



Datasheet

NetApp StorageGRID Webscale Object Storage Software

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

Key Benefits

Is Built for Web-Scale Data Repositories

Build massively scalable, globally distributed object stores that support industry-standard object APIs such as Amazon S3 and Swift.

Optimize Performance, Durability, and Cost

Protect data with layered erasure coding, which combines node-level and geo-distributed erasure coding to efficiently prevent data loss from disk, node, rack, or site outages.

Enables a Hybrid Cloud with Best-in-Class Flexibility

Seamlessly tier and replicate to public cloud storage while taking advantage of public cloud compute.

Uses a Metadata-Driven Policy Engine

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 not only propels massive data growth in unstructured data, but it also changes how data is stored and accessed. As data is created and consumed across many sites, versus a more traditional centralized data center, IT departments must reevaluate how to manage large amounts of data that is spread over many locations.

Modern application developers expect storage as a service — they no longer manage their own storage, but now expect a global name space. To take advantage of the benefits of serverless computing, organizations are looking for cost effective ways to combine on-premises security with public cloud efficiencies.

Users demand 24/7 access from any location and device; at the same time, IT must be able to guarantee the integrity and security of the data. In many cases, business and compliance requirements mandate that this data outlive the underlying storage infrastructure, sometimes by many generations.

To be able to store this data and meet requirements for durability, availability, and performance, all while containing costs, many IT organizations have turned to cloud-based software such as object storage. However, new questions have arisen: What happens if requirements change? Can customers dynamically reevaluate existing data storage policies? By choosing one solution, is vendor lock-in created? How can customers maintain the flexibility to use both on-premises and public cloud solutions while maintaining control?

The Solution

NetApp® StorageGRID® Webscale is a software-defined object-based storage solution that provides intelligent policy-driven data management. The ability to manage data while optimizing durability, protection, performance, and even physical placement across multiple geographies is key to meeting business requirements while reducing costs.

Deploying NetApp StorageGRID Webscale with a NetApp engineered appliance creates an enterprise-grade turnkey object storage appliance that is easy to implement. StorageGRID Webscale deploys nodes via Docker containers across bare metal, virtual machines, or an enterprise engineered hardware appliance. Customers can use any combination of node types and rapidly deploy PBs of storage.

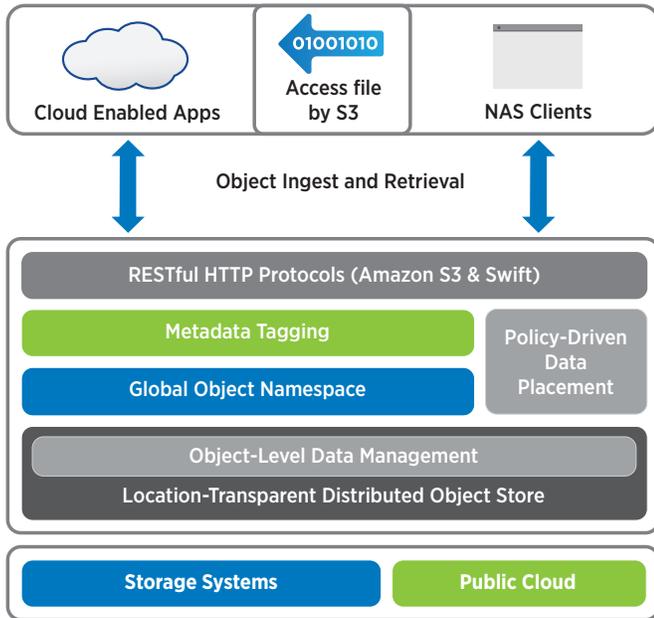


Figure 1) The NetApp StorageGRID Webscale object storage software offers massive scalability while providing policy-driven data management to meet your organization's requirements.

Many critical workloads 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 Simple Storage Service (Amazon S3) protocol. 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 NetApp StorageGRID Webscale layered erasure coding. By protecting at the node level and with geo-distributed coding, you can create policy-driven data protection with multiple levels of granularity. You can choose a combination of full copies and geo-distributed erasure-coded copies to balance performance needs and cost savings between different sets of data or during the object's lifecycle.

Enable the Hybrid Cloud

Achieve new levels of cost savings by enabling cloud-to-cloud data management. StorageGRID Webscale can manage and store objects not only within its own globally distributed infrastructure, but also in Amazon Web Services (AWS). You can add Amazon S3 storage as a storage tier, increasing data protection with an external cloud while reducing costs by performing more expensive operations against locally managed copies.

StorageGRID Webscale 11.0 introduces new Platform Services. Storage tenants can configure replication of buckets to S3 public cloud. Take advantage of public cloud compute and resource via notifications. Drive more value from your data with metadata search. Tenants can configure granular policies to index metadata via on-prem or AWS Elasticsearch.

StorageGRID Webscale provides industry-leading Amazon S3 API compatibility with advanced Amazon S3 features, including object versioning, multipart upload, and AWS Identity and Access Management-styled access policies. With Active Directory and LDAP identity federation for Amazon S3 and Swift users, and for administrative users, StorageGRID Webscale helps you bridge the gap between enterprise IT and cloud semantics.

Designed for Always-On Operations

NetApp StorageGRID Webscale 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 Webscale is suitable for single data centers or multi-data center deployments with many sites across the globe.

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

Reduce Complexity

Software-defined storage gives you the choice of deploying StorageGRID Webscale 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 Webscale 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 WEBSCALE 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 Swift

NAS access:

- CIFS and NFS
- File object duality

Grid\system management:

- Management API: administration of tenants, system tasks, and monitoring including Prometheus
- Tenant API: management of users, credentials, usage, and quotas

Compression and encryption

Platform Services – Tenant configurable hybrid cloud

- Simple Notification Services
- Elasticsearch
- Cloud Mirror bucket replication

Advanced security and encryption capabilities:

- Store objects with lossless compression
- Get support for AES-256 and SHA-256 encryption
- Get mixed-mode AES-256 and SHA-256 support for strong encryption 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

- Bare-metal\Software only servers via Docker containers
- VMs:
 - VMware ESXi and vCenter
 - Other hypervisor supported via SW only VMs
- Hardware appliance:
 - NetApp StorageGRID Webscale SG5660
 - NetApp StorageGRID Webscale SG5612

Service-level objective and performance monitoring

- Get comprehensive performance feeds:
 - Access throughput
 - Replication throughput
 - 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, enabling billing, or chargeback

