REFERENCE ARCHITECTURE

Building Collaborative Production Environments with Quantum CatDV and StorNext

Abstract

This document defines a reference architecture for an integrated Quantum StorNext® and CatDV workflow. The storage infrastructure is implemented with Quantum software and appliances.

Quantum

Table of Contents

Quantum

Executive Summary

The objective of this paper is to present a detailed, repeatable reference architecture for deploying asset management and archiving, with collaborative, shared storage solutions using the Quantum CatDV media management and orchestration platform integrated into a Quantum StorNext environment including Storage Manager. This combined solution leverages a key point of integration between the file system metadata services and Quantum CatDV to deliver dramatically efficient collaborative operations, especially regarding detection of on file moves, copy and re-name operations. The architecture demonstrated can scale into the Petabyte range and beyond, protects assets with a very high degree of resiliency, offers retrievability of assets with excellent performance, all while making the entire library available to and tightly integrated with the Quantum CatDV asset management platform.

By implementing the Media Asset Manager (MAM) as the creative user's primary interface into their workflow and archive, customers will be able to index and ingest master files, automatically generate links back to all the raw sources of those masters, which are then automatically retrieved and presented to the user or solution when needed. The customer can organize their assets by organization, library, catalog, project in any combination, monitor asset locations and asset levels of protection, and have a database-driven, searchable interface to quickly find content across their entire deployment. By configuring powerful workflow automation engines users can order and speed their workflow steps including managing ingest of raw or source files, production management and planning, craft editing and finishing, review and approval, asset enhancement and monetization, and archive workflow steps.

At a very high level, Figure 1 below provides a view of the components and options that go into a Quantum Solution:

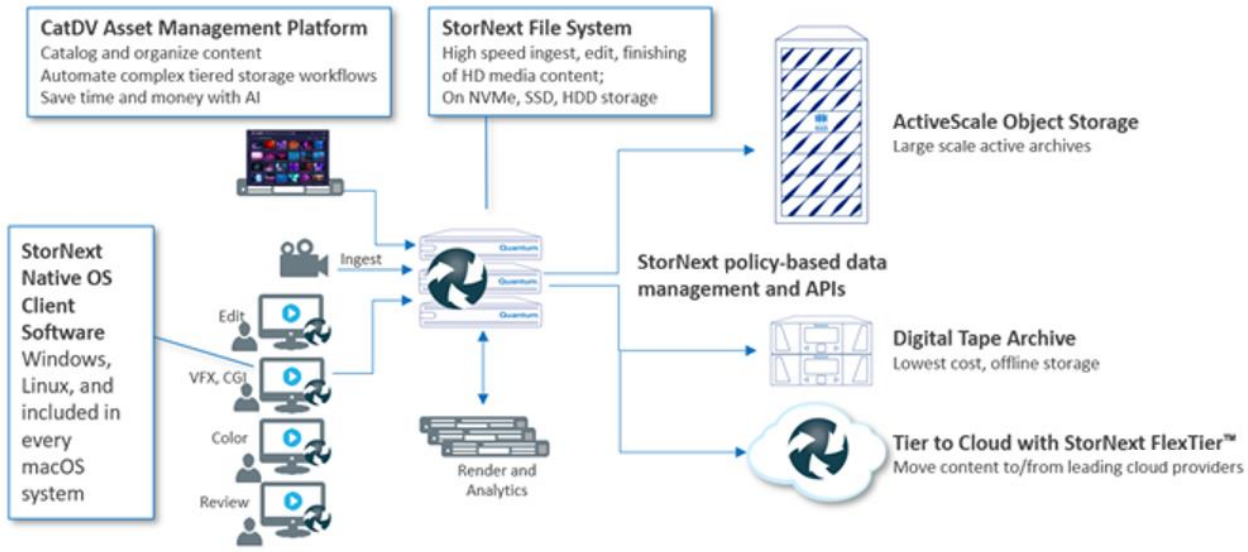

Figure 1 – Quantum Solutions for Media Workflows

The Need for This Solution

With skyrocketing demand for content in all its forms as the basis for highly effective communications for companies of all kinds, the rapidly increasing media resolutions and delivery formats, and the need to consolidate ever larger content libraries to speed content production and support extended monetization efforts, there is an urgent need for content library storage solutions that can seamlessly support asset management workflows. The challenge facing workflow architects today is to design a solution that can track, store, and retrieve these assets even while the asset library and creation workflows are growing dynamically and beyond traditional storage architectures. The number of collected assets, multiple ongoing projects, serving users in geo-dispersed sites and remote operators all make it extraordinarily challenging for organizations struggling to maintain order in media management and prevent the loss of irreplaceable content. Over the last year, facilities have been required to deliver highly reliable workflows despite a near-total remote work force. With the surge of new content and multiple content staging areas such as tape, nearline, and other disk storage, tracking all this data and managing their multiple locations manually is both error-prone and inefficient.

By implementing CatDV as the Media Asset Management (MAM) solution that presents a unified view of asset inventory across local and object storage that automates file movement and proxy generation, a single canonical view of content inventory is achieved while also allowing creatives and asset managers to be more efficient and work with content as an intrinsic part of their workflow rather than a separate effort to manage content, or storage, directly. The platform delivers a wide variety of media-centric capabilities, including traditional MAM, PAM and DAM, sophisticated workflow automation, and fully customized applications. CatDV deployments create transparency and visibility, ensure best practices are followed, and increases team productivity. Using artificial intelligence and machine learning CatDV makes it easier for businesses to catalog and analyze digital assets. It allows complete searches across StorNext® File systems and private or public cloud pools. Quantum's StorNext® collaborative shared file system, with a complete range of workflow-fit storage, CatDV asset management, and orchestration solution is ready to match your creative vision, deliver performance and economics exactly where you need it, help you manage your entire workflow with confidence, and deliver your best work on deadline. Employing StorNext APIs and multiple partner plugins, a wide range of platforms are available for archiving, from Nearline storage, Tape, Object Store in Private cloud to public cloud, all the while maintaining accounts and records of these media assets. CatDV unlocks value in our client's media, saving time, money, and stress, whether deployed on-premise or in the cloud, across traditional and object storage tiers.

When implemented in this way, content creation and archival workflows become a virtuous circle of media management and production. In the diagram below, we observe the workflow steps of asset and content creation through archive becomes a cyclical pattern, completing the circle with the ability to retrieve archived assets that can, in turn, be used to create more assets or select for the next project combined with newly created media.

