



White Paper

Consolidate Video Surveillance Storage with E-Series

See more. Do more.

Tish Best, Jeff Fowler, Lester Cheng, Kameron Kemnitz, NetApp
March 2020 | WP-7313

Abstract

The latest camera technologies are causing the amount of data generated by video surveillance workloads to skyrocket. Most video surveillance storage solutions require multiple appliances to meet performance and storage demands, and the architectures can be unreliable and complex to monitor and manage. NetApp® E-Series enterprise SAN storage has the performance to capture video from more cameras with better quality while enabling powerful analytics that create value from this immense collection of data. The simplicity with which E-Series SAN storage can scale offers the flexibility to meet today's needs and handle the technology of tomorrow.

TABLE OF CONTENTS

1	Video Surveillance Is Hard on Storage	4
1.1	Camera Coverage	4
1.2	Image Quality	4
1.3	Video Retention	5
1.4	Compliance and Total Growth	6
2	Consolidate on E-Series Storage	6
2.1	Capacity Scale	7
2.2	Performance Scale	7
2.3	Simple and Manageable	8
3	Get Uninterrupted Recordings on E-Series Storage	9
3.1	Reliable and Available	9
3.2	Performance During Drive Failure	10
3.3	Proven Solution	10
3.4	Enterprise Features	10
4	Future Proof with E-Series Storage	11
4.1	Future Scalability	11
4.2	Extending Video Surveillance Storage	12
4.3	Containerization	12
4.4	Cloud Ready	12
	Where to Find Additional Information	12

LIST OF TABLES

Table 1)	Storage requirements based on camera counts, at 2.34Mbps for 60-day retention (2.34Mbps based on 1080p, low complexity, 13 fps).	4
Table 2)	Storage requirements based on resolution; 500 cameras for 60-day retention (data rate using low complexity, 13 fps).	5
Table 3)	Storage requirements based on frame rates; 500, 1080p cameras for 60-day retention (data rate using low complexity).	5
Table 4)	Storage requirements based on change in retention; 1080p cameras at 2.34Mbps (data rate using low complexity, 13 fps).	5
Table 5)	Number of NVRs to support data rate of medium motion camera at 25 fps.	8
Table 6)	Capacity, 8-drive RAID 6 NVR compared to E-Series DDP, 8TB drives.	9
Table 7)	Performance and rebuild duration, 8TB drive, sustained 800Mbps write workload.	10

LIST OF FIGURES

Figure 1) Data rate growth as fps changes for 1080p cameras.....	6
Figure 2) Capacity growth with increasing camera counts.	7
Figure 3) Data rate growth with increasing camera counts.....	8

1 Video Surveillance Is Hard on Storage

Recording capabilities are rapidly evolving to provide a complete picture of everything that is happening around us. The ability to see everything offers better security, and it has opened new possibilities for using video. Manufacturing processes use video surveillance to monitor quality and productivity, shopping patterns are tracked to increase retail sales, and stadiums use video images for targeted marketing. Video surveillance can be mobile, assisting law enforcement with body-worn and dashboard cameras, and video is an integral part of decision making in autonomous vehicles. But when cameras capture video at the quality needed for these use cases, they generate enormous amounts of data compared to typical video surveillance solutions. For these solutions to work, all of that data needs to be received and stored reliably. Video surveillance solutions today require much more robust storage systems than ever before.

1.1 Camera Coverage

Before camera analytics, it made sense to limit the number of cameras to what could be manually monitored and reviewed. Now that event detection is more sophisticated, and camera analytics are being used for monitoring, more cameras can be continuously monitored, resulting in better security. However, even the best cameras can't record what they can't see. Site surveys recommend higher camera counts for better coverage. Overlapping fields of view create redundancy to ensure that images are available when they're needed and being able to see a person from every angle can be very useful. Table 1 shows the data rates and storage capacities for various camera counts.

Table 1) Storage requirements based on camera counts, at 2.34Mbps for 60-day retention (2.34Mbps based on 1080p, low complexity, 13 fps).

Camera Count	Data Rate (Mbps)	Capacity (TB)
50	120	80
100	240	160
300	710	460
500	1,170	760
1,000	2,340	1,520
5,000	11,700	7,590
10,000	23,400	15,170

As shown in Table 1, capacity and data rate clearly increase proportionally with camera count. So, although more cameras give better coverage, they also increase the storage requirement.

