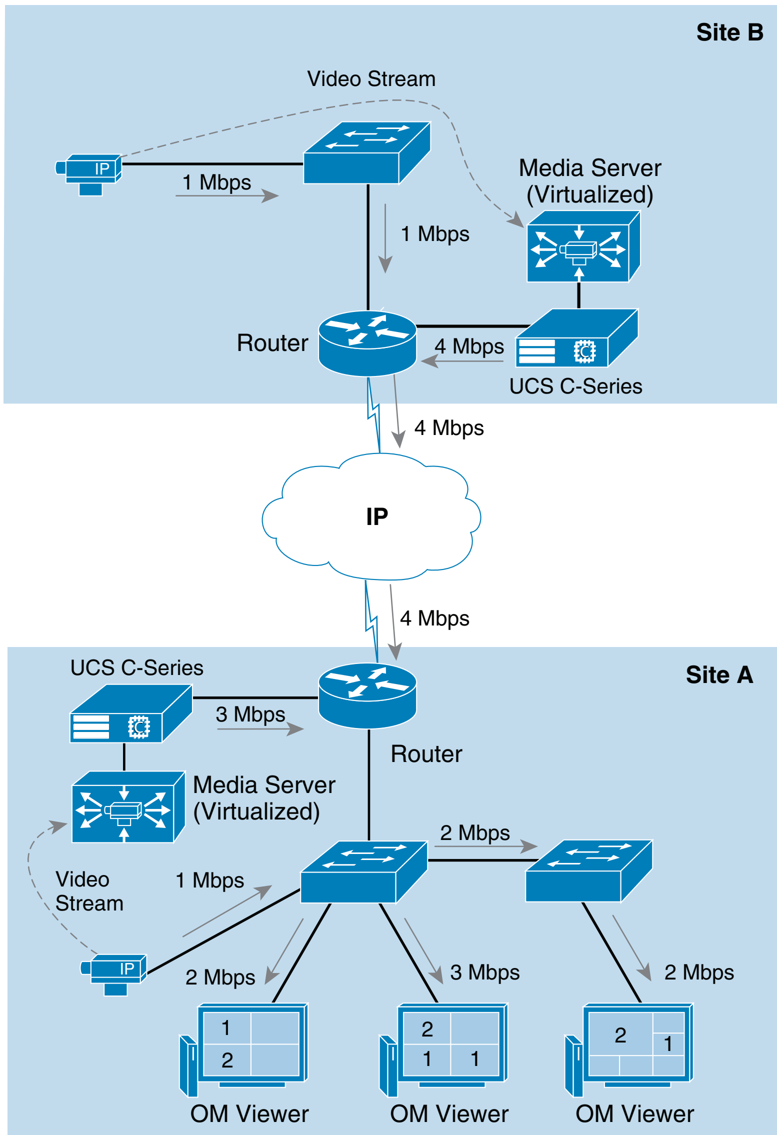

Note For simplicity, the Operations Manager has been removed from this graphic.

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

Example 2—Sites with Remote Storage

Figure 7 displays how Media Servers can be deployed at different WAN locations to minimize the bandwidth requirements. By deploying the Media Servers close to viewers and edge devices, the network traffic remains local to each site. Archiving video streams at each location is also an attractive solution to minimize the network traffic between sites.