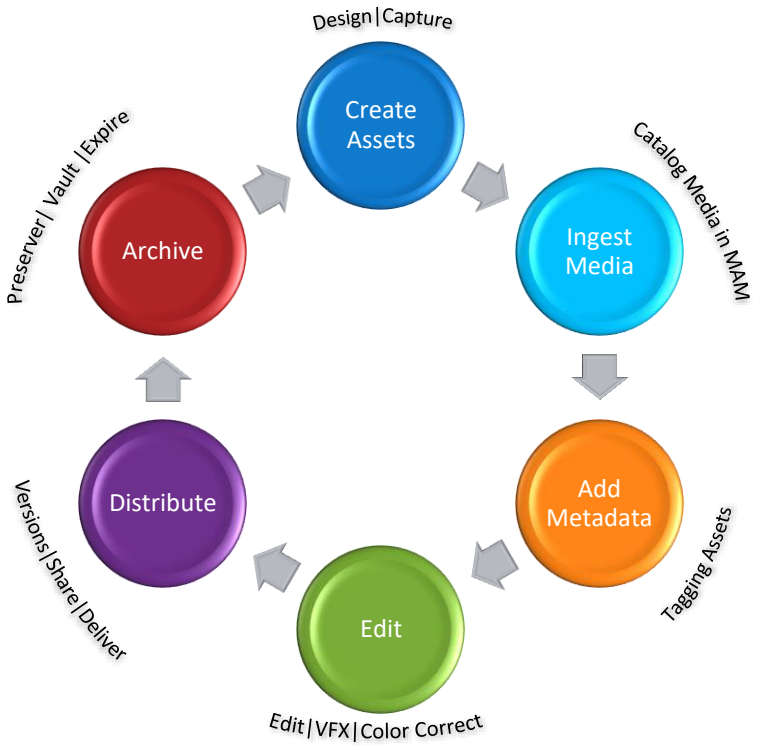

Figure 2 – Media Asset Life Cycle

Customer Scenarios

Following are a few examples of workflows that highlight the applications of the components of this Reference Architecture.

Example 1: Post-production Workflow

In a typical post-production workflow for film, episodic, or documentary work, a collaborative team is focused on a single project that can be highly complex, requiring the management of a large number of high-resolution, native shots and assets that need to be quickly graded, marked, and annotated to assemble an editorial cut, then support ongoing fine editorial as the remaining elements of the production are added including sound, color grading and finish, and special effects. Throughout this process, numerous editorial decisions need to be made requiring close collaboration and communication across the different departments, and even require keeping a record of notation and shot acceptance. Once the production is ready for delivery, the raw, unused assets, and all components of the finished production are archived together for ready retrieval to product new assets such as extended cuts, trailers, or special features. In this scenario, CatDV's strengths in collaboration, tagging, marking, and annotating help the team focus on the best shots for a production quickly, protect the assets throughout even extended productions, and archive the finished and component assets securely and ready for fast retrieval.

Example 2: Broadcast Production

In a typical broadcast production workflow, its common to have a large number of prepared packages ready for broadcast during an event, while also needing to be flexible enough to customize packages on the fly to adapt to live or current events. In this scenario, a team will have a large number of shorter packages prepared but may quickly change out component shots before delivering for broadcast. Similarly, if the library of assets is well managed and tagged, and even enriched with key metadata information such as transcription of audio with corresponding timecodes, and object and person identification, then assets can be rapidly found in even the largest libraries to create new packages and respond to changing events highly dynamic broadcast situations. These customers make full use of both StorNext for high tempo editing of HD and above editing, and deep use of CatDV's highly customizable asset tagging, commenting frameworks for customizable monetization and enrichment workflows.

Example 3: Corporate Video

Corporate Video users cover a wide range of scenarios, from highly disciplined marketing, audience engagements, and content branding operations, to highly technical training and orientation operations, to other regular rich content production as a growing part of business operations. Even more than Broadcast Production customers, these customers often need to serve a global internal customer base and have a nuanced content approval and tagging strategy in place. These customers can freely manage not only video content, but also high-resolution image content as well often needed for technical work. These customers need to ensure that content is versioned correctly for tight brand control and to ensure that the correct and latest version of the content, logos, and trademarks are used in a highly coordinated fashion, but also to ensure that this content is adapted to each region both for translation and localization, but also to adapt to each region's cultural needs. These customers find that they make full use of CatDV's flexible libraries and catalogs, and especially in the review and approval and chat and messaging frameworks.

Example 4: Content Curation

Content librarians and curators may have a specific service mission such as for local, regional, or departmental government needs, or for educational organizations of every size that need to organize content for research, reference, or preservation. These organizations can easily extend web-based access to a worldwide network of subject matter experts or researchers to enrich and document content while protecting the content itself in a centrally managed workflow. These customers will likely make the heaviest and most detailed use of metadata versioning and specific metadata sets and tagging that content librarians and managers can readily access and extend as needed.

Solution Overview

With the historic demand for digital assets and rich media content production in higher resolutions in all forms of communication, every industry is realizing the need for resilient, highly scalable, and reliable solutions to properly track, store, and retrieve these assets, even while the asset library and creation workflows are growing dynamically. The number of collected assets, multiple ongoing projects, global diversity of sites and locations and more remote operators and producers all contribute to the challenges of organizations struggling to maintain order in media management. Facilities have been tasked with providing highly reliable workflows despite a near-total remote workforce. As companies have seen, focused content and data-centric lifecycle with shared accessibility across worksites and remote workers has proven highly effective and offers a model for new levels of reliability and capability. Businesses that have adopted this model can transform their operations and capabilities by rising to evolving challenges of data and content collection, production, archiving, and retrieving of assets in a timely manner.

With all this content and the proliferation of multiple staging areas such as tape, nearline and other storage volumes, tracking all this data and its location with a spreadsheet or notes catalog is not efficient, or safe. Implementing a high-performance shared storage system in conjunction with an asset management solution gives organizations a more reliable and faster way to store these assets short or long term. By implementing CatDV as the organization's Media Asset Management (MAM) system presents a single, canonical view of all content in a single place. With StorNext and Storage Manager the target points for short-term needs for work in progress, or long-term archives are varied. But ultimately this solution supports the users' desired mix of performance, connectivity, and cost at each operation stage of working with content and provides multiple archive endpoints that are capable of Petabytes of storage capabilities.

