



StorageGRID solutions and resources

StorageGRID solutions and resources

NetApp

December 12, 2025

Table of Contents

StorageGRID solutions and resources	1
Steps to access StorageGRID evaluation software	2
Register for an account	2
Download StorageGRID	2
Validated third-party solutions	3
Validated third-party solutions: Overview	3
StorageGRID 12.0 validated third-party solutions	3
Third-party solutions validated on StorageGRID	3
Third-party solutions validated on StorageGRID with object lock	5
Third-party solutions supported on StorageGRID	5
Key managers supported on StorageGRID	5
StorageGRID 11.9 validated third-party solutions	6
Third-party solutions validated on StorageGRID	6
Third-party solutions validated on StorageGRID with object lock	7
Third-party solutions supported on StorageGRID	8
Key managers supported on StorageGRID	8
StorageGRID 11.8 validated third-party solutions	9
Third-party solutions validated on StorageGRID	9
Third-party solutions validated on StorageGRID with object lock	11
Third-party solutions supported on StorageGRID	11
Key managers supported on StorageGRID	12
StorageGRID 11.7 validated third-party solutions	12
Third-party solutions validated on StorageGRID	12
Third-party solutions validated on StorageGRID with object lock	14
Third-party solutions supported on StorageGRID	14
Key managers supported on StorageGRID	15
StorageGRID 11.6 validated third-party solutions	15
Third-party solutions validated on StorageGRID	15
Third-party solutions validated on StorageGRID with object lock	17
Third-party solutions supported on StorageGRID	17
StorageGRID 11.5 validated third-party solutions	17
Third-party solutions validated on StorageGRID	18
Third-party solutions validated on StorageGRID with object lock	19
Third-party solutions supported on StorageGRID	19
StorageGRID 11.4 validated third-party solutions	19
Third-party solutions validated on StorageGRID	20
Third-party solutions supported on StorageGRID	21
StorageGRID 11.3 validated third-party solutions	21
Third-party solutions validated on StorageGRID	21
Third-party solutions supported on StorageGRID	22
StorageGRID 11.2 validated third-party solutions	23
Third-party solutions validated on StorageGRID	23
Third-party solutions supported on StorageGRID	24

Product feature guides	25
Achieving zero RPO with StorageGRID - A Comprehensive Guide to Multi-Site Replication	25
StorageGRID Overview	25
Requirements for Zero RPO with StorageGRID	30
Synchronous Deployments across multiple sites	30
A Single Grid Multi-site deployment	30
A multi-site multi-grid deployment	34
Conclusion	36
Create Cloud Storage Pool for AWS or Google Cloud	36
Create Cloud Storage Pool for Azure Blob Storage	37
Use a Cloud Storage Pool for backup	38
Configure StorageGRID search integration service	39
Introduction	39
Create tenant and enable platform services	39
Search integration services with Amazon OpenSearch	40
Platform services endpoint configuration	44
Search integration services with on premises Elasticsearch	46
Platform services endpoint configuration	49
Bucket search integration service configuration	51
Where to find additional information	55
Node Clone	55
Node clone considerations	55
Node clone Performance estimates	55
Grid site relocation and site-wide network change procedure	58
Considerations before site relocation	58
Migrating object-based storage from ONTAP S3 to StorageGRID	62
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID	62
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID	62
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID	74
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID	86
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID	95
Tool and application guides	101
Use Cloudera Hadoop S3A connector with StorageGRID	101
Why use S3A for Hadoop workflows?	101
Configure S3A connector to use StorageGRID	101
Test S3A connection to StorageGRID	104
Use S3cmd to test and demonstrate S3 access on StorageGRID	107
Install and configure S3cmd	107
Initial configuration steps	107
Basic command examples	108

Vertica Eon mode database using NetApp StorageGRID as communal storage	108
Introduction	108
NetApp StorageGRID recommendations	110
Installing Eon Mode on-premises with communal storage on StorageGRID	111
Where to find additional information	121
Version history	121
StorageGRID log analytics using ELK stack	121
Requirements	122
Sample files	122
Assumption	122
Instruction	122
Additional resources	126
Use Prometheus and Grafana to extend your metrics retention	127
Introduction	127
Federate Prometheus	127
Install and configure Grafana	136
Use F5 DNS to globally load balance StorageGRID	143
Introduction	143
F5 BIG-IP multi-site StorageGRID configuration	143
Conclusion	158
Datadog SNMP configuration	158
Configure Datadog	159
Use rclone to migrate, PUT, and DELETE objects on StorageGRID	162
Install and configure rclone	162
Basic command examples	170
StorageGRID best practices for deployment with Veeam Backup and Replication	173
Overview	173
Veeam configuration	174
StorageGRID configuration	175
Implementation key points	178
Monitoring StorageGRID	183
Where to find additional information	186
Configure Dremio data source with StorageGRID	186
Configure Dremio data source	186
Instruction	186
NetApp StorageGRID with GitLab	189
Object Storage connection example	189
Procedures and API examples	191
Test and demonstrate S3 encryption options on StorageGRID	191
Server Side Encryption (SSE)	191
Server Side Encryption with Customer provided keys (SSE-C)	192
Bucket Server Side Encryption (SSE-S3)	193
Test and demonstrate S3 object lock on StorageGRID	194
Legal hold	194
Compliance mode	195

Default retention	196
Test deleting an object with a defined retention	197
Policies and permissions in StorageGRID	199
The structure of a policy	199
Using the AWS policy generator	201
Group Policies (IAM)	208
Bucket Policies	213
Bucket lifecycle in StorageGRID	215
What is a lifecycle configuration	215
Structure of a lifecycle policy	216
Apply lifecycle configuration to bucket	218
Example lifecycle policies for standard (non-versioned) buckets	218
Example lifecycle policies for versioned buckets	218
Conclusion	222
Technical reports	223
Introduction to StorageGRID technical reports	223
NetApp StorageGRID and big data analytics	223
NetApp StorageGRID use cases	223
Why StorageGRID for data lakes?	224
Benchmarking Data Warehouses and Lakehouses with S3 Object Storage: A Comparative Study	225
Hadoop S3A tuning	228
What is Hadoop?	228
Hadoop HDFS and S3A connector	228
Hadoop S3A connector tuning	229
TR-4871: Configure StorageGRID for backup and recovery with Commvault	234
Backup and recover data using StorageGRID and Commvault	234
Tested solution overview	236
StorageGRID sizing guidance	238
Run a data protection job	240
Review baseline performance tests	248
Bucket consistency level recommendation	249
TR-4626: Load balancers	250
Use third-party load balancers with StorageGRID	250
Use StorageGRID load balancers	251
Learn how to implement SSL certificates for HTTPS in StorageGRID	252
Configure trusted third-party load balancer in StorageGRID	253
Learn about local traffic manager load balancers	253
Learn about few use cases for StorageGRID configurations	256
Validate SSL connection in StorageGRID	259
Understand global load balancing requirements for StorageGRID	259
TR-4645: Security features	260
Secure StorageGRID data and metadata in an object store	260
Data access security features	262
Object and metadata security	270
Administration security features	272

Platform security features	276
Cloud integration	278
TR-4921: Ransomware defense	278
Protect StorageGRID S3 objects from ransomware	278
Ransomware defense using object lock	279
Ransomware defense using replicated bucket with versioning	282
Ransomware defense using versioning with protective IAM policy	285
Ransomware investigation and remediation	288
TR-4765: Monitor StorageGRID	289
Introduction to StorageGRID monitoring	289
Use the GMI dashboard to monitor StorageGRID	290
Use alerts to monitor StorageGRID	291
Advanced monitoring in StorageGRID	292
Access metrics using cURL in StorageGRID	295
View metrics using the Grafana dashboard in StorageGRID	296
Use traffic classification policies in StorageGRID	297
Use audit logs to monitor StorageGRID	300
Use the StorageGRID app for Splunk	300
TR-4882: Install a StorageGRID bare metal grid	300
Introduction to installing StorageGRID	300
Prerequisites to install StorageGRID	301
Install Docker for StorageGRID	311
Prepare node configuration files for StorageGRID	311
Install StorageGRID dependencies and packages	315
Validate the StorageGRID configuration files	315
Start the StorageGRID host service	317
Configure the Grid Manager in StorageGRID	317
Add StorageGRID license details	319
Add sites to StorageGRID	320
Specify grid network subnets for StorageGRID	321
Approve grid nodes for StorageGRID	322
Specify NTP Server details for StorageGRID	327
Specify DNS server details for StorageGRID	328
Specify the system passwords for StorageGRID	329
Review configuration and complete StorageGRID install	330
Upgrade bare-metal nodes in StorageGRID	332
TR-4907: Configure StorageGRID with veritas Enterprise Vault	333
Introduction to configuring StorageGRID for site failover	333
Configure StorageGRID and veritas Enterprise Vault	334
Configure StorageGRID S3 Object Lock for WORM storage	339
Configure StorageGRID site failover for disaster recovery	343
Steps to access StorageGRID evaluation software	347
Register for an account	347
Download StorageGRID	347
NetApp StorageGRID Blogs	348

NetApp StorageGRID documentation	350
Legal notices	351
Copyright	351
Trademarks	351
Patents	351
Privacy policy	351
Open source	351

StorageGRID solutions and resources

Steps to access StorageGRID evaluation software

This instruction is for NetApp sales, partners, and prospects engaged with NetApp.

Register for an account

1. Register for an account on the [NetApp Support site](#) using your business email.
 - a. Make sure you are not signed in with the newly created account.
 - b. If you already have an account, make sure you are not signed in and proceed with the next step.
2. Create a non-technical support case to elevate access levels to "prospect." To do this, click on the "[Report an Issue](#)" link in the footer of the website.
3. Select "Registration Issue" as the feedback category.
4. In the comments section, write: "My account email address is *Your-Email-Address*. I would like to get prospect access to download the StorageGRID evaluation software."
 - a. Mention the name of the NetApp internal person who suggested the request for prospect access.

Download StorageGRID

1. After your support case has been reviewed and approved, NetApp support will notify you via email that your account has been granted prospect access.
2. Download the [StorageGRID evaluation software](#).



The Eval license file is located within the zip file. It is StorageGRID-Webscale-<version>\vsphere\NLF000000.txt once unzipped.



Downloading the software is a process that involves trade compliance measures to adhere to legal requirements. To ensure compliance, users are required to create an account and open a support case before gaining access. This process helps us maintain proper control and documentation while providing prospects with the production-ready software they need.

We provide the "production-ready" version of StorageGRID, which is not an open-source or alternative version. It is important to note that **support is not provided** unless the prospect upgrades to a production license.

Please contact StorageGRID.Feedback@netapp.com for any trouble with the above steps.

Validated third-party solutions

Validated third-party solutions: Overview

NetApp, in collaboration with our partners, has validated these solutions for use with StorageGRID. Review the information in this section to learn what solutions have been validated, and to obtain additional instructions if applicable.

Join forces with NetApp to accelerate portfolio innovation, expand market awareness and increase sales when you create tested, best-of-breed NetApp solutions. [Become an alliance partner today.](#)

StorageGRID 12.0 validated third-party solutions

The following third-party solutions have been validated for use with StorageGRID 12.0. If the solution you are looking for is not listed, please contact your NetApp account representative.

Third-party solutions validated on StorageGRID

These solutions have been tested in collaboration with the respective partners.

- Actifio
- Alluxio
- Apache Kafka
- AWS Mountpoint
- Bridgestor
- Cantemo
- Citrix Content Collaboration
- Collibra (Minimum Collibra Data Quality version 2024.02)
- Commvault 11
- Couchbase Enterprise Analytics 2.0
- Ctera Portal 6
- Dalet
- Datadobi
- Data Dynamics StorageX
- DefendX
- Diskover Data
- Dremio
- Elasticsearch Snapshot (including frozen tier)
- eMAM
- FujiFilm Object Archive
- GitHub Enterprise Server

- IBM Filenet
- IBM Storage Protect
- Interica
- Komprise
- Microsoft SQL Server Big Data Clusters
- Model9
- Modzy
- Moonwalk Universal
- NICE
- Nasuni
- OpenText Documentum 16.4
- OpenText Documentum 21.4
- OpenText InfoArchive 16 EP7
- OpenText Media Management 16.5 with CyanGate Cloud
- Panzura
- PixitMedia ngenea
- PoINT Archival Gateway 2.0
- PoINT Storage Manager 6.4
- Primestream
- Quantum StorNext 5.4.0.1
- Reveille v10 build 220706 or above
- Rubrik CDM
- s3a
- Signiant
- Snowflake
- Spectra Logic On-Prem Glacier
- Splunk Smartstore
- Starburst
- Storage Made Easy
- Trino
- Varnish Enterprise 6.0.4
- Veeam 12
- Veritas Enterprise Vault 15.1
- Veritas NetBackup 10.1.1 and later
- Vertica 10.x
- Vidispine
- Virtualica StorageFabric

- Weka v3.10 or later

Third-party solutions validated on StorageGRID with object lock

These solutions have been tested in collaboration with the respective partners.

- Commvault 11 Feature Release 26
- IBM Filenet
- IBM Storage Protect
- OpenText Documentum 21.4
- Rubrik
- Veeam 12
- Veritas Enterprise Vault 15.1
- Veritas NetBackup 10.1.1 and later

Third-party solutions supported on StorageGRID

These solutions have been tested.

- Archiware
- Axis Communications
- Congruity360
- DataFrameworks
- EcoDigital DIVA platform
- Encoding.com
- FujiFilm Object Archive
- GE Centricity Enterprise Archive
- Gitlab
- Hyland Acuo
- IBM Aspera
- Milestone Systems
- OnSSI
- Reach Engine
- SilverTrak
- SoftNAS
- QStar
- Velasea

Key managers supported on StorageGRID

These solutions have been tested.

- Entrust Cryptographic Security Platform v10.4.5
- Entrust KeyControl 10.2
- Hashicorp Vault 1.20.2
- Thales CipherTrust Manager 2.20

StorageGRID 11.9 validated third-party solutions

The following third-party solutions have been validated for use with StorageGRID 11.9. If the solution you are looking for is not listed, please contact your NetApp account representative.

Third-party solutions validated on StorageGRID

These solutions have been tested in collaboration with the respective partners.

- Actifio
- Alluxio
- Apache Kafka
- AWS Mountpoint
- Bridgestor
- Cantemo
- Citrix Content Collaboration
- Collibra (Minimum Collibra Data Quality version 2024.02)
- Commvault 11
- Couchbase Enterprise Analytics 2.0
- Ctera Portal 6
- Dalet
- Datadobi
- Data Dynamics StorageX
- DefendX
- Diskover Data
- Dremio
- Elasticsearch Snapshot (including frozen tier)
- eMAM
- FujiFilm Object Archive
- GitHub Enterprise Server
- IBM Filenet
- IBM Storage Protect
- Interica
- Komprise

- Microsoft SQL Server Big Data Clusters
- Model9
- Modzy
- Moonwalk Universal
- NICE
- Nasuni
- OpenText Documentum 16.4
- OpenText Documentum 21.4
- OpenText InfoArchive 16 EP7
- OpenText Media Management 16.5 with CyanGate Cloud
- Panzura
- PixitMedia ngenea
- PoINT Archival Gateway 2.0
- PoINT Storage Manager 6.4
- Primestream
- Quantum StorNext 5.4.0.1
- Reveille v10 build 220706 or above
- Rubrik CDM
- s3a
- Signiant
- Snowflake
- Spectra Logic On-Prem Glacier
- Splunk Smartstore
- Starburst
- Storage Made Easy
- Trino
- Varnish Enterprise 6.0.4
- Veeam 12
- Veritas Enterprise Vault 15.1
- Veritas NetBackup 10.1.1 and later
- Vertica 10.x
- Vidispine
- Virtualica StorageFabric
- Weka v3.10 or later

Third-party solutions validated on StorageGRID with object lock

These solutions have been tested in collaboration with the respective partners.

- Commvault 11 Feature Release 26
- IBM Filenet
- IBM Storage Protect
- OpenText Documentum 21.4
- Rubrik
- Veeam 12
- Veritas Enterprise Vault 15.1
- Veritas NetBackup 10.1.1 and later

Third-party solutions supported on StorageGRID

These solutions have been tested.

- Archiware
- Axis Communications
- Congruity360
- DataFrameworks
- EcoDigital DIVA platform
- Encoding.com
- FujiFilm Object Archive
- GE Centricity Enterprise Archive
- Gitlab
- Hyland Acuo
- IBM Aspera
- Milestone Systems
- OnSSI
- Reach Engine
- SilverTrak
- SoftNAS
- QStar
- Velasea

Key managers supported on StorageGRID

These solutions have been tested.

- Entrust Cryptographic Security Platform v10.4.5
- Entrust KeyControl 10.2
- Hashicorp Vault 1.15.0
- Thales CipherTrust Manager 2.0
- Thales CipherTrust Manager 2.1

- Thales CipherTrust Manager 2.2
- Thales CipherTrust Manager 2.3
- Thales CipherTrust Manager 2.4
- Thales CipherTrust Manager 2.8
- Thales CipherTrust Manager 2.9
- Thales CipherTrust Manager 2.10
- Thales CipherTrust Manager 2.11
- Thales CipherTrust Manager 2.12
- Thales CipherTrust Manager 2.13
- Thales CipherTrust Manager 2.14
- Thales CipherTrust Manager 2.15
- Thales CipherTrust Manager 2.16
- Thales CipherTrust Manager 2.20

StorageGRID 11.8 validated third-party solutions

The following third-party solutions have been validated for use with StorageGRID 11.8. If the solution you are looking for is not listed, please contact your NetApp account representative.

Third-party solutions validated on StorageGRID

These solutions have been tested in collaboration with the respective partners.

- Actifio
- Alluxio
- Apache Kafka
- AWS Mountpoint
- Bridgestor
- Cantemo
- Citrix Content Collaboration
- Collibra (Minimum Collibra Data Quality version 2024.02)
- Commvault 11
- Ctera Portal 6
- Dalet
- Datadobi
- Data Dynamics StorageX
- DefendX
- Diskover Data
- Dremio

- Elasticsearch Snapshot (including frozen tier)
- eMAM
- FujiFilm Object Archive
- GitHub Enterprise Server
- IBM Filenet
- IBM Storage Protect
- Interica
- Komprise
- Microsoft SQL Server Big Data Clusters
- Model9
- Modzy
- Moonwalk Universal
- NICE
- Nasuni
- OpenText Documentum 16.4
- OpenText Documentum 21.4
- OpenText InfoArchive 16 EP7
- OpenText Media Management 16.5 with CyanGate Cloud
- Panzura
- PixitMedia ngenea
- PoINT Archival Gateway 2.0
- PoINT Storage Manager 6.4
- Primestream
- Quantum StorNext 5.4.0.1
- Reveille v10 build 220706 or above
- Rubrik CDM
- s3a
- Signiant
- Snowflake
- Spectra Logic On-Prem Glacier
- Splunk Smartstore
- Starburst
- Storage Made Easy
- Trino
- Varnish Enterprise 6.0.4
- Veeam 12
- Veritas Enterprise Vault 15.1

- Veritas NetBackup 10.1.1 and later
- Vertica 10.x
- Vidispine
- Virtualica StorageFabric
- Weka v3.10 or later

Third-party solutions validated on StorageGRID with object lock

These solutions have been tested in collaboration with the respective partners.

- Commvault 11 Feature Release 26
- IBM Filenet
- IBM Storage Protect
- OpenText Documentum 21.4
- Rubrik
- Veeam 12
- Veritas Enterprise Vault 15.1
- Veritas NetBackup 10.1.1 and later

Third-party solutions supported on StorageGRID

These solutions have been tested.

- Archiware
- Axis Communications
- Congruity360
- DataFrameworks
- EcoDigital DIVA platform
- Encoding.com
- FujiFilm Object Archive
- GE Centricity Enterprise Archive
- Gitlab
- Hyland Acuo
- IBM Aspera
- Milestone Systems
- OnSSI
- Reach Engine
- SilverTrak
- SoftNAS
- QStar
- Velasea

Key managers supported on StorageGRID

These solutions have been tested.

- Entrust KeyControl 10.2
- Hashicorp Vault 1.15.0
- Thales CipherTrust Manager 2.0
- Thales CipherTrust Manager 2.1
- Thales CipherTrust Manager 2.2
- Thales CipherTrust Manager 2.3
- Thales CipherTrust Manager 2.4
- Thales CipherTrust Manager 2.8
- Thales CipherTrust Manager 2.9
- Thales CipherTrust Manager 2.10
- Thales CipherTrust Manager 2.11
- Thales CipherTrust Manager 2.12
- Thales CipherTrust Manager 2.13
- Thales CipherTrust Manager 2.14

StorageGRID 11.7 validated third-party solutions

The following third-party solutions have been validated for use with StorageGRID 11.7. If the solution you are looking for is not listed, please contact your NetApp account representative.

Third-party solutions validated on StorageGRID

These solutions have been tested in collaboration with the respective partners.

- Actifio
- Alluxio
- Apache Kafka
- AWS Mountpoint
- Bridgestor
- Cantemo
- Citrix Content Collaboration
- Collibra (Minimum Collibra Data Quality version 2024.02)
- Commvault 11
- Ctera Portal 6
- Dalet
- Datadobi

- Data Dynamics StorageX
- DefendX
- Diskover Data
- Dremio
- Elasticsearch Snapshot (including frozen tier)
- eMAM
- FujiFilm Object Archive
- GitHub Enterprise Server
- IBM Filenet
- IBM Spectrum Protect Plus
- IBM Storage Protect
- Interica
- Komprise
- Microsoft SQL Server Big Data Clusters
- Model9
- Modzy
- Moonwalk Universal
- NICE
- Nasuni
- OpenText Documentum 16.4
- OpenText Documentum 21.4
- OpenText InfoArchive 16 EP7
- OpenText Media Management 16.5 with CyanGate Cloud
- Panzura
- PixitMedia ngenea
- PoINT Archival Gateway 2.0
- PoINT Storage Manager 6.4
- Primestream
- Quantum StorNext 5.4.0.1
- Reveille v10 build 220706 or above
- Rubrik CDM
- s3a
- Signiant
- Snowflake
- Spectra Logic On-Prem Glacier
- Splunk Smartstore
- Storage Made Easy

- Trino
- Varnish Enterprise 6.0.4
- Veeam 12
- Veritas Enterprise Vault 14
- Veritas NetBackup 10.1.1 and later
- Vertica 10.x
- Vidispine
- Virtualica StorageFabric
- Weka v3.10 or later

Third-party solutions validated on StorageGRID with object lock

These solutions have been tested in collaboration with the respective partners.

- Commvault 11 Feature Release 26
- IBM Filenet
- IBM Storage Protect
- OpenText Documentum 21.4
- Rubrik
- Veeam 12
- Veritas Enterprise Vault 14.2.2
- Veritas NetBackup 10.1.1 and later

Third-party solutions supported on StorageGRID

These solutions have been tested.

- Archiware
- Axis Communications
- Congruity360
- DataFrameworks
- EcoDigital DIVA platform
- Encoding.com
- FujiFilm Object Archive
- GE Centricity Enterprise Archive
- Gitlab
- Hyland Acuo
- IBM Aspera
- Milestone Systems
- OnSSI
- Reach Engine

- SilverTrak
- SoftNAS
- QStar
- Velasea

Key managers supported on StorageGRID

These solutions have been tested.

- Thales CipherTrust Manager 2.0
- Thales CipherTrust Manager 2.1
- Thales CipherTrust Manager 2.2
- Thales CipherTrust Manager 2.3
- Thales CipherTrust Manager 2.4
- Thales CipherTrust Manager 2.8
- Thales CipherTrust Manager 2.9

StorageGRID 11.6 validated third-party solutions

The following third-party solutions have been validated for use with StorageGRID 11.6. If the solution you are looking for is not listed, please contact your NetApp account representative.

Third-party solutions validated on StorageGRID

These solutions have been tested in collaboration with the respective partners.

- Actifio
- Alluxio
- Apache Kafka
- Bridgestor
- Cantemo
- Citrix Content Collaboration
- Commvault 11
- Ctera Portal 6
- Dalet
- Datadobi
- Data Dynamics StorageX
- DefendX
- Diskover Data
- Dremio
- eMAM

- FujiFilm Object Archive
- GitHub Enterprise Server
- IBM Filenet
- IBM Spectrum Protect Plus
- Interica
- Komprise
- Microsoft SQL Server Big Data Clusters
- Model9
- Modzy
- Moonwalk Universal
- NICE
- Nasuni
- OpenText Documentum 16.4
- OpenText Documentum 21.4
- OpenText InfoArchive 16 EP7
- OpenText Media Management 16.5 with CyanGate Cloud
- Panzura
- PixitMedia ngenea
- PoINT Archival Gateway 2.0
- PoINT Storage Manager 6.4
- Primestream
- Quantum StorNext 5.4.0.1
- Reveille v10 build 220706 or above
- Rubrik CDM
- s3a
- Signiant
- Snowflake
- Spectra Logic On-Prem Glacier
- Splunk Smartstore
- Storage Made Easy
- Trino
- Varnish Enterprise 6.0.4
- Veeam 12
- Veritas Enterprise Vault 14
- Veritas NetBackup 8.0
- Vertica 10.x
- Vidispine

- Virtualica StorageFabric
- Weka v3.10 or later

Third-party solutions validated on StorageGRID with object lock

These solutions have been tested in collaboration with the respective partners.

- Commvault 11 Feature Release 26
- IBM Filenet
- OpenText Documentum 21.4
- Veeam 12
- Veritas Enterprise Vault 14.2.2
- Veritas NetBackup 10.1.1 and later

Third-party solutions supported on StorageGRID

These solutions have been tested.

- Archiware
- Axis Communications
- Congruity360
- DataFrameworks
- EcoDigital DIVA platform
- Encoding.com
- FujiFilm Object Archive
- GE Centricity Enterprise Archive
- Gitlab
- Hyland Acuo
- IBM Aspera
- Milestone Systems
- OnSSI
- Reach Engine
- SilverTrak
- SoftNAS
- QStar
- Velasea

StorageGRID 11.5 validated third-party solutions

The following third-party solutions have been validated for use with StorageGRID 11.5. If the solution you are looking for is not listed, please contact your NetApp account representative.

Third-party solutions validated on StorageGRID

These solutions have been tested in collaboration with the respective partners.

- Actifio
- Alluxio
- Bridgestor
- Cantemo
- Citrix Content Collaboration
- Commvault 11
- Ctera Portal 6
- Dalet
- Datadobi
- Data Dynamics StorageX
- DefendX
- Interica
- Komprise
- Moonwalk Universal
- NICE
- Nasuni
- OpenText Documentum 16.4
- OpenText Documentum 21.4
- OpenText InfoArchive 16 EP7
- OpenText Media Management 16.5 with CyanGate Cloud
- Panzura
- PoINT Archival Gateway 2.0
- PoINT Storage Manager 6.4
- Primestream
- Quantum StorNext 5.4.0.1
- Rubrik CDM
- s3a
- Signiant
- Splunk Smartstore
- Trino
- Varnish Enterprise 6.0.4
- Veeam 11
- Veritas Enterprise Vault 11
- Veritas Enterprise Vault 12

- Veritas NetBackup 8.0
- Vertica 10.x
- Vidispine
- Virtualica StorageFabric

Third-party solutions validated on StorageGRID with object lock

These solutions have been tested in collaboration with the respective partners.

- OpenText Documentum 21.4
- Veeam 11

Third-party solutions supported on StorageGRID

These solutions have been tested.

- Archiware
- Axis Communications
- Congruity360
- DataFrameworks
- EcoDigital DIVA platform
- Encoding.com
- FujiFilm Object Archive
- GE Centricity Enterprise Archive
- Gitlab
- Hyland Acuo
- IBM Aspera
- Milestone Systems
- OnSSI
- Reach Engine
- SilverTrak
- SoftNAS
- QStar
- Velasea

StorageGRID 11.4 validated third-party solutions

The following third-party solutions have been validated for use with StorageGRID 11.4. If the solution you are looking for is not listed, please contact your NetApp account representative.

Third-party solutions validated on StorageGRID

These solutions have been tested in collaboration with the respective partners.

- Actifio
- Bridgestor
- Cantemo
- Citrix Content Collaboration
- Commvault 11
- Ctera Portal 6
- Dalet
- Datadobi
- Data Dynamics StorageX
- DefendX
- Interica
- Komprise
- NICE
- Nasuni
- OpenText Documentum 16.4
- OpenText InfoArchive 16 EP7
- OpenText Media Management 16.5 with CyanGate Cloud
- Panzura
- PoINT Archival Gateway 2.0
- PoINT Storage Manager 6.4
- Primestream
- Quantum StorNext 5.4.0.1
- Rubrik CDM
- Signiant
- Splunk Smartstore
- Varnish Enterprise 6.0.4
- Veeam 9.5.4
- Veritas Enterprise Vault 11
- Veritas Enterprise Vault 12
- Veritas NetBackup 8.0
- Vertica 10.x
- Vidispine

Third-party solutions supported on StorageGRID

These solutions have been tested.

- Archiware
- Axis Communications
- Congruity360
- DataFrameworks
- EcoDigital DIVA platform
- Encoding.com
- FujiFilm Object Archive
- GE Centricity Enterprise Archive
- Hyland Acuo
- IBM Aspera
- Milestone Systems
- OnSSI
- Reach Engine
- SilverTrak
- SoftNAS
- QStar
- Velasea

StorageGRID 11.3 validated third-party solutions

The following third-party solutions have been validated for use with StorageGRID 11.3. If the solution you are looking for is not listed, please contact your NetApp account representative.

Third-party solutions validated on StorageGRID

These solutions have been tested in collaboration with the respective partners.

- Actifio
- Bridgestor
- Cantemo
- Citrix Content Collaboration
- Commvault 11
- Ctera Portal 6
- Dalet
- Datadobi
- Data Dynamics StorageX

- DefendX
- Interica
- Komprise
- NICE
- Nasuni
- OpenText Documentum 16.4
- OpenText Media Management 16.5 with CyanGate Cloud
- Panzura
- PoINT Archival Gateway 2.0
- PoINT Storage Manager 6.4
- Primestream
- Quantum StorNext 5.4.0.1
- Rubrik CDM 5.0.1 p1-1342
- Signiant
- Splunk Smartstore
- Varnish Enterprise 6.0.4
- Veeam 9.5.4
- Veritas Enterprise Vault 11
- Veritas Enterprise Vault 12
- Veritas NetBackup 8.0
- Vidispine

Third-party solutions supported on StorageGRID

These solutions have been tested.

- Archiware
- Axis Communications
- Congruity360
- DataFrameworks
- EcoDigital DIVA platform
- Encoding.com
- FujiFilm Object Archive
- GE Centricity Enterprise Archive
- Hyland Acuo
- IBM Aspera
- Milestone Systems
- OnSSI
- Reach Engine

- SilverTrak
- SoftNAS
- QStar
- Velasea

StorageGRID 11.2 validated third-party solutions

The following third-party solutions have been validated for use with StorageGRID 11.2. If the solution you are looking for is not listed, please contact your NetApp account representative.

Third-party solutions validated on StorageGRID

These solutions have been tested in collaboration with the respective partners.

- Actifio
- Bridgestor
- Cantemo
- Citrix Content Collaboration
- Commvault 11
- Ctera Portal 6
- Dalet
- Datadobi
- Data Dynamics StorageX
- DefendX
- Interica
- Komprise
- NICE
- Nasuni
- OpenText Documentum 16.4
- OpenText Media Management 16.5 with CyanGate Cloud
- Panzura
- PoINT Archival Gateway 2.0
- PoINT Storage Manager 6.4
- Primestream
- Quantum StorNext 5.4.0.1
- Rubrik CDM 5.0.1 p1-1342
- Signiant
- Splunk Smartstore
- Varnish Enterprise 6.0.4

- Veeam 9.5.4
- Veritas Enterprise Vault 11
- Veritas Enterprise Vault 12
- Veritas NetBackup 8.0
- Vidispine

Third-party solutions supported on StorageGRID

These solutions have been tested.

- Archiware
- Axis Communications
- Congruity360
- DataFrameworks
- EcoDigital DIVA platform
- Encoding.com
- FujiFilm Object Archive
- GE Centricity Enterprise Archive
- Hyland Acuo
- IBM Aspera
- Milestone Systems
- OnSSI
- Reach Engine
- SilverTrak
- SoftNAS
- QStar
- Velasea

Product feature guides

Achieving zero RPO with StorageGRID - A Comprehensive Guide to Multi-Site Replication

This technical report provides a comprehensive guide to implementing StorageGRID replication strategies to achieve a Recovery Point Objective (RPO) of zero in the event of a site failure. The document details various deployment options for StorageGRID, including multi-site synchronous replication and multi-grid asynchronous replication. It explains how StorageGRID Information Lifecycle Management (ILM) policies can be configured to ensure data durability and availability across multiple locations. Additionally, the report covers performance considerations, failure scenarios, and recovery processes to maintain uninterrupted client operations. The goal of this document is to provide information to ensure that data remains accessible and consistent, even in the event of a complete site failure, by leveraging both synchronous and asynchronous replication techniques.

StorageGRID Overview

NetApp StorageGRID is an object-based storage system that supports the industry-standard Amazon Simple Storage Service (Amazon S3) API.

StorageGRID provides a single namespace across multiple locations with variable levels of service driven by information lifecycle management policies (ILM). With these lifecycle policies you can optimize where your data lives throughout its lifecycle.

StorageGRID allows for configurable durability and availability of your data in local and geo-distributed solutions. Whether your data is on premises or in a public cloud, integrated hybrid cloud workflows allow your business to leverage cloud services like Amazon Simple Notification Service (Amazon SNS), Google Cloud, Microsoft Azure Blob, Amazon S3 Glacier, Elasticsearch, and more.

StorageGRID scale

A minimal StorageGRID deployment consists of an Admin node and 3 Storage nodes in a single site. A single grid can grow up to 220 nodes.

StorageGRID can be deployed as a single site or extended to 16 sites.

The Admin node contains the management interface, a central point for metrics and logging, and maintains the configuration of the StorageGRID components. The Admin node also contains an integrated load balancer for S3 API access.

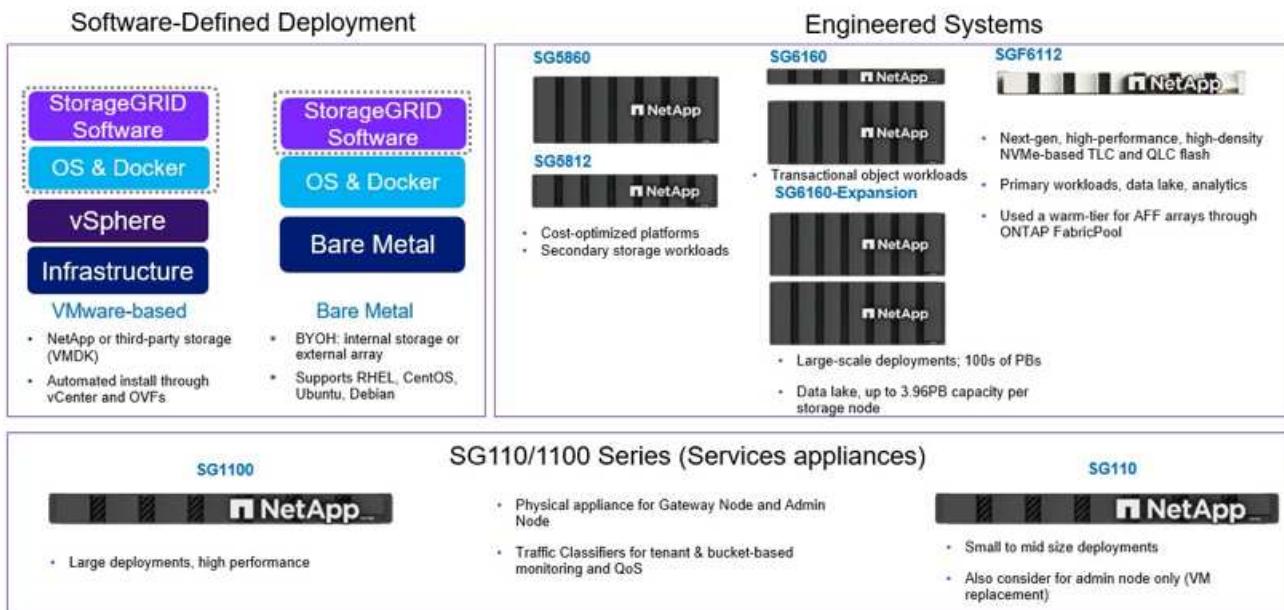
StorageGRID can be deployed as software-only, as VMware virtual machine appliances, or as purpose-built appliances.

A Storage node can be deployed as:

- A metadata only node maximizing object count
- An object storage only node maximizing object space
- A combined metadata and object storage node adding both object count and object space

Each Storage node can scale to multi-petabyte capacity for object storage allowing for a single namespace in the hundreds of petabytes. StorageGRID also provides an integrated load balancer for S3 API operations called a gateway node.

Delivery paths for any workload



StorageGRID consists of a collection of nodes placed into a site topology. A site in StorageGRID can be a unique physical location or reside in a shared physical location as other sites in the grid as a logical construct. A StorageGRID site should not span multiple physical locations. A site represents a shared local area network (LAN) infrastructure and failure domain.

StorageGRID and failure domains

StorageGRID contains multiple layers of failure domains to be considered in deciding how to architect your solution, how to store your data and where your data should be stored to mitigate the risks of failures.

- Grid level - A grid consisting of multiple sites can have site failures or isolation and the accessible site(s) can continue operating as the grid.
- Site level - Failures within a site may impact operations of that site but will not impact the rest of the grid.
- Node level - A node failure will not impact the operation of the site.
- Disk level - a disk failure will not impact operation of the node.

Object data and metadata

With object storage, the unit of storage is an object, rather than a file or a block. Unlike the tree-like hierarchy of a file system or block storage, object storage organizes data in a flat, unstructured layout. Object storage decouples the physical location of the data from the method used to store and retrieve that data.

Each object in an object-based storage system has two parts: object data and object metadata.

- Object data represents the actual underlying data, for example, a photograph, a movie, or a medical record.

- Object metadata is any information that describes an object.

StorageGRID uses object metadata to track the locations of all objects across the grid and to manage each object's lifecycle over time.

Object metadata includes information such as the following:

- System metadata, including a unique ID for each object, the object name, the name of the S3 bucket, the tenant account name or ID, the logical size of the object, the date and time the object was first created, and the date and time the object was last modified.
- The current storage location of each object's replicate copy or erasure-coded fragment.
- Any custom user metadata key-value pairs associated with the object.
- For S3 objects, any object tag key-value pairs associated with the object
- For segmented objects and multipart objects, segment identifiers, and data sizes.

Object metadata is customizable and expandable, making it flexible for applications to use. For detailed information about how and where StorageGRID stores object metadata, go to [Manage object metadata storage](#).

StorageGRID's Information lifecycle management (ILM) system is used to orchestrate the placement, duration, and ingest behavior for all object data in your StorageGRID system. ILM rules determine how StorageGRID stores objects over time using replicas of the objects or erasure coding the object across nodes and sites. This ILM system is responsible for the object data consistency within a grid.

Erasure coding

StorageGRID provides the ability to erasure code data at node level and the drive level. With StorageGRID appliances we erasure code the data stored on each node across all the drives within the node providing local protection against multiple disk failures causing data loss or interruptions. Rebuilds from drive failures are local to the node and do not require data replicated over the network.

Additionally, StorageGRID appliances use erasure coding schemes to store object data across the nodes within a site or spread across 3 or more sites in the StorageGRID system though StorageGRID's ILM rules protecting against node failure.

Erasure coding provides a storage layout that is resilient to node and site failures with a lower overhead than replication. All StorageGRID erasure coding schemes are deployable in a single site provided the minimum number of nodes required to store the data chunks are met. This means for an EC scheme of 4+2 there needs to be a minimum of 6 nodes available to receive the data.

Erasure-coding scheme ($k+m$)	Minimum number of deployed sites	Recommended number of Storage Nodes at each site	Total recommended number of Storage Nodes	Site loss protection?	Storage overhead
4+2	3	3	9	Yes	50%
6+2	4	3	12	Yes	33%
8+2	5	3	15	Yes	25%
6+3	3	4	12	Yes	50%
9+3	4	4	16	Yes	33%
2+1	3	3	9	Yes	50%
4+1	5	3	15	Yes	25%
6+1	7	3	21	Yes	17%
7+5	3	5	15	Yes	71%

Metadata consistency

In StorageGRID, metadata is typically stored with three replicas per site to ensure consistency and availability. This redundancy helps maintain data integrity and accessibility even in the event of a failure.

The default consistency is defined at a grid wide level. Users can change the consistency at the bucket level at any time.

The bucket consistency options available in StorageGRID are:

- **All:** Provides the highest level of consistency. All nodes in the grid receive the data immediately, or the request will fail.
- **Strong-global:**
 - **Legacy Strong Global:** Guarantees read-after-write consistency for all client requests across all sites.
 - This is the default behavior for all systems upgraded from 11.9 or older to 12.0 without manually changing to the new Quorum Strong Global.
 - **Quorum Strong-global:** Guarantees read-after-write consistency for all client requests across all sites. Offers consistency for multiple nodes or even a site failure if metadata replica quorum is achievable.
 - This is the default behavior for all systems newly installed at 12.0 or higher.
 - QUORUM consistency is defined as a quorum of Storage Node metadata replicas, where each site has 3 metadata replicas. It may be calculated as follows: $1+((N*3)/2)$ where N is the total number of sites
 - For example, a minimum of 5 replicas must be made from a 3-site grid with a maximum of 3 replicas within a site.
- **Strong-site:** Guarantees read-after-write consistency for all client requests within a site.
- **Read-after-new-write(default):** Provides read-after-write consistency for new objects and eventual

consistency for object updates. Offers high availability and data protection guarantees. Recommended for most cases.

- **Available:** Provides eventual consistency for both new objects and object updates. For S3 buckets, use only as required (for example, for a bucket that contains log values that are rarely read, or for HEAD or GET operations on keys that don't exist). Not supported for S3 FabricPool buckets.

Object data consistency

While metadata is automatically replicated within and across sites, object data storage placement decisions are up to you. Object data can be stored in replicas within and across sites, erasure coded within or across sites, or a combination of replicas and erasure coded storage schemes. ILM rules can apply to all objects, or be filtered to only apply to certain objects, buckets, or tenants. ILM rules define how objects are stored, replicas and/or erasure coded, how long objects are stored in those locations, if the number of replicas or erasure coding scheme should change, or locations should change over time.

Each ILM rule will be configured with one of three ingest behaviors for protecting objects: Dual commit, balanced or strict.

The dual commit option will make two copies on any two different storage nodes in the grid immediately and return the request to be successful to the client. The node selection will be tried within the site of the request but may use nodes of another site in some circumstances. The object is added to the ILM queue to be evaluated and placed according to the ILM rules.

The balanced option evaluates the object against the ILM policy immediately and places the object synchronously before returning the request to be successful to the client. If the ILM rule cannot be met immediately due to an outage or inadequate storage to meet the placement requirements, then dual commit will be used instead. Once the issue is resolved, ILM will automatically place the object based on the defined rule.

The strict option evaluates the object against the ILM policy immediately and places the object synchronously before returning the request to be successful to the client. If the ILM rule cannot be met immediately due to an outage or inadequate storage to meet the placement requirements, then the request will fail, and the client will need to retry.

Load balancing

StorageGRID can be deployed with client access through the integrated gateway nodes, an external 3rd party load balancer, DNS round robin, or directly to a storage node. Multiple gateway nodes can be deployed in a site and configured in high availability groups providing automated failover and fail back in the event of a gateway node outage. You can combine load balancing methods in a solution to provide a single point of access for all sites in a solution.

The gateway nodes will balance the load between the storage nodes in the site where the gateway node resides by default. StorageGRID can be configured to allow the gateway nodes to balance load using nodes from multiple sites. This configuration would add the latency between those sites to the response latency to the client's requests. This should only be configured if the total latency is acceptable to the clients.

Ensuring an RTO of zero can be achieved with a combination of local and global load balancing. Ensuring uninterrupted client access requires load balancing of client requests. A StorageGRID solution can contain many gateway nodes and high availability groups in each site. To provide uninterrupted access for clients in any site even in a site failure, you should configure an external load balancing solution in combination with StorageGRID Gateway nodes. Configure Gateway node high availability groups that manage the load within each site and use the external load balancer to balance the load across the high availability groups. The external load balancer must be configured to perform a health check to ensure requests are sent only to

operational sites. For more information on load balancing with StorageGRID please see the [StorageGRID load balancer technical report](#).

Requirements for Zero RPO with StorageGRID

To achieve zero Recovery Point Objective (RPO) in an object storage system, it is crucial that at the time of failure:

- Both metadata and object contents are in sync and considered consistent
- Object content remain accessible despite the failure.

For a multi-site deployment, Quorum Strong Global is the preferred consistency model to ensure metadata is synchronized across all sites, making it essential for meeting the zero RPO requirement.

Objects in the storage system are stored based on Information Lifecycle Management (ILM) rules, which dictate how and where data is stored throughout its lifecycle. For synchronous replication, one can consider between Strict execution or Balanced Execution.

- Strict execution of these ILM rules is necessary for zero RPO because it ensures that objects are placed in the defined locations without any delay or fallback, maintaining data availability and consistency.
- StorageGRID's ILM balance ingest behavior provides a balance between high availability and resiliency, allowing users to continue ingesting data even in the event of a site failure.

Synchronous Deployments across multiple sites

Multi-site solutions: StorageGRID allows you to replicate objects across multiple sites within the grid synchronously. By setting up Information Lifecycle Management (ILM) rules with balance or strict behavior, objects are placed immediately in the specified locations. Configuring bucket consistency level to Quorum Strong Global will ensure synchronous metadata replication as well. StorageGRID uses a single global namespace, storing object placement locations as metadata, so every node knows where all copies or erasure coded pieces are located. If an object can't be retrieved from the site where the request was made, it will be automatically retrieved from a remote site without needing failover procedures.

Once the failure is resolved, no manual failback efforts are required. The replication performance depends on the site with the lowest network throughput, highest latency, and lowest performance. A site's performance is based on the number of nodes, CPU core count and speed, memory, drive quantity, and drive types.

Multi-grid solutions: StorageGRID can replicate tenants, users, and buckets between multiple StorageGRID systems using Cross-Grid replication (CGR). CGR can extend select data to more than 16 sites, increase the usable capacity of your object store, and provide disaster recovery. The replication of buckets with CGR includes objects, object versions, and metadata, and can be bi-directional or one-way. The recovery point objective (RPO) depends on the performance of each StorageGRID system and the network connections between them.

Summary:

- Intra-grid replication includes both synchronous and asynchronous replication, configurable using ILM ingest behavior and metadata consistency control.
- Inter-grid replication is asynchronous only.

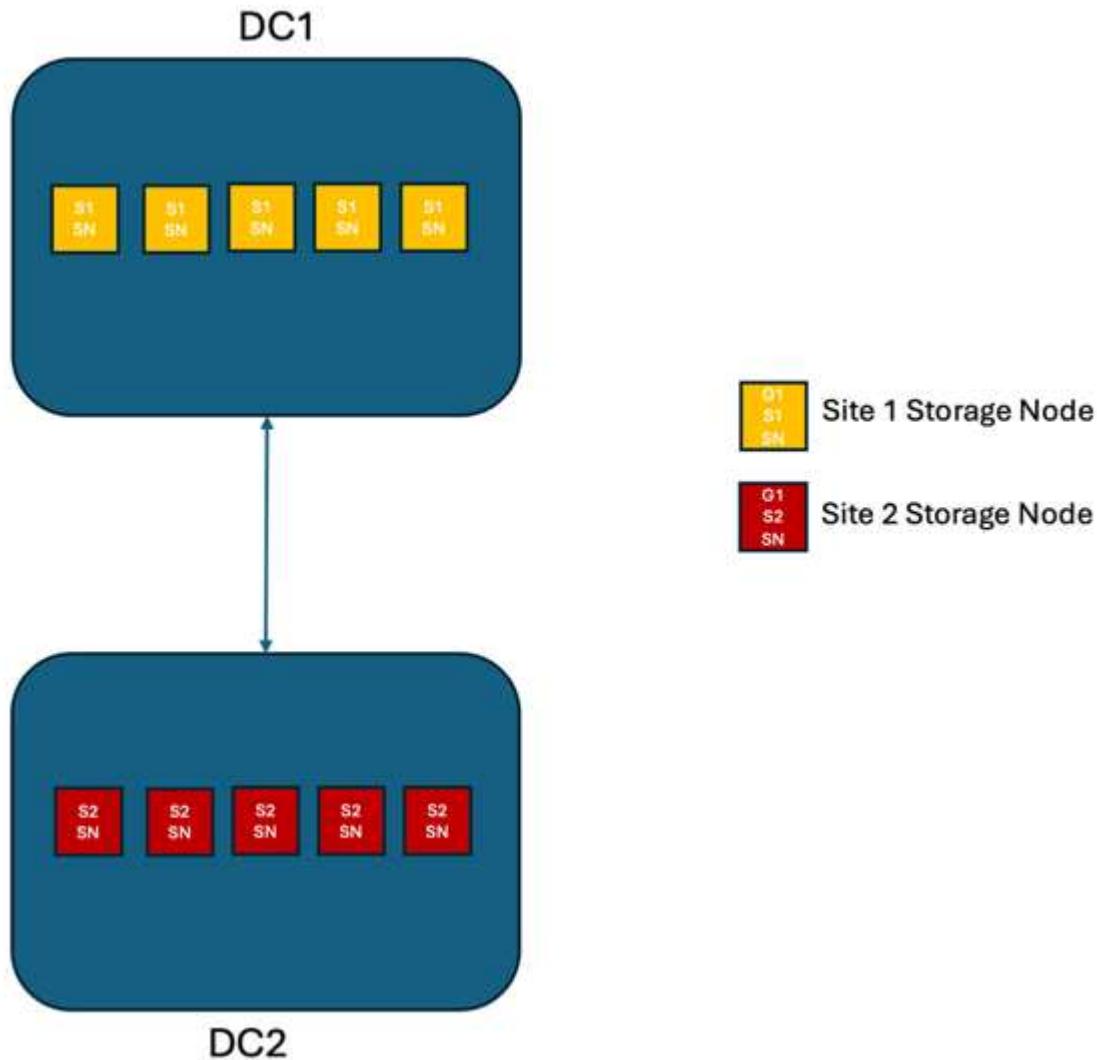
A Single Grid Multi-site deployment

In the following scenarios the StorageGRID solutions are configured with an optional external load balancer

managing requests to the integrated load balancer high availability groups. This will achieve an RTO of zero in addition to an RPO of zero. ILM is configured with Balanced ingest protection for synchronous placement. Each bucket is configured with the Quorum version of Strong Global consistency model for grids of 3 or more sites and the Legacy version of Strong Global consistency for 2 sites.

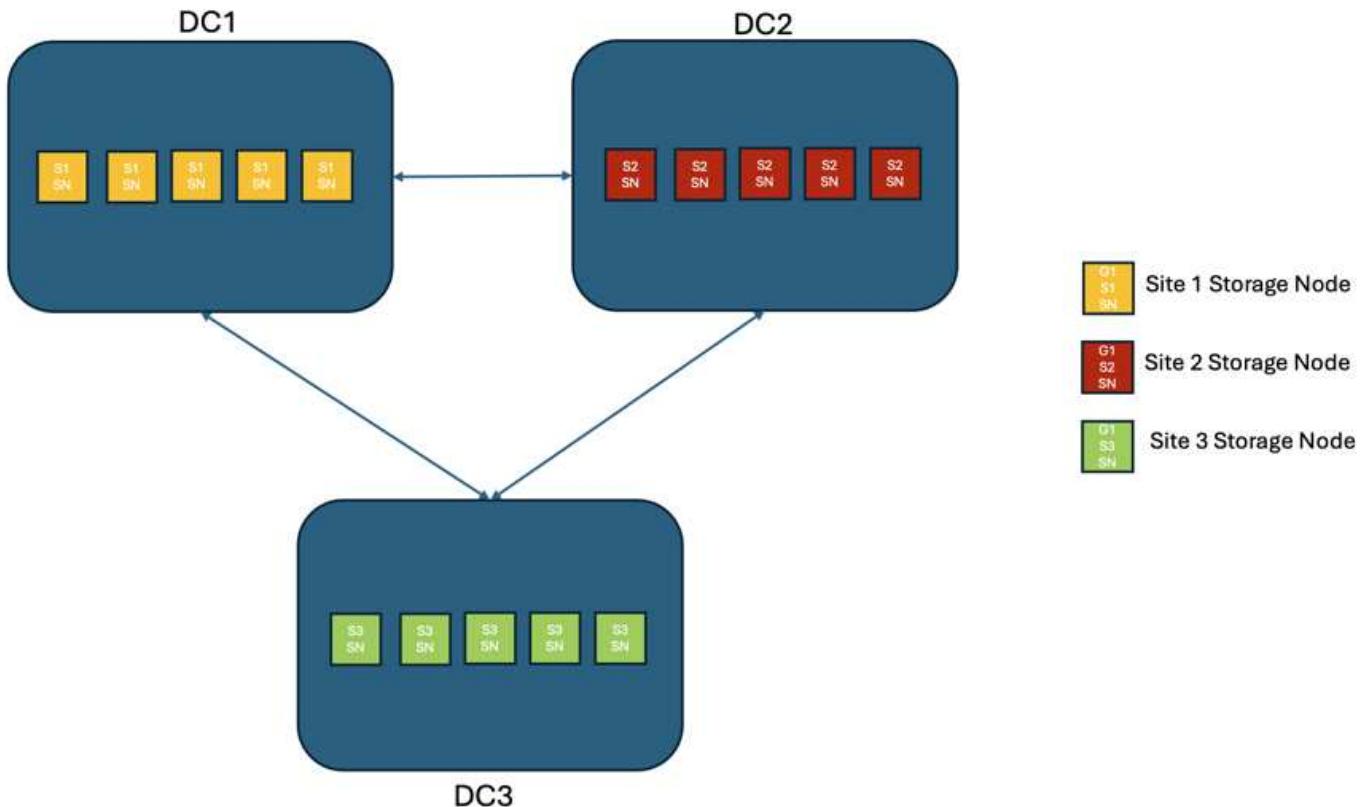
Scenario 1:

In a two site StorageGRID solution, there are at least two replicas of every object and 6 replicas of all metadata. Upon failure recovery, updates from the outage will synchronize to the recovered site/nodes automatically. With only 2 sites it is not likely to achieve a zero RPO in failure scenarios beyond a full site loss.



Scenario 2:

In a StorageGRID solution of three or more sites, there are at least 3 replicas or 3 EC chunks of every object and 9 replicas of all metadata. Upon failure recovery, updates from the outage will synchronize to the recovered site/nodes automatically. With three or more sites it is possible to achieve a zero RPO.



Multi-site failure scenarios

Failure	2-site Outcome Legacy Strong Global	3 or more sites outcome Quorum Strong Global
Single node drive failure	Each appliance uses multiple disk groups and can sustain at least 1 drive per group failing without interruption or data loss.	Each appliance uses multiple disk groups and can sustain at least 1 drive per group failing without interruption or data loss.
Single node failure in one site	No interruption to operations or data loss.	No interruption to operations or data loss.
Multiple node failure in one site	Disruption to client operations directed to this site but no data loss. Operations directed to the other site remain uninterrupted and no data loss.	Operations are directed to all other sites and remain uninterrupted and no data loss.

Failure	2-site Outcome Legacy Strong Global	3 or more sites outcome Quorum Strong Global
Single node failure at multiple sites	<p>No disruption or data loss if:</p> <ul style="list-style-type: none"> • At least one replicate copy exists in the grid • Sufficient EC chunks exist in the grid <p>Operations disrupted and risk of data loss if:</p> <ul style="list-style-type: none"> • No replicate copies exist • Insufficient EC chunks exist 	<p>No disruption or data loss if:</p> <ul style="list-style-type: none"> • At least a single replicate copy exists in the grid • Sufficient EC chunks exist in the grid <p>Operations disrupted and risk of data loss if:</p> <ul style="list-style-type: none"> • No replicate copies exist • Insufficient EC chunks exist to retrieve the object
Single site failure	<p>Some client operations will be interrupted until the failure is resolved.</p> <p>GET and HEAD operations will continue without interruption.</p> <p>Reduce bucket consistency to read-after-new-write or lower to continue operations uninterrupted in this failure state.</p>	No interruption to operations or data loss.
Single site plus single node failures	<p>Some client operations will be interrupted until either the failure is resolved.</p> <p>HEAD operations will continue without interruption.</p> <p>GET operations will continue without interruption if a replicate copy or sufficient EC chunks exist.</p> <p>Reduce bucket consistency to read-after-new-write or lower to continue operations uninterrupted in this failure state.</p>	<p>No interruption to operations or data loss.</p> <p>Possible data loss depending on the number of replicate copies.</p> <p>Local Erasure coding can prevent data loss.</p>

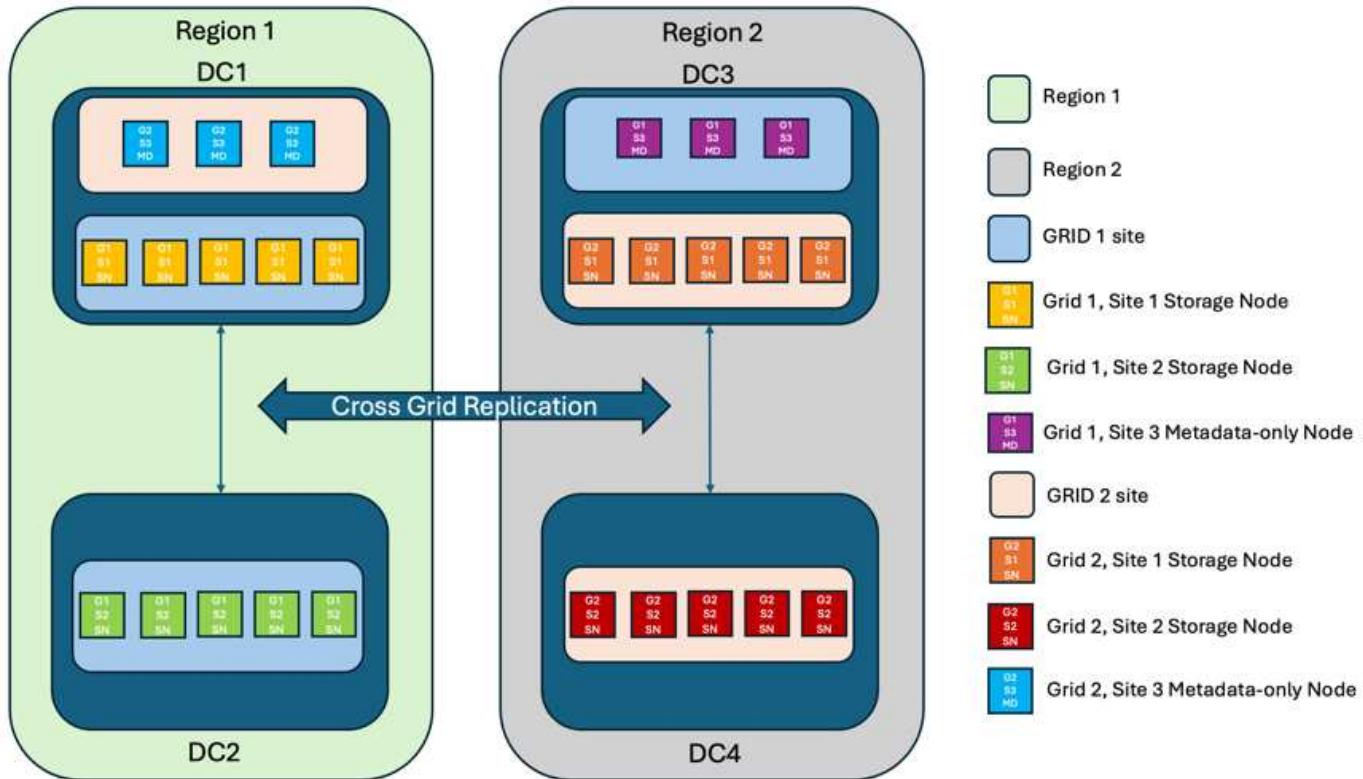
Failure	2-site Outcome Legacy Strong Global	3 or more sites outcome Quorum Strong Global
Single site plus a node from each remaining site	<p>Only two sites exist.</p> <p>See: Single site plus a single node.</p>	<p>Operations will be disrupted If metadata replica quorum cannot be met.</p> <p>Reduce bucket consistency to read-after-new-write or lower to continue operations uninterrupted in this failure state.</p> <p>Possible data loss for permanent failure depending on the number of replicate copies.</p> <p>Local Erasure coding can prevent data loss.</p>
Multi-site failure	<p>No operational sites remain.</p> <p>Data will be lost if at least one site cannot be recovered in its entirety.</p>	<p>Operations will be disrupted If metadata replica quorum cannot be met.</p> <p>Reduce bucket consistency to read-after-new-write or lower to continue operations uninterrupted in this failure state.</p> <p>Possible data loss for permanent failure if not enough erasure coded chunks remain.</p> <p>Local Erasure coding or replicate copies can prevent data loss.</p>
Network isolation of a site	<p>client operations will be interrupted until either the failure is resolved.</p> <p>Reduce bucket consistency to read-after-new-write or lower to continue operations uninterrupted in this failure state.</p> <p>No data loss</p>	<p>Operations will be disrupted for the isolated site, but no data loss.</p> <p>Reduce bucket consistency to read-after-new-write or lower to continue operations uninterrupted in this failure state.</p> <p>No disruption to operations in the remaining sites and no data loss.</p>

A multi-site multi-grid deployment

To add an extra layer of redundancy, this scenario will employ two StorageGRID Clusters and use cross-grid replication to keep them in sync. For this solution each StorageGRID clusters will have three sites. Two sites will be used for object storage and metadata while the third site will be used solely for metadata. Both systems will be configured with a balanced ILM rule to synchronously store the objects using erasure coding in each of the two data sites. Buckets will be configured with the Quorum Strong Global consistency model. Each grid will be configured with bi-directional cross-grid replication on every bucket. This provides the asynchronous replication between the regions. Optionally a global load balancer can be implemented to manage requests to

the integrated load balancer high availability groups of both StorageGRID systems to achieve a zero RPO.

The solution will use four locations equally divided into two regions. Region 1 will contain the 2 storage sites of grid 1 as the primary grid of the region and the metadata site of grid 2. Region 2 will contain the 2 storage sites of grid 2 as the primary grid of the region and the metadata site of grid 1. In each region the same location can house the storage site of the primary grid of the region as well as the metadata only site of the other regions grid. Using metadata only nodes as the third site will provide the consistency required for the metadata and not duplicate the storage of objects in that location.



This solution with four separate locations provides complete redundancy of two separate StorageGRID systems maintaining an RPO of 0 and will make use of both multi-site synchronous replication, and multi-grid asynchronous replication. Any single site can fail while maintaining uninterrupted client operations on both StorageGRID systems.

In this solution, there are four erasure coded copies of every object and 18 replicas of all metadata. This allows for multiple failure scenarios without client operations impact. Upon failure recovery updates from the outage will synchronize to the failed site/nodes automatically.

Multisite, multi-grid failure scenarios

Failure	Outcome
Single node drive failure	Each appliance uses multiple disk groups and can sustain at least 1 drive per group failing without interruption or data loss.
Single node failure in one site in a grid	No interruption to operations or data loss.
Single node failure in one site in each grid	No interruption to operations or data loss.
Multiple node failure in one site in a grid	No interruption to operations or data loss.

Failure	Outcome
Multiple node failure in one site in each grid	No interruption to operations or data loss.
Single node failure at multiple sites in a grid	No interruption to operations or data loss.
Single node failure at multiple sites in each grid	No interruption to operations or data loss.
Single site failure in a grid	No interruption to operations or data loss.
Single site failure in each grid	No interruption to operations or data loss.
Single site plus single node failures in a grid	No interruption to operations or data loss.
Single site plus a node from each remaining site in a single grid	No interruption to operations or data loss.
Single location failure	No interruption to operations or data loss.
Single location failure in each grid DC1 & DC3	Operations will be disrupted until either the failure is resolved, or the bucket consistency is lowered; each grid has lost 2 sites All data still exists at 2 locations
Single location failure in each grid DC1 & DC4 or DC2 & DC3	No interruption to operations or data loss.
Single location failure in each grid DC2 & DC4	No interruption to operations or data loss.
Network isolation of a site	Operations will be disrupted for the isolated site but no data will be lost No disruption to operations in the remaining sites or data loss.

Conclusion

Achieving zero Recovery Point Objective (RPO) with StorageGRID is a critical goal for ensuring data durability and availability in the event of site failures. By leveraging StorageGRID's robust replication strategies, including multi-site synchronous replication and multi-grid asynchronous replication, organizations can maintain uninterrupted client operations and ensure data consistency across multiple locations. The implementation of Information Lifecycle Management (ILM) policies and the use of metadata-only nodes further enhance the system's resilience and performance. With StorageGRID, businesses can confidently manage their data, knowing that it remains accessible and consistent even in the face of complex failure scenarios. This comprehensive approach to data management and replication underscores the importance of meticulous planning and execution in achieving zero RPO and safeguarding valuable information.

Create Cloud Storage Pool for AWS or Google Cloud

You can use a Cloud Storage Pool if you want to move StorageGRID objects to an external S3 bucket. The external bucket can belong to Amazon S3 (AWS) or Google Cloud.

What you'll need

- StorageGRID 11.6 has been configured.
- You have already set up an external S3 bucket on AWS or Google Cloud.

Steps

1. In the Grid Manager, navigate to **ILM > Storage Pools**.
2. In the Cloud Storage Pools section of the page, select **Create**.

The Create Cloud Storage Pool pop-up appears.

