I SPY WITH MY NETWORKED EYE

It's not a matter of *if* enterprise IT managers will be responsible for video surveillance, but *when*. Are you—and your network—ready?

The federal government is implementing video surveillance to help combat everything from traffic offenses to terrorist activity to illegal immigration. The U.S. Department of Homeland Security's FY06 budget, for example, earmarked $51.1 million for the America's Shield Initiative, which enhances electronic-surveillance capabilities along our borders; that's an increase of $19.8 million over FY05. In August 2005, New York's transit authority awarded Lockheed Martin a $212 million contract, which includes installation of 1,000 cameras and related equipment in city subway stations, bridges and tunnels.

Such projects, however, are few and far between. Compared with that in Europe, the level of surveillance in U.S. enterprises and municipalities is pitifully low. The sophisticated systems we see in movies may exist in large gaming houses, financial firms and a few select transportation installations, but most organizations still use VCR tapes to store low-resolution video camera output—if they store video at all. Many simply pipe feeds to banks of monitors for real-time viewing. But is anyone watching?

BY PETE TENEREILLO

The worldwide Internet protocol surveillance market will grow to $6.48 billion in 2012 from last year's $435.8 million, according to a report by Frost and Sullivan. The technology is advanced: IP-enabled high-resolution cameras and encoders, video storage and search capabilities, and object-recognition software have long been available. Digital video surveillance—that is, using computers and networks to store, play back and analyze surveillance video—is the future. IT pros who understand this can prepare now for the mandate that surveillance video share the IT infrastructure.

Meanwhile, video consumers are becoming more sophisticated. Enterprise decision-makers can watch decent-quality video, with amazing ease of use and reasonably advanced features, for free on YouTube and Google Video. They'll accept nothing less from a system their company paid millions of dollars to deploy.

And the lines between storage and networking of digital video are beginning to blur. Just as firewalls, application switching and Web caching moved from software on the server to appliances, video delivery technology is also migrating (see "Move Inward, Young Technology," at *www.nwc.com/1715/1715toc.html*). In enterprise surveillance networking, the future is analytics at the edge, video switching, caching and intelligent archival.

Make Room

In select cases—several airports come to mind—entirely parallel networks are built for digital surveillance, but the vast majority of corporate budgets dictate resource pooling. And this goes far beyond simply sharing network plumbing: The requirements for storing, sharing, managing and, most important, securing surveillance video overlap with the requirements for other types of video, including videoconferencing, video messaging and corporate communications.

"Just as the telecom department that managed the PBXs eventually merged into the IT department, a cultural convergence of the IT, physical security and even audio-video departments is inevitable," says Richard Mavrogeanes, founder and CTO of VBrick.

Large vendors are making a bid for this market, as evidenced by Cisco's April 2006 acquisition of SyPixx Networks, a maker of video-surveillance gear that lets analog surveillance systems operate as part of an IP network. EMC entered the market in August 2005 with the announcement of its Surveillance Analysis and Management System (SAMS), a combination of EMC video-server software; EMC storage platforms, including Clariion and Centera; and EMC storage-management software, including PowerPath and OnCourse.

There are, though, subtle but important differences between the requirements for surveillance video and other networked video.

READ MORE NETWORK INFRASTRUCTURE NEWS, REVIEWS AND PRODUCT ANALYSIS ON OUR INFRASTRUCTURE CHANNEL: *WWW.NWC.COM/CHANNELS/NETWORKINFRASTRUCTURE*

Surveillance video may be stored 24/7 from every camera and archived, but viewed only when a physical intrusion or emergency occurs, for example. In these cases, however, live video must be available in real time, and the consequences of loss of stored video may be dramatic.

In contrast, for corporate communications, a single video clip, say, a CEO address, is likely to be distributed over every segment of the enterprise network. If the storage and distribution system fails, the clip can be easily restored. Furthermore, corporate communications rarely require real-time delivery, so the simple buffering and pause/FF/rewind functions provided by clients such as Windows Media Player may suffice.

New Twist on Spyware

Advances such as IP cameras with Web servers, networked video storage and centralized power management provide an enormous level of architectural flexibility, an exciting new challenge for the IT manager—and a green-field opportunity for attackers. Security must be implemented at the source, the network, the storage system and the client. If someone can penetrate the security on the camera, he can watch you, and everything going on in your facility. If the video is stored on an NVR, he can access a history of people who've entered the facility, and possibly determine what they said or even typed on cipher-lock keypads or computer keyboards. If he can break the security on the POE switch, he can shut down all video. If he can penetrate the security on the surveillance control system, he can not only watch video, but also learn how the physical security department uses it. For these

Executive Summary

VIDEO SURVEILLANCE

Eventually, an entire new industry will emerge around enterprise video—not just surveillance video, but all types of video. Today, there are products well-suited to videoconferencing, corporate communications, training, streaming and surveillance. Soon, a new breed of enterprise video-distribution appliance will materialize, one that will accommodate all these purposes well.

