

Quantum®

WHITE PAPER

# Reinventing Large-Scale Digital Libraries With Object Storage Technology

## CONTENTS

Introduction .....	3
Hitting The Limits Of RAID .....	3
Object Storage Technology for Scalability and Flexibility .....	4
Lattus with StorNext for Digital Content Libraries .....	7
Conclusion .....	9

## INTRODUCTION

It's a cruel irony that at a time when there are more ways to monetize content than ever before, keeping digital assets accessible to production teams is only becoming more challenging. Much of this challenge is related to providing enough storage capacity, beginning with the initial capture and ingest of raw content. Live events now shoot with more cameras, cameras are less likely to be shut off between takes, and higher-resolution formats are filling up ingest systems faster than you can say "8K 120 fps."

This heavy storage consumption continues throughout the workflow with transcoders spitting out more distribution formats for more connected devices and content owners creating more second-screen content for both live and on-demand markets. The result is that content owners are straining to keep rapidly growing digital assets accessible to their content creators as well as for consumers. Digital libraries are growing exponentially to the petabyte level and beyond.

The all-too-common strategy is to store as much content on high-performance online storage as budgets allow, then move older content to offline tape archives as storage fills. In many facilities, unused raw footage is simply deleted after the project is complete. These are risky strategies given there are now new avenues for remonetizing content, such as distributing the content on new platforms, and reusing the content in new works to shorten production time. Content needs to be readily available to production teams to be effective; otherwise, it's a waste of storage space to save content that's unlikely to ever be used.

With the right storage architecture, content owners can capitalize on enormous revenue opportunities without requiring budgets that scale at the same rate as the content. Content owners need a storage solution that provides disk-speed access from multiple locations, protects content from data loss, scales indefinitely, and stays within modest budgets. That's where Quantum Lattus™, a new category of extended online storage built on object storage technology, comes in.

## HITTING THE LIMITS OF RAID

Disk-based object storage delivers performance yet avoids the pitfalls of using standard disk architectures for large-scale content libraries, starting with RAID, the technology driving most disk storage systems. RAID uses checksums or mirroring to protect data, and spreads the data and checksums across a group of disks referred to as a RAID volume. Using the multi-terabyte disk drives now available, it's common to manage data sets on a single, consistent logical RAID volume of 4-12 drives with a total usable capacity of about 30 terabytes.

The growth of content, though, has outpaced the technology of disk drives. Petabyte-sized data sets either require use of disk volumes larger than 12 disks, which increases risk of data loss from hardware failure; or they require dividing data across multiple RAID volumes, which becomes increasingly inefficient and increases the cost and complexity of managing data consistency and integrity across the multiple volumes.

### **Long Rebuild Times for Large RAID Volumes**

The higher capacity disks used to build these large RAID volumes for digital libraries with high-resolution content also dramatically lengthen the rebuild time when failures occur. The failure of a 3TB or larger drive can result in increased risk and degraded performance for 24 hours or more while the RAID volume rebuilds a replacement drive. Dual and triple parity RAID can mitigate this risk of data corruption, but adding additional parity checks is costly and reduces RAID performance, so it's less and less attractive as the storage scales.

### **Replication Increases Protection, but Decreases Capacity**

Without replication, RAID offers limited protection against node-level failures, and no protection against site-level disasters. Replication improves data integrity, recoverability, and accessibility by copying to a distant secondary location for disaster recovery or to improve data access. But it's expensive. It doubles the amount of storage required, requires expensive high-bandwidth networking between the sites, and introduces new operational complexities to ensure the replication is performing properly or to keep both sites accurately synchronized.

### **Painful Migrations with RAID Upgrades**

Finally, storage upgrades to denser disks typically require building a new RAID volume and migrating data from old to new data volumes. Since RAID requires disks in a volume to have a consistent size and layout, the whole volume must be upgraded at once, and these upgrades can require significant coordination and downtime. The larger the volume, the longer the downtime and the more disruptive the migration is to a production environment.

