

NETWORKED VIDEO

APP
NOTE

CAPTURE, TRANSPORT, VIEW AND STORE QUALITY VIDEO FROM
NETWORKED CAMERA SYSTEMS...ANYTIME, ANYWHERE

1 MOVING FROM ANALOG VIDEO TO NETWORKED VIDEO	4
1.1 LIKELY APPLICATIONS	4
1.1.1 Airport surveillance	4
1.1.2 Urban monitoring	5
1.1.3 Industrial applications	5
1.1.4 Corporate surveillance	5
1.1.5 Government applications	5
2 ISSUES FOR NETWORKED VIDEO IMPLEMENTATION	5
2.1 SYSTEM-LEVEL ISSUES FOR NETWORKED VIDEO IMPLEMENTATION	5
2.1.1 Quality versus compression: maximizing bandwidth and storage	5
2.1.2 Networked video recorder	6
2.2 COST EFFICIENT STORAGE ARCHITECTURE	7
2.2.1 Current problems of inefficiency	7
2.2.2 Efficient archive solutions: optimizing network video archive storage	7
2.2.3 Short-term video archiving	7
2.2.4 Long-term video archiving	8
2.2.5 Managing long-term archive storage with ease	9
2.3 STORAGE MANAGEMENT BENEFITS	9
2.3.1 Digital video archive benefits at a glance	10
2.3.2 Seamless integration	10
2.4 BEYOND STORAGE: DIGITAL PROCESSING ADDS UNIQUE FUNCTIONALITY	10
2.4.1 The challenge of image quality	11
2.4.2 Expanding video use while reducing storage needs	11
3 DIGITAL VIDEO ARCHIVE ARCHITECTURE	12
3.1 MISSION-CRITICAL FEATURES	12
3.1.1 Storage area network	12
3.1.2 Fault-tolerant architecture	12
3.1.3 Mirrored write cache	12
3.1.4 Hot spare storage drive replacement	13
3.1.5 High availability	13
3.1.6 Early failure detection	13
3.1.7 Storing video for evidence	13
3.1.8 Data confidentiality	14
3.2 DATA PROTECTION	14
3.2.1 Secure administration	14
3.2.2 Disaster recovery	14

3.3 IMPLEMENTING A NETWORKED VIDEO SYSTEM 15

 3.3.1 Small site 15

 3.3.2 Medium site 16

 3.3.3 Enterprise site with a SAN 17

4 CONCLUSIONS 18

5 GLOSSARY 19

ABOUT THE AUTHORS

J. Steven Rice

Security Surveillance Practice Manager, StorageTek

Steve Rice is the Security Surveillance Solution Business Manager and member of the Solutions Group at StorageTek world headquarters in Louisville, Colorado. He has more than 22 years of experience in the high-tech industry in both software and hardware manufacturing. As the Security Surveillance Solution Business Manager at StorageTek, Steve’s current responsibilities include business development, go-to-market strategies, strategic marketing, corporate alliance development and overall solution management for StorageTek’s security surveillance digital video archive solution.

Andrew D. Pickles

Technical Account Manager, StorageTek

Andy Pickles is a Senior Technical Account Manager for StorageTek’s UK Field Sales Organization. After receiving his Geophysics and Planetary Physics degree from the University of Newcastle upon Tyne, UK, Andy pursued a career in the oil and gas industries. His work dealt with a high degree of computational data analysis. He has also held a wide range of roles within the IT industry, from Software Analyst and IT Manager, to Network Security Consultant and his current position as Technical Account Manager at StorageTek. Andy’s responsibilities at StorageTek include systems engineering and support for the UK sales organization, as well as the development of viable storage solutions with StorageTek alliance partners.

1 MOVING FROM ANALOG VIDEO TO NETWORKED VIDEO

Networked video has grown well beyond its origins in primitive Web cameras. In many applications, networked video is now able to provide a viable alternative to conventional analog video products. In fact, not only can networked video technology parallel the performance and convenience of analog closed circuit television (CCTV), it can provide far greater flexibility and cost savings.

Upon examining the differences between analog CCTV and networked video solutions, it becomes obvious that networked video holds significant advantages for the user. First, analog CCTV systems require the monitoring location be in close proximity to the cameras, whereas networked video systems have no distance constraints. Once the video has been digitized, it can travel over cable to anywhere on the globe like any other digital information, without significant degradation. Thus, the monitoring site is not tied geographically to the location of the cameras.