Meanwhile, we focus here on aspects of digital video surveillance pertinent to the IT manager and network architect. Sooner or later, IT groups will need to integrate surveillance video into their networks and storage architectures. In "I Spy With My Networked Eye," we examine architectural and security considerations and explain the various elements of video-surveillance setups.

As high-resolution cameras grow in popularity, the data streams and management headaches will begin. Here's how to prepare.

reasons, IT must be empowered to treat surveillance video like any other sensitive data on the network.

"IP networked cameras became available about 10 years ago. The main selling point was the ability to monitor remotely," says Fredrik Nilsson, general manager of camera maker Axis Communications North America. "About three years ago, most DVR manufacturers added a network port. The surveillance professional said, 'Great, now the DVR has a network port.' The IT manager said, 'What is this new server on my network? Either integrate it into the IT security policy or get it off.'"

Most networked DVRs ship with customized versions of Windows and require preinstalled applications. The platform cannot run other software, such as a firewall or antivirus package, much less patches. Of course, connecting any system to a network without a firewall, virus protection and patch management is begging for trouble.

Cisco experienced this firsthand: A few years ago the company installed NDVRs from Verint in its data centers for remote monitoring and video storage. The NDVRs were not considered servers. In 2002, the NDVRs were infected with the Nimda virus. Cisco's IT department mandated that the NDVRs be subject to the same patch upgrades and Windows images, including antivirus software, as all the company's other workstations and servers. The NDVRs required custom software, though, and couldn't be refitted, so they were unplugged. Cisco and Verint had no comment on the specific event. Not surprisingly, Steve Collen, director of marketing for Cisco, told us the company is focusing its product development on ensuring that video traffic and the video infrastructure are protected. Verint currently suggests the use of its surveillance software on servers that are protected by standard firewall and antivirus software.

NVRs are sold as software for Windows or Linux servers (EMC) or as appliances (VBrick and SteelBox Networks). If the NVR is installed as software, make sure the server can be managed and secured by the same software and processes you use with other servers on your network. If it's installed as an appliance, it must be hardened and manageable.

The market for software-based NVRs is much stronger than for appliance-based devices, according to Axis' Nilsson. This is an interesting contradiction to, say, the enterprise firewall market, where software-only solutions have all but failed. We expect the NVR market to gravitate toward the appliance model as the space matures and becomes more tightly integrated with IT.

Get Smart All Over

A successful video-surveillance implementation will deploy intelligence in the camera, network and storage system. The future is brightest for systems that implement analytics both near the edge, to make system and network usage more efficient and scalable, and near the storage system, for more computationally intensive analysis, such as that requiring comparison of video from multiple sources (facial comparisons for tracking a suspect from building to building, say) and evidentiary purposes. Here's a breakdown of what goes where:

» Intelligence in the camera: Although we firmly believe network-centric functions, such as security, access control and stream management, will gravitate away from the camera or encoder, much innovation will continue at that endpoint. Basic examples are cameras that send video only when motion is detected or sound an alarm

Pieces of the Surveillance Puzzle

Before going shopping at your local spy shop, learn the lingo	
Digital video recorder (DVR)	A DVR is a direct replacement for a VCR. Coaxial cables from analog cameras are connected to frame-grabber cards in a PC running the DVR software. The DVR encodes analog video to a digital format, such as MPEG-4 or Motion JPEG, and stores it. Stored video can be viewed using a GUI running on the DVR computer. DVRs are typically Intel-based PCs running Windows or Linux. Typical capacity is 16 to 32 analog connections. Storage can be via internal disk drives, direct-attached RAID, NAS or even SAN. There are at least 200 brands of DVRs available.
Networked digital video recorder (NDVR)	An NDVR is simply a DVR that is network-enabled, allowing the stored video to be shared with other NDVRs, or viewed over the network using popular software such as Windows Media Player. NDVR software from one vendor is typically not interoperable with that from other vendors.
Network video recorder (NVR)	An NVR is an appliance, or software running on a server, that stores video that has been previously encoded. The NVR topology is fundamentally different, and vastly more scalable, than DVR/NDVR topology. For new installations, video is often encoded in the camera itself or by an encoder (see below). Video is then sent to the NVR over the network. The NVR stores the video on a DAS, NAS or SAN system. NVRs have a substantial scalability advantage over DVR/NDVRs because the processing overhead of encoding is handled by the camera or external encoder device.
Encoder/decoder (codec)	A codec is a device that encodes video inputs from analog cameras to a digital format, such as MPEG-4 or Motion JPEG, and sends it over the network. Products range from single-input devices about the size of a Linksys home gateway to industrial-strength, modular encoders with 40 or more analog inputs—think of a rackmount device that looks similar to a Cisco CAT6K switch. Encoders are typically located close to cameras. Codecs can be used in conjunction with NVRs, but many legacy installations use codecs simply to extend the range of pure analog systems. In these legacy installations, analog cameras are connected to a codec that performs the encoding function at one end; the video is sent over the network to a codec performing the decoding function; and the resulting analog video is viewed on analog monitors and optionally recorded using a VCR. This is a very common topology in large facilities such as airports.
IP networked camera	An IP networked camera has a built-in codec and wired or wireless network capability. Modern IP networked cameras are quite advanced, with multimegapixel resolution, built-in motion and object detection, built-in Web servers (so that any PC with Windows Media Player can connect to and view video directly from the camera), and Power Over Ethernet capabilities.

