



Technical reports

StorageGRID solutions and resources

NetApp
December 12, 2025

This PDF was generated from <https://docs.netapp.com/us-en/storagegrid-enable/technical-reports/index.html> on December 12, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Technical reports	1
Introduction to StorageGRID technical reports	1
NetApp StorageGRID and big data analytics	1
NetApp StorageGRID use cases	1
Why StorageGRID for data lakes?	2
Benchmarking Data Warehouses and Lakehouses with S3 Object Storage: A Comparative Study	3
Hadoop S3A tuning	6
What is Hadoop?	6
Hadoop HDFS and S3A connector	6
Hadoop S3A connector tuning	7
TR-4871: Configure StorageGRID for backup and recovery with Commvault	12
Backup and recover data using StorageGRID and Commvault	12
Tested solution overview	14
StorageGRID sizing guidance	16
Run a data protection job	18
Review baseline performance tests	26
Bucket consistency level recommendation	27
TR-4626: Load balancers	28
Use third-party load balancers with StorageGRID	28
Use StorageGRID load balancers	29
Learn how to implement SSL certificates for HTTPS in StorageGRID	30
Configure trusted third-party load balancer in StorageGRID	31
Learn about local traffic manager load balancers	31
Learn about few use cases for StorageGRID configurations	34
Validate SSL connection in StorageGRID	37
Understand global load balancing requirements for StorageGRID	37
TR-4645: Security features	38
Secure StorageGRID data and metadata in an object store	38
Data access security features	40
Object and metadata security	48
Administration security features	50
Platform security features	54
Cloud integration	56
TR-4921: Ransomware defense	56
Protect StorageGRID S3 objects from ransomware	56
Ransomware defense using object lock	57
Ransomware defense using replicated bucket with versioning	60
Ransomware defense using versioning with protective IAM policy	63
Ransomware investigation and remediation	66
TR-4765: Monitor StorageGRID	67
Introduction to StorageGRID monitoring	67
Use the GMI dashboard to monitor StorageGRID	68
Use alerts to monitor StorageGRID	69

Advanced monitoring in StorageGRID	70
Access metrics using cURL in StorageGRID	73
View metrics using the Grafana dashboard in StorageGRID	74
Use traffic classification policies in StorageGRID	75
Use audit logs to monitor StorageGRID	78
Use the StorageGRID app for Splunk	78
TR-4882: Install a StorageGRID bare metal grid	78
Introduction to installing StorageGRID	78
Prerequisites to install StorageGRID	79
Install Docker for StorageGRID	89
Prepare node configuration files for StorageGRID	89
Install StorageGRID dependencies and packages	93
Validate the StorageGRID configuration files	93
Start the StorageGRID host service	95
Configure the Grid Manager in StorageGRID	95
Add StorageGRID license details	97
Add sites to StorageGRID	98
Specify grid network subnets for StorageGRID	99
Approve grid nodes for StorageGRID	100
Specify NTP Server details for StorageGRID	105
Specify DNS server details for StorageGRID	106
Specify the system passwords for StorageGRID	107
Review configuration and complete StorageGRID install	108
Upgrade bare-metal nodes in StorageGRID	110
TR-4907: Configure StorageGRID with veritas Enterprise Vault	111
Introduction to configuring StorageGRID for site failover	111
Configure StorageGRID and veritas Enterprise Vault	112
Configure StorageGRID S3 Object Lock for WORM storage	117
Configure StorageGRID site failover for disaster recovery	121

Technical reports

Introduction to StorageGRID technical reports

NetApp StorageGRID is a software-defined object storage suite that supports a wide range of use cases across public, private, and hybrid multicloud environments.

StorageGRID offers native support for the Amazon S3 API and delivers industry-leading innovations such as automated lifecycle management to store, secure, protect, and preserve unstructured data cost effectively over long periods.

StorageGRID provides documentation to cover best practices and recommendations for several StorageGRID features and integrations.

NetApp StorageGRID and big data analytics

NetApp StorageGRID use cases

NetApp StorageGRID object storage solution offers scalability, data availability, security, and high performance. Organizations of all sizes and across various industries use StorageGRID S3 for a wide range of use cases. Let's explore some typical scenarios:

Big data analytics: StorageGRID S3 is frequently used as a data lake, where businesses store large amounts of structured and unstructured data for analysis using tools like Apache Spark, Splunk Smartstore and Dremio.

Data Tiering: NetApp customers use ONTAP's FabricPool feature to automatically move data between a high-performance local tier to StorageGRID. Tiering frees up expensive flash storage for hot data while keeping cold data readily available on low-cost object storage. This maximizes performance and savings.


Data backup and disaster recovery: Businesses can use StorageGRID S3 as a reliable and cost-effective solution for backing up critical data and recovering it in case of a disaster.

Data storage for applications: StorageGRID S3 can be used as a storage backend for applications, enabling developers to easily store and retrieve files, images, videos, and other types of data.

Content delivery: StorageGRID S3 can be used to store and deliver static website content, media files, and software downloads to users around the world, leveraging StorageGRID's geo distribution and global namespace for fast and reliable content delivery.

Data Archive: StorageGRID offers different storage types and supports tiering to public long term low-cost storage options, making it an ideal solution for archiving and long-term retention of data that needs to be retained for compliance or historical purposes.

Object storage use cases




Among the above, big data analytics is one of the topmost use cases and its usage is trending upward.

Why StorageGRID for data lakes?

- Increased collaboration - Massive shared multi-site, multi-tenancy w/industry standard API access
- Decreased operational costs - Operational simplicity of a single, self-healing, automated scale-out architecture
- Scalability - Unlike traditional Hadoop and data warehouse solutions, StorageGRID S3 object storage decouples storage from compute and data, allowing business to scale their storage needs as they grew.
- Durability and reliability - StorageGRID provides 99.999999999% durability, meaning that data stored is highly resistant to data loss. It also offers high availability, ensuring that data is always accessible.
- Security - StorageGRID offers various security features, including encryption, access control policy, data lifecycle management, object lock and versioning to protect data stored in S3 buckets

StorageGRID S3 Data Lakes



NetApp

Benchmarking Data Warehouses and Lakehouses with S3 Object Storage: A Comparative Study

This article presents a comprehensive benchmark of various data warehouse and lakehouse ecosystems using NetApp StorageGRID. The goal is to determine which system performs best with S3 object storage. Refer to this

[Apache Iceberg: The Definitive Guide](#) to learn more about datawarehouse/lakehouse architectures and table format (Parquet and Iceberg).

- Benchmark Tool - TPC-DS - <https://www.tpc.org/tpcds/>
- Big data ecosystems
 - Cluster of VMs, each with 128G RAM and 24 vCPU, SSD storage for system disk
 - Hadoop 3.3.5 with Hive 3.1.3 (1 name node + 4 data nodes)
 - Delta Lake with Spark 3.2.0 (1 master + 4 workers) and Hadoop 3.3.5
 - Dremio v25.2 (1 coordinator + 5 executors)
 - Trino v438 (1 coordinator + 5 workers)
 - Starburst v453 (1 coordinator + 5 workers)
- Object storage
 - NetApp® StorageGRID® 11.8 with 3 x SG6060 + 1x SG1000 load balancer
 - Object protection - 2 copies (result is similar with EC 2+1)
- Database size 1000GB
- Cache was disabled across all ecosystems for each query test using the Parquet format. For the Iceberg format, we compared the number of S3 GET requests and total query time between cache-disabled and cache-enabled scenarios.

TPC-DS includes 99 complex SQL queries designed for benchmarking. We measured the total time taken to execute all 99 queries and conducted a detailed analysis by examining the type and number of S3 requests.

Our tests compared the efficiency of two popular table formats: Parquet and Iceberg.

TPC-DS query result with Parquet table format

Ecosystem	Hive	Delta Lake	Dremio	Trino	Starburst
TPCDS 99 queries total minutes	1084 ¹	55	36	32	28
S3 Requests breakdown					
GET	1,117,184	2,074,610	3,939,690	1,504,212	1,495,039
observation: all range GET	80% range get of 2KB to 2MB from 32MB objects, 50 - 100 requests/sec	73% range get below 100KB from 32MB objects, 1000 - 1400 requests/sec	90% 1M byte range get from 256MB objects, 2500 - 3000 requests/sec	Range GET size: 50% below 100KB, 16% around 1MB, 27% 2MB-9MB, 3500 - 4000 requests/sec	Range GET size: 50% below 100KB, 16% around 1MB, 27% 2MB-9MB, 4000 - 5000 request/sec
List objects	312,053	24,158	120	509	512
HEAD (non-existent object)	156,027	12,103	96	0	0
HEAD (existent object)	982,126	922,732	0	0	0
Total requests	2,567,390	3,033,603	3,939,906	1,504,721	1,499,551

¹ Hive unable to complete query number 72

TPC-DS query result with Iceberg table format

Ecosystem	Dremio	Trino	Starburst
TPCDS 99 queries total minutes (cache disabled)	22	28	22
TPCDS 99 queries total minutes ² (cache enabled)	16	28	21.5
S3 Requests breakdown			
GET (cache disabled)	1,985,922	938,639	931,582

Ecosystem	Dremio	Trino	Starburst
GET (cache enabled)	611,347	30,158	3,281
observation: all range GET	Range GET size: 67% 1MB, 15% 100KB, 10% 500KB, 3500 - 4500 requests/sec	Range GET size: 42% below 100KB, 17% around 1MB, 33% 2MB-9MB, 3500 - 4000 requests/sec	Range GET size: 43% below 100KB, 17% around 1MB, 33% 2MB-9MB, 4000 - 5000 requests/sec
List objects	1465	0	0
HEAD (non-existent object)	1464	0	0
HEAD (existent object)	3,702	509	509
Total requests (cache disabled)	1,992,553	939,148	932,071

² Trino/Starburst performance is bottlenecked by compute resources; adding more RAM to the cluster reduces the total query time.

As shown in the first table, Hive is significantly slower than other modern data lakehouse ecosystems. We observed that Hive sent a large number of S3 list-objects requests, which are typically slow on all object storage platforms, especially when dealing with buckets containing many objects. This significantly increases the overall query duration. Additionally, modern lakehouse ecosystems can send a high number of GET requests in parallel, ranging from 2,000 to 5,000 requests per second, compared to Hive's 50 to 100 requests per second. The standard filesystem mimicry by Hive and Hadoop S3A contributes to Hive's slowness when interacting with S3 object storage.

Using Hadoop (either on HDFS or S3 object storage) with Hive or Spark requires extensive knowledge of both Hadoop and Hive/Spark, as well as an understanding of how the settings from each service interact. Together, they have over 1,000 settings, many of which are interrelated and cannot be changed independently. Finding the optimal combination of settings and values requires a tremendous amount of time and effort.

Comparing the Parquet and Iceberg results, we notice that the table format is a major performance factor. The Iceberg table format is more efficient than the Parquet in terms of the number of S3 requests, with 35% to 50% fewer requests compared to the Parquet format.

The performance of Dremio, Trino, or Starburst is primarily driven by the computing power of the cluster. Although all three use the S3A connector for S3 object storage connection, they do not require Hadoop, and most of Hadoop's fs.s3a settings are not used by these systems. This simplifies performance tuning, eliminating the need to learn and test various Hadoop S3A settings.

From this benchmark result, we can conclude that big data analytic system optimized for S3-based workloads is a major performance factor. Modern lakehouses optimize query execution, efficiently utilize metadata, and provide seamless access to S3 data, resulting in better performance compared to Hive when working with S3 storage.

Refer to this [page](#) to configure Dremio S3 data source with StorageGRID.

Visit the links below to learn more about how StorageGRID and Dremio work together to provide a modern and efficient data lake infrastructure and how NetApp migrated from Hive + HDFS to Dremio + StorageGRID to dramatically enhance big data analytic efficiency.

- Boost performance for your big data with NetApp StorageGRID
- Modern, powerful, and efficient data lake infrastructure with StorageGRID and Dremio
- How NetApp is Redefining the Customer Experience with Product Analytics

Hadoop S3A tuning

By *Angela Cheng*

Hadoop S3A connector facilitates seamless interaction between Hadoop-based applications and S3 object storage. Tuning the Hadoop S3A Connector is essential to optimize performance when working with S3 object storage. Before we go into tuning details, let's have a basic understanding of Hadoop and its components.

What is Hadoop?

Hadoop is a powerful open-source framework designed to handle large-scale data processing and storage. It enables distributed storage and parallel processing across clusters of computers.

The three core components of Hadoop are:

- **Hadoop HDFS (Hadoop Distributed File System):** This handles storage, breaking data into blocks and distributing them across nodes.
- **Hadoop MapReduce:** Responsible for processing data by dividing tasks into smaller chunks and executing them in parallel.
- **Hadoop YARN (Yet Another Resource Negotiator):** [Manages resources and schedules tasks efficiently](#)

Hadoop HDFS and S3A connector

HDFS is a vital component of the Hadoop ecosystem, playing a critical role in efficient big data processing. HDFS enables reliable storage and management. It ensures parallel processing and optimized data storage, resulting in faster data access and analysis.

In big data processing, HDFS excels at providing fault-tolerant storage for large datasets. It achieves this through data replication. It can store and manage large volumes of structured and unstructured data in a data warehouse environment. Moreover, it seamlessly integrates with leading big data processing frameworks, such as Apache Spark, Hive, Pig, and Flink, enabling scalable and efficient data processing. It is compatible with Unix-based (Linux) operating systems, making it an ideal choice for organizations that prefer using Linux-based environments for their big data processing.

As the volume of data has grown over time, the approach of adding new machines to the Hadoop cluster with their own compute and storage has become inefficient. Scaling linearly creates challenges for using resources efficiently and managing the infrastructure.

To address these challenges, the Hadoop S3A connector offers high-performance I/O against S3 object storage. Implementing a Hadoop workflow with S3A helps you leverage object storage as a data repository and enables you to separate compute and storage, which in turn enables you to scale compute and storage independently. Decoupling compute and storage also enable you to dedicate the right amount of resources for your compute jobs and provide capacity based on the size of your data set. Therefore, you can reduce your overall TCO for Hadoop workflows.

Hadoop S3A connector tuning

S3 behaves differently from HDFS, and some attempts to preserve the appearance of a file system are aggressively suboptimal. Careful tuning/testing/experimenting is necessary to make the most efficient use of S3 resources.

Hadoop options in this document are based on Hadoop 3.3.5, refer to [Hadoop 3.3.5 core-site.xml](#) for all available options.

Note – the default value of some Hadoop fs.s3a settings are different in each Hadoop version. Be sure to check out the default value specific to your current Hadoop version. If these settings are not specified in Hadoop core-site.xml, default value will be used. You can override the value at run time using Spark or Hive configuration options.

You must go to this [Apache Hadoop page](#) to understand each fs.s3a options. If possible, test them in non-production Hadoop cluster to find the optimal values.

You should read [Maximizing Performance when working with the S3A Connector](#) for other tuning recommendations.

Let's explore some key considerations:

1. Data compression

Do not enable StorageGRID compression. Most of big data systems use byte range get instead of retrieving the entire object. Using byte range get with compressed objects degrade the GET performance significantly.

2. S3A committers

In general, magic s3a committer is recommended. Refer to this [common S3A committer options page](#) to get a better understanding of magic committer and its related s3a settings.

Magic Committer:

The Magic Committer specifically relies on S3Guard to offer consistent directory listings on the S3 object store.

With consistent S3 (which is now the case), the Magic Committer can be safely used with any S3 bucket.

Choice and Experimentation:

Depending on your use case, you can choose between the Staging Committer (which relies on a cluster HDFS filesystem) and the Magic Committer.

Experiment with both to determine which best suits your workload and requirements.

In summary, the S3A Committers provide a solution to the fundamental challenge of consistent, high-performance, and reliable output commitment to S3. Their internal design ensures efficient data transfer while maintaining data integrity.

Option	Meaning	Default (Hadoop 3.3.5)
fs.s3a.committer.name	Committer to create for output to S3A, one of: "file", "directory", "partitioned", "magic".	file
fs.s3a.buffer.dir	Local filesystem directory for data being written and/or staged.	\${env.LOCAL_DIRS:-\${hadoop.tmp.dir}}/s3a
fs.s3a.committer.magic.enabled	Enable "magic committer" support in the filesystem.	true
fs.s3a.committer.abort.pending/uploads	list and abort all pending uploads under the destination path when the job is committed or aborted.	true
fs.s3a.committer.threads	Number of threads in committers for parallel operations on files.	8
fs.s3a.committer.generate.uuid	Generate a Job UUID if none is passed down from Spark	false
fs.s3a.committer.require.uuid	Require the Job UUID to be passed down from Spark	false
mapreduce.fileoutputcommitter.mapsuccessfuljobs	Write a _SUCCESS file on the successful completion of the job.	true
mapreduce.outputcommitter.factory.scheme.s3a	The committer factory to use when writing data to S3A filesystems. If mapreduce.outputcommitter.factory.class is set, it will override this property. (This property is set in mapred-default.xml)	org.apache.hadoop.fs.s3a.commit.S3ACommitterFactory

3. Thread, connection pool sizes and block size

- Each **S3A** client interacting with a single bucket has its own dedicated pool of open HTTP 1.1 connections and threads for upload and copy operations.
- [You can tune these pool sizes to strike a balance between performance and memory/thread usage.](#)
- When uploading data to S3, it is divided into blocks. The default block size is 32 MB. You can customize this value by setting the fs.s3a.block.size property.
- Larger block sizes can improve performance for large data uploads by reducing the overhead of managing multipart parts during upload. Recommended value is 256 MB or above for large data set.

Option	Meaning	Default (Hadoop 3.3.5)
fs.s3a.threads.max	The total number of threads available in the filesystem for data uploads *or any other queued filesystem operation*.	64
fs.s3a.connection.maximum	Controls the maximum number of simultaneous connections to S3. This must be bigger than the value of fs.s3a.threads.max so as to stop threads being blocked waiting for new HTTPS connections. Why not equal? The AWS SDK transfer manager also uses these connections.	96
fs.s3a.max.total.tasks	The number of operations which can be queued for execution. This is in addition to the number of active threads in fs.s3a.threads.max.	32
fs.s3a.committer.threads	Number of threads in committers for parallel operations on files (upload, commit, abort, delete...)	8
fs.s3a.executor.capacity	The maximum number of submitted tasks which is a single operation (e.g. rename(), delete()) may submit simultaneously for execution -excluding the IO-heavy block uploads, whose capacity is set in "fs.s3a.fast.upload.active.blocks" All tasks are submitted to the shared thread pool whose size is set in "fs.s3a.threads.max"; the value of capacity should be less than that of the thread pool itself, as the goal is to stop a single operation from overloading that thread pool.	16
fs.s3a.fast.upload.active.blocks (see also related fs.s3a.fast.upload.buffer option)	Maximum Number of blocks a single output stream can have active (uploading, or queued to the central FileSystem instance's pool of queued operations. This stops a single stream overloading the shared thread pool.	4
fs.s3a.block.size	Block size to use when reading files using s3a: file system. A suffix from the set {K,M,G,T,P} may be used to scale the numeric value.	32MB (tested 1TB data set with 256MB and 512MB block size shows significant improvement in both read and write)

4. Multipart upload

s3a committers **always** use MPU (multipart upload) to upload data to s3 bucket. This is needed to allow for: task failure, speculative execution of tasks, and job aborts before commit. Here are some key specifications related to multipart uploads:

- Maximum object size: 5 TiB (terabytes).
- Maximum number of parts per upload: 10,000.
- Part numbers: Ranging from 1 to 10,000 (inclusive).
- Part size: Between 5 MiB and 5 GiB. Notably, there is no minimum size limit for the last part of your multipart upload.

Using a smaller part size for S3 multipart uploads has both advantages and disadvantages.

Advantages:

- Quick Recovery from Network Issues: When you upload smaller parts, the impact of restarting a failed upload due to a network error is minimized. If a part fails, you only need to re-upload that specific part rather than the entire object.

- Better Parallelization: More parts can be uploaded in parallel, taking advantage of multi-threading or concurrent connections. This parallelization enhances performance, especially when dealing with large files.

Disadvantage:

- Network overhead: Smaller part size means more parts to upload, each part requires its own HTTP request. More HTTP requests increase overhead of initiating and completing individual requests. Managing a large number of small parts can impact performance.
- Complexity: Managing the order, tracking parts, and ensuring successful uploads can be cumbersome. If the upload needs aborted, all the parts that already uploaded need to be tracked and purged.

For Hadoop, 256MB or above part size is recommended for `fs.s3a.multipart.size`. Always set the `fs.s3a.multipart.threshold` value to 2 x `fs.s3a.multipart.size` value. For example if `fs.s3a.multipart.size = 256M`, `fs.s3a.multipart.threshold` should be 512M.

Use larger part size for large data set. It is important to choose a part size that balances these factors based on your specific use case and network conditions.

A multipart upload is a [three-step process](#):

1. The upload is initiated, StorageGRID returns an upload-id.
2. The object parts are uploaded using the upload-id.
3. Once all the object parts are uploaded, sends complete multipart upload request with upload-id. StorageGRID constructs the object from the uploaded parts, and client can access the object.

If the complete multipart upload request isn't sent successfully, the parts stay in StorageGRID and will not create any object. This happens when jobs are interrupted, failed, or aborted. The parts remain in the Grid until multipart upload completes or is aborted or StorageGRID purges these parts if 15 days elapsed since upload was initiated. If there are many (few hundreds thousand to millions) in-progress multipart uploads in a bucket, when Hadoop sends 'list-multipart-uploads' (this request does not filter by upload id), the request may take a long time to complete or eventually time out. You may consider set `fs.s3a.multipart.purge` to true with an appropriate `fs.s3a.multipart.purge.age` value (e.g. 5 to 7 days, do not use default value of 86400 i.e. 1 day). Or engage NetApp support to investigate the situation.