1.2 Image Quality

Users expect high-definition images when viewing any video. Video surveillance must provide that same quality, or better, not just for people but also for automated analysis. Video surveillance is moving from simple detection of people to recognition and even identification. According to [Axis Communications](#), moving from detection to recognition typically requires at least 5 times as many pixels. Video quality is affected by both resolution and frame rate. Lower resolutions provide blurry images that are good only for detection, and lower frame rates can make motion appear choppy, missing important details. Table 2 shows the data rate and storage capacity of various resolutions, and Table 3 shows the same data for varying frame rates.

Table 2) Storage requirements based on resolution; 500 cameras for 60-day retention (data rate using low complexity, 13 fps).

Resolution	Per Camera Data Rate (Mbps)	Total Data Rate (Mbps)	Capacity (TB)
720p	1.04	520	337
1080p	2.34	1,170	758
4K	9.98	4,990	3,240

Table 3) Storage requirements based on frame rates; 500, 1080p cameras for 60-day retention (data rate using low complexity).

Frames Per Second	Per Camera Data Rate (Mbps)	Total Data Rate (Mbps)	Capacity (TB)
5	0.76	380	250
13	1.98	990	650
20	3.05	1,530	990
25	3.81	1,910	1,234
30	4.57	2,285	1,481

As expected, higher-quality images and higher frame rates generate more data.

1.3 Video Retention

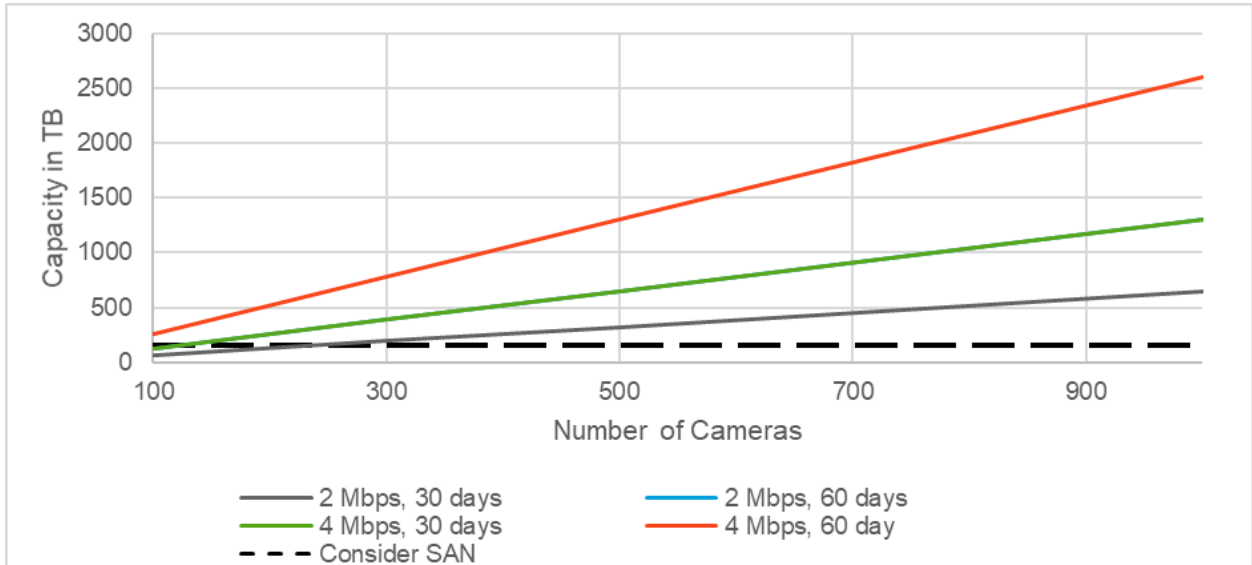
There are several reasons to store video after it's recorded. Often, events aren't reported until after they have occurred; safety and security solutions \ require storing video to include that window. Also, many policies for government, law enforcement, corporations, and gaming require video to be retained longer.

Video can also be a valuable asset for companies. Especially with the power of analytics, companies can monetize their data by looking for trends. Machine learning and artificial intelligence applications are constantly improving, and it can be valuable to have large datasets available for long periods of time. For the use cases of evidence, compliance, and analytics, high video quality is often required. However, increasing the retention rate linearly increases the amount of storage capacity required.

Table 4) Storage requirements based on change in retention; 1080p cameras at 2.34Mbps (data rate using low complexity, 13 fps).

