Cisco Systems IP Network-Centric Video Surveillance

Overview

Video surveillance has been a key component of many organization's safety and security group for decades. As an application, video surveillance has demonstrated its value and benefits countless times by:

- Providing real-time monitoring of a facility's environment, people, and assets
- Recording events for subsequent investigation, proof of compliance / audit purposes

As security risks increase, the need to visually monitor and record events in an organization's environment has become even more important. Moreover, the value of video surveillance has grown significantly with the introduction of motion, heat, and sound detection sensors as well as sophisticated video analytics. As a result, many nontraditional groups have also found value in video monitoring and recording. In transportation, video surveillance systems monitor traffic congestion. In retail, video can be helpful in identifying customer movement throughout a store, or serve to alert management when the number of checkout lines should be changed. Some video analytics packages even offer the ability to identify a liquid spill and generate an alert enabling faster response by custodial services, thus avoiding a slip and fall situation. Product and package shipment operations can use recorded video to help track and validate the movement of cargo and help to locate lost packages. Additionally, video surveillance can be integrated with and complement access control policies, providing video corroboration of access credential use.

Video surveillance has evolved not only in its application, but also in its deployment. This paper reviews the evolution of video surveillance, including the emergence of the fourth generation of video surveillance systems. These systems are realized through an open, standards-based,

IP-network-centric functional and management architecture. As a network-centric company, Cisco Systems® has enabled the migration of many applications and systems onto a converged infrastructure, such as mainframe computing and telephony. As a global enterprise organization, Cisco® has developed and adopted a network-centric system architecture that meets the extensive requirements for a world-class video surveillance system. Hence, Cisco is in a position to offer recommendations on how to build third- and fourth-generation video surveillance systems. This document shares the business case for adopting such architectures.

The Cisco video surveillance architecture provides several benefits:

- Increased reliability
- Higher system availability
- Greater utility (any camera to any monitoring or recording device for any application, anywhere)
- Increased accessibility and mobility
- Multivendor video surveillance system “best of breed” interoperability
- The ability to enhance other building management system capabilities through improved interoperability

- Reduced capital and operational expenditures

Video Surveillance, One of Many Silo'd Building Management Systems

Video Surveillance, similar to several other Building Management Systems, has typically been deployed as a silo'd system that is separately designed, procured, supplied, deployed, and supported by its own application/user group. These disparate systems do not communicate with each other, despite synergistic opportunities. As a result, system owners and operators suffer from a lack of operational consistency, interoperability, and capability that result in higher capital and operational costs as well as limit the return on their system investments.

Consider the amount of capital and operating expenses for an organization that needs to deploy and maintain independent, disparate video surveillance, access control, intrusion and alarm, HVAC, power, lighting, and life safety systems in each of its many facilities (Figure 1).

Figure 1. Disparate Building Management Systems

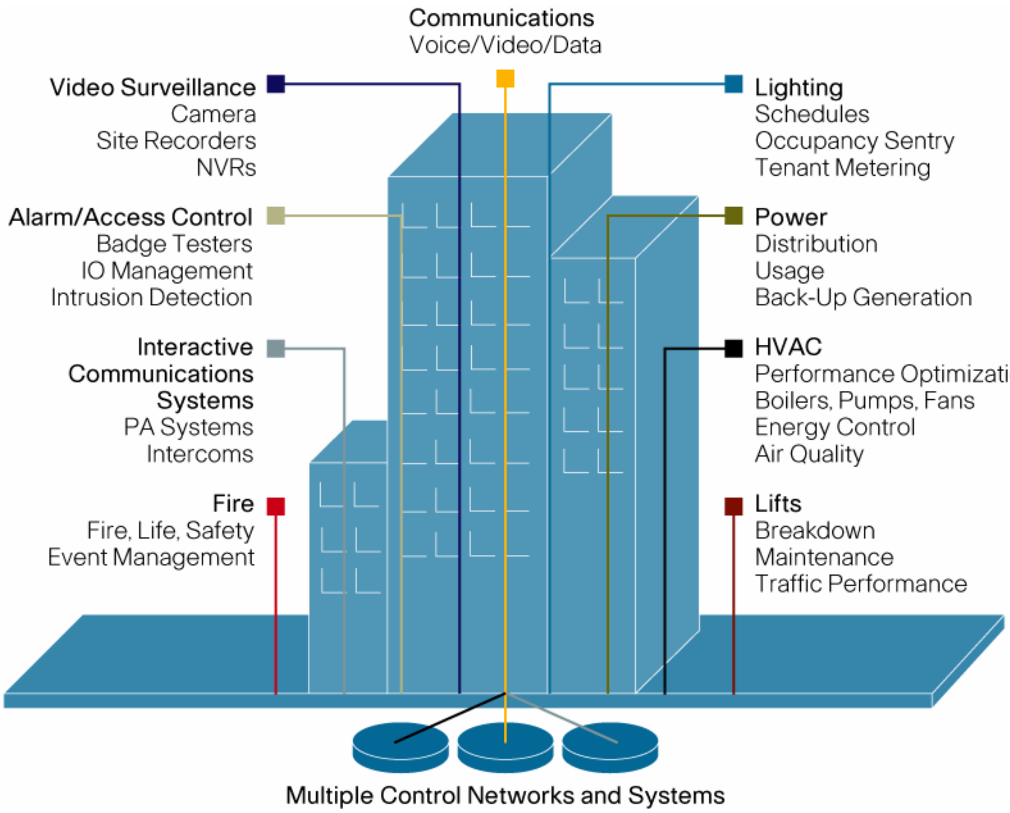

Each of these systems has its own cable plant, proprietary communications protocols, and management platforms. Interoperability between disparate systems is not available—or it requires an expensive and lengthy customization process.

