Quantum

WHITE PAPER

SELF-PROTECTING STORAGE: QUANTUM XCELLIS SCALE-OUT NAS

Integrated Data Protection for Unstructured Data

CONTENTS

INTRODUCTION

It's a modern version of Plutarch's classic causality dilemma about the chicken and the egg. Is exploding data driving demand for storage, or is affordable storage driving the accumulation of data? As with poultry, the answer is probably a bit of both. What is clear is that the mountains of data that modern organizations collect need to be stored, shared, and protected. The larger the volume and higher the velocity, the more difficult these tasks become. And the stakes are high: For high-value workflows where data is the product, getting it right—right now—can be the difference between profit and loss, growth or decline, falling behind or getting ahead.

Quantum Xcellis® Scale-out NAS is purpose built to store, share and protect the massive amounts of information data-driven organizations require. Xcellis Scale-out NAS appliances are deployed today in the most demanding data environments in government, media & entertainment, life sciences, scientific research, geospatial imaging, HPC and more. This paper will reveal how Xcellis Scale-out NAS can provide multi-dimensional data protection for the unstructured data you rely on.

FIRST PRINCIPALS OF DATA PROTECTION

Data protection used to be simple. Every night after everyone went home, a full backup of all corporate assets would run and be sent off-site for safekeeping. But as data grew, things got complicated. When it was too much to back up everything every night, incremental and differential backups became standard. When better Recovery Time Objectives and Recovery Point Objectives (RTOs and RPOs) were demanded, snapshots filled the gap. Disk-based backup made recovery easier but cost a lot, so deduplication was invented. The weekend got too short for a full backup; thus synthetic full backups, and so on. Unfortunately, none of these techniques work well for the truly huge data repositories that are common today.

It's time to step off the technology treadmill and return to first principals:

- Data growth is accelerating.
- The day isn't getting any longer.
- Data storage is getting cheaper, but not cheap enough or quickly enough.
- Time spent thinking about data protection would be better spent thinking about the business or the mission.

The ideal solution for data storage, then, is infinitely scalable in both performance and capacity, includes integrated, real-time data protection, intelligent options for controlling the cost of storage, and doesn't require months of vendor professional services to install or an army of consultants to maintain.

These truths aren't new. For the past twenty years, many of the world's biggest producers and consumers of data have used StorNext®—the software engine that drives Xcellis Scale-out NAS—to face them. Thanks to this technology heritage, Xcellis Scale-out NAS comes closer to the ideal than any other solution.

THE DATA THREAT LANDSCAPE

Threats to data come in many different forms, but broadly fall into two types. There are things that threaten *the very existence* of data—threats of data loss. Then there are threats of *unwanted exposure* of data—threats to data security. This document primarily focuses on the former, but will briefly touch on data security toward the end.

Table 1 summarizes the threats to data:

Threat Category	Description	Mitigation Strategy
User Error	Accidental deletion or modification of files and directories	Tools to recover deleted files and rollback to previous versions
Hardware Failure	Failure or destruction of storage media or related hardware needed to access the media	Deploy resilient/redundant hardware that will continue operating through common failures
Data Corruption	File, file system, or metadata corruption due to hardware or software failure, media degredation over time ('bit rot'), cosmic rays or quantum effects	Provide means to automatically detect corruption and store enough additional data (parity or full copies) to automatically recover from corruption
Site Loss	Catastrophic loss or extended unavailability of all the data at a site	Store additional copies of data at one or more alternate sites
Ransomware	Data encrypted and held hostage by malicious actors	Maintain an 'air-gapped' off-line copy of data
Technology Obsolescence	Data stranded on obsolete storage technology	Use an open platform that provides data migration functions and has a history of embracing new storage technology

Table 1: Threats to Data

The good news is that, with a little forethought and planning, it is possible to protect against every one of these threats, and it doesn't have to be overly complicated or obscenely expensive to do so. The unique features of Xcellis Scale-out NAS make it easy to protect even multi-petabyte data stores while controlling cost, minimizing complexity, and maintaining flexibility. The range of options provided allows easy customization to match data protection capabilities to need, and evolve the configuration over time.