Option	Meaning	Default (Hadoop 3.3.5)
fs.s3a.multipart.size	How big (in bytes) to split upload or copy operations up into. A suffix from the set {K,M,G,T,P} may be used to scale the numeric value.	64M
fs.s3a.multipart.threshold	How big (in bytes) to split upload or copy operations up into. This also controls the partition size in renamed files, as rename() involves copying the source file(s). A suffix from the set {K,M,G,T,P} may be used to scale the numeric value.	128M
fs.s3a.multipart.purge	True if you want to purge existing multipart uploads that may not have been completed/aborted correctly. The corresponding purge age is defined in fs.s3a.multipart.purge.age. If set, when the filesystem is instantiated then all outstanding uploads older than the purge age will be terminated -across the entire bucket. This will impact multipart uploads by other applications and users. so should be used sparingly, with an age value chosen to stop failed uploads, without breaking ongoing operations.	false
fs.s3a.multipart.purge.age	Minimum age in seconds of multipart uploads to purge on startup if "fs.s3a.multipart.purge" is true	86400

5. Buffer write data in memory

To enhance performance, you can buffer write data in memory before uploading it to S3. This can reduce the number of small writes and improve efficiency.

Option	Meaning	Default (Hadoop 3.3.5)
fs.s3a.fast.upload.buffer	The buffering mechanism to for data being written. Values: disk, array, bytebuffer. "disk" will use the directories listed in fs.s3a.buffer.dir as the location(s) to save data prior to being uploaded. "array" uses arrays in the JVM heap "bytebuffer" uses off-heap memory within the JVM. Both "array" and "bytebuffer" will consume memory in a single stream up to the number of blocks set by: fs.s3a.multipart.size * fs.s3a.fast.upload.active.blocks. If using either of these mechanisms, keep this value low The total number of threads performing work across all threads is set by fs.s3a.threads.max, with fs.s3a.max.total.tasks values setting the number of queued work items.	disk

Remember that S3 and HDFS work in distinct ways. Careful tuning/test/experiment is necessary to make the

most efficient use of S3 resources.

TR-4871: Configure StorageGRID for backup and recovery with Commvault

Backup and recover data using StorageGRID and Commvault

Commvault and NetApp have partnered to create a joint data protection solution combining Commvault Complete Backup and Recovery for NetApp software with NetApp StorageGRID software for cloud storage. Commvault Complete Backup and Recovery and NetApp StorageGRID provide unique, easy-to-use solutions that work together to help you meet demands of rapid data growth and increasing regulations around the world.

Many organizations want to migrate their storage to the cloud, scale their systems, and automate their policy for long-term retention of data. Cloud-based object storage is known for its resilience, ability to scale, and operational and cost efficiencies that make it a natural choice as a target for your backup. Commvault and NetApp jointly certified their combined solution in 2014 and since then have engineered deeper integration between their two solutions. Customers of all types around the world have adopted the Commvault Complete Backup and Recovery and StorageGRID combined solution.

About Commvault and StorageGRID

Commvault Complete Backup and Recovery software is an enterprise-level, integrated data and information management solution, built from the ground up on a single platform and with a unified code base. All of its functions share back-end technologies, bringing the unparalleled advantages and benefits of a fully integrated approach to protecting, managing, and accessing your data. The software contains modules to protect, archive, analyze, replicate, and search your data. The modules share a common set of back-end services and advanced capabilities that seamlessly interact with each other. The solution addresses all aspects of data management in your enterprise, while providing infinite scalability and unprecedented control of data and information.

NetApp StorageGRID as a Commvault cloud tier is an enterprise hybrid-cloud object-storage solution. You can deploy it across many sites, either on a purpose-built appliance or as a software-defined deployment. StorageGRID enables you to establish data management policies that determine how data is stored and protected. StorageGRID collects the information you need to develop and enforce policies. It examines a wide range of characteristics and needs, including performance, durability, availability, geographic location, longevity, and cost. Data is fully maintained and protected as it moves between locations and as it ages.

The StorageGRID intelligent policy engine helps you choose either of the following options:

- To use erasure coding to back up data across multiple sites for resilience.
- To copy objects to remote sites to minimize WAN latency and cost.

When StorageGRID stores an object, you access it as one object, regardless of where it is or how many copies exist. This behavior is crucial for disaster recovery, because with it, even if one backup copy of your data is corrupted, StorageGRID is able to restore your data.

Retaining backup data in your primary storage can be expensive. When you use NetApp StorageGRID, you free up space on your primary storage by migrating inactive backup data into StorageGRID while you benefit from the numerous capabilities of StorageGRID. The value of backup data changes over time, as does the cost of storing it. StorageGRID can minimize the cost of your primary storage while increasing the durability of your

data.

Key features

Key features of the Commvault software platform include:

- A complete data protection solution supporting all major operating systems, applications, and databases on virtual and physical servers, NAS systems, cloud-based infrastructures, and mobile devices.
- Simplified management through a single console: You can view, manage, and access all functions and all data and information across the enterprise.
- Multiple protection methods including data backup and archiving, snapshot management, data replication, and content indexing for e-discovery.
- Efficient storage management using deduplication for disk and cloud storage.
- Integration with NetApp storage arrays such as AFF, FAS, NetApp HCI, and E-Series arrays and NetApp SolidFire® scale-out storage systems. Integration also with NetApp Cloud Volumes ONTAP software to automate the creation of indexed, application-aware NetApp Snapshot™ copies across the NetApp storage portfolio.
- Complete virtual infrastructure management that supports leading on-premises virtual hypervisors and public cloud hyperscaler platforms.
- Advanced security capabilities to limit access to critical data, provide granular management capabilities, and provide single-sign-on access for Active Directory users.
- Policy-based data management that allows you to manage your data based on business needs—not physical location.
- A cutting-edge end-user experience, empowering your users to protect, find, and recover their own data.
- API-driven automation, allowing you to use third-party tools like vRealize Automation or Service Now to manage your data protection and recovery operations.

For details on supported workloads, visit [CommVault's supported technologies](#).

Backup options

When you implement Commvault Complete Backup and Recovery software with cloud storage, you have two backup options:

- Back up to a primary disk target and also back up an auxiliary copy to cloud storage.
- Back up to cloud storage as the primary target.

In the past, cloud or object storage was considered to be too low-performing to be used for primary backup. The use of a primary disk target allowed customers to have faster backup and restore processes and to keep an auxiliary copy on the cloud as a cold backup. StorageGRID represents the next generation of object storage. StorageGRID is capable of high performance and massive throughput as well as performance and flexibility beyond what other object-storage vendors offer.

The following table lists the benefits of each backup option with StorageGRID:

	Primary Backup to Disk and an Auxiliary Copy to StorageGRID	Primary Backup to StorageGRID
Performance	Fastest recovery time, using live mount or live recovery: best for Tier0/Tier1 workloads.	Cannot be used for live mount or live recovery operations. Ideal for streaming restore operation and for long-term retention.
Deployment architecture	Uses all flash or a spinning disk as a first backup landing tier. StorageGRID is used as a secondary tier.	Simplifies the deployment by using StorageGRID as the all-inclusive backup target.
Advanced features (live restore)	Supported	Not supported

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- StorageGRID 11.9 Documentation Center
<https://docs.netapp.com/us-en/storagegrid-119/>
- NetApp Product Documentation
<https://docs.netapp.com>
- Commvault documentation
<https://documentation.commvault.com/2024/essential/index.html>

Tested solution overview

The tested solution combines Commvault and NetApp solutions to make a powerful joint solution.

Solution setup

In the lab setup, the StorageGRID environment consisted of four NetApp StorageGRID SG5712 appliances, one virtual primary Admin node and one virtual Gateway node. The SG5712 appliance is the entry level option—a baseline configuration. Choosing higher performance appliance options such as the NetApp StorageGRID SG5760 or SG6060 can provide significant performance benefits. Consult your NetApp StorageGRID solution architect for sizing assistance.

For its data protection policy, StorageGRID uses an integrated lifecycle management (ILM) policy to manage and protect data. ILM rules are evaluated in a policy from top to bottom. We implemented the ILM policy shown in the following table:

ILM Rule	Qualifiers	Ingest Behavior
Erasure Coding 2+1	Objects over 200KB	Balanced
2 Copy	All objects	Dual Commit

The ILM 2 Copy rule is the default rule. The Erasure Coding 2+1 rule was applied for this testing to any object 200KB or larger. The default rule was applied to objects smaller than 200KB. Application of the rules in this way is a StorageGRID best practice.

For technical details about this test environment, read the Solution Design and Best Practices section in the [NetApp Scale-out Data Protection with Commvault](#) technical report.

StorageGRID hardware specifications

The following table describes the NetApp StorageGRID hardware used in this testing. The StorageGRID SG5712 appliance with 10Gbps networking is the entry-level option and represents a baseline configuration. Optionally the SG5712 can be configured for 25Gbps networking.

Hardware	Quantity	Disk	Usable Capacity	Network
StorageGRID SG5712 appliances	4	48 x 4TB (near-line SAS HDD)	136TB	10Gbps

Choosing higher-performance appliance options such as the NetApp StorageGRID SG5760, SG6060, or all flash SGF6112 appliances can provide significant performance benefits. Consult your NetApp StorageGRID solution architect for sizing assistance.

Commvault and StorageGRID software requirements

The following tables list the software requirements for the Commvault and NetApp StorageGRID software installed on VMware software for our testing. Four MediaAgent data transmission managers and one CommServe server were installed. In the test, 10Gbps networking was implemented for the VMware infrastructure. The following table

The following table lists Commvault software total system requirements:

Component	Quantity	Datastore	Size	Total	Total Required IOPS
CommServe Server	1	OS	500GB	500GB	n/a
		SQL	500GB	500GB	n/a
MediaAgent	4	Virtual CPU (vCPU)	16	64	n/a
		RAM	128GB	512	n/a
		OS	500GB	2TB	n/a
		Index Cache	2TB	8TB	200+
		DDB	2TB	8TB	200-80,000K

In the test environment, one virtual primary Admin node and one virtual Gateway node were deployed on VMware on a NetApp E-Series E2812 storage array. Each node was on a separate server with the minimum production environment requirements described in the following table:

The following table list requirements for StorageGRID virtual Admin nodes and Gateway nodes:

Node type	Quantity	vCPU	RAM	Storage
Gateway node	1	8	24GB	100GB LUN for the OS
Admin node	1	8	24GB	100GB LUN for the OS 200GB LUN for Admin node tables 200GB LUN for the Admin node audit log

StorageGRID sizing guidance

Consult your NetApp data protection specialists for specific sizing for your environment. NetApp data protection specialists can use the Commvault Total Backup Storage Calculator tool to estimate the backup infrastructure requirements. The tool requires Commvault Partner Portal access. Sign up for access, if needed.

Commvault sizing inputs

The following tasks can be used to perform discovery for sizing of the data protection solution:

- Identify the system or application/database workloads and corresponding front-end capacity (in terabytes [TB]) that will need to be protected.
- Identify the VM/file workload and similar front-end capacity (TB) that will need to be protected.
- Identify short-term and long-term retention requirements.
- Identify the daily % change rate for the datasets/workloads identified.
- Identify projected data growth over the next 12, 24, and 36 months.
- Define the RTO and RPO for data protection/recovery according to business needs.

When this information is available, the backup infrastructure sizing can be done resulting in a breakdown of required storage capacities.

StorageGRID sizing guidance

Before you perform NetApp StorageGRID sizing, consider these aspects of your workload:

- Usable capacity
- WORM mode

- Average object size
- Performance requirements
- ILM policy applied

The amount of usable capacity needs to accommodate the size of the backup workload you have tiered to StorageGRID and the retention schedule.

Will WORM mode be enabled or not? With WORM enabled in Commvault, this will configure object lock on StorageGRID. This will increase the object storage capacity required. The amount of capacity required will vary based on the retention duration and number of object changes with each backup.

Average object size is an input parameter that helps with sizing for performance in a StorageGRID environment. The average object sizes used for a Commvault workload depend on the type of backup.

The following table lists average object sizes by type of backup and describes what the restore process reads from the object store:

Backup Type	Average Object Size	Restore Behavior
Make an auxiliary copy in StorageGRID	32MB	Full read of 32MB object
Direct the backup to StorageGRID (deduplication enabled)	8MB	1MB random-range read
Direct the backup to StorageGRID (deduplication disabled)	32MB	Full read of 32MB object

In addition, understanding your performance requirements for full backups and incremental backups helps you determine sizing for the StorageGRID storage nodes. StorageGRID information lifecycle management (ILM) policy data protection methods determine the capacity needed to store Commvault backups and affect the sizing of the grid.

StorageGRID ILM replication is one of two mechanisms used by StorageGRID to store object data. When StorageGRID assigns objects to an ILM rule that replicates data, the system creates exact copies of the objects' data and stores the copies on storage nodes.

Erasure coding is the second method used by StorageGRID to store object data. When StorageGRID assigns objects to an ILM rule that is configured to create erasure-coded copies, it slices object data into data fragments. It then computes additional parity fragments and stores each fragment on a different storage node. When an object is accessed, it is reassembled using the stored fragments. If a data fragment or a parity fragment becomes corrupt or is lost, the erasure-coding algorithm can re-create that fragment using a subset of the remaining data and parity fragments.

The two mechanisms require different amounts of storage, as these examples demonstrate:

- If you store two replicated copies, your storage overhead doubles.
- If you store a 2+1 erasure-coded copy, your storage overhead increases by 1.5 times.

For the solution tested, an entry-level StorageGRID deployment on a single site was used:

- Admin node: VMware virtual machine (VM)

- Load balancer: VMware VM
- Storage nodes: 4x SG5712 with 4TB drives
- Primary Admin node and Gateway node: VMware VMs with the minimum production workload requirements



StorageGRID also supports third-party load balancers.

StorageGRID is typically deployed in two or more sites with data protection policies that replicate data to protect against node and site-level failures. By backing up your data to StorageGRID, your data is protected by multiple copies or by erasure coding that separates and reassembles data dependably through an algorithm.

You can use the sizing tool [Fusion](#) to size your grid.

Scaling

You can expand a NetApp StorageGRID system by adding storage to storage nodes, adding new grid nodes to an existing site, or adding a new data center site. You can perform expansions without interrupting the operation of your current system.

StorageGRID scales performance by using either higher performance nodes for storage nodes or the physical appliance which runs the load balancer and the admin nodes or by simply adding additional nodes.



For more information about expanding the StorageGRID system, see [StorageGRID 11.9 Expansion Guide](#).

Run a data protection job

To configure StorageGRID with Commvault Complete Backup and Recovery for NetApp, the following steps were performed to add StorageGRID as a cloud library within the Commvault software.

Step 1: Configure Commvault with StorageGRID

Steps

1. Log in to the Commvault Command Center. On the left panel, click Storage > Cloud > Add to see and respond to the Add Cloud dialog box:

Add cloud

X

Name

Type

NetApp StorageGRID

▼

MediaAgent

Select MediaAgent

▼



Server host

<ip-address-or-host-name>:<port>

Bucket

<Name-of-the-bucket-in-SG>

Credentials

Use saved credentials

Name

Select credentials

▼



Use deduplication

Deduplication DB location



Cancel

Save

2. For Type, select NetApp StorageGRID.
3. For MediaAgent, select all that are associated with the cloud library.
4. For Server Host, enter the IP address or the host name of the StorageGRID endpoint and the port number.

Follow the steps in StorageGRID documentation on [how to configure a load balancer endpoint \(port\)](#). Make sure you have an HTTPS port with a self-signed certificate and the IP address or the domain name of the StorageGRID endpoint.

5. If you want to use deduplication, turn on this option and provide the path to the deduplication database location.
6. Click Save.

Step 2: Create a backup plan with StorageGRID as the primary target

Steps

1. On the left panel, select Manage > Plans to see and respond to the Create Server Backup Plan dialog box.

Create server backup plan i

X

Plan name

Backup destinations

Add copy

Name

Storage

Retention period ↓

Primary

storageGRID final test

30

Primary

RPO

i

Backup frequency

Runs every 4 Hours

Add full backup

Backup window

Monday through Sunday : All day

Full backup window

Monday through Sunday : All day

Folders to backup i

▼

Snapshot options i

▼

Database options i

▼

Override restrictions

▼

Cancel

Save

2. Enter a plan name.
3. Select the StorageGRID Simple Storage Service (S3) storage backup destination that you created earlier.
4. Enter the backup retention period and recovery point objective (RPO) that you want.
5. Click Save.

Step 3: Start a backup job to protect your workloads

Steps

1. On the Commvault Command Center, navigate to Protect > Virtualization.
2. Add a VMware vCenter Server hypervisor.
3. Click the hypervisor that you just added.
4. Click Add VM group to respond to the Add VM Group dialog box so that you can see the vCenter environment that you plan to protect.

Add VM group i

Name

Browse and select VMs Hosts and clusters ▼

Search VMs ▼

[Select all](#) [Clear all](#)

- ▼ GDL1
- ▶ AOD
- ▼ SG
 - ▶ 10.193.92.169
 - ▶ 10.193.92.170
 - ▶ 10.193.92.171
 - ▶ 10.193.92.203
 - ▶ 10.193.92.227
 - ▶ 10.193.92.97
 - ▶ 10.193.92.98
 - ▶ 10.193.92.99
 - ▶ Ahmad
 - ▶ Arpita
 - ▶ Ask Ahmad before screwing around :)
 - ▶ Baremetal-VM-hosts
 - ▶ CVLT HCI POD
 - ▶ DO-NOT-TOUCH
 - ▶ Felix
 - ▶ Jonathan
 - ▶ JosephKJ
 - ▶ NAS Bridge Migration Test
 - ▶ steve
 - ▶ Yahoo Japan Test
 - Cloned-GW
 - GroupA-GW1
 - John

Backup configuration

Use backup plan

Plan to SG- No dedup ▼

Cancel Save

5. Select a datastore, a VM, or a collection of VMs, and enter a name for it.
6. Select the backup plan that you created in the previous task.
7. Click Save to see the VM group you created.
8. In the upper-right corner of the VM group window, select Backup:

The screenshot shows the CommVault Command Center interface. The left sidebar is a navigation menu with sections like 'Virtualization', 'Protect', 'Orchestrate', and 'Storage'. The main content area is titled 'VM groups / -DSN-server'. It has two tabs: 'Overview' (selected) and 'Configuration'. The 'Overview' tab displays a summary of VMs: 0 PROTECTED, 0 NOT PROTECTED, and 0 BACKED UP WITH ERRORS. Below this is a 'Summary' section with fields for Hypervisor name (sg-vc-gd1.rtp.openenglab.netapp.com-retry), Region (Not set), and Data management (a toggle switch). To the right is a 'Content' section showing a single backup job named 'DNS-server-Bootcamp'. In the top right corner of the main content area, there is a blue 'Backup' button. The top right of the interface also shows the user 'admin'.

9. Select Full as the backup level, (optionally) request an email when the backup is finished, then click OK to have your backup job start:

Select backup level

X

Full

Incremental

Synthetic full

When the job completes, [notify me via email](#)

Cancel

OK

10. Navigate to the job summary page to view the job metrics:

Review baseline performance tests

In the Aux Copy operation, four Commvault MediaAgents backed up data to a NetApp AFF A300 system and an auxiliary copy was created on NetApp StorageGRID. For details on the test setup environment, read the Solution Design and Best Practices section in the [NetApp Scale-out Data Protection with Commvault](#) technical report.

The tests were performed with 100 VMs and 1000 VMs, both tests with a 50/50 mix of Windows and CentOS VMs. The following table shows the results from our baseline performance tests:

Operation	Backup Speed	Restore Speed
Aux Copy	2 TB/hour	1.27 TB/hour
Direct to and from object (Deduplication On)	2.2 TB/hour	1.22 TB/hour

To test age-off performance, 2.5 million objects were deleted. As shown in Figures 2 and 3, the delete run was completed in less than 3 hours and freed up more than 80TB of space. The delete run started at 10:30 AM.

Figure 1: Deletion of 2.5 million (80TB) objects in less than 3 hours.





Figure 2: Freeing up 80TB of storage in less than 3 hours.



Bucket consistency level recommendation

NetApp StorageGRID allows the end user to select the consistency level for operations performed on the objects in Simple Storage Service (S3) buckets.

Commvault MediaAgents are the data movers in a Commvault environment. In most cases, MediaAgents are configured to write locally into a primary StorageGRID site. For this reason, a high consistency level within a local primary site is recommended. Use the following guidelines when you set the consistency level on Commvault buckets created in StorageGRID.



If you have a Commvault version earlier than 11.0.0 - Service Pack 16, consider upgrading Commvault to the newest version. If that is not an option, be sure to follow the guidelines for your version.

- Commvault versions earlier than 11.0.0 - Service Pack 16.* In versions earlier than 11.0.0 - Service Pack 16, Commvault performs S3 HEAD and GET operations on non-existent objects as part of restore and pruning process. Set the bucket consistency level to Strong-site to achieve the optimal consistency level for Commvault backups to StorageGRID.
- Commvault versions 11.0.0 - Service Pack 16 and later.* In versions 11.0.0 - Service Pack 16 and later, the number of S3 HEAD and GET operations performed on non-existent objects are minimized. Set the default

bucket consistency level to Read-after-new-write to ensure high consistency level in the Commvault and StorageGRID environment.

TR-4626: Load balancers

Use third-party load balancers with StorageGRID

Learn about the role of a third-party and global load balancers in an object storage systems like StorageGRID.

General guidance for implementing NetApp® StorageGRID® with third-party load balancers.


Object storage is synonymous with the term cloud storage, and, as you would expect, applications that leverage cloud storage address that storage through a URL. Behind that simple URL, StorageGRID can scale capacity, performance, and durability in a single site or over geo-distributed sites. The component that makes this simplicity possible is a load balancer.

The purpose of this document is to educate StorageGRID customers about load balancer options and provide general guidance for the configuration of third-party load balancers.

Load balancer basics

Load balancers are an essential component of an enterprise grade object storage system such as StorageGRID. StorageGRID consists of multiple storage nodes, each of which can present the entire Simple Storage Service (S3) name space for a given StorageGRID instance. Load balancers create a highly available endpoint behind which we can place StorageGRID nodes. StorageGRID is unique among S3-compatible object storage systems in that it provides its own load balancer, but it also supports third-party or general-purpose load balancers such as F5, Citrix Netscaler, HA Proxy, NGINX, and so on.

