



Technical Report

Video Surveillance Solutions with NetApp E-Series Storage

Planning and Design Considerations

James Laing, Frank Poole, NetApp
December 2017 | TR-4197

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	Training Offerings	4
2	Planning and Design Overview	5
2.1	Bare Metal Servers	5
2.2	Virtualized Servers	5
2.3	File System	5
3	Storage Planning with E-Series	6
3.1	Dynamic Disk Pools	7
3.2	RAID Levels	8
3.3	Hot Spares	8
3.4	Workflow	8
3.5	Deployment Example	8
3.6	I/O Characteristics	9
3.7	High Availability	9
3.8	Multipath Overview	10
3.9	E-Series Certified Multipath Drivers	11
4	Network Planning	11
4.1	Networking Example	12
4.2	Network Interfaces	13
4.3	VMware vSphere Networking Configuration	13
4.4	Uplink Connectivity (Layer 2)	14
4.5	Uplink Connectivity (Layer 3)	15
4.6	Network Management Caveats	16
4.7	Network Design Rationale	16
5	Server Planning	18
5.1	Hardware Recommendations	18
5.2	CPU 18	
5.3	Server Manufacture	19
5.4	Memory	19
5.5	General Design Criteria	19
6	Design Checklist	19

Summary	20
Definitions	20
References	22
Version History	23

LIST OF TABLES

Figure 1) E-Series disk structure.....	6
Figure 2) DDP usable capacity (6TB drives shown).....	7
Figure 3) High-availability design.	10
Figure 4) Architectural topology overview.....	12
Figure 5) Cisco Nexus topology overview.	13
Figure 6) VMware vSphere networking configuration.	14
Figure 7) Uplink connectivity (layer 2).	15
Figure 8) Uplink connectivity (layer 3).	16

LIST OF FIGURES

Figure 1) E-Series disk structure.....	6
Figure 2) DDP usable capacity (6TB drives shown).....	7
Figure 3) High-availability design.	10
Figure 4) Architectural topology overview.....	12
Figure 5) Cisco Nexus topology overview.	13
Figure 6) VMware vSphere networking configuration.	14
Figure 7) Uplink connectivity (layer 2).	15
Figure 8) Uplink connectivity (layer 3).	16

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 an introduction to video surveillance for those who sell, design, or implement such solutions based on NetApp E-Series storage. It describes the comprehensive functional components 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.

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

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—from 200 up to more than 5,000 cameras (1-2Mbps each saved for 30 days of archiving)—with retention periods of at least 30 days and primarily using HDTV/megapixel resolution cameras.

1.3 Training Offerings

There are a number of web-based and instructor-led training opportunities to enable successful deployment of the NetApp E-Series storage array. The classes listed in the [NetApp University Customer Learning Map](#) under E-Series are recommended end-user training classes.

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

Table 1) Training offerings.

Class Description	Duration (Hours)	Delivery
Product Foundation		
Technical Overview of NetApp E-Series Storage Systems	01:00	Web based
Technical Overview of NetApp EF-Series All-Flash Array	0:30	Web based
Hands-On Skills Development		

Class Description	Duration (Hours)	Delivery
Configuring and Monitoring E-Series and EF-Series Storage Systems	32:00	Instructor led
Certification		
NetApp Certified Implementation Engineer: SAN Specialist, E-Series	01:30	Exam

2 Planning and Design Overview

The remainder of this document discusses planning and design aspects the physical security integrator must consider when implementing an E-Series storage array in a video surveillance deployment.

NetApp E-Series storage arrays work with both bare metal and virtualized server installs. Bare metal installs are more similar to the dedicated DVR appliances familiar to physical security integrators. Bare metal may be more cost effective for installs of fewer than 500 cameras per server when using SAS direct attach. Virtualized recording server instances are often favored by IT teams who want to minimize server footprints.

2.1 Bare Metal Servers

Some physical security integrators prefer to deploy relatively inexpensive 1RU servers in a nonvirtual environment to eliminate the costs of purchasing, installing, and maintaining a hypervisor. Low-end, dual-socket 1RU servers with 4GB of RAM and two Gigabit Ethernet (GbE) or 10 Gigabit Ethernet (10GbE) interfaces can easily handle 100 to 200 cameras. Servers with dual 8-core CPUs and 16GB of RAM can scale to more than 500 cameras. The NetApp E-Series can support multiple servers and scale to more than 5,000 cameras. This deployment model is particularly attractive for IP SAN/iSCSI deployments. One server network interface is used for video ingress, and the second interface provides connectivity to the storage array.

2.2 Virtualized Servers

Server virtualization is widely accepted in the enterprise data center because it provides the same logical segmentation of servers that was previously accomplished using separate physical servers. Data center servers that are not constrained by resource consumption (memory or CPU) are ideal candidates for virtualization.

Deployments that require the higher performance characteristics of the E5700 controller and dual-port Fibre Channel host bus adapter (FC HBA) are more likely candidates for implementing the recording server component as virtual machines on a higher performance 1RU, 2RU, or 4RU server. In this configuration, the FC HBA is shared by all the virtual machines running on the physical chassis. As an example, the E5700 can be attached to the server FC HBA through the native multipath drivers of VMware ESXi. Raw disk mapping (RDM) is used to present one or more volumes as logical unit numbers (LUNs) directly to the recording server virtual machine. A four-port GbE adapter can be defined as a port channel configuration to the network switch, or a single 10GbE interface can be installed. Four or more recording server virtual machines per physical server can be supported in this configuration.

Virtualization is an ideal choice to deliver high availability and excellent throughput for recording servers, while reducing the number of physical machines that must be deployed and managed.

2.3 File System

The majority of open platform-based video management software (VMS) solutions use Windows Server 2012 R2 or higher as the operating system and the NTFS or ReFS file system with an allocation unit size of 64KB.

Parallel file systems such as StorNext or Lustre are not typically deployed for video surveillance. Some VMS applications implement a tiered approach to storage, allowing the VMS administrator to define multistage storage architectures. As video archive files are moved from one level of hierarchy to another, grooming to reduce the frame rate is an option. Encryption of the video archives might also be an option.

