



Technical Report

Video Surveillance Solutions with NetApp E-Series Storage

Performance Considerations

James Laing, Jürgen Türk, Frank Poole NetApp
December 2017 | TR-4198

Abstract

Video surveillance solutions using E-Series storage offer the physical security integrator 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

1	Introduction	4
1.1	Publication Scope	4
1.2	Audience	4
1.3	Video Surveillance Storage Solutions	4
2	I/O Characteristics	4
2.1	Viewing Archived Video	5
2.2	Performance of Tiered Storage	6
3	Server Planning	7
3.1	Hardware Recommendations	7
3.2	CPU	7
4	Product Selection and Performance	8
4.1	Overview	8
4.2	Operational Considerations	8
4.3	E-Series Storage Array	9
4.4	Configurable Performance Options	11
4.5	E-Series Performance Checklist	13
5	Example One: E-Series Storage Array E2800	14
5.1	General Performance Considerations	14
5.2	Typical Data Rates for This Example	14
5.3	E2800 Performance Summary for Example One	15
6	Example Two: E-Series Storage Array E5700	15
7	General Performance Considerations	16
7.1	Typical Data Rates	16
7.2	E5700 Performance Summary	17
8	Performance Validation of Tiered Storage	17
8.1	Archive Function	17
8.2	Recording Server	18
8.3	I/O Latency	20
9	Other Performance Considerations	21
10	Summary	22

Definitions	23
References	24
Version History	25

LIST OF TABLES

Figure 1) 64-camera transition from 1x to 16x.....	6
Figure 2) E-Series raw storage capacity.....	10
Figure 3) E2700 hardware and software components.	14
Figure 4) E5700 hardware and software components.	16
Figure 5) Storage configuration RACK-SVR-37.....	18
Figure 6) Write rate during archive.....	19
Figure 7) I/O latency RACK-SVR-37.....	21

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 to aid in forensic analysis or investigation of persons or events within the field of view of a single camera or 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 is intended to provide information for those who need to determine, analyze, or troubleshoot performance of a video surveillance system that uses NetApp E-Series storage.

1.2 Audience

This publication is intended to provide guidance to physical security integrators, video surveillance management software engineers, network and storage system engineers, and architects responsible for integrating NetApp E-Series storage systems into existing video surveillance deployments or designing and implementing new deployments.

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

Targeted Deployments

The targeted deployments for this introduction are large, 200–2,000 cameras or more, (2Mbps) with retention periods of at least 30 days and primarily using HDTV/megapixel resolution cameras and beyond.

1.3 Video Surveillance Storage Solutions

The information about video surveillance performance in this document refers to deployments that use either the E2800 or E5700 controller and DE460C disk shelves with large-capacity (8TB 7.2K NL-SAS or larger) disk drives.

2 I/O Characteristics

The video surveillance workload in many deployments is characterized as 99% write workload and 1% read workload. In these deployments, video is archived to disk either continuously or based on motion detection and is not reviewed unless there is an incident that requires analysis. The education market is one example where archives are viewed infrequently.

The write workload is typically a constant workload per volume (LUN) based on the number of cameras per server.

Read workload is based on the frequency and number of viewing stations reviewing archived video. Most video management systems implement analysis tools that enable the operator to fast-forward video. There are also features to intelligently search archived video for motion or objects in a particular area of the field of view of the camera. These search utilities might examine all archived video between two time periods or every 10th frame. Additionally, video archives from multiple cameras can be time-of-day synchronized and fast-forwarded.

This read workload might generate I/O requests at many times the rate the video was originally written to disk. Write workload is relatively easy to characterize, whereas read workload is less predictable.

The architecture and configuration of the video management system also affect the workload to the storage array. Systems that implement tiered storage schedule a copy from one volume or directory to another at a recurring interval (such as hourly or daily), and during the copy function the IOPS of the storage array might increase by a factor of eight or more. This function generates both read and write I/O.


While examining workflow and performance data, video surveillance deployments must first measure the baseline write performance and then consider the frequency that video is read or copied following the initial write.

2.1 Viewing Archived Video

The frequency and number of camera archives viewed concurrently determine the read I/O to the storage array. Video archives may be viewed at the normal speed or at an increased playback speed. The application might support increasing the playback speed over 1,000 times the recorded speed.

Figure 1 demonstrates the performance characteristics for viewing 64 HDTV cameras at normal playback speed and transitioning to 16x playback speed. For this test, writing of video was suspended so that changes in read I/O were easy to see.

Figure 1) 64-camera transition from 1x to 16x.



In this example, the average data rate increased from approximately 140Mbps (2.2Mbps per camera) to 400Mbps (6.2Mbps per camera) when transitioning from 1x to 16x playback. Note that the data rate increase is not a linear relationship; a 16x speed-up increased the I/O rate by a factor of 3x.

The performance characteristics associated with the forensic capabilities of the video management software (VMS) viewing client are implementation-specific and might vary between releases. Both video walls and client-viewing workstations affect the performance characteristics of the storage array but not the capacity required. Investigative activities by the client workstations and incident reporting files might need to be considered when sizing the storage array.

2.2 Performance of Tiered Storage

During the normal processing of video feeds from cameras and writing to the video recording storage location, the data rate is relatively constant between the server and storage array. When the archive process is initiated, the I/O characteristics change, based on the source and destination of the archive. If the administrator has configured separate volumes (seen by recording servers as LUNs) for the recording location and archive location (for example, E:\RECORDING and F:\ARCHIVE), then the volume containing the recording location incurs read I/O, and the archive location incurs write I/O. The archive function in this example changes the I/O characteristics of the recording location from primarily writes to both reads and writes for the duration of the archive process.

