Quantum

Video Surveillance Shared Storage Reference Architecture

March 2021 v1.0

Contents

Executive Summary

Data created by video surveillance (VS) is exploding. More data is created by video surveillance cameras than any other application or industry, worldwide. 1080P full-HD cameras have become the baseline, and 4K cameras are increasingly popular. Millions of additional cameras are deployed each year. VS operators increasingly need to retain more video online for quick access and real-time analytics, and more video nearline or offline for historical analysis and legal liability protection. Video storage can easily consume one-third of the budget for a surveillance project. For these reasons and more, video storage is a huge challenge for all but the smallest VS operations.

For over 20 years, Quantum has been building storage systems specifically tailored for video content. Quantum software and hardware are used to store and process video at the world's largest broadcasters, movie studios, government agencies, and Fortune 100 corporations. Now with the broadest security infrastructure portfolio from a global vendor, Quantum can provide VS solutions at any scale.

Based on recently published SPEC-SFS 2014 testing with the VDA workload (a video surveillance workload simulation), the Quantum StorNext File System (SNFS) is officially the fastest file system in the world for video. Not only was a new record of 7450 streams into a single file system achieved, but the previous record was shattered using significantly less – and less expensive – hardware, proving that StorNext solutions for video surveillance are also the best value by a wide margin.

This document describes a range of configurations for VS storage systems that include HA server pairs to host the VMS application and act as recording servers, a scalable shared primary storage infrastructure that eliminates video storage silos, and optional shared archive storage for cost-effective long-term retention. The configurations illustrated are modular and expandable using standard Quantum components. Every organization has a unique set of requirements, but these configurations can serve as important reference points for system design.

Objectives

The objectives of this paper are to:

1. Make the case for using shared storage in VS applications
2. Present a general, scalable architecture for shared VS primary storage
3. Outline a range of sample configurations, sized for specific camera counts and days of retention
4. Describe the components in the solution stack
5. Describe the design considerations dictated by VS workloads in general, and those related to the specific technologies and components used in the reference configurations
6. List the assumptions used to create the reference configurations, and call out factors that may significantly impact system sizing
7. Discuss options for shared archive storage that integrate with the shared primary storage architecture
8. Document additional guidelines, tips, and best practices to maximize success

The Case for VS Shared Storage

Experience tells us that video surveillance installations almost never shrink, and they almost always grow. The growth happens for many reasons, including the plummeting cost and increasing image quality of cameras, and the realization that VS isn't just for security any longer. The availability of sophisticated analytics applications enables more value to be extracted from captured video, and the more video you analyze, the more accurate the insights derived from it.

Growth creates an obvious problem. Where to put all this stuff? If you aren't careful, incremental growth in camera count leads to a profusion of storage systems scattered around the facility, the campus, or the enterprise. Multiple storage repositories take more time to manage, more effort to secure, and cost more to acquire and maintain. Scattered storage repositories also introduce extra hurdles for analytics, requiring data to be accessed across the network, leading to congestion, slowdowns, or expensive network upgrades.

For mid-sized, and especially for large installations (in terms of data retained), centralized shared VS storage is a better solution. A single, robust storage system costs less per TB to buy and expand and is easy to manage and secure. Analytics applications may run in close proximity to the storage, eliminating the need to drag video content hither and yon over general-purpose networks. In addition, the right shared primary storage system can act as a centralized gateway to archival storage targets, such as object storage, public cloud, or data tape. Smart integration of primary and archive storage lowers long-term storage costs, or taken from another view, enables storing more footage for longer with the same budget.

Solution Overview

The purpose of the solution detailed below is to provide a scalable shared primary storage solution for video surveillance camera ingest. An escalating series of configurations are charted, with varying camera counts (500 – 2000) and days of retention (30 – 365). Relevant portions of the configuration were lab tested to validate the sizing. Although Milestone XProtect® VMS software was used to validate these configurations, the results are expected to apply to other VMS applications, though implementation details may differ.

Additional information is provided on optional shared archive storage that may be deployed in conjunction with a shared primary storage system. Shared archive storage involves migrating older footage to lower-cost storage media. This enables a reduction in VS storage costs, or the storage of more days of camera output for a given budget.

