

Cloud Aqua for Media Applications

Introduction

With the explosive growth of digital content at record pace, more and more content providers broadcast operators are deploying multiple, geographically dispersed facilities, connected via a WAN (Wide Area Network), for their content creation, production, and delivery workflows. At present each of these facilities is an isolated island, with its own media applications, local storage, and content management. Operators rely on home grown, manually operated tools to manage and distribute content across these multiple facilities. Not only does this ad hoc approach incur significant operating cost and WAN consumption, it also does not allow operators to achieve their ultimate goal, which is to access any content in real time from anywhere regardless where the content is located. Multi user, multi system, and multi site access also introduce a lot of management complexity leaving companies not much of a choice but to spend more in managing the overall infrastructure not just at the onset but throughout the life of the company itself. It's increasingly clear that a new type of global infrastructure is required to consolidate today's isolated islands into one reliable, scalable, global content repository.

This white paper discusses the requirements and architecture of such a global infrastructure, Private Cloud Storage, and why XOR Media, with the storage technology and years of experience building media related IT infrastructure for the broadcast and cable operators, is in the best position to deliver such a solution.

Cloud Storage Requirements

To realize the promises of accessing any content from anywhere and intelligently managing content, a cloud storage solution needs to meet the following requirements:

Application transparency

Traditional media applications are only capable of accessing content on local storage via standard file system interface. A cloud storage solution needs to provide a virtualization layer that allows these media applications, without application level modifications, to access content on geographically dispersed storage systems. A cloud storage solution without such virtualization layer will experience limited adoption due to the impracticality of requiring changes to all

mainstream media applications.

Universal access

An application should be able to access cloud storage using the most appropriate access method specific to that application, be it low cost NAS or high performance SAN, on Windows, Linux, or Mac OS, via 1G/10G IP or Fibre Channel network. In other words, it should not require significant changes to the current infrastructure when cloud storage is introduced to replace the isolated local storage islands.

Locally reliable, scalable, tiered, globally expandable

Cloud storage does not eliminate the need for reliable and scalable local storage systems. On the contrary, a scalable cloud storage solution is built on top of reliable and scalable local storage. A cloud storage solution should allow the operator to deploy reliable and scalable local storage today and migrate to cloud storage infrastructure tomorrow without ripping and replacing the existing investment. A cloud storage solution should also add another level of reliability and scalability, especially in terms of protection against individual site failures. Cloud storage is also naturally suited to support tiered storage, where frequently accessed content is placed on high performance SSD based storage while relatively cold content on low cost HDD based storage.

Policy-driven self management

Cloud storage should be self-managed without much day to day manual intervention. Content should be cached, replicated, migrated, tiered automatically based on content access pattern, pre-defined policy, content-aware intelligence, and the ultimate goal of ensuring quality of service while reducing WAN consumption.

XOR Media Cloud Storage: Cloud Aqua

With decades of experience and strong technology background in media storage, XOR Media is in a good position to deliver the next generation cloud storage solution to the media industry. Cloud Aqua, the private cloud object-based storage solution provided

by XOR Media, is optimized for media operation. It enables real-time sharing and intelligent management of media files across networks or geographically dispersed locations.

Scalable global storage pool

Cloud Aqua consolidates storage by providing a virtual file system overlay pooling available storage resources into a scalable global storage pool. Virtual file system allows media applications to access geographically dispersed cloud storage without knowing the location of the content. Internally virtual file system relies on a global namespace to locate media content and access the content either locally or remotely over the WAN.

Cloud Aqua employs a federated architecture where underlying storage systems or hardware can be on different networks or geographically-dispersed locations. This makes Cloud Aqua also globally scalable without disrupting existing connections or access where for example, adding capacity to an existing storage node in Cloud Aqua or adding a new storage node, will also scale the entire cloud object store.

Through virtualization, Cloud Aqua facilitates comprehensive management of storage resources (from multi-user management to access control and quota settings, to authentication via integration with LDAP or Active Directory, etc), offering tremendous performance, scalability, manageability, and concurrent high-speed access to heterogeneous clients or applications, all across a single global namespace.

Object-based architecture

Cloud Aqua employs an object-based architecture, which means media files are encapsulated together with metadata (attributes that define or describe the associated content/data) into a single object. Unlike files which have standard attributes defined (such as file creation date, time, type of file, etc), objects are self-describing. They contain user-defined metadata – attributes that allow applications to further describe or define media content in

an object. Metadata makes objects self-describing. This is what brings about intelligence at the storage level.

Intelligent, policy-driven media processing

In Cloud Aqua's object-based architecture, the storage understands what is being stored - what the media is and what the metadata is. Cloud Aqua is aware of both content and its attributes, opening many different possibilities in optimization, searching, and performing automated actions at the storage level. Cloud Aqua provides a way for applications to communicate more efficiently with the storage systems.

Through the use of metadata, Cloud Aqua empowers automated media processing. Based on user pre-defined policies, Cloud Aqua can automatically take actions on a media content depending on its metadata (i.e. folder, tag, creation date/time) or during an event (such as an object creation, object update).