Second, analog infrastructure is complex and, therefore, difficult to alter or expand. Analog CCTV requires individual transmission paths between cameras and monitors for video, audio, alarm detection, camera telemetry (PTZ — pan/tilt/zoom) and other functions. Alternatively, digitized video data from any source can be carried over a single cable. Laying the infrastructure is simply a matter of installing either a one-cable ring around the site or a radial network of single cables serving individual cameras or clusters of cameras. In other words, networked video taps into the existing Internet Protocol (IP) computer network, utilizing existing technology that is robust, universally adopted, easy to install and inexpensive.

The ultimate result of networked video is reduced cost. By eliminating the need for line amplification and simplifying and reducing onsite cabling, installation and management costs are significantly reduced. PC-based systems management, administration and recording eliminate the need for peripherals such as video recorders, multiplexers and sequencers. And, as networked video technology takes hold in the marketplace, market forces will continue to drive down digital video costs.

Beyond cost savings, the almost limitless flexibility of IP-based solutions means networked video can deliver benefits that would have been beyond imagination with analog CCTV.

1.1 LIKELY APPLICATION AREAS

There are surveillance applications for which the limitations of analog CCTV are prohibitive. If the advanced application requirements for the CCTV system include a broad monitoring audience, centralization of operations, remote monitoring or flexibility of infrastructure and/or size, networked video may be the only realistic option. Some areas that may require networked video include:

1.1.1 Airport surveillance

There are many potential users of a surveillance system within an airport. These can be either onsite users such as security guards, baggage handlers, managers and retailers or offsite users such as security and emergency services. To provide good quality video to all of these users, a large network covering a large geographical area is required. This is an ideal application for networked video solutions.

1.1.2 Urban monitoring

City center monitoring schemes are an increasingly attractive way of providing public security and peace of mind at a realistic cost. Such schemes often entail monitoring locations spread across a broad area. Traditional analog video transmission over long distances is expensive and impractical. In contrast, networked video is unlimited with respect to distance and, therefore, ideal for this application.

1.1.3 Industrial applications

In the industrial and manufacturing arenas, adherence to health and safety legislation is the responsibility of the employer. In many companies, local area networks (LANs) are already in place, providing an existing infrastructure base for implementing networked video monitoring.

1.1.4 Corporate surveillance

Most large corporations have already invested heavily in their information technology (IT) infrastructure and have extensive LANs and wide area networks (WANs) with high bandwidth capabilities. While these resources are heavily utilized during working hours, they are often underutilized during off-peak hours, a time during which security surveillance is a major concern. Corporate surveillance using networked video is ideal for many industries including finance, healthcare, pharmaceuticals, petrochemicals, energy and central stations.

1.1.5 Government applications

Many government organizations are revising their security surveillance in terms of functionality and archive requirements, necessitating a reevaluation of their current systems and how IP solutions will address these new requirements. Networked video systems can also enable government agencies to link CCTV video information to other types of data in the archives, extending the functionality of the security information.

2 ISSUES FOR NETWORKED VIDEO IMPLEMENTATION

2.1 System-level issues for networked video implementation

As networked video transforms surveillance video from its traditional analog form to the digital realm, the issue becomes how to optimize the quality, transport and recording of video data and store it efficiently and cost-effectively. Once video data is integrated into an optimal transport, recording and storage scheme, it can be processed and analyzed to provide far greater value over simple image gathering.

2.1.1 Quality versus compression: maximizing bandwidth and storage

With digital images, the higher the image quality, the more data is required to describe the image. Therefore, there is a trade-off between image quality and the capacity of the network to carry the digital information (bandwidth). And the amount of bandwidth used drives the amount of archive storage needed to retain the information.

A number of different video compression techniques (codecs) can be used to compress analog video into digital data and deliver the maximum video quality for the minimum bandwidth use. These fall into two main categories: still image encoding schemes (like JPEG) or motion encoding schemes like H.26x and MPEG-4.

JPEG and M-JPEG

JPEG (Joint Photographic Expert Group) is probably the most popular of still image compression algorithms and is widely used in the still digital camera market. Motion JPEG (or M-JPEG) is a successive application of JPEG images into a “moving” picture, i.e., a series of picture files attached end to end, like a cartoon strip, to give the impression of a moving image. While each image is compressed, it is not filtered to remove, for example, static background images that are repeated in each frame.