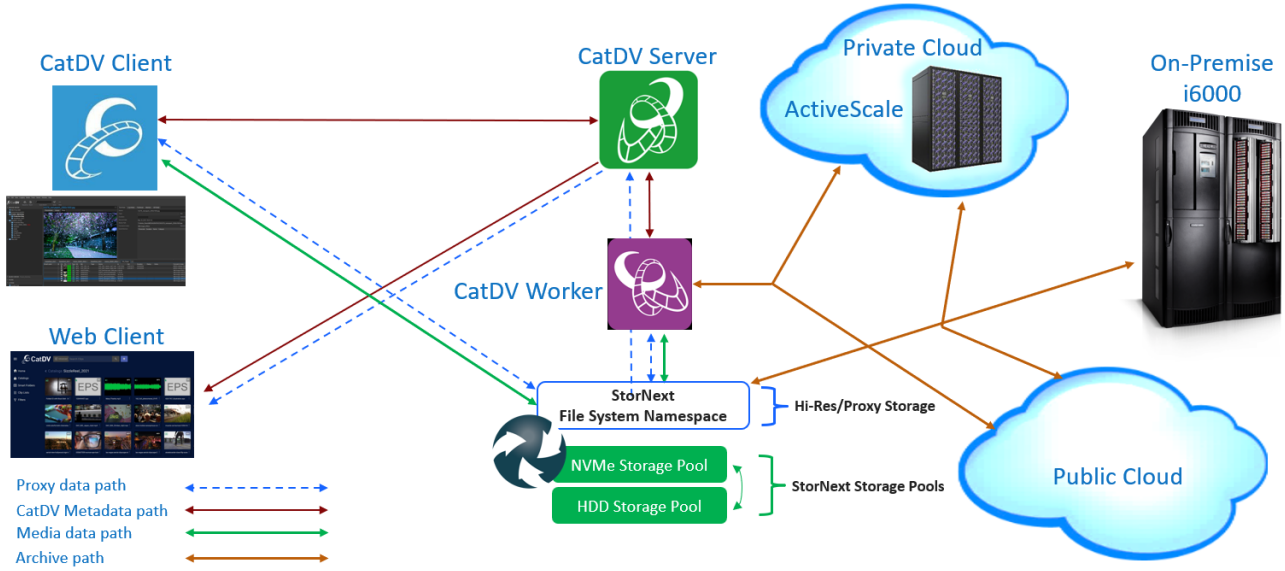

Figure 3 – High-Level CatDV and StorNext Architecture

Technology Summary

The table below lists the technology components that make up the Quantum CatDV and StorNext reference architecture outlined in this document. The paragraphs that follow the table provide more detail on the function of these components in the solution.

Solution Components	Item	Notes	
StorNext	StorNext Credentials	Typical access to StorNext Web UI	
CatDV	CatDV Credentials	Ability to login to CatDV, user or Administrator	
StorNext Protocol	**Version (s)**	**Inbound/Outbound**	**Notes**
NAS,DLC,Direct Attached	Ethernet and Fiber	Inbound/Outbound	SNB/NFS/DLC, Fiber Channel, iSER/RDMA
CatDV Protocol	Version (s)	Inbound/Outbound	Notes
Amazon S3 Signature	2	Inbound/Outbound	Provide S3 connectivity
REST API			HTTP-based API
HTML 5/CSS3/JavaScript			
Software	**Version (s)**		**Notes**
StorNext	7.0.1		CentOS Linux release 7.7.1908
CatDV Server	10.0.1		Windows 10 Server
CatDV Worker	9.0.1		Windows 10 Worker
CatDV Pegasus Client	14.01		macOS Big Sur Pegasus Client

TABLE-1 – TECHNOLOGY SUMMARY

Quantum CatDV Media Asset Management Platform

The figure-4 below lists the server versions available and the specific features with each.

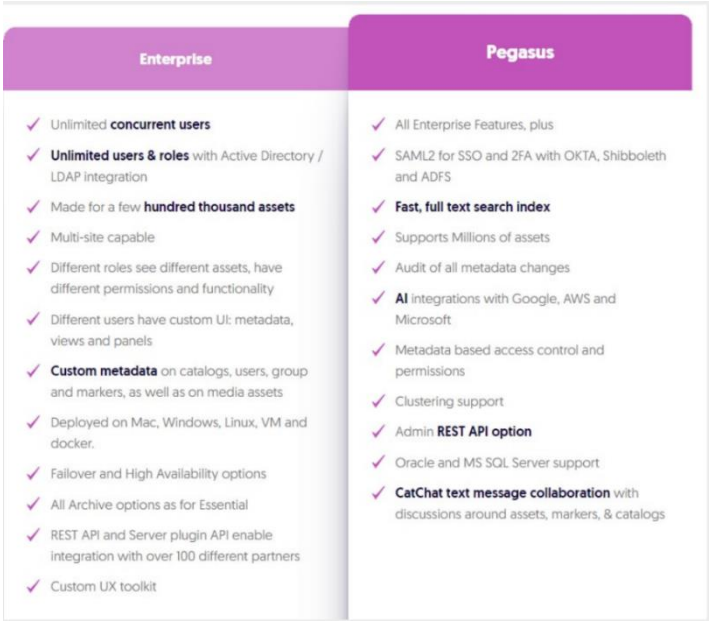

Figure 4 – Quantum CatDV Server Versions

Quantum CatDV is an agile asset management and workflow orchestration platform that provides powerful asset management, automation, and collaboration tools for organizations managing large volumes of digital media. The platform delivers a wide variety of media-centric capabilities, including traditional PAM, MAM, and DAM, sophisticated workflow automation, and fully customized applications. CatDV deployments create transparency and visibility, ensure that best practices are followed, and increase team productivity. CatDV unlocks value buried in large media libraries, saving time, money, and stress. Whether deployed on-premises or in the cloud, across traditional and object storage tiers, customer benefits include:

- Communicating and collaborating more effectively
- Automating workflows across tiered storage products
- Unlocking value in large content archives
- Saving time and money by leveraging Artificial Intelligence

CatDV is used today in post-production, corporate video, sports, government, and education markets, and has potential to expand to other markets using specifically designed plug-ins for expanded use cases such as genomics research, autonomous vehicle design, geospatial exploration, and any use case dealing with large unstructured data. For a closer look at what CatDV with Quantum follow the link here to find the latest in what Asset Management from Quantum brings you.

Quantum StorNext File System and Storage Manager

Quantum StorNext is a shared storage software platform – a high-speed file system that can span different types of storage in a single namespace with integrated data lifecycle management – that is ideally suited for large video files, digital imagery, and other forms of unstructured data. StorNext is used by large broadcasters, studios, and post-production houses for their collaborative video workflows, it is used by government agencies around the world to study the planet and used by scientific research organizations to store, protect, and archive research data.