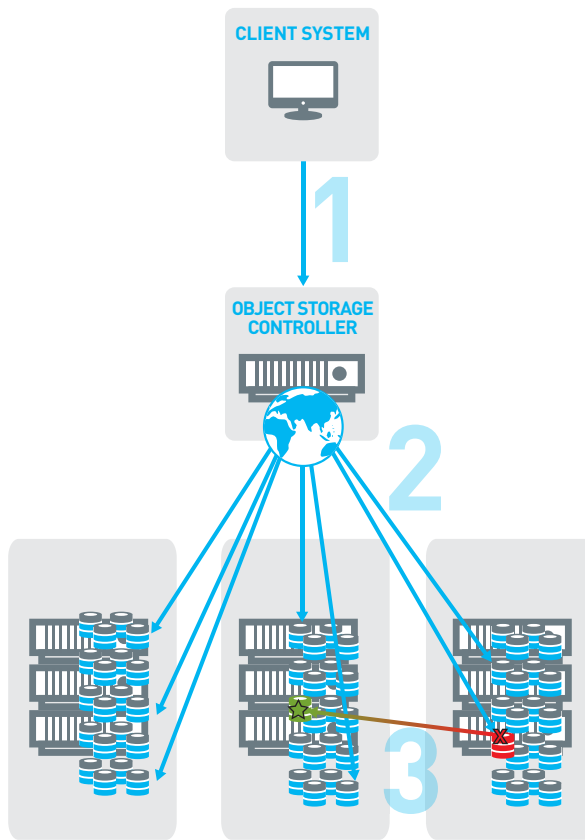
## **OBJECT STORAGE TECHNOLOGY FOR SCALABILITY AND FLEXIBILITY**

Object storage takes a fundamentally different approach by presenting a namespace of simple key and value pairs in contrast to traditional storage systems that organize files in a hierarchy of folders and files mapped to blocks on disk. By using a flat namespace and abstracting the data addressing from the physical storage, object storage-based systems offer more flexibility in how and where content is stored and preserved. Because object storage leverages the scale-out capabilities of IP networks, this addressing allows digital data sets to scale indefinitely. The result is a much more fluid and highly durable environment that is not tightly bound to individual hardware, and in fact seamlessly adapts to changes in the underlying components.

### **Erasur Codes Protect Data Efficiently**

Erasur code algorithms were developed decades ago to ensure transmission integrity of streaming data for space communications. Where RAID slices data into a fixed number of data blocks and checksums, and writes each chunk or checksum onto an independent disk in the volume, erasure coding transforms data objects into a series of codes. Each code contains the equivalent of both data and checksum redundancy. The codes are then dispersed across a large pool of storage devices, which can be independent disks, independent network-attached storage nodes, or any other storage medium. While each of the codes is unique, a random subset of the codes can be used to retrieve the data. (Figure 1)

Figure 1. Erasure coding breaks data into data objects and distributes them across multiple independent nodes, disks, or sites.



1. Client system writes data to object storage.
2. Object storage controller breaks data into data objects and distributes them over multiple independent nodes, disks, and sites according to policy.
3. If one of the data objects is lost because of a disk failure, the object storage will rebuild that data object on another storage node.

## A Wide Range of Protection Policies

Data protection with erasure codes is expressed as a durability policy, a ratio of two numbers: the minimum number of codes over which the data is dispersed, and the maximum number of codes that can be lost without losing data integrity. For example, with a durability policy of 20/4, the object storage will encode each object into 20 unique codes and distribute those codes over 20 storage nodes, often from a pool of a hundred or more nodes. Since the object storage only requires 16 codes to decode the original object, data is still accessible after loss of 4 of the 20 nodes.