The features of grooming and encryption, however, affect the performance of both the I/O and CPU. If the grooming is configured to move files from one volume (LUN) to a second, files must be read from the source LUN and groomed and written to the target volume (LUN). The effect on performance must be considered when implementing tiered storage.

3 Storage Planning with E-Series

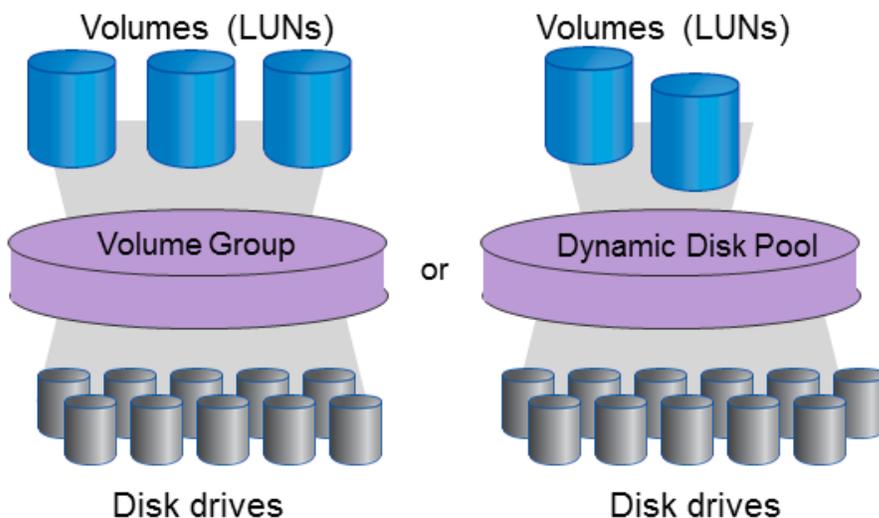
Each video recording server requires one or more volumes (LUNs) for archiving video files. The SANtricity® System Manager is used to configure the E-Series storage array. Individual hard disks are allocated to Dynamic Disk Pools (DDP) or a volume group using the Create Disk Pool/Create Volume Group wizard. In almost all cases, DDP should be the preferred configuration choice. It is easy to configure and manage, offers much faster rebuild performance, and is very simple to support.

The minimum number of disks in a DDP is 11. Thirty or more disks in a DDP are preferred. The minimum number of disks for a volume group depends on the RAID level. The maximum number of disks for RAID 5 or RAID 6 is 30. The limit for DDP is the total population of physical disk drives in the array.

During the definition step, the RAID level for all physical disks assigned to the volume group is selected. DDP uses RAID 6 stripes, with each stripe allocated over 10 drives in the pool. Stripes are staggered across the pool to distribute data and parity evenly across all drives. For traditional volume groups, the supported RAID levels are 0, 1, 10, 5, and 6. Each physical disk has a 512MB area for storing the array configuration database and optional space for dynamically changing the segment size.

Individual volumes (LUNs) are created and mapped to a host following the DDP or volume group definition. Each volume can be individually configured for segment size, modification priority, cache settings, and media scan frequency. The logical relationships of disk pools, volumes, and volume groups are shown in Figure 1.

Figure 1) E-Series disk structure.



The number of physical disks per disk pool or volume group and the number of volumes per group or pool are determined by the performance and sizing requirements of the video recording server and the application software.

3.1 Dynamic Disk Pools

DDP is a feature available on the E-Series to maintain a consistent level of performance delivery even in the event of drive failure and reconstruction. DDP is one of the features that differentiates E-Series from other storage systems. No other vendor offers a storage option that is as easy to install or that offers the benefits of DDP. The performance drop is minimized during rebuild, and the rebuild completes more quickly than with a traditional RAID rebuild. Because of the shorter rebuild time with DDP, the exposure to data loss from several drive failures is minimized. A single pool may be defined that includes all disks in the system, or multiple pools may be defined to the system. There are no idle hot spares when DDP is used; spare capacity is incorporated into the pool.

The minimum number of disk drives in a DDP is 11, though NetApp prefers 30 to 180 drive pools, maximizing usable capacity. Each stripe of data spans 10 drives in the pool, and an extra drive is needed to provide redundancy across all drives in the pool. DDP uses RAID 6. The storage administrator can configure a mixture of both traditional volumes and DDP. Traditional volumes with RAID 10 can be created for maximum performance, and DDP can be configured for capacity volumes.

As an example, the free (usable) capacity for volume sizes commonly deployed for video recording servers using 10TB drives is shown in Figure 2.

Figure 2) DDP free (usable) capacity (10TB drives shown).

Free Capacity (GiB)	Total Drives	Secure-Capable	DA Capable	Shelf Loss Protection	Drawer Loss Protection
129384.00	20	No	Yes	No	Yes
122196.00	19	No	Yes	No	No
115008.00	18	No	Yes	No	No
107820.00	17	No	Yes	No	No
100632.00	16	No	Yes	No	No
93444.00	15	No	Yes	No	Yes
86256.00	14	No	Yes	No	No
79068.00	13	No	Yes	No	No
71880.00	12	No	Yes	No	No
71876.00	11	No	Yes	No	No

For video surveillance solutions, smaller single-volume disk pools are optimal for bandwidth and provide performance comparable to that of RAID 6 with rebuilds that are twice as fast. This is why video surveillance management software performance testing is measured when the system has filled the volume to capacity and is in file-deletion mode. DDP performance does not degrade like RAID 6 and is more consistent throughout the lifecycle. Thirty drive pools are a conservative configuration, balancing performance, usable capacity, and simplicity.

3.2 RAID Levels

The Nevada Gaming Commission standards specify that the storage array must not lose data in the event of the failure of a single component. Although RAID 6 provides better fault tolerance because it can tolerate two disk failures, RAID 5 is often deployed instead because of lower costs while still adhering to the standards.

