Quantum

WHITE PAPER

QXS Q-TIER:

Optimizing Storage Performance and Costs for Animation Workflows

CONTENTS

ABSTRACT

Quantum QXS™ hybrid flash-disk storage arrays with Q-Tier™ real-time tiering provides the most efficient and cost-effective storage solution for animation facilities. Specifically, Q-Tier works within Quantum's StorNext® shared storage and data management platform to provide animators with the I/O performance and scalability they need for a fraction of the cost of an all-flash solution. This white paper describes the comprehensive testing of StorNext-powered QXS hybrid flash-disk storage with Q-Tier real-time tiering to validate Q-Tier benefits for animation workflow applications.

STORAGE CHALLENGES FOR ANIMATION WORKFLOWS

Next-generation animation applications offer creative professionals a host of exciting new possibilities. As the range of creative choices increases, and the cost of producing and consuming technologies declines, organizations across industries are using animation in new ways. In addition to producing compelling feature films and games, media companies are incorporating more exciting effects and experiences into everything from sporting events to corporate videos.

What does all this change mean for animation studios and post-production facilities? Growth, as clients increasingly demand engaging content that capitalizes on new technologies. As a result, many organizations need to overhaul their technology infrastructure. Finding the right storage solution is crucial to this task.

As animation artists create more shots, data sets are growing. Today, an animation-heavy film can easily consume more than a petabyte of data. Studios therefore need highly scalable storage solutions that can handle the tremendous capacity required for each project. Each animation and visual effects (VFX) project could require storage for hundreds of terabytes of large, high-resolution files.

Storage performance has also become critical. Teams need random I/O performance to support animation processes with large, high-resolution files, and later in the production cycle, exceptional sequential I/O performance for high-speed playout. Throughout production, animation teams require storage solutions that can provide shared access to content. Organizations need to facilitate collaboration across editing, computer-generated imagery (CGI), and rendering operations without delays that can slow productivity.

While producers and studios are often willing to invest heavily in animation and VFX, organizations must still keep costs under control. They need to make sure they are balancing performance and cost across storage tiers. In addition, they must keep administrative costs low: Production houses want to automate some tasks and streamline others, so they can stay focused on creative work.

XCELLIS WORKFLOW STORAGE

Xcellis boasts a unique architecture that integrates key media workflow storage components into an easy-to-manage, highly scalable hardware solution that supports both Fibre Channel and Ethernet clients, and hosts workflow applications while managing data across multiple storage tiers. Powered by the industry-leading StorNext platform with Q-Tier intelligent tiering, Xcellis allows users to optimize up to 64 separate file systems to meet the performance requirements of different departments. These characteristics provide a host of benefits to post-production facilities that share media across animation and editorial departments, including:

- Storage of all media files on a single system with a single namespace, eliminating the cost of managing and maintaining multiple storage silos
- Ethernet-based NAS and DLC access, providing the speed and availability animators need while direct Fibre Channel connections supply the editorial department with highly predictable performance for demanding high-resolution content
- Render farms that can connect to storage for direct access to content during rendering, eliminating the need for bulk data dumps and network shutdowns
- Fast, easy, and cost-effective system scaling through the simple addition of more storage, making it easy to accommodate growth and new technologies

Quantum VFX Workflow Reference Architecture

Quantum offers a storage solution designed expressly for the requirements of animation and visual effects production. The Quantum VFX Workflow Reference Architecture leverages StorNext software to power Xcellis® high-performance workflow storage and QXS hybrid flash-disk storage arrays with Q-Tier real-time tiering for highly efficient VFX and editing workflows. StorNext not only eliminates the need for separate storage silos, but also eases collaboration across these departments, and optimizes storage and media access according to the needs of each. With the resulting gains in capacity and productivity, animation and VFX facilities can take on larger and more demanding projects, and more projects overall.

The time- and resource-consuming task of transferring content between animation and editorial departments over a network is a key obstacle for post-production facilities seeking to maximize workflow efficiency. In some cases, entire facilities must essentially shut down for hours each day as they copy material over the network and dedicate storage systems on both sides of this divide to writing and reading data. Quantum's VFX Workflow Reference Architecture enables users to overcome these challenges, maintaining storage for both animation/VFX and editorial within a single shared environment that supplies the ideal level of connectivity across the end-to-end workflow.

