



ARCHIVE SOLUTIONS WITH DELL EMC ISILON SCALE-OUT NAS

ABSTRACT

This white paper demonstrates how a Dell EMC Isilon scale-out NAS cluster provides the scalability, efficiency, simplicity, and agility to fulfill the storage requirements of a large, centralized video surveillance system while reducing capital expenditures and operating expenses.

December 2016

The information in this publication is provided “as is.” DELL EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any DELL EMC software described in this publication requires an applicable software license.

DELL EMC², DELL EMC, the DELL EMC logo are registered trademarks or trademarks of DELL EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners. © Copyright 2016 DELL EMC Corporation. All rights reserved. Published in the USA. <12/16> <white paper> < H12546>

DELL EMC believes the information in this document is accurate as of its publication date. The information is subject to change without notice.

DELL EMC is now part of the Dell group of companies.

Table of Contents

EXECUTIVE SUMMARY	5
INTRODUCTION	5
A brief history of video surveillance technology	5
Streaming video data	7
Topologies for surveillance systems	7
STORAGE OPTIONS	8
The drawbacks of SAN and DAS	9
Scale-out NAS vs. scale-up NAS	9
Isilon scale-out NAS.....	10
SURVEILLANCE SYSTEM DESIGN AND DEPLOYMENT	10
A use case from a casino.....	13
Storage efficiency	13
Elasticity and agility	13
Failure handling	14
New technologies for new developments.....	15
CONCLUSION.....	15

EXECUTIVE SUMMARY

The global video surveillance industry is booming. A report on CNNMoney in April 2013 cites a research firm's estimate that the video surveillance market will expand from \$11.5 billion in 2008 to \$37.5 billion in 2015.¹ ABI Research, whose figures include storage, estimates that it will hit \$46 billion in 2013.² The rapid global expansion of video surveillance systems combined with longer retention times for video images and higher resolutions for video analytics require an ever-increasing amount of data storage, placing video streams and imagery squarely within the realm of Big Data.

The explosion in video data is also bringing changes in the design of video surveillance systems. Higher resolutions take more bandwidth to transmit the video streams and require more space to store the imagery. Larger deployments increase the level of virtualization required. And dynamic video systems demand storage solutions that scale—solutions that Dell EMC® Isilon® optimally provides.

Common data storage problems place limits on many centralized surveillance systems, especially those with more than 200 cameras. Traditional storage systems that statically map camera streams to volumes or logical unit numbers (LUNs) are time consuming to manage and difficult to expand. Storage-attached networks, direct-attached storage, and scale-up network-attached storage (NAS) are inefficient because they require the use of many disparate volumes, which increases management overhead. Aligning video streams with volumes reduces storage efficiency and drives up capital expenditures (CAPEX). The overall efficiency rate of traditional storage approaches that store surveillance video in LUNs or volumes is about 55 percent to 65 percent. In video surveillance systems, the management and storage of data contributes to about 30 percent of expenditures. Traditional storage systems seem attractive for use in video surveillance for their low entry cost but are marred by large operating expenses and a high rate of costs for underutilized capacity.

Dell EMC Isilon scale-out NAS solves these problems. With a distributed file system that presents a cluster of storage appliances as a global namespace, all camera streams in a video surveillance system point to a single, scalable volume of storage. When the addition of cameras or an increase in resolution requires more capacity or performance, an Isilon cluster seamlessly and non-disruptively scales out to more than 20 petabytes and 100 GB/s of throughput. The cluster ingests and distributes high-bandwidth video data by using standard application-layer network protocols. This white paper demonstrates how an Isilon scale-out NAS cluster fulfills the requirements of a large, centralized video surveillance system by providing simplicity, efficiency, agility, and scalability.

INTRODUCTION

Video surveillance systems are evolving to adopt IT-centric mechanisms such as the virtualization of servers, networking, and storage. Surveillance systems are increasingly segregating the data plane from the control plane, as the architectures of the systems are adapted to megapixel cameras, high-bandwidth video streams, and longer retention periods. The deployment strategies and network topologies of video surveillance systems lead to a number of options for storing video data. Many of these options have originated from the history of surveillance systems, the advent of different data transport protocols, the compression of video images, and the demands of emerging technology.

A brief history of video surveillance technology

In the mid-1980s, video surveillance systems were primarily analog systems that distributed video over coaxial cables. During the mid-1990s, surveillance systems migrated to Ethernet networks in order to distribute video data. Many of the system components, such as the cameras, recorders, and monitors, continue to employ analog technologies to this day.

During the first decade of the millennium, radical changes overtook older analog technologies. Distributed applications combined with networked infrastructure to shift the emphasis to purpose-built appliances for recording and viewing videos. The networked appliances were referred to as network video recorders (NVRs) that consisted of hardware packaged with

¹ "The great surveillance boom." CNN Money. April 26, 2013.
<http://management.fortune.cnn.com/2013/04/26/video-surveillance-boston-bombings/>.

