



Technical Report

# Video Surveillance Solutions with NetApp E-Series Storage

## Introduction to Video Surveillance

James Laing, Jürgen Türk, Frank Poole, NetApp  
December 2017 | TR-4196

### Abstract

Video surveillance solutions that use NetApp® E-Series storage offer physical security integrators a highly scalable repository for video management systems, supporting high camera counts, megapixel resolutions, high frame rates, and long retention periods. The architecture is designed to provide high reliability and availability to meet the demands of video surveillance deployments.

## TABLE OF CONTENTS

<b>1</b>	<b>Introduction</b>	<b>4</b>
1.1	Publication Scope	4
1.2	Audience	4
1.3	Why E-Series?	4
1.4	Training Offerings	5
<b>2</b>	<b>Overview and Best Practices</b>	<b>5</b>
2.1	Video Surveillance Market	6
2.2	Surveillance Cameras	6
2.3	Retention Periods	6
2.4	Converged Networks	7
2.5	Standards-Based Open Architectures	7
2.6	Solution Components	7
2.7	Video Management System (VMS) Software	8
2.8	Deployment Characteristics	9
2.9	Best Practices Guidelines	9
<b>3</b>	<b>Solution Components</b>	<b>11</b>
3.1	Deployment Models	11
3.2	Network Video Cameras	12
3.3	IP Network Infrastructure	13
3.4	Video Management Software	13
3.5	Viewing Workstation	13
3.6	Video Recording Server	14
3.7	Storage	14
<b>4</b>	<b>Summary</b>	<b>15</b>
<b>5</b>	<b>Definitions</b>	<b>16</b>
	<b>References</b>	<b>17</b>
	<b>Version History</b>	<b>17</b>

## LIST OF TABLES

Table 1)	Training offerings	5
Table 2)	Key characteristics of a solution target deployment.	9
Table 3)	Best practices guidelines	10
Table 4)	Glossary	16

## LIST OF FIGURES

Figure 1) Solution component overview. ....	7
Figure 2) Recording server logical topology. ....	11
Figure 3) Typical resolutions. Images are to scale between resolutions. ....	12
Figure 4) E-Series for video surveillance. ....	15

# 1 Introduction

NetApp E-Series storage arrays provide performance, efficiency, reliability, and enterprise-class support for large-scale video surveillance deployments.

All video surveillance management software shares the common feature of recording live video feeds to storage for subsequent replay. This approach aids in forensic analysis or investigation of persons or events within the field of view of a single camera or a group of cameras. These video feeds, generated by hundreds or thousands of cameras, are typically configured to record continuously, 24 hours per day, 7 days per week, with retention periods in the range of months to years.

## 1.1 Publication Scope

This document provides an introduction to video surveillance for those who sell, design, or implement such solutions. It describes the comprehensive functional components that are required to build a video surveillance solution based on NetApp E-Series storage that can reliably record and archive video from recording servers. This document identifies the major components and features of a video surveillance system.

Additional video surveillance resources are available on the [NetApp Field Portal](#).

## 1.2 Audience

This report provides guidance to people and teams who are responsible for integrating NetApp E-Series storage systems into existing video surveillance deployments or for designing and implementing new deployments. This audience includes physical security integrators, video surveillance management software engineers, network and storage system engineers, and architects.

The content in this report is presented with the expectation that these professionals can combine this information with their experience and the supporting documents to build an efficient, scalable, and highly available system.

## Targeted Deployments

The targeted deployments for this introduction are large—from 200 up to over 5,000 cameras or more (1Mbps to 2Mbps, each saved for 30 days of archiving). These deployments have retention periods of at least 30 days and primarily use HDTV or megapixel resolution cameras.

## 1.3 Why E-Series?

The E-Series architecture supports block-based protocols and can process real-time video applications with high reliability, performance, and availability. For these reasons, E-Series is the preferred choice for video surveillance solutions that use NetApp storage.

## Solution Benefits

NetApp E-Series provides the following benefits for large-scale video surveillance deployments:

- **Intuitive management.** NetApp SANtricity® software provides a graphical representation of the E-Series storage, with an easy-to-use interface.
- **Ease of provisioning.** All management tasks to the array are performed by SANtricity software without taking the array offline.
- **High availability.** Dual controllers mean nondisruptive controller firmware upgrades, host multipath support, and dual paths to expansion shelves.
- **High performance.** The E-Series controllers offer an excellent price-to-performance ratio.