Retention (Days)	Capacity (TB) 100 Cameras	Capacity (TB) 300 Cameras	Capacity (TB) 500 Cameras	Capacity (TB) 1,000 Cameras	Capacity (TB) 5,000 Cameras	Capacity (TB) 10,000 Cameras
15	40	120	190	380	1,900	3,790
30	80	330	480	760	3,790	7,580
60	160	460	760	1,520	7,580	15,160
120	310	910	1,520	3,040	15,160	30,330
365	930	2,780	4,620	9,230	46,120	92,240

Figure 1) Data rate growth as fps changes for 1080p cameras.



1.4 Compliance and Total Growth

Requirements for government, corporation, and gaming compliance continue to evolve. In many cases, the use of video has only just begun. Changes aren't isolated to only camera coverage, image quality, or video retention. Multiple requirements are likely to increase simultaneously, placing increasing pressure on the storage. Each change compounds the next, making the required performance and capacity expand faster than the changes illustrated in Figure 2.

2 Consolidate on E-Series Storage

In designing a video surveillance solution, the desired video quality, coverage, and retention determine the data rates and capacity that the storage solution needs to provide. It is difficult to meet these requirements using network video recorders (NVRs), which combine storage and video management services into a single device, because they lack the flexibility to match changing requirements. The increasing demands of video surveillance often require multiple NVRs to meet either the incoming data rate requirements, the capacity needs, or both. Because NVRs have fixed hardware resources, it is very challenging to design solutions with multiple NVR devices that don't result in unused hardware resources.


When implemented, video surveillance architectures with multiple NVRs are not flexible enough to support an increase in motion rates, camera coverage, frame rates, retention periods, or other expansions to the video surveillance capabilities that weren't part of the initial installation plan. Requirement changes involve not only adding more NVRs, they can also result in a cumbersome exercise in rebalancing cameras across all NVRs to use the hardware more efficiently and/or to organize the camera-to-NVR mappings.

These struggles have driven IT professionals to consolidate storage for several years now. If the architecture is not consolidated, the total cost of ownership of distributed licenses, rack and power, administration, and installation are higher. If the solution requires more than a few NVRs, consolidated SAN storage can solve these problems.

2.1 Capacity Scale

At low camera quantity and quality, NVRs are limited by their capacity. Table 4 shows how changing retention policies can dramatically increase capacity requirements, especially when compounded with changing camera quantity. Figure 2 shows capacity growth with increasing camera counts.

Figure 2) Capacity growth with increasing camera counts.



A single NetApp E-Series system scales to 6PB and can replace the storage capacity of more than 30 NVRs. The drives of most NVRs are not easily accessible, and they don't have vacant drive slots to add more drives for additional capacity. Therefore, additional capacity must come from additional hardware in the form of another NVR or attached storage. Although adding storage to an NVR can help scale the storage capacity, the additional storage behind the NVR does not increase the data rates at which the NVR can consume live streams, and more NVRs may still be needed.

2.2 Performance Scale

At higher camera quantities and quality, NVRs can be limited by their ingest rate. E-Series systems offer more than 30 times the data ingest capacity of NVRs. As data rates increase due to increasing image quality and coverage, the storage capacity of a single NVR becomes irrelevant because the NVR's ingest rate limit is encountered first. Most NVRs have one ingest port that has a maximum speed of 1Gbps, whereas E-Series SAN storage systems have four 10Gbps ports. Figure 3 shows the required data rates for various camera counts at different resolutions; notice that rates quickly reach 1Gbps and beyond.

Figure 3) Data rate growth with increasing camera counts.