² "Video Surveillance: A Market Poised for \$46 Billion of Explosive Growth." Digital Communities. March 18, 2008.
<http://www.digitalcommunities.com/articles/-Video-Surveillance-A-Market-Poised.html>.

software called a video management system (VMS). VMS was widely adopted for video surveillance systems. The primary function of the VMS is to map video spatiotemporally and to distribute it to viewing devices like video walls, monitors, mobile devices, and analytics systems.

Surveillance systems continue to mature. Administrators who maintain large surveillance systems are discarding purpose-built components to adopt virtual VMS servers and a more efficient and flexible infrastructure—a design that IMS Research refers to as a “private cloud.”³ During the past five years, administrators have been increasingly adopting real-time analytics and megapixel cameras. “Megapixel resolution network security cameras are predicted to out-sell standard resolution network security cameras in 2012 for the first time,” IMS Research says.³ In addition, a growing trend in forensic analytics is to gather intelligence from video libraries stored in archives. Some of the analytics is used for predictive purposes and is aligned with retail and health-care applications.

Older proprietary surveillance systems remain in wide use today. Some large-scale surveillance operations, like those used at casinos, are still supported by hundreds of digital video recorder (DVR) devices stored in racks that can number in the forties and fifties. Each DVR can support only a small number of cameras, and the racks housing the DVRs generate high operating expenses for power and cooling. The DVRs require multiple points of management, and such proprietary systems require tight coupling of hardware and software (server, VMS software, storage, and camera).

Older systems are also coming under pressure from expanding requirements, such as handling input from cameras that record high-definition video at 30 frames per second. High-definition video imaging demands between 1 and 5 terabytes of storage per camera for a standard 30-day retention period. Longer retention times increase the demand for more storage.

Surveillance system operators are also beginning to consider infrastructure as a service to create hybrid systems that manage some components locally while dispersing others into the cloud. The overarching trend to reposition information technology as a service is changing some aspects of video surveillance architectures when it makes economic and technological sense. Increasingly, open platforms with nonproprietary interfaces are decoupling software from the enabling infrastructure. For instance, many operators are replacing large systems of packaged hardware and software with platforms that combine servers running open software and storage platforms in an IP network.

Despite the interest in infrastructure as a service, typical surveillance architectures today remain closely tied to network topology because the system streams video between the cameras and the VMS over an IP network. Using Real-time Transport Protocol (RTP), the video is encoded with one of several codecs—MJPEG, MPEG-2, MPEG-4, H.264, and soon H.265.

Although block storage using Fibre Channel (FC), Fibre Channel over Ethernet (FCoE), or iSCSI remains common, surveillance systems are increasingly using more powerful application-layer protocols such as Server Message Block (SMB), network file system (NFS), and HTTP because they are highly efficient protocols that can support hundreds of cameras with each connection. Surveillance systems are switching to these protocols because they allow abstraction from the underlying physical hardware, whereas the previously mentioned block protocols do not.

The current state of video surveillance technology includes these dominant themes:

- Data management and storage is a vital element of most surveillance systems.
- Open, loosely coupled software and hardware are creating the need for virtualization.
- Mobile devices are becoming a de facto part of surveillance systems for viewing video and as sensors.
- Longer-term retention in video libraries and post-incident analytics are becoming more prevalent.

As of 2013, the primary trends in video surveillance include longer retention times, better graphics functions for clients, and higher resolutions. While higher resolutions increase the number of pixels for human and machine analytics, higher resolutions also require more processing power and bandwidth.

³ “Video Surveillance Trends for 2013.” IMS Research. December 18, 2012. http://www.imsresearch.com/press-release/Video_Surveillance_Trends_for_2013.

Streaming video data

Streaming video requires various amounts of bandwidth. When video streams from cameras to the VMS and then to the storage system, the resolution, frame rate, and scene dictate the average bandwidth. The information that follows associates average bandwidths with different resolutions to illustrate how much bandwidth each camera can consume with a busy scene.

Resolution	4CIF	1080p	3MPixel	5MPixel	10MPixel
Frame rate (FPS)	15	15	15	15	15
Average bandwidth (Mb/s)	1.5	6	7.7	9.6	21

Predicting the bandwidth that a camera requires, however, can be difficult. A camera's average bandwidth varies not only by resolution, frame rate, and scene but also by manufacturer and model due to the use of different video compression implementations. The primary video compression standard used in video surveillance, H.264, compresses video by sending frames that represent changed or changing pixels in a video sequence.