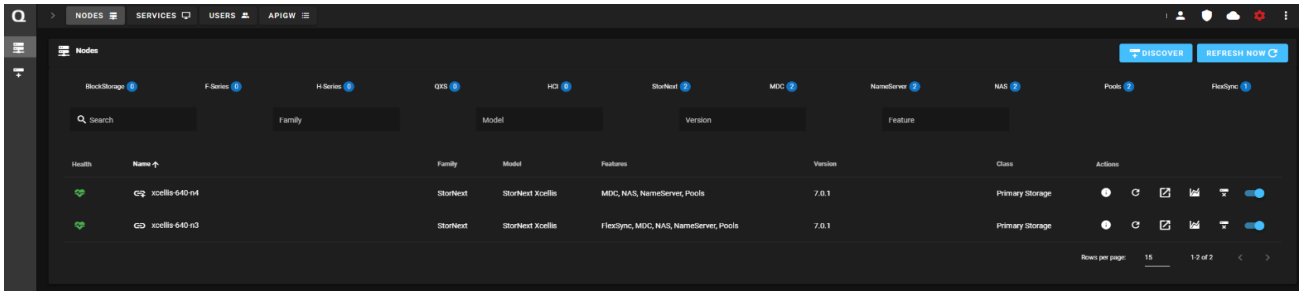

Figure 5 –Next-Generation Consolidated UI

Reference Architecture

Quantum CatDV Server and Worker nodes are server-side components typically reside in the datacenter, deployed on bare metal computers or VMs. Both applications can run on Mac, Windows, or Linux system, however, Windows and Linux are the predominant choices. It is recommended that the Server and Worker node are connected via 10Gig Ethernet to the file system. Both can also be connected via block level access over Fibre Channel or High-Performance Ethernet. Whichever deployment and connection selected; the worker node must have read/write access to the file system.

The following diagram represents the lab environment that was configured to show a typical implementation of the Reference Architecture and a combination of components used:

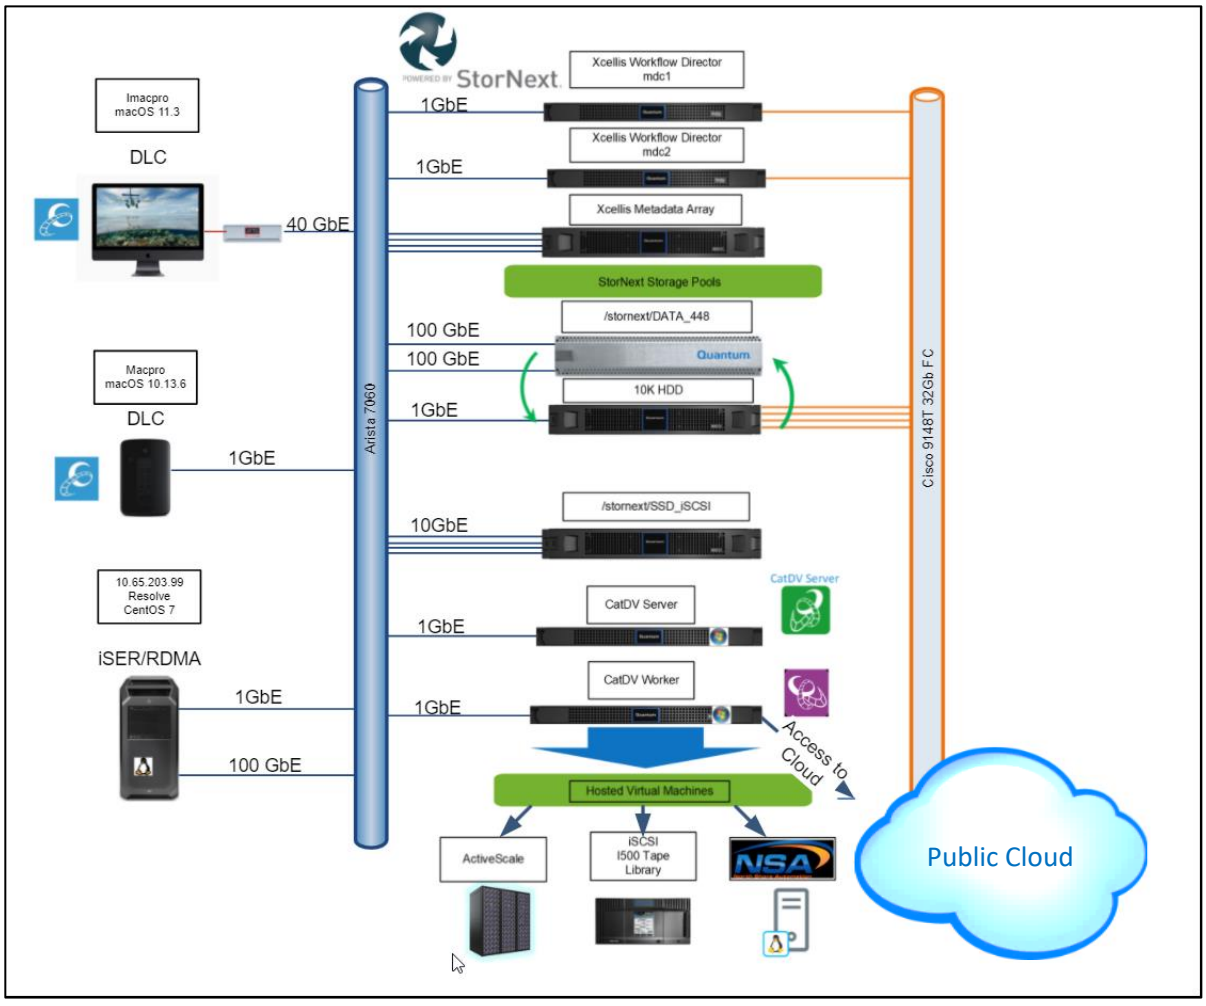

Figure 5 –CatDV and StorNext Reference Architecture

Requirements

File System

The file system needs to have the ability to playback footage at real-time speeds for various workloads up to 8K resolutions. Depending on the workflow, multiple streams of high resolution and high frame rate files will need to be supported. If the workflow is designed around proxy workflows the file system needs to have appropriate capacity to handle additional assets. Further, the file system should also be highly available for uninterrupted access to assets. The file system should be presented to the client's workstations NFS and SMB NAS protocols, in addition to direct access and seen as block storage devices with fibre channel and iSER. Distributed LAN Client (DLC) is another ethernet connectivity option but presented with a proprietary LAN client software. As the underlying file system, StorNext can provide these options in addition to multiple tiering options including S3 targets and single namespace tiering called pools.

Networking

The network requirements can vary depending on the facility and the specific workflow. For this Reference Architecture, we selected 100GbE networking from storage to specific Linux clients. There is also a limited amount of 32Gbit fibre channel storage connected to the file system. Clients in this RA are connected via 1, 10, 100GbE ethernet with limited 32Gb fibre channel connections.

Public Cloud Targets

Public Cloud targets will need appropriate credentials as part of the CatDV connections for private cloud locations. Multiple locations can be freely added to support advanced, hybrid environments, and multi-geo configurations.

Database

A database will be required for any CatDV solution, while a lightweight built-in database is provided during installation, performance is highly tuned for MySQL and Microsoft SQL.

