



Technical Report

# **NetApp E-Series for video surveillance best practices guide**

## **Deploy and configure E-Series for video surveillance**

Mitch Blackburn, NetApp  
October 2021 | TR-4825

### **Abstract**

Video surveillance solutions using NetApp® E-Series storage offer a highly scalable repository for video recording. This guide offers best practices for deploying E-Series arrays into video surveillance environments.

## TABLE OF CONTENTS

<b>Introduction .....</b>	<b>4</b>
Publication scope .....	4
Audience .....	4
Targeted deployments .....	4
Solution Benefits.....	4
<b>Reference architecture .....</b>	<b>5</b>
<b>VSS storage sizing and selection guide.....</b>	<b>7</b>
System requirements.....	7
NetApp E-Series storage sizing.....	7
NetApp E-Series virtualization sizing.....	13
<b>E-Series storage considerations .....</b>	<b>13</b>
Workload .....	14
Dynamic Disk Pools feature.....	14
RAID levels.....	14
I/O characteristics .....	15
High availability.....	15
<b>NetApp E-Series storage provisioning .....</b>	<b>16</b>
<b>Server considerations .....</b>	<b>26</b>
Multipath I/O device-specific module installation .....	27
E-Series LUN discovery and preparation .....	28
<b>Networking considerations.....</b>	<b>28</b>
Traffic management.....	29
<b>Where to find additional information .....</b>	<b>30</b>
<b>Version history.....</b>	<b>30</b>

## LIST OF TABLES

Table 1) VSS components.....	6
Table 2) NetApp VSS sizing tool parameters. ....	8
Table 3) DDP pool usable capacity versus number of drives. ....	12
Table 4) Supported controller and expansion enclosure combinations. ....	12
Table 5) Virtual machine common five-drive RAID 6 SSD usable capacities. ....	13
Table 6) DDP pool usable capacity for common virtualization.....	13
Table 7) Recording server characteristics that NetApp recommends.....	27

## LIST OF FIGURES

Figure 1) Video surveillance architecture with VMS and NetApp E-Series arrays.....	6
Figure 2) Parking lot camera group addition.....	9
Figure 3) Storage estimate for parking lot camera group. ....	10
Figure 4) Storage estimate for two camera groups. ....	10
Figure 5) Total storage estimate.....	11
Figure 6) SANtricity System Manager home screen.....	17
Figure 7) Pools & Volume Groups tile. ....	18
Figure 8) Create a DDP pool. ....	18
Figure 9) Name the pool and select drives for the pool. ....	19
Figure 10) Create a volume.....	20
Figure 11) Create a volume (continued). ....	21
Figure 12) Example of final volume configuration.....	22
Figure 13) iSCSI settings.....	23
Figure 14) iSCSI settings (continued).....	23
Figure 15) Create a host.....	24
Figure 16) Create a host (continued).....	24
Figure 17) Host definitions.....	25
Figure 18) Assign volumes to the host. ....	25
Figure 19) Select volumes to assign. ....	26
Figure 20) Host with assigned volumes.....	26
Figure 21) E-Series SANtricity Storage Manager installation. ....	27
Figure 22) Back view of the E2800 controller. ....	29
Figure 23) Back view of the E5700 controller. ....	29

# 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 replay helps with forensic analysis or with investigation of people or events that were 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 with retention periods in the range of months to years.

## Publication scope

This document provides architecture and deployment guidelines for video surveillance solutions to those who sell, design, or implement such solutions based on NetApp E-Series storage. 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 video and archive video from recording servers. This document identifies the major components and features of a video surveillance system.

## Audience

This publication is intended for IT professionals who are responsible for integrating NetApp E-Series storage systems into existing video surveillance deployments or who design and implement new video surveillance 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 use this information, combined 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 cameras or more) video surveillance installations with significant storage capacity requirements due to retention periods of at least 30 days or the use of HD/megapixel-resolution cameras.

## Solution Benefits

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

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

- **Easy management and monitoring.** The included NetApp SANtricity® System Manager software provides a graphical representation of the E-Series storage, with an easy-to-use interface.
- **Easy provisioning.** SANtricity System Manager software performs all management tasks for the array without taking the array offline.
- **High availability.** Dual controllers enable nondisruptive controller firmware upgrades, host multipath support, and dual paths to expansion shelves.
- **High performance.** E-Series controllers offer an excellent price-to-performance ratio.
- **Scalability.** With just a couple mouse clicks, you can easily add or expand capacity when you add more drives to your system. NetApp E5700 hybrid flash systems support up to 5.7PB of raw capacity (using 12TB drives) in an efficient footprint. The entry-level E2800 systems support up to 2.1PB of raw capacity (using 12TB drives).

- **Drive health monitoring.** E-Series systems provide proactive monitoring, background repair, and extensive diagnostic features for drives.
- **T10 Standard data integrity and media scan.** This scan detects and corrects data integrity issues that the recording server receives or that are caused by hardware failures on the drives.
- **Data protection.** E-Series systems support RAID levels 0, 1, 10, 3, 5, and 6 for volume groups and for Dynamic Disk Pools (DDP).
- **Certified interoperability.** E-Series systems are certified to be interoperable with multiple video management software (VMS) providers.

## Reference architecture

A typical video surveillance solution (VSS) with NetApp E-Series storage arrays consists of the following main components:

- IP cameras
- Camera network
- Management server or servers
- Recording server or servers
- Failover server or servers (optional)
- Viewing client or clients
- Storage network
- NetApp E-Series storage array or arrays

Figure 1 shows these components in a standard configuration.

**Figure 1) Video surveillance architecture with VMS and NetApp E-Series arrays.**

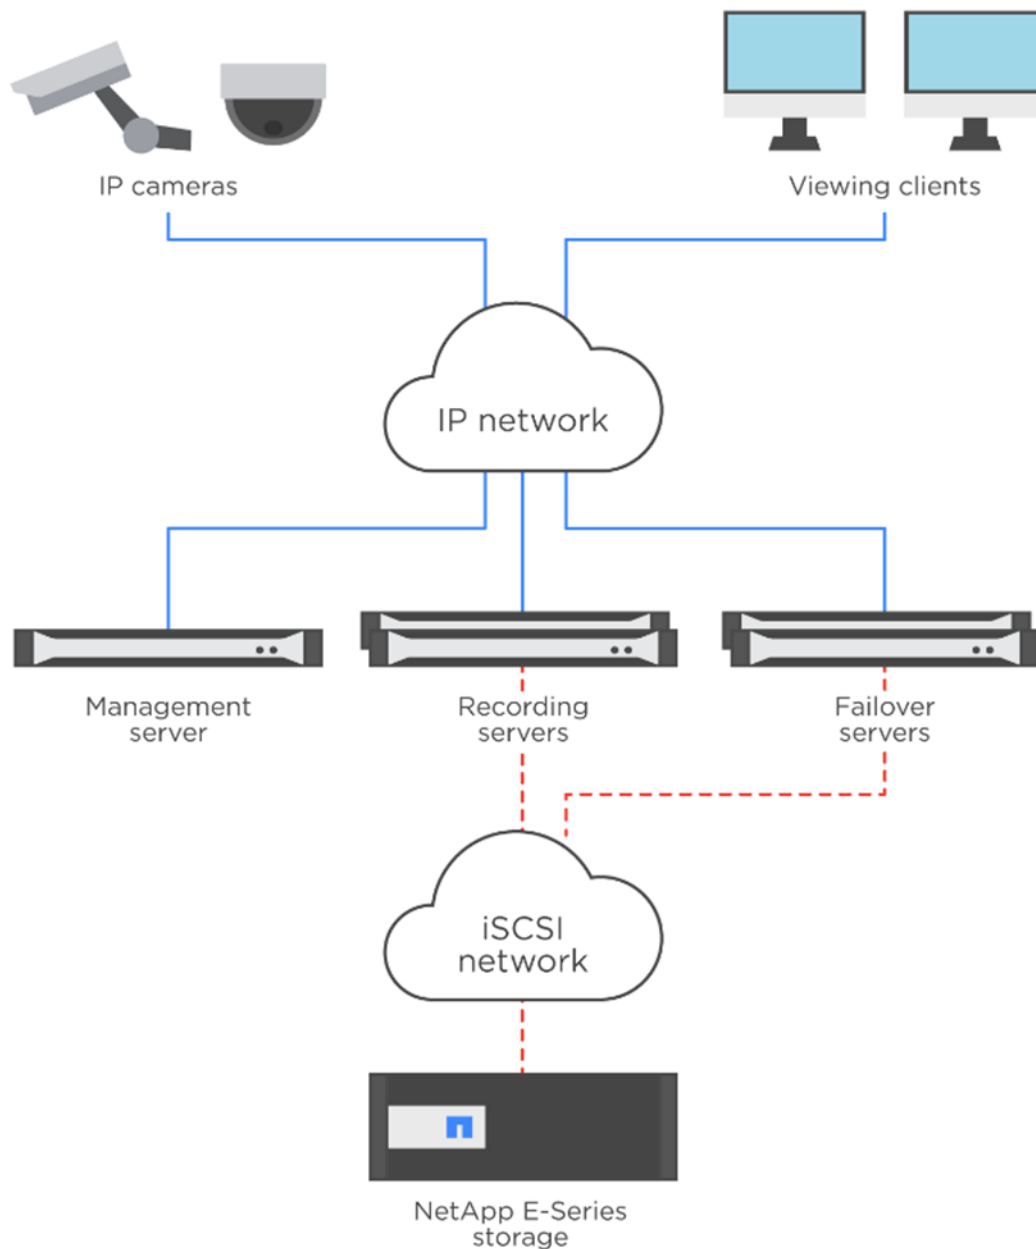


Table 1 provides a brief explanation of each of these components.