On lower-performing storage systems, RAID 10 is typically used for best read performance when combined with solid state disks (SSDs) or disks with higher (15K) rotational speed. RAID 5/RAID 6 is used for best write performance. The E-Series delivers both good read and write performance using DDP.

Some VMS vendors recommend using a combination of RAID 10 and RAID 5 in gaming deployments in which a high volume of forensic analysis occurs during the most recent minutes or hours of video archives. These designs use RAID 10 for the most recent archive and 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 analyses. The education market is one vertical in which reviewing archived video occurs only if an incident (for example, vandalism or an altercation between students) warrants analysis of the video.

Once again, the E-Series addresses these needs using DDP, eliminating the need for complicated storage configurations.

3.3 Hot Spares

Hot spares are disks that remain idle until needed. They are only needed with volume group configurations and are not used with DDP. Hot spares are used in place of a failed drive, allowing reconstruction of the data and parity across the number of drives in the volume group. Video surveillance performance is often measured during a disk rebuild because the system is under both read and write I/O during the rebuild process. NetApp recommends using a minimum of one hot spare per every 30 drives in the system when volume groups are used.

The amount of time required to rebuild a failed drive on a hot spare drive depends on the size of the drive and the number of drives in the volume group and might take hours or days.

3.4 Workflow

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

Optimizing IOPS is important when the disk array is used for small random I/O operations from multiple applications. Network packet-per-second performance is usually measured for small (64-byte) packets.

However, video surveillance deployments 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. Because the video ingress to the recording servers is over an IP network, and the data rate is typically calculated in Mbps for IP networks, many of the tables in this document list Mbps rather than megabytes per second (MBps).

3.5 Deployment Example

Examining the characteristics of an actual deployment helps put the workflow characteristics in perspective.

A single recording server is managing 46 network video cameras configured for continuous recording using H.264/RTP/UDP with 720p resolution at 30 frames per second. The network switch port interface statistics for the recording server report that the data rate to the server is 11,000 packets per second at 118Mbps.

From this information, the average packet size to the server is calculated at $((118\text{M} / 8 \text{ bits}) / 11,000)$, or 1,406 bytes per packet. The workload to the volume (LUN) defined to the recording server is reported by SANtricity at approximately 20Mbps at 44 IOPS. That rate is equivalent to 144Mbps with the average I/O size of 465KB. Video management systems commonly use a record size of either 256KB or 512KB.

This sample recording server receives 11 IP packets every millisecond (ms) and generates a write operation to the storage array every 22ms.

3.6 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. In video analysis deployments (data analytics, data forensics, and so on), the read workload can be significantly higher. Real-time analysis does not add to the read workload because the data is analyzed as it is written.

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 tenth 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 that the video was originally written to disk. Write workload is relatively easy to characterize, while read workload is less predictable.

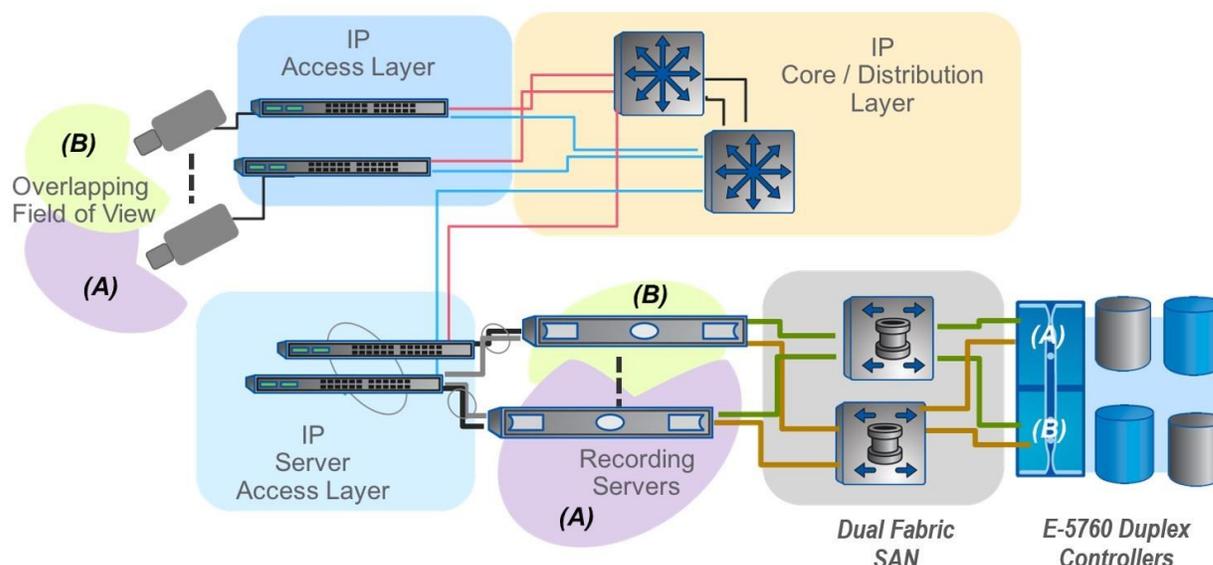
The architecture and configuration of the video management system also affect the workload to the storage array. Systems that implement tiered storage schedule a copy from one volume or directory to another at a recurring interval (such as hourly or daily), and during the copy function the IOPS of the 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 with which video is read or copied following the initial write.

3.7 High Availability

Real-time applications such as video provide a challenge for physical security integrators in that any outage or failure between a network video camera and the storage system means the record of events is lost and cannot be recovered. Implementing high availability for video surveillance begins with considering camera placement, the network infrastructure, server and video management software redundancy, and finally the storage array. These components are shown in Figure 3.

Figure 3) High-availability design.



For areas of critical importance, multiple cameras with overlapping fields of view should be implemented to maintain coverage in the event a single camera or access layer switch fails. Multiple cameras covering the critical area must be connected to separate access layer switches with redundant uplinks to the core/distribution layer switches. The IP network must implement high-availability network design principles, rapid convergence from link and/or 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 2012 and Windows Server 2016 supports link aggregation.