With H.264, the group of pictures (GOP) sequences (see Figure 1) define how the compression standard transcodes the video. A GOP sequence determines how well the transcoding handles packet losses and highly dynamic content in a scene. The nonstandard sequencing of the I, P, and B frames in the GOP contributes to the difficulty in predicting how much bandwidth a camera requires:

- I frame: Intraframe has no reference to any frame and is stand alone
- P frame: Interframe refers to an earlier P or I frame
- B frame: Bi-predictive interframe refers to earlier and future frames

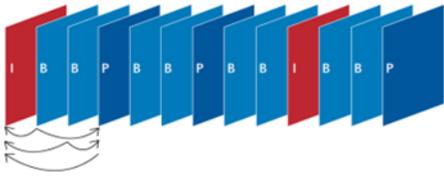


Figure 1: A GOP sequence

The camera manufacturer's implementation of H.264 causes a video stream's bandwidth to vary over time with changes in the scene. Given the same camera and different scenes shot in different light, the video bandwidth varies from 10 percent to 50 percent, making it difficult to predict the average network bandwidth of a camera's video stream. Streaming video from surveillance cameras represents a highly variable, high-volume data flow.

Topologies for surveillance systems

The demands of streaming video often dictate the topology of a surveillance network. Although the trend is to consolidate and simplify the components of a surveillance system's deployment, there are networking barriers that continue to block consolidation and simplification. Many private WAN connections, for instance, do not exceed 10 Mb/s, which is effectively equal to the bandwidth required by 5 to 10 standard definition IP cameras.

Despite such technologies as multiprotocol label switching, data center bandwidth becomes a linear scalar. Surveillance systems that are constrained by the bandwidth of WAN connections benefit from VMS appliances at each site. The following scenarios for the server, software, and storage components of a surveillance system typify the solutions:

- VMS appliance or NVR
- Bare metal servers with VMS software

- Virtual machines with VMS

In an ideal world, the scenario governs the technology that you deploy:

VMS appliance: When there are many distributed video sites, limited network access can create congestion. This scenario leads many system architects to deploy VMS appliances. At each site, the VMS software processes the video streams, which require some direct-attached storage for retention. To satisfy longer retention periods, the systems frequently transfer some of the video data to remote storage sites during network downtimes.

Bare metal servers with VMS software: Hardware procurement guidelines can force system administrators to host VMS applications on bare metal or commodity servers.

Virtual machines with VMS: During the late 2000s, the popularity of VMware and other hypervisors led VMS vendors to adopt virtualization to align with the IT-centric consolidation strategies taking place in data centers. At the same time, virtualization helped overcome the common limit of about 500 Mb/s per VMS server instance, or about 100 cameras per VMS server instance.⁴

Each of these approaches requires a storage system. Video piles up fast, and the volume and velocity of video imagery continues to accelerate. Table 1 provides an example of contemporary high-definition surveillance requirements.

Video scenario	Storage requirement
1080 pixels and 15 frames per second at 4 Mb/s	45 GB per day
A campus with 1,000 cameras	4.5 TB per day
Retention of 30 days	1.5 PB

Table 1: High-definition video storage requirements

Retention periods lead to tiering. Although a VMS server can implement tiers of storage and move data across the tiers, this legacy functionality stems from the days when VMS vendors had to develop a tiering solution because a lack of open APIs, standards, and data access protocols hindered the use of other, more efficient ways of managing the data.

Surveillance systems typically contain a video metadata database as well as the video files. The metadata database includes information mapping the video to a storage location, alerts, operator notes, temporal and spatial information, and other custom data. The video data is usually separate, and many VMS vendors tier the video data between the short term, often 14 days, and the long term, longer than 14 days.

Storage options

There are three primary methods for storing data from video management systems, which are shown in Table 2.

Storage option	Description
Direct-attached storage (DAS)	The VMS uses its on-server disks with an on-server RAID controller.
Storage area network (SAN)	The VMS connects to a SAN with Fibre Channel, Fiber Channel over Ethernet, or iSCSI.
Network-attached storage (NAS)	The VMS transmits the data to the NAS system with a protocol like SMB, NFS, or HTTP.

Table 2: The three primary storage methods for video management systems

⁴ The limit stems not only from CPU-intensive software that normalizes heterogeneous video streams but also from processing video in the user space of Microsoft Windows and Linux systems. With virtualization, a 2RU server, for example, can handle two to four VMS instances, which increases the number of cameras that a server can handle by 100 to 300 percent.