- **High capacity.** The E5700 systems support up to 5760TB of raw capacity (using 12TB disks) in an efficient footprint.
- **Drive health monitoring.** The E5700 provides proactive monitoring, background repair, and extensive diagnostic features of drives.
- **Data integrity.** Background media scans proactively check drives for defects and initiate repairs before defects can cause problems.
- **Data protection.** The E5700 supports Dynamic Disk Pools (DDP) technology or volume groups with RAID levels 0, 1, 10, 5, and 6.
- **Enterprise management.** The E5700 provides a single management view of all E-Series storage systems in the management domain.

## 1.4 Training Offerings

To enable successful deployment of the NetApp E-Series storage array, NetApp offers many web-based and instructor-led training opportunities. NetApp recommends the end-user training classes that are listed in the [NetApp University Customer Learning Map](#) under E-Series.

Table 1 lists the trainings offered, their duration, and the mode of delivery.

Table 1) Training offerings.

Class Description	Duration (Hours)	Delivery
<b>Product Foundation</b>		
Technical Overview of NetApp E-Series Storage Systems	01:00	Web-based
Technical Overview of NetApp EF-Series All-Flash Array	00:30	Web-based
<b>Hands-On Skills Development</b>		
Configuring and Monitoring E-Series and EF-Series Storage Systems	32:00	Instructor-led
<b>Certification</b>		
NetApp Certified Implementation Engineer: SAN Specialist, E-Series	01:30	Exam

## 2 Overview and Best Practices

Existing analog-based video surveillance systems are being replaced by network-based digital video surveillance equipment. IP technology is the preferred choice for new installations. This trend benefits end customers by addressing their physical security requirements with systems that offer more features at a lower cost.

This section provides an overview of:

- The video surveillance market
- Surveillance cameras
- Retention periods
- Converged networks
- Standards-based open architectures
- Solution components
- Video management system (VMS) software
- Deployment characteristics

- Best practices guidelines

## 2.1 Video Surveillance Market

The video surveillance market is characterized by several vertical markets at different stages of adoption. Gaming, manufacturing, transportation, education, and government/city surveillance are strong markets and have more aggressively implemented network-based digital video surveillance. Large enterprise manufacturing, service companies, and retail deployments lag, due in part to the physical dispersion of plants and facilities and the bandwidth requirements of networked video. Intelligent video analytics, including various forms of mobile surveillance that use body cams, dash cams, and drones, is an area of substantial growth.

Growth expectations for the industry, gleaned from financial reports of leading hardware and software suppliers of networked video systems, are estimated at approximately 17% compounded year over year and projected through 2022. The total addressable market (TAM) for video surveillance storage is more than \$8 billion annually. The market that is most suited to NetApp E-Series, 100TB or greater capacity with high availability, has a TAM in excess of \$2 billion globally.

Estimates for retail deployments indicate that storage is approximately 30% (and more for large projects) of the installation cost, with network video cameras and their installation at 25%. Servers, networking, and video management software make up the remainder. The market is strong and has good growth potential.

Retail enterprises, governments, and other public entities are faced with serious challenges regarding media storage. Traditional storage architectures are not designed for massive amounts of video content. And with more cameras, expanding retention periods, and rising camera bit rates, traditional standalone network video recording (NVR) solutions are becoming costly and inefficient. Customers who already use NVR systems are suffering from limited scalability and the limited space that they have in their data centers. As a result, most of them are converting to a server, storage, and VMS architecture for better return on investment and total cost of ownership.

## 2.2 Surveillance Cameras

Networked video surveillance cameras that offer more than 1 megapixel of resolution are now widely adopted because they offer at least four times the resolution of a standard-definition (4CIF) camera. Television broadcasting in HDTV resolution has changed end-user perception, and physical security managers are demanding the image clarity and higher resolution that HDTV/megapixel cameras and beyond provide. It is important for physical security integrators to manage end-user expectations. Even with the trend toward better resolution, lens quality, sharpness of focus, and lighting play a major role in determining image quality.

The increased resolution of networked video surveillance cameras contributes directly to an increased need for scalable storage.

