Data Storage for Video Surveillance

In the past, it was an expensive proposition to deploy a video camera. The camera itself was expensive, and it required a dedicated closed circuit television network to route the recorded information to its destination, usually a video tape recorder. These factors limited the number of cameras that could be deployed.

Today's IP video cameras are affordable, and since they are connected and transmit data via Ethernet, they are easy and inexpensive to deploy. They are also capable of higher frame rates and higher resolution. This higher frame rate and higher resolution are two of the factors driving the need for more storage capacity in video surveillance. Other factors are an increased number of cameras and longer retention times.

While some legacy systems still use tape storage, video data today is primarily stored on PCs or Servers using hard drives. Storing video data in this manner makes utilizing video with higher resolution and frame rates possible. Hard drives are also less expensive and more durable than tape. Digital video surveillance data storage is similar to IT data storage in some ways and different in others. They both use disks to store data, and data travels via Ethernet to the disk storage array. Both IT and surveillance data need some type of protection against disk failure, usually in the form of Redundant Array of Independent Disks (RAID). Both types of data need a disaster recovery strategy, such as a NAS with snapshots.

One major difference between IT storage and video storage is in the way the data is written. In an IT environment, data is being written many times and read many times as a matter of course. In a surveillance environment, data is constantly being written – because the cameras are constantly recording video. The surveillance data is rarely read. It is only when an incident occurs that review of the video is required and a disk read occurs. In most applications, that is a somewhat rare event. The data is usually kept for a specified period, then it is overwritten with new video, and the cycle repeats. The period the video data is kept may vary in length, but the routine of repeated writes and infrequent or nonexistent reads is the norm.

The level of risk associated with the loss of video data usually fall into three tiers: low, mid, and high. Low-tier applications are more concerned with cost, and have little to risk if frames are dropped or video data is lost. They can tolerate lower frame rates and lower resolution, and a drive failure is not a major concern. An example of a low-tier surveillance application is a small business or office with eight cameras or less.

Data Storage for Video Surveillance

Mid-tier Surveillance applications typically have more cameras, require higher resolution, and data is more important to them. While the objective is that no video data is lost, a minimal amount of dropped frames or lost data can be tolerated. Small to medium sized businesses and campuses fall under this tier.

High-tier applications have little or no ability to tolerate loss of video data. The cost of protecting and retaining the data is a consideration, but it is secondary. The retention period is typically longer than mid or low-tier, and they may have one hundred cameras or more in the network. There may be performance requirements that dictate which type of RAID must be used in order to guarantee a minimum level of performance when a disk has failed and performance is degraded. Prisons, airports and casinos fall under this tier. A casino may actually have to close if their video surveillance system is down.

The Tier a business falls under will dictate many of the characteristics of the storage system required.

There are a number of alternatives to choose from for storing video data. Some vendors offer proprietary solutions consisting of DVR devices. These can be very pricey and usually lock out all other vendor's hardware and software. Often these vendors offer some functionality or software that they claim makes their system superior. While there may be some small degree of benefit from proprietary features, it is usually offset by the higher cost of the hardware. Unfortunately for many who purchase proprietary systems, when it becomes necessary to upgrade they discover that they face a forklift proposition, that is, all equipment must be replaced in order to upgrade to the vendors' new proprietary hardware and software.

Cloud storage is another option for video data storage. One reason that is often cited for choosing cloud storage is cost. A quick analysis of the cost of cloud storage vs. purchasing on-site storage in the form of NAS or SAN devices will quickly prove that cloud storage is not less expensive; in fact it is often a lot more expensive. There are a number of other considerations to consider regarding cloud storage of surveillance data:

- *Data must be transmitted via the Internet*. If there is an Internet outage, data is not being stored.

- *Data must be transmitted via the Internet*. This imposes limitations based on bandwidth. These include compromises in video resolution, frames per second, number of streams, and other factors that could cause bandwidth to be exceeded. It is easy to set up an internal network that is capable of 1GB or 10GB speeds, but 1GB and 10GB connections to the Internet are usually quite expensive.

- *Data must be transmitted via the Internet* and stored in an unknown place and subject to scrutiny by unknown personnel. There are many transmissions and servers involved, and it is hard to protect data against theft or deletion. Securing video content from accidental or malicious deletion is of utmost importance, and this cannot be done if the physical device on which the data is stored is not known – in most cases the location of the device is unknown.

Data Storage for Video Surveillance

> If security and integrity of surveillance video is a prime consideration, placing it on the Internet where it is subject to hackers and loss of control over those who have access to it is very problematic.

The recurring theme is that surveillance data that is transmitted via the Internet can be problematic for a number of reasons, all rooted in the basic question of “who has access to the data?”

Storing any kind of sensitive data in a manner where the chain of possession cannot be documented will certainly raise issues. This is especially true for data that is subject to compliance restrictions. If surveillance data is not critical (Low-tier or mid-tier), the cloud may be an option, albeit a more expensive option. If having possession of the data is a necessity, or subject to compliance restrictions, owning and controlling the physical devices where the data is stored is a much more secure solution.

Owning and keeping the storage servers in an access controlled environment means the enterprise has complete control over access and chain of possession, making the choice of on-site storage the most secure option.