The drawbacks of SAN and DAS

The storage area network (SAN) and DAS approaches share some common problems:

- The static mapping of video feeds to each volume and physical disks requires multiple volumes with various groupings of video feeds mapped to them. With a storage area network, for instance, 1.5 PB might take 96 16 TB volumes. The loss of a volume results in a loss of video for playback. Because the bandwidths of cameras vary, hot spots can create a loss of playback and viewing as video feeds are migrated from one volume to another.
- Drive failure frequency and larger drive capacities requiring longer rebuild times can increase the risk of multiple, concurrent drive failures, compromising the availability of video surveillance content.
- For DAS, every volume has a limit set by the operating system. Although the limit depends on the operating system, the limit usually varies between 2 TB and 16 TB. Microsoft® Windows® Server 2008, for instance, is limited to 16 TB. For large surveillance systems with storage requirements over 500 TB, the volume limit means that a large number of volumes must be managed. More issues arise because most VMS manufacturers require contiguous volumes for each camera, which entails that a camera feed can write only to a single volume, resulting in high levels of inefficiency. Similarly, the need to defragment the volumes results in additional overhead.

Typical video surveillance architectures continue to combine an NVR with a traditional approach to storage. In such cases, the architecture statically maps IP cameras through NVRs to volumes or LUNs. Each volume, however, must be preconfigured for dedicated capacity, retention times, and the camera bit rate. Each volume also entails multiple points of management, making them time consuming to manage and even harder to scale. The efficiency rate of traditional approaches to storage that use LUNs or volumes is about 55 to 65 percent.

In addition to the complexity of multiple volumes, traditional storage systems forfeit a significant amount of capacity to overhead. Each volume in a video surveillance system is set up for dedicated capacity plus 20 percent overhead to handle the variability of the video streams, including unpredicted growth and changes to retention times. Each volume has a baseline overhead from the disk parity of RAID 5 or RAID 6 of 20 percent. Each volume also typically sets aside about 5 percent to align video files on contiguous LUNs. In total, traditional storage systems storing video data can lose 45 percent of their capacity to overhead.

The hundreds of video feeds that are recorded by the VMS require high bandwidth to transmit them through the network and into a storage system. The advent of network protocols such as SMB, NFS, HTTP, and Representational State Transfer (REST)—which satisfy high-bandwidth requirements—has opened up the use of NAS in high-volume video environments.

Scale-out NAS vs. scale-up NAS

There are two types of NAS systems: scale-up and scale-out. A scale-up NAS system is typically a node-based architecture with separate modules for storage and processing; each module can be added to a cluster independently. In contrast, a scale-out NAS system is a node-based architecture that combines processing power and storage capacity in every node; each node adds both performance and capacity to a cluster.

Both types of NAS systems benefit video surveillance systems for several reasons:

- Standard network protocols like SMB, NFS, and HTTP enable all operating systems—Linux, UNIX, Mac OS X, and Microsoft Windows—to interface with the storage system.
- IP-based storage systems deliver a rapid ROI when they store data for multiple applications because you can consolidate workflows into a single storage system.
- NAS abstracts the data protection schemes and underlying structures by using standard protocols for access and modification.
- NAS abstracts the physical location of the data, making video mappings to NAS volumes much less dependent on the data's physical location.

The primary issue with scale-up NAS systems is that they mimic traditional storage systems: Scale-up NAS systems and clustered scale-up NAS systems contain many volumes and require you to manually migrate data among components.

Isilon overcomes the problems that undermine traditional NAS systems by combining the three traditional layers of storage architecture—file system, volume manager, and data protection—into a scale-out NAS cluster with a distributed file system.

Isilon scale-out NAS

A Dell EMC Isilon cluster scales out, not up. In contrast to scale-up NAS, an Isilon cluster does not require volumes. Instead, the distributed Dell EMC Isilon OneFS® operating system combines the memory, I/O, CPUs, and disks of its nodes into a single, cohesive storage unit. Every node adds capacity and performance to the cluster. As nodes are added, the file system expands dynamically and redistributes data, eliminating the work of partitioning disks and creating volumes.

Isilon scale-out NAS adds capabilities that simplify storage for video surveillance systems:

- **A single volume.** An Isilon cluster presents all the nodes in the cluster as a single, global namespace. A VMS server can access the same volume and same data from any node in the cluster.
- **Seamless scalability.** An Isilon cluster's file system dynamically and non-disruptively expands to as much as 20 PB in a single volume.
- **Reliable, efficient data protection.** OneFS stripes data to guard it with parity blocks instead of parity disks. A larger cluster can lose up to four 36-drive nodes without a loss of video playback and without disrupting video recordings.
- **High-velocity, high-bandwidth video ingestion.** For performance, OneFS load balances network connections among all the nodes to eliminate bottlenecks. Surveillance networks can establish either dual 1 GbE or dual 10 GbE connections to the nodes.

An Isilon scale-out cluster's single volume is ideally suited to work with large, centralized surveillance systems. Scale-out NAS radically simplifies a surveillance system's design, operation, and augmentation. An Isilon cluster also protects against failure, delivers better ROI than other solutions, and supports open standards like REST for cloud-based access. For more information on the OneFS operating system, see "[Dell EMC Isilon OneFS: A Technical Overview.](#)"

Surveillance system design and deployment

A single-volume approach to designing a surveillance system simplifies the work for you as well as your system integrators and VMS vendors. But the single-volume approach must also let you scale the volume so that capacity and bandwidth scale almost linearly.

