Moving to an IP-Based Video Surveillance System

Table of Contents

Roger Hall
CSSG

Version 3.1
06/02/17

Moving to an IP-Based Video Surveillance System

About Seagate Technology

Seagate Technology is the leading expert in scalable storage solutions, developing robust products that enable people and businesses around the world to create, share and preserve their most critical memories and business data. Seagate has been at the cutting edge of storage solutions for over 38 years, starting with the development of the 5-megabyte ST506 5-¼" hard disk drive released in 1980 to a 60-terabyte solid-state drive (SSD) announced in 2016. Seagate is also renowned for its cutting-edge SAN solutions intended to solve today's most challenging storage problems where performance, availability, cost, and ease of use are paramount.

Seagate creates space for the human experience by innovating how data is stored, shared, and used. Learn more at http://www.seagate.com.

Executive Summary

This report focuses on the reference architecture behind an IP-based video surveillance system that meets and exceeds today's demanding video surveillance challenges in a security-conscious world. Whereas traditional analog systems struggle to manage the sheer volume of ultra-high definition video, demanding forensic and analytic applications, and long-term archival requirements, IP-based video surveillance easily surpasses these objectives and offers a scalable solution. Yet the investment in analog infrastructure can be preserved with a hybrid solution of both an analog and IP-based video surveillance architecture with robust, performance-driven storage. In selecting a video surveillance storage architecture, it is paramount to choose one that encompasses:

1. Performance: A robust video surveillance system with medium to large storage requirements of 64 cameras or more per server is dependent on the ability of storage systems to support iSCSI at 10Mbps or 8 to 16Mbps fibre channel SAN fabrics. And storage systems need to be able to accept video traffic without incurring frame loss, including high definition video at 29.97 frames per second.
2. Availability and Reliability: The storage configuration must be highly available with limited to no failure downtime.
3. Scalability: As the number of cameras increase, the storage architecture should also scale to ensure optimal performance and provide sufficient archival capabilities.
4. Interconnectivity: 8 to 16Gbps Fibre Channel (FC) or 10GigE or higher iSCSI speeds along with 12Gbps SAS to expansion disk arrays (EBODS) interconnectivity
5. Robustness: Even with the demands of high-definition video in conjunction with standard compression, the storage system must deliver a high-performance solution with low storage processor utilization.
6. Total Cost of Ownership: The storage system Must deliver an advantageous price for performance and reduce cost of operation over its lifetime.
7. Flexibility: An optimal storage system that takes advantage of technology changes to deliver value without requiring forklift changes.

To understand why an IP-based or hybrid solution requires dependable performance-driven storage, this paper examines the infrastructure behind modern video surveillance systems.

Implications of Analog to Digital Video Surveillance

With the rapid emergence of IP-based video surveillance, there is now a convergence of relationships between physical security and information technology professionals. Today's video surveillance systems resemble more of an IT-based environment than a physical security environment where the physical security department might be concerned that IT is trying to take over their domain.

Resistance by physical security specialists is not totally without merit. The limitations in analog systems are well managed, the technology is easy to learn and use, retention does not require expensive hardware in response to elevated expectations, and frame loss due to bandwidth limitations does not exist (because of the multi-homed nature of the analog video surveillance network, where each camera is connected to a DVR via coaxial cable).

Nevertheless, IP-based video surveillance capabilities far surpass those in traditional analog-based systems, especially relating to the demand for long-term retention of ultra-high definition (UHD) 3840 x 2160 pixel (4K UHD) video, the emergence of video analytics as an essential tool for security specialists. And, unlike many other technological transformations, the upgrade to IP-based video surveillance can occur gradually, using existing, sunk-cost infrastructure, including the use of established, expensive network infrastructure such as coaxial cable, and analog cameras.

This demand for performance, usefulness of data, and retention is driving the evolution of surveillance technology. There are four primary aspects of today's video surveillance systems pushing the envelope of what is possible:

1. IP-based camera design up to 4K UHD, and with localized storage and analytics
2. Video management system (VMS) software development, including increasingly sophisticated analytics and forensic analysis
3. Storage as it relates to performance, scalability, and reliability
4. Melding intrusion detection and other sensor capabilities with advances in video surveillance and analytics.