**Table 1) VSS components.**

Component	Description
IP cameras	Provide audio and video streams for live viewing and for recording for later playback by using the VMS.
IP network	Enables communication among camera streams and various VMS components.

Component	Description
Management server	Acts as the central point for configuring all the VMS components.
Recording servers	Are responsible for recording and playback of audio and video streams as per the configuration that the management server dictates.
Failover servers	Provide redundancy against recording server failures (optional).
Viewing clients	Provide live viewing of camera streams or playback of recorded audio and video by the recording servers.
Storage network	Enables communication among the recording and failover servers and the NetApp E-Series storage arrays. This network is typically iSCSI, but FC and SAS protocols are also supported.
NetApp E-Series storage arrays	Provide highly available storage for video and audio files from the recording and failover servers.

**Note:** For detailed information about all the supported configurations and additional components, refer to the specific VMS provider that you are using.

The following sections provide general guidelines on how to size and tune VSS components, especially the NetApp E-Series storage arrays, for optimal results. You should use these guidelines along with the VMS provider recommendations.

## VSS storage sizing and selection guide

This section highlights various aspects of the VSS that affect its storage needs, along with example estimation and model selection.

### System requirements

The system requirements are specified in a request for proposal or a quote that is developed by either the end customer or a physical-security integrator who is in contract with the end customer.

The physical-security integrator must work with the physical-security manager to accurately assess specific requirements, including:

- Number, location, and type of cameras; resolution; frame rate; and so on
- Number of recording servers
- Number of cameras per recording server
- Virtualization requirements
- Video management software (VMS) type
- Continuous recording or record on motion
- Retention period and archiving requirements
- Failover design requirements

**Note:** To verify these requirements, the VMS provider's sizer should be referenced.

### NetApp E-Series storage sizing

Each recording server has specific storage needs that are based on the number and type of cameras and the recording parameters that it must support. The following subsections provide guidelines on how to estimate storage for one recording server, and you should follow them for all the recording servers in a VSS.

## Storage estimation for live and archived recording

NetApp provides an online [VSS storage capacity sizing tool](#) that you can use along with the VMS provider's storage sizing tool.

Table 2 lists the parameters that you need for the NetApp VSS sizing tool.

**Table 2) NetApp VSS sizing tool parameters.**

Parameter	Description
Camera group/number of cameras	VSS site divided into camera groups, with the number of cameras installed in each group; for example, parking lot, office space.
Camera resolution	Average camera resolution of the camera group; for example, 1280x720, 1920x1080.
Compression type	Average compression type of the camera group; for example, H.264-30, H.265-20.
Frames per second (FPS)	Average FPS of the camera group; for example, 15, 30
Motion activity	The likelihood that motion will occur in the field of view of the camera group. For example, low (stairway, emergency exit), high (intersection, subway).
Bit rate	Average bit rate of the camera group in Kbps (if known). This parameter supersedes the resolution, compression type, FPS, and motion activity parameters.
Recorded hours per day	Average hours per day that the camera group will be recording; for example, 12, 24.
Days to retain	Maximum number of days that the video will be saved for the camera group, either under live recording, archived recording, or both.

Figure 2 and Figure 3 show an example of a storage estimate for a parking lot camera group with the following parameters:

- Number of cameras: 25
- Resolution: 1920x1080
- Compression type: H.264-30
- FPS: 20
- Motion activity: below average
- Recorded hours per day: 24
- Days to retain: 30



Figure 2) Parking lot camera group addition.

Add Camera Group

Group

Name ?

Parking Lot

Number of Cameras ?

-

25

+

Camera Specifications

Characteristics

Bitrate (Kbps)

Resolution ?

1920 x 1080 (Full HD)

Compression Type ?

H.264-30 (Average Quality)

Frames per Second ?

-

20

+

Motion Activity ?

40% - Below Average

Retention Settings

Recorded Hours per Day ?

-

24

+

Days to Retain ?

-

30

+

Figure 3) Storage estimate for parking lot camera group.

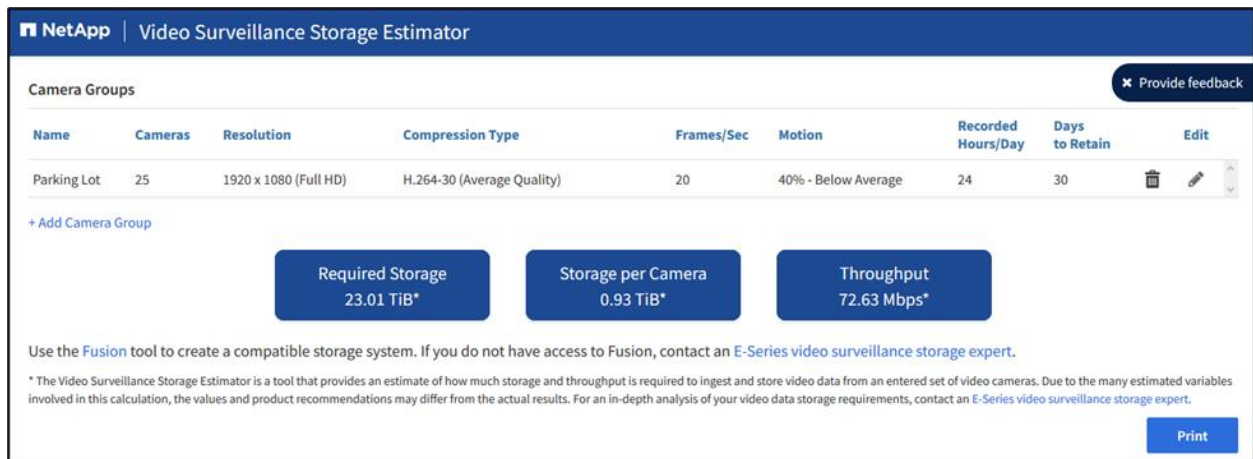


Figure 4 shows a storage estimation increase when a second camera group, called Office Space, is added, with the following parameters:

- Number of cameras: 50
- Bit rate: 2000Kbps
- Recorded hours per day: 24
- Days to retain: 30

Figure 4) Storage estimate for two camera groups.

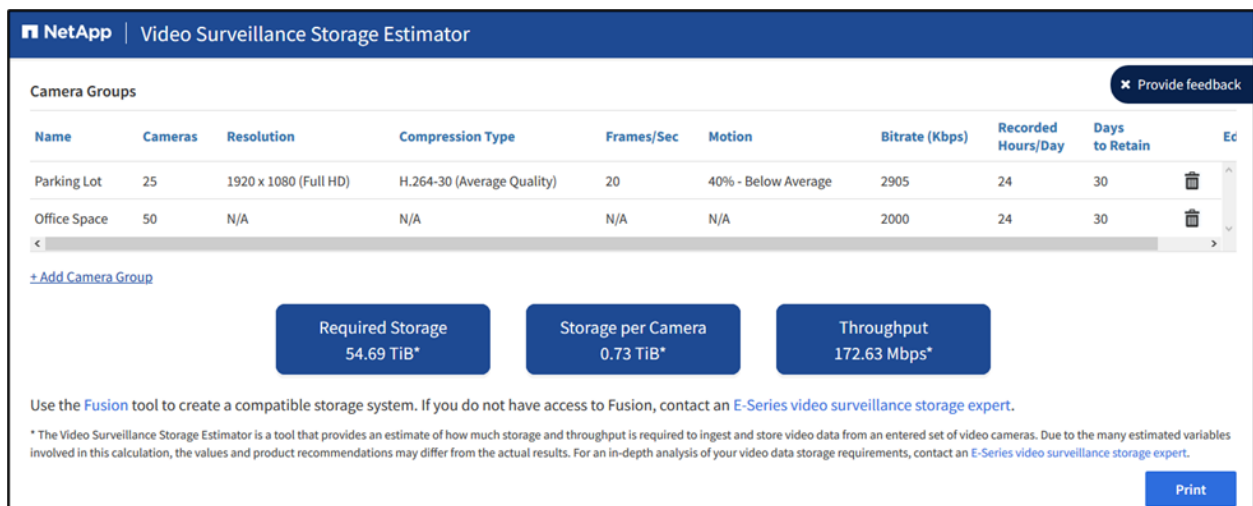
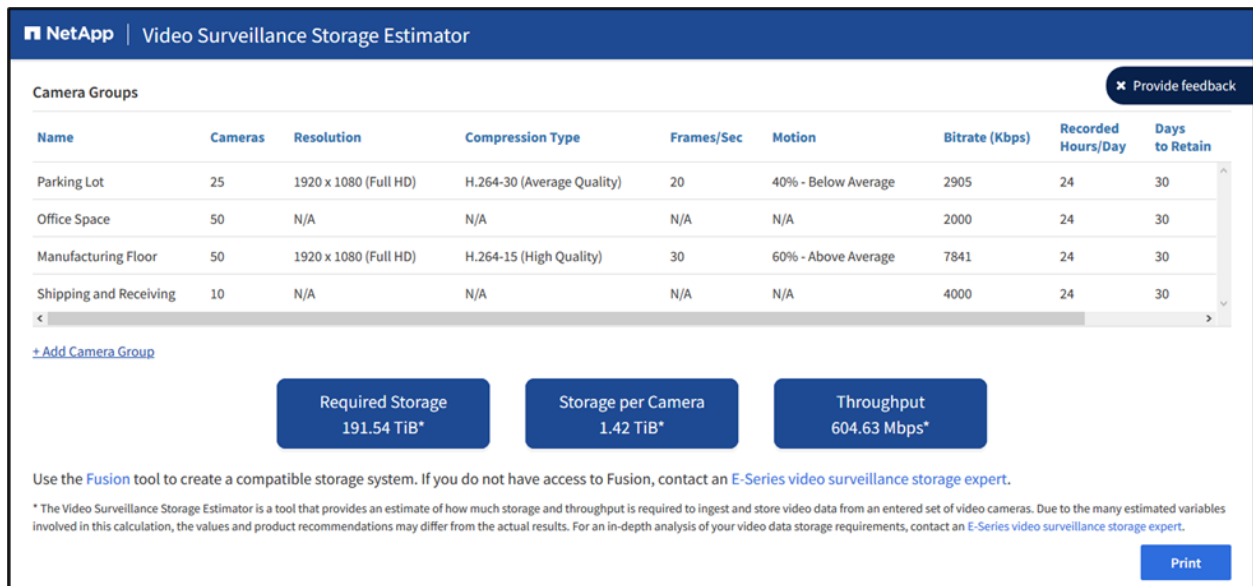


Figure 5 shows additional camera groups and the final storage estimate for this example.

