

Efficient hybrid cloud workflows for Media and Entertainment

Introduction

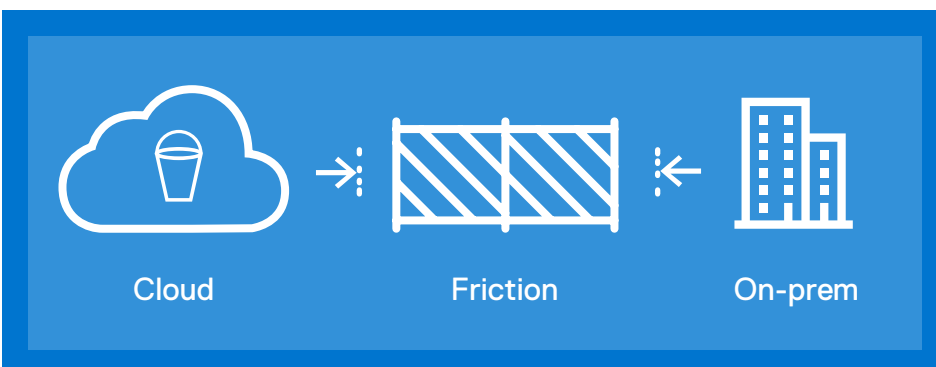
In the Media and Entertainment industry, production uses a distributed model with multiple studios contributing content from different physical locations. For example, a movie is a group of files organized by scene. Multiple editors work on the same movie and each editor can see all scenes in their own working directory. With the availability of cloud based tools for media production like rendering, post-production and AI-powered video, these workflows are shifting to a hybrid cloud operating model leveraging on-premise and Cloud media services. To make such a hybrid workflow possible production teams need to overcome multiple challenges.

Challenges of distributed workflows

The main challenge with this distributed model is that a single source of truth is never in one location. Which copy of the media files is the master? Moving content between locations becomes a task that falls on editors and content creators. To create a single source of truth and maintain a just-in-time distribution of content in a multi-location distributed team environment, a data orchestration solution needs to shuttle data from on-premise storage to the Cloud and aggregate object data from multiple locations into a single file system.

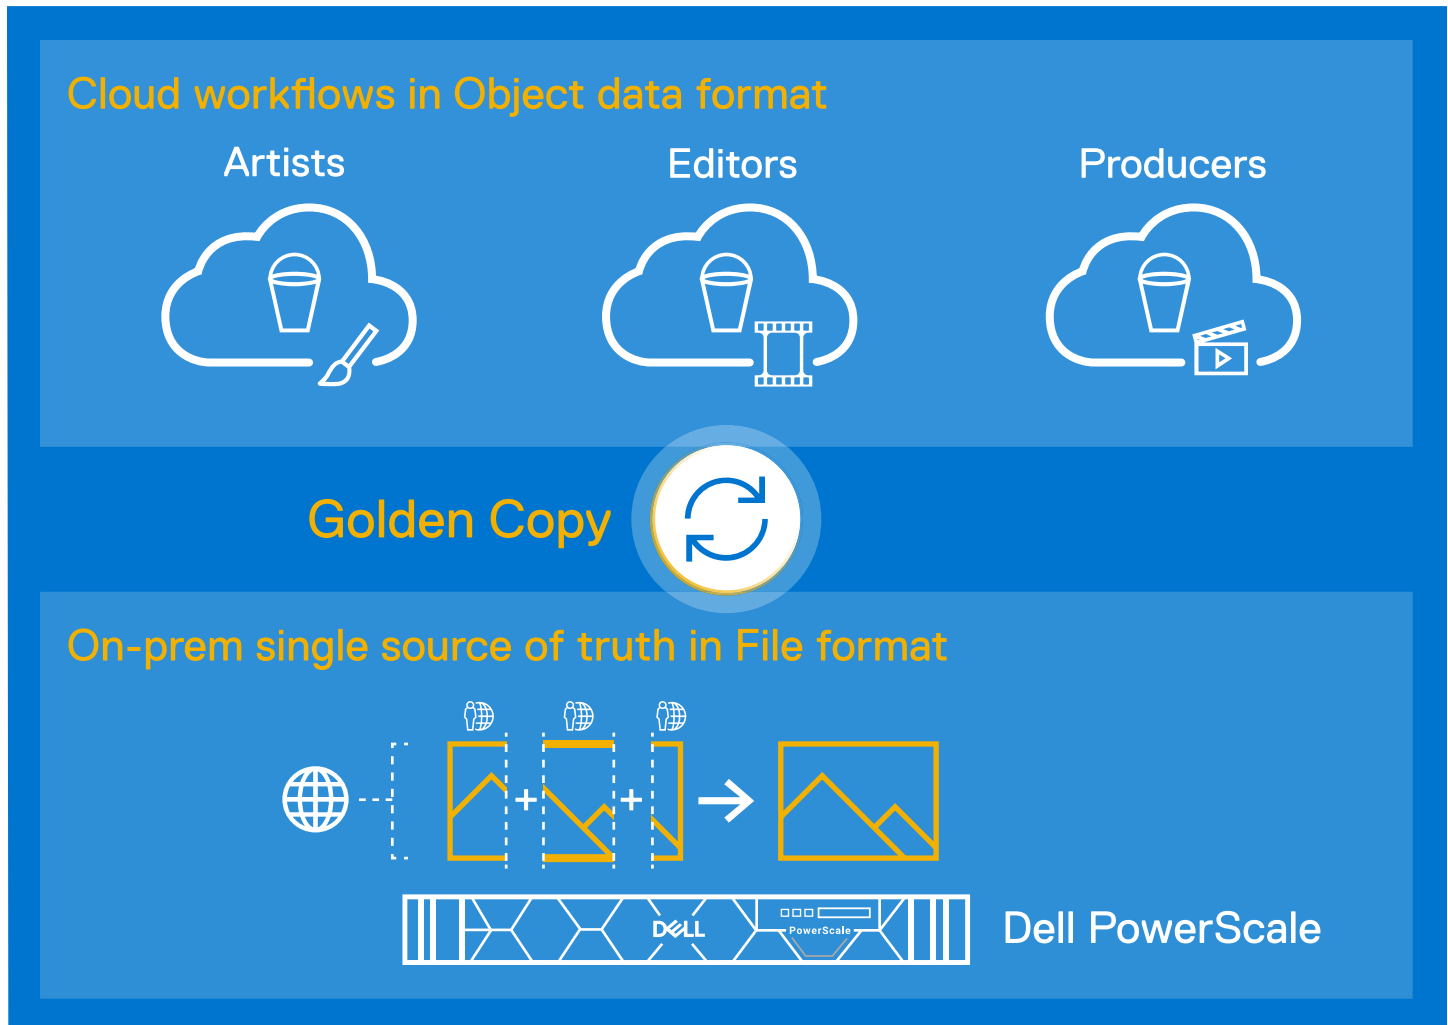
Files and objects have very little in common. The internet Cloud storage is dominated by S3 or object storage whereas on-premise video editing platforms are based on file systems. Dell Powerscale is a platform that offers high speed file access over SMB or NFS and it's scale out file system is ideal for video editing solutions. The problem with these different data storage formats is that metadata capabilities are different, permissions are different and access protocols are also different. Object access using https transport and specialized tools like Cyberduck or S3 Browser to access objects using an object key. Windows or Linux are used to access file systems over SMB or NFS. Simple things like date time stamps that were used to track last modified and creation dates of files don't exist in object storage with only a single date stamp used. Data management uses folder structures to group related files but in object storage the folder concept does not exist, allowing a flat name space. These differences complicate data management between cloud storage and file system storage.

