

Xcellis Technical Overview: A deep dive into the latest hardware designed for StorNext 5

ABSTRACT

Xcellis™ represents the culmination of over 15 years of file system and data management development combined with over one hundred thousand clients to arrive at a hardware platform perfectly matched to the software environment. StorNext® customers range from small video post-production facilities to DNA sequencing research institutes to 30,000 camera surveillance deployments. They choose StorNext because of its ability to provide any number of collaborators with extremely high-performance access to the largest files as if they were stored directly on their workstation.

The hardware platform on which StorNext is deployed is crucial to derive the maximum performance, user access, flexibility, scalability, and reliability. Xcellis was designed specifically for the job. This paper covers the technical details behind the features that allow Xcellis customers to get more out of StorNext than ever before.

TABLE OF CONTENTS

Introduction	3
Background	3
Xcellis Workflow Director Base Configuration and Options	4
Xcellis Storage Array Configuration and Options	4
XLS-Config and Connectivity	5
XLS-LUNs and Disk Configuration	7
NAS Connectivity	8
Conclusion	10

INTRODUCTION

Xcellis™ workflow storage systems are designed specifically for StorNext 5 data management software. The combination of the two creates a storage environment capable of supporting many users with access to the fastest streaming file delivery in computing today. As files have become larger and workflows more complicated, the demands on storage to be more flexible and accessible have driven the development of both hardware and software systems. The engineers behind Xcellis set out to push the hardware platform beyond current requirements to create a system that will provide many years of superior service as technology and workflows continue to advance.

BACKGROUND

The history of StorNext and its related hardware options are well documented. See these articles on Wikipedia for information on the origins of StorNext.

- StorNext File System: https://en.wikipedia.org/wiki/StorNext_File_System
- Xsan: <https://en.wikipedia.org/wiki/Xsan>
- Metadata controller (stub): https://en.wikipedia.org/wiki/Metadata_controller

More recently, and just prior to the introduction of Xcellis, Quantum has been providing metadata controller (MDC) appliances and storage arrays on which to deploy StorNext. These systems provided all the functionality required by StorNext to operate as a complete system. Yet there was still room for improvement. Customers of the Quantum M-series metadata controllers could choose from a range of memory and performance options that predetermined the degree to which the system could scale. Occasionally, customers would contend with the need to redeploy a higher-grade MDC in order to scale the system beyond the constraints of their existing implementation.

Because of the expense and direct connectivity associated with Fibre Channel connections between first-order clients and the storage system, the number of collaborators in a given StorNext environment was limited. Even those connecting over Ethernet needed to install a Distributed LAN Controller (DLC) client to access the storage. Demanding workflows have spawned an increased need for collaboration between expanded team members. Xcellis adds standard file sharing protocol support, including SMB 1, 2, and 3, as well as NFS 3 to allow collaborators even with low bandwidth to directly participate in the workflow and keep all the data in one system.

The technical capabilities integrated into Xcellis generally fall into three categories:

1. XLS-Config and connectivity
2. XLS-LUNs and storage
3. NAS access

XCELLIS WORKFLOW DIRECTOR BASE CONFIGURATION AND OPTIONS

The Xcellis Workflow Director is a pair of devices known as Xcellis Workflow Director Nodes. They are deployed as a high-availability pair to guarantee uninterrupted operation. The Xcellis Workflow Director controls the file system, responds to data access requests, and provides the connectivity for clients of all types. The table below lists the basic configuration of the Xcellis Workflow Director:

Component	Configuration
Physical	Dual 1 RU server-class computers
Operating System	CentOS 7 Linux
Central Processing Unit	Dual six-core processors per server
Memory	64 GB RAM for up to 32 virtual file systems
Internal Storage	Pair of 500 GB disks (mirrored) for operating system and configuration information
LAN Networking	Four 1 GbE connections for management, metadata, and service
Storage Networking	Dual 16 Gbps optical Fibre Channel
Management	Lights Out Management (LOM)

