Datasheet

NetApp StorageGRID Webscale Object Storage Software

Software-defined object storage for web applications and rich content

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.

Supports NAS Protocols
Support SMB and NFS clients by using a native protocol bridge with the ability to access ingested files through Amazon S3 for future-proof data access.

Balances Performance, Durability, and Cost
Protect data with hierarchical 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
Leverage combinations of engineered appliances and software-only nodes on virtual machines and heterogeneous storage. Seamlessly tier between on-premises and public cloud storage with policy-based data movement.

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, as opposed to a more traditional centralized data center, IT departments must reevaluate how to manage large amounts of data that is spread over several locations.

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, in some cases 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 platform 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 NetApp E-Series storage creates an enterprise-grade turnkey object storage appliance that is easy to deploy. Customers have the option to deploy Webscale nodes in any combination of virtual machines (VMs) on VMware or on OpenStack or as physical appliances.

NetApp®

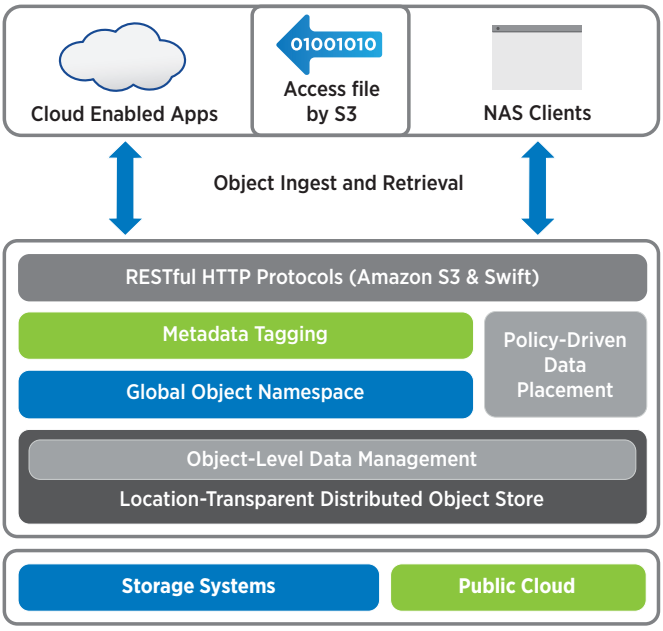

Figure 1) The NetApp StorageGRID Webscale object storage software offers massive scalability while providing policy-driven data management to meet customer requirements.

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. Customers can support their current workload while proactively supporting next-generation applications that natively support object protocols.

Efficient, Durable, and Flexible

Reduce costs without sacrificing durability with NetApp StorageGRID Webscale hierarchical erasure coding. By protecting at the node level and with geo-distributed coding, customers can create policy-driven data protection with multiple levels of granularity. They can choose a combination of full copies and/or geo-distributed 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). Customers 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 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, as well as administrative users, StorageGRID Webscale helps customers 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, even 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, allowing customers to design grids to maximize throughput and capacity. A centralized process that maintains configuration control and speeds deployments manages installation. Nodes, storage, and even entire sites 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 E-Series storage, customers can be confident that they are building on a rock-solid foundation for their web data repositories, data archives, and media repositories.

StorageGRID is a 10th-generation object store with 15 years of production deployments in the most demanding industries. E-Series dependability has been demonstrated with nearly 1 million systems shipped and 20 years of product hardening. By using advanced features such as the NetApp AutoSupport® tool for proactive, immediate response and backed by NetApp's world-class support organization, StorageGRID Webscale is a platform that customers trust with their most vital data assets.

Reduce Complexity

Software-defined storage gives customers the choice of deploying StorageGRID Webscale nodes as VMs, as optimized hardware-based appliances, or as a combination of both. In all cases, designing, deploying, and managing Webscale are a centrally managed and streamlined process.

About NetApp

Leading organizations worldwide count on NetApp for software, systems and services to manage and store their data. Customers value our teamwork, expertise and passion for helping them succeed now and into the future.
www.netapp.com

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 or hardware appliance
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 API Services • Management API: administration of tenants and system tasks
Compression and encryption	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	VMs: • VMware ESXi and vCenter • Kernel-Based Virtual Machine QEMU and OpenStack Hardware appliance: • NetApp StorageGRID Webscale SG5660 • NetApp StorageGRID Webscale SG5612
Service-level objective and performance monitoring	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, and gateway times
Management and monitoring	• Centralized 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