**Figure 5) Total storage estimate.**



**Note:** In addition to required storage, the throughput estimate also influences the optimal storage model selection for a VSS. For more details, see the NetApp E-Series Storage Selection section.

## Storage estimation for reserved capacity

When you define the retention period policy, you must consider the amount of reserved capacity. As a best practice to prevent the “disk full” error condition, each volume that is attached to a recording server should be maintained at approximately 80% utilization.

By using the storage estimate example from Figure 5, following is the reserved capacity estimate:

- Required storage: 192TiB
- Reserved storage:  $(192 \times 0.2) = 38\text{TiB}$

## Storage estimation for the failover server volume

NetApp recommends that you size the volume for each failover server to retain a minimum of 3 to 5 days of video from the recording server. To calculate the size of the failover server volume, divide the capacity of the recording server volume by the site retention period and then multiply by the number of days to retain during failover.

By using the storage estimate example from Figure 5, following is the failover server volume capacity estimate:

- Required storage: 192TiB
- Reserved storage:  $[(192 / 30) \times 5] = 32\text{TiB}$

## Sizing for archiving

The archiving process copies data from one location to another. Therefore, for every initial write, you must read and then write the data again. Also, archiving typically runs on a schedule, so it must be able to archive a day’s worth of data in less than a day. Many users choose to reduce the frame rate for the video that they archive, which saves some percentage of the initial size of the video. Depending on the format, it might be linear savings (MJPEG) or marginal savings (H.264). This approach reduces the amount of

bandwidth that is required to archive on the target volume and reduces the amount of storage capacity that is needed to reserve and to store the video.

Generally, you want to ensure that your archiving runs at no less than 1.5x the rate at which the initial recording is providing new data. For example, if your system ingests 100MBps of video data, it should be able to archive at 150MBps or more. This rate provides a margin of safety so that your video can be successfully archived before the drive capacity runs out. If archiving never completes in a timely manner, the VMS is eventually forced to delete data before the configured data expiration. To avoid drive contention, NetApp also recommends that you archive to a different volume on a different DDP pool instead of archiving to the live recording volume.

## NetApp E-Series storage selection

NetApp offers various E-Series storage models that can match typical VSS price, performance, and capacity needs. Following are some of the details of these models:

- E2812, E2860, and E5760 are dual-controller arrays.
- 4TB, 8TB, 10TB, and 12TB NL-SAS HDDs are supported.
- Traditional RAID volume groups and (RAID 6–based) DDP pools are supported

**Note:** The “E-Series Storage Considerations” section provides the pros and cons of using DDP pools over traditional RAID volume groups, as well as the optimal volume layout for recording and failover servers.

Table 3 provides a reference for estimating the number of different capacity drives that are required with DDP to satisfy specific usable capacity needs. You can use this table to extrapolate the additional number of drives that you might need.

**Table 3) DDP pool usable capacity versus number of drives.**

Number of Drives	DDP (4TB)	DDP (8TB)	DDP (10TB w/Encryption)	DDP (12TB)	DDP (18TB)
12	28.62TiB	56.31TiB	70.16TiB	84.2TiB	141.12TiB
24	62.95TiB	123.89TiB	154.34TiB	185.24TiB	282.25TiB
30	80.13TiB	157.67TiB	196.44TiB	235.76TiB	359.23TiB
48	130.08TiB	255.97TiB	318.9TiB	382.74TiB	583.2TiB
60	164.77TiB	324.23TiB	403.94TiB	484.8TiB	738.72TiB
90	248.6TiB	489.19TiB	609.45TiB	731.46TiB	1.09PiB
120	335.32TiB	659.83TiB	822.05TiB	986.62TiB	1.47PiB
180	502.98TiB	989.75TiB	1.2PiB	1.45PiB	2.2PiB

Table 4 provides a reference to determine the supported controller and expansion enclosure combinations that can satisfy specific usable capacity needs in conjunction with Table 3.

**Table 4) Supported controller and expansion enclosure combinations.**

	E2812	E2860	E5760
Form factor	2U/12 drives	4U/60 drives	4U/60 drives
Maximum drives	180	180	480
Controller shelf	1	1	1
Maximum expansion shelves	3	2	7
Total (maximum) number of drive shelves	4	3	8

**Note:** For a comprehensive list of storage options, go to [Fusion](#) (the NetApp storage sizing tool) or contact an [E-Series video surveillance storage expert](#). The Fusion tool considers the throughput requirements of the solution in addition to capacity needs.

## NetApp E-Series virtualization sizing

Virtualization is useful in deployments where resource utilization can be pooled and shared. It enables your system to provision resources faster and requires fewer physical servers to manage, all in a simplified interface. Your video surveillance deployment can use virtualization to scale compute, memory, and storage to meet the demands of many types of expansions. For example, you can increase retention, increase frame rate, and add cameras seamlessly.

NetApp E-Series systems are an excellent choice for the storage behind these virtual environments. E-Series systems give you the option for hybrid storage. With a hybrid approach, virtual machines can use high-performing SSDs and also satisfy VSS requirements with inexpensive, larger-capacity NL-SAS HDDs, all within the same storage system.

This section considers sizing requirements for virtualized video surveillance deployments that differ from bare-metal deployments. [TR-4818](#) provides a comprehensive deployment guide for virtualized video surveillance solutions.

## Storage estimation for virtualization

Use flash media to meet the latency requirements of common virtual infrastructures such as the ESXi virtual machine and the VMS application and OS. Because virtual machines are critical for the health of your overall VSS deployment, NetApp recommends a RAID 6 dual-drive parity. NetApp E-Series storage requires a minimum of five drives when you deploy RAID 6.

**Table 5) Virtual machine common five-drive RAID 6 SSD usable capacities.**

5x 800GB SSD	5x 1.6TB SSD
2.17TiB	4.35TiB

Virtual machine storage reduces the total usable capacities for VSS recording by a minimum of five drives, depending on the capacity that is required for the virtual machines. Table 6 lists usable VSS capacities for common deployments; these capacities include the reduction for common virtualization required storage.

**Table 6) DDP pool usable capacity for common virtualization.**

Number of Drives	DDP (4TB)	DDP (8TB)	DDP (10TB w/Encryption)	DDP (12TB)	DDP (18TB)
55	150.32TiB	295.79TiB	368.50TiB	442.28TiB	673.92 TiB
85	234.14TiB	460.75TiB	574.02TiB	688.93TiB	1.03 PiB
115	320.86TiB	631.39TiB	786.62TiB	944.09TiB	1.4 PiB
175	488.52TiB	961.31TiB	1.17PiB	1.40PiB	2.14 PiB

## E-Series storage considerations

Each video recording server requires one or more volumes (LUNs) to be defined to the OS to archive video files. To configure the E-Series storage array, use NetApp SANtricity System Manager to allocate individual hard drives to a disk pool or to a volume group.

The minimum number of drives in a disk pool is 11, but the minimum number of drives for a volume group depends on the RAID level. The maximum number of drives for RAID 5 or RAID 6 is 30. The limit for a disk pool is the total population of drives in the array.

During the volume group definition step, you select the RAID level for all drives that are assigned to the volume group. The supported levels of RAID 0, 1, 10, 3, 5, and 6 are for traditional volume groups, and DDP uses RAID 6 stripes that are allocated over 10 of the drives in the pool.

Individual volumes (LUNs) are created and mapped to a host after you define the disk pool or volume group. The performance and sizing requirements of your video recording server and your application software determine the number of drives per volume group or per pool and the number of volumes per group or per pool.

## **Workload**

The performance of storage systems is characterized by I/O operations per second (IOPS) and by throughput in megabytes per second. Network performance is measured in packets per second and the throughput is measured in megabits per second.

When the storage array is used for small random I/O operations from multiple applications, it is important to optimize IOPS. Also, in that use case, network packet-per-second performance is usually measured in small (64-byte) packets.

Video surveillance deployments, however, are more concerned with throughput performance than with IOPS. Network video cameras generate large IP packets to the recording servers and write relatively large records to the storage array. Video ingress to the recording servers is over an IP network, and the data rate is typically calculated in megabits per second (Mbps) for IP networks. Therefore, many of the tables in this document list Mbps rather than megabytes per second (MBps).

## **Dynamic Disk Pools feature**

To maintain a consistent level of performance even in the event of drive failure and reconstruction, the Dynamic Disk Pools (DDP) feature is available on E-Series systems. The DDP feature minimizes the performance drop during rebuild, and the rebuild completes more quickly than with a traditional RAID rebuild. Because of the shorter rebuild time with DDP, your exposure to data loss from several drive failures is minimized.

With DDP, you can also add capacity incrementally without having to create new volume groups. You can define a single pool that includes all the drives in your system, or you can define multiple pools for your system. A typical deployment has 30 to 60 drives per pool. Because of these factors, NetApp recommends DDP as the optimal choice for video surveillance.