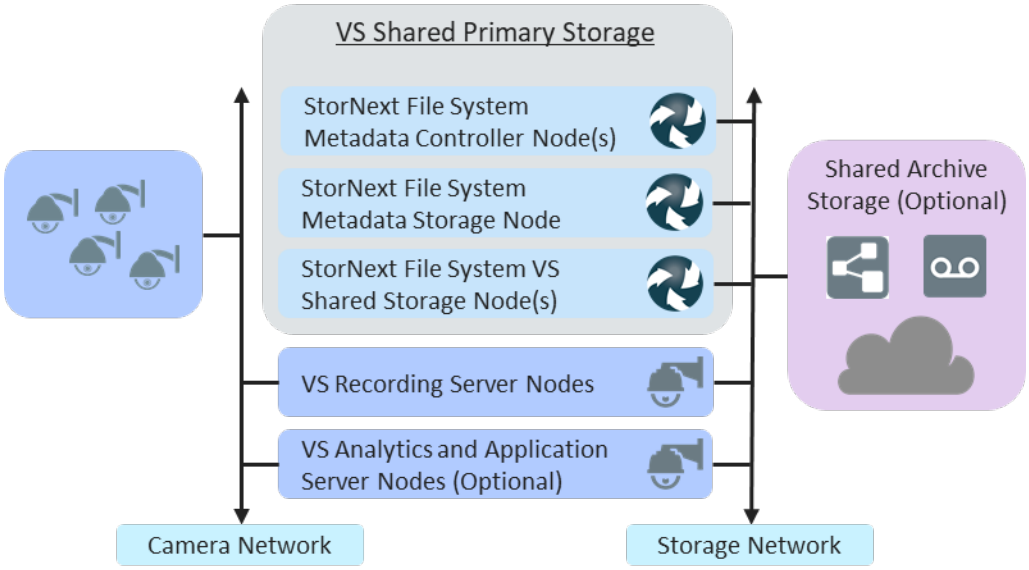

FIGURE 1 - VS SHARED STORAGE ARCHITECTURE - HIGH LEVEL

The shared primary storage solution centers around a StorNext File System implementation using Quantum compute and storage appliances. The StorNext system consists of several components:

- Xcellis® Workflow Directors (a pair for HA configurations) act as metadata controllers for SNFS
- A QXS™-424 SSD storage node, used to store SNFS metadata and journals
- Several QXS-484 HDD storage nodes, used to store the footage received from the cameras
- SNFS client software, used by the video recording servers to mount the SNFS file system
- Several Ethernet networks

Data flows from the cameras, through the video recording servers, onto the shared storage. Each recording server writes into its own directory structure, but the storage pool is shared. Every VMS server has access to every TB of the storage. There is no need to partition the storage or attempt to balance it.

Technology Summary – Shared Primary Storage

The table below lists the technology components that make up the shared primary storage reference architectures outlined in this document. The paragraphs that follow the table provide more detail on the function of these components in the solution.

Item	Notes
Quantum StorNext File System Software	V6.4.0 or later High-performance shared file system software Integrated data management
Quantum Xcellis Workflow Director Node	SNFS Metadata Controllers - Two recommended for HA
Quantum QXS-424 12G RAID Node	SNFS Metadata Storage SSD equipped, iSCSI or Fibre Channel

Quantum QXS-484 12G RAID Node Quantum QXS-x84 12G Expansion Node	SNFS VS Data Storage HDD equipped, iSCSI or Fibre Channel
Quantum VS1110-A Application Server Node	Acts as a video recording server with installation of appropriate Video Management Server (VMS) software [Optional] Additional application server
VMS software	Customer or integrator supplied, not sold by Quantum. Validation testing was performed with Milestone XProtect Corporate 2020 R3 and the StableFPS tool
[Optional] Quantum VS2108-A Video Surveillance Analytics Node	GPU-enhanced analytics server node
[Optional] VS Analytics Software	Customer or integrator supplied, not sold by Quantum
Camera Network	Front-end Ethernet network connecting cameras to video recorder and optional analytics nodes
SNFS Client Network	Ethernet network connecting video recording nodes to the SNFS
Storage Network	Back-end Ethernet / iSCSI or FC network connecting all server nodes with shared primary storage
Management Network	Back-end Ethernet for component management only

TABLE 1 - TECHNOLOGY COMPONENTS: PRIMARY STORAGE

StorNext Software

StorNext® is Quantum's high-performance file system. It provides high-speed block storage services, able to support thousands of cameras without dropping frames. This makes it perfect for VS shared storage. In addition, StorNext can manage multiple types of storage – both primary and secondary – within a single file system namespace. Optional policies move and copy data between storage types as required, optimizing cost and making StorNext file systems self-protecting without additional backup hardware or software.

This policy-based data movement enables a StorNext system to be used as an archive gateway if desired. Data initially lands on primary storage (HDD or SSD) and from there is copied to one or more secondary storage targets. As additional data lands on the primary storage, the file system self manages its fill level, 'truncating' older files. Truncation removes a file's data blocks from primary storage after they have been protected with secondary copies but retains the corresponding inode. This enables a user viewing the file system to see all files as if they still reside on primary storage, even those whose data blocks only exist on secondary storage.

StorNext file systems are scalable and flexible, able to grow in both capacity and performance as needed. For a technical overview of the StorNext file system's architecture and data services, refer to the whitepaper linked in the References table later in this document.

StorNext Storage Appliance

Though StorNext software may run on any servers and storage, using StorNext appliances is a convenient and cost-effective way to deploy the system. The need for extensive tuning and testing is eliminated, performance characteristics are well-defined, and centralized support and design assistance is provided by Quantum both before and after a purchase. StorNext storage appliances used in the VS shared storage reference architecture include the following node types:

Xcellis Workflow Director (XWD)

XWD nodes are the traffic cops for StorNext file systems, ensuring coherency and security of all data. Small systems may run a single XWD, but dual XWD nodes are typical, to provide redundancy. When required, additional StorNext data services may run on XWD nodes.

