



Datasheet

NetApp StorageGRID Object Storage Software

Software-defined object storage for web applications, rich content, data lakes, and backup and archival workloads

Key Benefits

Deliver Distributed Data Repositories for Objects and Files

Deliver easy-to-use, massively scalable, globally distributed unstructured data repositories for object (Amazon S3, OpenStack Swift) and file (SMB, NFS) applications.

Enable Global Data Governance

Protect data and comply with regulations with geo-distributed replication and layered erasure coding, write once, ready many (WORM) retention, access control policies, encryption, and audit trails.

Enable Hybrid Cloud Data Pipelines

Empower users to leverage the public cloud to process, transform, and analyze objects and metadata while securing your data in an on-premises object store.

Data at the Right Place, at the Right Time, with Policy-Driven Information Lifecycle Management

Optimize data availability, performance, geo-distribution, retention, protection, and storage cost with metadata-driven policies, and adjust them dynamically as the business value of data evolves.

The Challenge

The Internet of Things (IoT) propels unprecedented growth in unstructured data, presenting a massive opportunity to the enterprise with the potential of uncovering new customer engagements and revenue streams. But IT must overcome the challenge of keeping up with not only the volume of data, but also with changes in how data is stored and accessed. Users want IT to support both traditional and cloud-based applications, and to access data across many locations: inside the data centers, remote offices, and the public cloud. Consequently, IT organizations in data-intensive enterprises must reevaluate how to manage an increasingly decentralized and heterogeneous data and application landscape.

Application developers today expect to consume storage as a service. They want to connect to an endpoint and have their data available in a global namespace, as they would in the public cloud. Developers also want the freedom to leverage the elastic computing resources and rich data services APIs from the public cloud, such as on-demand computing, media services, analytics, machine learning, and serverless computing. IT must guarantee data integrity, security, sovereignty, and compliance while balancing user demands and public cloud efficiencies. In many cases, business and compliance requirements mandate that data must outlive the underlying storage infrastructure, sometimes by many generations.

To manage the growth of unstructured data and satisfy requirements for durability, availability, and performance, all while containing costs, many IT organizations have turned to cloud-based data management software such as object storage. However, new questions have arisen: What happens if requirements change? Can customers respond with agility when certain data suddenly becomes valuable? Does choosing one solution create vendor lock-in? How can customers maintain the flexibility to use both on-premises and public cloud solutions while maintaining control of their data?

The Solution

NetApp® StorageGRID® is a software-defined object-based storage solution that supports industry-standard object APIs such as Amazon Simple Storage Service (S3) API and OpenStack Swift API. You can create multiple service levels with metadata-driven object lifecycle policies, optimizing durability, protection, performance, and location across multiple geographies. You can adjust your policies and realign your data landscape as requirements change.

