



Technical FAQ

StorageGRID Webscale 11.0

Brian Atkins, Ronnie Chan, Aditya Kalyanakrishnan, Felix Lau, and Steven Pruchniewski

NetApp

October 2017

Abstract

This document answers frequently asked questions about NetApp® StorageGRID® Webscale 11.0 regarding the software and the appliance.

TABLE OF CONTENTS

1	Overview.....	8
1.1	Point of Contact.....	8
	If I need help, who should be my first point of contact?	8
1.2	Quick Start Guide	8
	How do I install StorageGRID Webscale?	8
2	StorageGRID Webscale 11.0	8
2.1	New Features and Upgrade	8
	What's new in StorageGRID Webscale 11.0?	8
	Is it possible to upgrade existing StorageGRID Webscale 10.4 customers to 11.0?	9
2.2	General Availability.....	9
	When will StorageGRID Webscale 11.0 be available?	9
3	Sales	9
3.1	Qualification.....	9
	What questions do I ask to qualify a NetApp StorageGRID Webscale proposal?	9
	Can a customer buy StorageGRID Webscale software if NetApp storage is already in place?	10
3.2	Licensing	10
	How is StorageGRID Webscale licensed?.....	10
	How do customers download the StorageGRID Webscale software license?.....	10
3.3	Perpetual Software License	10
	What is the StorageGRID Webscale perpetual software license?	10
	How do I calculate the perpetual software license required for storage nodes that use the StorageGRID appliance?.....	10
	How do I calculate the perpetual software license required for storage nodes that use other NetApp storage or third-party storage?.....	10
	Is there an evaluation license?	10
	How do I consider Amazon S3 (Simple Storage Service) or tape archiving when I calculate licensing capacity?	11
	Do I need to obtain a feature product variance request (FPVR) when I sell StorageGRID Webscale licenses to customers who want to run virtual storage nodes on third-party storage?	11
	Can I convert an evaluation license to a fully supported production license?	11
3.4	Software Subscription License	11
	What is the StorageGRID Webscale software subscription license?.....	11
	Is there a minimum capacity and subscription term?.....	11
3.5	Quoting.....	11
	How do I quote StorageGRID Webscale?	11
3.6	Demonstration and Proof of Concept	12
	How can I demo StorageGRID Webscale?.....	12
4	API support.....	12

4.1	Supported APIs	12
	What RESTful APIs does StorageGRID Webscale support?	12
4.2	S3 API Support.....	12
	What S3 API features are implemented in StorageGRID Webscale?	12
	What new S3 API features are added to StorageGRID Webscale 11.0?	12
4.3	SWIFT API Support.....	12
	What Swift API features are added in StorageGRID Webscale 11.0?.....	12
5	Erase Coding.....	13
5.1	Feature Product Variance Request (FPVR) Requirement	13
	Do I need to obtain an FPVR to enable erasure coding on StorageGRID Webscale?	13
5.2	Definitions.....	13
	What is erasure coding?	13
	What is geo-distributed erasure coding?	13
	What is layered erasure coding?	14
5.3	Value Proposition	15
	What is the value proposition of layered erasure coding in StorageGRID Webscale?	15
5.4	Erase Coding Schemes	15
	What EC schemes are supported by StorageGRID Webscale?.....	15
	What is the limit on the number of erasure coded objects?	15
	What is the minimum object size supported for EC?	16
6	Infrastructure	16
6.1	Virtualization.....	16
	What are the virtual machine requirements for StorageGRID Webscale?.....	16
	What are the storage requirements for virtualized component nodes in StorageGRID Webscale?	16
	What hypervisors can StorageGRID Webscale be deployed on?	16
	What operating system is StorageGRID Webscale deployed on?.....	16
6.2	Bare-Metal Deployment.....	17
	Can I install StorageGRID Webscale without a hypervisor directly on an x86 server?	17
	What are the system requirements for deploying StorageGRID Webscale on RHEL or CentOS?	17
	What types of storage can StorageGRID nodes use when deployed as a bare-metal deployment?	17
	Is any automation available for deploying StorageGRID as a bare-metal deployment?	17
6.3	Monitoring.....	17
	How do I monitor StorageGRID Webscale?	17
6.4	OpenStack Support	18
	How can I use StorageGRID Webscale in an OpenStack cloud?	18
	How can StorageGRID be deployed in an OpenStack environment?	18
	Does StorageGRID Webscale support OpenStack Keystone as an identity source?	18
6.5	Networking	18

In a logical StorageGRID data center, do the admin node, gateway node, and all storage nodes need to be in the same VLAN/IP subnet?	18
What Is the Grid, Admin, Client Network Model?	19
Does the StorageGRID appliance have VLAN support?	19
Does the StorageGRID appliance have Link Aggregation Control Protocol (LACP) support?	19
Are MTU 9000 (jumbo frames) required?	19
7 Design.....	20
7.1 Grid Architecture	20
How do I size NetApp StorageGRID Webscale?	20
What does a basic StorageGRID deployment consist of?	20
Should there be separate grids for specific operational models, such as production, nonproduction, test, backup, and so on?	20
Can a customer merge multiple grids into a single grid?	20
In a deployment with multiple admin nodes, does a secondary node take over the role of a primary node in if a primary node fails?	20
How many objects can a single storage node handle?	20
What are the maximum and minimum object sizes supported?	21
7.2 Expansion.....	21
Can a customer convert a single data center deployment to a multi data center deployment?	21
What are some expansion scenarios?	21
7.3 Tiering	21
Is there support for tiering out to the public cloud?	21
What is the capacity that can be tiered to the cloud using the archive node?	21
Is a capacity license required to cover capacity on a tape or cloud tier?	21
8 NAS Bridge.....	21
8.1 Protocols	21
What protocols does the NAS Bridge support?	21
Can I use both NFS and SMB protocols at the same time?	22
Can I use the NAS Bridge with any S3-based storage?	22
8.2 Deployment	22
How is the NAS Bridge deployed?	22
What are the virtualization requirements for the NAS Bridge?	22
How many instances of the NAS Bridge can I deploy?	22
Because all files are stored in StorageGRID Webscale, can two NAS Bridge instances access the same files?	22
8.3 Workloads	22
What workloads are appropriate for the NAS Bridge?	22
When the NAS Bridge updates an existing file, does it have to read the entire associated object into cache first, then write it back out?	22
What file sizes does the NAS Bridge supports?	22
Does the NAS Bridge support file-object duality?	23