QXS hybrid flash-disk storage with Q-Tier technology provides artists with the random I/O performance they need for a fraction of the cost of an all-flash solution. To optimize system performance, QXS Q-Tier technology scores, scans, and sorts data every five seconds, and moves it between flash and spinning disk tiers as access patterns change and evolve.

Because StorNext storage solutions are already certified with leading VFX applications, including the Autodesk and Adobe suites, and animation management tools such as Shotgun, facilities can seamlessly integrate the new reference architecture into existing environments. As a result, users on the animation and editorial sides of the business can continue working with their preferred software tools while realizing a notable increase in performance and much improved ability to work collaboratively.

TEST PLAN AND CONFIGURATION

The following section describes the specific configuration used for testing the Quantum VFX Workflow Reference Architecture to validate Q-Tier benefits for animation workflow applications.

Test Plan Overview

Unlike other hybrid storage arrays that tier data daily to avoid performance degradation during peak hours, QXS arrays have dedicated processing and dual active-active controllers for utilizing Q-Tier—constantly promoting active workloads to the fastest tier available with no reduction in performance. This provides the most needed data with the highest possible I/O per second (IOPS), and also reduces bottlenecks.

Q-Tier is automated, real-time storage tiering that ensures that the most frequently used data is in the highest performing storage when it is needed. The technology's intelligent algorithm continuously scores, scans, and sorts the data. This process occurs every five seconds, continuously moving data between the tiers as the data being accessed changes. As a result, users can achieve flash performance for data sets that need it without over-provisioning expensive flash storage infrastructure.

Because Q-Tier intelligent tiering can be of critical importance for enabling producers and studios to balance performance and costs while keeping costs low, the underlying focus of the testing was on corroborating Q-Tier benefits using rigorous, real-life animation workloads. Specifically, tests were conducted to answer the following questions pertaining to specific assumptions regarding the benefits that Q-Tier can provide for an animation workflow:

- Does Q-Tier meet its underlying expectation (i.e., does frequently, randomly accessed data get promoted to the performance tier)?
- Do rendered frames first get written to the storage performance tier and then (depending upon performance tier utilization metrics) to the archive tier?
- Does data get promoted or demoted between tiers while the QXS hybrid array is serving I/O requests and, if so, how rapidly?
- Will an animation video sequence written to a QXS array while it is also serving read requests (e.g., from editors) get written to the performance pool? What is the performance impact of the write operation on random read requests?

Render Farm Environment

The render farm configuration consisted of three Dell server platforms. Each server supported VMware ESXi configured to take advantage of the available number of CPU cores and memory expansion capability.

- Dell R510 system hosting four main virtual machines (VMs) used for management purposes:
  - The VMware vCenter Server running on Windows 2012 R2.
  - A general-purpose Windows 2012 R2 VM that bridged public and private VM networks.
  - The license server VM that serves Autodesk Maya and Arnold licenses. (Maya and Arnold are 3D rendering applications offered by Autodesk.)
  - The Deadline render manager VM. (Deadline is an administration and compute management toolkit that offers a wide-range of management options for render farms.)
- Dell R730 system supporting 40 virtual render nodes, each using a single core, 8 GB of RAM, and a single 1-GbE interface.
- Dell R720 supporting 20 virtual render nodes, each using a single core, 8 GB of RAM, and a single 1-GbE interface.

Rendering operations were orchestrated by the Deadline queue manager. Maya projects were loaded and frames were all rendered out of order using Arnold, from and to the QXS Q-Tier volume. (The QXS storage configuration is described in the following section).