In addition to the configuration listed above, a number of upgrade options customers may add to support larger configurations:


Component	Configuration
Tape Archive Connectivity	Quad 8 Gbps optical fiber
IP Archive Connectivity	Quad 1 GbE or Dual 10 GbE Ethernet
NAS and DLC Connectivity	Quad 1 GbE or Dual 10 GbE Ethernet
Memory	Additional 64 GB RAM for up to 64 virtual file systems

XCELLIS STORAGE ARRAY CONFIGURATION AND OPTIONS


One of the significant differences between Xcellis and earlier StorNext appliances is the ability to configure storage arrays to include both metadata and user data without sacrificing scalability or performance. Because metadata scales as new storage arrays are added, the system can be scaled continuously, and essentially endlessly, without the addition ancillary acceleration hardware.

The standard configuration of converged user and metadata storage also has the advantage of being fully managed by the system. This makes Xcellis customer installable through an automated storage configuration and deployment module built into StorNext Connect.

Xcellis systems are available in four initial storage configurations:



Quantum QXS-412
48TB or 72TB raw



Quantum QXS-456
224TB or 336TB raw

All of these configurations include the ability to scale in any increment and to any degree. As new disk capacities or performance levels become available, they can be added to the system and used to create dedicated volumes or integrated into the larger pool of storage.

In the event that the customer would like to utilize 3rd party storage, such as NetApp E-Series for example, Xcellis can support it. This situation can also arise from customers with existing StorNext deployments based on 3rd party storage they would like to upgrade to Xcellis.


In this case, an Xcellis Metadata Array is required to enable the storage. These arrays are specifically designed to provide reliable, high-performance system operation in a mixed vendor environment and include enough performance headroom to support capacity and performance scaling. Customers in this situation should work with their Quantum representative to determine the most appropriate option for their requirements. Xcellis Metadata Arrays are available in the following 2 RU, SSD-based array configurations:

- Quantum QXS™-424, 12 x 400 GB SSD
- Quantum QXS-424, 12x 800 GB SSD
- Quantum QXS-424, 24x 800 GB SSD

XLS-CONFIG AND CONNECTIVITY

A significant new function of the Xcellis system, and more specifically, the Xcellis Workflow Director is XLS-Config. Based on the modern Linux “slot/port” device mapping standard rather than “Ethx,” this is a simple and virtually automatic way to provision or modify the networking connectivity in an Xcellis system using the three available connectivity slots. XLS-Config discovers networking interfaces when they are installed in the Workflow Director, regardless of slot and protocol. The system also supports multiple instances of the same type of connector. As new cards for additional protocols become available over time, they will be easy to install and incorporate into the workflow.

Once installed, management of the cards, including configuration and bonding is performed through the StorNext GUI. The illustration below shows the rear panel of the Xcellis Workflow Director and the location of the XLS-Config slots. See the table below for the current possible configurations for Xcellis connectivity. Xcellis supports up to three populated slots per node, with two nodes per Xcellis Workflow Director. Slots are most easily thought of as providing a particular service to the system.



- Slot 1: IP Storage for Lattus™ Object Storage, Q-Cloud™ Archive and Q-Cloud Vault connections. Not required. Can be Quad 1GbE or Dual 10GbE.
- Slot 2: Collaborative Sharing for NAS and DLC client connectivity. Not required. Can be Quad 1GbE or Dual 10GbE.
- Slot 3: Fibre Channel connectivity for metadata and disk arrays and tape systems. This function is required. Can be Dual 16Gb or Quad 8Gb.

