Quantum®

WHITE PAPER

THE QUANTUM CLOUD STORAGE PLATFORM

An Architectural Overview

CONTENTS

INTRODUCTION

The Quantum Cloud Storage Platform (Quantum CSP) is a software-defined storage platform specifically designed for storing video and other forms of high-resolution content. Like storage technologies used in the cloud, the Quantum CSP is software-defined, can be deployed on bare metal, as a VM, as part of a hyperconverged infrastructure, or in other 'cloud-ready' infrastructures.

Unlike other software-defined storage technologies, the Quantum CSP was designed specifically for video and other forms of high-resolution content—engineered for extremely low latency, and to maximize streaming performance of large files to storage.

Among other things, the Quantum CSP provides block storage services at the core of several Quantum products, including:

- The Quantum F-Series NVMe storage array used for studio editing and rendering
- Quantum VS-Series for surveillance recording and industrial IoT
- Future hyperconverged and software-defined storage products on Quantum's roadmap

This white paper explains the Quantum Cloud Storage Platform: its purpose, details on the architecture, and its capabilities.

A STORAGE PLATFORM DESIGNED FOR VIDEO AND OTHER HIGH-RES CONTENT

Video and video-like data—or large unstructured data—is growing exponentially in every industry. It is growing much faster than databases or other forms of structured data, and already represents the vast majority of the data on the planet.

Some examples of this type of data include:

- **Movie and TV production** – where resolution of the content is increasing exponentially from 1K HD to 4K, 8K, and beyond.
- **Corporate video** – For many companies, video is becoming the primary medium for communication, in marketing, employee training, and more.
- **Video surveillance** – Surveillance footage is the biggest source of video on the planet, in fact surveillance cameras are the "IoT devices" that generate more data than all other IoT devices put together.
- **Autonomous vehicle design** – A self-driving car is outfitted with many sensors and cameras, including video, GPS, RADAR, LIDAR, and more, all of which generate continuous streams of high-resolution unstructured data.
- **Life sciences research** – Genomics analysis pipelines and imaging technologies such as light sheet and cryo-electron microscopy can generate tens of TB per instrument per day.
- **Geospatial research** – Satellite imagery consists of extremely high-resolution images in multiple wavelengths as well as LIDAR data.
- **Military and defense** – The data coming from satellites and drones is predominantly video and high-resolution images.
- **Manufacturing** – Manufacturers use high-resolution images and other sensor data for factory automation and quality assurance purposes.

This type of data presents a unique set of challenges:

- This data is typically 50x larger than the average corporate database.
- This data can't be compressed or deduplicated.
- It is difficult to 'snap' or replicate this data due to the size of the files.
- It is difficult technically, and not feasible economically, to transport and store this data in the public cloud.
- Because this data is 'unstructured' – it is difficult to search and analyze, despite the tremendous value contained in this data.

In short, many of the data services that Enterprise storage products have built over the years simply don't apply to this type of data.

In response to these requirements, Quantum built a storage platform to meet the needs of tomorrow's data—the Quantum Cloud Storage Platform. By stripping out inessential data services, we've engineered a platform to deliver low latency and maximum streaming performance for these data sets, while maintaining a core set of services that are required by our customers in these markets.

OUR HISTORY WITH VIDEO & SOFTWARE-DEFINED ARCHITECTURES

Quantum has over 20 years of experience as a technology leader dealing with digital video. Since the late 1990s, the StorNext® File System has been sharing high-resolution digital video, and our tape storage products were always used to archive images, video, and other 'large file' data.

Quantum's StorNext shared file system has long been the standard for studio editing in movie and TV production and post-production, and our experience in these applications has made us experts in ingesting, storing, and sharing video content.

In addition, with StorNext, Quantum was an early pioneer in software-defined storage. Over its 22-year history, StorNext has been deployed with a massive variety of server hardware, operating systems, block storage, tape storage, object storage, and now cloud—and StorNext is able to store and manage data across all those tiers within a single namespace.

In developing the Quantum Cloud Storage Platform, we have applied our years of experience working with video, plus our years of experience deploying and supporting software-defined storage architectures, to create this new storage platform.

QUANTUM CLOUD STORAGE PLATFORM ARCHITECTURE

To explain the architecture of the Quantum Cloud Storage Platform, it is helpful to think of it as consisting of several layers that reside between the raw storage devices and the application, shown in Figure 1. Each layer provides specific functions and interfaces to the layers above and below. Layers may be entirely software-defined or take advantage of hardware assist as appropriate.