Figure 7 Sites with Remote Storage

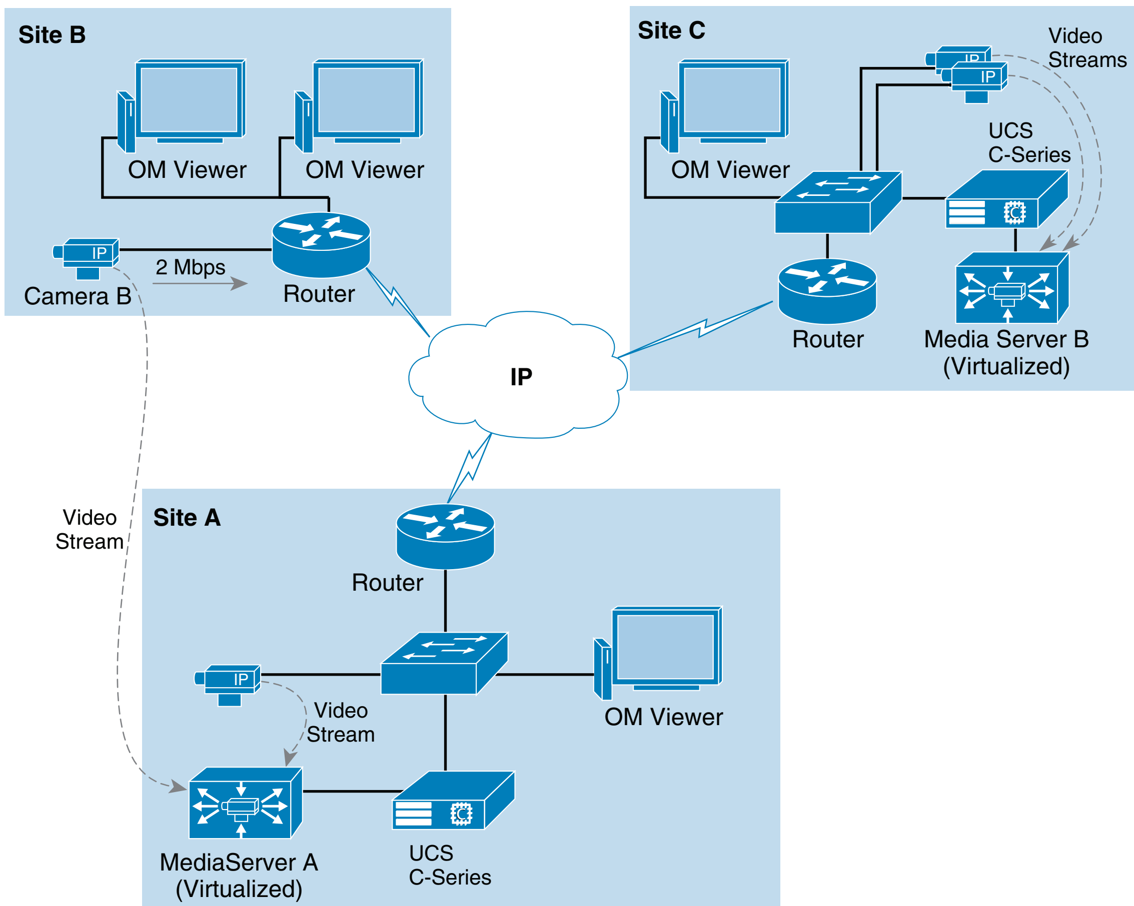

In this example, Site A and Site C have Media Servers acting as direct proxies and archives for the IP cameras. Because both sites archive and distribute video to the OM viewers locally, the network traffic remains local to each site.

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

Example 3—Virtual Matrix Scenario

Figure 8 displays an example that includes a Virtual Matrix server and virtual machine (VM) monitors located at two different sites. The server on Site A acts as the Media Server, Operations Manager, and Virtual Matrix for the environment. To reduce bandwidth traffic, Media Servers are also installed on Site C and Site D.

A single Operations Manager and a single Virtual Matrix are adequate to support this scenario. Because the cameras are located on Site C and Site D, they are able to serve the local OM viewers at those sites. The Media Server on Site A can also be configured with child feeds that come from the remote media servers and provide those feeds locally to viewers and monitors on Site A.

