

SOLUTION BRIEF

NetApp StorageGRID for media workflows

Get your growing media data
under control



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 (MAM), 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 data and their unique metadata into objects that can be coherently accessed by different MAM systems, file delivery systems, and other production and distribution applications.

StorageGRID is a software-defined scale-out object storage solution that supports industry-standard object APIs like the Amazon Simple Storage Service (S3) API. The system is designed to support multiple nodes across internet-connected sites in multiple geographies, and up to 16 data centers. With the StorageGRID intelligent policy engine, you can choose either erasure-coding objects across sites for geo-resilience or object replication between remote sites to minimize WAN access latency, optimizing performance and data durability and availability.

Key benefits

- **Simplify the creation and maintenance of content repositories and distribution**
NetApp® StorageGRID® object-based storage solution creates a single namespace of content assets that can span a facility or multiple sites.
- **Flexible choices for data protection with self-healing data durability**
StorageGRID data protection options maintain access to content without performance degradation.
- **Deliver the right performance and scalability**
Start small and grow nondisruptively to scale capacity and performance.
- **Support multiple workloads, departments, and geographies**
StorageGRID is designed to support multiple users and departments to act as a private repository for a myriad of datasets.
- **Enable hybrid cloud workflows with the hyperscalers**
Take advantage of the public cloud to process, transform, and analyze your assets.

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 data among volumes, systems, and sites.

StorageGRID enables you to establish highly granular, flexible data management policies that determine how data is stored and protected throughout its lifecycle. 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

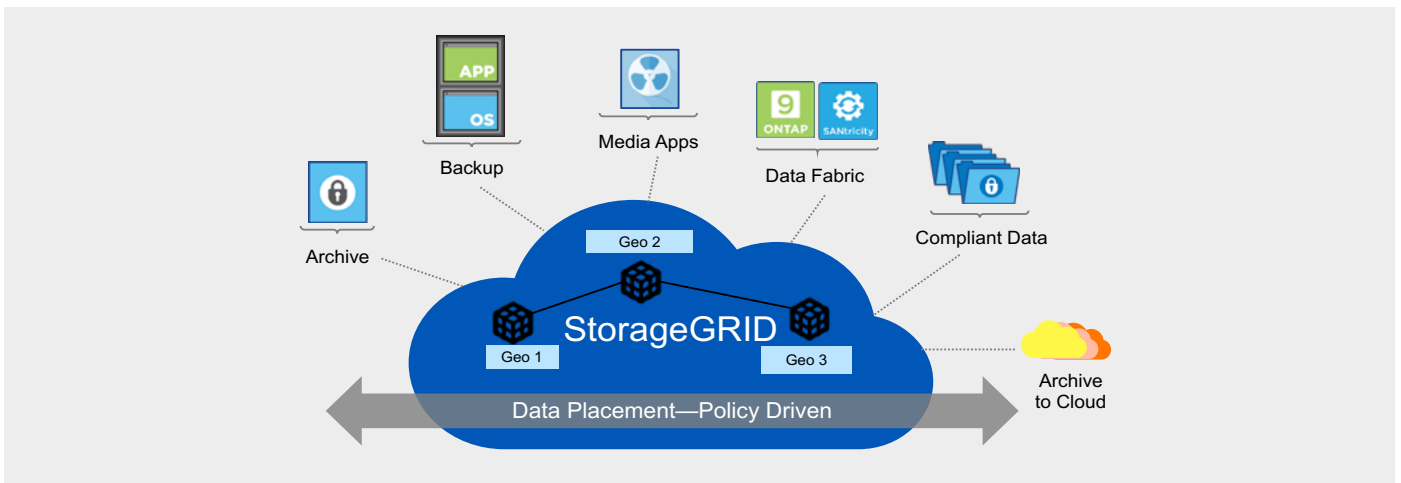


Figure 1: StorageGRID policies can be administered from any location to manage data on a local, regional, and global basis. Users and applications see StorageGRID as a single namespace.

sophisticated policy-driven functionality to support your MAM systems. Policies can be enforced at ingest, at rest, on read, after object attribute update, at object disposition request, or when information lifecycle management (ILM) policies change. StorageGRID policies can initiate object replication, erasure coding, tiering, and archiving to the cloud.

This storage management, monitoring, and migration functionality delivers an improved experience for all the workflows in your media enterprise. MAM 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.

By erasure coding, moving, or copying objects to sites to meet production and distribution access demands, StorageGRID relieves your MAM 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.

“We tested a number of object storage type products – the performance, ease of use, management capability within StorageGRID met all of our needs.”