## **RAID levels**

Although DDP is recommended for video surveillance deployments, RAID 5, RAID 6, or RAID 10 are commonly deployed in the industry. The Nevada Gaming Commission standards, for example, specify that the storage array must not lose data if a single component fails. Although RAID 6 and DDP provide better fault tolerance because they can tolerate two simultaneous drive failures, RAID 5 is often deployed instead because it costs less and still adheres to the standards. Also, if a VSS deployment uses virtualization and the virtual machine storage resides externally, it likely needs a RAID deployment for an all-flash volume group.

RAID 10 is typically used for optimal read performance when it is combined with SSDs. RAID 5 or RAID 6 are used for optimal write performance. On E-Series systems, you implement RAID 10 by selecting RAID 1 with four or more drives.

Some VMS vendors recommend a combination of RAID 10 and RAID 5 in gaming deployments, where a high volume of forensic analysis occurs, and during the most recent minutes or hours of video archives.

These designs use RAID 10 for the most recent archive, then, with the tiered storage feature, move video to a RAID 5 volume group for the duration of the retention period.

This design consideration might not be required in environments that have infrequent forensic analysis or where the performance level is such that the RAID 5 or RAID 6 volume group provides acceptable read performance. The education market is one vertical in which archived video is reviewed only if an incident (for example, vandalism or an altercation between students) warrants analysis of the video.

**Note:** When you work with RAID 5 and RAID 10 volume groups, which can handle only one drive failure, you should consider allocating hot spares to reduce your risk of data loss if an additional drive fails.

## I/O characteristics

In many deployments, the video surveillance workload is characterized as exceeding a 90% write workload. In these deployments, video is archived to drive either continuously or based on motion detection and is not reviewed unless an incident occurs that requires analysis. The education market is one example in which archives are viewed infrequently.

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

The read workload is based on the frequency and the number of viewing stations that review archived video. Most video management systems implement analysis tools that enable the operator to fast-forward video. They also include features to intelligently search archived video for motion or for 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. Also, 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 that the video was originally written to a drive. Write workload is relatively easy to characterize, but 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). During the copy function, the IOPS of the storage array might increase by a factor of 8 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 after the initial write.

As an additional note, two types of SCSI offload are available: Windows offloaded data transfer (ODX) and VMware XCOPY, which are both fully supported by NetApp E-Series controllers. If the OS and VMS application both support these SCSI offload operations, then performance for activities such as archiving between two volumes on the E-Series controller is increased over traditional read and write requests.

## High availability

Real-time applications such as video provide a challenge for physical-security integrators. Any outage or failure between a network video camera and the storage system means that the record of events is lost and cannot be recovered. Implementation of high availability for video surveillance begins with considering the camera placement, the network infrastructure, server and VMS redundancy, and the storage array.

For critical areas, to maintain coverage if a single camera or access layer switch fails, you should implement multiple cameras with overlapping fields of view. Multiple cameras that cover the critical area must be connected to separate access layer switches with redundant uplinks to the core or distribution layer switches. The IP network must implement high-availability network design principles, rapid

convergence from link and switch failures, deterministic traffic recovery, and sufficient capacity to adequately service traffic during failures.

VMS features that use local storage in the network video camera, failover recording servers, and a redundant management server protect the availability of the video archives. Hypervisors such as VMware ESXi have native support for link aggregation. For nonvirtual deployments, the Microsoft Failover Cluster Virtual Adapter for Windows Server 2019 supports link aggregation.

The failover drivers are at the center of providing path failure recovery between the server and the storage array. In general, failover drivers implement the following functions:

- Identify redundant I/O paths.
- Reroute I/O to an alternate controller when the controller or data path fails.
- Check the state of the paths to a controller.
- Provide the status of the controller or bus.

For Windows, the failover drivers are a combination of Microsoft Multipath I/O (MPIO) and the NetApp SANtricity Storage Manager host installation device-specific module (DSM). E-Series systems support the native multipath feature of VMware ESXi. For more information, see the “Multipath I/O Device-Specific Module Installation” section.

In addition to the reliability that pools and volume groups provide, E-Series arrays add an extra level of high availability with their hardware. With redundant controllers, power supplies, ports, and fans, E-Series arrays are built to withstand failure in the rare case that a failure occurs.

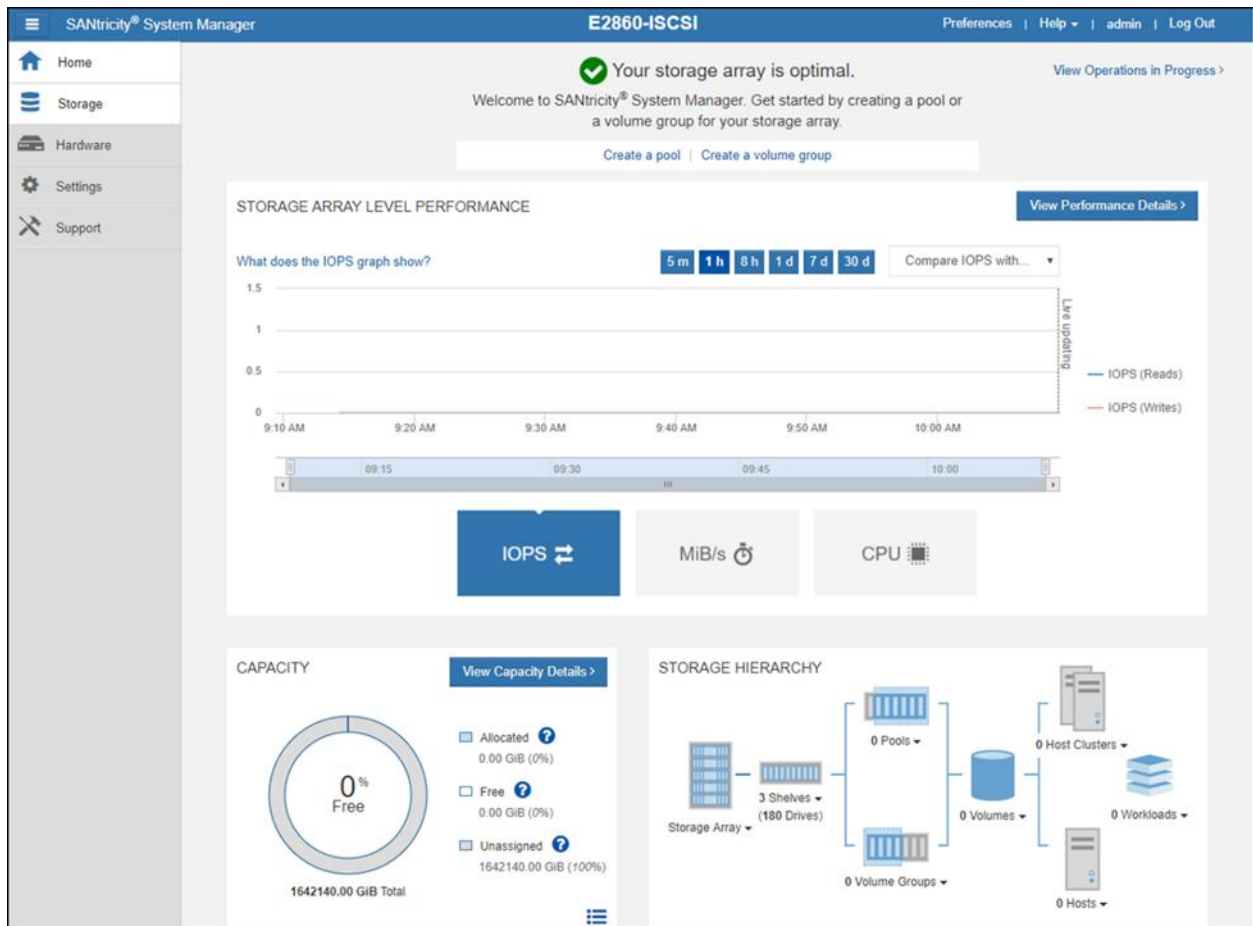
## NetApp E-Series storage provisioning

The NetApp SANtricity OS is extremely easy to use and is quite intuitive. The example that starts with Figure 6 shows the process of creating a pool, creating volumes, and mapping those volumes. This example uses an E2800 system with two drive shelves, totaling 180 12TB NL-SAS drives.

Figure 6 shows the home screen of SANtricity System Manager after you log in to the system from a web browser.

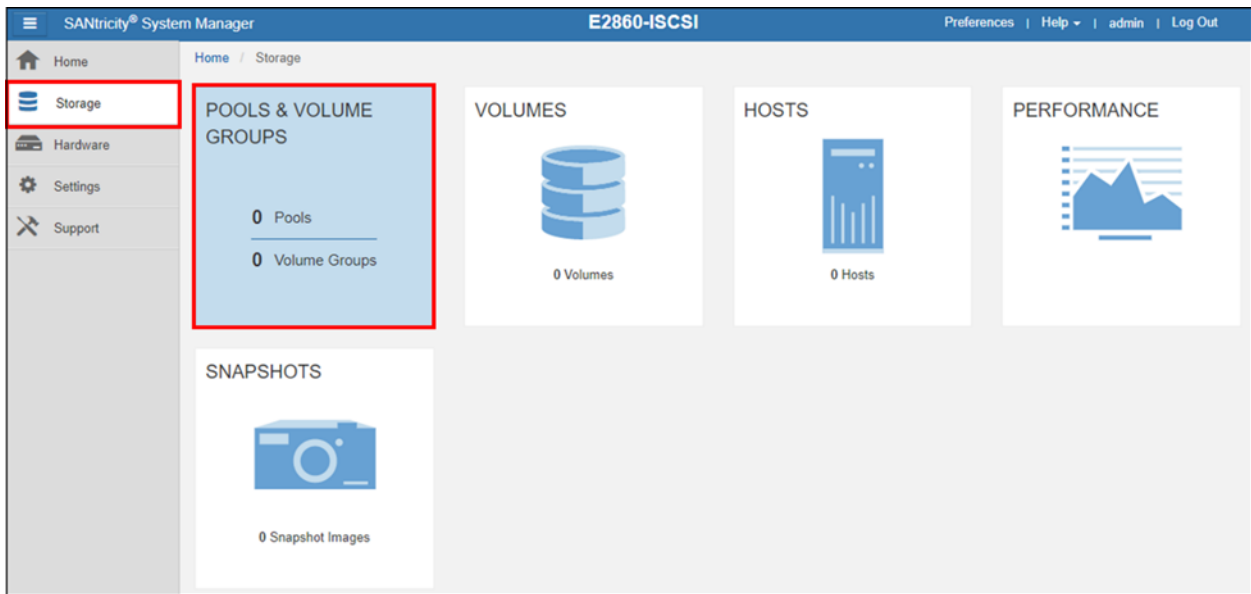


Figure 6) SANtricity System Manager home screen.