Figure 8 Virtual Matrix Scenario

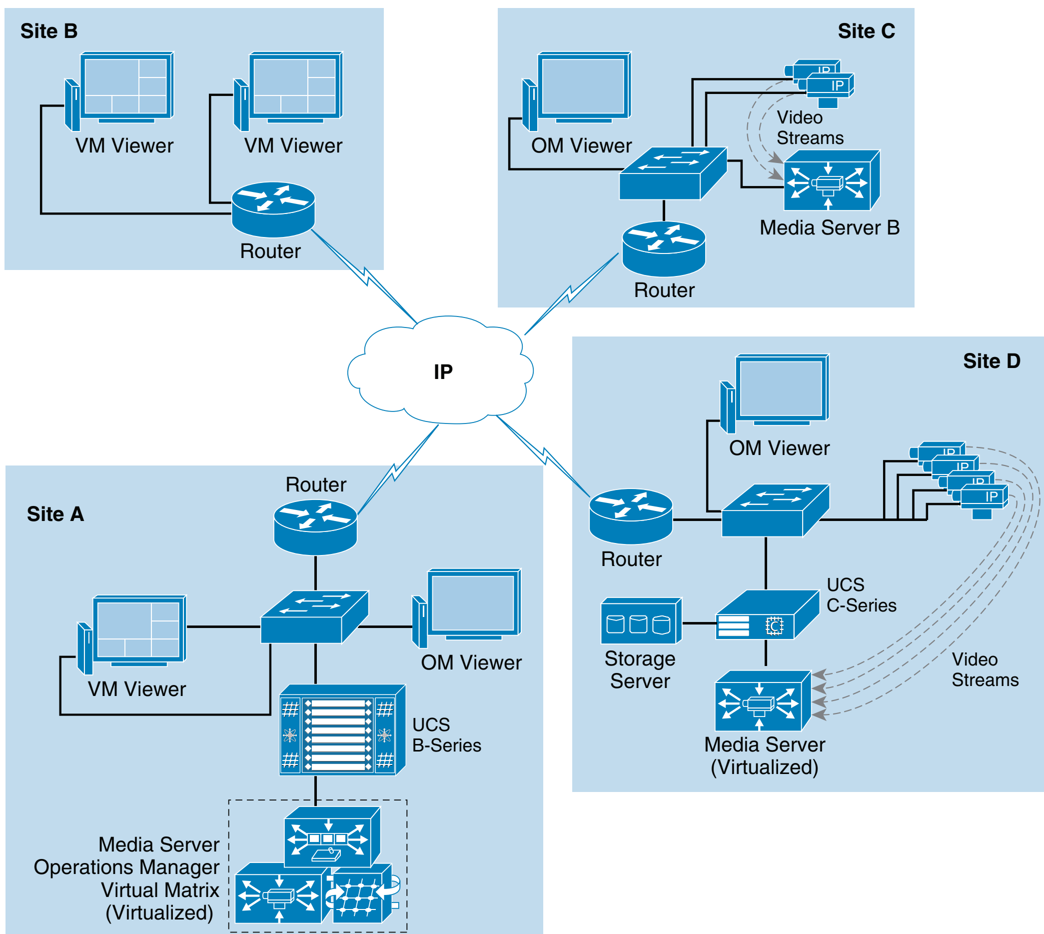

Example 4—Distributed Media Servers

Figure 9 displays a deployment with several remote sites, each with a local Media Server acting as the direct proxy and archive for local IP cameras. In this scenario, all recording occurs at the remote sites and live video streams are viewed by OM viewers and VM monitors (video walls) at the headquarters.

