NetApp StorageGRID

Overview

NetApp StorageGRID is an object storage solution that complements NetApp file and block systems including ONTAP FAS, ASA, and AFF, and E-Series. StorageGRID delivers a distributed content repository using object storage software with metadata-controlled data management executing on servers in virtual machines or containers and block storage attached to the physical servers or as an appliance. The metadata is stored in a distributed NoSQL database with distributed access. NetApp StorageGRID is capable of storing up to 640 PB of data and 300 billion objects in a single namespace. StorageGRID also provides migration of objects to tape and supports 200 PB of data on tape. StorageGRID supports third party storage for the software-defined distribution or NetApp appliances. Also offered is a software only version of StorageGRID as a VM for VMware, as a Docker container for bare metal installations, or for deploying in a client provided VM on a different hypervisor than VMware.

The origin of NetApp StorageGRID comes from the acquisition of Bycast, which had been a provider of federated systems with successes in the healthcare records storage and management area. NetApp has created a scalable object storage solution that uses an ILM engine to manage data across geographies or in a single data center environment.

StorageGRID is a mature object storage solution with a number of advanced features that have been added over time. Advanced features include the ability to make replicated copies at 16 locations, cross-namespace replication, geographic dispersion to multiple locations, cloud tiering to AWS, GCP, and Azure, and support as a target for FabricPool from NetApp ONTAP systems.

Highlights

- Scaling of nodes and storage across geographies
- Advanced metadata-based data management
- Data stored as objects
- S3 object access
- S3 Select native implementation
- ISV integration for vertical markets
- VMware and purpose-built appliances support
- S3 compatible versioning support for objects
- Docker container for bare metal installation
- Target for FabricPool
- Cloud tiering to AWS, Azure, and Google clouds
- Data encryption with KMIP external key management
- Object Lock immutability

Usage and Deployment

NetApp StorageGRID is used as a content repository by cloud service providers, in vertical market industries, and in large IT enterprises. Objects are supported with S3 over HTTP/REST, a custom StorageGRID RESTful API. The focus for StorageGRID is the metadata management and usage through ISV partners. Experience and credibility in the healthcare market will make that a continued opportunity for StorageGRID.

- Characteristics
  - Scale – StorageGRID can scale to hundreds of nodes and supports 640 PB of capacity and 300 billion objects. Exceeding the "soft limits" requires approval from NetApp.
  - Protection / Durability / Resiliency – Multisite distribution of data across nodes provides protection from data element or node failure. The system can be configured to tolerate multiple element failures. Individual device protection uses storage-system node-level erasure coding. Node protection is accomplished by replication of data at up to 16 locations and/or with geo-distributed erasure coding – providing a two-level erasure code protection. Grid federation is also supported with active/active cross namespace replication and account clone. Versioning is supported.
  - Index and Search – Embedded metadata indexing is included and searching using third party search tools is supported.
  - Performance – Performance of StorageGRID is based on distribution and parallel access to data. The system provides performance reporting including access throughput, response time, and replication throughput. Read caching acceleration was enhanced as part of the performance improvement. SSDs are used to improve performance alongside HDDs in the SG6060 model. The SGF6024 and SGF6112 models use all SSDs.
  - Access Methods – S3 for API access to objects over HTTP/REST is included as well as a StorageGRID custom RESTful API for management. S3 Select is supported with a native implementation and internal execution engine. S3 Object Lock is also supported with both compliance and governance modes.
  - Geographic Access – Distribution of data across geographic areas is a basic capability of StorageGRID.
  - Security and Compliance – Multi-tenancy is supported with multiple administrative and user access roles. Encryption is included using AES-256 algorithms and support for KMIP external key management. Audit trails are enabled with audit repositories. Compliance capabilities for retention controls and event-based notifications are also included. Multi-factor authentication is implemented using SAML. Object Lock functionality provides immutability for compliance and ransomware protection.