Some NAS systems do not scale linearly. When a single volume cannot transparently aggregate the capacity and bandwidth across all the storage units, you must fragment the VMS to properly segment the volumes. The NAS system ends up with multiple volumes, resembling a traditional DAS or SAN approach. The result is many storage volumes for each VMS instance—and even more for the whole surveillance system. The system bogs down in complexity because you must group video feeds by volume, as Figure 2 illustrates.

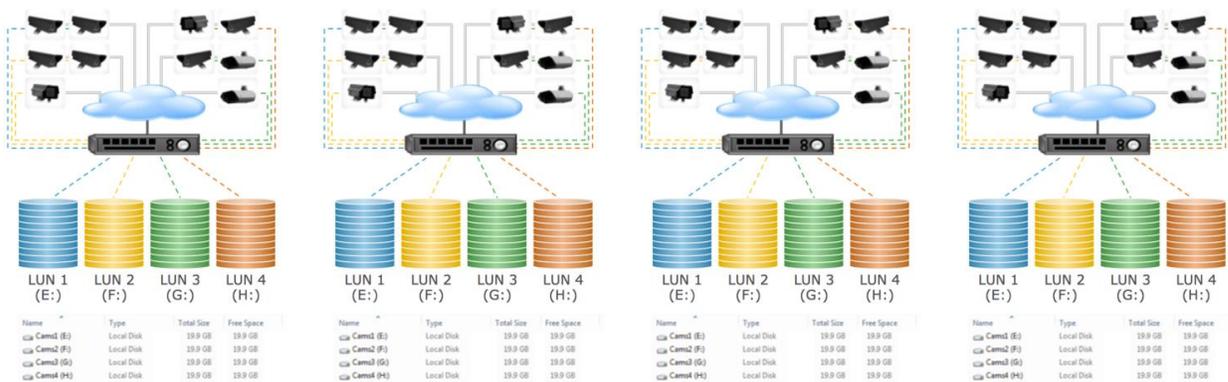


Figure 2: Surveillance system deployment with a traditional storage system

In contrast, the elegance of the Isilon solution lets you point all the video feeds at a single volume, which the VMS servers can access by connecting to any node. The single volume lets you size the storage system's aggregate bandwidth and capacity without resorting to complex calculations involving the number of volumes.

The cameras deployed in large surveillance systems often require more bandwidth than the average bandwidth specified in the system's design. Actual bandwidth tends to exceed the average expected bandwidth because scenes change and because H.264 compression varies by camera and manufacturer. The bandwidth that a camera requires can shift between 10 and 50 percent, depending on the scene.

With traditional approaches to storage, the variation in the bandwidth of different video feeds unevenly distributes video files among volumes, as can be seen in Figure 3. In video systems with smaller volumes and fewer cameras, for example, the variation in throughput can quickly overwhelm the 20 percent overhead built in to the storage capacity unless the capacity planning assumed extremely high average bandwidths. When the bandwidth associated with a video feed exceeds expectations, the storage volume requires more capacity to meet the same retention times.

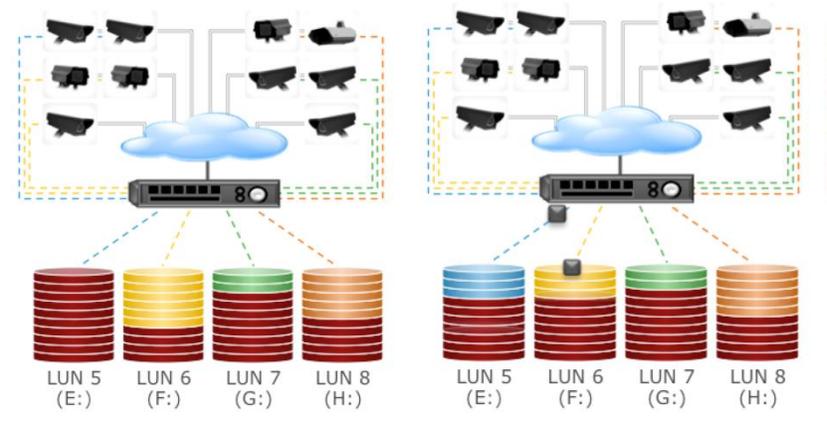


Figure 3: The distribution of video surveillance traffic with a traditional storage system

When a volume reaches its full capacity, the VMS deletes the oldest video until it frees up more space. With traditional storage systems, you end up having to migrate video cameras from one volume to another—which forces the VMS to move the video feed data among the volumes. The migration strains the I/O of the VMS servers, affecting video playback to such an extent that frames are lost. Such a data migration can take anywhere from a few hours to more than 12 hours, depending on the VMS vendor and the amount of stored data.