Port / Slot Identification	Onboard			Slot 1	Slot 2	Slot 3	
Function	Service	GUI	Metadata	Object/Cloud IP	IP Clients (Sharing Ports)	Disk FC	Tape FC
Fibre Channel Disk Only	1GbE iDRAC (Connect)	1GbE	Dual 1GbE	None	None	Dual 16 Gb FC	None
Fibre Channel Disk and Tape				None	None	Quad 8Gb FC	
Fibre Channel Disk and Tape 8 x 1GbE				Quad 1GbE	Quad 1GbE		
Fibre Channel Disk and Tape 4 x 1GbE 2 x 10GbE				Dual 10GbE			
Fibre Channel Disk 2 x 10GbE				None	Dual 10GbE	Dual 16Gb FC	None
Fibre Channel Disk and Tape 2 x 10GbE				None		Quad 8Gb FC	
Fibre Channel Disk and Tape 4 x 1GbE 2 x 10GbE				Quad 1GbE			
Fibre Channel Disk and Tape 4 x 10GbE				Dual 10GbE			
Fibre Channel Disk 4 x 1GbE				None	Quad 1GbE	Dual 16Gb FC	None
Fibre Channel Disk 8 x 1GbE				Quad 1GbE			
Fibre Channel Disk 2 x 10GbE 4 x 1GbE				Dual 10GbE			
Fibre Channel Disk 2 x 10GbE				None	Dual 10GbE	Dual 16Gb FC	None
Fibre Channel Disk 4 x 1GbE 2 x 10GbE				Quad 1GbE			
Fibre Channel Disk 4 x 10GbE				Dual 10GbE			

XLS-LUNS AND DISK CONFIGURATION

Storage configurations can be complicated and time consuming. Xcellis includes XLS-LUNs, an API-based system to create RAID sets and LUNs in a “blank” system through either StorNext Connect or the Xcellis service menu. It’s designed to make disk configuration easy, automatic and reliable. Storage experts will find it quick to deploy storage without the loss of flexibility and control. Less experienced customers will be able to use XLS-LUNs themselves to deploy recommended configurations.

Standard Configurations: Combined Metadata Storage and User Storage

XLS-LUNs identifies the array configuration uses a pre-defined RAID6 layout. Metadata, HAFS, and user data LUNs are created out of each RAID. These functions take up about 3% of the total capacity. Systems are designed to provisioning enough metadata to support an entire array filled with 200KB files. This allows Xcellis to scale continuously to the largest possible capacities.

At the time this brief was written, customer installation is support through StorNext Connect for up to 16 disk chassis. Configurations of more than 16 chassis require Quantum professional services or a Quantum ASI partner to perform the install.

The following table describes the range of configuration options for Xcellis storage with combined metadata and user data.

Combined Metadata and User Data				
Model / Drive	Disk Group & Format	File Systems per Shelf	LUN sizes	File count per shelf
QXS-412 4 TB HDD	10+2	1	HAFS: 800 GB / MD: 400 GB / Data 38.8 TB	195,312,500
QXS-412 6 TB HDD	10+2	1	HAFS: 1.2 TB / MD: 600 GB / Data 58.2 TB	292,968,750
QXS-456 4 TB HDD	4x 12+2	1-4	HAFS: 3.84 TB / MD: 1.92 TB / Data 186.24 TB	937,500,000
QXS-456 6 TB HDD	4x 12+2	1-4	HAFS: 5.76 TB / MD: 2.88 TB / Data 279.36 TB	1,406,250,000

Optional Configurations: Stand-Alone Metadata Arrays for Use With 3rd-Party Storage

XLS-LUNs can be used to simplify the layout of storage for an Xcellis Metadata Array. These arrays are available in either 12-drive or 24-drive SSD configurations, organized in RAID6 10+2 for each set of drives. Each of the RAID sets is configured into a QXS Virtual Pool and HAFS is a static capacity per RAID.