if they are blocked. A more advanced case is the "tripwire." In a retail store, for example, there may be an off-limits product podium. A camera can be configured with a tripwire, essentially a user-defined section of the field of view. If motion is detected within that space, an alarm message is sent over the network to the video-management system, which may display video to a special monitor. Some cameras even have audio-out capabilities—"Sir, get off the podium, please!"

On the topic of performance and scalability, network cameras by definition off-load video encoding—as much as 80 percent of the heavy lifting is done at the edge. The NVR is left only with the work of storing and playing back video. In contrast, NDVRs with frame-grabber cards must use system resources to encode video, at the expense of other NVR functions.

"Eventually most of the intelligence will be pushed to the camera level, since this is the only way to scale an intelligent video system," Nilsson says. Maybe. Cameras that can support application-specific intelligence, such as people-counting or license-plate recognition, are coming, but first, improvements are needed on processor capacity, algorithms and data collection, including frame rate and resolution. Axis cameras support multiple simultaneous streams with different resolutions and frame rates, useful in a fully meshed architecture where NVRs connect directly to servers. For large-scale apps, a switched architecture, with stream-splitting done closer to the user, is the better choice.

» Intelligence in the network: Imagine a scenario where a suspected shoplifter is spotted in a large retail warehouse. Security professionals in the video-control room at that location scan the aisles for individuals that match the description, in real time, from camera to camera, but they also must rewind stored video to view the crime as it happened, and to confirm the suspect's appearance. Security guards on foot want to view both live and stored video on wireless handhelds. Staff at the retailer's central location may also wish to monitor events, as may law enforcement. Thus, the same video must be intelligently and securely switched throughout the LAN, MAN and WAN. Such aggregate demand will create a large spike in traffic, and if the network fails at this critical moment, the IT manager will be in the hot seat.

Now consider the benefits, and also the network architecture implications, of an off-site, outsourced security service. Clearly, there's a case for technology to make all this as efficient—and therefore economical—as possible, while maintaining quality of service and a sufficient resolution.

The limitation is not technology, but rather budget. Surveillance departments could install high-definition cameras with 1,920x1,080 pixel resolution, at 30 fps (frame per second) streams, then store all that data, but in most cases this will be overkill. We'd settle for NTSC quality (648x486). In reality, most IP video deployments use lower resolutions (352x240) and frame rates (5 fps) because to do otherwise is still cost-prohibitive

