



White Paper

Storage Takes the Leap to IP, the Choices and Challenges

This executive briefing for chief security officers (CSOs), directors and managers centers on the implications of digital video on the storage needs and image management with a focus on the operation and business bottom line.

With the growth of video surveillance and the coexistence of analog and digital as a measured step to the transition to digital video surveillance continuing, there is another leap to take involving storage, retrieval and image management. Storage, it turns out has both a price and a growing role to play in evolving systems. There are different uses and purposes of stored video, various locations where images can be stored, a myriad of storage types and business and operations issues to consider.

Once again, we'll share the thoughts and experiences of CSO Terry Jones and Helena Smith, his second-in-command, who work for a mid-sized enterprise. They now face the options of storage for their current security video solutions, as well as anticipation of future technological changes and management needs.

By the way, the first paper in this white paper series examined in an overview way the coexistence strategy at the heart of a cost-effective move from analog to digital security video. That white paper also overviewed the impact on infrastructure including sharing the enterprise data network, bandwidth and compression/decompression schemes. A second white paper explored cameras – analog to IP-based as well as megapixel and high definition. The third white paper examined benefits and ways that fiber optics enhances the operation and business bottom line of surveillance solutions. These previous white papers are available for download at www.infinova.com

Terry's memory of security video storage goes back years to a Las Vegas trip at a major industry conference. With other colleagues, he toured a casino security operation and was shown their videotape room, filled with what seemed like thousands of tapes labeled by date and camera location and managed by a dozen or so employees who rotated the tapes constantly. The cost of doing business is what the casino security manager told the assembled group.

Terry and Helena both know storage and video solutions have come a long way. They also know that their current storage needs and especially future IP video developments, using megapixel cameras, have purchase, maintenance and operations costs, too.

In many ways it is a matter of capacity, when it comes to storage.

That's what it is all about. In the world of computers, where digital security video resides, there are three types of memory: video, storage and RAM. Storage memory is on your hard drive and is used to store everything that's software or virtual on your computer, such as programs, files, settings, etc. There are bits, kilobytes, megabytes, gigabytes, terabytes and petabytes. A bit is the smallest increment of data on a computer. A bit can hold only one of two values: 0 or 1, corresponding to the electrical values of off or on, respectively. Bits are usually assembled into a group of eight to form a byte. A byte contains enough information to store a single character, like "m." A kilobyte is 1,024 bytes, not one thousand bytes as might be expected, because computers use a binary (base two) math, instead of a decimal (base ten) system.

A medium-sized novel contains about 1 megabyte of information or 1,024 kilobytes, or 1,048,576 (1024x1024) bytes, not one million bytes. Similarly, one 1 gigabyte is 1,024 megabytes, or 1,073,741,824 (1024x1024x1024) bytes. A terabyte is 1,024 gigabytes; 1 terabyte is about the same amount of information as all of the books in a large library, or roughly 1,610 CDs worth of audio or video data. A petabyte is 1,024 terabytes. Google, for example, processes about 24 petabytes of data per day, much of it video through YouTube.



Author
Mark S. Wilson
Vice President, Marketing
Infinova

Within security applications, it is apparent that with more cameras, more megapixel and high definition cameras, coupled with the need to look at more details for investigations and liability reasons, there are more security video bytes to place in storage.

Unit	Equivalent
1 kilobyte (KB)	1,024 bytes
1 megabyte (MB)	1,048,576 bytes
1 gigabyte (GB)	1,073,741,824 bytes
1 terabyte (TB)	1,099,511,627,776 bytes
1 petabyte (PB)	1,125,899,906,842,624 bytes

Source: University Information Technology Services, Indiana University

Security Video Storage

Intrusion detection, fire alarm and electronic access control systems have vastly less storage needs when compared to what a video and audio content security system produces. Few things can create data bits more rapidly than a digital video or analog signals converted to digital information. Compression can allow storage devices to hold more data as well as let transmission infrastructure handle more information within the limits of bandwidth.

During a planning meeting on security video storage, Helena points out to Terry that storing and recalling even highly compressed video involves data rates of millions of bits per second. And the video, with the addition of metadata – really labels for light or heavy intelligent retrieval and analytics – is often compressed for practical storage. Each purpose spawns its own set of technical and economic requirements, costs and approaches. Archived security video, for instance, can involve video at several quality levels.

