SANs in Video Surveillance

Partner Solution Brief

Overview

Arxys and Seagate provide storage systems that are designed specifically for the unique demands of video surveillance, with capabilities to manage high-resolution video retention, provide high throughput capacities, and the ability to scale easily to meet current and future retention needs. High-availability SANs deliver advanced management, superb scalability, and reliable redundancy to larger video surveillance environments where retention, protection, and performance matter.

In the last ten years, we have witnessed the rapid transition from analog to digital recording technologies that utilize all internet protocol (IP) cameras. These IP cameras have made it possible to record images with higher resolutions than ever before by using advanced compression technologies and commercial-grade image processors. These high-resolution cameras, coupled with the demand for longer video retention times, have driven the need for increased storage. Today's security systems must be designed to scale for increased camera count, higher retention requirements, and evolving camera and surveillance technologies.

The technology cycle is not going to slow down; in fact, it will only accelerate. Customers stand to benefit by rolling out these advancements quickly—primarily through increased organizational safety, protection from expensive lawsuits, uninterrupted business continuity, and the valuable insights they can gain from video that will allow them to drive their businesses forward. However, they will also face obstacles as they try to use existing infrastructure, which is built for a finite number of cameras with much lower resolution.

While digital video surveillance systems offer an exceeding number of rich features compared to their analog ancestors, they still require thoughtful design and deployment to ensure maximum cost effectiveness. Among the various criteria that impact surveillance networks, storage capacity can be particularly significant and easily misunderstood.

Modern video surveillance environments include:

- Higher image resolution
- More active cameras
- Longer archival periods
- Intelligent video recognition

A critical component of video surveillance infrastructure is data storage, which is arguably the biggest investment made in an overall IP security system. Security directors have the responsibility to ensure that the storage architecture they choose today has the ability to scale for tomorrow as technology continues to evolve, and it should do so without the need for a forklift upgrade. With so many storage technologies available for video security, it's easy to choose an option that appears to be the best relative to its cost, but the total cost of ownership may not make sense when considering the big picture.

SANs in Video Surveillance

Abundant storage capacity is a fundamental enabler of such capabilities. At its core, video surveillance has two fundamental aspects:

1. Video capture and analysis
2. Video recording and retention

Every video surveillance installation has these two aspects in some combination. There are many ways to provide both. In this application note, we will focus on the recording/retention portion of video surveillance, highlighting best practices and guidance for where and when to best utilize storage area networks (SAN) in video surveillance.

How do you know if you should use a SAN for your video surveillance system?

24×7 video streams are the very lifeblood of video surveillance systems. To deliver superior performance and efficiency in a given security environment, these systems must have sufficient storage capacity to address three fundamental video parameters:

- **Quantity**—the number and time duration of the video streams
- **Quality**—the image quality of the video streams, expressed in terms of frame resolution (for example, 1920×1080 pixels, also called 1080P) and frames per second (fps)
- **Archiving**—the length of time the video streams will be stored

Every video surveillance system needs storage, but not every video system needs largescale, shared network storage. Many video surveillance environments are ably served by high-performance network video recorders (NVR) that combine video capture, analysis, and recording all in one device. However, there are installations for which a video surveillance optimized SAN brings new levels of flexibility, scalability, and protection. SAN is an excellent fit for a security application that needs to be scalable and flexible as business grows, retention requirements increase, and camera technology evolves.

SAN Benefits:

- A *one-to-one* relationship with NVR
- Exceptional manageability
- High throughput capacity
- Easy expansion
- Supported by most video management system (VMS) platforms
- Built-in redundancy

Large security systems for enterprise-class environments generally call for large numbers of high-resolution cameras and extended video retention times—usually more than 30 days. We frequently see these types of systems in industry verticals such as casinos, city surveillance, airports, and grow facilities; however, due to cost effectiveness and ability to scale, they are quickly becoming more popular in mid-tier installations, as well. A storage solution for this type of camera environment exceeds the capabilities and cost effectiveness of a typical NVR solution, offering a more robust architecture with the ability to scale over time. This is why SANs are an excellent fit for security applications that must scale and adapt as business grows and camera technology evolves.

A SAN allows for multiple NVR units to connect to a shared storage network at the block level for video recording. Connecting the recording appliances to the SAN is accomplished by iSCSI or Fibre Channel connectivity through a switch. The recording appliances will see the storage allocated by the SAN as a local hard drive as if it was directly connected, and the file system would be maintained by that recording appliance. A graphic user interface (GUI) on a web browser from any server that's connected to the storage system network allows for centralized management of the SAN. This creates efficiency and ease of management.