Baseline Video Surveillance Functions

Typical video surveillance systems, especially those found in higher-security environments such as airports, casinos, military sites, correctional facilities, and many corporate headquarters, have the basic system component functions shown in Figure 2.

Figure 2. Analog Video Surveillance System

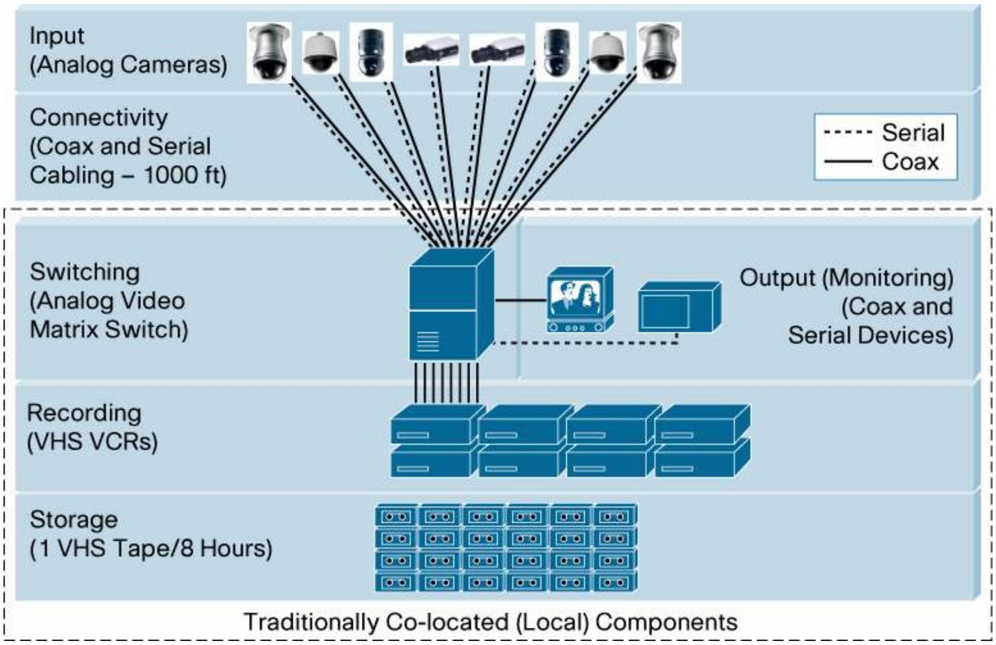

- **Input devices:** Cameras (fixed and/or pan tilt and zoom [PTZ]) that are available in either black and white or color are covertly or openly deployed.
- **Connectivity:** Multiple, parallel cable plants are necessary to deploy video surveillance solutions: coaxial cables for NTSC/PAL video transmission; low-voltage, ring-oriented RS-485 cable plants for serial PTZ command and control; and dedicated fiber-optic cabling for video transmission and PTZ command and control between buildings in a campus setting. Some installations convert the coaxial media to and from unshielded twisted pair (UTP) cabling or use wireless transmission systems, such as microwave.
- **Switching (video stream management):** Real-time central command and control monitoring of video streams is provided. Monitoring station personnel can direct a transmission or aggregation device (commonly referred to as a matrix switch) to switch from one camera feed to another in order to display the scene on a monitor. Switching is traditionally supported by matrix switches that come in many sizes, scaling from tens of cameras to thousands of cameras.
- **Monitoring:** Viewing live video. Operators select the desired video feed and specify where the video is to be displayed. For larger installations, a special-purpose keyboard controls which camera video feed is displayed over an RS-232 connection that sends vendor-specific or proprietary commands to the matrix switch. The requested video stream is delivered to the monitor over a coaxial connection that supports the analog (NTSC/PAL) video signal. Unlike a typical PC keyboard, the layout and operation of the video surveillance keyboard is specific to the video surveillance market. This special-purpose keyboard references cameras by simple numbering schemes (01 = camera 1, 104 = camera 104, etc.). In some installations, PCs can be used instead of special-purpose keyboards and displays but many operators prefer the special purpose monitoring stations and keyboard/joystick controls.
- **Recording:** Independent from monitoring functions, recording has been historically accomplished using videocassette recorders (VCRs) or, more recently, digital video recorders (DVRs).

- **Storage:** Based upon regulatory and other organization requirements, recorded video may be archived for a few days, weeks, or months. This facilitates the investigation of events that may have occurred or need to be correlated with other events.

Most manufacturers of cameras, fiber-optic transmission equipment, matrix switches, and monitoring keyboards have their own proprietary communications protocols and languages to interconnect these systems. This approach has locked the customer into a single-vendor solution, increasing equipment costs and decreasing the customer’s ability to pick best-in-class solutions.

Video Surveillance System Evolution

With proprietary solutions, the video surveillance market has seen limited innovation compared to the open, standards-based IT industry. However, various technologies have been so compelling in their ability to solve significant video surveillance user challenges, such that this field has begun to evolve, catching up with many other systems and applications.

Figure 3 shows second- and third-generation video surveillance system deployments, which incorporate newer methods of video recording and storage.