Retrieval of video from storage is a key function and one that has grown more important thanks to the use of megapixel and high definition cameras which has spurred forensics applications in which more details allow security or law enforcement to see more of what happened. There are many different measures for evaluating the performance of video and image retrieval systems but all require a collection of data in storage and a query, often through a video management system or physical security information management system (PSIM). Every piece of an image is known to be either relevant or non-relevant to a particular query but there are shades of relevancy.

Total cost of ownership of storage involves a diversity of elements. There is the cost of the devices themselves as well as their functionalities. Different types of storage are differently priced. So that once Terry and Helena decide the number of cameras, image size and resolution, and whether cameras will be recording continuously or motion-activated, calculating the necessary storage capacity can begin.

There are tools to determine storage. For example, another Infinova white paper -- How to Estimate the Network Bandwidth & Storage Needed for IP Surveillance Systems – explains the complexities of measuring security video for storage.

For instance, in a typical surveillance project, there are several factors that affect the amount of storage required by installed security cameras including:

- Type of camera, analog or IP
- Resolution, 1.3 megapixel, 3 megapixel, or others
- CODEC, MJPEG, H.264, or others
- Frame rate, 5 fps, 30 fps, or others
- Dual streaming, separate streams for viewing and recording
- Motion in the scene
- Recording periods during the day
- Levels of activity through the week

Storing bits is like operating any complex filing system. The size of the container must fit the application. But, as Helena points out, no practical universal storage medium exists, and all forms of storage have some drawbacks. She says that, when evaluating storage options, core characteristics are volatility, mutability, accessibility, and addressability. And high density, low cost, high reliability, and fast read and write times are all needed for many video applications.

Future Proofing the Storage

Scalability is another storage issue. Security video systems tend to grow and change, especially in regard to coexistence of analog and IP video, when migrating to all digital video and with megapixel and high definition cameras. Data storage at the edge can provide useful when facing bandwidth challenges or in building in redundancy. Digital video recorders (DVRs) are among the most used security video storage devices. Still, when it comes to scaling, DVRs can add to cost since they handle a set number of cameras within a set geography.

DVRs are generally available with 4, 8, 16 or 32 inputs while network video recorders may have limits based upon licenses. In both cases the amount of storage space needed may affect selection. For example, Terry has some 16-camera DVRs with only eight cameras so that he can take advantage of the larger hard drive size.

Some DVRs are limited in the number or capacities of hard drives they can handle. Some require two to three minutes of rebooting when installing new hard drives. Others have hot-swappable hard drives, which mean no rebooting is necessary. The drives can be installed while another hard drive in the DVR is being used. The ability to add external hard drives also is an important factor.

Overall, however, as Terry points out during the storage options meeting with Helena, Moore's Law predicts increasing capacity and decreasing cost of storage. Moore's law describes a long-term trend in the history of computing hardware, in which the number of transistors that can be placed on an integrated circuit has doubled approximately every two years. It often refers to processing speed, memory capacity, sensors and even the number and size of pixels in digital cameras. The law is named after Intel co-founder Gordon E. Moore. But, due to limitations imposed by the laws of physics, the curve that plots Moore's Law will eventually flatten out. Still, modern engineering has advanced recording techniques much further and much faster than developers predicted just a few years ago.

At the meeting, the two security executives review the specific types of storage.

Videotape and time-lapse recorders are an historic means of image storage for consumers and security operations, including that casino Terry visited way back when. Videotape is certainly not yet dead. It is cheap, ubiquitously available, portable and highly reliable, not to mention quite familiar. But, when compared to other techniques, it suffers from limited shelf life and long seek times that make nonlinear read/write impractical.

Flash memory is a non-volatile computer storage that can be electronically erased and reprogrammed for memory cards and USB flash drives for storage and transfer of data between computers and other digital devices. Flash memory stores information in an array of memory cells made from floating-gate transistors. Flash memory costs far less than other technical approaches and is the dominant technology wherever a significant amount of non-volatile, solid state storage is needed in digital cameras, for instance.

Since flash memory is non-volatile, no power is needed to maintain the information stored in the chip. In addition, flash memory offers fast read access times and better shock resistance than, say, hard disks. It is useful for embedded applications. Yet, flash has storage size limitations (though there are 256 gigabyte units) and has a finite number of erase-write cycles.