## 2.3 Retention Periods

The retention period is the length of time that video is retained on storage for viewing and analysis. This parameter is regulated by a government agency, such as the State of Nevada Gaming Control Board; by a corporate policy; or by the necessities of costs and the availability of disk space. Retention periods are getting longer almost across the board. The typical retention period of 30 days is quickly becoming 90 days. Body cams are likely to have a minimum retention of 180 days, and that minimum affects fixed surveillance as well. Some situations already require year or multiyear retention.

Physical security managers generally prefer the longest retention period possible given efficient and cost-effective storage.

## 2.4 Converged Networks

Just as IP telephony deployments moved from disparate networks to a common IP network, the surveillance industry is also moving to a converged IP network. Modern physical security deployments are more than simple IP-based cameras. Most video management software also supports access control systems and integrates video and access control events. Building management systems and energy management systems are also internetworked and might generate alarms for abnormal temperature changes or when sensors detect water infiltration.

Although many deployments use dedicated access-layer Ethernet switches to support networked video cameras, switch selection should align with corporate standards for Ethernet LAN switching to leverage the support and expertise of the network management staff. At some point in the network topology, the network devices are interconnected, whether or not a fully converged network of voice, video, and data is implemented or some physical segmentation is present.

## 2.5 Standards-Based Open Architectures

There are two competing video surveillance standards organizations: the Physical Security Interoperability Alliance (PSIA) and the Open Network Video Interface Forum (ONVIF). Both of these organizations promote standards-based information exchange between networked video devices. The standards address concepts such as device discovery, media streaming, and exchange of metadata. Implementation of these standards facilitates the integration of video management software vendors, cameras, and other IP-based network devices that are sourced from different manufacturers.

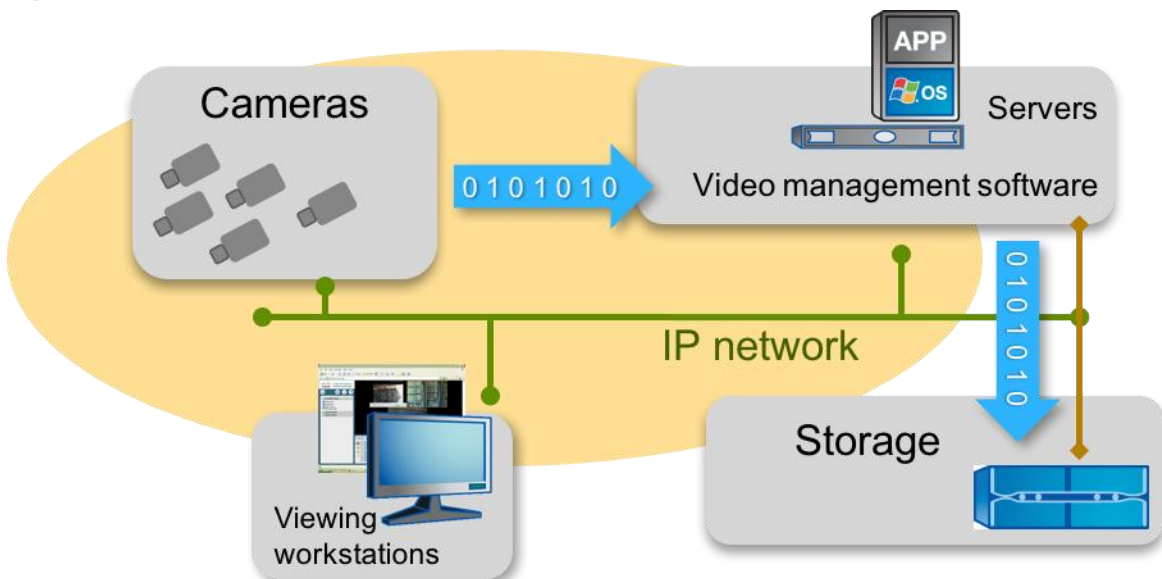
## 2.6 Solution Components

The typical video surveillance deployment consists of:

- IP network video cameras
- IP network infrastructure
- Servers and video management software
- Viewing workstations and other mobile viewing devices
- Storage

These components are shown in Figure 1.

Figure 1) Solution component overview.



In proprietary systems, all components are sourced from a single manufacturer. Open-platform systems allow the physical security integrator to select the IP cameras, network routers and switches, servers and workstations, and video management software and storage. The integrator's goal is to provide the best price, performance, and reliability to meet the specifications of the end user.

