Storage Technologies for Video Surveillance

AVIGILON
THE BEST EVIDENCE™

The surveillance industry continues to transition from analog to digital. This transition is taking place on two fronts — how the images are captured and how they are stored. The way surveillance images are stored has also changed from analog-based recording solutions — like VHS tapes and film — to digital storage on hard drives. As a result, most surveillance systems in place today make use of hard disks to store surveillance images digitally. With hard drive capacity doubling approximately every two years and the cost per unit shrinking by 50% surveillance storage is less expensive, more efficient and progressively powerful.

However, the part of the analog to digital surveillance transition — how images are captured — is placing unique demands on surveillance storage technology. These demands require integrators and customers to have a deeper understanding of the storage technologies available and the impact that each storage option has on their surveillance system.

The unique demands that result from the transition from analog to digital image capture are largely driven by the fact that digital images can exceed the resolution limits of analog cameras. Higher resolution and higher definition images mean more detail and better evidence in surveillance video but also mean each camera has the potential to generate larger images that require appropriate technology solutions to efficiently store, search, playback and manage.

This white paper explores available surveillance video storage technologies and the impact of each solution on video management software performance to arm security and IT professionals with the information to make the best storage decisions when designing a surveillance system.

Storage Technologies

Disk Drives

The hard disk drive (HDD) is the central component to the digital storage of surveillance video. HDD technology has evolved incrementally from the same basic concept pioneered in 1956 — a stack of spinning platters and scanning heads. A conventional mechanical hard disk operates using a series of spinning rigid plates that store information using magnetism. The data is read from the spinning platters using scanning heads positioned over the spinning disks. Though the fundamentals of HDD technology haven't changed, the ability to increase the density in which data is stored has allowed HDD size to shrink tremendously. This change resulted in a dramatic capacity increase while costs continued to decrease. For example, the typical price per gigabyte of raw HDD storage has decreased from $56.30 per gigabyte in 1998 to $0.053 per gigabyte in 2011.

Innovations in speed and interfaces have also allowed HDDs to handle larger amounts of incoming data and quickly provide that data back to applications enabling advances like HD and megapixel surveillance.

One shortcoming of HDD technology is its reliance on mechanical moving parts. Moving parts always come with some risk of failure over time. Because of this, most surveillance systems use multiple disks in redundant configurations (see the next section on RAID for more detail). Solid state drives (SSDs) are another option. SSDs rely on a large amount of solid-state memory, essentially microchips, and have no moving parts to store information. Along with the reliability and resiliency improvements that are gained with the elimination of moving parts, SSDs also offer substantial increases in speed for both reading and writing data over mechanical HDDs. The tradeoff for improved speed and reliability is increased cost. In 2011, the cost per gigabyte of SSD technology was $1.50. Though declining rapidly — at 30 times more than current HDD prices — the cost of SSDs are too expensive for the large storage capacities needed for most surveillance systems.

Redundancy – RAID

The fundamental technology of digital surveillance storage — the hard drive — is only the beginning of the hardware involved in a storage system. The most important technology that will impact any storage system involving multiple hard drives is called RAID — redundant array of independent disks. Hard drives have the possibility of failure — especially when using moving components — so redundancy is a critical consideration to ensure data won't be lost and systems will continue to operate in the event of a drive failure. Beyond redundancy and fault tolerance, implementing RAID also increases the performance of a storage system by increasing the throughput beyond what could be accomplished with a single hard drive.

RAID provides redundancy by grouping multiple hard drives and employing different techniques, mirroring and striping, to distribute data across multiple hard drives. Mirroring is the simplest form of redundancy and simply copies the data on two drives — as any data is written, it's created in both places. With striping, data can be spread out at a low level across multiple hard drives. Striping improves performance of a RAID system by writing data to a large number of drives simultaneously so that the duration of a single write operation is greatly reduced. Striping with parity extends the process of breaking up a file and storing it across multiple drives to include a calculated value — parity information — in addition to the original data that rebuilds data if some portion of it is later lost.

To help understand these concepts better its best to review how they're implemented in real surveillance systems and the different RAID levels used. The most common RAID levels in a surveillance system are 0, 1, 5 and 6. RAID 0 is the simplest form and provides no redundancy