The proliferation of video surveillance in several key applications by industry and government are further driving demand, including:

1. Physical security – access control, Occupational Safety and Health Administration (OSHA) conformance, Health Insurance Portability and Accountability (HIPAA) conformance, intrusion detection, event capture, and asset protection
2. Analytics – shopper count, traffic flow, transaction analysis
3. Marketing - business intelligence including shopper preferences and competitive analysis
4. Manufacturing and process/quality control
5. Risk management – Fraud and "slip and fall" protection; insurance fraud-related incidents cost companies an estimated $30 billion a year in the U.S. alone

Expanded physical security requirements, along with new business opportunities and increased efficiencies, have prompted the need for higher retention rates, higher camera resolution, and faster forensic analysis methodologies, all of which drive storage requirements.

Camera Evolution: Streaming Video to Video as Data

Analog cameras capture an analog video signal and transfer that signal via coaxial cable to a digital video recorder (DVR). The DVR then converts the analog signal to digital, compresses it, and stores it on one or more hard drives. In contrast, digital cameras have processors that not only can conduct video analytics, but are capable of converting and compressing analog video content before transferring it over a local area network (LAN) to a network video recorder (NVR), or a server running VMS software. Due to advances in camera technology, digital cameras can now support resolutions

much higher than analog cameras. Analog cameras cannot support screen resolutions higher than 0.4 megapixels. In contrast, digital cameras are capable of supporting up to 8 megapixels and can capture a much clearer image when objects are in motion.

In fact, the primary driver for video surveillance systems today is the volume of surveillance content being created and processed by cameras, especially high-performance digital cameras. Transferring continuous, activity-based content through a video management system and onto a storage platform creates changing requirements in storage sizing which much take into consideration:

1. Number and type of cameras
2. Camera resolution in pixel count
3. Recording frame rates
4. Transmission of video surveillance footage based on activity and events
5. Compression methods
6. For some compression methods, scene complexity and movement between frames

Storage and bandwidth requirements are directly impacted by resolution, the number of cameras, continuous streaming, high frame rates, and scene complexity. These factors result in a higher aggregate amount of data being processed and stored by a given surveillance system.

While storage and bandwidth requirements continue to grow, advances in compression standards have tempered storage and bandwidth requirements while improving image quality. Moving Picture Experts Group (MPEG) version 4 part 10 Advanced Video Coding (AVC), also known as H.264, and Motion Joint Photographic Experts Group (MJPEG) are the two primary compression standards used today in video surveillance, though H.265 is destined to replace H.264 as the standard in the not too distant future.

MJPEG streams live video frame-by-frame compression with frame rates correlating directly to the number of JPEG images being streamed. No inter-frame compression occurs. In contrast, H.264 accomplishes lower data volumes via inter-frame compression methods that take into account frame-to-frame motion by basing compression on variance to reference frames – difference data that can be reconstructed at its destination. The result is high compression with lower bandwidth and storage requirements compared to MJPEG, but additionally with lower fidelity than MJPEG.

Many video surveillance cameras today can support both H.264 and MJPEG, with H.264 streaming at rest state, converting to MJPEG streaming under event/alarm conditions. Nevertheless, H.264 is by far the dominant compression standard used today, not only in video surveillance, but also in video content delivery to the desktop and to handheld devices, such as smartphones.

Many VS cameras today can support both H.264 and MJPEG, with H.264 streaming at rest state, converting to MJPEG streaming under event/alarm conditions. Nevertheless, H.264 is by far the dominant compression standard used today, not only in video surveillance, but also in video content delivery to the desktop and to handheld devices, such as smartphones.

Advancements in compression methods continue with the development of the next generation of MPEG compression. High efficiency video coding (HEVC), known as H.265 or MPEG-H Part 2 offers double the data compression ratio while maintaining the same level of quality, or substantially improved video quality at the same bit rate. Cameras are already available that support this standard.

However, H.265 adoption has been slow due to the payment of royalties to licensors of HEVC patents. This is why some of the largest tech companies, such as Google, Amazon, Cisco, and NVIDA have joined the Alliance for Open Media with an intent to release a royalty-free alternative format.

Nevertheless, various H.265 tests have showed a dramatic bitrate reduction of as much as 66% with comparable, subjective video quality when compared to H.264. This bitrate reduction could make consistent UHD camera use in video surveillance a reality.

Video Surveillance Network Design Considerations

Perhaps the most contentious issue with IP-based video surveillance is the amount of bandwidth multiple video streams might require in an otherwise standard, IT-centric local area network (LAN). Bandwidth usage and security is one of the first collision points between traditional video surveillance and IT groups. The natural tendency of physical security experts might be to design and install a completely separate video surveillance network to ensure bandwidth requirements are met where there is no chance of intra and inter-network intervention.