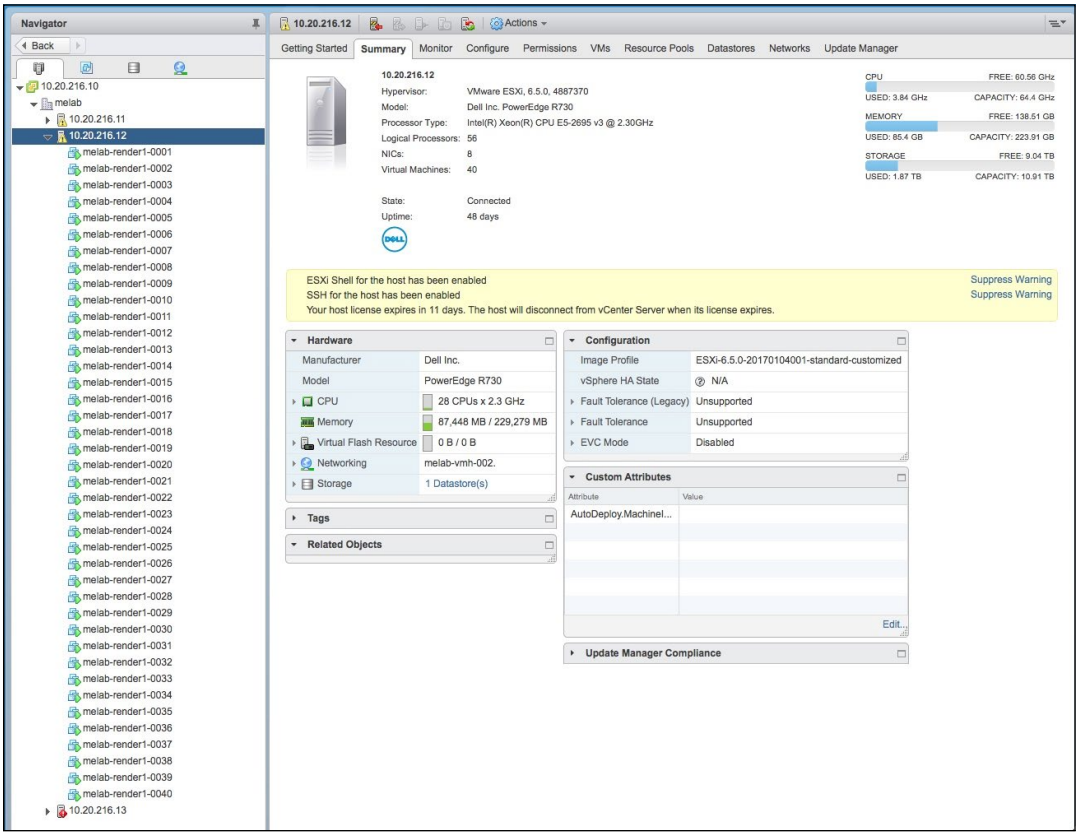

Figure 1: Forty virtual render nodes defined

Virtual storage utilized a repurposed storage unit to accommodate each server's data store (separate from the QXS-456 array used for rendering media storage). Every VM instance resided on these data stores, which hosted their virtual system drive, used for the node's operating system, render software, and various temporary files.

Each VM used a single virtual 1-GbE link to connect to the Q-Tier volume via the network file system (NFS). The 1-GbE links were all created out of the Dell servers' physical 10-GbE interfaces using virtual switches. This allowed each one of the virtual machines to have access to a single, dedicated 1-GbE interface. An Xcellis Workflow Director node shared the volume using the StorNext NAS stack over a 10-GbE link, over the NFS protocol.

A Dell PowerConnect 8024F 10-GbE switch was used to connect all physical 10-GbE links for the render farm servers and the Xcellis nodes.

Quantum QXS Configuration

The Quantum QXS hybrid storage portfolio optimizes flash and disk to significantly lower operating and capital costs. Powerful features include real-time, automated and intelligent tiering, plus one-button configuration to meet the demands of specific applications and workflows.

The capacity-optimized QXS-456 hybrid-storage configuration provided the storage for the rendering test. QXS-456 is a 4U system that can support up to 56 4-TB high-capacity, 3.5-inch large form factor (LFF) hard disk drives (HDDs) as well as 12-inch 400-GB solid-state drives (SSDs). As a result, QXS-456 provides a much higher IOPS performance compared to a storage configuration utilizing only HDDs.

Specifically, the QXS-456 storage configuration used for the test consisted of:

- "Performance" tier using 12 400-GB SSDs, or a total of 4.8 TB flash storage
- "Archive" tier using a total of 40 4-TB LFF HDDs, or a total of 160 TB hard disk storage
- Four 4-TB LFF HDDs as hot spares

Figure 2: QXS disk configuration

The QXS-456 SSD storage is leveraged by Q-Tier tiering to ensure the highest IOPS performance for frequently accessed data through the use of algorithms that run on QXS controllers. This allows StorNext to provide the highest performance for client requests, with tiering occurring in the background.

The combination of optimum capacity, performance, and cost make QXS-456 with Q-Tier an ideal animation workflow storage solution, where typically multiple versions of the same project exist on the storage system and thus require much more capacity than the final product, and a high level of simultaneous, random access requests occur due to the nature of modern animation projects.

For example, hundreds of workstations may be working on the same project simultaneously while rendering occurs. This random data access pattern is very different from the sequential I/O-heavy data access typically seen in 4K video post-production workflows.