The duration of the archive process is a function of the amount of data that must be moved. It is important to understand that the recording server writes video to storage at approximately the rate of arrival from the networked video cameras. The archive process reads and writes data as rapidly as the recording server can read, process, and write the files to the destination. In testing, the I/O rates have been observed to increase by eight times or more during the archive process.

3 Server Planning

Physical security integrators have traditionally looked at servers as a commodity item, and deploying the lowest cost server that meets the performance requirements of the video management software is the primary design consideration. The idea behind deploying open platform systems is to allow flexibility in selecting the best component for the task. Best may be defined as least expensive while meeting the performance criteria.

A common design point for physical security integrators is to deploy relatively low-end 1RU recording servers without virtualization, as a rack-and-stack means of cost savings. This design is most advantageous when the host interface to the storage array is a relatively inexpensive 1Gbps or 10Gbps iSCSI. Autosensing can be used to adapt to different host capabilities. When dual-port Fibre Channel HBAs or dual-port direct connect SAS host interfaces are used, the cost of the HBA, or the limits to scalability with a direct connection, preclude a rack-and-stack approach. Deploying fewer high-end servers with virtualization might be a more cost-effective choice.

3.1 Hardware Recommendations

The viewing workstation and recording servers must meet the minimum hardware requirements of the video management software vendor. For example, OnSSI lists its hardware recommendations at <http://www.onssi.com/hardware-recommendations>. These are general recommendations that don't usually specify an exact model or clock rate. (For example, a recommendation might specify dual-core Intel Xeon (quad core recommended) or Intel Core i5 or better.) CPU specifications change too frequently and have far too many derivations for exhaustive testing of each model.

3.2 CPU

In validation testing, NetApp has tested a variety of CPUs. From a design standpoint, faster CPUs with more cores can support more cameras per recording server. Or alternately, the CPU utilization is higher for the same number of cameras with lower performing CPUs.

One advantage of deploying recording servers as virtual machines is the ability to take advantage of unused CPU cycles by adding recording server virtual machines to a physical machine. Although there might be additional costs associated with licenses for the hypervisor, they can be offset by more efficient use of resources.

4 Product Selection and Performance

This chapter discusses the video surveillance storage product selection and performance evaluation and provides results, recommendations, and conclusions that can be used as design parameters when planning and implementing the solution.

4.1 Overview

The primary objective of a video management recording server is to receive video feeds over an IP network from video surveillance cameras and record all or portions of this video content to disk for a given retention period. The workload for this function is primarily write I/O at a relatively constant data rate on a server-to-server basis.

The secondary objective is to allow surveillance operators to view, search, and analyze the video written to the archive of the video server. This workload is primarily read I/O at a rate that might be substantially higher than the rate at which it was originally written. This workload might be infrequent and transient and depends on the deployment model. For example, public school deployments might only view recorded video once or twice per week, whereas gaming deployments utilize scores of operators to continuously analyze activities on the casino floor.

A tertiary workload is the management of the video files by the VMS. Video files are deleted when they exceed the configured retention period or if the volume reaches the minimum free space, which is a configurable parameter. Some VMS implementations write video feeds to a temporary directory for a few minutes and then copy these files to a permanent directory. Other implementations store the first 24 hours of video in a live directory location and then move these files to one or more archive locations on a periodic, configurable basis.

The three video workloads are:

- Recording
- Viewing
- Management

Video recording is constant write I/O, viewing is transient read I/O, and management is both read and write I/O on a periodic basis.

4.2 Operational Considerations

Virus-Scanning Software

Virus-scanning software might have a negative effect on performance because of the system resources consumed, and the software might temporarily lock files during scanning. These file locks might affect performance or cause file corruption. It is recommended not to use virus scanning on recording or archiving directories of the recording servers or on the management servers in general.

User Access and Third-Party Software

Video surveillance recording servers should not have third-party software such as DVD-burning software installed because these packages might have a detrimental effect on

performance. Also, using Windows Explorer or other applications to view (open) archive files might have the same performance and file corruption issue as virus-scanning software.

Disk Full Conditions

Most VMS applications define both a retention period for video files and a software-defined maximum size of the archive location. For example, the volume (LUN) defined for video storage might be 29.3TB in usable capacity, but the maximum size defined in the application is 28TB. When the storage location reaches the 28TB mark, the application begins to delete the oldest video files regardless of the configured retention period. When the tiered storage approach is used with both OnSSI Ocularis and Milestone XProtect, the recording server attempts to move files from the initial location to the archive location prior to the scheduled archive time to free space.

Disk full conditions affect performance by adding additional workload to the archive function. These emergency archive functions occur (as observed every two minutes) outside the normal schedule.

In deployments that do not implement the archive function, some additional workload occurs as a result of the emergency file deletion, but the performance implications should be minimal.

As a best practice, accurately sizing and configuring the application to maintain adequate performance and free space provides more deterministic performance.

Database Corruption and Repair

The file structure of an archive location might become corrupted in the event of a recording server failure or ungraceful shutdown. If failover recording servers are implemented, their function is to assume the video archiving function while the primary recording server is out of service.

When the primary recording server is restored, the database structure must be repaired. The workload might change significantly during this recovery because the corrupted database files might be moved to subfolders and repaired in the background. Additionally, video files stored on the failover servers must be moved from the failover server to the primary recording server. The repair and recovery process might take 30 minutes or more, and the additional workload might alter normal system performance.