As the camera system grows, storage servers or expansion units can be easily added over time to extend capacity without requiring the administrator to re-architect the camera to NVR network. This allows for complete scalability of the storage system. In addition to scalability, the tight coupling of storage and software is effectively removed, allowing for the flexibility to change VMS software without the need for a forklift upgrade of the hardware. Finally, this type of architecture allows for greater redundancy, as it eliminates the single point of failure that is found in a typical NVR environment.

SANs in Video Surveillance

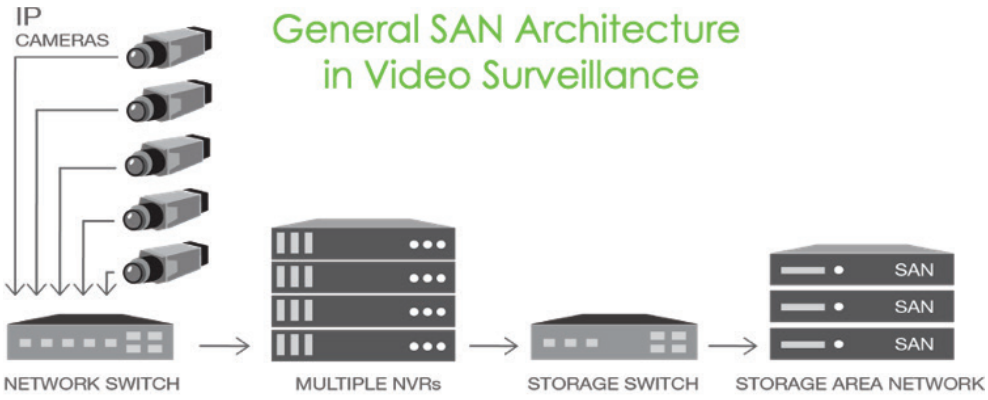

Environment A: Extended Archiving

Video surveillance is moving from a physical security backstop to a front-end security business necessity. In many video surveillance installations, emerging statutes, regulations, and other legal requirements are mandating ever longer video retention times. Whereas most video retention times average around 14 to 30 days, there's a large and growing requirement for much longer video archiving. Many large school districts and state government agencies are moving toward 60- to 90-day retention in light of increasing lawsuits and legal challenges for incidents. Detention centers across the country are averaging 90 days, with larger facilities moving toward one-year video retention. Combining these extended archiving times with Full HD 15 fps video feeds and full video motion detection analytics is changing the nature of video surveillance storage.

Increasing the retention rate of even mid-range video surveillance specs makes for exponential growth in storage requirements. For example, using mainstream parameters of 1080p and 15 fps continuous recording for only 200 cameras with 90-day retention requires over 540TB of usable storage. Extend that archiving to 365 days and you need a minimum of 2.3 petabytes of storage.

In cases of long retention and archiving, high-capacity network storage SANs are an excellent fit. If your video surveillance system needs long retention for critical video to meet regulation or legal requirements, building your system with a SAN-attached storage array delivers highly available, high-capacity storage for all of your video recorder appliances.

When your video surveillance system has a requirement for long retention, a good configuration is to have two to six recording appliances with high GHz, built-in graphics processing units (GPU), and solid state drives (SSD) for recording the live video. Those recording appliances then archive to the SAN on a set schedule to optimize storage and capacity. This enables a robust surveillance system, with recorders acting as the primary storage and the SAN fulfilling the long archiving retention requirements.

Typical surveillance environments:

- 100 to 300 cameras
- 2MP and up resolution
- 90- to 365-day retention
- Single site/installation
- Legally-mandated retention

Environment B: High Camera Quantities

Environments with more stringent security requirements (for example, schools, public buildings, and airports) rely on 24×7 video streams with higher resolution and more frames per second to better identify suspicious persons and activity. Coupled with broader camera deployment, the enhanced video quality of such streams gives security personnel a more detailed, comprehensive view of the areas under surveillance.

Not surprisingly, the high camera counts, along with the higher resolution, artificial intelligence (AI) meta tagging, and frame rates of such 24×7 video streams, can significantly increase the capacity and resilience requirements of a video surveillance system.

When building a large-scale video surveillance network that consists of hundreds of cameras, there are two approaches that make the most sense: 1) build using many independent recording appliances with built-in storage, or 2) build using many recording appliances with no archive storage and use network storage like a SAN to build video retention storage.

Both have benefits and drawbacks. Here, we will focus on building a large-scale video surveillance network that uses the recording appliances connected to a large SAN for retention and failover.

For those surveillance installations that reach thousands of cameras, a high-availability, scalable SAN is a strong choice. Building a network of recording appliances that does all of the video processing, video motion detection, and analytics while video archiving to the shared SAN brings scalability and flexibility.