For Fibre Channel or direct-connect SAS connectivity between the server and the E-Series, dual-port HBAs are installed that provide redundant paths to each controller. For iSCSI deployments, multiple Ethernet NICs connecting to dual IP SANs also provide for high availability to the E-Series controllers.

The failover drivers are critical to providing path failure recovery between server and 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 paths to a controller
- Provide status of controller/bus

For Windows, the failover drivers are a combination of Microsoft MPIO plus the SANtricity host installation device-specific module (DSM). E-Series supports the native multipath feature of VMware ESXi.

3.8 Multipath Overview

Hosts identify devices based on their initiator port, the target port, and the LUN number. Hosts with redundant IP SAN interfaces (iSCSI), dual-port SAS interfaces, or dual-port HBA adapters (Fibre Channel) connected to a duplex E-Series controller have redundant paths to their LUNs. The host installation option of the SANtricity installation utility must be installed on physical recording servers for Windows deployments to implement the multipath driver necessary to direct I/O through the correct path

to the LUN. Windows Server running in guest virtual machines does not require the E-Series multipath drivers to be installed; the native multipath drivers for ESXi are used instead.

In addition to providing multiple path discovery and configuration, multipath drivers manage I/O load balancing across multiple paths and manage controller, path failover, and failback. Using all available paths, for example, by selecting a round robin or least queue depth option, is most effective for increasing the throughput between host and storage controller for the relatively slower host interfaces' connectivity. Deploying iSCSI over 10GbE or 25GbE interfaces might encounter substantially better throughput by load balancing the multiple paths than deploying a single 16Gb or 32Gb Fibre Channel connection.

3.9 E-Series Certified Multipath Drivers

The E-Series certified multipath drivers for Windows Server 2012 and Windows Server 2016 are the Windows MPIO component and the SANtricity host installation that loads the appropriate DSM. By default, Windows supports four paths per controller, with a maximum of 32 paths. Windows supports up to 255 volumes (LUNs) per host.

For VMware, the VMware native multipathing plug-in (NMP) is certified. When running Windows Server in virtual machines under VMware ESXi, only the SANtricity host utility should be installed. The utility SMdevices is installed as part of the SANtricity host utility installation and is a useful troubleshooting tool for identifying the attached storage array name and volume information.

The sample topology illustrated in Figure 3 has redundant paths between the video recording server and the storage array. For Fibre Channel deployments, multiple active and standby paths might exist in the topology, depending on the number of ports in use. For iSCSI, configuring multiple active and standby paths is a manual process in the Microsoft iSCSI initiator.

To recap, SANtricity is installed on each Windows recording server:

- For VMware ESXi guests, select custom installation and install only the utilities.
- For nonvirtual Windows deployments, selecting the host option to install the utilities and the DSM provides multipathing support for high availability.

4 Network Planning

Video surveillance deployments require a network infrastructure that addresses these requirements:

- Provides sufficient available capacity (bandwidth) to transport video
- Exhibits very low/no loss of IP video packets
- Features network latency within the range suitable for the transport protocol (TCP or UDP) of the video feed
- Provides high availability through network redundancy and best practices in network design
- Meets the network security and services requirements

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

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

converged on an existing network infrastructure, the physical security department or integrator must work with the IT department to implement network equipment consistent with the existing infrastructure.

Leading network vendors as well as leading integrators offering voice and video network implementation services can assist with network readiness assessments for IP video surveillance deployments.

4.1 Networking Example

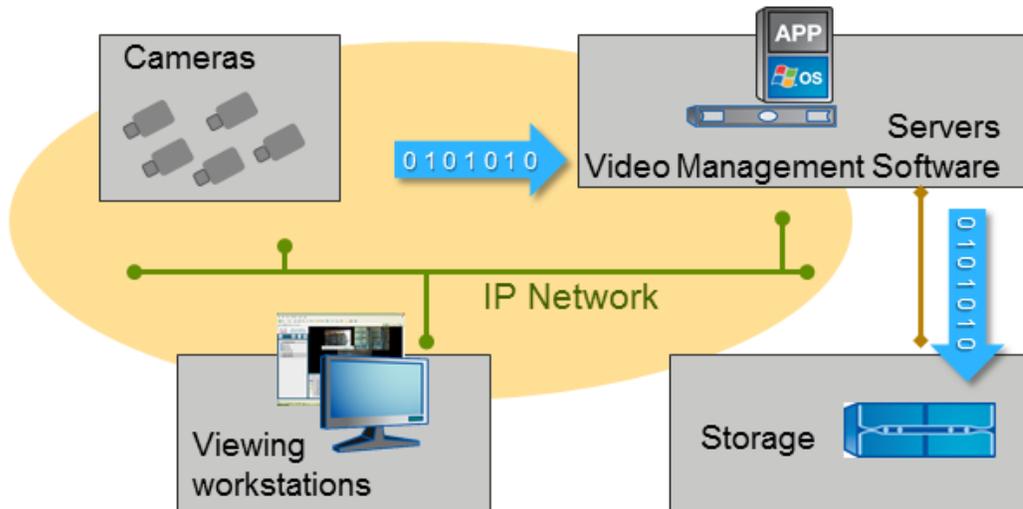
This section discusses how the network switches provide top-of-rack connectivity to the recording servers and integrate with the core/distribution layer switches. A sample configuration built and tested in NetApp's RTP labs is used to illustrate the specifics of networking for video surveillance.

A sample video surveillance solution has been validated for the deployment, consisting of NetApp E-Series E2800 storage, up to four Cisco UCS C-series servers, and two Cisco Nexus 3000 series top-of-rack integrated layer 2/3 switches. VMware ESXi is installed on each Cisco UCS server. Each server is configured with four virtual machines, each running Windows Server. The testing described in this section was performed with a NetApp E2860 storage system.

The Cisco Nexus 3000 series switches provide the server access-layer switching infrastructure to connect the video surveillance system to the end-customer IP network infrastructure. In most deployments, the network video cameras and viewing workstations are connected to existing or new network routers, and switches are installed as part of the physical security deployment.

