



Reducing Costs While Enhancing Reliability of Video Surveillance Storage Systems

By John Minasyan, Senior Product Manager Pelco

Introduction

The migration of video recording media from analog VCR tapes to digital hard disk drives has been motivated by a desire to reduce costs while maximizing the number of cameras and days of retention achieved on a finite amount of storage capacity. The industry as a whole has experimented with ways to manage this cost by reducing the size of recorded video. Today, a great deal of energy is being expended on leveraging existing IT infrastructures and potential idle capacity that most IT departments have to record surveillance video. The question is whether there is a more efficient way of achieving the same goal.

A properly designed storage system built for the unique requirements of IP video surveillance can offer a more robust, scalable and cost-effective solution. Additionally, a properly designed storage sub-system can help customers take full advantage of the benefits of higher resolution and real-time video produced by today's advanced IP cameras.

This white paper explores the unique requirements that surveillance video places on the recording sub-system as well as the elements that make up the total cost of ownership of the system. It concludes by examining how Pelco's second-generation recording system addresses these major concerns while delivering a cost-optimized solution for IP video surveillance.

The Need for a New Paradigm

The security industry has had decades to understand and learn how to accommodate the massive storage requirements for video data. In an effort to mitigate its costs, the industry has taken various approaches that have mainly revolved around reducing the quality of the recorded video or the frequency with which video frames are recorded on the storage media. In recent years, however, the advent of and attraction to higher resolution cameras, combined with regulations that require extended retention periods of high-resolution, real-time footage, have made these practices useless and obsolete.

The use of VCR tapes for surveillance footage popularized time-lapse recording. In order for a single VCR tape to hold more content, users reduced their expectations for both video quality and frame rate. The goal was simple: record up to 16 cameras over a 24-hour period on a single tape. This was achieved by slowing down how often an image was recorded on the tape (60 fields/sec for standard 2-hour tapes, 10 fields/sec for 12-hour mode, and 5 fields/sec for 24-hour mode) and ultimately by multiplexing each of the 16 cameras to record to the same tape. However, people soon realized that even with the use of multiplexers, recording video at a very low frame rate yielded footage that simply had too much gap between subsequent images from the same camera to be of any value for investigators.

The era of recording on digital hard disk drives brought additional elements into the mix. The concept of increasing frame rate upon an event or alarm became standard when recording to a digital video recorder. However, the cost of hard disk drives and the speed with which video filled those drives continued to force users to accept less-than-optimal video quality. Additionally, video prior to and after the event was frequently needed to piece together the entire story of what occurred. Even with the ability to increase frame rate on alarm, vital information was often times not available. And this concept typically missed those events that, while important, did not meet the established criteria for high frame rate alarm recording.

Contents

- 2 Introduction
- 2 The Need for a New Paradigm
- 3 The Need for Purpose-Built Solutions
- 7 The NSM5200: Designed for the Rigors of Video Surveillance Storage

Today's IP cameras can produce a significantly better picture than their recent analog predecessors. And with the emergence of megapixel cameras, users have begun to reevaluate their willingness to compromise on quality and frame rate. (The picture produced by a 2.1 megapixel camera displayed on a 1080p Full HD monitor is truly remarkable, especially compared to playing back video that was captured at CIF resolution, at a time-lapsed frame rate.) Thankfully, a storage sub-system with the proper architecture can now enable users to take advantage of these capabilities without breaking the budget.

Technology is just one element driving the need for a new paradigm. New regulations have begun mandating minimally acceptable quality, frame rate, and retention periods. The backlog in criminal cases in California courts mandate that video evidence be archived for 365 days. Digital video recording standards for gaming applications in the state of Nevada require that video footage for all gaming cameras be captured at 30 images per second at a minimum of 4CIF (704x480) resolution for seven days. France's recently enacted public safety laws require 4CIF (704x576) resolution, real-time video for all public safety and city surveillance applications.

End-user demand is also increasing for higher quality video footage, driving a need to deliver a cost-effective recording solution for today's IP and megapixel cameras. Combined with aforementioned regulatory agency requirements, traditional recording and storage practices are becoming increasingly outdated and irrelevant. As a result, today, more than ever, the industry is faced with finding a cost-effective, highly reliable, scalable storage architecture to satisfy these needs.

[The Need for Purpose-Built Solutions](#)

Digital video storage is unique and differs from traditional data in the following ways:

- Database Size
- Write Intensive Operations
- Intolerance of System Latencies
- Constant FIFO of Recorded Footage with Sustained Bit-Rate Streaming
- Operating Environments

[Database Size](#)

Network administrators are experts at managing storage farms. However, most IT professionals are not accustomed to dealing with the massive capacities that video storage requires.

