

A WHITE PAPER FROM

CENTER FOR
DIGITAL
GOVERNMENT

Video Surveillance and Mission-Critical Data Storage: Getting the Whole Picture



Introduction

The use of video surveillance is an increasingly common fixture in public spaces as state agencies, local governments and school districts seek to more efficiently and effectively protect people and property. The news is filled with reports of shootings, terrorist activity and other crimes that were (or could have been) averted or more easily investigated and prosecuted with video surveillance footage. For example, in 2012, New York police used a live video feed to identify and arrest two suspects within minutes of a homicide¹ and in 2013, video footage helped investigators identify suspects in the infamous Boston Marathon bombings.² More recently, high-profile police officer-related shootings have prompted calls for video surveillance worn by law enforcement personnel.

Police and prosecutors are not the only officials relying on video footage. Video surveillance has also been a powerful aid during natural disasters, allowing responders to identify imperiled areas, manage emergency response teams more effectively, keep the public informed and document rescue missions for future training purposes. Video surveillance during Hurricane Sandy, for example, allowed the New York Port Authority to monitor subway stations for flooding and thereby keep passengers out of harm's way.³

As budgets tighten, populations increase and threats proliferate, video surveillance cameras will become an even more important resource for state and local government — but they require careful implementation. While digital and analog cameras are vital components of a video surveillance system, they are only one piece of a mission-critical surveillance solution. To extract the greatest value from video surveillance technology, organizations need a storage infrastructure that can reliably capture, retain and retrieve images of the activity that cameras monitor.

Given the unique requirements of video surveillance, traditional storage solutions do not always have the capacity, performance or reliability to accommodate ever-expanding volumes of data associated with high-resolution cameras and longer data retention periods. In weighing the costs of a storage solution for video surveillance, it's important to factor in the value of the data

This paper provides background on the unique storage challenges associated with video surveillance and discusses key considerations for selecting a storage solution that ensures public sector organizations have the images they need for mission-critical operations.

being stored and the price of losing that data in the event of failure. This paper provides background on the unique storage challenges associated with video surveillance and discusses key considerations for selecting a storage solution that ensures public sector organizations have the images they need for mission-critical operations.

Image Is Everything

Until recently, many video surveillance cameras could not provide the level of detail and clarity needed to effectively deter and investigate crimes or other threats. Today, high-resolution digital cameras and the capability to efficiently record higher frame rates (e.g., when alarms or alerts are triggered) are providing the image clarity and temporal context needed to capture an accurate picture of a particular event. As a result, officials are mounting more cameras than ever to keep an eye on mass transit systems, parks, schools, highways, stadiums, government buildings and other public spaces. According to one estimate, the compound annual growth rate (CAGR) of IP-based video surveillance will be 24.2 percent between 2013 and 2019 due to increasing demand for surveillance cameras with higher-quality video.⁴

State and local leaders are expanding their use of video surveillance to not only improve public safety but also reduce the costs associated with crime. The direct and indirect cost of crime in the United States amounts to \$3.2 trillion per year⁵ and includes expenses incurred by law enforcement, the justice system, corrections and victims. Investments in surveillance pay off — especially when active monitoring is part of the surveillance solution.



A surveillance program in Chicago, for example, saved about \$815,000 per month on criminal justice and victims' costs, with about \$4.30 saved for every dollar spent on the surveillance system.⁶

Maintaining an Image: Petabytes of Surveillance Data

Although better surveillance technology creates an opportunity to improve public safety and reduce crime-related costs, it also presents significant challenges in terms of data storage. In a traditional surveillance environment, cameras are connected to video feeds that continuously record data either directly to storage or, more commonly, to video recording servers that record footage to the storage array. With better video resolution (i.e., HDTV and megapixel cameras), higher frame rates and more cameras, the volume of video surveillance data in an organization can easily be greater than the capacity and performance capabilities of storage infrastructure designed for traditional operational needs. Longer data retention periods further compound the challenge. Depending on internal policies and regulatory

compliance requirements, organizations may have data retention periods that span weeks, months or years.