The following figure uses the example URL/ fully qualified domain name (FQDN) “s3.company.com”. The load balancer creates a virtual IP (VIP) that resolves to the FQDN through DNS, then directs any requests from applications to a pool of StorageGRID nodes. The load balancer performs a health check on each node and only establishes connections to healthy nodes.



The figure shows the StorageGRID provided load balancer, but the concept is the same for third-party load balancers. Applications establish an HTTP session using the VIP on the load balancer and the traffic passes through the load balancer to the storage nodes. By default, all traffic, from application to load balancer, and from load balancer to storage node is encrypted through HTTPS. HTTP is a supported option.

Local and global load balancers

There are two types of load balancers:

- **Local Traffic Managers (LTM).** Spreads connections over a pool of nodes in a single site.
- **Global Service Load Balancer (GSLB).** Spreads connections over multiple sites, effectively load balancing LTM load balancers. Think of a GSLB as an intelligent DNS server. When a client requests a StorageGRID endpoint URL, the GSLB resolves it to the VIP of an LTM based on availability or other factors (for example, which site can provide lower latency to the application). While an LTM is always required, a GSLB is optional depending on the number of StorageGRID sites and your application requirements.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- NetApp StorageGRID Documentation Center
<https://docs.netapp.com/us-en/storagegrid/>
- NetApp StorageGRID Enablement
<https://docs.netapp.com/us-en/storagegrid-enable/>
- StorageGRID f5 load balancer design considerations
<https://www.netapp.com/blog/storagegrid-f5-load-balancer-design-considerations/>
- Loadbalancer.org—Load balancing NetApp StorageGRID
<https://www.loadbalancer.org/applications/load-balancing-netapp-storagegrid/>
- Kemp—Load balancing NetApp StorageGRID
<https://support.kemptechnologies.com/hc/en-us/articles/360045186451-NetApp-StorageGRID>

Use StorageGRID load balancers

Learn about the role of a StorageGRID Gateway Node load balancer.

General guidance for implementing NetApp® StorageGRID® Gateway Nodes.

StorageGRID Gateway Node load balancer versus third-party load balancer

StorageGRID is unique among S3-compatible object storage vendors in that it provides a native load balancer available as a purpose-built appliance, VM, or container. The StorageGRID provided load balancer is also referred to as a Gateway Node.

For customers that do not already own a load balancer such as F5, Citrix, and so on, implementation of a third-party load balancer can be very complex. The StorageGRID load balancer greatly simplifies load balancer operations.

The Gateway Node is an enterprise grade, highly available, and high-performance load balancer. Customers can choose to implement the Gateway Node, third-party load balancer, or even both, in the same grid. The Gateway Node is a local traffic manager versus a GSLB.

The StorageGRID load balancer provides the following advantages:

- **Simplicity.** Automatic configuration of resource pools, health checks, patching, and maintenance, all managed by StorageGRID.

- **Performance.** The StorageGRID load balancer is dedicated to StorageGRID, can provide high performance caching, and you do not compete with other applications for bandwidth.
- **Cost.** The virtual machine (VM) and container versions are provided at no additional cost.
- **Traffic classifications.** The Advanced Traffic Classification feature allows for StorageGRID-specific QoS rules along with workload analytics.
- **Future StorageGRID specific features.** StorageGRID will continue to optimize and add innovative features to the load balancer over upcoming releases.

As an integrated node of StorageGRID, the local traffic manager has the ability to use advanced health checking to distribute requests based on Storage Node health, load and resource availability. In addition it has the ability to distribute the load across multiple sites when the StorageGRID link costs are set to "0" between the sites. In the event the Storage Nodes are unavailable but the Gateway Node is available in a site, the load will automatically be directed to another site in the grid.

The Load balancer caching feature of the Gateway Node is intended to provide a substantial performance improvement for certain workloads (such as AI training) which re-read a data set multiple times as part of processing that data.

Caching gateway nodes can also be deployed physically distant from the rest of the grid enabling better performance and lower WAN network utilization in some workloads. The cache operates in a read back mode where writes are not cached and do not modify the state of the cache. Each Caching Gateway Node operates independently of any other caching Gateway node.

For details about deploying the StorageGRID Gateway Node, see the [StorageGRID documentation](#).

Learn how to implement SSL certificates for HTTPS in StorageGRID

Understand the importance and the steps to implement of SSL certificates in StorageGRID.

If you are using HTTPs, you must have a Secure Sockets Layer (SSL) certificate. The SSL protocol identifies the clients and endpoints, validating them as trusted. SSL also provides encryption of the traffic. The SSL certificate must be trusted by the clients. To accomplish this, the SSL certificate can be from a globally trusted Certificate Authority (CA) such as DigiCert, a private CA running in your infrastructure, or a self-signed certificate generated by the host.

Using a globally trusted CA certificate is the preferred method as there is no additional client-side actions required. The certificate is loaded into the load balancer or StorageGRID, and the clients trust and connect to the endpoint.

Using a private CA requires the root and all subordinate certificates be added to the client. The process to trust a private CA certificate can vary by client operating system and applications. For example, in ONTAP for FabricPool, you must upload each certificate in the chain individually (root certificate, subordinate certificate, endpoint certificate) to the ONTAP cluster.

Using a self-signed certificate requires the client to trust the provided certificate without any CA to verify the authenticity. Some applications might not accept self-signed certificates and have no ability to ignore verification.

The placement of the SSL certificate in the client load balancer StorageGRID path depends on where you need the SSL termination to be. You can configure a load balancer to be the termination endpoint for the client, and then re-encrypt or hot encrypt with a new SSL certificate for the load balancer to StorageGRID connection. Or you can pass through the traffic and let StorageGRID be the SSL termination endpoint. If the load balancer is the SSL termination endpoint, the certificate is installed on the load balancer and contains the subject name

for the DNS name/URL and any alternative URL/DNS names for which a client is configured to connect to the StorageGRID target through the load balancer, including any wild card names. If the load balancer is configured for pass through, the SSL certificate must be installed in StorageGRID. Again, the certificate must contain the subject name for the DNS name/URL, and any alternative URL/DNS names for which a client is configured to connect to the StorageGRID target through the load balancer, including any wild card names. Individual Storage Node names do not need to be included on the certificate, only the endpoint URLs.

```
Subject DN: /C=US/postalCode=94089/ST=California/L=Sunnyvale/street=495 East Java Dr/O=NetApp, Inc./OU=IT1/OU=Unified Communication  
s/CN=webscaledemo.netapp.com  
Serial Number: 37:4C:6B:51:61:84:50:F8:7A:29:D9:83:24:12:36:2C  
Issuer DN: /C=GB/ST=Greater Manchester/L=Salford/O=Sectigo Limited/CN=Sectigo RSA Organization Validation Secure Server CA  
Issued On: 2019-05-23T00:00:00.000Z  
Expires On: 2021-05-22T23:59:59.000Z  
Alternative Names: DNS:webscaledemo.netapp.com  
DNS:*.webscaledemo-rtp.netapp.com  
DNS:*.webscaledemo.netapp.com  
DNS:webscaledemo-rtp.netapp.com  
SHA-1 Fingerprint: 60:91:44:E5:4F:7E:25:6B:B5:A0:19:87:D1:F2:8C:DD:AD:3A:88:CD  
SHA-256 Fingerprint: FE:21:5D:BF:08:D9:5A:E5:09:CF:F6:3F:D3:5C:1E:9B:33:63:63:CA:25:2D:3F:39:0B:6A:B8:EC:08:BC:57:43
```

Configure trusted third-party load balancer in StorageGRID

Learn how to configure trusted third-party load balancer in StorageGRID.

If you are using one or more external layer 7 load balancers, and an S3 bucket or group policies that are IP based, StorageGRID must determine the real sender's IP address. It does this by looking at the X-Forwarded-For (XFF) header, which is inserted into the request by the load balancer. As the XFF header can be easily spoofed in requests sent directly to the Storage Nodes, StorageGRID must confirm that each request is being routed by a trusted layer 7 load balancer. If StorageGRID cannot trust the source of the request, it will ignore the XFF header. There is a Grid Management API to allow a list of trusted external layer 7 load balancers to be configured. This new API is private and is subject to change in future StorageGRID releases. For the most up to date information, see the KB article, [How to configure StorageGRID to work with third-party Layer 7 load balancers](#).

Learn about local traffic manager load balancers

Explore the guidance for local traffic manager load balancers and determine the optimal configuration.

The following is presented as general guidance for configuration of third-party load balancers. Work with your load balancer administrator to determine the optimal configuration for your environment.

Create a resource group of Storage Nodes

Group StorageGRID Storage Nodes into a resource pool or service group (the terminology might differ with specific load balancers).

StorageGRID Storage Nodes present the S3 API on the following ports:

- S3 HTTPS: 18082
- S3 HTTP: 18084

Most customers choose to present the APIs on the virtual server through the standard HTTPS and HTTP ports (443 and 80).



Each StorageGRID site requires a default of three Storage Nodes, two of which must be healthy.

Health check

Third-party load balancers require a method to determine the health of each node and its eligibility to receive traffic. NetApp recommends the HTTP OPTIONS method to perform the health check. The load balancer issues HTTP OPTIONS requests to each individual Storage Node and expects a 200 status response.

If any Storage Node does not provide a 200 response, that node is not able to service storage requests. Your application and business requirements should determine the timeout for these checks and the action your load balancer takes.

For example, if three of four Storage Nodes in data center 1 are down, you might direct all traffic to data center 2.

The recommended polling interval is once per second, marking the node offline after three failed checks.

S3 health check example

In the following example, we send OPTIONS and check for 200 OK. We use OPTIONS because Amazon S3 does not support unauthorized requests.

```
curl -X OPTIONS https://10.63.174.75:18082 --verbose --insecure
* Rebuilt URL to: https://10.63.174.75:18082/
*   Trying 10.63.174.75...
* TCP_NODELAY set
* Connected to 10.63.174.75 (10.63.174.75) port 18082 (#0)
* TLS 1.2 connection using TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384
* Server certificate: webscale.stl.netapp.com
* Server certificate: NetApp Corp Issuing CA 1
* Server certificate: NetApp Corp Root CA
> OPTIONS / HTTP/1.1
> Host: 10.63.174.75:18082
> User-Agent: curl/7.51.0
> Accept: /
>
< HTTP/1.1 200 OK
< Date: Mon, 22 May 2017 15:17:30 GMT
< Connection: KEEP-ALIVE
< Server: StorageGRID/10.4.0
< x-amz-request-id: 3023514741
```

File or content-based health checks

In general, NetApp does not recommend file-based health checks. Typically, a small file — `healthcheck.htm`, for example — is created in a bucket with a read-only policy. This file is then fetched and evaluated by the load balancer. This approach has several disadvantages:

- **Dependent on a single account.** If the account that owns the file is disabled, the health check fails, and

no storage requests are processed.

- **Data protection rules.** The default data protection scheme is a two-copy approach. In this scenario, if the two storage nodes hosting the health check file are unavailable, the health check fails, and storage requests are not sent to healthy storage nodes, rendering the grid offline.
- **Audit log bloat.** The load balancer fetches the file from every storage node every X minutes, creating many audit log entries.
- **Resource intensive.** Fetching the health check file from every node every few seconds consumes grid and network resources.

If a content-based health check is required, use a dedicated tenant with a dedicated S3 bucket.

Session persistence

Session persistence, or stickiness, refers to the time a given HTTP session is allowed to persist. By default, sessions are dropped by Storage Nodes after 10 minutes. Longer persistence can lead to better performance because applications do not have to reestablish their sessions for every action; however, holding these sessions open consumes resources. If you determine that your workload will benefit, you can reduce the session persistence on a third-party load balancer.

Virtual hosted-style addressing

Virtual hosted-style is now the default method for AWS S3, and while StorageGRID and many applications still support path style, it is best practice to implement virtual hosted-style support. Virtual hosted-style requests have the bucket as part of the host name.

To support virtual hosted-style, do the following:

- Support wildcard DNS lookups: *.s3.company.com
- Use an SSL certificate with subject alt names to support wildcard: *.s3.company.com

Some customers have expressed security concerns around the use of wildcard certificates. StorageGRID continues to support path style access, as do key applications such as FabricPool. That said, certain S3 API calls fail or behave improperly without virtual hosted support.

SSL termination

There are security benefits to SSL termination on third-party load balancers. If the load balancer is compromised, the grid is compartmentalized.

There are three supported configurations:

- **SSL pass-through.** The SSL certificate is installed on StorageGRID as a custom server certificate.
- **SSL termination and re-encryption (recommended).** This might be beneficial if you are already doing SSL certificate management on the load balancer rather than installing the SSL certificate on StorageGRID. This configuration provides the additional security benefit of limiting the attack surface to the load balancer.
- **SSL termination with HTTP.** In this configuration, SSL is terminated on the third-party load balancer and communication from the load balancer to StorageGRID is nonencrypted to take advantage of SSL off-load (with SSL libraries embedded in modern processors this is of limited benefit).

Pass through configuration

If you prefer to configure your load balancer for pass through, you must install the certificate on StorageGRID.

Go to **Configuration > Server Certificates > Object Storage API Service Endpoints Server Certificate**.

Source client IP visibility

StorageGRID 11.4 introduced the concept of a trusted third-party load balancer. In order to forward the client application IP to StorageGRID, you must configure this feature. For more information, see [How to configure StorageGRID to work with third-party Layer 7 load balancers](#).

To enable the XFF header to be used to view the IP of the client application, follow these steps:

Steps

1. Record the client IP in the audit log.
2. Use `aws:SourceIp` S3 bucket or group policy.

Load balancing strategies

Most load balancing solutions offer multiple strategies for load balancing. The following are common strategies:

- **Round robin.** A universal fit but suffers with few nodes and large transfers clogging single nodes.
- **Least connection.** A good fit for small and mixed object workloads, resulting in an equal distribution of the connections to all nodes.

The choice of algorithm becomes less important with an increasing number of Storage Nodes to choose from.

Data path

All data flows through local traffic manager load balancers. StorageGRID does not support direct server routing (DSR).

Verifying distribution of connections

To verify that your method is distributing the load evenly across Storage Nodes, check the established sessions on each node in a given site:

- **UI Method.** Go to **Support > Metrics > S3 Overview > LDR HTTP Sessions**
- **Metrics API.** Use `storagegrid_http_sessions_incoming_currently_established`

Learn about few use cases for StorageGRID configurations

Explore few use cases for StorageGRID configurations implemented by customers and NetApp IT.

The following examples illustrate configurations as implemented by StorageGRID customers, including NetApp IT.

F5 BIG-IP local traffic manager health check monitor for S3 bucket

To configure the F5 BIG-IP local traffic manager health check monitor, follow these steps:

Steps

1. Create a new monitor.

- a. In the Type field, enter HTTPS.
- b. Configure the interval and timeout as desired.
- c. In the Send String field, enter OPTIONS / HTTP/1.1\r\n\r\n.
\r\n are carriage returns; different versions of BIG-IP software require zero, one, or two sets of \r\n sequences. For more information, see <https://support.f5.com/csp/article/K10655>.
- d. In the Receive String field, enter: HTTP/1.1 200 OK.

Local Traffic > Monitors > New Monitor...

General Properties	
Name	https_storagegrid
Description	
Type	HTTPS
Parent Monitor	https
Configuration: Basic	
Interval	5 seconds
Timeout	16 seconds
Send String	OPTIONS / HTTP/1.1\r\n\r\n
Receive String	HTTP/1.1 200 OK
Receive Disable String	
Cipher List	DEFAULT+SHA+3DES+KEDH
User Name	
Password	
Reverse	<input type="radio"/> Yes <input checked="" type="radio"/> No
Transparent	<input type="radio"/> Yes <input checked="" type="radio"/> No
Alias Address	* All Addresses
Alias Service Port	* All Ports
Adaptive	<input type="checkbox"/> Enabled


2. In Create Pool, create one pool for each port required.
 - a. Assign the health monitor that you created in the previous step.
 - b. Select a load-balancing method.
 - c. Select service port: 18082 (S3).
 - d. Add nodes.

Citrix NetScaler

Citrix NetScaler creates a virtual server for the storage endpoint and refers to StorageGRID Storage Nodes as Application Servers, which are then grouped into Services.

Use the HTTPS-ECV health check monitor to create a custom monitor to perform the recommended health check by using the OPTIONS request and receiving 200. HTTP-ECV is configured with a send string and validates a receive string.

For more information, see the Citrix documentation, [Sample configuration for HTTP-ECV health check monitor](#).



The screenshot shows the 'Monitors' configuration screen in Citrix NetScaler. At the top, there are buttons for 'Add Binding', 'Edit Binding', 'Unbind', and 'Edit Monitor'. A table lists monitors, with one entry for 'STORAGE-GRID-TCP-ECV-MON' which is marked as 'Active' (indicated by a checkmark in the 'State' column). The 'Configure Monitor' section below shows the following configuration details:

- Name:** STORAGE-GRID-TCP-ECV-MON
- Type:** TCP-ECV
- Basic Parameters:**
 - Interval:** 5 seconds
 - Response Timeout:** 2 seconds
- Send String:** OPTIONS / HTTP/1.1/r/r/r/n
- Receive String:** HTTP/1.1 200 OK
- SSL Profile:** receive

Loadbalancer.org

Loadbalancer.org has conducted their own integration testing with StorageGRID and has an extensive configuration guide: https://pdfs.loadbalancer.org/NetApp_StorageGRID_Deployment_Guide.pdf.

Kemp

Kemp has conducted their own integration testing with StorageGRID and has an extensive configuration guide: <https://kemptechnologies.com/solutions/netapp/>.

HAProxy

Configure HAProxy to use the OPTIONS request and check for a 200 status response for the health check in haproxy.cfg. You can change the bind port in the front end to a different port, such as 443.

The following is an example for SSL termination on HAProxy:

```

frontend s3
    bind *:443 crt /etc/ssl/server.pem ssl
    default_backend s3-serve
rs
backend s3-servers
    balance leastconn
    option httpchk
    http-check expect status 200
    server dc1-s1 10.63.174.71:18082 ssl verify none check inter 3000
    server dc1-s2 10.63.174.72:18082 ssl verify none check inter 3000
    server dc1-s3 10.63.174.73:18082 ssl verify none check inter 3000

```

The following is an example for SSL pass-through:

```

frontend s3
    mode tcp
    bind *:443
    default_backend s3-servers
backend s3-servers
    balance leastconn
    option httpchk
    http-check expect status 200
    server dc1-s1 10.63.174.71:18082 check-ssl verify none inter 3000
    server dc1-s2 10.63.174.72:18082 check-ssl verify none inter 3000
    server dc1-s3 10.63.174.73:18082 check-ssl verify none inter 3000

```

For full examples of configurations for StorageGRID, see [Examples for HAProxy Configuration](#) on GitHub.

Validate SSL connection in StorageGRID

Learn how to validate the SSL connection in StorageGRID.


After your load balancer is configured, you should validate the connection using tools such as OpenSSL and the AWS CLI. Other applications, such as S3 Browser, might ignore SSL misconfiguration.

Understand global load balancing requirements for StorageGRID

Explore the design considerations and requirements for global load balancing in StorageGRID.

Global load balancing requires integrating with DNS to provide intelligent routing across multiple StorageGRID sites. This function falls outside of the StorageGRID domain and must be provided by a third-party solution such as the load balancer products discussed previously and/or a DNS traffic control solution such as Infoblox. This top level load balancing provides smart routing to the closest destination site in the namespace, as well as outage detection and redirection to the next site in the namespace. A typical GSLB implementation consists of the top level GSLB with site pools containing site-local load balancers. The site load balancers contain pools of

the local site Storage Nodes. This can include a combination of third-party load balancers for GSLB functions and StorageGRID providing the site-local load balancing, or a combination of third parties, or many of the third parties discussed previously can provide both GSLB and site-local load balancing.



TR-4645: Security features

Secure StorageGRID data and metadata in an object store

Discover the integral security features of the StorageGRID object storage solution.

This is an overview of the many security features in NetApp® StorageGRID®, covering data access, objects and metadata, administrative access, and platform security. It has been updated to include the newest features released with StorageGRID 12.0.

Security is an integral part of the NetApp StorageGRID object storage solution. Security is particularly important because many types of rich content data that are well suited for object storage are also sensitive in nature and subject to regulations and compliance. As StorageGRID capabilities continue to evolve, the software makes available many security features that are invaluable for protecting an organization's security posture and helping the organization adhere to industry best practices.

This paper is an overview of the many security features in StorageGRID 12.0, divided into five categories:

- Data access security features
- Object and metadata security features
- Administration security features
- Platform security features
- Cloud integration

This paper is intended to be a security datasheet—it does not detail how to configure the system to support the security features enumerated within that are not configured by default. The [StorageGRID Hardening Guide](#) is available on the official [StorageGRID Documentation](#) page.

In addition to the capabilities described in this report, StorageGRID follows the [NetApp Product Security Vulnerability Response and Notification Policy](#). Reported vulnerabilities are verified and responded to according to the product security incident response process.

NetApp StorageGRID provides advanced security features for highly demanding enterprise object storage use cases.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- NetApp StorageGRID: SEC 17a-4(f), FINRA 4511(c) and CFTC 1.31(c)-(d) Compliance Assessment
<https://www.netapp.com/media/9041-ar-cohasset-netapp-storagegrid-sec-assessment.pdf>
- NetApp StorageGRID NIST FIPS 140-3 Kernel Crypto Certification
<https://csrc.nist.gov/projects/cryptographic-module-validation-program/certificate/5097>
- NetApp StorageGRID NIST SP 800-90B Entropy Certification
<https://csrc.nist.gov/projects/cryptographic-module-validation-program/entropy-validations/certificate/223>
- NetApp StorageGRID Canadian Centre for Cyber Security Common Criteria Certification
<https://www.commoncriteriaportal.org/nfs/ccpfiles/files/epfiles/565-LSS%20CT%20v1.0.pdf>
- StorageGRID Documentation page
<https://docs.netapp.com/us-en/storagegrid/>
- NetApp Product Documentation
<https://www.netapp.com/support-and-training/documentation/>

Terms and acronyms

This section provides definitions for the terminology used in the document.