Monitoring Q-Tier Statistics

QXS-456 offers multiple ways to monitor Q-Tier statistics:

1. The QXS v3 web browser–based GUI enables viewing Q-Tier statistics in real time when the cursor hovers over the bottom of the interface, more precisely over the bar graph area.

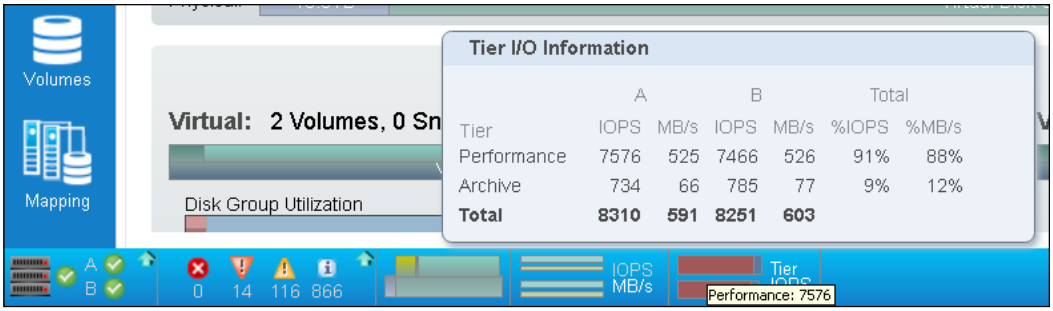

Figure 3: Real-time Q-Tier statistics

2. Regular QXS performance statistics in the GUI show IOPS and megabytes-per-second performance for virtual tiers, along with the number of pages being allocated/deallocated by the minute.

3. Using the QXS CLI, specific commands can show the same information. For example, the following command can be used to determine the capacity that is allocated and available for a specific tier: # show tiers tier all

Pool	Tier	% of Pool	Disks	Total Size	Alloc Size	Available Size	Affinity Size
A	Performance	3	6	1996.2 GB	1996.2 GB	0 B	0 B
A	Archive	96	20	63.9 TB	886.7 GB	63.1 TB	0 B
B	Performance	3	6	1996.2 GB	1996.2 GB	0 B	0 B
B	Archive	96	20	63.9 TB	886.9 GB	63.1 TB	0 B

The following command will show performance information for all tiers, similar to what is provided using the QXS web browser-based GUI: # show tier-statistics tier all

Q-TIER TEST RESULTS AND ANALYSIS

This section describes the comprehensive testing that was performed using the test configuration described in the Test Plan and Configuration section to validate the benefits of Q-Tier for animation workflows.

Test 1: Tiering of Frequently Accessed Animation Asset Data

An important aspect of the Q-Tier value proposition for animation workflows is its ability to automatically (and in real time) migrate data between the SSD-based performance tier and the HDD-based archive tier based on the frequency of random data access. The first set of tests had the goal of validating this Q-Tier capability using animation asset data for the test case.

Specifically, test 1 consisted of the following steps:

1. Copying 14 GB of animation assets to an asset directory on QXS-456. The initial data transfer was to the performance tier of the QXS array.

2. Following a one-week time interval, only 15% (i.e., 2 GB) of the assets data resided on the performance tier of QXS-456 while 85% of the data had migrated to HDD-based archive tier.

3. A rendering process using the 60-node rendering configuration was then started while Q-Tier statistics were monitored using the StorNext GUI. As the assets' data was being read by the render nodes (or transferred to QXS-456 storage to render nodes), the IOPS shown by the performance tier went up marginally.

4. By the time all render nodes had read the assets once, approximately 30% (or 4.2 GB) of the assets had been promoted to the performance tier.

5. The rendering process was stopped before any frames could be generated and restarted to determine if additional data migration could take place during later tests.

6. A script that automated the transfer of the asset data 10 times in succession was run simultaneously on the 60-render-node configuration. It effectively caused 8.4 TB of random access requests to be executed on the QXS-456 system in two hours.

At the end of the fifth restart of the script referred to in step 6 above, approximately 80% of the asset data was found to have migrated to the performance tier over time. After the fifth restart, the StorNext GUI displayed Q-Tier statistics as follows:

Tier I/O Information

	A		B		Total	
Tier	IOPS	MB/s	IOPS	MB/s	%IOPS	%MB/s
Performance	7718	509	7984	522	100%	100%
Archive	0	0	0	0	0%	0%
Total	**7718**	**509**	**7984**	**522**		

At this point, the performance tier was effectively serving 100% of the read requests. This clearly validated the Q-Tier operation in migrating data subject to a high frequency of random read requests from the slower archive tier to an SSD-based performance tier, since 85% of the asset data resided on the archive tier at the start of the test.