Skottie Miller
Technology Fellow and VP of Platform and Services
Architecture, Dreamwoks Animation

With object store support, MAM 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.”

Leverage data protection methods

Erasure coding is a method of data protection in which data is broken into fragments and encoded by mathematically combining the fragments into a greater number of data pieces composed of multiple copies of the original fragments. It is then dispersed across different locations or storage media. Erasure coding offers better overall storage efficiency than replication, with equivalent durability and availability. In an object store with geo-dispersed erasure coding, data can be retrieved from a subset of the data pieces.

With ILM policies, every object can use a different erasure coding scheme. StorageGRID offers a variety of erasure coding schemes that you can choose or change to at any time.

StorageGRID appliances leverage an additional layer of erasure coding. The appliance uses Dynamic Disk Pools (DDP) technology to provide drive failure resilience at the node level. This 2-level erasure coding allows object access even during temporary or permanent failure of one or more drives, nodes, or sites. It also reduces network consumption during rebuilds, giving StorageGRID up to 9 9s of availability and 15 9s of durability.

Minimize the cost of cradle-to-grave data management

Achieving ubiquitous content access often requires storing copies of the same asset in multiple places. This approach increases complexity, and








































Cloud Apps	Media Workflow	Backup, Archive, EFSS
    	                	                

Figure 2: Growing StorageGRID Media ecosystem. Media asset management systems, production, distribution, backup, and archival applications access StorageGRID natively by using the Amazon S3 interface.

it also requires additional storage capacity. With StorageGRID, you can minimize the capacity that is needed while optimizing for data resilience, production access, and distribution requirements through policydriven 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 a set period 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 tap into hybrid cloud by moving that content to less expensive public cloud storage. Management and maintenance of file copies, migrations, and deletions are driven by the StorageGRID ILM policy engine.

Meet the demands of multiple media use cases StorageGRID is built to fit into a multiworkload, multitenant architecture. Whether it's using S3 as a backup target, active archiving from your MAM, or capacity tiering via FabricPool or other tiering solutions, StorageGRID enables you to painlessly run multiple applications. In fact, StorageGRID offers robust multitenant features for security, access control, and traffic limiting, allowing everything from financial, DevOps, and media archives to coexist in the same

namespace while maintaining secure, authorized access.

As your use cases scale, capacity can be easily allocated and put into service immediately. Maintenance and expansion procedures can be done without interruption or downtime, across generations of StorageGRID nodes.

The latest StorageGRID nodes are built with object as primary storage in mind. This means high-performing ingest and egress of your media files in seconds, not minutes.

The content in your repository will live longer than the media that it is stored on, and StorageGRID can 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.

The list of media tools that implement object storage capability is large and growing. With deployment in the most demanding industries since 2002, StorageGRID is the system you can trust with your growing media repository.

To learn more about StorageGRID, go to <https://www.netapp.com/data-storage/storagegrid>

Key StorageGRID Features	Description
Scalability and flexibility	<ul style="list-style-type: none"> • Multiple geo-distributed sites • Support for multiple storage tiers • Geo-erasure coding and geo-replication • Deploy on VMs, hardware appliances, or bare-metal servers • 640PB of raw capacity per grid • 200 billion objects
Application interfaces	<ul style="list-style-type: none"> • Integrated load balancing • Object access protocols: Amazon S3 and OpenStack Swift • NFS and CIFS access via NetApp ONTAP® • APIs for grid and tenant management
Deployment options	<ul style="list-style-type: none"> • Physical or virtual servers via Docker containers • Virtual appliance • Hardware appliances: <ul style="list-style-type: none"> - NetApp StorageGRID SGF6024 for high-performance primary object storage workloads - NetApp StorageGRID SG6060 for large-scale capacity with expansion shelf options - NetApp StorageGRID SG5712 and SG5760 for secondary capacity object storage workloads - NetApp StorageGRID SG100 and SG1000 services appliances for simplified operations include admin node software and load balancing
Service-level object and performance monitoring	<ul style="list-style-type: none"> • Comprehensive system monitoring and alerts • Adjust data policies with flexible ILM rules • QoS rate limiting to manage workload performance
Management	<ul style="list-style-type: none"> • Centralized management for automated installation and expansions • Rolling upgrades without downtime • Node cloning for easy hardware refreshes • Event-based audit messages
Data services	<ul style="list-style-type: none"> • Platform services for hybrid cloud integration • Write once, read many (WORM) retention with Object Lock • Advanced security and encryption capabilities