Term or acronym	Definition
S3	Simple Storage Service.
Client	An application that can interface with StorageGRID either through the S3 protocol for data access or HTTP protocol for management.
Tenant admin	The administrator of the StorageGRID tenant account
Tenant user	A user within a StorageGRID tenant account
TLS	Transport Layer Security
ILM	Information Lifecycle Management
LAN	Local Area Network
Grid administrator	The administrator of the StorageGRID system
Grid	The StorageGRID system
Bucket	A container for objects stored in S3
LDAP	Lightweight Directory Access Protocol

Term or acronym	Definition
SEC	Securities and Exchange Commission; regulates exchange members, brokers, or dealers
FINRA	Financial Industry Regulatory Authority; defers to the format and media requirements of SEC Rule 17a-4(f)
CFTC	Commodity Futures Trading Commissions; regulates commodity futures trading
NIST	National Institute of Standards and Technology

Data access security features

Learn about the data access security features in StorageGRID.

Feature	Function	Impact	Regulatory compliance
Configurable Transport Layer Security (TLS)	<p>TLS establishes a handshake protocol for communication between a client and a StorageGRID gateway node, storage node, or load balancer endpoint.</p> <p>StorageGRID supports the following cipher suites for TLS:</p> <ul style="list-style-type: none"> • TLS_AES_256_GCM_SHA384 • TLS_AES_128_GCM_SHA256 • ECDHE-ECDSA-AES256-GCM-SHA384 • ECDHE-RSA-AES256-GCM-SHA384 • ECDHE-ECDSA-AES128-GCM-SHA256 • ECDHE-RSA-AES128-GCM-SHA256 • TLS_AES_256_GCM_SHA384 • DHE-RSA-AES128-GCM-SHA256 • DHE-RSA-AES256-GCM-SHA384 • AES256-GCM-SHA384 • AES128-GCM-SHA256 • TLS_CHACHA20_POLY1305_SHA256 • ECDHE-ECDSA-CHACHA20-POLY1305 • ECDHE-RSA-CHACHA20-POLY1305 <p>TLS v1.2 & 1.3 supported.</p> <p>SSLv3, TLS v1.1 and earlier are not supported.</p>	<p>Enables a client and StorageGRID to identify and authenticate each other and communicate with confidentiality and data integrity. Ensures use of a recent TLS version. Ciphers are now configurable under the Configuration/Security settings</p>	—

Feature	Function	Impact	Regulatory compliance
Configurable Server Certificate (Load Balancer Endpoint)	Grid administrators can configure Load Balancer Endpoints to generate or use a server certificate.	Enables the use of digital certificates signed by their standard trusted certificate authority (CA) to authenticate object API operations between grid and client per Load Balancer Endpoint.	—
Configurable Server Certificate (API endpoint)	Grid administrators can centrally configure all StorageGRID API endpoints to use a server certificate signed by their organization's trusted CA.	Enables the use of digital certificates signed by their standard, trusted CA to authenticate object API operations between a client and the grid.	—

Feature	Function	Impact	Regulatory compliance
Multitenancy	<p>StorageGRID supports multiple tenants per grid; each tenant has its own namespace. A tenant provides S3 protocol; by default, access to buckets/containers and objects is restricted to users within the account. Tenants can have one user (for example, an enterprise deployment, in which each user has their own account) or multiple users (for example, a service provider deployment, in which each account is a company and a customer of the service provider). Users can be local or federated; federated users are defined by Active Directory or Lightweight Directory Access Protocol (LDAP). StorageGRID provides a per-tenant dashboard, where users log in using their local or federated account credentials. Users can access visualized reports on tenant usage against the quota assigned by the grid administrator, including usage information in data and objects stored by buckets. Users with administrative permission can perform tenant-level system administration tasks, such as managing users and groups and access keys.</p>	<p>Allows StorageGRID administrators to host data from multiple tenants while isolating tenant access, and to establish user identity by federating users with an external identity provider, such as Active Directory or LDAP.</p>	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)
Nonrepudiation of access credentials	<p>Every S3 operation is identified and logged with a unique tenant account, user, and access key.</p>	<p>Allows Grid administrators to establish what API actions are performed by which individuals.</p>	—

Feature	Function	Impact	Regulatory compliance
Disabled anonymous access	<p>By default, anonymous access is disabled for S3 accounts. A requester must have a valid access credential for a valid user in the tenant account to access buckets, containers, or objects within the account.</p> <p>Anonymous access to S3 buckets or objects can be enabled with an explicit IAM policy.</p>	Allows Grid administrators to disable or control anonymous access to buckets/containers and objects.	—
Compliance WORM	<p>Designed to meet the requirements of SEC Rule 17a-4(f) and validated by Cohasset. Customers can enable compliance at the bucket level. Retention can be extended but never reduced. Information lifecycle management (ILM) rules enforce minimum data protection levels.</p>	Allows tenants with regulatory data retention requirements to enable WORM protection on stored objects and object metadata.	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)
WORM	<p>Grid administrators can enable grid-wide WORM by enabling the Disable Client Modify option, which prevents clients from overwriting or deleting objects or object metadata in all tenant accounts.</p> <p>S3 Tenant admins can also enable WORM by tenant, bucket, or object prefix by specifying IAM policy, which includes the custom S3: PutOverwriteObject permission for object and metadata overwrite.</p>	Allows Grid administrators and tenant admins to control WORM protection on stored objects and object metadata.	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)

Feature	Function	Impact	Regulatory compliance
KMS host server encryption key management	Grid administrators can configure one or more external key management servers (KMS) in the Grid Manager to provide encryption keys to StorageGRID services and storage appliances. Each KMS host server or KMS host server cluster uses the Key Management Interoperability Protocol (KMIP) to provide an encryption key to the appliance nodes at the associated StorageGRID site.	Data-at-rest encryption is achieved. After the appliance volumes are encrypted, you cannot access any data on the appliance unless the node can communicate with the KMS host server.	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)
Automated failover	StorageGRID provides built-in redundancy and automated failover. Access to tenant accounts, buckets, and objects can continue even if there are multiple failures, from disks or nodes to entire sites. StorageGRID is resource-aware and automatically redirects requests to available nodes and data locations. StorageGRID sites can even operate in islanded mode; if a WAN outage disconnects a site from the rest of the system, reads and writes can continue with local resources, and replication resumes automatically when the WAN is restored.	Enables Grid administrators to address uptime, SLA, and other contractual obligations and to implement business continuity plans.	—
S3-specific data access security features			
AWS Signature Version 2 and Version 4	Signing API requests provides authentication for S3 API operations. Amazon supports two versions of Signature Version 2 and Version 4. The signing process verifies the identity of the requester, protects data in transit, and protects against potential replay attacks.	Aligns with AWS recommendation for Signature Version 4 and enables backward compatibility with older applications with Signature Version 2.	—

Feature	Function	Impact	Regulatory compliance
S3 Object Lock	The S3 Object Lock feature in StorageGRID is an object-protection solution that is equivalent to S3 Object Lock in Amazon S3.	Allows tenants to create buckets with S3 Object Lock enabled to comply with regulations that require certain objects to be retained for a fixed amount of time or indefinitely.	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)
Secured storage of S3 credentials	S3 access keys are stored in a format that is protected by a password hashing function (SHA-2).	Enables secure storage of access keys by a combination of key length (a 10^{31} randomly generated number) and a password hashing algorithm.	—
Time-bound S3 access keys	When creating an S3 access key for a user, customers can set an expiration date and time on the access key.	Gives Grid administrators the option to provision temporary S3 access keys.	—
Multiple access keys per user account	StorageGRID enables multiple access keys to be created and simultaneously active for a user account. Because each API action is logged with a tenant user account and access key, nonrepudiation is preserved despite multiple keys being active.	Enables clients to non-disruptively rotate access keys and allows each client to have its own key, discouraging key sharing across clients.	—
S3 IAM access policy	StorageGRID supports S3 IAM policies, enabling Grid administrators to specify granular access control by tenant, bucket, or object prefix. StorageGRID also supports IAM policy conditions and variables, allowing more dynamic access control policies.	Allows Grid administrators to specify access control by user groups for the whole tenant; also enables tenant users to specify access control for their own buckets and objects.	—

Feature	Function	Impact	Regulatory compliance
S3 Security Token Service API AssumeRole	<p>StorageGRID supports the S3 STS API AssumeRole to provide temporary security credentials (access key ID, secret access key, session token) with downscoped permissions and limited duration.</p> <p>Inline session policies to further restrict permissions during the session are supported as part of the AssumeRole API.</p>	Allows Tenant administrators to provide secure temporary access to object data.	—
Simple Notification Service	<p>StorageGRID supports sending notification on object access. The following event types are supported:</p> <ul style="list-style-type: none"> • s3:ObjectCreated: • s3:ObjectCreated:Put • s3:ObjectCreated:Post • s3:ObjectCreated:Copy • s3:ObjectCreated:Complete MultipartUpload • s3:ObjectRemoved: • s3:ObjectRemoved:Delete • s3:ObjectRemoved:Delete MarkerCreated • s3:ObjectRestore:Post 	Allows Tenant administrators to monitor access to objects	—
Server-side encryption with StorageGRID-managed keys (SSE)	StorageGRID supports SSE, allowing multitenant protection of data at rest with encryption keys managed by StorageGRID.	<p>Enables tenants to encrypt objects.</p> <p>Encryption key is required to write and retrieve these objects.</p>	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)

Feature	Function	Impact	Regulatory compliance
Server-side encryption with customer-provided encryption keys (SSE-C)	<p>StorageGRID supports SSE-C, enabling multitenant protection of data at rest with encryption keys managed by the client.</p> <p>Although StorageGRID manages all object encryption and decryption operations, with SSE-C, the client must manage the encryption keys themselves.</p>	<p>Enables clients to encrypt objects with keys they control.</p> <p>Encryption key is required to write and retrieve these objects.</p>	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)

Object and metadata security

Explore the object and metadata security features in StorageGRID.

Feature	Function	Impact	Regulatory compliance
Advanced Encryption Standard (AES) Server-Side Object Encryption	StorageGRID provides AES 128- and AES 256-based server-side encryption of objects. Grid administrators can enable encryption as a global default setting. StorageGRID also supports the S3 x-amz-server-side-encryption header to allow enabling or disabling encryption on a per-object basis. When enabled, objects are encrypted when stored or in transit between grid nodes.	Helps secure storage and transmission of objects, independent of the underlying storage hardware.	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)
Built-in Key Management	When encryption is enabled, each object is encrypted with a randomly generated unique symmetric key, which is stored inside StorageGRID with no external access.	Enables encryption of objects without requiring External Key Management.	
Federal Information Processing Standard (FIPS) 140-2 compliant encryption disks	The SG5812, SG5860, SG6160, and SGF6024 StorageGRID appliances offer the option of FIPS 140-2 compliant encryption disks. Encryption keys for the disks can be optionally managed by an external KMIP server.	Enables secure storage of system data, metadata, and objects. Also provides StorageGRID software-based object encryption, which secures storage and transmission of objects.	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)

Feature	Function	Impact	Regulatory compliance
Federal Information Processing Standard (FIPS) 140-3 compliant encryption for nodes	<p>The SG5812, SG5860, SG6160, SGF6112, SG1100, and SG110 StorageGRID appliances offer the option of FIPS 140-3 compliant node encryption.</p> <p>Encryption keys for the nodes are managed by an external KMIP server.</p>	<p>Enables secure storage of system data, metadata, and objects. Also provides StorageGRID software-based object encryption, which secures storage and transmission of objects.</p>	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)
Background Integrity Scan and Self-Healing	<p>StorageGRID uses an interlocking mechanism of hashes, checksums, and cyclic redundancy checks (CRCs) at the object and sub-object level to protect against data inconsistency, tampering, or modification, both when objects are in storage and in transit.</p> <p>StorageGRID automatically detects corrupt and tampered objects and replaces them, while quarantining the altered data and alerting the administrator.</p>	<p>Enables Grid administrators to meet SLA, regulations, and other obligations regarding data durability. Helps customers detect ransomware or viruses attempting to encrypt, tamper, or modify data.</p>	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)
Policy-based object placement and retention	<p>StorageGRID enables Grid administrators to configure ILM rules, which specify object retention, placement, protection, transition, and expiration. Grid administrators can configure StorageGRID to filter objects by their metadata and to apply rules at various levels of granularity, including grid-wide, tenant, bucket, key prefix, and user-defined metadata key-value pairs.</p> <p>StorageGRID helps to ensure that objects are stored according to the ILM rules throughout their lifecycles, unless they are explicitly deleted by the client.</p>	<p>Helps enforce data placement, protection, and retention. Helps customers achieve SLA for durability, availability, and performance.</p>	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)

Feature	Function	Impact	Regulatory compliance
Background metadata scanning	StorageGRID periodically scans object metadata in the background to apply changes in object data placement or protection as specified by ILM.	Helps discover corrupted objects.	
Tunable consistency	Tenants can select consistency levels at the bucket level to ensure that resources such as multisite connectivity are available.	Provides the option to commit writes to the grid only when a required number of sites or resources are available.	

Administration security features

Discover the administration security features in StorageGRID.

Feature	Function	Impact	Regulatory compliance
Server Certificate (Grid Management Interface)	Grid administrators can configure the Grid Management Interface to use a server certificate signed by their standard, trusted CA to authenticate management UI and API access between a management client and the grid.	Enables the use of digital certificates signed by their standard, trusted CA to authenticate management UI and API access between a management client and the grid.	—
Administrative user authentication	Administrative users are authenticated using username and password. Administrative users and groups can be local or federated, imported from the customer's Active Directory or LDAP. Local account passwords are stored in a format protected by bcrypt; command-line passwords are stored in a format protected by SHA-2.	Authenticates administrative access to the management UI and APIs.	—

Feature	Function	Impact	Regulatory compliance
SAML support	StorageGRID supports single sign-on (SSO) using the Security Assertion Markup Language 2.0 (SAML 2.0) standard. When SSO is enabled, all users must be authenticated by an external identity provider before they can access the Grid Manager, the Tenant Manager, the Grid Management API, or the Tenant Management API. Local users cannot sign in to StorageGRID.	Enables additional levels of security for grid and tenant administrators such as SSO and multifactor authentication (MFA).	NIST SP800-63
Granular permission control	Grid administrators can assign permissions to roles and assign roles to administrative user groups, which enforces which tasks administrative clients are allowed to perform by using both the management UI and APIs.	Allows Grid administrators to manage access control for admin users and groups.	—

Feature	Function	Impact	Regulatory compliance
Distributed audit logging	<p>StorageGRID provides a built-in, distributed audit logging infrastructure, scalable to hundreds of nodes across up to 16 sites. StorageGRID software nodes generate audit messages, which are transmitted through a redundant audit relay system and ultimately captured in one or more audit log repositories. Audit messages capture events at an object-level granularity such as client-initiated S3 API operations, object lifecycle events by ILM, background object health checks, and configuration changes made from the management UI or APIs.</p> <p>Audit logs can be exported by syslog allowing audit messages to be mined by tools such as Splunk and ELK. There are four types of audit messages:</p> <ul style="list-style-type: none"> • System audit messages • Object storage audit messages • HTTP protocol audit messages • Management audit messages <p>Audit logs can be stored in a S3 bucket for long term retention and application access.</p>	<p>Provides Grid administrators with a proven and scalable audit service and enables them to mine audit data for various objectives. Such objectives include troubleshooting, auditing SLA performance, client data access API operations, and management configuration changes.</p>	—
System audit	<p>System audit messages capture system-related events, such as grid node states, corrupt object detection, objects committed at all specified locations per ILM rule, and progress of system-wide maintenance tasks (grid tasks).</p>	<p>Helps customers troubleshoot system issues and provides proof that objects are stored according to their SLA. SLAs are implemented by StorageGRID ILM rules and are integrity-protected.</p>	—

Feature	Function	Impact	Regulatory compliance
Object storage audit	<p>Object storage audit messages capture object API transaction and lifecycle-related events. These events include object storage and retrieval, grid-node to grid-node transfers, and verifications.</p>	<p>Helps customers audit the progress of data through the system and whether SLA, specified as StorageGRID ILM, are being delivered.</p>	—
HTTP protocol audit	<p>HTTP protocol audit messages capture HTTP protocol interactions related to client applications and StorageGRID nodes. In addition, customers can capture specific HTTP request headers (such as X-Forwarded-For and user metadata [x-amz-meta-*]) into audit.</p>	<p>Helps customers audit data access API operations between clients and StorageGRID and trace an action to an individual user account and access key. Customers can also log user metadata into audit and use log mining tools, such as Splunk or ELK, to search on object metadata.</p>	—
Management audit	<p>Management audit messages log admin user requests to the management UI (Grid Management Interface) or APIs. Every request that is not a GET or HEAD request to the API logs a response with the username, IP, and type of request to the API.</p>	<p>Helps Grid administrators establish a record of system configuration changes made by which user from which source IP and which destination IP at what time.</p>	—
TLS 1.3 support for management UI and API access	<p>TLS establishes a handshake protocol for communication between an admin client and a StorageGRID admin node.</p>	<p>Enables an administrative client and StorageGRID to identify and authenticate each other and communicate with confidentiality and data integrity.</p>	—

Feature	Function	Impact	Regulatory compliance
SNMPv3 for StorageGRID monitoring	<p>SNMPv3 provides security by offering both strong authentication and data encryption for privacy. With v3, protocol data units are encrypted, using CBC-DES for its encryption protocol.</p> <p>User authentication of who sent the protocol data unit is provided by either the HMAC-SHA or HMAC-MD5 authentication protocol.</p> <p>SNMPv2 and v1 are still supported.</p>	Helps Grid administrators monitor the StorageGRID system by enabling an SNMP agent on the Admin Node.	—
Client certificates for Prometheus metrics export	Grid administrators can upload or generate client certificates which can be used to provide secure, authenticated access to the StorageGRID Prometheus database.	Grid administrators can use client certificates to monitor StorageGRID externally using applications such as Grafana.	—

Platform security features

Learn about the platform security features in StorageGRID.

Feature	Function	Impact	Regulatory compliance
Internal public-key infrastructure (PKI), node certificates, and TLS	StorageGRID uses an internal PKI and node certificates to authenticate and encrypt internode communication. Internode communication is secured by TLS.	Helps secure system traffic over the LAN or WAN, especially in a multisite deployment.	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)
Node firewall	StorageGRID automatically configures IP tables and firewalling rules to control incoming and outgoing network traffic, as well as closing unused ports.	Helps protect the StorageGRID system, data, and metadata against unsolicited network traffic.	—

Feature	Function	Impact	Regulatory compliance
OS hardening	The base operating system of StorageGRID physical appliances and virtual nodes is hardened; unrelated software packages are removed.	Helps minimize potential attack surfaces.	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)
Periodic platform and software updates	StorageGRID provides regular software releases that include operating system, applications binaries, and software updates.	Helps keep the StorageGRID system updated with current software and application binaries.	—
Disabled Root Login Over Secure Shell (SSH)	Root login over SSH is disabled on all StorageGRID nodes. SSH access uses certificate authentication.	Helps customers protect against potential remote password cracking of the root login.	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)
Automated time synchronization	StorageGRID automatically synchronizes system clocks of each node against multiple external time Network Time Protocol (NTP) servers. At least four NTP servers of Stratum 3 or later are required.	Ensures the same time reference across all nodes.	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)
Separate networks for client, admin, and internal grid traffic	StorageGRID software nodes and hardware appliances support multiple virtual and physical network interfaces, so that customers can separate client, administration, and internal grid traffic over different networks.	Allow Grid administrators to segregate internal and external network traffic and deliver traffic over networks with different SLAs.	—
Multiple virtual LAN (VLAN) interfaces	StorageGRID supports configuring VLAN interfaces on your StorageGRID client and grid networks.	Allow Grid administrators to partition and isolate application traffic for security, flexibility, and performance.	—
Untrusted Client Network	The Untrusted Client Network interface accepts inbound connections only on ports that have been explicitly configured as load-balancer endpoints.	Ensures that interfaces exposed to untrusted networks are secured.	—
Configurable Firewall	Manage open and closed ports for Admin, Grid, and client networks.	Allow grid administrators to control access on ports and manage approved device access to the ports.	—

Feature	Function	Impact	Regulatory compliance
Enhanced SSH behavior	disable SSH by default prior to installation. In the default state, SSH access is only enabled on the link-local management ports address. The admin and root user passwords are set to the appliance compute controller serial number. Login is only allowed on serial console and graphical console (BMC KVM). SSH on any network port is disabled.	Enhances network access protection.	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)
Node encryption	As part of the new KMS host server encryption feature, a new Node Encryption setting is added to the StorageGRID Appliance Installer.	This setting must be enabled during the hardware configuration stage of appliance installation.	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)

Cloud integration

Understand how StorageGRID integrates with cloud services.

Feature	Function	Impact
Notifications-based virus scanning	StorageGRID platform services support event notifications. Event notifications can be used with external cloud computing services to trigger virus scanning workflows on the data.	Allows tenant administrators to trigger virus scanning of data using external cloud computing services.

TR-4921: Ransomware defense

Protect StorageGRID S3 objects from ransomware

Learn about ransomware attacks and how to protect data with StorageGRID security best practices.

Ransomware attacks are on the rise. This document provides some recommendations on how to protect your object data on StorageGRID.

Ransomware today is the ever-present danger in the data center. Ransomware is designed to encrypt data and make it unusable by the users and applications that rely on it. Protection starts with the usual defenses of hardened networking and solid user security practices, and we need to follow through with data access security practices.

Ransomware is one of today's largest security threats. The NetApp StorageGRID team is working with our

customers to keep ahead of these threats. With the use of object lock and versioning, you can protect against unwanted alterations and recover from malicious attacks. Data security is a multi-layer venture, with your object storage being just one part in your data center.