Figure 3. Video Surveillance Evolution

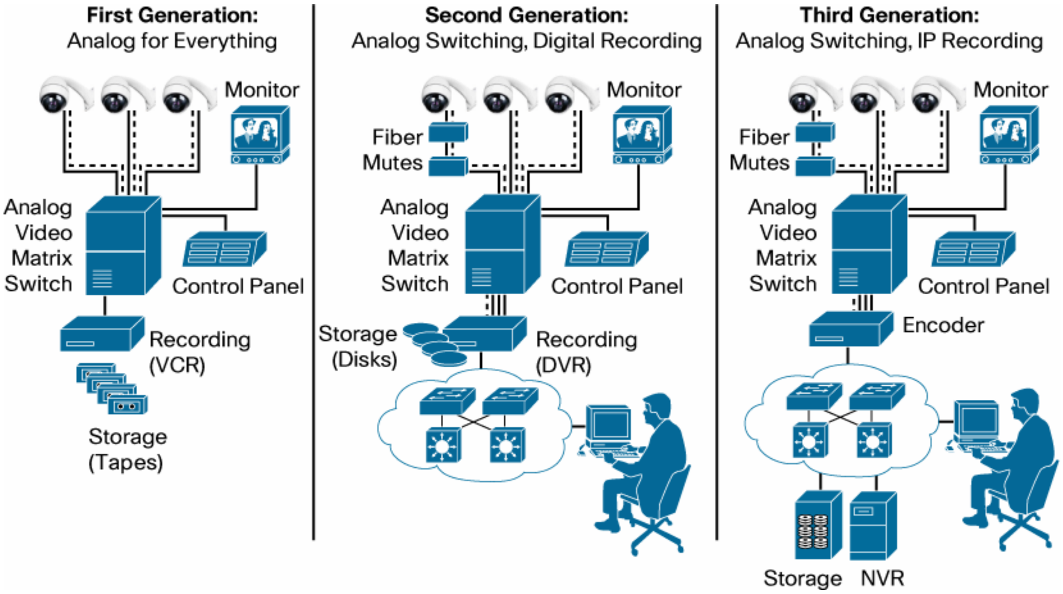

First-Generation Video Surveillance

First-generation video surveillance systems are entirely analog. Cameras are controlled and transmitted video in an analog format. These video streams are aggregated, switched, and dispersed to monitoring displays using analog matrix switching technology. The matrix switch also provides the video stream to analog VCRs for recording purposes.

While these analog devices provide basic monitoring and recording capabilities, they do have several operational drawbacks. VCR-based recording does not facilitate simultaneous recording and playback of video; separate record and playback (review) components are required in order to record video during the investigation process. Moreover, the recording process is prone to human error: replacing blank media or ensuring that recording was activated, for example. From a reliability and system availability perspective, any failure of the recording system can go undetected for an extended period.

Storage and access are also issues. Because videos can be required for future investigation, tapes must be manually stored and indexed unless used in a jukebox type VCR device. These consume a significant amount of size and power, and generate quite a bit of heat.

The viewing of live or recorded video is limited to specific operations and investigation centers. To review recorded video from a remote location requires the appropriate tape to be located and sent to the investigation center.

In virtually all cases, video surveillance system operations are based on proprietary signaling and format protocols; best-in-class multi-vendor component interoperability is not an option for video surveillance customers without extensive and costly customization.

Second-Generation Video Surveillance

Second-generation systems are also based on analog camera (input), fiber or coaxial connectivity, with video switching provided by an analog video matrix switch. However, recording functions are enhanced.

Second-generation systems primarily focus on addressing recording and storage problems. DVRs replace analog VCRs. DVRs convert the analog video feeds into a digital format and save the resulting digitized video on internal hard disk drives or on locally direct attached storage (for example, digital tapes, disk drives, or DVDs). Thus, many manual efforts associated with VCRs are eliminated or reduced in frequency. Additionally, the DVR’s internal database reduces video retrieval time during investigations.

While DVRs offer longer operation life than VCRs, they can pose recording system availability problems. In the event of a DVR failure, the DVR has to be replaced, generally resulting in a loss of video unless an N+1 redundancy system was offered, such as the failover capability enabled in Cisco Systems’ Stream Manager Storage Administration Failover Software. Some DVRs use personal computer operating systems which can be subject to tampering and virus propagation; thus, DVRs should be included in a prophylactic maintenance program with regular virus protection and security mechanism configuration. Moreover, since many DVRs tie the software-based video stream/storage management value into a hardware- (vendor-) specific platform, a generic server/storage device with a considerably lower price may not be available.

Frequently, the DVR software is accessed and controlled by a vendor-specific user interface often running as a set of administrator and operator applications on a personal computer (PC). As such, second-generation DVRs frequently require a PC viewing client. The use of client software offers some trade-offs; it can limit access to recorded video on a local basis, which may be desirable, but it also can impose problems in emergency situations where remote viewing, over an IP network may be helpful.

Some DVRs can be accessed via a network-connected PC to further reduce the time associated with video archiving and retrieval. On-demand access to archived video accelerates evidence review and improves evidence control. It also saves time and effort; investigators do not have to travel to other facilities to perform investigations. To preserve remote-location WAN bandwidth, the video can be pulled over the network on an on-demand basis.

Third-Generation Video Surveillance

As with first- and second-generation video surveillance deployments, third-generation deployments are primarily based on analog camera (input), fiber or coaxial connectivity, and video switching is provided by an analog video matrix switch. However, accessibility of live and recorded video is enhanced.