Block Storage

Block storage nodes house user data (in this case VS camera output files and any associated analytics outputs) and StorNext file system metadata. A wide variety of storage node types are available, from SATA HDD to SSD to NVMe flash. Connectivity options include iSCSI, Fibre Channel, and RDMA Ethernet.

Because file system metadata performance is critical to overall file system performance, the reference architecture uses an array of SSDs solely for file system metadata and journals, and separate arrays of HDDs for video.

VS1110-A Application Server

VS1110-A nodes are 1U servers designed as application servers. In the reference architecture these systems host the VMS software. Using an embedded hypervisor, each VS1110-A can run multiple application VMs simultaneously. Quantum's recommendation is to deploy application servers in active/passive HA pairs, using the hypervisor to manage VM failover.

Optionally, VS1110-A nodes may be deployed to run additional, related applications, such as access control, environmental, or other building management systems.

VMS Software

Video Management System software is the heart of a video surveillance system and is typically supplied by the system integrator. Quantum does not sell VMS software. All testing referenced in this document was performed with Milestone XProtect and Milestone's StableFPS stream simulation tool.

[Optional] VS2108-A Analytics Node

The VS2108-A is a GPU-enhanced server node designed to run VS analytics applications. Configurations are available with one, three, or five GPUs. All configurations ship with mirrored SSD boot/OS drives, and a range of SSD capacities are offered for auxiliary onboard storage.

[Optional] VS Analytics Software

VS analytics applications that may be run on VS2108-A analytics nodes are customer or integrator-supplied. Quantum does not sell VS analytics software.

Networks

Ethernet networks are used in the shared primary storage reference architecture for all network functions. Several distinct networks are used; a front-end camera network, one or two StorNext client networks, a back-end storage network (which could be FC if desired), and a management network. The Ethernet networks may be separate physical networks or simply vLANs. The camera, StorNext client, and storage networks are typically 10GbE or faster and use jumbo frames, while the management network can be 1GbE.

Technology Summary – Shared Archive Storage

The table below lists the *additional* technology components that may be used for the shared archive storage option. The paragraphs that follow the table provide more detail on the function of these components in the solution.

Item	Notes
[Optional] Quantum Xcellis Workflow Extender Node	For StorNext 7 Systems
[Optional] Quantum Scalar LTO Tape Library	Shared archive cold storage option
[Optional] Quantum ActiveScale Object Storage	S3-based shared archive storage option
[Optional] Public Cloud Storage Subscription	S3-based shared archive storage option Customer-supplied, not sold by Quantum
Storage Network	Back-end Ethernet or Fibre-Channel network connecting StorNext XWD and XWE nodes with shared archive storage targets

TABLE 2 - TECHNOLOGY COMPONENTS: ARCHIVE STORAGE

StorNext Storage Appliance

Depending on the details of the specific customer solution design, one additional StorNext appliance node type may be required.

[Optional] Xcellis Workflow Extender (XWE)

XWE nodes are used to scale out the capabilities of a StorNext system. XWEs may run any StorNext data service. In the shared archive portion of the reference architecture XWEs are envisioned as data movers for secondary storage tiers, known as Distributed Data Movers, or DDMs. Whether XWEs would be required, and the number of XWEs needed, will depend on the characteristics of the specific installation.

[Optional] Scalar LTO Tape Library

LTO Tape remains, by far, the lowest-cost option for cold storage of video. Quantum Scalar i3 and Scalar i6 automated tape libraries use LTO digital tape technology to store data safely and cost-effectively.

Quantum-proprietary data verification functionality known as Enterprise Data Lifecycle Management (EDLM) is included in all Scalar tape libraries. EDLM periodically tests the quality of stored footage, and alerts the operator if media begins to degrade. If a vaulted copy is required, StorNext policies may be configured to automatically create a second copy of content on a separate set of media. These tapes may be exported from the library for secure storage elsewhere.

LTO-WORM media is supported for compliance applications where footage must be preserved with absolute assurance that it cannot be modified.

[Optional] ActiveScale Object Storage

ActiveScale is scalable, cost-effective “private cloud” storage. Content is stored by StorNext software into ActiveScale via an S3 interface to erasure-coded hard disk storage. To provide site disaster protection, a single ActiveScale system may be deployed across multiple sites, a practice known as geo-spreading.

Though on-prem object storage such as ActiveScale has a higher cost per TB than tape, it has the advantage that all archived content is online and instantly available. Content stored on tape can take up to a few minutes to retrieve. Disaster recovery protection via geo-spreading is another advantage vs. tape. Geo-spreading is automatic and completely "hands-off", while sending tapes off-site requires operator action.

ActiveScale is also notably less expensive than public cloud storage for large and/or active archives, as there are no monthly storage fees or penalties for retrieval.

[Optional] Public Cloud Storage Subscription

StorNext software also supports all popular public cloud providers as archive targets. Public cloud can be an ideal target for smaller archives, or where archived data is very seldomly accessed. It is also useful when an off-site copy is desired, but the organization resides at a single site.

Because StorNext policies can be configured to make up to four copies of files, all of which may be on different media types, mixed configurations are possible. For example, one copy of the archive may reside on-site in a Scalar tape library, with a DR copy in public cloud storage. Or two copies may exist at two separate cloud providers. Whether a particular archive requires multiple copies, and where those copies should be located, depends on the value of the content, individual requirements for disaster protection, compliance retention, and other factors.