4.3 E-Series Storage Array

Video surveillance solutions can be implemented by either of the following controller options, both based on the DE460C disk enclosure using 8TB, 10TB or 12TB NL-SAS drives:

- **E5700 controller** with 32Gb Fibre Channel (FC) host interface card (HIC) and 32GB cache per system (128GB of cache per system available as an option)
- **E2800 controller** with the host systems connected to directly attached 12Gb SAS with 8GB or 16GB cache per system

The NetApp E-Series storage array is targeted at the video surveillance market through its price and performance characteristics. Figure 2 provides an overview of the number of cameras and disk shelves supported by the E2860 and E5760 storage arrays.

Figure 2) E-Series raw storage capacity.



The E5760 and E2860 storage arrays include patented mechanical engineering, providing dense, scalable, and highly reliable bandwidth and capacity. The disk controller firmware supports an optimal mix of high-bandwidth, large-block streaming and small-block random I/O.

- Controllers.** The controller for this solution is the E5760 or E2860. The E5760 is targeted at FC deployments, and the E2860 target deployment is direct SAS attachment. The solution includes dual controllers for high availability. All components of the E-Series are hot swappable; firmware upgrades can be completed while the system is operational. Both controllers have a data path to all shelves and drives in the array. Both controller models include cache memory for read and write buffering.
- Disk shelves.** The DE460C is a 4-rack unit (RU) shelf holding up to 60 3.5" 8TB, 10TB or 12TB NL-SAS drives. The E5760 configuration can support the controller shelf plus 7 expansion shelves for a total of 480 drives. The E2860 can support the controller shelf plus 2 expansion shelves for a total of 180 drives. Using 12TB drives, this gives a total raw capacity of 5,670TB and 2,160TB, respectively. This is represented in tabular format in Table 1.

Table 1) E-Series controllers and disk shelves.

Category	E5760	E2860
Form factor	4U/60 drives	4U/60 drives
Maximum disk drives	480	180
Controller shelf	1	1

Category	E5760	E2860
Maximum expansion shelves	7	2
Total maximum number of disk shelves	8	3

- **Disk drives.** Each shelf can be populated with near-line SAS (NL-SAS), SAS, or solid state drives (SSDs) of varying sizes and rotational speeds. Because of the continuous workload placed on drives used to store video surveillance archives, enterprise-class drives are required for this solution. Enterprise-class drives are designed to be vibration tolerant and rated for 24/7 duty with an optional five-year warranty.

4.4 Configurable Performance Options

The following items are performance-related recommendations common to all operating systems, hypervisors, and VMS packages. A checklist is provided at the end of this document to assist in the implementation of the recommended values.

Data Assurance

Data assurance is an optional configuration option supported on certain disk drives. It adds a checksum to every block of data written to the volume. The feature incurs a performance penalty, which might be acceptable for some applications, but is considered unnecessary for most video surveillance applications.

Read and Write Cache

Both read and write cache should be enabled. These parameters must be configured on a volume-by-volume basis. Dynamic read prefetch should be enabled. The read prefetch is more commonly called read-ahead. The prefetch or read-ahead function might increase read throughput by preloading cache with data anticipated to be requested in the future.

NetApp recommends that the value of the write cache without batteries be set to disabled. In the event of a power failure, the battery on the controller maintains power to the controller to flush the write cache to an onboard flash memory. When power is restored, the I/O in the write cache can then be completed to disk. Failure of the controller battery is logged in the event log and should be corrected as soon as is practical.

Cache Mirroring

NetApp recommends that the value of cache mirroring be set to disabled. Cache mirroring effectively decreases the available cache by 50% because I/O is mirrored on both controllers. Cache mirroring incurs a performance penalty for the mirror operation, in addition to the reduction of available cache. Because the failure of a controller incurs a loss of video recording for seconds to minutes regardless of the value of cache mirroring, NetApp advises that cache mirroring be disabled.

Cache Block Size

The common pool of cache for each controller is organized into blocks of a configurable size. The allowable sizes are 4KB, 8KB, 16KB, and 32KB. All volumes share the common pool of cache for the controller, and thus the size is constant for all volumes. All I/O in the system must pass through the cache, and the block size determines how many blocks are required to hold each I/O. If the server issues an I/O that is 12KB in size and the cache block size is configured at 16KB, two blocks are allocated, and the second block has 4K of wasted space.

Because the I/O size of VMS packages is typically greater than 256KB, NetApp recommends using a 32KB cache block size.

Cache Flushing

The E-Series manages the pool of cache based on demand and time. Cache is used for both read and write I/O. By default, the cache blocks containing write I/O are flushed at 10 seconds or more frequently if the cache gets filled. The demand parameter is a high-/low-watermark value that NetApp recommends to be at 80% initially. These values instruct the algorithm to attempt to maintain the cache utilization at 80%.

Media Scan

The media scan feature provides error detection before the condition disrupts read and write activity to the disk. NetApp recommends enabling this feature with a frequency of 30 days and recommends disabling redundancy check. This option is configured on a volume-by-volume basis. If errors are detected, the condition is recorded in the event log for storage administrator action.

Segment Size (Dynamic Disk Pools)

With DDP, segment size is not configurable as it is with a traditional volume. The segment size is 128K, and 4,096 segments are written to a disk (512MB) before writing to the next disk. See TR-4197: Video Surveillance Solutions with NetApp E-Series Storage: Planning and Design Considerations for more information about DDP.

A pool works best with a minimum of 30 drives up to the maximum the system supports. The larger the DDP, the easier it is to install and administer and the faster reconstruction occurs should a disk fail. DDP derives the most benefit from allocating all the disks in a storage array to a single pool and then creating individual volumes out of the pool.