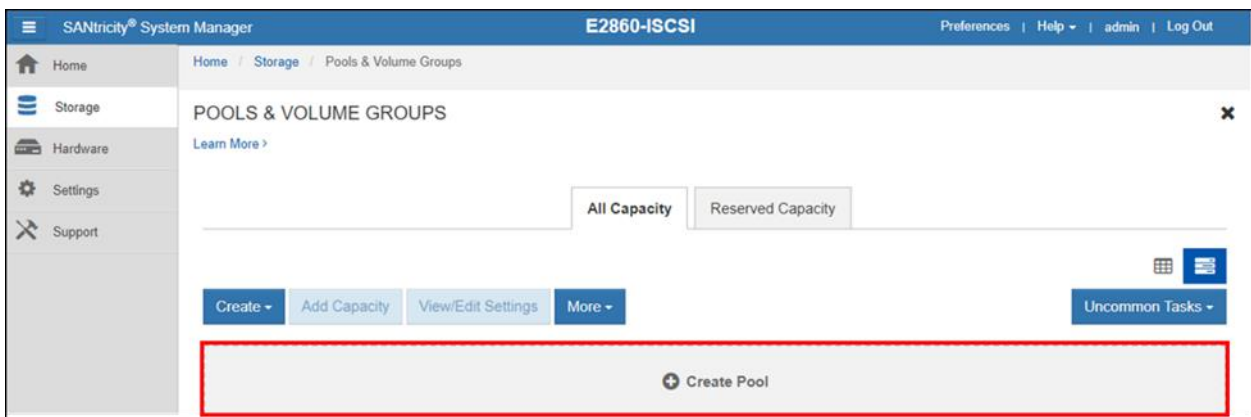
To begin creating your pool, click the Storage tab to the left and then click the Pools & Volume Groups tile as shown in Figure 7.

**Figure 7) Pools & Volume Groups tile.**



Next, click either the Create drop-down menu or the Create Pool button, shown in Figure 8.

**Figure 8) Create a DDP pool.**



From here, you see a Create Pool dialog box as shown in Figure 9. You must name your DDP pool and select the drives that you want to use for this pool. After you name your pool and select your drives, click the Create button at the bottom right. Your pool has now been created, and you are ready to begin creating your volumes.

Figure 9) Name the pool and select drives for the pool.

Create Pool

What is shelf loss protection and drawer loss protection?

Name ?

VMS\_Pool1

Select a capacity for your pool ...

Free Capacity (GiB)	Total Drives	Secure-Capable	Enable Security?	DA Capable	Shelf Loss Protection	Drawer Loss Protection
201152.00	30	Yes - FIPS	Key Required	No	No	Yes
193964.00	29	Yes - FIPS	Key Required	No	No	No
186784.00	28	Yes - FIPS	Key Required	No	No	No
179600.00	27	Yes - FIPS	Key Required	No	No	No
172412.00	26	Yes - FIPS	Key Required	No	No	No
165232.00	25	Yes - FIPS	Key Required	No	No	No
158048.00	24	Yes - FIPS	Key Required	No	No	No
150860.00	23	Yes - FIPS	Key Required	No	No	No
143680.00	22	Yes - FIPS	Key Required	No	No	No
136496.00	21	Yes - FIPS	Key Required	No	No	No

« < | Page 4 of 4 | > »

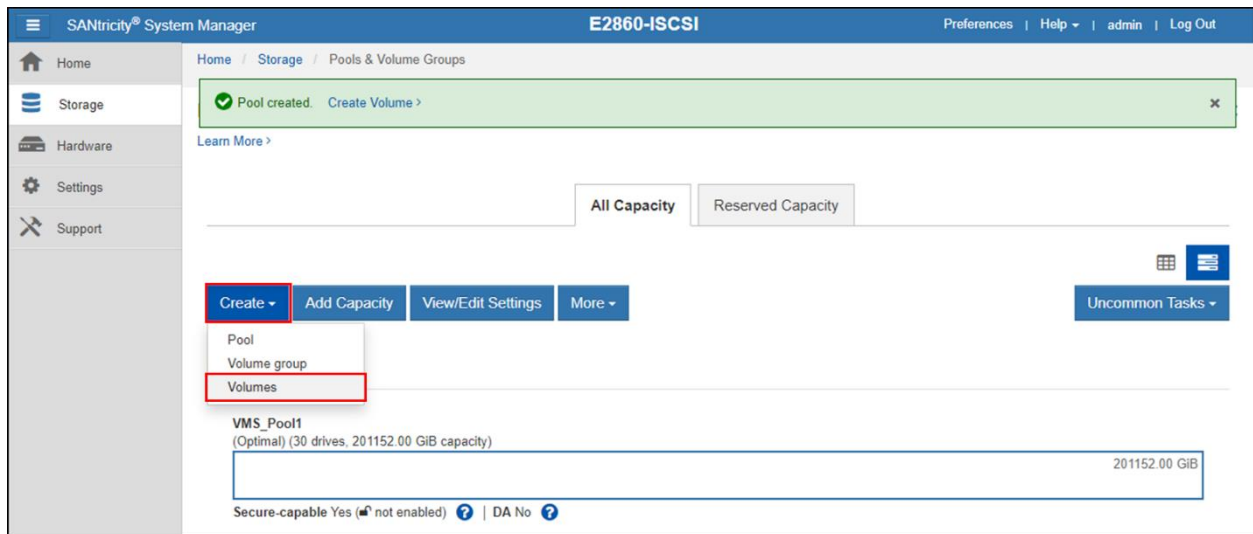
Showing 50 rows per page

Create

Cancel

After you create your pool, the dialog box in Figure 10 appears. Here, you can see that the pool called VMS\_Pool1 has 201152.00GiB free, so now you can create some volumes. Click the Create drop-down menu, and click Volumes.

**Figure 10) Create a volume.**



On E-Series systems, be sure to create at least two volumes on each system. Having two volumes enables both controllers in the E-Series array to actively contribute to the performance of the system. Figure 11 shows the creation of two volumes on VMS\_Pool1. These volumes are named VMS\_Vol1 and VMS\_Vol2.

Figure 11) Create a volume (continued).

Create Volumes

1 Select Host

2 Select Workload

3 Add/Edit Volumes

4 Review

I want to create volumes from the following pool(s) and/or volume group(s)...

VMS\_Pool1  
(Optimal) (30 drives, 201152.00 GiB capacity)

Allocated

Proposed

Free

201152.00 GiB

Secure-capable Yes (not enabled) | DA No

<input type="checkbox"/>	Volume Name	Reported Capacity		
<input type="checkbox"/>	VMS_Vol1	100576	GiB	✕
<input type="checkbox"/>	VMS_Vol2	100576	GiB	✕