IT specialists in particular are concerned about video surveillance traffic, especially when multiple camera streams traverse a LAN to access network attached storage (NAS) devices. This is a primary reason why SAN storage is now preferred over NAS.

For example, one storage manufacturer's IT department would not allow permanent access to their managed network for two video surveillance cameras in that company's video surveillance lab because of traffic congestion fears. One of the two cameras were later installed on a dedicated LAN within the lab itself and used to simulate over five hundred 8 Mbps video streams on that network where a SAN storage array was used.

The optimal IP-based design approach is a closed, dedicated network where only VS and security-based sensor traffic is allowed. However, network engineers can design networks that incorporate both video surveillance and IT requirements while using much of the same infrastructure. To facilitate this, consider the following design considerations:

1. Implement multicast instead of unicast and broadcast.
2. Determine the volume of video surveillance required based on camera footage, and the extent of live video display requirements, including forensic analysis
3. Make use of compression and frame rate methods that take advantage of video analytics at the camera; increase frame rates and/or initiate streaming directly with alert/alarm events
4. Implement Quality of Service (QoS) reminiscent of asynchronous transfer mode (ATM) and now available through multi-protocol label switching (MPLS)

If simultaneously streaming video to multiple locations is a requirement, multicast transmission is the best choice because it transmits a single video stream from a camera to specific locations, without sending multiple streams to each location (unicast), or broad casting streams to every recipient regardless of need.

To clarify, a unicast transmission/stream sends IP packets to a single recipient. A multicast transmission sends packets to a group of recipients requesting those packets. A broadcast transmission sends packets to every recipient on a network with no exceptions. Therefore, multicast limits traffic only to those recipients requesting service.

Video surveillance traffic bandwidth requirements are heavily dependent on compression technologies, frame rates, the number of cameras used, and how networks are designed around them. Bandwidth requirements are also dependent on video display requirements, such as live video wall views in multiple locations.

Another effective method for managing bandwidth requirements is to utilize smart cameras that only send heightened transmissions when an event such as detected motion occurs. In other words, heightened transmission of video surveillance traffic only occurs when there is a reason to transmit.

ATM has been the most effective QoS protocol ever since its inception in the late 1990s, and its continued use in telecommunications is testament to its robustness and reliability. However, a lack of major scalability compared to Ethernet, along with esoteric complexity in an increasingly IP-centric world, has doomed ATM to a footnote in the annals of networking history. MPLS, an IP-based protocol intended to replace ATM, is capable of providing some QoS, but not at the level of ATM. MPLS's saving grace has been the massive increase of available bandwidth over time; a 10Gbps network is commonplace and a 100Gbps network is now possible.

When designing an IP-based video surveillance system, it is critical to select a high-throughput and scalable storage solution capable of supporting high reliability metrics, and maintaining a performance level that minimizes or eliminates frame loss and latency concerns. An optimal storage solution should also be able to support multiple transport protocols, including iSCSI, Fibre Channel and SAS. Response to redundant array of independent disks (RAID) recovery after a disk failure is also a consideration.

SEAGATE

Storage Design Challenges for Video Surveillance

Video surveillance has significantly progressed from the days of analog magnetic tape used in conjunction with various types of magnetic tape recorders that support low-resolution, low-frame-rate, black-and-white analog camera streams; an expensive system beset with high capital investments and maintenance costs. With advances in storage technologies, coupled with improvements in both analog and digital camera design, the process became simpler, which then increased the demand for additional storage over time.

MarketsAndMarkets recently projected that the video surveillance storage market would be $18.28 billion US by 2020 with a compound annual growth rate of 22.41% during the forecast period.[1] This 2015 estimate is considered by some to be conservative.

To meet this growing demand, planning efforts must support the following design elements:

1. Performance
2. Availability and Reliability
3. Scalability
4. Interconnectivity
5. Robustness
6. Total Cost of Ownership
7. Flexibility

---

1. Performance

Storage I/O performance is a major challenge because multiple, live video stream writes continue to grow both in the number and quality of streams.

From a storage performance perspective, video surveillance system test results indicate a relationship between IOPs, write and read latency, and frame loss in conjunction with storage processor utilization and bandwidth. The more recording servers (and cameras per server) a SAN storage array supports, the more the ensuing traffic influences storage performance where these parameters come into play. A high availability, well-designed SAN storage array supporting a 8 Gbps Fibre Channel (FC) network fabric has been extrapolated to support as many as 1,000 high definition cameras with storage processor utilization of less than 50%, and with little to no frame loss.