The physical security integrator might standardize on servers and workstations, network equipment, and storage for most of its business opportunities. However, networked video cameras and video management software are often selected based on end-customer requirements. Typically, the majority of cameras are from a primary vendor, but it is common to have the cameras of several vendors implemented in a single deployment. Additionally, analog-to-digital encoders might be used to include legacy analog cameras. It is uncommon, however, to see more than one video management software package implemented in a single deployment. It is important for the video management software and the storage array to work together seamlessly.

## **2.7 Video Management System (VMS) Software**

This document describes video surveillance solutions that are founded on open platform-based video management software. For example, Milestone XProtect Corporate, OnSSI Ocularis ES, and Genetec Omnicast/Security Center are qualified for use on the NetApp E-Series storage systems. NetApp has also worked with many other VMS providers such as Genetec and Virent to make sure of seamless compatibility with E-Series storage systems.



## 2.8 Deployment Characteristics

Target deployments for E-Series solutions are characterized by the key items that are described in Table 2.

Table 2) Key characteristics of a solution target deployment.

Element	Description
High camera counts	Rack space savings of up to 60% over competitive offerings can be achieved because of the maximum storage density of the E-Series using 60-drive 4U disk shelves.
Long retention periods	Video can be maintained for months to years by using high-capacity NL-SAS drives in the E-Series combined with the video-grooming technology of VMS.
HDTV/megapixel deployments	E-Series is well suited for the increased storage demands of HDTV and megapixel camera deployments because of its high storage density and performance.
High availability	A deployment should be designed and validated to provide high availability at the application, network, and storage system levels. Fault tolerance is a key component of all video surveillance solutions.
NetApp validation testing	Solutions have been validated with several video management system software offerings, the Axis Communications megapixel network video cameras, and the Axis virtual camera simulator. This validation incorporated thousands of video feeds in the recording servers of NetApp's video surveillance system technology partners. Frame rates of up to 30 fps from HDTV 720p and 1080p validate the performance of the solution.
Ease of use	NetApp SANtricity management software provides an enterprise view of all the storage arrays in the domain. Management of the arrays is not limited to the local network; storage arrays can be managed from one or more workstations with IP connectivity to the management interfaces of the arrays.
High performance	Validation testing demonstrates that the E-Series has the performance capabilities to support the requirements of video surveillance workloads. The throughput of the E-Series controllers is not the limiting factor in typical deployments.
Serviceability	Controller firmware can be upgraded without taking the storage array offline: a feature of the E-Series duplex controllers. Additionally, power supplies, cooling fans, and disk drives can all be replaced without system downtime.
Data protection	The E-Series supports DDP technology and volume groups with RAID levels 0, 1, 5, 6, and 10.
Drive health monitoring	The health of the individual disk drives is monitored, and problems can be identified before a hard drive failure. When a hard drive fails, the system incorporates automatic drive failover and detection and rebuilds by using available spare capacity in a DDP or global hot spare drives.

## 2.9 Best Practices Guidelines

Table 3 represents general best practices guidelines for video surveillance solutions.

Table 3) Best practices guidelines.

Element	Description
Number of cameras per recording server	The number of cameras supported per server is based primarily on the aggregate data rate of the configured cameras. However, features such as server-side motion detection might substantially decrease the number of cameras per server.
Number of virtual machines per physical server	As a rule, NetApp can support as many virtual machines as required, up to the throughput limit of the physical server. Typical configurations have 300 to 500 cameras per physical server.
Implement the Network Time Protocol (NTP)	An accurate time source is critical for the proper functioning of all video management applications. Synchronize all components (including IP cameras) with several accurate and reliable NTP sources.
Implement DDP or provision hot spare drives	Disk drives fail over time. DDP automatically provides spare capacity in the pool and recovers from failure quickly. If volume groups are used, provision the recommended hot spare coverage. Immediately replace failed drives.
Monitor the operational state of the storage array	The SANtricity Enterprise Management window provides an overview of the operational health of all storage arrays in the domain. Address all nonoptimal array conditions before they become critical problems.
Use the Recovery Guru	Refer to the SANtricity Recovery Guru to resolve reported problems.
Provision adequate reserve capacity	As a general rule, size the system with 20% to 30% of the reserve capacity for the target retention period. This approach allows increased capacity to address future requirements.
Allow SANtricity to automatically select drives for volume groups	The system attempts to provide both drawer and shelf loss protection if possible.
Verify equal distribution of volumes across controllers	Verify that volumes are on the preferred path for optimal balanced performance after storage array service or outages.
Implement recommended performance tuning options	Verify that all recommended performance tuning parameters have been implemented.
Conduct a network assessment before implementation	Recording servers only archive video they receive. Any network impairment between cameras and recording servers is lost video. Verify adequate bandwidth with low packet loss and reasonable latency for transporting IP video. Third-party vendors also provide these service offerings.
Verify that all components are operational	This validated design implements redundancy for high availability. While implementing the system, verify that all redundant network paths, power supplies, fans, and so on are operational.
If using volume groups, implement RAID 6 when feasible	RAID 6 provides an extra measure of protection over RAID 5: two parity disks rather than one.
Follow the proper electrostatic discharge (ESD) protocol	ESD-related component degradation might affect the long-term reliability of the system. ESD-caused degradation might not manifest into a hard outage until after months or years of service.