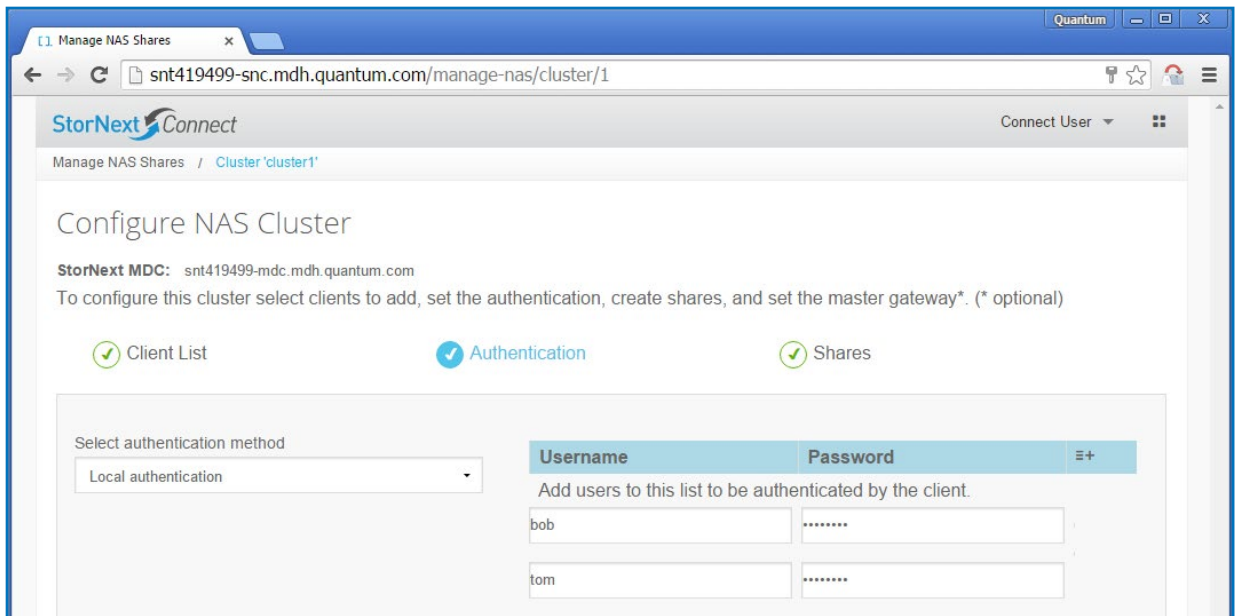
During setup, users request the desired size of the metadata LUNs through the StorNext Connect GUI using XLS-LUNs. The table below is an example that assumes all metadata LUNs are the same size.

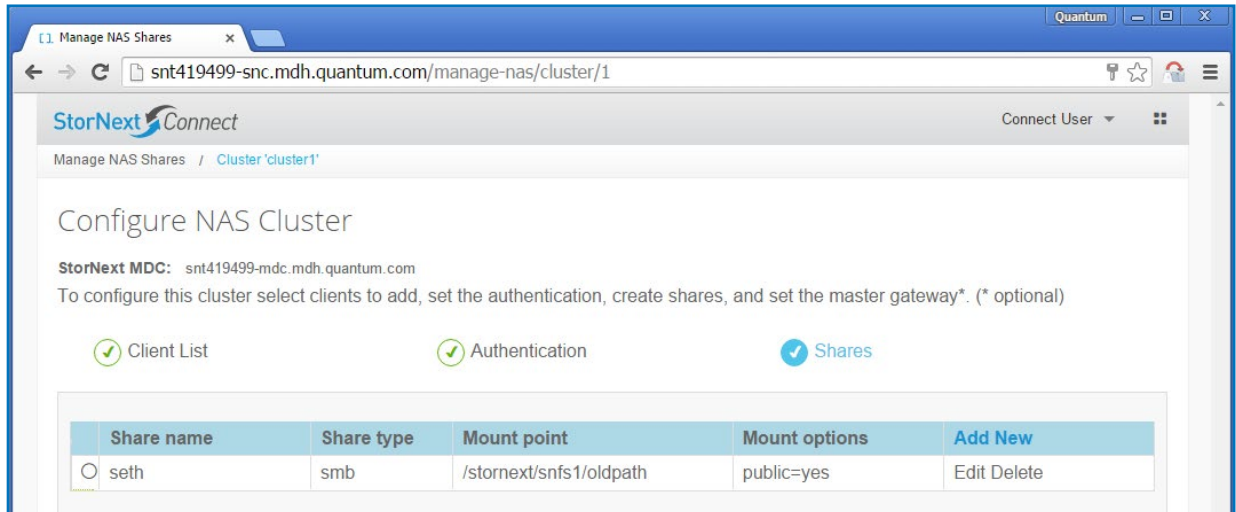
Dedicated Metadata Array: Provisioning and File Count					
Drives	Disk Group & Format	File Systems Supported	LUN Sizes	Managed File Count	Maximum File Count
QXS-424 (12) 400 GB SSD	1x 10x2 HAFS & metadata	1-32	HAFS: 1.437 TB MD: 2.563 TB	350,000,128	1,251,520,512
QXS-424 (12) 800 GB SSD	1x 10x2 HAFS & metadata	1-32	HAFS: 2.867 TB MD: 5.126 TB	700,000,256	2,503,069,696
QXS-424 (24) 800 GB SSD	2x 10x2 HAFS & metadata	1-64	HAFS: 5.734 TB MD: 10.253 TB	1,400,000,512	5,006,139,392