Networks

When a shared VS archive is added to a shared VS primary storage system, there are additional network factors to consider. Additional Ethernet ports on the management network will be required for archive targets such as ActiveScale and Scalar tape libraries. A connection to the Internet, sized appropriately, will be required for storage of content to public cloud destinations. ActiveScale geo-spreading requires additional Ethernet bandwidth between sites. These may be LAN links across a campus, or private WAN or VPN connections between sites that are separated by greater distances. For long distances, site to site network latency and expected levels of inter-site traffic must be understood to ensure that a solution design will be successful.

Scalar tape libraries require either additional ports on the Ethernet / iSCSI network, or a Fibre Channel network to connect the tape drives and robot to the StorNext XWD and (if used) XWE nodes. This usually requires a switched network, so that all XWD and XWE nodes can 'see' all tape resources. The exception is for small installations using only a single (non-redundant) XWD node, in which case point-to-point connections from the XWD node to the tape library will be sufficient.

Reference Architecture

This solution was designed to provide shared storage for video surveillance camera ingest, for environments with 500 – 2,000 cameras, and retention periods of 30 – 365 days

Solution Sizing

End-to-end Performance	3.55 to 14.29 Gb/s, 38.4 to 153.6 TB per day
Usable Storage Capacity	1,152 TB to 56,064 TB

TABLE 3 - SOLUTION SIZING

Solution Component Configuration Summary

StorNext Metadata Controller Nodes	Quantum Xcellis Workflow Director (XWD) single or dual-node HA 128GB RAM Quad 10GbE Network Adapter, iSCSI, OR Dual 25GbE Network Adapter, iSCSI
StorNext Metadata Storage Node	1x Quantum QXS-424 RAID node, iSCSI, 8 TB usable (12 x 0.8 TB), SSD Single 10+2 RAID6 LUN
StorNext Disk Storage Nodes	Quantum QXS-484 12G iSCSI RAID and expansion nodes Up to 84 x 16 TB HDD per node 10+2 RAID6 LUNs, up to 7 per node
StorNext File System	Single file system Single SNFS metadata stripe group One user data stripe group per QXS-484 array controller
Video Recording Server Nodes	Quantum VS1110-A Application Servers 2x, 4x, 6x, or 8x, deployed in active/passive HA pairs VMs configured with synchronous replication & hypervisor-level failover
Cameras	500 – 2000 1080p, H.264, 30 frames per second All cameras observing constant motion, normal level of scene activity All cameras recording 24hr per day
Networks	Camera Network: 1x 10GbE or 25GbE, Jumbo Frames (MTU 9000) StorNext Client Networks: 2x 10GbE or 25GbE, Jumbo Frames (MTU 9000) Storage Network: 1x 10GbE, Jumbo Frames (MTU 9000) or 16G FC Management: 1x 1GbE
[Optional] Application Server Nodes	As needed to run additional applications
[Optional] VS Analytics Server Nodes	As needed to run VS analytics applications