Owing to this “inefficiency,” JPEG and M-JPEG require more IP bandwidth and, therefore, are not optimal for video data compression. Because M-JPEG is not a standard in itself but rather an adaptation of another format, the interoperability between products may be problematic. It does, however, have low latency in comparison to the other compression techniques; that is, it allows near real-time compression of images even from controllable cameras.

H.261 and H.263

The H.26x standard was originally invented for use in video conferencing and video telephony applications, but it has recently attracted significant interest for IP-based security systems. H.26x is a motion codec. That is, in addition to encoding each individual image, it also encodes the video stream. The codec actually analyzes each image in a sequence and decides to send only those image parts that have changed since the last frame was transmitted.

This vastly reduces the amount of video data to be transported over an IP network in security situations that tend to monitor largely static scenes, making H.26x very attractive for most security applications.

MPEG-4

MPEG (Moving Picture Experts Group) is the name of a family of standards used for coding audio-visual information (e.g., movies, video, music) in a digital compressed format. MPEG-4 provides the standardized technological elements enabling the integration of the production, distribution and content access paradigms of the three fields: digital television, interactive graphics applications and interactive multimedia including content on the World Wide Web.

2.1.2 Networked video recorder

The networked video recorder (NVR) is a radical concept to the CCTV industry but not to the IT discipline. NVR application software resides on a server within the network and, once configured, can record from any IP camera across the network directly onto local or remote storage media. The NVR receives compressed video information from the IP network and allows the user to store hundreds of streams of video on a single server. Although the NVR server is a network device, it is well-suited to CCTV because of its ability to record video data on a predetermined schedule or on alarm as well as its ability to search and retrieve video from the digital video archive.

The networked video paradigm also yields a significant benefit from the NVR, because the NVR server can be physically located anywhere as long as it appears on the network. If required, the complete NVR can also be mirrored (duplicated in real time) anywhere on the network. By using standard PC server technology and a standard storage system, the NVR can become part of an overall IP storage infrastructure thereby leveraging existing IT investments.

2.2 COST-EFFICIENT STORAGE ARCHITECTURE

2.2.1 Current problems of inefficiency

IP video technologies are making rapid advancements that support network video recording. The dramatic increase in surveillance systems has led to a proliferation of cameras that, in turn, has vastly increased the storage requirements to support these systems. However, companies that want to rely more heavily on CCTV surveillance systems to secure and manage their operations are often discouraged by the storage costs.

Today, security and surveillance application vendors frequently install tens or even hundreds of digital video recorder (DVR) systems for a single customer to support a large number of cameras. Each of these DVRs has its own distinct storage requirement and is often unable to share storage with other DVRs, maximize available capacity or expand beyond a fixed number of devices. These limitations increase storage costs and management complexity.

In addition, traditional digital storage systems are built for transaction processing applications, which are applications centered on short bursts of randomly accessed information (high input/output per second or IOPS). Though data transfer (throughput) is fast, these short bursts actually impede effective streaming of large, linear video or audio files.

So traditional digital storage systems are not the answer for digital video archives. What is required instead is a storage architecture designed for bandwidth and continuous throughput rather than transaction-based IOPS. Digital video archives also demand storage devices that transfer large files concurrently to applications that measure streams in megabits per second and/or number of concurrent streams. Digital video archive storage choices impact longevity, access speed, cost, capacity and more.

2.2.2 Efficient archive solutions: optimizing network video archive storage

From caching to disk and/or automated tape, video archives can use many forms of physical storage devices to balance response time, performance, capacity and cost. Knowing archive requirements and having the ability to monitor video access can dramatically improve business value by creating a video archive suited for a short-term rolling archive or a long-term archive capable of handling a great deal more video. In either case, the ability to reduce storage costs and management requirements can be achieved.

2.2.3 Short-term video archiving

Enterprises with a short-term network video storage requirement (generally 60 to 90 days or less) prefer to store all of that video data in an online format (disk). In situations where the video ages for a certain period, it is then groomed from the archive. For many years, the decision regarding how much data to store online has been based on the economics of storage, resulting in decisions that compromise the amount of data stored online.

If storage and storage management were “free,” additional data would be stored online. But storage isn’t free and no matter how much video is stored in an online archive, there is the issue of ensuring continuous operation and quickly getting back into operation when there is a system failure.

StorageTek has been developing enterprise-class, high-availability storage systems for over 30 years. It now has a disk system specifically designed for short-term video archives for networked video systems at a break-through price comparable to analog VCR tape.