The IP network is a critical component in the architecture because it provides connectivity to all key components, as shown in Figure 4.

Figure 4) Architectural topology overview.



Note: The E-Series storage arrays are connected to the network switches to provide management access for workstations running SANtricity, even though the host attachment is over Fibre Channel links or direct SAS attachment.

Each Cisco Nexus 3000 series switch supports multiple 1/10/40 Gbps ports. Two ports on each switch are configured as 10Gbps virtual port channel (vPC) peer links; two ports on each switch should be used

for either layer 3 (routed) or layer 2 (switched) uplinks. The tested configuration utilized one 10Gbps SFP+ fibre on each switch for uplink connectivity and high availability.

The Cisco Nexus 3000 series connects to data and management interfaces on servers and the E-Series management ports. There are redundant power supplies and redundant fans in the fan tray.

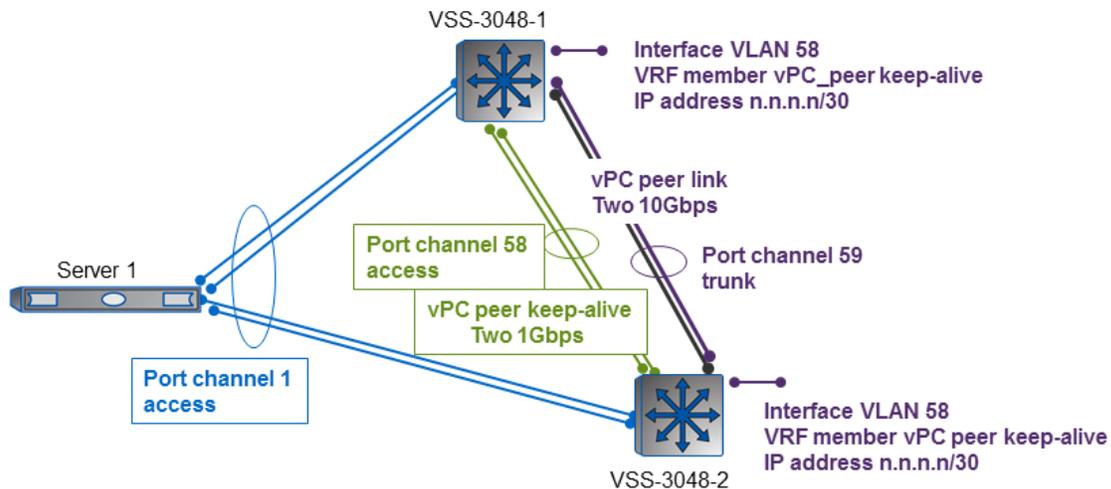
The Cisco Nexus 3000 series enables high availability through redundant power supplies and fans, redundant uplinks, and the vPC feature. Either of the Cisco Nexus switches in the video surveillance deployment can fail or be taken out of service without disrupting the ability of the video surveillance servers to capture and record video streams from networked video cameras.

4.2 Network Interfaces

The sample configuration tested by NetApp used Cisco UCS C220 servers, each with a quad-port Ethernet adapter. The four ports are aggregated into one logical link. This link provides video ingress to the servers from the network video cameras. Two of the member links are connected to one Cisco Nexus switch, and the other two member links are connected to the second Cisco Nexus switch. The Cisco Nexus switches are configured with two vPC peer keep-alive links (1Gbps) between switches and two 10Gbps vPC peer links between switches. The vPC peer keep-alive links carry only control plane traffic and are used to detect a peer failure. The vPC peer links are layer 2 trunks and transport both the device management VLAN traffic as well as traffic for the server port channel interfaces in certain failure situations. Although we tested with Cisco Nexus, any enterprise-class switch should be able to satisfy the networking requirement.

This network topology is illustrated in Figure 5.

Figure 5) Cisco Nexus topology overview.

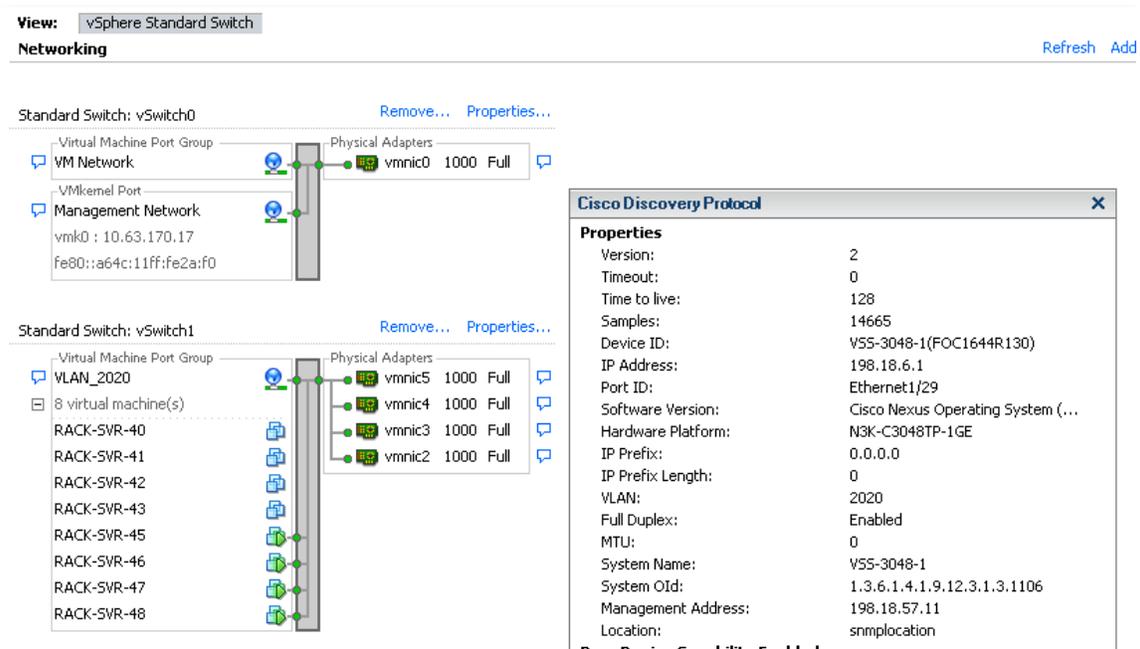


Note: Only one server (Server 1) is shown for clarity. All servers are similarly configured. The uplink connectivity is shown later in this document.

