Pivot3

FIRST IN HYPER-CONVERGENCE

BEST PRACTICES FOR VIDEO STORAGE INFRASTRUCTURE

CONTENTS

EXECUTIVE SUMMARY

Systems must be carefully thought out to ensure crucial data isn't lost.

Successfully recording video from hundreds of security cameras 24 hours a day, seven days a week without losing a single frame is a very complex challenge. If that wasn't tough enough, your system also has to allow for future growth, and show that it's reducing your claims and/or shortening response times.

Video surveillance is becoming more and more important as perceived and actual physical security threats increase worldwide. Hardware and solutions proliferate, even as budgets have flattened or turned downward. Whether you've been a security professional for decades, or your IT department just inherited video surveillance, there's a morass of technologies to wade through to find the right components.

This white paper focuses on how to specify video storage, it explains how video is unique in the world and why systems must be carefully thought out to ensure crucial data isn't lost. Central to the discussion is a review of the trio of storage technologies you're likely to run across:

DIRECT ATTACHED STORAGE (DAS) – hard drive storage usually found inside a dedicated network video recorder (NVR). Widely used for years as a simple replacement for analog recorders, it is designed for small facilities requiring a handful of cameras. Performance is high, because the data is close to where the user is, however, it cannot scale and storage capacity is fixed.

NETWORK ATTACHED STORAGE (NAS) – a storage device connected to a network. Unlike DAS, NAS was designed from the ground-up to enable groups of people to share work files (documents, email, PowerPoint presentations, etc.) over a computer network. NAS storage includes a file system. This extra file layer creates increased network traffic (good for creating, reading, and sharing documents in a general-purpose IT environment, but very bad for recording streaming video from hundreds of cameras).

STORAGE AREA NETWORK (SAN) – like NAS, a SAN is a storage device connected to a network. However, it works very differently, making the raw hard drives directly available for writing (recording) and reading data. This "block-level storage" is perfect for video, but you're also paying more for very high reliability, storage capacity, performance and data protection.

Before diving into the details of each of these technologies, the next few pages look at some broader considerations to have in mind while writing video surveillance system requirements for your building, complex, or campus.

IMPORTANT CONSIDERATIONS FOR BUILDING AND UPGRADING SURVEILLANCE SYSTEMS

In 2015, video surveillance spending in the US is expected to grow to $37.5 billion

You can spend hours Googling "surveillance systems" or an entire week at a trade show and easily come away confounded by the plethora of security hardware, software, and services. In this first section, we outline the problem and a few considerations to keep in mind while developing the requirements for your physical security system.

THE CONUNDRUM OF SURVEILLANCE

Thirty years ago, you simply went out and bought some cameras, coaxial cable, and a VCR. Now, all the components are digital. The surveillance conundrum is clear to anyone who follows the news:

- Threats (real and perceived) are growing
- In response, the public, private companies, and governments are demanding more and better physical security
- Surveillance options are growing in number and capability
- The growing camera population (with ever-higher resolution) is creating a flood of data
- Cameras never stop recording and what they "see" must be stored somehow
- Despite event-driven spikes, security budgets have generally declined in recent years.

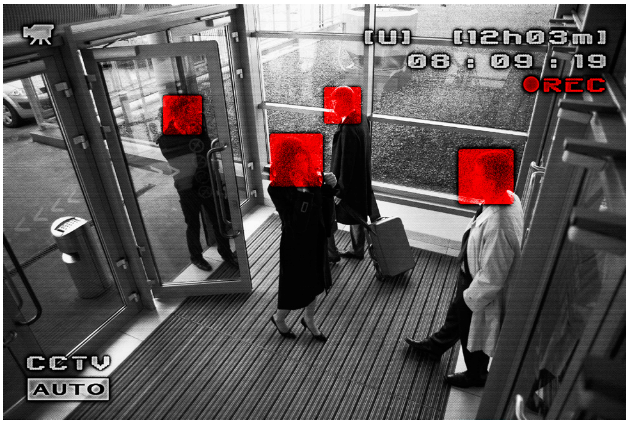

A growing demand for surveillance

In 2015, video surveillance spending in the US is expected to grow to $37.5 billion. There are more than 4,000 public area cameras in Manhattan. Chicago has 10,000. London? Half a million. Beyond public security, there's an ever-growing demand for video surveillance inside and around banks, casinos, school campuses, hospitals, hotels, transportation hubs and highways, railways, harbors, factories, power plants, and refineries. Of course, many of these systems can serve a dual purpose, such as speeding up ferry departures based on traffic conditions, remotely-monitoring trucks as they're being loaded, or alerting hotel staff of a VIP's arrival.