A file system has symbolic links or hard links that allow for more efficiency when multiple copies of common files are used in media workflows. Links create pointers to scenes and present what appears to be a complete copy of the entire movie while only consuming space of each scene once regardless of how many copies have been made. Given the lack of symbolic or hard links in object storage, any copy from a file system to an object store could inflate the cost of the object storage by an order of magnitude and it will also take longer to copy the data. If the data is copied from object storage back to a file system, the symbolic or hard links are gone and now the full size of the data exists in the file system as well. All the advantages of the file system have been stripped from the data just by copying it from file to object and object back to file.



The solution: Superna Golden Copy

Superna Golden Copy is an intelligent and efficient data replication technology built for Dell PowerScale that solves. The data replication bridges the gap between file and object formats while enabling content creators seamless participate in a distributed hybrid cloud workflow. These key issues for hybrid cloud workflows between file systems and cloud object stores.



Data Reduction

Data reduction is possible with the Golden Copy auto detection of symlinks and hard links and inodes. Golden Copy maintains file system symlink and hard links and inode references by storing the data in the object store in a manner that does not duplicate the data and preserves the link information.

The symbolic and hard links are transparently encoded with inode object references on upload and on recall to a file system the sym and hard link references are automatically maintained.



Metadata Integrity

Metadata integrity is key to allowing workflows that use a hybrid file system and object storage approach. Golden Copy encodes metadata (created, modified, last access, mode bits, owner, group and folder ACL's) into custom properties in the object store.

This enables recall of objects to files to maintain the key metadata needed by storage administrators to properly manage data security and data ownership.



Single Source of Truth

With Golden Copy media workstreams can have a working set of media data on-premise for use by editors and content creators and to have the master copy in the Cloud for remote studios to get the current view of the project. In addition, the Cloud is used to distribute contributions from distributed teams for review and integration into the project as the content can easily be uploaded to object storage from anywhere and from any time zone.

Workflow automation with Pipelines

The Golden Copy dropbox solution enables on-premise to cloud and cloud to on-premise automated movement of data. This is part of the Golden Copy Pipeline feature. In the on-premise to cloud data movement direction, new or modified data is auto detected leveraging the Dell Powerscale changelist API to detect file system changes. This offload to Dell Powerscale allows fast detection of new or modified content in the file system regardless of the size of the file system. The new or modified data is copied to the cloud in object format while maintaining the file system hierarchy in the object storage.

When multiple teams contribute content from remote studios into Cloud storage, this data needs to be synchronized back to the file system. The Golden Copy Pipeline feature allows monitoring of Cloud storage in various locations across any Cloud storage provider and syncing of this data into the file system using a mechanism that prevents copying data that already exists in the target file system.

Both directions of data synchronization support “move” data functionality that deletes data after the copy is completed so that duplicate data and data sprawl is avoided.

In a dynamic environment, data recovery and troubleshooting issues still require a method to recover data that was deleted. The Golden Copy workflow supports temporary file system snapshots with expiry to allow recovery of data that was deleted but should not have been deleted. In the cloud to file system direction the solution places deleted objects in a “trash can” bucket that can have deleted data available for a period of time to administrators.

Summary

Golden Copy is built to solve key challenges of a hybrid cloud media and entertainment workflows. It drastically simplifies collaboration among artists, editors and producers while minimizing the cost of storage. The solution has already become indispensable at some reputed media houses that take advantage of PowerScale's scale, performance and flexibility.

Discover more about PowerScale platform



[Learn more](#) about our PowerScale platform



[Follow](#) Dell Storage on Twitter



Contact a Dell Technologies Expert for [Sales or Support](#)