4.3 VMware vSphere Networking Configuration

From the VMware vSphere client of the Cisco UCS C220 server, the video ingress network is configured as a switch with four physical adapters. This configuration sample is shown as vSwitch1 (VLAN 2020) in Figure 6.

Figure 6) VMware vSphere networking configuration.



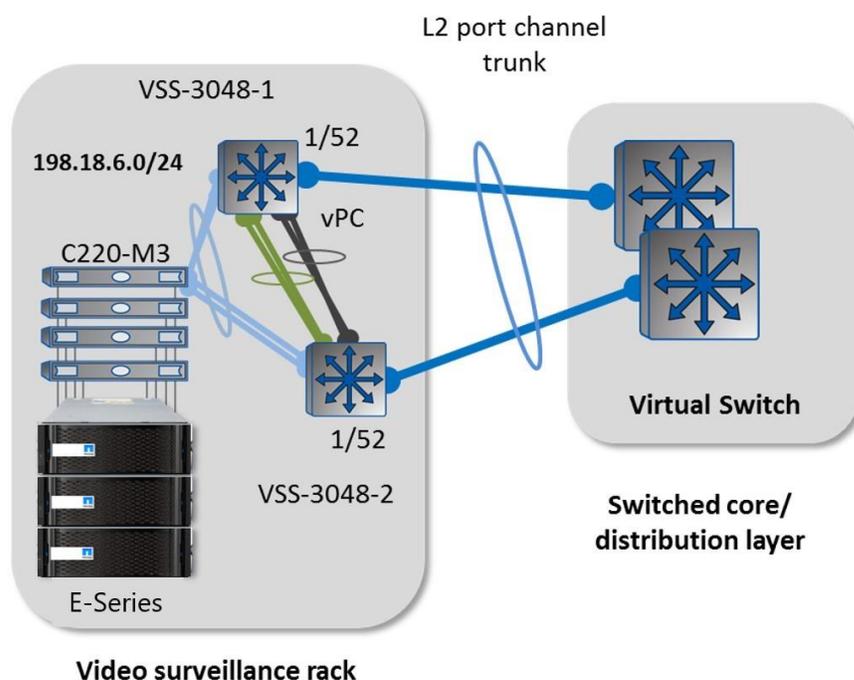
In addition to the video ingress VLAN, a device management VLAN is configured to provide management connectivity for the E-Series management ports; a Cisco integrated management controller port for physical server management; an ESXi VMkernel management network port; and a guest operating system management network port for SSH, Linux X-terminal, or Windows remote desktop connection connectivity.

4.4 Uplink Connectivity (Layer 2)

As tested, the design does not require any additional Cisco NX-OS software packages for the Cisco Nexus switch if only layer 2 services are used. The system default (no license required) includes features used in this solution: VLAN, IEEE 802.1Q trunking, vPC, Link Aggregation Control Protocol (LACP), Secure Shell Version 2 (SSHv2) access, and Cisco Discovery Protocol.

With this option, the uplink connections between the video surveillance system and the campus core/distribution switches are configured as layer 2 (switched) port channel trunks. This topology is shown in Figure 7.

Figure 7) Uplink connectivity (layer 2).



Given this assumption, the end-customer core/distribution switches must be configured to provide layer 2 and layer 3 features to support a video surveillance solution. These features include:

- Primary and secondary root spanning-tree bridge (Rapid Spanning-Tree Protocol [RSTP])
- Ethernet switch virtual interfaces (SVIs), for example, interface VLAN for video ingress and management VLAN
- Hot Standby Router Protocol (HSRP) or Virtual Router Redundancy Protocol (VRRP); virtual IP addresses for the video ingress and management VLAN

As part of the installation and implementation process, verify the high-availability configuration of the design by alternately reloading the Cisco Nexus 3000 series switches and validating connectivity and recovery.

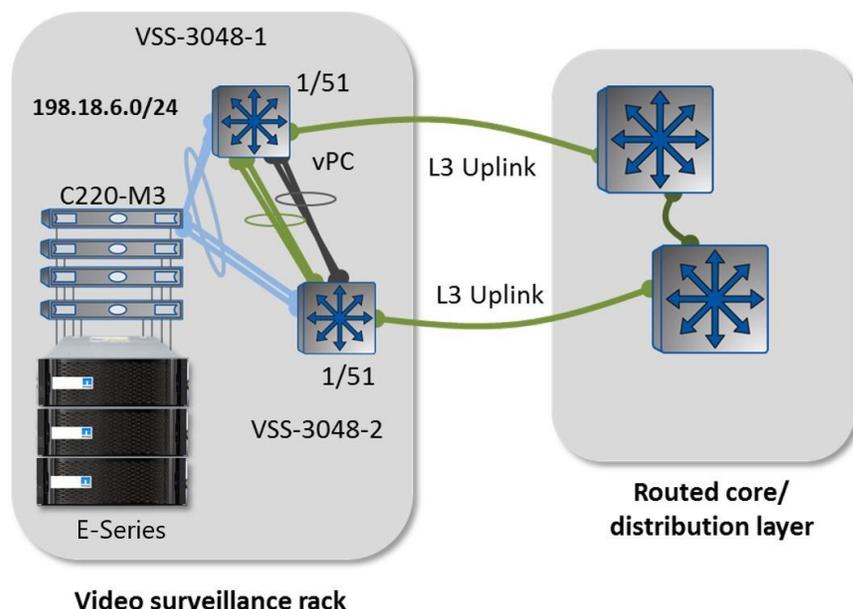
4.5 Uplink Connectivity (Layer 3)

The end customer might require additional features not supported in the NX-OS system default (no license) and can purchase additional Cisco NX-OS software packages, the base license, and the LAN enterprise license. These packages include features such as IP multicast support (IP PIM-SM) or advanced layer 3 routing such as OSPFv2 or EIGRP.