TABLE 4 - COMPONENT CONFIGURATION SUMMARY

Solution Diagram – Shared Primary Storage

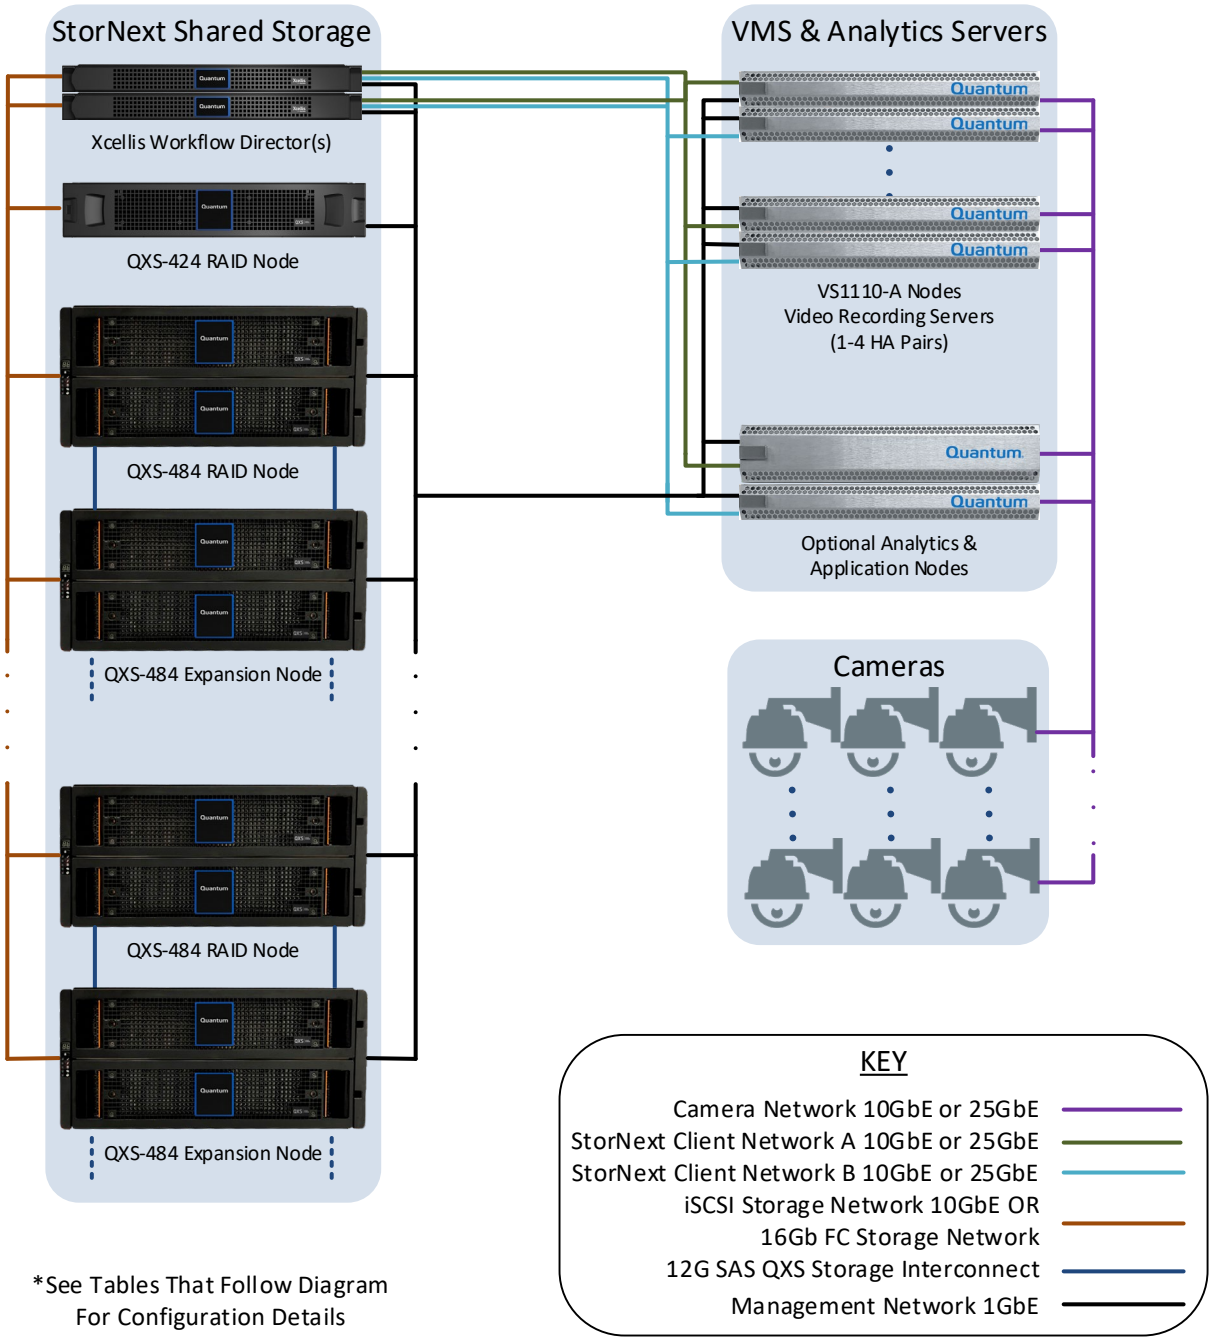

FIGURE 2 - REFERENCE ARCHITECTURE SOLUTION DIAGRAM

Solution Component Configuration Detail

Camera Count	VS1110-A Nodes	Xcellis and QXS Configuration	
		30 days	90 days
500	2	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE Adapter * 2x 10GbE ports: 1 cam net, 1 SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 1x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 1152 TB usable	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE Adapter * 2x 10GbE ports: 1 cam net, 1 SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 1x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 2x QXS-484 Expansion Node, 1152 TB (84x16TB), HDD * 3456 TB usable
1000	4	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE Adapter * 2x 10GbE ports: 1 cam net, 1 SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 1x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 1x QXS-484 Expansion Node, 1152 TB (84x16TB), HDD * 2304 TB usable	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE Adapter * 2x 10GbE ports: 1 cam net, 1 SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 2x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 4x QXS-484 Expansion Node, 1152 TB (84x16TB), HDD * 6912 TB usable
1500	6	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE or Dual 25GbE Adapter * 4x 10GbE ports: 2x cam net, 2x SNFS client net, OR * 2x 25GbE ports: 1x cam net, 1x SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 1x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 2x QXS-484 Expansion Node, 1152 TB (84x16TB), HDD * 3456 TB usable	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE or Dual 25GbE Adapter * 4x 10GbE ports: 2x cam net, 2x SNFS client net, OR * 2x 25GbE ports: 1x cam net, 1x SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 3x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 6x QXS-484 Expansion Node, 1152 TB (84x16TB), HDD * 10368 TB usable
2000	8	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE or Dual 25GbE Adapter * 4x 10GbE ports: 2x cam net, 2x SNFS client net, OR * 2x 25GbE ports: 1x cam net, 1x SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 1x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 3x QXS-484 Expansion Node, 1152 TB (84x16TB), HDD * 4608 TB usable	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE or Dual 25GbE Adapter * 4x 10GbE ports: 2x cam net, 2x SNFS client net, OR * 2x 25GbE ports: 1x cam net, 1x SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 3x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 9x QXS Expansion Node, 1152 TB (84x16TB), HDD * 13824 TB usable
NOTES:			
* All capacity figures are usable TB			
* All capacity calculations per Quantum VS sizing tool v1.04			
* Metadata storage requirements increase if optional shared archive storage is used			
* Additional important notes, caveats, and assumptions are contained elsewhere in this document			