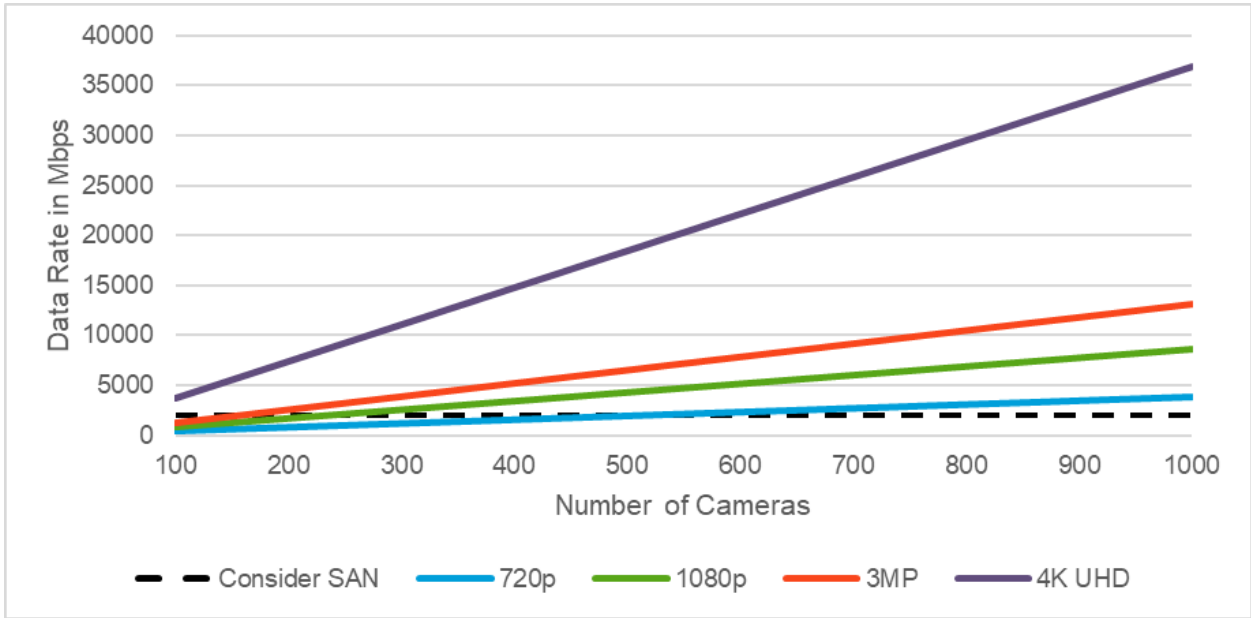


Table 5 shows how many NVRs it takes to support the necessary data rates for these installations. Most installations can support all of their recording requirements with a single E-Series system, but even solutions with greater than 6PB systems that require multiple E-Series systems can be managed from a single pane of glass, supporting up to 270PB of storage.

Table 5) Number of NVRs to support data rate of medium motion camera at 25 fps.

Resolution/ Cameras	720p 1280x720	1080p 1920x1080	3MP 2048x1536	4K UHD 4096x2160
100	1	1	2	4
300	2	3	4	12
500	2	5	7	19
1,000	4	9	14	37
5,000	20	44	66	185
10,000	39	87	132	370

2.3 Simple and Manageable

Because E-Series storage systems can handle such high data rates, architecting solutions requires only the storage capacity to be calculated. It's no longer necessary to balance data rates and calculate capacity needs for cameras with various capabilities and retention rates at each NVR location. There's no need to plan (and pay) for all the resources you might need on day 1 to avoid costly rearchitecting when plans change. E-Series systems can be installed with the capacity you need today with the confidence that it can meet any performance demands.

If capacity needs increase tomorrow, it's easy to add more storage to the existing system, in increments as small as a single drive or 60 drives all at once, with a single click. Drive additions are dynamic and are completed without disrupting the video surveillance workloads that are online. Consolidating video surveillance storage on one centralized system reduces complexity and improves scalability. NVRs scale

by incrementally adding independent, fixed-capacity NVRs. These incremental additions result in video data that is separated across NVRs, which creates many challenges. Some NVRs support clustering, which can be complex to configure and create issues with reliability when infrastructure between NVRs has outages. Adding storage behind an NVR can relieve the capacity scale, but it might not meet the data rate requirements.

Scaling capacity separately from video management services reduces waste of unused hardware, and fewer systems to secure and manage means reduced administrative costs. Consolidating multiple NVR architectures to a single E-Series storage system means that the solution requires fewer drives to monitor and fewer replacements of failed drives. In addition to the cost savings of the drives themselves, fewer drives to monitor and fewer replacements of failed drives results in further reduced costs. Table 6 shows the required NVRs and their drive counts for various usable storage capacities using RAID 6 compared to NetApp E-Series Dynamic Disk Pools (DDP).

Table 6) Capacity, 8-drive RAID 6 NVR compared to E-Series DDP, 8TB drives.