How do I find my file's key so I can get it directly from S3?	23
But I'm using SMB. How do I find my file's key?	23
Does the NAS Bridge use more than one object per file?	23
Are modifications of a file directly via S3 supported?	23
9 Multitenancy.....	23
9.1 Overview	23
What multitenancy capabilities does NetApp StorageGRID Webscale have?.....	23
What features are added in 11.0?.....	23
9.2 LDAP and Active Directory Support	24
How is LDAP/AD support implemented in StorageGRID Webscale?	24
What options do users have to import LDAP/AD groups?	24
How does StorageGRID Webscale synchronize with the LDAP server?.....	24
9.3 OpenStack Keystone support.....	24
9.4 Limits	24
What are the multitenancy limits?	24
Is there cross-API access between S3 and Swift accounts?	25
9.5 Credential Management	25
How are the various credentials stored and protected in StorageGRID Webscale?	25
9.6 Reporting and Chargeback	25
What reporting and chargeback features does StorageGRID Webscale have?.....	25
Can I see these metrics on the Grid Management Interface?	25
10 Information Lifecycle Management (ILM)	25
What is ILM?	25
What has been updated in ILM for 11.0?	26
How many ILM rules can a StorageGRID Webscale deployment have?.....	26
Can ILM policies be set for each tenant in a multitenant deployment?	26
Can I retrieve an object that has been deleted, but not purged from the grid?	26
11 Versioning	27
What is versioning?.....	27
12 Performance.....	27
What questions do I ask my customers to gather the correct information about the workload that is intended for a NetApp StorageGRID Webscale environment?	27
What performance should I expect from StorageGRID Webscale?.....	28
13 Partner integration	28
Where can I find information about partner solutions with NetApp StorageGRID Webscale?	28
Where can I find more information about AltaVault and StorageGRID Webscale integration?	28
14 StorageGRID Webscale Appliance	28

14.1 Overview	28
What is the NetApp StorageGRID Webscale appliance?	28
What is the value proposition of a StorageGRID Webscale appliance?	29
14.2 Configuration	29
What are the available models in the SG5700 family?	29
What are the available configurations?	29
How is the appliance configured?	29
14.3 Software	30
Which versions of StorageGRID Webscale are supported with the appliance?	30
What is the pre-grid environment (PGE)?	30
What are the enhancements in the new PGE 2.0?	30
How are software upgrades performed on the appliance?	30
How do I upgrade the firmware on the appliance?	30
14.4 Hardware Infrastructure.....	30
Are any physical servers required?.....	30
Do I still need a hypervisor?.....	30
The SG5660 8TB and 10TB configurations exceed the 256TB limit of a VM storage node. Is that OK? ..	31
How does the appliance communicate with the grid?	31
Is there a concern with saturating a single 10GbE connection on a node, given the network connection for the appliance?	31
Is the appliance highly available?	31
Does the appliance have support for LACP and VLAN?	31
How can I use the 7-segment display on the appliance to find out the IP address of the appliance?	31
What does the customer receive? Are there additional installation tasks?.....	31
Are SSD configurations available?	31
Are partially populated configurations supported for the SG5760?	31
14.5 Performance.....	32
What about performance?	32
Would the appliance be sized differently than a VM storage node?.....	32
14.6 Miscellaneous.....	32
Can a grid consist of a mix of appliances and VM storage nodes?	32
How is it StorageGRID licensed?.....	32
Are expansion shelves supported?.....	32
15 StorageGRID 9	32
15.1 StorageGRID 9 Comparison	32
What is the difference between NetApp StorageGRID 9 and StorageGRID Webscale?	32
What is the future of StorageGRID 9?	32
15.2 StorageGRID 9 Transition	32
Is it possible to upgrade existing StorageGRID 9 customers to StorageGRID Webscale?	32
What is the transition plan for StorageGRID 9 customers to StorageGRID Webscale?	33

Version History 34

LIST OF TABLES

Table 1) Internal mailing lists. 8

Table 2) Supported EC schemes..... 15

LIST OF FIGURES

Figure 1) Introduction to erasure coding. 13

Figure 2) Introduction to geo-distributed erasure coding. 14

Figure 3) Layered erasure coding: erasure coding at the software and the hardware level. 14

1 Overview

1.1 Point of Contact

If I need help, who should be my first point of contact?

Answer: For any questions about NetApp StorageGRID Webscale, partners should contact their NetApp sales engineer. You can also find useful information on the [StorageGRID Webscale product page](#) in the Field Portal.

NetApp field team: You can access the internal [StorageGRID Community](#) page for more information about technical content, third party-application integration, and how-to guides. Table 1 describes internal mailing lists that are available for further help.

Table 1) Internal mailing lists.

Mailing List	Audience
ng-storagegrid-ses@netapp.com	This mailing list should be the first point of contact for questions related to sales that include sizing, competition, StorageGRID partner information, and so on.
ng-storagegrid-tme@netapp.com	This mailing list includes technical marketing engineers and product management. Use it for technical information and feature requests.

1.2 Quick Start Guide

How do I install StorageGRID Webscale?

Answer: Refer to the StorageGRID Webscale installation guides in the [NetApp StorageGRID Webscale Resources](#) page, which give step-by-step instructions on how to install StorageGRID Webscale across various platforms.

2 StorageGRID Webscale 11.0

2.1 New Features and Upgrade

What's new in StorageGRID Webscale 11.0?