The StorageTek disk system offers reliability, scalability and enterprise features. Its performance enables replication, mirroring, online archiving and other features that can be implemented within budget and without compromise. In addition, when requested by the application, segments of video can be marked as “permanent” and held in the archive indefinitely on the disk. Copies can be created on CD ROM, optical or even VCR tape. This process delivers improved access, data protection and distribution in an online environment. (See the implementation section for small and medium sites on page 15 and 16 for further details).

2.2.4 Long-term video archiving

As archives move beyond 60 to 90 days, storage economics become a factor and use of lower-cost media and additional storage management software must be introduced into the archive to be cost effective. There will still be an online disk pool to rapidly capture, archive and retrieve current video, but as the video ages and access to video diminishes, storage management software actively migrates it to inexpensive media (automated tape), retaining the indexable, searchable data (metadata) for easy search and retrieval. This provides the capability to keep recently captured or recently retrieved video on fast, more expensive disk and low-use video on very inexpensive automated tape. With metadata maintained, retrieval of aged video is transparent to the networked video system, so older video can easily be retrieved for further review.

Disk storage has a wide range of architectures, each with its own connectivity, sharability, scalability, price/performance, management and reliability issues. But for networked video systems, scalability, price/performance, management and reliability generally top the priority list. These applications require superior storage bandwidth to serve many large streams and provide non-obtrusive scalability.

Using automated tape subsystems, the archive can store low-use video on inexpensive tape with the automated tape library functioning as the catalog and location manager to store and retrieve footage. Industrial-strength automated tape libraries can contain hundreds or thousands of cartridges (months to years of video) and retrieve them very quickly. Many libraries can also incorporate multiple tape media types and tape drives, so you can store video requiring relatively fast access on tape at a slightly higher cost and very low-use footage with slightly slower access at a very low cost. As with short-term archive, segments of video can be marked as “permanent,” and thus, held indefinitely in the archive. Copies can be created on CD ROM, optical or VCR tape.

As a rule, tape drives are designed around duty cycles (amount/frequency of use), access time and capacity, and have traditionally been used for backup applications which, while critical, do not require many duty cycles from the tape drive or media. But using digital tape for a digital video archive requires tape drives and media to be designed and engineered for regular use duty cycles. Here, the importance of duty cycles is obvious. The tape drive and media must be capable of performing close to continuous operation for several years.

As for access time, most tape drives designed specifically for backup take many seconds or even minutes to load and access data for transfer. However, StorageTek tape drives and media, designed to balance quick access and storage capacity, can often access data in under 60 seconds, stream it off at a high rate and store up to 25 days of video (depending on frame rate and codec) on a 200 gigabyte digital tape cartridge very inexpensively for many years.

2.2.5 Managing long-term archive storage with ease

Storage management software presents a single image of storage to the networked video system. The system simply accesses one highly scalable storage repository of mixed storage media that can include disk (online), tape (nearline), video residing on shelves (offline), or video in remote storage (far line).

When video is stored, storage management software seamlessly transfers a copy of the video, selecting the right storage media according to a specified level of service. The service level indicates the frequency of use or anticipated access time and helps optimize storage costs for less-used video. Service quality can vary over time, based on usage or age, and the storage manager can enable automatic video migration from online to nearline or nearline to offline. The storage manager also maintains a catalog of all video and, as network video applications process requests for video, the video is located on disk and/or tape. If the video is stored on tape, a copy is made available on disk for the network video application to access.

In addition, the storage manager, when requested by the network video application, can mark segments of video as “permanent,” to hold them in the archive indefinitely on disk, digital tape, CD ROM, optical or VCR tape. (See section on implementation of enterprise sites on page 17 for further details.)

2.3 STORAGE MANAGEMENT BENEFITS

In essence, the digital video archive is a central archive management system controlled by a networked video system or an operator. The archive is designed to control storage devices and migrate video files across those storage devices depending on the predetermined service levels required. A storage area network (SAN) can be used to share and dynamically allocate storage resources. To do that, a policy module interfaces with the networked video system and copies the file to the appropriate storage media. The choice of media is based on the required quality of service.

Multiple copies of archived objects can be created and stored on different media and/or in different locations. Not only does this allow for media of different performance levels, capacity and cost, but a user can store a copy of a video in a remote archive, providing a backup in case of disaster.