+ Add new volume

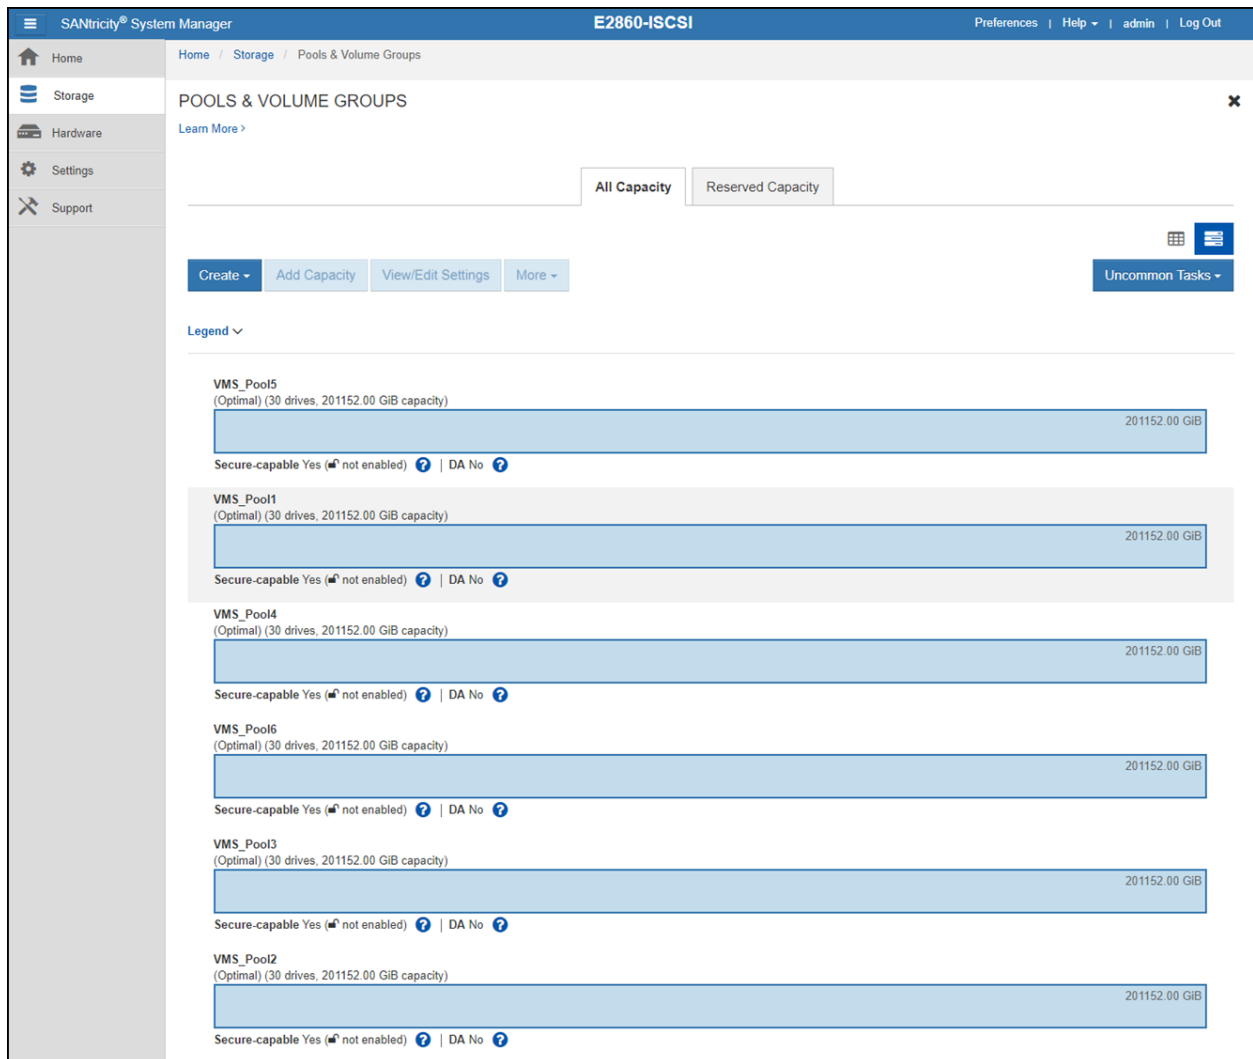
< Back

Cancel

Next >

Figure 12 shows what the system configuration looks like after you have allocated all the storage on the system.

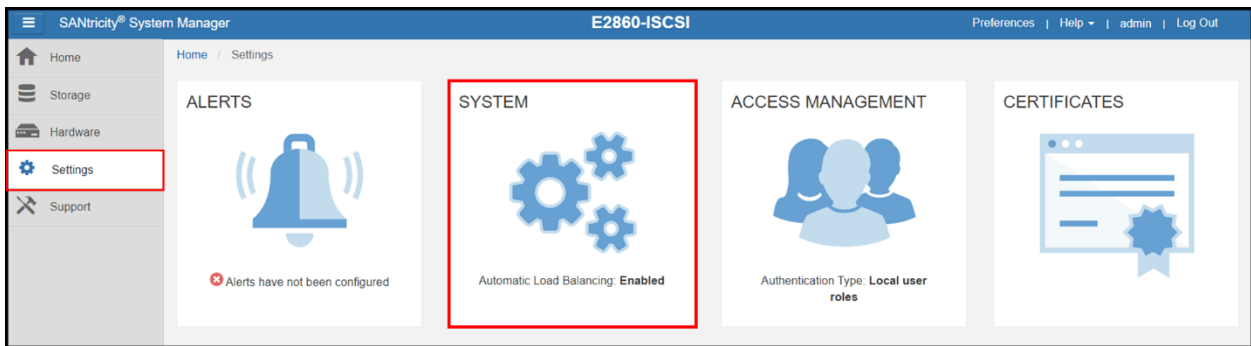
**Figure 12) Example of final volume configuration.**



**Note:** Before you can create hosts as shown in the following example, you must perform some setup on the attached server or servers. You can find those steps on the [E-Series and SANtricity 11 Documentation Center](#) under the software installation, configuration, and upgrade section for your OS.

Because iSCSI is used as the protocol in this example, Figure 13 and Figure 14 show where you configure your iSCSI ports. To enable communication with the E-Series storage array, discovery from the host side is required. To navigate to Configure iSCSI Ports (Figure 14), click the Settings tab and then click the System tile as shown in Figure 13.

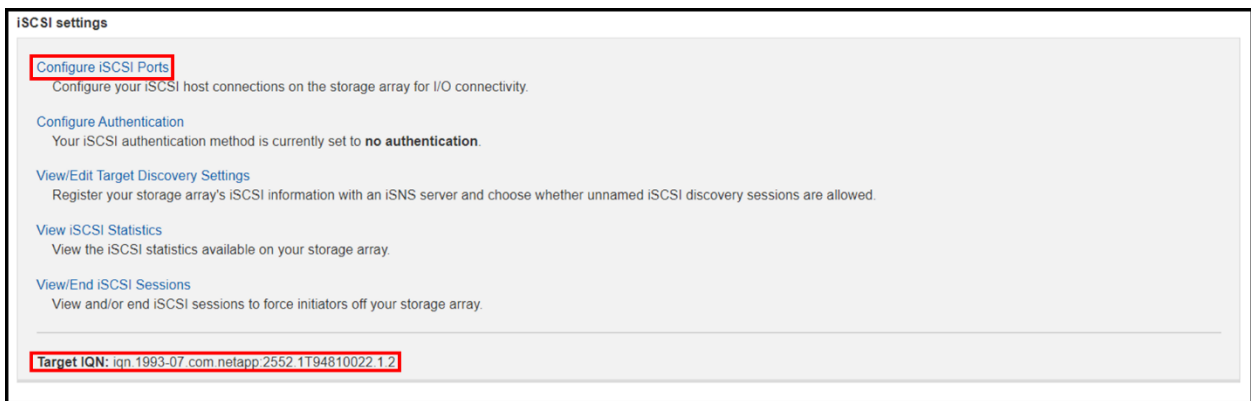
**Figure 13) iSCSI settings.**



In the System menu, scroll down to the iSCSI Settings section, which contains the Configure iSCSI Ports link that is shown in Figure 14. From there, you can configure your IPs for the iSCSI ports.

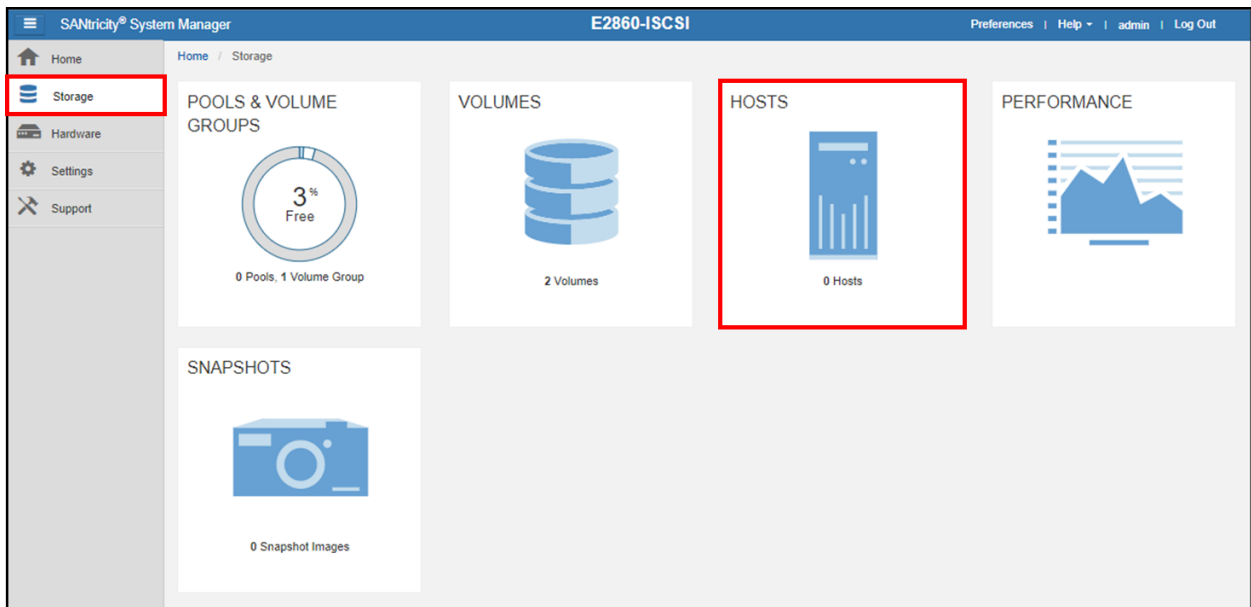
**Note:** Also outlined in Figure 14 is the target iSCSI Qualified Name (IQN). If you are connecting to multiple E-Series storage arrays and need to differentiate between arrays, this setting might come in handy.