To illustrate this, let's take a simple example of an installation with the following attributes:

- 100 cameras
- MJPEG compression with medium quality
- 640x480 resolution
- 15 images per second recording per camera
- 24 hours/day of recording
- 7 day retention target

A 640x480 resolution image at 15IPS should produce roughly 4Mbps of bit-rate that needs to be stored. Simple multiplication yields the following result:

$$4\text{Mbps} * 60 \text{ sec/min} * 60 \text{ min/hr} * 24 \text{ hr/day} * 7 \text{ days} * 1/8 \text{ byte/bit} = 302\text{GB}$$

In this scenario, each camera will require 302GB of storage to meet retention requirements. This may not seem like much, but when you add the storage requirements for all 100 cameras, the total capacity needed becomes 30.20TB.

To put this figure into perspective, the entire amount of storage used on the Pelco campus – with 1,500 PC users, software and databases required to sustain the needs of full-blown fabrication and assembly factories, inventory management and logistics systems to support five warehouses around the world, and storage capacity needed for sales, accounting, and CRM systems – is only 32TB. In fact, most enterprise-class IT storage systems are well under 100TB of total capacity.

The above example was for 100 cameras recording for seven days. Consider installations that have over 5,000 cameras or installations that require 365 days of on-line video storage. As one of the requirements of video surveillance systems is the ability to support the needs of different types and sizes of cameras, well-designed video storage systems must be capable of efficiently scaling from a few gigabytes to petabytes. What's more, storage systems designed for video must offer the same degree of data protection and safeguards as enterprise-class IT storage systems, but at a fraction of the price.

Write Intensive Operations

Surveillance cameras never cease streaming content – and when doing so, nearly 95% of the operations executed on a hard disk drive are write operations. In the grand scheme of things, retrieval of recorded content represents an extremely small percentage of the workload that video places on a storage sub-system.

This is a complete reversal of what most enterprise-class storage solutions are designed to do. For example, consider the workload placed on a server utilized by Youtube.com. Here, a video file only needs to be written once and can be served up millions of times as viewers download the clip. In fact, numerous other applications within a typical enterprise environment work on the same principle. For this reason, most storage systems are designed to support fast access times for read operations while compromising access time allocated for write operations.

The problem that a mismatched storage system introduces for video surveillance applications is unnecessary cost. If the storage sub-system is not tuned for massive amounts of write operations, the cost of deploying a storage solution increases dramatically. To accommodate a lack of write operation throughput, NVR servers need to buffer incoming streams. Not only do more servers increase capital expenditures, but they also consume additional power, cooling, network ports, and rack space that increase the on-going operational budget of the system.

Intolerance of System Latency

Video surveillance systems are designed to reduce end-to-end latency as much as possible. As such, components including the camera, recorder, decoder, etc. have very little buffer space available. Conversely, buffer and re-transmit schemes are used widely in enterprise storage systems to accommodate asynchronous requests from various servers. Therefore, while enterprise storage systems have very impressive throughput numbers, the ability to service an incoming request from a given server is not a deterministic activity. The asynchronous manner in which storage systems service incoming requests can wreak havoc on a video surveillance system.

Let's demonstrate what can happen if the recorder is non-deterministic in how quickly it serves write requests. Consider a system with the following parameters:

- Records to SAN storage
- Controller throughput of 1 Gigabit per second (most modern enterprise storage controllers have far higher throughput numbers than this, but we'll use 1Gbps to make the math easier to follow)
- 10 servers running NVR software connected to the same storage controller

The theoretical throughput available per server is 1000Mbps/10 servers or 100Mbps per server. On the surface, 100Mbps seems like sufficient bandwidth for recording several standard resolution IP cameras (assuming 5Mbps per camera yields 20 cameras per input). However, this throughput is not dedicated and it is not guaranteed. SAN controllers only guarantee that over a statistically relevant amount of time, each controller will make available, on average, 100Mbps of throughput. As a result, over a ten-second time interval, all ten servers would be able to transmit, on average, 100Mbps. Averaging, however, is what makes this approach dangerous for video surveillance storage. Server 1's allocation of 100Mbps can easily take the form of 1Gbps for a second and nothing for the next 9 seconds ($1000\text{Mbps}/10\text{ seconds} = 100\text{Mbps}$ on average). In this example, the storage controller is well within its specified operational parameters. But from a camera's perspective, something must now buffer nine seconds worth of video or the result will be dropped video frames.