HOW XCELLIS SCALE-OUT NAS PROTECTS DATA

No two organizations are alike when it comes to data protection needs, and even within the same organization there exists a range of data lifecycles and requirements. A storage system must be able to accommodate policies ranging from "throw it out when the job is done" processing to "keep it forever," and everything in between. The cost of storage demands control over where data lives and when, and also how many copies of it exist and where those copies reside over time.

Xcellis Scale-out NAS incorporates the broadest range of features for data protection, making it possible to configure it to fit today's—and tomorrow's—requirements. Table 2 below summarizes how specific Xcellis Scale-out NAS features protect against the various threat categories. Note that there are multiple remediation options for each threat category, and most of these features are configurable with directory-level granularity. The details of each feature are explained following the table.

	Threat Category					
Xcellis Scale-Out NAS Feature	**User Error**	**Hardware Failure**	**Data Corruption**	**Site Loss**	**Ransom-ware**	**Obsolete Technology**
RAID / DDP		■	■			
Erasure Coding		■	■	■		
Hardware Redundancy		■				
Checksums			■			
EDLM			■			
System Backup		■		■		
Copies		■	■	■	■	■
Versions	■		■		■	
FlexSync	■	■		■	■	
FlexTier				■	■	
Alternate Store Location		■	■	■	■	
Vaulting				■	■	
Open Storage Options						■

Table 2: Feature/Threat Matrix

RAID/DDP

RAID is the familiar method of striping data and parity blocks among a set of disks to achieve redundancy. This allows recovery from a disk failure by rebuilding the missing data onto a hot spare disk. Dynamic Disk Pooling (DDP), also known as distributed RAID, is a newer technology that achieves the same result in a different way. With DDP, data, parity, and hot spare capacity are striped across a set of disks. This can drastically shorten rebuild times vs. traditional RAID, especially for high-capacity disk drives. With RAID, a rebuild happens on a single hot spare disk, which becomes a serious bottleneck. In DDP, the rebuild process happens in parallel across all remaining drives in the array. RAID and DDP also protect against data corruption to a degree. If a block becomes corrupt, it will be detected and recovered from the parity information present on the remaining disks.

Xcellis Scale-out NAS appliances leverage RAID 1 (mirroring) for the operating system disks, and RAID 6, which can survive two disk failures, for the file system metadata. Quantum QXS™ arrays used for user data are typically also configured with RAID 6. Quantum QD and QS user data arrays offer the choice of either RAID or DDP. If these options aren't suitable, customers may choose from a broad selection of third-party storage arrays that work with Xcellis Scale-out NAS.

Erasure Coding

Erasure coding is a family of mathematical techniques commonly used in large-scale object storage applications. Data is broken into segments, and each segment is encoded into a series of equations. To decode a segment, you only need a subset of the equations. This enables storage systems to be built that have very high data durability (resistance to data loss or corruption), with much better storage efficiency than simpler techniques such as "store three copies of everything." Storage systems using erasure coding are able to continually scan for corruption, and can withstand many individual disk failures or corruptions before the risk of data loss. Erasure coding also enables an object store to be spread across geography—known as geo-spreading—providing protection against site loss.

Xcellis Scale-out NAS supports Quantum's erasure-code-based Lattus™ object store as a tier, and the FlexTier™ option supports all of the most popular third-party object stores to preserve customer choice.

Hardware Redundancy

The components that make up Xcellis Scale-out NAS include a generous level of hardware redundancy to protect the availability of the data stored within. Redundant power supplies and fans are standard, as are dual array controllers in Quantum-provided disk arrays. Redundant network interfaces are available. The system is normally run with a clustered pair of Xcellis Workflow Director (WFD) nodes managing the file system, and up to 16 Xcellis Workflow Extenders (WFE) serving as a cluster of NAS heads. If a large amount of data tiering movement is expected, additional WFEs may be configured as Distributed Data Movers (DDMs). Quantum realizes, however, that not everyone requires extreme availability and redundancy. For less-critical applications and more-constrained budgets, systems with a single WFD node are also available.

Checksums

When enabled, Xcellis Scale-out NAS will generate a 128-bit MD5 checksum for each file stored, and record the checksum in the database. Optionally, upon retrieval, a checksum will be calculated on the retrieved file and compared to the stored value. If the values don't match (indicating the retrieved file is corrupt), the Storage Manager policy and tiering engine will discard the corrupt copy, retrieve an alternate copy of the file, and notify the administrator of the failure. The party accessing the file will experience a momentary additional delay in retrieval, but the file access will succeed. Although modern storage media, devices, and networks are very reliable, checksums enable another layer of assurance that the data retrieved is identical to the data originally stored.