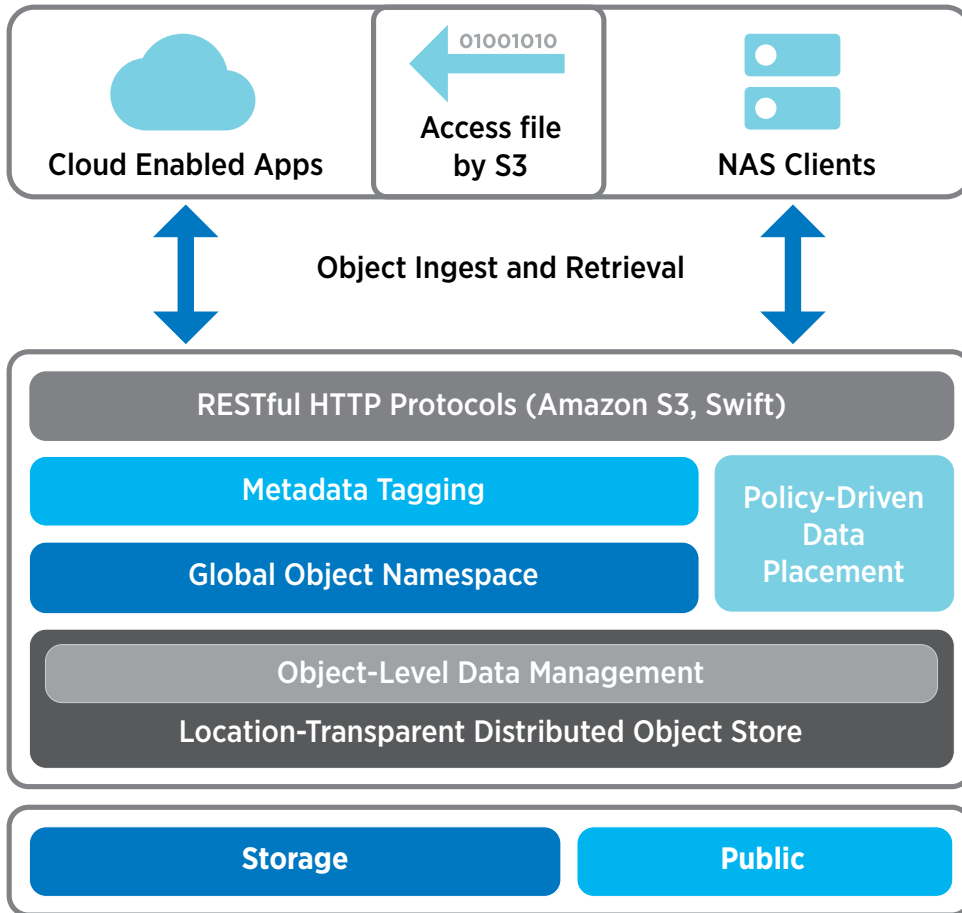


Figure 1) StorageGRID object storage software offers massive scalability while providing policy-driven data management to meet your organization's requirements.

Deploying NetApp StorageGRID appliances creates an enterprise-grade turnkey solution that is easy to implement. Customers can also deploy software-only StorageGRID nodes as containers on physical or virtual servers, leveraging heterogeneous storage underneath. Customers can use any combination of node types and rapidly deploy petabytes of storage.

Many unstructured data applications require NAS protocols. The StorageGRID NAS protocol bridge supports SMB and NFS access and at the same time enables object access to these files by using the Amazon S3 API. You can support your current workload while proactively supporting next-generation applications that natively support object protocols.

Improve Efficiency, Durability, and Flexibility

Reduce costs without sacrificing durability with StorageGRID layered erasure coding. Protect against failed disk drives and rapidly rebuild lost data segments with node-level coding and protect against site-level disasters with geo-distributed coding. You can combine replication and geo-distributed coding to balance performance needs and cost savings between different sets of data or during an object's lifecycle.

Enable the Hybrid Cloud

StorageGRID offers industry-leading hybrid cloud integration with user-controlled platform services. Storage tenants can configure mirroring of select objects to an S3-compatible public cloud. Trigger hybrid cloud workflows by integrating S3 notification of events in your on-premises buckets with Amazon Simple Notification Service (SNS). Drive further value with metadata search and analytics by streaming object metadata to an external Elasticsearch service, on the premises or in the public cloud.

Achieve new levels of cost savings by enabling cloud-to-cloud data management. StorageGRID can manage and store objects in its own globally distributed infrastructure, and also in Amazon S3 or S3-compatible object store or public cloud. You can add Amazon S3 storage as a storage tier, increasing data protection with an external cloud.

StorageGRID provides industry-leading Amazon S3 API compatibility with advanced features, such as object versioning, multi-part upload, Amazon Identity and Access Management-styled access policies, cross-origin resource sharing, and object tags. With Active Directory and LDAP identity federation for Amazon

S3 and OpenStack Swift users, and OpenStack Keystone authentication for Swift users, StorageGRID helps you bridge the gap between enterprise IT and cloud semantics.

Facilitate Compliance with Tamper-Proof Data Retention

StorageGRID offers many features to help you meet your regulatory obligations. Storage tenants can configure WORM retention and litigation hold for objects by buckets. You can configure StorageGRID so that compliance data is stored with duplicate copies or logical equivalents, such as erasure-coded objects. Secure your data with software-based encryption and built-in audit trails.