The server running the NVR software can typically buffer some of this video traffic, but as different cameras come into play and different storage topologies are introduced, the NVR server must either be significantly over-resourced (read higher cost) to accommodate the variation in load or have a strict limit as to the number and type of cameras that each server can support. As one of the elements of a well-designed storage solution is to reduce cost, requiring more servers to accommodate a given number of cameras is again in direct opposition to the effort to minimize TCO.

Constant FIFO of Recorded Footage with Sustained Bit-Rate Streaming

Most IT departments have excess capacity in their storage farms to accommodate future growth, which enables the IT manager to estimate a set amount of storage capacity per user or per application and size his needs accordingly. Video storage systems do not have excess capacity because video surveillance cameras never stop streaming data.

Most surveillance systems configure storage arrays based on set parameters such as bit-rate from the camera, recording schedule, number of cameras to record, and retention period needed.



Once the system is commissioned, video is recorded and the hard drives begin to fill. If initial storage calculations were done correctly, the drives will begin to overwrite once the retention period is reached, deleting the oldest video files to make room for new ones.

As anyone who has ever used a Windows PC can attest, as you delete information from your hard drive, fragments are left behind. This is why we are frequently forced to run the Windows Defrag application and wait while our drive gets defragmented. As a typical PC user doesn't regularly delete old files, we may have to run the Defrag utility perhaps once or twice a year. Now imagine a video storage array whose steady state is a constant overwrite scenario. The fragmentation you see build up when you decide to move your annual family vacation pictures from your PC's drive to a DVD in order to free up some room appears every second across millions of files in a video surveillance system. As fragmentation builds, it will eventually choke off the remaining capacity on the drives. Most NVR software is designed in such a way that if no additional storage capacity remains, the NVR will cease to write new files in an effort to protect the data that is already on the drives. Needless to say, ceasing to record video usually defeats the purpose of a video surveillance system.

Unfortunately, fragmentation is not something that can easily be avoided, but its effects can be mitigated. One such way is to leverage a file system that is less prone to fragmentation and can recover faster. XFS is an example of a file system built for performance and equipped with tools to easily recover from fragmentation issues. As video systems are constantly overwriting data, having a robust file system is critical to a properly designed video storage solution.

Sub-Standard Operating Environments

Hard disk drives are the most important and most sensitive component of a surveillance storage system. While the prolific increase in density combined with decreasing price points have been a key driver in the transition to digital video recording, the risk introduced by a hard disk drive failure has increased exponentially. Protecting these drives and the data they contain is paramount to any well-designed storage system.

Hard disk drives fail for multiple reasons. Heat, vibration, the quality of the drive itself, the specified bit-error rate, etc. all are contributing factors.

Temperature

With the amount of wiring involved in video surveillance applications and the fact that space in IT server rooms is difficult to come by, most surveillance equipment is prone to issues with operating temperatures. While there has been a great deal of debate as to what these numbers really mean, studies by Google and Carnegie Mellon University have shown a statistical correlation between hard disk drive failure rates and operating temperatures. Google's study actually confirmed that operating at elevated temperatures (above 40°C) has a direct correlation to premature drive failure. Additionally, an abnormally high failure rate was observed when drives operated at the low end of the temperature range (around 20°C).

The rule of thumb is that drives operating at either specified temperature extreme will have a significantly reduced life expectancy than those operating at 25°C. Here, it is worthwhile to understand what operating temperature means. As the drive is always in a server chassis,

25°C refers not to the room's temperature, but rather to the temperature of the hard disk drive inside the chassis. As the drive itself and other electronics in the chassis dissipate heat during operation, care must be taken by the hardware vendor to ensure that drives inside a given chassis are adequately ventilated to maintain a constant operating temperature. Additionally, care must be taken by the system integrator to ensure adequate ventilation in the rack design to keep drives cool.

Vibration

Vibration is another common threat to the life expectancy of hard disk drives. This is especially true when multiple drives are installed next to each other. Rotational vibration can easily lead to catastrophic failure when the spinning platters and heads inside the drive come in contact with each other. Should vibration become severe enough, the connection between the drive and the RAID controller card can easily become loose, leading to intermittent power and data transfer problems.

Bit-Error Rate

Another common, but often neglected failure point is the bit-error rate specified for the drive. All drives have bad sectors; it is an unavoidable consequence inherent in the manufacturing process. Manufacturers will screen drives to determine the amount of bad sectors on the platters as one of the aspects used in binning the drive during the manufacturing process. Bit-error rates of 10^{14} or 10^{15} are common on most drives today. While this rating was insignificant a few years ago, today's 1TB and 2TB drives make it almost certain that a write error will be encountered before the drive is full. How systems deal with these errors is vital for long-term dependability.