3. Enter a display name.
4. Select **Amazon S3** from the Provider Type drop-down list.

This provider type works for AWS S3 or Google Cloud.

5. Enter the URI for the S3 bucket to be used for the Cloud Storage Pool.

Two formats are allowed:

`https://host:port`

`http://host:port`

6. Enter the S3 bucket name.

The name you specify must exactly match the S3 bucket's name; otherwise, Cloud Storage Pool creation fails. You cannot change this value after the Cloud Storage Pool is saved.

7. Optionally, enter the Access Key ID and the Secret Access Key.
8. Select **Do Not Verify Certificate** from the drop-down.
9. Click **Save**.

Expected result

Confirm that a Cloud Storage Pool has been created for Amazon S3 or Google Cloud.

By Jonathan Wong

Create Cloud Storage Pool for Azure Blob Storage

You can use a Cloud Storage Pool if you want to move StorageGRID objects to an external Azure container.

What you'll need

- StorageGRID 11.6 has been configured.
- You have already set up an external Azure container.

Steps

1. In the Grid Manager, navigate to **ILM > Storage Pools**.
2. In the Cloud Storage Pools section of the page, select **Create**.

The Create Cloud Storage Pool pop-up appears.

3. Enter a display name.
4. Select **Azure Blob Storage** from the Provider Type drop-down list.
5. Enter the URI for the S3 bucket to be used for the Cloud Storage Pool.

Two formats are allowed:

`https://host:port`

`http://host:port`

6. Enter the Azure container name.

The name you specify must exactly match the Azure container name; otherwise, Cloud Storage Pool creation fails. You cannot change this value after the Cloud Storage Pool is saved.

7. Optionally, enter the Azure container's associated account name and account key for authentication.
8. Select **Do Not Verify Certificate** from the drop-down.
9. Click **Save**.

Expected result

Confirm that a Cloud Storage Pool has been created for Azure Blob Storage.

By Jonathan Wong

Use a Cloud Storage Pool for backup

You can create an ILM rule to move objects into a Cloud Storage Pool for backup..

What you'll need

- StorageGRID 11.6 has been configured.
- You have already set up an external Azure container.

Steps

1. In the Grid Manager, navigate to **ILM > Rules > Create**.
2. Enter a description.
3. Enter a criterion to trigger the rule.
4. Click **Next**.
5. Replicate the object to Storage Nodes.
6. Add a placement rule.
7. Replicate the object to the Cloud Storage Pool
8. Click **Next**.
9. Click **Save**.

Expected result

Confirm that the retention diagram shows the objects stored locally in StorageGRID and in a Cloud Storage

Pool for backup.

Confirm that, when the ILM rule is triggered, a copy exists in the Cloud Storage Pool and you can retrieve the object locally without doing an object restore.

By Jonathan Wong

Configure StorageGRID search integration service

This guide provides detailed instructions for configuring NetApp StorageGRID search integration service with either Amazon OpenSearch Service or on-premises Elasticsearch.

Introduction

StorageGRID supports three types of platform services.

- **StorageGRID CloudMirror replication.** Mirror specific objects from a StorageGRID bucket to a specified external destination.
- **Notifications.** Per-bucket event notifications to send notifications about specific actions performed on objects to a specified external Amazon Simple Notification Service (Amazon SNS).
- **Search integration service.** Send Simple Storage Service (S3) object metadata to a specified Elasticsearch index where you can search or analyze the metadata by using the external service.

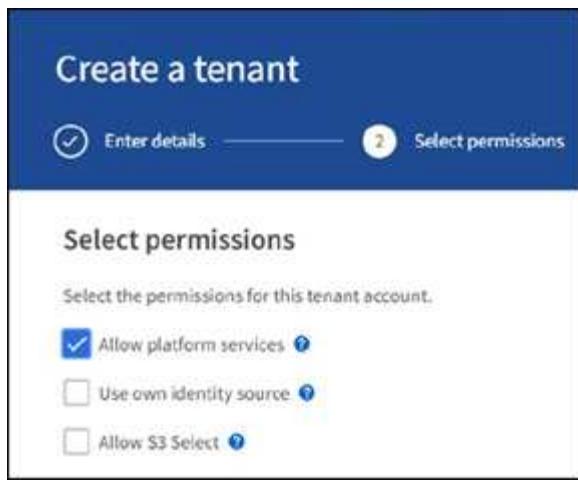
Platform services are configured by the S3 tenant through the Tenant Manager UI. For more information, see [Considerations for using platform services](#).

This document serves as a supplement to the [StorageGRID 11.6 Tenant Guide](#) and provides step by step instructions and examples for the endpoint and bucket configuration for search integration services. The Amazon Web Services (AWS) or on-premises Elasticsearch setup instructions included here are for basic testing or demo purposes only.

Audiences should be familiar with Grid Manager, Tenant Manager, and have access to the S3 browser to perform basic upload (PUT) and download (GET) operations for StorageGRID search integration testing.

Create tenant and enable platform services

1. Create an S3 tenant by using Grid Manager, enter a display name, and select the S3 protocol.
2. On the Permission page, select the Allow Platform Services option. Optionally, select other permissions, if necessary.



3. Set up the tenant root user initial password or, if identity federation is enabled on the grid, select which federated group has root access permission to configure the tenant account.
4. Click Sign In As Root and select Bucket: Create and Manage Buckets.

This takes you to the Tenant Manager page.

5. From Tenant Manager, select My Access Keys to create and download the S3 access key for later testing.

Search integration services with Amazon OpenSearch

Amazon OpenSearch (formerly Elasticsearch) service setup

Use this procedure for a quick and simple setup of the OpenSearch service for testing/demo purposes only. If you are using on-premises Elasticsearch for search integration services, see the section [Search integration services with on premises Elasticsearch](#).



You must have a valid AWS console login, access key, secret access key, and permission to subscribe to the OpenSearch service.

1. Create a new domain using the instructions from [AWS OpenSearch Service Getting Started](#), except for the following:
 - Step 4. Domain name: sgdemo
 - Step 10. Fine-grained access control: deselect the Enable Fine-Grained Access Control option.
 - Step 12. Access policy: select Configure Level Access Policy, select the JSON tab to modify the access policy by using the following example:
 - Replace the highlighted text with your own AWS Identity and Access Management (IAM) ID and user name.
 - Replace the highlighted text (the IP address) with the public IP address of your local computer that you used to access the AWS console.
 - Open a browser tab to <https://checkip.amazonaws.com> to find your public IP.

```
{  
  
    "Version": "2012-10-17",  
    "Statement": [  
        {  
            "Effect": "Allow",  
            "Principal":  
                {"AWS": "arn:aws:iam:: nnnnnnn:user/xyzabc"},  
            "Action": "es:*",  
            "Resource": "arn:aws:es:us-east-1:nnnnnn:domain/sgdemo/*"  
        },  
        {  
            "Effect": "Allow",  
            "Principal": {"AWS": "*"},  
            "Action": [  
                "es:ESHttp*"  
            ],  
            "Condition": {  
                "IpAddress": {  
                    "aws:SourceIp": [ "nnn.nnn.nn.n/nn"  
                ]  
            }  
        },  
        "Resource": "arn:aws:es:us-east-1:nnnnnn:domain/sgdemo/*"  
    ]  
}
```

Fine-grained access control

Fine-grained access control provides numerous features to help you keep your data secure. Features include document-level security, field-level security, read-only users, and OpenSearch Dashboards/Kibana tenants. Fine-grained access control requires a master user. [Learn more](#)

Enable fine-grained access control

SAML authentication for OpenSearch Dashboards/Kibana

SAML authentication lets you use your existing identity provider for single sign-on for OpenSearch Dashboards/Kibana. [Learn more](#)

Prepare SAML authentication

 To use SAML authentication, you must first enable fine-grained access control.

Amazon Cognito authentication

Enable to use Amazon Cognito authentication for OpenSearch Dashboards/Kibana. Amazon Cognito supports a variety of identity providers for username-password authentication. [Learn more](#)

Enable Amazon Cognito authentication

Access policy

Access policies control whether a request is accepted or rejected when it reaches the Amazon OpenSearch Service domain. If you specify an account, user, or role in this policy, you must sign your requests. [Learn more](#)

Domain access policy

- Only use fine-grained access control
Allow open access to the domain.
- Do not set domain level access policy
All requests to the domain will be denied.
- Configure domain level access policy

Visual editor

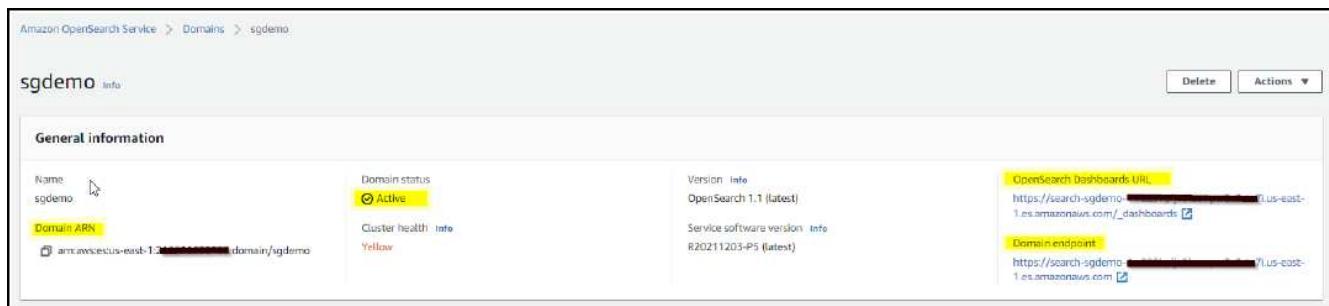
JSON

Import policy

Access policy

```
3+ "Statement": [
4+   {
5+     "Effect": "Allow",
6+     "Principal": {
7+       "AWS": "arn:aws:iam::210111111111:user/sgdemo"
8+     },
9+     "Action": "es:*",
10+    "Resource": "arn:aws:es:us-east-1:210111111111:domain/sgdemo/*"
11+  },
12+  {
13+    "Effect": "Allow",
14+    "Principal": {
15+      "AWS": "*"
16+    },
17+    "Action": [
18+      "es:ESHttp*"
19+    ],
20+    "Condition": {
21+      "IpAddress": [
22+        "aws:SourceIp": [
23+          "216.210.111.24"
24+
25+        ]
26+      ],
27+      "Resource": "arn:aws:es:us-east-1:210111111111:domain/sgdemo/*"
28+    }
29+  }
30+ ]
```

2. Wait 15 to 20 minutes for the domain to become active.



Amazon OpenSearch Service > Domains > sgdemo

sgdemo [Info](#)

General information

Name	sgdemo	Domain status	Active	Version	Info
Domain ARN	arn:aws:es:us-east-1:123456789012:domain/sgdemo	Cluster health	Yellow	OpenSearch 1.1 (latest)	OpenSearch Dashboards URL
				R20211203-PS (latest)	https://search- sgdemo -1.es.amazonaws.com/_dashboards
					Domain endpoint
					https://search- sgdemo -1.es.amazonaws.com

3. Click OpenSearch Dashboards URL to open the domain in a new tab to access the dashboard. If you get an access denied error, verify that the access policy source IP address is correctly set to your computer public IP to allow access to the domain dashboard.
4. On the dashboard welcome page, select Explore On Your Own. From the menu, go to Management → Dev Tools
5. Under Dev Tools → Console , enter `PUT <index>` where you use the index name 'sgmetadata' in the following example. Click the small triangle symbol to execute the PUT command. The expected result displays on the right panel as shown in the following example screenshot.



OpenSearch Dashboards

☰ Dev Tools

Console

History Settings Help

```
1 PUT sgmetadata
```

Execute (triangle icon)

```
1 {  
2   "acknowledged" : true,  
3   "shards_acknowledged" : true,  
4   "index" : "sgmetadata"  
5 }
```

6. Verify that the index is visible from Amazon OpenSearch UI under sgdomain > Indices.

Successfully updated to service software version 820211203-15

Amazon OpenSearch Service > Domains > sgdemo

sgdemo Info

General information

Name: sgdemo	Domain status: Active	Version info: OpenSearch 1.1 (latest)	OpenSearch Dashboards URL: https://search-sgdemo-1234567890.us-east-1.es.amazonaws.com/_dashboards
Domain ARN: arn:aws:es:us-east-1:211234567890:domain/sgdemo	Cluster health: Yellow	Service software version: 820211203-15 (latest)	Domain endpoint: https://search-sgdemo-1234567890.us-east-1.es.amazonaws.com

Indices [2]

Indices is the method by which search engines organize data by field names. Before you can search data, you must index it. Learn more [»](#)

Index	Document count	Size (byte)	Query total	Mapping type	Field mappings	Actions
aribon_1	1	5.08 kB	10	dynamic,meta,proper...		View
sgmedata	0	1.02 kB	0			View

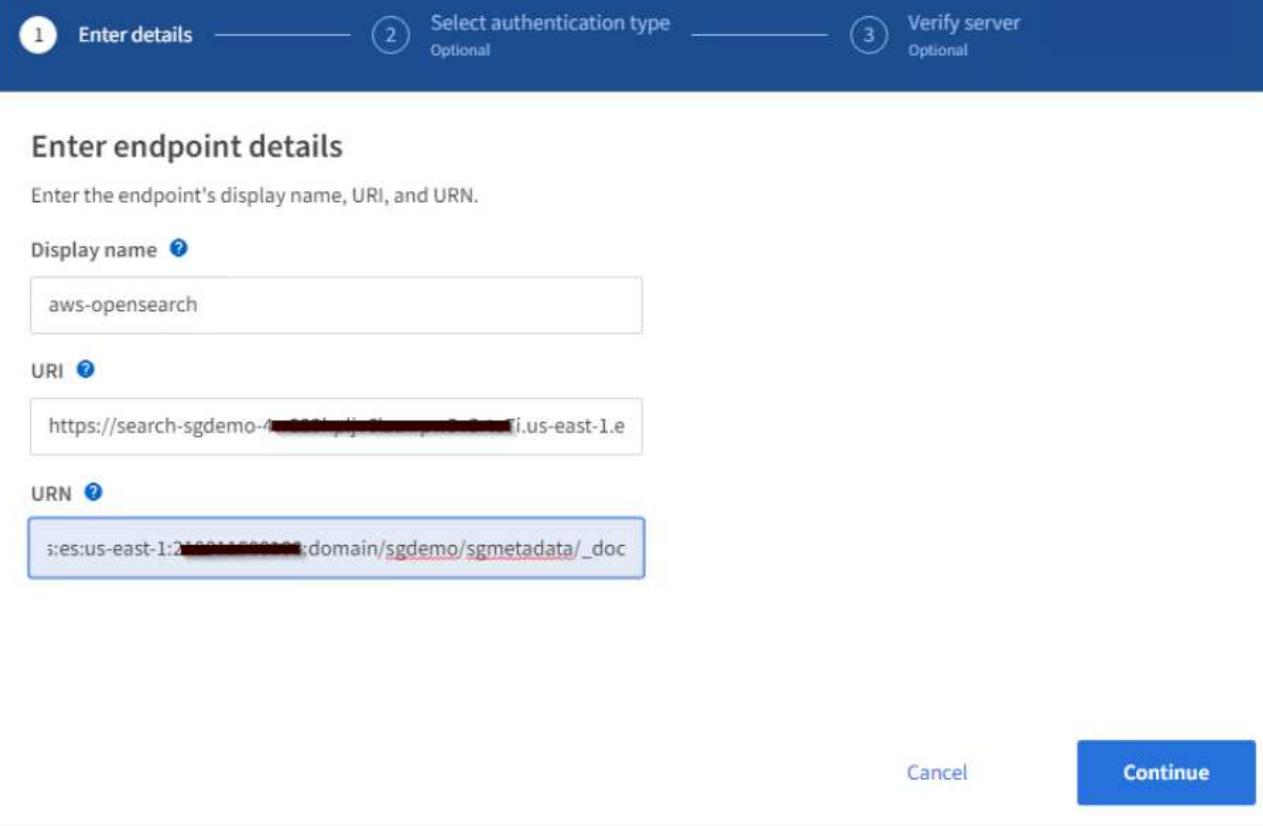
Platform services endpoint configuration

To configure the platform services endpoints, follow these steps:

1. In Tenant Manager, go to STORAGE(S3) > Platform services endpoints.
2. Click Create Endpoint, enter the following, and then click Continue:
 - Display name **example aws-opensearch**
 - The domain endpoint in the **example** screenshot under Step 2 of the preceding procedure in the **URI field**.
 - The domain ARN used in Step 2 of the preceding procedure in the **URN field** and add `/<index>/_doc` to the end of ARN.

In this example, URN becomes `arn:aws:es:us-east-1:211234567890:domain/sgdemo/sgmedata/_doc`.

Create endpoint



3. To access the Amazon OpenSearch sgdomain, choose Access Key as the authentication type and then enter the Amazon S3 access key and secret key. To go the next page, click Continue.

Create endpoint

4. To verify the endpoint, select Use Operating System CA Certificate and Test and Create Endpoint. If verification is successful, an endpoint screen similar to the following figure displays. If verification fails, verify that the URN includes /<index>/_doc at the end of the path and the AWS access key and secret key are correct.

Platform services endpoints

A platform services endpoint stores the information StorageGRID needs to use an external resource as a target for a platform service (CloudMirror replication, notifications, or search integration). You must configure an endpoint for each platform service you plan to use.

1 endpoint

[Create endpoint](#)

Delete endpoint	Display name	Last error	Type	URI	URN
<input type="checkbox"/>	aws-opensearch		Search	https://search-sgdemo-42221521123232.us-east-1.es.amazonaws.com/	arn:aws:es:us-east-1:23232323232323232323232323232323:domain/sgdemo/sgmetadata/_doc

Search integration services with on premises Elasticsearch

On premises Elasticsearch setup

This procedure is for a quick setup of on premises Elasticsearch and Kibana using docker for testing purposes only. If the Elasticsearch and Kibana server already exists, go to Step 5.

1. Follow this [Docker installation procedure](#) to install docker. We use the [CentOS Docker install procedure](#) in this setup.

```
sudo yum install -y yum-utils
sudo yum-config-manager --add-repo
https://download.docker.com/linux/centos/docker-ce.repo
sudo yum install docker-ce docker-ce-cli containerd.io
sudo systemctl start docker
```

- To start docker after reboot, enter the following:

```
sudo systemctl enable docker
```

- Set the `vm.max_map_count` value to 262144:

```
sysctl -w vm.max_map_count=262144
```

- To keep the setting after reboot, enter the following:

```
echo 'vm.max_map_count=262144' >> /etc/sysctl.conf
```

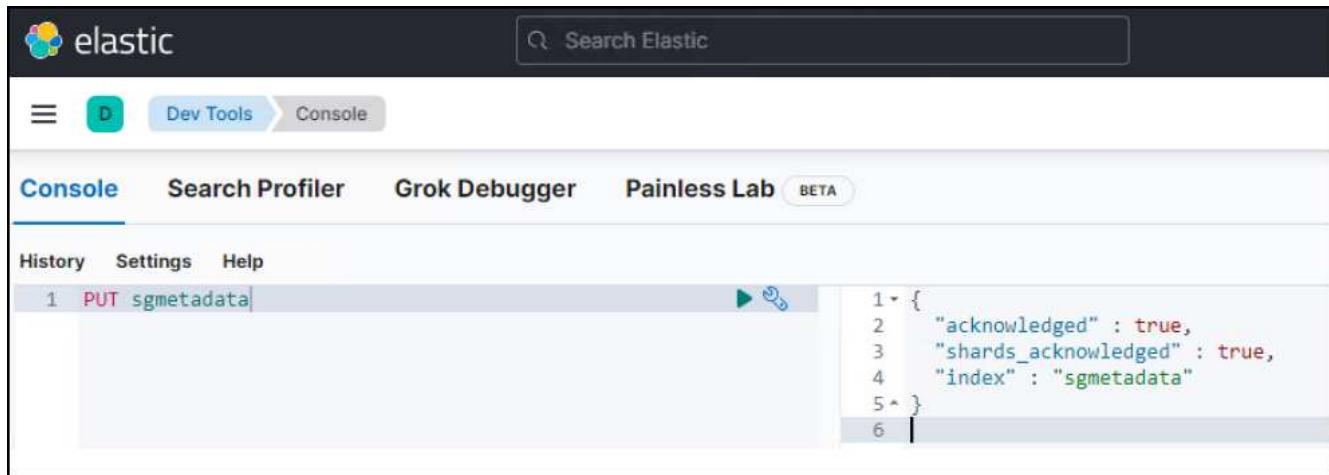
2. Follow the [Elasticsearch Quick start guide](#) self-managed section to install and run the Elasticsearch and Kibana docker. In this example, we installed version 8.1.



Note down the user name/password and token created by Elasticsearch, you need these to start the Kibana UI and StorageGRID platform endpoint authentication.

Elasticsearch Service	Self-managed
Install and run Elasticsearch <ol style="list-style-type: none">1. Install and start Docker Desktop.2. Run:<pre>docker network create elastic docker pull docker.elastic.co/elasticsearch/elasticsearch:8.1.0 docker run --name es-node01 --net elastic -p 9200:9200 -p 9300:9300 -it</pre>	
<p>When you start Elasticsearch for the first time, the following security configuration occurs automatically:</p> <ul style="list-style-type: none">• Certificates and keys are generated for the transport and HTTP layers.• The Transport Layer Security (TLS) configuration settings are written to <code>elasticsearch.yml</code>.• A password is generated for the <code>elastic</code> user.• An enrollment token is generated for Kibana. <p> NOTE You might need to scroll back a bit in the terminal to view the password and enrollment token.</p> <p>3. Copy the generated password and enrollment token and save them in a secure location. These values are shown only when you start Elasticsearch for the first time. You'll use these to enroll Kibana with your Elasticsearch cluster and log in.</p> <p> NOTE If you need to reset the password for the <code>elastic</code> user or other built-in users, run the <code>elasticsearch-reset-password</code> tool. To generate new enrollment tokens for Kibana or Elasticsearch nodes, run the <code>elasticsearch-create-enrollment-token</code> tool. These tools are available in the Elasticsearch bin directory.</p> <h3>Install and run Kibana</h3> <p>To analyze, visualize, and manage Elasticsearch data using an intuitive UI, install Kibana.</p> <ol style="list-style-type: none">1. In a new terminal session, run:<pre>docker pull docker.elastic.co/kibana/kibana:8.1.0 docker run --name kib-01 --net elastic -p 5601:5601 docker.elastic.co/kibana/kibana:8.1.0</pre>2. To access Kibana, click the generated link in your terminal.<ol style="list-style-type: none">a. In your browser, paste the enrollment token that you copied and click the button to connect your Kibana instance with Elasticsearch.b. Log in to Kibana as the <code>elastic</code> user with the password that was generated when you started Elasticsearch.	

3. After the Kibana docker container has started, the URL link `https://0.0.0.0:5601` displays in the console. Replace 0.0.0.0 with the server IP address in the URL.
4. Log in to the Kibana UI by using user name `elastic` and the password generated by Elastic in the preceding step.
5. For first time login, on the dashboard welcome page, select **Explore On Your Own**. From the menu, select **Management > Dev Tools**.
6. On the Dev Tools Console screen, enter `PUT <index>` where you use this index for storing StorageGRID object metadata. We use the index name `sgmetadata` in this example. Click the small triangle symbol to execute the PUT command. The expected result displays on the right panel as shown in the following example screenshot.



The screenshot shows the Elasticsearch Dev Tools Console interface. The top navigation bar includes the elastic logo, a search bar, and tabs for Dev Tools (which is selected), Console, Search Profiler, Grok Debugger, and Painless Lab (BETA). Below the tabs, there are History, Settings, and Help links. The main area has a code editor with the following content:

```

1 PUT sgmetadata
2
3
4
5
6

```

On the right, the results panel shows the response:

```

1 { "acknowledged" : true,
2   "shards_acknowledged" : true,
3   "index" : "sgmetadata"
4 }
5
6

```

Platform services endpoint configuration

To configure endpoints for platform services, follow these steps:

1. On Tenant Manager, go to **STORAGE(S3) > Platform services endpoints**
2. Click **Create Endpoint**, enter the following, and then click **Continue**:
 - Display name **example: elasticsearch**
 - **URI: https://<elasticsearch-server-ip or hostname>:9200**
 - **URN: urn:<something>:es:::<some-unique-text>/<index-name>/_doc** where the index-name is the name you used on the Kibana console.
Example: `urn:local:es:::sgmd/sgmetadata/_doc`

Create endpoint

1 Enter details 2 Select authentication type 3 Verify server

Enter endpoint details

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

Display name ?

URI ?

URN ?

[Cancel](#)

[Continue](#)

3. Select Basic HTTP as the authentication type, enter the user name `elastic` and the password generated by the Elasticsearch installation process. To go to the next page, click Continue.

Authentication type ?

Select the method used to authenticate connections to the endpoint.

Basic HTTP

Username ?

Password ?

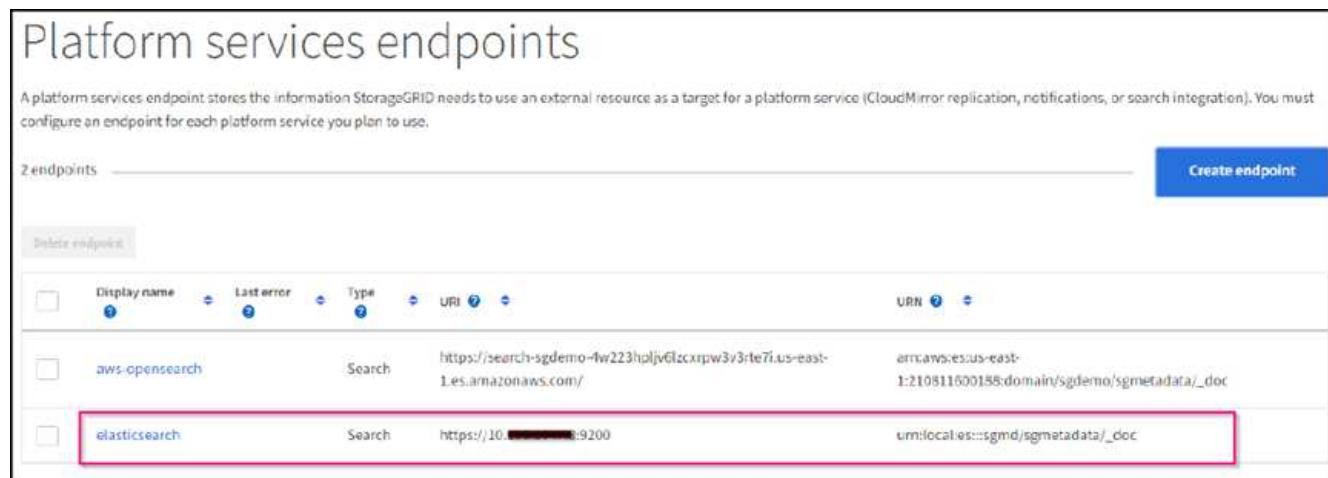


[Previous](#)

[Continue](#)

4. Select Do Not Verify Certificate and Test and Create Endpoint to verify the endpoint. If verification is

successful, an endpoint screen similar to the following screenshot displays. If the verification fails, verify the URN, URI, and username/password entries are correct.



Display name	Last error	Type	URI	URN
aws-opensearch		Search	https://search-sgdemo-fw223hpljv6lccxrpw3v3rte7i.us-east-1.es.amazonaws.com/	arn:aws:es:us-east-1:210811600158:domain/sgdemo/sgmetadata/_doc
elasticsearch		Search	https://10.0.0.1:9200	urn:local:es::sgmd/sgmetadata/_doc

Bucket search integration service configuration

After the platform service endpoint is created, the next step is to configure this service at bucket level to send object metadata to the defined endpoint whenever an object is created, deleted, or its metadata or tags are updated.

You can configure search integration by using Tenant Manager to apply a custom StorageGRID configuration XML to a bucket as follows:

1. In Tenant Manager, go to STORAGE(S3) > Buckets
2. Click Create Bucket, enter the bucket name (for example, sgmetadata-test) and accept the default us-east-1 region.
3. Click Continue > Create Bucket.
4. To bring up the bucket Overview page, click the bucket name, then select Platform Services.
5. Select the Enable Search Integration dialog box. In the provided XML box, enter the configuration XML using this syntax.

The highlighted URN must match the platform services endpoint that you defined. You can open another browser tab to access the Tenant Manager and copy the URN from the defined platform services endpoint.

In this example, we used no prefix, meaning that the metadata for every object in this bucket is sent to the Elasticsearch endpoint defined previously.

```

<MetadataNotificationConfiguration>
  <Rule>
    <ID>Rule-1</ID>
    <Status>Enabled</Status>
    <Prefix></Prefix>
    <Destination>
      <Urn> urn:local:es:::sgmd/sgmetadata/_doc</Urn>
    </Destination>
  </Rule>
</MetadataNotificationConfiguration>

```

6. Use S3 Browser to connect to StorageGRID with the tenant access/secret key, upload test objects to sgmetadata-test bucket and add tags or custom metadata to objects.

File	Size	Type	Last Modified	Storage Class
Koala.jpg	762.53 KB	JPG File	3/19/2022 12:39:52 AM	STANDARD
Lighthouse.jpg	548.12 KB	JPG File	3/19/2022 12:39:52 AM	STANDARD
test1.txt	45 bytes	Text Document	3/19/2022 12:39:52 AM	STANDARD
test2.txt	35 bytes	Text Document	3/19/2022 12:39:52 AM	STANDARD

Tasks (14) Permissions Http Headers Tags Properties Preview Versions EventLog

URL: <https://10.193.204.106:10445/sgmetadata-test/Koala.jpg> Copy

Key	Value
date	01-01-2020
owner	testuser
project	test
type	jpg

+ Add | Edit | Delete | Default Tags.. ✓ Apply changes | Reload

7. Use the Kibana UI to verify that the object metadata was loaded to sgmetadata's index.
 - From the menu, select Management > Dev Tools.
 - Paste the sample query to the console panel on the left and click the triangle symbol to execute it.

The query 1 sample result in the following example screenshot shows four records. This matches number of objects in the bucket.

```
GET sgmetadata/_search
{
  "query": {
    "match_all": {}
  }
}
```

```
1 GET sgmetadata/_search
2 {
3   "query": {
4     "match_all": {}
5   }
6 }
```

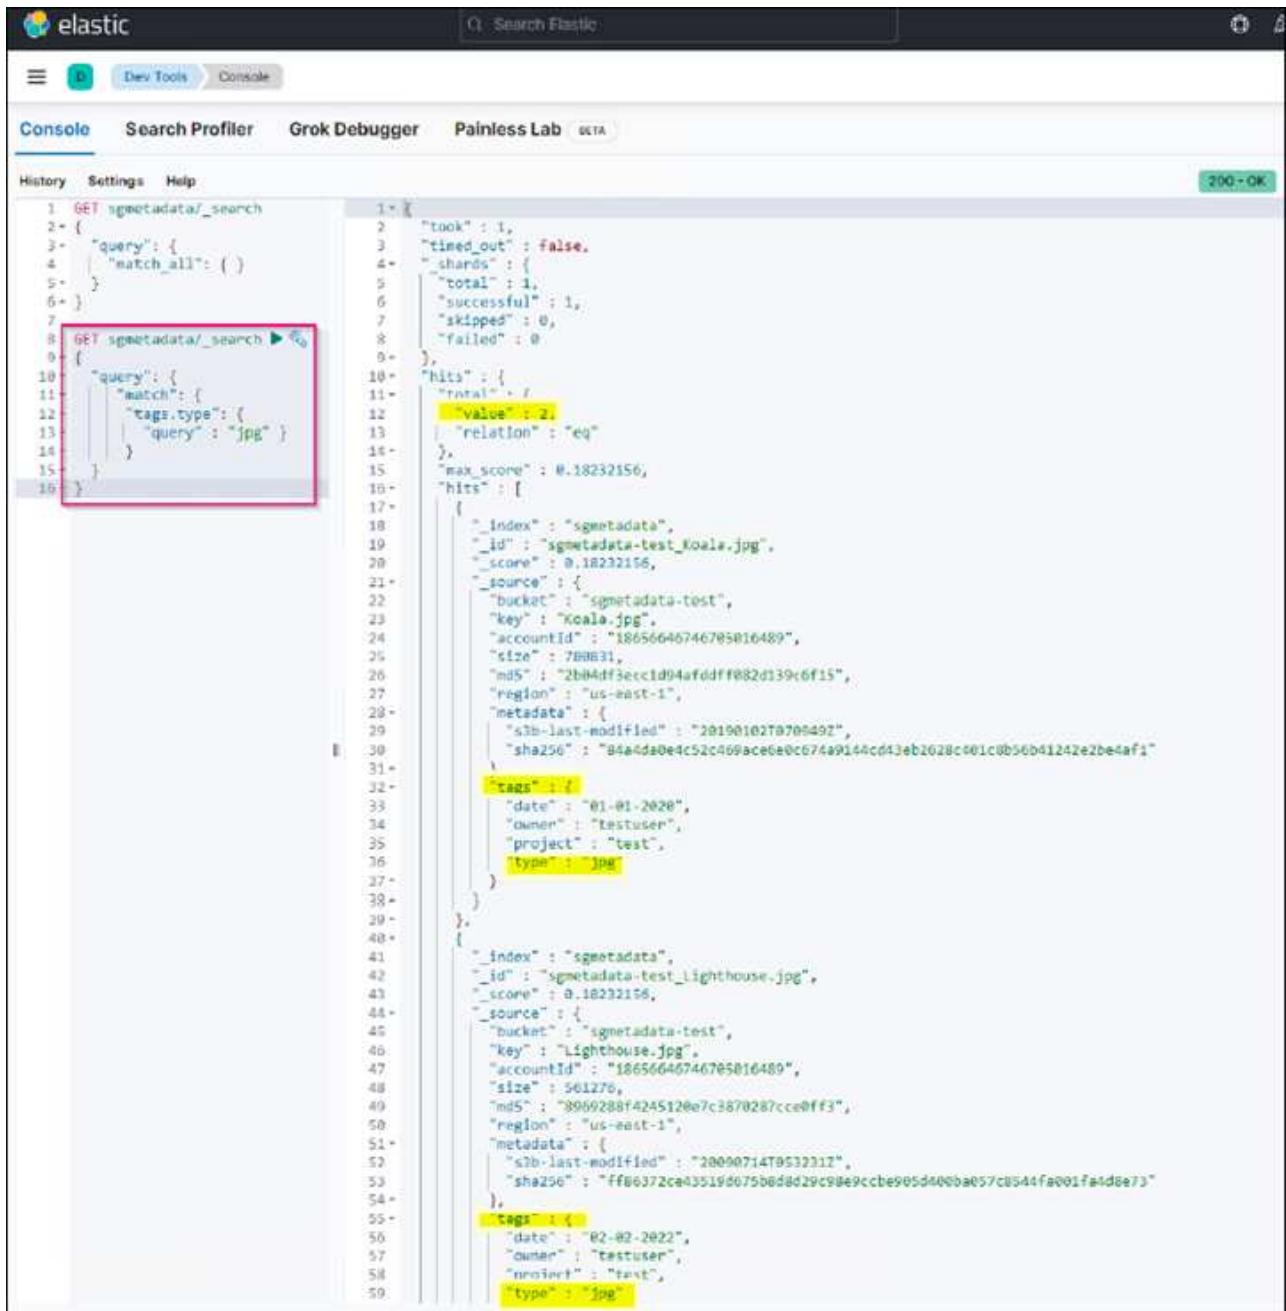
```
1 {
2   "took": 1,
3   "timed_out": false,
4   "_shards": {
5     "total": 1,
6     "successful": 1,
7     "skipped": 0,
8     "failed": 0
9   },
10  "hits": {
11    "total": {
12      "value": 4,
13      "relation": "eq"
14    },
15    "max_score": 1.0,
16    "hits": [
17      {
18        "_index": "sgmetadata",
19        "_id": "sgmetadata-test_test1.txt",
20        "_score": 1.0,
21        "source": {
22          "bucket": "sgmetadata-test",
23          "key": "test1.txt",
24          "accountId": "18656646746705816489",
25          "size": 45,
26          "md5": "36b194a8ac536f09a7061f024b97211e",
27          "region": "us-east-1",
28          "metadata": {
29            "s3b-last-modified": "2017042910102497",
30            "sha256": "6bf96e898615852c94f87015800d9a0299487f4cbe4429a1e1d7d7f427ab10f51"
31          },
32          "tags": {
33            "owner": "testuser",
34            "project": "test"
35          }
36        },
37      },
38      {
39        "_index": "sgmetadata",
40        "_id": "sgmetadata-test_Koala.jpg",
41        "_score": 1.0,
42        "source": {
43          "bucket": "sgmetadata-test",
44          "key": "Koala.jpg",
45          "accountId": "18656646746705816489",
46          "size": 788831,
47          "md5": "2b04df3ecc1d94afddff082d139c6f15",
48          "region": "us-east-1",
49          "metadata": {
50            "s3b-last-modified": "20190102100949Z",
51            "sha256": "84adda0e4c52c469acefe0c674a9144cd43eb2628c401c6b56b41242e2be4af1"
52          },
53          "tags": {
54            "date": "01-01-2020",
55            "owner": "testuser",
56            "project": "test",
57            "type": "jpg"
58          }
59        }
60      }
61    ]
62  }
63 }
```

The query 2 sample result in the following screenshot shows two records with tag type jpg.

```

GET sgmetadata/_search
{
  "query": {
    "match": {
      "tags.type": {
        "query" : "jpg" }
      }
    }
}

```



The screenshot shows the Elasticsearch Dev Tools Console interface. The top navigation bar includes 'Dev Tools' and 'Console'. The 'Console' tab is selected. The left panel shows the search query:

```

1 GET sgmetadata/_search
2 {
3   "query": {
4     "match_all": {}
5   }
6 }
7
8 GET sgmetadata/_search▶
9 {
10   "query": {
11     "match": {
12       "tags.type": {
13         "query" : "jpg" }
14       }
15     }
16 }

```

The right panel displays the search results. The status bar at the top right shows '200 OK'. The results are as follows:

```

1+ {
2   "took" : 1,
3   "timed_out" : false,
4   "shards" : {
5     "total" : 1,
6     "successful" : 1,
7     "skipped" : 0,
8     "failed" : 0
9   },
10  "hits" : {
11    "total" : 2,
12    "value" : 2,
13    "relation" : "eq"
14  },
15  "max_score" : 0.18232156,
16  "hits" : [
17    {
18      "_index" : "sgmetadata",
19      "_id" : "sgmetadata-test_koala.jpg",
20      "_score" : 0.18232156,
21      "_source" : {
22        "bucket" : "sgmetadata-test",
23        "key" : "Koala.jpg",
24        "accountId" : "18656646746705016489",
25        "size" : 288631,
26        "md5" : "2b04df3ecc1d94afddff882d139c6f15",
27        "region" : "us-east-1",
28        "metadata" : {
29          "s3b-last-modified" : "20190102T070549Z",
30          "sha256" : "84a4da0e4c52c469ace6e0c674a9144cd43eb1628c001c0b56b41242e2be4af1"
31        },
32        "tags" : [
33          "date" : "01-01-2020",
34          "owner" : "testuser",
35          "project" : "test",
36          "type" : "jpg"
37        ]
38      },
39      {
40        "_index" : "sgmetadata",
41        "_id" : "sgmetadata-test_lighthouse.jpg",
42        "_score" : 0.18232156,
43        "_source" : {
44          "bucket" : "sgmetadata-test",
45          "key" : "Lighthouse.jpg",
46          "accountId" : "18656646746705016489",
47          "size" : 561276,
48          "md5" : "8969288f4245120e07c3870287cce0ff3",
49          "region" : "us-east-1",
50          "metadata" : {
51            "s3b-last-modified" : "20090714T051231Z",
52            "sha256" : "ff86372ca435196075b6d8d29c98e9ccbe905da00ba057c8544fa001fa0d8e73"
53          },
54          "tags" : [
55            "date" : "02-02-2022",
56            "owner" : "testuser",
57            "project" : "test",
58            "type" : "jpg"
59          ]
60        }
61      }
62    }
63  }
64 }

```

Where to find additional information

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

- [What are platform services](#)
- [StorageGRID 11.6 Documentation](#)

By *Angela Cheng*

Node Clone

Node clone considerations and performance.

Node clone considerations

Node clone can be a faster method for replacing existing appliance nodes for a tech refresh, increase capacity, or increase performance of your StorageGRID system. Node clone can also be useful for converting to node encryption with a KMS, or changing a storage node from DDP8 to DDP16.

- The used capacity of the source node is not relevant to the time required for the clone process to complete. Node clone is a full copy of the node including free space in the node.
- The source and destination appliances must be at the same PGE version
- The destination node must always have larger capacity than the source
 - Make sure the new destination appliance has a larger drive size than the source
 - If the destination appliance has the same size drives and is configured for DDP8, you can configure the destination for DDP16. If the source is already configured for DDP16 then node clone will not be possible.
 - When going from SG5660 or SG5760 appliances to SG6060 appliances be aware that the SG5x60's have 60 capacity drives where the SG6060 only has 58.
- The node clone process requires the source node to be offline to the grid for the duration of the cloning process. If an additional node goes offline during this time client services may be impacted.
- 11.8 and below: A storage node can only be offline for 15 days. If the cloning process estimate is close to 15 days or will exceed 15 days, use the expansion and decommission procedures.
 - 11.9: The 15 day limit has been removed.
- For a SG6060 or SG6160 with expansion shelves, you need to add the time for the correct shelf drive size to the time of the base appliance time to get the full clone duration.
- The number of volumes in a target storage appliance must be greater than or equal to the number of volumes in the source node. You cannot clone a source node with 16 object store volumes (rangedb) to a target storage appliance with 12 object store volumes even if the target appliance has larger capacity than the source node. Most storage appliances have 16 object store volumes, except the SGF6112 storage appliance that has only 12 object store volumes. For example, you cannot clone from a SG5760 to a SGF6112.

Node clone Performance estimates

The following tables contain calculated estimates for node clone duration. Conditions vary so, entries in **BOLD** may risk exceeding the 15 day limit for a node down.

DDP8

SG5612/SG5712/SG5812 → Any

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	1 Day	2 Days	2.5 Days	3 Days	4 Days	4.5 Days	5.5 Days
25GB	1 Day	2 Days	2.5 Days	3 Days	4 Days	4.5 Days	5.5 Days

SG5660 → SG5760/SG5860

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	3.5 Day	7 Days	8.5 Days	10.5 Days	13.5 Days	15.5 Days	18.5 Days
25GB	3.5 Day	7 Days	8.5 Days	10.5 Days	13.5 Days	15.5 Days	18.5 Days

SG5660 → SG6060/SG6160

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	2.5 Day	4.5 Days	5.5 Days	6.5 Days	9 Days	10 Days	12 Days
25GB	2 Day	4 Days	5 Days	6 Days	8 Days	9 Days	10 Days

SG5760/SG5860 → SG5760/SG5860

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	3.5 Day	7 Days	8.5 Days	10.5 Days	13.5 Days	15.5 Days	18.5 Days
25GB	3.5 Day	7 Days	8.5 Days	10.5 Days	13.5 Days	15.5 Days	18.5 Days

SG5760/SG5860 → SG6060/SG6160

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	2.5 Day	4.5 Days	5.5 Days	6.5 Days	9 Days	10 Days	12 Days
25GB	2 Day	3.5 Days	4.5 Days	5.5 Days	7 Days	8 Days	9.5 Days

SG6060/SG6160 → SG6060/SG6160

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	2.5 Day	4.5 Days	5.5 Days	6.5 Days	8.5 Days	9.5 Days	11.5 Days
25GB	2 Day	3 Days	4 Days	4.5 Days	6 Days	7 Days	8.5 Days

DDP16

SG5760/SG5860 → SG5760/SG5860

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	3.5 Day	6.5 Days	8 Days	9.5 Days	12.5 Days	14 Days	17 Days
25GB	3.5 Day	6.5 Days	8 Days	9.5 Days	12.5 Days	14 Days	17 Days

SG5760/SG5860 → SG6060/SG6160

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	2.5 Day	5 Days	6 Days	7.5 Days	10 Days	11 Days	13 Days
25GB	2 Day	3.5 Days	4 Days	5 Days	6.5 Days	7 Days	8.5 Days

SG6060/SG6160 → SG6060/SG6160

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	3 Day	5 Days	6 Days	7 Days	9.5 Days	10.5 Days	13 Days
25GB	2 Day	3.5 Days	4.5 Days	5 Days	7 Days	7.5 Days	9 Days

Expansion shelf (add to above SG6060/SG6160 for each shelf on source appliance)

Network Interface speed	4TB Drive size	8TB Drive size	10TB Drive size	12TB Drive size	16TB Drive size	18TB Drive size	22TB Drive size
10GB	3.5 Day	5 Days	6 Days	7 Days	9.5 Days	10.5 Days	12 Days
25GB	2 Day	3 Days	4 Days	4.5 Days	6 Days	7 Days	8.5 Days

By Aron Klein

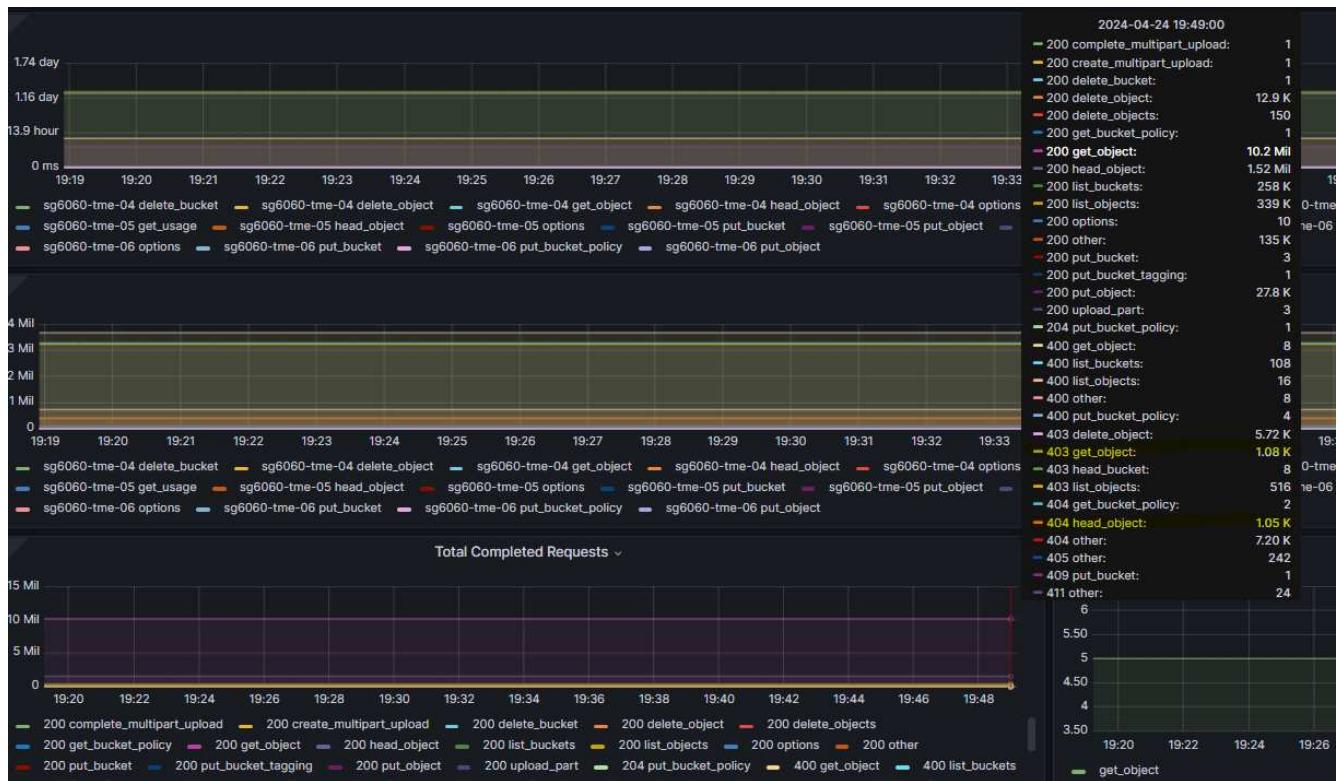
Grid site relocation and site-wide network change procedure

This guide describes the preparation and procedure for StorageGRID site relocation in a multi-sites Grid. You should have a complete understand of this procedure and plan ahead to ensure smooth process and minimize interruption to clients.

If you need to change the Grid network of entire Grid, see [Change IP addresses for all nodes in grid](#).

Considerations before site relocation

- Site move should be completed, and all nodes online within 15 days to avoid Cassandra database rebuild. [Recover Storage Node down more than 15 days](#)
- If any ILM rule in active policy is using strict ingest behavior, consider changing it to balance or dual commit if customer wants to continue to PUT objects into the Grid during site relocation.
- For storage appliances with 60 drives or more, never move the shelf with disk drives installed. Label each disk drives and remove them from storage enclosure before pack/move.
- Change StorageGRID appliance Grid network VLAN can be performed remotely over admin network or client network. Or else plan to be onsite to perform the change before or after the relocation.
- Check if customer application is using HEAD or GET nonexistence object before PUT. If yes, change the bucket consistency to strong-site to avoid HTTP 500 error. If you are not sure, check S3 overview Grafana charts **Grid manager > Support > Metrics**, mouse over the 'Total Completed Request' chart. If there is very high count of 404 get Object or 404 head object, likely one or more applications are using head or get nonexistence object. The count is accumulative, mouse over different timeline to see the difference.



Procedure to change Grid IP address before site relocation

Steps

1. If new Grid network subnet will be used at the new location, [add the subnet to Grid network subnet list](#)
2. Log in to the primary Admin Node, use change-ip to make Grid IP change, must **stage** the change before shutdown the node for relocation.
 - a. Select 2 then 1 for Grid IP change

```
Editing: Node IP/subnet and gateway
```

```
Use up arrow to recall a previously typed value, which you can then edit
Use d or 0.0.0.0/0 as the IP/mask to delete the network from the node
Use q to complete the editing session early and return to the previous menu
Press <enter> to use the value shown in square brackets
```

```
=====
Site: LONDON
=====
```

```
LONDON-ADM1 Grid IP/mask [ 10.45.74.14/26 ]: 10.45.74.24/26
LONDON-S1 Grid IP/mask [ 10.45.74.16/26 ]: 10.45.74.26/26
LONDON-S2 Grid IP/mask [ 10.45.74.17/26 ]: 10.45.74.27/26
LONDON-S3 Grid IP/mask [ 10.45.74.18/26 ]: 10.45.74.28/26
=====
```

```
LONDON-ADM1 Grid Gateway [ 10.45.74.1 ]:
LONDON-S1 Grid Gateway [ 10.45.74.1 ]:
LONDON-S2 Grid Gateway [ 10.45.74.1 ]:
LONDON-S3 Grid Gateway [ 10.45.74.1 ]:
=====
```

```
=====
Site: OXFORD
=====
```

```
OXFORD-ADM1 Grid IP/mask [ 10.45.75.14/26 ]:
OXFORD-S1 Grid IP/mask [ 10.45.75.16/26 ]:
OXFORD-S2 Grid IP/mask [ 10.45.75.17/26 ]:
OXFORD-S3 Grid IP/mask [ 10.45.75.18/26 ]:
=====
```

```
OXFORD-ADM1 Grid Gateway [ 10.45.75.1 ]:
OXFORD-S1 Grid Gateway [ 10.45.75.1 ]:
OXFORD-S2 Grid Gateway [ 10.45.75.1 ]:
OXFORD-S3 Grid Gateway [ 10.45.75.1 ]:
=====
```

```
Finished editing. Press Enter to return to menu. █
```

- b. select 5 to show changes

```
=====
Site: LONDON
=====
LONDON-ADM1 Grid IP      [ 10.45.74.14/26 ]: 10.45.74.24/26
LONDON-S1  Grid IP      [ 10.45.74.16/26 ]: 10.45.74.26/26
LONDON-S2  Grid IP      [ 10.45.74.17/26 ]: 10.45.74.27/26
LONDON-S3  Grid IP      [ 10.45.74.18/26 ]: 10.45.74.28/26
Press Enter to continue
```

c. select 10 to validate and apply the change.

```
Welcome to the StorageGRID IP Change Tool.

Selected nodes: all

1:  SELECT NODES to edit
2:  EDIT IP/mask and gateway
3:  EDIT admin network subnet lists
4:  EDIT grid network subnet list
5:  SHOW changes
6:  SHOW full configuration, with changes highlighted
7:  VALIDATE changes
8:  SAVE changes, so you can resume later
9:  CLEAR all changes, to start fresh
10: APPLY changes to the grid
0:  Exit

Selection: 10
```

d. Must select **stage** in this step.

```
Validating new networking configuration... PASSED.
Checking for Grid Network IP address swaps... PASSED.

Applying these changes will update the following nodes:

LONDON-ADM1
LONDON-S1
LONDON-S2
LONDON-S3

The following nodes will also require restarting:

LONDON-ADM1
LONDON-S1
LONDON-S2
LONDON-S3

Select one of the following options:

apply: apply all changes and automatically restart nodes (if necessary)
stage: stage the changes; no changes will take effect until the nodes are restarted
cancel: do not make any network changes at this time

[apply/stage/cancel]> stage
```

e. If primary admin node is included in above change, Enter 'a' to restart primary admin node manually

```
10.45.74.14 - PuTTY
Validating new networking configuration... PASSED.
Checking for Grid Network IP address swaps... PASSED.

Applying these changes will update the following nodes:
LONDON-ADM1
LONDON-S1
LONDON-S2
LONDON-S3

The following nodes will also require restarting:
LONDON-ADM1
LONDON-S1
LONDON-S2
LONDON-S3

Select one of the following options:
apply: apply all changes and automatically restart nodes (if necessary)
stage: stage the changes; no changes will take effect until the nodes are restarted
cancel: do not make any network changes at this time

[apply/stage/cancel]> stage

Generating new grid networking description file... PASSED.
Running provisioning... PASSED.
Updating network configuration on LONDON-S1... PASSED.
Updating network configuration on LONDON-S2... PASSED.
Updating network configuration on LONDON-S3... PASSED.
Updating network configuration on LONDON-ADM1... PASSED.

Finished staging network changes. You must manually restart these nodes for the changes to take effect:
LONDON-ADM1 (has IP 10.45.74.14 until restart)
LONDON-S1 (has IP 10.45.74.16 until restart)
LONDON-S2 (has IP 10.45.74.17 until restart)
LONDON-S3 (has IP 10.45.74.18 until restart)

Importing bundles... PASSED.
*****
*          IMPORTANT          *
*          *          *
* A new recovery package has been generated as a result of the *          *
* configuration change. Select Maintenance > Recovery Package *          *
* in the Grid Manager to download it.          *
*****
Network Update Complete. Primary admin restart required. Select 'continue' to restart this node immediately, 'abort' to restart manually.
Enter a to abort, c to continue [a/c]> a
```

f. Press enter to return to previous menu and exit from change-ip interface.

```
Network Update Complete. Primary admin restart required. Select 'continue' to restart this node immediately, 'abort' to restart manually.
Enter a to abort, c to continue [a/c]> a
Restart aborted. You must manually restart this node as soon as possible
Press Enter to return to the previous menu.
```

3. From Grid Manager, download the new recovery package. **Grid manager > Maintenance > Recovery package**
4. If VLAN change is required on StorageGRID appliance, see the section [Appliance VLAN change](#).
5. Shutdown all nodes and/or appliances at the site, label/remove disk drives if necessary, unrack, pack and move.
6. If you plan to change admin network ip and/or client VLAN and ip address, you can perform the change after the relocation.

Appliance VLAN change

The procedure below assume you have remote access to StorageGRID appliance's admin or client network to perform the change remotely.

Steps

1. Before shutdown the appliance,
[place the appliance in maintenance mode](#).

2. Using a browser to access the StorageGRID appliance installer GUI using <https://<admin-or-client-network-ip>:8443>. Cannot use Grid IP as the new Grid IP already in place once the appliance is boot into maintenance mode.
3. Change the VLAN for Grid network. If you are accessing the appliance over client network, you cannot change Client VLAN at this time, you can change it after the move.
4. ssh to the appliance and shutdown the node using 'shutdown -h now'
5. After the appliances are ready at new site, access to the StorageGRID appliance installer GUI using <https://<grid-network-ip>:8443>. Confirm the storage are in optimal state and network connectivity to other Grid nodes using ping/nmap tools in the GUI.
6. If plan to change client network IP, you can change the client VLAN at this stage. The client network is not ready until you update the client network ip using change-ip tool in later step.
7. Exit maintenance mode. From the StorageGRID Appliance Installer, select **Advanced > Reboot Controller**, and then select **Reboot into StorageGRID**.
8. After all nodes are up and Grid shows no connectivity issue, use change-ip to update the appliance admin network and client network if necessary.

Migrating object-based storage from ONTAP S3 to StorageGRID

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Migration Demo

This is a demonstration on migrating users and buckets from ONTAP S3 to StorageGRID.

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

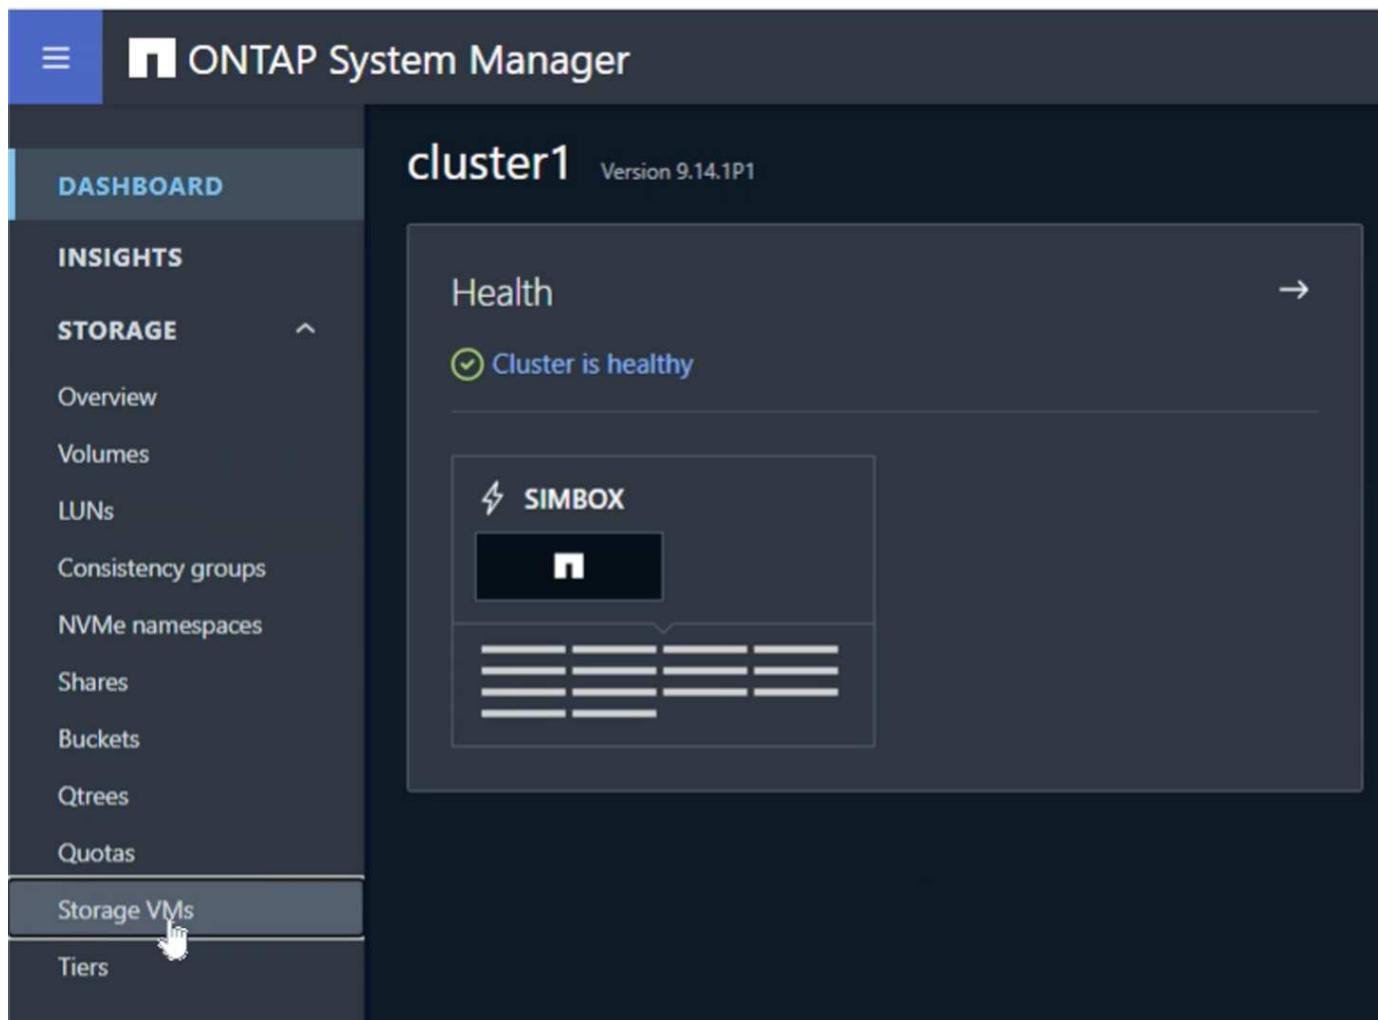
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Preparing ONTAP

For demonstration purposes we will create an SVM object store server, user, group, group policy and buckets.

Create the Storage Virtual Machine

In ONTAP System Manager, navigate to Storage VM's and add a new storage VM.



The screenshot shows the ONTAP System Manager interface. The left sidebar has a 'DASHBOARD' tab, followed by a 'STORAGE' section with 'INSIGHTS' and 'STORAGE' sub-options: Overview, Volumes, LUNs, Consistency groups, NVMe namespaces, Shares, Buckets, Qtrees, Quotas, Storage VMs (which is selected and highlighted in blue), and Tiers. The main content area is titled 'cluster1 Version 9.14.1P1'. It displays a 'Health' section with a green checkmark and the text 'Cluster is healthy'. Below this is a 'SIMBOX' icon with a lightning bolt, representing a storage virtual machine. The SIMBOX icon contains a small ONTAP logo and a grid of horizontal lines representing storage volumes.

Select the "Enable S3" and "Enable TLS" check boxes and configure the HTTP(S) ports. Define the IP, subnet mask and define the gateway and broadcast domain if not using the default or required in your environment.

Add storage VM

STORAGE VM NAME

Access protocol

SMB/CIFS, NFS, S3 iSCSI FC NVMe

Enable SMB/CIFS
 Enable NFS
 Enable S3

S3 SERVER NAME

Enable TLS

PORT

CERTIFICATE

Use system-generated certificate i
 Use external-CA signed certificate

Use HTTP (non-secure)

PORT

DEFAULT LANGUAGE i

NETWORK INTERFACE

Use multiple network interfaces when client traffic is high.

onPrem-01

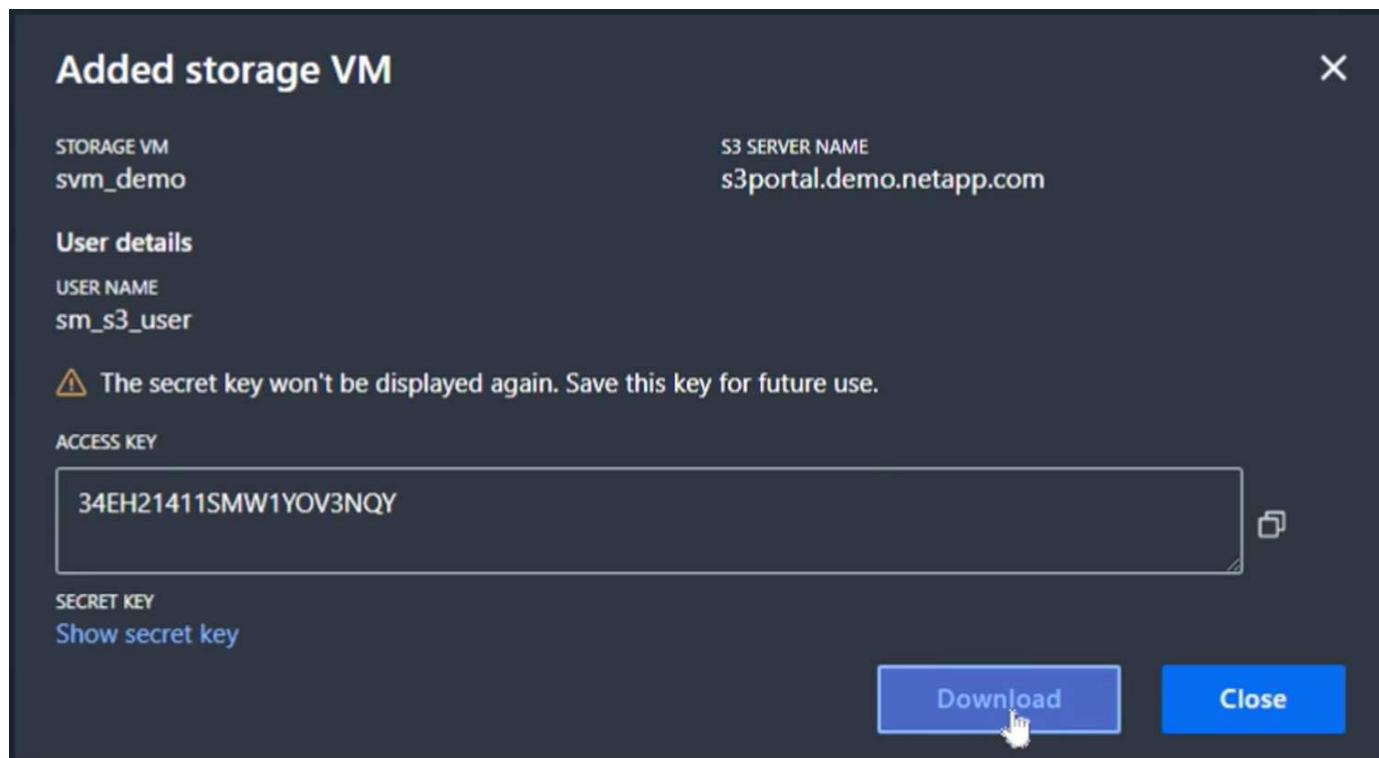
IP ADDRESS	SUBNET MASK	GATEWAY	BROADCAST DOMAIN AND PORT
192.168.0.200	24	Add optional gateway	<input type="text" value="Default"/>

Storage VM administration

Enable maximum capacity limit
The maximum capacity that all volumes in this storage VM can allocate. [Learn More](#)

Manage administrator account

As part of the SVM creation a user will be created. Download the S3 keys for this user and close the window.