TABLE 5 - CONFIGURATION DETAIL: 30 & 90 DAY RETENTION

Camera Count	VS1110-A Nodes	Xcellis and QXS Configuration	
		180 days	365 days
500	2	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE Adapter * 2x 10GbE ports: 1 cam net, 1 SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 2x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 4x QXS-484 Expansion Node, 1152 TB (84x16TB), HDD * 6912 TB usable	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE Adapter * 2x 10GbE ports: 1 cam net, 1 SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 4x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 8x QXS-484 Expansion Node, 1152 TB (84x16TB), HDD * 1x QXS-x84 12G Expansion Node, 384 TB (28x16TB), HDD * 14208 TB usable
1000	4	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE Adapter * 2x 10GbE ports: 1 cam net, 1 SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 3x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 9x QXS-484 Expansion Node, 1152 TB (84x16TB), HDD * 13824 TB usable	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE Adapter * 2x 10GbE ports: 1 cam net, 1 SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 7x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 17x QXS-484 Expansion Node, 1152 TB (84x16TB), HDD * 1x QXS-x84 12G Expansion Node, 384 TB (28x16TB), HDD * 28224 TB usable
1500	6	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE or Dual 25GbE Adapter * 4x 10GbE ports: 2x cam net, 2x SNFS client net, OR * 2x 25GbE ports: 1x cam net, 1x SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 5x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 13x QXS-484 Expansion Node, 1152 TB (84x16TB), HDD * 20736 TB usable	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE or Dual 25GbE Adapter * 4x 10GbE ports: 2x cam net, 2x SNFS client net, OR * 2x 25GbE ports: 1x cam net, 1x SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 9x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 27x QXS-484 Expansion Node, 1152 TB (84x16TB), HDD * 1x QXS-x84 12G Expansion Node, 768 TB (56x16TB), HDD * 42240 TB usable
2000	8	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE or Dual 25GbE Adapter * 4x 10GbE ports: 2x cam net, 2x SNFS client net, OR * 2x 25GbE ports: 1x cam net, 1x SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 6x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 18x QXS Expansion Node, 1152 TB (84x16TB), HDD * 27648 TB usable	* 1x Xcellis Workflow Director for SN7, one or two node * iSCSI, 128GB RAM, Quad 10GbE or Dual 25GbE Adapter * 4x 10GbE ports: 2x cam net, 2x SNFS client net, OR * 2x 25GbE ports: 1x cam net, 1x SNFS client net * 1x QXS-424 12G RAID Node, iSCSI, 9.6TB (12x0.8TB), SSD * 13x QXS-484 RAID Node, 1152 TB (84x16TB), HDD * 35x QXS Expansion Node, 1152 TB (84x16TB), HDD * 1x QXS-x84 12G Expansion Node, 768 TB (56x16TB), HDD * 56064 TB usable
NOTES:			
* All capacity figures are usable TB			
* All capacity calculations per Quantum VS sizing tool v1.04			
* Metadata storage requirements increase if optional shared archive storage is used			
* Files 16MB or smaller will require more metadata disk for the 2000x365 config			
* Additional important notes, caveats, and assumptions are contained elsewhere in this document			

TABLE 6 – CONFIGURATION DETAIL: 180 AND 365 DAY RETENTION

StorNext Client Tuning

StorNext software has many tuning parameters at both the server and the client level. The following client tuning parameters were determined to be helpful in obtaining the best performance in a Milestone XProtect environment. While tuning may vary for other VMS applications, the values listed are a good starting point.

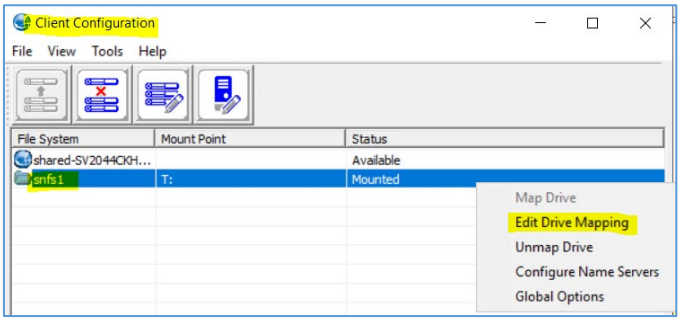

FIGURE 3 - STORNEXT CLIENT CONFIGURATION

On each VMS host, open the StorNext Client Configuration tool. Right-click on the StorNext file system entry in the table and select *Edit Drive Mapping* from the context menu.