With so many failure scenarios, it should be no surprise that hard disk drives do fail. In fact, given the correlation of hard disk drive failures to operating environment and vibration, it should be expected that if one drive fails, another drive failure in the same chassis cannot be far behind. As video surveillance systems depend on stored data for investigation and generating evidence for prosecution, data loss must be guarded against at all cost.

The NSM5200: Designed for the Rigors of Video Surveillance Storage

The NSM5200 is Pelco's second-generation IP video surveillance storage chassis, designed from the ground up to meet the unique requirements of video surveillance applications. The NSM5200 integrates the functions traditionally spread across three disparate servers – NVR functionality, storage, and storage pool manager – into one.

Hardware Architected for Reliability

To enhance reliability, the NSM5200 chassis has been specifically engineered to protect its most sensitive components while allowing for on-line maintenance and service. As previously stated, hard disk drives are the most critical component of a recorder and yet the component most likely to fail. To protect




the drives, Pelco has over-engineered the cooling system. Tests have proven that the NSM5200's 12 drives, under full load, experience a 7.8°C rise in temperature; according to Seagate's test labs, a temperature spike of 20°C is common in most enterprise RAID arrays. Such a decrease in operating temperature was accomplished by utilizing five fans to pull ambient air from the front and exhaust it through the back of the unit. Obstacles that create air dams inside the chassis have been minimized as have obstacles that can impede airflow out the back (wiring and cabling, for example have been minimized with the use of a cable management scheme designed into the chassis). The resulting airflow ensures that when installed in an environment where the ambient temperature (measured at the front bezel of the unit) is between 10°C and 35°C, the operating temperature of the drive will not approach its specified boundaries. As validated by the studies published by Google and Carnegie Mellon, avoiding either temperature extreme is a critical step in preventing early drive failure.

As vibration is another potential cause of early drive failure, the NSM5200's drive carriers and drive bays have been designed to minimize the harmful effects of rotational vibration. In testing at the same Seagate laboratory, the performance degradation due to rotational vibration was measured to be 1.7%. Seagate's norm for enterprise RAID arrays is 10%. The dampening effects of the drive carriers and chassis work to minimize the harmonics that can lead to catastrophic failures. A custom-designed SAS backplane with positive latching built into the drive carriers also works to ensure a secure mating connection between the drive and the backplane.



Even with the best engineering practices and technologies available, the reality is that as long as we continue using rotational media for storage, there will be failure. The physics of hard disk drives and the known contributors to disk failure suggest that if one drive fails, another cannot be too far behind. For that reason, Pelco decided to forego the traditional RAID5 array and deploy RAID6 as standard. RAID6 utilizes a double parity scheme such that two drives can fail before recorded data is exposed. The reality of 24/7 operations typically means that immediate drive replacement is not always possible. Hence, having double parity protection buys the surveillance tech time to replace a failed drive.

Another critical enhancement is that the new drive does not have to be completely striped into the array before data can be recorded to it. In fact, within a couple of minutes of being replaced, data starts to record to the drive. As data is recorded, that portion of the new drive is striped into the array. This enhancement avoids the nail biting moments as system administrators wait for a replaced drive to get rebuilt into the array, hoping that a second drive will not fail and



corrupt data in the meantime. The ability to swap out a failed drive with a new drive, and have the array automatically rebuild that new drive in the background makes maintenance for failed drives transparent. The RAID6 array continues to provide protection against a second drive failing while the array is being rebuilt on the new drive.

Drives are not the only things that can fail. Fans and power supplies are other common failure points in a chassis. As a result, the NSM5200 has been engineered with a custom rack rail system, mid-plane maintenance bays, and hot-swappable fan modules that allow mid-plane fans to be easily replaced. Should a fan fail, diagnostic alarms are sent to a designated administrator and LED indicators on the front of the chassis indicate a hardware failure while LEDs on each fan indicate the particular failure.

The rack rail system allows the chassis to be pulled half-way out of the rack in order to replace the fan without having to take the unit off-line. The exhaust fan can similarly be replaced from the back of the chassis along with the redundant, hot-swappable power supply modules. Additional diagnostic instrumentation, such as warnings as operating temperatures approach specified thresholds, allow administrators to perform maintenance work prior to failures occurring. The ability to maintain and service the NSM5200 without taking it off-line is a crucial element in providing a recording platform for a 24/7 surveillance operation.


Software Architected for Reliability

A key element of the NSM5200 software is that of storage pooling. Storage pooling provides a significantly more robust strategy for recorder failover while dramatically reducing energy consumption and total cost of ownership.