StorageGRID best practices

For StorageGRID, security best practices should include using HTTPS with signed certificates for both management and object access. Create dedicated user accounts for applications and individuals, and do not use the tenant root accounts for application access or user data access. In other words, follow the least privilege principle. Use security groups with defined Identity and Access Management (IAM) policies to govern user rights, and access accounts specific to the applications and users. With these measures in place, you still must ensure that your data is protected. In the case of Simple Storage Service (S3), when objects are modified to encrypt them, it is accomplished by an overwrite of the original object.

Methods of defense

The primary ransomware protection mechanism in the S3 API is to implement object lock. Not all applications are compatible with object lock, so there are two other options to protect your objects that are described in this report: replication to another bucket with versioning enabled and versioning with IAM policies.

Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- NetApp StorageGRID Documentation Center
<https://docs.netapp.com/us-en/storagegrid/>
- NetApp StorageGRID Enablement
<https://docs.netapp.com/us-en/storagegrid-enable/>
- NetApp Product Documentation
<https://www.netapp.com/support-and-training/documentation/>

Ransomware defense using object lock

Explore how object lock in StorageGRID provides a WORM model to prevent data deletion or overwrite, and how it meets regulatory requirements.

Object lock provides a WORM model to prevent objects from being deleted or overwritten. StorageGRID implementation of object lock is [Cohasset assessed](#) to help meet regulatory requirements, supporting legal hold, compliance mode, and governance mode for object retention, and default bucket retention policies. You must enable object lock as part of the bucket creation and versioning. A specific version of an object is locked, and if no version ID is defined, the retention is placed on the current version of the object. If the current version has the retention configured and an attempt is made to delete, modify, or overwrite the object, a new version is created with either a delete marker, or the new revision of the object as the current version, and the locked version is retained as a non-current version. For applications that are not yet compatible, you might still be able to make use of object lock and a default retention configuration placed on the bucket. After the configuration is defined, this applies an object retention to each new object put into the bucket. This works as long as the application is configured to not delete or overwrite the objects before the retention time has passed.

When creating a bucket in the Tenant management UI, you can enable object lock and configure a default retention mode and retention period. When configured this will set a minimum object lock retention on every object that is ingested to that bucket.

S3 Object Lock

Allows you to specify retention and legal hold settings for the objects ingested into a bucket. If you want to use S3 Object Lock, you must enable this setting when you create the bucket. You cannot add or disable S3 Object Lock after a bucket is created.

If S3 Object Lock is enabled, object versioning is enabled for the bucket automatically and cannot be suspended.

Enable S3 Object Lock

Default retention

Disable
 New objects added to the bucket will not be protected from being deleted or overwritten. Does not apply to objects already in the bucket or to objects that have their own retain-until-dates.

Enable
 New objects added to the bucket will be protected from being deleted or overwritten based on the default retention mode and period you specify below. Does not apply to objects already in the bucket or to objects that have their own retain-until-dates.

Default retention mode

Governance
Users with special permissions can change an object's retention settings or they can override these settings to delete the object.

Compliance
No users can overwrite or delete protected object versions during the retention period.

Default retention period [?](#)

90 Days ▼

Maximum retention period on this tenant: 100 years

Here are a few examples using the object lock API:

Object lock legal hold is a simple on/off status applied to an object.

```
aws s3api put-object-legal-hold --bucket mybucket --key myfile.txt --legal-hold Status=ON --endpoint-url https://s3.company.com
```

Setting the legal hold status does not return any value if successful, so it can be verified with a GET operation.

```
aws s3api get-object-legal-hold --bucket mybucket --key myfile.txt --endpoint-url https://s3.company.com
{
    "LegalHold": {
        "Status": "ON"
    }
}
```

To turn legal hold off, apply the OFF status.

```
aws s3api put-object-legal-hold --bucket mybucket --key myfile.txt --legal
-hold Status=OFF --endpoint-url https://s3.company.com
aws s3api get-object-legal-hold --bucket mybucket --key myfile.txt
--endpoint-url https://s3.company.com
{
  "LegalHold": {
    "Status": "OFF"
  }
}
```

Setting the object retention is done with a retain until timestamp.

```
aws s3api put-object-retention --bucket mybucket --key myfile.txt
--retention '{"Mode":"COMPLIANCE", "RetainUntilDate": "2022-06-
10T16:00:00"}' --endpoint-url https://s3.company.com
```

Again, there is no returned value on success, so you can verify the retention status similarly with a get call.

```
aws s3api get-object-retention --bucket mybucket --key myfile.txt
--endpoint-url https://s3.company.com
{
  "Retention": {
    "Mode": "COMPLIANCE",
    "RetainUntilDate": "2022-06-10T16:00:00+00:00"
  }
}
```

Putting a default retention on an object lock enabled bucket uses a retention period in days and years.

```
aws s3api put-object-lock-configuration --bucket mybucket --object-lock
-configuration '{ "ObjectLockEnabled": "Enabled", "Rule": {
  "DefaultRetention": { "Mode": "COMPLIANCE", "Days": 1 } } }' --endpoint-url
https://s3.company.com
```

As with most of these operations, no response is returned on success so, we can perform a GET for the configuration to verify.

```
aws s3api get-object-lock-configuration --bucket mybucket --endpoint-url https://s3.company.com
{
  "ObjectLockConfiguration": {
    "ObjectLockEnabled": "Enabled",
    "Rule": {
      "DefaultRetention": {
        "Mode": "COMPLIANCE",
        "Days": 1
      }
    }
  }
}
```

Next, you can put an object in the bucket with the retention configuration applied.

```
aws s3 cp myfile.txt s3://mybucket --endpoint-url https://s3.company.com
```

The PUT operation does return a response.

```
upload: ./myfile.txt to s3://mybucket/myfile.txt
```

On the retention object, the retention duration set on the bucket in the preceding example is converted to a retention timestamp on the object.

```
aws s3api get-object-retention --bucket mybucket --key myfile.txt
--endpoint-url https://s3.company.com
{
  "Retention": {
    "Mode": "COMPLIANCE",
    "RetainUntilDate": "2022-03-02T15:22:47.202000+00:00"
  }
}
```

Ransomware defense using replicated bucket with versioning

Learn how to replicate objects to a secondary bucket using StorageGRID CloudMirror.

Not all applications and workloads are going to be compatible with object lock. Another option is to replicate the objects to a secondary bucket either in the same grid (preferably a different tenant with restricted access), or any other S3 endpoint with the StorageGRID platform service, CloudMirror.

StorageGRID CloudMirror is a component of StorageGRID that can be configured to replicate the objects of a bucket to a defined destination as they are ingested into the source bucket and does not replicate deletes.


Because CloudMirror is an integrated component of StorageGRID, it cannot be turned off or manipulated by an S3 API-based attack. You can configure this replicated bucket with versioning enabled. In this scenario you need some automated cleanup of the replicated bucket's old versions that are safe to discard. For this, you can use the StorageGRID ILM policy engine. Create rules to manage the object placement based on non-current time for several days sufficient to have identified and recovered from an attack.

One downside to this approach is that it consumes more storage by having a complete second copy of the bucket plus multiple versions of the objects retained for some time. Additionally, the objects that were intentionally deleted from the primary bucket must be manually removed from the replicated bucket. There are other replication options outside of the product, such as NetApp CloudSync, that can replicate deletes for a similar solution. Another downside for the secondary bucket being versioning enabled and not object lock enabled is that there exists a number of privileged accounts that might be used to cause damage on the secondary location. The advantage is that it should be a unique account to that endpoint or tenant bucket and the compromise likely does not include access to accounts on the primary location or vice-versa.

After the source and destination buckets are created and the destination is configured with versioning, you can configure and enable replication, as follows:

Steps

1. To configure CloudMirror, create a platform services endpoint for the S3 destination.



Enter endpoint details

Enter the endpoint's display name, URI, and URN.

Display name ?

MyGrid

URI ?

<https://s3.company.com>

URN ?

arn:aws:s3:::mybucket

2. On the source bucket, configure replication to use the endpoint configured.

```
<ReplicationConfiguration>
  <Role></Role>
  <Rule>
    <Status>Enabled</Status>
    <Prefix></Prefix>
    <Destination>
      <Bucket>arn:aws:s3:::mybucket</Bucket>
      <StorageClass>STANDARD</StorageClass>
    </Destination>
  </Rule>
</ReplicationConfiguration>
```

3. Create ILM rules to manage the storage placement and version storage duration management. In this example, the non-current versions of the objects to store are configured.

Create ILM Rule Step 1 of 3: Define Basics

Name	MyTenant - version retention
Description	retain non-current versions for 30 days
Tenant Accounts (optional) <small>?</small>	mytenant (26261433202363150471) <small>X</small>
Bucket Name	contains <input type="text" value="mybucket"/>

Create ILM Rule Step 2 of 3: Define Placements

Configure placement instructions to specify how you want objects matched by this rule to be stored.

MyTenant - version retention
retain non-current versions for 30 days

A rule that uses Noncurrent Time only applies to noncurrent versions of 53 objects.
You cannot use this rule as the default rule in an ILM policy because it does not apply to current object versions.

Reference Time **Noncurrent Time**

Placements **Sort by start day**

From day 0 store for 30 days **Add** **Remove**

Type replicated Location site1 **Add Pool** Copies 2 **Temporary location** **Optional** **+** **-**

Retention Diagram **Refresh**

Trigger Day 0 Day 30

site1 F-Cluster

Duration 30 days

There are two copies in site 1 for 30 days. You also configure the rules for the current version of the objects based on using ingest time as reference time in the ILM rule to match the source bucket storage duration. The storage placement for the object versions can be erasure coded or replicated.

Ransomware defense using versioning with protective IAM policy

Learn how to protect your data by enabling versioning on the bucket and implementing IAM policies on user security groups in StorageGRID.

A method to protect your data without using object lock or replication is to enable versioning on the bucket and implement IAM policies on the user security groups to limit users' ability to manage versions of the objects. In the event of an attack, new bad versions of the data are created as the current version, and the most recent non-current version is the safe clean data. The accounts compromised to gain access to the data do not have access to delete or otherwise alter the non-current version protecting it for later restore operations. Just like the previous scenario, ILM rules manage the retention of the noncurrent versions with a duration of your choice. The downside is that there is still the possibility of privileged accounts existing for a bad actor attack, but all application service accounts and users must be configured with a more restrictive access. The restrictive group policy must explicitly allow each action you want the users or application to be capable of and explicitly deny any actions that you do not want them to be capable of. NetApp does not recommend using a wildcard allow because a new action might be introduced in the future, and you will want to control whether it is allowed or denied. For this solution, the deny list must include DeleteObjectVersion, PutBucketPolicy, DeleteBucketPolicy, PutLifecycleConfiguration, and PutBucketVersioning to protect the versioning configuration of the bucket and object versions from user or programmatic changes.

In StorageGRID The S3 group policy option "Ransomware Mitigation" makes implementing this solution easier. When creating a user group in the tenant, after selecting the group permissions, you can see this optional policy.

Create group

Choose a group type — Manage permissions — Set S3 group policy — Add users

Set S3 group policy

An S3 group policy controls user access permissions to specific S3 resources, including buckets. Non-root users have no access by default.

- No S3 Access
- Read Only Access
- Full Access
- Ransomware Mitigation 
- Custom

```
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "s3:CreateBucket",
        "s3:DeleteBucket",
        "s3:DeleteReplicationConfiguration",
        "s3:DeleteBucketMetadataNotification",
        "s3:GetBucketAcl",
        "s3:GetBucketCompliance",
```

Previous

Continue

The following is the content of the group policy that includes most of the available operations explicitly allowed and the minimum required denied.

```

        "s3:GetEncryptionConfiguration",
        "s3:GetLifecycleConfiguration",
        "s3>ListBucket",
        "s3>ListBucketVersions",
        "s3>ListAllMyBuckets",
        "s3>ListBucketMultipartUploads",
        "s3:PutBucketConsistency",
        "s3:PutBucketLastAccessTime",
        "s3:PutBucketNotification",
        "s3:PutBucketObjectLockConfiguration",
        "s3:PutReplicationConfiguration",
        "s3:PutBucketCORS",
        "s3:PutBucketMetadataNotification",
        "s3:PutBucketTagging",
        "s3:PutEncryptionConfiguration",
        "s3:AbortMultipartUpload",
        "s3>DeleteObject",
        "s3>DeleteObjectTagging",
        "s3>DeleteObjectVersionTagging",
        "s3:GetObject",
        "s3:GetObjectAcl",
        "s3:GetObjectLegalHold",
        "s3:GetObjectRetention",
        "s3:GetObjectTagging",
        "s3:GetObjectVersion",
        "s3:GetObjectVersionAcl",
        "s3:GetObjectVersionTagging",
        "s3>ListMultipartUploadParts",
        "s3:PutObject",
        "s3:PutObjectAcl",
        "s3:PutObjectLegalHold",
        "s3:PutObjectRetention",
        "s3:PutObjectTagging",
        "s3:PutObjectVersionTagging",
        "s3:RestoreObject",
        "s3:ValidateObject",
        "s3:PutBucketCompliance",
        "s3:PutObjectVersionAcl"
    ],
    "Resource": "arn:aws:s3:::/*"
},
{
    "Effect": "Deny",
    "Action": [
        "s3>DeleteObjectVersion",
        "s3>DeleteBucketPolicy",

```

```

        "s3:PutBucketPolicy",
        "s3:PutLifecycleConfiguration",
        "s3:PutBucketVersioning"
    ],
    "Resource": "arn:aws:s3:::/*"
}
]
}

```

Ransomware investigation and remediation


Learn how to investigate and remediate buckets after a possible ransomware attack with StorageGRID.

In StorageGRID 12.0, the new branch bucket feature has been added to extend the usefulness of versioning for ransomware defense. A branch bucket provides access to objects in a bucket as they existed at a certain time provided they still exist in the bucket. Branch buckets can only be created for versioning-enabled base buckets.

This means if you suspect a ransomware attack has occurred, you can create a read/write, or read-only branch bucket containing all objects and versions that existed prior to the initial attack time. You can use this branch bucket to compare against the base bucket contents to figure out what objects have changed and if the change was part of the attack or not. You could also use a branch bucket to continue client operations using the clean branch while investigating the attack.

Creating a Branch bucket

- Navigate to the base bucket details page and the Branches tab to create a branch bucket.



The screenshot shows the StorageGRID Tenant Manager interface. The left sidebar shows 'My access keys', 'Buckets' (selected), 'Groups', 'Users', and 'Identity federation'. The main content area shows a 'base-bucket' with details: Region: us-east-1, Date created: 2025-06-25 14:01:49 IST, Object count: 0, and metrics: Space used: 0 bytes, Capacity limit: —, Object count limit: —. Below these are buttons for 'Delete objects in bucket' and 'Delete bucket'. At the bottom of the main area, tabs for 'S3 Console', 'Bucket options', 'Bucket access', and 'Branches' are shown, with 'Branches' being the active tab. A red box highlights the 'Create branch bucket' button. Below the tabs, a section titled 'Branch buckets for base-bucket' contains a table with one result: branch-bucket-1, Read-write, 2025-06-25 14:05:21 IST, 2025-06-25 14:06:07 IST. A red box highlights the search bar 'Search branch bucket name'. A note at the bottom says: 'A branch bucket provides access to objects in a bucket as they existed at a certain time. A branch bucket provides access to protected data, but doesn't serve as a backup. To continue to protect data, use these features on base buckets: S3 Object Lock, cross-grid replication for base buckets, or bucket policies for versioned buckets to clean up old object versions.' At the bottom right, there are 'Previous' and 'Next' buttons.

- Once the Create branch bucket button is clicked, a popup will open with pre-filled details of the region associated with the base bucket.
- provide the branch bucket name, before time, and select what type of branch bucket to create.

Create branch bucket of base-bucket

1 Enter details 2 Manage settings Optional

Enter branch bucket details

Branch bucket name ?

Required

Region ?

us-east-1

Before time ?

6/25/2025 03 : 04 IST

Branch bucket type

Read-write
In the branch bucket, you can add or delete objects or object versions.

Read-only
In the branch bucket, you can't modify objects. In the user interface, bucket settings related to the modification of objects will be disabled.

Cancel Continue

TR-4765: Monitor StorageGRID

Introduction to StorageGRID monitoring

Learn how to monitor your StorageGRID system by using external applications, such as Splunk.

Effective monitoring of NetApp StorageGRID object-based storage enables administrators to quickly respond to urgent issues and to proactively add resources to handle growing workloads. This report provides general guidance about how to monitor key metrics and how to leverage external monitoring applications. It is meant to supplement the existing Monitoring and Troubleshooting guide.

A NetApp StorageGRID deployment typically consists of multiple sites and many nodes that operate to create a distributed and fault-tolerant object storage system. In a distributed and resilient storage system such as

StorageGRID, it is normal for error conditions to exist while the grid continues to operate normally. The challenge for you as an administrator is to understand the threshold at which error conditions (such as nodes down) present a problem that should be immediately addressed versus information that should be analyzed. By analyzing the data that StorageGRID presents, you can understand your workload and make informed decisions, such as when to add more resources.

StorageGRID provides excellent documentation that dives deep into the subject of monitoring. This report assumes that you are familiar with StorageGRID and that you have reviewed the documentation about it. Rather than repeating that information, we refer to the product documentation throughout this guide. StorageGRID product documentation is available online and in PDF format.

The goal of this document is to complement the product documentation and discuss how to monitor your StorageGRID system by using external applications, such as Splunk.

Data sources

To successfully monitor NetApp StorageGRID, it is important to know where to gather data about the health and operations of your StorageGRID system.

- **Web UI and Dashboard.** The StorageGRID Grid Manager presents a top-level view of the information that you as an administrator need to see in a logical presentation. As an administrator, you can also dig deeper into service-level information for troubleshooting and log collections.
- **Audit Logs.** StorageGRID keeps granular audit logs of tenant actions such as PUT, GET, and DELETE. You can also trace the lifecycle of an object from ingest to the application of data management rules.
- **Metrics API.** Underlying the StorageGRID GMI are open APIs, as the UI is API-driven. This approach enables you to extract data by using external monitoring and analysis tools.

Where to find additional information


To learn more about the information that is described in this document, review the following documents and/or websites:

- NetApp StorageGRID Documentation Center
<https://docs.netapp.com/us-en/storagegrid-118/>
- NetApp StorageGRID Enablement
<https://docs.netapp.com/us-en/storagegrid-enable/>
- NetApp Product Documentation
<https://www.netapp.com/support-and-training/documentation/>
- NetApp StorageGRID App for Splunk
<https://splunkbase.splunk.com/app/3898/#/details>

Use the GMI dashboard to monitor StorageGRID

The StorageGrid Grid Management Interface (GMI) dashboard provides a centralized view of the StorageGRID infrastructure, allowing you to oversee the health, performance, and capacity of the entire grid.

Use the GMI dashboard to examine each core component of the grid.



Information that you should monitor regularly

A previous version of this technical report listed the metrics to check periodically versus trends. That information is now included in the [Monitoring and Troubleshooting guide](#).

Monitor storage

A previous version of this technical report listed where to monitor important metrics, such as Object Storage Space, Metadata Space, Network Resources and so on. That information is now included in the [Monitoring and Troubleshooting guide](#).

Use alerts to monitor StorageGRID

Learn how to use the alerts system in StorageGRID to monitor issues, manage custom alerts, and extend alert notifications using SNMP or email.

Alerts provide critical information that allow you to monitor the various events and conditions within your StorageGRID system.

The alerts system is designed to be the primary tool for monitoring any issues that might occur in your StorageGRID system. The alerts system focuses on actionable problems in the system and provides an easy-to-use interface.

We provide a variety of default alerting rules that aim to help monitor and troubleshoot your system. You can further manage alerts by creating custom alerts, editing or disabling default alerts, and silencing alert notifications.

Alerts are also extensible through SNMP or email notification.

For more information on alerts, see the [product documentation](#) available online and in PDF format.

Advanced monitoring in StorageGRID

Learn how to access and export metrics to help troubleshoot issues.

View metrics API through a Prometheus query

Prometheus is an open-source software for collecting metrics. To access StorageGRID's embedded Prometheus through the GMI, go to [Support > Metrics](#).

Metrics

Access charts and metrics to help troubleshoot issues.

ⓘ The tools available on this page are intended for use by technical support. Some features and menu items within these tools are intentionally non-functional.

Prometheus

Prometheus is an open-source toolkit for collecting metrics. The Prometheus interface allows you to query the current values of metrics and to view charts of the values over time.

Access the Prometheus UI using the link below. You must be signed in to the Grid Manager.

- <https://webscalegmi.netapp.com/metrics/graph>

Grafana

Grafana is open-source software for metrics visualization. The Grafana interface provides pre-constructed dashboards that contain graphs of important metric values over time.


Access the Grafana dashboards using the links below. You must be signed in to the Grid Manager.

ADE
Account Service Overview
Alertmanager
Audit Overview
Cassandra Cluster Overview
Cassandra Network Overview
Cassandra Node Overview
Cloud Storage Pool Overview
EC Read (11.3) - Node
EC Read (11.3) - Overview

Grid
ILM
Ingests
Node
Node (Internal Use)
Platform Services Commits
Platform Services Overview
Platform Services Processing
Renamed Metrics

Replicated Read Path Overview
S3 - Node
S3 Overview
Site
Streaming EC - ADE
Streaming EC - Chunk Service
Support
Traces
Traffic Classification Policy
Virtual Memory (vmstat)

Alternatively, you can navigate to the link directly.




With this view, you can access the Prometheus interface. From there, you can search through available metrics and even experiment with queries.

To make a Prometheus URL query, follow these steps:

Steps

1. Start typing in the query text box. As you type, metrics are listed. For our purposes, only metrics that start with StorageGRID and Node are important.
2. To see the number of HTTP sessions for each node, type `storagegrid_http` and select `storagegrid_http_sessions_incoming_currently_established`. Click Execute and display the information in a graph or console format.



ⓘ Queries and charts that you build through this URL do not persist. Complex queries consume resources on the admin node. NetApp recommends that you use this view to explore available metrics.