Once the SVM has been created, edit the SVM and add the DNS settings.

Services

NIS ✗ ⚙

Not configured

Name service switch ✗ ✎

Services lookup order ⓘ

HOSTS
Files, then DNS

GROUP
Files

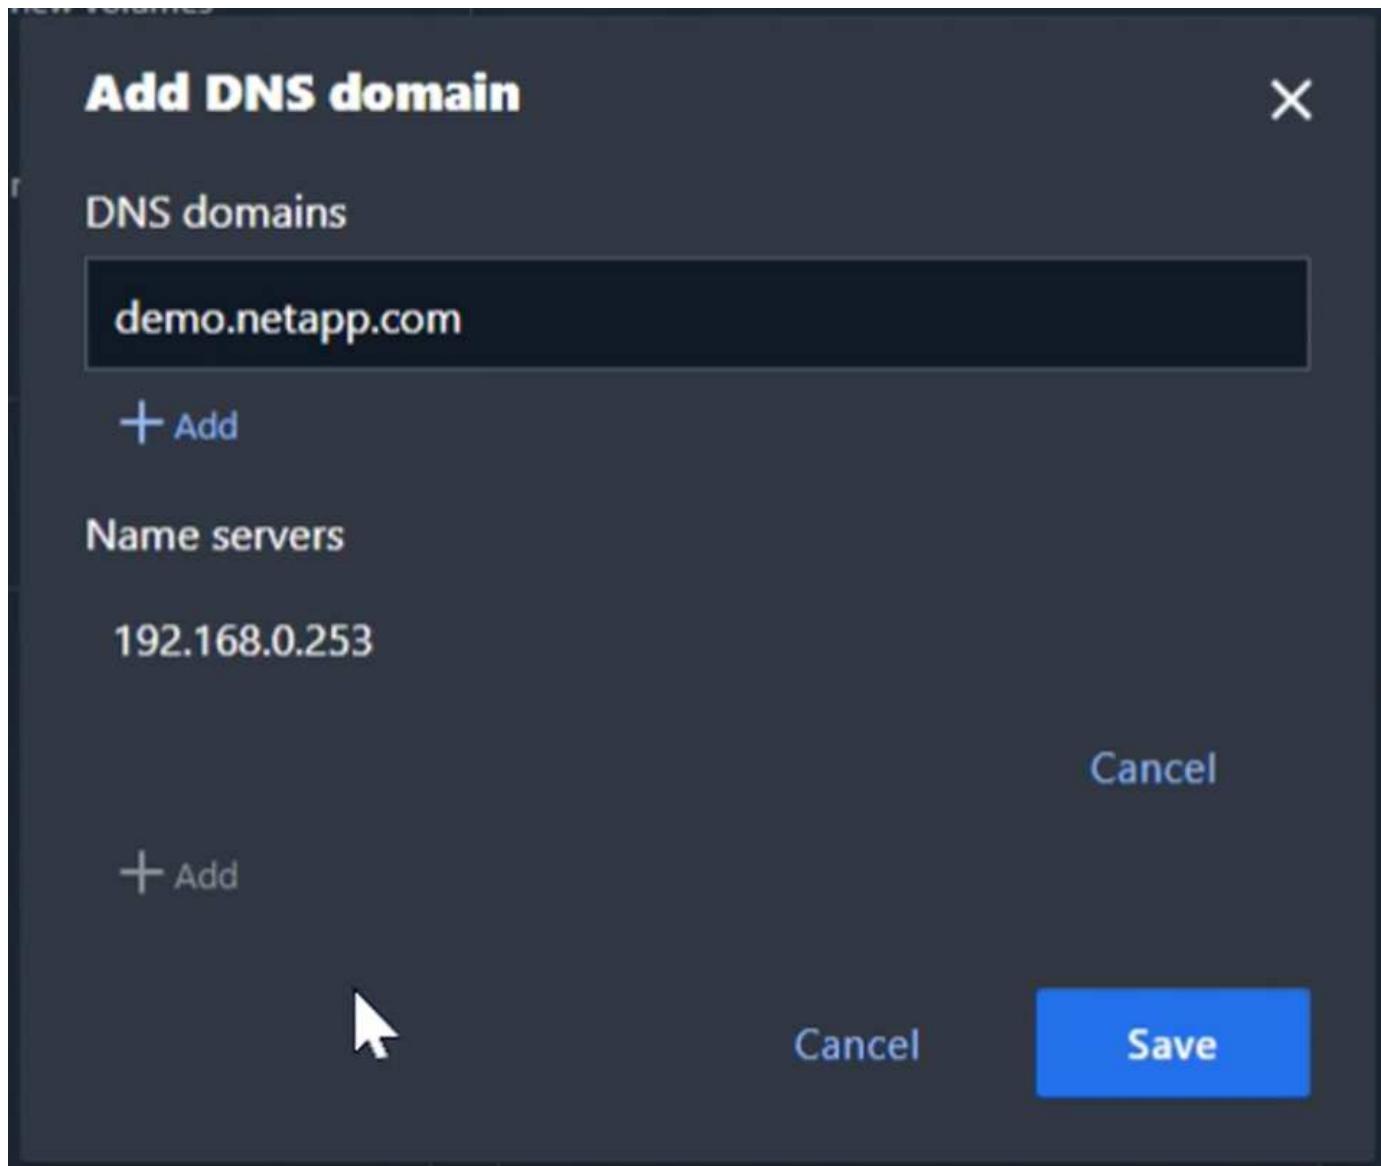
NAME MAP
Files

NETGROUP
Files

DNS ✗ Configure

Not configured

Define the DNS name and IP.



Create SVM S3 User

Now we can configure the S3 users and group. Edit the S3 settings.

Protocols

NFS Not configured	 	SMB/CIFS Not configured	 	iS N
NVMe Not configured	 	S3 STATUS  Enabled TLS Disabled HTTP Enabled	 	

Add a new user.

Storage VMs

+ Add : More

Name
svm_demo

S3 All settings

Enabled

Server [Edit](#)

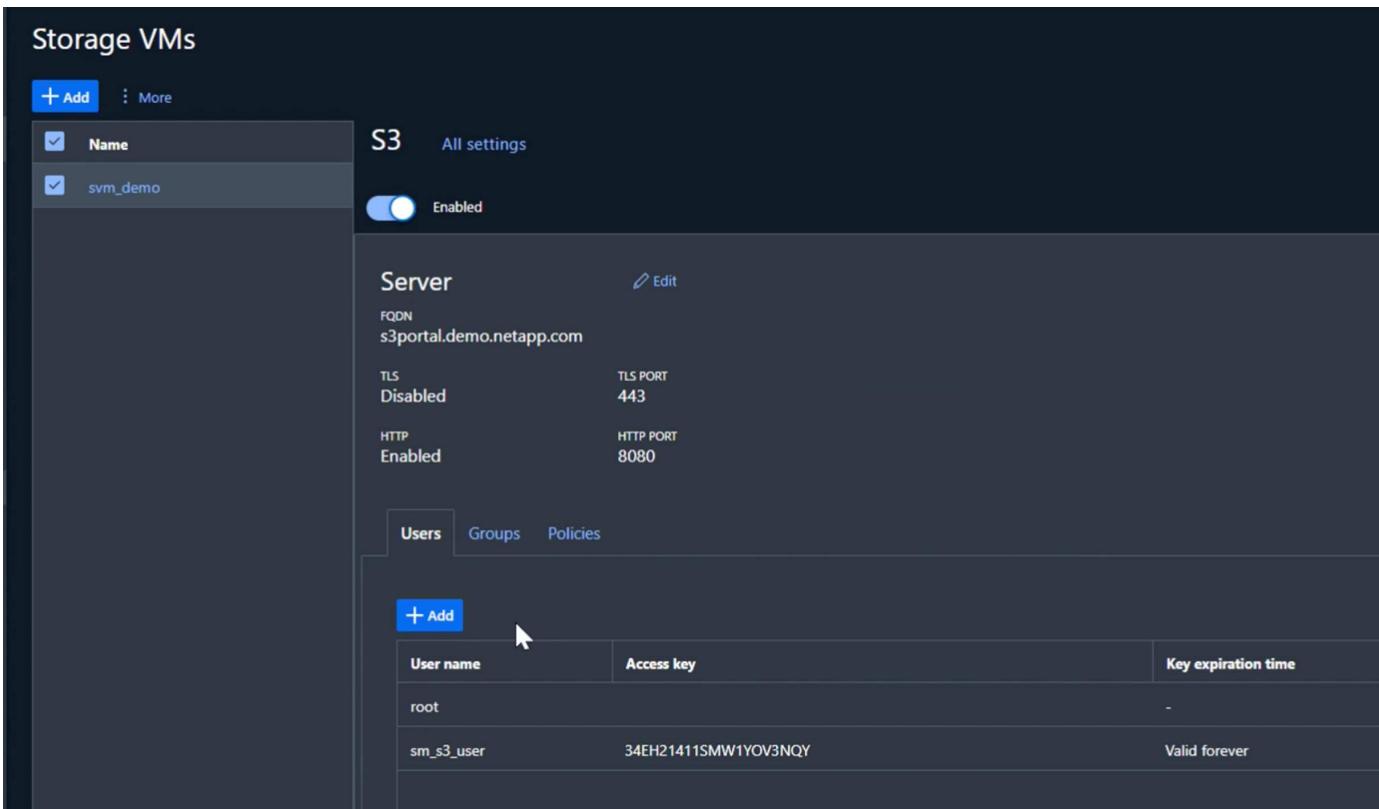
FQDN: s3portal.demo.netapp.com

TLS: Disabled TLS PORT: 443

HTTP: Enabled HTTP PORT: 8080

Users [Groups](#) [Policies](#)

User name	Access key	Key expiration time
root		-
sm_s3_user	34EH21411SMW1YOV3NQY	Valid forever



Input the user name and key expiration.

Storage VMs

+ Add : More

Name
svm_demo

S3 All settings

Enabled

Server [Edit](#)

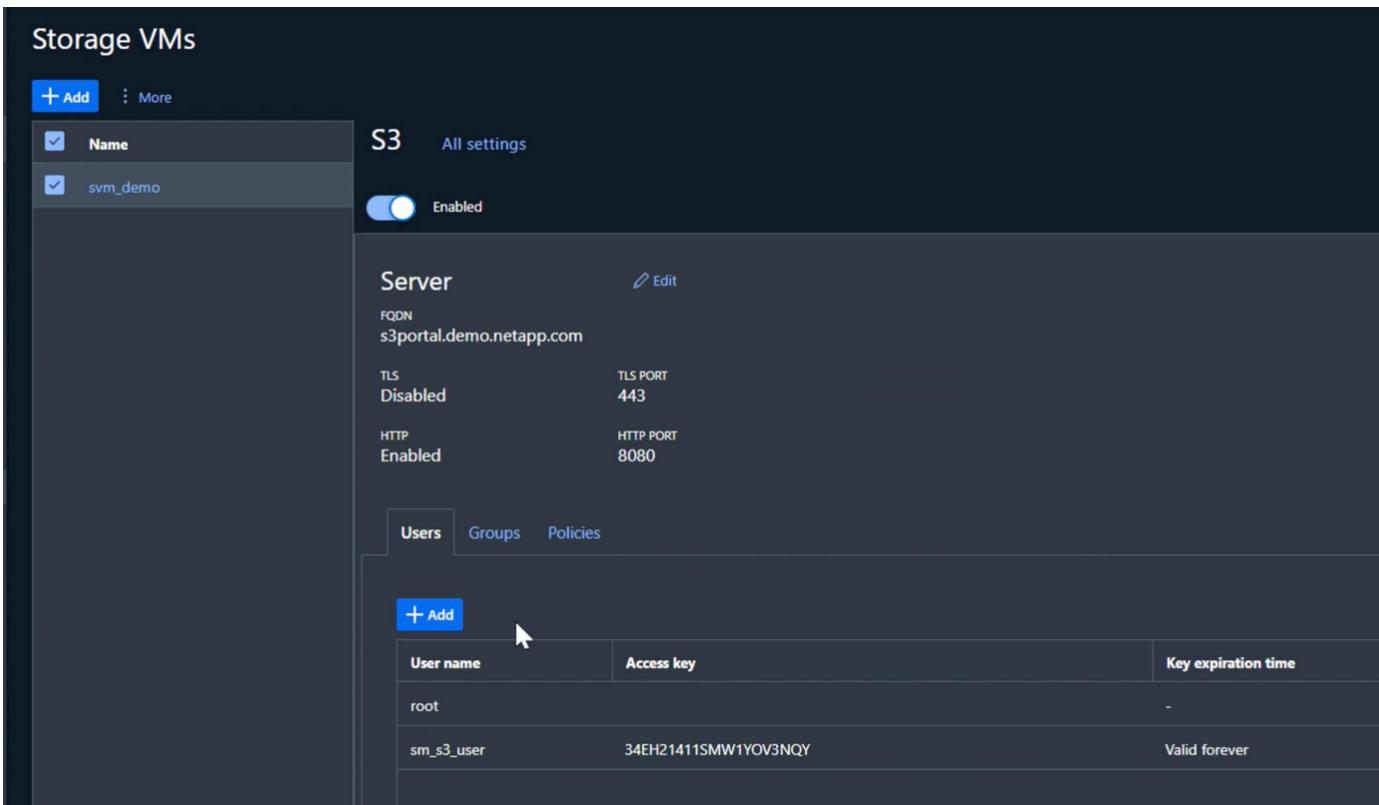
FQDN: s3portal.demo.netapp.com

TLS: Disabled TLS PORT: 443

HTTP: Enabled HTTP PORT: 8080

Users [Groups](#) [Policies](#)

User name	Access key	Key expiration time
root		-
sm_s3_user	34EH21411SMW1YOV3NQY	Valid forever



Download the S3 keys for the new user.

Added user

**USER NAME**

demo_s3_user

ACCESS KEY

3TVPI142JGE3Y7FV2KC0

**SECRET KEY**

75a1QqKBU4quA132twl4g4lC4Gg5PP30ncy0sPE8

[Hide secret key](#)**KEY EXPIRATION TIME**

Valid forever

 The secret key won't be displayed again. Save this key for future use.

[Download](#)[Close](#)

USER NAME

3TVPI142JGE3Y7FV2KC0

VALID FOREVER

Create SVM S3 group

On the Groups tab of the SVM S3 settings, add a new group with the user created above and FullAccess permissions.

Add group

X

NAME

demo_s3_group

USERS

demo_s3_user X

POLICIES

FullAccess X

Cancel

Save



Create SVM S3 buckets

Navigate to the Buckets section and click the "+Add" button.



DASHBOARD

INSIGHTS

STORAGE ^

Overview

Volumes

LUNs

Consistency groups

NVMe namespaces

Shares

Buckets

Qtrees

Quotas

Storage VMs

Tiers

Buckets



Name

Stora

Enter a name, capacity, and deselect the "Enable ListBucket access..." check box. and click on the "More options" button.

Add bucket

X

NAME

bucket

CAPACITY

100



GiB



Enable ListBucket access for all users on the storage VM "svm_demo".
Enabling this will allow users to access the bucket.

More options

Cancel

Save

In the "More options" section select the enable versioning check box. and click the "Save" button.

Add bucket

X

NAME

FOLDER (OPTIONAL)

[Browse](#)

Specify the folder to map to this bucket. [Know more](#)

CAPACITY

Use for tiering

If you select this option, the system will try to select low-cost media with optimal performance for the tiered data.

Enable versioning

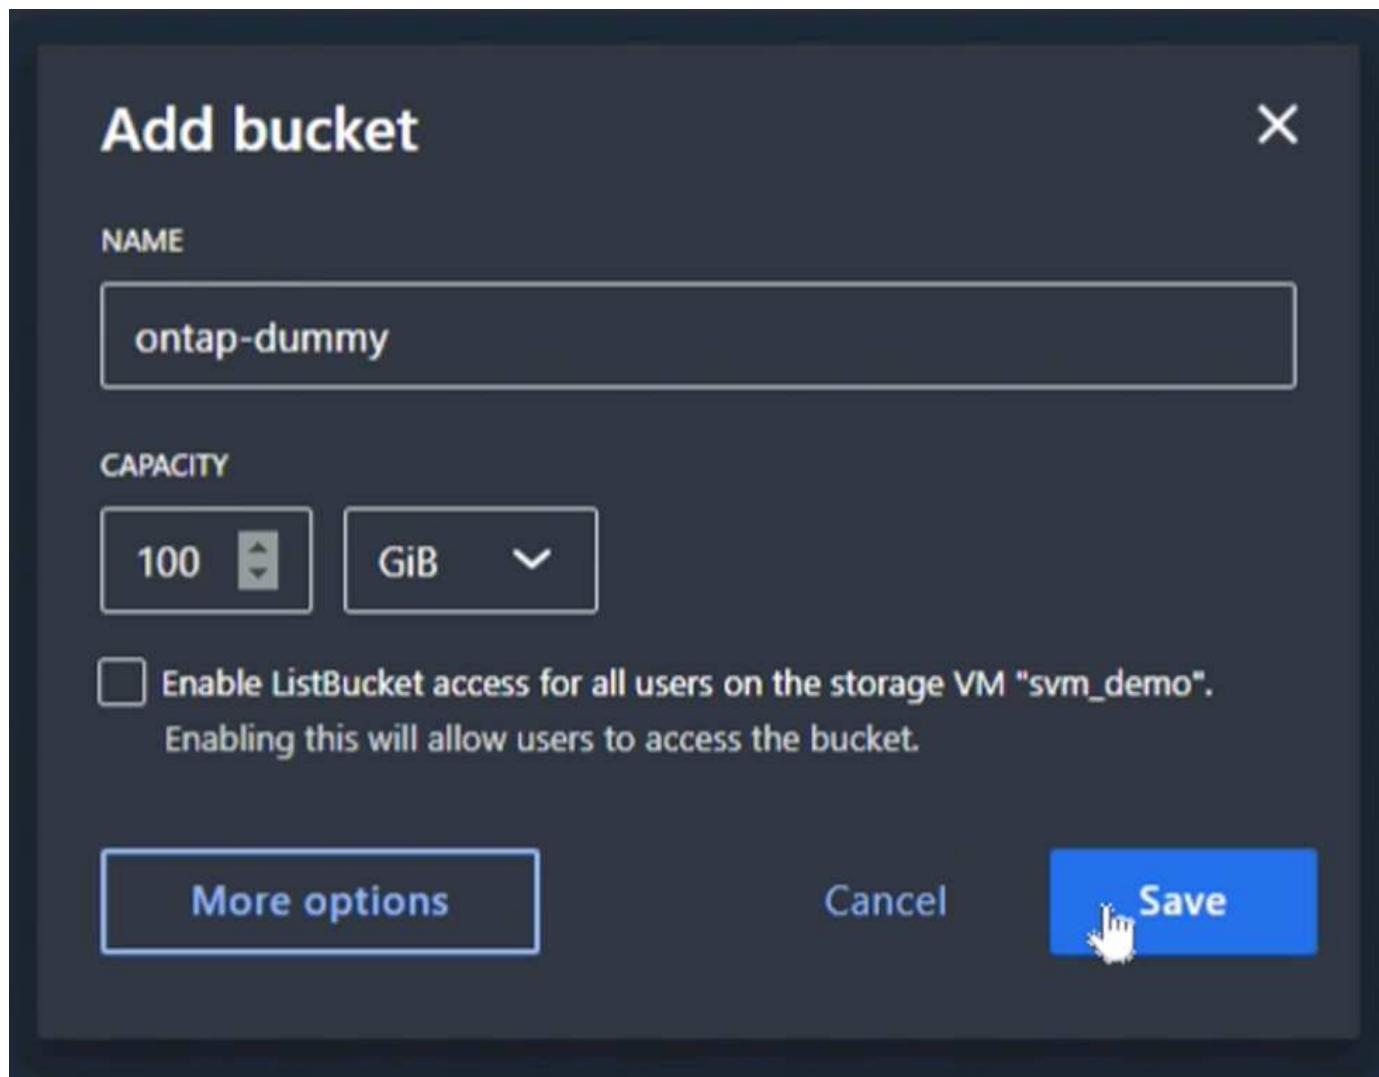
Versioning-enabled buckets allow you to recover objects that were accidentally deleted or overwritten. After versioning is enabled, it can't be disabled. However, you can suspend versioning.

PERFORMANCE SERVICE LEVEL

▼

Not sure? [Get help selecting type](#)

Repeat the process and create a second bucket without versioning enabled. Enter a name, the same capacity as bucket one, and deselect the "Enable ListBucket access..." check box. and click on the "Save" button.



By Rafael Guedes, and Aron Klein

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Preparing StorageGRID

Continuing the configuration for this demo we will create a Tenant, user, security group, group policy, and bucket.

Create the tenant

Navigate to the "Tenants" tab and click on the "create" button

NetApp | StorageGRID Grid Manager

Search by page title Search

DASHBOARD

ALERTS 0

NODES

TENANTS

ILM

CONFIGURATION

MAINTENANCE

SUPPORT

Tenants

View information for each tenant account. Depending on the timing of ingests, network connectivity, and node status, the usage data shown might be out of date. To view more recent values, select the tenant name.

date Export to CSV Actions Search tenants by name or ID Search

No results

	Name ? ▲	Logical space used ? ▲	Quota utilization ? ▲	Quota ? ▲	Object count ? ▲	Sign in/Copy URL ?
No tenants found.						
Create						

Fill in the details for the tenant providing a tenant name, select S3 for the client type, and no quota is required. No need to select platform services or allow S3 select. You can choose to use own Identity source if you choose. Set the root password and click on the finish button.

Click on the tenant name to view the tenant details. **You will need the tenant ID later so copy it off.** Click on the Sign in button. This will bring you to the tenant portal login. Save the URL for future use.

Tenants

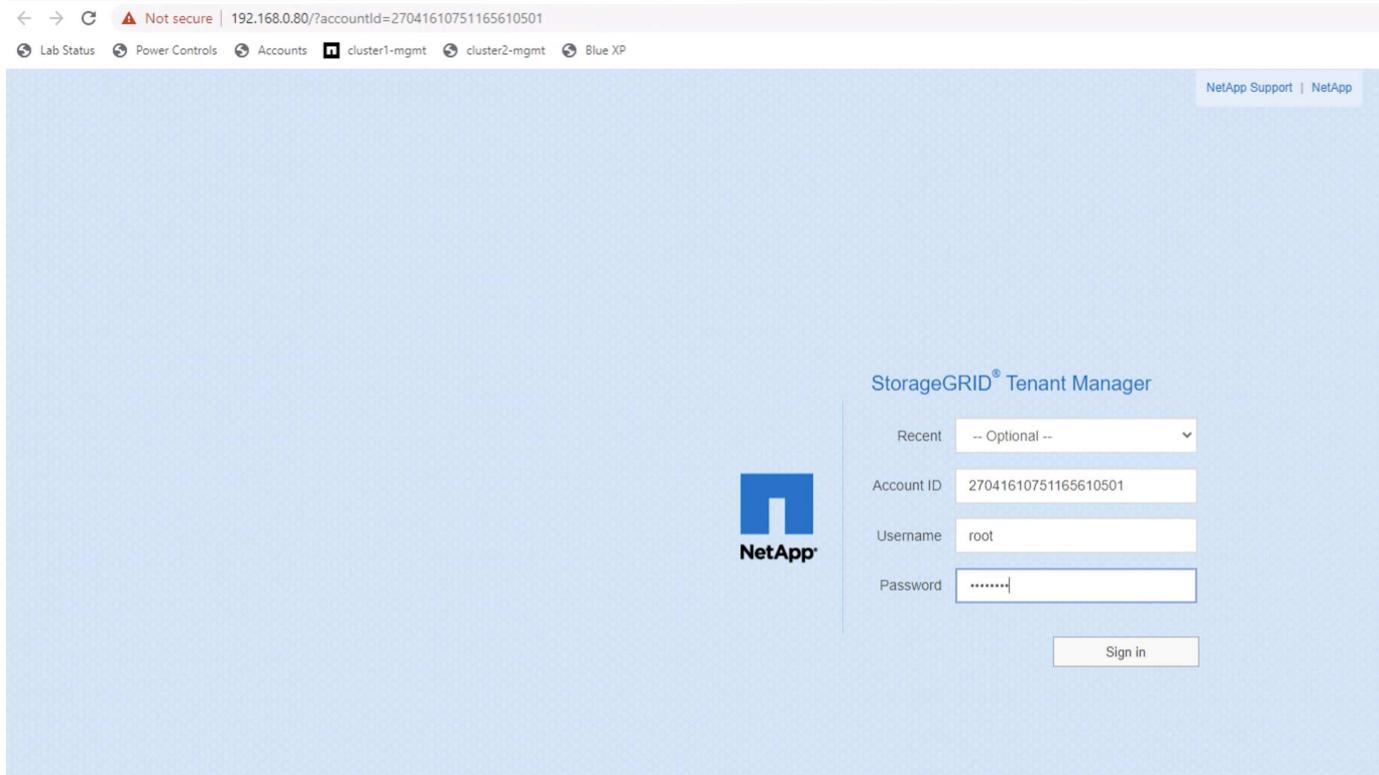
View information for each tenant account. Depending on the timing of ingests, network connectivity, and node status, the usage data shown might be out of date. To view more recent values, select the tenant name.

Create Export to CSV Actions Search tenants by name or ID Search Displaying one result

	Name ? ▲	Logical space used ? ▲	Quota utilization ? ▲	Quota ? ▲	Object count ? ▲	Sign in/Copy URL ?
<input type="checkbox"/>	tenant_demo	0 bytes	—	—	0	→ Copy

← Previous 1 Next →

This will bring you to the tenant portal login. Save the URL for future use, and enter the root user credentials.



Create the user

Navigate to the Users tab and create a new user.

Actions	Username	Full Name	Denied	Type
<input type="checkbox"/>	root	Root		Local

Optional

Enter user credentials

Create a new local user and configure user access.

Full name ?

Must contain at least 1 and no more than 128 characters

Username ?

Password

?

Must contain at least 8 and no more than 32 characters

Confirm password

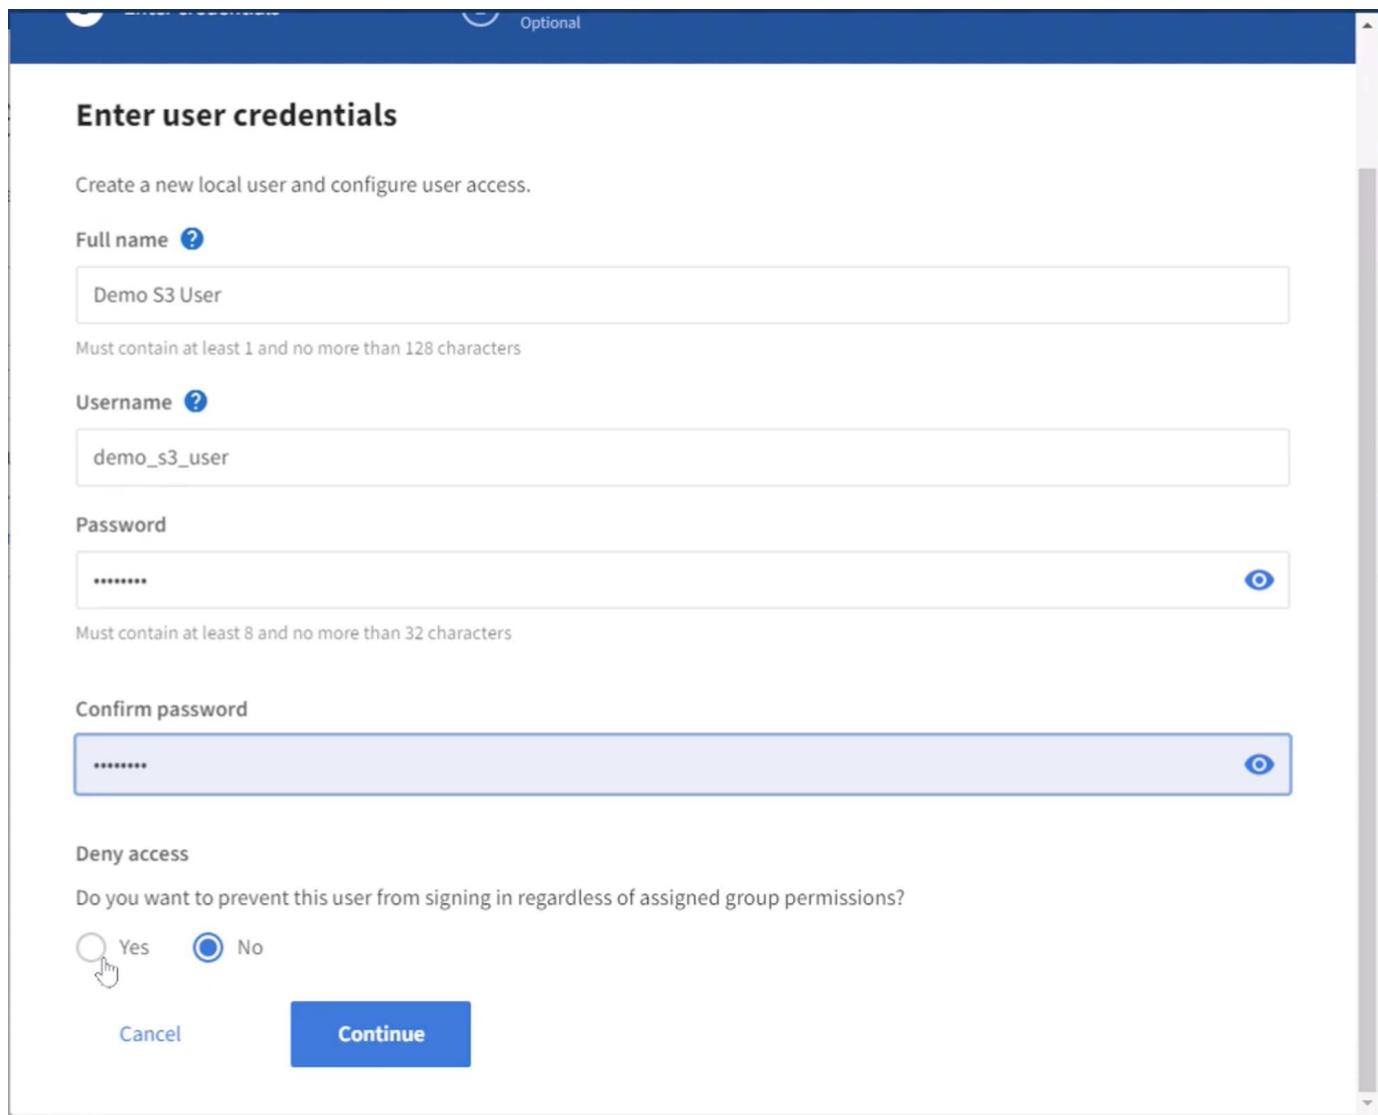
?

Deny access

Do you want to prevent this user from signing in regardless of assigned group permissions?

Yes No

[Cancel](#) [Continue](#)



Now that the new user has been created, click on the users name to open the details of the user.

Copy the user ID from the URL to be used later.

⚠ Not secure | <https://192.168.0.80/ui/#/users/ebc132e2-cfc3-42c0-a445-3b4465cb523c>

Power Controls Accounts cluster1-mgmt cluster2-mgmt Blue XP

NetApp | StorageGRID Tenant Manager

Users > Demo S3 User

Overview

Full name:	Demo S3 User edit
Username:	demo_s3_user
User type:	Local
Denied access:	Yes
Access mode:	No Groups
Group membership:	None

>Password [edit](#) Access [edit](#) Access keys [edit](#) Groups [edit](#)

Change password

Change this user's password.

***** [Change Password](#)

To create the S3 keys click on the user name.

NetApp | StorageGRID Tenant Manager

DASHBOARD STORAGE (S3) My access keys Buckets Platform services endpoints ACCESS MANAGEMENT Groups Users Identity federation

Users

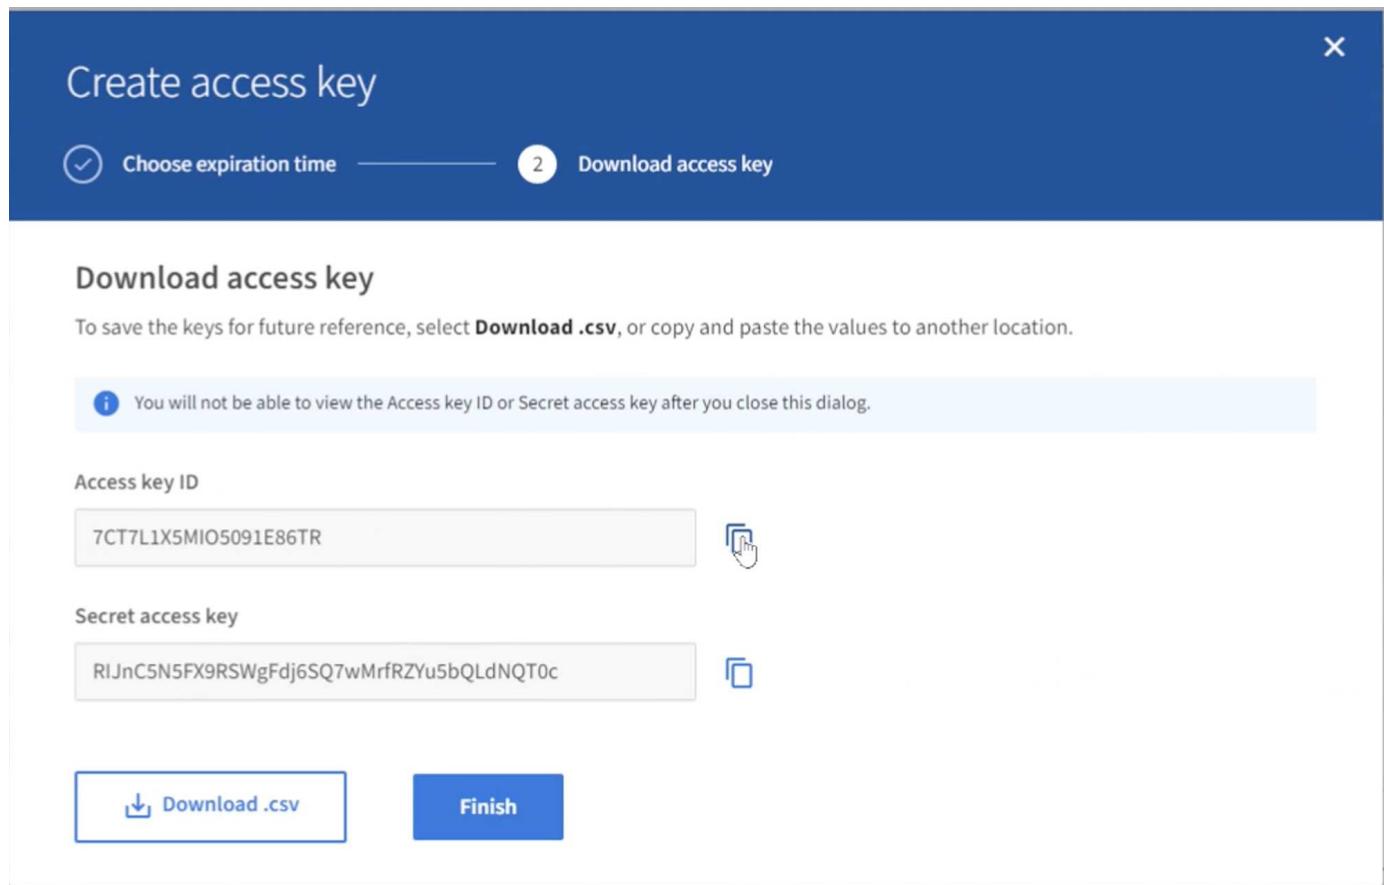
View local and federated users. Edit properties and group membership of local users.

2 users [Create user](#)

Actions	Username	Full Name	Denied	Type
<input type="checkbox"/>	root	Root		Local
<input type="checkbox"/>	demo_s3_user	Demo S3 User	✓	Local

← Previous 1 Next →

Select the "Access keys" tab and click on the "Create Key" button. There is no need to set an expiration time. Download the S3 keys as they cannot be retrieved again once the window is closed.



Create the security group

Now go to the Groups page and create a new group.

Create group

X

- 1 Choose a group type
- 2 Manage permissions
- 3 Set S3 group policy
- 4 Add users
Optional

Choose a group type ?

Create a new local group or import a group from the external identity source.

Local group

Federated group

Create local groups to assign permissions to any local users you defined in StorageGRID.

Display name

Demo S3 Group

Must contain at least 1 and no more than 32 characters

Unique name ?

demo_s3_group

Cancel

Continue



Set the group permissions to Read-Only. This is the Tenant UI permissions, not the S3 permissions.



Choose a group type



Manage permissions



Set S3 group policy

Add users
Optional

Manage group permissions

Select an access mode for this group and select one or more permissions.

Access mode [?](#)

Select whether users can change settings and perform operations or whether they can only view settings and features.



Read-write



Read-only

Group permissions [?](#)

Select the permissions you want to assign to this group.

[Root access](#)

Allows users to access all administration features. Root access permission supersedes all other permissions.

[Manage all buckets](#)

Allows users to change settings of all S3 buckets (or Swift containers) in this account.

[Manage endpoints](#)

Allows users to configure endpoints for platform services.

[Manage your own S3 credentials](#)

Allows users to create and delete their own S3 access keys.

[Previous](#)[Continue](#)

S3 permissions are controlled with the group policy (IAM Policy). Set the Group policy to custom and paste the json policy in the box. This policy will allow users of this group to list the buckets of the tenant and perform any S3 operations in the bucket named "bucket" or sub-folders in the bucket named "bucket".

```
{
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "s3>ListAllMyBuckets",
      "Resource": "arn:aws:s3:::*"
    },
    {
      "Effect": "Allow",
      "Action": "s3:*",
      "Resource": ["arn:aws:s3:::bucket", "arn:aws:s3:::bucket/*"]
    }
  ]
}
```

Create group

1 Choose a group type — 2 Manage permissions — 3 Set S3 group policy — 4 Add users Optional

Set S3 group policy ?

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

No S3 Access

Read Only Access

Full Access

Custom

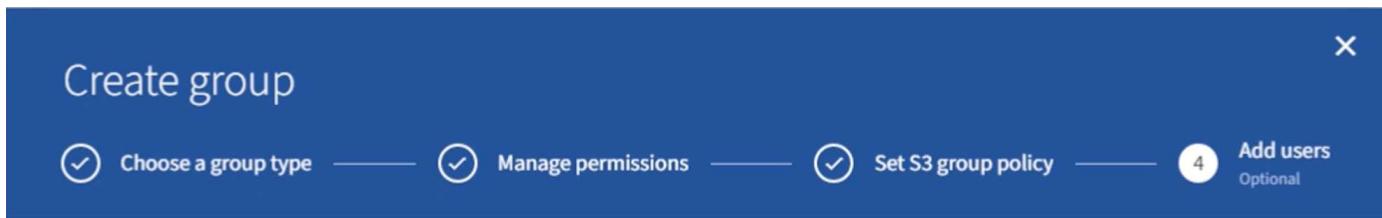
(Must be a valid JSON formatted string.)

```

    "Effect": "Allow",
    "Action": "s3>ListAllMyBuckets",
    "Resource": "arn:aws:s3:::*"
  },
  {
    "Effect": "Allow",
    "Action": "s3:*",
    "Resource": ["arn:aws:s3:::bucket", "arn:aws:s3:::bucket/*"]
  }
}
  
```

[Previous](#) [Continue](#)

Finally, add the user to the group and finish.



Add users

(This step is optional. If required, you can save this group and add users later.)

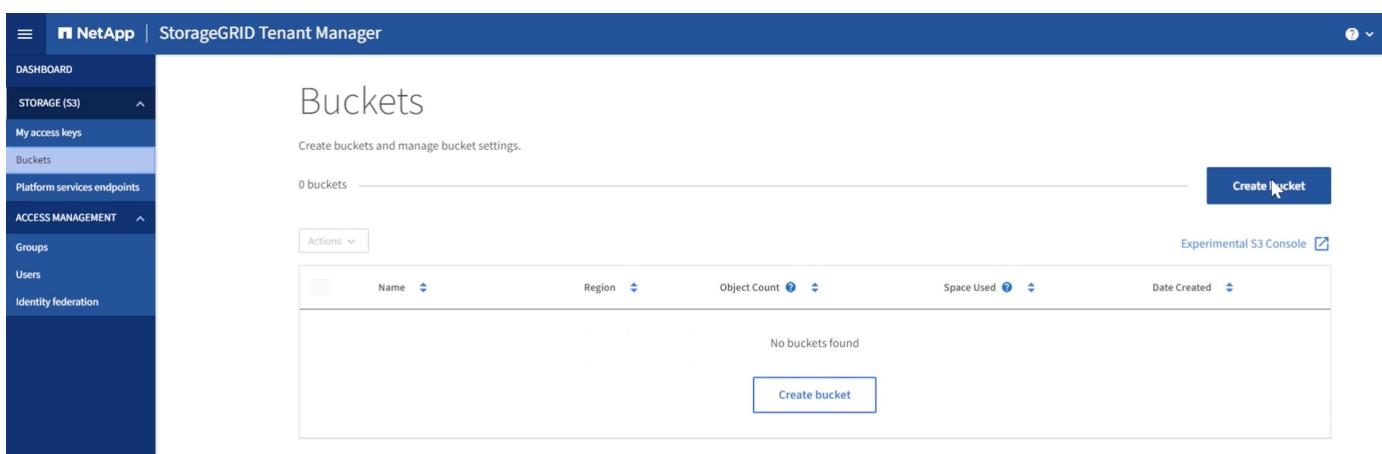
Select local users to add to the group **Demo S3 Group**.

Username	Full Name	Access Level
demo_s3_user	Demo S3 User	✓

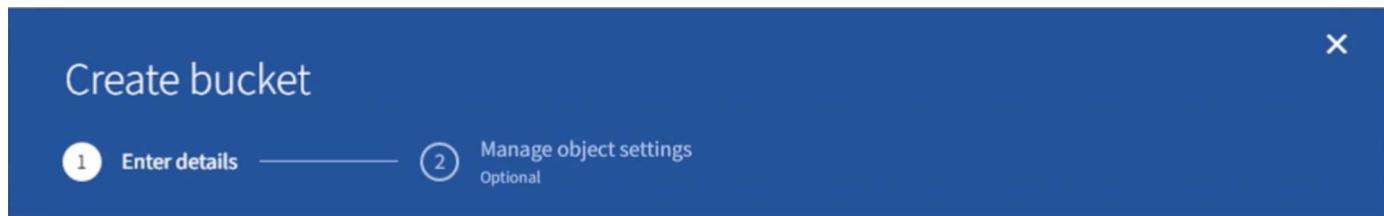
Previous Create group

Create two buckets

Navigate to the buckets tab and click on the Create bucket button.



Define the bucket name and region.



Enter bucket details

Enter the bucket's name and select the bucket's region.

Bucket name [?](#)

bucket

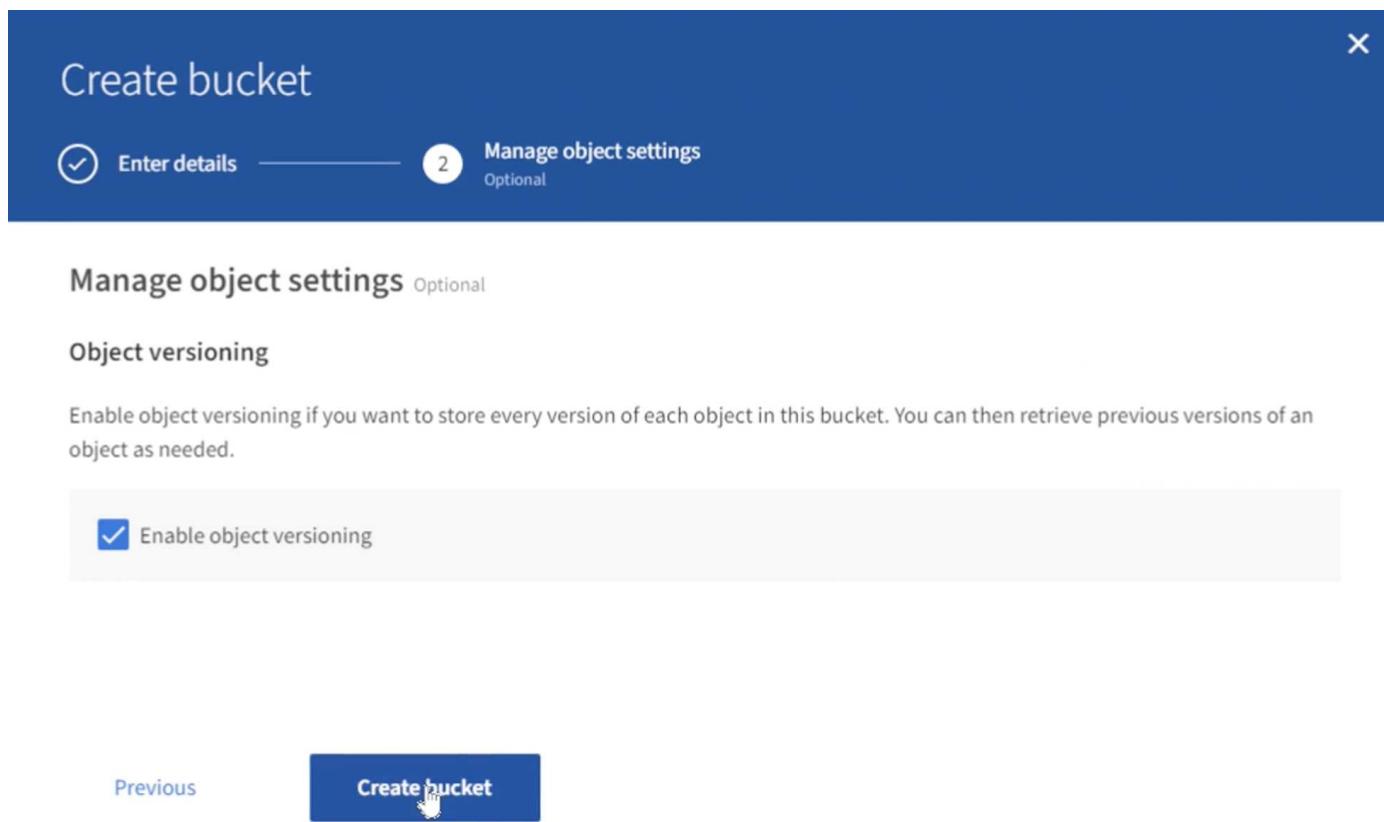
Region [?](#)

us-east-1

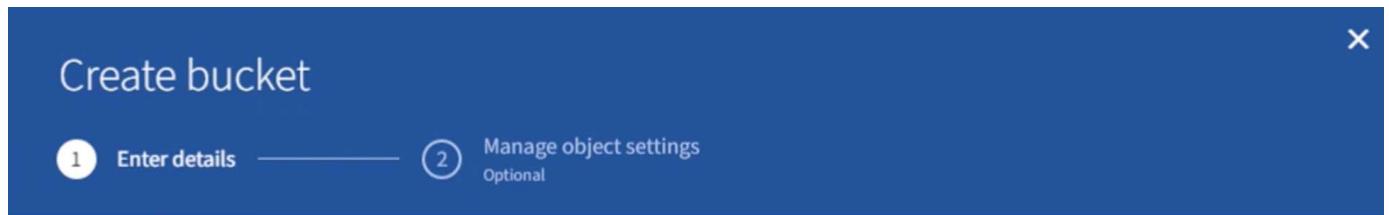
[Cancel](#)

[Continue](#)

On this first bucket enable versioning.



Now create a second bucket without versioning enabled.



Enter bucket details

Enter the bucket's name and select the bucket's region.

Bucket name ?

sg-dummy

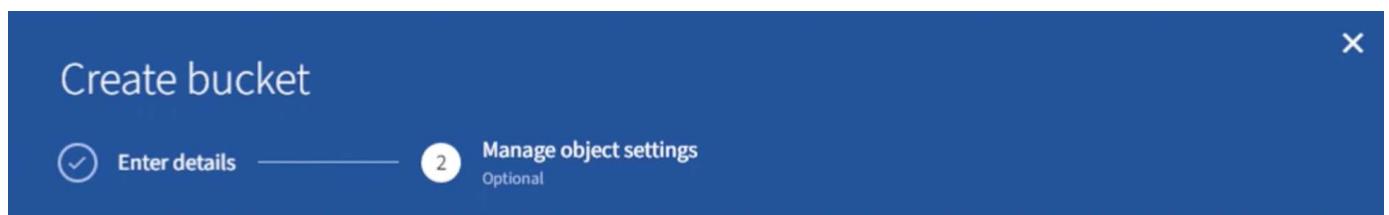
Region ?

us-east-1

Cancel

Continue

Do not enable versioning on this second bucket.



Previous

Create bucket

By Rafael Guedes, and Aron Klein

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Populate the Source Bucket

Lets put some objects in the source ONTAP bucket. We will use S3Browser for this demo but you could use any tool you are comfortable with.

Using the ONTAP user s3 keys created above, configure S3Browser to connect to your ontap system.

 Add New Account

Add New Account

Enter new account details and click Add new account

Display name:

Bucket (original and post-migration)

Assign any name to your account.

Account type:

S3 Compatible Storage

Choose the storage you want to work with. Default is Amazon S3 Storage.

REST Endpoint:

s3portal.demo.netapp.com:8080

Specify S3-compatible API endpoint. It can be found in storage documentation. Example: rest.server.com:8080

Access Key ID:

3TVPI142JGE3Y7FV2KC0

Required to sign the requests you send to Amazon S3, see more details at <https://s3browser.com/keys>

Secret Access Key:

Required to sign the requests you send to Amazon S3, see more details at <https://s3browser.com/keys>

Encrypt Access Keys with a password:

Turn this option on if you want to protect your Access Keys with a master password.

Use secure transfer (SSL/TLS)

If checked, all communications with the storage will go through encrypted SSL/TLS channel

[advanced settings..](#)

 [Add new account](#)  [Cancel](#)

Now lets upload some files to the versioning enabled bucket.

S3 Browser 11.6.7 - Free Version (for non-commercial use only) - Bucket (original and post-migration)

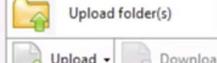
Accounts Buckets Files Bookmarks Tools Upgrade to Pro! Help

New bucket + Add external bucket Refresh

Path: /

Name	Size	Type	Last Modified	Storage Class

Upload file(s) 

Upload folder(s) 

Upload Download Delete New Folder Refresh

Tasks (1) Permissions Headers Tags Properties Preview Versions Eventlog

Task	Size	%	Progress	Status	Speed

S3 Browser 11.6.7 - Free Version (for non-commercial use only) - Bucket (original and post-migration)

Open

Downloads This PC Downloads

Search Downloads

Organize New folder

Name	Date modified	Type	Size
9141P1_q_image.tgz	3/22/2024 1:25 AM	TGZ File	2,641,058 KB
cluster1_demo_s3_user_s3_user.txt	3/23/2024 11:04 PM	Text Document	1 KB
cluster1_svm_demo_s3_details (1).txt	3/23/2024 11:03 PM	Text Document	1 KB
cluster1_svm_demo_s3_details.txt	3/23/2024 11:01 PM	Text Document	1 KB
hfs.exe	3/22/2024 1:24 AM	Application	2,121 KB
hotfix-install-11.6.0.14	3/23/2024 11:55 AM	14 File	717,506 KB
putty	7/18/2020 6:39 PM	Shortcut	2 KB
s3browser-11-6-7.exe	3/23/2024 12:36 PM	Application	9,807 KB

File name: "s3browser-11-6-7.exe" "cluster1_demo_s3_user_s3_user.txt" "cluster1_svm_demo_s3_details (1).txt" "cluster1_svm_demo_s3_details.txt" "hfs.exe" "putty"

Open Cancel

S3 Browser 11.6.7 - Free Version (for non-commercial use only) - Bucket (original and post-migration)

Accounts Buckets Files Bookmarks Tools Upgrade to Pro! Help

New bucket Add external bucket Refresh

Path: /

Name	Size	Type	Last Modified	Storage Class
cluster1 дем...	157 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
cluster1_svm...	211 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
cluster1_svm...	211 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
putty.exe	834.05 KB	Application	3/23/2024 11:23:25 PM	STANDARD
hfs.exe	2.07 MB	Application	3/23/2024 11:23:25 PM	STANDARD
s3browser-11...	9.58 MB	Application	3/23/2024 11:23:26 PM	STANDARD

Upload Download Delete New Folder Refresh

Tasks (1) Permissions Headers Tags Properties Preview Versions Eventlog

Now lets create some object versions in the bucket.

Delete a file.

S3 Browser 11.6.7 - Free Version (for non-commercial use only) - Bucket (original and post-migration)

Accounts Buckets Files Bookmarks Tools Upgrade to Pro! Help

New bucket Add external bucket Refresh

Path: /

Name	Size	Type	Last Modified	Storage Class
cluster1 дем...	157 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
cluster1_svm...	211 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
cluster1_svm...	211 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
putty.exe	834.05 KB	Application	3/23/2024 11:23:25 PM	STANDARD
hfs.exe	2.07 MB	Application	3/23/2024 11:23:25 PM	STANDARD
s3browser-11...	9.58 MB	Application	3/23/2024 11:23:26 PM	STANDARD

Upload Download Delete New Folder

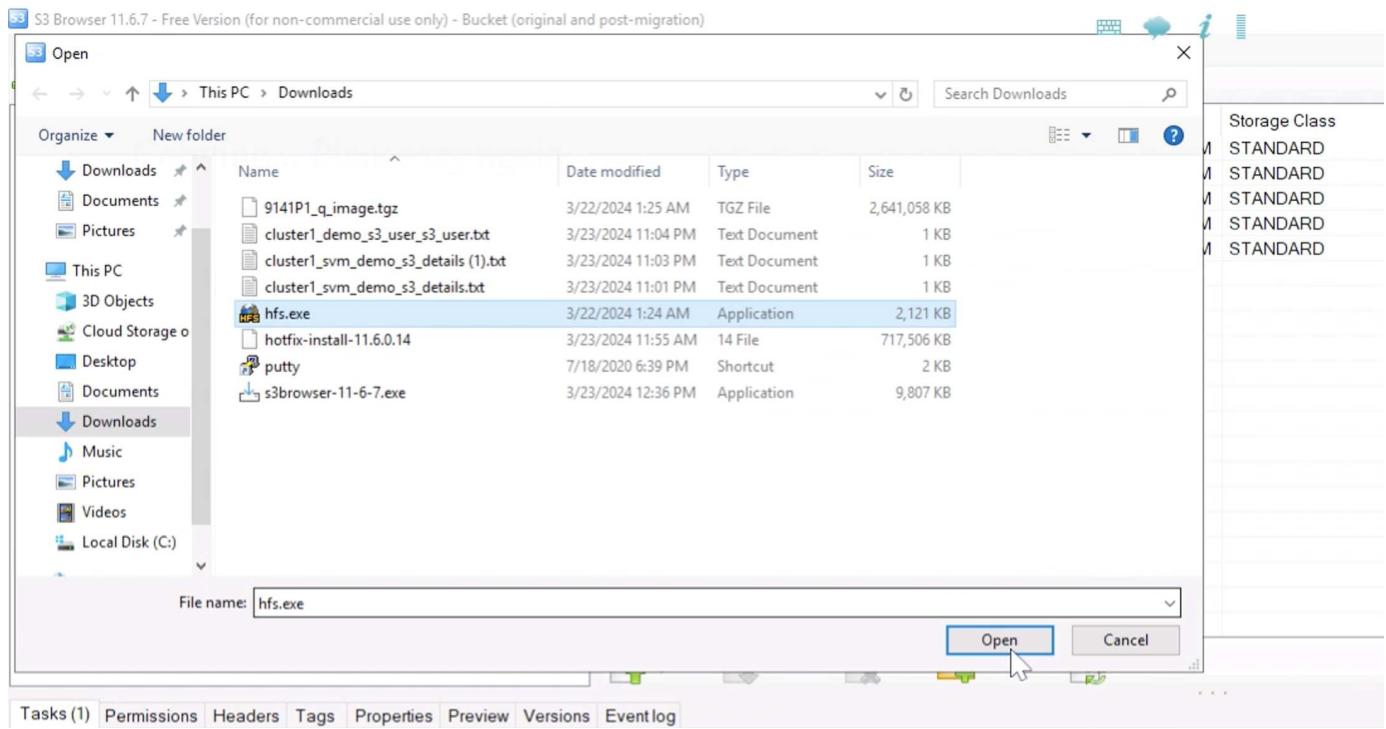
Confirm File Delete

Are you sure to delete 'putty.exe'?

Yes No

Tasks (1) Permissions Headers Tags Properties Preview Versions Eventlog

Upload a file that already exists in the bucket to copy the file over itself and create a new version of it.



In S3Browser we can view the versions of the objects we just created.

Establish the replication relationship

Lets start sending data from ONTAP to StorageGRID.

In ONTAP System Manager navigate to "Protection/Overview". Scroll down to "Cloud object stores". and click the "Add" button and select "StorageGRID".

Lets you select specific volumes for protection if you don't need to protect entire storage VMs.

Lets you select which volumes you want to be backed up to a cloud destination.

NetApp SnapCenter software simplifies backup, restore, and clone management for the applications hosted across ONTAP enabled platforms. [Use NetApp SnapCenter for application-consistent protection.](#)

Bucket protection

SnapMirror (local or remote)

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100% **2 of the 2 buckets aren't protected.**

Back up to cloud

0% 10% 20% 30% 40% 50% 60% 70% 80% 90% 100% **2 of the 2 buckets aren't backed up to the cloud.**

Protect buckets

Lets you select specific buckets for setting up SnapMirror protection to either an ONTAP destination or a cloud destination.

Object store providers

- StorageGRID
- ONTAP S3
- Amazon S3
- Others

+ Add

Input the StorageGRID information by providing a name, URL style (for this demo we will use Path-style URLs). Set the object store scope to "Storage VM".

Add cloud object store

NAME

URL STYLE

Path-style URL

OBJECT STORE SCOPE

Cluster Storage VM

USE BY

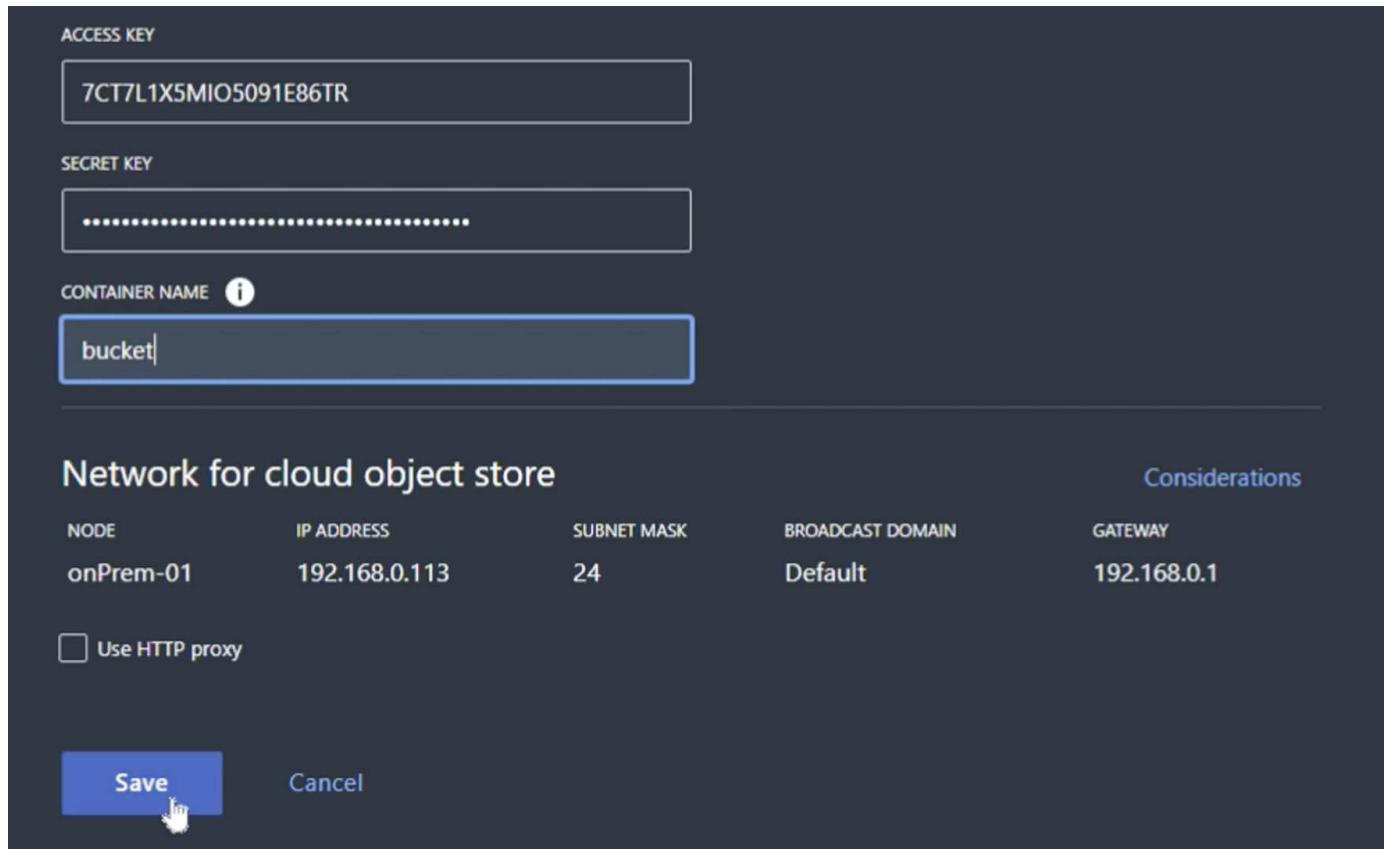
SnapMirror ONTAP S3 SnapMirror

SERVER NAME (FQDN)

If you are using SSL, set the load balancer endpoint port and copy in the StorageGRID endpoint certificate

here. otherwise uncheck the SSL box and input the HTTP endpoint port here.

Input the StorageGRID user S3 keys and bucket name from the StorageGRID configuration above for the destination.

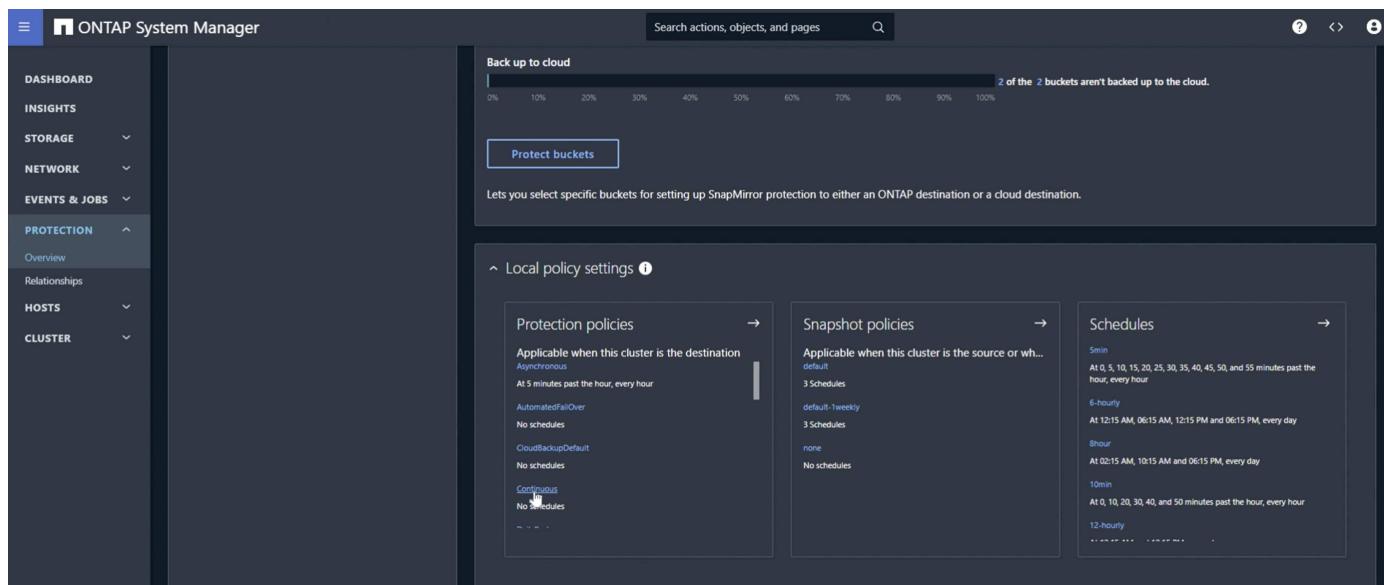


NODE	IP ADDRESS	SUBNET MASK	BROADCAST DOMAIN	GATEWAY
onPrem-01	192.168.0.113	24	Default	192.168.0.1

Use HTTP proxy

Save **Cancel**

Now that we have a destination target configured, we can configure the policy settings for the target. Expand "Local policy settings" and select "continuous".