Answer: The 11.0 release continues to improve on the StorageGRID Webscale object storage platform by adding a number of new features and enhancements to the existing feature set. New features include:

- Platform services that enable truly hybrid cloud workflows:
 - S3 Bucket replication. Mirror objects in native format to S3 destinations, such as AWS.
 - S3 Notification support. Publish notifications to AWS Simple Notification Service (SNS) to trigger cloud workflows.
 - Elastic search integration. Index object metadata in real time to elastic search clusters on premises or in the cloud.
- S3 Object tagging support. Classify objects with key-value pairs by using the S3 object tagging API.

- Location constraint support. ILM rules can evaluate location constraints on buckets created with a region parameter.
- Metrics management engine (powered by Prometheus). Build customized status and health dashboards by using new metrics management APIs.
- Log collection UI. Simplified and granular log collection on Grid Management Interface (GMI).
- StorageGRID appliance updated (release date November 2017):
 - New SG5700 appliance platform with more CPU and memory per appliance
 - Firmware update (PGE 2.0) for a better installation experience
- NAS Bridge 2.0.3:
 - Improved large file read and write performance

For more information, refer to the following resources [StorageGRID Webscale 11.0 SE presentation](#).

NetApp field team: You can also refer to the [11.0 field deep-dive presentation](#).

Is it possible to upgrade existing StorageGRID Webscale 10.4 customers to 11.0?

Answer: Yes. Existing 10.4 customers can upgrade to 11.0. Customers have the option either to use the automated upgrade process from the Grid Management Interface of the primary admin node or to upgrade manually if the automated upgrade is not applicable to their environment. For step-by-step instructions, refer to the upgrade and expansion section of the [NetApp StorageGRID Webscale Resources](#) page.

Notes:

1. Apply any applicable hotfixes to 10.4 before upgrading to 11.0.
2. Upgrade from 10.3 directly to 11.0 is not supported.
3. You must read the upgrade guide in full to understand the requirements and the steps that are needed to perform the upgrade.

2.2 General Availability

When will StorageGRID Webscale 11.0 be available?

StorageGRID Webscale 11.0 is generally available (GA) as of October 2017.

For all documentation, go to the [NetApp StorageGRID Webscale Resources](#) page on the NetApp Support site.

3 Sales

3.1 Qualification

What questions do I ask to qualify a NetApp StorageGRID Webscale proposal?

Answer: NetApp field team: See the blog [10 Qualifying Questions for StorageGRID Webscale](#) and the vertical and horizontal targeting resource [StorageGRID Account Targeting Strategy](#).

Can a customer buy StorageGRID Webscale software if NetApp storage is already in place?

Answer: Yes. StorageGRID Webscale can be deployed as a software-only deployment on existing NetApp storage.

StorageGRID Webscale can be flexibly deployed across heterogeneous hardware platforms, allowing customers to have a single deployment with a mix of StorageGRID Webscale appliance (SGA) nodes and software-only nodes.

3.2 Licensing

How is StorageGRID Webscale licensed?

Answer: StorageGRID Webscale software can be licensed in two ways:

- Perpetual software license, Per terabytes (TB) raw capacity. (SSP purchased separately.)
- Software subscription license. Per terabytes (TB) used capacity. (Includes SSP.)

For more information, go to the [NetApp licensing website](#).

How do customers download the StorageGRID Webscale software license?

Answer: StorageGRID Webscale uses the NetApp license file (NLF) for software licensing. The NLF contains software entitlement information such as quantity of software and software support terms, and is tied to a StorageGRID serial number. Customers can download their NLF by using the [NetApp license file generator](#).

For more information, go to the [StorageGRID software license guide](#).

3.3 Perpetual Software License

What is the StorageGRID Webscale perpetual software license?

Answer: The perpetual software license allows customers to use the StorageGRID Webscale software indefinitely. This license is calculated as a per TB raw capacity license and varies based on the underlying storage platform.

Note: The Software Support Plan (SSP) must be purchased separately.

How do I calculate the perpetual software license required for storage nodes that use the StorageGRID appliance?

Answer: The perpetual software license is licensed as raw TB based on total capacity of drives in the SGA.

How do I calculate the perpetual software license required for storage nodes that use other NetApp storage or third-party storage?

Answer: Virtual StorageGRID Webscale storage nodes (non-appliance nodes) run as virtual machines on VMware or as containerized nodes on approved host operating systems, and can use any compatible storage regardless of whether it is sold by NetApp.

In these cases, the licensing cost is calculated as $1.25 \times$ usable storage. To determine storage compatibility, consult the VMware hardware compatibility list.

Is there an evaluation license?

Answer: Yes. An evaluation-only NetApp license file for StorageGRID Webscale is included with the software.

How do I consider Amazon S3 (Simple Storage Service) or tape archiving when I calculate licensing capacity?

Answer: If a customer wants to use an archive node to archive data to Amazon S3 or tape, the StorageGRID licensing cost is calculated as $1.25 \times$ archived data.

In other words, if the customer wants to archive 50TB of data to Amazon S3 or tape, the licensing cost is calculated as $1.25 \times 50\text{TB}$. The underlying licensing principle is that all storage under the management of StorageGRID Webscale must be licensed.

Do I need to obtain a feature product variance request (FPVR) when I sell StorageGRID Webscale licenses to customers who want to run virtual storage nodes on third-party storage?

Answer: No, an FPVR is not required for third-party storage. Third-party storage must be RAID protected and compatible with VMware ESX.

Can I convert an evaluation license to a fully supported production license?

Answer: Yes. You can update the license information for your grid by uploading the updated NetApp license file that contains your license entitlement to the StorageGRID Grid Management Interface (GMI).

For details, refer to the *StorageGRID Webscale Administration Guide* in the [NetApp StorageGRID Webscale Resources](#) page.

3.4 Software Subscription License

What is the StorageGRID Webscale software subscription license?