ⓘ It is not recommended to directly interface to our Prometheus instance because this requires opening additional ports. Accessing metrics through our API is the recommended and secure method.

Export metrics through the API

You can also access the same data through the StorageGRID management API.

To export metrics through the API, follow these steps:

1. From the GMI, select **Help > API Documentation**.
2. Scroll down to Metrics and select **GET /grid/metric-query**.

GET	/grid/metric-labels/{label}/values	Lists the values for a metric label	
GET	/grid/metric-names	Lists all available metric names	
GET	/grid/metric-query	Performs an instant metric query at a single point in time	

The format of metric queries is controlled by Prometheus. See <https://prometheus.io/docs/querying/basics>

Parameters

Cancel

Name	Description
query <small>* required</small>	Prometheus query string (query) storagegrid_http_sessions_incoming_current
time	query start, default current time (date-time) (query) time - query start, default current time (date-ti)
timeout	timeout (duration) (query) 120s

Execute

Clear

The response includes the same information that you can obtain through a Prometheus URL query. You can again see the number of HTTP sessions that are currently established on each storage node. You can also download the response in JSON format for readability. The following figure shows sample Prometheus query responses.

Responses

Response content type: application/json

Curl

```
curl -X GET "https://10.193.92.230/api/v3/grid/metric-query?query=storagegrid_http_sessions_incoming_currently_established&timeout=120s" -H "accept: application/json" -H "X-Csrftoken: 0b94910621b19c120b4488d2e537e374"
```

Request URL

https://10.193.92.230/api/v3/grid/metric-query?query=storagegrid_http_sessions_incoming_currently_established&timeout=120s

Server response

Code	Details
200	Response body

```
{
  "responseTime": "2020-06-02T21:26:36.008Z",
  "status": "success",
  "apiVersion": "3.2",
  "data": {
    "resultType": "vector",
    "result": [
      {
        "metric": {
          "_name": "storagegrid_http_sessions_incoming_currently_established",
          "instance": "us-storage-1",
          "job": "storagegrid",
          "node_id": "a8cc5d26-b52a-4d78-95ec-0f21e76c61bd",
          "service": "ldn",
          "site_id": "fc56d838-cd56-423b-af07-ed8a3a2885d",
          "site_name": "us-east-fuse"
        },
        "value": [
          1591133196.007,
          "0"
        ]
      },
      {
        "metric": {
          "_name": "storagegrid_http_sessions_incoming_currently_established",
          "instance": "us-storage-2",
          "job": "storagegrid",
          "node_id": "8093353e-0fb9-49ca-b66b-b5744ad54bec",
          "service": "ldn",
          "site_id": "fc56d838-cd56-423b-af07-ed8a3a2885d",
          "site_name": "us-east-fuse"
        },
        "value": [
          1591133196.007,
          "0"
        ]
      }
    ]
  }
}
```

Download



The advantage of using the API is that it enables you to perform authenticated queries


Access metrics using cURL in StorageGRID

Learn how to access metrics through the CLI using cURL.

To perform this operation, you must first obtain an authorization token. To request a token, follow these steps:

Steps

1. From the GMI, select **Help > API Documentation**.
2. Scroll down to Auth to find operations on authorization. The following screenshot shows the parameters for the POST method.



The screenshot shows the GMI API Documentation for the auth /authorize POST method. The 'body' parameter is highlighted, showing a JSON example with 'username', 'password', 'cookie', and 'csrfToken' fields. The 'Parameter content type' is set to 'application/json'.

3. Click Try It Out and edit the body with your GMI username and password.
4. Click Execute.
5. Copy the cURL command that is provided in the cURL section and paste it in a terminal window. The command looks like the following:

```
curl -X POST "https://<Primary_Admin_IP>/api/v3/authorize" -H "accept: application/json" -H "Content-Type: application/json" -H "X-Csrf-Token: dc30b080e1ca9bc05ddb81104381d8c8" -d "{\"username\": \"MyUsername\", \"password\": \"MyPassword\", \"cookie\": true, \"csrfToken\": false}" -k
```



If your GMI password contains special characters, remember to use \ to escape special characters. For example, replace ! with \!

6. After you run the preceding cURL command, the output gives you an authorization token like the following example:

```
{"responseTime": "2020-06-03T00:12:17.031Z", "status": "success", "apiVersion": "3.2", "data": "8a1e528d-18a7-4283-9a5e-b2e6d731e0b2"}
```

Now you can use the authorization token string to access metrics through cURL. The process to access metrics is similar to the steps in section [Advanced monitoring in StorageGRID](#). However, for demonstration purposes, we show an example with GET /grid/metric-labels/{label}/values selected in the Metrics category.

7. As an example, the following cURL command with the preceding authorization token will list the site names in StorageGRID.

```
curl -X GET "https://10.193.92.230/api/v3/grid/metric-labels/site_name/values" -H "accept: application/json" -H "Authorization: Bearer 8a1e528d-18a7-4283-9a5e-b2e6d731e0b2"
```

The cURL command will generate the following output:

```
{"responseTime": "2020-06-03T00:17:00.844Z", "status": "success", "apiVersion": "3.2", "data": ["us-east-fuse", "us-west-fuse"]}
```

View metrics using the Grafana dashboard in StorageGRID

Learn how to use the Grafana interface to visualize and monitor your StorageGRID data.

Grafana is an open-source software for metric visualization. By default, we have preconstructed dashboards that provide useful and powerful information regarding your StorageGRID system.

These preconstructed dashboards are not only useful for monitoring but also for troubleshooting an issue. Some are intended for use by technical support. For example, to view the metrics of a storage node, follow these steps.

Steps

1. From the GMI, **Support > Metrics**.
2. Under the Grafana section, select the Node dashboard.

Grafana

Grafana is open-source software for metrics visualization. The Grafana interface provides pre-constructed dashboards that contain graphs of important metric values over time. Access the Grafana dashboards using the links below. You must be signed in to the Grid Manager.

ADE	Grid	Replicated Read Path Overview
Account Service Overview	ILM	S3 - Node
Alertmanager	Identity Service Overview	S3 Overview
Audit Overview	Ingests	Site
Cassandra Cluster Overview	Node	Streaming EC - ADE
Cassandra Network Overview	Node (Internal Use)	Streaming EC - Chunk Service
Cassandra Node Overview	Platform Services Commits	Support
Cloud Storage Pool Overview	Platform Services Overview	Traffic Classification Policy
EC Read - Node	Platform Services Processing	
EC Read - Overview	Renamed Metrics	

3. In Grafana, set the hosts to whichever node you want to view metrics on. In this case, a storage node is selected. More information is provided than the following screenshot captures.



Use traffic classification policies in StorageGRID

Learn how to set up and configure traffic classification policies to manage and optimize network traffic in StorageGRID.

Traffic Classification Policies provide a method to monitor and/or limit traffic based on a specific tenant, buckets, IP subnets, or load balancer endpoints. Network connectivity and bandwidth are especially important metrics for StorageGRID.

To configure a Traffic Classification Policy, follow these steps:

Steps

1. On the GMI, navigate to **Configuration > System Settings > Traffic Classification**.
2. Click **Create +**
3. Enter a name and description for your policy.

4. Create a matching rule.

Create Matching Rule


Matching Rules

Type ? Tenant

Tenant Jonathan.Wong (22497137670163214190) Change Account

Inverse Match ?

Cancel Apply



5. Set a limit (optional).

Create Limit

Limits (Optional)


Type ? -- Choose One --

Value ? Choose One

Value ? Aggregate Bandwidth In
Aggregate Bandwidth Out
Concurrent Read Requests
Concurrent Write Requests
Per-Request Bandwidth In
Per-Request Bandwidth Out
Read Request Rate
Write Request Rate

Traffic that matches any rule

Cancel Apply



6. Save your policy

Create Traffic Classification Policy

Policy

Name ? Match a Tenant

Description (optional)

Matching Rules

Traffic that matches any rule is included in the policy.

+ Create Edit Remove			
Type	Inverse Match	Match Value	
<input checked="" type="radio"/> Tenant		Jonathan.Wong (22497137670163214190)	
Displaying 1 matching rule.			

Limits (Optional)

+ Create Edit Remove			
Type	Value	Units	
		No limits found.	

Cancel

Save

To view the metrics associated to your Traffic Classification Policy, select your policy and click Metrics. A Grafana dashboard is generated displaying information such as Load Balancer Request Traffic and Average Request Duration.



Use audit logs to monitor StorageGRID

Learn how to use the StorageGRID audit log for detailed insights into tenant and grid activity, and how to leverage tools like Splunk for log analysis.

The StorageGRID audit log enables you to collect detailed information about tenant and grid activity. The audit log can be exposed for analytics through NFS. For detailed instructions on how to export the audit log, see the Administrator's Guide.

After the audit has been exported, you can use log analysis tools such as Splunk or Logstash + Elasticsearch to understand tenant activity or to create detailed billing and chargeback reports.

Details about audit messages are included in StorageGRID documentation. See [Audit messages](#).

Use the StorageGRID app for Splunk

Learn about the NetApp StorageGRID app for Splunk that allows you to monitor and analyze your StorageGRID environment within the Splunk platform.

Splunk is a software platform that imports and indexes machine data to provide powerful search and analysis features. The NetApp StorageGRID app is an add-on for Splunk that imports and enriches data leveraged from StorageGRID.

Instructions on how to install, upgrade and configure the StorageGRID add-on can be found here:
<https://splunkbase.splunk.com/app/3895/#/details>

TR-4882: Install a StorageGRID bare metal grid

Introduction to installing StorageGRID

Learn how to install StorageGRID on bare metal hosts.

TR-4882 provides a practical, step-by-step set of instructions that produces a working installation of NetApp StorageGRID. The installation could be either on bare metal or on virtual machines (VMs) running on Red Hat Enterprise Linux (RHEL). The approach is to perform an “opinionated” installation of six StorageGRID containerized services onto three physical (or virtual) machines in a suggested layout and storage configuration. Some customers might find it easier to understand the deployment process by following the example deployment in this TR.

For a more in-depth understanding about StorageGRID and the installation process, see
<https://docs.netapp.com/us-en/storagegrid-118/landing-install-upgrade/index.html> [Install, upgrade, and hotfix StorageGRID] in the product documentation.

Before you start your deployment, let's examine the compute, storage, and networking requirements for NetApp StorageGRID software. StorageGRID runs as a containerized service within Podman or Docker. In this model, some requirements refer to the host operating system (the OS that hosts Docker, which is running the StorageGRID software). And some of the resources are allocated directly to the Docker containers running within each host. In this deployment, in order to maximize hardware usage, we are deploying two services per physical host. For more information, continue on to the next section, [Prerequisites to install StorageGRID](#).

The steps outlined in this TR result in a working StorageGRID installation on six bare metal hosts. You now have a working grid and client network, which are useful in most testing scenarios.

Where to find additional information

To learn more about the information that is described in this TR, review the following documentation resources:

- NetApp StorageGRID Documentation Center
<https://docs.netapp.com/us-en/storagegrid-118/>
- NetApp StorageGRID Enablement
<https://docs.netapp.com/us-en/storagegrid-enable/>
- NetApp Product Documentation
<https://www.netapp.com/support-and-training/documentation/>

Prerequisites to install StorageGRID

Learn about the compute, storage, network, docker, and node requirements to deploy StorageGRID.

Compute requirements

The table below lists the supported minimum resource requirements for each type of StorageGRID node. These are the minimum resources required for StorageGRID nodes.

Type of node	CPU cores	RAM
Admin	8	24GB
Storage	8	24GB
Gateway	8	24GB

In addition, each physical Docker host should have a minimum of 16GB of RAM allocated to it for proper operation. So, for example, to host any two of the services described in the table together on one physical Docker host, you would do the following calculation:

$$24 + 24 + 16 = 64\text{GB RAM}$$

and

$$8 + 8 = 16 \text{ cores}$$

Because many modern servers exceed these requirements, we combine six services (StorageGRID containers) onto three physical servers.

Networking requirements

The three types of StorageGRID traffic include:

- **Grid traffic (required).** The internal StorageGRID traffic that travels between all nodes in the grid.
- **Admin traffic (optional).** The traffic used for system administration and maintenance.
- **Client traffic (optional).** The traffic that travels between external client applications and the grid, including all object storage requests from S3 and Swift clients.

You can configure up to three networks for use with the StorageGRID system. Each network type must be on a separate subnet with no overlap. If all nodes are on the same subnet, a gateway address is not required.

For this evaluation, we will deploy on two networks, which contain the grid and client traffic. It is possible to add

an admin network later to serve that additional function.

It is very important to map the networks consistently to the interfaces throughout all of the hosts. For example, if there are two interfaces on each node, ens192 and ens224, they should all be mapped to the same network or VLAN on all hosts. In this installation, the installer maps these into the Docker containers as eth0@if2 and eth2@if3 (because the loopback is if1 inside the container), and therefore a consistent model is very important.

Note on Docker networking

StorageGRID uses networking differently from some Docker container implementations. It does not use the Docker (or Kubernetes or Swarm) provided networking. Instead, StorageGRID actually spawns the container as `--net=none` so that Docker doesn't do anything to network the container. After the container has been spawned by the StorageGRID service, a new macvlan device is created from the interface defined in the node configuration file. That device has a new MAC address and acts as a separate network device that can receive packets from the physical interface. The macvlan device is then moved into the container namespace and renamed to be one of either eth0, eth1, or eth2 inside the container. At that point the network device is no longer visible in the host OS. In our example, the grid network device is eth0 inside the Docker containers and the Client Network is eth2. If we had an admin network, the device would be eth1 in the container.

The new MAC address of the container network device might require promiscuous mode to be enabled in some network and virtual environments. This mode allows the physical device to receive and send packets for MAC addresses that differ from the known physical MAC address.



If running in VMWare vSphere, you must accept promiscuous mode, MAC address changes, and forged transmits in the port groups that will serve StorageGRID traffic when running RHEL. Ubuntu or Debian works without these changes in most circumstances.

Storage requirements

The nodes each require either SAN-based or local disk devices of the sizes shown in the following table.

The numbers in the table are for each StorageGRID service type, not for the entire grid or each physical host. Based on the deployment choices, we will calculate numbers for each physical host in [Physical host layout and requirements](#), later in this document.



The paths or file systems marked with an asterisk will be created in the StorageGRID container itself by the installer. No manual configuration or file system creation is required by the administrator, but the hosts need block devices to satisfy these requirements. In other words, the block device should appear by using the command `lsblk` but not be formatted or mounted within the host OS.

Node type	LUN purpose	Number of LUNs	Minimum size of LUN	Manual file system required	Suggested node config entry
All	Admin Node system space /var/local (SSD helpful here)	One for each Admin Node	90GB	No	<code>BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/ADM-VAR-LOCAL</code>

Node type	LUN purpose	Number of LUNs	Minimum size of LUN	Manual file system required	Suggested node config entry
All nodes	Docker storage pool at /var/lib/docker for container pool	One for each host (physical or VM)	100GB per container	Yes – ext4	NA – format and mount as host file system (not mapped into the container)
Admin	Admin Node audit logs (system data in Admin container) /var/local/audit/export	One for each Admin Node	200GB	No	BLOCK_DEVICE_AUDIT_LOGS =/dev/mapper/ADM-OS
Admin	Admin Node tables (system data in Admin container) /var/local/mysql_ibdata	One for each Admin Node	200GB	No	BLOCK_DEVICE_TABLES = /dev/mapper/ADM-MySQL
Storage nodes	Object storage (block devices) /var/local/rangedb0 (SSD helpful here) /var/local/rangedb1 /var/local/rangedb2	Three for each storage container	4000GB	No	BLOCK_DEVICE_RANGEDB_000 = /dev/mapper/SN-Db00 BLOCK_DEVICE_RANGEDB_001 = /dev/mapper/SN-Db01 BLOCK_DEVICE_RANGEDB_002 = /dev/mapper/SN-Db02

In this example, the disk sizes shown in the following table are needed per container type. The requirements per physical host are described in [Physical host layout and requirements](#), later in this document.

Disk sizes per container type

Admin container

Name	Size (GiB)
Docker-Store	100 (per container)
Adm-OS	90
Adm-Audit	200
Adm-MySQL	200

Storage container

Name	Size (GiB)
Docker-Store	100 (per container)
SN-OS	90
Rangedb-0	4096
Rangedb-1	4096
Rangedb-2	4096


Gateway container

Name	Size (GiB)
Docker-Store	100 (per container)
/var/local	90

Physical host layout and requirements

By combining the compute and network requirements shown in table above, you can get a basic set of hardware required for this installation of three physical (or virtual) servers with 16 cores, 64GB of RAM, and two network interfaces. If higher throughput is desired, it is possible to bond two or more interfaces on the grid or Client Network and use a VLAN-tagged interface such as bond0.520 in the node config file. If you expect more intense workloads, more memory for both the host and the containers is better.

As shown in the following figure, these servers will host six Docker containers, two per host. The RAM is calculated by providing 24GB per container and 16GB for the host OS itself.



Total RAM required per physical host (or VM) is $24 \times 2 + 16 = 64$ GB.
 The following tables list the required disk storage for hosts 1, 2, and 3.

Host 1	Size (GiB)
Docker Store	
/var/lib/docker (File system)	200 (100 x 2)
Admin container	
BLOCK_DEVICE_VAR_LOCAL	90
BLOCK_DEVICE_AUDIT_LOGS	200
BLOCK_DEVICE_TABLES	200

Host 1	Size (GiB)
Storage container	
SN-OS	90
/var/local (Device)	
Rangedb-0 (Device)	4096
Rangedb-1 (Device)	4096
Rangedb-2 (Device)	4096

Host 2	Size (GiB)
Docker Store	
/var/lib/docker (Shared)	200 (100 x 2)
Gateway container	
GW-OS */var/local	100
Storage container	
*/var/local	100
Rangedb-0	4096
Rangedb-1	4096
Rangedb-2	4096

Host 3	Size (GiB)
Docker Store	
/var/lib/docker (Shared)	200 (100 x 2)

Host 3	Size (GiB)
Gateway container	
*/var/local	100
Storage container	
*/var/local	100
Rangedb-0	4096
Rangedb-1	4096
Rangedb-2	4096

The Docker Store was calculated by allowing 100GB per /var/local (per container) x two containers = 200GB.

Preparing the nodes

To prepare for the initial installation of StorageGRID, first install RHEL version 9.2 and enable SSH. Set up network interfaces, Network Time Protocol (NTP), DNS, and the host name according to best practices. You need at least one enabled network interface on the grid network and another for the Client Network. If you are using a VLAN-tagged interface, configure it as per the examples below. Otherwise, a simple standard network interface configuration will suffice.

If you need to use a VLAN tag on the grid network interface, your configuration should have two files in /etc/sysconfig/network-scripts/ in the following format:

```
# cat /etc/sysconfig/network-scripts/ifcfg-enp67s0
# This is the parent physical device
TYPE=Ethernet
BOOTPROTO=none
DEVICE=enp67s0
ONBOOT=yes
# cat /etc/sysconfig/network-scripts/ifcfg-enp67s0.520
# The actual device that will be used by the storage node file
DEVICE=enp67s0.520
BOOTPROTO=none
NAME=enp67s0.520
IPADDR=10.10.200.31
PREFIX=24
VLAN=yes
ONBOOT=yes
```

This example assumes that your physical network device for the grid network is enp67s0. It could also be a bonded device such as bond0. Whether you are using bonding or a standard network interface, you must use the VLAN-tagged interface in your node configuration file if your network port does not have a default VLAN or if the default VLAN is not associated with the grid network. The StorageGRID container itself does not untag Ethernet frames, so it must be handled by the parent OS.

Optional storage setup with iSCSI

If you are not using iSCSI storage, you must ensure that host1, host2, and host3 contain block devices of sufficient size to meet their requirements. See [Disk sizes per container type](#) for host1, host2, and host3 storage requirements.

To set up storage with iSCSI, complete the following steps:

Steps

1. If you are using external iSCSI storage such as NetApp E-Series or NetApp ONTAP® data management software, install the following packages:

```
sudo yum install iscsi-initiator-utils
sudo yum install device-mapper-multipath
```

2. Find the initiator ID on each host.

```
# cat /etc/iscsi/initiatorname.iscsi
InitiatorName=iqn.2006-04.com.example.node1
```

3. Using the initiator name from step 2, map LUNs on your storage device (of the number and size shown in the [Storage requirements](#) table) to each storage node.
4. Discover the newly created LUNs with `iscsiadm` and log in to them.

```
# iscsiadm -m discovery -t st -p target-ip-address
# iscsiadm -m node -T iqn.2006-04.com.example:3260 -l
Logging in to [iface: default, target: iqn.2006-04.com.example:3260,
portal: 10.64.24.179,3260] (multiple)
Login to [iface: default, target: iqn.2006-04.com.example:3260, portal:
10.64.24.179,3260] successful.
```



For details, see [Creating an iSCSI Initiator](#) on the Red Hat Customer Portal.

5. To show the multipath devices and their associated LUN WWIDs, run the following command:

```
# multipath -ll
```

If you are not using iSCSI with multipath devices, simply mount your device by a unique path name that will persist device changes and reboots alike.

```
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0
```

Simply using `/dev/sdx` device names could cause issues later if devices are removed or added.



If you are using multipath devices, modify the `/etc/multipath.conf` file to use aliases as follows.



These devices might or might not be present on all nodes, depending on layout.

```

multipaths {
multipath {
wwid 36d039ea00005f06a000003c45fa8f3dc
alias Docker-Store
}
multipath {
wwid 36d039ea00006891b000004025fa8f597
alias Adm-Audit
}
multipath {
wwid 36d039ea00005f06a000003c65fa8f3f0
alias Adm-MySQL
}
multipath {
wwid 36d039ea00006891b000004015fa8f58c
alias Adm-OS
}
multipath {
wwid 36d039ea00005f06a000003c55fa8f3e4
alias SN-OS
}
multipath {
wwid 36d039ea00006891b000004035fa8f5a2
alias SN-Db00
}
multipath {
wwid 36d039ea00005f06a000003c75fa8f3fc
alias SN-Db01
}
multipath {
wwid 36d039ea00006891b000004045fa8f5af
alias SN-Db02
}
multipath {
wwid 36d039ea00005f06a000003c85fa8f40a
alias GW-OS
}
}

```

Before installing Docker in your host OS, format and mount the LUN or disk backing `/var/lib/docker`. The other LUNs are defined in the node config file and are used directly by the StorageGRID containers. That is, they do not show up in the host OS; they appear in the containers themselves, and those file systems are handled by the installer.