### 3 Solution Components

This section summarizes the overall architecture of a typical video surveillance deployment. It describes both the target deployment model and the individual components. The following concepts are discussed:

- Deployment models
- Network video cameras
- IP network infrastructure
- Video management software
- Viewing workstation
- Video recording server
- Storage

#### 3.1 Deployment Models

IP network-based video surveillance deployments are characterized by two deployment models:

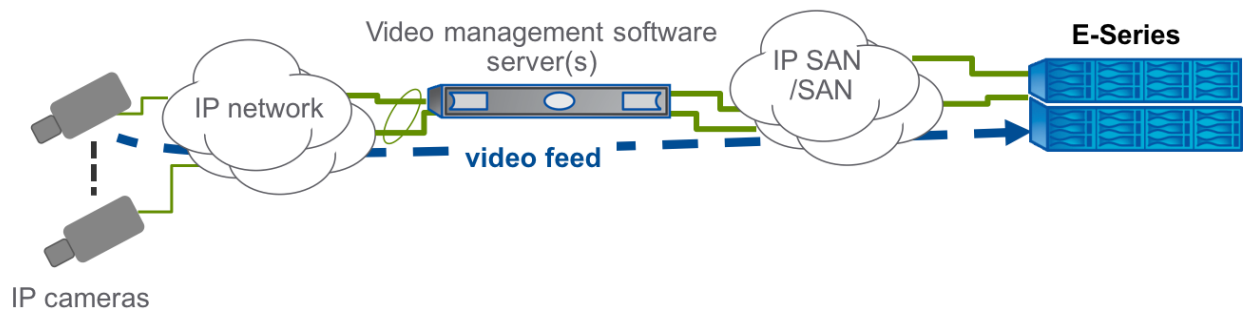
- Cameras streaming video to recording servers
- Cameras recording directly to storage

Implementations of cameras that record directly to storage include Bosch (iSCSI), MOBOTIX (NFS/CIFS), and IQInvision (NFS/CIFS). These implementations may have a server-based management platform for the control plane, but the media plane is direct from camera to storage. This deployment model is not discussed in this document. For more information about Bosch Security Systems, visit <http://www.boschsecurity.us/en-us>. Bosch offers several NetApp E-Series OEM configurations for [IP recording](#).

The target deployment focus for this document is the camera-to-recording server model. In this model, the recording servers have a control plane to the IP cameras and through the media plane receive one or more video feeds over the IP network by unicast and/or multicast packets. The media stream can be connectionless (H.264/UDP/RTP) or connection oriented (MJPEG/TCP or H.264/RTP and RTSP interleaved over TCP).

The recording server model is the more common of the two models and is supported by a wide array of open-system video management software vendors. A high-level diagram of the logical topology is shown in Figure 2.

Figure 2) Recording server logical topology.



The video management software market is predominantly based on Microsoft Windows Server 2012 R2 or later releases. Most software vendors support both NAS (NFS/CIFS) and SAN (iSCSI), as long as there is acceptable read and write throughput performance.

### 3.2 Network Video Cameras

IP surveillance cameras generate video feeds for both live viewing and video archiving by the video recording server. Most networked video cameras run a subset of the Linux operating system (Red Hat, SUSE, Ubuntu, and so on) and implement TCP/IP services such as HTTP/HTTPS, SMTP, SNMP, FTP, Telnet, and so on. Cameras increasingly include local storage either as internal flash memory or through the insertion of a secure digital (SD) nonvolatile memory card.