Segment Size (Traditional Volumes)

The segment size parameter in E-Series is the amount of data written to one disk drive before moving to the next disk in the volume group. The default value is 128KB, which is suitable for most video surveillance deployments.

E-Series traditional volume groups have configurable segment sizes of 8KB, 16KB, 32KB, 64KB, 128KB, 256KB, and 512KB. SANtricity® provides a means to increase or decrease the segment size up or down by one increment at a time. For example, if the

current segment size of the volume group is 128K, the segment size can be migrated down to 64K or up to 256K.

Note: Changing the segment size takes a long time to complete and cannot be cancelled after it is started.

As an example, assume a volume group is configured for RAID 6 using 14 disks (12+2 configuration) in the volume group, and the segment size is 128KB. A full stripe write would be (12 x 128KB) = 1536KB, or 1.536MB.

4.5 E-Series Performance Checklist

Table 2 describes the storage array global parameters.

Table 2) Storage array global parameters.

Parameter	Recommended Value
Start cache flushing at (in percentage)	80% (default)
Stop cache flushing at (in percentage)	80% (default)
Cache block size (in KB)	32KB
Media scan frequency (in days)	30 days
Failover alert delay	5 minutes (default)

Table 3 describes the volume and volume group parameters.

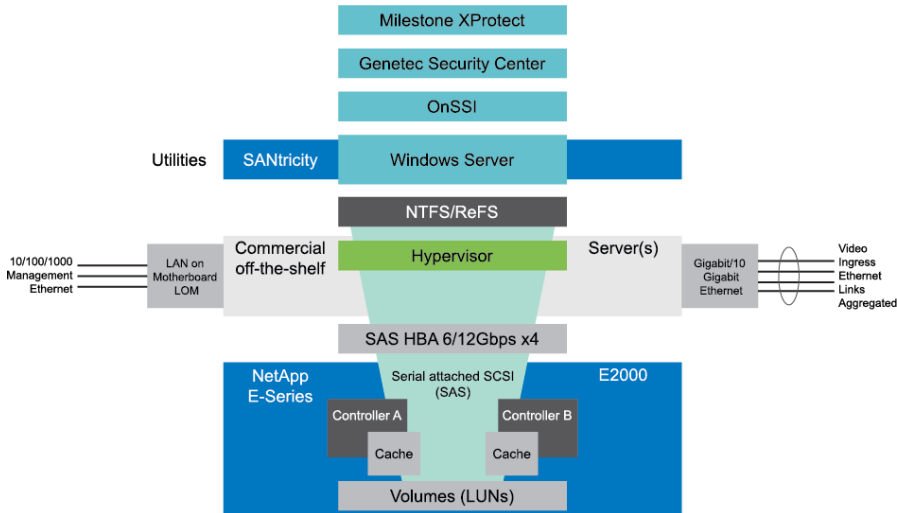
Table 3) Volume and volume-group parameters.

Parameter	Recommended Value
Data assurance (DA) enabled	No
Segment size	128KB (default)
Capacity reserved for future segment size changes	No
Maximum future segment size	Not applicable
Modification priority	Lowest
Read cache	Enabled (default)
Write cache	Enabled (default)
Write cache without batteries	Disabled (default)
Write cache with mirroring	Disabled
Flush write cache after (in seconds)	10 seconds (default)
Dynamic cache read prefetch	Enabled
Enable background media scan	Enabled
Pre-read redundancy check	Disabled

5 Example One: E-Series Storage Array E2800

An example illustrates the performance principles in this document. Here is a sample video surveillance environment using the E2800 and SAS host interfaces. (The E2800 also supports iSCSI and Fibre Channel attach, but direct-attach SAS is preferable from a cost standpoint for smaller configurations.) This example provides a two- to four-physical server configuration for up to approximately 640 network video cameras. VMware ESXi is utilized to provide up to 16 virtual machines in this example. The hardware and software components are shown graphically in Figure 3.

Figure 3) E2800 hardware and software components.



5.1 General Performance Considerations

The E2800 controller supports up to four SAS interfaces per controller, or eight per duplex controller storage array. Up to four physical servers, each with a dual-port SAS HBA, can be directly attached to each controller, providing one active and one redundant path to the array.

The E2800 controller can be ordered with either 16GB or 64GB of cache memory per controller. The theoretical maximum write performance is approximately 2GB/sec. A typical deployment scenario for this configuration entails video ingress from 0.5Gbps to 2Gbps.

5.2 Typical Data Rates for This Example

The deployments included:

- 640 Axis M3204 network video cameras
- 1280 x 720 (HDTV format) resolution
- 12 frames per second
- 30% compression
- H.264 UDP/RTP transport

Each camera generates a data rate from 0.8Mbps to 1.2Mbps or an aggregate data rate from 512Mbps to 768Mbps. This deployment has been validated with 10 recording

servers recording 64 video cameras per server. Using the Axis Design tool, this deployment requires a minimum of 270TB of storage for a 30-day retention.

Given the recording servers are running in virtual machines on four physical hosts, the SAS HBAs from each of the four hosts and four 12Gbps SAS ports on each E2800 controller, there is ample capacity in both the host interfaces and controller throughput to support the implementation.

5.3 E2800 Performance Summary for Example One

Each physical server had four virtual machines. Server 4 has four active recording servers during normal operations and would be the busiest of the four servers. Server 2 and 3 each contain a failover server that would only have four active recording servers during a failure recovery.