Edit the continuous policy and change the "Recovery point objective" from "1 Hours" to "3 Seconds".

Policies Protection overview

Name	Description	Policy type	Scope
Continuous	Policy for S3 bucket mirroring.	(All)	Cluster
THROTTLE Unlimited	RECOVERY POINT OBJECTIVE 1 Hours		

Now we can configure snapmirror to replicate the bucket.

```
snapmirror create -source-path sv_demo: /bucket/bucket -destination-path sgws_demo: /objstore -policy Continuous
```

```
cluster1-mgmt
Using username "admin".
Using keyboard-interactive authentication.
Password:

Last login time: 3/24/2024 00:02:00
cluster1::> snapmirror create -source-path sv_demo:/bucket/bucket -destination-path sgws_demo:/objstore -policy Continuous
[Job 220] Job is queued: Create an S3 SnapMirror relationship between bucket "sv_demo:bucket" and bucket "objstore/sgws_demo"..
cluster1::> [green bar]
```

The bucket will now show a cloud symbol in the bucket list under protection.

Name	Storage VM	Lifecycle rules	Capacity (available total)	Protection	Path
bucket	svm_demo	0	100 GiB 100 GiB	Cloud	-
ontap-dummy	svm_demo	0	100 GiB 100 GiB	Cloud	-

If we select the bucket and go to the "SnapMirror (ONTAP or Cloud)" tab we will see the snapmirror relationship status.

The screenshot shows the 'Buckets' page in the NetApp StorageGRID web interface. A table displays a SnapMirror relationship between a source bucket ('svm_demo/bucket/bucket') and a destination ('sgws_demo/objstore'). The relationship is set to 'Continuous' and is marked as 'Healthy' and 'Mirrored'.

The replication details

We now have a successfully replicating bucket from ONTAP to StorageGRID. But what is actually replicating? Our source and destination are both versioned buckets. Do the previous versions also replicate to the destination? If we look at our StorageGRID bucket with S3Browser we see that the existing versions did not replicate and our deleted object does not exist, nor does a delete marker for that object. Our duplicated object only has 1 version in the StorageGRID bucket.

The screenshot shows the S3 Browser interface. The left sidebar shows buckets 'bucket' and 'sg-dummy'. The main area displays a table of objects in the 'sg-dummy' bucket. The 'hfs.exe' file is selected, showing its details: Size 2.07 MB, Type Application, Last Modified 3/24/2024 12:13:53 AM, Storage Class STANDARD. Below the table are buttons for Upload, Download, Delete, New Folder, and Refresh. At the bottom, a table shows the version history for 'hfs.exe'.

Key	Last Modified	ETag	Size	Storage Class	Owner	Version Id
hfs.exe	3/24/2024 12:13:53 AM	"9e0557e98ed1269372ff0ace91d6347"	2.07 MB	STANDARD	tenant_demo (27041610751...	NJUSRDRhCNDIIR...
revision #, 1 (current)						

In our ONTAP bucket, let's add a new version to our same object that we used previously and see how it replicates.

Name	Size	Type	Last Modified	Storage Class
cluster1_demo_157 bytes	157 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
cluster1_svm_211 bytes	211 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
cluster1_svm_211 bytes	211 bytes	Text Document	3/23/2024 11:23:25 PM	STANDARD
putty.exe	834.05 KB	Application	3/23/2024 11:23:25 PM	STANDARD
hfs.exe	2.07 MB	Application	3/24/2024 12:14:52 AM	STANDARD
s3browser-11...	9.58 MB	Application	3/23/2024 11:23:26 PM	STANDARD

Key	Last Modified	ETag	Size	Storage Class	Owner	Version Id
cluster1_demo_s3_user_s3_user.txt	3/23/2024 11:23:25 PM	acf4c9543e97e0678b2b6ed5a60e1bc	157 bytes	STANDARD	Unknown (Unknown)	Mzg0MjQ1MDAwL...
cluster1_svm_demo_s3_details (1).txt	3/23/2024 11:23:25 PM	407753b646a6cfef19de71eefb5ff04	211 bytes	STANDARD	Unknown (Unknown)	NDg0MjQ1MDAw...
cluster1_svm_demo_s3_details.txt	3/23/2024 11:23:25 PM	17d20651856480a587a039ecc10e2	211 bytes	STANDARD	Unknown (Unknown)	NTU2NzI0MDAwL...
hfs.exe	revision # 3 (current)	9e8557e98ed1269372fface91d63477	2.07 MB	STANDARD	Unknown (Unknown)	NTY0NDg0MDAw...
revision # 2	3/23/2024 11:23:36 PM	9e8557e98ed1269372fface91d63477	2.07 MB	STANDARD	Unknown (Unknown)	NzQ1OTE4MDAw...
revision # 1	3/23/2024 11:23:25 PM	9e8557e98ed1269372fface91d63477	2.07 MB	STANDARD	Unknown (Unknown)	NjK2ODI0MDAwL...
putty.exe	revision # 1 (current)	54cb91395daad9d7882533c21fe0e9	834.05 KB	STANDARD	Unknown (Unknown)	NzE2NzEyMDAwL...
s3browser-11-6-7.exe	revision # 1 (current)	ae36fb97054782962d6597c5d082d2	9.58 MB	STANDARD	Unknown (Unknown)	NDY2ODcwMDEu...

If we look on the StorageGRID side we see that a new version has been created in this bucket too, but is missing the initial version from before the snapmirror relationship.

Name	Size	Type	Last Modified	Storage Class
cluster1_demo_157 bytes	157 bytes	Text Document	3/24/2024 12:13:53 AM	STANDARD
cluster1_svm_211 bytes	211 bytes	Text Document	3/24/2024 12:13:53 AM	STANDARD
cluster1_svm_211 bytes	211 bytes	Text Document	3/24/2024 12:13:53 AM	STANDARD
putty.exe	834.05 KB	Application	3/24/2024 12:14:28 AM	STANDARD
hfs.exe	2.07 MB	Application	3/24/2024 12:14:56 AM	STANDARD
s3browser-11...	9.58 MB	Application	3/24/2024 12:13:53 AM	STANDARD

Key	Last Modified	ETag	Size	Storage Class	Owner	Version Id
hfs.exe	revision # 2 (current)	9e8557e98ed1269372fface91d63477	2.07 MB	STANDARD	tenant_demo (27041610751...	OE14RjYNDgIRT...
revision # 1	3/24/2024 12:13:53 AM	9e8557e98ed1269372fface91d63477	2.07 MB	STANDARD	tenant_demo (27041610751...	NjU5RDhCNDI0T...

This is because the ONTAP SnapMirror S3 process only replicates the current version of the object. This is why we created a versioned bucket on the StorageGRID side to be the destination. This way StorageGRID can maintain a version history of the objects.

By Rafael Guedes, and Aron Klein

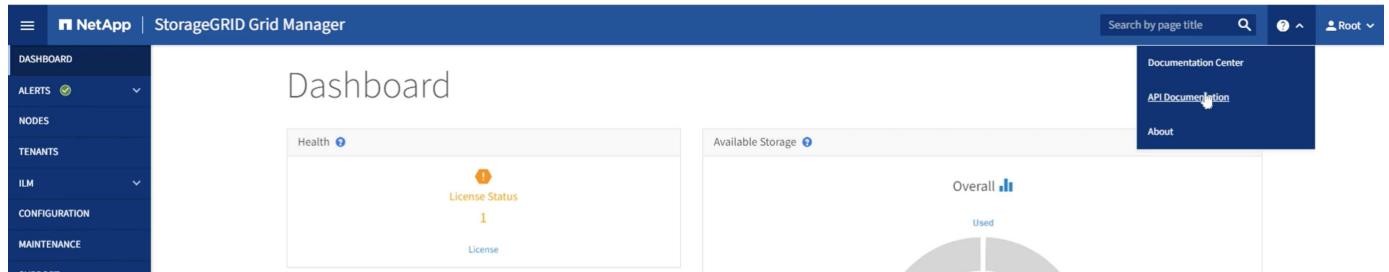
Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

Enabling enterprise-grade S3 by seamlessly migrating object-based storage from ONTAP S3 to StorageGRID

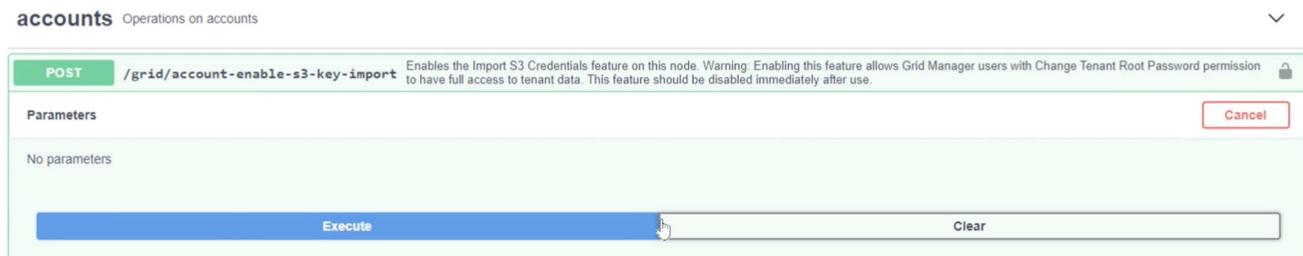
Migrate S3 Keys

For a migration, most of the time you will want to migrate the credentials for the users rather than generate new credentials on the destination side. StorageGRID provides api's to allow s3 keys to be imported to a user.

Logging into the StorageGRID management UI (not the tenant manager UI) open the API Documentation swagger page.



Expand the "accounts" section, select the "POST /grid/account-enable-s3-key-import", click the "Try it out" button, then click on the execute button.



Now scroll down still under "accounts" to "POST /grid/accounts/{id}/users/{user_id}/s3-access-keys"

Here is where we are going to input the tenant ID and user account ID we collected earlier. fill in the fields and the keys from our ONTAP user in the json box. you can set the expiration of the keys, or remove the ", "expires": 123456789" and click on execute.

POST`/grid/accounts/{id}/users/{user_id}/s3-access-keys` Imports S3 credentials for a given user in a tenant account**Parameters**

Name

Description

id * requiredstring
(path)

ID of Storage Tenant Account

27041610751165610501

user_id * requiredstring
(path)

ID of user in tenant account.

ebc132e2-cfc3-42c0-a445-3b4465cb523c

body * required[Edit Value](#) | Model

(body)

```
{  
  "accessKey": "3TVPI142JGE3Y7FV2KC0",  
  "secretAccessKey": "75a10qKBU4quA132twI4g41C4Gg5PP30ncy0sPE8"  
}
```

Once you have completed all of your user key imports you should disable the key import function in "accounts" "POST /grid/account-disable-s3-key-import"

POST`/grid/account-disable-s3-key-import` Disables the Import S3 Credentials feature on this node.[Cancel](#)**Parameters**

No parameters

[Execute](#)**Responses**

Response content type

application/json



If we look at the user account in the tenant manager UI, we can see the new key has been added.

Overview

Full name: ?	Demo S3 User 
Username: ?	demo_s3_user
User type: ?	Local
Denied access: ?	Yes
Access mode: ?	Read-only
Group membership: ?	Demo S3 Group

- [Password](#)
- [Access](#)
- [Access keys](#)
- [Groups](#)

Manage access keys

Add or delete access keys for this user.

Create key	Actions 
<input type="checkbox"/>	Access key ID  
<input type="checkbox"/>	*****86TR   Expiration time  
<input type="checkbox"/>	*****2KC0   None

The final cut-over

If the intention is to have a perpetually replicating bucket from ONTAP to StorageGRID, you can end here. If this is a migration from ONTAP S3 to StorageGRID, then its time to put an end to it and cut over.

Inside ONTAP system manager, edit the S3 group and set it to "ReadOnlyAccess". This will prevent the users from writing to the ONTAP S3 bucket anymore.

Edit group

X

NAME

demo_s3_group

USERS

demo_s3_user x

POLICIES

ReadOnlyAccess x

Cancel

Save

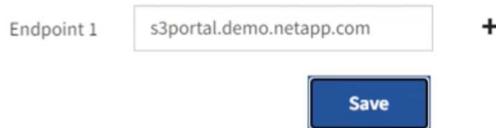


All that is left to do is configure DNS to point from the ONTAP cluster to the StorageGRID endpoint. Make sure your endpoint certificate is correct and if you need virtual hosted style requests then add the endpoint domain names in storageGRID

Endpoint Domain Names

Virtual Hosted-Style Requests

Enable support of S3 virtual hosted-style requests by specifying API endpoint domain names. Support is disabled if this list is empty. Examples: s3.example.com, s3.example.co.uk, s3-east.example.com



Endpoint 1 s3portal.demo.netapp.com + **Save**

Your clients will either need to wait for the TTL to expire, or flush DNS to resolve to the new system so you can test that everything is working. All that is left is to clean up the initial temporary S3 keys we used to test the StorageGRID data access (NOT the imported keys), remove the snapmirror relationships, and remove the ONTAP data.

By Rafael Guedes, and Aron Klein

Tool and application guides

Use Cloudera Hadoop S3A connector with StorageGRID

By *Angela Cheng*

Hadoop has been a favorite of data scientists for some time now. Hadoop allows for the distributed processing of large data sets across clusters of computers using simple programming frameworks. Hadoop was designed to scale up from single servers to thousands of machines, with each machine possessing local compute and storage.

Why use S3A for Hadoop workflows?

As the volume of data has grown over time, the approach of adding new machines 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 client 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 enables 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.

Configure S3A connector to use StorageGRID

Prerequisites

- A StorageGRID S3 endpoint URL, a tenant s3 access key, and a secret key for Hadoop S3A connection testing.
- A Cloudera cluster and root or sudo permission to each host in the cluster to install the Java package.

As of April 2022, Java 11.0.14 with Cloudera 7.1.7 was tested against StorageGRID 11.5 and 11.6. However, the Java version number might be different at the time of a new install.

Install Java package

1. Check the [Cloudera support matrix](#) for the supported JDK version.
2. Download the [Java 11.x package](#) that matches the Cloudera cluster operating system. Copy this package to each host in the cluster. In this example, the rpm package is used for CentOS.
3. Log into each host as root or using an account with sudo permission. Perform the following steps on each host:
 - a. Install the package:

```
$ sudo rpm -Uvh jdk-11.0.14_linux-x64_bin.rpm
```

- b. Check where Java is installed. If multiple versions are installed, set the newly installed version as default:

```

alternatives --config java

There are 2 programs which provide 'java'.

Selection      Command
-----
+1            /usr/java/jre1.8.0_291-amd64/bin/java
 2            /usr/java/jdk-11.0.14/bin/java

Enter to keep the current selection[+], or type selection number: 2

```

c. Add this line to the end of `/etc/profile`. The path should match the path of above selection:

```
export JAVA_HOME=/usr/java/jdk-11.0.14
```

d. Run the following command for the profile to take effect:

```
source /etc/profile
```

Cloudera HDFS S3A configuration

Steps

1. From the Cloudera Manager GUI, select Clusters > HDFS, and select Configuration.
2. Under CATEGORY, select Advanced, and scroll down to locate Cluster-wide Advanced Configuration Snippet (Safety Valve) for `core-site.xml`.
3. Click the (+) sign and add following value pairs.

Name	Value
fs.s3a.access.key	<i><tenant s3 access key from StorageGRID></i>
fs.s3a.secret.key	<i><tenant s3 secret key from StorageGRID></i>
fs.s3a.connection.ssl.enabled	[true or false] (default is https if this entry is missing)
fs.s3a.endpoint	<i><StorageGRID S3 endpoint:port></i>
fs.s3a.impl	org.apache.hadoop.fs.s3a.S3AFileSystem
fs.s3a.path.style.access	[true or false] (default is virtual host style if this entry is missing)

Sample screenshot

Cluster-wide Advanced Configuration Snippet (Safety Valve) for core-site.xml

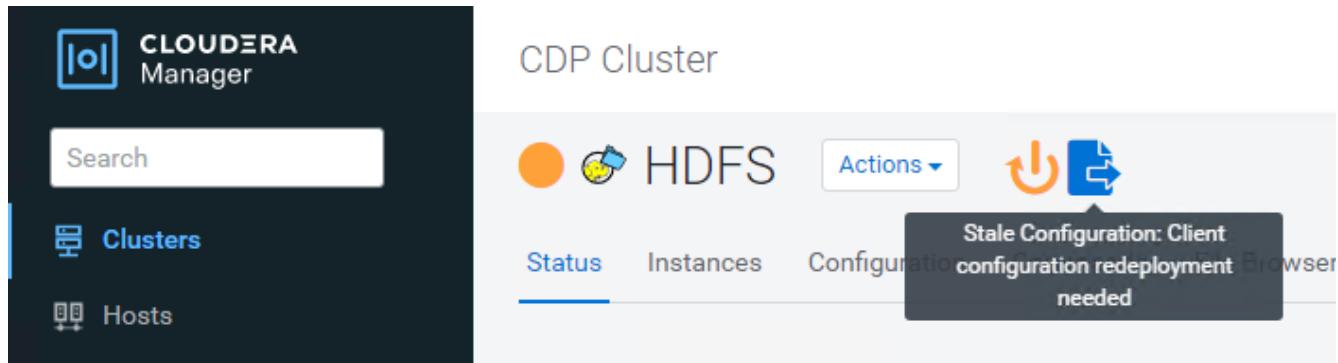
HDFS (Service-Wide)

Name	fs.s3a.endpoint	
Value	sgdemo.netapp.com:10443	
Description	StorageGRID s3 load balancer endpoint	
<input checked="" type="checkbox"/> Final		
Name	fs.s3a.access.key	
Value	OMC...BAN	
Description	SG CDP S3 access key	
<input checked="" type="checkbox"/> Final		
Name	fs.s3a.secret.key	
Value	mapz...Qfc	
Description	SG CDP S3 secret key	
<input checked="" type="checkbox"/> Final		
Name	fs.s3a.impl	
Value	org.apache.hadoop.fs.s3a.S3AFileSystem	
Description		
<input checked="" type="checkbox"/> Final		
Name	fs.s3a.path.style.access	
Value	true	
Description		
<input checked="" type="checkbox"/> Final		

Cluster-wide Advanced Configuration Snippet (Safety Valve) for core-site.xml

4. Click the Save Changes button. Select the Stale Configuration icon from the HDFS menu bar, select

Restart Stale Services on the next page, and select Restart Now.



The screenshot shows the Cloudera Manager interface. On the left, there's a sidebar with a search bar, 'Clusters' (which is selected and highlighted in blue), and 'Hosts'. The main area is titled 'CDP Cluster' and shows an 'HDFS' icon with a yellow circle and a blue gear. To the right of the icon are 'Actions' (with a dropdown arrow), a power icon, and a refresh icon. Below these are tabs for 'Status', 'Instances', 'Configuration', and 'Browser'. A tooltip box is overlaid on the 'Configuration' tab, containing the text: 'Stale Configuration: Client configuration redeployment Browser needed'.

Test S3A connection to StorageGRID

Perform basic connection test

Log into one of the hosts in the Cloudera cluster, and enter `hadoop fs -ls s3a://<bucket-name>/`.

The following example uses path style with a pre-existing hdfs-test bucket and a test object.

```
[root@ce-n1 ~]# hadoop fs -ls s3a://hdfs-test/
22/02/15 18:24:37 WARN impl.MetricsConfig: Cannot locate configuration:
tried hadoop-metrics2-s3a-file-system.properties,hadoop-
metrics2.properties
22/02/15 18:24:37 INFO impl.MetricsSystemImpl: Scheduled Metric snapshot
period at 10 second(s).
22/02/15 18:24:37 INFO impl.MetricsSystemImpl: s3a-file-system metrics
system started
22/02/15 18:24:37 INFO Configuration.deprecation: No unit for
fs.s3a.connection.request.timeout(0) assuming SECONDS
Found 1 items
-rw-rw-rw- 1 root root 1679 2022-02-14 16:03 s3a://hdfs-test/test
22/02/15 18:24:38 INFO impl.MetricsSystemImpl: Stopping s3a-file-system
metrics system...
22/02/15 18:24:38 INFO impl.MetricsSystemImpl: s3a-file-system metrics
system stopped.
22/02/15 18:24:38 INFO impl.MetricsSystemImpl: s3a-file-system metrics
system shutdown complete.
```

Troubleshooting

Scenario 1

Use an HTTPS connection to StorageGRID and get a `handshake_failure` error after a 15 minute timeout.

Reason: Old JRE/JDK version using outdated or unsupported TLS cipher suite for connection to StorageGRID.

Sample error message

```
[root@ce-n1 ~]# hadoop fs -ls s3a://hdfs-test/
22/02/15 18:52:34 WARN impl.MetricsConfig: Cannot locate configuration:
tried hadoop-metrics2-s3a-file-system.properties,hadoop-
metrics2.properties
22/02/15 18:52:34 INFO impl.MetricsSystemImpl: Scheduled Metric snapshot
period at 10 second(s).
22/02/15 18:52:34 INFO impl.MetricsSystemImpl: s3a-file-system metrics
system started
22/02/15 18:52:35 INFO Configuration.deprecation: No unit for
fs.s3a.connection.request.timeout(0) assuming SECONDS
22/02/15 19:04:51 INFO impl.MetricsSystemImpl: Stopping s3a-file-system
metrics system...
22/02/15 19:04:51 INFO impl.MetricsSystemImpl: s3a-file-system metrics
system stopped.
22/02/15 19:04:51 INFO impl.MetricsSystemImpl: s3a-file-system metrics
system shutdown complete.
22/02/15 19:04:51 WARN fs.FileSystem: Failed to initialize filesystem
s3a://hdfs-test/: org.apache.hadoop.fs.s3a.AWSClientIOException:
doesBucketExistV2 on hdfs: com.amazonaws.SdkClientException: Unable to
execute HTTP request: Received fatal alert: handshake_failure: Unable to
execute HTTP request: Received fatal alert: handshake_failure
ls: doesBucketExistV2 on hdfs: com.amazonaws.SdkClientException: Unable to
execute HTTP request: Received fatal alert: handshake_failure: Unable to
execute HTTP request: Received fatal alert: handshake_failure
```

Resolution: Make sure that JDK 11.x or later is installed and set to default the Java library. Refer to the [Install Java package](#) section for more information.

Scenario 2:

Failed to connect to StorageGRID with error message Unable to find valid certification path to requested target.

Reason: StorageGRID S3 endpoint server certificate is not trusted by Java program.

Sample error message:

```
[root@hdp6 ~]# hadoop fs -ls s3a://hdfs-test/
22/03/11 20:58:12 WARN impl.MetricsConfig: Cannot locate configuration:
tried hadoop-metrics2-s3a-file-system.properties,hadoop-
metrics2.properties
22/03/11 20:58:13 INFO impl.MetricsSystemImpl: Scheduled Metric snapshot
period at 10 second(s).
22/03/11 20:58:13 INFO impl.MetricsSystemImpl: s3a-file-system metrics
system started
22/03/11 20:58:13 INFO Configuration.deprecation: No unit for
fs.s3a.connection.request.timeout(0) assuming SECONDS
22/03/11 21:12:25 INFO impl.MetricsSystemImpl: Stopping s3a-file-system
metrics system...
22/03/11 21:12:25 INFO impl.MetricsSystemImpl: s3a-file-system metrics
system stopped.
22/03/11 21:12:25 INFO impl.MetricsSystemImpl: s3a-file-system metrics
system shutdown complete.
22/03/11 21:12:25 WARN fs.FileSystem: Failed to initialize filesystem
s3a://hdfs-test/: org.apache.hadoop.fs.s3a.AWSClientIOException:
doesBucketExistV2 on hdfs: com.amazonaws.SdkClientException: Unable to
execute HTTP request: PKIX path building failed:
sun.security.provider.certpath.SunCertPathBuilderException: unable to find
valid certification path to requested target: Unable to execute HTTP
request: PKIX path building failed:
sun.security.provider.certpath.SunCertPathBuilderException: unable to find
valid certification path to requested target
```

Resolution: NetApp recommends using a server certificate issued by a known public certificate signing authority to make sure that the authentication is secure. Alternatively, add a custom CA or server certificate to the Java trust store.

Complete the following steps to add a StorageGRID custom CA or server certificate to the Java trust store.

1. Backup the existing default Java cacerts file.

```
cp -ap $JAVA_HOME/lib/security/cacerts
$JAVA_HOME/lib/security/cacerts.orig
```

2. Import the StorageGRID S3 endpoint cert to the Java trust store.

```
keytool -import -trustcacerts -keystore $JAVA_HOME/lib/security/cacerts
-storepass changeit -noprompt -alias sg-lb -file <StorageGRID CA or
server cert in pem format>
```

Troubleshooting tips

1. Increase the hadoop log level to DEBUG.

```
export HADOOP_ROOT_LOGGER=hadoop.root.logger=DEBUG,console
```

2. Execute the command, and direct the log messages to error.log.

```
hadoop fs -ls s3a://<bucket-name>/ &>error.log
```

By Angela Cheng

Use S3cmd to test and demonstrate S3 access on StorageGRID

By Aron Klein

S3cmd is a free command line tool and client for S3 operations. You can use s3cmd to test and demonstrate S3 access on StorageGRID.

Install and configure S3cmd

To install S3cmd on a workstation or server, download it from [command line S3 client](#). s3cmd is pre-installed on each StorageGRID node as a tool to aid in troubleshooting.

Initial configuration steps

1. s3cmd --configure
2. Provide only access_key and secret_key, for the rest keep the defaults.
3. Test access with supplied credentials? [Y/n]: n (bypass the test as it will fail)
4. Save settings? [y/N] y
 - a. Configuration saved to '/root/.s3cfg'
5. In .s3cfg make fields host_base and host_bucket empty after the "=" sign :
 - a. host_base =
 - b. host_bucket =



If you specify host_base and host_bucket in step 4, you don't need to specify an endpoint with --host in the CLI. Example:

```
host_base = 192.168.1.91:8082
host_bucket = bucketX.192.168.1.91:8082
s3cmd ls s3://bucketX --no-check-certificate
```

Basic command examples

- **Create a bucket:**

```
s3cmd mb s3://s3cmdbucket --host=<endpoint>:<port> --no-check-certificate
```

- **List all buckets:**

```
s3cmd ls --host=<endpoint>:<port> --no-check-certificate
```

- **List all buckets and their contents:**

```
s3cmd la --host=<endpoint>:<port> --no-check-certificate
```

- **List objects in a specific bucket:**

```
s3cmd ls s3://<bucket> --host=<endpoint>:<port> --no-check-certificate
```

- **Delete a bucket:**

```
s3cmd rb s3://s3cmdbucket --host=<endpoint>:<port> --no-check-certificate
```

- **Put an object:**

```
s3cmd put <file> s3://<bucket> --host=<endpoint>:<port> --no-check-certificate
```

- **Get an object:**

```
s3cmd get s3://<bucket>/<object> <file> --host=<endpoint>:<port> --no-check-certificate
```

- **Delete an object:**

```
s3cmd del s3://<bucket>/<object> --host=<endpoint>:<port> --no-check-certificate
```

Vertica Eon mode database using NetApp StorageGRID as communal storage

By *Angela Cheng*

This guide describes the procedure to create a Vertica Eon Mode database with communal storage on NetApp StorageGRID.

Introduction

Vertica is an analytic database management software. It is a columnar storage platform designed to handle large volumes of data, which enables very fast query performance in a traditionally intensive scenario. A Vertica database runs in one of the two modes: Eon or Enterprise. You can deploy both modes on-premises or in the cloud.

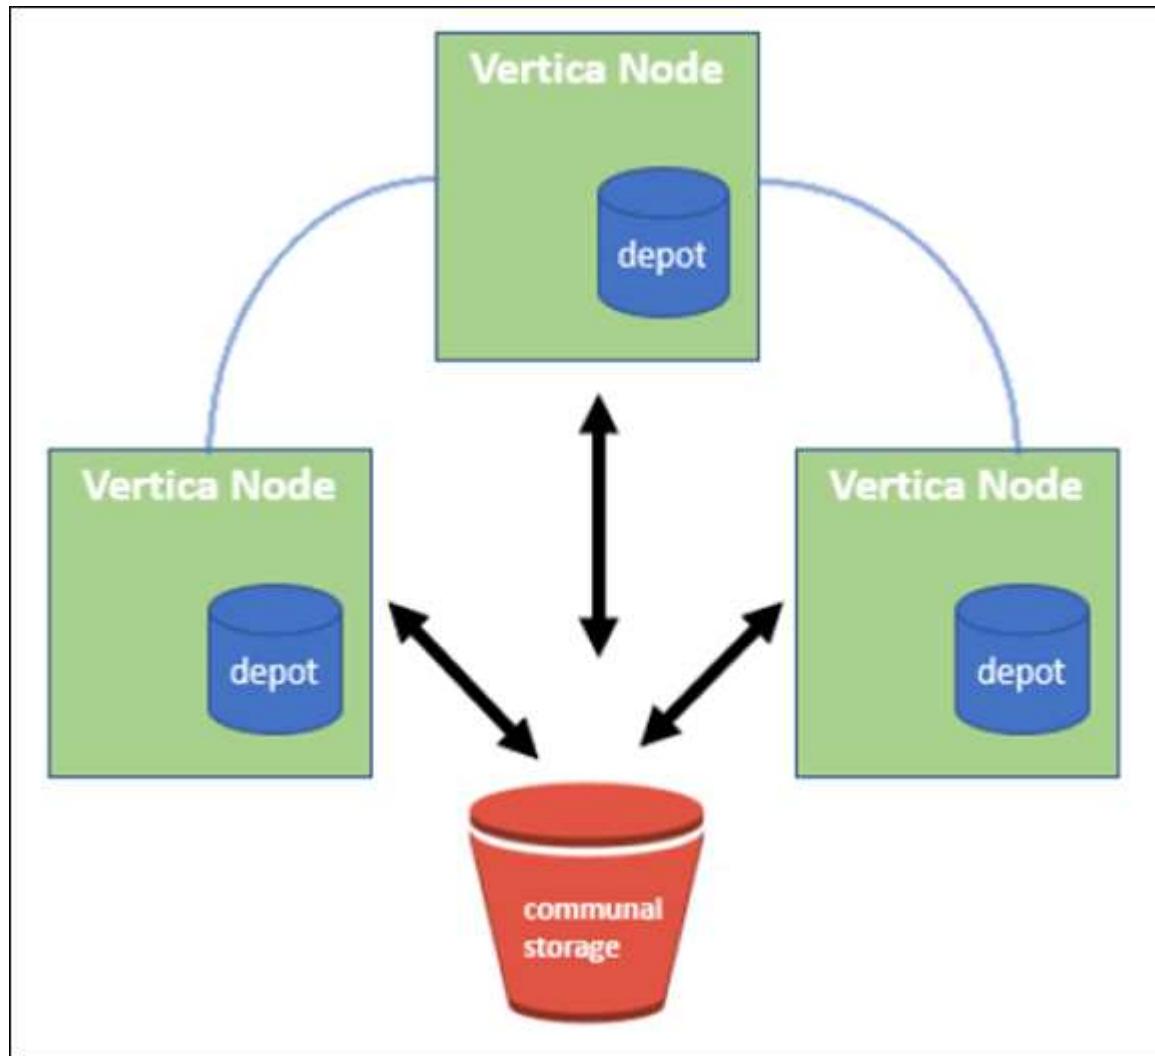
Eon and Enterprise modes primarily differ in where they store data:

- Eon Mode databases use communal storage for their data. This is recommended by Vertica.
- Enterprise Mode databases store data locally in the file system of nodes that make up the database.

Eon Mode architecture

Eon Mode separates the computational resources from the communal storage layer of the database, which allows the compute and storage to scale separately. Vertica in Eon Mode is optimized to address variable workloads and isolate them from one another by using separate compute and storage resources.

Eon Mode stores data in a shared object store called communal storage—an S3 bucket, either hosted on premises or on Amazon S3.



Communal storage

Instead of storing data locally, Eon Mode uses a single communal storage location for all data and the catalog (metadata). Communal storage is the database's centralized storage location, shared among the database nodes.

Communal storage has the following properties:

- Communal storage in the cloud or on-premises object storage is more resilient and less susceptible to data loss due to storage failures than storage on disk on individual machines.

- Any data can be read by any node using the same path.
- Capacity is not limited by disk space on nodes.
- Because data is stored communally, you can elastically scale your cluster to meet changing demands. If the data were stored locally on the nodes, adding or removing nodes would require moving significant amounts of data between nodes to either move it off nodes that are being removed, or onto newly created nodes.

The depot

One drawback of communal storage is its speed. Accessing data from a shared cloud location is slower than reading it from local disk. Also, the connection to communal storage can become a bottleneck if many nodes are reading data from it at once. To improve data access speed, the nodes in an Eon Mode database maintain a local disk cache of data called the depot. When executing a query, the nodes first check whether the data it needs is in the depot. If it is, then it finishes the query by using the local copy of the data. If the data is not in the depot, the node fetches the data from communal storage, and saves a copy in the depot.

NetApp StorageGRID recommendations

Vertica stores database data to object storage as thousands (or millions) of compressed objects (observed size is 200 to 500MB per object). When a user runs database queries, Vertica retrieves the selected range of data from these compressed objects in parallel using the byte-range GET call. Each byte-range GET is approximately 8KB.

During the 10TB database depot off user queries test, 4,000 to 10,000 GET (byte-range GET) requests per second were sent to the grid. When running this test using SG6060 appliances, though the CPU% utilization % per appliance node is low (around 20% to 30%), 2/3 of CPU time is waiting for I/O. A very small percentage (0% to 0.5%) of I/O wait is observed on the SGF6024.

Due to the high demand of small IOPS with very low latency requirements (the average should be less than 0.01 seconds), NetApp recommends using the SGF6024 for object storage services. If the SG6060 is needed for very large database sizes, the customer should work with the Vertica account team on depot sizing to support the actively queried dataset.

For the Admin Node and API Gateway Node, the customer can use the SG100 or SG1000. The choice depends on the number of users' query requests in parallel and database size. If the customer prefers to use a third-party load balancer, NetApp recommends a dedicated load balancer for high performance demand workload. For StorageGRID sizing, consult the NetApp account team.

Other StorageGRID configuration recommendations include:

- **Grid topology.** Do not mix the SGF6024 with other storage appliance models on the same grid site. If you prefer to use the SG6060 for long term archive protection, keep the SGF6024 with a dedicated grid load balancer in its own grid site (either physical or logical site) for an active database to enhance performance. Mixing different models of appliance on same site reduces the overall performance at the site.
- **Data protection.** Use replicate copies for protection. Do not use erasure coding for an active database. The customer can use erasure coding for long term protection of inactive databases.
- **Do not enable grid compression.** Vertica compresses objects before storing to object storage. Enabling grid compression does not further save storage usage and significantly reduces byte-range GET performance.
- **HTTP versus HTTPs S3 endpoint connection.** During the benchmark test, we observed about 5% performance improvement when using an HTTP S3 connection from the Vertica cluster to the StorageGRID load balancer endpoint. This choice should be based on customer security requirements.

Recommendations for a Vertica configuration include:

- **Vertica database default depot settings are enabled (value = 1) for read and write operations.** NetApp strongly recommends keeping these depot settings enabled to enhance performance.
- **Disable streaming limitations.** For configuration details, see the section [Disabling streaming limitations](#).

Installing Eon Mode on-premises with communal storage on StorageGRID

The following sections describe the procedure, in order, to install Eon Mode on-premises with communal storage on StorageGRID. The procedure to configure on-premises Simple Storage Service (S3) compatible object storage is similar to the procedure in the Vertica guide, [Install an Eon Mode Database on-premises](#).

The following setup was used for the functional test:

- StorageGRID 11.4.0.4
- Vertica 10.1.0
- Three virtual machines (VMs) with Centos 7.x OS for Vertica nodes to form a cluster. This setup is for the functional test only, not for the Vertica production database cluster.

These three nodes are set up with a Secure Shell (SSH) key to allow SSH without a password between the nodes within the cluster.

Information required from NetApp StorageGRID

To install Eon Mode on-premises with communal storage on StorageGRID, you must have the following prerequisite information.

- IP address or fully qualified domain name (FQDN) and port number of the StorageGRID S3 endpoint. If you are using HTTPS, use a custom certificate authority (CA) or self-signed SSL certificate implemented on the StorageGRID S3 endpoint.
- Bucket name. It must pre-exist and be empty.
- Access key ID and secret access key with read and write access to the bucket.

Creating an authorization file to access the S3 endpoint

The following prerequisites apply when creating an authorization file to access the S3 endpoint:

- Vertica is installed.
- A cluster is set up, configured, and ready for database creation.

To create an authorization file to access the S3 endpoint, follow these steps:

1. Log in to the Vertica node where you will run `admintools` to create the Eon Mode database.
The default user is `dbadmin`, created during the Vertica cluster installation.
2. Use a text editor to create a file under the `/home/dbadmin` directory.
The file name can be anything you want, for example, `sg_auth.conf`.
3. If the S3 endpoint is using a standard HTTP port 80 or HTTPS port 443, skip the port number. To use HTTPS, set the following values:

- awsenablehttps = 1, otherwise set the value to 0.
- awsauth = <s3 access key ID>:<secret access key>
- awsendpoint = <StorageGRID s3 endpoint>:<port>

To use a custom CA or self-signed SSL certificate for the StorageGRID S3 endpoint HTTPS connection, specify the full file path and filename of the certificate. This file must be at the same location on each Vertica node and have read permission for all users. Skip this step if StorageGRID S3 Endpoint SSL certificate is signed by publicly known CA.

- awscafile = <filepath/filename>

For example, see the following sample file:

```
awsauth = MNVU4OYFAY2xyz123:03vu04M4KmdfwffT8nqnBmnMVTr78Gu9wANabcxyz
awsendpoint = s3.england.connectlab.io:10443
awsenablehttps = 1
awscafile = /etc/custom-cert/grid.pem
```



In a production environment, the customer should implement a server certificate signed by a publicly known CA on a StorageGRID S3 load balancer endpoint.

Choosing a depot path on all Vertica nodes

Choose or create a directory on each node for the depot storage path.

The directory you supply for the depot storage path parameter must have the following:

- The same path on all nodes in the cluster (for example, /home/dbadmin/depot)
- Be readable and writable by the dbadmin user
- Sufficient storage

By default, Vertica uses 60% of the file system space containing the directory for depot storage. You can limit the size of the depot by using the --depot-size argument in the `create_db` command. See [Sizing Your Vertica Cluster for an Eon Mode Database](#) article for general Vertica sizing guidelines or consult with your Vertica account manager.

The `admintools create_db` tool attempts to create the depot path for you if one does not exist.

Creating the Eon on-premises database

To create the Eon on-premises database, follow these steps:

1. To create the database, use the `admintools create_db` tool.

The following list provides a brief explanation of arguments used in this example. See the Vertica document for a detailed explanation of all required and optional arguments.

- `-x <path/filename of authorization file created in “Creating an authorization file to access the S3 endpoint” >`.

The authorization details are stored inside database after successful creation. You can remove this file to avoid exposing the S3 secret key.

- --communal-storage-location <s3://storagegrid bucketname>
- -s <comma-separated list of Vertica nodes to be used for this database>
- -d <name of database to be created>
- -p <password to be set for this new database>.

For example, see the following sample command:

```
admintools -t create_db -x sg_auth.conf --communal-storage
           -location=s3://vertica --depot-path=/home/dbadmin/depot --shard
           -count=6 -s vertica-vm1,vertica-vm2,vertica-vm3 -d vmart -p
           '<password>'
```

Creating a new database takes several minutes duration depending on number of nodes for the database. When creating database for the first time, you will be prompted to accept the License Agreement.

For example, see the following sample authorization file and `create_db` command:

```
[dbadmin@vertica-vm1 ~]$ cat sg_auth.conf
awsauth = MNVU4OYFAY2CPKVXVxxxx:03vu04M4KmdfwffT8nqnBmnMVTr78Gu9wAN+xxxx
awsendpoint = s3.england.connectlab.io:10445
awsenablehttps = 1

[dbadmin@vertica-vm1 ~]$ admintools -t create_db -x sg_auth.conf
--communal-storage-location=s3://vertica --depot-path=/home/dbadmin/depot
--shard-count=6 -s vertica-vm1,vertica-vm2,vertica-vm3 -d vmart -p
'xxxxxxxx'
Default depot size in use
Distributing changes to cluster.
Creating database vmart
Starting bootstrap node v_vmart_node0007 (10.45.74.19)
Starting nodes:
    v_vmart_node0007 (10.45.74.19)
Starting Vertica on all nodes. Please wait, databases with a large
catalog may take a while to initialize.
Node Status: v_vmart_node0007: (DOWN)
Node Status: v_vmart_node0007: (DOWN)
Node Status: v_vmart_node0007: (DOWN)
Node Status: v_vmart_node0007: (UP)
Creating database nodes
Creating node v_vmart_node0008 (host 10.45.74.29)
Creating node v_vmart_node0009 (host 10.45.74.39)
Generating new configuration information
```

```
Stopping single node db before adding additional nodes.
Database shutdown complete
Starting all nodes
Start hosts = ['10.45.74.19', '10.45.74.29', '10.45.74.39']
Starting nodes:
  v_vmart_node0007 (10.45.74.19)
  v_vmart_node0008 (10.45.74.29)
  v_vmart_node0009 (10.45.74.39)
Starting Vertica on all nodes. Please wait, databases with a large
catalog may take a while to initialize.
  Node Status: v_vmart_node0007: (DOWN) v_vmart_node0008: (DOWN)
v_vmart_node0009: (DOWN)
  Node Status: v_vmart_node0007: (UP) v_vmart_node0008: (UP)
v_vmart_node0009: (UP)
Creating depot locations for 3 nodes
Communal storage detected: rebalancing shards

Waiting for rebalance shards. We will wait for at most 36000 seconds.
Installing AWS package
  Success: package AWS installed
Installing ComplexTypes package
  Success: package ComplexTypes installed
Installing MachineLearning package
  Success: package MachineLearning installed
Installing ParquetExport package
  Success: package ParquetExport installed
Installing VFunctions package
  Success: package VFunctions installed
Installing approximate package
  Success: package approximate installed
Installing flextable package
  Success: package flextable installed
Installing kafka package
  Success: package kafka installed
Installing logsearch package
  Success: package logsearch installed
Installing place package
  Success: package place installed
Installing txtindex package
  Success: package txtindex installed
```

Installing voltagesecure package

Success: package voltagesecure installed

Syncing catalog on vmart with 2000 attempts.

Database creation SQL tasks completed successfully. Database vmart created successfully.

Object size (byte)	Bucket/object key full path
61	s3://vertica/051/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a07/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a07_0_.dfs
145	s3://vertica/2c4/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a3d/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a3d_0_.dfs
146	s3://vertica/33c/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a1d/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a1d_0_.dfs
40	s3://vertica/382/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a31/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a31_0_.dfs
145	s3://vertica/42f/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a21/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a21_0_.dfs
34	s3://vertica/472/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a25/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a25_0_.dfs
41	s3://vertica/476/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a2d/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a2d_0_.dfs
61	s3://vertica/52a/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a5d/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a5d_0_.dfs

Object size (byte)	Bucket/object key full path
131	s3://vertica/5d2/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a19/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a19_0_0.dfs
91	s3://vertica/5f7/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a11/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a11_0_0.dfs
118	s3://vertica/82d/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a15/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a15_0_0.dfs
115	s3://vertica/9a2/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a61/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a61_0_0.dfs
33	s3://vertica/acd/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a29/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a29_0_0.dfs
133	s3://vertica/b98/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a4d/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a4d_0_0.dfs
38	s3://vertica/db3/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a49/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a49_0_0.dfs
38	s3://vertica/eba/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a59/026d63ae9d4a33237bf0e2c2cf2a794a00a000000021a59_0_0.dfs
21521920	s3://vertica/metadata/VMart/Libraries/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000215e2/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000215e2.tar

Object size (byte)	Bucket/object key full path
6865408	s3://vertica/metadata/VMart/Libraries/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021602/026d63ae9d4a33237bf0e2c2cf2a794a00a00000000021602.tar
204217344	s3://vertica/metadata/VMart/Libraries/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021610/026d63ae9d4a33237bf0e2c2cf2a794a00a00000000021610.tar
16109056	s3://vertica/metadata/VMart/Libraries/026d63ae9d4a33237bf0e2c2cf2a794a00a00000000217e0/026d63ae9d4a33237bf0e2c2cf2a794a00a000000000217e0.tar
12853248	s3://vertica/metadata/VMart/Libraries/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021800/026d63ae9d4a33237bf0e2c2cf2a794a00a00000000021800.tar
8937984	s3://vertica/metadata/VMart/Libraries/026d63ae9d4a33237bf0e2c2cf2a794a00a000000002187a/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000002187a.tar
56260608	s3://vertica/metadata/VMart/Libraries/026d63ae9d4a33237bf0e2c2cf2a794a00a00000000218b2/026d63ae9d4a33237bf0e2c2cf2a794a00a000000000218b2.tar
53947904	s3://vertica/metadata/VMart/Libraries/026d63ae9d4a33237bf0e2c2cf2a794a00a00000000219ba/026d63ae9d4a33237bf0e2c2cf2a794a00a000000000219ba.tar
44932608	s3://vertica/metadata/VMart/Libraries/026d63ae9d4a33237bf0e2c2cf2a794a00a00000000219de/026d63ae9d4a33237bf0e2c2cf2a794a00a000000000219de.tar
256306688	s3://vertica/metadata/VMart/Libraries/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021a6e/026d63ae9d4a33237bf0e2c2cf2a794a00a00000000021a6e.tar

Object size (byte)	Bucket/object key full path
8062464	s3://vertica/metadata/VMart/Libraries/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021e34/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021e34.tar
20024832	s3://vertica/metadata/VMart/Libraries/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021e70/026d63ae9d4a33237bf0e2c2cf2a794a00a0000000021e70.tar
10444	s3://vertica/metadata/VMart/cluster_config.json
823266	s3://vertica/metadata/VMart/nodes/v_vmart_node0016/Catalog/859703b06a3456d95d0be28575a673/Checkpoints/c13_13/chkpt_1.cat.gz
254	s3://vertica/metadata/VMart/nodes/v_vmart_node0016/Catalog/859703b06a3456d95d0be28575a673/Checkpoints/c13_13/complete.d
2958	s3://vertica/metadata/VMart/nodes/v_vmart_node0016/Catalog/859703b06a3456d95d0be28575a673/Checkpoints/c2_2/chkpt_1.cat.gz
231	s3://vertica/metadata/VMart/nodes/v_vmart_node0016/Catalog/859703b06a3456d95d0be28575a673/Checkpoints/c2_2/completed
822521	s3://vertica/metadata/VMart/nodes/v_vmart_node0016/Catalog/859703b06a3456d95d0be28575a673/Checkpoints/c4_4/chkpt_1.cat.gz
231	s3://vertica/metadata/VMart/nodes/v_vmart_node0016/Catalog/859703b06a3456d95d0be28575a673/Checkpoints/c4_4/completed
746513	s3://vertica/metadata/VMart/nodes/v_vmart_node0016/Catalog/859703b06a3456d95d0be28575a673/Txnlogs/txn_14_g14.cat

Object size (byte)	Bucket/object key full path
2596	s3://vertica/metadata/VMart/nodes/v_vma rt_node0016/Catalog/859703b06a3456d95d0 be28575a673/Txnlogs/txn_3_g3.cat.gz
821065	s3://vertica/metadata/VMart/nodes/v_vma rt_node0016/Catalog/859703b06a3456d95d0 be28575a673/Txnlogs/txn_4_g4.cat.gz
6440	s3://vertica/metadata/VMart/nodes/v_vma rt_node0016/Catalog/859703b06a3456d95d0 be28575a673/Txnlogs/txn_5_g5.cat
8518	s3://vertica/metadata/VMart/nodes/v_vma rt_node0016/Catalog/859703b06a3456d95d0 be28575a673/Txnlogs/txn_8_g8.cat
0	s3://vertica/metadata/VMart/nodes/v_vma rt_node0016/Catalog/859703b06a3456d95d0 be28575a673/tiered_catalog.cat
822922	s3://vertica/metadata/VMart/nodes/v_vma rt_node0017/Catalog/859703b06a3456d95d0 be28575a673/Checkpoints/c14_7/chkpt_1.cat.gz
232	s3://vertica/metadata/VMart/nodes/v_vma rt_node0017/Catalog/859703b06a3456d95d0 be28575a673/Checkpoints/c14_7/completed
822930	s3://vertica/metadata/VMart/nodes/v_vma rt_node0017/Catalog/859703b06a3456d95d0 be28575a673/Txnlogs/txn_14_g7.cat.gz
755033	s3://vertica/metadata/VMart/nodes/v_vma rt_node0017/Catalog/859703b06a3456d95d0 be28575a673/Txnlogs/txn_15_g8.cat
0	s3://vertica/metadata/VMart/nodes/v_vma rt_node0017/Catalog/859703b06a3456d95d0 be28575a673/tiered_catalog.cat
822922	s3://vertica/metadata/VMart/nodes/v_vma rt_node0018/Catalog/859703b06a3456d95d0 be28575a673/Checkpoints/c14_7/chkpt_1.cat.gz

Object size (byte)	Bucket/object key full path
232	s3://vertica/metadata/VMart/nodes/v_vma_rt_node0018/Catalog/859703b06a3456d95d0be28575a673/Checkpoints/c14_7/completed
822930	s3://vertica/metadata/VMart/nodes/v_vma_rt_node0018/Catalog/859703b06a3456d95d0be28575a673/Txnlogs/txn_14_g7.cat.gz
755033	s3://vertica/metadata/VMart/nodes/v_vma_rt_node0018/Catalog/859703b06a3456d95d0be28575a673/Txnlogs/txn_15_g8.cat
0	s3://vertica/metadata/VMart/nodes/v_vma_rt_node0018/Catalog/859703b06a3456d95d0be28575a673/tiered_catalog.cat

Disabling streaming limitations

This procedure is based on the [Vertica guide for other on-premises object storage](#) and should be applicable to StorageGRID.

1. After creating the database, disable the `AWSStreamingConnectionPercentage` configuration parameter by setting it to 0.
This setting is unnecessary for an Eon Mode on-premises installation with communal storage. This configuration parameter controls the number of connections to the object store that Vertica uses for streaming reads. In a cloud environment, this setting helps avoid having streaming data from the object store use up all the available file handles. It leaves some file handles available for other object store operations. Due to the low latency of on-premises object stores, this option is unnecessary.
2. Use a `vsq1` statement to update the parameter value.
The password is the database password that you set in “[Creating the Eon on-premises database](#)”.
For example, see the following sample output:

```
[dbadmin@vertica-vm1 ~]$ vsq1
Password:
Welcome to vsq1, the Vertica Analytic Database interactive terminal.
Type:  \h or \? for help with vsq1 commands
      \g or terminate with semicolon to execute query
      \q to quit
dbadmin=> ALTER DATABASE DEFAULT SET PARAMETER
AWSStreamingConnectionPercentage = 0; ALTER DATABASE
dbadmin=> \q
```

Verifying depot settings

Vertica database default depot settings are enabled (value = 1) for read and write operations. NetApp strongly

recommends keeping these depot settings enabled to enhance performance.

```
vsq1 -c 'show current all;' | grep -i UseDepot
DATABASE | UseDepotForReads | 1
DATABASE | UseDepotForWrites | 1
```

Loading sample data (optional)

If this database is for testing and will be removed, you can load sample data to this database for testing. Vertica comes with sample dataset, VMart, found under `/opt/vertica/examples/VMart_Schema/` on each Vertica node.

You can find more information about this sample dataset [here](#).

Follow these steps to load the sample data:

1. Log in as dbadmin to one of the Vertica nodes: `cd /opt/vertica/examples/VMart_Schema/`
2. Load sample data to the database and enter the database password when prompted in substeps c and d:
 - a. `cd /opt/vertica/examples/VMart_Schema`
 - b. `./vmart_gen`
 - c. `vsq1 < vmart_define_schema.sql`
 - d. `vsq1 < vmart_load_data.sql`
3. There are multiple predefined SQL queries, you can run some of them to confirm test data are loaded successfully into the database.
For example: `vsq1 < vmart_queries1.sql`

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 11.7 Product Documentation](#)
- [StorageGRID data sheet](#)
- [Vertica 10.1 Product Documentation](#)

Version history

Version	Date	Document version history
Version 1.0	September 2021	Initial release.

By Angela Cheng

StorageGRID log analytics using ELK stack

By Angela Cheng

With the StorageGRID syslog forward feature, you can configure an external syslog server to collect and analyze StorageGRID log messages. ELK (Elasticsearch, Logstash, Kibana) has become one of the most popular log analytics solutions. Watch the [StorageGRID log analysis using ELK video](#) to view a sample ELK configuration and how it can be used to identify and troubleshoot failed S3 requests.

StorageGRID 11.9 supports exporting load balancer endpoint access log to external syslog server. Watch this [Youtube Video](#) to learn more about this new feature.

This article provides sample files of Logstash configuration, Kibana queries, charts and dashboard to give you a quick start for StorageGRID log management and analytics.

Requirements

- StorageGRID 11.6.0.2 or higher
- ELK (Elasticsearch, Logstash and Kibana) 7.1x or higher installed and in operation

Sample files

- [Download the Logstash 7.x sample files package](#)
md5 checksum 148c23d0021d9a4bb4a6c028746deab
sha256 checksum f51ec9e2e3f842d5a7861566b167a561beb4373038b4e7bb3c8be3d522adf2d6
- [Download the Logstash 8.x sample files package](#)
md5 checksum e11bae3a662f87c310ef363d0fe06835
sha256 checksum 5c670755742cfdf5aa723a596ba087e0153a65bcaef3934afdb682f61cd278d
- [Download the Logstash 8.x sample files package for StorageGRID 11.9](#)
md5 checksum 41272857c4a54600f95995f6ed74800d
sha256 checksum 67048ee8661052719990851e1ad960d4902fe537a6e135e8600177188da677c9

Assumption

Readers are familiar with StorageGRID and ELK terminology and operations.

Instruction

Two sample versions are provided due to differences in names defined by grok patterns.

For example, the SYSLOGBASE grok pattern in Logstash config file defines field names differently depending on the installed Logstash version.

```
match => { "message" => '<%{POSINT:syslog_pri}>%{SYSLOGBASE}  
%{GREEDYDATA:msg-details}' }
```

Logstash 7.17 sample

Field	Value
<code>t _id</code>	7C1MaYEVRH8UbfKnIls8
<code>t _index</code>	sgrid2-2022.06.15
<code># _score</code>	-
<code>t _type</code>	_doc
<code>file @timestamp</code>	Jun 15, 2022 @ 17:36:46.038
<code>t host</code>	grid2-site2-s1
<code>t logsource</code>	SITE2-S1
<code>t msg-details</code>	Reloading syslog service
<code>t pid</code>	628
<code>t program</code>	update-sysl
<code>t syslog_pri</code>	37
<code>t timestamp</code>	Jun 15 21:36:46

Logstash 8.23 sample

Table	JSON
<input type="text"/> Search field names	
Actions	Field
...	<code>t _id</code>
...	yuh0iIEBVP6KX4EwqcyU
...	<code>t _index</code>
...	sgrid-2022.06.21
...	<code># _score</code>
...	-
...	<code>file @timestamp</code>
...	Jun 21, 2022 @ 18:07:45.444
...	<code>t event.original</code>
...	<28>Jun 21 22:07:45 SITE2-S3 ADE: syslog messages being dropped
...	<code>t host.hostname</code>
...	SITE2-S3
...	<code>t msg-details</code>
...	syslog messages being dropped
...	<code>t process.name</code>
...	ADE
...	<code>t syslog_pri</code>
...	28
...	<code>t timestamp</code>
...	Jun 21 22:07:45

Steps

1. Unzip the provided sample based on your installed ELK version.
The sample folder includes two Logstash config samples:
sglog-2-file.conf: this config file outputs StorageGRID log messages to a file on Logstash without data transformation. You can use this to confirm Logstash is receiving StorageGRID messages or to help understand StorageGRID log patterns.
sglog-2-es.conf: this config file transforms StorageGRID log messages using various pattern and filters. It includes example drop statements, which drop messages based on patterns or filter. The output is sent to Elasticsearch for indexing.
Customize the selected config file according to the instruction inside the file.
2. Test the customized config file:

```
/usr/share/logstash/bin/logstash --config.test_and_exit -f <config-file-path/file>
```

If the last line returned is similar to the below line, the config file has no syntax errors:

```
[Logstash::Runner] runner - Using config.test_and_exit mode. Config Validation Result: OK. Exiting Logstash
```

3. Copy the customized conf file to the Logstash server's config: /etc/logstash/conf.d
If you have not enabled config.reload.automatic in /etc/logstash/logstash.yml, restart the Logstash service.
Otherwise, wait for the config reload interval to elapse.

```
grep reload /etc/logstash/logstash.yml
# Periodically check if the configuration has changed and reload the pipeline
config.reload.automatic: true
config.reload.interval: 5s
```

4. Check /var/log/logstash/logstash-plain.log and confirm there are no errors starting Logstash with the new config file.
5. Confirm TCP port is started and listening.
In this example, TCP port 5000 is used.

```
netstat -ntpa | grep 5000
tcp6      0      0 :::5000          :::*
LISTEN    25744/java
```

6. From the StorageGRID manager GUI, configure external syslog server to send log messages to Logstash.
Refer to the [demo video](#) for details.
7. You need to configure or disable firewall on the Logstash server to allow StorageGRID nodes connection to the defined TCP port.

8. From Kibana GUI, select Management → Dev Tools. On the Console page, run this GET command to confirm new indices are created on Elasticsearch.

```
GET /_cat/indices/*?v=true&s=index
```

9. From Kibana GUI, create index pattern (ELK 7.x) or data view (ELK 8.x).

10. From Kibana GUI, enter 'saved objects' in the search box which is located in the top center. On the Saved Objects page, select Import. Under Import options, select 'Request action on conflict'

Import saved objects

Select a file to import



Import

Import options



Check for existing objects



Automatically overwrite conflicts

Request action on conflict

Create new objects with random IDs



Import elk<version>-query-chart-sample.ndjson.

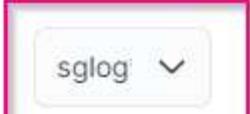
When prompted to resolve the conflict, select the index pattern or data view you created in step 8.

Import saved objects

X

⌚ Data Views Conflicts

The following saved objects use data views that do not exist. Please select the data views you'd like re-associated with them. You can [create a new data view](#) if necessary.

ID	Count	Sample of aff...	New data view
594f91a0-d192-11ec-b30f-09f67aedd1d9	2		
60cf3620-e5fa-11ec-af71-8f6e980d6eb0	1		

The following Kibana objects are imported:

Query

- * audit-msg-s3rq-orlm
- * bycast log s3 related messages
- * loglevel warning or above
- * failed security event
- * nginx-gw endpoint access log (available only in elk8-sample-for-sg119.zip)

Chart

- * s3 requests count based on bycast.log
- * HTTP status code
- * audit msg breakdown by type
- * average s3 response time

Dashboard

- * S3 request dashboard using the above charts.

You are now ready to perform StorageGRID log analysis using Kibana.

Additional resources

- [syslog101](#)

- [What is the ELK stack](#)
- [Grok patterns list](#)
- [A beginner's guide to Logstash: Grok](#)
- [A practical guide to Logstash: syslog deep dive](#)
- [Kibana guide – Explore the document](#)
- [StorageGRID audit log messages reference](#)

Use Prometheus and Grafana to extend your metrics retention

By Aron Klein

This technical report provides detailed instructions for configuring NetApp StorageGRID with external Prometheus and Grafana services.

Introduction

StorageGRID stores metrics using Prometheus and provides visualizations of these metrics through built in Grafana dashboards. The Prometheus metrics can be accessed securely from StorageGRID by configuring client access certificates and enabling prometheus access for the specified client. Today, the retention of this metric data is limited by the storage capacity of the administration node. To gain a longer duration and an ability to create customized visualizations of these metrics we will deploy a new Prometheus and Grafana server, configure our new server to scrape the metrics from StorageGRIDs instance, and build a dashboard with the metrics that are important to us. You can get more information on the Prometheus metrics collected in the [StorageGRID documentation](#).

Federate Prometheus

Lab details

For the purposes of this example, I will be using all virtual machines for StorageGRID 11.6 nodes, and a Debian 11 server. The StorageGRID management interface is configured with a publicly trusted CA certificate. This example will not go through the installation and configuration of the StorageGRID system or Debian linux installation. You can use any Linux flavor you wish that is supported by Prometheus and Grafana. Both Prometheus and Grafana can install as docker containers, build from source, or pre-compiled binaries. In this example I will be installing both Prometheus and Grafana binaries directly on the same Debian server. Download and follow the basic installation instructions from <https://prometheus.io> and <https://grafana.com/grafana> respectively.

Configure StorageGRID for Prometheus Client access

In order to gain access to StorageGRIDs stored prometheus metrics you must generate or upload a client certificate with private key, and enable permission for the client. The StorageGRID management interface must have an SSL certificate. This certificate must be trusted by the prometheus server either by a trusted CA, or manually trusted if it is self-signed. To read more, please visit the [StorageGRID documentation](#).

1. In the StorageGRID management interface, select "CONFIGURATION" on the bottom left hand side, and in the second column under "Security" click on Certificates.
2. On the Certificates page select the "Client" tab and click on the "Add" button.

3. Provide a name for the client that will be granted access and use this certificate. Click on the box under "Permissions", in front of "Allow Prometheus" and click the Continue button.

The screenshot shows a web-based configuration interface for adding a client certificate. At the top, a blue header bar displays the title "Add a client certificate" and a progress indicator showing "1 Enter details" and "2 Enter details". The main content area is titled "Certificate details". It contains a "Certificate name" field with the value "prometheus" and a "Permissions" section with a checked checkbox labeled "Allow prometheus".

4. If you have a CA signed certificate you can select the radio button for "Upload certificate", but in our case we are going to let storageGRID generate the client certificate by selecting the radio button for "Generate Certificate". The required fields will be displayed to be filled in. Enter the FQDN for the client server, the IP of the server, the subject, and Days valid. Then click the "Generate" button.

Add a client certificate

1 Enter details 2 Enter details

Certificate type

Upload certificate Generate certificate

Domain name [?](#)
prometheus.grid.local

[Add another domain](#)

IP [?](#)
192.168.0.10

[Add another IP address](#)

Subject [?](#)
/CN=Prometheus

Days valid [?](#)
730

[Generate](#)

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



Be mindful of the certificate days valid entry as you will need to renew this certificate in both StorageGRID and the Prometheus server before it expires to maintain uninterrupted collection.

1. Download the certificate pem file, and the private key pem file.

Generate

Certificate details

[Download certificate](#) [Copy certificate PEM](#)

Subject DN: /CN=Prometheus
Serial Number: 72:D9:6E:D7:04:CC:4F:29:66:0A:CA:53:24:79:1B:09:49:3A:BC:56
Issuer DN: /CN=Prometheus
Issued On: 2022-08-22T17:54:33.000Z
Expires On: 2024-08-21T17:54:33.000Z
SHA-1 Fingerprint: 10:47:6E:FD:67:D8:53:E7:6E:E5:D8:8A:DF:BD:45:94:04:53:47:1E
SHA-256 Fingerprint: 74:23:C2:02:3A:D9:08:C0:EE:C1:F8:59:8A:7C:AE:18:AB:B0:7D:21:31:F3:EB:AF:BF:4F:9E:C7:90:C9:FA:E7
Alternative Names: DNS:prometheus.grid.local
IP Address: 192.168.0.10

Certificate private key ⓘ

⚠ You will not be able to view the certificate private key after you close this dialog. To save the keys for future reference, copy and paste the values to another location.

[Download private key](#) [Copy private key](#)

```
-----BEGIN RSA PRIVATE KEY-----  
MIIEpaIBAAKCAQEa3bIcyIEpMWPk5ritVpMkmIDKLIjaTM3ertq23VcAALwxziaU  
-----END RSA PRIVATE KEY-----
```



This is the only time you can download the private key, so make sure you do not skip this step.

Prepare the Linux server for Prometheus installation

Before installing Prometheus, I want to get my environment prepared with a Prometheus user, the directory structure, and configure the capacity for the metrics storage location.

1. Create the Prometheus user.

```
sudo useradd -M -r -s /bin/false Prometheus
```

2. Create the directories for Prometheus, client certificate, and metrics data.

```
sudo mkdir /etc/Prometheus /etc/Prometheus/cert /var/lib/Prometheus
```

3. I formatted the disk I am using for metrics retention with an ext4 filesystem.

```
mkfs -t ext4 /dev/sdb
```

4. I then mounted the filesystem to the Prometheus metrics directory.

```
sudo mount -t auto /dev/sdb /var/lib/prometheus/
```

5. Obtain the uuid of the disk you are using for your metrics data.

```
sudo ls -al /dev/disk/by-uuid/  
lrwxrwxrwx 1 root root 9 Aug 18 17:02 9af2c5a3-bfc2-4ec1-85d9-  
ebab850bb4a1 -> ../../sdb
```

6. Adding an entry in /etc/fstab/ making the mount persist across reboots using the uuid of /dev/sdb.

```
/etc/fstab  
UUID=9af2c5a3-bfc2-4ec1-85d9-ebab850bb4a1 /var/lib/prometheus ext4  
defaults 0 0
```

Install and configure Prometheus

Now that the server is ready, I can begin the Prometheus installation and configure the service.