The following example provides some idea of the magnitude of the challenge when meeting storage needs for video surveillance:

For a single network video camera at 1920 x 1080 (full HDTV format) resolution, 30 percent compression, a recording rate of 30 frames per second per camera (for an aggregate data rate of 3.1 Mbps), 24/7 recording and a 30-day retention period, the minimum amount of storage needed is 1.15 terabytes (TB).⁷

Top Considerations for Mission-Critical Video Storage

Given the unique requirements of video surveillance, what has always worked for an organization's traditional storage needs may not be up to the task of surveillance-related data storage. To capture, retain and retrieve video surveillance data, organizations must carefully consider write speeds, reliability, storage footprint, bandwidth requirements and more.

Performance — Read/Write Speeds

In a typical storage environment, files are “written” once to a disk and then “read” by users as needed; the bulk of the server’s work is in enabling these read operations. With surveillance systems, the opposite is true. Because surveillance data is continually ingested into the system, the storage server’s write workload can be as high as 99 percent of the workload.

The read workload for video surveillance is small most of the time because stored data usually is not accessed unless some sort of incident requires review. When an incident does occur, however, the read workload can spike significantly. The storage system must be able to go into overdrive, allowing reviewers to rapidly access and retrieve content, fast forward video, intelligently search for motion or objects within an image, or search within a specific time frame.

The speed at which servers can process a read/write workload (i.e., a given amount of data) in a given period of time is called throughput. It is measured in megabits (Mb) per second. Because video surveillance systems must be able to manage not only the exceptionally high write workload, but also the read workload, it is important to ensure storage system throughput is not a limiting factor. Individual servers should be capable of read throughputs of 500 Mb/sec and write throughputs of 800 Mb/sec. These higher speeds help ensure frames are not corrupted or dropped during recording. They also enable a greater number of concurrent video streams, which helps organizations use bandwidth as efficiently as possible.

Availability

Unlike traditional storage architectures, storage for video surveillance does not include tape backups or other forms of duplication. It is simply too costly and complex to back up surveillance data. Because there is only one opportunity to record a particular moment in time and only one stored record of events, a video surveillance system must be extremely fault tolerant.

This means every component of the storage solution — from video management applications to IP networks and storage arrays — must operate at the highest levels of reliability and availability.

What about the Cloud?

Although cloud storage is a cost-effective and increasingly popular option for traditional storage scenarios, the high cost of continually transmitting video data to the cloud usually precludes cloud storage for high-bandwidth video surveillance environments.

When evaluating and/or designing a data storage solution for video surveillance, state and local governments and school systems should ensure the technology has not only undergone rigorous validation testing (with 99.999 percent availability for system hardware), but has also been proven in the field. In addition, the solution should include the following safeguards:

- **Built-in redundancy** — All storage components, including servers, disks and storage controllers should have an identical counterpart that can be swapped in immediately in the event of failure. Network paths, power supplies, fans and all other critical components of the solution should also be redundant. Redundant array of independent disks (RAID) configurations (e.g., RAID 5, RAID 10) can be implemented to further improve data availability; however, it’s important to keep in mind that various RAID configurations have trade-offs in terms of cost and fault tolerance.
- **Automatic, multi-path failover** — In the event a hard drive fails, the system should automatically detect the failure and re-route processes to another hard drive. It should also have the capability to rebuild failed drives by using either always-ready standby (i.e., “hot”) spare drives or available storage capacity in a dynamic disk pool.
- **Health monitoring and performance reporting** — The system should continually monitor the health of individual disk drives so defects can be proactively identified and corrected before a hard drive fails.

Storage Capacity, Footprint and Scaling

IT professionals face a number of challenges when containing the high volume of data generated by video



The best surveillance equipment in the world is useless if the network does not have the bandwidth and network configuration to support high-quality data transmission from cameras to back-end storage devices and monitors.

surveillance. Among these challenges are having sufficient storage capacity, controlling the storage footprint (and associated costs) and preparing for future storage needs. In some cases, solutions may benefit from a tiered storage architecture, where data is moved to the most cost-efficient type of storage (e.g., temporary vs. archival) according to actual requirements for performance, capacity or some other variable. Reporting and analytic tools are essential to determine the optimal distribution across tiers.