The video ingress network interface for each physical server is composed of a quad-port GbE adapter with links aggregated across two physical Cisco Nexus 3048 switches in a virtual port channel. This test solution has been validated to function with only one of the four member links active. The video ingress network is not expected to present a performance bottleneck.

The physical servers ran the ESXi hypervisor, which can be configured to implement a virtual machine environment that exceeds the VMS recommended hardware specifications.

The E-Series host interfaces were SAS, demonstrated through test tools and solution validation to exceed the performance requirements of the solution by a factor of three to four times, even during periods when previously recorded video moved from a recording volume to one or more archive volumes.

The Cisco Nexus 3000 Series switches have sufficient backplane and uplink capacity when properly configured to transport IP video traffic without packet loss.

The hardware and software components in this configuration met or exceeded the performance requirements of the VMS software packages tested.

6 Example Two: E-Series Storage Array E5700

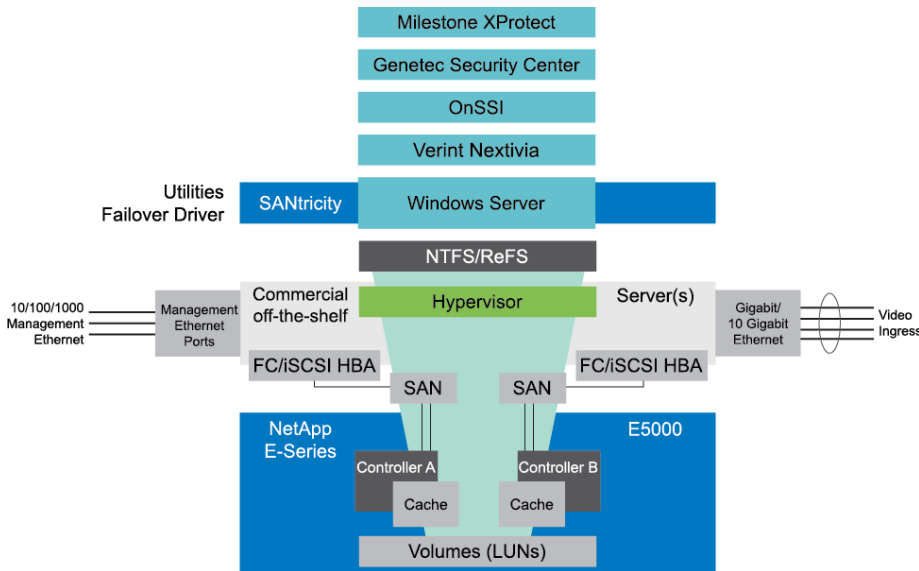
The E5700 configuration differs from the E2800 configuration in the following ways:

- The storage array is an E5760 with autosensing 32Gb FC host interfaces.
- Common off-the-shelf servers were deployed with or without a hypervisor.

Additionally, the E5760-based configuration offers substantially higher storage capacity than the solutions using the E2800, because it supports up to seven DE420C shelves with a total capacity of 480 drives. Using 12TB drives, this solution scales to 5,670TB of raw disk capacity.

The hardware and software components of the video surveillance solution using the E5700 are illustrated in Figure 4.

Figure 4) E5700 hardware and software components.



Note: Genetec Omnicast/Security Center, Verint Nextivia, OnSSI Ocularis, and Milestone XProtect video management software applications have been tested on the E-Series.

7 General Performance Considerations

The E5700 controller supports up to 8 Fibre Channel interfaces, a maximum of four per controller. Although this configuration could support four dual-attached servers, a more typical deployment would be to implement a dual-fabric SAN. The NetApp recommended configuration for a dual-fabric SAN is to connect each server to the fabric with dual-port Fibre Channel host bus adapters (HBAs), one port to each switch, and connect at least two Fibre Channel ports from each controller to separate switches. This configuration provides two active and two standby paths to the storage array. Optimally, four Fibre Channel ports may be used from each controller.

The E5700 controller can be ordered with 32GB or 128GB of cache memory per controller. The theoretical maximum write performance is approximately 9GBps.

7.1 Typical Data Rates

The deployment included:

- 640 Axis P1346 network video cameras
- 1920 x 1080 (full HDTV format) resolution
- 30 frames per second
- 30% compression
- H.264 UDP/RTP transport
- A stairway scene complexity

Each camera generates 4.9Mbps or an aggregate data rate of 3.1Gbps. This deployment is estimated to require 20 recording servers recording 32 video cameras

per server. Using the Axis Design Tool to calculate the storage requirement, this deployment requires a minimum of 983TB of storage for a 30-day retention period.

Using 16 or 32Gbps FC HBAs from each of the 20 hosts and four FC ports on each E5700 controller, there is ample capacity in both the host interfaces and controller throughput to support the implementation.

This analysis demonstrates that E-Series performance throughput is not a limiting factor for high-frame rate HDTV/megapixel deployments. Rather, the total disk capacity required to meet the video retention policy is the limiting factor.

7.2 E5700 Performance Summary

A video surveillance solution using the E5700 offers the capability for larger deployments with higher camera counts because it supports five additional DE420C disk shelves compared to an E2800-based video surveillance solution. The E5700 also has a theoretical maximum throughput almost three times higher.

Either solution provides substantially higher performance than required by the video management applications deployed in these examples.

8 Performance Validation of Tiered Storage

Both Milestone XProtect Corporate and OnSSI Ocularis ES provide the option for a tiered storage approach where video is initially written to a recording volume and then optionally written to a separate volume or directory (within the same volume) for a configured retention period. This feature enables using different RAID levels or disk types for the different storage tiers.