Extended Data Lifecycle Management (EDLM)

EDLM is a capability unique to Quantum Scalar® tape libraries that is integrated with Xcellis Scale-out NAS. One of the main benefits of using tape for cold storage is that the media itself is completely off-line when not actively being read or written. In that state, it takes no power or cooling, only a little bit of space, and offers up to a 30-year rated archival shelf life. But all data storage media—including tape—degrades over time. The rate of decline depends on many factors, including usage pattern and storage conditions. The only way to ensure that your files are safe is to periodically test the media, analyze its condition, and replace degraded media proactively before data is lost. This is the purpose of EDLM.

The system administrator configures EDLM policies to scan media periodically or based upon certain error events, selects a scan depth, and chooses whether Xcellis Scale-out NAS should automatically recover or simply raise an alert. Based on analysis of tape alerts, soft (recoverable) errors, and other low-level data from the tape drives, media that is beginning to degrade is flagged. If automatic recovery is selected, the policy engine within Xcellis Scale-out NAS automatically copies the files to fresh media, which doesn't even have to be of the same type. For example, files from several failing LTO-5 cartridges could be migrated to a single LTO-8 cartridge. The copy may even span two libraries.

System Backup

The purpose of the system backup function is to protect the system configuration and metadata to ease recovery in case of disaster. By default, a full backup of the entire system database, all configuration files, and the complete file system metadata archive is performed once per week. A partial backup containing database and file system journal files as well as configuration files is run on all other days of the week. The schedule and type of backup may be modified by the administrator. Backup files are written to any configured storage tier, which can include object storage, cloud, replicated disk, or tape for off-site protection.

While it is possible to protect data stored on an unmanaged Xcellis Scale-out NAS using an old-school batch backup application, it is far more cost effective and time efficient to instead leverage built-in policy functions such as copies and versions.

Copies

A copy is the fundamental unit of data protection, and key to many of the protection functions within Xcellis Scale-out NAS. Up to four copies of each file written may be generated. Copies are entirely "behind the scenes" and invisible behind a single namespace. By default, they are made after a file has been inactive for five minutes, but this parameter is tunable. This provides near-real-time protection as files are changed. Copies may be written to any defined storage tier, including disk, tape, object storage and public cloud. When files are deleted, the copies are retained for a configurable period, enabling recovery from accidents. A copy written to an off-site target provides disaster recovery protection against storage device or total site loss. An off-line copy (such as a tape in a vault) protects against ransomware. Copies are used to automatically recover from data corruption detected by Xcellis Scale-out NAS, and may even be leveraged to migrate data off an older storage system onto a new one, protecting data against technology obsolescence.

Versions

While copies enable accidentally deleted files to be recovered, versions take recoverability to another level. Up to 45 versions of a file may be retained, providing point-in-time rollback to a previous state. This is useful when files have been accidentally changed, but is essential to protect against ransomware. If files are maliciously encrypted, reverting to the previous version saves you from paying bitcoin ransom to the criminals.

FlexSync

FlexSync™ is the Xcellis Scale-out NAS asynchronous replication and synchronization engine, designed to efficiently create local or remote replicas of file system data and metadata. It can be used to protect an entire file system, a specific directory, or anything in between. It supports one-to-one, one-to-many, and many-to-one topologies. Because FlexSync is integrated with StorNext, it knows instantly when changes are made to protected files, and incrementally synchronizes the changes to the destination system. If the destination is remote, delta block compression is used to send only the changed blocks, increasing speed and reducing bandwidth requirements. If versioning is enabled, the destination system retains a configurable history of changes to enable easy recovery from user error or ransomware attack. Multi-threading and multi-streaming make FlexSync incredibly fast. Cruder tools like rsync rely on scanning the file system to detect changes, which robs performance and doesn't scale. They also don't protect critical metadata such as Windows ACLs, extended attributes, and named streams. FlexSync is complete, fast, efficient, scalable, and easy to manage via GUI or CLI.