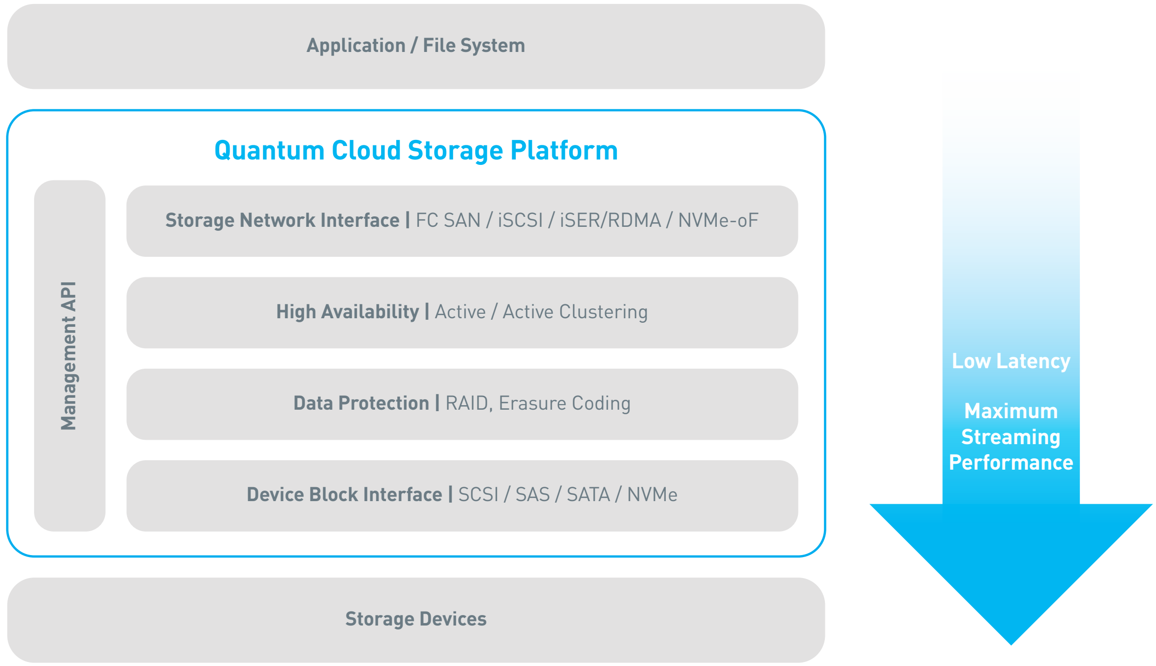

Figure 1 - Quantum Cloud Storage Platform Architecture

What follows is a thorough description of these functional layers and technology options for each one, starting from the bottom and moving upward.

Storage Devices

Below Quantum CSP are the storage devices themselves. The key task for a storage system is to serve up available storage performance and capacity to applications in a consumable way, while not getting in the way. Quantum CSP can manage any common form of primary storage device, including hard disk (HDD), solid-state disk (SSD), and persistent memory (PM) devices attached via PCIe, SAS, or SATA. Device drivers make the individual storage devices visible to the layer above. Since Quantum CSP supports multiple storage device types, it is easily optimized for specific use cases. If high performance is needed in a small footprint, PM or NVMe flash may be selected. For applications where high capacity, moderate performance, and lower cost are important, traditional SAS HDDs are usually the best fit.

Device Block Interface Layer

Applications that consume block storage do so with the help of a block layer in the operating system. "Application" in this context refers to the software that is the direct consumer of the block storage. Most often this is a file system, but not always. Some applications can write to raw block storage without going through a file system. Below the block layer reside drivers for individual storage device types. Communication with the storage devices is via a standard interface such as SATA, SAS, or NVMe. Quantum CSP may use traditional SCSI block devices or newer NVMe technology devices.

Data Protection Layer

Above the device block interface is the data protection layer. Data protection is a critically important function required of a storage system. If you are not certain that you can access the data that you have stored, the system is useless. Another key task for a storage platform, therefore, is to provide data reliability and redundancy. Storage devices themselves, whether spinning HDD or flash SSDs, have built-in safeguards to prevent data loss and corruption. Data is verified with CRCs and errors are corrected using ECC where possible. Bad areas of the storage are remapped on the fly to a backup pool of reserved space. This takes care of the day-to-day data integrity issues, but eventually individual storage devices will fail unrecoverably, sometimes with no warning. This is where the Quantum CSP data protection layer fits in.

Protection of data across multiple individual storage devices means using traditional RAID techniques or erasure coding. The idea is to spread data and parity information across several devices so the data is safe from one or more device failures. The best approach depends on the storage devices chosen (HDD vs. SSD), the amount of protection required, and the desired performance.