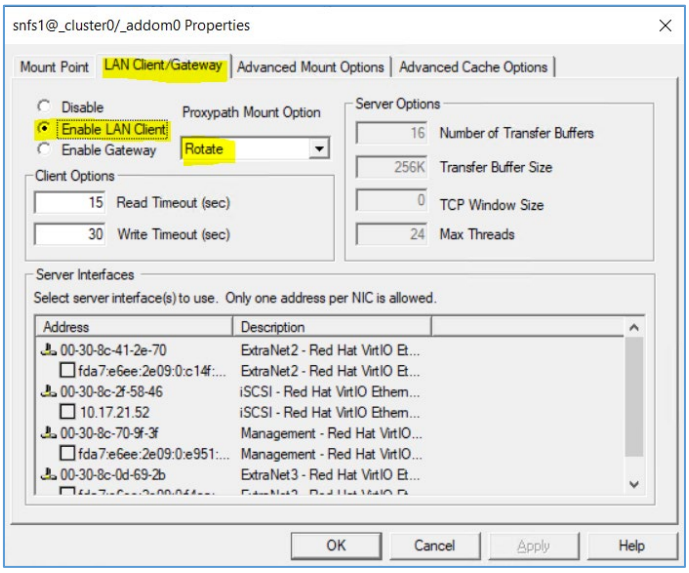

FIGURE 4 - LAN CLIENT / GATEWAY OPTIONS

On the *LAN Client/Gateway* tab, ensure that *Enable LAN Client* is selected, and in the *Proxypath Mount Option* dropdown select *Rotate*.

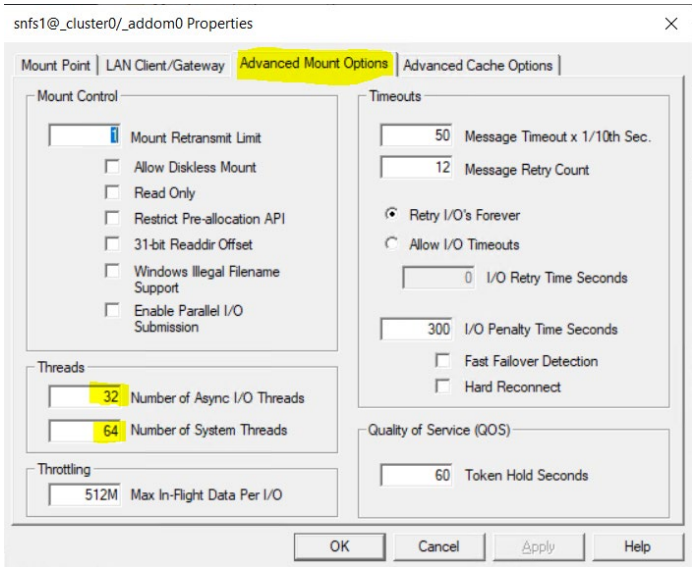

FIGURE 5 - ADVANCED MOUNT OPTIONS

On the *Advanced Mount Options* tab, in the *Threads* section, set the following values:

- Number of Async I/O Threads: 32
- Number of System Threads: 64

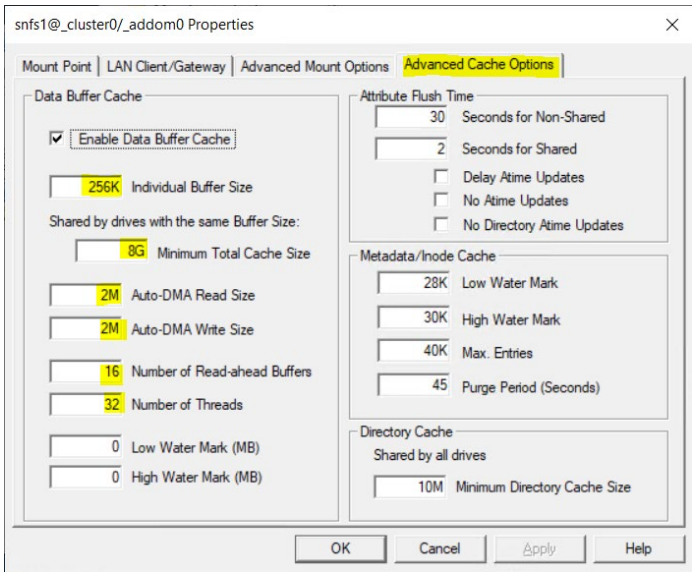

FIGURE 6 - ADVANCED CACHE OPTIONS

Finally, on the *Advanced Cache Options* Tab set the following Data Buffer Cache values:

- Individual Buffer Size: 256K
- Minimum Total Cache Size: 8G
- Auto DMA Read Size: 2M
- Auto DMA Write Size: 2M
- Number of Threads: 32

Considerations & Caveats

As with any complex technical solution, there is no 'one size fits all.' Every deployment will be different, not only in terms of hardware, but also the nature of the data. Many factors are involved in determining the best configuration for a specific solution, and many things determine the ultimate performance that is achievable. For large deployments a proof-of-concept or limited deployment is recommended first to validate assumptions and hone the final configuration. The Quantum Pre-Sales Engineering team is uniquely qualified to assist with solution design and sizing and is always ready to help.

This section lists a few of the things to consider when designing a video surveillance shared storage solution. Refer to the documentation linked in the References section for more.