Networked video cameras are machine-to-machine (M2M) endpoints that are under the control of the recording server, which issues commands and responses by a combination of HTTP and the Real Time Streaming Protocol (RTSP). The initial configuration consists of assigning IP addresses, configuring NTP servers and the local time zone, and entering the camera name and descriptive information on the video overlay. It also includes adjusting the physical characteristics of focus, white balance, and color correction. Network video camera manufacturers design their cameras for easy installation to help reduce implementation costs for physical security integrators. Features such as autoback focus and power over Ethernet (PoE) provide installation efficiency.

Networked video cameras support a wide range of resolutions; the most common are standard definition CIF, HDTV, and megapixel. Both HDTV resolutions (1920 x 1080 and 1280 x 720) are megapixel resolutions, but not all megapixel resolutions are HDTV formats. Typical resolutions are shown in Figure 3.

Figure 3) Typical resolutions. Images are to scale between resolutions.



Because HDTV and megapixel cameras generate larger volumes of video data compared with standard definition, solutions that are based on NetApp E-Series storage are typically targeted at HDTV/megapixel deployments and beyond. Axis Communications is a market-leading networked video camera manufacturer and a NetApp partner, as are Sony and Bosch. Other predominant manufacturers include Hikvision and Panasonic.

### 3.3 IP Network Infrastructure

Video surveillance deployments require a network infrastructure that addresses the following requirements:

- Provide sufficient available capacity (bandwidth) to transport video.
- Exhibit very low or no loss of IP video packets.
- Feature network latency within the range that is suitable for the transport protocol (TCP or UDP) of the video feed.
- Provide high availability through network redundancy and best practices in network design.
- Meet network security and services requirements.

Video may be transported between endpoints using either UDP or TCP. Image quality problems (loss of frames) can occur in both transport methods. Although TCP is a connection-oriented protocol, TCP transport is the first to give up its bandwidth during congestion. In that case, real-time applications such as video might arrive too late and have to be discarded by the receiver because the playout time has passed.

Although IP network-based video surveillance deployments share many of the same SLAs as voice over IP (VoIP), the bandwidth requirements of video are substantially higher than VoIP. Additionally, each network camera streams video over the network constantly (24/7), whereas an IP phone uses fewer network resources unless there is an active call. Implementing network-based video on an existing network requires network quality of service (QoS) for data, VoIP, and video.

The physical security department or integrator must work with the IT department to implement network equipment that is consistent with the existing infrastructure. This consistency is important regardless of whether a physically separate network is implemented for video surveillance or video is converged on an existing network infrastructure.

Leading network vendors, as well as leading integrators who offer voice and video network implementation services, can help with network-readiness assessments for IP video surveillance deployments.

### 3.4 Video Management Software

The video management software supported with NetApp video surveillance solutions includes internal NetApp tested and partner self-certified software. The tested VMS partners include Milestone, Genetec, and OnSSI.

### 3.5 Viewing Workstation

One or more workstations capable of viewing live or archived video are a basic requirement of any deployment. The workstations must meet or exceed the hardware specifications of the VMS. Viewing video at higher resolutions and frame rates typically requires a high-end workstation with a video-gaming performance-class video card. Not implementing the minimum hardware that is required for a viewing workstation is a common deployment mistake and leads to issues with end-user satisfaction.

Low-resolution video might be viewable on laptops or smartphone applications when mobile or remote access to the video stream is more important than displaying the highest resolution. For example, Milestone XProtect Mobile is a free application for smartphones and tablets that works with XProtect video management software. The marketplace is moving toward bring your own device (BYOD).

### 3.6 Video Recording Server

The video recording server represents one or more instances of the hardware and software that are used to record live video to the storage array. The software can run on a physical machine or as a guest on a virtual machine. The guest virtual machine must have the same virtual memory and virtual CPU as specified by the video management system software requirements for a physical machine. The physical machine must include, at a minimum:

- One Gigabit Ethernet (GbE) interface for video ingress from network video cameras
- Either a dual-port FC host bus adapter (HBA), dual-port SAS HBAs, or dual Gigabit/10 Gigabit Ethernet (10GbE) iSCSI interfaces for connectivity to the storage array

The number of networked video cameras per recording server and the resulting data rate are determined by the architecture and best practices that are documented by the VMS provider. The server must meet or exceed the hardware specifications of the VMS.