An Isilon cluster preempts this problem before it occurs. OneFS ingests traffic from video management systems through a single volume and then automatically balances the data among the nodes in the cluster, as Figure 4 illustrates. To address the variability of video streams, OneFS also continuously rebalances data to optimize capacity.

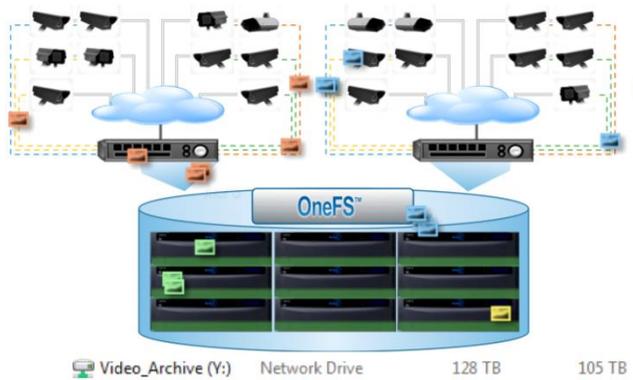


Figure 4: The distribution of video surveillance traffic with an Isilon cluster

The digital footage captured by the video surveillance cameras is encoded by the VMS server, including those hosted on VMware virtualized servers, and is then transmitted over 10 GbE network connections to the Isilon storage cluster. With an Isilon cluster, every VMS instance can point to the same top-level directory—for example, \\video_repository. Because the OneFS distributed operating system imposes no constraints on the size of a volume, an Isilon cluster presents all the files on the cluster within a global namespace when a server connects to any node (see Figure 5).

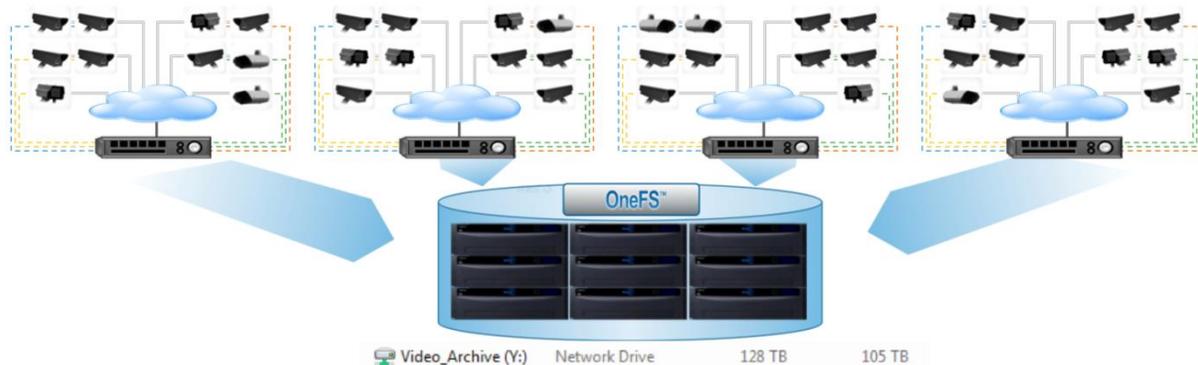


Figure 5: Surveillance system deployment with a scale-out NAS system

OneFS further simplifies the deployment of surveillance applications by automatically load balancing incoming SMB and NFS client connections across all the nodes. By default, the OneFS load-balancing module, called Dell EMC Isilon SmartConnect™, distributes client connections using a round-robin algorithm. In this way, an Isilon cluster is optimized, by default, for surveillance applications.

Because data writes with a large cluster can go to as many as 40 disks—which is many more disks than with a traditional file server—OneFS can write a video stream to disk with performance-enhancing parallel writes. The parallel writes that place a single, large video file on many disks have a multiplier effect on disk throughput that, when combined with the cluster’s aggregate RAM, optimizes the speed with which OneFS can write video files to disk.

Large video storage systems typically require defragmentation, which introduces overhead and reduces performance. OneFS, however, eliminates the need for defragmentation by distributing data among a cluster’s nodes and by restriping the data to optimize protection and data access. As a result, OneFS is commonly deployed in video surveillance systems that store as much as 20 PB of video and file data without the need for defragmentation.

Meanwhile, to retain data with longer retention times or to address changes in retention times, OneFS lets you set policies to automate tiering so that you can cost-effectively store data over time. The tiering policies automatically move files among tiers of storage according to the rules that you set.

A use case from a casino

A large casino in California recently deployed an Isilon cluster to solve a number of problems.⁵ The surveillance cameras at the casino record high-definition video at 30 frames per second. Before the Isilon cluster was installed, the surveillance operations were supported by about 400 digital video recording devices spread across 40 racks in the data center. The high power and cooling requirements needed for this deployment resulted in unacceptable operating costs. In addition, because each DVR could support only a small number of cameras, the system was complex to manage and difficult to scale.