1. Extract the Prometheus installation package

```
tar xzf prometheus-2.38.0.linux-amd64.tar.gz
```

2. Copy the binaries to /usr/local/bin and change the ownership to the prometheus user created earlier

```
sudo cp prometheus-2.38.0.linux-amd64/{prometheus,promtool}  
/usr/local/bin  
sudo chown prometheus:prometheus /usr/local/bin/{prometheus,promtool}
```

3. Copy the consoles and libraries to /etc/prometheus

```
sudo cp -r prometheus-2.38.0.linux-amd64/{consoles,console_libraries}  
/etc/prometheus/
```

4. Copy the client certificate and private key pem files downloaded earlier from StorageGRID to /etc/prometheus/certs

5. Create the prometheus configuration yaml file

```
sudo nano /etc/prometheus/prometheus.yml
```

6. Insert the following configuration. The job name can be anything you wish. Change the "-targets: []" to the

FQDN of the admin node, and if you altered the names of the certificate and private key file names, please update the `tls_config` section to match. then save the file. If your grid management interface, is using a self-signed certificate, download the certificate and place it with the client certificate with a unique name, and in the `tls_config` section add `ca_file: /etc/prometheus/cert/UICert.pem`

- In this example I am collecting all of the metrics that begin with alertmanager, cassandra, node, and storagegrid. You can see more information on the Prometheus metrics in the [StorageGRID documentation](#).

```
# my global config
global:
  scrape_interval: 60s # Set the scrape interval to every 15 seconds.
  Default is every 1 minute.

scrape_configs:
  - job_name: 'StorageGRID'
    honor_labels: true
    scheme: https
    metrics_path: /federate
    scrape_interval: 60s
    scrape_timeout: 30s
    tls_config:
      cert_file: /etc/prometheus/cert/certificate.pem
      key_file: /etc/prometheus/cert/private_key.pem
    params:
      match[]:
        -
        ' {__name__ =~ "alertmanager_.*|cassandra_.*|node_.*|storagegrid_.*" } '
        static_configs:
        - targets: ['sgdemo-rtp.netapp.com:9091']
```

If your grid management interface is using a self-signed certificate, download the certificate and place it with the client certificate with a unique name. In the `tls_config` section add the certificate above the client certificate and private key lines

i

```
ca_file: /etc/prometheus/cert/UICert.pem
```

- Change the ownership of all files and directories in `/etc/prometheus`, and `/var/lib/prometheus` to the `prometheus` user

```
sudo chown -R prometheus:prometheus /etc/prometheus/
sudo chown -R prometheus:prometheus /var/lib/prometheus/
```

- Create a prometheus service file in `/etc/systemd/system`

```
sudo nano /etc/systemd/system/prometheus.service
```

3. Insert the following lines, note the `--storage.tsdb.retention.time=1y` which sets the retention of the metric data to 1 year. Alternatively, you could use `--storage.tsdb.retention.size=300GiB` to base retention on storage limits. This is the only location to set the metrics retention.

```
[Unit]
Description=Prometheus Time Series Collection and Processing Server
Wants=network-online.target
After=network-online.target

[Service]
User=prometheus
Group=prometheus
Type=simple
ExecStart=/usr/local/bin/prometheus \
    --config.file /etc/prometheus/prometheus.yml \
    --storage.tsdb.path /var/lib/prometheus/ \
    --storage.tsdb.retention.time=1y \
    --web.console.templates=/etc/prometheus/consoles \
    --web.console.libraries=/etc/prometheus/console_libraries

[Install]
WantedBy=multi-user.target
```

4. Reload the systemd service to register the new prometheus service. then start and enable the prometheus service.

```
sudo systemctl daemon-reload
sudo systemctl start prometheus
sudo systemctl enable prometheus
```

5. Check the service is running properly

```
sudo systemctl status prometheus
```

- prometheus.service - Prometheus Time Series Collection and Processing Server


```
Loaded: loaded (/etc/systemd/system/prometheus.service; enabled; vendor preset: enabled)
      Active: active (running) since Mon 2022-08-22 15:14:24 EDT; 2s ago
        Main PID: 6498 (prometheus)
          Tasks: 13 (limit: 28818)
        Memory: 107.7M
          CPU: 1.143s
        CGroup: /system.slice/prometheus.service
                  └─6498 /usr/local/bin/prometheus --config.file
/etc/prometheus/prometheus.yml --storage.tsdb.path /var/lib/prometheus/
--web.console.templates=/etc/prometheus/consoles --web.con>
```

Aug 22 15:14:24 aj-deb-prom01 prometheus[6498]: ts=2022-08-22T19:14:24.510Z caller=head.go:544 level=info component=tsdb msg="Replaying WAL, this may take a while"
Aug 22 15:14:24 aj-deb-prom01 prometheus[6498]: ts=2022-08-22T19:14:24.816Z caller=head.go:615 level=info component=tsdb msg="WAL segment loaded" segment=0 maxSegment=1
Aug 22 15:14:24 aj-deb-prom01 prometheus[6498]: ts=2022-08-22T19:14:24.816Z caller=head.go:615 level=info component=tsdb msg="WAL segment loaded" segment=1 maxSegment=1
Aug 22 15:14:24 aj-deb-prom01 prometheus[6498]: ts=2022-08-22T19:14:24.816Z caller=head.go:621 level=info component=tsdb msg="WAL replay completed" checkpoint_replay_duration=55.57µs wal_rep>
Aug 22 15:14:24 aj-deb-prom01 prometheus[6498]: ts=2022-08-22T19:14:24.831Z caller=main.go:997 level=info fs_type=EXT4_SUPER_MAGIC
Aug 22 15:14:24 aj-deb-prom01 prometheus[6498]: ts=2022-08-22T19:14:24.831Z caller=main.go:1000 level=info msg="TSDB started"
Aug 22 15:14:24 aj-deb-prom01 prometheus[6498]: ts=2022-08-22T19:14:24.831Z caller=main.go:1181 level=info msg="Loading configuration file" filename=/etc/prometheus/prometheus.yml
Aug 22 15:14:24 aj-deb-prom01 prometheus[6498]: ts=2022-08-22T19:14:24.832Z caller=main.go:1218 level=info msg="Completed loading of configuration file" filename=/etc/prometheus/prometheus.y>
Aug 22 15:14:24 aj-deb-prom01 prometheus[6498]: ts=2022-08-22T19:14:24.832Z caller=main.go:961 level=info msg="Server is ready to receive web requests."
Aug 22 15:14:24 aj-deb-prom01 prometheus[6498]: ts=2022-08-22T19:14:24.832Z caller=manager.go:941 level=info component="rule manager" msg="Starting rule manager..."

6. You should now be able to browse to the UI of your prometheus server <http://Prometheus-server:9090> and see the UI

7. Under "Status" Targets you can see the status of the StorageGRID endpoint we configured in prometheus.yml

Endpoint	State	Labels	Last Scrape	Scrape Duration	Error
https://sgdemo-rtp.netapp.com:9091/federate match[]="(_name_ =~ \"alertmanager_ cassandra_ node_ storagegrid_\")"	UP	instance="sgdemo-rtp.netapp.com:9091" job="StorageGRID"	43.396s ago	280.876ms	

8. On the Graph page, you can execute a test query and verify the data is successfully being scraped. for example enter "storagegrid_node_cpu_utilization_percentage" into the query bar and click the Execute button.

storagegrid_node_cpu_utilization_percentage	3.4062500000005547	2023-09-27T10:00:00Z
storagegrid_node_cpu_utilization_percentage{instance="TD-SG-Adm01", job="node", node_id="fc1f00fc-d148-42b6-b9c4-72b34c2cd0c3", site_id="a3d223fd-cc25-4255-8987-771e7246ad35", site_name="Tera01"} 3.4062500000005547	3.4062500000005547	2023-09-27T10:00:00Z
storagegrid_node_cpu_utilization_percentage{instance="TD-SG-GW01", job="node", node_id="97b62a35-c5f0-4cccd-a1f8-24e6ddfc770b", site_id="a3d223fd-cc25-4255-8987-771e7246ad35", site_name="Tera01"} 3.264583333336901	3.264583333336901	2023-09-27T10:00:00Z
storagegrid_node_cpu_utilization_percentage{instance="TD-SG-SN01", job="node", node_id="17ba14f4-59fc-44fd-a0cc-96d2525c31db", site_id="a3d223fd-cc25-4255-8987-771e7246ad35", site_name="Tera01"} 12.991666666666641	12.991666666666641	2023-09-27T10:00:00Z
storagegrid_node_cpu_utilization_percentage{instance="TD-SG-SN02", job="node", node_id="b4343f55-16fd-4471-993c-1cd749867718", site_id="a3d223fd-cc25-4255-8987-771e7246ad35", site_name="Tera01"} 14.618749999999494	14.618749999999494	2023-09-27T10:00:00Z
storagegrid_node_cpu_utilization_percentage{instance="TD-SG-SN03", job="node", node_id="77313bb8-0300-45af-b748-98cd128dd39d", site_id="a3d223fd-cc25-4255-8987-771e7246ad35", site_name="Tera01"} 10.620833333423812	10.620833333423812	2023-09-27T10:00:00Z

Install and configure Grafana

Now that prometheus is installed and working, we can move on to installing Grafana and configuring a dashboard

Grafana Instalation

1. Install the latest enterprise edition of Grafana

```
sudo apt-get install -y apt-transport-https
sudo apt-get install -y software-properties-common wget
sudo wget -q -O /usr/share/keyrings/grafana.key
https://packages.grafana.com/gpg.key
```

2. Add this repository for stable releases:

```
echo "deb [signed-by=/usr/share/keyrings/grafana.key]
https://packages.grafana.com/enterprise/deb stable main" | sudo tee -a
/etc/apt/sources.list.d/grafana.list
```

3. After you add the repository.

```
sudo apt-get update
sudo apt-get install grafana-enterprise
```

4. Reload the systemd service to register the new grafana service. then start and enable the Grafana service.

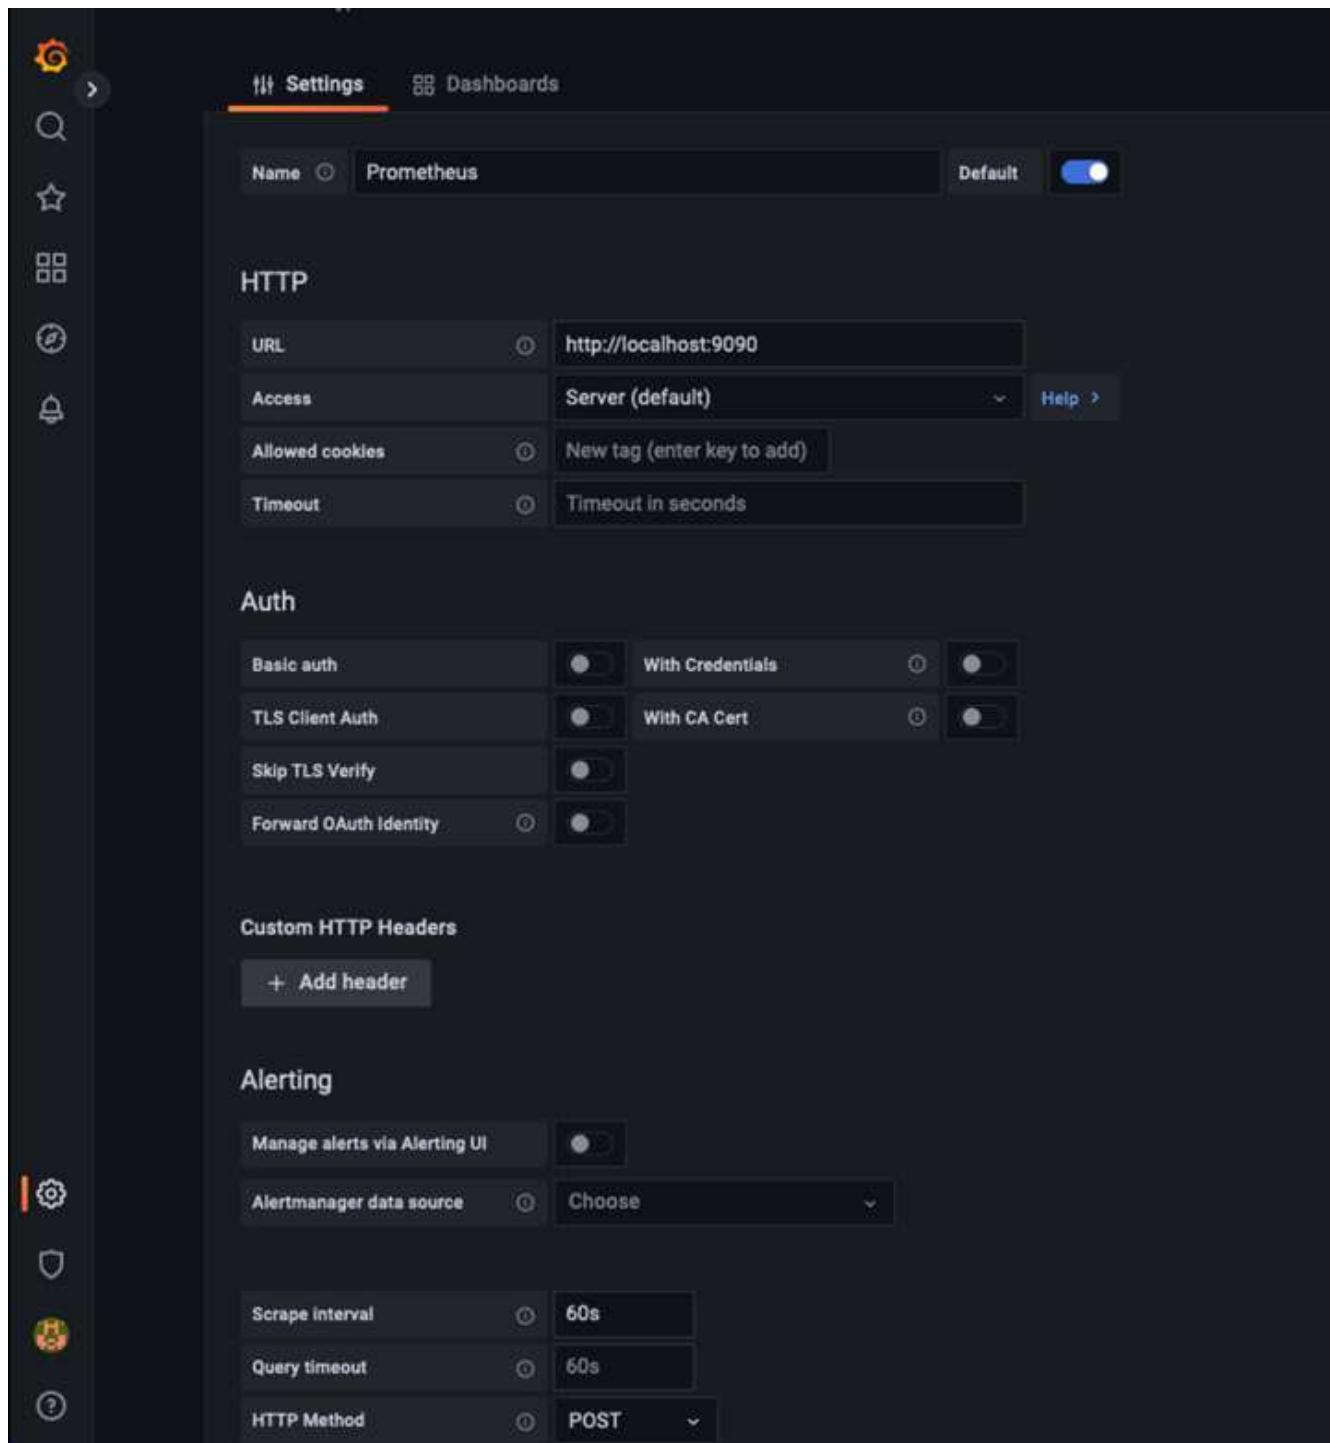
```
sudo systemctl daemon-reload
sudo systemctl start grafana-server
sudo systemctl enable grafana-server.service
```

5. Grafana is now installed and running. When you open a browser to `HTTP://Prometheus-server:3000` you will be greeted with the Grafana login page.
6. The default login credentials are `admin/admin`, and you should set a new password as it prompts you to.

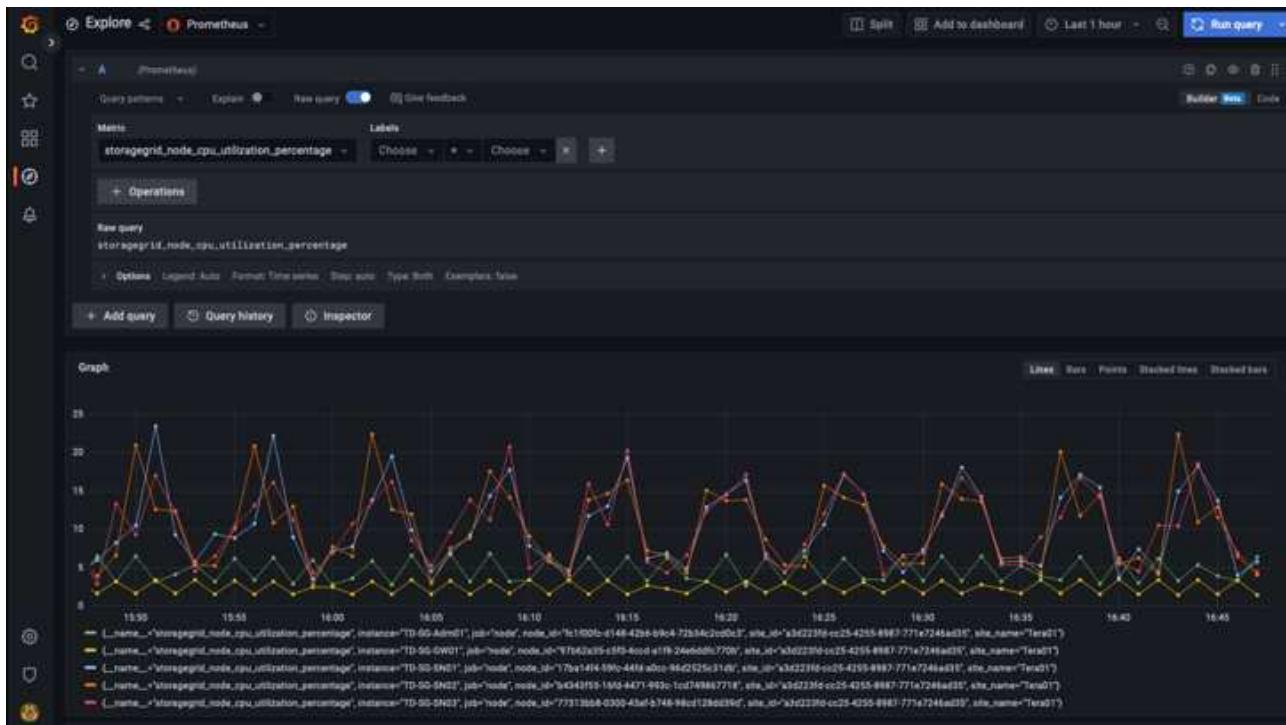
Create a Grafana dashboard for StorageGRID

With Grafana and Prometheus installed and running, now its time to connect the two by creating a data source and build a dashboard

1. On the left hand pane expand "Configuration" and select "Data sources", then click on the "Add Data source" button
2. Prometheus will be one of the top data sources to choose from. If it is not, then use the search bar to locate "Prometheus"
3. Configure the Prometheus source by entering the URL of the prometheus instance, and the scrape interval to match the Prometheus interval. I also disabled the alerting section as I did not configure the alert manager on prometheus.

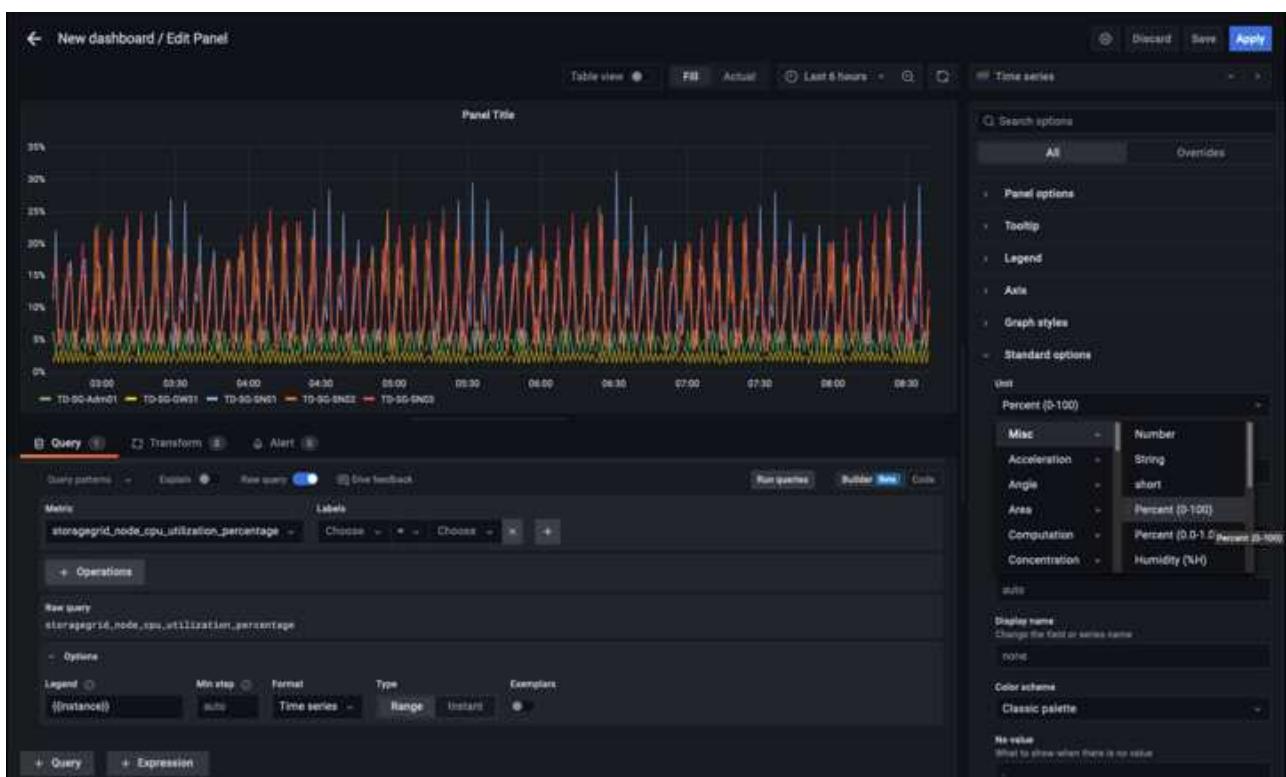


4. With the desired settings entered, scroll down to the bottom and click on "Save & test"
5. After the configuration test is successful, click on the explore button.
 - a. In the explore window you can use the same metric we tested Prometheus with "storagegrid_node_cpu_utilization_percentage", and click the "Run query" button



6. Now that we have the data source configured, we can create a dashboard.

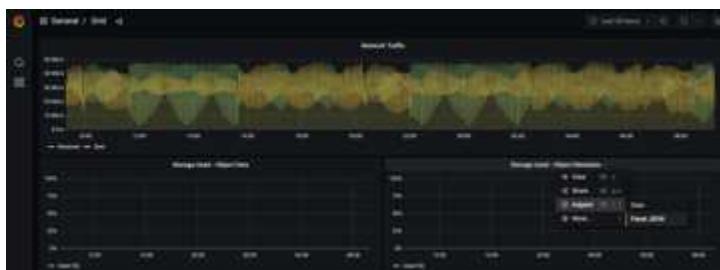
- On the left hand pane expand "Dashboards", and select "+ new Dashboard"
- Select "Add a new panel"
- Configure the new panel by selecting a metric, again I will use "storagegrid_node_cpu_utilization_percentage", Enter a title for the panel, expand "Options" at the bottom and for legend change to custom and enter "{{instance}}" to define the node names", and on the right pane under "Standard options" set "Unit" to "Misc/Percent(0-100)". Then click "Apply" to save the panel to the dashboard.



7. We could continue to build out our dashboard like this for each metric we want, but luckily StorageGRID already has dashboards with panels we can copy into our custom dashboards.

- From the StorageGRID management interface left hand pane, select "Support", and at the bottom of the "Tools" column click on "Metrics".
- Within metrics, I am going to select the "Grid" link on the top of the middle column.

- From the Grid dashboard, lets select the "Storage Used - Object Metadata" panel. Click the little down arrow and the end of the panel title to drop down a menu. From this menu select "Inspect" and "Panel JSON".



- Copy out the JSON code and close the window.

Inspect: Storage Used - Object Metadata

4 queries with total query time of 549 ms

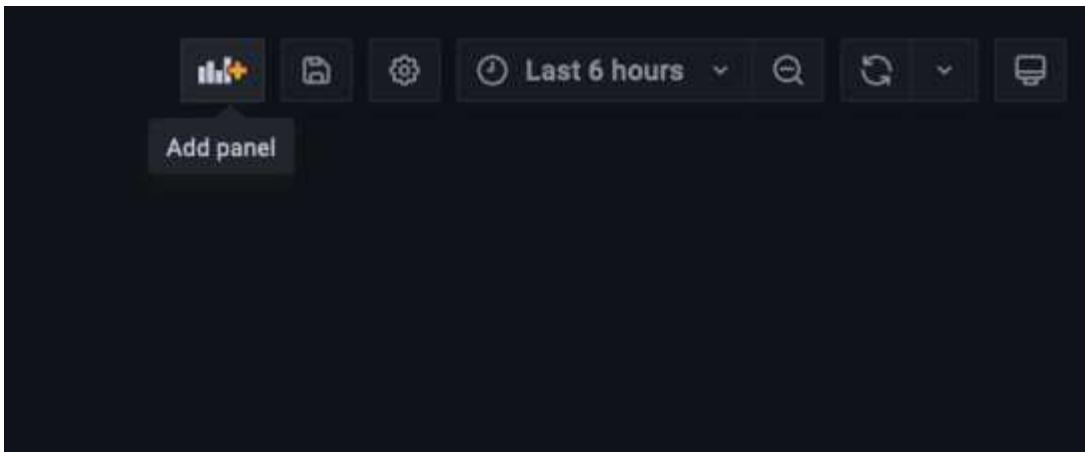
Data Stats JSON

Select source

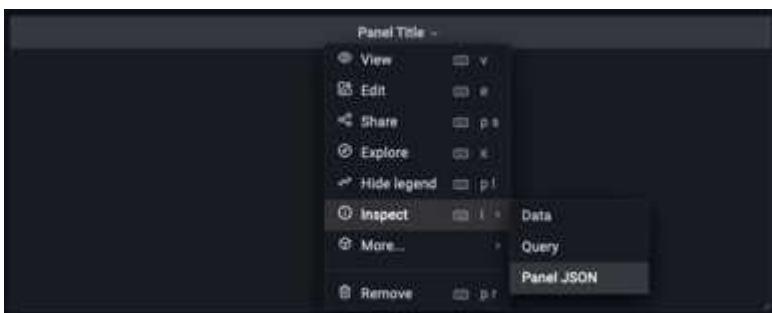
Panel JSON

```
1  {
2      "aliasColors": {},
3      "bars": false,
4      "dashLength": 10,
5      "dashes": false,
6      "datasource": "Prometheus",
7      "decimals": 2,
8      "fill": 1,
9      "fillGradient": 0,
10     "gridPos": {
11         "h": 7,
12         "w": 12,
13         "x": 12,
14         "y": 7
15     },
16     "id": 6,
17     "legend": {
18         "avg": false,
19         "current": false,
20         "max": false,
21         "min": false,
22         "show": true,
23         "total": false,
24         "values": false
25     },
26     "lines": true,
27     "linewidth": 1,
28     "links": [],
29     "nullPointMode": "null",
30     "options": {
31         "alertThreshold": true
32     },
33     "percentage": false,
34     "pointradius": 5,
35     "points": false,
36     "renderer": "flat",
37     "seriesOverrides": [
38         {
39             "alias": "Used",
```

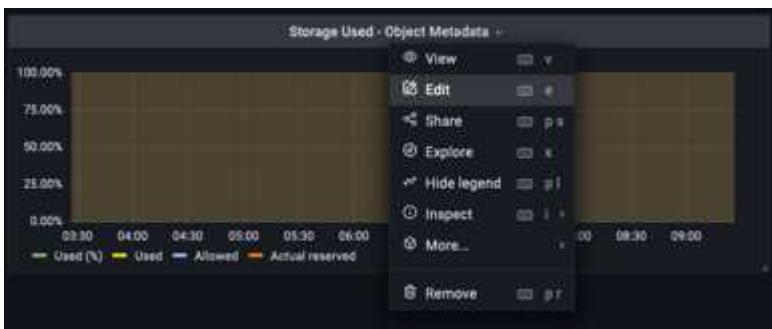
e. In our new dashboard, click on the icon to add a new panel.

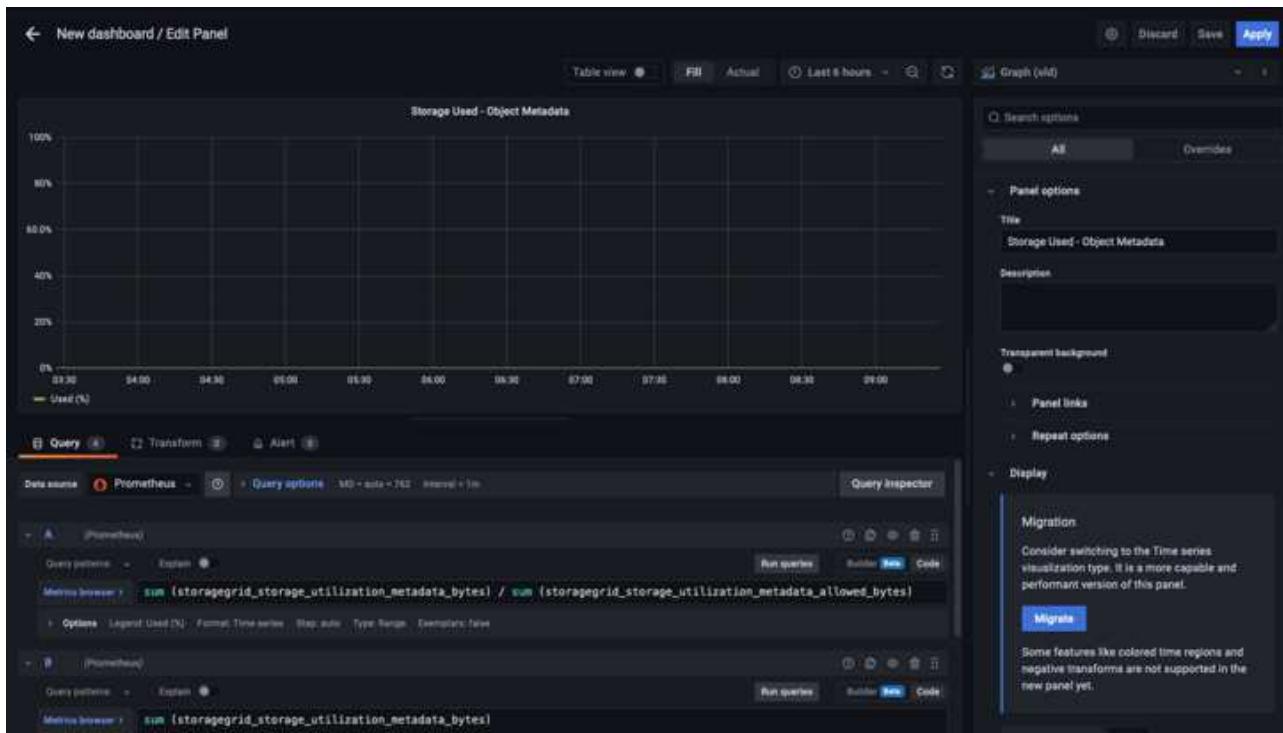


- f. Apply the new panel without making any changes
- g. Just like with the StorageGRID panel, inspect the JSON. Remove all JSON code and replace it with the copied code from the StorageGRID panel.



- h. Edit the new panel, and on the right hand side you will see a Migration message with a "Migrate" button. Click the button and then click the "Apply" button.





- Once you have all of the panels in place and configured as you like. Save the dashboard by clicking the disk icon in the upper right and give your dashboard a name.

Conclusion

Now we have a Prometheus server with customizable data retention and storage capacity. With this we can continue build out our own dashboards with the metrics that are most relevant to our operations. You can get more information on the Prometheus metrics collected in the [StorageGRID documentation](#).

Use F5 DNS to globally load balance StorageGRID

By Steve Gorman (F5)

This technical report provides detailed instructions for configuring NetApp StorageGRID with F5 DNS services for global load balancing to deliver better data availability, greater data consistency, and optimize S3 transaction routing when your grid is distributed across multiple sites and/or HA groups.

Introduction

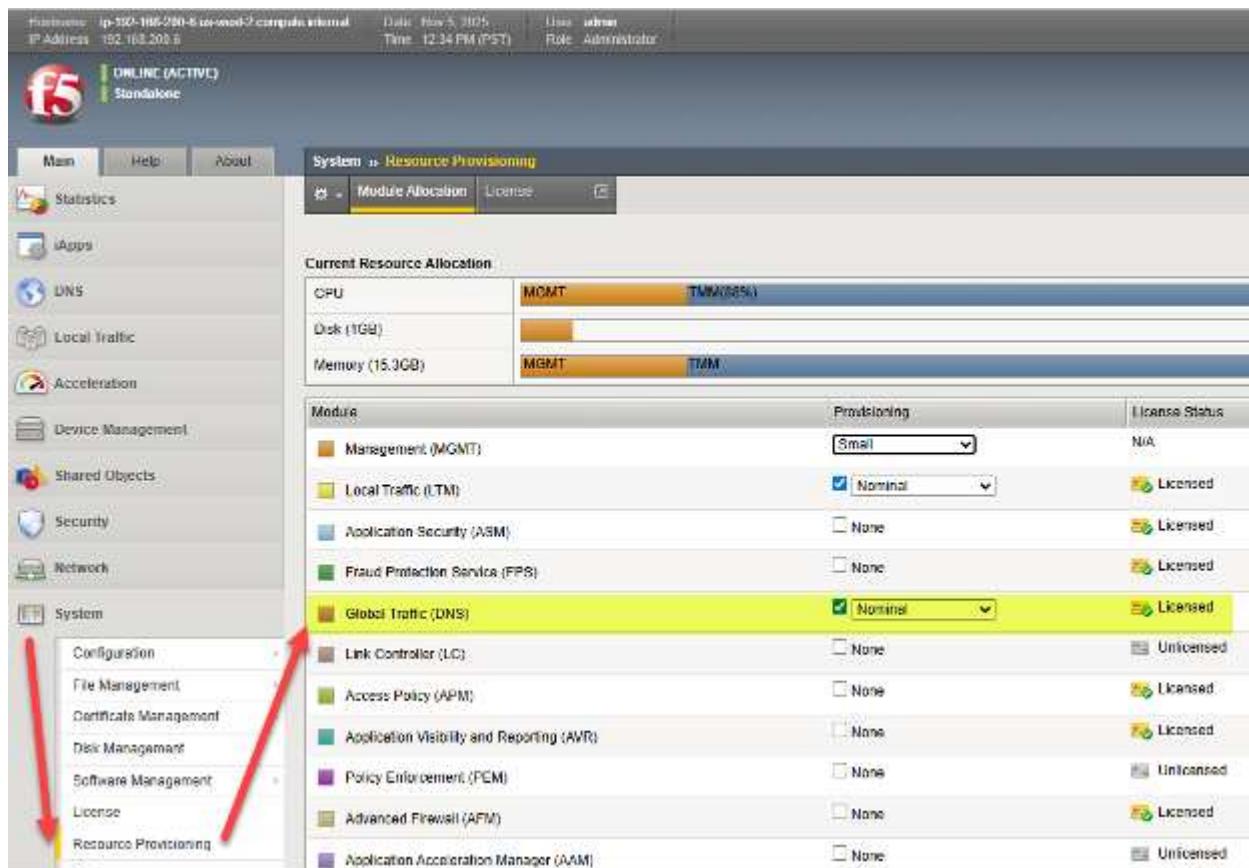
The F5 BIG-IP DNS solution formerly called BIG-IP GTM (Global Traffic Manager) and informally GSLB (Global Server Load Balancing) allows for seamless access across multiple active-active HA groups and active-active multi-site StorageGRID solutions to effectively be realized.

F5 BIG-IP multi-site StorageGRID configuration

Regardless of the number of StorageGRID sites to be supported, a minimum of two BIG-IP appliances, physical or virtual, must have the BIG-IP DNS module enabled and setup. The more DNS appliances, the further the degree of redundancy an enterprise will benefit from.

BIG-IP DNS - First Steps in Initial Set Up

Once the BIG-IP appliance has undergone at least initial provisioning, use a web browser to log into the TMUI (BIG-IP GUI) interface, and choose System → Resource Provisioning. As highlighted, ensure that the “Global Traffic (DNS)” module has a check mark and is shown to be licensed. Note, as in the image, it is common that “Local Traffic (LTM)” can be provisioned on the same appliance.

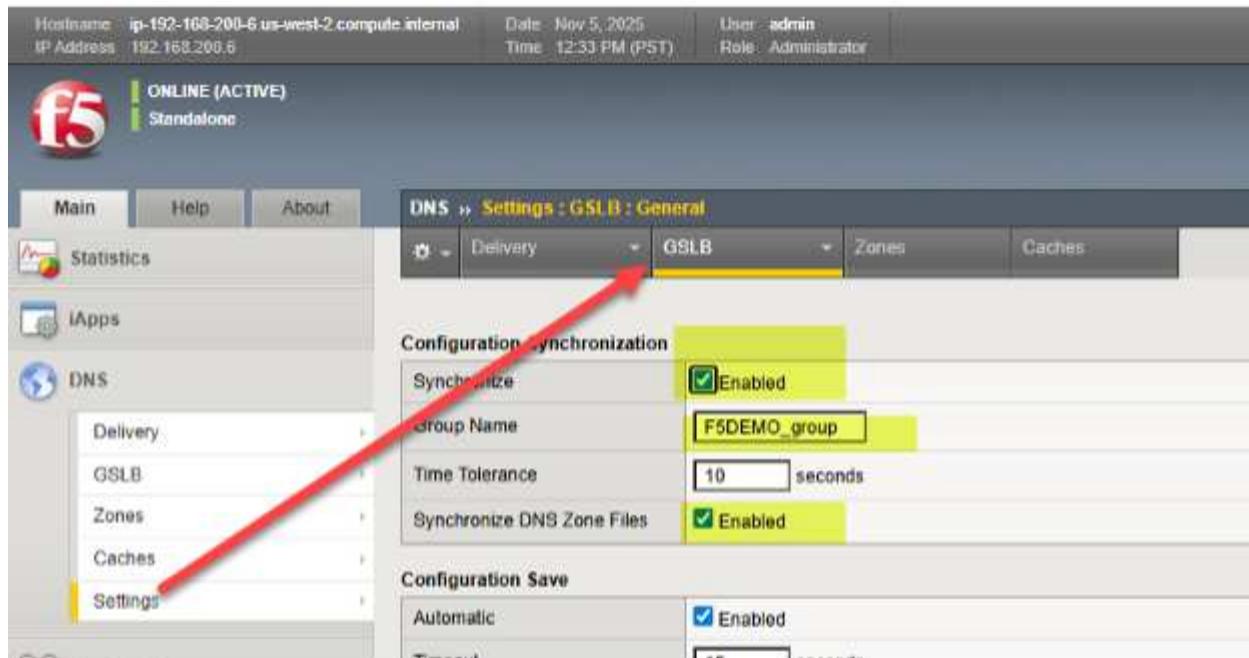


The screenshot shows the BIG-IP TMUI interface. The top bar displays the IP address (192.168.200.6), date (Nov 5, 2025), and time (12:34 PM (PST)). The user is logged in as 'admin' with the role 'Administrator'. The main navigation bar includes links for Main, Help, About, and System (with Resource Provisioning selected). The left sidebar lists various system management options. The 'Resource Provisioning' section is highlighted with a red arrow. The main content area shows the 'Current Resource Allocation' with CPU, Disk (1TB), and Memory (15.0GB) usage. Below this is a table of modules and their provisioning status. The 'Global Traffic (DNS)' module is highlighted with a yellow background and a red arrow, showing it is licensed. Other modules listed include Management (MGMT), Local Traffic (LTM), Application Security (ASM), Fraud Protection Service (FPS), Link Controller (LC), Access Policy (APM), Application Visibility and Reporting (AVR), Policy Enforcement (PEM), Advanced Firewall (AFM), and Application Acceleration Manager (AAM). The 'Global Traffic (DNS)' row shows 'Nominal' provisioning and a 'Licensed' status with a green icon.

Module	Provisioning	License Status
Management (MGMT)	Small	N/A
Local Traffic (LTM)	<input checked="" type="checkbox"/> Nominal	Licensed
Application Security (ASM)	<input type="checkbox"/> None	Licensed
Fraud Protection Service (FPS)	<input type="checkbox"/> None	Licensed
Global Traffic (DNS)	<input checked="" type="checkbox"/> Nominal	Licensed
Link Controller (LC)	<input type="checkbox"/> None	Unlicensed
Access Policy (APM)	<input type="checkbox"/> None	Licensed
Application Visibility and Reporting (AVR)	<input type="checkbox"/> None	Licensed
Policy Enforcement (PEM)	<input type="checkbox"/> None	Unlicensed
Advanced Firewall (AFM)	<input type="checkbox"/> None	Licensed
Application Acceleration Manager (AAM)	<input type="checkbox"/> None	Unlicensed

Configure DNS Protocol Foundational Elements

The first step towards global traffic management for StorageGRID sites is to choose the DNS tab, where virtually all the global traffic steering will be configured, and choose Settings→ GLSB. Enable the two synchronization options and choose a DNS group name that will be shared among participating BIG-IP appliances.



Next, navigate to DNS > Delivery > Profiles > DNS: Create and create a profile that will govern the DNS capabilities you wish to enable or disable. See the previous link for the DNS classroom guide if generation of specific DNS logs are of interest. Here is an example of a working DNS Profile, note the four highlights that represent settings that are important values. For awareness, each possible setting is explained at the following F5 KB (Knowledge Base) article [here](#).

iApps

DNS

- Delivery
- GSLB
- Zones
- Caches
- Settings

Local Traffic

Acceleration

Device Management

Shared Objects

Security

Network

System

General Properties

Name	f5demo.net_dns_profile
Partition / Path	Common
Parent Profile	dns

Denial of Service Protection

Rapid Response Mode	Disabled
Rapid Response Last Action	Drop

Hardware Acceleration

Protocol Validation	Disabled
Response Cache	Disabled

DNS Features

DNSSEC	Disabled
GSLB	Enabled
DNS Express	Disabled
DNS Cache	Disabled
DNS Cache Name	Select...
DNS IPv6 to IPv4	Disabled
Unhandled Query Actions	Drop
Use BIND Server on BIG-IP	Disabled
Insert Source Address into Client Subnet Option	Disabled

DNS Traffic

Zone Transfer	Disabled
DNS Security	Disabled
DNS Security Profile Name	Select...
Process Recursion Desired	Enabled

Logging and Reporting

Logging	Enabled
Logging Profile	f5demo_dns_logging_profile
AVR Statistics Sample Rate	<input type="checkbox"/>

Buttons

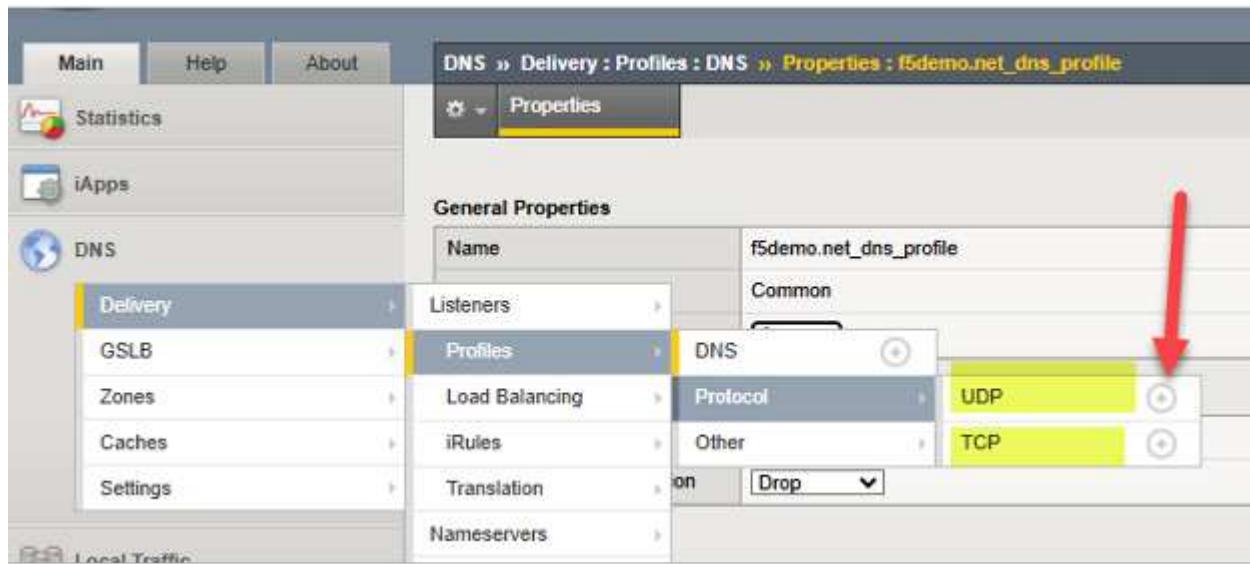
Update **Delete...**

At this point we can adjust the characteristics of the UDP and TCP protocols, through created “profiles”, which can both carry DNS traffic involving BIG-IP. Simply create one new profile for UDP and TCP. Presuming DNS traffic will cross WAN links, a good practice is simply to inherit UDP and TCP characteristics known to perform well in WAN environments.

To add each simply click the “+” icon beside each protocol, and set the parent profile to the following:

UDP → use “parent” profile “udp_gtm_dns”

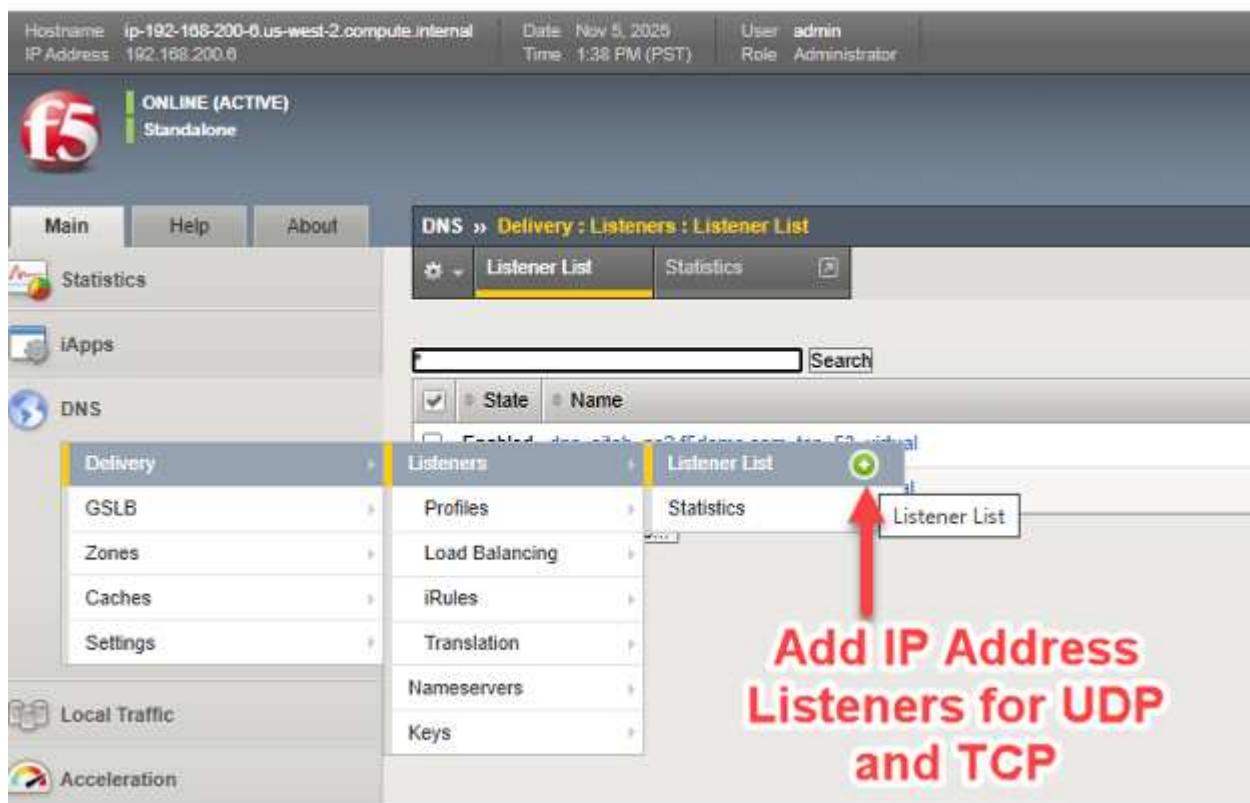
TCP → use “parent” profile “f5-tcp-wan”



The screenshot shows the 'Properties' tab for the 'f5demo.net_dns_profile'. The 'Protocol' dropdown is set to 'TCP'. A red arrow points to this dropdown.

Now, we simply need to assign an IP address for both UDP and TCP traffic involving the BIG-IP DNS. For those familiar with BIG-IP LTM, this is essentially the creation of DNS virtual servers, and virtual servers need “listener” IP addresses.

As in the screenshot, follow the arrows to create listener/virtual servers for DNS/UDP and DNS/TCP.



The screenshot shows the 'Listener List' tab. A red arrow points to the 'Listener List' button. A large red text overlay in the bottom right corner reads: 'Add IP Address Listeners for UDP and TCP'.

The following is one example from a live BIG-IP DNS, in it we see the TCP virtual server listener settings and can see how it ties together many of the previous steps. This includes referencing the DNS profile and protocol (TCP) profile, as well as configuring a valid IP address for DNS to use. As with all the objects one creates with BIG-IP, it is helpful to use a meaningful name which serves to self-identify what the object is, such as dns/siteb/TCP53 in the example name assigned.

This concludes the preliminary, typically “one time”, setup steps of a BIG-IP appliance with the DNS module enabled. At this point we are ready to transition to the specifics of setting up a global traffic management solution with our appliances, which will of course be linked to the characteristics of the StorageGRID sites.

Setting up Data Center Sites and Establishing Inter BIG-IP Communications in Four Steps

Step one: Create Data Centers

Each site which will house clusters of nodes to be locally load balanced by BIG-IP LTM, should be entered into BIG-IP DNS. This needs to be done on only one BIG-IP DNS, as we are creating a DNS synchronized group to support traffic management, as such this configuration will be shared among DNS members of the group.

Through the TMUI GUI, select DNS > GSLB > Data Centers > Data Center List and create an entry for each of the StorageGRID sites. If using a network setup aligned with Figure 1, DNS appliance located in other non-StorageGRID sites, add Data Centers for these sites in addition to storage sites. In this example sites a and b are created in Ohio and Oregon, the BIG-IPs are dual DNS and LTM appliances.

Detailed description of the screenshot: The screenshot shows the F5 BIG-IP Configuration interface. On the left, there's a navigation menu with options like 'Statistics', 'Apps', 'DNS', 'Delivery', 'GSLB', 'Zones', 'Caches', 'Settings', 'Local Traffic', 'Acceleration', 'Device Management', 'Shared Objects', 'Security', 'Network', and 'System'. The 'DNS' menu is selected. In the center, there's a 'Data Center List' table with two entries: 'west_datacenter' and 'east_datacenter'. A red arrow points from the 'west_datacenter' entry to the 'Properties' tab of its configuration dialog on the right. The 'Properties' tab has tabs for 'General Properties', 'Probes', 'Probes', 'Statistics', and 'Logs'. The 'General Properties' tab is active, showing fields for 'Name' (west_datacenter), 'Partition / Path' (Common), 'Description' (West Data Center), 'Location' (Customer Data), 'Contact' (None), 'Prober Preference' (Inside Data Center), 'Prober Interval' (Any Available), 'Availability' (Available (Enabled) - Available), and 'State' (Enabled). Another red arrow points from the 'Properties' tab to the 'Availability' section of the 'General Properties' table.

Step two: Create Servers (List of All BIG-IP Appliances in Solution)

We are now ready to connect the individual StorageGRID site clusters to the BIG-IP DNS setup. Recall, the BIG-IP appliance at each site will do the actual load balancing of S3 traffic, through the configuration of virtual servers that tie a “front-end” reachable IP address/port to a set of back end “pool” of Storage Node appliances, using “back-end” IP addresses/ports.

Should, as one example, all Storage Nodes in a pool be taken offline administratively, perhaps for a site decommissioning, or unexpectedly through real-time failed health checks, traffic will be directed to other sites through altering DNS query responses.

To tie the StorageGrid sites, specifically the local virtual servers, into the BIG-IP DNS configuration on each appliance, the setup need only be done once. The entire group of BIG-IP DNS appliances will have their setups synchronized, in an upcoming step.

In simple terms, we will create a list, referred to as a server list, of all of our BIG-IP appliances, whether licensed for DNS, LTM or both DNS and LTM. This master list will be sync'd with all BIG-IP DNS appliances upon completion of the list.

On one BIG-IP DNS licensed appliance, choose DNS > GSLB > Servers > Server List and choose the add button (+).

ONLINE (ACTIVE)
Standalone

Main Help About

DNS » GSLB: Servers : Server List

Server List Trusted Server Certificates

Statistics iApps DNS

Delivery GSLB Zones Caches Settings

Local Traffic Acceleration Device Management Shared Objects

Wide IPs Pools iRules Data Centers

Servers Server List

Links Prober Pools Monitors Topology Distributed Applications

Reconnect Reconnect All

SiteBBIGIP 192.168.1.10

Click to Add Server Entry

The four key elements when adding each BIG-IP include:

- * Selecting BIG-IP from the product pull down, other load balancers are possible but generally lack the real time visibility responsiveness when backend node health deteriorates at each site.
- * Add the IP address of the BIG-IP DNS appliance. likely, the first time adding a BIG-IP DNS appliance, the address will be the current GUI-accessed appliance, future appliances will be the other appliances in the solution.
- * Choose a health monitor, always use “BIG-IP” when the load balancer being added is BIG-IP appliance, for back-end StorageGRID node health consideration.
- * Optionally, request virtual server automatic discovery if the appliance is a dual DNS/LTM appliance.

In some situations, such as transient networking problems or firewall ACL rules between network locations, when adding a remote appliance at this stage, the virtual server discovery may not show entries for remote appliances with LTM configured. In such cases, after adding the new appliance ("server"), one can manually add the virtual servers as indicated below. If adding a BIG-IP DNS-only appliance, there will be no virtual servers to be discovered or added to that device.

We need to add these server entries for each appliance in our solution at all sites, including BIG-IP DNS appliances, BIG-IP LTM appliances, and any appliances serving the dual roles of both DNS and LTM units.

Step three: Establish Trust Between all BIG-IP Appliances

In the following example, four appliances have been added as servers, they are spread across two sites. Note each site has a dedicated BIG-IP DNS and BIG-IP LTM. However, all of the appliances, other than the one currently logged into, are showing blue icons in the "Status" column. This means a trust relationship has not yet been established with the other BIG-IP appliances.

Hostname: dns.sitea.15demo.com Date: Oct 9, 2023
 IP Address: 10.1.1.6 Time: 10:47 PM (CEST) User: admin
 Role: Administrator Partition: /

f5 ONLINE (ACTIVE)
 Standalone

Main Help About DNS » GSLB : Servers : Server List

Statistics iApps DNS

Delivery GSLB Zones Caches

Server List Trusted Server Certificates Statistics

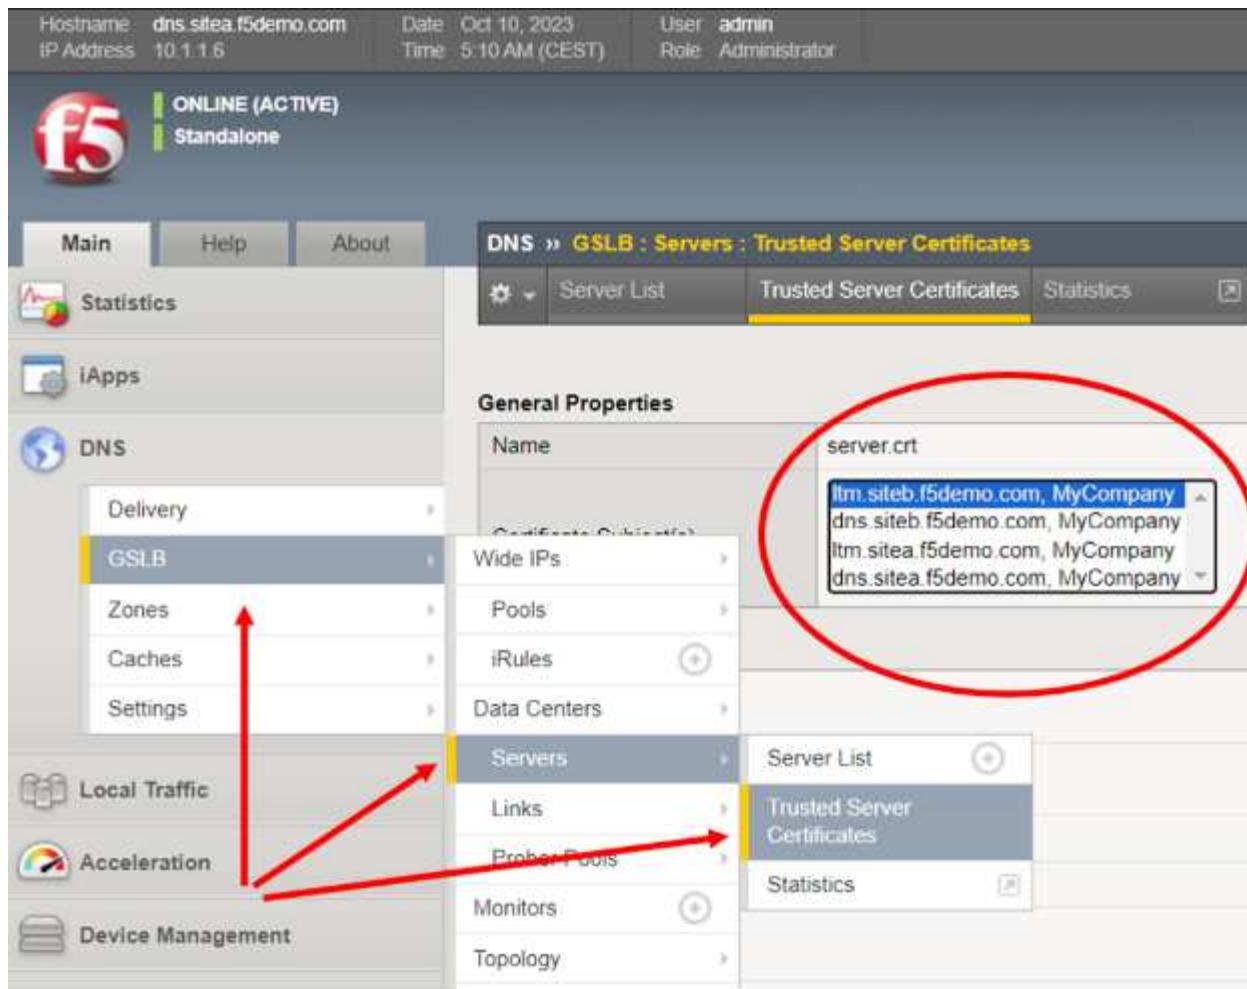
<input checked="" type="checkbox"/>	Status	Name	Devices	Address	Data Center	Virtual Server
<input type="checkbox"/>	Green	dns.sitea_server	1	10.1.10.100	sitea_datacenter	0
<input type="checkbox"/>	Blue	dns.siteb_server	1	10.1.60.100	siteb_datacenter	0
<input type="checkbox"/>	Blue	sitea_llm	1	10.1.10.50	sitea_datacenter	1
<input type="checkbox"/>	Blue	siteb_llm	1	10.1.60.50	siteb_datacenter	1

Search | Enable | Disable | Delete... | Reconnect | Reconnect All

To add trust, SSH into the BIG-IP where the configuration details have just been entered via the GUI, use the "root" account to access the BIG-IP command line interface.

Issue the following single command at the prompt: `bigip_add`

The "bigip_add" command pulls the management certificate from the destination BIGIP devices for use during the encrypted "iQuery" channel setup between GSLB servers in the cluster. iQuery, by default, runs using TCP port 4353 and is the heartbeat that allows BIG-IP DNS members to stay in a synchronized state. It makes use of xml and gzip in the encrypted channel. When running "bigip_add" without any options, the command will be run against all BIGIP devices in the GSLB Server list using the current username to connect to the endpoints. As a quick check of success, simply return to the BIG-IP GUI and confirm all servers now have certificates listed in the displayed pull-down menu.



Step four: Synchronize all BIG-IP DNS Appliances to the DNS Group

The final step will allow all BIG-IP DNS appliances to be fully configured by simply using the TMUI GUI of one single unit. In a sample case, where there are two StorageGRID sites, this means now using SSH to reach the command line of the **other** site's BIG-IP DNS.

After connecting as root, and ensuring that firewall policies/ACLs allow the two BIG-IP DNS devices to talk on TCP ports 22 (SSH), 443 (HTTPS) and 4354 (F5 iQuery protocol), issue this one command at the prompt:
`gtm_add <IP address of first site BIG-IP DNS, where all of the GUI steps were previously done>`

At this point all further DNS configuration work can be performed on any BIG-IP DNS appliance that has been added to the group. The above command, `gtm_add`, need not be applied on appliance members that are LTM only. Only appliances supporting DNS require this command to become part of the synchronized DNS group.

Setting up Data Center Sites and Establishing Inter BIG-IP Communications

At this point, all the steps to create the underlying, healthy BIG-IP DNS appliance group is complete. We can now get on with creating names, FQDNs, that point towards our distributed web/S3 services exposed at each StorageGRID datacenter.

These names are referred to as "Wide IPs", or WIPs for short, and they are normal DNS FQDNs with DNS A resource records. However, rather than pointing at a server like a traditional A resource record, they internally point at pools of BIG-IP virtual servers. Each pool, individually, can be made up of a set of one or more virtual servers. An S3 client requesting an IP address to name resolution will receive the address of the S3 virtual server at the optimal, policy-selected StorageGRID site.

Wide IPs, Pools and Virtual Servers in a Nutshell

To give a simple, fictitious example, a WIP for the name **storage.quantumvault.com** might see the BIG-IP DNS solution linked with two pools of potential virtual servers. The first pool might be made up of 4 sites in North America; the second pool might consist of 3 sites in Europe.

The selected pool might be arrived at from a range of policy decisions, perhaps a simple ratio of 5:1 could be used to have the bulk of traffic directed to North American StorageGRID sites. More probable perhaps, a topology-based choice where the pool is chosen where, for instance, all European sourced S3 traffic is directed to European sites, and the remainder of world S3 traffic is directed to North American data centers.

Once a pool is arrived at by BIG-IP DNS, let us assume the North American pool was selected, the actual DNS A Resource Record returned to resolve storage.quantumvault.com can be any one of the 4 virtual servers supported by BIG-IP LTM in any of the 4 North American sites. Again, which is chosen is policy driven, simple “static” approaches such as Round-Robin exist, whereas more advanced “dynamic” selections such as performance probes to measure each sites latency from local DNS resolvers is maintained and used as the criteria for site selection.

To set a Pool of virtual servers on a BIG-IP DNS, follow the menu path **DNS > GSLB > Pools > Pool List > Add (+)**.

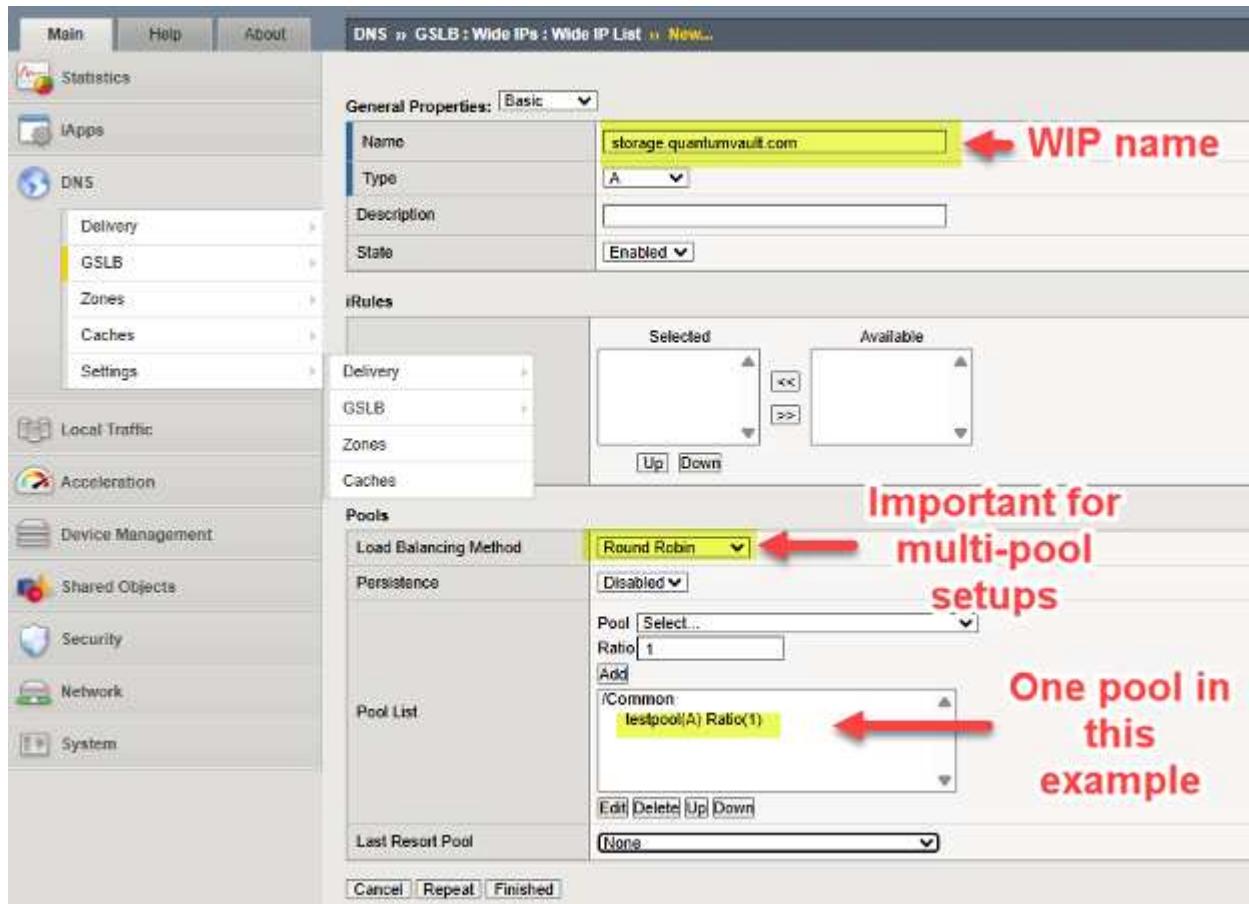
In this example, we can see various North American virtual servers are added to a pool and the preferred approach to load balancing, when this pool is selected, is chosen in a tiered fashion.