Usable Capacity (TB)	# of NVRs	NVR Drives	E-Series Drives	E-Series Savings
100	3	24	18	25%
200	5	40	35	13%
400	10	80	68	15%
800	19	152	135	11%

3 Get Uninterrupted Recordings on E-Series Storage

Video surveillance solutions demand more than just performance from their storage. Uninterrupted surveillance translates to enterprise reliability and availability, even when components fail. Hardware failures without redundancy can cause lengthy outages that customers simply cannot tolerate. NVRs typically contain consumer-grade components that are more likely to fail than enterprise-grade components. When a fan fails on an NVR that does not have redundant cooling, the NVR will probably begin to overheat and shut itself down. Cameras recording on that NVR can no longer record any video until the failed component is replaced (if it is field replaceable), or the entire NVR is replaced. Outages that cause data collection to be halted are costly for customers who are using that data for business decisions, and in some cases can put safety and security at risk.

Any hardware component can fail, and video surveillance solutions with tens or hundreds of hard drives can expect several drive failures per year. NVRs and other storage solutions that rely on consumer-grade drives are likely to experience even higher failure rates because consumer drives have a 50% greater failure rate than enterprise drives (according to [Backblaze](#)). Most systems use parity to provide redundancy for drive failures to prevent data loss. However, the performance of a given system is often determined with the assumption that all drives are optimal. Maintaining performance when a failed drive is replaced is challenging.

Systems often consume all their bandwidth using parity to rebuild the data that existed on the failed drive, causing recording and playback capabilities to be severely degraded. This degradation can cause data drops to the point where customers stop recording on cameras due to the unacceptable quality of the video collected. Furthermore, systems can take days to rebuild the data on a failed drive, creating a large window when the system is both degraded and at risk of a second failure.

3.1 Reliable and Available

E-Series systems are built with enterprise-class hardware, including dual-port enterprise drives with additional protection information (T10 Standard Data Integrity) and full-disk encryption (FDE) to ensure

data integrity during every read and every write and data security at rest. The storage arrays provide fully redundant, highly available, online-replaceable hardware and support online firmware upgrades with security updates. For example, when a component, such as a fan, fails, the redundant fan keeps the system cool so that video recording can continue without interruption. The NetApp SANtricity Storage Manager manages system alerts so the administrator is informed about component failures. When the replacement is available it can be hot swapped while the system is online and fully functional. No outages, no disruptions. E-Series drive reliability goes beyond reacting to a failed drive by proactively monitoring the drive to identify potential issues. Predicted drive failures will result in automatic recovery procedures that can prevent the drive from needing to be replaced before it is truly necessary, saving money and time. With Network Video Recorders (NVRs), reacting to a drive failure with a notification may not even exist and failure may go undetected for some time.

3.2 Performance During Drive Failure

Similar to NVRs, E-Series systems use parity to provide cost-effective redundancy for drive failures. However, E-Series systems contain a proprietary solution using Dynamic Disk Pools that allows the replaced drive to rebuild its data at a much faster rate without degrading the performance of the system. As shown in Table 7, video recording and playback performance are not affected by disk rebuilds, as with NVRs.

Table 7) Performance and rebuild duration, 8TB drive, sustained 800Mbps write workload.

System	Performance Difference During Drive Failure	Performance Difference During Drive Replacement	Rebuild Duration (Days)
E-Series	0%	0%	3.5
NVR	-15%	-25%	6

3.3 Proven Solution

E-Series storage systems are certified with major VMS solutions, including Milestone and Genetec. NetApp focuses on quality and interoperability of all video surveillance components to ensure the best complete solution. Certifications and integration with video surveillance are a constant focus of the NetApp Video Surveillance Center of Excellence.

3.4 Enterprise Features

E-Series storage takes reliability a step further with built-in intelligence, such as detecting misbehaving drives, a feature that saves money and administrative cost by preventing drive failures. Enterprise-grade drives are much more reliable than consumer-grade drives, but they still aren't bullet proof. Through years of drive failure analysis, E-Series systems can often solve intermittent drive issues with a power cycle. When the NetApp SANtricity® operating system encounters a drive it thinks is misbehaving, it power cycles the drive, probably revitalizing it and preventing the need for a full drive reconstruction and avoiding hours or days of increased risk to data integrity. NetApp E-Series Dynamic Disk Pools (DDP) is unique and goes beyond the standard data parity protections that are found in Network Video Recorders (NVRs) and other enterprise SANs. Dynamic Disk Pools have spare capacity designed into them. The spare capacity allows automatic rebuilds when drives fail, significantly reducing the risk of data loss when multiple drives fail in a short amount of time compared to standard RAID implementations. Network Video Recorders (NVRs) don't have hot capacity ready to repair itself and therefore must rely upon administrative action to locate, order, and manually replace the failed drive before the system can become optimal once again.