**Figure 14) iSCSI settings (continued).**



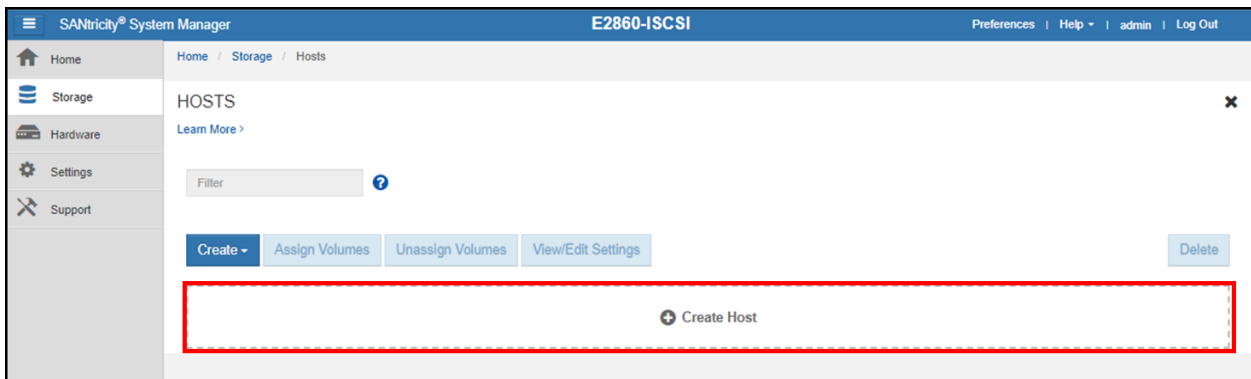
Finally, after you create your volumes (and after you set up iSCSI sessions if you use the iSCSI protocol), you can create hosts and map volumes to hosts. After you create your volumes, return to the home screen, click the Storage tab on the left, then click the Hosts tile as shown in Figure 15.

**Figure 15) Create a host.**



After you click the Hosts tile, the dialog box in Figure 16 appears. Simply click Create Host.

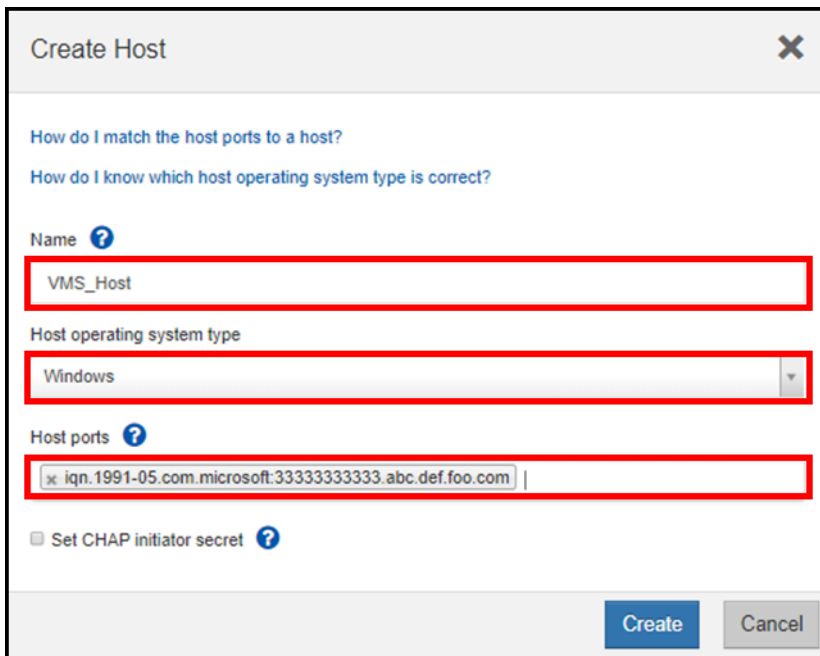
**Figure 16) Create a host (continued).**



In the following dialog box in Figure 17, you must provide a name for your host, select the OS, and select the IQN of the host. After you have completed that information, click Create, then your host is available for you to assign volumes to it.



**Figure 17) Host definitions.**




The 'Create Host' dialog box contains the following fields and options:

- Name:** A text input field containing 'VMS\_Host'.
- Host operating system type:** A dropdown menu with 'Windows' selected.
- Host ports:** A text input field containing 'iqn.1991-05.com.microsoft:3333333333.abc.def.foo.com'.
- Set CHAP initiator secret:** An unchecked checkbox.
- Buttons:** 'Create' and 'Cancel' buttons at the bottom right.

After you have defined the host, the dialog box in Figure 18 appears. Here, you click the Assign Volumes button and assign the volumes that you created previously.

**Figure 18) Assign volumes to the host.**



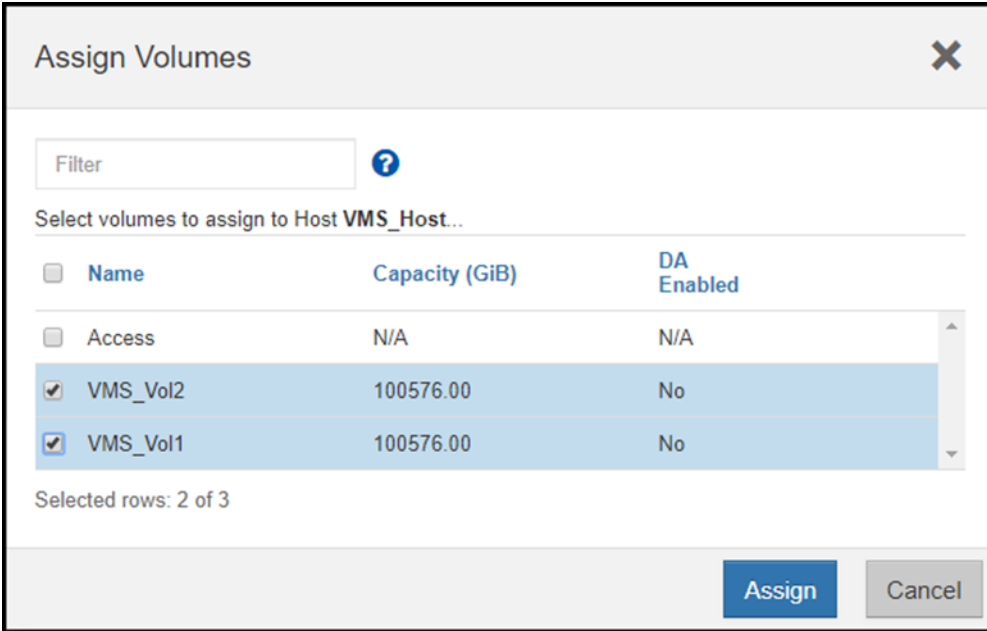
The SANtricity System Manager interface shows the 'HOSTS' page. The 'Assign Volumes' button is highlighted. Below it is a table with the following data:

Name	Total Assigned Volumes	Reported Capacity (GiB)	Host Type	Edit
VMS_Host	0	0.00	Windows	

Total rows: 1

Figure 19 shows the Assign Volumes dialog box. Here, you see the volumes that you have created. Check the boxes next to the volumes that you want to assign to the particular host that you selected in Figure 18.


Figure 19) Select volumes to assign.

A dialog box titled "Assign Volumes" with a close button (X) in the top right corner. It features a "Filter" input field with a help icon (?) to its right. Below the filter is the text "Select volumes to assign to Host VMS\_Host...". A table lists available volumes with columns for a selection checkbox, "Name", "Capacity (GiB)", and "DA Enabled". Two rows are visible: "Access" (N/A capacity, N/A DA Enabled) and "VMS\_Vol1" (100576.00 capacity, No DA Enabled). The "VMS\_Vol1" row is selected. Below the table, it says "Selected rows: 2 of 3". At the bottom right are "Assign" and "Cancel" buttons.

<input type="checkbox"/>	Name	Capacity (GiB)	DA Enabled
<input type="checkbox"/>	Access	N/A	N/A
<input checked="" type="checkbox"/>	VMS_Vol1	100576.00	No

Finally, in Figure 20, you can see that VMS Host now has two volumes assigned to it. From the server side, you can now use the volumes as storage.

Figure 20) Host with assigned volumes.

A screenshot of the SANtricity System Manager web interface. The top navigation bar shows "SANtricity System Manager", "E2860-ISCSI", and links for "Preferences", "Help", "admin", and "Log Out". A left sidebar contains icons for "Home", "Storage", "Hardware", "Settings", and "Support". The main content area is titled "HOSTS" and includes a "Filter" input field. Below the filter are buttons for "Create", "Assign Volumes", "Unassign Volumes", "View/Edit Settings", and "Delete". A table displays the "VMS\_Host" with columns for "Name", "Total Assigned Volumes", "Reported Capacity (GiB)", and "Host Type". The "Total Assigned Volumes" cell for "VMS\_Host" contains the number "2" and is highlighted with a red box. An "Edit" icon is visible at the end of the row. At the bottom left, it says "Total rows: 1".

Name	Total Assigned Volumes	Reported Capacity (GiB)	Host Type
VMS_Host	2	201152.00	Windows

## Server considerations

This section focuses on recording and failover server requirements.

The recording server represents one or more instances of the hardware and software that are used to record live video or to archive 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 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 your VMS provider. As a general principle, and depending on your system hardware specifications, the amount of video that any individual server can process ranges from 100Mbps to 600Mbps.

Table 7 lists the general recording server characteristics that NetApp recommends.