Designed for Always-On Operations

StorageGRID provides the foundation for global data availability anytime, anywhere, to facilitate nonstop operations. Configurations can be designed for resilience to one or multiple simultaneous failures, and even for resilience to entire site losses and regional disasters. StorageGRID is suitable for single data centers or multiple data center deployments with many sites around the globe.

StorageGRID is built on a modular architecture, so you can design grids to maximize throughput and capacity. A centralized process that manages installation maintains configuration control and speeds deployments. Storage nodes can be added to and removed from the grid without disruption.

Rely on Proven Software

Object stores must provide a solution for massive scale and long-term retention. With the proven track record of StorageGRID software and NetApp storage, you can be confident that you are building on a rock-solid foundation for your web data repositories, data archives, and media repositories.

StorageGRID is an 11th-generation object store with more than 15 years of production deployments in the most demanding industries. NetApp dependability has been demonstrated with more than 1 million systems shipped and more than 20 years of product hardening. With advanced features such as the NetApp Active IQ® monitoring tool for proactive, immediate response and with backing by NetApp's world-class support organization, StorageGRID is a solution that you can trust with your critical data assets.

Reduce Complexity

Software-defined storage gives you the choice of deploying StorageGRID nodes as virtual machines, as optimized hardware-based appliances, as bare-metal servers with Docker containers, or as a combination. In all cases, designing, deploying, and managing StorageGRID is a centrally managed and streamlined process.

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. #DataDriven

KEY FEATURES FOR OBJECT STORAGE

NETAPP STORAGEGRID PROVIDES

Massive scalability and flexible infrastructure

- Massive elastic content store
- Multiple geo-distributed sites
- Support for multiple storage tiers:
 - SSD, SAS, SATA, tape
 - Amazon S3
- Geo-erasure coding and geo-replication
- Deployment on VMs, hardware appliances, or bare-metal servers with Docker containers

Application interfaces

- Massively parallel transaction engine with:
- Integrated load balancing
 - Transaction multithread pipelining
- Object access protocols:
- Amazon S3 and OpenStack Swift
- NAS access:
- CIFS and NFS
 - File object duality
- System and account management:
- Management API: system installation, system administration, tenant management, maintenance tasks, and system monitoring including Prometheus
 - Tenant API: management of users, credentials, usage, and quotas

Data services

- Platform services – tenant configurable hybrid cloud integration:
- S3 event notification with Amazon Simple Notification Service
 - Cloud Mirror bucket replication with Amazon S3 or S3-compatible target
 - Metadata search and analysis with streaming metadata to external Elasticsearch
- Write once, ready many (WORM) retention:
- Reinforced data integrity with compliance-grade WORM
 - Litigation hold
- Advanced security and encryption capabilities:
- Store objects with lossless compression
 - Transport Security Layer (TSL) 1.2 and AES 256-bit encryption
 - Secure Hash Algorithm 2 (SHA-2) and CPU-efficient integrity protection

Metadata and content awareness

- Metadata-based data management:
- Content-aware self-healing maintains data protection even during network disruptions
 - Policies can be modified and applied retroactively to existing objects

Deployment options

- Physical or virtual servers via Docker containers
- Virtual appliance:
 - VMware ESXi and vCenter
- Hardware:
 - NetApp StorageGRID SG6060 for high-performance object storage workloads
 - NetApp StorageGRID SG5712 and SG5760 for capacity object storage workloads

Service-level objective and performance monitoring

- Get comprehensive performance feeds:
 - Access throughout
 - Replication throughout
 - Time to first byte
 - Time to policies achieved
- Get support for synthetic transactions
- Demonstrate SLAs
- Measure transaction round-trip time
- Separate WAN, storage, gateway times
- Advanced system monitoring via Prometheus

Management and monitoring

- Centralized and automatable installation and expansions
- Automated monitoring and tenant management through an API
- Rolling upgrades without downtime
- Comprehensive ad-hoc real-time, rolling-period, and historical-usage query capability
- 200+ predefined monitoring, usage, and performance reports
- Event-based audit messages for performance tracing, usage monitoring, and enabling billing or chargeback