Configuration

The following sections detail the configuration required for CatDV and StorNext for this reference architecture. The objective of this section is to provide enough setup details that the reader can get the products working together and create a basic workflow.

Quantum StorNext Product Configuration Details

- Configuring or modifying the file system
    a. Create filesystems or modify existing file systems to accommodate multiple users
    Select Configuration > File System > New

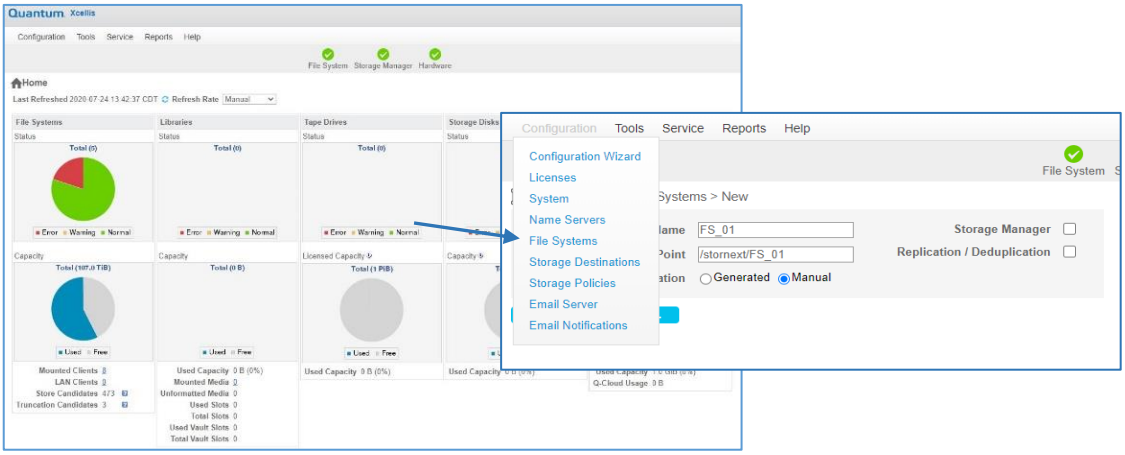

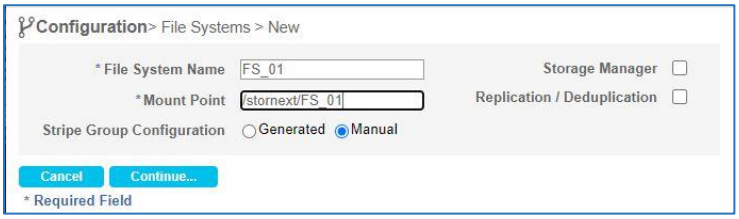

    b. Create a NAS presentation, SMB or NFS access. This is accomplished by starting up Connect on the Xcellis, as an example: https://xcellis2.quantum.com:442/

c. Select NAS on the left, then select “Create NAS cluster”

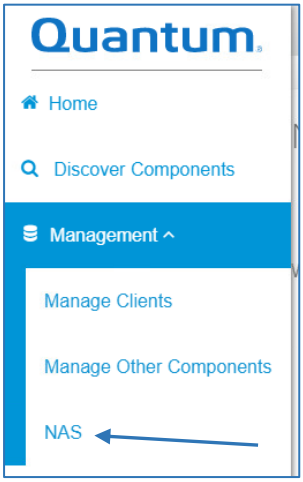

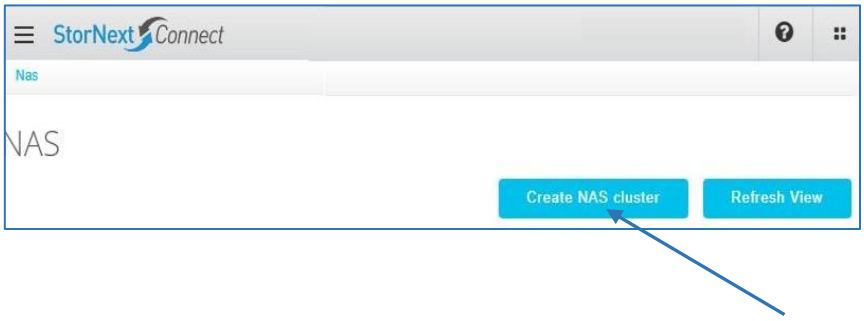

d. Configure a share or shares under the Shares Tab

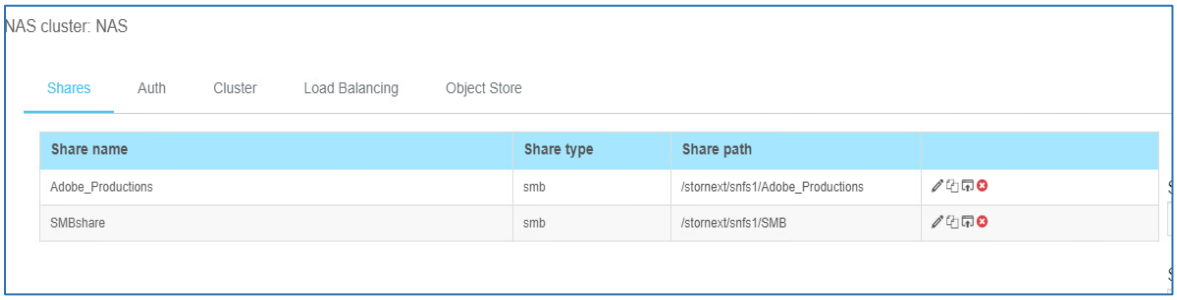

e. Configure users, under the Auth tab

- Chose Local, AD, OpenLDAP with Kerberos, Or OpenLDAP with Samba Ext.

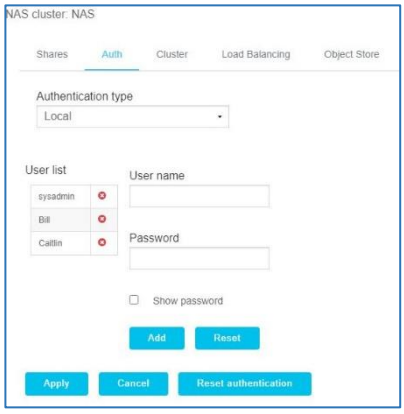

- Configure any Xcellis Workflow Extenders
    a. Configure any gateways for distributed LAN servers
    b. If a NAS is being presented from the Xcellis configure the gateways to ensure maximum bandwidth can be achieve to the storage and to the hosts

Storage Manager Configuration

- Define Storage Destinations
    a. Libraries – pick library connected to your infrastructure
    b. After addition of tape library, view configuration to see available drives