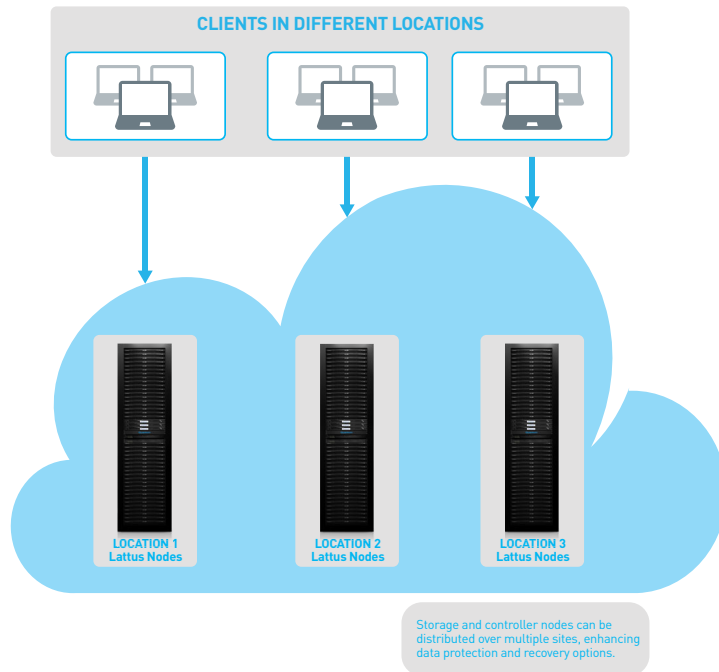
## Easier Migration to New Storage Technologies

The erasure code algorithms that spread the data to protect it also make upgrading to new storage technologies simpler. When drives fail, the object storage redistributes the protection codes without having to replace disks and without degrading user performance. Administrators are now free to schedule drive replacement at their convenience, and are free to upgrade to larger disk sizes.

## Protects Against Component Failure and Site Disasters

Unlike RAID, codes can be spread across racks or geographically dispersed sites without replication, allowing erasure coding to protect from disk, node, rack, or even site failures all on the same scalable system. The algorithms can also apply different durability policies within the same object storage so that critical data can be given greater data protection without segregating it at the hardware level. (Figure 2)

Figure 2: Object storage can spread data across multiple sites for content sharing and recovery options in case of site failure.



## Enhanced Object Storage with Both Object and File Access

More sophisticated object storage solutions also enable sharing of the storage between both the file system-based clients and applications that are engineered specifically to use object storage. This not only allows the storage pool to be shared across architectures, but it also provides seamless integration with workflow applications such as media asset management systems that use traditional file system access to share data with applications that are written specifically for HTTP-based object storage access. This guarantees the most flexible range of data use across the organization.

## LATTUS WITH STORNEXT FOR DIGITAL CONTENT LIBRARIES

Because it's built on object storage and erasure code technologies, Quantum Lattus is ideal for digital content libraries because it's scalable, secure, and cost-effective and it enables content to be accessible at disk access speeds from multiple locations. But object storage technology alone doesn't provide a complete digital library solution. By combining Lattus with a high-performance file access and management solution like StorNext® by Quantum, you get a complete, intelligent, ingest-to-archive workflow storage solution.

### StorNext 5 for Managing Workflow Storage

Newly engineered from the ground up, StorNext 5 combines file management technology with easy-to-deploy appliances to support complex and demanding digital media workflows. StorNext 5 supports content ingest, production, and distribution with the performance and reliability needed to meet tight production and delivery deadlines, plus archive integration that allows you to rapidly monetize content in digital libraries.

A complete, end-to-end StorNext 5 media workflow storage solution includes:

- StorNext 5 Metadata Appliances, which centralize content and enable simultaneous, seamless, and faster file sharing of very large numbers of very large files on high-speed SAN or LAN networks.
- StorNext 5 Q-Series Storage, which offers the unique performance capabilities required to manipulate large media files at very high rates.
- StorNext 5 Gateway Appliances, which give compute servers high-speed access to media files across LAN connections, especially for transcoding and proxy acceleration.
- StorNext Storage Manager, which automatically moves content and assets between high-speed disk, object storage, and/or tape storage, based on policies you set based on your workflow and budget requirements.
- Quantum LTFS Appliances, which read and write files in the open-standard, cross-platform LTFS tape format so you can access media files on LTO digital tape almost as if it were disk storage.