The performance of a typical video surveillance workload does not vary dramatically based on scheduled functions or usage based on time of day. However, when implementing tiered storage, the workload is not a constant throughout the day.

8.1 Archive Function

The archive function of the tiered storage approach can be configured to move video files from the recording volumes (LUNs) to the archive volumes on a periodic basis. The archive function can be scheduled to run every eight hours, four hours, or hourly. The duration of the archive function is determined by the number and size of video files that must be moved from tier to tier. In the validated test configurations described previously, the archive function was scheduled to initiate every hour at the top of the hour. The duration of the archive process was found to be 20 to 30 minutes typically.

Because of this configuration, the performance characteristics of the system are dramatically different between the first 30 minutes of each hour and the last 30 minutes of each hour. To contrast the performance characteristics using a load test tool that runs at a relatively constant data rate, the performance characteristics of a live deployment were examined.

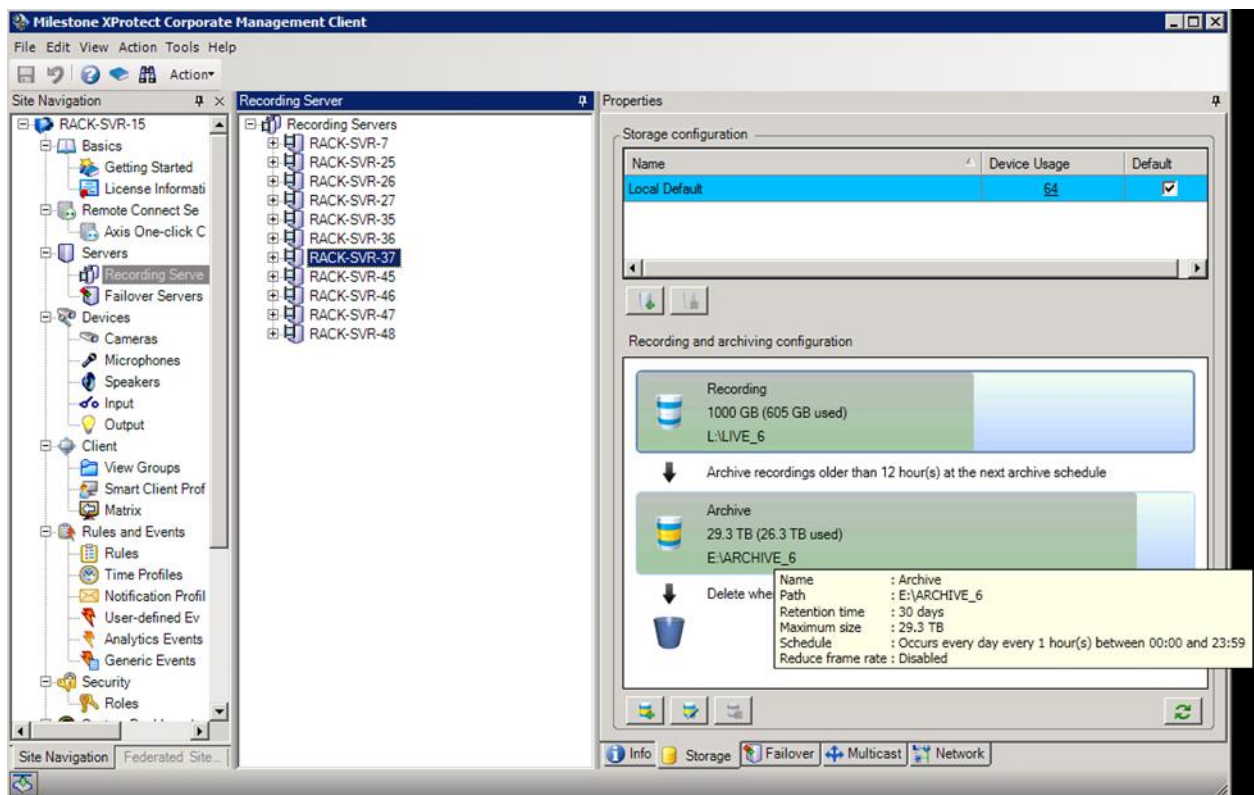
8.2 Recording Server

The characteristics of a single recording server from the test configuration RACK-SVR-37 were looked at first. This server managed 48 simulated Axis M3204 cameras configured for 1280x720p at 30 frames per second, 30% compression with RTP/UDP transport. The aggregate input video data rate for these cameras was approximately 128Mbps (16MB/sec) at 12,000 packets per second. The average data rate for each of the 48 cameras was approximately 2.6Mbps. This rate approximates the Axis design tool image scenario of Intersection Night Option (2.7Mbps).

The size of the archive volume (LUN) attached to this server was 29.3TB. At the observed data rate, the volume (LUN) would maintain approximately 22 days of video from the 48 cameras. If the goal is to meet a 30-day retention period, either the number of cameras supported by this server would need to be reduced to 35 or the archive volume would need to be increased to over 40TB.

The storage configuration of recording server RACK-SVR-37 is shown in Figure 5.

Figure 5) Storage configuration RACK-SVR-37.



The recording volume (LUN) is 1TB in size, and video recordings over 12 hours old are moved to the archive volume (LUN) hourly at the top of the hour. The archive volume is configured for 30-day retention, but as shown previously, the video files will need to be deleted after approximately 22 days at the observed data rate.