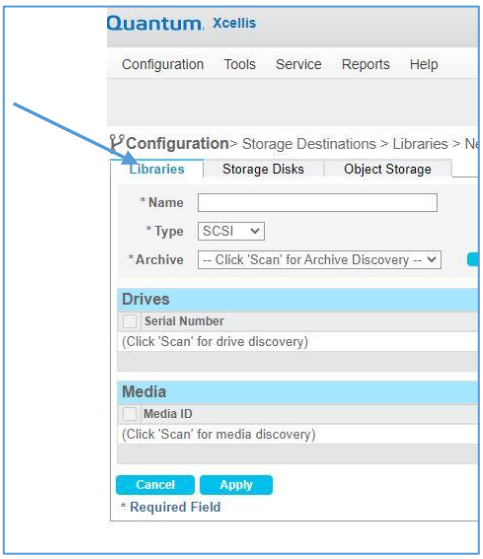

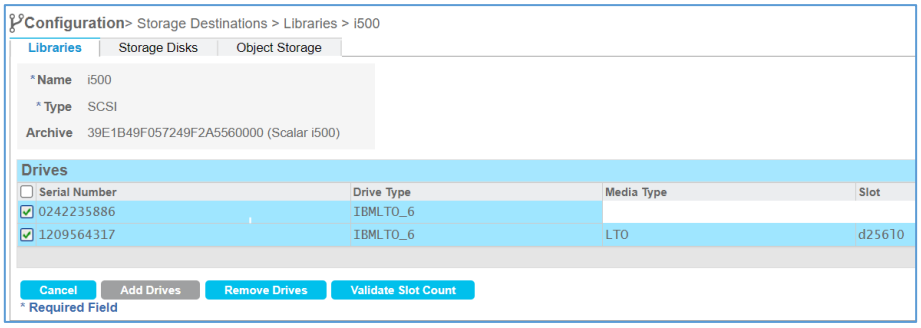

c. Storage Disks – used for near-line storage

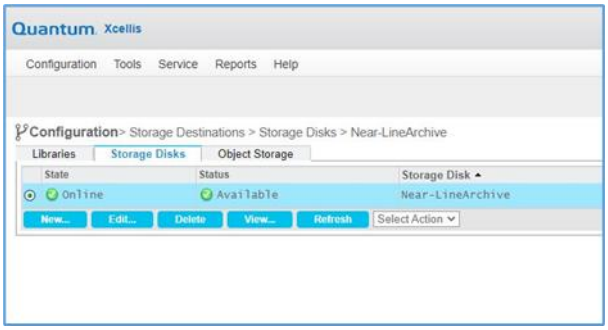

d. Object Storage – S3 Compatible on-prem or in the cloud (existing Object Store and Cloud account created)

- Top section Enter Name, Host, port, username and password to Object Store

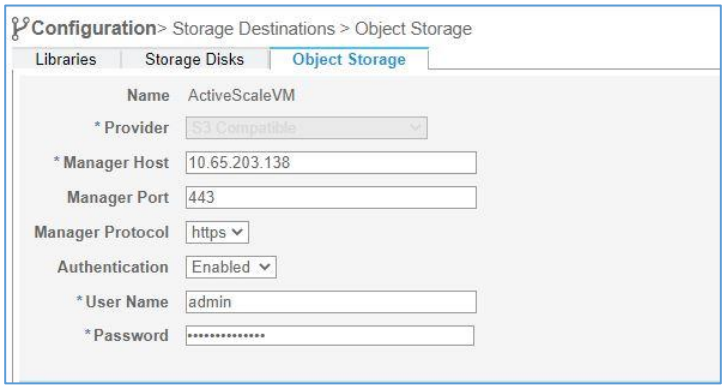

- Next Section, add a controller, by clicking add, most on-prem have multiple controllers

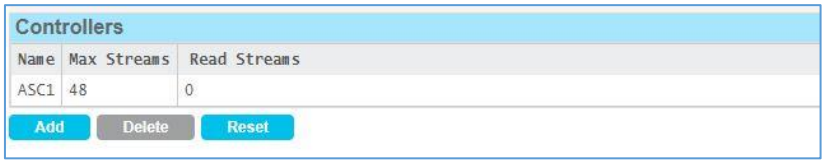

- Add I/O path/s, define path name, Controller name, media type, URL style, along with Object Access Protocol (https preferred), finally the Host:Port optional

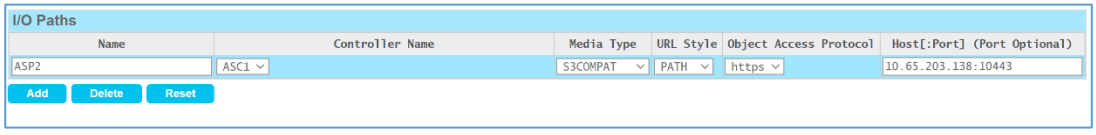

- Finally, the Containers section, click add and complete the fields:
    a. Containers: name of the OBJ bucket
    b. Media ID: unique name for this bucket
    c. Media Type: S3 Compatible
    d. Singing Type
    e. Storage class
    f. Authentication Type
    g. User Name: Access Key ID
    h. Password: Secret Access Key
    i. Policy Class

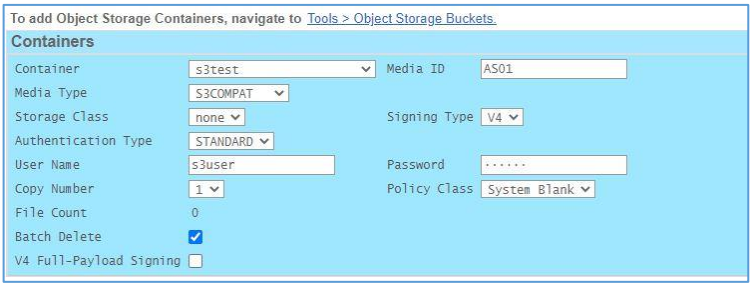

Quantum CatDV Product Configuration Details

1. **Server** - Download installer for specific operating system here
    a. Post install – Configure through Server Control Panel 5 tabs, instructions follow link above
        1. Installation – location of installation
        2. Licensing – install registration code
        3. Database – either the lightweight built-in database commonly used for workflow prototyping or small deployments, or more commonly MySQL or Microsoft SQL, for which CatDV is highly optimized. Additional database choices are possible but may require engineering assistance to ensure that the system's metadata operations can be properly supported.
        4. Web Server - CatDV deploys its own webserver as part of installation, though can be integrated with other web servers or existing web server implementations.
        5. CatDV Server – the CatDV application server is automatically setup during installation on a separate server IP address and port number

    b. CatDV Enterprise or Pegasus Server Software - minimum requirements are as follows:
        - Windows or Linux Servers (Can be virtualized)
        - 4 Cores
        - 16GB RAM (32 Recommended)
        - SSD (NVMe recommended) 40GB (More for larger databases)
        - 10GbE network connection