2.3.1 Digital video archive benefits at a glance

- > Integrate a complete turn-key solution easily into an existing storage infrastructure.
- > Save space and cost by effectively setting up a storage pool and using the complete hierarchy of storage for a digital video archive.
- > Centralize management of video and protect against loss.
- > Automate real-time copies of video archives.
- > Manage storage resources intelligently with dynamic allocation per operation.
- > Achieve a greater degree of redundancy with automatic allocation of another storage resource in case of failure.
- > Scale seamlessly with minimal impact on operations.
- > Integrate easily with most networked video recorders and applications for standard NT or UNIX servers.

2.3.2 Seamless integration

Despite all the tasks handled by the digital archive manager, its major value comes from its ability to seamlessly integrate with the networked video system and the operational environment. In security surveillance, where operations are managed from the networked video system, it is crucial for the archive manager to provide an interface that allows the applications to control and retrieve video from the archive. This allows them to move video from source servers to the archive or retrieve repository content and transfer it to destination servers.

2.4 BEYOND STORAGE: DIGITAL PROCESSING ADDS UNIQUE FUNCTIONALITY

For networked video systems to deliver a truly distinct benefit over conventional analog CCTV systems, they must be able to deliver additional value that could not possibly be realized using analog technology. The two key foreseeable areas of value are image processing and video quality. Once video data is stored in a centralized data store, valuable additional information can easily be obtained by background/offline processing of that video information.

Possible application areas include:

Facial recognition/image analysis

A single video image could be rapidly compared with a centralized data store.

People counting

Many marketing organizations and retailers could benefit from knowing which areas of the store attract customers and which go largely ignored.

Vehicle tracking

Many urban areas have traffic monitoring and congestion analysis systems, but if these individual systems could communicate with each other, a single threat could be tracked over a large distance to obtain valuable security information.

Security screening

Once video information is managed by a networked video system, it can be forwarded to distant locations for further review. One possible application of this feature would be sending video images of airline passengers forward to their destination airport while the plane is still in flight. This would enable customs/security review and screening of passengers prior to arrival.

2.4.1 The challenge of image quality

Image quality is the next big challenge for digital CCTV systems. Today, a large amount of analog video is viewed and a significant amount is recorded for future review. Unfortunately, the quality of this video is often insufficient to allow for recognition of individuals or conviction of criminals based on this information. If the quality of the video being stored could be improved, video information could definitely be more useful in forensic situations and potentially provide further reductions in video storage costs.

New video standards now exist to create DVD-quality video with much-reduced storage requirements. This means that large quantities of very high-quality video information can be stored at very reasonable costs. The most important new standard in this area is MPEG-4, a truly open standard that is now attracting interest from all across the video industry and providing the CCTV industry with a range of products suitable for their requirements.

2.4.2 Expanding video use while reducing storage needs

Image processing techniques have been in use for some time within the CCTV industry. Video motion detectors have also been available for some time. But we have only recently started to see this functionality integrated into the networked camera. By taking this next step, we can start to reduce the amount of video that has to be stored and analyzed so that only video in which motion occurs will be transmitted by the camera.

This level of “data abstraction” can be used in other ways. As previously described, it is possible to integrate other image processing functions with video information to provide useful data. If sufficient computing power were in the network camera itself, this data abstraction process could be carried out at the camera, thereby reducing demand on the network as well as storage costs.

To deliver these benefits, we must move to an infrastructure platform that can ensure both very high-quality video and the flexibility to perform higher-level data abstraction.

3 DIGITAL VIDEO ARCHIVE ARCHITECTURE

A technical explanation of networked video system implementation

3.1 MISSION-CRITICAL FEATURES

The main priority when designing a large networked video system is to balance cost of storage against the need to protect against video data loss or alteration. With this in mind, StorageTek's archive solution was architected for fault tolerance and designed to run in a SAN environment.

3.1.1 Storage area network

The SAN is a network architecture suited to alleviating many of the problems associated with the current proliferation of recording servers and storage. SANs consist of integrated hardware to provide a robust, high-speed fabric, enabling clusters of network video application servers to share storage arrays with exclusive data access or share data on common storage devices.

Advanced SAN features provide increased server performance, better optimization and high availability through features such as failover, load balancing and distributed applications. SANs further reduce the cost of online storage by allowing efficient data centralization and, therefore, less expensive administration.

3.1.2 Fault-tolerant architecture

The disk subsystem architecture consists of two major components: the disk controller and the disk arrays. The architecture is fully fault tolerant from dual processors, power supplies and host bus adapters (HBAs) to firmware features such as mirrored write cache and RAID level 0, 1 and 5 support.