In fact, frame loss of 1% or less is a gating factor limiting the amount of live video streams that can be introduced to a storage array. Frame loss vs bandwidth is the challenge in designing a video surveillance storage installation in an environment that also requires high availability and reliability.

2. Availability and reliability

The standard for professional network devices is the same as quality network devices requiring 99.999% high availability (HA). For storage, this means redundant and hot swappable controllers, fans, and power supplies along with hot swappable spinning hard drives or SSDs. Power supplies in particular should be 80Plus Gold Certified, capable of handling 110 to 240 Volts AC and 50 to 60 Hz.

3. Scalability

Because of enhancements in camera design supporting high definition and ultra-high definition video in multiple streams along with analytics at the edge, storage architecture scalability has become a major challenge especially for existing installations with NSRs where scalability is minimal to non-existent.

SAN and NAS storage arrays coupled with expansion chassis provides the scalability needed for modern mid to large scale video surveillance architectures. Further, storage arrays with tiering capability allows for cost savings using high speed SSDs and SAS HDDs in conjunction with lower cost SATA storage.

1. MarketsAndMarkets.com Video Surveillance Storage Market worth 18.28 Billion USD by 2020
http://www.marketsandmarkets.com/PressReleases/video-surveillance-storage.asp

4. Interconnectivity

In order to meet high bandwidth requirements inherent in today's standard high definition and higher resolution video surveillance cameras, storage arrays must support 8 to 16Gbps Fibre Channel (FC) or 10GigE or higher IP or iSCSI speeds. If expansion disk arrays are to be used, at the very least 6Gbps, preferably 12Gbps SAS should be the standard.

5. Robustness

Even with the demands of high-definition video in conjunction with standard compression, the storage system must deliver a high-performance solution with low storage processor utilization. The system needs to be open, robust, and cost-effective. Where governments are involved, effectiveness and reliability of video surveillance storage is likely to trump cost considerations, though it is likely cost will continue to be a strong influence.

6. Total Cost of Ownership

Storage systems must be competitively priced and able to integrate with existing video surveillance systems as needed. This includes the ability to combine disparate system components (such as analog and digital cameras, existing networks, and separable storage platforms) into a complete system, allowing for a roadmap of continual system design improvements and for a smooth transition to modern scalable, robust storage arrays that allow for cost savings using high speed SSDs and SAS HDDs in conjunction with lower cost SATA storage in a hybrid solution.

7. Flexibility

Storage systems need to be able to take advantage of new technologies without requiring wholesale, expensive hardware and software changes. Further, robust, flexible storage systems require an operating system that supports multiple hybrid storage technologies and RAID configurations. This includes easy storage tiering configuration changes based on component read/write performance as well as cost.

Capturing surveillance video in digital form provides the ability to conduct after-the-event (forensic) analysis. Once video surveillance content and its associated metadata is captured, the video surveillance system and its storage capabilities enable real-time and post processing, video analytics, and advanced search algorithms, which may include database integration of both structured (metadata) and unstructured (video) content.

This activity takes place while video surveillance cameras and video management systems continue to stream and write video content to the same storage array.

SAN versus NAS Storage in a Video Surveillance Architecture

SAN is the standard storage technology used today in video surveillance when requirements extend beyond the capability of NVRs and DVRs. It is a dedicated switched network or fabric, as it is sometimes called, that could include RAID arrays, tape libraries, and storage expansion arrays available to one or more servers.

A SAN supports block access to data much like direct attached storage (DAS) does. In other words, like an internal drive, a SAN can present storage to a server by volume or drive letter using session layer protocols, such as Fibre Channel or transport layer protocols, such as Gigabit Ethernet and iSCSI.

In contrast, network attached storage (NAS) is a method of remote file access to one or more servers using application layer protocols, including common internet file system (CIFS) and network file system (NFS). In other words, rather than a drive letter, a NAS device presents itself as a device; essentially a server with its own file system, operating system, and IP address. Consequently, the management overhead inherent in NAS devices, along with the requirement that traffic go through full ISO protocol stack handling of an Ethernet connection, results in slower read/write response times when compared to SAN devices and with the deployment of iSCSI where SAN network hardware costs are identical to NAS, SAN is a clear choice.

NAS is not without its advantages, however. One example is the use of NAS to allow for server failover, used by at least one VMS provider. It is possible SAN could not be used in the same manner, if the VMS were designed to support it.

RAID Arrays are Not All Created Equal

Most Redundant Array of Independent Disk (RAID) storage arrays now support RAID 0,1,3,5,6, 10, and 50. Each RAID type relates to ability, or inability, to recover in case of drive failure. It also relates to storage availability and performance.