Load Balancing Method of Virtual Servers in the Pool

Add Appropriate Virtual Servers

We add the WIP (Wide IP), the name of our service which will be resolved by DNS, to a deployment by following the **DNS > GSLB > Wide IPs > Wide IP List > Create (+)**. In the following example, we provide an

example WIP for an S3-enabled storage service.



Adjust DNS to Support Global Traffic Management

At this point all of our underlying BIG-IP appliances are ready to perform GSLB (global server load balancing). We simply need to adjust and assign the names used for S3 traffic flows to leverage the solution. The general approach is to delegate part of an existing DNS domain of an enterprise to the control of BIG-IP DNS. This is to say to “carve” a section of the name space, a sub-domain, and delegate control of this sub-domain to the BIG-IP DNS appliances. Technically, this is done by ensuring the BIG-IP DNS appliances have A DNS resource records (RRs) in the enterprise DNS and then making these names/addresses Name Server (NS) DNS resource records for the delegated domain.

There are various ways enterprises maintain DNS today, one method is a fully hosted solution. An example of this would be operating and managing DNS through Windows Server 2025. An alternative approach can be for an enterprise to leverage cloud-DNS providers like AWS Route53 or Squarespace.

Here is a fictitious example for illustration purposes. We have StorageGRID supporting object reads and writes via the S3 protocol with an existing domain managed by AWS Route53, the existing example domain is f5demo.net.

We would like to assign the sub-domain engineering.f5demo.net to the BIG-IP DNS appliances for global traffic management. To do this, we create a new NS (name server) resource record for engineering.f5demo.net and point that to the list of BIG-IP DNS appliance names. In our example, we have two BIG-IP DNS appliances, and as such we create two A resource records for them.

The sub-domain for our S3 StorageGRID-targetted traffic

Our BIG-IP DNS Appliance Addresses

We now, as an example, will set up a Wide IP (WIP) in our BIG-IP DNS, since DNS uses group synchronization, we only need to adjust using the GUI of one appliance. Within the BIG-IP DNS GUI, go to **DNS > GSLB > Wide IPs > Wide IP List (+)**. Recall, in a traditional DNS FQDN setup one would be entering one or more IPv4 addresses, in our case we simply point at one or more pools of StorageGRID virtual servers.

The screenshot shows the F5 BIG-IP Configuration interface. The top navigation bar includes 'Not secure' and the URL 'https://44.250.237.170:5443/iui/'. The main menu on the left lists 'Technology', 'Analytics', 'File', 'Related', 'File Related Links', 'File Support', 'File Help', 'File Specs', 'File S...', 'Local Links', 'login and get to w...', 'Link Providers (Obj...)', and 'Browse objects (Obj...)'.

The main content area shows a 'Pools' configuration page. The URL in the address bar is 'https://44.250.237.170:5443/iui/Local/Link/MyLinks?view=pool&selected=www&selectedTab=A'. The 'Pool' tab is selected. A red box highlights the 'Pool' tab in the top navigation bar, and a red arrow points from this box down to the 'Pool' tab in the main content area.

The 'Pool' configuration section includes:

- 'Local Resource Method' dropdown: 'Load Balancing'.
- 'Resistence' dropdown: 'Disabled'.
- 'Local Head Pool' dropdown: 'None'.
- 'Update' button.

The 'Pools' list shows one entry:

- Pool: 'www' (highlighted with a green circle).
- Address: 'www.vip.ingrediente.it'.
- Status: 'Not pool'.

Below the list are 'Create' and 'Remove' buttons.

In our example, we have generic web HTTPS servers located in both Ohio and Oregon sites. With a simple “round robin” approach, we should be able to see the global DNS respond to queries for the A resource record mappings for www.wip.engineering.f5demo.net with both virtual server IPs.

A simple test can be done with web browsers or, in the case of S3 using StorageGRID, perhaps graphical tools like S3Browser. Each DNS query will see the next data center site in the pool used as the target for ensuing traffic, due to our choice of Round Robin within the pool.

In our example setup, we can use dig or nslookup to quickly generate a series of two DNS queries and ensure BIG-IP DNS is indeed doing a round robin load balancing, resulting in both sites receiving traffic over time.

```
C:\Users\gorman>nslookup www.wip.engineering.f5demo.net
Server: dns.google
Address: 8.8.8.8

Non-authoritative answer:
DNS request timed out.
    timeout was 2 seconds.
Name: www.wip.engineering.f5demo.net
Address: 44.250.237.170
```

First Query


```
C:\Users\gorman>nslookup www.wip.engineering.f5demo.net
Server: dns.google
Address: 8.8.8.8

Non-authoritative answer:
DNS request timed out.
    timeout was 2 seconds.
Name: www.wip.engineering.f5demo.net
Address: 3.145.176.246
```

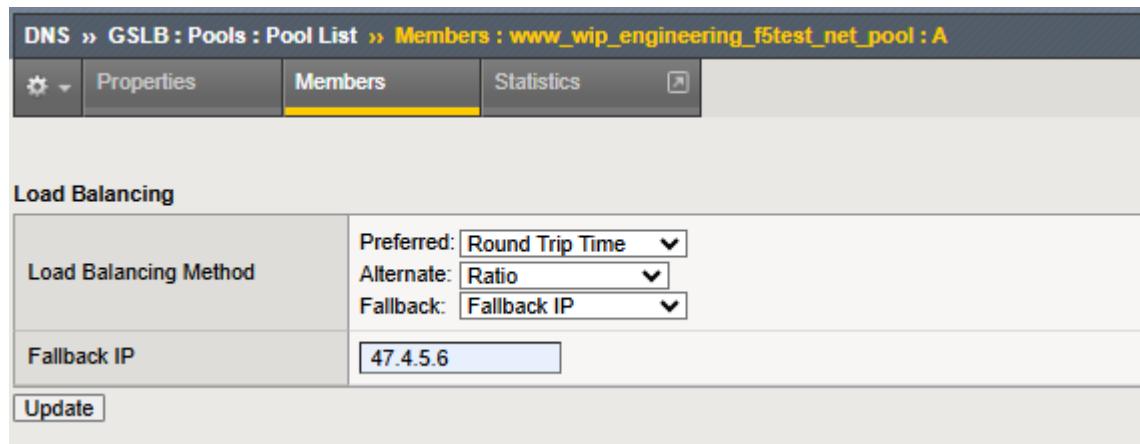
Second Query

Suggested exploration for more advanced techniques

One of many possible approaches, would use “Global Availability” mode as opposed to the simple “Round Robin” example given above. With Global Availability the sequenced order of pools, or virtual servers within just a single pool, can have traffic directed to it. In this way, all S3 traffic could by default, be directed towards, say, a New York City site.

If health checks indicate an issue with StorageGRID node availability at this site, traffic could at that point be directed to St. Louis. Should St. Louis encounter health concerns, a site in Frankfurt could in turn begin to receive S3 read or write transactions. Thus, global availability is one approach to S3 StorageGRID overall solution resiliency.

Another approach is to mix and match load balancing approaches, where a tiered approach is used.



DNS » GSLB : Pools : Pool List » Members : www_wip_engineering_f5test_net_pool : A

Properties Members Statistics

Load Balancing

Load Balancing Method	Preferred: Round Trip Time Alternate: Ratio Fallback: Fallback IP
Fallback IP	47.4.5.6

Update

In this example, a “dynamic” option is the first load balancing choice for the sites in the configured pool. In the example shown, an on-going measurement approach using active probing of local DNS resolver performance is maintained and the catalyst for site selection. Should this approach be unavailable, the individual sites can be selected by the ratio assigned to each. With ratio, larger, higher-bandwidth StorageGRID sites can receive more S3 transactions than smaller sites.

Finally, as perhaps a disaster recovery scenario, should all sites in the pool become unhealthy, the specified fallback IP is used as the site of last resort.

One of the more interesting load balancing methods of BIG-IP DNS is “Topology” whereby the incoming source of DNS queries, the S3 user’s local DNS resolver, is observed and using Internet topology information the seemingly “closest” site is selected from the pool.

Lastly, if sites span the entire globe, it may be worth considering using the dynamic “probe” technology discussed in detail in the F5 BIG-IP DNS manual. With probes, frequent sources of DNS queries can be monitored, take for example a business-to-business partner whose traffic generally uses the same local DNS resolver. BIG-IP DNS probes can be launched from the BIG-IP LTM in each site around the globe, to determine generally which potential site would likely offer the lowest latency for S3 transactions. As such, traffic from Asia might be better served by Asian StorageGRID sites than sites located in North America or Europe.

Conclusion

The integration of F5 BIG-IP with NetApp StorageGRID addresses technical challenges related to data availability and consistency across multiple sites and optimizing S3 transaction routing. Deploying this solution enhances storage resilience, performance, and reliability, making it ideal for enterprises seeking a robust, scalable, and flexible storage infrastructure.

To learn more, the official F5 documentation for BIG-IP DNS can be found at this [link](#).

A guided classroom style guide which provides step-by-step instructions on an example setup can also be found [here](#).

Datadog SNMP configuration

By Aron Klein

Configure Datadog to collect StorageGRID snmp metrics and traps.

Configure Datadog

Datadog is a monitoring solution providing metrics, visualizations, and alerting. The following configuration was implemented with linux agent version 7.43.1 on an Ubuntu 22.04.1 host deployed local to the StorageGRID system.

Datadog Profile and Trap files Generated from StorageGRID MIB file

Datadog provides a method for converting product MIB files into datadog reference files required to map the SNMP messages.

This StorageGRID yaml file for Datadog Trap resolution mapping generated following the instruction found [here](#).

Place this file in /etc/datadog-agent/conf.d/snmp.d/traps_db/ +

- [Download the trap yaml file](#) +
 - **md5 checksum** 42e27e4210719945a46172b98c379517 +
 - **sha256 checksum** d0fe5c8e6ca3c902d054f854b70a85f928cba8b7c76391d356f05d2cf73b6887 +

This StorageGRID profile yaml file for Datadog metrics mapping generated following the instruction found [here](#). Place this file in /etc/datadog-agent/conf.d/snmp.d/profiles/ +

- [Download the profile yaml file](#) +
 - **md5 checksum** 72bb7784f4801adda4e0c3ea77df19aa +
 - **sha256 checksum** b6b7fadd33063422a8bb8e39b3ead8ab38349ee0229926eadc8585f0087b8cee +

SNMP Datadog configuration for Metrics

Configuring SNMP for metrics can be managed in two ways. You can configure for auto-discovery by providing a network address range containing the StorageGRID system(s), or define the IP's of the individual devices. The configuration location is different based on the decision made. Auto-discovery is defined in the datadog agent yaml file. Explicit device definitions are configured in the snmp configuration yaml file. Below are examples of each for the same StorageGRID system.

Auto-discovery

configuration located in /etc/datadog-agent/datadog.yaml

```

listeners:
  - name: snmp
snmp_listener:
  workers: 100 # number of workers used to discover devices concurrently
  discovery_interval: 3600 # interval between each autodiscovery in
seconds
  loader: core # use core check implementation of SNMP integration.
recommended
  use_device_id_as_hostname: true # recommended
configs:
  - network_address: 10.0.0.0/24 # CIDR subnet
    snmp_version: 2
    port: 161
    community_string: 'st0r@gegrid' # enclose with single quote
    profile: netapp-storagegrid

```

Individual devices

/etc/datadog-agent/conf.d/snmp.d/conf.yaml

```

init_config:
  loader: core # use core check implementation of SNMP integration.
recommended
  use_device_id_as_hostname: true # recommended
instances:
- ip_address: '10.0.0.1'
  profile: netapp-storagegrid
  community_string: 'st0r@gegrid' # enclose with single quote
- ip_address: '10.0.0.2'
  profile: netapp-storagegrid
  community_string: 'st0r@gegrid'
- ip_address: '10.0.0.3'
  profile: netapp-storagegrid
  community_string: 'st0r@gegrid'
- ip_address: '10.0.0.4'
  profile: netapp-storagegrid
  community_string: 'st0r@gegrid'

```

SNMP configuration for traps

The configuration for SNMP traps is defined in the datadog configuration yaml file /etc/datadog-agent/datadog.yaml

```

network_devices:
  namespace: # optional, defaults to "default".
  snmp_traps:
    enabled: true
    port: 9162 # on which ports to listen for traps
    community_strings: # which community strings to allow for v2 traps
      - st0r@gegrid

```

Example StorageGRID SNMP configuration

The SNMP agent in your StorageGRID system is located under the configuration tab, Monitoring column. Enable SNMP and enter the desired information. If you wish to configure traps, select the "Traps Destinations" and Create a destination for the Datadog agent host containing the trap configuration.

SNMP Agent

You can configure SNMP for read-only MIB access and notifications. SNMPv1, SNMPv2c, SNMPv3 are supported. For SNMPv3, only User Security Model (USM) authentication is supported. All nodes in the grid share the same SNMP configuration.

Enable SNMP

System Contact

System Location lab

Enable SNMP Agent Notifications

Enable Authentication Traps

Community Strings

Default Trap Community st0r@gegrid

Read-Only Community

String 1 st0r@gegrid

Other Configurations

Trap Destinations (1)

Agent Addresses (0) **USM Users (0)** **Trap Destinations (1)**

+ Create **Edit** **X Remove**

Version	Type	Host	Port	Protocol	Community/USM User
SNMPv2C	Inform	10.193.92.241	9162	UDP	Default Community: st0r@gegrid

Use rclone to migrate, PUT, and DELETE objects on StorageGRID

By Siegfried Hepp and Aron Klein

rclone is a free command line tool and client for S3 operations. You can use rclone to migrate, copy, and delete object data on StorageGRID. rclone includes the capability to delete buckets even when not empty with a "purge" function as seen in an example below.

Install and configure rclone

To install rclone on a workstation or server, download it from rclone.org.

Initial configuration steps

1. Create the rclone configuration file by either running the config script or manually creating the file.
2. For this example I will use sgdemo for the name of the remote StorageGRID S3 endpoint in the rclone configuration.
 - a. Create the config file `~/.config/rclone/rclone.conf`

```
[sgdemo]
type = s3
provider = Other
access_key_id = ABCDEFGH123456789JKL
secret_access_key = 123456789ABCDEFGHIJKLMN0123456789PQRST+V
endpoint = sgdemo.netapp.com
```

- b. Run `rclone config`

```
# rclone config
```

```
2023/04/13 14:22:45 NOTICE: Config file
"/root/.config/rclone/rclone.conf" not found - using defaults
No remotes found - make a new one
n) New remote
s) Set configuration password
q) Quit config
n/s/q> n
name> sgdemo
```

Option Storage.

Type of storage to configure.

Enter a string value. Press Enter for the default ("").

Choose a number from below, or type in your own value.

```
1 / 1Fichier
  \ "fichier"
2 / Alias for an existing remote
  \ "alias"
3 / Amazon Drive
  \ "amazon cloud drive"
4 / Amazon S3 Compliant Storage Providers including AWS,
Alibaba, Ceph, Digital Ocean, Dreamhost, IBM COS, Minio,
SeaweedFS, and Tencent COS
  \ "s3"
5 / Backblaze B2
  \ "b2"
6 / Better checksums for other remotes
  \ "hasher"
7 / Box
  \ "box"
8 / Cache a remote
  \ "cache"
9 / Citrix Sharefile
  \ "sharefile"
10 / Compress a remote
  \ "compress"
11 / Dropbox
  \ "dropbox"
12 / Encrypt/Decrypt a remote
  \ "crypt"
13 / Enterprise File Fabric
  \ "filefabric"
14 / FTP Connection
```

```
  \ "ftp"
15 / Google Cloud Storage (this is not Google Drive)
    \ "google cloud storage"
16 / Google Drive
    \ "drive"
17 / Google Photos
    \ "google photos"
18 / Hadoop distributed file system
    \ "hdfs"
19 / Hubic
    \ "hubic"
20 / In memory object storage system.
    \ "memory"
21 / Jottacloud
    \ "jottacloud"
22 / Koofr
    \ "koofr"
23 / Local Disk
    \ "local"
24 / Mail.ru Cloud
    \ "mailru"
25 / Mega
    \ "mega"
26 / Microsoft Azure Blob Storage
    \ "azureblob"
27 / Microsoft OneDrive
    \ "onedrive"
28 / OpenDrive
    \ "opendrive"
29 / OpenStack Swift (Rackspace Cloud Files, Memset Memstore,
OVH)
    \ "swift"
30 / Pcloud
    \ "pcloud"
31 / Put.io
    \ "putio"
32 / QingCloud Object Storage
    \ "qingstor"
33 / SSH/SFTP Connection
    \ "sftp"
34 / Sia Decentralized Cloud
    \ "sia"
35 / Sugarsync
    \ "sugarsync"
36 / Tardigrade Decentralized Cloud Storage
    \ "tardigrade"
```

```
37 / Transparently chunk/split large files
  \ "chunker"
38 / Union merges the contents of several upstream fs
  \ "union"
39 / Uptobox
  \ "uptobox"
40 / Webdav
  \ "webdav"
41 / Yandex Disk
  \ "yandex"
42 / Zoho
  \ "zoho"
43 / http Connection
  \ "http"
44 / premiumize.me
  \ "premiumizeme"
45 / seafile
  \ "seafile"
```

Storage> 4

```
Option provider.
Choose your S3 provider.
Enter a string value. Press Enter for the default ("").
Choose a number from below, or type in your own value.
1 / Amazon Web Services (AWS) S3
  \ "AWS"
2 / Alibaba Cloud Object Storage System (OSS) formerly Aliyun
  \ "Alibaba"
3 / Ceph Object Storage
  \ "Ceph"
4 / Digital Ocean Spaces
  \ "DigitalOcean"
5 / Dreamhost DreamObjects
  \ "Dreamhost"
6 / IBM COS S3
  \ "IBMCOS"
7 / Minio Object Storage
  \ "Minio"
8 / Netease Object Storage (NOS)
  \ "Netease"
9 / Scaleway Object Storage
  \ "Scaleway"
10 / SeaweedFS S3
  \ "SeaweedFS"
11 / StackPath Object Storage
  \ "StackPath"
12 / Tencent Cloud Object Storage (COS)
  \ "TencentCOS"
13 / Wasabi Object Storage
  \ "Wasabi"
14 / Any other S3 compatible provider
  \ "Other"
provider> 14
```

```
Option env_auth.  
Get AWS credentials from runtime (environment variables or  
EC2/ECS meta data if no env vars).  
Only applies if access_key_id and secret_access_key is blank.  
Enter a boolean value (true or false). Press Enter for the  
default ("false").  
Choose a number from below, or type in your own value.  
1 / Enter AWS credentials in the next step.  
  \ "false"  
2 / Get AWS credentials from the environment (env vars or IAM).  
  \ "true"  
env_auth> 1
```

```
Option access_key_id.  
AWS Access Key ID.  
Leave blank for anonymous access or runtime credentials.  
Enter a string value. Press Enter for the default ("").  
access_key_id> ABCDEFGH123456789JKL
```

```
Option secret_access_key.  
AWS Secret Access Key (password).  
Leave blank for anonymous access or runtime credentials.  
Enter a string value. Press Enter for the default ("").  
secret_access_key> 123456789ABCDEFGHIJKLMN0123456789PQRST+V
```

```
Option region.  
Region to connect to.  
Leave blank if you are using an S3 clone and you don't have a  
region.  
Enter a string value. Press Enter for the default ("").  
Choose a number from below, or type in your own value.  
  / Use this if unsure.  
1 | Will use v4 signatures and an empty region.  
  \ ""  
  / Use this only if v4 signatures don't work.  
2 | E.g. pre Jewel/v10 CEPH.  
  \ "other-v2-signature"  
region> 1
```

Option endpoint.

Endpoint for S3 API.

Required when using an S3 clone.

Enter a string value. Press Enter for the default ("").

endpoint> sgdemo.netapp.com

Option location_constraint.

Location constraint - must be set to match the Region.

Leave blank if not sure. Used when creating buckets only.

Enter a string value. Press Enter for the default ("").

location_constraint>

Option acl.

Canned ACL used when creating buckets and storing or copying objects.

This ACL is used for creating objects and if bucket_acl isn't set, for creating buckets too.

For more info visit

<https://docs.aws.amazon.com/AmazonS3/latest/dev/acl-overview.html#canned-acl>

Note that this ACL is applied when server-side copying objects as S3

doesn't copy the ACL from the source but rather writes a fresh one.

Enter a string value. Press Enter for the default ("").

Choose a number from below, or type in your own value.

```
/ Owner gets FULL_CONTROL.  
1 | No one else has access rights (default).  
  \ "private"  
  / Owner gets FULL_CONTROL.  
2 | The AllUsers group gets READ access.  
  \ "public-read"  
  / Owner gets FULL_CONTROL.  
3 | The AllUsers group gets READ and WRITE access.  
  | Granting this on a bucket is generally not recommended.  
  \ "public-read-write"  
  / Owner gets FULL_CONTROL.  
4 | The AuthenticatedUsers group gets READ access.  
  \ "authenticated-read"  
  / Object owner gets FULL_CONTROL.  
5 | Bucket owner gets READ access.  
  | If you specify this canned ACL when creating a bucket,  
Amazon S3 ignores it.  
  \ "bucket-owner-read"  
  / Both the object owner and the bucket owner get FULL_CONTROL  
over the object.  
6 | If you specify this canned ACL when creating a bucket,  
Amazon S3 ignores it.  
  \ "bucket-owner-full-control"  
acl>
```

Edit advanced config?

y) Yes

n) No (default)

y/n> n

```
-----  
[sgdemo]  
type = s3  
provider = Other  
access_key_id = ABCDEFGH123456789JKL  
secret_access_key = 123456789ABCDEFGHIJKLMN0123456789PQRST+V  
endpoint = sgdemo.netapp.com:443  
-----  
y) Yes this is OK (default)  
e) Edit this remote  
d) Delete this remote  
y/e/d>
```

Current remotes:

Name	Type
=====	=====
sgdemo	s3

- e) Edit existing remote
- n) New remote
- d) Delete remote
- r) Rename remote
- c) Copy remote
- s) Set configuration password
- q) Quit config
- e/n/d/r/c/s/q> q

Basic command examples

- **Create a bucket:**

```
rclone mkdir remote:bucket
```

```
# rclone mkdir sgdemo:test01
```



Use --no-check-certificate if you need to ignore SSL certificates.

- **List all buckets:**

```
rclone lsd remote:
```

```
# rclone lsd sgdemo:
```

- **List objects in a specific bucket:**

```
rclone ls remote:bucket
```

```
# rclone ls sgdemo:test01
```

```
65536 TestObject.0
65536 TestObject.1
65536 TestObject.10
65536 TestObject.12
65536 TestObject.13
65536 TestObject.14
65536 TestObject.15
65536 TestObject.16
65536 TestObject.17
65536 TestObject.18
65536 TestObject.2
65536 TestObject.3
65536 TestObject.5
65536 TestObject.6
65536 TestObject.7
65536 TestObject.8
65536 TestObject.9
33554432 bigobj
102 key.json
47 locked01.txt
4294967296 sequential-read.0.0
15 test.txt
116 version.txt
```

- **Delete a bucket:**

```
rclone rmdir remote:bucket
```

```
# rclone rmdir sgdemo:test02
```

- **Put an object:**

```
rclone copy filename remote:bucket
```

```
# rclone copy ~/test/testfile.txt sgdemo:test01
```

- **Get an object:**

```
rclone copy remote:bucket/objectname filename
```

```
# rclone copy sgdemo:test01/testfile.txt ~/test/testfileS3.txt
```

- **Delete an object:**

```
rclone delete remote:bucket/objectname
```

```
# rclone delete sgdemo:test01/testfile.txt
```

- **Migrate objects in a bucket**

```
rclone sync source:bucket destination:bucket --progress
```

```
rclone sync source_directory destination:bucket --progress
```

```
# rclone sync sgdemo:test01 sgdemo:clone01 --progress
```

```
Transferred:          4.032 GiB / 4.032 GiB, 100%, 95.484 KiB/s, ETA
0s
Transferred:          22 / 22, 100%
Elapsed time:         1m4.2s
```



Use --progress or -P to display the progress of the task. Otherwise there is no output.

- **Delete a bucket and all object contents**

```
rclone purge remote:bucket --progress
```

```
# rclone purge sgdemo:test01 --progress
```

Transferred:	0 B / 0 B, -, 0 B/s, ETA -
Checks:	46 / 46, 100%
Deleted:	23 (files), 1 (dirs)
Elapsed time:	10.2s

```
# rclone ls sgdemo:test01
```

2023/04/14 09:40:51 Failed to ls: directory not found

StorageGRID best practices for deployment with Veeam Backup and Replication

By Oliver Haensel and Aron Klein

This guide focuses on the configuration of NetApp StorageGRID and partly Veeam Backup and Replication. This paper is written for storage and network administrators who are familiar with Linux systems and tasked with maintaining or implementing a NetApp StorageGRID system in combination with Veeam Backup and Replication.

Overview

Storage Administrators are looking to manage the growth of their data with solutions that will meet the availability, rapid recovery goals, scale to meet their needs and automate their policy for long-term retention of data. These solutions should also provide protection from loss or malicious attacks. Together, Veeam and NetApp have partnered to create a data protection solution combining Veeam Backup & Recovery with NetApp StorageGRID for on-premises object storage.

Veeam and NetApp StorageGRID provide an easy-to-use solution that work together to help meet the demands of rapid data growth and increasing regulations around the world. Cloud-based object storage is known for its resilience, ability to scale, operational and cost efficiencies that make it a natural choice as a target for your backups. This document will provide guidance and recommendations for the configuration of your Veeam Backup solution and StorageGRID system.

The object workload from Veeam creates a large number of concurrent PUT, DELETE, and LIST operations of small objects. Enabling immutability will add to the number of requests to the object store for setting retention and listing versions. The process of a backup job includes writing objects for the daily change then after the new writes are complete the job will delete any objects based on the retention policy of the backup. The scheduling of backup jobs will almost always overlap. This overlap will result in a large portion of the backup window consisting of 50/50 PUT/DELETE workload on the object store. Making adjustments in Veeam to the number of concurrent operations with the task slot setting, increasing the object size by increasing the backup job block size, reducing the number of objects in the multi-object delete requests, and choosing the maximum time window for the jobs to complete will optimize the solution for performance and cost.

Make sure to read the product documentation for [Veeam Backup and Replication](#) and [StorageGRID](#) before you begin. Veeam provides calculators for understanding the sizing of the Veeam infrastructure and capacity

requirements that should be used prior to sizing your StorageGRID solution. Please always check the Veeam-NetApp validated configurations at the Veeam Ready Program website for [Veeam Ready Object, Object Immutability, and Repository](#).

Veeam configuration

Recommended version

It is always recommended to stay current and apply the latest hotfixes for your Veeam Backup & Replication 12 or 12.1 system. Currently we recommend at a minimum installing Veeam 12 patch P20230718.

S3 Repository configuration

A scale-out backup repository (SOBR) is the capacity tier of S3 object storage. The capacity tier is an extension of the primary repository providing longer data retention periods and a lower cost storage solution. Veeam offers the ability to provide immutability through the S3 Object Lock API. Veeam 12 can use multiple buckets in a scale out repository. StorageGRID does not have a limit for the number of objects or capacity in a single bucket. Using multiple buckets may improve performance when backing up very large datasets where the backup data could get to petabyte scale in objects.

Limiting concurrent tasks may be required depending on the sizing of your specific solution and requirements. The default settings specify one repository task slot for each CPU core and for each task slot a concurrent task slot limit of 64. For example if your server has 2 CPU cores a total of 128 concurrent threads will be used for the object store. This is inclusive of PUT, GET, and batch Delete. It is recommended to select a conservative limit to the task slots to start with and tune this value once Veeam backups have reached a steady state of new backups and expiring backup data. Please work with your NetApp account team to size the StorageGRID system appropriately to meet the desired time windows and performance. Adjusting the number of task slots and the limit of tasks per slot may be required to provide the optimal solution.

Backup job configuration

Veeam backup jobs can be configured with different block size options that should be considered carefully. The default block size is 1MB and with the storage efficiencies Veeam provides with compression and deduplication creates object sizes of approximately 500kB for the initial Full backup and 100-200kB objects for the incremental jobs. We can greatly increase performance and scale down the requirements for the object store by choosing a larger backup block size. Though the larger block size makes great improvements in the object store performance it comes at the cost of potentially increased primary storage capacity requirement due to reduced storage efficiency performance. It is recommended for the backup jobs to be configured with a 4MB block size which creates approximately 2MB objects for the full backups and 700kB-1MB object sizes for incrementals. Customers may consider even configuring backup jobs using 8 MB block size, which can be enabled with assistance from Veeam support.

The implementation of immutable backups makes use of S3 Object Lock on the object store. The immutability option generates an increased number of requests to the object store for listing and retention updates on the objects.

As backup retentions expire the backup jobs will process the deletion of objects. Veeam sends the delete requests to the object store in multi-object delete requests of 1000 objects per request. For small solutions this may need to be adjusted to reduce the number of objects per request. Lowering this value will have the added benefit of more evenly distributing the delete requests across the nodes in the StorageGRID system. It is recommended to use the values in the table below as a starting point in configuring the multi object delete limit. Multiply the value in the table by the number of nodes for the chosen appliance type to get the value for the setting in Veeam. If this value is equal to or greater than 1000 there is no need to adjust the default value. If this value needs to be adjusted, please work with Veeam support to make the change.

Appliance Model	S3MultiObjectDeleteLimit per node
SG5712	34
SG5760	75
SG6060	200

Please work with your NetApp Account team for the recommended configuration based on your specific needs. The Veeam configuration settings recommendations will include:



- Backup job block size = 4MB
- SOBR task slot limit= 2-16
- Multi Object Delete Limit = 34-1000

StorageGRID configuration

Recommended version

NetApp StorageGRID 11.9 or 12.0 with the latest hotfix are the recommended versions for Veeam deployments. It is always recommended to stay current and apply the latest hotfixes for your StorageGRID system.

Load balancer and S3 endpoint configuration

Veeam requires the endpoint to be connected via HTTPS only. A non-encrypted connection is not supported by Veeam. The SSL certificate can be a self-signed certificate, private trusted certificate authority, or public trusted certificate authority. To ensure continuous access to the S3 repository it is recommended to use at least two load balancers in an HA configuration. The load balancers can be a StorageGRID provided integrated load balancer service located on every admin node and gateway node or third-party solution such as F5, Kemp, HAProxy, Loadbalancer.org, etc. Using a StorageGRID load balancer will provide the ability to set traffic classifiers (QoS rules) that can prioritize the Veeam workload, or limit Veeam to not impact higher priority workloads on the StorageGRID system.

S3 Bucket

StorageGRID is a secure multi-tenant storage system. It is recommended to create a dedicated tenant for the Veeam workload. A storage quota can be optionally assigned. As a best practice enable “use own identity source”. Secure the tenant root management user with an appropriate password. Veeam Backup 12 requires strong consistency for S3 buckets. StorageGRID offers multiple consistency options configured at the bucket level. For multi-site deployments with Veeam accessing the data from multiple locations, select “strong-global”. If Veeam backups and restores happen at a single site only, consistency level should be set to “strong-site”. For more information on bucket consistency levels please review the [documentation](#). To use StorageGRID for Veeam immutability backups, S3 Object Lock must be enabled globally and configured on the bucket during the bucket creation.

Lifecycle management

StorageGRID supports replication and erasure coding for object level protection across StorageGRID nodes and sites. Erasure Coding requires at least a 200kB object size. The default block size for Veeam of 1MB produces object sizes that can often be below this 200kB recommended minimum size after Veeam’s storage efficiencies. For the performance of the solution, it is not recommended to use an erasure coding profile spanning multiple sites unless the connectivity between the sites is sufficient to not add latency or restrict the

bandwidth of the StorageGRID system. In a multi-site StorageGRID system the ILM rule can be configured to store a single copy at each site. For ultimate durability a rule could be configured to store an erasure coded copy at each site. Using two copies local to the Veeam Backup servers is the most recommended implementation for this workload.

Delete performance

Veeam provides delete request rate tuning and scheduling of the backup delete process. To further tune delete performance you can disable synchronous deletes and let the ILM scanner manage the eventual deletion of objects.

Steps to Disabling Synchronous Deletes

1. Open the StorageGRID Grid Manager.
2. In top right corner, select the Question mark then API Documentation.
3. In top right corner, click on the Private API Documentation page link.
4. Expand `ilm-advanced`.
5. Select `GET ilm-advanced`.
6. Select Try it out, then Execute.
7. Check the response result.
 - a. If the values are null, then it means the default `ilm-advanced` values are in-use.
 - b. If the values are not null, then it means custom ILM advanced values are in-use. Copy all the output after "data" :, starting with the { up until the second to last }.
 - i. Save it in some text editor.

Example response:

Response body

```
{  
  "responseTime": "2025-09-19T15:01:28.142Z",  
  "status": "success",  
  "apiVersion": "4.2",  
  "data": {  
    "deletes": {  
      "synchronous": null,  
      "deleteQueueWorkers": null,  
      "asynchronousQueueRatio": null,  
      "synchronousTimeout": null,  
      "asyncILMDeletes": null,  
      "maxConcurrentUnlinkTruncateOps": null  
    },  
    "scanner": {  
      "ignoreTimeSinceLastClientOp": null,  
      "ignoreTimeSinceLastILMOp": null,  
      "scanRate": null,  
      "leakedUUIDCheckRatio": null,  
      "leakedUUIDMaxConcurrentWorkers": null,  
      "leakedUUIDIgnoreTimeSinceLastEvent": null,  
      "bucketDeleteObjectsMaxConcurrentWorkers": null  
    }  
  }  
}
```

8. Select PUT ilm-advanced.
9. Select Try it out to begin editing the API body.
 - a. By default, the API body will contain default values and not any custom values that were previously configured. This is the reason it is VERY important to execute steps 5-7.
10. If non-default values are found in step 5-7, then replace the API body with the output saved in step 7. . Otherwise if the values were null in step 5-7, then leave API body as is.
11. Adjust the following parameters in the API body box:
 - a. Set the synchronous value to false.

Example API body text:

```
{  
  "deletes": {  
    "synchronous": false,  
    "deleteQueueWorkers": null,  
    "asynchronousQueueRatio": 10,  
    "synchronousTimeout": 30,  
    "asyncILMDeletes": null,  
    "maxConcurrentUnlinkTruncateOps": null  
  },  
  "scanner": {  
    "ignoreTimeSinceLastClientOp": 3600,  
    "ignoreTimeSinceLastILMOp": 10800,  
    "scanRate": null,  
    "leakedUUIDCheckRatio": 10,  
    "leakedUUIDMaxConcurrentWorkers": 64,  
    "leakedUUIDIgnoreTimeSinceLastEvent": 3600,  
    "bucketDeleteObjectsMaxConcurrentWorkers": 64  
  }  
}
```

12. Once complete, select Execute

Implementation key points

StorageGRID

Ensure Object Lock is enabled on the StorageGRID system if immutability is required. Find the option in the management UI under Configuration/S3 Object Lock.

Configuration > S3 Object Lock

S3 Object Lock

ⓘ S3 Object Lock has been enabled for the grid and cannot be disabled.

Enable S3 Object Lock for your entire StorageGRID system if S3 tenant accounts need to satisfy regulatory compliance requirements when saving object data. After this setting is enabled, it cannot be disabled.

Before enabling S3 Object Lock, you must ensure that the default rule in the active ILM policy is compliant. A compliant rule satisfies the requirements of buckets with S3 Object Lock enabled.

- It must create at least two replicated object copies or one erasure-coded copy.
- These copies must exist on Storage Nodes for the entire duration of each line in the placement instructions.
- Object copies cannot be saved in a Cloud Storage Pool.
- Object copies cannot be saved on Archive Nodes.
- At least one line of the placement instructions must start at day 0, using Ingest Time as the reference time.
- At least one line of the placement instructions must be "forever".

Enable S3 Object Lock

Apply

When creating the bucket, select “Enable S3 Object Lock” if this bucket is to be used for immutability backups. This will automatically enable bucket versioning. Leave default retention disabled as Veeam will set object retention explicitly. Versioning and S3 Object Lock should not be selected if Veeam isn’t creating immutable backups.

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.

i 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

Default retention ?

Automatically protect new objects put into this bucket from being deleted or overwritten.

Disable

Enable

Once the bucket is created go to the details page of the bucket created. Select the consistency level.

Buckets > veeam12

veeam12

Region: us-east-1

S3 Object Lock: Enabled

Date created: 2023-09-21 08:01:38 GMT

Object count: 0

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

[Delete objects in bucket](#) [Delete bucket](#)

[Bucket options](#) [Bucket access](#) [Platform services](#)

Consistency level	Read-after-new-write (default)	▼
Last access time updates	Disabled	▼
Object versioning	Enabled	▼
S3 Object Lock	Enabled	▼

Veeam requires strong consistency for S3 buckets. So, for multi-site deployments with Veeam accessing the data from multiple locations, select “strong-global”. If Veeam backups and restores happen at a single site only, consistency level should be set to “strong-site”. Save the changes.

[Bucket options](#) [Bucket access](#) [Platform services](#)

Consistency level [Read-after-new-write \(default\)](#) [▲](#)

Change the consistency control for operations performed on the objects in the bucket. Consistency levels provide a balance between the availability of objects and the consistency of those objects across different Storage Nodes and sites.

In general, use the **Read-after-new-write** consistency level for your buckets. Then, if objects do not meet availability or consistency requirements, change the client application's behavior, or set the Consistency-Control header for an individual API request, which overrides the bucket setting.

All
Provides the highest guarantee of consistency. All nodes receive the data immediately, or the request will fail.

Strong-global
Guarantees read-after-write consistency for all client requests across all sites.

Strong-site
Guarantees read-after-write consistency for all client requests within a site.

Read-after-new-write (default)
Provides read-after-write consistency for new objects and eventual consistency for object updates. Offers high availability and data protection guarantees. Recommended for most cases.

Available
Provides eventual consistency for both new objects and object updates. For S3 buckets, use only as required (for example, for a bucket that contains log values that are rarely read, or for HEAD or GET operations on keys that do not exist). Not supported for FabricPool buckets.

[Save changes](#)

Last access time updates [Disabled](#) [▼](#)

StorageGRID provides an integrated load balancer service on every admin node and dedicated gateway nodes. One of the many advantages of using this load balancer is the ability to configure Traffic Classification

Policies (QoS). Though these are mainly used for limiting an applications impact on other client workloads or prioritizing a workload over others, they also provide a bonus of additional metrics collection to assist in monitoring.

In the configuration tab, select “Traffic Classification” and create a new policy. Name the rule and select either the bucket(s) or tenant as the type. Enter the name(s) of the bucket(s) or tenant. If QoS is required, set a limit, but for most implementations, we just want to add the monitoring benefits this provides so do not set a limit.

Create a traffic classification policy

You can create traffic classification policies to monitor the network traffic for specific buckets, tenants, IP addresses, subnets, or load balancer endpoints. You can optionally limit this traffic based on bandwidth, number of concurrent requests, or the request rate.

Enter policy name — Add matching rules — Set limits — **4** Review the policy

Review the policy

Policy Veeam
name:

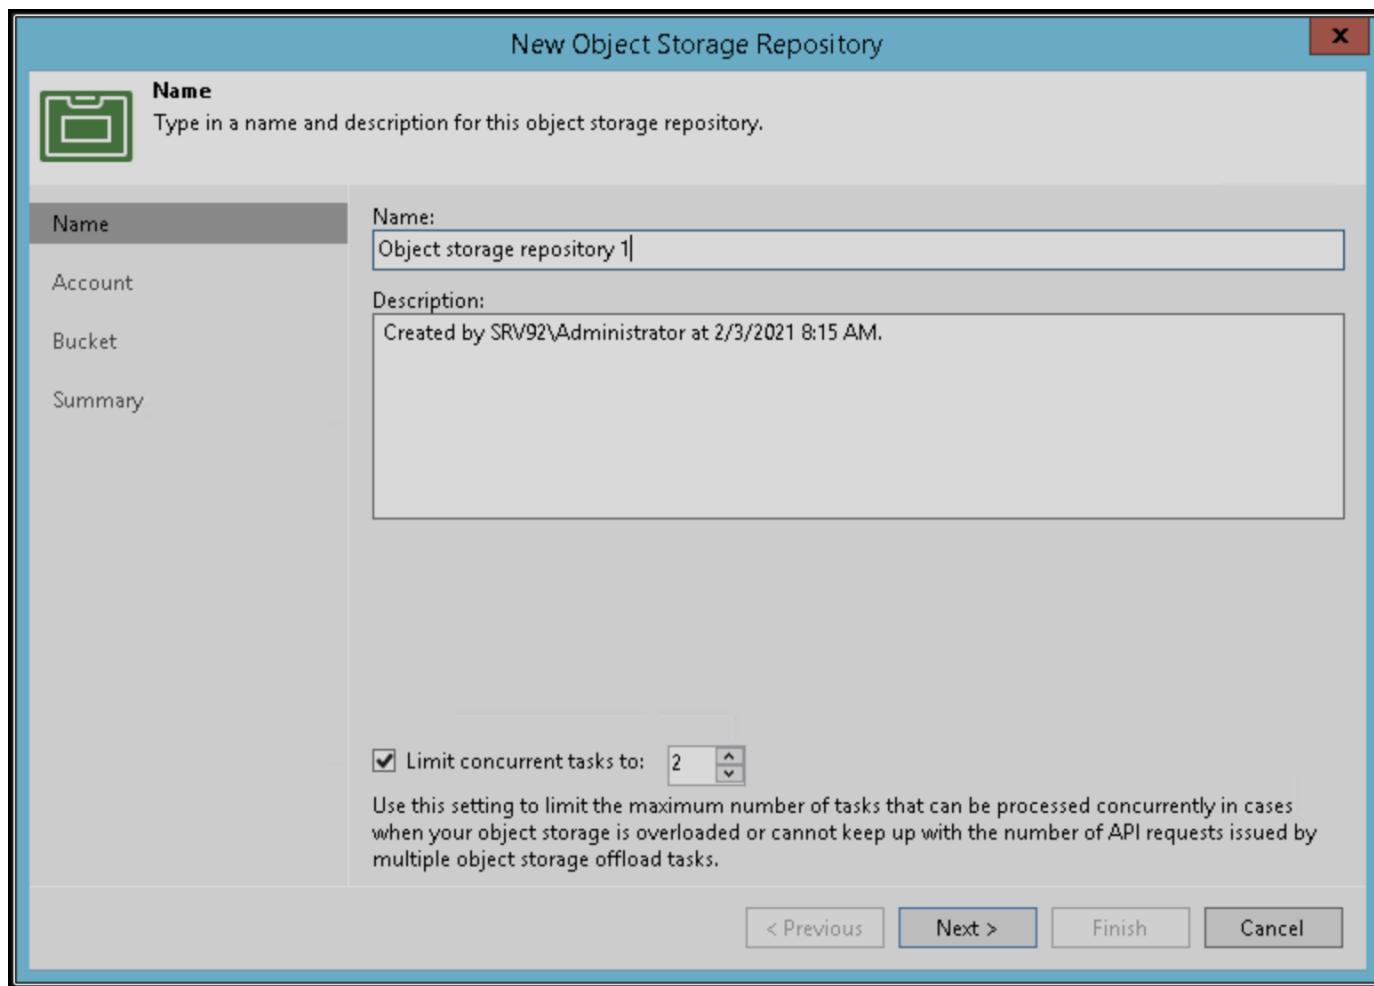
Description: Policy to monitor
Veeam bucket
traffic

Matching rules

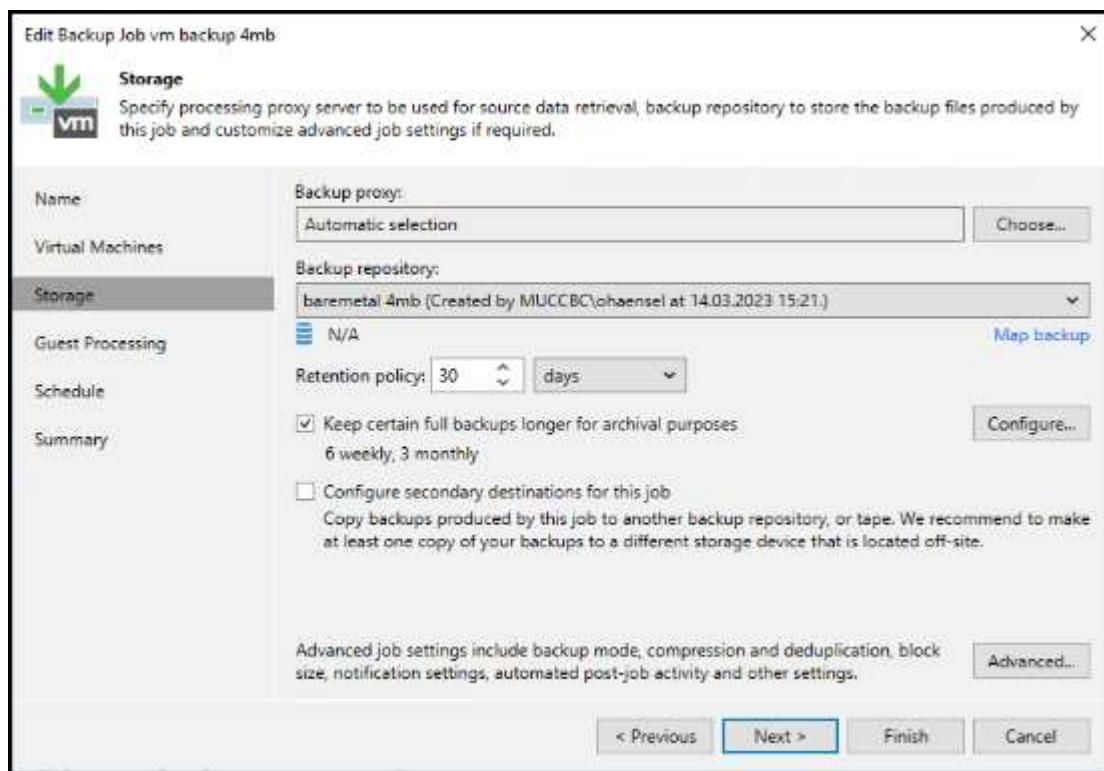
Type	Match value	Inverse match
Bucket	test	No

Veeam

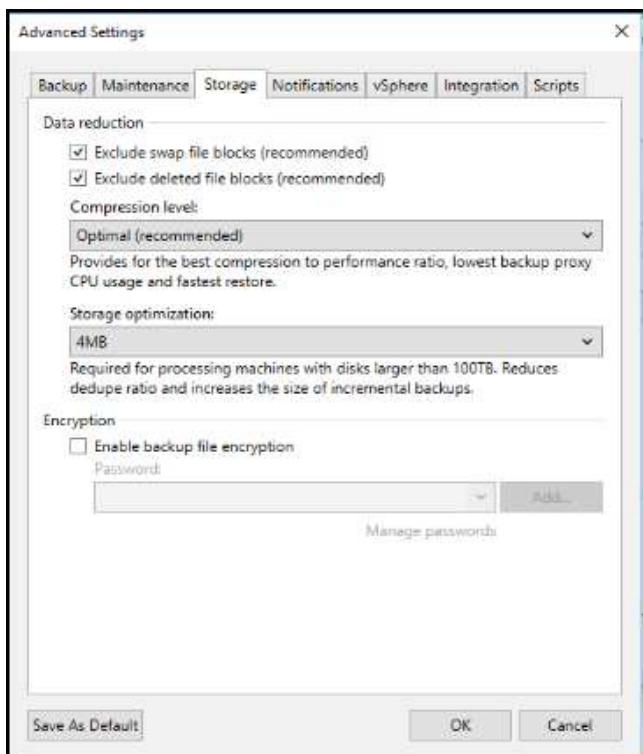
Depending on the model and quantity of StorageGRID appliances it may be necessary to select and configure a limit to the number of concurrent operations on the bucket.



Follow the Veeam documentation on backup job configuration in the Veeam console to start the wizard. After adding VMs select the SOBR repository.



Click Advanced settings and change storage optimization settings to 4 MB or larger. Compression and deduplication shall be enabled. Change guest settings according to your requirements and configure the backup job schedule.

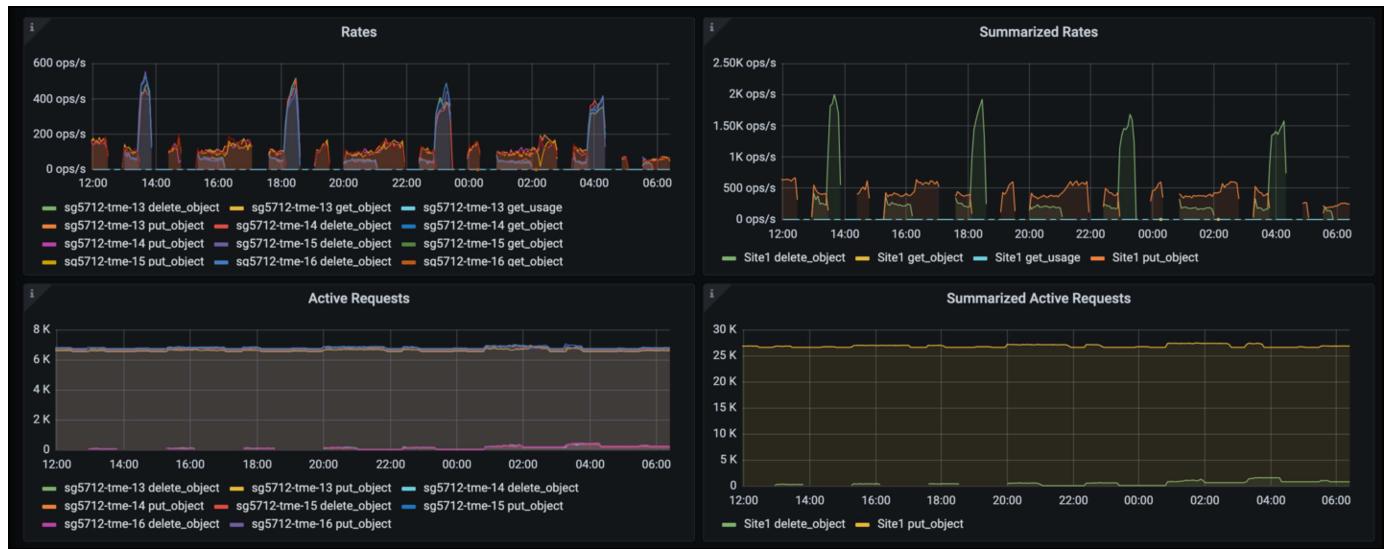


Monitoring StorageGRID

To get the full picture of how Veeam and StorageGRID are performing together you will need to wait until the retention time of the first backups have expired. Up until this point the Veeam workload consists primarily of PUT operations and no DELETEs have occurred. Once there is backup data expiring and cleanups are occurring you can now see the full consistent usage in the object store and adjust the settings in Veeam if needed.

StorageGRID provides convenient charts to monitor the operation of the system located in the Support tab Metrics page. The primary dashboards to look at will be the S3 Overview, ILM, and Traffic Classification Policy if a policy was created. In the S3 Overview dashboard you will find information on the S3 operation rates, latencies, and request responses.

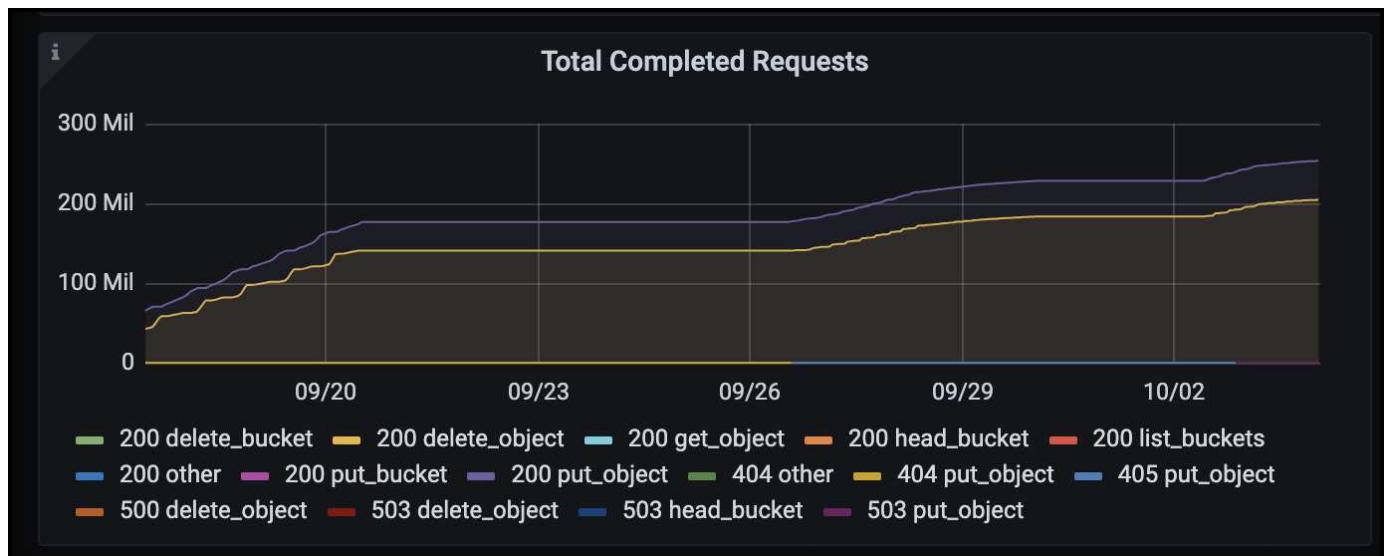
Looking at the S3 rates and active requests you can see how much of the load each node is handling and the overall number of requests by type.



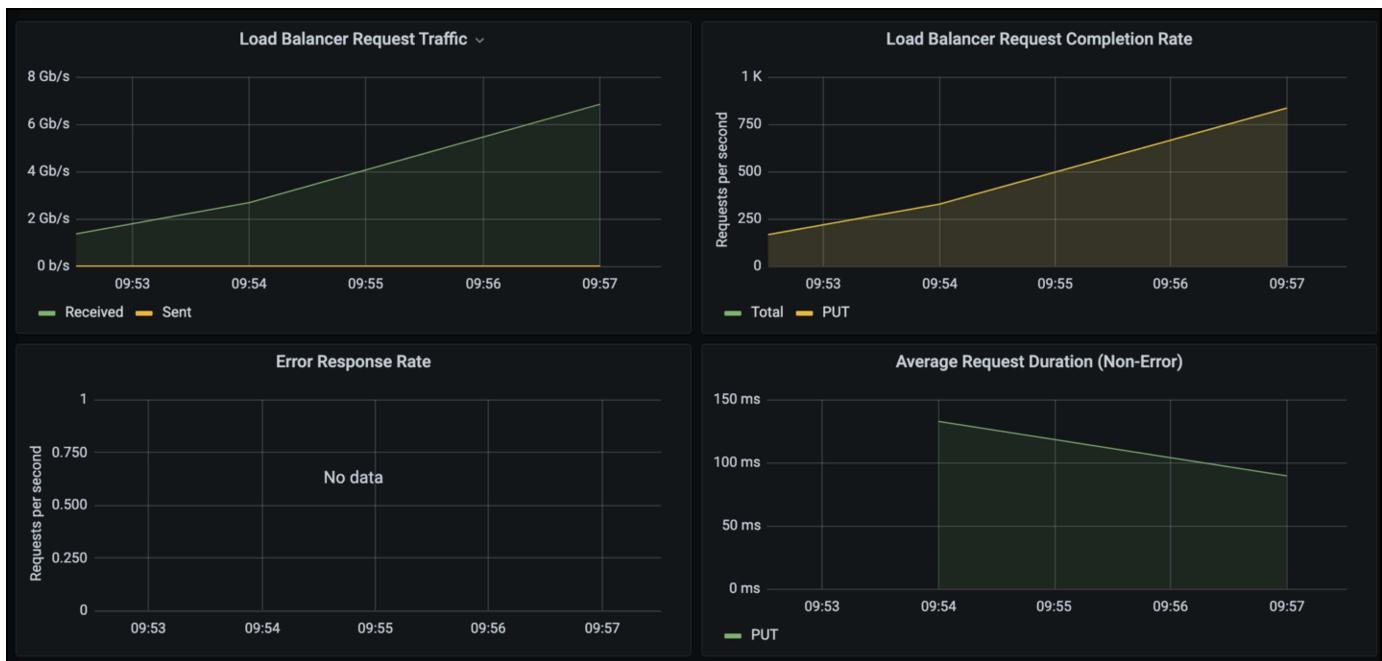
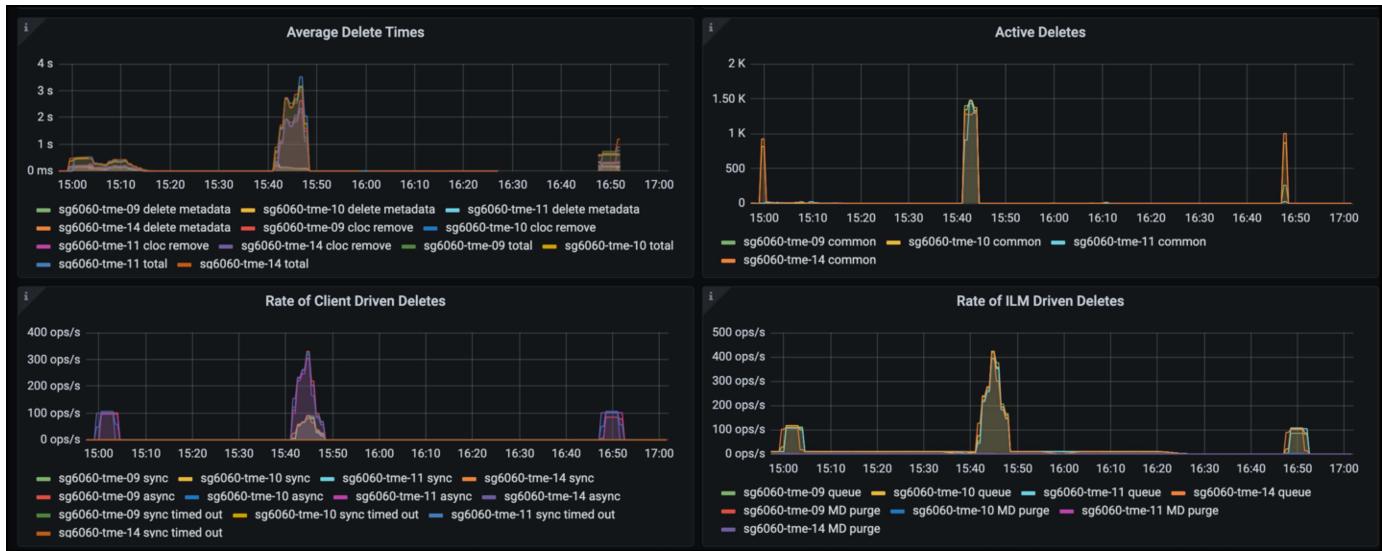
The Average Duration chart shows the average time each node is taking for each request type. This is the average latency of the request and may be a good indicator that additional tuning may be required, or there is room for the StorageGRID system to take on more load.



In the Total Completed Requests chart, you can see the requests by type and response codes. If you see responses other than 200 (Ok) for the responses this may indicate an issue like the StorageGRID system is getting heavily loaded sending 503 (Slow Down) responses and some additional tuning may be necessary, or the time has come to expand the system for the increased load.



In the ILM Dashboard you can monitor the Delete performance of your StorageGRID system. StorageGRID uses a combination of synchronous and asynchronous deletes on each node to try and optimize the overall performance for all requests.



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 Product Documentation](#)
- [Veeam Backup and Replication](#)

Configure Dremio data source with StorageGRID

By *Angela Cheng*

Dremio supports a variety of data sources, including cloud-based or on-premises object storage. You can configure Dremio to use StorageGRID as object storage data source.

Configure Dremio data source

Prerequisites

- A StorageGRID S3 endpoint URL, a tenant s3 access key ID, and secret access key.
- StorageGRID configuration recommendation: disable compression (disabled by default).
Dremio uses byte range GET to fetch different byte ranges from within the same object concurrently during query. Typical size for byte-range requests is 1MB. Compressed object degrades byte-range GET performance.

Dremio guide

[Connecting to Amazon S3 - Configuring S3-Compatible Storage.](#)

Instruction

1. On Dremio Datasets page, click + sign to add a source, select 'Amazon S3'.
2. Enter a name for this new data source, StorageGRID S3 tenant access key ID and secret access key.
3. Check the box 'Encrypt connection' if using https for connection to StorageGRID S3 endpoint.
If using self-signed CA cert for this s3 endpoint, follow Dremio guide instruction to add this CA cert into Dremio server's <JAVA_HOME>/jre/lib/security

Sample screenshot

General

[Advanced Options](#)
[Reflection Refresh](#)
[Metadata](#)
[Privileges](#)

 Amazon S3 Source

Name
parquet-1tb

Authentication

AWS Access Key
 EC2 Metadata
 AWS Profile
 No Authentication

All or allowlisted (if specified) buckets associated with this access key or IAM role to assume (if specified) will be available.

AWS Access Key

.....

AWS Access Secret

.....

IAM Role to Assume

.....

Encrypt connection

Public Buckets

Buckets
No public buckets added

⊕ Add bucket

4. Click 'Advanced Options', check 'Enable compatibility mode'
5. Under Connection properties, click + Add Properties and add these s3a properties.
6. `fs.s3a.connection.maximum` default is 100. If your s3 datasets include large Parquet files with 100 or more columns, must enter a value greater than 100. Refer to Dremio guide for this setting.

Name	Value
<code>fs.s3a.endpoint</code>	<code><StorageGRID S3 endpoint:port></code>
<code>fs.s3a.path.style.access</code>	<code>true</code>
<code>fs.s3a.connection.maximum</code>	<code><a value greater than 100></code>

Sample screenshot

General

<input checked="" type="checkbox"/>	Enable asynchronous access when possible
<input checked="" type="checkbox"/>	Enable compatibility mode
<input type="checkbox"/>	Apply requester-pays to S3 requests

Advanced Options

Reflection Refresh

<input checked="" type="checkbox"/>	Enable file status check
-------------------------------------	--------------------------

Metadata

<input type="checkbox"/>	Enable partition column inference
--------------------------	-----------------------------------

Privileges

<input type="checkbox"/>	Root Path
<input type="text" value="/"/>	

Server side encryption key ARN

<input type="text"/>	
----------------------	--

Default CTAS Format

<input type="text" value="PARQUET"/>	▼
--------------------------------------	--

Connection Properties

Name	Value	X
<input type="text" value="fs.s3a.path.style.access"/>	<input type="text" value="true"/>	<input type="button" value="X"/>
<input type="text" value="fs.s3a.endpoint"/>	<input type="text" value="sgdemo.netapp.com"/>	<input type="button" value="X"/>
<input type="text" value="fs.s3a.connection.maximum"/>	<input type="text" value="1000"/>	<input type="button" value="X"/>

[⊕ Add property](#)

Allowlisted buckets

No allowlisted buckets added

[⊕ Add bucket](#)

Cache Options

<input checked="" type="checkbox"/>	Enable local caching when possible
Max percent of total available cache space to use when possible	
<input type="text" value="100"/>	

7. Configure other Dremio options as per your organization or application requirements.
8. Click the Save button to create this new data source.
9. Once StorageGRID data source is added successfully, a list of buckets will be displayed on the left panel.

Sample screenshot

The screenshot shows the NetApp StorageGRID web interface. On the left, a sidebar has icons for Home, Spaces, and Sources. The main area has a search bar at the top. Under 'Datasets', there is a section for 'hdp-user' with 'Spaces (0)' and a '+'. Below it is a placeholder for 'No spaces yet' with an 'Add space' button. Under 'Sources', there is a section for 'Object Storage (2)' with 'StorageGRID' and a '0'. To the right, a 'StorageGRID' section lists various storage locations: apache-hive, cdp-cluster, cdp-tera, databrick-tpcds, delta-lake, dlcluster-tpcds, dremio-10g-csv, and dremio-csv.

NetApp StorageGRID with GitLab

By Angela Cheng

NetApp has tested StorageGRID with GitLab. See sample GitLab configuration below. Refer to [GitLab object storage configuration guide](#) for details.