While this synthetic test forced the data transfer process to occur rapidly, we can conclude that in a real-life scenario, where every render node may not read some or all of the assets' data, the tiering will not occur as deterministically and some less-accessed files may not be transferred to the SSD tier.

Test 2: Tiering for Data Transfer to Storage

The second test had the goal of validating the Q-Tier data tiering operation for transfer of render-frame data to the QXS hybrid array (i.e., to verify that initially frame data would be transferred to the performance tier and, depending on space constraints, would then be migrated to the archive tier). Test 2 consisted of a simulated render of 12,000 4K frames (or an equivalent of an eight-minute, 20-second 24-fps clip) comprised of the following steps:

1. Prior to the simulated render test, the QXS performance tier was filled up to 50% capacity.

2. A test script that simulated rendering operations was run on each of the 60 render nodes simultaneously. It read all assets, and after a five-minute interval (to simulate CPU processing), wrote a single 4K full-aperture frame having a 51-MB operation, and then repeated this operation 200 times.

3. Each of the 10-GbE links supported by the render nodes were found to be utilized to capacity (i.e., to a performance level of approximately 1.1 GB/sec).

After having the script run for one hour, Q-Tier statistics showed that the frames were being written primarily to the performance tier, approximately 70% of which was now utilized.

show tiers tier all

Pool	Tier	% of Pool	Disks	Total Size	Alloc Size	Available Size	Affinity Size
A	Performance	3	6	1996.2 GB	1413.3 GB	583.0 GB	0 B
A	Archive	96	20	63.9 TB	662.4 GB	63.2 TB	0 B
B	Performance	3	6	1996.2 GB	1413.0 GB	583.3 GB	0 B
B	Archive	96	20	63.9 TB	662.8 GB	63.2 TB	0 B

The render process was allowed to continue until the performance tier was fully utilized and additional frames were being written to the archive tier.

Pool	Tier	% of Pool	Disks	Total Size	Alloc Size	Available Size	Affinity Size
A	Performance	3	6	1996.2 GB	1996.2 GB	0 B	0 B
A	Archive	96	20	63.9 TB	886.7 GB	63.1 TB	0 B
B	Performance	3	6	1996.2 GB	1996.2 GB	0 B	0 B
B	Archive	96	20	63.9 TB	886.9 GB	63.1 TB	0 B

Test 3: Data Promotion/Demotion Across Tiers Concurrent With I/O

The third set of tests was designed to validate Q-Tier functionality for promotion and demotion of data between the performance and archive tiers concurrent with the QXS array supporting I/O operations and to measure the rate at which such data movement took place.

Test 3 consisted of the following steps:

1. When the performance tier was fully utilized at the end of the render process conducted in test 2, and frames were being written directly to the archive tier, testing was intentionally paused to determine if Q-Tier would start demoting data from the performance tier to the archive tier.

2. Since no data was being written to the array, frames started to be demoted from the performance tier to the archive tier.

3. Four discrete measurements were taken regarding the rate at which data was being migrated from the performance tier to the archive tier during the 1.5-hour period following stoppage of the render process. The results of the four readings, which averaged 815 MB/min, are shown below.

Time of Day	Time Delta	Free Space on Performance Tier A (MB)	MB Delta	Migration Rate (MB/min)	Free Space on Performance Tier B (MB)	MB Delta	Migration Rate (MB/min)
6:31:45	N/A	8,736.7	N/A	N/A	8,728.3	N/A	0
6:42:58	0:11:13	16,900	8,163.3	728	16,900	8,171.7	729
7:34:10	0:51:12	58,400	41,500	811	58,500	41,600	813
7:58:06	0:23:56	79,100	20,700	865	79,200	20,700	865

4. The render test in test 2 was restarted until it completed about two days later. At the end of the test, the complete sequence of 12,000 frames was found to have been fully transferred to the renders directory, with the performance tier fully utilized.