The casino deployed an Isilon scale-out cluster with 50 Dell EMC Isilon X200 nodes to store the enormous amounts of footage captured by the hundreds of casino floor video surveillance cameras. The Isilon cluster consolidated the storage system from 400 machines on 40 racks to 4 racks of Isilon nodes that provide 1.5 PB of storage capacity, significantly reducing operational expenses associated with power, cooling, and floor space. The cluster's scalability enables the casino to increase retention periods and to easily and cost-effectively scale the cluster to meet ever-growing data storage requirements—requirements that are not hindered by bandwidth because each node has two 10 GbE network connections.

The casino's success with Isilon scale-out NAS is not an isolated case. According to an IDC survey titled "Quantifying the Business Benefits of Scale-Out Network Attached Storage Solutions," companies using Isilon "were able to drop their related facilities costs by 40%, saving \$42 per TB. Facilities savings were driven by the more efficient head architecture, which requires less rack space. Depending on the architecture being replaced, Isilon could use anywhere from 25% to 75% fewer racks."⁶

The IDC survey also found that companies reduced their power requirements by 13 percent. "The companies in our study also paid less per TB for their Isilon storage, which, when combined with the efficiency advantages, means that they lowered their storage costs by 37%, saving \$1,527 per TB. This advantage means that companies can budget more accurately for their storage requirements and more efficiently scale to meet new or fluctuating demand."

Storage efficiency

An Isilon cluster eliminates much of the overhead that traditional storage systems require. OneFS evenly distributes, or stripes, data among a cluster's nodes with layout algorithms that maximize storage efficiency and performance. The system continuously reallocates data to conserve space.

At the same time, OneFS protects data with forward error correction, or FEC—a highly efficient method of reliably protecting data. The capacity overhead for data protection is about 20 percent in aggregate, not per individual volume. The additional capacity overhead of 20 percent required by DAS and SAN is not needed with the Isilon cluster. In practice, an Isilon cluster runs at 90 percent of full capacity, saving as much as 35 percent of capacity over traditional storage systems.

Isilon's efficient use of storage space reduces CAPEX. "Capex reduction is driven by Isilon's much more efficient use of storage compared with traditional storage," an IDC survey of companies found. "For every \$1.00 spent on raw storage, Isilon delivers \$0.72 worth of storage in use compared with \$0.58 for other traditional storage. One key reason is that the utilization rate of Isilon storage is typically much higher than that of traditional storage. With traditional storage, IT administrators often have to set aside significant discretionary reserves for snapshots and performance headroom so that their storage can continue to operate."⁷

Elasticity and agility

Most video surveillance systems are dynamic. As the systems mature, adding or modifying the camera configurations is a common method used to cover other lines of sight. Resolutions are increased to enlarge the field of view or to capture more

⁵ "Pechanga Resort and Casino: EMC Isilon consolidates surveillance storage for high reliability and quick ROI." <http://www.emc.com/collateral/customer-profiles/h11272-pechanga-casino-cp.pdf>, last visited September 2013.

⁶ "Quantifying the Business Benefits of Scale-Out Network Attached Storage Solutions." IDC white paper. <http://www.emc.com/collateral/analyst-reports/idc-ar-business-benefits-of-scaleout.pdf>, last visited September 2013.

⁷ "Quantifying the Business Benefits of Scale-Out Network Attached Storage Solutions." IDC white paper. <http://www.emc.com/collateral/analyst-reports/idc-ar-business-benefits-of-scaleout.pdf>, last visited September 2013.

visual data for analytics. With a VMS server, it takes little more than a mouse click to change a network camera's resolution from 1080 pixels to 3 megapixels.

Traditional storage systems like the one shown in Figure 6 lack agility. When changes to the video stream affect the volume of a traditional storage system, the changes can consume overhead and cause the VMS to warn you that you need to migrate the cameras to another volume. You then need to spend time adding disks, configuring the array with an additional volume that has the right LUN, and formatting for RAID 6. Afterward, the VMS needs to accommodate associating the video feeds to the new LUN and volume, which can take up to 12 hours and affect playback.