Object Storage connection example

For Linux Package installations, this is an example of the connection setting in the consolidated form. Edit `/etc/gitlab/gitlab.rb` and add the following lines, substituting the values you want:

```

# Consolidated object storage configuration
gitlab_rails['object_store']['enabled'] = true
gitlab_rails['object_store']['proxy_download'] = true
gitlab_rails['object_store']['connection'] = {
  'provider' => 'AWS',
  'region' => 'us-east-1',
  'endpoint' => 'https://<storagegrid-s3-endpoint:port>',
  'path_stype' => 'true',
  'aws_access_key_id' => '<AWS_ACCESS_KEY_ID>',
  'aws_secret_access_key' => '<AWS_SECRET_ACCESS_KEY>'
}
# OPTIONAL: The following lines are only needed if server side encryption
# is required
gitlab_rails['object_store']['storage_options'] = {
  'server_side_encryption' => 'AES256'
}
gitlab_rails['object_store']['objects']['artifacts']['bucket'] = 'gitlab-
artifacts'
gitlab_rails['object_store']['objects']['external_diffs']['bucket'] =
'gitlab-mr-diffs'
gitlab_rails['object_store']['objects']['lfs']['bucket'] = 'gitlab-lfs'
gitlab_rails['object_store']['objects']['uploads']['bucket'] = 'gitlab-
uploads'
gitlab_rails['object_store']['objects']['packages']['bucket'] = 'gitlab-
packages'
gitlab_rails['object_store']['objects']['dependency_proxy']['bucket'] =
'gitlab-dependency-proxy'
gitlab_rails['object_store']['objects']['terraform_state']['bucket'] =
'gitlab-terraform-state'
gitlab_rails['object_store']['objects']['pages']['bucket'] = 'gitlab-
pages'

```

Procedures and API examples

Test and demonstrate S3 encryption options on StorageGRID

By Aron Klein

StorageGRID and the S3 API offer a number of different ways to encrypt your data at rest. To learn more, see [Review StorageGRID encryption methods](#).

This guide will demonstrate the S3 API encryption methods.

Server Side Encryption (SSE)

SSE allows the client to store an object and encrypt it with a unique key that is managed by StorageGRID. When the object is requested, the object is decrypted by the key stored in storageGRID.

SSE Example

- PUT an object with SSE

```
aws s3api put-object --bucket <bucket> --key <file> --body "<file>"  
--server-side-encryption AES256 --endpoint-url https://s3.example.com
```

- HEAD the object to verify encryption

```
aws s3api head-object --bucket <bucket> --key <file> --endpoint-url  
https://s3.example.com
```

```
{  
  "AcceptRanges": "bytes",  
  "LastModified": "2022-05-02T19:03:03+00:00",  
  "ContentLength": 47,  
  "ETag": "\"82e8bfb872e778a4687a26e6c0b36bc1\"",  
  "ContentType": "text/plain",  
  "ServerSideEncryption": "AES256",  
  "Metadata": {}  
}
```

- GET the object

```
aws s3api get-object --bucket <bucket> --key <file> <file> --endpoint  
-url https://s3.example.com
```

Server Side Encryption with Customer provided keys (SSE-C)

SSE allows the client to store an object and encrypt it with a unique key that is provided by the client with the object. When the object is requested, the same key must be provided in order to decrypt and return the object.

SSE-C Example

- For testing or demonstration purposes you can create an encryption key
 - Create an encryption key

```
openssl enc -aes-128-cbc -pass pass:secret -P`
```

```
salt=E9DBB6603C7B3D2A
key=23832BAC16516152E560F933F261BF03
iv =71E87C0F6EC3C45921C2754BA131A315
```

- Put an object with the generated key

```
aws s3api put-object --bucket <bucket> --key <file> --body "file" --sse
--customer-algorithm AES256 --sse-customer-key
23832BAC16516152E560F933F261BF03 --endpoint-url https://s3.example.com
```

- Head the object

```
aws s3api head-object --bucket <bucket> --key <file> --sse-customer
--algorithm AES256 --sse-customer-key 23832BAC16516152E560F933F261BF03
--endpoint-url https://s3.example.com
```

```
{
    "AcceptRanges": "bytes",
    "LastModified": "2022-05-02T19:20:02+00:00",
    "ContentLength": 47,
    "ETag": "\"f92ef20ab87e0e13951d9bee862e9f9a\"",
    "ContentType": "binary/octet-stream",
    "Metadata": {},
    "SSECustomerAlgorithm": "AES256",
    "SSECustomerKeyMD5": "rjGuMdjLpPV1eRuotNaPMQ=="
}
```

 If you do not provide the encryption key, you will receive an error "An error occurred (404) when calling the HeadObject operation: Not Found"

- Get the object

```
aws s3api get-object --bucket <bucket> --key <file> <file> --sse
--customer-algorithm AES256 --sse-customer-key
23832BAC16516152E560F933F261BF03 --endpoint-url https://s3.example.com
```



If you do not provide the encryption key, you will receive an error "An error occurred (InvalidRequest) when calling the GetObject operation: The object was stored using a form of Server Side Encryption. The correct parameters must be provided to retrieve the object."

Bucket Server Side Encryption (SSE-S3)

SSE-S3 allows the client to define a default encryption behavior for all objects stored in a bucket. The objects are encrypted with a unique key that is managed by StorageGRID. When the object is requested, the object is decrypted by the key stored in storageGRID.

Bucket SSE-S3 Example

- Create a new bucket and set a default encryption policy
 - Create new bucket

```
aws s3api create-bucket --bucket <bucket> --region us-east-1
--endpoint-url https://s3.example.com
```

- Put bucket encryption

```
aws s3api put-bucket-encryption --bucket <bucket> --server-side
--encryption-configuration '{"Rules": [
  {"ApplyServerSideEncryptionByDefault": {"SSEAlgorithm": "AES256"} }]} --endpoint-url https://s3.example.com
```

- Put an object in the bucket

```
aws s3api put-object --bucket <bucket> --key <file> --body "file"
--endpoint-url https://s3.example.com
```

- Head the object

```
aws s3api head-object --bucket <bucket> --key <file> --endpoint-url
https://s3.example.com
```

```
{  
  "AcceptRanges": "bytes",  
  "LastModified": "2022-05-02T20:16:23+00:00",  
  "ContentLength": 47,  
  "ETag": "\"82e8bfb872e778a4687a26e6c0b36bc1\"",  
  "ContentType": "binary/octet-stream",  
  "ServerSideEncryption": "AES256",  
  "Metadata": {}  
}
```

- GET the object

```
aws s3api get-object --bucket <bucket> --key <file> <file> --endpoint  
-url https://s3.example.com
```

Test and demonstrate S3 object lock on StorageGRID

By Aron Klein

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 and compliance mode for object retention, and default bucket retention policies.

This guide will demonstrate the S3 Object Lock API.

Legal hold

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

```
aws s3api put-object-legal-hold --bucket <bucket> --key <file> --legal  
-hold Status=ON --endpoint-url https://s3.company.com
```

- Verify it with a GET operation.

```
aws s3api get-object-legal-hold --bucket <bucket> --key <file>  
--endpoint-url https://s3.company.com
```

```
{  
  "LegalHold": {  
    "Status": "ON"  
  }  
}
```

- Turn legal hold off

```
aws s3api put-object-legal-hold --bucket <bucket> --key <file> --legal  
-hold Status=OFF --endpoint-url https://s3.company.com
```

- Verify it with a GET operation.

```
aws s3api get-object-legal-hold --bucket <bucket> --key <file>  
--endpoint-url https://s3.company.com
```

```
{  
  "LegalHold": {  
    "Status": "OFF"  
  }  
}
```

Compliance mode

- The object retention is done with a retain until timestamp.

```
aws s3api put-object-retention --bucket <bucket> --key <file>  
--retention '{"Mode":"COMPLIANCE", "RetainUntilDate": "2025-06-  
10T16:00:00"}' --endpoint-url https://s3.company.com
```

- Verify the retention status

```
aws s3api get-object-retention --bucket <bucket> --key <file> --endpoint  
-url https://s3.company.com  
+
```

```
{
  "Retention": {
    "Mode": "COMPLIANCE",
    "RetainUntilDate": "2025-06-10T16:00:00+00:00"
  }
}
```

Default retention

- Set the retention period in days and years verses a retain until date defined with the per object api.

```
aws s3api put-object-lock-configuration --bucket <bucket> --object-lock-configuration '{"ObjectLockEnabled": "Enabled", "Rule": { "DefaultRetention": { "Mode": "COMPLIANCE", "Days": 10 }}}' --endpoint-url https://s3.company.com
```

- Verify the retention status

```
aws s3api get-object-lock-configuration --bucket <bucket> --endpoint-url https://s3.company.com
```

```
{
  "ObjectLockConfiguration": {
    "ObjectLockEnabled": "Enabled",
    "Rule": {
      "DefaultRetention": {
        "Mode": "COMPLIANCE",
        "Days": 10
      }
    }
  }
}
```

- Put an object in the bucket

```
aws s3api put-object --bucket <bucket> --key <file> --body "file" --endpoint-url https://s3.example.com
```

- The retention duration set on the bucket is converted to a retention timestamp on the object.

```
aws s3api get-object-retention --bucket <bucket> --key <file> --endpoint-url https://s3.company.com
```

```
{  
  "Retention": {  
    "Mode": "COMPLIANCE",  
    "RetainUntilDate": "2022-03-02T15:22:47.202000+00:00"  
  }  
}
```

Test deleting an object with a defined retention

Object Lock is built on top of versioning. The retention is defined on a version of the object. If an attempt is made to delete an object with a retention defined, and no version is specified, a delete marker is created as the current version of the object.

- Delete the object with retention defined

```
aws s3api delete-object --bucket <bucket> --key <file> --endpoint-url https://s3.example.com
```

- List the objects in the bucket

```
aws s3api list-objects --bucket <bucket> --endpoint-url https://s3.example.com
```

- Notice the object is not listed.

- List versions to see the delete marker, and the original locked version

```
aws s3api list-object-versions --bucket <bucket> --prefix <file> --endpoint-url https://s3.example.com
```

```
{
  "Versions": [
    {
      "ETag": "\"82e8bfb872e778a4687a26e6c0b36bc1\"",
      "Size": 47,
      "StorageClass": "STANDARD",
      "Key": "file.txt",
      "VersionId": "RDVDMjYwMTQtQkNDQS0xMUVDLThGOEUtNjQ3NTAwQzAxQTk1",
      "IsLatest": false,
      "LastModified": "2022-04-15T14:46:29.734000+00:00",
      "Owner": {
        "DisplayName": "Tenant01",
        "ID": "56622399308951294926"
      }
    }
  ],
  "DeleteMarkers": [
    {
      "Owner": {
        "DisplayName": "Tenant01",
        "ID": "56622399308951294926"
      },
      "Key": "file01.txt",
      "VersionId": "QjVDQzgzOTAtQ0FGNi0xMUVDLThFMzgtQ0RGMjAwQjk0MjM1",
      "IsLatest": true,
      "LastModified": "2022-05-03T15:35:50.248000+00:00"
    }
  ]
}
```

- Delete the locked version of the object

```
aws s3api delete-object --bucket <bucket> --key <file> --version-id
"<VersionId>" --endpoint-url https://s3.example.com
```

An error occurred (AccessDenied) when calling the DeleteObject operation: Access Denied

Policies and permissions in StorageGRID

Here are example Policies and permissions in StorageGRID S3.

The structure of a policy

In StorageGRID, group policies are the same as AWS user (IAM) S3 service policies.

Group policies are required in StorageGRID. A user with S3 access keys but not assigned to a user group, or assigned to a group without a policy granting some permissions, will not be able to access any data.

Bucket and group policies share most of the same elements. Policies are built in json format and can be generated using the [AWS policy generator](#)

All policies will define the effect, action(s), and resource(s). Bucket policies will also define a principal.

The **Effect** will be either to Allow or Deny the request.

The Principal

- Only applies for bucket policies.
- The principal is the account(s)/user(s) being granted or denied the permissions.
- Can be defined as:
 - A wildcard `"*"`

```
"Principal": "*"
```

```
"Principal": { "AWS": "*" }
```

- A tenant ID for all users in a tenant (equivalent to AWS account)

```
"Principal": { "AWS": "27233906934684427525" }
```

- A user (local or federated from within the tenant the bucket resides, or another tenant in the grid)

```
"Principal": { "AWS": "arn:aws:iam::76233906934699427431:user/tenant1user1" }
```

```
"Principal": { "AWS": "arn:aws:iam::27233906934684427525:federated-user/tenant2user1" }
```

- A group (local or federated from within the tenant the bucket resides, or another tenant in the grid).

```
"Principal": { "AWS":  
"arn:aws:iam::76233906934699427431:group/DevOps" }
```

```
"Principal": { "AWS": "arn:aws:iam::27233906934684427525:federated-  
group/Managers" }
```

The **Action** is the set of S3 operations being granted or denied to the user(s).



For Group policies, the s3>ListBucket action Allowed is required for users to perform any S3 actions.

The **Resource** is the bucket or buckets the principals being granted or denied the ability to perform the actions on.

Optionally there can be a **Condition** for when the policy action is valid.

The format of the Json policy will look like this:

```
{  
  "Statement": [  
    {  
      "Sid": "Custom name for this permission",  
      "Effect": "Allow or Deny",  
      "Principal": {  
        "AWS": [  
          "arn:aws:iam::tenant_ID:user/User_Name",  
          "arn:aws:iam::tenant_ID:federated-user/User_Name",  
          "arn:aws:iam::tenant_ID:group/Group_Name",  
          "arn:aws:iam::tenant_ID:federated-group/Group_Name",  
          "tenant_ID"  
        ]  
      },  
      "Action": [  
        "s3>ListBucket",  
        "s3:Other_Action"  
      ],  
      "Resource": [  
        "arn:aws:s3:::Example_Bucket",  
        "arn:aws:s3:::Example_Bucket/*"  
      ],  
      "Condition": {}  
    }  
  ]  
}
```

Using the AWS policy generator

The AWS policy generator is a great tool to assist with getting the json code with the correct format and information you are trying to implement.

To generate the permissions for a StorageGRID group policy:

- * choose the IAM policy for the type of policy.
- * Select the button for the desired effect - Allow or Deny. It is a good practice to start your policies with the deny permissions and then add the allow permissions
- * In the actions drop-down click the box next to as many of the S3 actions you want to include in this permission or the "All Actions" box.
- * type in the bucket paths in the Amazon Resource Name (ARN) box. include "arn:aws:s3:::" before the bucket name. ex. "arn:aws:s3:::example_bucket"



AWS Policy Generator

The AWS Policy Generator is a tool that enables you to create policies that control access to Amazon Web Services (AWS) products and resources. For more information about creating policies, see [key concepts in Using AWS Identity and Access Management](#). Here are sample policies.

Step 1: Select Policy Type

A Policy is a container for permissions. The different types of policies you can create are an [IAM Policy](#), an [S3 Bucket Policy](#), an [SNS Topic Policy](#), a [VPC Endpoint Policy](#), and an [SQS Queue Policy](#).

Select Type of Policy For group policy choose IAM Policy

Step 2: Add Statement(s)

A statement is the formal description of a single permission. See [a description of elements](#) that you can use in statements.

Effect Allow Deny

AWS Service Choose Amazon S3 service

Actions Select the S3 actions to allow or deny

Amazon Resource Name (ARN) arn:aws:s3:::Bucket_Name

ARN should follow the following format: arn:aws:s3:::\${BucketName}/\${KeyName}.
Use a comma to separate multiple values.

Add Conditions (Optional)

No Action selected. You must select at least one Action

Step 3: Generate Policy

A [policy](#) is a document (written in the [Access Policy Language](#)) that acts as a container for one or more statements.

Add one or more statements above to generate a policy.

To generate the permissions for a bucket policy:

- * choose the S3 Bucket Policy for the type of policy.
- * Select the button for the desired effect - Allow or Deny. It is a good practice to start your policies with the deny permissions and then add the allow permissions
- * Type in the user or group information for the Principal.
- * In the actions drop-down click the box next to as many of the S3 actions you want to include in this permission or the "All Actions" box.
- * type in the bucket paths in the Amazon Resource Name (ARN) box. include "arn:aws:s3:::" before the bucket name. ex. "arn:aws:s3:::example_bucket"

AWS Policy Generator

The AWS Policy Generator is a tool that enables you to create policies that control access to Amazon Web Services (AWS) products and resources. For more information about creating policies, see [key concepts in Using AWS Identity and Access Management](#). Here are [sample policies](#).

Step 1: Select Policy Type

A Policy is a container for permissions. The different types of policies you can create are an [IAM Policy](#), an [S3 Bucket Policy](#), an [SNS Topic Policy](#), a [VPC Endpoint Policy](#), and an [SQS Queue Policy](#).

Select Type of Policy S3 Bucket Policy ← For bucket policy choose S3 Bucket Policy

Step 2: Add Statement(s)

A statement is the formal description of a single permission. See [a description of elements](#) that you can use in statements.

Effect Allow Deny

Principal arn:aws:iam::Tenant_ID:user/User_Name ← arn:aws:iam::Tenant_ID:user/User_Name

AWS Service Amazon S3 ← arn:aws:s3:::Bucket_Name All Services (*)

Actions -- Select Actions -- ← Select the S3 actions to allow or deny All Actions (*)

Amazon Resource Name (ARN) arn:aws:s3:::Bucket_Name ← arn:aws:s3:::Bucket_Name

Add Conditions (Optional)

Add Statement

Step 3: Generate Policy

A *policy* is a document (written in the [Access Policy Language](#)) that acts as a container for one or more statements.

Add one or more statements above to generate a policy.

For an example, if you wanted to generate a bucket policy to allow all users to perform GetObject operations on all objects in the bucket, while only users belonging the group "Marketing" in the specified account are allowed full access.

- Select S3 Bucket Policy as the policy type.
- Choose the Allow effect
- Enter the Marketing group information - arn:aws:iam::95390887230002558202:federated-group/Marketing
- Click the box for "All Actions"
- Enter the bucket information - arn:aws:s3:::example_bucket,arn:aws:s3:::example_bucket/*



AWS Policy Generator

The AWS Policy Generator is a tool that enables you to create policies that control access to **Amazon Web Services (AWS)** products and services. For more information about creating policies, see [key concepts in Using AWS Identity and Access Management](#). Here are [sample policies](#).

Step 1: Select Policy Type

A Policy is a container for permissions. The different types of policies you can create are an **IAM Policy**, an **S3 Bucket Policy**, an **SNS To Queue Policy**.

Select Type of Policy

Step 2: Add Statement(s)

A statement is the formal description of a single permission. See [a description of elements](#) that you can use in statements.

Effect Allow Deny

Principal ← arn:aws:iam::95390887230002558202:federated-group/Marketing
Use a comma to separate multiple values.

AWS Service ← arn:aws:s3:::examplebucket All Services (*)
Use multiple statements to add permissions for more than one service.

Actions ← arn:aws:s3:::examplebucket,arn:aws:s3:::examplebucket/* All Actions (*)

Amazon Resource Name (ARN) ← arn:aws:s3:::examplebucket,arn:aws:s3:::examplebucket/*
ARN should follow the following format: arn:aws:s3:::\${BucketName}/\${KeyName}.
Use a comma to separate multiple values.

Add Conditions (Optional)

Add Statement

- click the "Add Statement" button

You added the following statements. Click the button below to Generate a policy.

Principal(s)	Effect	Action	Resource	Conditions
• arn:aws:iam::95390887230002558202:federated-group/Marketing	Allow	s3:*	• arn:aws:s3:::examplebucket • arn:aws:s3:::examplebucket/*	None

- Choose the Allow effect
- Enter the asterisk * for everyone
- Click the box next to GetObject and ListBucket actions"

1 Action(s) Selected

- GetMultiRegionAccessPointRoutes
- GetObject
- GetObjectAcl
- GetObjectAttributes
- GetObjectLegalHold
- GetObjectRetention
- GetObjectTagging
- GetObjectTorrent

2 Action(s) Selected

- ListAccessPointsForObjectLambda
- ListAllMyBuckets
- ListBucket
- ListBucketMultipartUploads
- ListBucketVersions
- ListCallerAccessGrants
- ListJobs

- Enter the bucket information - arn:aws:s3:::example_bucket,arn:aws:s3:::example_bucket/*



AWS Policy Generator

The AWS Policy Generator is a tool that enables you to create policies that control access to Amazon Web Services (AWS) products and services. For more information about creating policies, see [key concepts in Using AWS Identity and Access Management](#). Here are [sample policies](#).

Step 1: Select Policy Type

A Policy is a container for permissions. The different types of policies you can create are an [IAM Policy](#), an [S3 Bucket Policy](#), an [SNS Topic Policy](#), and an [SQS Queue Policy](#).

Select Type of Policy

Step 2: Add Statement(s)

A statement is the formal description of a single permission. See [a description of elements](#) that you can use in statements.

Effect Allow Deny

Principal Use a comma to separate multiple values.

AWS Service All Services (*) Use multiple statements to add permissions for more than one service.

Actions All Actions (*)

Amazon Resource Name (ARN) ARN should follow the following format: arn:aws:s3:::\${BucketName}/\${KeyName}. Use a comma to separate multiple values.

Add Conditions (Optional)

- click the "Add Statement" button

You added the following statements. Click the button below to Generate a policy.

Principal(s)	Effect	Action	Resource	Conditions
* arn:aws:iam::95390887230002558202:federated-group/Marketing	Allow	s3:*	* arn:aws:s3:::examplebucket * arn:aws:s3:::examplebucket/*	None
* *	Allow	s3:GetObject s3>ListBucket	* arn:aws:s3:::examplebucket * arn:aws:s3:::examplebucket/*	None

- Click on the button "Generate Policy" and a pop-up window will appear with your generated policy.

Policy JSON Document X

Click below to edit. To save the policy, copy the text below to a text editor.
Changes made below will not be reflected in the policy generator tool.

```
1  {
  "Id": "Policy1744399292233",
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1744399152830",
      "Action": "s3:*",
      "Effect": "Allow",
      "Resource": [
        "arn:aws:s3:::examplebucket",
        "arn:aws:s3:::examplebucket/*"
      ],
      "Principal": {
        "AWS": [
          "arn:aws:iam::95390887230002558202:federated-group/Marketing"
        ]
      }
    },
    {
      "Sid": "Stmt1744399280838",
      "Action": [
        "s3:PutObject"
      ],
      "Resource": [
        "arn:aws:s3:::examplebucket/*"
      ],
      "Principal": [
        "arn:aws:iam::95390887230002558202:federated-group/Marketing"
      ]
    }
  ]
}
```

Close

- Copy out the complete json text that should look like this:

```
{
  "Id": "Policy1744399292233",
  "Version": "2012-10-17",
  "Statement": [
    {
      "Sid": "Stmt1744399152830",
      "Action": "s3:*",
      "Effect": "Allow",
      "Resource": [
        "arn:aws:s3:::example_bucket",
        "arn:aws:s3:::example_bucket/*"
      ],
      "Principal": {
        "AWS": [
          "arn:aws:iam::95390887230002558202:federated-group/Marketing"
        ]
      }
    },
    {
      "Sid": "Stmt1744399280838",
      "Action": [
        "s3:GetObject",
        "s3>ListBucket"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:s3:::example_bucket",
        "arn:aws:s3:::example_bucket/*"
      ],
      "Principal": "*"
    }
  ]
}
```

This json can be used as is, or you can remove the ID and Version lines above the "Statement" line and you can customize the Sid for each permission with a more meaningful title for each permission or these can be removed as well.

For example:

```
{
  "Statement": [
    {
      "Sid": "MarketingAllowFull",
      "Action": "s3:*",
      "Effect": "Allow",
      "Resource": [
        "arn:aws:s3:::example_bucket",
        "arn:aws:s3:::example_bucket/*"
      ],
      "Principal": {
        "AWS": [
          "arn:aws:iam::95390887230002558202:federated-group/Marketing"
        ]
      }
    },
    {
      "Sid": "EveryoneReadOnly",
      "Action": [
        "s3:GetObject",
        "s3>ListBucket"
      ],
      "Effect": "Allow",
      "Resource": [
        "arn:aws:s3:::example_bucket",
        "arn:aws:s3:::example_bucket/*"
      ],
      "Principal": "*"
    }
  ]
}
```

Group Policies (IAM)

Home Directory style bucket access

This group policy will only allow users to access objects in the bucket named the users username.

```
{  
  "Statement": [  
    {  
      "Sid": "AllowListBucketOfASpecificUserPrefix",  
      "Effect": "Allow",  
      "Action": "s3>ListBucket",  
      "Resource": "arn:aws:s3:::home",  
      "Condition": {  
        "StringLike": {  
          "s3:prefix": "${aws:username}/*"  
        }  
      }  
    },  
    {  
      "Sid": "AllowUserSpecificActionsOnlyInTheSpecificUserPrefix",  
      "Effect": "Allow",  
      "Action": "s3:*Object",  
      "Resource": "arn:aws:s3:::home/?/?/${aws:username}/*"  
    }  
  ]  
}
```

Deny object lock bucket creation

This group policy will restrict users from creating a bucket with object lock enabled on the bucket.



This policy is not enforced in the StorageGRID UI, it is only enforced by S3 API.

```
{
  "Statement": [
    {
      "Action": "s3:*",
      "Effect": "Allow",
      "Resource": "arn:aws:s3:::*"
    },
    {
      "Action": [
        "s3:PutBucketObjectLockConfiguration",
        "s3:PutBucketVersioning"
      ],
      "Effect": "Deny",
      "Resource": "arn:aws:s3:::*"
    }
  ]
}
```

Object lock retention limit

This Bucket policy will restrict Object-Lock retention duration to 10 days or less

```
{
  "Version": "2012-10-17",
  "Id": "CustSetRetentionLimits",
  "Statement": [
    {
      "Sid": "CustSetRetentionPeriod",
      "Effect": "Deny",
      "Principal": "*",
      "Action": [
        "s3:PutObjectRetention"
      ],
      "Resource": "arn:aws:s3:::testlock-01/*",
      "Condition": {
        "NumericGreaterThan": {
          "s3:object-lock-remaining-retention-days": "10"
        }
      }
    }
  ]
}
```

Restrict users from deleting objects by versionID

This group policy will restrict users from deleting versioned objects by versionID

```
{  
  "Statement": [  
    {  
      "Action": [  
        "s3:DeleteObjectVersion"  
      ],  
      "Effect": "Deny",  
      "Resource": "arn:aws:s3:::*"  
    },  
    {  
      "Action": "s3:*",  
      "Effect": "Allow",  
      "Resource": "arn:aws:s3:::/*"  
    }  
  ]  
}
```

Restrict a group to single subdirectory (prefix) with read-only access

This policy allows members of the group to have read-only access to a subdirectory (prefix) within a bucket. The bucket name is "study" and the subdirectory is "study01".

```
{  
  "Statement": [  
    {  
      "Sid": "AllowUserToSeeBucketListInTheConsole",  
      "Action": [  
        "s3>ListAllMyBuckets"  
      ],  
      "Effect": "Allow",  
      "Resource": [  
        "arn:aws:s3:::/*"  
      ]  
    },  
    {  
      "Sid": "AllowRootAndstudyListingOfBucket",  
      "Action": [  
        "s3>ListBucket"  
      ],  
      "Effect": "Allow",  
      "Resource": [  
        "arn:aws:s3::: study"  
      ]  
    }  
  ]  
}
```

```

] ,
  "Condition": {
    "StringEquals": {
      "s3:prefix": [
        "",
        "study01/"
      ],
      "s3:delimiter": [
        "/"
      ]
    }
  }
},
{
  "Sid": "AllowListingOfstudy01",
  "Action": [
    "s3>ListBucket"
  ],
  "Effect": "Allow",
  "Resource": [
    "arn:aws:s3:::study"
  ],
  "Condition": {
    "StringLike": {
      "s3:prefix": [
        "study01/*"
      ]
    }
  }
},
{
  "Sid": "AllowAllS3ActionsInstudy01Folder",
  "Effect": "Allow",
  "Action": [
    "s3>GetObject"
  ],
  "Resource": [
    "arn:aws:s3:::study/study01/*"
  ]
}
]
}

```

Bucket Policies

Restrict bucket to single user with read-only access

This policy allows a single user to have read-only access to a bucket and explicitly denies access to all other users. Grouping the Deny statements at the top of the policy is a good practice for faster evaluation.

```
{  
  "Statement": [  
    {  
      "Sid": "Deny non user1",  
      "Effect": "Deny",  
      "NotPrincipal": {  
        "AWS": "arn:aws:iam::34921514133002833665:user/user1"  
      },  
      "Action": [  
        "s3:*"  
      ],  
      "Resource": [  
        "arn:aws:s3:::bucket1",  
        "arn:aws:s3:::bucket1/*"  
      ]  
    },  
    {  
      "Sid": "Allow user1 read access to bucket bucket1",  
      "Effect": "Allow",  
      "Principal": {  
        "AWS": "arn:aws:iam::34921514133002833665:user/user1"  
      },  
      "Action": [  
        "s3:GetObject",  
        "s3>ListBucket"  
      ],  
      "Resource": [  
        "arn:aws:s3:::bucket1",  
        "arn:aws:s3:::bucket1/*"  
      ]  
    }  
  ]  
}
```

restrict a bucket to a few users with read-only access.

```

{
  "Statement": [
    {
      "Sid": "Deny all S3 actions to employees 002-005",
      "Effect": "deny",
      "Principal": {
        "AWS": [
          "arn:aws:iam::46521514133002703882:user/employee-002",
          "arn:aws:iam::46521514133002703882:user/employee-003",
          "arn:aws:iam::46521514133002703882:user/employee-004",
          "arn:aws:iam::46521514133002703882:user/employee-005"
        ]
      },
      "Action": "*",
      "Resource": [
        "arn:aws:s3:::databucket1",
        "arn:aws:s3:::databucket1/*"
      ]
    },
    {
      "Sid": "Allow read-only access for employees 002-005",
      "Effect": "Allow",
      "Principal": {
        "AWS": [
          "arn:aws:iam::46521514133002703882:user/employee-002",
          "arn:aws:iam::46521514133002703882:user/employee-003",
          "arn:aws:iam::46521514133002703882:user/employee-004",
          "arn:aws:iam::46521514133002703882:user/employee-005"
        ]
      },
      "Action": [
        "s3:GetObject",
        "s3:GetObjectTagging",
        "s3:GetObjectVersion"
      ],
      "Resource": [
        "arn:aws:s3:::databucket1",
        "arn:aws:s3:::databucket1/*"
      ]
    }
  ]
}

```

Restrict user deletes of versioned objects in a bucket

This bucket policy will restrict a user(identified by userID "56622399308951294926") from deleting versioned objects by versionID

```
{  
  "Statement": [  
    {  
      "Action": [  
        "s3:DeleteObjectVersion"  
      ],  
      "Effect": "Deny",  
      "Resource": "arn:aws:s3:::verdeny/*",  
      "Principal": {  
        "AWS": [  
          "56622399308951294926"  
        ]  
      }  
    },  
    {  
      "Action": "s3:*",  
      "Effect": "Allow",  
      "Resource": "arn:aws:s3:::verdeny/*",  
      "Principal": {  
        "AWS": [  
          "56622399308951294926"  
        ]  
      }  
    }  
  ]  
}
```

Bucket lifecycle in StorageGRID

You can create an S3 lifecycle configuration to control when specific objects are deleted from the StorageGRID system.

What is a lifecycle configuration

A lifecycle configuration is a set of rules that are applied to the objects in specific S3 buckets. Each rule specifies which objects are affected and when those objects will expire (on a specific date or after some number of days).

Each object follows the retention settings of either an S3 bucket lifecycle or an ILM policy. When an S3 bucket lifecycle is configured, the lifecycle expiration actions override the ILM policy for objects matching the bucket lifecycle filter. Objects that do not match the bucket lifecycle filter use the retention settings of the ILM policy. If an object matches a bucket lifecycle filter and no expiration actions are explicitly specified, the retention

settings of the ILM policy are not used and it is implied that object versions are retained forever.

As a result, an object might be removed from the grid even though the placement instructions in an ILM rule still apply to the object. Or, an object might be retained on the grid even after any ILM placement instructions for the object have lapsed.

StorageGRID supports up to 1,000 lifecycle rules in a lifecycle configuration. Each rule can include the following XML elements:

- **Expiration:** Delete an object when a specified date is reached or when a specified number of days is reached, starting from when the object was ingested.
- **NoncurrentVersionExpiration:** Delete an object when a specified number of days is reached, starting from when the object became noncurrent.
- **Filter (Prefix, Tag)**
- **Status**
*ID

StorageGRID supports the use of the following bucket operations to manage lifecycle configurations:

- `DeleteBucketLifecycle`
- `GetBucketLifecycleConfiguration`
- `PutBucketLifecycleConfiguration`

Structure of a lifecycle policy

As the first step in creating a lifecycle configuration, you create a JSON file that includes one or more rules. For example, this JSON file includes three rules, as follows:

1. **Rule 1** applies only to objects that match the prefix `category1/` and that have a `key2` value of `tag2`. The `Expiration` parameter specifies that objects matching the filter will expire at midnight on 22 August 2020.
2. **Rule 2** applies only to objects that match the prefix `category2/`. The `Expiration` parameter specifies that objects matching the filter will expire 100 days after they are ingested.



Rules that specify a number of days are relative to when the object was ingested. If the current date exceeds the ingest date plus the number of days, some objects might be removed from the bucket as soon as the lifecycle configuration is applied.

3. **Rule 3** applies only to objects that match the prefix `category3/`. The `Expiration` parameter specifies that any noncurrent versions of matching objects will expire 50 days after they become noncurrent.

```
{
  "Rules": [
    {
      "ID": "rule1",
      "Filter": {
        "And": {
          "Prefix": "category1/",
          "Tags": [
            {
              "Key": "key2",
              "Value": "tag2"
            }
          ]
        }
      },
      "Expiration": {
        "Date": "2020-08-22T00:00:00Z"
      },
      "Status": "Enabled"
    },
    {
      "ID": "rule2",
      "Filter": {
        "Prefix": "category2/"
      },
      "Expiration": {
        "Days": 100
      },
      "Status": "Enabled"
    },
    {
      "ID": "rule3",
      "Filter": {
        "Prefix": "category3/"
      },
      "NoncurrentVersionExpiration": {
        "NoncurrentDays": 50
      },
      "Status": "Enabled"
    }
  ]
}
```

Apply lifecycle configuration to bucket

After you have created the lifecycle configuration file, you apply it to a bucket by issuing a PutBucketLifecycleConfiguration request.

This request applies the lifecycle configuration in the example file to objects in a bucket named testbucket.

```
aws s3api --endpoint-url <StorageGRID endpoint> put-bucket-lifecycle-configuration  
--bucket testbucket --lifecycle-configuration file://bktjson.json
```

To validate that a lifecycle configuration was successfully applied to the bucket, issue a GetBucketLifecycleConfiguration request. For example:

```
aws s3api --endpoint-url <StorageGRID endpoint> get-bucket-lifecycle-configuration  
--bucket testbucket
```

Example lifecycle policies for standard (non-versioned) buckets

Delete objects after 90 days

Use case: This policy is ideal for managing data that is only relevant for a limited time, such as temporary files, logs, or intermediate processing data.

Benefit: Reduce storage costs and ensure the bucket is clutter-free.

```
{  
  "Rules": [  
    {  
      "ID": "Delete after 90 day rule",  
      "Filter": {},  
      "Status": "Enabled",  
      "Expiration": {  
        "Days": 90  
      }  
    }  
  ]  
}
```

Example lifecycle policies for versioned buckets

Delete noncurrent versions after 10 days

Use case: This policy helps manage the storage of non-current version objects, which can accumulate over time and consume significant space.

Benefit: Optimize storage usage by keeping only the most recent version.

```
{
  "Rules": [
    {
      "ID": "NoncurrentVersionExpiration 10 day rule",
      "Filter": {},
      "Status": "Enabled",
      "NoncurrentVersionExpiration": {
        "NoncurrentDays": 10
      }
    }
  ]
}
```

Keep 5 noncurrent versions

Use case: Useful where you want to retain a limited number of previous versions for recovery or audit purposes

Benefit: Retain enough non-current versions to ensure sufficient history and recovery points.

```
{
  "Rules": [
    {
      "ID": "NewerNoncurrentVersions 5 version rule",
      "Filter": {},
      "Status": "Enabled",
      "NoncurrentVersionExpiration": {
        "NewerNoncurrentVersions": 5
      }
    }
  ]
}
```

Remove delete markers when no other versions exist

Use case: This policy helps manage the delete markers left over after all non-current versions are removed which can accumulate over time.

Benefit: Reduce unnecessary clutter.

```
{  
  "Rules": [  
    {  
      "ID": "Delete marker cleanup rule",  
      "Filter": {},  
      "Status": "Enabled",  
      "Expiration": {  
        "ExpiredObjectDeleteMarker": true  
      }  
    }  
  ]  
}
```

Delete current versions after 30 days, delete non-current versions after 60 days, and remove the delete markers created by the current version delete once no other versions exist.

Use case: Provide a complete lifecycle for current and non-current versions including the delete markers.

Benefit: Reduce storage costs and ensure the bucket is clutter-free while retaining sufficient recovery points and history.

```
{
  "Rules": [
    {
      "ID": "Delete current version",
      "Filter": {},
      "Status": "Enabled",
      "Expiration": {
        "Days": 30
      }
    },
    {
      "ID": "noncurrent version retention",
      "Filter": {},
      "Status": "Enabled",
      "NoncurrentVersionExpiration": {
        "NoncurrentDays": 60
      }
    },
    {
      "ID": "Markers",
      "Filter": {},
      "Status": "Enabled",
      "Expiration": {
        "ExpiredObjectDeleteMarker": true
      }
    }
  ]
}
```

remove delete markers that have no other versions, Retain 4 non-current versions and at least 30 days worth of history for objects with the "accounts_ prefix" and keep 2 versions and at least 10 days worth of history for all other object versions.

Use case: Provide unique rules for specific objects along side other objects to manage the complete lifecycle for current and non-current versions including the delete markers.

Benefit: Reduce storage costs and ensure the bucket is clutter-free while retaining sufficient recovery points and history to meet a mix of client requirements.

```
{
  "Rules": [
    {
      "ID": "Markers",
      "Filter": {},
      "Status": "Enabled",
      "Expiration": {
        "ExpiredObjectDeleteMarker": true
      }
    },
    {
      "ID": "accounts version retention",
      "Filter": {"Prefix": "account_"},
      "Status": "Enabled",
      "NoncurrentVersionExpiration": {
        "NewerNoncurrentVersions": 4,
        "NoncurrentDays": 30
      }
    },
    {
      "ID": "noncurrent version retention",
      "Filter": {},
      "Status": "Enabled",
      "NoncurrentVersionExpiration": {
        "NewerNoncurrentVersions": 2,
        "NoncurrentDays": 10
      }
    }
  ]
}
```

Conclusion

- Regularly review and update lifecycle policies and align them with ILM and data management goals.
- Test policies in a non-production environment or bucket before applying them broadly to ensure they work as intended
- Use Descriptive ID's for rules to make it more intuitive, as the logic structure can get complex
- Monitor the impact of these bucket lifecycle policies on storage usage and performance to make necessary adjustments.

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.

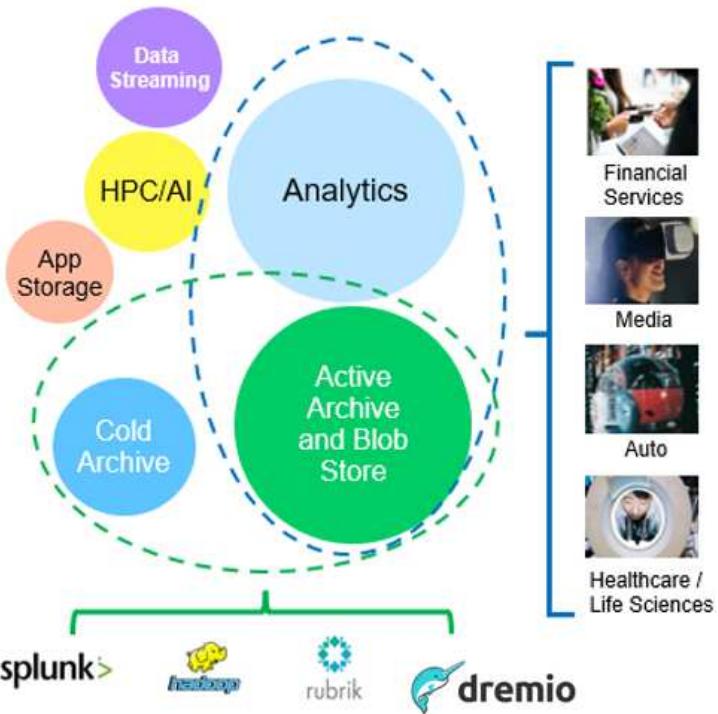
Data backup and disaster recovery: Businesses can use StorageGRID S3 as a reliable and cost-effective solution for backing up critical data and recovering it in case of a disaster.

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

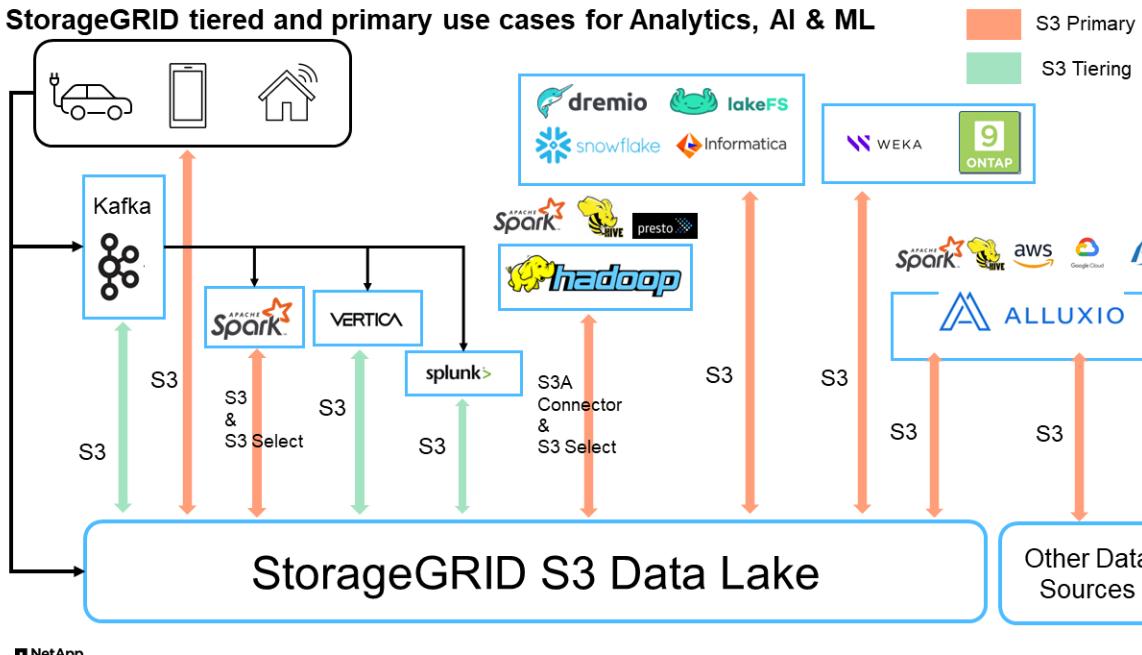


Among the above, big data analytics is one of the topmost use cases and its usage is trending upward.

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 titled 'Virtualization' and includes sections for File servers, Databases, Laptops, Cloud apps, Activate, Orchestrate, Jobs, Reports, and Monitoring. The main content area is titled 'VM groups / -DSN-server'. It has tabs for 'Overview' (which is selected) and 'Configuration'. The 'Overview' section shows the following statistics: 0 PROTECTED, 0 NOT PROTECTED, and 0 BACKED UP WITH ERRORS. The 'Content' section shows a single item: 'DNS-server-Bootcamp'. In the top right, there are buttons for 'Backup' and 'View jobs'. The top right corner of the interface 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.

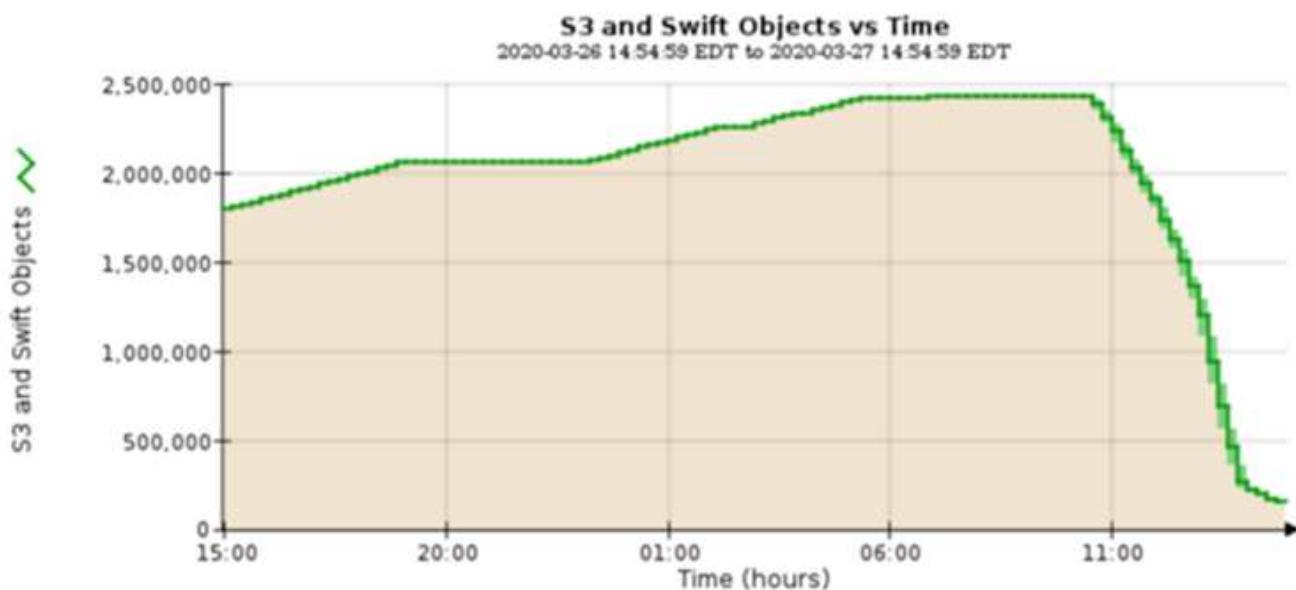
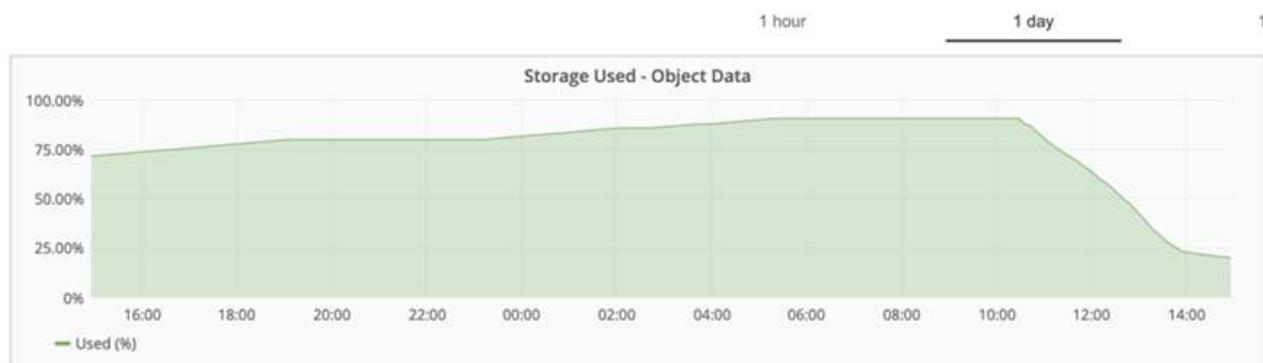


Figure 2: Freeing up 80TB of storage in less than 3 hours.



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.

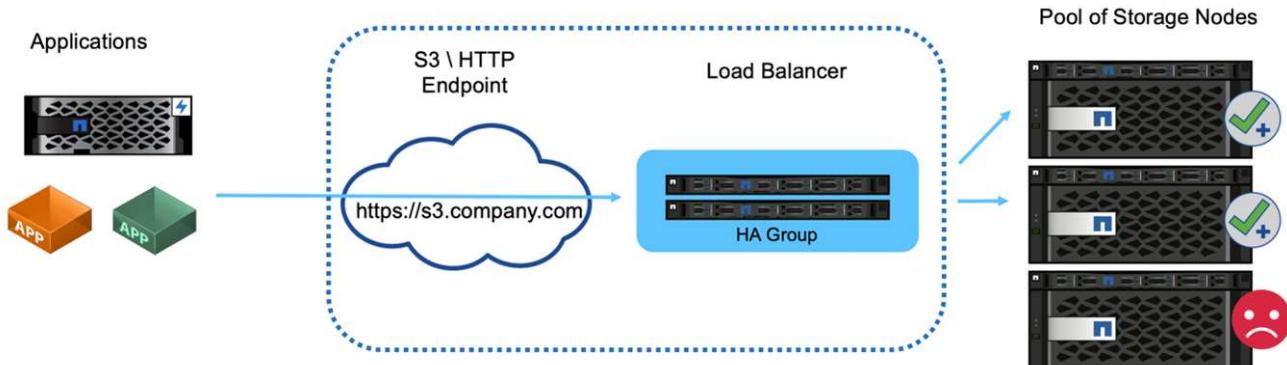
Object storage is synonymous with the term cloud storage, and, as you would expect, applications that leverage cloud storage address that storage through a URL. Behind that simple URL, StorageGRID can scale capacity, performance, and durability in a single site or over geo-distributed sites. The component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

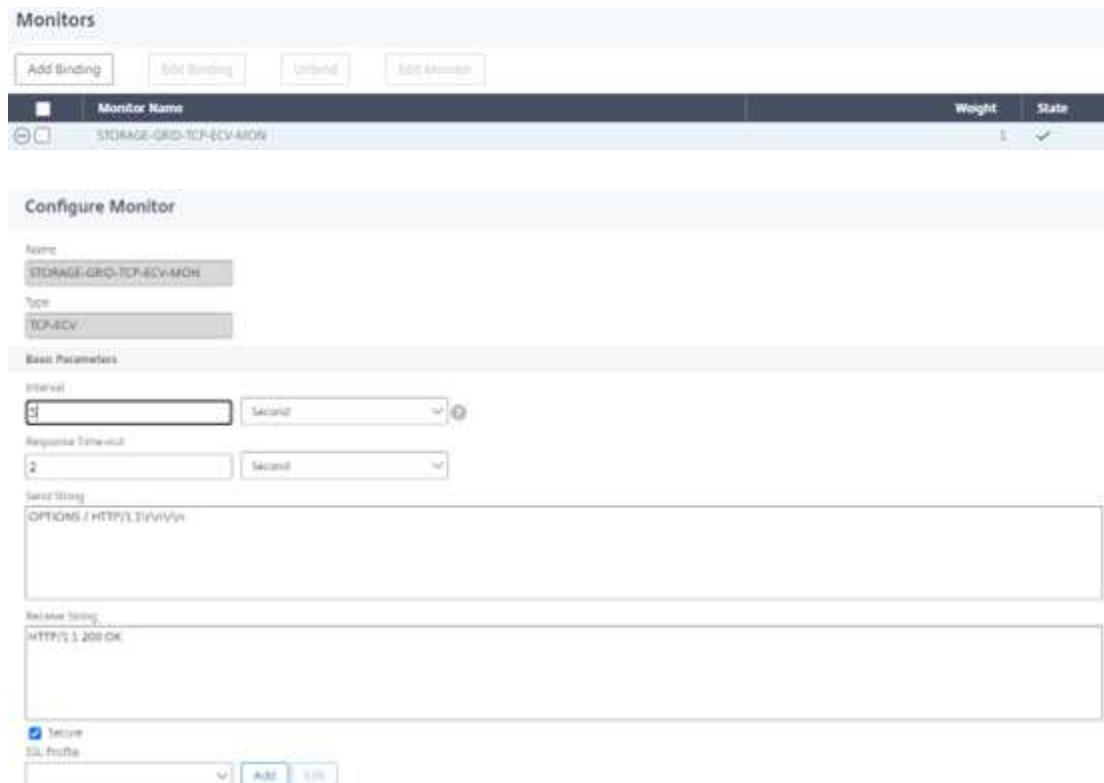
2. In Create Pool, create one pool for each port required.
 - a. Assign the health monitor that you created in the previous step.
 - b. Select a load-balancing method.
 - c. Select service port: 18082 (S3).
 - d. Add nodes.

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 interface in Citrix NetScaler. At the top, there are buttons for 'Add Binding', 'Edit Binding', 'Unbind', and 'Edit Monitor'. A table lists a single monitor entry:

Monitor Name	Weight	State
STORAGE-GRID-TCP-ECV-MON	1	✓

Below the table is a 'Configure Monitor' section with the following fields:

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

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.

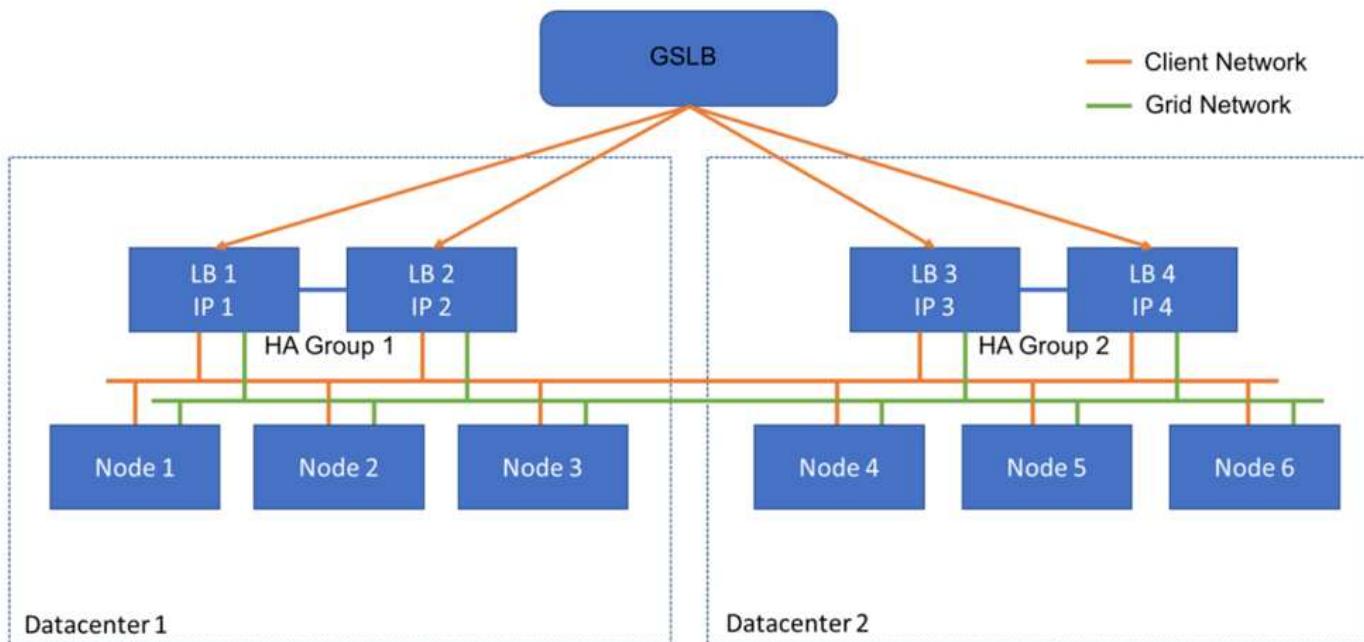
After your load balancer is configured, you should validate the connection using tools such as OpenSSL and the AWS CLI. Other applications, such as S3 Browser, might ignore SSL misconfiguration.

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	<p>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.</p>	<p>Helps secure storage and transmission of objects, independent of the underlying storage hardware.</p>	SEC Rule 17a-4(f) CTFC 1.31(c)-(d) (FINRA) Rule 4511(c)
Built-in Key Management	<p>When encryption is enabled, each object is encrypted with a randomly generated unique symmetric key, which is stored inside StorageGRID with no external access.</p>	<p>Enables encryption of objects without requiring External Key Management.</p>	
Federal Information Processing Standard (FIPS) 140-2 compliant encryption disks	<p>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.</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)

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.