If you are using an iSCSI-backed LUN, place something similar to the following line in your `fstab` file. As noted,

the other LUNs do not need to be mounted in the host OS but must show up as available block devices.

```
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0 /var/lib/docker ext4
defaults 0 0
```

Preparing for Docker installation

To prepare for Docker installation, complete the following steps:

Steps

1. Create a file system on the Docker storage volume on all three hosts.

```
# sudo mkfs.ext4 /dev/sd?
```

If you are using iSCSI devices with multipath, use /dev/mapper/Docker-Store.

2. Create the Docker storage volume mount point:

```
# sudo mkdir -p /var/lib/docker
```

3. Add a similar entry for the docker-storage-volume-device to /etc/fstab.

```
/dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0 /var/lib/docker ext4
defaults 0 0
```

The following `_netdev` option is recommended only if you are using an iSCSI device. If you are using a local block device `_netdev` is not necessary and `defaults` is recommended.

```
/dev/mapper/Docker-Store /var/lib/docker ext4 _netdev 0 0
```

4. Mount the new file system and view disk usage.

```
# sudo mount /var/lib/docker
[root@host1]# df -h | grep docker
/dev/sdb 200G 33M 200G 1% /var/lib/docker
```

5. Turn off swap and disable it for performance reasons.

```
$ sudo swapoff --all
```

6. To persist the settings, remove all swap entries from /etc/fstab such as:

```
/dev/mapper/rhel-swap swap defaults 0 0
```



Failing to disable swap entirely can severely lower performance.

7. Perform a test reboot of your node to ensure that the `/var/lib/docker` volume is persistent and that all disk devices return.

Install Docker for StorageGRID

Learn how to install Docker for StorageGRID.

To install Docker, complete the following steps:

Steps

1. Configure the yum repo for Docker.

```
sudo yum install -y yum-utils
sudo yum-config-manager --add-repo \
https://download.docker.com/linux/rhel/docker-ce.repo
```

2. Install the needed packages.

```
sudo yum install docker-ce docker-ce-cli containerd.io
```

3. Start Docker.

```
sudo systemctl start docker
```

4. Test Docker.

```
sudo docker run hello-world
```

5. Make sure that Docker runs on system start.

```
sudo systemctl enable docker
```

Prepare node configuration files for StorageGRID

Learn how to prepare the node configuration files for StorageGRID.

At a high level, the node configuration process includes the following steps:

Steps

1. Create the /etc/storagegrid/nodes directory on all hosts.

```
sudo [root@host1 ~]# mkdir -p /etc/storagegrid/nodes
```

2. Create the needed files per physical host to match the container/node type layout. In this example, we created two files per physical host on each host machine.



The name of the file defines the actual node name for installation. For example, dc1-adm1.conf becomes a node named dc1-adm1.

— Host1:

```
dc1-adm1.conf  
dc1-sn1.conf
```

— Host2:

```
dc1-gw1.conf  
dc1-sn2.conf
```

— Host3:

```
dc1-gw2.conf  
dc1-sn3.conf
```

Preparing the node config files

The following examples use the /dev/disk/by-path format. You can verify the correct paths by running the following commands:

```
[root@host1 ~]# lsblk  
NAME MAJ:MIN RM SIZE RO TYPE MOUNTPOINT  
sda 8:0 0 90G 0 disk  
└─sda1 8:1 0 1G 0 part /boot  
└─sda2 8:2 0 89G 0 part  
└─rhel-root 253:0 0 50G 0 lvm /  
└─rhel-swap 253:1 0 9G 0 lvm  
└─rhel-home 253:2 0 30G 0 lvm /home  
sdb 8:16 0 200G 0 disk /var/lib/docker  
sdc 8:32 0 90G 0 disk  
sdd 8:48 0 200G 0 disk  
sde 8:64 0 200G 0 disk  
sdf 8:80 0 4T 0 disk  
sdg 8:96 0 4T 0 disk  
sdh 8:112 0 4T 0 disk  
sdi 8:128 0 90G 0 disk  
sr0 11:0 1 1024M 0 rom
```

And these commands:

```
[root@host1 ~]# ls -l /dev/disk/by-path/
total 0
lrwxrwxrwx 1 root root 9 Dec 21 16:42 pci-0000:02:01.0-ata-1.0 ->
../../sr0
lrwxrwxrwx 1 root root 9 Dec 21 16:42 pci-0000:03:00.0-scsi-0:0:0:0 ->
../../sda
lrwxrwxrwx 1 root root 10 Dec 21 16:42 pci-0000:03:00.0-scsi-0:0:0:0-part1
-> ../../sda1
lrwxrwxrwx 1 root root 10 Dec 21 16:42 pci-0000:03:00.0-scsi-0:0:0:0-part2
-> ../../sda2
lrwxrwxrwx 1 root root 9 Dec 21 16:42 pci-0000:03:00.0-scsi-0:0:1:0 ->
../../sdb
lrwxrwxrwx 1 root root 9 Dec 21 16:42 pci-0000:03:00.0-scsi-0:0:2:0 ->
../../sdc
lrwxrwxrwx 1 root root 9 Dec 21 16:42 pci-0000:03:00.0-scsi-0:0:3:0 ->
../../sdd
lrwxrwxrwx 1 root root 9 Dec 21 16:42 pci-0000:03:00.0-scsi-0:0:4:0 ->
../../sde
lrwxrwxrwx 1 root root 9 Dec 21 16:42 pci-0000:03:00.0-scsi-0:0:5:0 ->
../../sdf
lrwxrwxrwx 1 root root 9 Dec 21 16:42 pci-0000:03:00.0-scsi-0:0:6:0 ->
../../sdg
lrwxrwxrwx 1 root root 9 Dec 21 16:42 pci-0000:03:00.0-scsi-0:0:8:0 ->
../../sdh
lrwxrwxrwx 1 root root 9 Dec 21 16:42 pci-0000:03:00.0-scsi-0:0:9:0 ->
../../sdi
```

Example for primary Admin node

Example file name:

```
/etc/storagegrid/nodes/dcl-adm1.conf
```

Example file contents:



Disk paths can follow the examples below or use /dev/mapper/alias style naming. Do not use block device names such as /dev/sdb because they can change on reboot and cause great damage to your grid.

```
NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Primary
MAXIMUM_RAM = 24g
BLOCK_DEVICE_VAR_LOCAL = /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:2:0
BLOCK_DEVICE_AUDIT_LOGS = /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:3:0
BLOCK_DEVICE_TABLES = /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:4:0
GRID_NETWORK_TARGET = ens192
CLIENT_NETWORK_TARGET = ens224
GRID_NETWORK_IP = 10.193.204.43
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.193.204.1
CLIENT_NETWORK_CONFIG = STATIC
CLIENT_NETWORK_IP = 10.193.205.43
CLIENT_NETWORK_MASK = 255.255.255.0
CLIENT_NETWORK_GATEWAY = 10.193.205.1
```

Example for a storage node

Example file name:

```
/etc/storagegrid/nodes/dcl-sn1.conf
```

Example file contents:

```
NODE_TYPE = VM_Storage_Node
MAXIMUM_RAM = 24g
ADMIN_IP = 10.193.174.43
BLOCK_DEVICE_VAR_LOCAL = /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:9:0
BLOCK_DEVICE_RANGEDB_00 = /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:5:0
BLOCK_DEVICE_RANGEDB_01 = /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:6:0
BLOCK_DEVICE_RANGEDB_02 = /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:8:0
GRID_NETWORK_TARGET = ens192
CLIENT_NETWORK_TARGET = ens224
GRID_NETWORK_IP = 10.193.204.44
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.193.204.1
```

Example for gateway node

Example file name:

```
/etc/storagegrid/nodes/dcl-gw1.conf
```

Example file contents:

```
NODE_TYPE = VM_API_Gateway
MAXIMUM_RAM = 24g
ADMIN_IP = 10.193.204.43
BLOCK_DEVICE_VAR_LOCAL = /dev/disk/by-path/pci-0000:03:00.0-scsi-0:0:1:0
GRID_NETWORK_TARGET = ens192
CLIENT_NETWORK_TARGET = ens224
GRID_NETWORK_IP = 10.193.204.47
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 10.193.204.1
CLIENT_NETWORK_IP = 10.193.205.47
CLIENT_NETWORK_MASK = 255.255.255.0
CLIENT_NETWORK_GATEWAY = 10.193.205.1
```

Install StorageGRID dependencies and packages

Learn how to install StorageGRID dependencies and packages.

To install the StorageGRID dependencies and packages, run the following commands:

```
[root@host1 rpms]# yum install -y python-netaddr
[root@host1 rpms]# rpm -ivh StorageGRID-Webscale-Images-*.rpm
[root@host1 rpms]# rpm -ivh StorageGRID-Webscale-Service-*.rpm
```

Validate the StorageGRID configuration files

Learn how to validate the content of the configuration files for StorageGRID.

After you create the configuration files in /etc/storagegrid/nodes for each of your StorageGRID nodes, you must validate the contents of those files.

To validate the contents of the configuration files, run the following command on each host:

```
sudo storagegrid node validate all
```

If the files are correct, the output shows PASSED for each configuration file:

```
Checking for misnamed node configuration files... PASSED
Checking configuration file for node dc1-adm1... PASSED
Checking configuration file for node dc1-gw1... PASSED
Checking configuration file for node dc1-sn1... PASSED
Checking configuration file for node dc1-sn2... PASSED
Checking configuration file for node dc1-sn3... PASSED
Checking for duplication of unique values between nodes... PASSED
```

If the configuration files are incorrect, the issues are shown as WARNING and ERROR. If any configuration errors are found, you must correct them before you continue with the installation.

```
Checking for misnamed node configuration files...
WARNING: ignoring /etc/storagegrid/nodes/dc1-adm1
WARNING: ignoring /etc/storagegrid/nodes/dc1-sn2.conf.keep
WARNING: ignoring /etc/storagegrid/nodes/my-file.txt
Checking configuration file for node dc1-adm1...
ERROR: NODE_TYPE = VM_Foo_Node
        VM_Foo_Node is not a valid node type.  See *.conf.sample
ERROR: ADMIN_ROLE = Foo
        Foo is not a valid admin role.  See *.conf.sample
ERROR: BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-gw1-var-local
        /dev/mapper/sgws-gw1-var-local is not a valid block device
Checking configuration file for node dc1-gw1...
ERROR: GRID_NETWORK_TARGET = bond0.1001
        bond0.1001 is not a valid interface.  See `ip link show`
ERROR: GRID_NETWORK_IP = 10.1.3
        10.1.3 is not a valid IPv4 address
ERROR: GRID_NETWORK_MASK = 255.248.255.0
        255.248.255.0 is not a valid IPv4 subnet mask
Checking configuration file for node dc1-sn1...
ERROR: GRID_NETWORK_GATEWAY = 10.2.0.1
        10.2.0.1 is not on the local subnet
ERROR: ADMIN_NETWORK_ESL = 192.168.100.0/21,172.16.0foo
        Could not parse subnet list
Checking configuration file for node dc1-sn2... PASSED
Checking configuration file for node dc1-sn3... PASSED
Checking for duplication of unique values between nodes...
ERROR: GRID_NETWORK_IP = 10.1.0.4
        dc1-sn2 and dc1-sn3 have the same GRID_NETWORK_IP
ERROR: BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/sgws-sn2-var-local
        dc1-sn2 and dc1-sn3 have the same BLOCK_DEVICE_VAR_LOCAL
ERROR: BLOCK_DEVICE_RANGEDB_00 = /dev/mapper/sgws-sn2-rangedb-0
        dc1-sn2 and dc1-sn3 have the same BLOCK_DEVICE_RANGEDB_00
```

Start the StorageGRID host service

Learn how to start the StorageGRID host service.

To start the StorageGRID nodes and ensure that they restart after a host reboot, you must enable and start the StorageGRID host service.

To start the StorageGRID host service, complete the following steps.

Steps

1. Run the following commands on each host:

```
sudo systemctl enable storagegrid
sudo systemctl start storagegrid
```



The start process might take some time on the initial run.

2. Run the following command to ensure the deployment is proceeding:

```
sudo storagegrid node status node-name
```

3. For any node that returns a status of **Not-Running** or **Stopped**, run the following command:

```
sudo storagegrid node start node-name
```

For example, given the following output you would start the `dc1-adm1` node:

```
[user@host1]# sudo storagegrid node status
Name Config-State Run-State
dc1-adm1 Configured Not-Running
dc1-sn1 Configured Running
```

4. If you have previously enabled and started the StorageGRID host service (or if you aren't sure whether the service has been enabled and started), also run the following command:

```
sudo systemctl reload-or-restart storagegrid
```

Configure the Grid Manager in StorageGRID

Learn how to configure the Grid Manager in StorageGRID on the primary admin node.

Complete the installation by configuring the StorageGRID system from the Grid Manager user interface on the primary Admin Node.

High-level steps

Configuring the grid and completing the installation involves the following tasks:

Steps

1. [Navigate to Grid Manager](#)
2. [Specify the StorageGRID license information](#)
3. [Add sites to StorageGRID](#)
4. [Specify grid network subnets](#)
5. [Approve pending grid nodes](#)
6. [Specify NTP server information](#)
7. [Specify domain name system server information](#)
8. [Specify the StorageGRID system passwords](#)
9. [Review your configuration and complete installation](#)

Navigate to Grid Manager

Use Grid Manager to define all of the information required to configure your StorageGRID system.

Before you begin, the primary Admin Node must be deployed and have completed the initial startup sequence.

To use Grid Manager to define information, complete the following steps.

Steps

1. Access Grid Manager at the following address:

```
https://primary_admin_node_grid_ip
```

Alternatively, you can access Grid Manager on port 8443.

```
https://primary_admin_node_ip:8443
```

2. Click **Install a StorageGRID System**.
The page used to configure a StorageGRID grid is displayed.

1
License2
Sites3
Grid Network4
Grid Nodes5
NTP6
DNS7
Passwords8
Summary

License

Enter a grid name and upload the license file provided by NetApp for your StorageGRID system.

Grid Name

License File

Add StorageGRID license details

Learn how to upload the StorageGRID license file.

You must specify the name for your StorageGRID system and upload the license file provided by NetApp.

To specify the StorageGRID license information, complete the following steps:

Steps

1. On the License page, in the Grid Name field, enter a name for your StorageGRID system. After installation, the name is displayed as the top level in the grid topology tree.
2. Click Browse, locate the NetApp License File (`NLF-unique-id.txt`), and click Open. The license file is validated, and the serial number and licensed storage capacity are displayed.



The StorageGRID installation archive includes a free license that does not provide any support entitlement for the product. You can update to a license that offers support after installation.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords


8 Summary

Sites

In a single-site deployment, infrastructure and operations are centralized in one site.

In a multi-site deployment, infrastructure can be distributed asymmetrically across sites, and proportional to the needs of each site. Typically, sites are located in geographically different locations. Having multiple sites also allows the use of distributed replication and erasure coding for increased availability and resiliency.

Site Name 1



3. Click Next.

Add sites to StorageGRID


Learn how to add sites to StorageGRID to increase reliability and storage capacity.

When you are installing StorageGRID, you must create at least one site. You can create additional sites to increase the reliability and storage capacity of your StorageGRID system.

To add sites, complete the following steps:

Steps

1. On the Sites page, enter the site name.
2. To add additional sites, click the plus sign next to the last site entry and enter the name in the new Site Name text box.
Add as many additional sites as required for your grid topology. You can add up to 16 sites.



3. Click Next.

Specify grid network subnets for StorageGRID

Learn how to configure the grid network subnets for StorageGRID.

You must specify the subnets that are used on the grid network.

The subnet entries include the subnets for the grid network for each site in your StorageGRID system, in addition to any subnets that must be reachable through the grid network (for example, the subnets hosting your NTP servers).

If you have multiple grid subnets, the grid network gateway is required. All grid subnets specified must be reachable through this gateway.


To specify grid network subnets, complete the following steps:

Steps

1. In the Subnet 1 text box, specify the CIDR network address for at least one grid network.

2. Click the plus sign next to the last entry to add an additional network entry.

If you have already deployed at least one node, click Discover Grid Networks Subnets to automatically populate the grid network subnet list with the subnets reported by grid nodes that have registered with Grid Manager.



3. Click Next.

Approve grid nodes for StorageGRID

Learn how to review and approve any pending grid nodes that join the StorageGRID system.

You must approve each grid node before it joins the StorageGRID system.



Before you begin, all virtual and StorageGRID appliance grid nodes must be deployed.

To approve pending grid nodes, complete the following steps:

Steps

1. Review the Pending Nodes list and confirm that it shows all of the grid nodes you deployed.



If a grid node is missing, confirm that it was deployed successfully.

2. Click the radio button next to a pending node that you want to approve.

NetApp® StorageGRID®

Help ▾

Install

1 License 8
2 Sites
3 Grid Network
4 Grid Nodes
5 NTP
6 DNS
7 Passwords
Summary

Grid Nodes

Approve and configure grid nodes, so that they are added correctly to your StorageGRID system.

Pending Nodes

Grid nodes are listed as pending until they are assigned to a site, configured, and approved.

		<input type="button" value="Approve"/>	<input type="button" value="Remove"/>	Search	<input type="button" value=""/>	
		Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address
<input checked="" type="radio"/>	f6:8a:36:44:c4:80	dc1-adm1	Admin Node	CentOS Container	10.193.204.43/24	
<input type="radio"/>	46:5a:b6:7a:6d:97	dc1-sn1	Storage Node	CentOS Container	10.193.204.44/24	
<input type="radio"/>	ba:e5:f7:6e:ec:0b	dc1-sn3	Storage Node	CentOS Container	10.193.204.46/24	
<input type="radio"/>	c6:89:e5:bf:8a:47	dc1-gw1	API Gateway Node	CentOS Container	10.193.204.47/24	
<input type="radio"/>	fe:91:ad:e1:46:c0	dc1-gw2	API Gateway Node	CentOS Container	10.193.204.98/24	

3. Click Approve.
4. In General Settings, modify the settings for the following properties, as necessary.

Admin Node Configuration

General Settings

Site	<input type="text" value="New York"/>
Name	<input type="text" value="dc1-adm1"/>
NTP Role	<input type="text" value="Automatic"/>

Grid Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="10.193.204.43/24"/>
Gateway	<input type="text" value="10.193.204.1"/>

Admin Network

Configuration DISABLED

This network interface is not present. Add the network interface before configuring network settings.

IPv4 Address (CIDR)	<input type="text"/>
Gateway	<input type="text"/>
Subnets (CIDR)	<input type="text"/>

Client Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="10.193.205.43/24"/>
Gateway	<input type="text" value="10.193.205.1"/>

— **Site:** The system name of the site for this grid node.

— **Name:** The host name that will be assigned to the node, and the name that will be displayed in Grid Manager. The name defaults to the name you specified during node deployment, but you can change the name as needed.

— **NTP role:** The NTP role of the grid node. The options are Automatic, Primary, and Client. Selecting the Automatic option assigns the Primary role to Admin Nodes, Storage nodes with Administrative Domain Controller (ADC) services, Gateway Nodes, and any grid nodes that have nonstatic IP addresses. All other grid nodes are assigned the client role.



Make sure that at least two nodes at each site can access at least four external NTP sources. If only one node at a site can reach the NTP sources, timing issues will occur if that node goes down. In addition, designating two nodes per site as primary NTP sources ensures accurate timing if a site is isolated from the rest of the grid.

— **ADC service (storage nodes only):** Select Automatic to let the system determine whether the node requires the ADC service. The ADC service keeps track of the location and availability of grid services. At least three storage nodes at each site must include the ADC service. You cannot add the ADC service to a node after it is deployed.

5. In Grid Network, modify the settings for the following properties as necessary:

— **IPv4 address (CIDR):** The CIDR network address for the grid network interface (eth0 inside the container). For example, 192.168.1.234/24.

— **Gateway:** The grid network gateway. For example, 192.168.0.1.



If there are multiple grid subnets, the gateway is required.



If you selected DHCP for the grid network configuration, and you change the value here, the new value is configured as a static address on the node. Make sure that the resulting IP address is not in a DHCP address pool.

6. To configure the admin network for the grid node, add or update the settings in the Admin Network section as necessary.

Enter the destination subnets of the routes out of this interface in the subnets (CIDR) text box. If there are multiple admin subnets, the admin gateway is required.



If you selected DHCP for the admin network configuration, and you change the value here, the new value is configured as a static address on the node. Make sure that the resulting IP address is not in a DHCP address pool.

Appliances: For a StorageGRID appliance, if the admin network was not configured during the initial installation using the StorageGRID Appliance Installer, it cannot be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- a. Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**. Rebooting can take several minutes.
- b. Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- c. Select **Configure Networking > IP Configuration** and configure the enabled networks.
- d. Return to the Home page and click Start Installation.
- e. In Grid Manager: If the node is listed in the Approved Nodes table, reset the node.
- f. Remove the node from the Pending Nodes table.
- g. Wait for the node to reappear in the Pending Nodes list.
- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page. For additional information, see the installation and maintenance instructions for your appliance model.


7. If you want to configure the Client Network for the grid node, add or update the settings in the Client Network section as necessary. If the Client Network is configured, the gateway is required, and it becomes the default gateway for the node after installation.

Appliances: For a StorageGRID appliance, if the Client Network was not configured during the initial installation using the StorageGRID Appliance Installer, it cannot be configured in this Grid Manager dialog box. Instead, you must follow these steps:

- a. Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**. Rebooting can take several minutes.
- b. Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- c. Select **Configure Networking > IP Configuration** and configure the enabled networks.
- d. Return to the Home page and click Start Installation.
- e. In Grid Manager: If the node is listed in the Approved Nodes table, reset the node.
- f. Remove the node from the Pending Nodes table.
- g. Wait for the node to reappear in the Pending Nodes list.
- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page. For additional information, see the installation and maintenance instructions for your appliance.