SD Cards in the Camera

A type of flash memory, secure digital (SD) is a card format originally developed by Panasonic, SanDisk, and Toshiba for use in portable devices. Standard SD cards have a maximum of 4 gigabyte capacity. SDHC high capacity

cards can often hold 32 gigabytes while even newer SDXC extended capacity cards handle two terabytes of storage. SD is there to trigger recording when there is an alarm, among other techniques.

Video cameras equipped with an SDHC card can often store a couple of days worth of video recorded at 30 frames per second. In the event that greater recording time is needed, some video cameras have a connection for a hard drive or other external storage unit.

A hard disk drive (HDD) is a non-volatile storage device that stores data on rapidly rotating rigid (hence the hard part) platters with magnetic surfaces. HDDs are in personal and desktop computers and some now can hold up to two terabytes of storage. Laptop units hold much less. A typical hard drive has two electric motors, one to spin the disks and one to position the read/write head assembly; and because of that, there is a lifespan for all HDDs.

When it comes to optical devices, there are CDs and DVDs and their read/writers, often used to capture security video of incidents for forensics and law enforcement use. The best DVDs hold about 9 or so gigabytes. Newer Blu-ray technology handles high definition.

Embedded and Network DVRs

A digital video recorder (DVR) is a device that records video in a digital format. It has software and personal computer hardware, which enables video capture and playback to and from disk. DVRs can be stand-alone but more often these days, points out Helena, are part of an overall security video network or hang from the overall corporate network infrastructure. Thanks to such consumer applications as TiVo and satellite/cable set-top boxes, when it comes to security applications, many DVRs have a myriad of advanced functions such as searches by event, time, date and camera. There is also much more control over quality and frame rate allowing disk space use to be optimized and the DVR can also be set to overwrite the oldest security footage should the disk become full. Some include compression capabilities and/or video analytics firmware.

Security DVRs are either PC-based or embedded. A PC-based DVR's architecture is a classical personal computer with video capture cards. An embedded device is specifically designed to have its own operating system and application software contained in firmware or read only memory. Embedded operating systems may be less susceptible to an outside attack. If the recorder uses a non-embedded operating system, protection software, such as anti-virus, may be installed.

A network video recorder or NVR is an Internet protocol-based device that sits on a security or enterprise network. NVRs can be managed remotely via a local area network or over the Internet. NVRs are used in IT-based systems to expand capacity beyond DVRs.

Most DVRs and NVRs have the ability to be remotely accessed either over a LAN, WAN, the Internet either by using proprietary licensed software or through a Web browser.

From DAS to NAS to SAN

At this point, Helena provided a quick technology briefing on direct and networked storage.

Direct-attached storage (DAS) is traditional mass storage, Helena reports, that does not use any network.

Network-attached storage (NAS) is mass storage attached to a computer which another computer can access at file level over a local area network, a private wide area network, or, in the case of online file storage, over the Internet.

Storage area network (SAN) is specialized network storage that provides other computers with storage capacity. The crucial difference between NAS and SAN is, NAS presents and manages file systems to client computers, while SAN provides access at block-addressing (raw) level, leaving it to attaching systems to manage data or file systems within the provided capacity. SAN is commonly associated with fiber networks. In other words, NAS solutions are computers dedicated to data storage. SANs are networks of storage devices that could include servers and digital tape libraries that may be remotely located but can be accessed as if they are locally attached.

When it comes to storage options, Helena thought it best to help Terry with some other terms.

Serial Attached Small Computer System Interface (SCSI or SAS) is a computer bus, which moves data to and from computer storage devices such as hard drives. SAS depends on a point-to-point serial protocol that replaces the parallel SCSI bus technology from the 1980s. Serial ATA or SATA is also a computer bus interface for connecting host bus adapters to mass storage devices such as hard disk drives and optical drives.

SAS is full-duplex, SATA is half-duplex. SATA follows the ATA command set and thus supports hard drives and CD/DVD drives.

In addition, Internet Small Computer System Interface (iSCSI) is an IP-based storage networking standard for linking data storage facilities. By carrying SCSI commands over IP networks, iSCSI facilitates data transfers over intranets and manages storage over long distances. It's used to transmit data over local area networks, wide area networks, or the Internet and can enable location-independent data storage and retrieval. Unlike traditional fiber channel, with its special-purpose cabling, iSCSI runs over long distances using existing network infrastructure.

RAID or Mirroring?