NetApp E-Series Storage System is managed by the NetApp SANtricity System Manager. The NetApp SANtricity System Manager comes on-board with the E-Series Storage System. It enables users to

manage all E-Series Storage Systems anytime and anywhere. Launchable directly from browsers, the SANtricity Storage Manager features includes the following features.

- Automated workflows and intelligent provisioning defaults
- Application tagging to enable filtered views of volumes and performance information
- Enhanced performance monitoring and access to 30 days of performance data
- Performance-tuning actions
- Support for no-downtime reconfiguration, maintenance, and software and firmware upgrades
- Recovery Guru with step-by-step instructions

NetApp AutoSupport® technology provides risk assessment and communication via At Risk System (ARS) to help prevent problems or mitigate their impact. It leverages sophisticated monitoring for faster incident management, and more importantly the ability to resolve issues before they cause a problem. It also provides automated hardware replacement for selected components such as drives to simplify management and enhance reliability.

In addition to data reliability and availability, NetApp is committed to data security and confidentiality. A major vulnerability to data security is through the management of the storage solution which ultimately controls access to the video data. The NetApp SANtricity System Manager includes the following security features.

- Administer access with role-based access control and audit log
- Manage digital certificates (import, export, CSR)
- Set up LDAP/LDAPS server for user authentication
- Manage data-at-rest encryption with FDE or FIPS drives, including Secure Erase
- Configure multifactor authentication via SAML .20
- Common Criteria certified (The Common Criteria certification is an international standard (ISO/IEC 15408) for IT Security Evaluation. For more details, go to the [NetApp Security Certification](#) page.

These and many more enterprise features provide a complete solution to meet the reliability demands of video surveillance.

4 Future Proof with E-Series Storage

Analytics and other opportunities to expand the value of video surveillance offer endless ways in which marketing, retail, and manufacturing, to name a few, can monetize previously recorded video data. Analyzing how shoppers navigate through a store, or where productivity or quality faltered in the manufacturing process, can lead to new insights to increase monetization. If the video streams are spread out across multiple NVRs, this level of analytics would be difficult or impossible. Performing analytics on stored video data creates a demand for storage systems with a very different workload than systems that store video streams alone. Storage solutions need to support analytics workloads without disrupting surveillance and playback capabilities.

4.1 Future Scalability

E-Series systems offer the performance and capacity scalability to support 4K cameras, compliance with longer retention policies, increased camera coverage, and other major changes to video surveillance requirements. A single E-Series system scales to 6PB with a 3.7GBps maximum data rate. The 6PB can be added in increments as small as a single drive or 60 drives all at once. As capacity is added, no additional license is required, making scaling cost efficient. As more analytic capabilities become available, today's solutions will need to support recording higher resolution cameras along with the read performance for analytics on previously recorded data. Customers can be confident that a consolidated E-Series storage system will meet their future demands for video surveillance performance.

4.2 Extending Video Surveillance Storage

As the demands and importance of video surveillance data increase, new customer needs constantly arise. E-Series storage currently has options for cloud archiving, backup, and disaster recovery solutions and is continuing to innovate to solve new challenges in the market. E-Series systems can go beyond current video surveillance challenges; they can also provide capacity for any other applications that can benefit from high-performing, low-cost, reliable storage.

4.3 Containerization

High-quality video images have made data an asset in video surveillance, opening new opportunities to create value. Typical video applications run on a camera with limited resources or on an external server that accesses the video images remotely. The SANtricity operating system supports Docker containers, allowing applications to eliminate the additional data movement and have direct local access to the video recordings of all cameras. This solution enables innovation that can truly differentiate systems in video surveillance.

4.4 Cloud Ready

With NetApp Cloud Connector, data that is saved on an E-Series system can be easily backed up and restored to an existing Amazon S3, NetApp StorageGRID®, or NetApp Cloud Backup account. This solution offers customers an integrated software alternative to begin their cloud integration journey.

Where to Find Additional Information

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

- NetApp Video Surveillance Storage (VSS) Solutions
<https://www.netapp.com/vss>
- NetApp Product Documentation
<https://docs.netapp.com/>

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 © 2019–2020 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.

WP-7313-0220