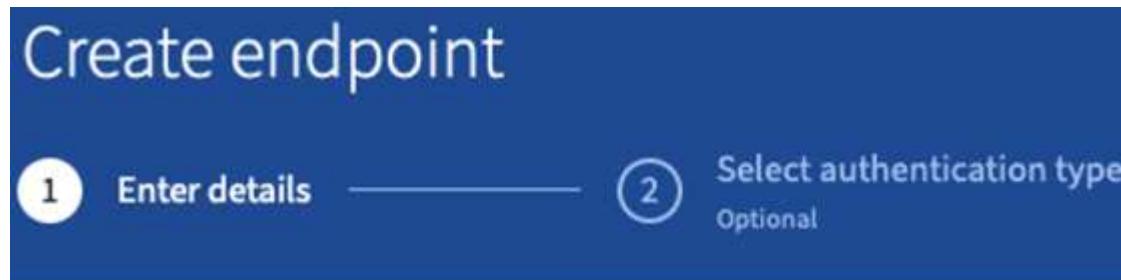
Because CloudMirror is an integrated component of StorageGRID, it cannot be turned off or manipulated by an S3 API-based attack. You can configure this replicated bucket with versioning enabled. In this scenario you need some automated cleanup of the replicated bucket's old versions that are safe to 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>
```

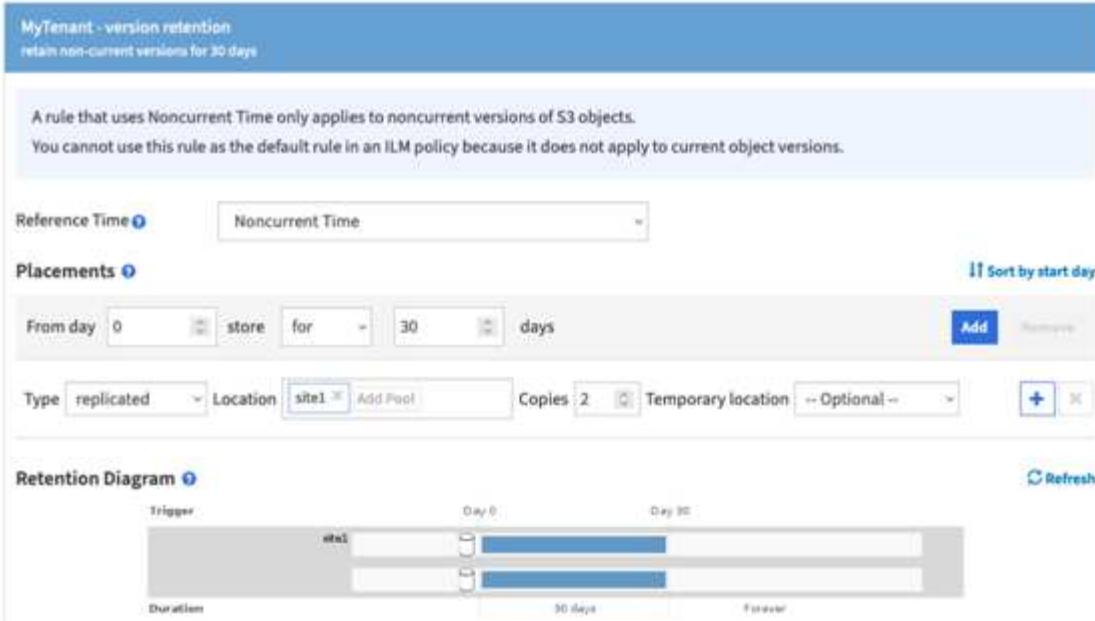
3. Create ILM rules to manage the storage placement and version storage duration management. In this example, the non-current versions of the objects to store are configured.

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

Type: replicated Location: site1 Add Pool: Copies: 2 Temporary location: Optional

Retention Diagram: Refresh

Trigger: site1 Day 0 Day 30 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

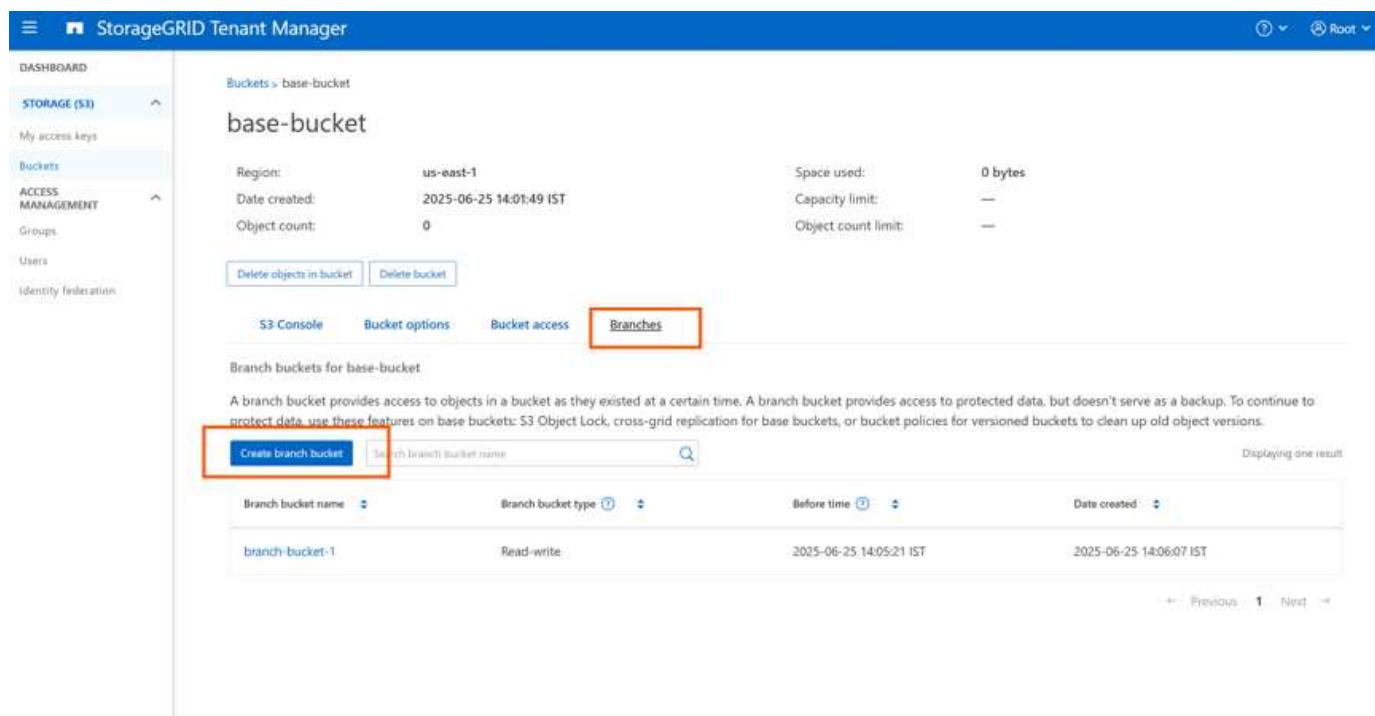
Learn how to investigate and remediate buckets after a possible ransomware attack with StorageGRID.

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 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 for Space used: 0 bytes, Capacity limit: —, and Object count limit: —. Below these are buttons for 'Delete objects in bucket' and 'Delete bucket'. The 'Branches' tab is highlighted with a red box. A sub-section titled 'Branch buckets for base-bucket' explains the purpose of branch buckets and shows a table with one result. The table has columns: Branch bucket name, Branch bucket type, Before time, and Date created. The result is 'branch-bucket-1' with 'Read-write' type, '2025-06-25 14:05:21 IST' before time, and '2025-06-25 14:06:07 IST' date created. A search bar for 'Search branch bucket name' is also highlighted with a red box. Navigation buttons for 'Previous' and 'Next' are at the bottom.

- Once the Create branch bucket button is clicked, a popup will open with prefilled 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

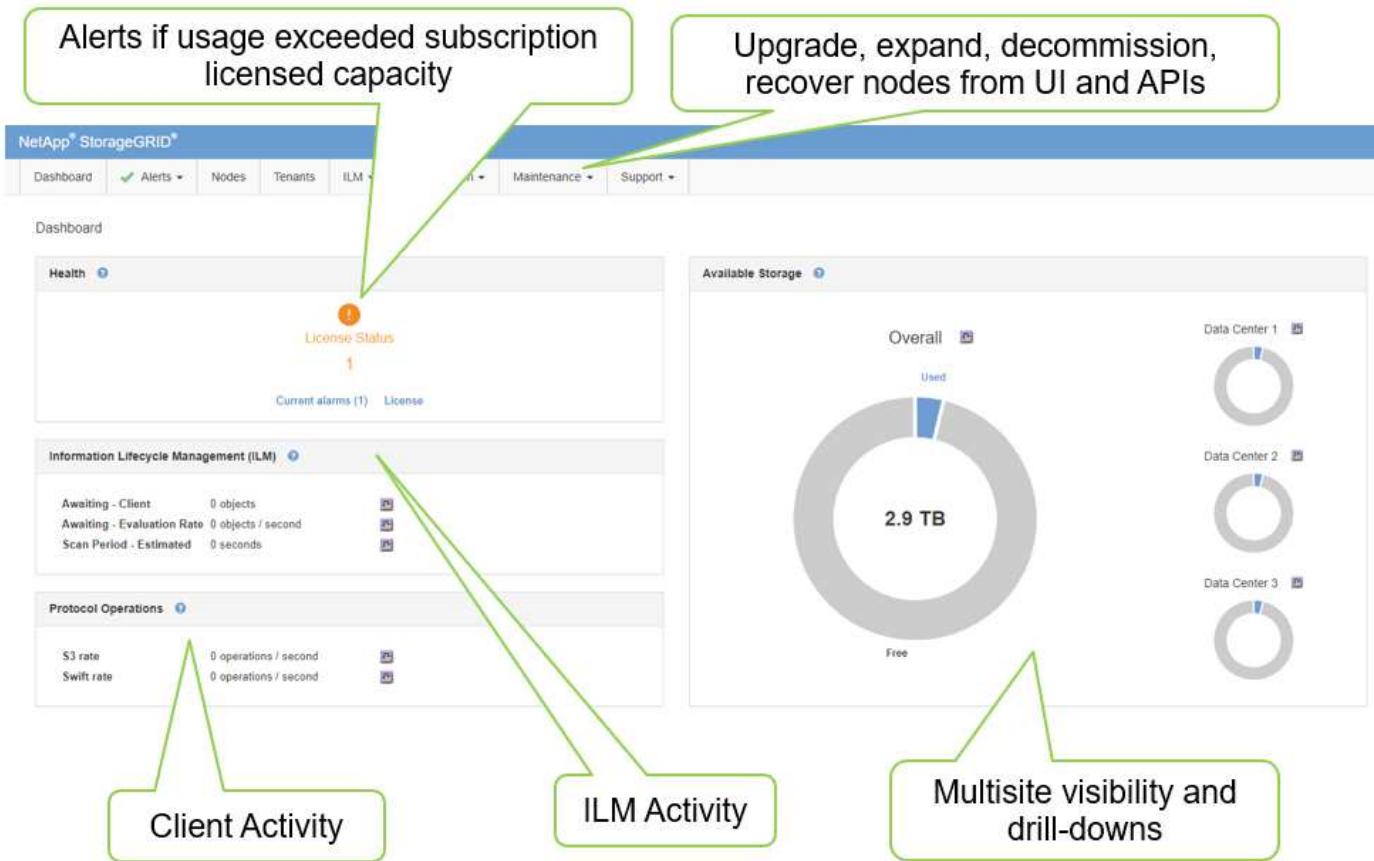
To learn more about the information that is described in this document, review the following documents and/or websites:

- 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.

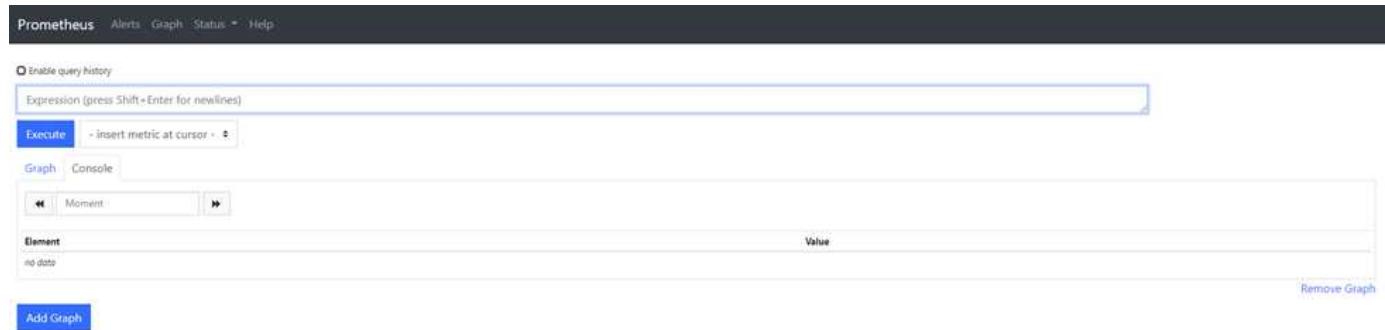
Access the Grafana dashboards using the links below. You must be signed in to the Grid Manager.

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
Identity Service Overview
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.

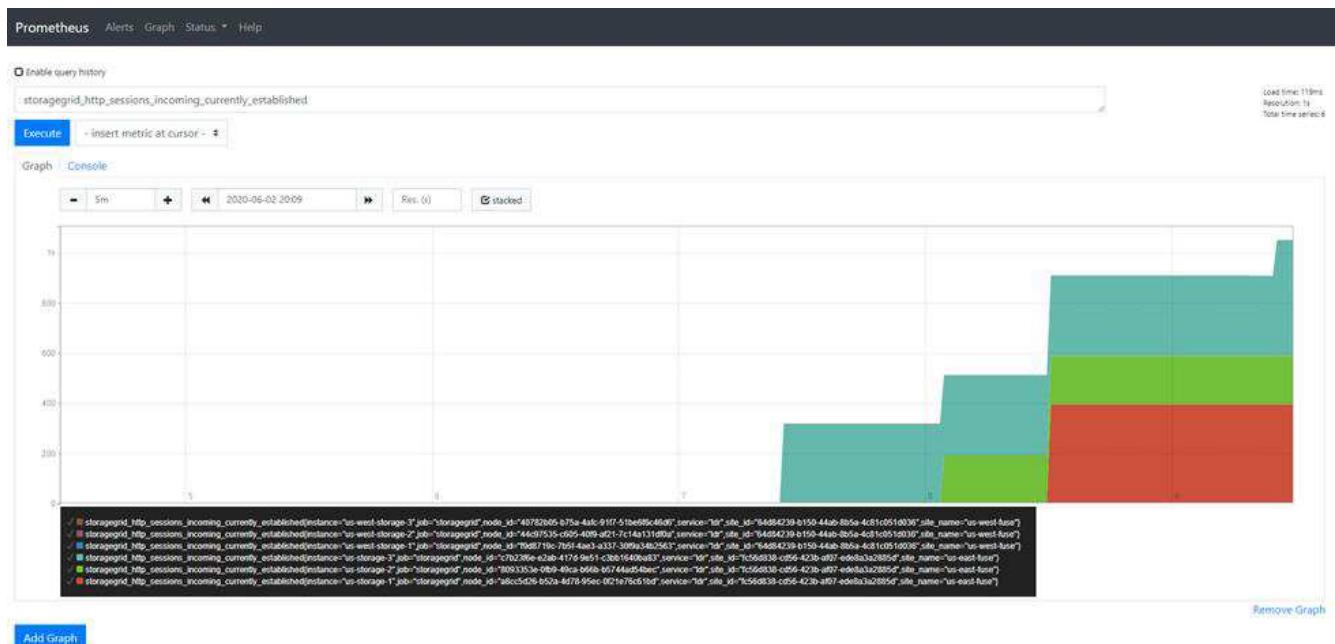


With this view, you can access the Prometheus interface. From there, you can search through available metrics and even experiment with queries.

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)
time	query start, default current time (date-time) (query)
timeout	timeout (duration) (query)

storagegrid_http_sessions_incoming_current
time - query start, default current time (date-t)
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

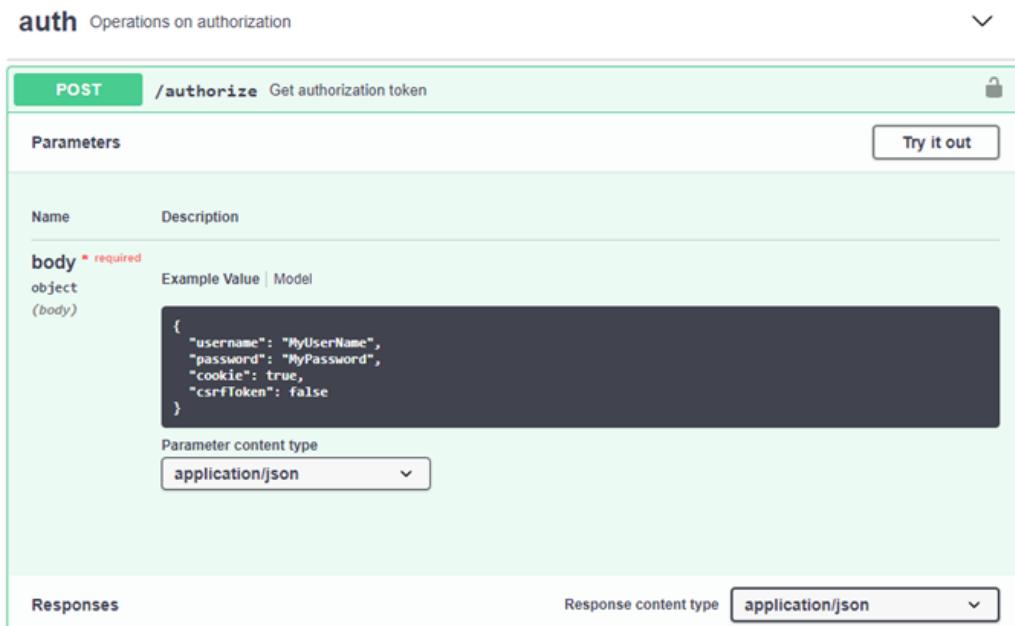
Access metrics using cURL in StorageGRID

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' section. The 'auth /authorize' POST method is selected. 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 Account Service Overview Alertmanager Audit Overview Cassandra Cluster Overview Cassandra Network Overview Cassandra Node Overview Cloud Storage Pool Overview EC Read - Node EC Read - Overview	Grid ILM Identity Service Overview 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 Traffic Classification Policy
---	--	---

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

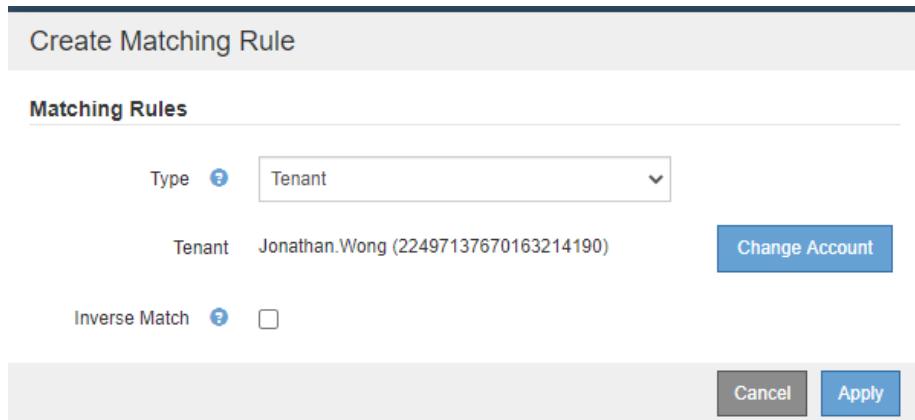
Matching Rules

Type ? Tenant

Tenant Jonathan.Wong (22497137670163214190) Change Account

Inverse Match ?

Cancel Apply



5. Set a limit (optional).

Create Limit

Limits (Optional)

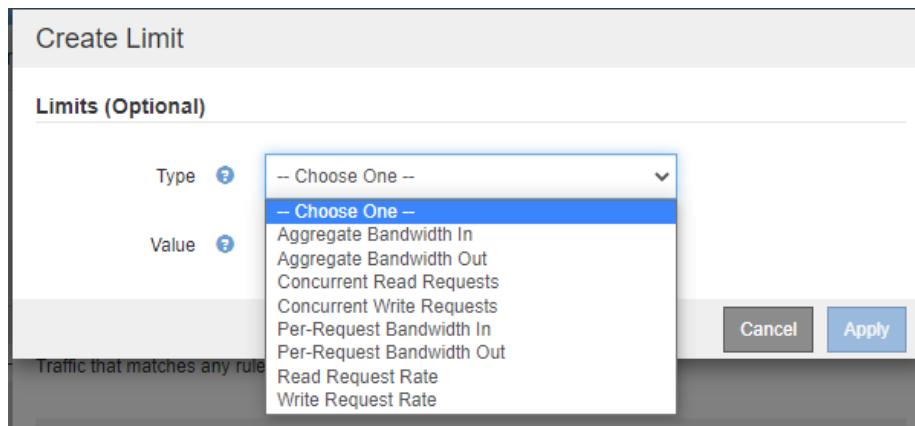
Type ? -- Choose One --

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

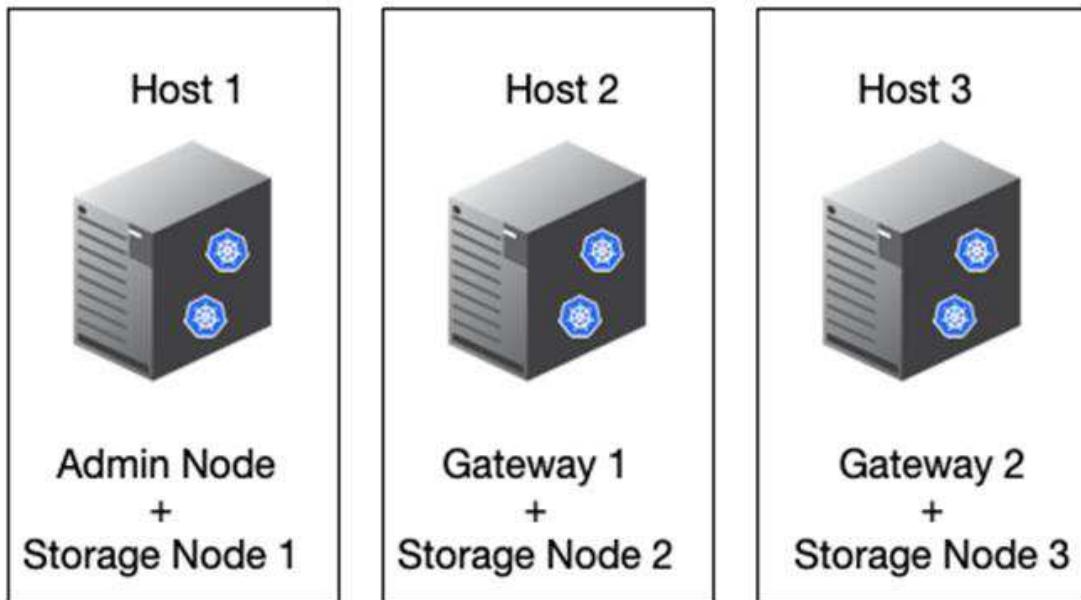
Gateway container

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

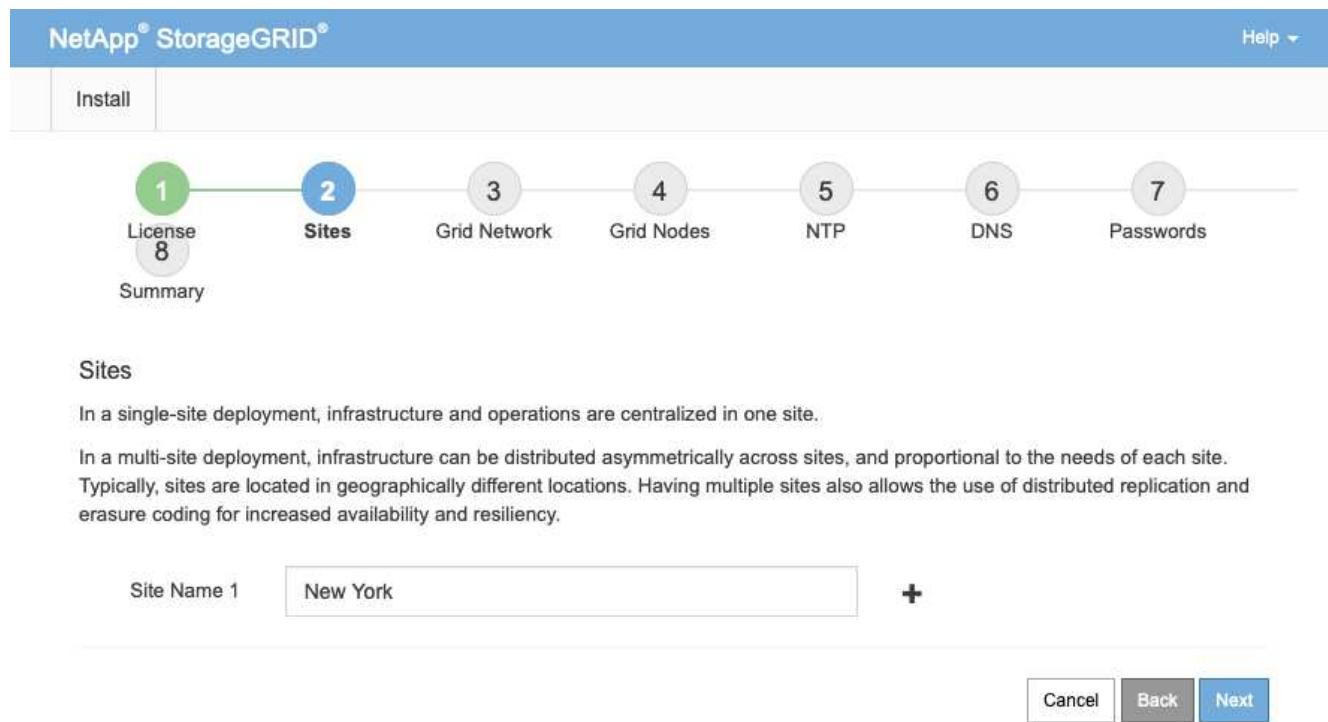
8 Summary

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

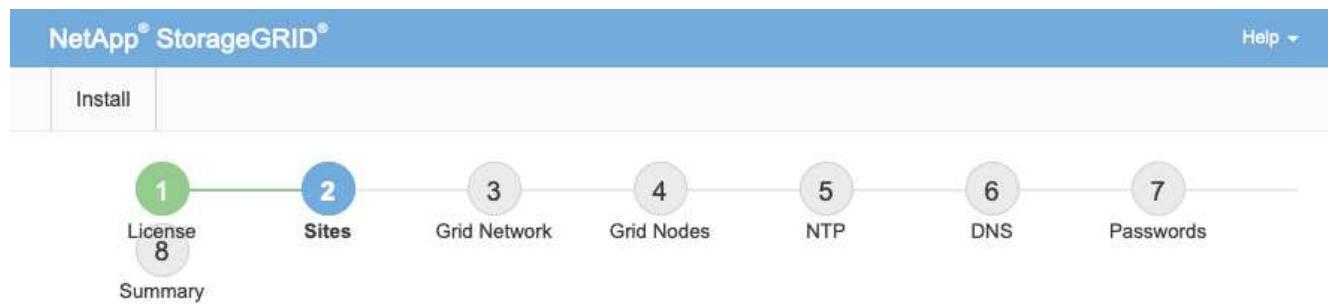
Learn how to add sites to StorageGRID to increase reliability and storage capacity.

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.



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

New York



Cancel

Back

Next

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.

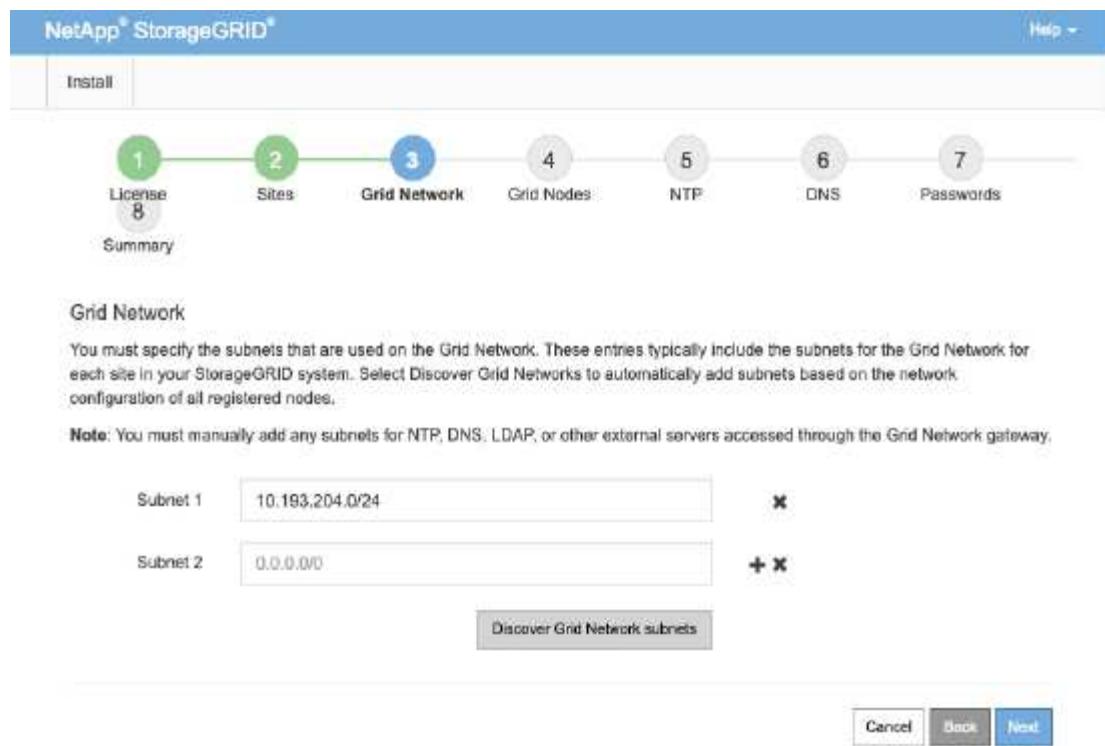
To specify grid network subnets, complete the following steps:

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.

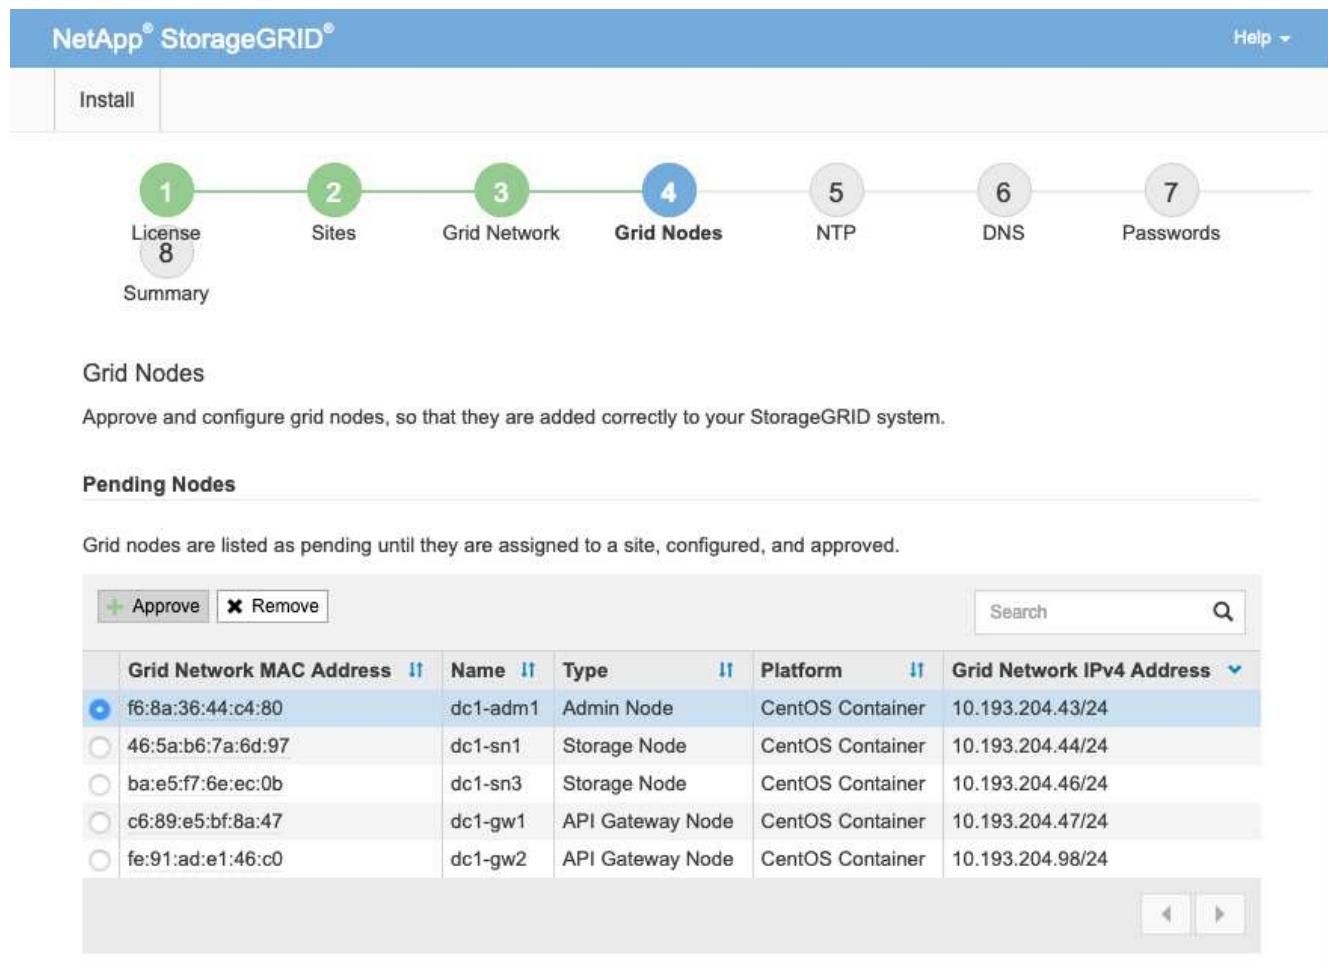
7. If you want to configure the Client Network for the grid node, add or update the settings in the Client Network section as necessary. If the Client Network is configured, the gateway is required, and it becomes the default gateway for the node after installation.

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 setup 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 not yet started (grey), Step 6 (DNS) is not yet started (grey), and Step 7 (Passwords) is not yet started (grey). Step 8 (Summary) is also completed (green). The main area is titled 'Grid Nodes' and contains a sub-section 'Pending Nodes'. A note says: 'Approve and configure grid nodes, so that they are added correctly to your StorageGRID system.' Below this, a table lists pending nodes:

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

At the bottom of the table are buttons for 'Approve' and 'Remove', and a 'Search' field with a magnifying glass icon. Navigation arrows are also present at the bottom of the table area.

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.

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

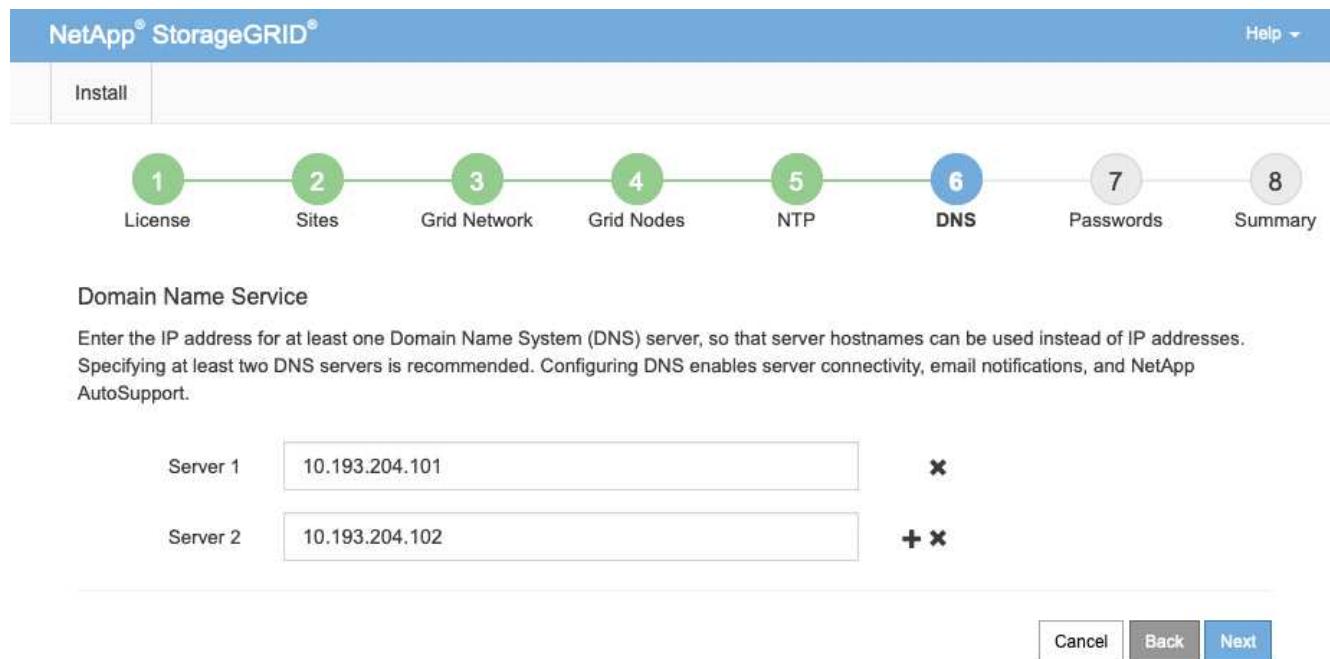
1 License 2 Sites 3 Grid Network 4 Grid Nodes 5 NTP 6 DNS 7 Passwords 8 Summary

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.

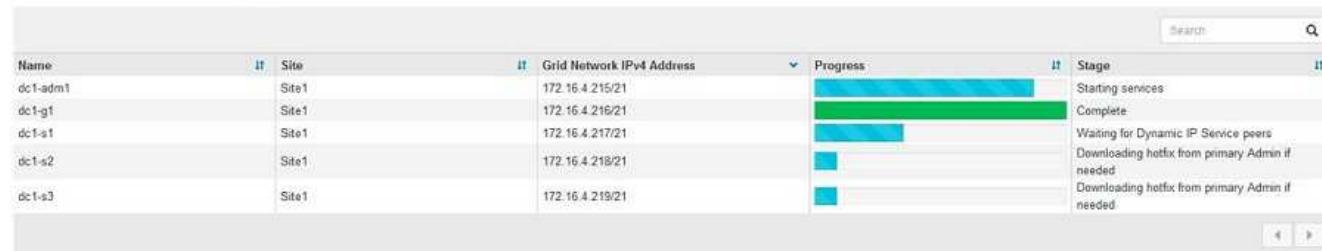
Download Recovery Package

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-a1, dc1-s2, and dc1-a3. 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-a1		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-a3		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

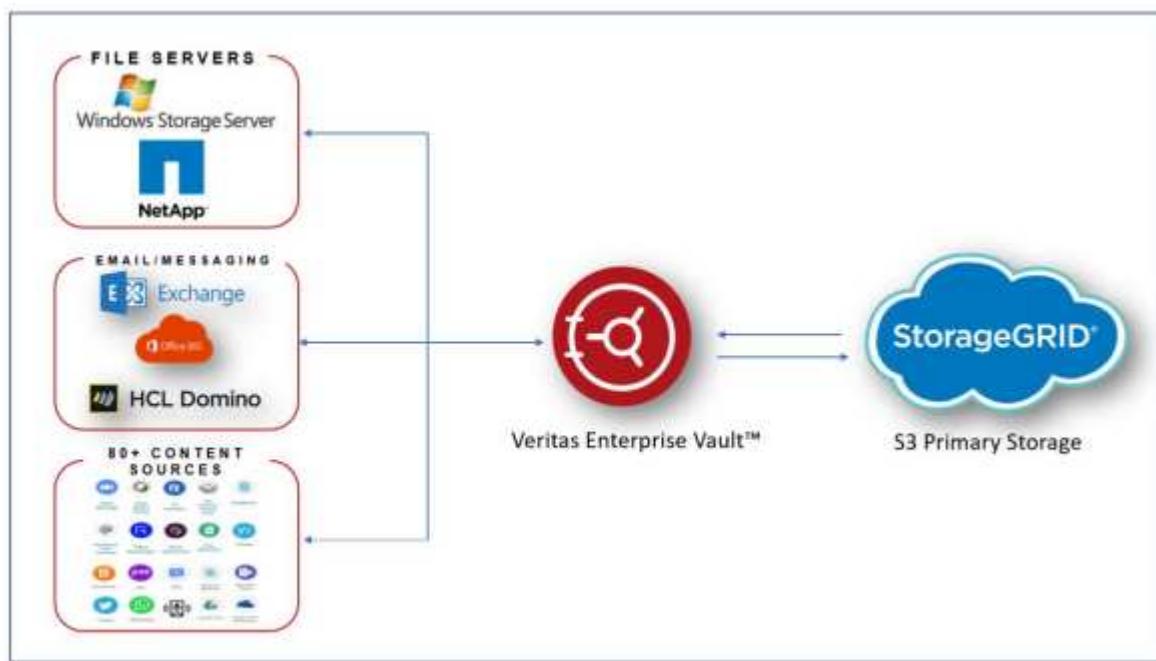
Introduction to configuring StorageGRID for site failover

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.

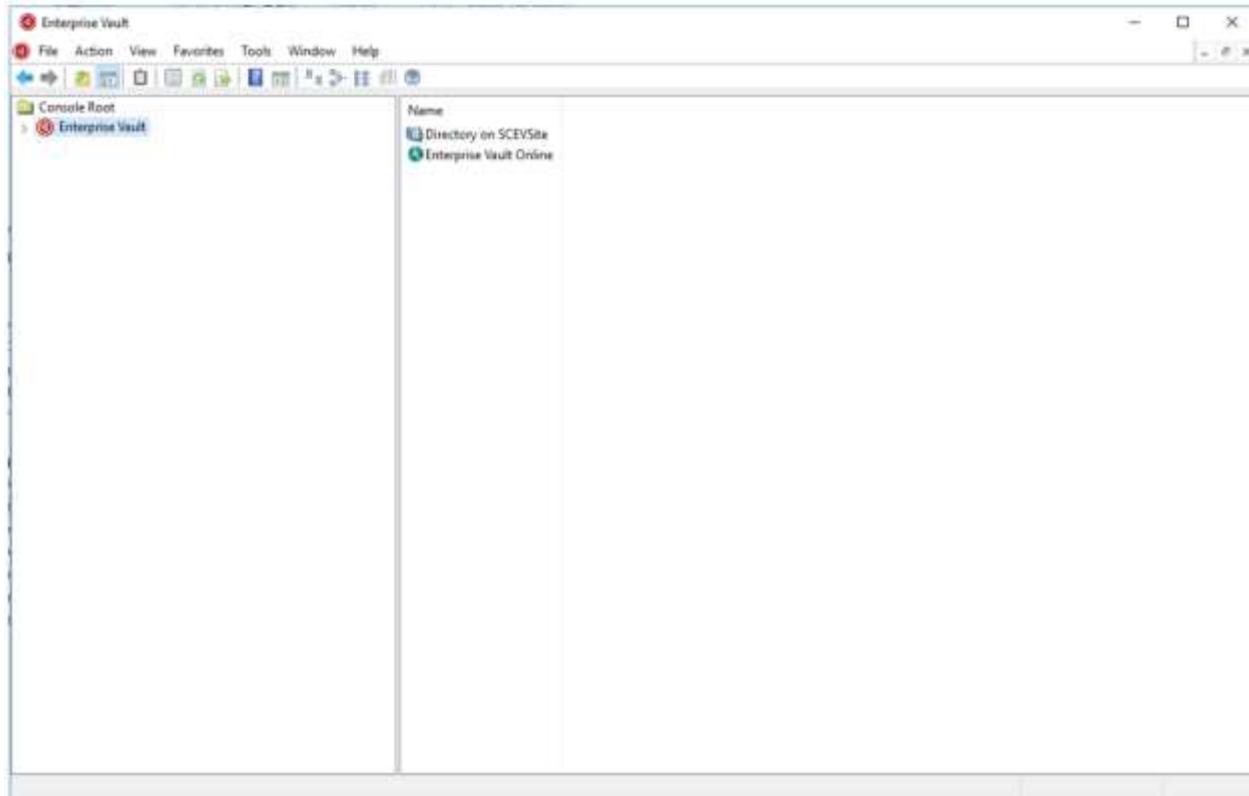
- Vault store groups and a vault store has been created.
For more information, see the veritas Enterprise Vault Administration Guide.
- A StorageGRID tenant, access key, secret key and bucket have been created.
- A StorageGRID load balancer endpoint has been created (either HTTP or HTTPS).
- If using 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

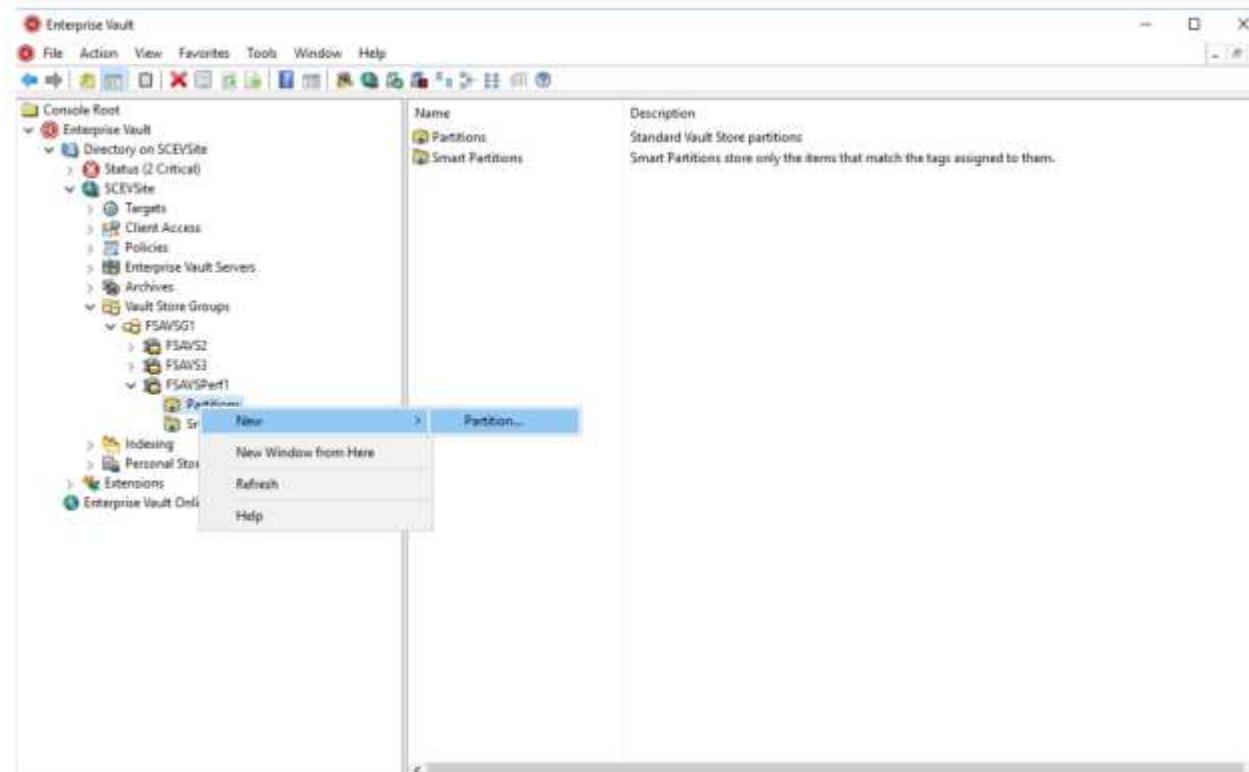
To configure StorageGRID with veritas Enterprise Vault, complete the following steps:

Steps

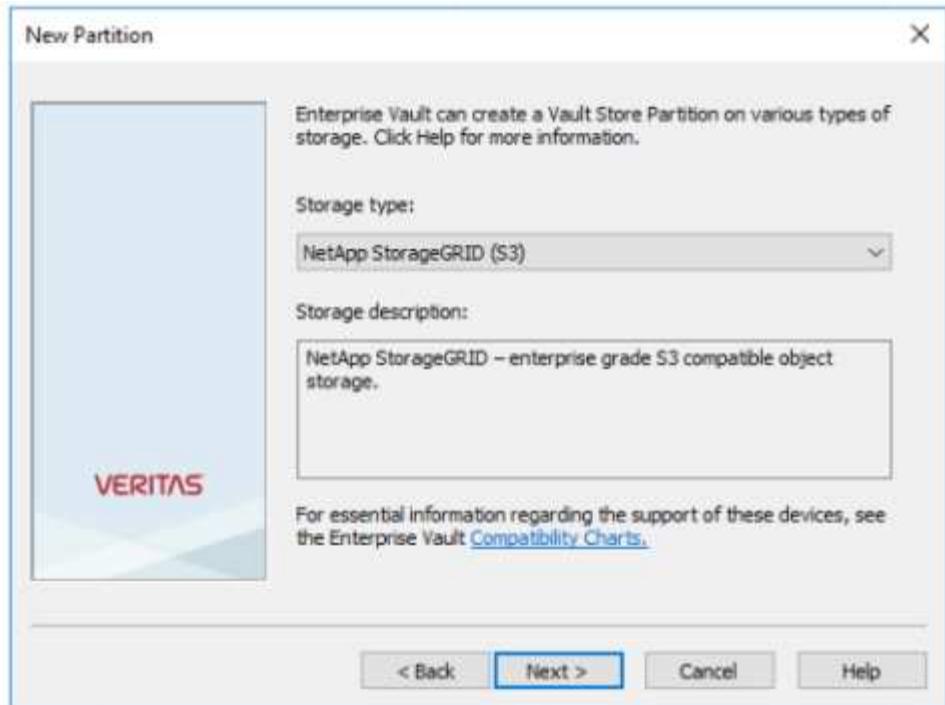
1. Launch the Enterprise Vault Administration console.



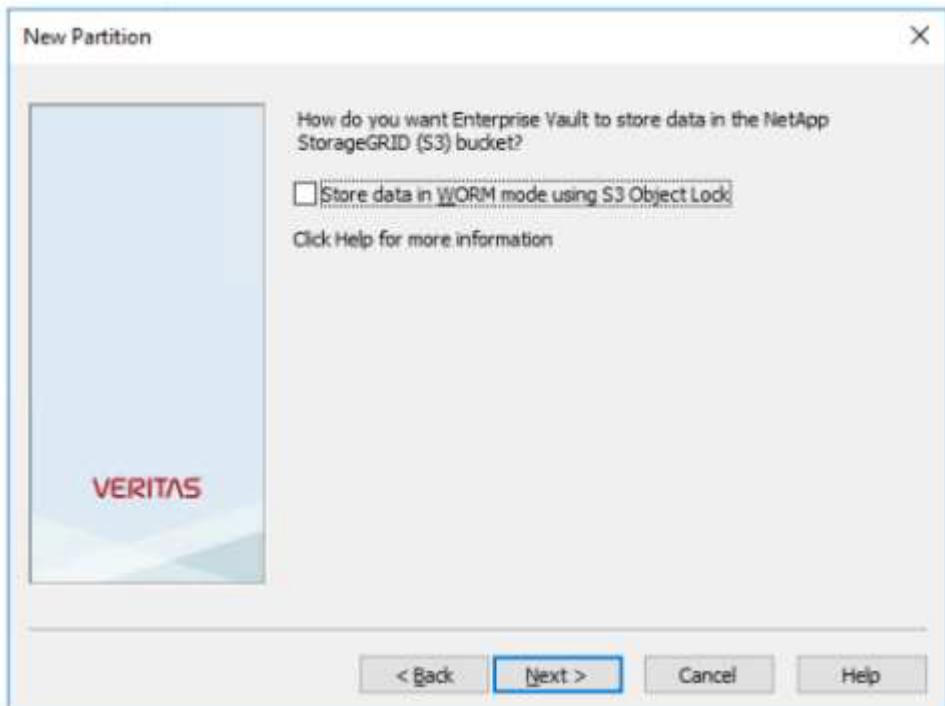
2. Create a new vault store partition in the appropriate vault store. Expand the Vault Store Groups folder and then the appropriate vault store. Right-click Partition and select **New > Partition**.



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



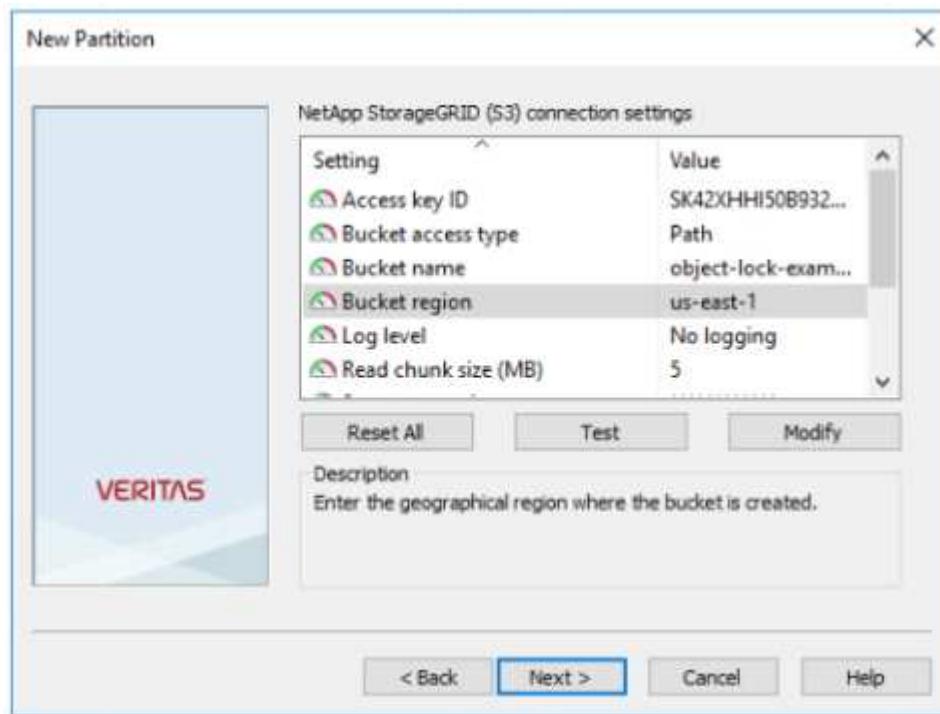
4. Leave the Store Data in WORM Mode Using S3 Object Lock option unchecked. Click Next.



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.

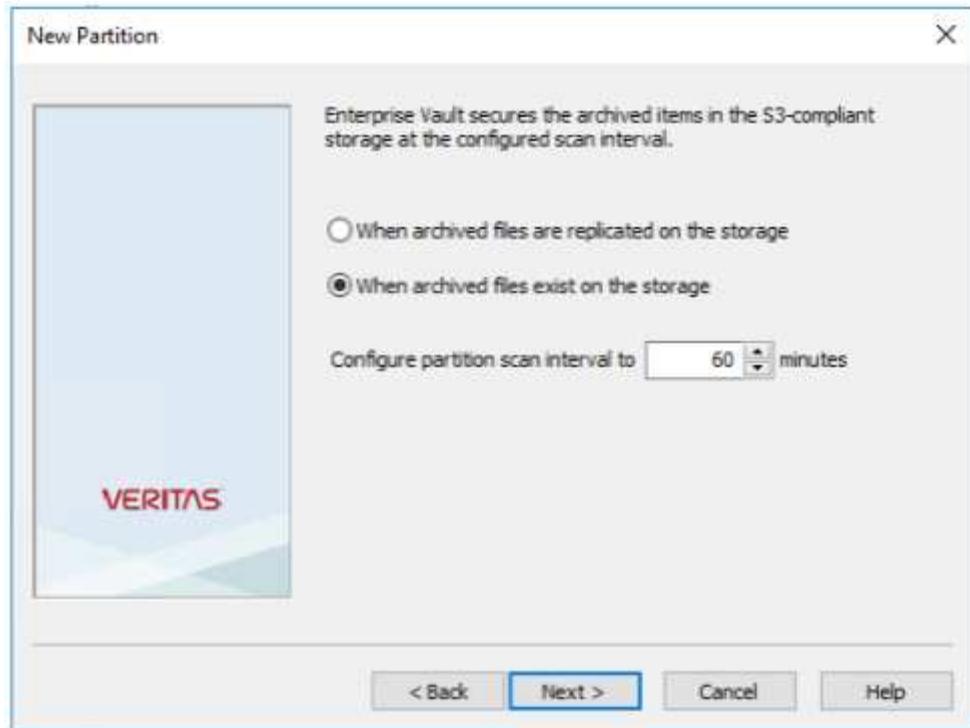
- Bucket region: us-east-1 is the default value.



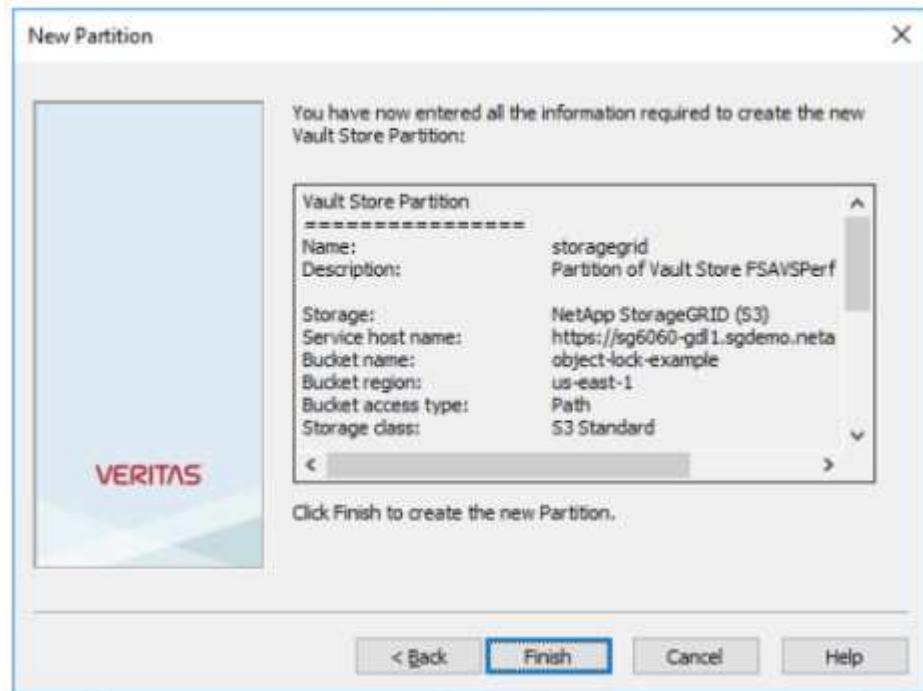
6. To verify the connection to the StorageGRID bucket, click Test. Verify that the connection test was successful. Click OK and then Next.



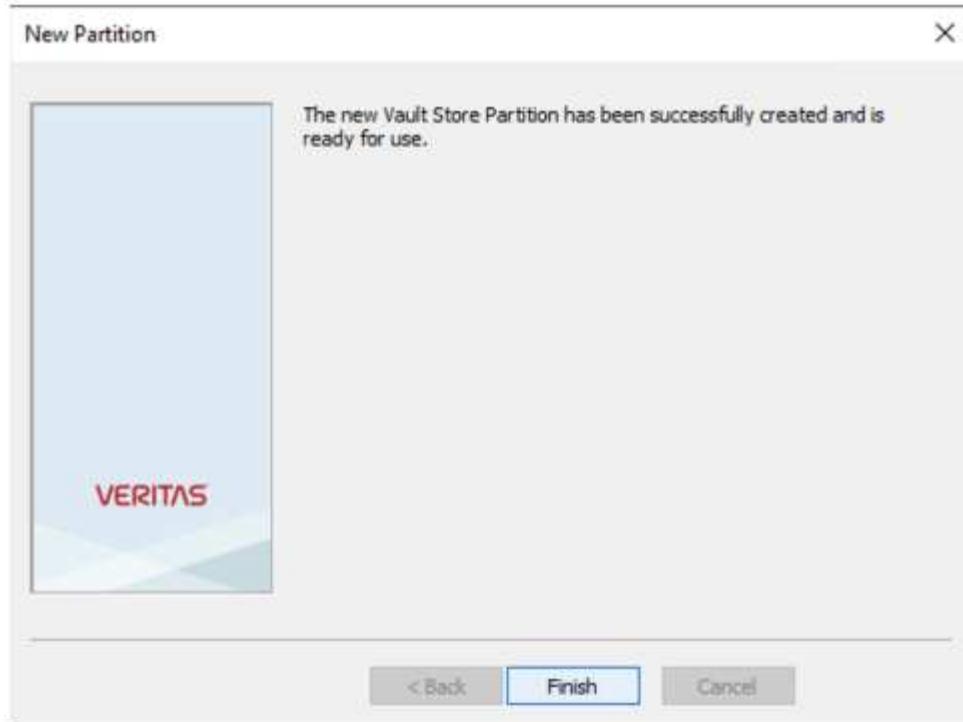
7. StorageGRID does not support the S3 replication parameter. To protect your objects, StorageGRID uses Information Lifecycle Management (ILM) rules to specify data protection schemes - multiple copies or erasure coding. Select the When Archived Files Exist on the Storage Option and click Next.



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

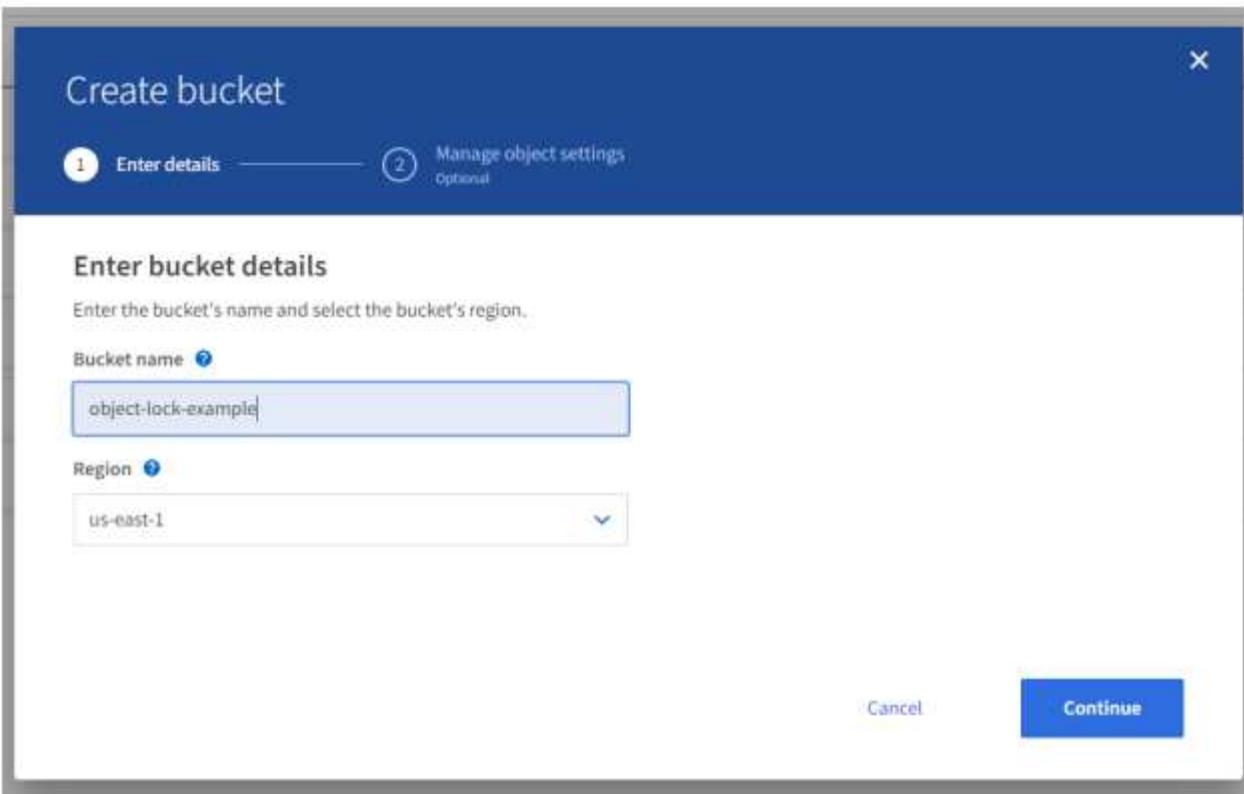
For WORM storage, StorageGRID uses S3 Object Lock to retain objects for compliance. This requires StorageGRID 11.6 or higher, where S3 Object Lock default bucket retention was introduced. Enterprise Vault also requires version 14.2.2 or higher.

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 Read-after-new-write (default)

Last access time updates Disabled

Object versioning Enabled

S3 Object Lock Enabled

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 Enabled

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
1 Days

[Save changes](#)

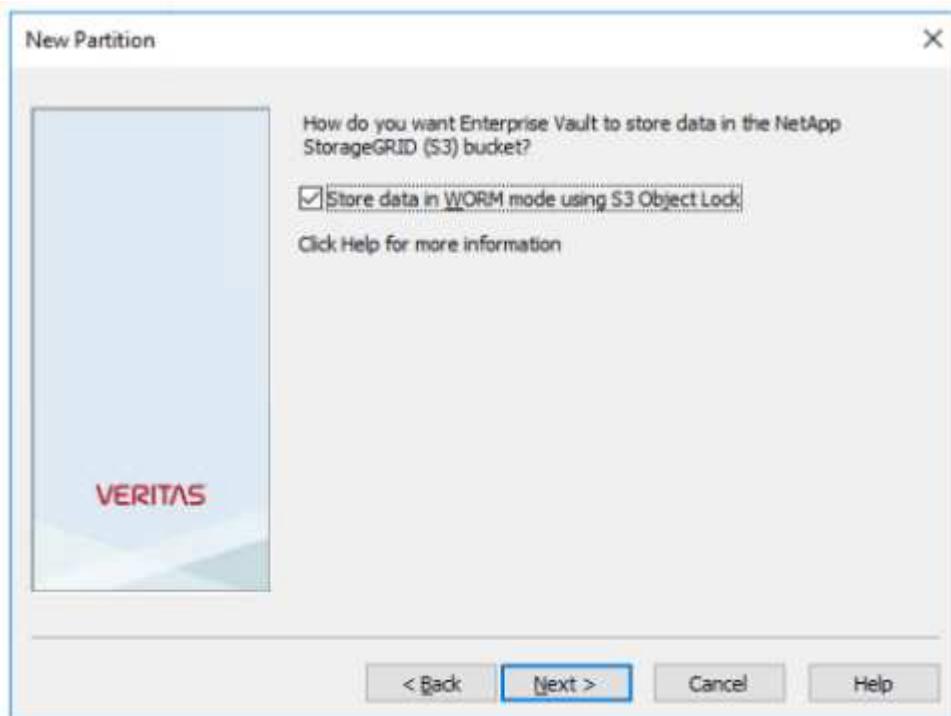
The bucket is now ready to be used by Enterprise Vault to store WORM data.

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	205.0.24	<input checked="" type="radio"/>
SITE2-ADM1	eth2	205.0.24	<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	+
----------------------	--------	---

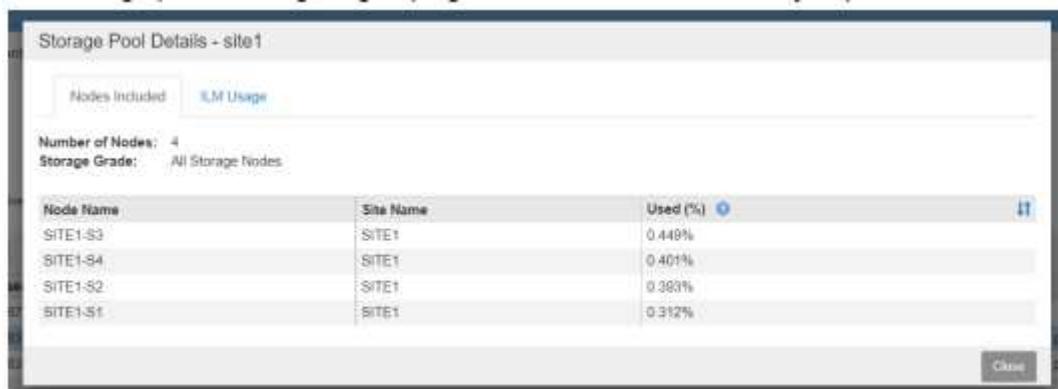
Buttons

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

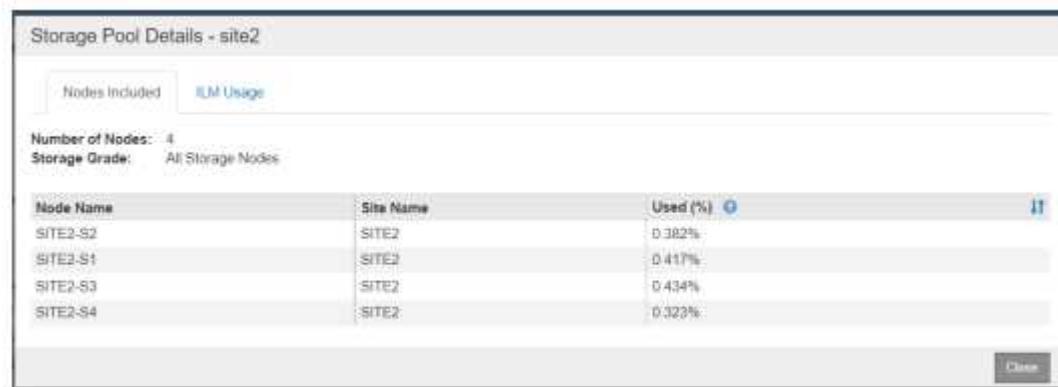


Storage Pool Details - site1

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:
Ingest Behavior: Balanced
Ingest Time: Ingest Time
Filter Criteria: Matches all objects

Retention Diagram:

Stage	Days	Days
Initial	0	1
Final	0	Forever

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.

Steps to access StorageGRID evaluation software

This instruction is for NetApp sales, partners, and prospects engaged with NetApp.

Register for an account

1. Register for an account on the [NetApp Support site](#) using your business email.
 - a. Make sure you are not signed in with the newly created account.
 - b. If you already have an account, make sure you are not signed in and proceed with the next step.
2. Create a non-technical support case to elevate access levels to "prospect." To do this, click on the "[Report an Issue](#)" link in the footer of the website.
3. Select "Registration Issue" as the feedback category.
4. In the comments section, write: "My account email address is *Your-Email-Address*. I would like to get prospect access to download the StorageGRID evaluation software."
 - a. Mention the name of the NetApp internal person who suggested the request for prospect access.

Download StorageGRID

1. After your support case has been reviewed and approved, NetApp support will notify you via email that your account has been granted prospect access.
2. Download the [StorageGRID evaluation software](#).



The Eval license file is located within the zip file. It is StorageGRID-Webscale-<version>\vsphere\NLF000000.txt once unzipped.



Downloading the software is a process that involves trade compliance measures to adhere to legal requirements. To ensure compliance, users are required to create an account and open a support case before gaining access. This process helps us maintain proper control and documentation while providing prospects with the production-ready software they need.

We provide the "production-ready" version of StorageGRID, which is not an open-source or alternative version. It is important to note that **support is not provided** unless the prospect upgrades to a production license.

Please contact StorageGRID.Feedback@netapp.com for any trouble with the above steps.

NetApp StorageGRID Blogs

You can find some great NetApp StorageGRID blogs here:

- Feb 16 2024: [Introducing StorageGRID 11.8: Enhanced security, simplicity, and user experience](#)
- Feb 16 2024: [Introducing StorageGRID 11.8](#)
- Feb 2 2024: [Announcing the StorageGRID + lakeFS Solution Brief](#)
- Dec 12 2023: [Big data analytics on StorageGRID: Dremio performs 23 times faster than Apache Hive](#)
- Nov 7 2023: [Spectra Logic On-Prem Glacier with StorageGRID](#)
- Oct 17 2023: [Moving on from Hadoop: Modernizing Data Analytics with Dremio and StorageGRID](#)
- Sep 1 2023: [Leveraging Cloud Insights to Monitor and Collect Logs Using Fluent Bit](#)
- Aug 30 2023: [Mountpoint for Amazon S3 File System is Now GA](#)
- May 16 2023: [Introducing StorageGRID 11.7 and the new all-flash object storage appliance SGF6112](#)
- May 16 2023: [What's new in the StorageGRID object storage family](#)
- Mar 30 2023: [Mountpoint for Amazon S3 alpha release with StorageGRID](#)
- Mar 30 2023: [Use BlueXP to protect Epic EHR with a 3:2:1 -compliant backup policy](#)
- Mar 14 2023: [How to back up Epic Systems EHR databases with one command in a 3:2:1-compliant architecture](#)
- Feb 14 2023: [What do chocolate, skiing, watches, and mainframes have in common?](#)
- Jan 18 2023: [StorageGRID S3 Object Lock validated for veritas NetBackup](#)
- Jan 16 2023: [StorageGRID renews NF203 and ISO/IEC 25051 compliance certification](#)
- Dec 6 2022: [StorageGRID achieves KPMG compliance certification](#)
- Nov 23 2022: [Explainable AI with MLOps powered by NetApp and Modzy](#)
- Nov 7 2022: [StorageGRID and ONTAP S3 support: Differences, similarities, and integration](#)
- Oct 5 2022: [NetApp Cloud Insights adds StorageGRID gallery dashboards](#)
- Oct 5 2022: [Defrost your data on StorageGRID for Snowflake](#)
- Sep 26 2022: [NetApp StorageGRID for service providers](#)
- Sep 19 2022: [DataLock and Ransomware Protection Support for StorageGRID](#)
- Sep 1 2022: [Take these Metrics and Graph it](#)
- Aug 23 2022: [Build your data lake on StorageGRID](#)
- Aug 17 2022: [It all starts with Object Locking... Building a S3 storage ecosystem for critical backup applications](#)
- Aug 16 2022: [Integrating StorageGRID with the open-source ELK stack to enhance customer experience](#)
- Aug 5 2022: [NetApp StorageGRID earns Common Criteria security certification](#)
- July 26 2022: [Check out the growing list of validated partner solutions for StorageGRID](#)
- June 9 2022: [Use Cloudera Hadoop S3A connector with StorageGRID](#)
- May 26 2022: [StorageGRID: storing and managing the on-premises backup and replication data](#)
- May 24 2022: [Modernize your analytics workloads with NetApp and Alluxio](#)

- May 10 2022: [Lab on demand is your best sales tool for StorageGRID](#)

NetApp StorageGRID documentation

You can find the complete documentation for each NetApp StorageGRID release here:

- [StorageGRID appliances](#)
- [StorageGRID software 11.5 - 12.0](#)

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

<https://www.netapp.com/company/legal/copyright/>

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<https://www.netapp.com/company/legal/trademarks/>

Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/pdf.html?item=/media/11887-patentspage.pdf>

Privacy policy

<https://www.netapp.com/company/legal/privacy-policy/>

Open source

Notice files provide information about third-party copyright and licenses used in NetApp software.

https://library.netapp.com/ecm/ecm_download_file/2879263

https://library.netapp.com/ecm/ecm_download_file/2881511

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.