

## Solution Brief

# NetApp StorageGRID for Media Workflows

Manage your data across space and time

### Key Benefits

**Simplify the creation and maintenance of content repositories and distribution**  
StorageGRID creates a single namespace of content assets that can span a facility or the globe.

**Flexible choices for data protection with self-healing data durability mean 100% uptime**

StorageGRID erasure coding and replicating data protection options maintain access to content without the performance degradation or downtime associated with RAID protection.

**Integrate with existing MAM and production applications to extend access to content for all departments**

StorageGRID easily integrates with applications through standard protocols like S3 and HTTP/Rest. The list of media tools implementing object storage capability is large and growing.

### The Challenge

Media is ubiquitous, and complexity abounds. Social media sites alone store billions of video clips and photos. Broadcast networks, movie studios, and media streaming services are now global enterprises. The challenge of managing, moving, and monetizing media assets is overwhelming. The cost of managing exponentially growing media repositories is outstripping the cost of data storage capacity.

At the same time, media enterprises are battling the complexities of dispersed production and distribution operations across remote sites. Multiple sites with separate media asset management, production, and distribution applications are causing a proliferation of file copies with little or no coordination among workflows.

Media companies are being forced to reevaluate how to effectively manage and migrate large amounts of content spread over many locations. They need to confirm that content is stored in the right location, on the right tier, at the right time, reduce duplication of content files, and identify and delete copies that are no longer needed. How can enterprises improve workflows across media asset management domains?

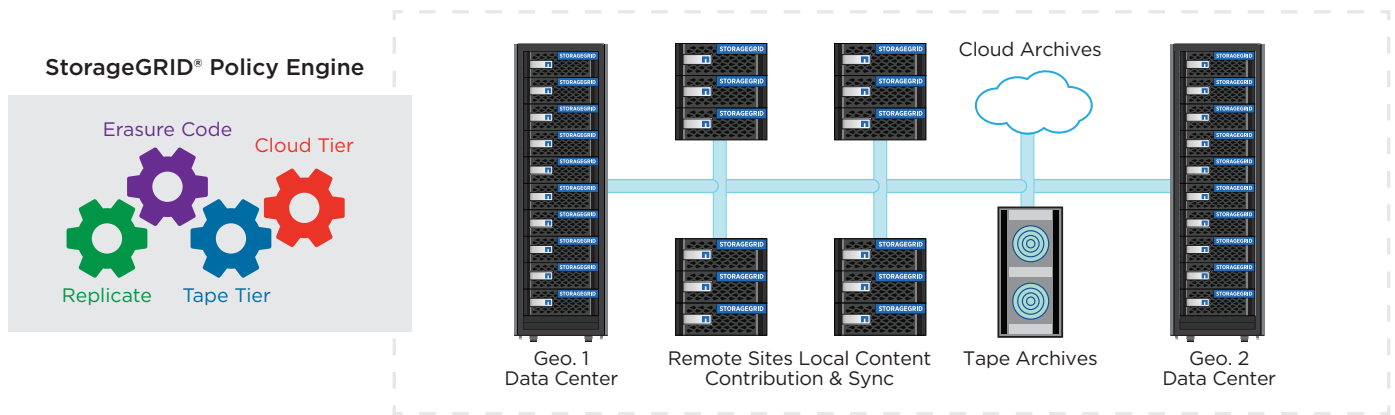
### The Solution

To support vast media repositories, broadcasters and enterprise cloud providers use NetApp® StorageGRID® object storage. Unlike file-based storage systems that rely on limited metadata for access, object stores combine files and their unique metadata into objects that can be coherently accessed by different media asset management systems, file delivery systems, and other production and distribution applications.

StorageGRID is a scale-out system designed to support multiple nodes across internet-connected sites, regardless of distance. With the intelligent policy engine of StorageGRID, you can choose erasure-coding objects across sites for geo-resiliency or object replication between remote sites to minimize WAN access latency.

As content ages, policies can transition from replicas to erasure coding or move it to tape or to the cloud automatically. This flexibility eliminates cloud vendor lock-in and the need to manually move or delete millions of files, greatly improving the economics of your media repository.

Figure 1) StorageGRID policies can be administered from any location to manage data on a local, regional, and global basis. Users and applications see StorageGRID as a single namespace. Always presenting exactly the expected version of localized content makes StorageGRID the ideal repository for content delivery.



### Reduce the Complexity of Media Repositories

With the StorageGRID single global namespace, it doesn't matter where or how an object is stored and protected, it's available to any user at any site at any time. With massive scalability across this namespace, StorageGRID frees your media applications from the task of finding and moving files among volumes, systems, and sites.