For example, a routed server access layer can be implemented with these optional licenses as described in the Cisco document [High Availability Campus Network Design—Routed Access Layer using EIGRP or OSPF](#).

If the end customer requires layer 3 connectivity to the video surveillance deployment, the topology is as shown in Figure 8.

Figure 8) Uplink connectivity (layer 3).



4.6 Network Management Caveats

The Cisco Nexus NX-OS does not include software support for the domain name system (DNS) server or Dynamic Host Configuration Protocol (DHCP) server. These services must be provided by the end-customer network management systems if desired.

As a best practice, the Cisco Nexus 3000 series switches should be configured to log to a syslog server and respond to SNMP queries as well as send SNMP traps to a network management workstation.

4.7 Network Design Rationale

The sample Cisco Nexus switch configuration implements these best practice network design concepts, as described in Table 2.

Table 2) Best practice network design concepts.

Design Decision	Explanation
VLAN security best practices	This application note provides details about VLAN security best practices. The concepts are incorporated in this design.
VLAN 1	In many deployments, VLAN 1 spans multiple switches and is not bounded by pruning from trunk ports between switches. Ports in this deployment are not assigned to VLAN 1, and the VLAN 1 SVI is shut down.
VLAN 2 unused ports	In the sample topology, VLAN 2 is defined and configured for all unused ports on the switches. VLAN 2 is not permitted on trunk ports. If a rogue user attaches to a switch port that is unused, the connected device does not have ready access to the network outside the local switch. Additionally, disabling unused ports is recommended.

Design Decision	Explanation
VLAN 3 native VLAN	VLAN 3 is designated as the native VLAN for this deployment. A native VLAN is the untagged VLAN on an 802.1q trunked port. The native VLAN in this topology is only configured on trunked ports. VLAN 3 has no edge ports.
VLAN 7 device management	This deployment utilizes a VLAN designated for managing the E-Series controller ports and the three management interfaces of each server. There are also available unused ports configured in VLAN 7 for service personnel to attach a laptop to a port for initial installation and ongoing troubleshooting. This VLAN is trunked using a layer 2 uplink or through layer 3 connectivity to the network core.
VLAN 58 virtual port channel peer keep-alive VLAN	VLAN 58 is designated as the virtual port channel (vPC) peer keep-alive VLAN. A port channel 58 interface and SVI interface are configured on both switches with two 1G member ports. An IP network address is assigned to the SVI interfaces and used as the source and destination IP addresses for vPC keep-alives. The switch management interface (mgmt0) is not used for vPC keep-alives, allowing the end user to connect the management interface to other devices in the network topology to manage the switches out of band.
VLAN 2020 video ingress	VLAN 2020 is designed to transport IP video surveillance network traffic from video surveillance cameras to the recording servers. The video management software management server virtual machines are also assigned interfaces on this VLAN. Each ESXi host has four 1Gbps Ethernet aggregated lines configured on a virtual switch and connected to a port channel on the Cisco Nexus switches with four member ports. Two of the four links are attached to each Cisco Nexus switch and are associated by a common vPC number.
Virtual port channel	A vPC allows links that are physically connected to the two Cisco Nexus switches to appear as a single port channel to a third device. The third device in the deployment is a Cisco UCS server with a quad-port Broadcom Ethernet adapter. Additionally, if the deployment utilizes layer 2 uplinks, these are also configured as vPCs.
Virtual port channel peer links	The vPC peer link is a port channel with two 10Gbps member interfaces, per Cisco best practices. The vPC peer link carries control traffic between two vPC switches, multicast, broadcast, and in some instances unicast traffic.
Port channel load balancing	The default port channel load-balancing hash uses the source and destination IP to determine the load-balancing algorithm used across the interfaces in the port channel. The default configuration is suitable for most deployments. This load-balancing algorithm may be changed as required.

Design Decision	Explanation
Routed server access layer	<p>This configuration illustrates using either layer 2 or layer 3 uplinks to the network core. Layer 3 features require additional license files to be purchased and installed on each switch. The vPC is part of the system default license. The base license (N3K-BAS1K9) includes limited layer 3 and IP multicast that has some application in video surveillance deployments. The LAN enterprise license (N3K-LAN1K9) includes all the layer 3 routing features plus virtual routing and forwarding lite (VRF-Lite).</p> <p>Installing the LAN enterprise license on the video surveillance solution switches allows a more defined demarcation between the rack and the core/distribution network switches. The same VLAN numbering scheme can be used on all rack switches because of the layer 3 demarcation. Troubleshooting network connectivity problems might also be easier because the default gateway addresses (HSRP virtual addresses) are configured on the Cisco Nexus switches and not on the network core/distribution switches. The layer 2 spanning-tree domain only encompasses the two Cisco Nexus switches.</p> <p>As a best practice, NetApp recommends implementing a routed access layer. For more information, see High Availability Campus Network Design—Routed Access Layer using EIGRP or OSPF.</p>

5 Server Planning

Physical security integrators have traditionally looked at servers as a commodity item. 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 iSCSI connection. 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.

5.1 Hardware Recommendations

At a minimum, 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 usually general recommendations, for example, dual-core Intel Xeon (quad core recommended) or Intel Core i5 or better, rather than specifying an exact model or clock rate. CPU specifications change too frequently and have far too many derivations for exhaustive testing of each model.

5.2 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. The CPU utilization will be 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 to more fully utilize

CPU cycles. Although there might be additional costs associated with licenses for the hypervisor, these costs might be offset by more efficient use of resources.

5.3 Server Manufacture

For solution validation testing, Cisco UCS C220 rack servers as well as Fujitsu PRIMERGY servers have been used.

One advantage of selecting a preferred system from the same manufacturer is the support synergy. Each manufacturer has its own management interface; for example, Fujitsu ServerView Remote Management iRMC or Cisco integrated management controller is used for management and monitoring of the system. Implementing systems from multiple vendors means additional support costs associated with learning multiple management interfaces.

