STORAGE TECHNOLOGIES

FOR

VIDEO SURVEILLANCE

Video surveillance is rapidly moving from analog to IP systems. As a result, storage of surveillance images has changed from analog-based recording solutions, like VHS tapes, to digital storage on hard drives. Further, hard drive capacity is doubling approximately every two years. At the same time the cost per unit is declining rapidly. **Video surveillance storage on hard drives is less expensive, more efficient and progressively powerful.**

HDD

The transition from analog to digital image also drives the demand for higher resolution and higher definition images (more detail and better evidence in surveillance video). This requires appropriate technology solutions to efficiently store, search, play back and manage the stored images.

Storage Technologies

The hard disk drive (HDD) is the central component of digital storage of surveillance video. HDD technology has evolved incrementally from the basic concept pioneered over 50 years ago. 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. The fundamentals of HDD technology have not changed much. However the increased density of data storage has resulted in a dramatic reduction in the size of the HDD. This change has resulted in an exponential increase in capacity with decreased cost. **The typical price per gigabyte of raw HDD storage is, today, a mere fraction (1/1000th) of what it was ten years ago.**

Innovations in speed and interfaces enable HDDs to handle larger amounts of incoming data. They also enable that data to be quickly provided back to applications, enabling advances like HD and megapixel surveillance. **A shortcoming of HDD technology is its reliance on mechanical moving parts. Moving parts always carry some risk of failure. To counter this, most surveillance systems use multiple disks in redundant configurations (RAID).** 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. The elimination of moving parts results in greater reliability and resilience. SSDs also offer substantial increases in speed for both reading and writing data. However, as of now, SSDs are expensive for large scale deployment for storage needed for most surveillance systems.

Redundancy – RAID

RAID is a **r**edundant **a**rray of **i**ndependent **d**isks. **Hard drives have the possibility of failure so redundancy is critical to ensure that stored data is not lost and that systems will continue to operate in the event of a drive failure.** Further, implementing RAID also increases the performance of a storage system by increasing the throughput, far 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. **Mirroring is the simplest form of redundancy and simply copies the data on two drives.** As any data is written, it is created in both the drives. **With striping, the data is 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.
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. RAID 0 uses striping without parity or mirroring. This enables 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.
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 is typically used only for operating systems or programs, and rarely as the RAID level for large amounts of storage like surveillance video.

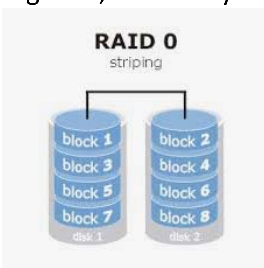

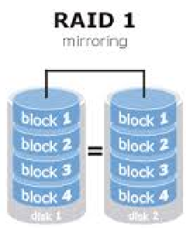

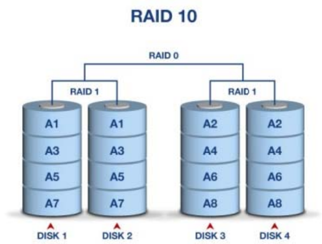

RAID levels 5 and 6 are more commonly used for large amounts of storage for surveillance video.
RAID 5 uses striping with distributed parity. This means that data is both spread out over multiple drives and additional information is stored to allow data to be rebuilt after failure. For example if you 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 will not always be written to the same drive with every file.
In the event, any one of the eight drives fail, 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.

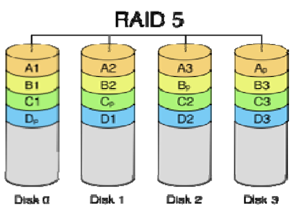

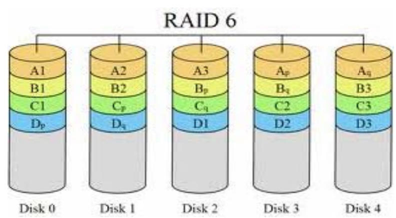

RAID is implemented either at the hardware or software level within a server or storage enclosure. For the performance needs of surveillance and truly enterprise class storage, a hardware RAID controller is a necessity for any surveillance video storage. Implementing RAID at the software level relies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

Connectivity of Storage with the Surveillance system.

The three commonly deployed methods are (i) directly attached storage (DAS), (ii) storage area network (SAN) and (iii) 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 is done via a serial attached SCSI (SAS) cable from a dedicated hardware RAID card.

SAN is typically used for very large amounts of storage that require access from multiple servers. A Storage Area Network (SAN) is a high-speed sub-network of shared storage devices. A storage device is a machine that contains nothing but a disk or disks for storing data. A SAN's architecture works in a way that makes all storage devices available to all servers on a LAN or WAN. As more storage devices are added to a SAN, they too will be accessible from any server in the larger network. In this case, the server merely acts as a pathway between the end user and the stored data. Because stored data does not reside directly on any of a network's servers, server power is utilized for business applications, and network capacity is released to the end user.

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, network-attached storage (NAS), is a server that is dedicated to nothing more than file sharing. NAS does not provide any of the activities that a server in a server-centric system typically provides, such as e-mail, authentication or file management. NAS allows more hard disk storage space to be added to a network that already utilizes servers without shutting them down for maintenance and upgrades. With a NAS device, the server handles all of the processing of data but the NAS device delivers the data to the user. A NAS device does not need to be located within the server but can exist anywhere in a LAN and can be made up of multiple networked NAS devices.

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. 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. 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. However performance needs typically 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 is part of a surveillance system.

Relational databases play a key role in storing and indexing data for video surveillance systems. They also optimize features like search, playback, backups, and export.

Summary

- The increased density of data storage has resulted in a dramatic reduction in the size of the HDD.
- There is an exponential increase in HDD capacity with a dramatic decrease in cost.
- Video surveillance storage on hard drives is less expensive, more efficient and becoming progressively powerful.
- Advances in storage technology are accelerating the transition from analog to IP in surveillance and enabling security enhancing technologies like HD surveillance systems.
- RAID implementation is critical to ensure that stored data is not lost and that systems will continue to operate in the event of a drive failure.
- SAN products are making the transition from Fibre Channel to the same IP-based approach NAS uses.
- Rapid improvements in disk storage technology in today's NAS devices, offer capacities and performance that once were only possible with SAN.
- There is a partial convergence of NAS and SAN approaches to network storage.
- The boundaries between NAS and SAN are expected to blur, with developments like SCSI over IP and Open Storage Networking (OSN). Under the OSN initiative, many vendors are working to combine the best of NAS and SAN into one coherent data management solution.