Answer: The software subscription license is a new way to consume the StorageGRID Webscale software and SSP. In this licensing model, customers pay per TB of *used* capacity for a fixed subscription term. This pricing model allows customers to scale flexibly as they grow their storage footprint, with minimal up-front capital investment.

You can find more information here: [Product Bulletin: NetApp StorageGRID Webscale Software Subscription](#) and [Software subscription pricing blog](#).

NetApp field team: You can also review the [Subscription Pricing field deep-dive presentation](#).

Is there a minimum capacity and subscription term?

Answer: NetApp recommends a minimum 10TB, 6-month subscription license.

This license allows customers to store up to 10TB of data over the 6-month period. Customers can purchase terms of any length in 6-month increments (6, 12, 18, 24, and so on) and any capacity in multiples of 10TB.

NetApp field team: You can also review the [Subscription Pricing field deep-dive presentation](#).

3.5 Quoting

How do I quote StorageGRID Webscale?

Answer: Read [StorageGRID Webscale Quoting guide](#) to learn about quoting the perpetual license.

Read the [Product Bulletin: NetApp StorageGRID Webscale Software Subscription](#) to learn about quoting the software subscription license.

NetApp field team: You can also refer to [StorageGRID Webscale Quoting guide](#) and [Quoting video](#) in the StorageGRID Community page.

3.6 Demonstration and Proof of Concept

How can I demo StorageGRID Webscale?

Answer: See the [StorageGRID Webscale proof of concept guide and test cases](#) for details on how to demo StorageGRID Webscale.

4 API support

4.1 Supported APIs

What RESTful APIs does StorageGRID Webscale support?

Answer: StorageGRID Webscale supports S3 and Swift RESTful APIs.

Note: CDMI and the legacy StorageGRID API (SGAPI) are no longer supported.

4.2 S3 API Support

What S3 API features are implemented in StorageGRID Webscale?

Answer: This information is available in the *StorageGRID Webscale S3 Implementation Guide* in the [NetApp StorageGRID Webscale Resources](#) page.

NetApp field team: See this internal-only link for details:

https://wikid.netapp.com/w/NGS_NPI/TOI/Content_Repository/StorageGRID_11.0

What new S3 API features are added to StorageGRID Webscale 11.0?

Answer: StorageGRID Webscale 11.0 adds the following new APIs:

- S3 bucket replication. This feature enables users to mirror the contents of an S3 bucket in StorageGRID Webscale to a destination S3 bucket (for example, in AWS S3 or another StorageGRID instance). Objects mirrored to the destination S3 bucket are in native format, allowing users to access the object and its metadata transparently across S3 targets.
- S3 event notification. Users can now publish notifications about events on a StorageGRID S3 bucket to the AWS Simple Notification Service (SNS). Notifications sent to SNS can then be used as triggers for hybrid cloud workflows.
- Object tagging. StorageGRID Webscale 11.0 allows users to add tags (key value pairs) to objects. Similar to adding custom metadata, tagging can help users classify objects. Users can create ILM rules based on object tags. Tags can also be searched and analyzed in Elasticsearch.
Note: Search requires integrating to an Elasticsearch cluster.
- Location constraint. Grid administrators can define S3 regions, which are values that an S3 user can specify as location constraints for their buckets. ILM rules can then be created to specify object placement based on location constraint.

4.3 SWIFT API Support

What Swift API features are added in StorageGRID Webscale 11.0?

Answer: Swift API features remain the same as in 10.4. For details, refer to the *StorageGRID Webscale Swift Implementation Guide* in the [NetApp StorageGRID Webscale Resources](#) page.

5 Erasure Coding

5.1 Feature Product Variance Request (FPVR) Requirement

Do I need to obtain an FPVR to enable erasure coding on StorageGRID Webscale?


Answer: No, erasure coding no longer requires an FPVR.

5.2 Definitions

What is erasure coding?

Answer: Erasure coding (EC) is a data protection technology that takes an object, slices it into data fragments, and computes additional parity fragments.

Figure 1) Introduction to erasure coding.




When an object is accessed, the data fragments are retrieved and the object is decoded. The parity fragments provide data protection and can be used to recreate lost data. If either a data fragment or a parity fragment becomes corrupt or lost, the erasure coding algorithm can recreate the fragment by using a set of the remaining data and parity fragments.

What is geo-distributed erasure coding?

Answer: StorageGRID Webscale features geo-distributed EC, or the use of an information dispersion algorithm. It takes the data and parity fragments that are produced by applying an erasure code algorithm to an object and distributes them across many nodes or sites. This process provides data protection for node failures or site loss. Each fragment is stored to a separate storage node. An object can be retrieved or repaired even if a node or a site is lost.

Figure 2) Introduction to geo-distributed erasure coding.




What is layered erasure coding?

Answer: StorageGRID Webscale features layered EC (previously known as hierarchical EC), which provides multiple levels of EC to protect object data. At the software level, StorageGRID Webscale geo-distributed EC provides data protection across nodes and sites.

At the hardware level, also called the node level, NetApp E-Series provides erasure coding protection at the storage array level by using E-Series Dynamic Disk Pools (DDP) technology.

Figure 3) Layered erasure coding: erasure coding at the software and the hardware level.



NetApp field team: For more information, go to the internal [StorageGRID Community](#)

5.3 Value Proposition

What is the value proposition of layered erasure coding in StorageGRID Webscale?

Answer: Layered EC is best suited for the following use cases:

- Distributed content across three or more sites
- High data availability and durability
- Failure tolerance: protection against complete site and node failures
- Storage efficiency: reduced cost per GB used
- Providing single-copy data protection at a single site

Layered EC uses the StorageGRID Webscale distributed architecture to provide high availability and durability for objects, while reducing the use of network and compute resources for disk-level repair operations.

5.4 Erasure Coding Schemes

What EC schemes are supported by StorageGRID Webscale?

Answer: The schemes in Table 2 are provided out of the box.

Table 2) Supported EC schemes.