The disk controller design employs two service processors (SP) configured for high availability. If one SP fails, cache is turned off and the healthy cache is flushed to disk. The I/O to the failed SP will be automatically re-routed to the healthy SP. Also, as part of the disk controller design, an integrated software driver component provides fault tolerance and recovery for the data path.

The disk array uses dual-ported disk drives, so that any drive is accessible through a secondary loop if the primary loop fails. All drives in the disk subsystem are available on two pairs of redundant Fibre Channel Arbitrated Loops (FC-AL), providing both rapid and highly available access to data.

If any one component fails, the second component will dynamically take over for the failed component. This design provides unsurpassed subsystem availability.

3.1.3 Mirrored write cache

An added advantage inherent to the disk controller design is the implementation of mirrored write cache. Competitive offerings provide mirrored cache, but the two copies reside on the same cache card, creating a single point of failure. The disk controller provides write cache mirroring on two separate cache cards, protecting against a cache card failure. The system also provides "battery-backed" mirrored cache that is flushed to disk in the event of a power outage.

3.1.4 Hot spare storage drive replacement

Using the system management console, any number of global hot spare disks can be configured. These spare drives automatically rebuild the disk array to a consistent state in the event of one or more disk drive failures, allowing the disk subsystem to continue storing video. During this period of reduced robustness, the system will alert the StorageTek service organization so that it can be restored to full health again without operator intervention.

Along with the archive's tape subsystem, intelligent software monitors and collects library statistics. Streamlined management optimizes space, increases availability and boosts access and archive performance. The administrator is also able to access and monitor information remotely using a standard Web-based browser.

Tape drives, fans and power supplies are hot-swapped (replaced without shutting down the system). Uninterruptible and redundant power supplies are also available. No routine maintenance is required.

3.1.5 High availability

The StorageTek disk system is designed to run continuously, ensuring that no security surveillance video is ever lost during periods of maintenance and service. This system can be dynamically changed on the fly in the following areas:

- > Additional capacity
- > Volume expansion
- > Firmware upgrades
- > Performance parameters and RAID levels
- > Performance enhancement (adding more bandwidth or increasing response times)

3.1.6 Early failure detection

Monitoring software oversees the subsystem's environmental aspects and routinely performs status checks. If the health of the system is ever impaired, the software is capable of posting a Simple Network Management Protocol (SNMP) trap and diagnosing the problem. The monitoring software has the ability to automatically detect any hardware component error. Components that log errors in excess of preset thresholds are flagged for action. Any error detected within the disk subsystem generates an SNMP alert and an e-mail message or a page is sent to a monitoring station.

3.1.7 Storing video for evidence

Along with providing a reliable and cost-effective digital video archive, it is important to protect the integrity of the video stored in it. On occasion, video is used as evidence in a court of law. Evidence of tampering or insufficient security in the system may prevent the video from being used successfully in court. It is the job of the network video application and digital video archive — collaborating with established manual and automated security processes in place for the appropriate procedures and mechanisms — to minimize the risk that any video captured and stored could be compromised.

From a storage management point of view, there are three critical areas for securing video data integrity:

- > Video transport over the network
- > Storage management software
- > Hardware used to physically store the video

3.1.8 Data confidentiality

How data is protected from unauthorized access and the security measures in place to make the data as tamper-proof as possible, are of paramount importance in legal cases. All systems use the following three methods to identify users, verify that they have the correct access and that their activity is logged.

Authentication.

To accurately identify a user trying to access a system or network, servers are protected at the operating system level and application layers, using different password systems to protect each layer.

Authorization.

Once identified, the systems and applications give each user only the level of privileges required to perform tasks particular to their business operation. Typically, an administrator will have a different set of privileges than an operator.

Audit trails.

Critical to verifying the confidentiality of the system, audit trails record all activity associated with any access or modification to the archive. Times are synchronized so auditors can compare and correlate the audit trail logs. In the event of unlawful behavior or tampering, the audit logs of the various systems are unlikely to correlate and the intrusion will clearly be highlighted.

3.2 DATA PROTECTION

3.2.1 Secure administration

StorageTek disk subsystems are managed using SAN management software. This software provides an easy-to-use, browser-based GUI management facility that monitors the array's performance, provides tools for automatic RAID volume configuration and enables remote visibility and management of the array. The simplicity and flexibility are also potential vulnerabilities, so access to the management interface is narrowly restricted to key administrators.