- **Storage capacity** — Depending on camera quantity, data retention period, frame rates and other factors, an organization may require hundreds of terabytes (1,000 gigabytes) or petabytes (1,000 terabytes) of raw storage capacity. For large deployments, where multiple factors add to the complexity of capacity planning, it is usually prudent to configure reserve capacity of 20 to 30 percent. Without this reserve, organizations risk losing data once the system reaches its maximum capacity.
- **Storage footprint** — Rack space is at a premium in data centers. To save rack space and reduce costs associated with the physical storage space (e.g., square footage, cooling and power), school districts and state and local government organizations should aim for industry-standard rack configurations that maximize storage density.

- **Storage scalability** — Scalability allows an organization to expand capabilities such as performance and capacity in small, manageable increments. A scalable solution typically has a modular design that allows growth where and when it is needed instead of forcing organizations to over-build at the beginning. With a scalable storage architecture, organizations can more easily phase in surveillance cameras over time (e.g., for budgeting purposes), respond agilely to new threats, and leverage existing investments in cameras and other technology to accommodate new use cases.

Support for a High-Bandwidth Environment

The best surveillance equipment in the world is useless if the network does not have the bandwidth and network configuration to support high-quality data transmission from cameras to back-end storage devices and monitors. Today's video surveillance environments increasingly rely on IP networks to transmit data. The bandwidth and quality of service (QoS) requirements for video surveillance are significantly higher than those for other IP-based functions such as Voice over IP (VoIP).

Organizations should conduct a network assessment to ensure the IP network can adequately handle a high-bandwidth environment. In particular, organizations should aim for minimal packet loss and latency.

To maximize investments and ensure success, state and local governments and education systems can work with a reputable technology partner on strategic tasks such as planning, component sourcing, technology integration, system configuration, life cycle management, training and regulatory compliance.

- **Packet loss** refers to video frames that are lost in transmission. This may occur, for example, when the transport mechanism (i.e., protocol) gives up its bandwidth to another process when the network becomes congested.
- **Latency** refers to the time it takes for the data from the video feed to reach its destination. Latency, combined with bandwidth, affects data throughput (speed). If latency is too high, video frames could be dropped before all the data reaches storage from the network.

Ease of Use for Non-IT Professionals

It takes considerable skill to manage, maintain and update a high-performance, high-availability storage infrastructure. The challenge is compounded by the sheer volume of data to be managed. A video surveillance solution can easily entail petabytes of storage management. To give some perspective on the staffing needs (and operational cost) of the storage environment, industry best practice used to be approximately 1 full-time employee for each 150 TB of data.⁸

System designs and automated storage management advances have recently made it practical for large video surveillance archives of 700TB+ to be effectively managed by non-IT operators on an exception-only basis. The automation manages the day-to-day functions, including calling home for service when needed. In a perfect world, education institutions and state and local governments would have the budget to hire experienced storage professionals to store, protect and manage valuable data; however, many organizations must rely on internal teams that are already stretched thin and may not have the skillsets needed to effectively manage storage complexity. To minimize the burden on the internal IT team, organizations should look for a solution with an intuitive

storage management interface, a single, centralized view of all storage systems (arrays) in the organization, and the capability to manage arrays from one or more IP-connected workstations. In addition, administrators should be able to update an array and perform other management tasks without taking the array offline.

A Picture is Worth...

As might be expected, costs go up for high-performance, high-reliability and low-maintenance storage solutions. In determining the best storage solution, state and local agencies face a tough balancing act between the costs of data storage and the less easily quantified costs of losing mission-critical surveillance data due to poor performance or system failure.

When considering budgets and investments in storage technology, it's important to factor in the cost of failure. How much is a record of events worth in terms of public safety, trial evidence, crime prevention, risk management or some other imperative? What happens if images are lost? The cost of failure is different for every organization, but as an example, consider the impact of an upgraded surveillance system at one county jail in the United States. With higher-quality images covering all of the jail, the county was able to reduce the cost of lawsuits from \$12 million to \$3 million per year. In addition, inmate-on-inmate and inmate-on-staff violence decreased by 92 percent.