**Table 7) Recording server characteristics that NetApp recommends.**

Characteristic	Description
Form factor	1 RU for space savings
CPU	Quad-core in the 2.0GHz to 2.9GHz frequency range
RAM	8GB or higher
Network adapters	Integrated Ethernet adapters and PCI-based quad-port 1Gbps/10Gbps Ethernet for video ingress and, optionally, IP SAN connectivity
Internal drives	Dual RAID 1 (internal RAID controller) for a high-availability boot drive
OS	Windows 2008 R2 or later
Recording volume file system	NTFS (allocation unit size: 64KB)

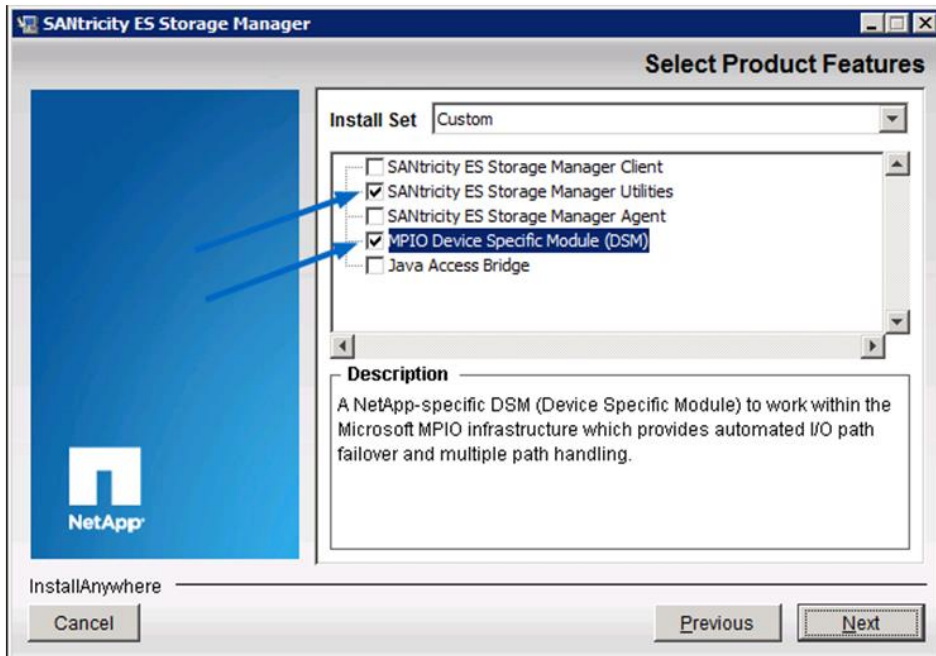
**Note:** For the latest hardware and software requirements, go to your VMS provider's website.

## Multipath I/O device-specific module installation

As described in the “High Availability” section, NetApp E-Series arrays support multiple paths to the LUNs from a server. To manage these paths, you must download the appropriate E-Series SANtricity Storage Manager software from the [NetApp Support site](#), and you must install the software on each server.

When you install the SANtricity Storage Manager software, use the Custom installation option and then select the items that are shown in Figure 21. This step installs the host MPIO DSM and the utilities files. If you are asked to start a background monitor process or agent, do not select this option, or select No.

**Figure 21) E-Series SANtricity Storage Manager installation.**



**Note:** For more help, go to the [E-Series and SANtricity 11 Documentation Center](#).

## E-Series LUN discovery and preparation

After the appropriate LUNs have been mapped to the recording server as shown in the “NetApp E-Series Storage Provisioning” section, you should perform the following steps:

1. Navigate to:

```
C:\Program Files (x86)\Storage Manager\util
```

2. Run the command `SMdevices`; it should list all the E-Series LUNs that are mapped to the server and display various information, including the current and preferred controllers.
3. Use the information that `SMdevices` reports to help you in mapping E-Series LUNs to drive letters in the Windows Disk Management tool. The Windows Disk Management tool identifies an E-Series LUN as a drive that must be initialized, formatted, and mapped to a drive letter before I/O can be issued to that drive. The tool is used to view and to set details, such as the configuration of drive type, volume name, and allocation unit size. For video surveillance implementation, NetApp recommends an allocation unit size of 64KB.

**Note:** For more information about Windows Disk Management, read the [overview from Microsoft](#).

## Networking considerations

The network infrastructure for video surveillance deployments must meet the following requirements:

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

To meet the preceding requirements, for NetApp E-Series systems, NetApp recommends a 10Gb Ethernet network at a minimum. The E2800 and the E5700 systems have two onboard host ports per controller that support either 10Gb iSCSI or 16Gb FC. Figure 22 and Figure 23 show these host ports on the top left of each controller.

For iSCSI deployments, the use of multiple Ethernet network interface cards (NICs) connecting to dual IP SANs also provides high availability to the E-Series controllers. For other protocols, such as FC or SAS between the server and the E-Series controllers, dual-port host bus adapters (HBAs) provide redundant paths to each controller.

Figure 22) Back view of the E2800 controller.

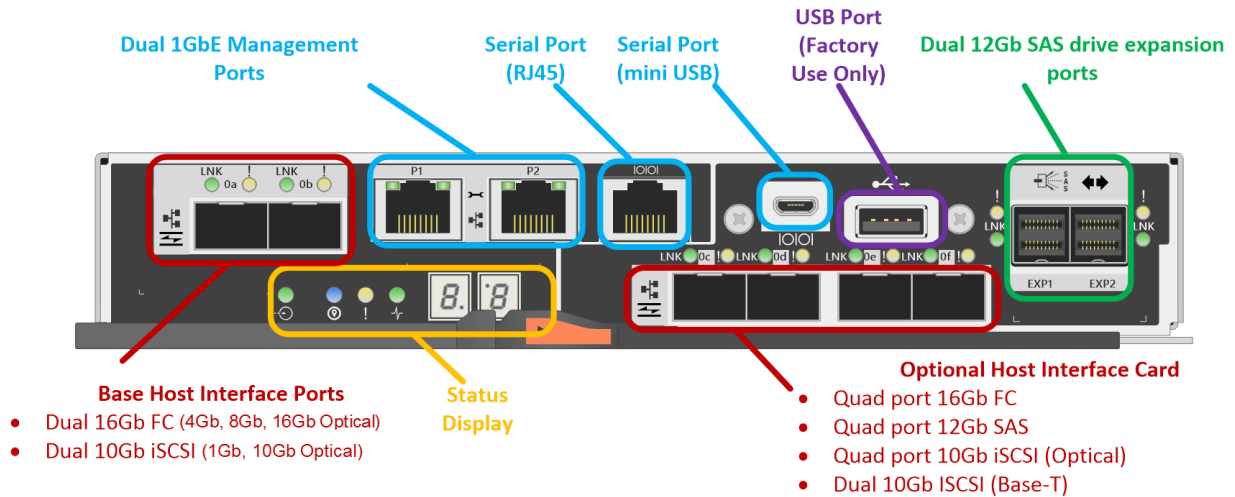
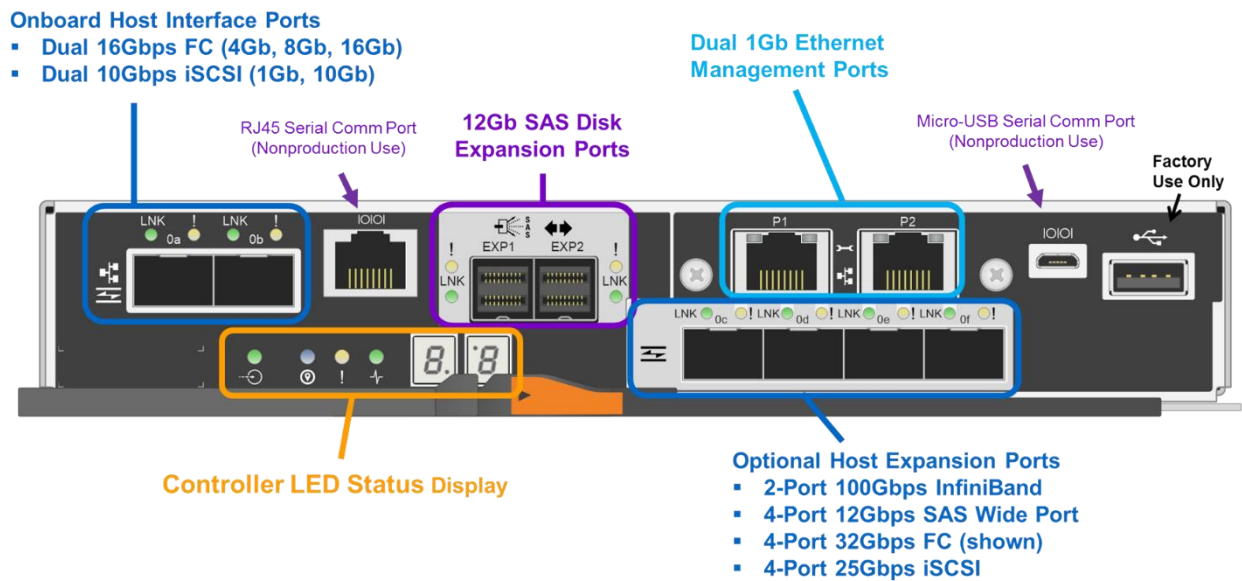


Figure 23) Back view of the E5700 controller.



## Traffic management

To maintain a high-performing network, use multiple NICs to separate the camera network, the client network, and SANs. By separating these networks, you gain the following benefits:

- **Increased performance.** By separating the traffic, you eliminate the impact on recording performance that a high load on a client network might have.
- **Stability.** With separated networks, interference on the client network does not affect the camera network, which promotes predictable performance.
- **Increased security.** No accidental or intentional interference with camera operations occurs. By isolating the camera network, you eliminate the possibility of devices sending information through the internet without your knowledge or permission.

- **Improved management.** Management is also easier because the load is independent to each network. These independent loads make it easier to calculate and to measure the bandwidth usage on each network.

**Note:** As a best practice, be sure to eliminate single points of failure. An example is to implement a secondary SAN switch. If a switch fails, your recording servers have an additional path to your NetApp E-Series storage and cameras can continue recording.

## Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- Video Management Systems Using VMware Virtualization with NetApp E-Series Storage  
<https://www.netapp.com/us/media/tr-4818.pdf>
- NetApp E-Series Documentation Center  
<https://docs.netapp.com/ess-11/index.jsp>

## Version history

As an option, use the NetApp Table style to create a Version History table. Do not add a table number or caption.

Version	Date	Document version history
Version 1.0	March 2020	Initial release of best practices guide.
Version 1.1	October 2021	Updated to add larger drive sizes.

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 © 2021 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

### **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-4825-1021