or fault tolerance. Using striping without parity or mirroring, RAID 0 allows a group of hard drives to act as a single usable storage array with increased performance of distributing load across all hard drives. RAID 1 mirrors all data to two drives and uses no striping or parity, although multiple pairs of mirrored hard drives can be striped together in a RAID 10 or a RAID 1+0 array. This is the most fault tolerant RAID setup since all data is mirrored fully, however since it requires twice as many hard drives to provide the needed capacity it's typically used only for operating systems or programs, and rarely as the RAID level for large amounts of storage like surveillance video. RAID levels 5 and 6 are more commonly used for large amounts of storage for surveillance video. RAID 5 uses striping with distributed parity, this means that data is both spread out over multiple drives and additional information is stored to allow data to be rebuilt after failure. For example if you had eight hard drives in a RAID 5 setup each file would be spread out over seven drives and parity information about the file written to the eighth drive. Since RAID 5 distributes the parity information, the parity information won't always be written to the same drive with every file. In the event any one of the eight drives fails, no data is lost and the system can rebuild the contents of the failed drive once it's replaced by using the data on the other seven drives. RAID 5 provides fault tolerance to a single drive failure, increased performance of spreading disk operations over multiple hard drives and only costs an additional hard drive in each array of disks. RAID 6 extends RAID 5 by distributing parity information for each file to two drives; this allows the system to tolerate the failure of two drives in an array without losing data but costs an additional drive per array. RAID 6 becomes more important as array size grows and to ensure storage isn't vulnerable in the event of a rebuild process. Rebuild time increases with array size, so RAID 6 ensures data is not vulnerable during the rebuild process.

RAID is implemented either at the hardware or software level within a server or storage enclosure. For the performance needs of surveillance and truly enterprise class storage a hardware RAID controller is a requirement for any surveillance video storage. Implementing RAID at the software level relies on system resources and introduces additional overhead and delays that impact performance of your storage and make it unsuitable for video surveillance storage.

Storage Communication

We've already discussed the basic technologies commonly used for surveillance storage — a collection of hard disk drives most likely grouped in a RAID 5 or 6 configuration. The next important topic is how the surveillance system communicates with that storage. Generally this communication will fall into one of three categories — directly attached storage (DAS), storage area network (SAN), or network attached storage (NAS).

DAS is the simplest and most economical for storage that needs to be accessed by only a single server. As the name implies DAS provides storage directly attached to a server. The methods of attachment vary but for enterprise class storage it's done via a serial attached SCSI (SAS) cable from a dedicated hardware RAID card.

SAN will typically be used for very large amounts of storage that require access from multiple servers while centralizing management, resources, and scalability. A wide variety of technologies are available for communication in SANs, however modern SANs used for surveillance most often use either Fiber Channel or iSCSI. The choice between the two protocols is based on the needs of the network in terms on inter-connectivity and performance. iSCSI is more cost effective because it uses standard network cabling and switches when compared to the more expensive and complex cabling and switching that Fiber Channel requires.

NAS is the final category for connecting storage to a surveillance system. An important thing that differentiates the storage communication technologies is whether they provide block level or file level access to the storage device from the connected servers and applications. The difference between block and file level access translates primarily to performance: block level access provides higher performance by allowing lower level access, while file level access limits performance but can provide easier concurrent access to multiple users in non-surveillance applications like file sharing. NAS provides only file level access to a storage volume while DAS and SAN provide block level access. In surveillance, block level access is required by most systems, including Avigilon Control Center (ACC). As a result, file level access and the use of network attached storage is typically limited to applications with very few cameras.

Storage Architecture

The final layer between a surveillance application and storage is the software technology and architecture used to structure and index stored video, configuration data, and events. The simplest form of software used by any surveillance application is a file system. A fundamental operating system feature, the file system provides the basic functionality to read, write, and organize files. A surveillance system could make use of a simple file system for everything from configuration files to the video itself — typically however performance needs require the use of a relational database in addition to a basic file system to provide the necessary performance of writing, indexing, and managing the complex event and video related information that's part of a surveillance system. Relational databases play a key role storing and indexing data for a video surveillance systems. They also allow features like search, playback, backups, and export to happen efficiently. In the following section we'll explore how video surveillance software stores video in more detail.

VMS Storage and Impact on Surveillance Design

A typical surveillance system generates four types of data — system-related configuration data, events, metadata and surveillance video — that need to be stored and pose unique challenges. Video management software (VMS) simplifies the process by storing data between files and databases based on its nature.

The smallest and simplest type of data is system-related configuration data. This can be information on recording schedules, users and groups, rules, alarms, and system logs. Usually it's a very small amount of information, accessed infrequently and does not present high demands on speed of retrieval. For these reasons, a VMS will typically store this data within a structured database. Configuration data does pose some unique demands because it has to be encrypted and only modified when proper security checks have been performed. A VMS will also use a database to provide these features.

Another data type generated by a VMS for storage is events. Events will typically be small in size but can be high in volume and place significant demands on a system to store, index and search. The type of events generated and stored will range from user audit events, motion events, general system events, device input and output events, or external events from an integrated system. The high volume needs of managing event storage are met well by a database. Events also require rapid indexing and searchability — functionality that is also best supported by a database.

The most important and significant type of data generated by a VMS is the surveillance video itself. The amount of video generated by a single server will vary based on number of cameras, resolution, and amount of motion, but can be several terabytes per day. Typically the data generated by video is much too large for storage within a database so it is usually stored separately in a file system and indexed in a database. The video itself poses the largest challenge to a storage system. As a result, how a VMS stores video significantly impacts the performance and features of a surveillance system.