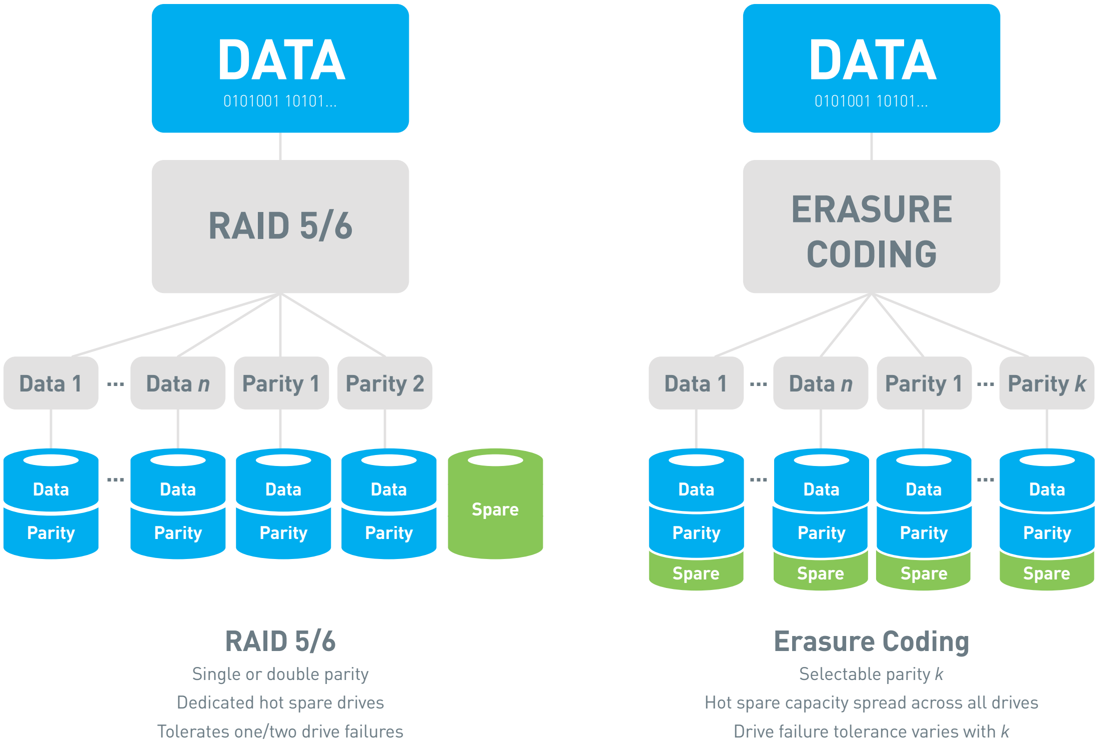

Figure 2 - RAID 5/6 vs. Erasure Coding

Device type, protection level, and performance are all interdependent. For example, when a drive fails, it takes longer to rebuild a RAID set with HDD vs. SSD, so for HDD, RAID6 (dual parity) is preferred to minimize the window of time when additional failures would result in data loss. Triple parity may even be required for HDD as drive capacities grow above 15 TB. For SSD, RAID5 (single parity) is usually enough because rebuild times are much shorter than with HDD. RAID5 also offers higher storage efficiency and thus lower cost per TB than RAID6. Higher performance requirements during device failure and rebuild make SSD or HDD with erasure coding more attractive than HDD with RAID. Storage efficiency, massive scalability, and the option to spread data across a cluster (even a geographically dispersed cluster) also favor erasure coding over RAID.

There is another important aspect to consider when choosing a protection scheme—performance. Particularly with HDD, RAID is used not just to provide protection, but to aggregate device performance by striping I/O across multiple devices. Different RAID levels have different performance characteristics, which differ—sometimes dramatically—for read vs. write operations. Nested RAID arrangements such as RAID10 or RAID60 may be needed to achieve the proper blend of protection and performance for a given workload.

These protection algorithms may be run entirely in software using the main CPUs, or on a purpose-built hardware adapter such as a RAID controller. Running RAID parity calculations on the main CPU enables fully software-defined configurations, but there are cost considerations if the CPU must be upsized for this additional load. Depending on the hardware configuration and workload, it may make more sense to deploy a dedicated hardware RAID controller. RAID controllers also offer the option of battery or supercap-backed write cache, which is important for some applications. Erasure coding is essentially always software-defined, though it may make use of CPU-specific accelerated instruction sets.

Quantum CSP can be deployed in a software-defined configuration or can take advantage of hardware assist for the protection layer, as appropriate, and can support both RAID and erasure coding options.

High-Availability Layer

It's important to note that data availability is not a feature, it's a goal that is reached by designing a resilient system with the intended use in mind. Availability features may be implemented in various places in the architecture depending on the use case. For example, a highly transactional database system with lots of small random I/O might be well served by array-based block-level replication or snapshots. A system with mostly large files accessed sequentially (such as 4K video) can be better served by performing these functions in the file system, unburdening the array controllers to focus on providing high-performance I/O.

Data availability always requires hardware redundancy. A single system unit may house redundant compute, storage, power, and network interfaces. Clustering is another approach, where multiple nodes exist, each one totally without internal redundancy, but sharing duties to together provide the required level of data availability.