  - Metadata – System and user metadata for objects are maintained in a separate database and used for data management including retention controls. Management of information in the repository is based on the metadata.
  - Integrity and Verification – An automatic integrity check is performed using the digital fingerprint of the object on every access and as a background process. Administrators may also invoke a verification of integrity.
  - Longevity of Object Data – Bulk data movement is policy controlled allowing for migration of objects to transfer to new technologies. Nodes can have data drained and moved or redistributed automatically.
  - Billing and Chargeback – More than 200 reports are available that provide usage, QoS, and other information that is used for billing and chargeback. ISV's provide additional reporting based on integration with StorageGRID.
- Applications
  - Content repositories such as healthcare records with multi-tenancy and compliance requirements, media and entertainment video, oil & gas data, and for other large capacity vertical industry areas.
  - Service providers object storage including archive and backup.
  - Big data analytics storage
  - Enterprise IT as another tier of storage or for private/hybrid cloud environments
  - Target for data protection and tiering of data
- System environments
  - Any environment with web access – either through private cloud or cloud service providers
- Deployment and Administration
  - StorageGRID uses a browser-based GUI for installation and administration.
  - Built-in ILM functions including a policy engine, policy simulator and data movers
  - Management via NetApp BlueXP
  - Roles available for system administrator and tenant administrators.
  - Management API available for integration into applications by ISVs.
  - EC Rebalance redistributes data for node expansions
  - Deployment options:
    - As appliance on NetApp appliances
    - SG1000 and SG100 load balancer nodes
    - Software to run in a VM
    - Software in a Docker container for bare-metal installation

Key Capabilities

Architecture and Deployment

StorageGRID is software running on a set of nodes, which are servers with storage attached. Objects are broken into chunks and stored in StorageGRID nodes. StorageGRID uses a Cassandra database for storing and managing the chunks of data. Small objects (<200K) are replicated with two copies of data stored on different nodes. Larger objects can be replicated or erasure coded.

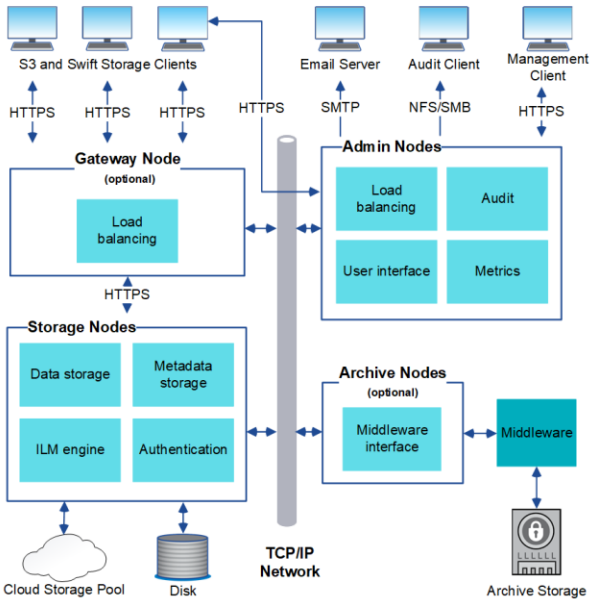

Figure 1: StorageGRID Architecture (Source: NetApp)

Grid Nodes

Grid Nodes are the foundation of the architecture of StorageGRID. Each node can be thought of as a server (or storage system with added software) that is executing StorageGRID software to perform specific functions. The individual Grid Nodes are:

Admin Nodes – as the name implies, Admin Nodes are used for the administrative functions of configuration, monitoring, and system logging.

Storage Nodes – store and manage the data for object and metadata. Protection of the data is also performed in the storage nodes.

Archive Nodes – object data can be archived to external storage such as cloud using S3 or to tape with the inclusion of IBM Spectrum Protect (TSM) software. Archive nodes are optional in the system.

API Gateway Nodes – the gateway nodes provide the load-balancing interface into StorageGRID. Gateway nodes are optional.

Deployment Options

StorageGRID may be deployed by a number of means, shown in the following diagram.