Analytics: the shift to digital tranforms watchers into actors

Sci-fi and action films may offer windows into the future of surveillance, but today you can now tap into the wealth of information found in the real world. When tied to biometric readers (such as iris scanners[3]) and using behavioral analysis algorithms[4], video surveillance systems can now monitor numerous real-time scenes and automatically respond with, say, a coupon to a shopper who shows interest in a particular shirt, or an alert to unusual activity in a subway station. With automation[5], human eyes aren't needed for the mind-numbing task of watching a bank of video monitors[6]. Personnel can focus instead on stopping a bad guy in the act, correcting a problem, anticipating a need, or providing a service. In fact, advanced users are turning their surveillance data from cost centers into cash.

A virtual machine is a software emulation of part or all of a computer. Today, hardware is so powerful that you can run several virtual machines on a single physical computer.

DIGITAL EQUIPMENT AND SOFTWARE

The shift to digital has also changed the way that video surveillance systems are built. Instead of endless "home runs" of coaxial and power cables from a control room to each camera, IP (internet protocol) cameras and monitors can be networked just like computers. Cameras can even be powered by the same Ethernet cabling that transmits their video data.

Today's video surveillance systems typically have at least one computer server running video management software (VMS). The VMS enables users to control the cameras and monitors, as well as search archived "footage" in storage. Storage can either be inside the VMS server (as DAS in a NVR) or in a separate storage device on the network (NAS or SAN). Computer processing and storage infrastructure software underlies the VMS application layer, ensuring that all your equipment is working as it should, with little or no administrative burden. You can also run all the software and storage on virtual machines.

Ever-improving cameras mean ever-growing data streams

Whether you are securing a small office or large factory campus, now that cameras are digital, you're able to take advantage of Moore's Law and watch prices drop as sophistication soars. However, a lower price also suggests the temptation of buying more. Better capabilities (like high-definition) offer better detail in a wider range of light conditions. With a 180˚ or 360˚ view, one camera[1] can do the work of several analog eyes. For example, the wide angle can enable you to watch an entire parking lot, then pan or zoom electronically to read a license plate or see a face. The downside is such cameras require a huge amount of network bandwidth and storage.[2]

A 1.3 megapixel camera can generate 18 gigabytes of data every eight hours[†]. Imagine how much data a system with hundreds or thousands of cameras would produce.

As the name suggests, each IP camera has its own IP address and connects to the network with a standard RJ-45 jack. Often, it has a built-in web server, email client, FTP client and supports Power over Ethernet (PoE) standards. As IP cameras become more sophisticated, they're able to stream to more than one destination, perform more processing and analytics, and make adjustments for changing environmental conditions (such as rain or fog), lighting changes, and reduce frame rates if a scene is unchanged (thereby lowering bandwidth and storage loads).

System performance is measured in terms of how many cameras or video streams can be delivered without dropping frames. Camera and video management vendors use the combination of:

Number of channels supported: a "channel" is one camera viewed by one or more people

Image resolution: 1.3 megapixel, 2 megapixel (1080p), etc.

Frames per second (fps): the higher your fps, the better your video quality

Compression method: MPEG-4, H.264[7], H.265

Network bandwidth[‡]	Number of megapixel (MP) cameras streaming simultaneously[§]		
	1.3MP	3MP	5MP
10Mbps	2	1	1
100Mbps (fast Ethernet)	23	13	10
1Gbps	230	130	105
10Gbps	2,300	1,300	1,040

A camera with a 180° or 360° view contains several lenses and sensors. Onboard software stitches multiple images together to look like one.

A 180° camera with five lenses can generate five times the data of a single-lens camera.

[†] *"A 1.3MP camera generates 18GB per 8-hour day" assumes a single camera with H.264 compression recording constantly for eight hours.*

[‡] *Mbps is megabits per second. Gbps is gigabits per second.*

[§] *Assuming all H.264 cameras are recording continuously and simultaneously and are set to 15 frames per second (half the frame rate of full-motion video).*

HOW IS VIDEO DIFFERENT?