Intelligent Switching and Storage of Video

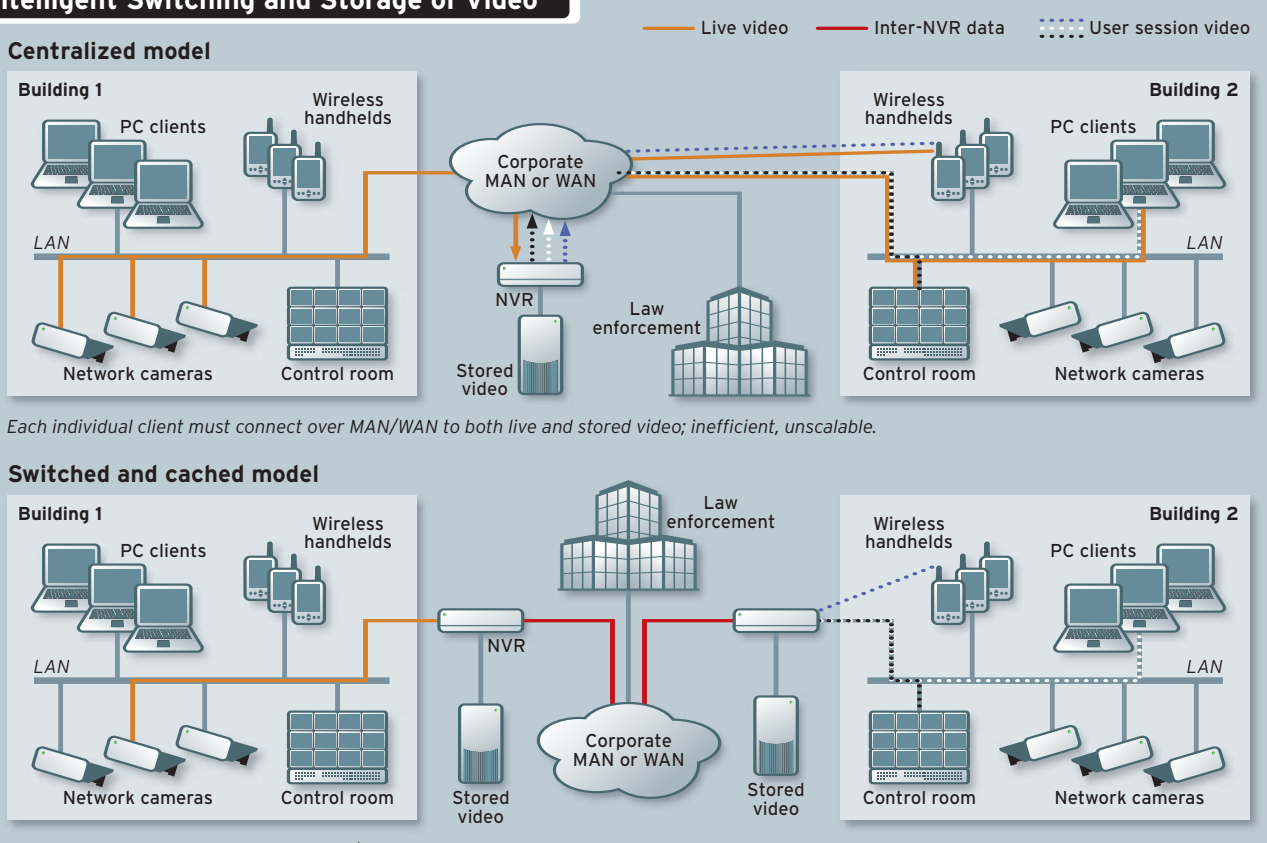

in terms of hardware, bandwidth and storage.

Luckily, there's much innovation in this space. We found three examples of hardened network appliances with features that move intelligence into the network. Cisco's (formerly SyPixx) products make extensive use of IP multicast, originating the multicast at the encoder, and letting the network switch and route the video.

"The heart of Cisco's development efforts in the surveillance area are in product integration; that is, making the products work well in a Cisco network," says Cisco's Collen. "The goal is to set QoS as far out to the edge as possible, because that's where you get the most benefit." IP multicast allows multiple viewers of live video, with streams being simultaneously directed to NVR storage. Collen says Cisco has customers using multicast to deliver live video to as many as 20 locations, 24/7.

VBrick's products also make extensive use of IP multicast. "Multicast is clearly the superior way to send live video to more than a handful of viewers," VBrick's Mavrogeanes says. "But multicast is not deployed everywhere, and sometimes there are pockets of older network technology that may not have multicast enabled in a campus network." VBrick WM appliances support "multicast rollover," which attempts to deliver live video to Windows Media Player using multicast. If multicast fails, the player reverts to unicast; the unicast may come directly from the VBrick appliance or from a reflector server.

SteelBox takes a different approach, a sort of surveillance-specific content delivery network appliance. "IP multicast is poorly suited to digital surveillance video, because rarely do multiple consumers require exactly the same experience," says Richard Howes, CEO of SteelBox. In the case of a possible perpetrator in a building on an enterprise campus, for example, a security guard in the local building, another in the central HQ across the MAN/WAN, and yet another on foot with handheld wireless, would almost certainly each want independent ability to pause, fast forward and rewind both live video and video from the time of the actual breach. Each would also require different data rates. Howes says multicast does not suit this purpose, as it delivers the same stream to each client that has joined the multicast group. He also cites performance and security implications of multicast.