8. Click Save.

The grid node entry moves to the Approved Nodes list.



The screenshot shows the StorageGRID Grid Manager interface. At the top, there is a navigation bar with a 'Help' dropdown. Below the bar, a progress bar indicates the installation process: Step 1 (License) is completed (green), Step 2 (Sites) is in progress (yellow), Step 3 (Grid Network) is completed (green), Step 4 (Grid Nodes) is in progress (blue), Step 5 (NTP) is in progress (yellow), Step 6 (DNS) is in progress (yellow), and Step 7 (Passwords) is in progress (yellow). Step 8 (Summary) is shown as a link below the progress bar.

Grid Nodes

Approve and configure grid nodes, so that they are added correctly to your StorageGRID system.

Pending Nodes

Grid nodes are listed as pending until they are assigned to a site, configured, and approved.

		Approve	Remove	Search	
	Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address
<input checked="" type="radio"/>	16:8a:36:44:c4:80	dc1-adm1	Admin Node	CentOS Container	10.193.204.43/24
<input type="radio"/>	46:5a:b6:7a:6d:97	dc1-sn1	Storage Node	CentOS Container	10.193.204.44/24
<input type="radio"/>	ba:e5:f7:6e:ec:0b	dc1-sn3	Storage Node	CentOS Container	10.193.204.46/24
<input type="radio"/>	c6:89:e5:bf:8a:47	dc1-gw1	API Gateway Node	CentOS Container	10.193.204.47/24
<input type="radio"/>	fe:91:ad:e1:46:c0	dc1-gw2	API Gateway Node	CentOS Container	10.193.204.98/24

9. Repeat steps 1-8 for each pending grid node you want to approve.

You must approve all nodes that you want in the grid. However, you can return to this page at any time before you click **Install** on the Summary page. To modify the properties of an approved grid node, click its radio button and then click **Edit**.

10. When you have finished approving grid nodes, click **Next**.

Specify NTP Server details for StorageGRID

Learn how to specify the NTP configuration information for your StorageGRID system so that operations performed on separate servers can be kept synchronized.

To prevent issues with time drift, you must specify four external NTP server references of Stratum 3 or higher.

 When specifying the external NTP source for a production-level StorageGRID installation, do not use the Windows Time (W32Time) service on a version of Windows earlier than Windows Server 2016. The time service on earlier versions of Windows is not sufficiently accurate and is not supported by Microsoft for use in demanding environments like StorageGRID.

The external NTP servers are used by the nodes to which you previously assigned the primary NTP roles.

 The Client Network is not enabled early enough in the installation process to be the only source of NTP servers. Make sure that at least one NTP server can be reached over the grid network or admin network.

To specify NTP server information, complete the following steps:

Steps

1. In the Server 1 to Server 4 text boxes, specify the IP addresses for at least four NTP servers.
2. If necessary, click the plus sign next the last entry to add more server entries.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Network Time Protocol

Enter the IP addresses for at least four Network Time Protocol (NTP) servers, so that operations performed on separate servers are kept in sync.

Server 1	10.193.204.1
Server 2	10.193.204.1
Server 3	10.193.174.249
Server 4	10.193.174.250

+

Cancel Back Next

3. Click Next.

Specify DNS server details for StorageGRID

Learn how to configure the DNS server for StorageGRID.

You must specify the DNS information for your StorageGRID system so that you can access external servers using host names instead of IP addresses.

Specifying DNS server information allows you to use fully qualified domain name (FQDN) host names rather than IP addresses for email notifications and NetApp AutoSupport® messages. NetApp recommends specifying at least two DNS servers.



You should select DNS servers that each site can access locally in the event of network islanding.

To specify DNS server information, complete the following steps:

Steps

1. In the Server 1 text box, specify the IP address for a DNS server.
2. If necessary, click the plus sign next to the last entry to add more servers.

NetApp® StorageGRID®

Help ▾

Install


1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Domain Name Service

Enter the IP address for at least one Domain Name System (DNS) server, so that server hostnames can be used instead of IP addresses. Specifying at least two DNS servers is recommended. Configuring DNS enables server connectivity, email notifications, and NetApp AutoSupport.

Server 1 10.193.204.101

Server 2 10.193.204.102



3. Click Next.

Specify the system passwords for StorageGRID

Learn how to secure your StorageGRID system by setting the provisioning passphrase and the Grid Management root user password.

To enter the passwords to use to secure your StorageGRID system, follow these steps:

Steps

1. In Provisioning Passphrase, enter the provisioning passphrase that will be required to make changes to the grid topology of your StorageGRID system. You should record this password in a secure place.
2. In Confirm Provisioning Passphrase, reenter the provisioning passphrase.
3. In Grid Management Root User Password, enter the password to use to access Grid Manager as the root user.
4. In Confirm Root User Password, reenter the Grid Manager password.



Passwords

Enter secure passwords that meet your organization's security policies. A text file containing the command line passwords must be downloaded during the final installation step.

Provisioning Passphrase	<input type="text" value="*****"/>
Confirm Provisioning Passphrase	<input type="text" value="*****"/>
Grid Management Root User Password	<input type="text" value="*****"/>
Confirm Root User Password	<input type="text" value="*****"/>

Create random command line passwords.

5. If you are installing a grid for proof of concept or demo purposes, deselect the Create Random Command Line Passwords option.

For production deployments, random passwords should always be used for security reasons. Deselect the Create Random Command Line Passwords option only for demo grids if you want to use default passwords to access grid nodes from the command line using the root or admin account.



When you click **Install** on the Summary page, you are prompted to download the Recovery Package file (`sgws-recovery-packageid-revision.zip`). You must download this file to complete the installation. The passwords to access the system are stored in the `Passwords.txt` file, contained in the Recovery Package file.

6. Click **Next**.

Review configuration and complete StorageGRID install

Learn how to validate the grid configuration information and complete the StorageGRID install process.

To make sure that the installation completes successfully, carefully review the configuration information you have entered. Follow these steps.

Steps

1. View the Summary page.

NetApp® StorageGRID®

Help ▾

Install

1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

Summary

Verify that all of the grid configuration information is correct, and then click **Install**. You can view the status of each grid node as it installs. Click the **Modify** links to go back and change the associated information.

General Settings

This is an unsupported license and does not provide any support entitlement for this product.

Grid Name	North America	Modify License
Passwords	StorageGRID demo grid passwords.	Modify Passwords

Networking

NTP	10.193.204.101 10.193.204.102 10.193.174.249 10.54.17.30	Modify NTP
DNS	10.193.204.101 10.193.204.102	Modify DNS
Grid Network	10.193.204.0/24	Modify Grid Network

Topology

Topology	New York	Modify Sites	Modify Grid Nodes
	dc1-adm1 dc1-gw1 dc1-gw2 dc1-sn1 dc1-sn2 dc1-sn3		

Buttons

[Cancel](#) [Back](#) [Install](#)

2. Verify that all of the grid configuration information is correct. Use the **Modify** links on the **Summary** page to go back and correct any errors.
3. Click **Install**.



If a node is configured to use the Client Network, the default gateway for that node switches from the grid network to the Client Network when you click **Install**. If you lose connectivity, make sure that you are accessing the primary Admin Node through an accessible subnet. For more information, see "Network Installation and Provisioning."

4. Click **Download Recovery Package**.

When the installation progresses to the point where the grid topology is defined, you are prompted to download the Recovery Package file (.zip) and confirm that you can access the contents of this file. You must download the Recovery Package file so that you can recover the StorageGRID system in case one or more grid nodes fail.

Verify that you can extract the contents of the .zip file and then save it in two safe, secure, and separate locations.



The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

5. Select the I Have Successfully Downloaded and Verified the Recovery Package File option and then click Next.

Download Recovery Package

Before proceeding, you must download the Recovery Package file. This file is necessary to recover the StorageGRID system if a failure occurs.

When the download completes, open the .zip file and confirm it includes a "gpt-backup" directory and a second .zip file. Then, extract this inner .zip file and confirm you can open the passwords.txt file.

After you have verified the contents, copy the Recovery Package file to two safe, secure, and separate locations. The Recovery Package file must be secured because it contains encryption keys and passwords that can be used to obtain data from the StorageGRID system.

 The Recovery Package is required for recovery procedures and must be stored in a secure location.


Download Recovery Package

I have successfully downloaded and verified the Recovery Package file.

If the installation is still in progress, the Installation Status page opens. This page indicates the progress of the installation for each grid node.

Installation Status

If necessary, you may [Download the Recovery Package file again](#).



The screenshot shows a table with columns for Name, IP, Site, Grid Network IPv4 Address, Progress, and Stage. The table lists five grid nodes: dc1-adm1, dc1-g1, dc1-t1, dc1-s2, and dc1-s3. The progress bar for dc1-adm1 is nearly full (blue), while the others are mostly empty (blue). The stages are: Starting services (Complete), Waiting for Dynamic IP Service peers, Downloading hotfix from primary Admin if needed, and Downloading hotfix from primary Admin if needed.

Name	IP	Site	Grid Network IPv4 Address	Progress	Stage
dc1-adm1		Site1	172.16.4.215/21	<div style="width: 90%;"></div>	Starting services
dc1-g1		Site1	172.16.4.216/21	<div style="width: 100%;"></div>	Complete
dc1-t1		Site1	172.16.4.217/21	<div style="width: 0%;"></div>	Waiting for Dynamic IP Service peers
dc1-s2		Site1	172.16.4.218/21	<div style="width: 0%;"></div>	Downloading hotfix from primary Admin if needed
dc1-s3		Site1	172.16.4.219/21	<div style="width: 0%;"></div>	Downloading hotfix from primary Admin if needed

When the Complete stage is reached for all grid nodes, the sign-in page for Grid Manager opens.

6. Sign in to Grid Manager as the root user with the password that you specified during the installation.

Upgrade bare-metal nodes in StorageGRID

Learn about the upgrade process for bare-metal nodes in StorageGRID.

The upgrade process for bare-metal nodes is different than that for appliances or VMware nodes. Before performing an upgrade of a bare-metal node, you must first upgrade the RPM files on all hosts before running the upgrade through the GUI.

```
[root@host1 rpms]# rpm -Uvh StorageGRID-Webscale-Images-*.rpm
[root@host1 rpms]# rpm -Uvh StorageGRID-Webscale-Service-*.rpm
```

Now you can proceed to the software upgrade through the GUI.

TR-4907: Configure StorageGRID with veritas Enterprise Vault


Introduction to configuring StorageGRID for site failover

Learn how veritas Enterprise Vault uses StorageGRID as a primary storage target for disaster recovery.

This configuration guide provides the steps to configure NetApp® StorageGRID® as a primary storage target with veritas Enterprise Vault. It also describes how to configure StorageGRID for site failover in a disaster recovery (DR) scenario.

Reference architecture

StorageGRID provides an on-premises, S3-compatible cloud backup target for veritas Enterprise Vault. The following figure illustrates the veritas Enterprise Vault and StorageGRID architecture.



Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- NetApp StorageGRID Documentation Center
<https://docs.netapp.com/us-en/storagegrid-118/>
- NetApp StorageGRID Enablement
<https://docs.netapp.com/us-en/storagegrid-enable/>
- NetApp Product Documentation
<https://www.netapp.com/support-and-training/documentation/>

Configure StorageGRID and veritas Enterprise Vault

Learn how to implement basic configurations for StorageGRID 11.5 or higher and veritas Enterprise Vault 14.1 or higher.

This configuration guide is based on StorageGRID 11.5 and Enterprise Vault 14.1. For write once, read many (WORM) mode storage using S3 Object Lock, StorageGRID 11.6 and Enterprise Vault 14.2.2 was used. For more detailed information about these guidelines, see the [StorageGRID Documentation](#) page or contact a StorageGRID expert.

Prerequisites to configure StorageGRID and veritas Enterprise Vault

- Before you configure StorageGRID with veritas Enterprise Vault, verify the following prerequisites:



For WORM storage (Object Lock), StorageGRID 11.6 or higher is required.

- veritas Enterprise Vault 14.1 or higher is installed.



For WORM storage (Object Lock), Enterprise Vault version 14.2.2 or higher is required.


- Vault store groups and a vault store has been created.
For more information, see the veritas Enterprise Vault Administration Guide.
- A StorageGRID tenant, access key, secret key and bucket have been created.
- A StorageGRID load balancer endpoint has been created (either HTTP or HTTPS).
- If using a self-signed certificate, add the StorageGRID self-signed CA certificate to the Enterprise Vault Servers. For more information, see this [veritas Knowledge Base article](#).
- Update and apply the latest Enterprise Vault configuration file to enable supported storage solutions such as NetApp StorageGRID. For more information, see this [veritas Knowledge Base article](#).

Configure StorageGRID with veritas Enterprise Vault


To configure StorageGRID with veritas Enterprise Vault, complete the following steps:

Steps


1. Launch the Enterprise Vault Administration console.




2. Create a new vault store partition in the appropriate vault store. Expand the Vault Store Groups folder and then the appropriate vault store. Right-click Partition and select **New > Partition**.



3. Follow the New Partition creation wizard. From the Storage Type drop-down menu, select NetApp StorageGRID (S3). Click Next.




4. Leave the Store Data in WORM Mode Using S3 Object Lock option unchecked. Click Next.




5. On the connection settings page, provide the following information:
 - Access key ID
 - Secret access key
 - Service host name: Ensure to include the load balancer endpoint (LBE) port configured in StorageGRID (such as https://<hostname>:<LBE_port>)
 - Bucket name: Name of the pre created target bucket. veritas Enterprise Vault does not create the

bucket.


- Bucket region: us-east-1 is the default value.




6. To verify the connection to the StorageGRID bucket, click Test. Verify that the connection test was successful. Click OK and then Next.




7. StorageGRID does not support the S3 replication parameter. To protect your objects, StorageGRID uses Information Lifecycle Management (ILM) rules to specify data protection schemes - multiple copies or erasure coding. Select the When Archived Files Exist on the Storage Option and click Next.



8. Verify the information on the summary page and click Finish.



9. After the new vault store partition has been successfully created, you can archive, restore, and search data in Enterprise Vault with StorageGRID as the primary storage.



Configure StorageGRID S3 Object Lock for WORM storage

Learn how to configure StorageGRID for WORM storage using S3 Object Lock.

Prerequisites to configure StorageGRID for WORM storage


For WORM storage, StorageGRID uses S3 Object Lock to retain objects for compliance. This requires StorageGRID 11.6 or higher, where S3 Object Lock default bucket retention was introduced. Enterprise Vault also requires version 14.2.2 or higher.

Configure StorageGRID S3 Object Lock default bucket retention

To configure the StorageGRID S3 Object Lock default bucket retention, complete the following steps:

Steps

1. In StorageGRID Tenant Manager, create a bucket and click Continue



2. Select the Enable S3 Object Lock option and click Create Bucket.

Create bucket

1 Enter details 2 Manage object settings Optional

Manage object settings Optional

Object versioning

Enable object versioning if you want to store every version of each object in this bucket. You can then retrieve previous versions of an object as needed.

Info Object versioning has been enabled automatically because this bucket has S3 Object Lock enabled.

Enable object versioning

S3 Object Lock

S3 Object Lock allows you to specify retention and legal hold settings for the objects ingested into a bucket. If you want to use S3 Object Lock, you must enable this setting when you create the bucket. You cannot add or disable S3 Object Lock after a bucket is created.

If S3 Object Lock is enabled, object versioning is enabled for the bucket automatically and cannot be suspended.

Enable S3 Object Lock

[Previous](#) [Create bucket](#)

3. After the bucket is created, select the bucket to view the bucket options. Expand the S3 Object Lock drop-down option.

Overview

Name:	object-lock-example
Region:	us-east-1
S3 Object Lock:	Enabled
Date created:	2022-06-24 14:44:54 PDT

[View bucket contents in Experimental S3 Console](#)

Bucket options **Bucket access** **Platform services**

Consistency level

Last access time updates

Object versioning

S3 Object Lock

S3 Object Lock allows you to specify retention and legal hold settings for the objects ingested into a bucket. If you want to use S3 Object Lock, you must enable this setting when you create the bucket. You cannot enable or disable S3 Object Lock after a bucket is created.

After S3 Object Lock is enabled for a bucket, you can't disable it. You also can't suspend object versioning for the bucket.

S3 Object Lock
Enabled

Default retention Disable Enable

[Save changes](#)

4. Under Default Retention, select Enable and set a default retention period of 1 day. Click Save Changes.

S3 Object Lock

S3 Object Lock allows you to specify retention and legal hold settings for the objects ingested into a bucket. If you want to use S3 Object Lock, you must enable this setting when you create the bucket. You cannot enable or disable S3 Object Lock after a bucket is created.

After S3 Object Lock is enabled for a bucket, you can't disable it. You also can't suspend object versioning for the bucket.

S3 Object Lock
Enabled

Default retention Disable Enable

Default retention mode
Compliance
No users can overwrite or delete protected object versions during this retention period.

Default retention period

[Save changes](#)


The bucket is now ready to be used by Enterprise Vault to store WORM data.

Configure Enterprise Vault

To configure Enterprise Vault, complete the following steps:

Steps

1. Repeat steps 1-3 in the [Basic configuration](#) section, but this time select the Store data in WORM Mode Using S3 Object Lock option. Click Next.



2. When entering your S3 Bucket connection settings, make sure you are entering the name of an S3 bucket that has the S3 Object Lock Default retention enabled.
3. Test the connection to verify the settings.

Configure StorageGRID site failover for disaster recovery

Learn how to configure StorageGRID site failover in a disaster recovery scenario.

It is a common for a StorageGRID architecture deployment to be multisite. Sites can either be active-active or active-passive for DR. In a DR scenario, make sure that veritas Enterprise Vault can maintain connection to its primary storage (StorageGRID) and continue to ingest and retrieve data during a site failure. This section provides high-level configuration guidance for a two-site, active-passive deployment. For detailed information about these guidelines, see the [StorageGRID Documentation](#) page or contact a StorageGRID expert.

Prerequisites to configure StorageGRID with veritas Enterprise Vault

Before you configure StorageGRID site failover, verify the following prerequisites:

- There is a two-site StorageGRID deployment; for example, SITE1 and SITE2.
- An admin node running the load balancer service or a gateway node, at each site, for load balancing has

been created.

- A StorageGRID load balancer endpoint has been created.

Configure StorageGRID site failover

To configure StorageGRID site failover, complete the following steps:

Steps

1. To ensure connectivity to StorageGRID during site failures, configure a high-availability (HA) group. From StorageGRID Grid Manager Interface (GMI), click Configuration, High Availability Groups, and + Create.
[vertias/veritas-create-high-availability-group]
2. Enter the required information. Click Select Interfaces and include both SITE1 and SITE2's network interfaces where SITE1 (the primary site) is the preferred master. Assign a virtual IP address within the same subnet. Click Save.

Edit High Availability Group 'site1-HA'

High Availability Group

Name	site1-HA
Description	site1-HA

Interfaces

Select interfaces to include in the HA group. All interfaces must be in the same network subnet.

Select Interfaces			
Node Name	Interface	IPv4 Subnet	Preferred Master
SITE1-ADM1	eth2	205.0.0/24	<input checked="" type="radio"/>
SITE2-ADM1	eth2	205.0.0/24	<input type="radio"/>

Displaying 2 interfaces.

Virtual IP Addresses

Virtual IP Subnet: 10.193.205.0/24. All virtual IP addresses must be within this subnet. There must be at least 1 and no more than 10 virtual IP addresses.

Virtual IP Address 1	205.43	+
----------------------	--------	---


Buttons

Cancel Save

3. This virtual IP (VIP) address should be associated to the S3 host name used during veritas Enterprise Vault's partition configuration. The VIP address resolves traffic to SITE1—and during SITE1 failure, the VIP address transparently reroutes traffic to SITE2.
4. Make sure the data is replicated to both SITE1 and SITE2. That way if SITE1 fails, the object data is still available from SITE2. This is done by first configuring the storage pools.

From StorageGRID GMI, click ILM, Storage Pools, and then + Create. Follow the wizard to create two storage pools: one for SITE1 and another for SITE2.

Storage pools are logical groupings of nodes used to define object placement




Storage Pool Details - site1

Number of Nodes: 4
Storage Grade: All Storage Nodes

Node Name	Site Name	Used (%)
SITE1-S3	SITE1	0.448%
SITE1-S4	SITE1	0.401%
SITE1-S2	SITE1	0.383%
SITE1-S1	SITE1	0.312%

Close



Storage Pool Details - site2

Number of Nodes: 4
Storage Grade: All Storage Nodes

Node Name	Site Name	Used (%)
SITE2-S2	SITE2	0.382%
SITE2-S1	SITE2	0.417%
SITE2-S3	SITE2	0.434%
SITE2-S4	SITE2	0.323%

Close

5. From StorageGRID GMI, click ILM, Rules, and then + Create. Follow the wizard to create an ILM rule specifying one copy to be stored per site with an ingest behavior of Balanced.



1 copy per site

Description: 1 copy per site

Ingest Behavior: Balanced Ingest Time

Retention Time: 1 day 1 week

Filtering Criteria: Matches all objects

Retention Diagram:

Stage	Days	Days
Initial	1	1
Final	1	1

6. Add the ILM rule into an ILM policy and activate the policy.

This configuration results in the following outcome:

- A virtual S3 endpoint IP where SITE1 is the primary and SITE2 is the secondary endpoint. If SITE1 fails, the VIP fails over to SITE2.
- When archived data is sent from veritas Enterprise Vault, StorageGRID ensures one copy is stored in SITE1 and another DR copy is stored in SITE2. If SITE1 fails, Enterprise Vault continues to ingest and retrieve from SITE2.



Both of these configurations are transparent to veritas Enterprise Vault. The S3 endpoint, bucket name, access keys, and so on are the same. There is no need to reconfigure the S3 connection settings on the veritas Enterprise Vault partition.

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.