For example, RAID 0 provides data read and write transfer rates higher than individual transfer rates (estimated rate increases are based on the number of disks where n disks result in disk rates up to n times higher than individual disk rates), but with no data redundancy at all and no way to recover if one disk fails. In contrast, RAID 1 disk performance is unchanged, since the content of each disk is mirrored by a second drive associated with it. The advantage is optimal reliability, but with a 50% loss of storage availability to the end-user.

In order to maximize availability and reliability while keeping an eye on cost, the predominant RAID type used today in video surveillance is level 6, also known as double-parity RAID where two parity stripes are used on each disk, which allows for 2 disk failures rather than one disk failure allowed in a RAID 5 configuration.

Besides I/O performance, reliability, scalability, and flexibility, time to recover from disk failure in a RAID 6 configuration is dependent on disk size, on the number of disks, and the design of the RAID controller itself. Further, not all RAID controllers are created equal. RAID controller design consisting of dedicated ASICs have proven to provide superior performance at lower cost compared to today's typical processor-based RAID controller design. This higher performance manifests in high IOPs, write latency, optimal storage processor utilization, drive loss recovery time, and low to no video frame loss.

An alternative to hardware RAID is software-based RAID. While a hardware-based system manages RAID independently from a host, software RAID relies on the host's CPU to transform disk sets introduced to the host, to RAID sets configured and managed BY the host. Performance is entirely dependent on software design reliability and on the host's CPU the software relies on. While software RAID may optimal for cost considerations in a DAS configuration, software RAID in a SAN involving multiple hosts will not meet the performance requirements today in a video surveillance solution.

Conclusion

The video surveillance industry is experiencing unprecedented growth driven by technological innovations, by world events in an increasingly risky and dynamic world, and by changes in the way stakeholders approach security to examine possibilities beyond traditional practices.

The culmination of advancements in camera design and performance has stretched the imagination to the point where 4K cameras in a video surveillance world is not only possible, but becoming a not-too-distant future expectation. Monumental improvements in video compression have made it possible to stream complex, high-frame-rate, high-resolution video content without taxing available bandwidth. Strides in network protocols focused on increasing the quality of service have helped to develop a comfortable reliance on best-effort network protocol design.

Storage itself, ranging from disk and solid-state drive design to RAID controller design in 5 9's (99.999% reliability) high-availability systems, has repeatedly broken through not one but many glass expectation ceilings.

In 1985, 10-megabyte disk drives were the standard. Many believed this amount of drive space was all the world would ever need. Storing video content on a computer system was not a consideration because it was patently impossible. Before the inroads into new compression standards, it was estimated a two-hour movie would consume roughly 80 terabytes of disk space, and the storage would fill a warehouse, which was incomprehensible at the time.

Today a two-hour, H.264-compressed movie at 25 frames per second and 720P HD resolution can require as little as 480 megabytes of disk space. While this is well beyond the storage capabilities of the first 10-megabyte 5-¼" disk drive, this is a mere trifle for hard disk drives of 10 terabytes or more.

These astonishing innovations in disk drive design, where 16 terabytes on a single 3.5" disk drive is now viable, have made continuous recording of extended-retention, high-definition video surveillance traffic in a storage array not only a possibility, but now an expectation, and one that will engender growth for the foreseeable future.

Yet video surveillance systems today can encompass existing investments in network infrastructure, such as coaxial cable and analog cameras. Whether a hybrid solution of both analog and digital infrastructure or a purely digital architecture, today's video surveillance industry demands robust storage with high performance to keep frame rate loss to a minimum while maximizing camera count per system. To manage the sheer volume of multiple, high-rate data stream throughput with little to no frame loss, the storage array must deliver a scalable, highly available, and reliable solution. Video surveillance systems risk impeding throughput with sub-par storage that acts as a bottleneck, and encumbers expedient results for forensic and analytic applications. By properly designing a video surveillance system with performance-driven storage, today's and future needs can be confidently met.

seagate.com

AMERICAS: Seagate Technology LLC, 10200 South De Anza Boulevard, Cupertino, California 95014, United States, 408-658-1000
ASIA/PACIFIC: Seagate Singapore International Headquarters Pte. Ltd. 7000 Ang Mo Kio Avenue 5, Singapore 569877, 65-6485-3888
EUROPE, MIDDLE EAST AND AFRICA: Seagate Technology SAS, 16 – 18, rue du Dôme, 92100 Boulogne-Billancourt, France, 33 1-4186 10 00