NAS CONNECTIVITY

With the introduction of Xcellis, StorNext provides access to content for Ethernet-based clients using SMB and NFS protocols. These connections allows users who don't require direct Fibre Channel performance and connectivity to their files the ability to participate directly in the workflow without the need for separate storage, file transfers, or sneaker-net. By using NAS connectivity in a StorNext environment, all clients, including those on Ethernet, have visibility of files in tape and cloud archives without administrative support.

Configuration and management of NAS clients within the StorNext environment is done through StorNext Connect. The screen images below illustrate the provisioning of user authentication and NAS shares.





Xcellis supports a variety of NAS protocols including SMB 1 (CIFS), SMB 2, SMB 3, and NFS 3 with multiple shares per StorNext file system. Supported authentication models include Microsoft Active Directory, Open LDAP with Kerberos support, Apple Open Directory, and users created locally through StorNext Connect. In a Windows environment, Xcellis provides ACL compatibility between SMB, Windows SAN client, Windows DLC clients. SMB protocols support offline and archive bits so that users have visibility to files that have been archived and removed from primary storage.

Additional features include batch import of shares for large environments, system logging, and easy creation of support data. The following table lists the current capabilities of Xcellis NAS support.

Maximum Number of Connected Users	100 directly on Xcellis 500 per dedicated gateway
Maximum Number of Concurrent Data Streams	100
Supported NAS Protocols	SMB 1 SMB 2 SMB 3 NFS 3
Supported User Authentication Schemes	Microsoft Active Directory 2008 and 2012 <ul style="list-style-type: none"> • With Unix extensions (RFC2307) • With RID UID mapping OpenLDAP with Samba 3.2 Schema extensions OpenLDAP with Kerberos Apple Open Directory Local users created on the Artico appliance <i>Note: only one authentication scheme can be used; Apple Open Directory requires manual configuration via CLI</i>
Client Support	Windows, MacOS, Linux

CONCLUSION

Quantum Xcellis represents an unprecedented advancement in shared, scale-out storage. The combination of expanded connectivity, increased scalability, improved efficiency and simplified operations reaffirms StorNext as the storage environment of choice for demanding workflows and establishes Xcellis as the pre-eminent platform for StorNext deployments.

For more information, visit www.quantum.com/xcellis

ABOUT QUANTUM

Quantum is a leading expert in scale-out storage, archive, and data protection, providing solutions for capturing, sharing, and preserving digital assets over the entire data lifecycle. From small businesses to major enterprises, more than 100,000 customers have trusted Quantum to address their most demanding data workflow challenges. Quantum's end-to-end, tiered storage foundation enables customers to maximize the value of their data by making it accessible whenever and wherever needed, retaining it indefinitely and reducing total cost and complexity. See how at www.quantum.com/customerstories.