- **Surveillance cameras never stop streaming content. Ever.** *This reality is a complete reversal of what most traditional storage solutions are designed to accommodate.*
- **Video content is constantly being written to disc (recording).** *Most other types of digital content are read-intensive.*
- **Write-speed is crucial in video storage.** *In "normal" IT environments, storage devices are optimized for quick reads (providing files to users). Video playback is infrequent, occurring only when there's been an incident that requires review.*
- **Video's volume is huge and fluctuates with activity.**
- **A large amount of bandwidth is crucial in video surveillance systems.** *It must be able to handle huge traffic spikes (called pixel storms) caused by an increase in scene activity (such as a classroom change).*
- **When a video surveillance system becomes overwhelmed, it simply drops frames or entire video streams with no warning.** *Why? Because the content never stops coming. In a normal IT datacenter environment, when the system becomes overwhelmed, users experience slow response times. With video, the only options are to drop the frames/video or to not store it. Users become aware of this silent form of data loss only when they are trying to retrieve the content days later. How can your organization respond without this vital evidence?*

Pixel storms *occur when there's movement in a digital camera's field of view. (Imagine the changing of classes at a school.) A camera can immediately double its frames per second (or more).*

Streaming video from here to there: networks

IP networks are indifferent to the kind of data that moves through them. Like a highway, a network is all about volume: how much traffic can you move back and forth? Performance is measured in bandwidth: what's the throughput capacity available to move data?

Like traffic lights, IP networks also have routers and switches. Performance is measured by the number of packets per second a device can direct and what the latency (delay) is for each packet as it moves through that device.

Overall system performance is measured by the number of cameras or video streams the surveillance system can deliver to live monitors and storage without losing frames or an entire stream.

Beyond pure bandwidth and latency considerations, just how are you going to power your remote cameras? Using the Power over Ethernet (PoE) standard on a 1Gbps wired network, you can kill the networking and power birds with one stone.

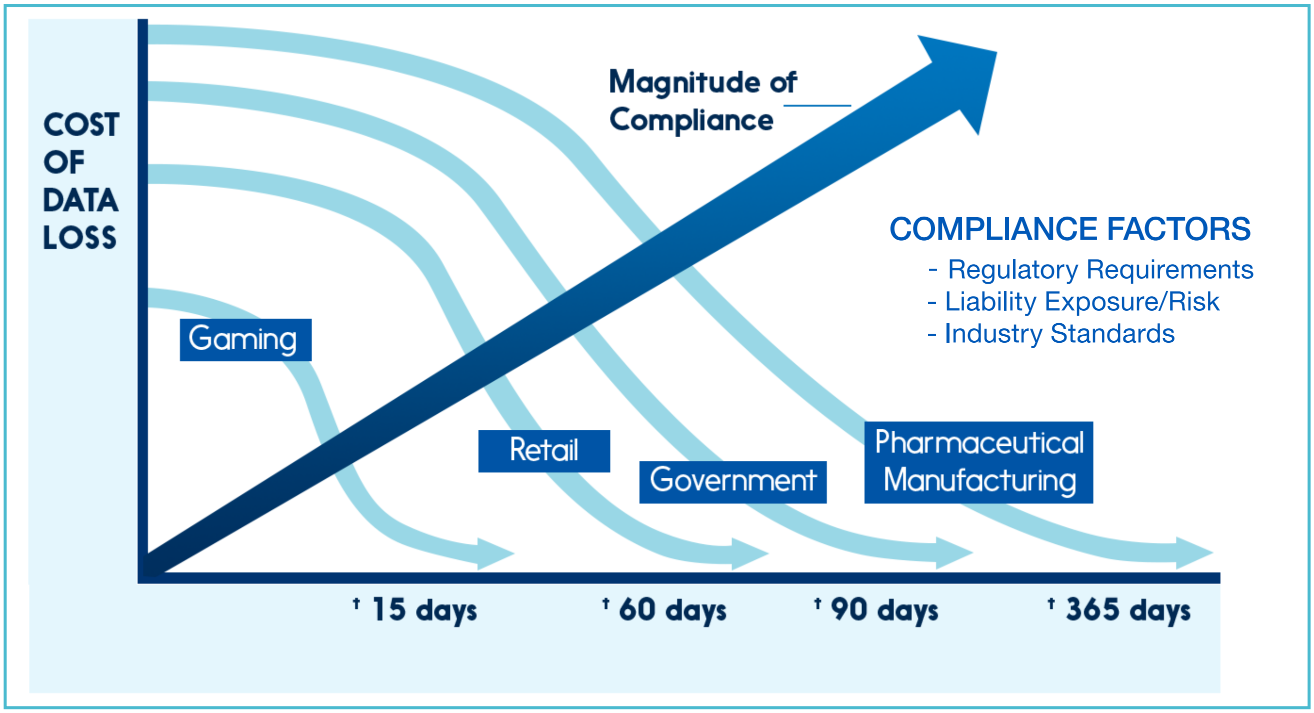

The value of video evidence is highest in the first 24-48 hours

When there's an event, you want to be sure that your system captures it, and that your team has immediate access to it. Be sure to specify a system that has failover capabilities measured in minutes, not hours. The longer your system (or a crucial component) is down, the less likely you'll be able to apprehend a suspect.