Basic hardware availability features such as redundant power supplies and fans reside below the Quantum CSP layers. Above the hardware, Quantum CSP provides clustering capabilities that support single and multi-node active/active and active/passive architectures today. Full shared-nothing cluster support is on the roadmap. When needed, replication occurs above Quantum CSP at the file system or application level.

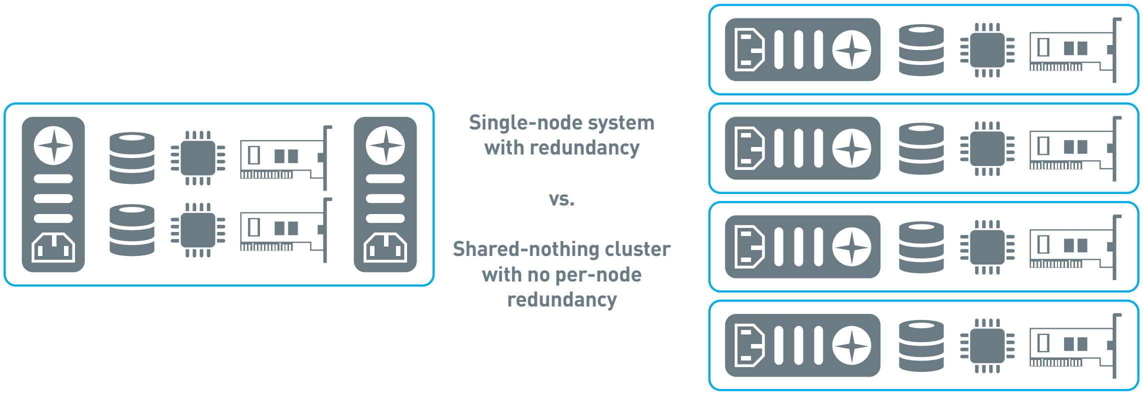

Storage Network Interface Layer

The highest layer in the stack is the storage network interface layer. There are a lot of different networking technologies that can be used to create a storage network, including TCP/IP, Fibre Channel (FC), and InfiniBand (IB). Newer options include RDMA-capable Ethernet transports such as Internet Wide-Area RDMA Protocol (iWARP) and RDMA over Converged Ethernet (RoCE). On top of these transports run SCSI, iSCSI, iSCSI extensions for RDMA (iSER) and NVMe over Fabrics (NVMe-oF). In addition to these standards-based approaches are a handful of proprietary, semi-proprietary, and lesser-known options.

Protocol	Transport	Network
iSCSI	TCP/IP	Ethernet
iSER	RDMA (RoCE, iWARP)	Ethernet
NVMe-oF	RDMA (RoCE, iWARP), FC, IB, TCP/IP	Ethernet, FC, IB
SCSI	FC	FC

The best storage network option for a new project varies, but often depends as much on existing investments in both infrastructure and skills as it does on future strategy. Organizations with large and successful investments in FC may find it easiest to stay that course, especially if they have recently upgraded to 32G or 64G FC technology. If an upgrade is due from older generation 16G or 8G FC, however, the lower cost of Ethernet equipment vs. FC may sway the decision in favor of a converged Ethernet.

For networking low-latency storage technology such as NVMe flash, especially in new deployments without a legacy of FC or IB, RDMA-capable Ethernet is a natural choice. Implementing an iWARP or RoCE foundation enables the use of iSER or NVMe-oF for very low-latency communication. iSCSI is the choice for hyperconverged architectures where applications run on the same physical system alongside the storage.

The Quantum Cloud Storage Platform supports a broad variety of storage networking interfaces, including iSCSI, FC, and iSER, with NVMe-oF on the roadmap. Additional or new networking technologies may be easily incorporated in the future.

SUMMARY

Quantum's Cloud Storage Platform is the first software-defined storage platform specifically designed for video and other forms of high-resolution content. By building a storage platform that provides very low latency and high streaming performance, as well as architectural flexibility to be deployed in a variety of ways, the Quantum Cloud Storage Platform is a key building block that is both in the field today, and will be at the core of Quantum's innovation tomorrow.

Quantum®

ABOUT QUANTUM

Quantum technology and services help customers capture, create, and share digital content—and preserve and protect it for decades at the lowest cost. Quantum's platforms provide the fastest performance for high-resolution video, images, and industrial IoT, with solutions built for every stage of the data lifecycle, from high-performance ingest to real-time collaboration and analysis and low-cost archiving. Every day the world's leading entertainment companies, sports franchises, research scientists, government agencies, enterprises, and cloud providers are making the world happier, safer, and smarter on Quantum. See how at **www.quantum.com**.

www.quantum.com • 800-677-6268

WP00245A-v01 April 2019