As observed, second-generation DVRs typically require video to be viewed by PC, which affects video surveillance operator efficiency. Some vendors, including Cisco, offer IP-to-analog video gateway decoders (IP gateway decoders) as part of a third-generation video surveillance solution that allows operators to view recorded video from their analog monitoring stations. By using familiar video surveillance PTZ joystick controls, operators can select the video associated with a specific camera, rewind the video, and review it over analog monitors. This enables faster response and investigation of events, eliminating the need for a PC and the associated delay. Moreover, in multi-display environments, the operator can continue to monitor other camera video while investigating a recorded event. Cisco decoders provide a higher degree of vendor interoperability; operators can use their preferred keyboard and joystick from one vendor while viewing video using another vendor’s cameras. Cisco IP gateways provide any-to-any vendor interoperability, and protect investments in analog video surveillance cameras and monitoring stations.

Many third-generation systems frequently unbundle the DVR; discrete encoders or high-density, rack-mountable, chassis-based encoders provide the conversion from analog to digital and use the network to a greater extent. Thus recording becomes a separate function from video digital encoding.

Encoders serve as analog-to-IP gateways and as a connection point to the network. The IP network transports the video streams to monitoring and recording locations. Encoders digitize analog video; typically, they compress the digital video using various compression algorithms, including the same ones used for production-quality motion picture DVDs, and transmit the compressed digital video over a frame-based (Ethernet) or packet-based (IP) network.

Some encoders, such as Cisco IP Gateway Encoders, provide additional features that allow them to operate with a wide variety of analog cameras. This gives video surveillance operators more control over their analog vendor camera selection process by offering a greater degree of multi-vendor keyboard/camera interoperability. This aspect becomes even more important when PTZ cameras are used, many of which have proprietary camera control signaling.

Encoders can also be differentiated by the latency induced by the digitization and compression algorithm implementation. The lowest-latency, high-video-quality encoders generally have less than 200 ms of latency. A lag of more than 200 ms can be problematic for video surveillance operators using PTZ cameras—they commonly overshoot the intended item to monitor (zoom in too far or pass the given object). Cisco IP Gateway encoders are hardware- and digital signal processing (DSP)-based platforms with high-quality, low-latency compression algorithm implementations; latencies are negligible for most operations environments.

Another benefit from unbundling the DVR and using encoders is that the recording (stream and storage management) function, sometimes referred to as a network video recorder (NVR), can be fully independent of storage. The NVR can be located anywhere on the network, often in the data center with other server systems. Moreover, the NVR software can run on lower-cost, common off the shelf (COTS) servers.

In first- and second-generation deployments, surveillance cameras must be within 1000 feet of the recording device when connecting over coaxial, or require fiber connections for longer distances. Now that encoding occurs in a separate device, the NVR can be located anywhere on the network—at an organization's headquarters, for example, or using servers in two data centers—to simplify management and increase availability. Physically separating the encoding device from the server has another advantage as well: the server no longer needs to devote compute cycles to managing video cards and compression. Cisco corporate physical security operations personnel observed that by moving to third-generation NVR technology, each server can manage 32 cameras compared to the 8 to 16 they previously managed, reducing server hardware requirements from more than 300 to 172, or by 40 percent.

Many organizations have resorted to maintaining a separate database for each remote DVR. However, when using NVRs that can be deployed anywhere on the network, it is possible to centralize the closed circuit television CCTV database into fewer distinct geographic database environments that can be replicated back to the organization's central safety and security operations center. This partitioning and semi- centralization of databases further simplifies video surveillance system management and reduces equipment costs.

This partitioning of system functions also helps improve operational efficiency. An organization's IT group can be tasked with the responsibility of maintaining the video surveillance servers and storage, as well as protecting them along with other mission-critical servers. This allows security personnel to focus on security issues, not maintenance of storage devices. As a result, it is possible to reduce not only redundant capital infrastructure investment by using the network for transport and access of the video, but also optimize operational roles and responsibilities.

It should be recognized this model of separate but complementary functional responsibilities is quite common in most organizations today. For example, human resources is responsible for personnel issues, yet uses the power and flexibility of the IT-supported network to run networked human resource applications. The same is true of other mission-critical and business-sensitive applications, including finance, engineering development, and sales. The organization's IT group ensures that edge devices are properly connected to the network and servers are properly maintained (including virus protection), and provides constant monitoring of these networked assets. Moreover, IT works with these user groups to ensure appropriate security policies are in place to provide appropriate access to restricted resources.

NVR deployments offer several other advantages compared to second-generation deployments using DVRs. Recording and storage component availability is further increased—the failure of a storage device can be almost instantly remedied by having the NVR direct the video stream to another network-connected server or storage device. The use of superior long life (higher MTBF) storage devices also helps increase video surveillance system availability.

As mentioned, NVRs offer both video stream management and video stream storage management. Storage management can be an important factor for users with high 24-hour "record everything" storage requirements. NVRs that can prune stored video based on motion or other criteria (i.e. first in first out) can further minimize regular maintenance tasks and potentially reduce the amount of storage needed to meet long-term retention requirements.