Given that the SANtricity utilities have been installed on the virtual machine, the LUN numbers of the two volumes can be determined by executing the `smdevices` command as shown:

```

C:\Program Files (x86)\StorageManager\util>smdevices
SANtricity ES Storage Manager Devices, Version 10.00.30.24
Built Tue Aug 28 04:07:49 CDT 2016
Copyright (C) 1999-2012 NetApp, Inc. All Rights Reserved.

  \\.\PHYSICALDRIVE1 [Storage Array st1E2860-33_34, Volume VOL_ARCHIVE_6, LUN 6, Volume ID
<60080e50002e3192000009d550644b55>, Prefe
rred Path (Controller-A): Owing controller - Active/Optimized]
  \\.\PHYSICALDRIVE2 [Storage Array st1E2860-33_34, Volume VOL_LIVE_6, LUN 16, Volume ID
<60080e50002e3192000009de50644e09>, Preferr
ed Path (Controller-A): Owing controller - Active/Optimized]

```

As a best practice, the volume names represent the function of the volume (LUN). It is obvious that the function of the recording (live) volume is LUN 16, and the archive volume is LUN 6.

Because the testing was done in a VMware ESXi environment, VMware performance analysis tools were used. Using the VMware vSphere client, the physical server was highlighted, and the performance tab, chart options, and storage path were selected in real time. All HBA storage was selected, and the read and write rate parameters were checked. The write rates for the active path of LUNs 6 and 16 were selected to highlight these lines in the graph. These values represent the write data rate for the two LUNs encompassing the archive function beginning at the top of the hour. This graph is shown in Figure 6.

Figure 6) Write rate during archive.



The write data rate to the recording (live) LUN 16 is a generally constant rate at a maximum write rate of 24,530KBps (196Mbps), while the archive LUN 6 is active for approximately 30 minutes and reaches a peak write rate of 63,723KBps (509Mbps).

An additional observation from reviewing the graph is that the write rate varies more during the archive function due to the increased I/O and workload on the recording server during the archive.

8.3 I/O Latency

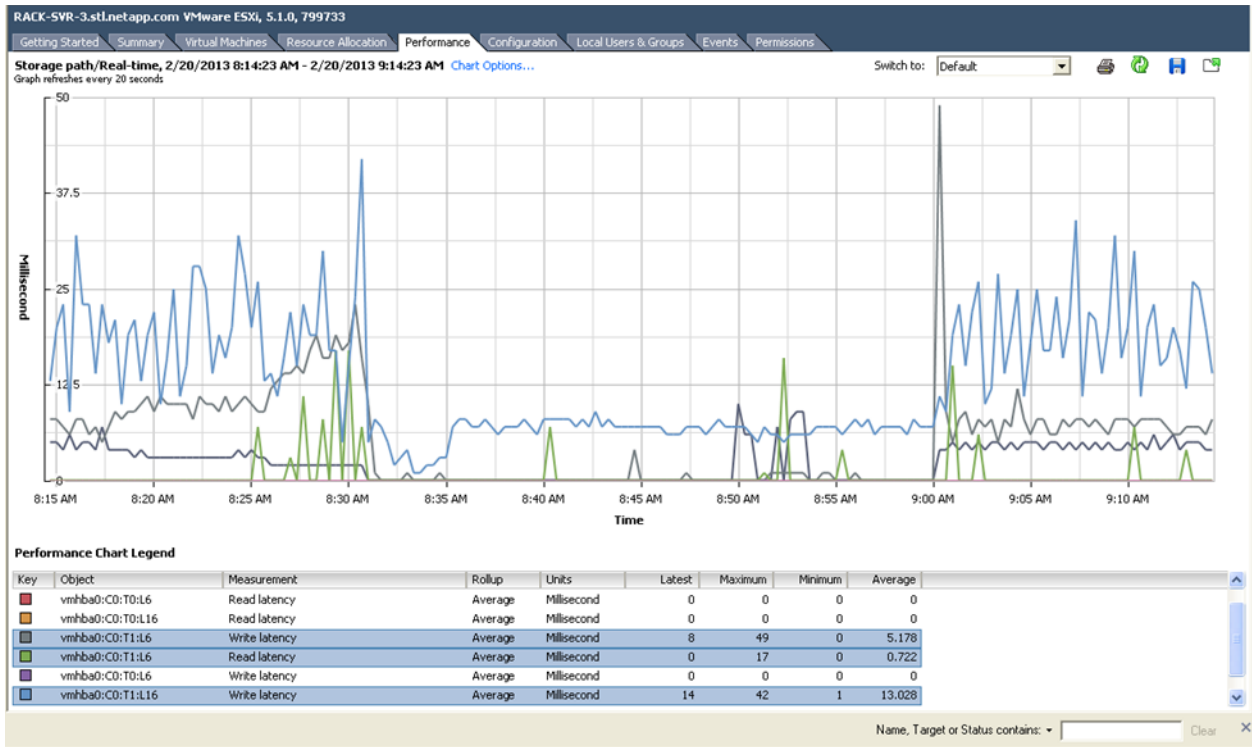
Observed latency for reads and writes to the respective volumes (LUNs) varies at different times of the day due to the changes in the workload when tiered storage is implemented

Milestone utilizes CONNEX International as a third-party testing agency for validating storage vendors with the XProtect product. The CONNEX test plan stated a goal of less than 0.1% of frame loss and write latency less than 200ms between the recording server and the storage array.

The recording server RACK-SVR-37 managed 48 simulated Axis M3204 cameras configured for 1280x720p at 30 frames per second, 30% compression with RTP/UDP transport. The average data rate for each of the 48 cameras was approximately 2.6Mbps.

The VMware vSphere client was used to monitor the real-time data for the storage path for LUNs 6 (VOL_ARCHIVE_6) and 16 (VOL_LIVE_6) for this recording server. The one-hour length of time on the x-axis of the chart includes the entire bottom half of the hour, when only recording and viewing are active, as well as portions of the top half of two hours, when the archive function is active. This is shown in Figure 7.