2. **Worker** - Download installer for the specific operating system from here
    a. Install client per instructions, follow the link above
        1. Install License
        2. Enter Server hostname or IP under the CatDV tab

3. Proceed to “watch actions, click + to add new action. This part takes some investigation and planning. Review “Configure the Worker Node” in the manual
4. Another good section “Hints and tips”

b. CatDV Enterprise or Pegasus Worker node Software – minimum requirements
   - Windows or Linux Servers (Can be virtualized)
   - 12 Cores
   - 24GB RAM
   - SSD 40GB
   - 10GbE network connection

This example was designed with 4 simultaneous Worker node threads running at any time. More may be needed depending on the customer environment

3. Web-Client

b. Compatible with most modern web browsers are compatible, the interface uses HTML 5/CSS3/JavaScript
c. Licenses should be obtained along with Server license
d. Can be deployed in a DMZ with SSL certificate
e. Administrator or created username, then opens the CatDV Web3 interface

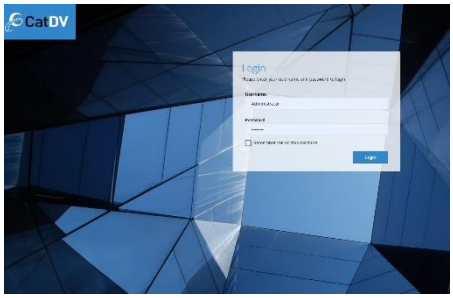

4. Desktop Client - Download installer for specific operating system from here

a. Install the client and following instructions
b. License the application with registration code
c. Connect to server with proper credentials
d. Overview of main features
   1. Import clips
   2. View clips
   3. Outputting clips

5. CatDV Active Scale Archive Plugin

a. Installed on CatDV Server
b. Can be deployed on 1 or more dedicated Archive servers for load balancing and redundancy

6. Plugins

a. Server-side
   1. CatDV Project Template Plugin
      - Allows definition of one more Project templates

2. Some licenses for worker plugins will be licensed on the server

b. Worker-side
   1. Download worker plugins from here
   2. Install plugins per instructions, follow link above, each plugin has different instructions, but the plugin usually resides in:
      - Windows - C:\ProgramData\Square Box\Extensions
      - Linux: /usr/local/catdvWorker/extensions
      - Mac: /Library/Application Support/Square Box/Extensions

7. **Optional CatDV Tools**
   a. CatDV Enterprise or Pegasus Client

**Pegasus Client is required for complex cinematic media types to playback natively in the application examples - RED, R3D, Arri, etc...*

8. **CatDV Adobe Panels and FCPX Extensions**
   a. Installation instructions here
9. **Dedicated CatDV Archive Server**
   a. Final option, have a dedicated server to manage archive tasks

Validation Activities

The following procedures are needed to validate the StorNext integration with CatDV. Testing conducted solely by Quantum staff.

1. Configure CatDV to catalog assets on StorNext file system
2. Worker setup to process workflows to and from file system
3. Archive to an archive target defined in Storage Manager using QSM plugin, monitor CatDV ability to follow assets
4. Restore from Archive target using QSM plugin to the File system, then updating CatDV to reflect that restore.

StorNext and CatDV Workflow Testing

1. Show assets being ingested into CatDV
    a. By selecting the + sign you can drag and drop assets into CatDV

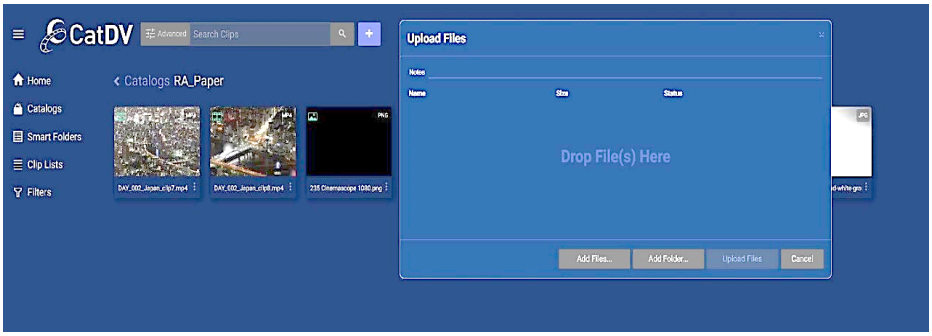

2. Archive to Storage Manager with CatDV using NSA integration, QSM-supports two types of archives
    1. HSM style archiving using Storage manager policy folders
    2. Direct in-place archiving
        i. First steps no matter what method
        i. Download and install VM on Worker node from NorthShore Automation
            1. NSA PS installation required after download
        ii. Worker-actions are defined in the configuration menu on the worker.
            1. The worker actions point to the NSA VM running on the worker-node
                a. QSM – Archive, Archive and Truncate, Restore, Truncate

3. Archive asset in watch folder to Tape, disk or S3 target
    - QSM - Archive worker-action

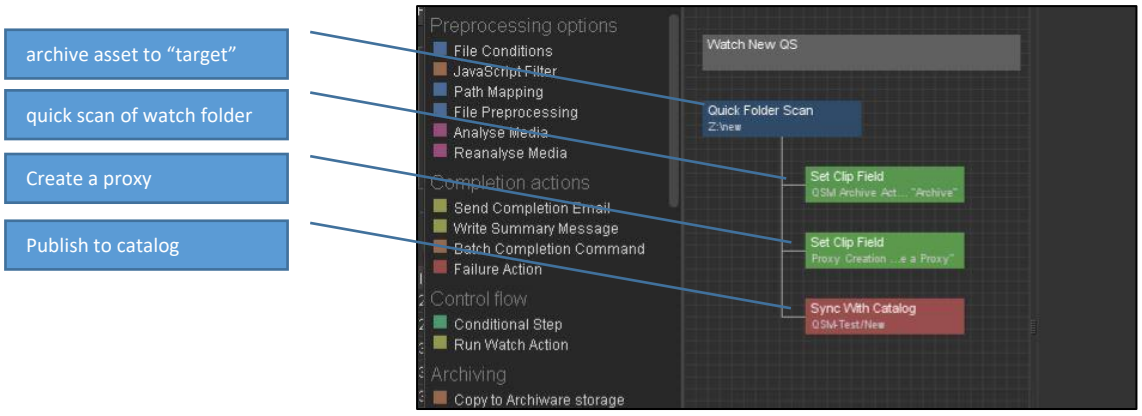

- Below is a list of other plugin options within the worker-actions on the worker