Likewise, the value of video surveillance evidence declines over time when it's required for regulatory or legal compliance. Make sure that the system you specify protects adequately against data loss for the retention period your organization is required to follow.

Strict budgets demand flexible tool

When you're reviewing the cost of your video surveillance system, look at the total cost of ownership over a five-year period. How flexible will the system be as your needs will change over time? Can you incrementally add to the system or will you be required to replace major components? Can you justify a higher capital expenditure if the system costs less to maintain, is easier to use, or ensures that when you're looking for footage a month after an incident, all the frames will be there?

IOPS (input/output operations per second) is a term you'll frequently hear in storage

IOPS

STORAGE TECHNOLOGIES

Fully one-third to one-half of the total cost of your video surveillance system will be storage equipment and administration.

There are several ways that you can capture streaming video for safekeeping. This section describes the three very different storage technologies you're likely to run across in your research and discussions. First, however, a brief history on the development of the boxes called "recorders."**

A brief history of video recorders

Videocassette recorders (VCRs) record analog video and audio on a magnetic tape. Television stations began using them in the mid-1950s, and the first home video recorders were introduced a decade later. By the late '70s, they had become a mass-market product. With improved equipment, longer recording times, and lower costs, VCRs were quickly adopted by the surveillance profession, and reigned supreme for more than two decades.[8]

DVRs are VCRs with captive digital storage

By the early-2000s, digital video recorders (DVRs) had dropped in price enough to begin outselling VCRs. DVRs control one or more cameras and record an analog video and audio feed on some kind of digital media (DVDs, hard drives, and/or flash drives). (The analog-to-digital conversion is performed in the box.[9]) Despite their widespread popularity (they're relatively cheap), DVRs have a number of shortcomings:

- **Bandwidth is fixed.** When you run out of physical ports, you'll have to buy a second box.
- **Storage capacity is fixed, and captive.** Data can not be shared with other DVRs.
- **Access is restricted.** To search for and review stored content, you must be at the DVR.[10]
- **Reliability is very limited.** When a DVR breaks down, there's no automatic failover. It simply stops recording.

When the system is specifically designed for video surveillance, it turns out to be particularly robust, flexible, scalable, and reliable.

** *Also called 'archivers'.*

NVRs: computer servers running Windows

A network video recorder (NVR) can be thought of as a souped-up DVR. NVRs work with IP cameras and offer more features, more throughput, and can handle a greater number of cameras than DVRs can. With the falling prices of IP cameras and NVRs, both have been gaining in popularity in the last few years. An NVR can be a box (thus sharing the same 1980s-era limitations as its DVR cousin), or it can be software-based: loaded on a Windows-based computer server, and connected to external storage. However, given the higher cost and complexity of software-based NVRs, most "solutions" come in the form of self-contained boxes with internal, direct attached storage (DAS).

Software-based recorders offer reliability and scale

In the late 2000's, several vendors began to take advantage of the obvious shortcomings of DVRs and NVRs and introduced competing software-based video processing and storage systems. This software is loaded on virtual machines, which in turn are hosted on commodity computers. The approach takes advantage of the substantial advances enterprise IT has made over the last two decades, by allowing many users and devices to share resources over a network. When the system is specifically designed for video surveillance, it turns out to be particularly robust, flexible, scalable, and reliable. For example, when a single hard drive fails (which they do!), your system can be built to ensure your critical evidence doesn't go with it. If you want to add more cameras, you can. When there's a blizzard of activity, the system has a much better chance than a DVR or NVR of capturing every frame amid the pixel storm.

Now that you've reviewed the different video recording form factors (boxes), take a look at the following sections and the varied approaches to actually storing a video stream.

DAS

Direct Attached Storage.

Designed for small facilities with fixed storage capacities.

NAS

Network Attached Storage.

Offers widespread network access to data. NAS can also scale much better than DAS.

SAN

Storage Area Network.

Have seamless consolidation and sharing of storage space, making them more efficient than NAS.