As servers become more powerful and as storage gets denser and more powerful, the limiting factor for the number of cameras per server is the size of the failure domain. E-Series is fault tolerant, but servers are not. The storage is a nonissue in terms of performance, reliability, and throughput.

Integrators must look at the most likely point of failure. A server might be capable of connecting to 1,000 cameras, but if it fails, 1,000 cameras go dark. The current trend is to configure 250 to 300 cameras for any single physical server or virtual machine for this reason.

Video loss can occur between the network camera and the workstation or between the workstation and the storage array. On ingress, missing packets can be detected by gaps in the RTP sequence numbers. On egress, missing packets cause the video management server software to log archive-queue-full errors, media-overflow errors, or similar warnings. It might also display the number of records that are queued for I/O.

### 3.7 Storage

NetApp E-Series high-performance storage systems support the following block-based SAN protocols:

- E5700: FC (32Gbps), InfiniBand (100Gbps), iSCSI (25Gbps), and SAS (12Gbps)
- E2800: SAS (12Gbps), FC (16Gbps), and iSCSI (1/10Gbps)

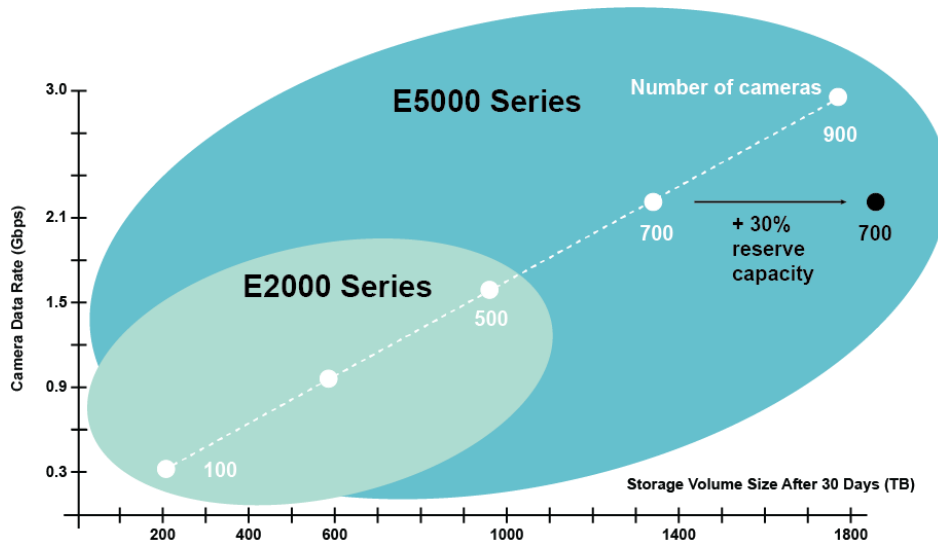
Video surveillance solutions have been validated with iSCSI, SAS, and FC host connectivity on the E-Series. All the components on storage array models are redundant, providing automated path failover. Online administration is accomplished through the NetApp SANtricity management client.

The E-Series is well suited for video surveillance archiving because it incorporates:

- **High throughput.** Up to 21GB per second for the E5700 controller.
- **Space efficiency.** 60 drives in 4RU (720TB raw storage with 12TB drives) and up to 5.7PB in 32RU.
- **Reliability.** Fault tolerance and redundancy are built in; all components are hot swappable.
- **Maintainability.** Firmware updates occur on one controller while the second controller handles all I/O.

Depending on the controller model, the E-Series scales from one 2RU shelf to up to three or eight 4RU shelves. Deployments can encompass from hundreds to thousands of cameras, depending on the camera data rate, retention period, and available free space for reserve capacity. The breadth of the solution is illustrated in Figure 4.

Figure 4) E-Series for video surveillance.



The target for E-Series storage in the physical security market is open VMS deployments, which enable the physical security integrator to design a solution that provides best-in-class network video cameras, servers, software, and NetApp storage.

The innovative functionality of the E-Series storage system makes it an optimal solution for large video surveillance deployments that use high-resolution cameras and have long retention requirements.

## 4 Summary

Video surveillance deployments that use NetApp E-Series storage offer physical security integrators a highly scalable repository for video management systems that support high camera counts, megapixel resolutions, high frame rates, and long retention periods. The architecture is designed to provide high reliability and availability to meet the demands of video surveillance deployments.