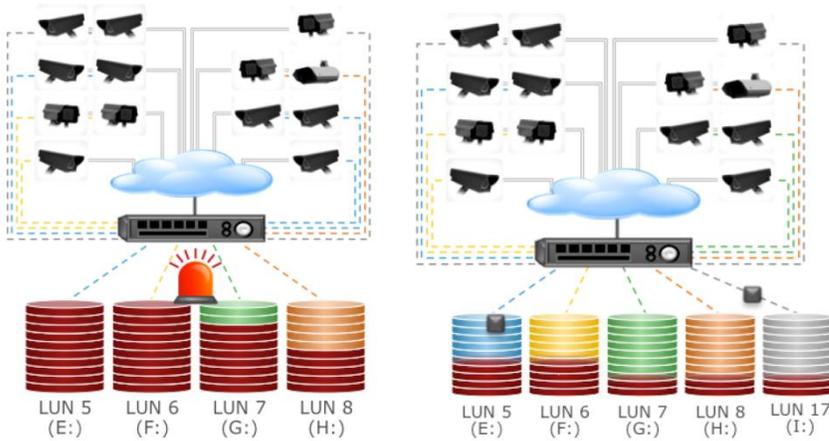


Figure 6: Traditional storage volume distribution and migration

An Isilon cluster handles change with agility. Because an Isilon cluster contains a single, global volume, the problem of one volume filling up faster than another does not occur. A single volume also reduces the likelihood that you will need more capacity because there is no static association between a small set of disks and a volume.

If the cluster does require additional capacity to handle more cameras, it takes less than two minutes to have a node added to a cluster and appear as part of the single volume in a usable manner. When you add a node, OneFS automatically rebalances the data across the new node, with no affect on playback and recording, as verified in tests at Dell EMC labs. More importantly, no changes are necessary on the VMS, where the target volume does not have to change and remains `\\video_repository`.

An Isilon cluster's ease of use and ability to rapidly scale out saves time and money. "IT productivity gains reduced the operating cost per TB by 3.5 FTE hours or 48%," IDC found in its survey for general IT system administration.⁸ For video surveillance systems, many of which undergo constant change, the operating costs of Isilon scale-out NAS may be even lower.

Failure handling

There are two primary points of infrastructure failure in a video surveillance system: VMS and storage. Most video management systems include a failover mechanism that transfers the video feeds from a VMS that fails to a secondary VMS, which receives the feeds. With a traditional storage system, when the VMS remaps the camera feeds to the new server, the recordings for the cameras associated with the failed VSM's volumes become unreachable. With an Isilon storage system, there is no loss of playback for the recorded video.

Disks can present another point of failure in the storage system. If two or more disks fail in the volume of a traditional system, the cameras associated with the volume lose the capability to make new recordings and play back existing content. Even

⁸ "Video Surveillance Trends for 2013." IMS Research. December 18, 2012. http://www.imsresearch.com/press-release/Video_Surveillance_Trends_for_2013.

when one disk fails in a volume, frame loss typically occurs because rebuilds affect the surveillance system's bandwidth and processing.

By comparison, an Isilon storage system can handle two simultaneous disk failures as well as the failure of a node with as many as 36 disks without affecting new video recordings and without affecting playback. New recordings are unaffected if a node fails because the SmartConnect module's automated load balancing of video streams minimizes frame loss in-flight.

New technologies for new developments

An Isilon cluster also protects your investment with emerging techniques that can transfer video data and move it into the cloud.

In one trend, camera endpoints are becoming more robust in CPUs, memory, and mass storage. Many endpoints can connect directly to storage systems with network protocols like NFS and HTTP. With such direct connections, VMS servers can bypass processing the video because standards such as those of the Open Network Video Interface Forum (ONVIF) define the data format for storing the video data. Although ONVIF and other standards reduce the amount of hardware that you need to normalize video data, the surveillance system must still coordinate the interface between the VMS and the storage subsystem through standards like ONVIF Profile G. Dell EMC is a member of [ONVIF](#).

Another trend is moving data management from VMS applications to open-standard application programming interfaces (APIs) that third-party systems handle. Such APIs let the VMS access data from any component, such as a video wall or a surveillance system sensor.

OneFS includes a RESTful API that lets an application migrate data to the cloud while removing the I/O burden of moving files from one storage tier to another. The result is a scale-out, standards-based approach to NAS that ensures a high return on investment through efficient data management while seamlessly scaling to match the ever-increasing amount of surveillance video.

Conclusion

A Dell EMC Isilon cluster solves many common storage problems for video surveillance systems by providing capabilities that traditional storage systems lack:

- **A single volume.** A Dell EMC Isilon cluster presents all the nodes in the cluster as a single, global namespace. VMS servers access the same data by connecting to any node in the cluster.
- **Seamless scalability.** The Isilon clustered file system dynamically expands up to 20 PB in a single volume and over 100 GB/s of aggregate throughput.
- **Efficient data protection.** OneFS distributes data across the cluster to guard the data with parity blocks instead of parity disks. You can lose up to four 36-drive nodes without a loss of video playback and without disrupting video recording.
- **High-velocity video ingestion.** For performance, OneFS load balances 1 GbE and 10 GbE network connections among all the nodes in a cluster to eliminate bottlenecks and to accept video streams using network protocols such as NFS, SMB, HTTP, and REST.

For large, centralized video surveillance systems, an Isilon cluster is a highly scalable storage solution that delivers quick ROI and reduces both CAPEX and OPEX.