The concept of pooled or clustered storage is not new. Critical IT infrastructure has utilized and proven the effectiveness of this concept in ensuring that distributed IT systems can respond to peaks in demand as well as keep critical systems running in the event of a server failure. However, the notion of pooled storage is new to the surveillance industry. To safeguard against a catastrophic failure, the surveillance industry has long deployed a hot standby. The hot standby required a mini replica of the storage sub-system – the recorder server, storage array, and in some cases encoders – to be deployed at the edge of the network. This additional node would be powered on, consuming energy and requiring cooling, take up rack space and network ports, but perform no useful work until a failure occurred in one of the main recorders. The concept behind it was that should a recorder fail, some intelligent entity in the system (analog matrix in the old days or the system manager in today's IP systems) would route the cameras from the failed unit to the standby unit for recording.

Aside from today's reality that energy conservation is an intricate part of corporate citizenship and financial strategy, the concept that the failover system would be available and error-free when it was called into action was fraught with risk.

Endura 2.0



The NSM5200 uses distributed load balancing and active failover to ensure that in the event a recorder is taken off-line, all cameras continue to be recorded. The concept first exposes the capacity of all pooled nodes to the designated pool controller. The cameras are presented to the controller as well. The controller then assigns cameras to each node. Once recording has started, the controller monitors the storage consumption for each node in the pool and automatically rotates cameras to ensure that storage consumption is kept at an equilibrium. This means that each recorder in the pool is doing an equal amount of work at all times. Therefore, no wasted energy, no wasted cooling requirements, and no wasted network ports or rack space. As each recorder is also on-line and recording cameras, the system also knows of any potential problems with all the recorders in the pool. Should a unit fail or require being taken off-line, the pool manager simply routes the cameras that were being recorded to the off-line unit and distributes them evenly across the remaining members in the pool. This is the concept of active failover.

Additional software elements built into the NSM5200 specifically for video surveillance include automatic database repair utilities and defrag utilities that run every time upon boot up. Diagnostics that monitor packet loss, video loss, etc. are critical for surveillance operations. These diagnostics are exposed to the Endura system as well as to SNMP monitoring systems via custom MIBs. Finally, the NVR software itself has been designed to accommodate the variability of bit-rates, file sizes, etc. that are now becoming commonplace as different types of cameras are utilized throughout a facility.


Consolidation Yields Cost Benefits

The NSM5200 is designed to replace up to three disparate servers common in most surveillance applications. The purpose-built server and software combines the functionality of the NVR server, storage array, and storage pool manager into a single box. Additionally, the throughput of the NSM5200 has been optimized to support significantly more write operations than read operations.

The server consolidation yields significant savings in operational expenditures as a single server draws far less power and generates far less heat than the three servers it replaces. In turn, the reduction in heat dissipation further reduces HVAC requirements for both the rack and server room. To get a sense of the savings, let's take the following example:

- NVR server – Dell PowerEdge 860 - 238W of power and 812BTU/h
- Storage array – Dell PowerVault MD1000 - 430W of power and 1,468BTU/h
- Storage pool manager – Dell PowerEdge 860 - 238W of power and 812BTU/h
- 24/7 operation for 365 days = 8,760 hours of operation
- Global average of \$0.32/KW for energy cost
- 64 cameras, each producing 2Mbps of data for storage
- 7 days of retention is required

Assuming 64 cameras per server and configuring the MD1000 with 15 1TB SATA drives, the multi-server approach would cost \$3,833 per year in energy consumption. The NSM5200, which integrates the functions of all three components, consumes just 262W



for the entire chassis and requires only 895BTU/h of cooling. Replacing the Dell servers and storage with a single NSM5200 would save \$2,361 – or 61% - per year in energy consumption, and the savings obviously mount as the number of nodes installed increases.

In addition to the savings realized from consolidating servers, the NSM5200 has been tuned to provide optimum throughput for recording streams. Supporting 250Mbps of guaranteed write throughput, each NSM can support over 100 2Mbps streams. This means fewer servers to handle a given quantity of cameras. In fact, when compared to the NVR5100, Pelco's first generation storage solution, the NSM5200 delivers twice the performance at half the power dissipation.

Summary

Storing video surveillance footage is a difficult task. The unique characteristics of streaming video and the characteristics of security installations require careful planning to achieve the desired results. The NSM5200 has been architected from the ground up to meet these challenges while drastically lowering the total cost of ownership.



by Schneider Electric

The recognized worldwide leader in video and security systems, Pelco boasts the most comprehensive array of products, services and expertise available in today's marketplace. And now as a member of the Schneider Electric family, Pelco brings a network of assets backed by the strength of a Fortune 500 company to help you define and achieve your business objectives.

www.pelco.com