Sl. No	EC Scheme	Sites (Single site supported)
1	2 + 1	3
2	4 + 2	3
3	6 + 3	3
4	6 + 2	4
5	9 + 3	4
6	8 + 2	5

Notes:

1. For a three-site deployment requiring single-site loss protection, NetApp recommends the 4 + 2 over the 6 + 3 because fewer compute resources are required for ingests, retrievals, and repairs.
2. Although 2 + 1 is supported for three-site deployment desiring single-site loss protection, durability is diminished because 2 + 1 is only capable of recovering from a single node loss.
3. Any of the EC schemes can be used for single-site or two-site grid deployments, but they do not provide site loss protection. To have site loss protection, the deployment must have a minimum of three sites.

What is the limit on the number of erasure coded objects?

Answer: There are no separate limits on number of erasure coded objects. A node-level limit of 500 million objects per node applies.

What is the minimum object size supported for EC?

Answer: EC is supported for objects greater than 200KB in size. However, for optimal performance, NetApp recommends that EC be used on objects 1MB and larger.

6 Infrastructure

6.1 Virtualization

What are the virtual machine requirements for StorageGRID Webscale?

Answer: Every production node is required to have 8 vCPUs (equivalent to 4 physical cores when hyperthreading is enabled) and 24GB RAM.

What are the storage requirements for virtualized component nodes in StorageGRID Webscale?

Answer: Storage requirements include support for any block-based storage, with VMware interoperability support from all required vendors.

RangeDB disks are recommended to be RDM attached, but VMFS/VMDK is also supported.

Here are the best practices and recommended storage for StorageGRID Webscale:

- NetApp E-Series storage system with capacity-optimized NL-SAS in DDP
- SSDs for storage node system disks (100GB per node) to hold a Cassandra commit log and StorageGRID audit disk queue*
- SSDs for all admin node disks (500GB per node) to hold a MySQL commit log, audit disk queue, and database*
- All block-based storage with protected write-back caching (for example, battery-backed write cache)

* SAS drives are also supported if they are sized correctly for performance.

Note: Since 10.3, the admin node storage requirement has increased from 300GB per node to 500GB per node. This change is because of the increase in audit and database disk requirements by 100GB each.

What hypervisors can StorageGRID Webscale be deployed on?

Answer: StorageGRID Webscale can be deployed on VMware ESXi, with dependencies on VMware vCenter for virtual machine deployment orchestration.

Other hypervisors, such as OpenStack KVM, Hyper-V, and Xen, are also supported using customer-provided virtual machines running supported host operating systems (RHEL, CentOS, Debian, or Ubuntu). Virtual machines are treated as virtual hosts, similar to physical bare-metal hosts in containerized deployments.

For details, go to the [NetApp Interoperability Matrix](#) (IMT).

What operating system is StorageGRID Webscale deployed on?

Answer: StorageGRID Webscale is deployed on the Debian operating system (version 8), as a Docker container, for VMware deployments.

StorageGRID Webscale is also supported on customer-provided VMs running RHEL 7, CentOS 7, Ubuntu 16.04 (new in 11.0), or Debian 8 (new in 11.0) host operating systems, as either VMs or physical hosts.

For more details, go to the [NetApp Interoperability Matrix](#) (IMT).

6.2 Bare-Metal Deployment

Can I install StorageGRID Webscale without a hypervisor directly on an x86 server?

Answer: Yes. StorageGRID Webscale nodes can be installed on RHEL 7, CentOS 7, Ubuntu 16.04 (new in 11.0), or Debian 8 (new in 11.0) hosts, as either virtual machines or physical hosts.

On the appliance, StorageGRID Webscale is already directly installed on an x86 server.

What are the system requirements for deploying StorageGRID Webscale on RHEL or CentOS?

Answer: Host system requirements:

Processor: Intel/AMD x86-64 architecture
Host CPUs must support SSE4.2
Intel Nehalem and newer; AMD Bulldozer/Jaguar/Piledriver and newer
OS: RHEL 7, CentOS 7, Ubuntu 16.04, or Debian 8
Docker: Version 1.10.3 or later

Notes:

1. For production deployments, to minimize impact on availability, there should be only one storage node in a single grid instance running on a given bare-metal or virtual machine host.
2. Non storage nodes (admin, API gateway, archive) can run on the same host as a storage node.

What types of storage can StorageGRID nodes use when deployed as a bare-metal deployment?

Answer: As long as the underlying storage is protected by RAID 6 or higher, any storage that can be given a block device appearance in the supported operating systems can be used (JBODs and bare HDDs/SSDs are not supported). This includes, but is not limited to, the following:

- Local (DAS) hardware RAID 6
- Local (DAS) software RAID 6
- iSCSI shared storage (RAID 6 or higher)
- Fibre Channel shared storage (RAID 6 or higher)

For details, see the *StorageGRID Webscale Software Installation Guide* for your target environment in the [NetApp StorageGRID Webscale Resources](#) page.

Is any automation available for deploying StorageGRID as a bare-metal deployment?

Yes. Sample Ansible deployment artifacts are included in the StorageGRID software bundle.

For details on bare metal deployment, see the *StorageGRID Webscale Software Installation Guide* for your target environment in the [NetApp StorageGRID Webscale Resources](#) page.

NetApp field team: See this internal-only link for more details:

https://wikid.netapp.com/w/NGS_NPI/TOI/Content_Repository/StorageGRID_10.4/RHEL

6.3 Monitoring

How do I monitor StorageGRID Webscale?

Answer: There are multiple ways to monitor StorageGRID Webscale:

- Grid Management Interface (GMI): Web interface (previously known as the NMS).