StorageGRID enables you to establish highly granular, flexible data management policies that determine how data is stored and protected. When developing and enforcing policies, the object store examines a wide range of characteristics and needs, including performance, durability, availability, geographic location, longevity, and cost.

### Optimize Media Workflows

StorageGRID is not a media asset management system. It's a data management system that offers sophisticated policy-driven functionality to support your media asset management systems. Policies can be enforced at ingest, at rest, on read, after object attribute update, at object disposition request, or when information lifecycle management policies change. StorageGRID policies can initiate object replication, erasure coding, tiering, and archiving to tape or to the cloud.

This storage management, monitoring, and migration functionality delivers an improved experience for all the workflows in your media enterprise. Media asset management systems support specific production or delivery workflows in an enterprise. An object store confirms that the data that those systems need is in the right place at the right time, and that it's easy to access that data through the global namespace.

“StorageGRID is allowing us to offer more advanced cloud media storage functionality at a lower cost.”

—Andrew Sjoquist, Founder of ASE IT

By erasure coding, moving, or copying objects to sites to meet production and distribution access demands, StorageGRID relieves your media asset management systems, users, and administrators from manually performing these tasks. When an object is stored, it is seen and accessed as one object by all clients, regardless of where it is or how many copies exist.

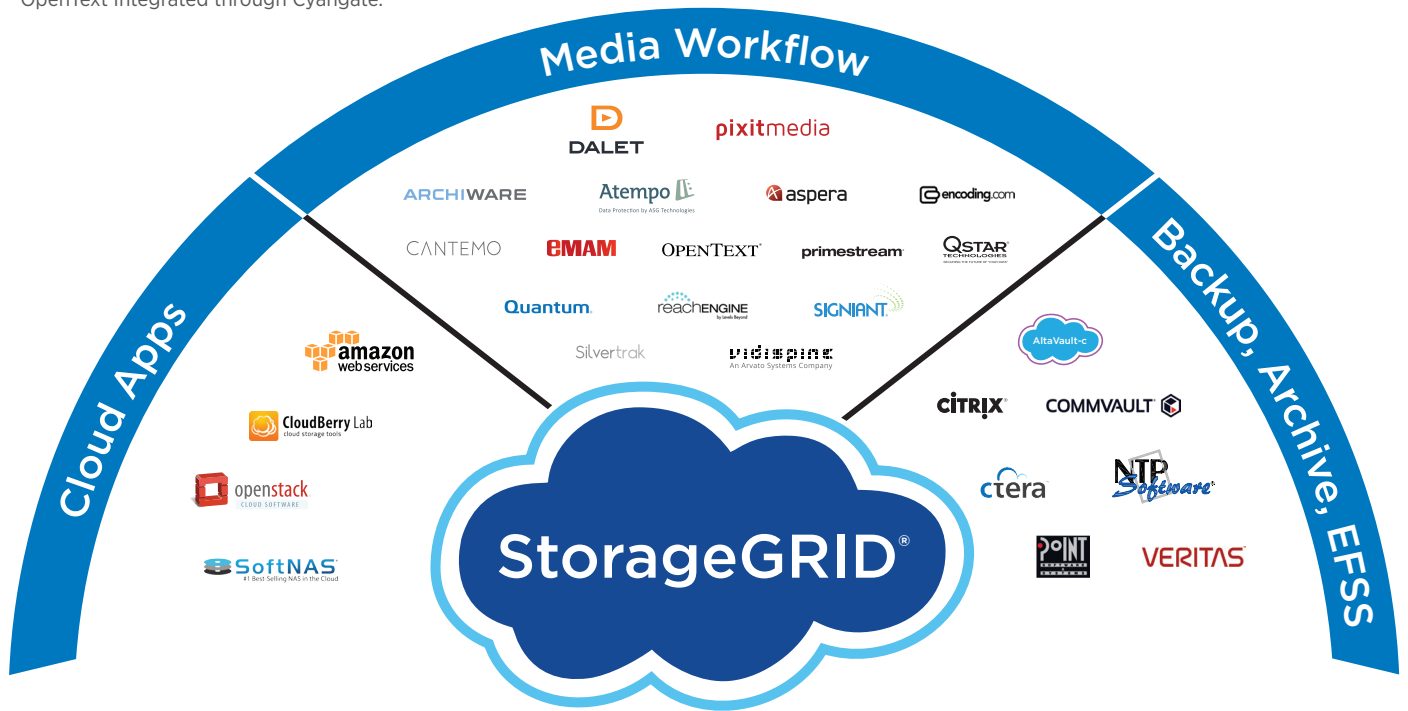
With object store support, media asset management systems can pass key metadata along to other systems and workgroups by storing it with the object along with the media file data, or “payload.”