The ability for the NVR to ingest IP video also enables IP camera video to be recorded in addition to the video coming from analog video encoders. While IP cameras are discussed in greater detail in other Cisco documents and whitepapers, it should be recognized that IP cameras offer several advantages to analog cameras / and analog encoders. These benefits include:

- Compact, single video capture form factor (as compared to an analog camera plus an encoder)

- No separate power source required when Power over Ethernet is provided by the IP network switch, which in many cases has battery back-up in the event of a power failure
- Ease of deployment using wireless LAN technology
- Lower cost deployment using Category 5 structured cabling

One cautionary note: many current video surveillance component and system vendor “network-connected” products tunnel their proprietary communications over Ethernet to maintain a strong, single vendor-lock on customer deployments. These components cannot harness the intelligence or true interoperability of the standards-based IP network infrastructure.

Fourth-Generation Video Surveillance

Collapsing video switching functions onto an existing Ethernet switched environment further reduces the complexity and lowers the cost of deploying video surveillance. It also provides video surveillance system owners with the flexibility to design solutions tailored to their unique requirements. Furthermore, as part of an open network, operators can create policies allowing the inherent value of the video, as a source of information, to be used by other safety and security applications, as well as other nontraditional business applications.

Cisco’s own experience in converging various applications and technologies on the IP network and in operating a large global video surveillance system over the network shows that it is feasible to adopt such deployments. Fourth-generation video surveillance provides additional benefits and advantages over preceding generations (Figure 4). It expands and extends the capability of video surveillance gateways (enhanced encoders and decoders) and the NVR, which allows the matrix switch to be replaced by standard and typically lower-cost Ethernet switching platforms.

Figure 4. Fourth-Generation IP Network Centric Video Surveillance

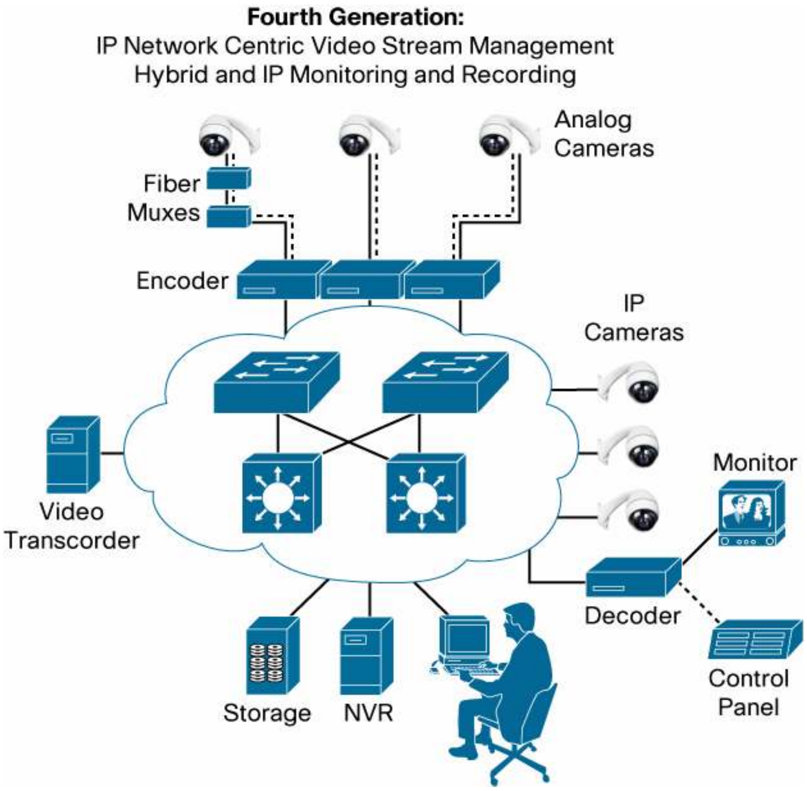

When used with PCs for monitoring and reviewing video, some NVRs offer matrix switch-like functions, allowing the matrix switch to be eliminated. The switching is provided by the network infrastructure with the video stream management provided by the NVR. Without the matrix switch, encoders can be either centralized in multiport configurations to support home-run cabling schemes or located closer to the camera. By situating the encoder closer to the camera, the encoder can use the pervasive IP network cabling infrastructure, further reducing the cost of redundant cabling infrastructures.

For pre-existing long-range camera deployments, fiber multiplexers and distribution amplifiers can continue to coexist with encoders. Of course, for new deployments it may be possible to eliminate the need for fiber multiplexers and distribution amplifiers entirely, thereby further reducing deployment costs.

Cisco Stream Manager software supported on Cisco video surveillance IP Gateway encoders and IP Gateway decoders as well as Services Platforms can provide true matrix switch functionality that supports not only analog cameras and PCs but also the highly specialized video surveillance keyboard controllers and monitors. This true matrix switch capability provides full multivendor best of breed mix and match interoperability.

Alternately, Cisco Video Surveillance Manager (VSM) software which runs on standard computer servers running the Linux operating system provides a browser-based user interface to collect, manage, record / archive and distribute video from multiple third party video encoders and IP cameras. The Web interface enables operators and other users to easily access live or recorded video using a PC, or various other browser equipped devices. As a result, the video can be viewed in remote and mobile environments. Additionally, Video Surveillance Manager allows for easier integration with other network applications including third party command and control software. Additionally, much like analog based systems, VSM video can be directed to digital video walls based on various pre-defined events.

Video surveillance system gateways convert or translate proprietary vendor-specific video signals and formats into a common format and then to the same or other vendor-specific formats. This level of interoperability provides the ability to share video information with other systems via the common format. For example, this enables the integration of video surveillance with access control and intrusion detection without the need for a centralized server, which would represent a single point of failure. The ability to integrate, or unify, the surveillance system with alarm systems would increase the effectiveness of security operations personnel and cut the expense of responding to false alarms. This unification would also reduce the false alarm rate, which exceeds 90 percent in most organizations. For instance, with video surveillance, a security officer could determine that the source of a “door-forced alarm” was a gust of wind, not an intruder.