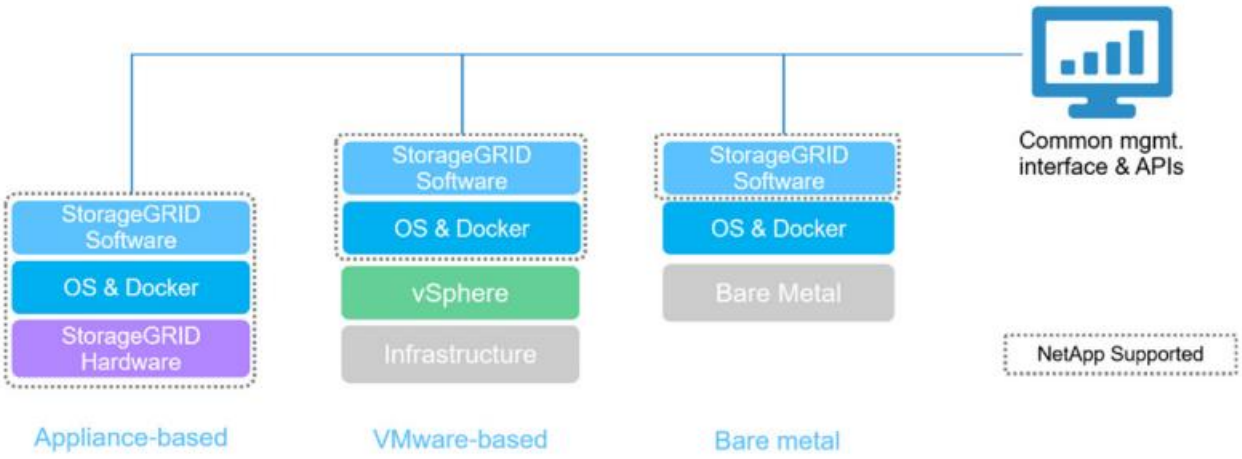

Figure 1: StorageGRID Deployment Options (Source: NetApp)

StorageGRID appliances include the SG5000 series, focused on cost optimization and secondary workloads, and the SG6000 series targeted at midrange to primary workloads. The majority of StorageGRID appliance uses an E-Series 60 or 12 drive storage system with one controller executing the SANtricity storage software and the other node executing the StorageGRID software in a non-HA configuration. NetApp has also released the SGF6112 appliance that is not based on E-Series hardware.
The SGF6112 model is an all-flash system that can house 12 NVMe SSDs in 1U. The use of all SSD flash devices, such as in the SGF6024 and SGF6112, provides an object storage system with a significant increase in performance and greater handling of small objects.

NetApp SG1000 and SG100 appliances are also available for deployment of StorageGRID gateway and admin nodes as physical appliances.

Software Architecture

The StorageGRID software executes on Debian Linux operating system on each node. Objects are transferred to StorageGRID and are broken into chunks and stored with metadata on individual nodes.

The services executing in the StorageGRID nodes are encapsulated as Docker containers, running on appliances, or running in virtual machines.

StorageGRID Services are the software that run on the Grid Nodes and each node may run a different set of services. The services, their function, and the nodes where they execute are described in the table below.

Service	Function	Executing Nodes
Administrative Domain Controller - ADC	Maintains the system topology and provides access authentication	Storage Node
Audit Management System - AMS	Tracks system activity and events	Admin Node
Archive – ARC	Manages the archiving function by moving data to cloud or tape through other services such as Spectrum Protect	Archive Node
Load Balancer Service	Does the load balancing function for access to StorageGRID routing service (LDR)	Admin Node API Gateway Node
Configuration Management Node – CMN	Manages the configuration of the system, connections, and the grid tasks	Admin Node (primary)
Distributed Data Storage – DDS	Monitors the Cassandra database and manages the object data placement.	Storage Node
Local Distribution Router – LDR	Manages the storing of object to the devices in a storage node	Storage Node
Network Management System – NMS	Monitors the system status and allows configuration of the system	Admin Node
Server Status Monitor – SSM	Hardware performance monitor, capturing metrics such as operating system and network	Admin Node Storage Node Archive Node API Gateway Node

Information Lifecycle Management (ILM)

ILM is the name for the policies applied to manage object data in StorageGRID. At the time an object is stored, a set of ILM policies are applied. Based on the matching of the policy to the object, called filtering in StorageGRID documentation, the rules are applied regarding data protection and storage location. Default policies can be established or bucket-specific policies assigned. The ILM rules establish:

1. Where the data is stored – which storage pool which is the grouping of Storage Nodes
2. Type of storage used (devices within a Storage Node or Archival storage)
3. The number and type of copies made – either replicated or erasure coded
4. Time-based management factors – location and protection change over time

Device and Node Failure Protection

A device or node failure uses either the multiple copies of data written (determined by the ILM policy) or the correcting erasure code depending on ILM setting for protection. There is no packing or containerization of small objects. Object less than 200K bytes in size are only eligible to be replicated, due to the overhead involved with erasure coding/

Replication

If the ILM rule for Content Placement is set to make replicated copies of object data, the LDR service of Storage Nodes will make copies of data. The DDS service controls the placement of the data and verifies the number of copies and location. The default ILM setting is to make two copies.

Erasure Coding Forward Error Correction

If the ILM rule for Content Placement is set to make erasure coded copies of object data, an erasure coding algorithm breaks the object into data and parity fragments and distributes the fragments across the Storage Nodes in the selected storage pool.