Direct attached storage (DAS)

DAS is, as it sounds, the storage that is found inside a DVR or NVR box. (In fact, an NVR is often called a "DAS recorder.") Without question, DAS is by far the most dominant storage type used in video surveillance. Why? Because onboard storage is a simple concept and bundled storage inside a DVR or NVR appears, at first, to be inexpensive. Performance is very high, because the data is close to where the user is. (Latency is very low.) However:

- **DAS is designed for small facilities requiring only a handful of cameras.** Some folks attempt to circumvent this inherent restriction by placing several DAS systems in the same control room, with each machine tied to a different collection of cameras. However, this approach creates far more problems than it solves.††

- **Storage capacity is fixed, and captive.** Upgrades or expansions usually require replacing systems, rather than simply adding to them.[11]

If you have a small facility that only requires a handful of cameras and you know that your storage capacity requirements are fixed, then DAS may be the way to go. The remaining questions are:

- How valuable is your data? [12]

- How accessible do you need your data to be?

Network attached storage (NAS)

NAS was developed in the early 1980s as a way to share remote file access with a number of networked client computers. Businesses and governmental agencies alike quickly saw the benefits of sharing information and resources (infrastructure, servers, printers, etc.) among their employees. Over the next two decades, the use of NAS systems proliferated as IT departments stitched together ever-growing and changing organizations and teams of employees.

NAS systems are perfect for IT environments: they're designed to record data in a file format, and they perform very well for typical office applications like word processing, email, and accounting. To understand why this core design feature is problematic for video, it's necessary to dive a bit below the surface to see how NAS works.

File fragmentation in NAS impacts playback and recording performance over time.

- Streaming surveillance videos are saved on a hard drive for a time, then deleted.

- A single video file may be spread in segments across several locations on the disc.

- As writes and deletes are repeated over weeks, files become more and more fragmented, which results in declining playback and recording performance.

- The solution is to run a de-fragment program. While the defragmentation process can run in the background, fragmentation can occur almost as quickly (in a video surveillance environment). In essence, this creates a nearly constant defrag state, severely affecting system performance.

†† Such as "stranded capacity" (one unit has extra storage available, but other systems cannot use it because DAS machines are not networked). This is to say nothing of the difficulties associated with trying to monitor or synchronize the footage on more than one standalone machine. For a multiple-DAS use case, see how the U.S. Park Police used to manage security video at the Statue of Liberty:

http://www.securitysystemsnews.com/article/total-recall-secures-lady-liberty

> *NAS was designed for read-intensive, general-purpose IT workloads,* ***not write-intensive video surveillance.***

A dominant protocol for communicating with NAS storage devices is the Common Internet File System (CIFS, usually for Windows-based machines). This network protocol enables a client computer (say, an NVR, PC, or workstation) to manipulate a storage device's files as though they were on the local machine.

CIFS works by sending packets from the client to the server. The server then checks to see if the request is legal, verifies that the client has the appropriate file permissions, executes the request, then returns a response packet to the client. The client then reviews the response packet to determine whether the initial request was successful.

All of this back-and-forth traffic is necessary for users who are creating, editing, sending, receiving, and deleting files such as documents and email. In the case of video recording, however, the additional traffic takes up crucial bandwidth and only impedes performance.

NAS requires that all writes to storage access an additional layer (the storage device's file system). In addition to creating increased network traffic, this extra layer has another hidden cost: file fragmentation (see box, preceding page).

In surveillance systems, numerous video streams are continuously writing to the NVR, forcing the file system to constantly reallocate space on the hard drives. The result? Fragmentation is much worse than what normally occurs in a normal IT environment. While some systems allow continued operations while defragmenting processes are underway, they nevertheless experience major performance issues which can result in dropped frames or video streams.

Despite these drawbacks, NAS offers some important advantages over DAS:

- Video data is more easily available to you and your colleagues (since it's on a network)
- NAS can also scale much better than DAS

Storage area network (SAN)

While NAS offered IT users widespread network access to data, each application server still had its own internal storage device. In effect, this created "islands" of storage that proved cumbersome to access. In the mid-1990s, NASA began to research clustering several application servers together on a network with a shared pool of storage.[13] By 1999, several vendors were offering SANs to the commercial IT market. Thought leaders began defining a SAN as a specialized high-speed network[14] that creates universal storage connectivity for all the storage devices, servers and client computers.[15]

This behind-the-scenes network enabled SANs to have seamless consolidation (and sharing) of storage space, making them much more efficient than their NAS competitors.[16]

A SAN server also makes storage available at a lower block level. (The file system is only found on the clients.) This design is perfect for capturing streaming video, enabling the data to go from the camera through the network and directly to a generalized pool of storage, without interruption.

