



Block (SAN) versus File (NAS) Storage for Video Surveillance

Learn the differences so you can protect and optimize your video infrastructure

Background

To improve their efficiency and extend their capabilities, security teams worldwide are taking advantage of higher-resolution video surveillance cameras and analytical applications. The captured media is being used not only for security, but also for policy compliance and other critical business needs, and the result is longer retention and ultimately more storage. These advancements create a new series of challenges at the data storage layer, precipitating an increase in bandwidth, write speeds, and storage capacity.

Two architectures dominate the storage space for enterprise applications:

- NAS, which stores data as files and presents those files to the application as a network “drive letter”
- SAN, which looks like local storage, presenting capacity for the operating system to manage

This solution brief contrasts the two approaches in the context of video surveillance applications.

The Challenge

Video surveillance is now an enterprise-grade application. High-resolution cameras and analytics have added strenuous workloads, and the old appliance-based approach can no longer get the job done. Organizations now need enterprise storage technologies that can handle this rapid increase in video data, and that can provide the protection required for this high-value content.

Whereas the data storage workloads of the past were comprised of video streams from low resolution cameras, today’s surveillance is comprised of a more strenuous profile:

- Continuous high-performance (random) write activity from camera video streams
- Unpredictable random read operations to serve video on request
- High-speed (I/O) transactional access patterns for database activity
- High-speed memory manipulation of smaller chunks of data for analytics

Beyond handling the diversity and intensity of these storage workloads, the ability to scale capacity is critical to accommodate new camera zones and continuous upgrades in resolution.

The Solution

To enable the system to grow and change, while minimizing the system management workload, larger surveillance implementations are moving beyond NVRs and server-based storage options. Today’s video surveillance demands an enterprise-grade infrastructure, where the servers and storage are separate. This layered approach enables an increase in processing power, rapid I/O request processing, bandwidth, and capacity.

Among the choices for enterprise storage topologies, the SAN approach has emerged as the preferred option for video surveillance. IHS projections show the SAN market growing by over 15% from 2020 to 2022, whereas the NAS segment is dropping from 5% to around 2% annual growth. NAS technology is a solid fit in many areas, but databases and analytics workloads demand performance that mandates a direct-attached or SAN approach.

For this reason, video management software providers recommend local or SAN-connected storage.*

SAN For Virtual Environments

Although the preferred approach depends on each integrator's best practices, many video surveillance implementations run in virtualized server environments. In these cases, each virtual video management server (instance) needs high-performance storage, not only for video footage, but to run the operating system, applications, and databases. SAN is the only architecture with the performance to handle these workloads at the extreme reliability levels required.

Other security applications such as alarm management or access control can be consolidated on the SAN, simplifying monitoring and management.

Reduce complexity to increase availability with SAN

Both SAN and NAS are simple to use. They involve many of the same steps for deployment, as they both require Ethernet-based connectivity.

NAS solutions provide a shared file system so that files and directories can be accessed from multiple systems. These solutions must use file locking to prevent multiple systems from modifying files simultaneously. Because video management systems do not require shared access to video files, all this file locking and shared file system complexity is unnecessary overhead that limits performance and adds layers of software to maintain and secure.

Pay only for features you use with SAN

Deduplication and compression, also offered by many NAS systems, are not needed for video surveillance solutions. By choosing a solution with these features, additional expense is incurred for technology that isn't used. The unused features still have integrated software that negatively affects product performance and requires maintenance for security and reliability.

Storage tiering

Storing data on different performance tiers can be useful for video surveillance deployments. However, the video management software already manages the tiers, creating separate storage for databases, live recording, and archiving. Because the data is managed by the video management software, there is no need for a storage system to dynamically move data between tiers. Therefore, data tiering or auto-tiering as a storage feature is unnecessary and adds risk and complexity.

Better storage efficiency and scalability with SAN

Some scale-out file systems require multiple (server) platforms or nodes to scale. Multiple-node solutions require a back-end network that can pose problems: Each write operation creates a series of data transfers across the back-end network, limiting performance. The multi-node interconnect creates more potential failure points, and this approach can make adding or replacing capacity more complex. To achieve the same levels of hardware redundancy that SANs offer, these solutions end up delivering lower capacity.

Consider the E-Series—SAN-Attached Storage

Unlike scale-out (NAS) file systems, NetApp® E-Series systems can scale to hundreds of petabytes within the same environment. SANtricity, the array management utility for E-Series, enables the administrator to manage over 90 arrays – more than 270 PB from a single pane of glass.

Graphic displays enable you to monitor, fine-tune, and optimize system performance. Using Dynamic Disk Pools, it is easy to grow capacity without adding overhead. For example, when Tufts University doubled its video security infrastructure, E-Series systems enabled the university's IT department to add 238TB of video data with no additional headcount. According to Tufts Storage Specialist Cope Frazier, "The storage has been the easiest part."

Beyond scale, E-Series systems offer unprecedented flexibility. As a performance leader, the technology is designed to excel at handling both high-bandwidth and high-transaction workloads. This versatility makes it an excellent fit for the high-resolution cameras and analytics technologies rapidly coming to market, while providing a storage environment for any future upgrades or changes.

About Netapp

NetApp is the leader in cloud data services, empowering global organizations to change their world with data. Together with our partners, we are the only ones who can help you build your unique data fabric. Simplify hybrid multicloud and securely deliver the right data, services and applications to the right people at the right time. Learn more at www.netapp.com.

* From Milestone, *XProtect Storage Architecture and Recommendations* (January 31, 2019): "With the XProtect VMS a NAS may only be used for the archive databases. The reason for this is that the recording database needs uninterrupted block-level access to the storage system to ensure high performance and continuous recording of live media data. Because a NAS is connected to a standard IT network and uses file-level access, direct uninterrupted disk access cannot be guaranteed as needed, which means that even small delays or gaps in the communications will cause performance issues and loss of recordings."