Examples of actions that can be automatically triggered upon execution of pre-defined policies would be:

- Create data file replica(s) in different regions or storage clusters
- Transcode media content into different formats
- Distribute transcoded content to expected site(s)

Such policy-driven automated actions can be especially useful in content distribution, media file processing, backup creation, versioning, storage optimization, content caching (on demand local copy creation), on-demand automated content placement in Hierarchical Storage Management (HSM), data migration, and many other uses that enable higher levels of efficiency.

Cloud Aqua also has a built-in media workflow engine that picks up from automated processing abilities and allows one to setup workflows that incorporate a set of Cloud Aqua processing functionalities specifically for media files. Such set of actions include transcod-

ing, encryption, low-bit proxy generation, and content replication/distribution that may be triggered upon a media file being created, updated, or transferred (metadata updated). Based on pre-defined policies, these actions can affect efficiencies in different media workflows. Given these capabilities, Cloud Aqua is made integral to various video/media systems outside Cloud Aqua: play-to-air systems, VOD, CDN, Digital Media Archiving, Production, and a lot more.

Multi-tenancy Infrastructure

Cloud Aqua allows flexible resource allocation and cloud services through its tenant-based infrastructure. This means a global storage pool can be provisioned to multiple groups of users (such as branches within a company), providing potential not only for flexibility in resource allocation, but possible improvements to economies of cost from consolidating resources and streamlined management of resource provisioning.

Cloud Aqua provisions storage (from its entire global storage pool) in a "leasing model" to multiple tenants that manage and share media files within their own domains.

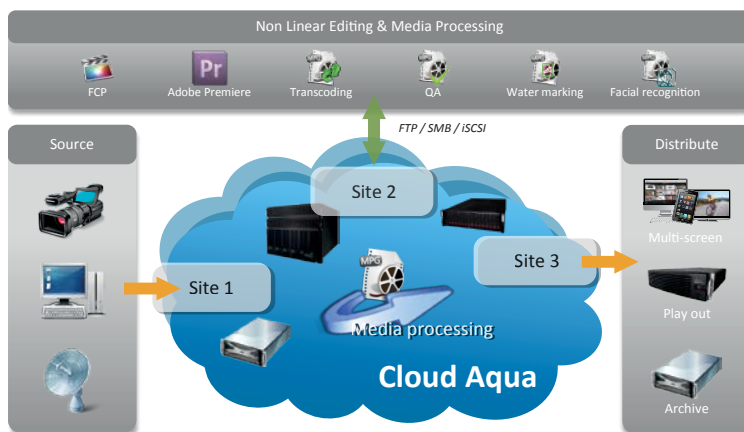
WAN optimized transfers and content distribution

Traditional data transfer technology such as FTP is not optimized for media content distribution over the WAN. Data transfer rate by FTP over the WAN is relatively low due to TCP's limitation in the presence of long network latency. Cloud Aqua is an optimized infrastructure situated in different networks and/or physical locations connected over wide area networks. It implements reliable UDP based peer to peer data transfer technology to optimize transfers and streaming even over long distances, allow advanced features such as stream through and access while transfer, and further improve quality of service and reducing WAN consumption.

Solutions intended for media applications

Cloud Aqua provides features that bring about many possibilities in managing, handling, and transferring media contents in a storage cloud. These features make Cloud Aqua ideal for different media-related cloud storage solutions, empowering creation and operation of centralized media libraries, content transformation and content distribution, content and metadata processing workflows, media playout, and many others.

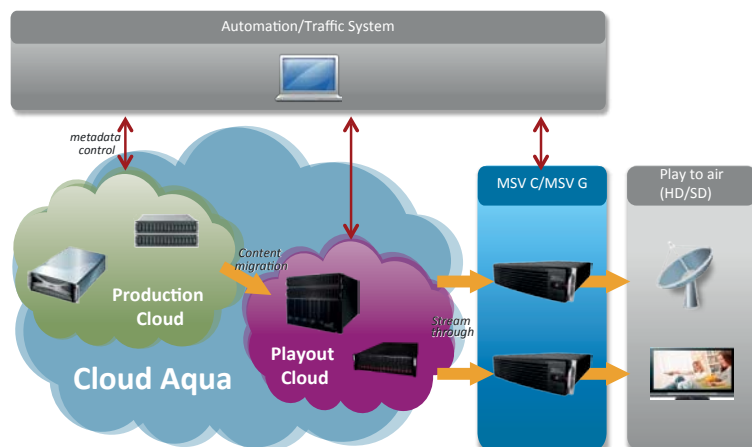
Cloud Aqua for media production



Benifits

- Heterogeneous storage hardware.
- CIFS SMB2 iSCSI performance tuning for Adobe, FCP, etc.
- Integrated media processing modules.
- Collaborate across organizations and geographically-dispersed locations.
- Extend a single MAM system across multiple locations.
- Intelligent asset management throughout its lifecycle.
- Increased reliability and scalability.

Cloud Aqua for media delivery



Benifits

- Automated content management based on playlist.
- Integrates with automation or traffic systems.
- Automated content migration from near-line to online storage.
- Deletes content from PTA storage based on playlist.
- Playout cloud with richer features and lower cost.
- Provides N+1 scalability.
- Robust content availability.
- Cost-effective solution.