Test 4: Tiering of Frequently Accessed Animation Frame Data

Test 4 used animation frame data to validate Q-Tier functionality for migrating data between the SSD-based performance tier and the HDD-based archive tier based on the frequency of random data access. Test 4 consisted of the following steps:

1. An interval of five days following test 3 was allowed, during which no I/O activity happened on the QXS-456 array.

2. Render frame data on the array was read concurrently by the 60 render nodes using random data access (the same script that was used in test 1 to read asset data from the array to the render nodes was modified to read frame data).

3. At the start of step 3, the ratio of IOPS served by the performance tier versus the archive tier was 84% to 16%. This validated the aspect of Q-Tier functionality that the extent of demotion of data from the performance tier to archive data depended on the extent to which the QXS array was serving I/O operations.

Tier I/O Information

	A		B		Total	
Tier	IOPS	MB/s	IOPS	MB/s	%IOPS	%MB/s
Performance	5212	391	5153	385	84%	79%
Archive	967	103	1067	108	16%	21%
Total	**6179**	**494**	**6220**	**493**		

4. Following a one-hour interval, the I/O statistics were measured again and the ratio of IOPS served by the performance tier versus the archive tier was found to be 91% to 9%. This signifies that promotion of data is primarily driven by I/O requests supported by the QXS array (and not by the array being idle), which validates the active, intelligent nature of the Q-Tier data tiering algorithm.

Tier I/O Information

	A		B		Total	
Tier	IOPS	MB/s	IOPS	MB/s	%IOPS	%MB/s
Performance	7576	525	7466	526	91%	88%
Archive	734	66	785	77	9%	12%
Total	**8310**	**591**	**8251**	**603**		

Test 5: Writing of Video Sequence Concurrent With Random Access of Asset Data

The focus of the final test was to validate that Q-Tier tiering functionality allows an animation video sequence to be written to the performance tier of the QXS array while concurrently serving random read requests and to determine the performance impact of the write operation on the read requests.

The test consisted of the following steps:

1. An asset read from all render nodes from all 60 render nodes generating random I/O.

2. Concurrently with step 1, a Fibre Channel connected initiator host started writing a sequential stream of 10,000 uncompressed 4K frames at a 1.6 GB/sec data rate to the QXS array.

3. When the initiator host started the write process, the frames were found to be directly written to the performance tier until it was fully utilized, as shown by the tier statistics displayed by the QXS GUI and CLI.

4. Concurrently, the render nodes were seen to be reading the assets at the nearline rate via the QXS 10-GbE link.

Test results show that the SSD-based performance tier is helpful in enabling overall high random read IOPS performance by render nodes while also supporting sequential write of a large 4K video file at line rate by a single node. The performance tier contributes a high proportion of overall system IOPS and, under high-stress conditions, a relatively small proportion of SSD capacity for supporting "hot" data, which can provide a critical benefit.

Capacity Ratios

The SSD space versus HDD ratio was relatively low in the test configuration, 3%. Quantum recommends that the SSD space versus HDD ratio for animation environments should be between 5% and 10%, achieved ideally using 800-GB SSD drives. Determining the optimum SSD versus HDD capacity ratio (e.g., by determining the size of a typical customer project) is critical for overall array performance, as the tests previously described show how the performance tier can easily get filled to capacity.

TEST CONCLUSIONS

This white paper described comprehensive testing based on the StorNext-powered QXS-456 hybrid flash-disk storage with Q-Tier real-time tiering to validate Q-Tier benefits for animation workflow applications. This testing provided the following key conclusions:

- Q-Tier automatically (and in real time) migrated animation asset data between the QXS-456 SSD-based performance tier and HDD-based archive tier based on frequency of random data access.
- For transfer of render frame data to the QXS array, Q-Tier initially transferred frame data to the performance tier and, depending upon space constraints (as the SSD tier reaches its 80% capacity limit), demoted it to the archive tier.
- Q-Tier functionality for automatic demotion of data from the performance tier to the archive tier if the array was not supporting any I/O operations was validated.
- Q-Tier capability for migrating animation frame data between the SSD-based performance tier and the HDD-based archive tier based on frequency of random data access was confirmed.
- Q-Tier enabled an animation video sequence to be written to the performance tier of the QXS array while concurrently serving random read requests without any performance impact of the write operation on the read requests.

These test conclusions confirm that the Quantum VFX Workflow Reference Architecture with Xcellis high-performance hybrid flash-disk storage with Q-Tier real-time tiering provides a robust and exceptionally flexible architecture that can deliver optimal storage capabilities for animation and editorial operations—streamlining workflows and boosting overall efficiency and productivity.

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

WP00237A-v01 Jan 2018