All StorNext 5 solutions are built on the foundation of providing high-performance collaboration to leading post and broadcast for over a decade.

### Lattus with StorNext 5 for Nearline Asset and Content Inventory

For online asset inventory with disk-speed access, Lattus can be integrated directly into a media asset manager or other digital library application through its HTTP-based cloud interface. Alternatively, it can be deployed with a file system layer that allows both application and direct user access through the familiar file and directory structure.

Regardless of the access method, Lattus provides lower latency than tape archives, for more predictable restore times for workflow applications and production teams. With quicker and more direct access, production teams can easily access digital assets to monetize content or make content available to distribution partners and subscribers.

## Lattus with StorNext 5 for Workflow Storage Automation

Policy-based storage automation migrates content between storage types based on the content's performance and access requirements, from high-performance disk storage to long-term storage like object storage. Simplifying the process, policies can be defined to automate migration based on last access date of the files, file type, size, or file locations. (Figure 3)

Figure 3: For broadcasters, object storage provides economical storage of large-scale, long-term, distributed content repositories.



Because it offers disk-speed access, Lattus can extend broadcast facilities' online storage to significantly reduce their highest-performance storage requirements. In some cases, media workflows have been set up to successfully migrate content after 10 days since last access, with the assurance that migrated content can be quickly accessed from the Lattus extended online storage.

## Lattus with StorNext 5 for Archive on Ingest

Lattus is ideal for preserving raw footage on ingest. With a policy-based storage manager, broadcasters can stream live video onto a traditional high-performance file system, and those files can automatically be copied onto an object storage system for long-term preservation, while keeping the content accessible on high-performance storage. Workflow applications such as media asset management systems can use existing capabilities—including timecode-based data retrieval. At the same time, new applications can be built to directly leverage the data from the HTTP-based object storage—like a website that may provide consumers direct access to archived broadcast video. This opens new possibilities to monetize content.



## CONCLUSION

At a time where the digital footprint of content is growing exponentially due to higher-resolution formats, additional distribution formats, and more cameras capturing more footage, the opportunities for monetizing that content have never been greater. The trick is keeping that content readily available and easily accessible for the production teams to do their magic and deliver those new revenue streams from that content.

A digital library built with Quantum Lattus can extend online storage by providing disk-speed access to all content without the capital and operating expenses associated with traditional disk-based storage, and provide direct file system access and automation through its integration with StorNext 5. With erasure coding algorithms for extreme scalability, data integrity and efficiency, plus file system support for open access and policy-based automation, Quantum Lattus and StorNext 5 can support the needs of petabyte-scale digital libraries today, and can support the workflow of tomorrow.



**BE CERTAIN**

#### **ABOUT QUANTUM**

Quantum is a leading expert in scale-out storage, archive and data protection. Its StorNext 5 solutions power modern workflows, enabling high-performance, real-time collaboration and keeping content readily accessible for future use and re-monetization. More than 100,000 customers have trusted Quantum to address their most demanding content workflow needs, including top studios, major broadcasters and new, cutting-edge content creators. With Quantum, customers can Be Certain™ they have the end-to-end storage platform to manage assets from ingest through finishing, and into delivery and long-term preservation. See how at [www.stornext.com](http://www.stornext.com).

[www.stornext.com](http://www.stornext.com) • 800-677-6268

©2014 Quantum Corporation. All rights reserved. Quantum, the Quantum logo, Lattus and StorNext are either registered trademarks or trademarks of Quantum Corporation and its affiliates in the United States and/or other countries. All other trademarks are the property of their respective owners.

**Quantum**®

WP00194A-v02 Oct 2014