Typical surveillance environments:

- 500 to 3,000 cameras
- 2MP and up resolution
- 30- to 90-day retention
- Mutiple sites or buildings
- Build multiple SANs to cover various scenarios

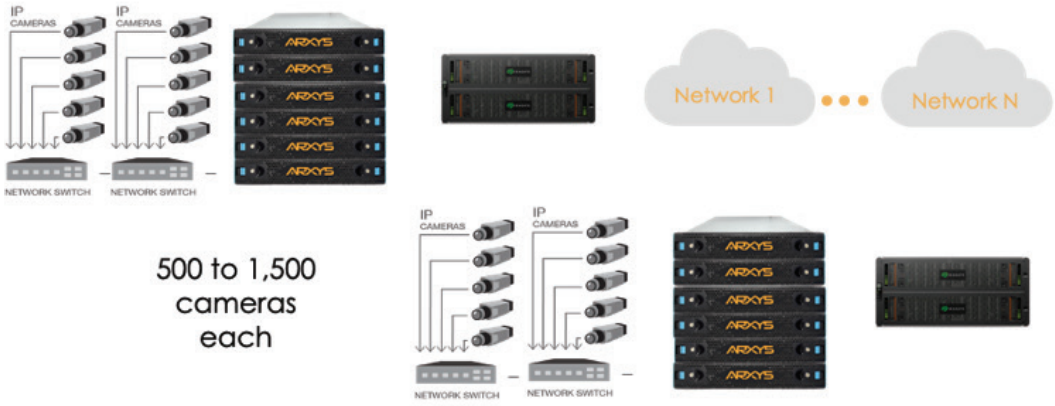

Environment C: High-Availability Scaling

Scalability is a growing focus in today's video surveillance market. The explosion of new IP cameras, Internet of Things (IoT) devices, access control, audio devices, and more is driving intense new scalability requirements for video surveillance systems. Additionally, the mainstreaming of Full HD resolution (1080P) and 15 to 30 fps and higher recording is being exacerbated by a seemingly endless march toward longer mandated retention times.

Even smaller-sized security systems that start out with 30 to 50 cameras are growing by leaps and bounds in terms of camera count, video motion processing, and video retention. Going from a single all-in-one appliance with 20 to 200 cameras to a distributed system comprised of many hundreds of cameras at exponentially higher bandwidth and processing rates is not an easy task.

Scaling Up a Server Versus Scaling Out a Network

Infrastructure scalability handles the changing needs of an application by statically adding or removing resources to meet changing application demands as needed. In most cases, this is handled by scaling up (vertical scaling) and/or scaling out (horizontal scaling).

Scale up is done by adding more resources to an existing system to reach a desired state of performance. For example, a database or web server needs additional resources to continue performance at a certain level. More compute, memory, storage, or network can be added to that system to keep the performance at desired levels.

Scale out is usually, but not always, associated with distributed architectures. Scaling out is essentially adding infrastructure capacity in pre-packaged blocks of infrastructure or nodes. Think multiplying appliances rather than enlarging a single monolithic server.

Both approaches work well and have various pros and cons. Video surveillance, however, has its own unique challenges and data flow architecture that make it lean toward scaling out appliances rather than attempting to cram it all into one monster box.

Scale Recording Appliances and Retention Appliances Separately

All of the leading global VMS companies have somewhat similar system design architecture. All of them have some sort of database performance or capacity limit—i.e. the limits of free database apps. All VMSs attempt to be flexible and open,

SANs in Video Surveillance

and all allow customers to choose the resolution, bandwidth, fps, and analytics of each video stream. This means there are wildly varying performance implications for every single video surveillance installation. When you add in growth, a flexible, scale-out architecture works best.

Using a scale-out system design, you can seamlessly scale your video surveillance system with ease. Scaling out allows you to: add more cameras, devices, resolutions, and frame rates; ensure every video stream is running video motion detection with analytics metadata; and empower smart search across every video stream. Scale up VideoX Appliances for video processing and live database recording, and scale up DataX SAN Appliances for video retention and higher availability with enhanced failover.

The other benefit of a distributed, scale-out topology is that taking advantage of the advanced failover features of Milestone Corporate and Expert, ONSSI Enterprise, and other top-end VMS versions delivers improved security and availability—and it does so at a much lower cost than traditional scale-up designs.

In video surveillance, best practice is to use a pool of failover servers to serve a group of active recording servers (see above). To improve reliability, the pool should have N+1 failover servers, with N equalling the highest number of planned concurrent outages due to regular maintenance. The failover servers are sized to support whichever active recording server in the group is the largest; therefore, any failover can easily take over for and support the full load of any active recording server from the group.

Typical Surveillance environments:

- 500-10,000 cameras
- 2MP and up resolution
- 30- to 90-day retention
- Mutiple sites or buildings
- High availability SANs for protection and scale
- Scale cameras, appliances and storage independently

seagate.com