5.4 Memory

In validation testing, memory is not a limiting factor. 16GB or less of RAM per recording server is sufficient, as recommended by the video management software.

5.5 General Design Criteria

The following items represent general design criteria when selecting a server hardware platform:

- Sufficient main memory to support the virtual machine requirements (for example, 8GB RAM per virtual machine; using four virtual machines, 32GB total RAM)
- Quad-core or better CPU per recording server/virtual machine
- Integrated Ethernet adapters and PCI-based quad-port 10Gbps Ethernet or 25Gbps Ethernet for video ingress and optionally IP SAN connectivity
- Dual internal disk drives configured as a RAID 1 virtual drive (internal RAID controller) for high-availability boot drive
- Form factor: 1RU for space savings
- Dual power supplies for resiliency
- Embedded server management to provide a remote virtual KVM and power cycle/reset capabilities

6 Design Checklist

To design the system properly, a myriad of factors must be considered to address customer requirements in an efficient and cost-effective manner. Table 3 represents some of the high-level considerations that must be examined to select the best components.

Table 3) Design checklist components.

Design Element	Description
Aggregate video data rate	The number of cameras and the resulting aggregate data rate must be determined to estimate the number of recording servers and the type and size of the storage array.
Video management system software	The architecture of the VMS determines the workload requirements of the storage array.
End-user requirements	Systems implemented for public sector deployments might have dramatically different workloads. Deployments with a high percentage of viewing video might require more servers and different volume layouts than systems with little forensic review of video.

Design Element	Description
Local support	This refers to the geographical location and local support staff. Readiness of onsite support staff might determine the number of hot spare disk drives or influence the decision to deploy traditional volume groups or Dynamic Disk Pools.
High availability	The costs associated with download or video loss might be more of a consideration in some deployments than others. Implementing a highly available design mitigates outages, but increases the cost and complexity of the deployment.
Host interface considerations	The number of servers required influences the choice of host interface to the storage array. Direct connect serial-attached SCSI (SAS) provides high throughput but is limited by distance and scalability. Fibre Channel is costlier, but provides high throughput and reasonable cabling flexibility. iSCSI provides acceptable throughput at low interface costs and has no practical distance limitation.
Retention requirements	The video retention policy is a key component to sizing the system and has a direct influence on the performance characteristics of the system.
Network requirements	The additional network routers and switches required to support the implementation must address the high-availability requirements, the type of host interface connectivity (IP SAN requirements), and the readiness of the existing customer network. Video that is lost between camera and server is never archived.
Type of servers	Deployments that implement recording servers in a virtual machine have different server requirements compared to deployments in which the host operating system is installed on a physical server.

Summary

Video surveillance deployments using 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. Detailed planning and careful system design are necessary for successful implementation.

Definitions

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

Table 4) Glossary.

Term	Definition
Controller	The controller is composed of the hardware board and firmware that manage the physical disk drives and present that capacity to a computer as logical units (LUNs).

Term	Definition
Dynamic Disk Pools (DDP)	DDP 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, making sure of 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)	The 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.
HDTV	High-definition TV defines resolutions of 1920x1080 and 1280x720 pixels along with other criteria, including aspect ratio.
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 is a small form factor IP networked Linux host that encodes and transports video over an IP network.
Layer 2	The data link layer in the open systems interconnection (OSI) model. Primarily associated with LAN switching network functions.
Layer 3	The network layer in the OSI model. Primarily associated with routing network functions.
LUN	The logical unit number is an address number for how the server identifies different hard drives or, in the case of storage systems, different volumes. Most operating systems show the LUN as properties of the SCSI hard drives discovered.
Megapixel	This is any video resolution of 1 million pixels or more. The HDTV resolution of 1280x720 is 921,600 pixels but is commonly referred to as a megapixel resolution.
Port channel	Also known as EtherChannel link aggregation. A port channel bundles individual Ethernet interfaces into a logical group. It increases bandwidth by load sharing on the member links. The port channel is operational with only one active member link.
RAID	RAID is an acronym for redundant array of independent disks, and it determines how data is protected from hard drive failures.
RAID 5	A striped disk with parity, RAID 5 combines three or more disks in a way that protects data against 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
RAID 10	RAID 10 provides high availability by combining 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.
SAS	Serial-attached SCSI (SAS) is a computer bus used to move data to and from computer storage devices such as hard drives and tape drives. SAS depends on a point-to-point serial protocol that replaces the parallel SCSI bus technology.
Storage array	The storage array is a collection of both physical components and logical components for storing data. Physical components include drives, controllers, fans, and power supplies. Logical components include volume groups and volumes. The storage management software manages these components.
Viewing station	This is a high-end workstation for displaying live or archived camera feeds on a locally attached monitor.
Virtual routing and forwarding (VRF)	Refers to multiple instances of a routing table in a layer 3 switch or router. It is a virtual routing table. When logged on the switch or router, the administrator might need to specify the VRF to use for commands such as ping or traceroute.
VMS server	Video management system server is also referred to as network DVR recording server. 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 of 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-4198: Video Surveillance Solutions with NetApp E-Series Storage: Performance Considerations
<http://www.netapp.com/us/media/TR-4198.pdf>
- TR-4199: Video Surveillance Solutions with NetApp E-Series Storage: Sizing Considerations
<http://www.netapp.com/us/media/TR-4199.pdf>
- Guided Solution Sizing
<https://fieldportal.netapp.com/content/204292>
- Video Surveillance Storage Solution Page
<https://fieldportal.netapp.com/content/211536?assetComponentId=211635>

- Cisco Design Zone
<http://www.cisco.com/go/designzone>
- Network Readiness Assessment for IP Video Surveillance
http://www.cisco.com/en/US/docs/solutions/Enterprise/Video/IPVS/IPVS_Network_Assessment.html
- OnSSI Hardware Recommendations
www.onssi.com/hardware-recommendations

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 disk models

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

Copyright Information

Copyright © 2013-2017 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.