Ensuring Success

To maximize investments and ensure success, state and local governments and education systems can work with a reputable technology partner on strategic tasks such as planning, component sourcing, technology integration, system configuration, life cycle management, training and regulatory compliance.

When selecting a technology partner, decision-makers should keep in mind the following considerations:

- **Stability/longevity** – Video surveillance technology has a long life cycle. It's important to choose a company that has the commitment, financial resources and track record to be a long-term partner.
- **History of providing high-capacity, high-performance solutions** – As complexity increases, a vendor with hands-on experience, proven processes and established suppliers becomes indispensable.
- **Expertise in the entire technology supply chain** – Video surveillance environments have many moving parts. In addition to storage expertise, a vendor should have in-depth knowledge about surveillance cameras, video management software, Internet cabling, IP networks and other variables affecting the overall solution.
- **Support** – Given the mission-critical nature of video surveillance, organizations need a partner that offers industry-leading support in terms of response times, event escalation and problem resolution.

By combining proven technology with industry expertise, state and local governments and education institutions can ensure they'll have the video images they need to protect their communities and accomplish their missions.

Endnotes

1. "NYPD Cameras Capture Deadly Shooting, Help Cops Nab Suspects," NBC News, March 2012, www.nbcnewyork.com/news/local/NYPD-Camera-Murder-Shooting-Jamaica-Houses-Queens-VIPER-141306653.html
2. "FBI Scanned 13,000 Videos, 120,000 Photos In Boston Marathon Bombings Probe," '60 Minutes,' CBS Local News, March 2014, <http://boston.cbslocal.com/2014/03/21/60-minutes-fbi-scanned-13000-videos-120000-photos-in-boston-marathon-bombings-probe/>
3. "Video shows Hurricane Sandy turning Hoboken and Exchange Place PATH subway stations into tunnels of roaring whitewater," NY Daily News, November 2012, www.nydailynews.com/new-york/video-shows-sandy-turning-path-stations-tunnels-roaring-whitewater-article-1.1203055
4. "Global Video Surveillance and VSaaS Market is Expected to Reach USD 42.81 Billion Globally in 2019," Transparency Market Research, October 2013, www.prnewswire.com/news-releases/global-video-surveillance-and-vsaaS-market-is-expected-to-reach-usd-4281-billion-globally-in-2019-transparency-market-research-229836881.html
5. "Study by Professor David Anderson Estimates the Annual Cost of Crime to Be \$3.2 Trillion," Centre College, 2012, www.centre.edu/study-by-professor-david-anderson-estimates-the-annual-cost-of-crime-to-be-3-2-trillion/
6. "Evaluating the Use of Public Surveillance Cameras for Crime Control and Prevention – A Summary," Urban Institute – Justice Policy Center, September 2011
7. Video Surveillance Solutions with NetApp E-Series Storage, NetApp, July 2013.
8. "Defining and Measuring IT Efficiency: Maximizing Return on Storage Investments," Forrester, 2011, www.arrowecs.co.uk/ArrowECS/media/Main-Library/Storage/NetApp/SyncSort/wp-defining-and-measuring-it-efficiency.pdf



ALL IMAGES PROVIDED BY: SHUTTERSTOCK.COM



Underwritten by:



Avnet Technology Solutions

A global IT solutions distributor, Avnet Technology Solutions collaborates with its customers and suppliers to create and deliver services, software and hardware solutions that address their end-user customers' business needs locally and around the world. Avnet Technology Solutions is an operating group of Avnet, Inc. For more information, visit www.ats.avnet.com.

NetApp

State and local government agencies of all levels count on NetApp for software, systems and services to manage and store their data. Top cities, states and counties value our teamwork, expertise and passion for helping them succeed now and into the future. For more information visit www.netapp.com/SLG.



The Center for Digital Government, a division of e.Republic, is a national research and advisory institute on information technology policies and best practices in state and local government. Through its diverse and dynamic programs and services, the Center provides public and private sector leaders with decision support, knowledge and opportunities to help them effectively incorporate new technologies in the 21st century.

www.centerdigitalgov.com