3.2.2 Disaster recovery

A key differentiator of StorageTek's digital video archive is its automation of day-to-day storage management that allows it to run in a "lights-out" environment. The digital video archive is automatically backed up to protect stored video assets. Multiple copies of the content can be cloned to one or more locations and to StorageTek's automated tape libraries. The digital video archive facilitates rapid recovery in the event of a disaster or complete loss of a control center. Copies of the video can be replicated to remote sites for immediate failover to a secondary command center.

3.3 IMPLEMENTING A NETWORKED VIDEO SYSTEM

By linking the inherent flexibility of networked video components to StorageTek's storage solution model, a large number of implementations can be constructed. These can scale from small- and medium-size solutions to large enterprise solutions. In addition, the networked video system can easily be integrated with the existing system to support the use of traditional CCTV devices.

In each site configuration:

- > Video servers connect analog cameras to the IP network.
- > Video servers connect to the Ethernet switch via 100 megabits-per-second links.
- > Digital cameras connect directly to the IP network.
- > A standard PC runs the NVR server application plus viewing and administration stations.
- > The NVR server connects to the switch via one 1000 megabits per second link; while a 100 megabits-per-second link would be more than adequate for 20 video streams, the 1000 megabits-per-second link removes any possible bottleneck.
- > End-user PCs running viewing software can also be connected to the network and operated as usual.
- > 10/100/1000 megabits-per-second Ethernet switch is used.
- > The StorageTek disk storage solution connects to the networked video system via a Fibre Channel interface.

3.3.1 Small site

A networked video recoding solution like this might be used for a small site employing a LAN infrastructure with video feeds from 20 cameras.




Figure 1. Networked video components for small-site implementation.

3.3.2 Medium site

This infrastructure is amazingly easy to scale using video over IP. As a network grows to hundreds of cameras, NVRs and storage are easily added. Below, another NVR server is included. Just by adding more storage and a network switch to the existing disk subsystem, the new NVR can now share this same subsystem.




Figure 2. Networked video implementation scaled to accommodate a medium sized site (hundreds of cameras).

3.3.3 Enterprise site with a SAN

Even as networked video systems scale to thousands of cameras with long-term archive requirements, cameras and NVRs are added as before. However, a SAN is added to the archive to interface with NVR servers and viewing stations, providing for a transparently shared, mixed-media storage pool. Storage management software, running on a digital video archive server, places and migrates video.




Figure 3. Networked implementation scaled to the enterprise level (thousands of cameras and long-term archive).

4 CONCLUSIONS

It is clear that networked video solutions are starting to replace existing analog technology in numerous security application areas. This seems to be a direct result of increased flexibility, reliability and cost effectiveness that improves a business's bottom line when compared to conventional analog solutions. When implementing a security system using networked video components, there are some critical success factors that must be addressed to achieve the desired business value:

- > Design and implement the system to meet the organization's operational requirements.
 - Work through the cost/quality/reliability tradeoffs.
 - Assess how much bandwidth is required and what is already available.
 - Identify functional requirements.
 - Determine how long and where the video information should be stored.
- > Select the appropriate video codec to suit the chosen application. H.26x offers unrivaled quality that uses minimal bandwidth. However, MPEG-4 is quickly becoming the defacto standard by combining picture quality with incredible flexibility.
- > Make sure the networked video system's individual components (cameras, switches, network, NVR application and storage) are certified to work together.
- > Make sure that all partners chosen to implement a security system have the required IP skills to succeed.

The storage solution is critical to networked video applications in the security market. There are many architectural and design issues to consider when creating a digital video storage archive. With the right choices, it is possible to cost-effectively meet growth demands, provide high availability, and maintain video integrity.

We must also consider what the future holds for networked video systems. It seems very likely that picture quality for both live and recorded video data will become a more significant issue in the near future. By moving away from the "closed" system model of CCTV into a new enterprise paradigm, the opportunities for delivering additional value through video data are almost limitless.

Networked video — See all, store all.

5 GLOSSARY

Bandwidth. A measure of the traffic volume that can be handled by a network or other physical carrier. Thinking of bandwidth as a data “pipe,” you can see that the larger the pipe, the more data that can pass through the network in a given period of time.

Cache. Pronounced “cash,” this special high-speed storage mechanism can be either a reserved section of main memory or an independent high-speed storage device.

Codec. Short for compressor/decompressor, a codec is any technology that compresses and decompresses data.

Fabric. The hardware that connects workstations and servers to storage devices in a storage area network.

Failover. A backup operation that automatically switches to a standby database, server or network if the primary system fails or is temporarily shut down for servicing.