As you can guess, putting all of your data storage eggs in one basket means that SANs are engineered to be:

- *highly available (reliable): must be able to withstand component failures without interrupting access or losing data*
- *highly scalable: capacity and input/output performance (bandwidth) must be able to grow as you add greater load (cameras) and SAN storage*

Rather than creating several islands of information as in NAS:

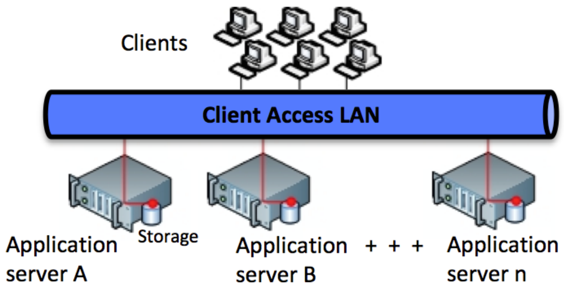

SAN presents an entire pool of storage to every server and client:

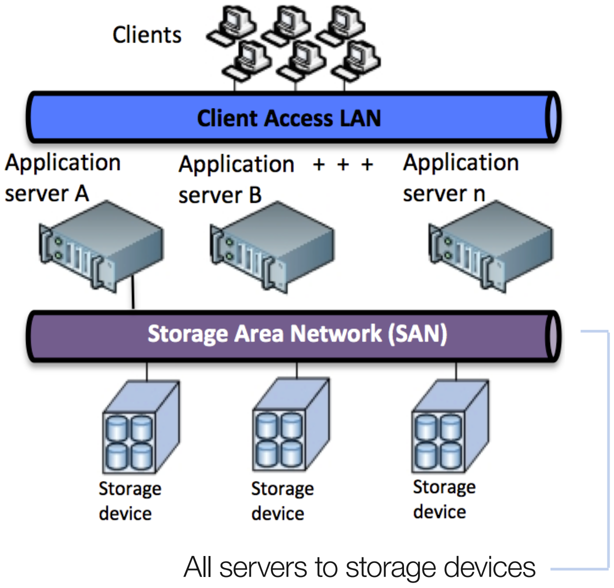

No matter how many storage devices you add to a SAN, cameras and other devices 'see' them as one destination.

Fibre Channel SANs

In answer to the highly-scalable requirement, the first SANs used a Fibre Channel (FC) infrastructure. Its fiber-optic cabling and switching provided a much faster and more reliable storage access than NAS's higher-level file system protocols. Indeed, FC SANs proved to be superb tools for applications like online transaction processing and providing access to big databases.

Unfortunately, the cost and complexity of FC SANs (not to mention the advanced storage administration skills and certifications needed to install and configure them) largely restricts them to the datacenters of Fortune 1000 companies and large governmental agencies.

iSCSI SANs

The basic SAN concept was enticing to a much wider market: after all, who wouldn't want the kind of scalability and reliability that the technology offered? In the mid- to late- 2000s, several start-ups developed iSCSI SAN, which used IP over a one gigabit Ethernet network to interconnect the storage devices.

Today, iSCSI SANs on 10 gigabit per second Ethernet (10GbE) offer the same benefits (very high reliability, storage capacity, and bandwidth) without the high hardware and administration costs of FC SANs. As a result, iSCSI SANs have made huge inroads into the IT datacenter market, displacing FC SANs at the high end as well as NAS in the small and medium-sized business market.

However, their inherent design (perfect for a general-purpose IT environment with lots of small, random inputs and outputs) requires you to greatly over-provision them in order to manage the constant barrage of incoming video from hundreds or thousands of sources. Fortunately, a newer, much more flexible technology is available that offers the pooled-resource benefits of an iSCSI SAN without requiring the sizable capital investment and operating expenses of a six-foot tall, 1,500-pound machine. *(See Emerging technologies on page 16.)*

> ***A packet** is typically a computer's request for an action of some kind, such as: "open file," "read file," "close file."*

Snapshot comparison of storage technologies

So, how do the different technologies compare in a video streaming environment?

Consideration	DAS	NAS	SAN
Reliability	Low	Moderate	High
Level of fault-tolerance protection	Low	Moderate	High
Storage	Captive to the box	Storage device “islands”	Pooled storage
Capacity	Fixed	Moderate	High
Scalability	Requires hardware replacement	Add additional NAS devices	Add new hard drive trays, with additional control units *(see next section)*[§§]
Input/output bandwidth			
Capacity	Fixed	Moderate	High
Scalability	Requires hardware replacement	Moderate	Moderate
File system fragmentation	N/A	Yes	N/A

Datacenter SANs vs. SANs optimized for video

If you’ve decided to take a look at SANs for your surveillance system, keep in mind that your business use case is very different from what most IT departments are trying to solve.