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

Figure 9 Distributed Media Servers

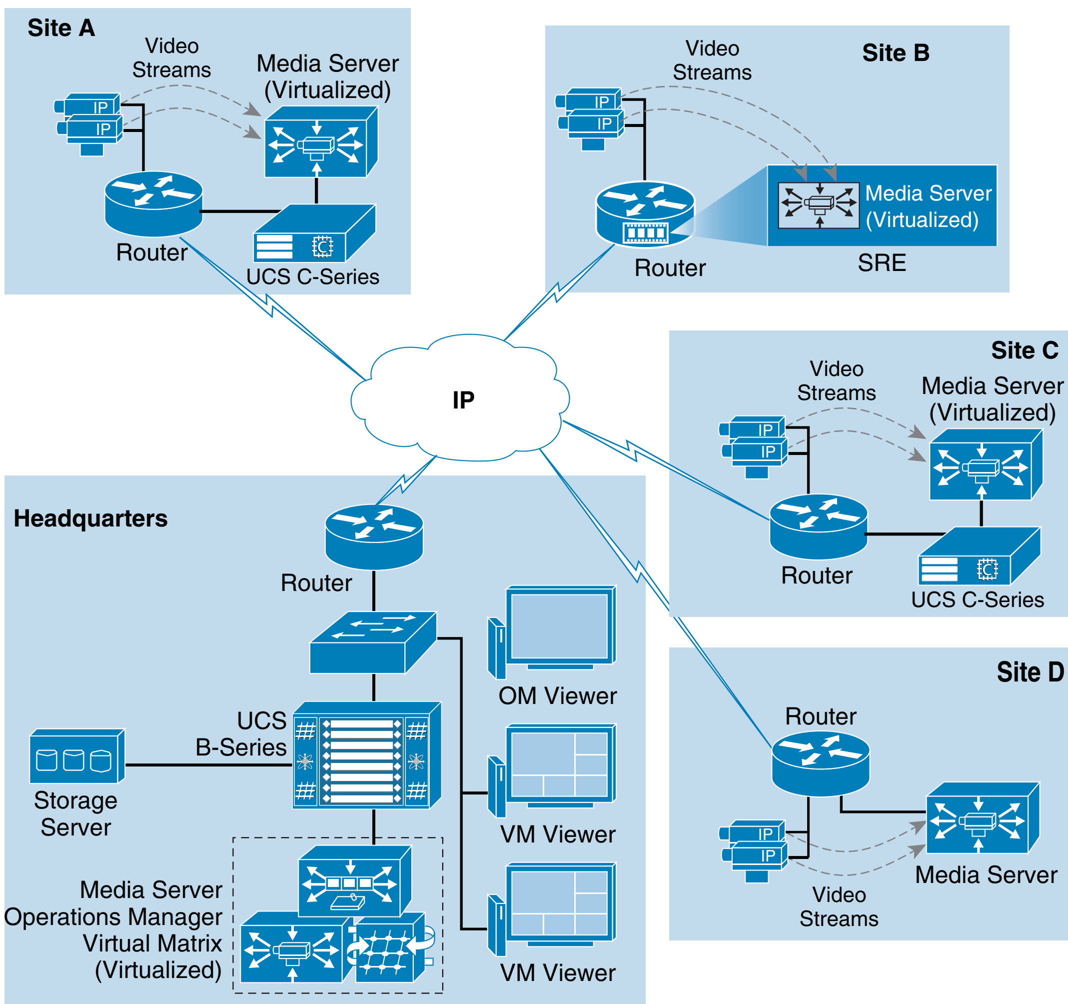

Design Checklist

The following design checklist facilitates pre-implementation planning and the decision process:

- Estimate the number of IP cameras required at each location.
- Using a floor plan or exterior survey, determine cameras that can be powered by Power over Ethernet (PoE) and those requiring power supplies.
- Survey existing IP or analog cameras and determine if these cameras are to be replaced or migrated.
- Estimate the codec, resolution, and frame rate or bit rate requirements for cameras at each location.
- Determine the retention period requirements for cameras at each location.
- Survey existing LAN switches for necessary features and available capacity.
- Based on the number of cameras per location, determine server requirements.
- Using the estimate on the number of servers required, calculate the storage requirements for video archives based on the retention period analysis.
- Analyze the IP addressing requirements and VLAN assignments for IP cameras, Media Servers, routers, switches, and other systems.
- Determine if suitable Network Time Protocol (NTP) sources exist in the current network.
- Investigate what network management servers and software are currently available for services, such as Syslog and SNMP traps and Trivial File Transfer Protocol/File Transfer Protocol (TFTP/FTP) for firmware download and storage.
- Consider implementing network management servers for performance, fault, and capacity planning such as the CiscoWorks Internetwork Performance Monitor (end-to-end network performance), Cisco Secure Access Control Server for Windows (authentication: Terminal Access Controller Access Control System+/Remote Authentication Dial-In User Service [TACACS+/RADIUS] server), CiscoWorks Device Fault Manager (reporting of device faults), and CiscoNetFlow Collector (NetFlow analysis for capacity planning).
- Analyze the existing QoS policies and configuration on routers and switches and incorporate the IP video surveillance requirements into the plan.
- Determine requirements for external users to access video feeds. Analyze what level of encryption or access-control is required to meet the end-user requirements and to align with the corporate network security posture.
- Discuss with the physical security manager and network manager the need for segmentation features such as VRF-Lite, VLANS, firewalls, and access lists to limit access to end nodes.
- Determine the inherent redundancy in the existing network and develop a plan for meeting the physical security needs in the event of a line card or access switch failure.
- Consult with the physical security manager to determine the live viewing requirements. Determine what cameras must be viewed live and the viewing stations locations in the network topology.
- Determine the existing staff skill set and estimate training requirements for physical security installers, operators, and managers in basic internetworking. Consider involving the network staff in day-to-day operations of the physical security operations staff.