Remote Protection

Geographic dispersion is supported as well as asynchronous remote replication. If there is a temporary site outage where the remote connection or the entire site has failed, the writes become eventually consistent with a queue for access being restored.

StorageGRID 11.7 added additional support for grid-to-grid active/active replication

Versioning

S3 compatible versioning is supported. Objects can be updated without the prior version being deleted or overwritten with the new one. Multiple versions are supported with access allowed to the prior versions as a subset to the object ID. Version information is kept in the metadata.

Compliance

Compliance settings include WORM mode, audit trails of administrative actions, SSL settings for transfer of data, and encryption for data at rest.

S3 Object Lock

S3 Object Lock is a WORM mode that provides functionality equivalent to the Amazon S3 Object Lock feature. When using Object Lock, objects may not be modified or deleted during the retention period. A retain until date can be specified, or objects can be placed under a legal hold as needed. StorageGRID supports S3 Object Lock in both Compliance Mode and Governance Mode.

Data Reduction

Compression of data is an optional setting. Object data that is not already compressed (detected by a pre-scan of data) will be compressed on ingest.

Reports

Over 200 pre-defined reports are available and are in the form of text or as line, area, or state graphs. Additionally, data may be exported for use by other reporting software.

Cloud Storage Pools

Release 11.2 introduced Cloud Storage Pools which allow objects to be stored outside of StorageGRID with the use of ILM rules. Cloud Storage Pools can be used to tier objects to Amazon Glacier, Azure Blob, or Google Cloud in order to free up on-premises storage for higher performance workloads.

S3 Select

Release 11.6 introduced support for S3 Select with a native implementation. StorageGRID S3 Select supports CSV and Parquet formats. S3 Select requires the use of a StorageGRID Admin or Gateway node, either physical or software defined.

FabricPool

NetApp FabricPool is a functionality within ONTAP systems that provides automated data movement between tiers based on data access patterns. StorageGRID is supported as private cloud target for FabricPool, allowing cold data to be tiered off from ONTAP systems to StorageGRID.

BlueXP

NetApp BlueXP is a SaaS based management tool for NetApp on-premises and cloud solutions. Discovery and management of StorageGRID systems is supported with BlueXP. More information on BlueXP features and capabilities is available in Evaluator Group's coverage of NetApp BlueXP.

Significant Announcements

- May 2023 – StorageGRID 11.7 release
  - SGF6112 All Flash Appliance
  - Cross Grid Replication
  - Enhanced Security and Usability

Futurum Group EvaluScale – Object Storage

The Futurum Group product review methodology "EvaluScale" assesses each product within a specific technology area. The evaluation of each product is based on its capabilities, with capabilities for each technology segment grouped into distinct categories. The products are evaluated based on the following 4 criteria categories:

- Performance / Capacity
- Basic Functionality
- Advanced Capabilities
- Ability to Execute

The full Object Storage EvaluScale can be found here.

The Futurum Group Opinion and Outlook for NetApp StorageGRID

NetApp StorageGRID has been deployed in a number of vertical environments, including healthcare with integration from independent software vendors. The metadata management capabilities could be exploited by more ISVs to create high value-add offering for content repositories for specific verticals, all built on the StorageGRID foundation. The capabilities of the system will be very useful as a repository, not only vertical market industries but as a repository for enterprise IT environments.

The solution provides flexibility for IT organizations with several different deployment options, as well as a range of hardware appliances. NetApp also provides flexibility in purchasing with the option of a consumption-based pricing model, which is an advance for selling based on OpEx economics. This will be very valuable for many customers as they scale their environments.

The SGF6024 and SGF6112 all-flash models are another differentiator for NetApp StorageGRID. SSDs offer significant performance increases over traditional HDDs; however, they are not commonly found in object storage systems. All flash systems are especially beneficial for the handling of small objects and use of object as primary storage.

StorageGRID offers a mature object storage solution, with several features that have been added over time. StorageGRID is well suited as part of a NetApp-based hybrid cloud strategy due to its support as a target for FabricPool as well as support for common hybrid-cloud management in NetApp BlueXP.

While several features have been added, discontinuing of the NAS Bridge functionality for file access leaves customers with the need to also store files without that capability. Using ONTAP with the S3 protocol is one option but it requires a different solution than StorageGRID. While this S3 support has been added to ONTAP, StorageGRID remains a key part of NetApp's portfolio as the main object storage offering, and will likely remain so for the foreseeable future.