Consideration	Datacenter SAN	Video surveillance SAN
Workload optimized for	Short-block, 80% ***reads***, maximizing cache-hit ratios and fast IOPs.	Streaming data ingest and 95% backend ***writes***.[‡‡‡]
Bandwidth scalability	Speed is critical to fast reads, so the system should be able to scale ***up***.	Aggregate bandwidth is important, so the system should be able to scale ***out***.
Disc drives	Rotation (read) speed is important.	Spindle count (number of drives) is important.
Performance for video?	Poor	Excellent

[§§] *Scaling-out a SAN is incremental until you reach its cabinet’s physical capacity to contain hard drive trays (new storage) and their attendant control units (additional bandwidth). At that point, you have to add a second cabinet.*

[‡‡‡] *Write-throughput is all about the number of hard drives you have. In fact, real-world use shows that a larger group of slower-spinning, larger-capacity SATA hard drives can outperform their much more expensive 15,000 rpm Fibre Channel cousins.*

Scale up vs. scale out

The point about bandwidth scalability (on the previous page) is worth further discussion. A traditional IT SAN is a scale-up not a scale-out model. What that means is that the SAN has a single control unit (which may include some redundancy) with a fixed amount of bandwidth. That's why IT storage SANs have "knee-of-the-curve" performance issues, meaning that they can't balance the bandwidth available on that one controller with the capacity demands of a growing storage backend. A video-optimized scale-out SAN adds bandwidth with every new set of hard drives. This feature is critical for handling all the streams your hundreds or thousands of cameras are sending.

An aged file system is one in which all the storage capacity has been allocated for use, then deallocated a few times.

Simply put, make sure that the SANs you are researching are optimized for video surveillance.

Benchmarking performance

If you want to compare the performance of the various technologies that are in the marketplace, be sure the tools and configuration you use are designed to measure and mimic the unique requirements of a video surveillance environment. As you might guess, many of the benchmark tests measure aspects that are important to IT datacenters (such as IOPs), while overlooking problems that occur after a few weeks of normal video surveillance use (such as heavy disc fragmentation).[17]

For video surveillance systems, some important benchmarks include:

- Before beginning your tests, the system must have a failed drive or storage component. A degraded system sets the maximum input/output you can expect in a real-world situation.

- What is the number of frames per second recorded as cameras are added to the system?

- At what point does the system being tested begin dropping video frames or streams?

- How well is bandwidth utilized? Can you add bandwidth as you add cameras and storage?

- What happens after the disc storage system is "aged?" Do the discs become fragmented?

- What is the expected failure rate for the system's disc drives?

- What happens when more than one hard drive fails in an array? Does the system still operate at its rated performance?

Emerging technologies

As mentioned in Software-based recorders offer reliability and scale (page 10), virtualization software allows you to host several virtual CPUs on a single box. The technology has enabled administrators to consolidate enterprise IT resources considerably, saving equipment and power costs.[18]

One emerging technology, called hyperconvergence, takes virtualization a step further. It abstracts the CPU and the other computer infrastructure components wihin a single machine: RAM, storage, and network cards, for example. [19] What if you could virtualize computing, networking, memory, and storage resources across more than one physical machine? As your surveillance needs grew, you could incrementally grow your general pool of video storage one 3.4-inch tall (2U) server appliance at a time.

The main reasons that enterprise IT adopted external network storage technologies in the 1990s were: a lower total cost of ownership (TCO), a much higher performance and reliability, and a lower administrative burden. With the right network storage, properly tuned to the unique properties of video, all of these advantages are available to the surveillance community. To learn more about hyperconvergence and what it can mean for physical security systems, see the "Always-available Surveillance Video" white paper.

RECOMMENDATIONS

This white paper touches on the growing demand for and sophistication of video surveillance, and how those facts have completely transformed the surveillance landscape. DVRs and NVRs with DAS proliferated because they were simple "digital" replacements for disappearing VCRs. Yet, those systems cannot offer the scalability, bandwidth, and reliability that surveillance professionals demand. At first glance, it seems clear that NAS's design is perfect for IT datacenters (and day-to-day business applications like word processing and email). However, that doesn't mean it's a good fit for video. NAS's "islands of information" and inherent file system layer creates additional traffic on the very network you're trying to optimize for video data recording.

To ensure you spec the kind of video system that you and your surveillance team can depend on:

- Begin with the physical requirements: what rooms, buildings, perimeters, and campuses do you need to cover?

- How critical is your video data to your mission? If you need high reliability, scalability, and data protection, study the available iSCSI SAN options.