Topic	Notes
Camera Type	The reference architecture sizing assumes a single sensor per camera, capturing video. Cameras with multiple sensors (e.g. video & audio, multiple image sensors), will generate additional data, changing throughput and storage requirements.
Video Stream	The reference architecture sizing assumes 1080p video at 30 frames per second (fps) using an H.264 codec. Using a different resolution, fps, or codec can change the throughput and storage requirements significantly.
Motion Detection	The reference architecture sizing assumes a moderate scenario, where all cameras are observing motion (and thus recording at 30fps) 24x7x365. If all cameras were not observing constant motion, throughput and storage requirements would decrease. If using motion detection, however, CPU requirements on the recording servers would increase. Validation testing indicates that video with 20% motion results in CPU utilization of approximately 75% on the recording servers, vs. 50% utilization with constant motion / no motion detection.
Scene Motion	The reference architecture sizing assumes a moderate level of motion in view of each camera at all times. This is not necessarily typical. In a live deployment, scene motion varies based on camera location, time of day / day of week, and other factors that are unique to the specific installation. These factors may not be fully known or knowable in advance, and they will have a significant impact on throughput and storage requirements.
Playback	The reference architecture sizing does not consider the playback of video occurring concurrently with recording. If significant amounts of playback or other read activity (such as for analytics) are expected, these factors must be accounted for when sizing a solution.
File Size	The load on the shared file system, particularly the load on the metadata subsystem, varies based on the number of file operations per unit time. For a given amount of data, smaller files create a larger load. During the validation testing, the Milestone recording servers were configured to output 1GB files, significantly larger than the default 16MB. This reduced the load on the SNFS metadata subsystem to a minimum. If it is not possible or desirable to record using large files, higher-performance metadata storage may be needed, and in very large deployments, additional metadata capacity may be required.

Topic	Notes
VMS Software	Solution validation testing was performed using Milestone XProtect Corporate 2020 R3. Results may vary, and sizing requirements may differ, with other VMS applications or versions. VMS housekeeping functions (such as Milestone's "table split" activity) can impact system performance, and cause dropped frames if not accounted for.
Antivirus / Antimalware	Security applications such as those that monitor for computer viruses and malware can have a significant, deleterious impact on VMS server performance, causing dropped frames during scans. Windows Defender was disabled during validation testing for this reason. Whether or not it is acceptable to disable security applications depends on specific system and network configuration (e.g. are the VMS systems completely isolated from other users and business applications?) and organizational security policy. If disabling these applications is not possible, system sizing may be impacted. In Quantum's experience, running Antivirus / Antimalware applications on the VMS servers can reduce the number of cameras supported by as much as 30%.
VS1110-A Failover	The VS1110-A Application Server Nodes used in this reference architecture to run the VMS software support VM failover. The VMS application VMs are synchronously replicated to a standby VS1110-A. In the event a VM fails, it will be restarted on the same node. If the running VS1110-A node fails, all VMs will be failed over to the standby VS1110-A node. During a VM restart or node failover, several (single digit) minutes of video is lost from the cameras on the affected recording server VMs. Because these events should be extremely rare, this loss was deemed acceptable.
Shared Archive	Adding optional shared archive components to the shared VS primary storage reference architecture will change the sizing and configuration requirements, potentially in a substantial fashion. Quantum should be consulted for assistance with the design of any system that includes a shared archive.
Jumbo Frames	Solution validation testing found the best performance was obtained when using jumbo frames (MTU 9000) on all interfaces in the camera network, StorNext client networks, and iSCSI storage network.
StorNext Client Networks	Larger configurations utilizing 10GbE networking should designate two separate networks for StorNext client communication with the Xcellis Workflow Directors. Splitting the traffic across two networks will reduce congestion and improve ability to handle peak loads.
VMS Service Startup Delay	On StorNext client machines, such as the VMS recording servers, VMS service startup must be delayed to allow the StorNext client services to start first. If the VMS services start before the StorNext client, the shared storage will not be visible. A delay of 30 seconds is sufficient. Consult the Microsoft Windows Server OS documentation for instructions on the relevant registry keys to change.

TABLE 7 - CONSIDERATIONS & CAVEATS

References

The documents below contain more information on the Quantum components referenced in this document. Familiarizing yourself with these resources will ensure you obtain maximum value from the information contained in this reference architecture.

Document Title	URL
Quantum Product Documentation Center	https://www.quantum.com/documentation
StorNext Technical Whitepaper	Link
SPEC-SFS 2014 VDA Test – Quantum Configuration	Link
SPEC-SFS 2014 VDA Test – All Results	Link
Quantum Surveillance & Physical Security Page	Link
VS1110-A Application Server Page	Link
StorNext File System Page	Link
StorNext Xcellis Appliances Page	Link
QXS Block Storage Page	Link

TABLE 8 - REFERENCES

Document Feedback

For feedback, questions, and suggestions for improvements to this document contact the authors.

Revision History

Version	Authors	Notes
1.0	Dan Duperron (dan.duperron@quantum.com) Matt Stoller (matt.stoller@quantum.com) Steve Haberman (steve.haberman@quantum.com)	Initial Release 03/02/2021

TABLE 9 - REVISION HISTORY