- Metrics API (new in 11.0): Query the new Prometheus metrics engine in real time by using the management API and build customized dashboards.
- E-mail: See “Configuring Notifications” in the StorageGRID Webscale 10.x administrator guide.
- SNMP: See “Configure SNMP Monitoring” in the StorageGRID Webscale 10.x administrator guide.
- Syslog: The bycast.log file can be forwarded by using SNMP:
 - For information about configuring syslog-ng, refer to this document: <https://fieldportal.netapp.com/content/205434>
- NetApp AutoSupport tool
- Elasticsearch, Logstash, and Kibana (ELK)
 - Go to <https://fieldportal.netapp.com/content/205434>

6.4 OpenStack Support

How can I use StorageGRID Webscale in an OpenStack cloud?

Answer: Customers can use StorageGRID Webscale object storage services in an OpenStack cloud environment by using Swift or S3 API. Apart from being a general-purpose object storage service, StorageGRID Webscale can also be used as a backup target for Cinder volume backups and snapshots.

How can StorageGRID be deployed in an OpenStack environment?

Answer: StorageGRID Webscale can be deployed on OpenStack KVM virtual machines running any supported operating system (RHEL, CentOS, Debian, or Ubuntu). KVM virtual machines are treated as virtual hosts, similar to physical hosts in bare-metal deployments.

Customers can either deploy a whole grid on OpenStack, or they can mix StorageGRID appliances and nodes in OpenStack together in the same grid.

Note: OpenStack deployment of StorageGRID Webscale no longer requires an FPVR.

Does StorageGRID Webscale support OpenStack Keystone as an identity source?

Answer: Yes. StorageGRID Webscale supports OpenStack Keystone as an identity source in addition to native and LDAP-based authentication. Keystone v3 authentication is supported for Swift API users. Each Keystone project maps to a Swift account, and these accounts are automatically created on StorageGRID Webscale after Keystone is setup. StorageGRID Webscale also provides a token caching feature that is optimized for multisite deployments. (UUID and Fernet token are currently supported.)

Note: Keystone integration requires an FPVR.

6.5 Networking

In a logical StorageGRID data center, do the admin node, gateway node, and all storage nodes need to be in the same VLAN/IP subnet?

Answer: No. However, it is much easier and less complicated to have all the nodes in the same network. It would be ideal to create a separate VLAN just for StorageGRID and collect all data center components in the same network. The VLAN would then need to be allowed to route traffic to clients and other data centers in the infrastructure.

What Is the Grid, Admin, Client Network Model?

Answer: Starting in StorageGRID Webscale 10.3, a Grid, Client, Admin (GAC) network model is implemented. The GAC model provides a simple method for configuring three different networks to meet customer requirements for segregating StorageGRID network traffic. The three networks are defined as follows:

- Grid network. Globally routable so that all nodes can talk to all other nodes. This network is required for internal StorageGRID traffic, such as replication.
- Admin network. A "closed" network that is used for maintenance and administration.
- Client network. An "open" network that is used for S3 or Swift API access.

For more information, see the *StorageGRID Webscale Primer* in the [NetApp StorageGRID Webscale Resources](#) page.

NetApp field team: See this internal-only link for more details:

https://wikid.netapp.com/w/NGS_NPI/TOI/Content_Repository/StorageGRID_10.3/Webscale#Network_Model

Does the StorageGRID appliance have VLAN support?

Answer: Yes. StorageGRID Webscale includes StorageGRID appliance firmware version 2.0, which adds VLAN support to the StorageGRID appliance. The VLAN support provides operational and deployment flexibility that minimize dependency on network administrators.

For more information, see the Installing and setting up section in the [NetApp StorageGRID Webscale Resources](#) page.

NetApp field team: See this internal-only link for more details:

https://wikid.netapp.com/w/NGS_NPI/TOI/Content_Repository/StorageGRID_10.4/Topic_-_SGA

Does the StorageGRID appliance have Link Aggregation Control Protocol (LACP) support?

Answer: Yes. StorageGRID Webscale 10.4 includes StorageGRID appliance firmware version 2.0, which adds LACP support to the StorageGRID appliance (SGA). LACP enables better use of available bandwidth by bundling several physical ports together to form a single logical channel. Administrators can now bond 10G ports into 20G, 30G, or 40G logical channels that allow higher-bandwidth connections to the Grid and Client networks. LACP also provides better resiliency of communication with the SGA in the presence of network faults.

For more information, see the Installing and setting up section in the [NetApp StorageGRID Webscale Resources](#) page.

NetApp field team: See this internal-only link for more details:

https://wikid.netapp.com/w/NGS_NPI/TOI/Content_Repository/StorageGRID_10.4/Topic_-_SGA

Are MTU 9000 (jumbo frames) required?

Answer: No, although MTU 9000 is preferable if your workload consists of large objects.

7 Design

7.1 Grid Architecture

How do I size NetApp StorageGRID Webscale?

Answer: The [StorageGRID Webscale sizer tool](#) is available on the Field Portal.

Also see the [Sizer user guidelines](#).

What does a basic StorageGRID deployment consist of?

Answer: At minimum, a StorageGRID deployment requires three storage nodes, one admin node, and a minimum of one site. You can optionally add a gateway node and an archive node.

For more information on StorageGRID architecture, see the *StorageGRID Webscale Primer* in the [NetApp StorageGRID Webscale Resources](#) page.

Should there be separate grids for specific operational models, such as production, nonproduction, test, backup, and so on?

Answer: This depends on the customer's requirement. The best approach is to build a single large grid (similar to a cloud) and to use the information lifecycle management (ILM) policies to specify how data is replicated, protected, and made available within the grid. For example, ILM can be configured so that production data can have more robust data durability than nonproduction data does, while all data is still located on a single grid.

One potential use of a test environment would be to configure and test ILM policies before implementing them on a production grid.

Note: Currently, there is no support to export or import configurations to and from a grid.

Can a customer merge multiple grids into a single grid?

Answer: This is currently not supported. To merge grids, customers must migrate data from disparate grids into a single grid.

Note: Because a single grid can tier out to any Amazon S3 compatible storage, one StorageGRID Webscale grid can tier out to another grid by using the archive node. This approach could be used to connect multiple disparate grids.