SteelBox's Digital Matrix Storage Switch (DMSS), an appliance similar in architecture to the Cisco PIX firewall, is a combined Layer 7 video switch, video cache and NVR. In fact, the founding team of SteelBox is the same team that developed the Cisco PIX Firewall and LocalDirector. Video from the source, be it live or stored, is delivered from DMSS to DMSS, intelligently switching and caching along the way. This caching scheme lets each client have full rewind/FF/pause capabilities. Clients connect to the closest DMSS. If that DMSS does not have requested video cached locally, it gets the video from the DMSS that does have it, delivers it to the client and caches the video for other users. The DMSS also allows real-time frame-rate reduction, so that a single stream can be simultaneously split and groomed for a diverse set of devices. Although the DMSS can support multicast, switching achieves the bandwidth efficiencies of IP multicast, without security issues and without requiring multicast to be enabled on routers, according to Howes.

The jury is still out on multicast for surveillance applications, and the topic is hotly contested among network appliance vendors. Axis told us that fewer than 10 percent of its customers enable multicast on Axis cameras. We think it's a matter appropriateness and implementation. Multicast helps only when multiple consumers want to watch the exact same stream at exactly the same time at exactly the same resolution/frame rate. This isn't often the case with surveillance, but quite useful with corporate communications or

US AND THEM

We asked Kevin Marier, editor in chief of *IP Video Security* and someone with an extensive background in video surveillance, what architecture he would propose for simultaneous viewing and storage of IP surveillance video over an IT network in an installation with, say, 100 cameras. Marier says modern IP cameras have Web servers that can accept as many as 20 simultaneous connections, so both clients, such as PCs running Windows Media Player or some other software, and NVR software running on servers, can connect to cameras directly.

"The LAN is fine; the real issue is power management," says Marier, referring to Power Over Ethernet to the cameras.

Although we agree that power management is important, Marier's response shows the chasm between the IP video professional and the network architect. Clients and NVR servers connecting directly to cameras over Web servers in a fully meshed fashion, all viewing the same data over redundant, independent streams that originate at the source, is a great example of what network architects call "a full mess"—inefficient, unmanageable and unscalable.

Now, IT clearly needs domain experts, like Marier, who understands things IT network professionals don't. Do you know the difference between CIF and 2CIF? The minimum frame rate required to identify the face of a person walking down a hallway? We didn't think so. Our job as networking professionals is to deliver video as efficiently, securely and cost-effectively as possible. Oh, and at high enough quality to be useful in investigations. So if you're embarking on a video project and need a resource for specialized IP video-surveillance information, check out IP Video Security magazine at *www.ipvs.com*.

training. With an all-Cisco network, multicast is likely to perform flawlessly. With lower-end switches, we recommend testing before investing.

» Intelligence in storage: Should you archive video, and for how long? There are several schools of thought. What's certain is that, as multimegapixel cameras become more popular, huge amounts of data must be stored and managed.

Most intelligence and analysis is applied to archived video, says Dick O'Leary, senior director of the global solutions group at EMC and responsible for EMC's video surveillance offerings. At some point in the future, computers will be able to scan human faces, possibly even bone structures, in real time. But we're not there yet. Although simple motion and object recognition can occur in real time at the edge, O'Leary says the bulk of intelligence for analysis should reside on or near the storage system because most analysis of surveillance video happens after there's been an incident.

Centralized storage is easiest to manage, but is not the most efficient setup. O'Leary says IT managers should ask themselves, "Do I have the bandwidth to store all surveillance video in one spot?" In the diagram on page 60, we illustrate why it's smart to store surveillance video close to the source, and use intelligence in the network to distribute it as needed. The dashed lines show users who each require their own TiVo-type session, with various frame rates and resolutions.

You also need to develop an intelligent archival plan that meets applicable policies and regulations. Video storage requirements may change over time and vary by subject. "Some customers do not want any video stored for more than 30 days," O'Leary says. "Other customers want high-resolution, high-frame-rate storage for 45 days, followed by long-term archival at a lower resolution and/or frame rate."

Some storage requirements are tied to financial transactions; for example, cash-register video is often stored for 60 days. If a customer challenges a credit-card statement, a retailer may wish to view the video of a given purchase, so security policy may require point-of-sale transactions be stored for a set number of days for fraud detection, then deleted. That same retailer may wish to store high-resolution video of the store aisles for only 24 hours to gather evidence related to shoplifting, then have that same video scaled back to a lower frame rate and/or lower resolution for the following 60 days for detecting fraudulent insurance claims. NWC

PETE TENEREILLO was one of the developers of the first commercial server load balancer and the first firewall appliance. He is currently an independent consultant. Write to him at *pete@tenereillo.com*. Post a comment or question on this story at *www.nwc.com/go/ask.html*.