Pre-Configuration Work

- Prepare a high-level network topology diagram that displays the place of VSMS, VSOM, and Video Surveillance Virtual Matrix (VSVM) servers in the network, network link speeds, and network connectivity between various sites.
- Note the serial numbers for the servers, cameras, and encoders.
- Document the following data for the VSMS server (this information is used to configure the new server):
  - IP address
  - Hostname
  - Domain Name System (DNS) name
  - Domain name
  - Time zone
  - NTP server
  - Serial number
  - New Linux credentials for the root user
  - New virtual server management console (VSMC) password
- Document the following data for each camera/encoder (this information is used to configure the cameras on the VSMS server):
  - IP address
  - Credentials to use
  - VSMS server
- Consult with the customer about creating naming conventions for servers, cameras, encoders, regular and loop archives, camera groups, views, monitors, events, and user roles.
- Discuss with the customer expected feed configurations. In a lab environment, configure the camera feeds in VSOM (for all expected feed configurations) and show the video to the customer.

  Note Changes in the feed configuration affect the actual video and disk space utilization.

- Document expected feed configurations for all cameras: camera type, resolution, media type, and frame/bit rate, and so on.
- Document configurations all archives: type of archives (loop vs. regular), duration, expiration, and so on.
  - [Best Practice] The archive expiration should be greater than or equal to its duration.
  - [Recommended] Do NOT use “Never Expires” as expiration value.
  - [Recommended] To back up the archives via VSOM, ensure that the loop archive duration is at least two days. Do NOT configure one-day loop archives.
- Identify basic use cases for a configuration:
  - Use of archives (duration)
  - Motion detection events
  - Device triggers, soft triggers

- Views (static vs. rotating)
- Users, roles, and permissions
- Schedules
- Archive backups

Determine Bandwidth and Storage Estimates on each VSMS

- Calculate per-camera bandwidth estimates based on the expected feed media type, resolution, and frame/bit rate information.
- Calculate bandwidth requirement estimates per Media Server based on expected number of feeds in the Media Server.
- Calculate storage requirement estimates per Media Server based on expected number of archives in the Media Server, their expected duration and feed information. Include 7-minute loop archives in motion detection event setups.
- Based on the bandwidth and storage estimates, determine if the VSMS hardware can handle the expected proxy and archive configuration. Consider future expansion plans, if any.
  - [Recommended] Consider disk space requirements for server side clips.

    Note Even if the clip duration is (for example, 15 seconds), it occupies 5 to 10 minutes worth of disk space.

  - [Recommended] Consider disk space requirements for archive backups.
  - [Recommended] Consider future expansion plans and disk space requirements for new feeds and archives.

Network Links

- The VSMS, VSOM, VSVM servers, and VSM Client workstations are connected to a 1 Gb full-duplex, network link. Verify the link speed and duplexity from switch, server, or Client workstations.
- The IP cameras and encoders are connected to a 100-Mbps full-duplex link.
- The cable distance between switch port and servers, client workstations, IP cameras, or encoders is within the maximum allowed distance for the link speed and cable type.
- The IP cameras, encoders, servers, and client workstations are connected to individual switch ports; they are not connected to a hub.

Network Reachability

- The IP cameras, encoders, servers, client workstations are reachable on the network.
- The VSMS servers are reachable from the VSOM server, as well as client workstations using both hostname and IP address.
- The VSOM server is reachable from the VSMS servers using a hostname and IP address.

More Information

For more information about Cisco-related products, see the following resources:

Cisco Physical Security product information:
http://www.cisco.com/go/physec/

Cisco UCS Manager Configuration Guide:
http://www.cisco.com/en/US/products/ps10281/products_installation_and_configuration_guides_list.html

Cisco UCS B-Series Blade Servers Data Sheet and Literature:
http://www.cisco.com/en/US/products/ps10280/prod_literature.html

Cisco Video Surveillance Media Server Software Install and Upgrade Guides:
http://www.cisco.com/en/US/products/ps9152/prod_installation_guides_list.html

Cisco IP Video Surveillance Design Guide:
http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/IPVS/IPVS_DG/IPVS_DG.pdf