- If initial equipment cost is your top priority, DAS looks like a sound investment. However, DAS presents significant limitations, such as fixed, captive storage capacity and an inability to scale. Look instead at your total cost of ownership budget over a period of five years. (Equipment, additions, licensing, operating expenses, maintenance, and administration.)

- Find an integrator who understands your business and is willing to work with you to assemble the best, most cost-effective system for the long haul.

- When comparing and benchmarking systems, make sure that you replicate your own video streaming environment as closely as possible.

ENDNOTES

[1] Proctor, Keith. “The great surveillance boom.” Fortune. April 26, 2013. http://fortune.com/2013/04/26/the-great-surveillance-boom/, accessed November 29, 2014.

[2] Ibid.

[3] Dolmetsch, Chris. “’Minority Report’ May Come to Real World with Iris Recognition.” Bloomberg News. February 1, 2011. http://www.bloomberg.com/news/articles/2011-02-02/-minority-report-may-come-to-real-world-with-iris-recognition, accessed November 29, 2014.

[4] Levine, Barry. “Minority Report’ vision tech – coming to a smartphone or store near you.” VentureBeat. November 21, 2014. http://venturebeat.com/2014/11/21/minority-report-vision-tech-coming-soon-to-a-smartphone-or-store-near-you/, accessed November 29, 2014.

[5] Markoff, John. “Researchers Announce Advance in Image-Recognition Software.” New York Times. November 17, 2014. http://www.nytimes.com/2014/11/18/science/researchers-announce-breakthrough-in-content-recognition-software.html, accessed November 19, 2014.

[6] Sengupta, Somini. “Privacy Fears Grow as Cities Increase Surveillance.” New York Times. October 13, 2013. http://www.nytimes.com/2013/10/14/technology/privacy-fears-as-surveillance-grows-in-cities.html, accessed November 29, 2014.

[7] Video compression technology is fundamental to high-definition video surveillance. Without it, it would be impossible to transport and store the enormous amount of data that a single camera sensor generates, much less hundreds or thousands of cameras running concurrently. H.264’s strong compression enables high-definition video to be transmitted at low bit rates. Not only is it important for the surveillance world, but also for video streaming services like Netflix.

[8] “Videocassette recorder.” Wikipedia. http://en.wikipedia.org/wiki/Videocassette_recorder, accessed Nov. 27, 2014.

[9] To confuse matters further, some DVRs support both analog and IP cameras.

[10] One Source Security. “4 Differences between DVR’s and NVR’s.” http://www.onesourcesecurity.com/portal25/one-source-security-blog/entry/4-differences-between-dvr-s-and-nvr-s August 22, 2013. Accessed November 27, 2014.

[11] In multi-DAS solutions, it’s difficult to achieve consistent retention periods across all the cameras (because each camera records at a slightly different rate). This fact adds another complication when you attempt to scale beyond a handful of cameras: you have to “balance” them to ensure you can concurrently retire the footage on all your DVRs or NVRs.

[12] Both DAS and NAS storage devices typically include a redundant array of independent discs (RAID), which preserves critical data when a hard drive fails.

[13] Jordan, James. Storage Consolidation: Moving from DAS to SAN/NAS. Dell White Paper. August, 2002, page 5.

[14] Neema, Farid J. SAN Applications – classic article: What is SAN? StorageSearch.com. May, 1999. http://www.storagesearch.com/periphcart.html, accessed December 19, 2014.

[15] Storage Networking Industry Association. “What is a Storage Area Network.” http://www.snia.org/education/storage_networking_primer/san/what_san, accessed November 29, 2014.

[16] “Storage area network.” Wikipedia. http://en.wikipedia.org/wiki/Storage_area_network#cite_note-2, accessed November 28, 2014.

[17] While somewhat dated, the side-by-side NAS vs. iSCSI SAN benchmark tests performed by Intransa are still a good guide in this debate. See Intransa. “Video Surveillance and External IP Storage Solutions.” http://intransabrand.com/pdfs/wp/Video-Storage-IP-SAN-vs-NAS-Whitepaper.pdf, accessed Nov. 29, 2014.

[18] Marshall, David. “Top 10 benefits of server virtualization.” InfoWorld. Nov. 2, 2011. http://www.infoworld.com/article/2621446/server-virtualization/top-10-benefits-of-server-virtualization.html, accessed Dec. 11, 2014.

[19] “Virtualization.” Wikipedia. http://en.wikipedia.org/wiki/Virtualization#Desktop_virtualization, accessed Dec. 11, 2014.