Disk fragmentation is a key concern. As outlined earlier, storage is performed on hard disks comprised of rapidly rotating disks. The physical position of data on disks impacts performance tremendously. For example, if related data is all physically located in one place the speed to read access the data is greatly reduced. Over time — as different amounts of data are written, deleted and overwritten on a disk — fragmentation can occur and related data is no longer located on the disk close to other related data. When fragmentation is not managed properly in a video surveillance environment, data from a camera can become widely spread out over a disk. Searching and playback of that data can become very slow.

In addition, if the free space or regions being written are spread far apart, writing video may also take longer which can limit system capacity. Fragmentation is ignored by most conventional computer applications and systems rely either on user-initiated applications or periodic tasks to defragment a disk. Unfortunately — in the case of surveillance systems — the storage system is required 24x7 to write high volumes of data generated by video. Therefore, the downtime needed to defragment a disk after fragmentation has occurred is unacceptable. For this reason, it's important that the VMS be aware of and manage fragmentation as it writes and overwrites video on a disk.

Tightly related to the video and generating at an equal rate but in smaller size is metadata. Metadata is any additional information that describes or adds to the content of the video. For motion detection and analytics applications this will include the location of motion within video as well as object location and direction. Because this data must be stored in an indexed fashion that's rapidly searchable, it must be stored in a database. Two important demands must be carefully accounted for by the VMS when storing metadata: speed of searching and time synchronization to the video. Metadata that's inaccurately sequenced relative to video becomes useless. As well, if the data can't be easily and rapidly searched — it will not enable the effective handling of investigations.

How Does Avigilon Store Video?

Now that we've covered the technologies, implementations, and challenges related to surveillance video storage, it's important to understand how Avigilon Control Center (ACC) has been architected to minimize storage requirements and maximize performance despite the higher storage demands of high-definition surveillance.

HD video places unique demands on a storage system as the video generated is higher in resolution, potentially higher in data rate, and contains more information than conventional video leading to users doing more searching and reviewing as the value of their stored video increases.

As outlined previously, the database plays a key role in a surveillance system. For this reason Avigilon Control Center makes use of a high performance embedded database that is tailored for HD video surveillance storage. The use of a database embedded into the VMS is unique to Avigilon and delivers considerable advantages. Other VMS will typically use an external database, such as Microsoft SQL Server, MySQL, or Postgres SQL, which has several disadvantages. First, performance becomes limited — external databases require data be passed between processes from the main VMS process to the process running the database. Embedding the database means data and queries are quickly handled and database resources can be more actively managed by the VMS process.

Second, external databases are subject to separate management of licensing, upgrades, and in large corporate environments may require involvement of other departments tasked with managing all external databases on site. An external database can easily be patched incorrectly, corrupted by another application, or accessed insecurely than one embedded and managed by the VMS.

The type of database used and how it's managed also impacts the speed of playback and searching. The high performance embedded database used in ACC is tuned to ensure the industry's fastest HD searching and a responsive timeline that allows intuitive jog and shuttle playback controls. These features greatly reduce investigation time and allow security professionals to quickly access all the additional information captured by HD and multi-megapixel IP surveillance cameras. Conventional off the shelf database technologies are unable to provide the efficiency and power required to quickly search and playback HD video. This results in cumbersome investigations and limits an operator's ability to realize the value of high-resolution cameras.

Another core technology Avigilon Control Center uses is a unique file system architecture to manage storage efficiently for HD video surveillance system. The file system used to store the video content in ACC makes use of files broken into different intelligent buckets whose content and structure are managed by ACC and indexed within the embedded high performance database. The use of these buckets allows ACC to strictly manage fragmentation while enabling technologies like data aging and backup of HD video in highly efficient ways not possible with other systems. Data aging, a high definition stream management (HDSM) technology, allows JPEG and JPEG2000 compressed video to be automatically managed to store video at half the original frame rate after a set period, and then at a quarter the original frame rate after a further point of time. The unique file system used by Avigilon allows data aging to be applied at different times for different cameras on a single server allowing finer control of storage allocation then other VMS that only rarely offers any aging settings and typically only has a global value.

Summary

Advances in storage technology are accelerating the transition from analog to digital in surveillance and enabling security enhancing technologies like HD surveillance systems. A full understanding of these technologies empowers security professionals to understand the options available when purchasing or designing video surveillance. Beyond understanding the technologies involved, knowing how a VMS uses storage is necessary to understanding how the VMS can affect the quality of evidence a surveillance system stores. Avigilon Control Center's purpose built approach to using the best available storage technologies and architectures provides users with a system that is highly scalable and provides rapid access to the best evidence and reduced investigation times.