## 5 Definitions

Table 4 contains a glossary of the terms that are used in this document.

Table 4) Glossary.

Term	Definition
Controller	The controller includes the hardware board and the firmware that manage the physical disk drives and present that capacity to a computer as logical unit numbers (LUNs).
Dynamic Disk Pools (DDP)	DDP technology distributes data, parity information, and spare capacity across a pool of drives. Its intelligent algorithm (seven patents pending) defines which drives are used for segment placement, providing full data protection. DDP dynamic rebuild technology uses every drive in the pool to rebuild a failed drive, enabling exceptional performance under failure.
HDTV	High-definition TV defines resolutions of 1920 x 1080 and 1280 x 720 pixels along with other criteria, including aspect ratio.
LUN	The logical unit number is a unique number that the server uses to identify different hard drives or, in the case of storage systems, different volumes. Most operating systems show the LUN as properties of the SCSI hard drives that are discovered.
Megapixel	Any video resolution of 1 million pixels or more. The HDTV resolution of 1280 x 720 is 921,600 pixels but is commonly referred to as a <i>megapixel resolution</i> .
RAID	RAID is an acronym for “redundant array of independent disks.” RAID determines how data is protected from hard drive failures.
RAID 10	RAID 10 provides high availability by combining the features of RAID 0 and RAID 1. RAID 0 increases performance by striping volume data across numerous disk drives. RAID 1 provides disk mirroring, which duplicates data between two disk drives. By combining the features of RAID 0 and RAID 1, RAID 10 provides a second optimization for fault tolerance.
RAID 5	A striped disk with parity, RAID 5 combines three or more disks in a way that protects data against the loss of any one disk. The protected storage capacity of the volume group is reduced by one disk from the raw capacity.
RAID 6	Striped disks with dual parity, RAID 6 can recover from the loss of up to two disks. The protected storage capacity of the volume group is reduced by two disks from the raw capacity.
RTP	Real-Time Transport Protocol is a connectionless protocol for transporting voice and video over an IP network.
RTSP	The Real Time Streaming Protocol is used for establishing and controlling media sessions between endpoints. Optionally, media is embedded (interleaved) to transport media in the same TCP session.
SAS	Serial-attached SCSI (SAS) is a computer bus that is used to move data to and from computer storage devices such as hard drives and tape drives. SAS depends on a point-to-point serial protocol that replaces the parallel SCSI bus technology.
Storage array	The storage array is a collection of both physical components and logical components for storing data. Physical components include drives, controllers, fans, and power supplies. Logical components include volume groups and volumes. The storage management software manages these components.
TCP	Transmission Control Protocol is a connection-oriented IP protocol. RTSP is commonly implemented as TCP port 554. Motion JPEG is typically transported over TCP.
Volume group	A volume group is a set of drives that the controller logically groups together to provide one or more volumes to an application host. All the drives in a volume group must have the same media type and interface type.



## References

The following references were used in this technical report:

- NetApp E-Series Storage for Video Surveillance: The Advantages of Simple, Reliable Block Storage in Video Surveillance Environments  
<http://www.netapp.com/us/media/wp-7240.pdf>
- TR-4197: Video Surveillance Solutions with NetApp E-Series Storage: Planning and Design Considerations  
<http://www.netapp.com/us/media/tr-4197.pdf>
- TR-4198: Video Surveillance Solutions with NetApp E-Series Storage: Performance Considerations  
<http://www.netapp.com/us/media/tr-4198.pdf>
- TR-4199: Video Surveillance Solutions with NetApp E-Series Storage: Sizing Considerations  
<http://www.netapp.com/us/media/tr-4199.pdf>
- [Guided Solution Sizing](#)
- <https://fieldportal.netapp.com/content/204292>
- [Video Surveillance Storage Solution Page](#)  
<https://fieldportal.netapp.com/content/211536?assetComponentId=211635>
- Bosch Security Systems North America  
<http://www.boschsecurity.us/en-us>

## Version History

Version	Date	Document Version History
Version 1.0	June 2013	Initial release.
Version 2.0	November 2014	Updated with new controller models.
Version 3.0	December 2016	Updated with new controller models.
Version 4.0	December 2017	Updated with new controller and disk models.

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

### **Copyright Information**

Copyright © 1994–2017 NetApp, Inc. All rights reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

### **Trademark Information**

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

TR-4196-1217