FlexTier

Another member of the 'Flex' family of features for Xcellis Scale-out NAS is FlexTier. FlexTier enables Xcellis Scale-out NAS to use object stores and cloud destinations as storage tiers. Virtually any target that speaks S3 may be used, including all the most popular object storage systems and public clouds. Because FlexTier is integrated into the Xcellis Scale-out NAS policy engine, it works together with the copy and versioning features to provide protection against site loss and ransomware attack. Writing a copy to a cloud target or geo-spread object store gets data off-site, while adding versioning to the mix enables rollback to a pre-attack point in time.

Alternate Store/Retrieve Location

The Alternate Store Location (ASL) feature provides an automated ability to copy all or part of a file system from a local Xcellis Scale-out NAS to a remote system. The remote copy is made while other copies are made from the local Xcellis Scale-out NAS to local storage tiers or cloud. Once at the remote site, files are managed according to the policies on the target Xcellis Scale-out NAS system. This enables a variety of protection architectures. For example, make one copy from Xcellis Scale-out NAS at site A to cloud with FlexTier, a second copy from site A Xcellis Scale-out NAS to site B Xcellis Scale-out NAS using ASL, and a third copy from site B Xcellis Scale-out NAS to a tape library at site B. Once copied, files landing at the remote site are not monitored by the originating system. This makes ASL suitable for "pushing" content to a second site for later use and modification there.

Alternate Retrieve Location (ARL) enables copies made with ASL to be leveraged as "copies of last resort." With ARL, if all local copies of a file are corrupt or unavailable, the local Xcellis Scale-out NAS system will reach out to the remote Xcellis Scale-out NAS system and pull a copy back over the network. When using ARL this way, the remote copy is normally configured as read-only.

Vaulting

Vaulting is the concept of taking a copy of data completely off-line. This can literally mean tapes that are stored in a safe at another site, or an "Active Vault" partition within a Quantum Scalar i6000 tape library. A vaulted copy may be physically out of sight, but the Xcellis Scale-out NAS policy engine remains aware of it so administrators always know where their data is, even when it's on a tape in a safe. If a file that only exists on a vaulted cartridge is requested, the system notifies the administrator to insert the cartridge into a library to be read.

Vaulting provides a critical "air-gapped" copy of data—meaning one that can't be accessed automatically. This is the most secure form of protection against ransomware, viruses, and similar threats, even for the Scalar i6000 Active Vault. An Active Vault can operate in a remote or lights-out data center, as no human has to physically touch the media or the library to use this feature. By design, moving media from the Active Vault partition back to the Xcellis-managed partition still requires an explicit administrator action. This has the unique advantage of providing an air-gapped copy of data, while all cartridges remain safely locked within the Scalar i6000 library.

Open Storage Options

The StorNext software that drives Xcellis Scale-out NAS has always been hardware and protocol agnostic. It is able to easily adapt and grow as storage technology changes, and even migrate data forward onto new technology. Today, the core Xcellis Scale-out NAS components are sold as appliances to simplify integration and provide predictable performance, but customer data may be stored on a wide range of Quantum and third-party storage targets. For the current list of supported disk, tape, object storage, and cloud targets, reference the compatibility guides on the Quantum web site here: https://www.quantum.com/serviceandsupport/compatibilityguides/index.aspx

DATA SECURITY

As mentioned at the beginning, data security is not the same as data protection but is an important related topic not covered in depth here. Xcellis Scale-out NAS offers a range of methods and options to enable and enhance data security, including:

- Cross-platform ACLs and immutable file support
- Encryption of data at rest, including Self-Encrypting Disk (SED) options and encrypted LTO tape
- Encryption of data in transit
- Support for LTO WORM media for immutable storage of data
- Audit logs
- Active Directory/LDAP integration
- Integration with third-party security products including those from Varonis and Vormetric.

For more on the security-related features of Xcellis Scale-out NAS, refer to the Documentation Center on Quantum.com.

CONCLUSION

Threats to data come in many forms—some old, some new. The old models of data protection don't work for modern, data-driven organizations. To maximize efficiency and minimize cost, data protection must be integrated within high-performance workflow storage itself—not be a difficult-to-manage, bolt-on afterthought. Because Xcellis Scale-out NAS includes a wide range of built-in data protection features, it's easy to construct a customized data protection scheme to fit any need, and adjust it when needs change.

Quantum®

ABOUT QUANTUM

Quantum is a leading expert in scale-out tiered 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**.

www.quantum.com • 800-677-6268

WP00240A-v01 Apr 2018