Storage needs for an IP video system will be determined by a number of factors. These include:

- Number of cameras
- Resolution of cameras
- Frame rate of cameras
- The length of storage time
- Will video be recorded continuously or upon a trigger (Motion or alarm)?

All of these factors affect the amount of storage that will be required. As with any installation of digital storage or computers, it is a good idea to plan for the future, and to plan for expansion. A surveillance plan that has been checked and double-checked for completeness often has another camera or two added before or shortly after deployment.

The first step in choosing a storage system should be selecting a NAS that is capable of scaling out, because it is a safe bet that additional video streams (cameras) will be added, and future cameras will be capable of higher resolution. Factors such as these will result in increased demands for storage space, and a NAS system capable of scaling out will accommodate these needs.

In some cases the needs driving the surveillance system change. It is a good idea in the planning stages to anticipate changes of the purpose for which surveillance video is used. For example, an application that was intended to capture low-resolution security images may evolve into a tool for determining demographics of people frequenting an establishment. A low-res camera may not provide the quality of image to gather the desired data. A crime or lawsuit that is decided or dismissed because of insufficient or poor video data could certainly change the requirements of the system. In extreme cases, a video surveillance system may evolve into a facial recognition system, or feed some other video analysis system. Any of these evolutions of the original system would require higher resolution, higher frame rates, and more storage capacity.

DISK DRIVES

A dentist once said one only needs to floss the teeth one intends to keep. In a way, this applies to digital video data stored on disks as well. One only needs to purchase high quality drives to store the video one wishes to keep. Of course, surveillance systems would not be purchased, installed, and utilized if the video data being gathered was not important. It is only reasonable to purchase and install disk drives of sufficient quality and construction to reliably store the data.

Hard disk usage in video applications is different than hard disk usage in IT environments. In IT environments, the 80 /20 rule applies – 80% read and 20% write. In video surveillance applications, it's often 0 / 100%. Data is rarely, if ever read (the video is not reviewed unless an incident has occurred to warrant it), but the cameras are constantly writing. Or at least writing when there is motion, in the case of motion triggered systems.

The size of hard drives is increasing, and the cost per terabyte is dropping. This makes on-site storage of video data affordable. Hard drives can be broken into several groups. Consumer class hard drives are designed to be used in stand alone PC's, and to operate in short duty cycle applications. They have no protection against the vibration caused by multiple disk drives in a rack installation. They should not be used in commercial video surveillance applications.

At the opposite end of the spectrum, enterprise class drives are designed for the most demanding applications, including 24/7/365 applications. They have vibration-dampening technology that protects the drive from vibration caused by multiple rack-mount devices filled with hard drives. Enterprise class drives are the most rugged hard drives available and if cost is no object, they will offer the greatest data protection. But cost is always an object, and enterprise class drives may deliver more ruggedness than is required for surveillance video applications.

There is a recent addition to the classes of hard drives called surveillance grade drives. These are drives that are purpose-built for surveillance applications. These drives typically do not have the vibration dampening technology that enterprise class drives have. Care should be exercised when choosing surveillance class drives to make sure that they are appropriate for the application. When choosing surveillance grade drives pay particular attention to the specifications concerning the physical mounting of the drives. Be sure the drives are compatible with the intended installation by answering the following questions:

- Is it approved for rack-mount devices?
- What is the maximum supported number of drives per chassis

Failure to observe and follow these specifications can result in premature drive failure due to vibrations beyond the tolerance of the drive.

Data Storage for Video Surveillance

RAID STRIPING

Since video surveillance is write intensive, and since most applications will use a type of RAID that involves striping, there is an opportunity to configure the RAID controller to optimize the storage system for the best video storage performance. The setup will be a bit different than for IT storage applications. Select a RAID device that allows the configuration of longer stripes of data – 128KB or even 256KB. This will help achieve maximum performance from the storage array.

Being able to produce a video record of an incident may be the deciding factor in an expensive lawsuit, a crime of property, or even identifying a situation that could determine life or death. Making the best choice in video surveillance storage systems and configuring them optimally may make the difference in capturing needed (or mandatory) video or discovering the video data was dropped or lost.

NFINA SOLUTIONS

Nfina Technologies recommends the following systems for surveillance applications:

For systems with fewer than 30 cameras

Nfina 214i2 rack-mount server
Featuring an Intel® Xeon® v3 processor, up to four hot-swap drives, and 32GB of memory, this versitle server offers high performance and energy efficient design.

Nfina 418i2 desktop tower server
Featuring an Intel® Xeon® v3 processor, up to eight hot-swap drives, and 32GB of memory, this desktop server is ideal for small business, sercurity applications.

For systems with 30 cameras or more

Nfina 724i20 rack-mount server
Featuring dual Intel® Xeon® v4 processors, 12Gb/s SAS connectivity, and LSI® hardware RAID, this powerful server has up to twelve hot-swap drives and 768GB of memory and is ideal for large, rugged, security applications.

Nfina 428i2 desktop tower server
Featuring dual Intel® Xeon® v4 processors, 12Gb/s SAS connectivity, and LSI® hardware RAID, this eight bay, sever offers high performance operation in a compact desktop package. This server is ideal for large, enterprise, security applications.