### Erasure Coding

Erasure coding is a method of data protection in which data is broken into fragments, encoded by mathematically combining the fragments into a greater number of data pieces composed of multiple copies of the original fragments, and dispersed across different locations or storage media. Erasure coding provides improved overall storage efficiency versus replication, with equivalent overall durability and availability. In an object store with geo-dispersed erasure coding, data can be retrieved from a subset of the data pieces. This feature allows object access even during temporary or permanent failure of one or more drives, nodes, or sites.

Figure 2) Growing StorageGRID media ecosystem. Media asset management systems, production, distribution, backup, and archival applications access StorageGRID natively by using the Amazon Simple Storage Service (Amazon S3) interface.

\* OpenText Integrated through Cyangate.



**Minimize the Cost of Cradle-to-Grave Data Management**

Providing ubiquitous content access often requires storing copies of the same asset in multiple places. This approach not only increases complexity, it also requires additional storage capacity. With StorageGRID, you can minimize the capacity that is needed while optimizing for data resiliency, production access, and distribution requirements through policy-driven automation.

The policies that you set to move and copy objects can also include deletion criteria. For example, copies of objects can be automatically deleted after set periods of inactivity or after their distribution rights have expired, reducing capacity requirements.

The value of content changes over time, as does the cost of storing it. High-value content requires high-performance storage to keep it readily available. As content ages and is accessed less often, policies can automatically reduce the number of copies, move to an erasure-coded protection scheme, or move that content to less expensive disk arrays, tape, or public or private clouds. Management and maintenance of file copies, migrations, and deletions are driven by the StorageGRID policy engine.

The content in your repository will live longer than the media that it is stored on, and StorageGRID will greatly reduce migration challenges and costs in the years to come. The policy engine will drive the process of moving millions of objects from aging media to newer, lower-cost, higher-capacity media in the future. Upgrades, expansions, decommissioning, and refreshing activities are all nondisruptive in the StorageGRID environment.

With deployment in the most demanding industries since 2002, StorageGRID is the system you can trust with your growing media repository.

**About NetApp**

NetApp is the data authority for hybrid cloud. We provide a full range of hybrid cloud data services that simplify management of applications and data across cloud and on-premises environments to accelerate digital transformation. Together with our partners, we empower global organizations to unleash the full potential of their data to expand customer touchpoints, foster greater innovation and optimize their operations. For more information, visit [www.netapp.com](http://www.netapp.com). #DataDriven

FEATURE	DESCRIPTION
Erasure coding and replication	Replicate copies to remote sites to accelerate global workflows. Erasure coding across sites provides geo-distributed and cost-effective data integrity and protection
Deployment options	<ul style="list-style-type: none"> <li>• Software-defined storage that runs on VMware, Docker containers, and deployments, and that manages either NetApp or third-party storage arrays</li> <li>• Multiple NetApp appliances with node-level erasure coding and Dynamic Disk Pools technology, and the NetApp AutoSupport® tool, providing 99.999% availability at the storage-array level: <ul style="list-style-type: none"> <li>• SG5660—4U, 60-drive with near-line SAS drives</li> <li>• SG5612—2U, 12-drive with 4TB, 6TB, or 8TB near-line SAS drives</li> <li>• SG6060—5U/9U/13U with 58/118/178 near-line SAS drives in multiple density options</li> <li>• SGF6024 3U with 24 SSDs all-flash object storage appliance</li> </ul> </li> </ul>
Data availability	<ul style="list-style-type: none"> <li>• Fault-tolerant architecture supports nondisruptive operations, upgrades, and infrastructure refreshes</li> <li>• Load balancing automatically distributes workloads during normal operations and during failures</li> <li>• NetApp AutoSupport technology automatically alerts NetApp Supportengineers for proactive issue resolution</li> <li>• Node-level erasure coding further improves single-node availability (with NetApp E-Series Dynamic Disk Pools technology)</li> </ul>
Data integrity	<ul style="list-style-type: none"> <li>• Creates a digital fingerprint as data is ingested</li> <li>• Offers multiple interlocking layers of integrity protection, including hashes, checksums, and authentications</li> <li>• Provides data object integrity verification on ingest, retrieval, replication, and migration, and at rest; suspect objects are automatically regenerated</li> </ul>
Scalability	<ul style="list-style-type: none"> <li>• 100 billion objects</li> <li>• 400PB of capacity across up to 16 sites</li> </ul>
Interfaces	RESTful HTTP APIs include Amazon S3 and OpenStack Swift. Standard network protocols through a NAS bridge include NFS and CIFS

Table 1) NetApp StorageGRID features.