In a deployment with multiple admin nodes, does a secondary node take over the role of a primary node in if a primary node fails?

Answer: Because a grid can have only one configuration management node located in the primary admin node, a secondary admin node cannot be designated as a primary node. You can only rebuild the failed node and then copy the database from another admin node to it. Copying the database is optional, depending on whether the historical data and audit logs need to be restored into the primary admin node.

During this time, the grid might still be operational, but eventually there will be issues. These issues arise because the primary admin node allocated a batch of unique IDs (UUIDs) to each storage node. The storage node, in turn, allocated these UUIDs to ingested objects. After the storage node runs out of UUIDs, it cannot service any new object ingest requests.

How many objects can a single storage node handle?

Answer: Each storage node can support up to 500 million objects.

What are the maximum and minimum object sizes supported?

Answer: Maximum object size: 5TB (tested limit)

Minimum object size: Erasure coded objects – 1MB

Replicated object – 0KB (metadata only)

7.2 Expansion

Can a customer convert a single data center deployment to a multi data center deployment?

Answer: Yes. StorageGRID is designed to allow grid expansion of this sort. This topic is also covered in the expansion guides in the [NetApp StorageGRID Webscale Resources](#) page.

What are some expansion scenarios?

Answer: See the expansion section in the [NetApp StorageGRID Webscale Resources](#) page.

7.3 Tiering

Is there support for tiering out to the public cloud?

Answer: Yes. You can tier out objects to the AWS cloud through the archive node. Objects that are tiered out using the archive node are StorageGRID objects and can't be accessed by external applications through the public cloud.

Alternatively, in StorageGRID Webscale 11.0, users can use bucket replication to mirror objects in StorageGRID buckets to a destination S3 bucket (for example, AWS S3 or another StorageGRID instance). Objects mirrored to the destination S3 bucket are in native format, allowing users to access the object and its metadata transparently across S3 targets.

What is the capacity that can be tiered to the cloud using the archive node?

Answer: The software does not have any hard capacity limits for the archive node, other than the total StorageGRID object count limit. However, to avoid having a single archive node as a single point of failure, it is a good practice to have multiple archive nodes for large amounts of data.

Is a capacity license required to cover capacity on a tape or cloud tier?

Answer: Yes. Tiering to tape or cloud requires a capacity license. Refer to the [Licensing](#) section of this document for details.

Note: (1) Bucket replication does not require additional StorageGRID capacity licenses.

(2) Additional charges may apply for mirrored objects, depending on the S3 destination where the objects are replicated.

8 NAS Bridge

8.1 Protocols

What protocols does the NAS Bridge support?

Answer: The NAS Bridge currently supports NFSv3 and SMBv2.1 and 3.0.

Can I use both NFS and SMB protocols at the same time?

Answer: The NAS Bridge maps an NFS export or an SMB share to an S3 bucket. An S3 bucket can be exported (or shared) as either NFS or SMB but not both. However, the NAS Bridge supports multiple S3 buckets, and each of these buckets can be an NFS export or an SMB share.

Can I use the NAS Bridge with any S3-based storage?

Answer: No, the NAS Bridge uses extended S3 APIs that are specific to NetApp StorageGRID Webscale.

8.2 Deployment

How is the NAS Bridge deployed?

Answer: The NAS Bridge can be deployed on a VMware ESX 5.x or ESX 6.0 compatible virtual machine.

Note: NAS Bridge deployments (proof of concept and production) are currently under FPVR.

What are the virtualization requirements for the NAS Bridge?

The NAS Bridge requires 16 vCPUs, 64GB of RAM, and 40GB of disk space for the OS. Cache device size is based on customer workload requirements.

NetApp field team: For more information on supported workloads and performance requirements, see the [NAS Bridge best practices guide](#).

How many instances of the NAS Bridge can I deploy?

Answer: You can deploy as many as necessary. The NAS Bridge is included with StorageGRID Webscale as a separate downloadable open virtual appliance. Each instance exports and manages its own NFS or SMB file systems.

Because all files are stored in StorageGRID Webscale, can two NAS Bridge instances access the same files?

Answer: Only a single active NAS Bridge for each file system is supported. However, you can shut down a NAS Bridge at one site and bring up a new one up at a different site to provide access to the original data. That way, you can move an active NAS Bridge to a site where it is needed and can protect against site loss.

8.3 Workloads

What workloads are appropriate for the NAS Bridge?

Answer: The NAS Bridge is targeted at archive workloads, which are large datasets that must be accessible but are not frequently updated or deleted.

When the NAS Bridge updates an existing file, does it have to read the entire associated object into cache first, then write it back out?

Answer: No, the NAS Bridge only pushes changes to an object. However, this approach translates into an S3 `multi-part upload - copy` command to StorageGRID Webscale, which copies the unchanged data to a new object. The changed data is then uploaded and added. Random write workloads result in much higher disk usage and should be avoided.

What file sizes does the NAS Bridge supports?

Answer: The NAS Bridge supports file sizes from 10MB to 1TB.

Does the NAS Bridge support file-object duality?

Answer: Yes. Files ingested through the NAS Bridge can be accessed directly via S3. However, you must know the file's key to retrieve it.

NetApp field team: For more information, go to [Object to inode conversion](#).

How do I find my file's key so I can get it directly from S3?

Answer: For NFS you can find a file's key by running 'stat' or 'ls -li' on the file. The returned inode number is the key for the associated object.

But I'm using SMB. How do I find my file's key?

Answer: The NAS Bridge manages the namespace and ingested files separately. For SMB, you need to walk the namespace to find the key of a specific file.

NetApp field team: For more information, go to [Object to inode conversion](#).

Does the NAS Bridge use more than one object per file?

Answer: Yes. When a file is ingested, up to three objects are created. A namespace object represents the namespace for the ingested file. An inode object is the file itself. An extended attribute object is used for any extended attributes. Extended attributes are currently used only for SMB. So, there are two or three objects per ingested file. You should take this factor into consideration when planning your installation.