Fibre Channel Arbitrated Loop (FC-AL). A serial data transfer architecture developed by a consortium of computer and mass storage device manufacturers and now being standardized.

H.26x. A video compression standard developed primarily to support video phones and video conferencing.

I/O (input/output). The term I/O is used to describe any program, operation or device that transfers data to or from a computer and to or from a peripheral device. Every transfer is an output from one device and an input into another.

IP (Internet Protocol). IP pertains to how information is sent across the Internet.

JPEG. Short for Joint Photographic Experts Group and pronounced “jay-peg.” JPEG is a compression technique for color images that can reduce file sizes to about five percent of their normal size.

LUN (Logical Unit Number). A unique identifier used on a SCSI bus to distinguish between devices that share the same bus.

Mirroring. The act of copying data from one location to a storage device in real time. Because the data is copied in real time, the information stored from the original location is always an exact copy of the data from the production device. Data mirroring is useful in the speedy recovery of critical data after a disaster.

M-JPEG. Stands for motion-JPEG, which extends the JPEG standard by supporting videos. In motion-JPEG, each frame in the video is stored with the JPEG format.

MPEG-4. Short for Moving Picture Experts Group and pronounced “m-peg,” the term refers to a digital video compression that achieves a high compression rate by storing only the changes from one frame to another, instead of each entire frame.

RAID (Redundant Array of Independent [or Inexpensive] Disks). A category of disk drives that employs two or more drives in combination for fault tolerance and performance.

Replication. The process of creating and managing duplicate versions of a database. Replication not only copies a database but also synchronizes a set of replicas so that changes made to one replica are reflected in all the others.

SAN (storage area network). A high-speed subnetwork of shared storage devices, a SAN's architecture works in a way that makes all storage devices available to all servers.

SNMP (Simple Network Management Protocol). A set of protocols for managing complex networks. SNMP works by sending messages, called protocol data units, to different parts of a network. SNMP-compliant devices store data about themselves and return this data to the SNMP requesters.

World Headquarters

Storage Technology Corporation
One StorageTek Drive
Louisville, Colorado 80028 USA
Phone: 1.800.877.9220

About StorageTek®

Storage Technology Corporation (NYSE: STK), a \$2 billion worldwide company with headquarters in Louisville, CO, has been delivering a broad range of storage management solutions designed for IT professionals for over 30 years. StorageTek offers solutions that are easy to manage, integrate well with existing infrastructures, and allow universal access to data across servers, media types and storage networks. StorageTek's practical and safe storage solutions for tape automation, disk storage systems and storage integration, coupled with a global services network, provide IT professionals with confidence and know-how to manage their entire storage management ecosystem today and in the future.

StorageTek products are available through a worldwide network. For more information, visit www.storagetek.com, or call 1.800.786.7835.

International Offices

Australia: 61.2.9438.4844	Chile: 56.2.201.0490	Hong Kong: 852.8200.0791	Malaysia: 603.772.41125	Singapore: 65.6774.9248	United Arab Emirates: 971.4.3900162
Austria: 0800.20.16.31	China: 86.10.8251.1911	India: 91.80.226.7272	Mexico: 52.55.9177.1800	South Africa: 0800.99.5820	United Kingdom: 0800.731.8852
Belgium: 0800.75.327	Denmark: 44836111	Ireland: 1800.55.33.54	The Netherlands: 0800.022.8496	Spain: 900.99.33.66	
Brazil: 55.11.3044.4599	Finland: 08001.13361	Italy: 167.790.852	Norway: 800.11.220	Sweden: 020.798.711	
Canada: 905.602.5586	France: 0800.82.83.57	Japan: 81.3.3746.9711	Poland: 48.22.5757.157	Switzerland: 0800.83.87.65	
	Germany: 0800.181.6238	Korea: 82.2.2191.1100		Taiwan: 886.2.6600.9840	

Specifications/features may change without notice.
© 2003 Storage Technology Corporation, Louisville, CO. All rights reserved. Printed in USA. StorageTek and the logo are trademarks of Storage Technology Corporation. Other product names mentioned may be trademarks of Storage Technology Corporation or other vendors/manufacturers.

StorageTek equipment is manufactured from new parts, or new and used parts. In some cases, StorageTek equipment may not be new and may have been previously installed. Regardless, StorageTek's standard warranty terms (as set forth at www.storagetek.com/warranties/) apply, unless the equipment is specifically identified by StorageTek as "used" or "refurbished." Replacement parts provided under warranty or any service offering may be either new or equivalent-to-new, at StorageTek's option.