Another "buzzy" word, according to Helena, is RAID and aims at storage reliability and input/output performance. Redundant Array of Inexpensive Disks (RAID) permits high levels of storage reliability from low-cost and less reliable PC-class disk drives. Various RAID data storage schemes divide and replicate data among multiple hard disk drives. Data can be distributed across multiple disks, while the array is seen by the user and operating system as one single disk. Some systems create backups automatically, called mirroring. Different types of RAID recording are specified by numbers. A popular RAID recording method practiced in the security industry for video surveillance is RAID 5. Terry and Helena's integrator has recommended RAID 5 instead of mirroring. Mirroring is using one additional hard drive to make a backup copy instead of several. When several hard drives are used for backup, they can all be writing portions of the data at the same time. The risk of loss of the recorded security video data also is spread over several or more drives. Because mirroring uses a single hard drive for backup it can be slower in comparison.

Near the end of the tech briefing, Terry asks about the cloud for storing and retrieving security video.

A cloud example he is familiar with is Gmail from Google. Gmail for business, for example, offers 25 gigabytes of storage per user. Cloud computing is Internet-based computing, in which shared resources, software and information are provided to computers and other devices on-demand, like a public utility.

Among those things that both Terry and Helena agree is to forget about how many terabytes to buy. Instead they will use proposal terminology such as "system to record 210 cameras at 4 CIF RAID 5 for no less than 20 days.

Practical Needs on top of Technological Ones

There are practical approaches to security video storage beyond the various technological ones. When you can store seemingly unlimited quantities of video, Terry figures, the question becomes how much do you really need and at what frame speed and resolution.

It may be that larger hard drives are not the one answer. Often in smaller installations or in remote locations, a smaller DVR will do the job, especially if the frames per second recorded is small. Sometimes an integrator will use a consumer grade PC to build a recorder to handle cameras but there are obvious dangers in this approach. In the case of IP and high definition cameras, storage needs change to external options. Here DVRs can feature multiple, reasonably inexpensive SATA hard drive bays. And with MPEG4 and H.264, a lot of video data can be stored at real-time image rates. Simplicity of set up and operation for the end user means a one-box embedded solution can be a winner for some.

Motion detection and lower frame rate can make a storage difference. Picking higher quality cameras and digital noise control can also impact storage requirements. And good cabling also helps. In such cases, there are fewer pixels of distortion to store.

Preferences, corporate rules, liability concerns and government regulations all may impact how long, how easy to retrieve and in what condition images are stored. Harkening back to his casino visit, Terry remembers that now-a-days digital storage in the gaming industry is determined by federal, state and Native American gaming regulations on a case-by-case, state-by-state basis.

Isolating Specific DVRs

In such large camera installations, where often there are two systems – for instance, one for gaming and one for hotel, parking lot and retail, there may be simultaneous recording and playback of multiple cameras in real-time, which allows security personnel the ability to isolate specific DVRs for dedicated operations without disrupting the balance of the system. Terry and Helena still mull over the length of time they wish to store security video. Helena wishes they could keep everything forever. But Terry says that's a price tag he could not afford. Burning a DVD of an incident could allow a more affordable permanent record. Most requests for proposals set a 30-day storage period.

Effect of Image Size on Storage Requirements

Image size in Kb	Frames per second	Mb per hour
10	1	36
15	1	54
30	1	108

Actual storage space may vary.

For longer term storage, if it is really needed, Terry sees value in digital tapes and even off-site storage. It is cheaper and more reliable than 24/7 hard drive operations. I like the idea of having multiple storage levels and places, he tells Helena.

In another way to look at it, different types of storage can be placed at different locations, for example, storage at the edge of a system in hard drive or flash memory in cameras or on servers in network closets to save bandwidth. Helena points out that some storage in the camera could be a day of storage, in the closet a week, and in a SAN or data center a month of storage for backup and redundancy.

What about Resolution and Frame Rate?

Quality of images also affects storage needs and pricing, although less so as storage capacity increases while prices drop. Helena again points out that with analog cameras co-existing with a digital recording system, if we don't do it right, there will be noise which will generate additional storage requirements. Measures of video resolution, such as 480 horizontal lines, 640 x 480 pixels, or 4 CIF, may not provide a complete picture. Helena suggests they measure the visual details that remain in the picture during playback after recording, which is totally different.