Figure 7) I/O latency RACK-SVR-37.



One key concept derived from this chart is that average latency does not provide a useful metric when the range includes both the top half and bottom half of the hour. The workloads were very different during those two time periods, and therefore the average latency value was skewed.

However, the chart shows the maximum latency. For both LUNs 6 and 16, the reported write latency value was less than 50ms. During this interval, no buffer overflows representing frame loss were reported in the XProtect manager system log. Given this validation, it can be seen that the E-Series latency performance is well within the partner specifications for latency under a real-time work load.

9 Other Performance Considerations

The workload changes when the archive function is active compared to the normal recording and playback workload. This archive function is specific to OnSSI Ocularis and Milestone XProtect because they share the same recording server code base.

Other video management applications such as Verint Nextivia and Genetec Omnicast do not implement a multistage storage architecture. Video files are written to a volume (LUN) and then subsequently deleted at the expiration of the retention period or if the defined storage reaches a full condition.

Although the tiered storage design is commonly deployed, it is not a requirement for OnSSI Ocularis ES and Milestone XProtect Corporate.

The [Milestone XProtect System Migration Guide: Migration from XProtect Enterprise to XProtect Corporate](#) states the following:

“Basically, archiving is not necessarily a must when using XProtect Corporate. In case the hard disks you have allocated for the live database are fast enough and able to contain the expected amount of data, the system can run without archiving. This is possible due to the automatic 1 hour segment division of the live database, which keeps a potential database repair after a failure as short as possible, as only the last (hour) segment of the database needs to be repaired.”

The advantages of not implementing a separate live database and one or more archive locations are more efficient use of the storage array and simplicity in deployment. A video surveillance configuration using E2860 was tested successfully with 10,000 RPM SAS drives for the recording volume (LUN) and high-capacity 7,200 RPM NL-SAS drives for the archive volumes (LUNs). (It's becoming increasingly common to deploy solid state drives [SSDs] in place of high-performance SAS drives.)

Another video surveillance configuration using E5760 was tested and performed successfully using high-capacity 7200 RPM NL-SAS drives only, with the recording and archive location on the same volume (LUN).

Using solid state drives (SSDs) or 10,000 RPM SAS drives for the recording volume (LUN) might be desired when deploying the solution in the gaming market, where there is a high degree of forensic analysis. The physical security integrator might, however, choose to implement the tiered storage approach to take advantage of the ability to implement the additional features of archiving: digital signing, encryption, and grooming of video. Grooming is the reduction in data rate, through reductions in frame rate, compression, and other parameters, depending on the VMS features.

10 Summary

Video surveillance solutions based on NetApp E-Series storage offer the physical security integrator 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.

Definitions

Table 4 contains the glossary of terms used throughout this document.

Table 4) Glossary.

Term	Definition
Controller	The controller comprises 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.
FC host bus adapter	FC HBA is a Fibre Channel adapter on the host machine that acts as an initiator in a SAN environment to provide connectivity between storage system LUNs and the host operating system. Each HBA has a unique worldwide name (WWN), which is similar to an Ethernet MAC address.
H.264	A video codec—ISO/IEC Moving Picture Experts Group (Part 10)—and more efficient than MPEG-4 and commonly used by HDTV and megapixel cameras.
HDTV	High-definition TV defines resolutions of 1920 x 1080 and 1280 x 720 pixels along with other criteria, including aspect ratio.
HICs	Host interface cards. Typically seen on E5700 and E2800 controllers. There are different types of HICs available, including SAS, FC, InfiniBand, and iSCSI.
Host bus adapter (HBA)	The host bus adapter is usually a separate card, for example, PCI-express, that is installed in the server to allow communication with the storage system.
IP video surveillance camera	A digital video camera or network video camera using a small form factor IP networked Linux host that encodes and transports video over an IP network.
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 will 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> .
MPEG-4	A video codec—ISO/IEC Moving Picture Experts Group (Part 2)—and the predecessor to H.264.
NL-SAS	Near-line-SAS, or NL-SAS, drives are enterprise SATA drives with a SAS interface, head, media, and rotational speed of traditional enterprise-class SATA drives with the fully capable SAS interface typical for classic SAS drives.
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.

Term	Definition
SAS	Serial attached SCSI (SAS) is a computer bus that is used to move data to and from computer storage devices such as disk drives and tape drives. SAS depends on a point-to-point serial protocol that replaces the parallel SCSI bus technology.
SAS host bus adapter	SAS HBA is a channel adapter on the host machine that directly attaches to the SAS host interface on the controller. SAS HBAs provide a low-cost, high-performance alternative to FC HBAs, but are limited in scalability due to cable lengths and direct attachment to the storage array.
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.
Viewing station	A high-end workstation for displaying live or archived camera feeds on a locally attached monitor.
VMS server	A video management system server is also referred to as a <i>network DVR server</i> . It manages IP camera video feeds and storage media.
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 document:

- 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-4196: Video Surveillance Solutions with NetApp E-Series Storage: Introduction to Video Surveillance
<http://www.netapp.com/us/media/tr-4196.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-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>
- Axis Communications
<http://www.axis.com>
- Optimizing NTFS
<http://technet.microsoft.com/en-us/library/cc767961.aspx>
- Overview of Disk Management
<http://technet.microsoft.com/en-us/library/cc754936.aspx>

Version History

Version	Date	Document Version History
Version 1.0	July 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 drive 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.