A common format for video and control signals that is transmitted across the IP network also provides the ability to add new functions such as video analytics anywhere in the network. Video analytics offer the ability to automatically monitor surveillance video for violations (that is, package left alone, presence after building closure, going in the wrong direction). Therefore, it is a tool for prevention and early detection. No matter where this function is deemed “best” to run—at the edge of the network, embedded in the camera, in the encoder, or centralized in monitoring center—a common format enables the same video analytics program to be used or varied based on specific circumstances. Most video analytics vendors suggest the optimal frame rate for current programs is no more than 15 frames per second. This frame rate is excessive:

- There is very little difference in the movement of an object between frames sent and analyzed every 1/15 of second versus 1/30 of a second.
- Given current server CPU performance, it would be more desirable to analyze more video streams than to waste CPU cycles on images with very little difference between them.

When the NVR also supports a Web-browser-based graphical user interface and is sometimes complemented with video transcoding capabilities, video surveillance monitoring and reviewing can be become highly mobile activities that provide first responders and their centralized safety and security command infrastructure with an unparalleled level of information and collaborative capabilities.

Video Analytics

Conventional analog CCTV systems rely solely on human operators to identify events and determine what action should be taken (further investigation, generate an alarm, etc.). U.S. Army studies of video surveillance operators found that a significant number of actionable events were not identified by operators after watching as little as a few hours of video. The efficacy of real-time monitoring of video is greatly diminished if key events are missed.

One of the more interesting and promising benefits of digital video is the ability for computer processing and analysis of video, also referred to as video analytics. Much of the early research and application of computer analysis of video, or “computer vision”, was done under the auspices of the U.S. Defense Advanced Research Projects Agency (DARPA), which is the central research and development organization for the Department of Defense (DoD). However, many educational institutions and businesses have also worked on video analytics and advanced its capabilities.

By using a set of computer algorithms that scrutinize the change in the digital image at the pixel level by comparing one frame or image of video with the previous frame, computers can identify movement, recognize objects or people as a grouping of related pixels, and determine the size of an object. As a result, computers can alert operators or generate alarms based on specific events, such as people entering the field of view, the direction of an object, or the removal of an item from the field of view.

Common applications for video analytics include:

- Motion sensing
- Object left behind
- Object trip wire
- Object moving in the wrong direction
- Line counting/crowding—i.e., how many people are in the field of view

Video analytics can also tag video, making a visual annotation such as highlighting an object of interest with a circle or marking the object with a certain color. Video analytics can also make non-visual annotations or index marks on video to facilitate faster identification of an event when reviewing recorded video.

Video analytics can be used in a wide range of applications, not just to alert operators or investigators to an event, but also to highlight common patterns such as traffic flow of people or cars. Video analytics also create new opportunities for the use of video surveillance by non-traditional users, who are not generally focused on safety or security. These users may be in other parts of an organization, such as in marketing in the retail industry or manufacturing and production control.

Administrators can use these analytics to trigger automated alarms and other responses. For example, by using video analytics, an airport or train station video surveillance system can sense the presence of a bag left in its field of view and page security staff immediately. Casinos can monitor crowding around gaming tables and, when needed, alert the floor staff to quickly open another table. Retail businesses can monitor checkout areas to detect delays as well as understand customer movement for optimized product placement in a far less invasive manner. Analyzing video can even help freight companies validate the movement of cargo and help to locate lost packages. Thus, video analytics can create new uses for video and transform it into a business tool that contributes to the productivity or sales growth of a business or organization.

Video analytics capabilities can be embedded into DSPs or run on microprocessors that can be deployed in devices such as cameras and DVRs/NVRs (third- or fourth-generation video surveillance systems), and on dedicated computer servers that have access to video. The flexibility of where video analytics can be deployed and who can use this new tool is enhanced when they are used on an IP network. The network provides the video to be analyzed and generate reports that can be distributed anywhere the network goes.

Video analytics continue to evolve. While they are helpful and deliver reasonably good results, video analytics may not always be accurate in identifying a given event; false alarms do occur. In some applications, the algorithms to identify certain items or events are fairly nascent and do not deliver acceptable accuracy. As an example, in facial recognition, video analytics algorithms will need to improve to be used in situations where the number of new or different faces (and the resulting database) is large. The state of the art continues to progress to address these issues.

Bandwidth Challenges—Sorting Fact and Fiction

There are numerous misconceptions in the video surveillance market regarding bandwidth requirements of video surveillance running over a digital IP network infrastructure. Several aspects should be considered:

- Compression and DSP technology improvements
- Typical LAN link bandwidth is 100 Mbps to 1000 Gbps
- Number of video streams coming over a single link
- Where are these video streams being aggregated—is this aggregation point a choke point based on device performance?
- Is the video running over a dedicated or shared link?
- What is the available WAN bandwidth?
- What kind of streaming is being used: multicast or unicast?