Image or frame rates also affect storage. No doubt, full-motion video at 30 fps is not used in many applications because less may be sufficient. It may also depend on how many cameras get recorded by a DVR, with the ability to adjust frames per second for each camera based on its location and scenes. In that gaming example, a casino camera would be 30 fps, as opposed to a backroom operations camera, which would record at 15 fps or lower.

Effect of Frame Rate on Storage Requirements

Image size	Frames	Mb
------------	--------	----

in Kb	per second	per hour
15	1	54
15	5	270
15	30	1,620

Actual storage space may vary.

Terry wants it both ways: low frame rate when nothing is happening, but when motion detection or an alarmed event occurs, recording is at high resolution and maximum frame rate.

Compression Can Play an Important Role

Video storage needs are in part determined by the security video compression method. Newer compression tech such as MPEG-4 and its variant H.264 refresh only the parts of the screen that change from one frame to the next, helpful for maximum recording speed and longer storage capacity.

Another element in the security video storage equation is the organization's IT department, which may have standards and hardware testing before a storage device can hang from the network. Some security executives welcome the involvement and ownership by IT of various parts of a video surveillance system while others see value in owning most of their gear.

In the next white paper: how to perform a site security survey and building of a risk/vulnerability matrix.

Embedded and Network Recorders, CD/DVD Writers

Infinova has embedded digital video recorders, network digital video recorders and CD/DVD writers. For example, the V3009 Series Embedded Digital Video Recorder has H.264 video and OggVorbis audio compression technology, NTSC/PAL supported, 1 BNC primary output, 1 BNC matrix output, 1 VGA output, up to 16 video looping outputs, 4 audio inputs (only for 4/16-channel models), 1 audio output (only for 4/16-channel models), up to 16 alarm inputs and 4 alarm outputs, supports video stream and composite stream (for 4/16 models with 1 to 4 channel), 4 SATA interfaces and supports 4 SATA hard disks, among other specifications.

Infinova has a stable of technology that emphasizes the coexistence strategy which can gradually and intelligently migrate from analog to IP video. There are analog fixed cameras and fixed minidomes, PTZ domes and PTZ cameras, and IR illuminated cameras. There are IP fixed cameras and fixed minidomes as well as PTZ domes and cameras.

Infinova[®]

By helping channel partners provide their customers with complete, affordable, best-in-class, large and small video surveillance solutions, Infinova helps integrators generate more business more profitably. Leveraging a manufacturing process certified to ISO 9001:2000 standards and over 250 engineers with a list of video industry firsts, Infinova channel partners provide their end-users with industry-acknowledged product reliability and technical leadership.

So that Infinova channel partners can create complete solutions, Infinova provides IP surveillance cameras and components, CCTV analog cameras, DVRs and components, camera accessories, monitors, power supplies and fiber optics communications devices. Infinova also has the technical ability and manufacturing flexibility to let integrators propose customized solutions. In addition, Infinova will partner with other manufacturers making other surveillance equipment and software to help its channel partners create turnkey solutions. Contrary to most other companies, Infinova will back-up their partners' products as well as its own to assure both the integrator and its customers that one call – to Infinova only – takes care of everything.

Infinova works diligently to assure its channel partners can provide cost-conscious solutions. With Infinova's hybrid systems, channel partners can propose systems that protect a customer's investment in its already-installed analog surveillance system but that also put them on a dynamic migration pathway to IP systems.

Infinova is lauded for its exceptional maintenance programs. A major highlight is the company's 24-hour advanced replacement policy in which a substitute product is shipped immediately upon notice of a problem.

With such customer focus, Infinova is often referred to as "the integrators' manufacturer."

Global Contact Information



United States

Infinova

51 Stouts Lane

Monmouth Junction, NJ, 08852

United States

Phone: +1 732-355-9100

+1 888-685-2002 (toll-free)

Fax: +1 732-355-9101

Email: Sales@infinova.com

Latin America

Miami: +1-954-990-0787

Mexico: +52-55-5392-1735

Venezuela: +58-212-336-0661

Brazil: +55-11-7479-5640

Email: Sales-LAR@infinova.com

Europe

Phone: +40 2 6841 5582

Email: Sales-EUR@infinova.com

Middle East

Phone: +965 247 5678

Email: Sales-ME@infinova.com

India

Sales: +0091 9980728579

(South and East India)

Email: Sales-IND@infinova.com

Hong Kong

Phone: +852 2795 6540

Email: Sales-HK@infinova.com