4. Initiating actions in Web client
   - QSM Archive in CatDV – QSM Archive Tab

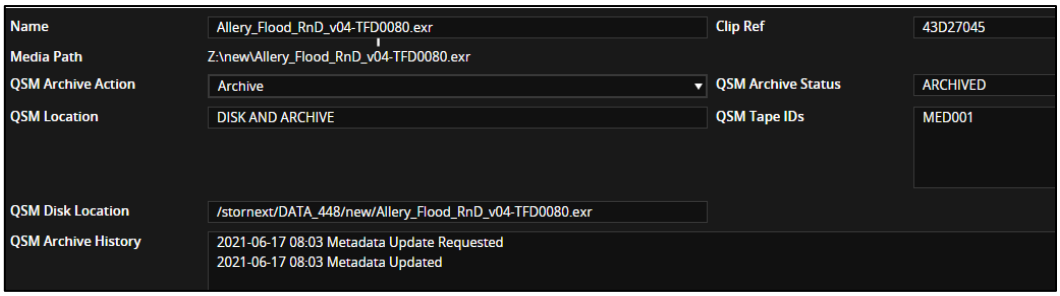

   - QSM Truncate in CatDV – QSM Archive Tab

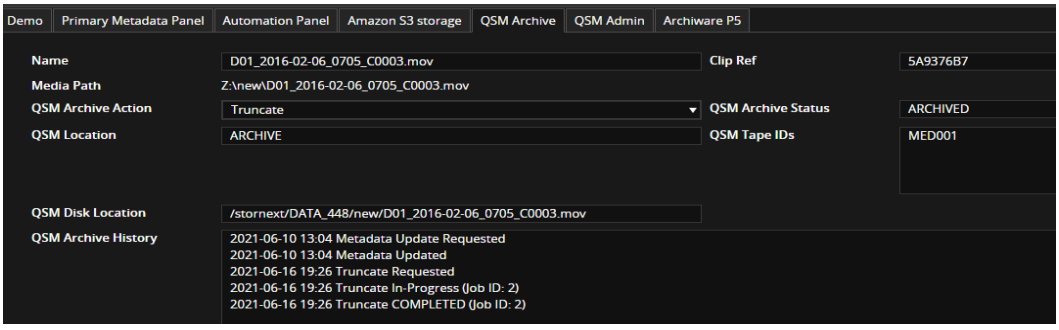

- QSM Restore in CatDV – QSM Archive Tab

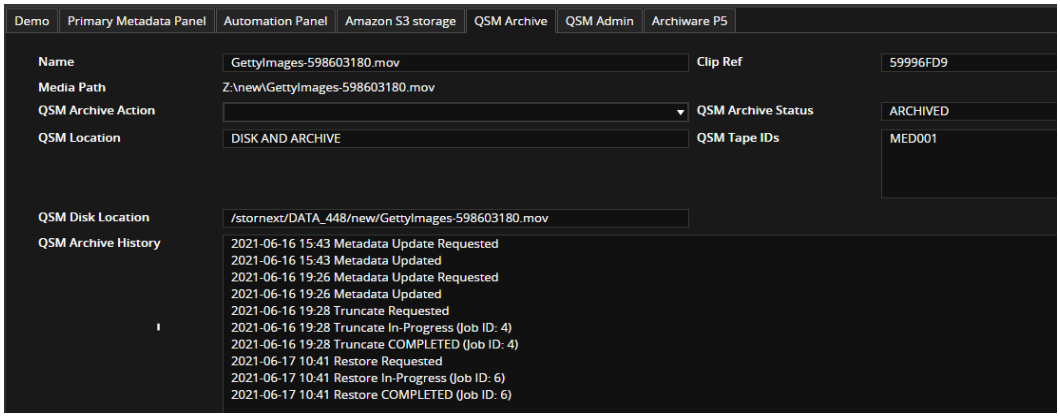

- QSM Update in CatDV – QSM Archive Tab
    i. With Update, we are going out and doing a query of the file looking for any information that you would typically get with "fsfileinfo" Then taking that output and populating the fields in the QSM Archive tab

b.

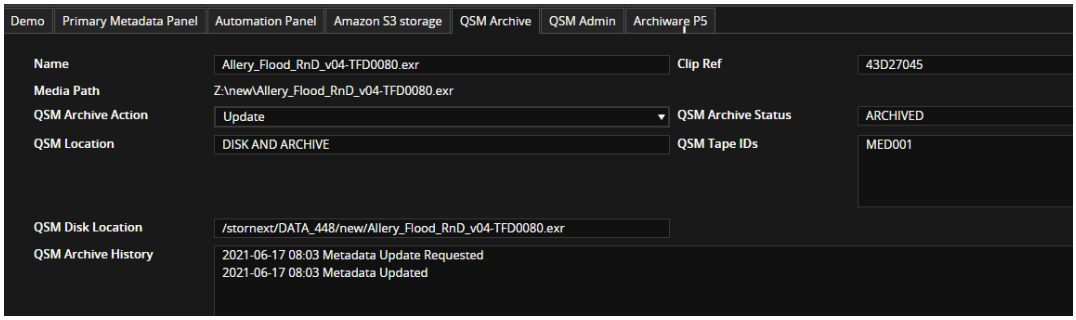

Considerations

When considering implementation of this workflow by sales, presales, and solution architects, it is recommended to conduct a thorough inventory the customer's unique workflow desired, current and future needs, and the existing systems and network in place. What does the capacity and speed of the current installation look like? What is the growth rate of content, project demands? For new installations, be prepared to architect a solution for immediate concerns but know the growth and sizing steps likely to be needed up to 6 years of projected future demand. What does the landscape look like in terms of expandability of storage capacity, performance requirements, asset management growth? Will the production facilities of tomorrow look different – perhaps emphasizing more dispersed, local facilities in different regions rather than very large, consolidated production centers? Will there be a higher emphasis on NVMe storage locally, and object storage everywhere else? Will cloud delivered content analysis solutions and API-level reporting or solution integration be needed? Look for future opportunities to expand capacity and performance as the operation grows and seek to build the infrastructure capability for the future today, with a flexible mindset, and platforms that can help you respond to ever-changing needs.

Summary

A solution based on Quantum StorNext with Storage Manager along with CatDV for ongoing asset management and long-term Archive abilities provides a wide range of options and a highly flexible framework to meet a customer's asset management and protection needs for the tremendous amount of assets being gathered today. With CatDV and Storage Manager delivering multiple options for archiving targets and asset management, you can easily deliver an on-premise highly robust solution to manage the entire life cycle of your assets, including access to Public or Private cloud offerings not only for archive but as an off-site disaster recovery plan.

References

The documents below are references to configure the software and the systems for functional testing

Document Title	Download URL
StorNext Documentation Center	Link Here
CatDV Server installation	Link Here
CatDV Worker Manual	
CatDV Client Manual	

Version History

Version	Notes	Date
1.0	Initial Release CatDV and StorNext Reference Architecture	January 2021

Quantum