Video Compression and DSP Evolution

The rapid advancement of compression algorithms has been helpful in mitigating video bandwidth demand on the network. While just a few years ago MPEG-2 created a 4 to 5 Mbps video stream from a 30 frames per second D1 720x480 resolution (NTSC) video camera, the cost- effective and more recently introduced MPEG-4 algorithms cut the video stream down to approximately 3.5 Mbps. This improvement in compression ratios comes with no discernable change from MPEG-2 compression. Other compression technologies promise to further reduce bandwidth requirements with the same or greater video fidelity of today’s technologies. Moreover, thanks to cheaper and more powerful semiconductors, the necessary DSPs are very affordable. H.264 compression offers to further reduce video bandwidth requirements.

Network Bandwidth

Contrary to misconception, there is typically plenty of bandwidth in a LAN connection. For the last 10 years, organizations have deployed switched (point-to-point) 100-Mbps links to user desktops and other edge-connected devices. The effective throughput of each link is approximately 80 Mbps. Given these compression ratios, a single LAN-edge link can support 40, 2-Mbps video surveillance streams. Many enterprises and some other organizations have been purchasing switches that support 1000 Mbps or 1 Gigabit to the desktop for several years, given “Gigabit per port” prices have approached 100-Mbps price per port. In these cases, with overhead, the throughput of a single 1000-Mbps link would support 400, 2-Mbps video streams.

In practical terms, most video surveillance deployments only carry hundreds of video streams when aggregated across many network switches as the video traffic traversed the core of the network where multiple Gigabit or 10-Gigabit links are now deployed. Thus, it is unlikely that video surveillance would put a strain on a typical organization’s LAN. Figure 5 shows the number of video streams supported using various compression algorithms as opposed to common LAN and WAN link speeds.

Figure 5. Popular IP Network Switch Bandwidth

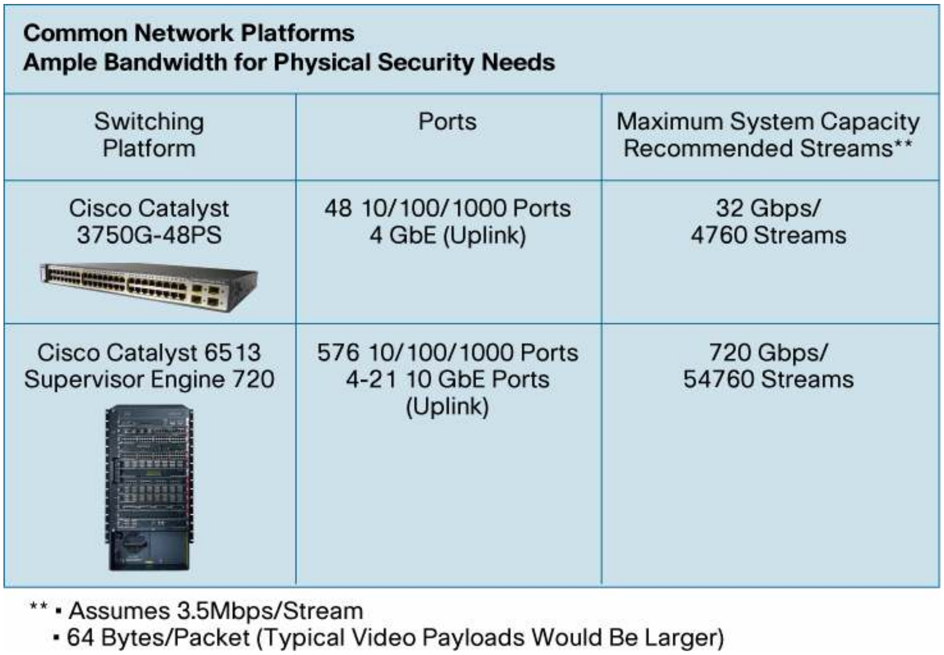

Common Network Platforms
Ample Bandwidth for Physical Security Needs

Switching Platform	Ports	Maximum System Capacity Recommended Streams**
Cisco Catalyst 3750G-48PS	48 10/100/1000 Ports 4 GbE (Uplink)	32 Gbps/ 4760 Streams
Cisco Catalyst 6513 Supervisor Engine 720	576 10/100/1000 Ports 4-21 10 GbE Ports (Uplink)	720 Gbps/ 54760 Streams

** ▪ Assumes 3.5Mbps/Stream
▪ 64 Bytes/Packet (Typical Video Payloads Would Be Larger)

When looking at WAN connectivity from remote locations, video may pose some challenges. However, as observed in third- and fourth- generation deployments, several approaches are available to minimize bandwidth consumption. These include:

- Onsite recording (does not affect WAN speeds as video is recorded locally, over the LAN)
- On-demand-only monitoring (video stream requests)
- Video transmission increased or transmitted only upon prescribed events (motion, perimeter alarm, video analytics rule violation, etc.)

Multicast Saves Bandwidth

A single video camera produces a single “unicast” video stream. However, when the stream must be monitored and recorded, two streams, which may or may not be at the same frame rate, must be provided; for example, two unicast streams are provided. If more than one person wishes to view the same video stream simultaneously, this would be yet a third unicast stream. Sending multiple copies of video across a network is inefficient. And today, many IP cameras or encoders have a limited amount of power and cannot support more than one or two high- quality video streams.

This problem was solved many years ago in networking when the same information was to be sent or offered to multiple users or devices using “multicast” technology. Multicast provides a single copy that can be offered by way of subscription to as many users as desired. The benefit for video surveillance is that a camera or encoder must only produce a single multicast stream, regardless of the number of devices for all users who wish to view the video. The network infrastructure, preferably at the point closest to each subscriber, handles the replication of the video to multiple devices, thereby minimizing bandwidth over shared links. Transcoding of frame rates can be provided at the IP camera, encoder, or once again, at the point closest to the subscriber, maximizing the value of a single multicast stream.

In summary, several practical solutions can address any legitimate video surveillance bandwidth challenges, which would only be a concern across remote WAN links.

Business Case for Deploying Network-Centric Video Surveillance

While no two video surveillance deployments are identical, Cisco’s safety and security group has migrated to a large-scale, network-centric third- and fourth-generation video surveillance deployment. It has a mix of analog and IP cameras; most are analog. The cameras are located in more than 300 facilities around the world. The safety and security group has reported impressive results; much of their efforts preceded general availability of third- or fourth-generation video surveillance components.

Results

Migrating to IP network-based video surveillance has yielded the following benefits for Cisco:

- Reduced storage requirements by 60 percent, representing US$500,000 in savings. The new system reduces storage volume requirements because it has the intelligence to store video only if motion is detected.
- Reduced number of servers by 40 percent, representing $200,000 in savings. Servers can support nearly double the number of cameras they did previously because they no longer need to devote compute cycles to video encoding.
- Improved video quality. At four frames per second, the ability to recognize faces is vastly improved over the previous system’s two frames per second.
- Gained ability to unify the CCTV system with other security systems, such as alarm detection and access control systems.
- Reduced false alarms in areas covered by video surveillance cameras by an anticipated 90 percent by giving security personnel the ability to view the associated video in real time.
- Mitigated risk by expediting maintenance and repair. When a proprietary box broke, a third-party technician was required. Many of these technicians were unfamiliar with Cisco site requirements and working with in-house video system technicians. As a result, remediation

that used to take five days or longer has been reduced to a 2-hour response and 24-hour turnaround for repair.

- Less time required to investigate security incidents. Security operations center and other authorized safety and security personnel can view stored or real-time CCTV video from any camera around the world, responding more quickly and appropriately to incidents. Investigation is further accelerated because investigators can retrieve more video at one time. Second-generation DVRs limited the amount of video Cisco could pull on demand to one hour's worth. Now, Cisco can easily pull 24 hours' worth of video at a time. To speed investigations further, the NVR software allows Cisco security operations to highlight only the changes to a particular area of the video footage, such as a desktop.
- Reduced maintenance costs by 20 percent. Cisco IT has economies of scale and spends less time monitoring and maintaining servers than when Cisco's physical security team maintained their own servers
- Increased security. The standards-based system is more secure. Network protection and virus definitions are implemented as soon as they become available.
- Return on investment. The system has paid for itself more than once by enabling the safety and security department to apprehend equipment thieves and then recover the stolen property.

Lessons Learned

The chief lessons learned from the transition to digital CCTV pertain to making the best use of Cisco IT resources. In order for physical security to become a mission-critical application that operates on the IP network, both physical security and IT groups must work together for maximum benefit. When Cisco's physical security managed the servers, a hardware or software problem was a serious issue for the department. Now physical security operations personnel generate a case and work order, and IT uses its technical resources and expertise to resolve the issue. Cisco physical security had to shift its culture and allow IT do the work and run through its own processes. For a truly successful partnership, IT has to fully understand and agree with the project goals.

Finally, a successful CCTV over IP deployment involves planning, and agreeing on physical security operational responsibilities and IT responsibilities. For global deployments, it is helpful to work with someone with the authority to approve the project worldwide, thus avoiding the need to negotiate with various regional entities—but recognize that each regional entity must support the video surveillance deployment initiative.

All parties at Cisco agree that the culture change required to partner with IT yielded positive results. Reliability and accuracy increased, and security staff could focus on their main responsibility—the safety of people and the security of the organization's assets. The project was successful from a financial standpoint as well -- the savings of the fourth-generation video surveillance solution completely offset the costs of its deployment.

Moving Forward

Network-centric video surveillance will continue to evolve. Pure IP-based products are already available. These products may be attractive for new construction, "greenfield" deployments. However, many video surveillance deployments are already in place and can be enhanced with innovations enabled by open standards-based products, such as those offered by Cisco. Rather than looking at a massive "forklift upgrade" to the video surveillance system, many organizations will have hybrid deployments for many years to come, which will have a mix of analog and IP-based video surveillance products. Cisco any-to-any-for-any video surveillance and network infrastructure products provide a high degree of vendor interoperability to realize best-in-class deployments. As a result, greater accessibility to video is provided from almost anywhere.

When implementing physical security as a network application, these resources must be properly segmented and secured; accessible only by authorized parties. Cisco's wide range of network platforms (routers, switches, firewalls, intrusion prevention) can provide the appropriate level of security.

Cisco delivers a networked and unified video surveillance solution. Given the role of the IP network in most organizations today, and its presence in video surveillance vendors' product roadmaps, Cisco expertise and experience in network convergence provides an obvious mediation point for legacy and new video surveillance devices and applications. It is important to work with vendors that truly understand what it means to provide world-class, network-centric solutions. Through innovation and internal R&D process, shared across all of its technology groups, Cisco will continue to evolve a standards-based intelligent and resilient network that customers can rely on for their most mission- and business-critical operations. This level of intelligence allows customers to rapidly deploy new devices and introduce new business processes and applications that are supported by policies governing end-user access, device connection, and stream management.

For more information on Cisco Video Surveillance products, visit http://www.cisco.com/go/videosurveillance.

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 (USA) Pte. Ltd.
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 C11-449363-00 12/07