Are modifications of a file directly via S3 supported?

Answer: No. All changes to a file must go through the NAS Bridge.

9 Multitenancy

9.1 Overview

What multitenancy capabilities does NetApp StorageGRID Webscale have?

Answer: Multitenancy capabilities have been available in StorageGRID Webscale since the 10.2 release. Multitenancy allows multiple groups of users to share resources of a single StorageGRID instance, while keeping their data isolated and secure from each other.

Grid administrators can create tenants and configure permissions and features per tenant. For identity federation, grid administrators have the flexibility to either integrate with a single LDAP/AD server for the entire grid or allow tenants to configure their own identity federation. They can also restrict tenant data usage by enforcing capacity quotas.

Tenant users can access StorageGRID Webscale by using the Tenant Management UI portal. They can use this portal to configure their tenant account, provision S3 credentials, set access management policies for their data, manage settings on buckets, and more.

What features are added in 11.0?

Answer: In 11.0, tenants have a greater ability to self-manage their accounts and configure how their data is stored and managed. New features include a bucket management UI that allows tenants to configure bucket settings such as:

- Consistency level. Set tradeoffs between availability and consistency of objects across sites.
- Last access time. Allow the use of ILM rules based on last access time of content in a bucket.

- Event notifications. Set notifications based on events occurring in a bucket.
- Bucket replication. Mirror objects between a StorageGRID bucket and a destination S3 bucket such as AWS S3.
- Search integration. Send object metadata to external search endpoints such as Elasticsearch.

For more information, see the *StorageGRID Webscale Tenant Administration Guide* in the [NetApp StorageGRID Webscale Resources](#) page.

9.2 LDAP and Active Directory Support

How is LDAP/AD support implemented in StorageGRID Webscale?

Answer: In StorageGRID Webscale, customers can use LDAP or Active Directory (AD) to manage (1) administrator access and permission for the Grid Management Interface (GMI), and (2) storage end user access and permission for tenants.

Grid administrators have the option to either configure and manage identity federation for the tenants, or delegate this responsibility to the tenants themselves. Identity federation for tenants is now configured in the tenant UI, and each tenant can configure their own LDAP or AD server. Tenants can also create local users in case they don't want to use federated LDAP or AD users.

What options do users have to import LDAP/AD groups?

Answer: To import the groups to StorageGRID Webscale, the administrator can either use the GMI or the management API.

How does StorageGRID Webscale synchronize with the LDAP server?

Answer: StorageGRID does a forced synchronization with the LDAP server periodically to update any changes to the groups.

On login attempt, user credentials are forwarded to the identity service (internal grid component) and from there to the LDAP server. If LDAP returns a success, then the user is authenticated by StorageGRID. All groups (and the groups of which those groups are a subgroup) that a user is a member of are also imported when a user successfully authenticates. Additionally, no user credentials are stored or cached on StorageGRID.

9.3 OpenStack Keystone support

Does StorageGRID Webscale have OpenStack Keystone support?

Answer: StorageGRID Webscale supports OpenStack Keystone as an identity source in addition to native and LDAP-based authentication. Keystone v3 authentication is supported for Swift API users. Each Keystone project maps to a Swift account, and these accounts are created automatically. StorageGRID Webscale also provides a token caching feature that is optimized for multisite deployments. (UUID and Fernet token are currently supported.)

Note: Keystone integration requires an FPVR.

9.4 Limits

What are the multitenancy limits?

Answer:

Tenant accounts:	API: Over 10k
	UI: Up to 1,000

S3 access keys:	Soft limit is 100 per user per tenant
Groups:	Soft limit 100 per account
Users:	No practical limit
LDAP servers:	For grid admin: One per grid (for GMI access) For tenants: Scales with number of tenant accounts
Buckets/containers:	Hard limit: 1,000 per account (increased from 100 buckets per account in 10.4); globally unique names; Swift and S3

Is there cross-API access between S3 and Swift accounts?

Answer: No. This is not yet available.

9.5 Credential Management

How are the various credentials stored and protected in StorageGRID Webscale?

Answer: StorageGRID Webscale protects the various types of credentials and passwords in the following ways:

- LDAP/AD passwords. These credentials are not stored on StorageGRID, but are forwarded to the LDAP/AD server for authentication.
- Tenant accounts. Access to tenant account credentials is protected by Transport Layer Security (TLS) with client certification authentication. Local users associated with tenant accounts are encrypted using bcrypt.
- S3 access keys and secret keys. Hash using SHA-256 (Secure Hash Algorithm 256).

9.6 Reporting and Chargeback

What reporting and chargeback features does StorageGRID Webscale have?

Answer: Tenants can view and monitor their usage on the tenant dashboard. The dashboard shows the tenant's usage against any quota enforced and the top 10 buckets that are using the most storage. All tenant-level usage can also be reported by using the management API, through which users can report the amount that is used. The usage can be retrieved either by the grid administrator or by the tenants themselves. The usage reported is granular and can be reported at a per-bucket level. Additional metrics will be added in future releases.

Can I see these metrics on the Grid Management Interface?

Answer: There is no per-tenant reporting in the GMI. This feature will be added in future releases. However, each tenant can view their own usage on their tenant dashboard.

10 Information Lifecycle Management (ILM)

What is ILM?

Answer: ILM provides a powerful and flexible policy engine that allows users to define how objects are stored and managed in StorageGRID. Users can outline criteria to filter objects and then define where these objects are placed on the grid. Scanning ILM was introduced in 10.4, which continuously scans objects and brings them into compliance with the current active ILM policy. This feature made ILM truly

dynamic, because administrators no longer have to reevaluate ILM after making changes to the ILM policy.

For more information, see the ILM section of the *StorageGRID Webscale Administration Guide* in the [NetApp StorageGRID Webscale Resources](#) page.

What has been updated in ILM for 11.0?

Answer: StorageGRID Webscale now supports object tagging, allowing users to add up to 10 custom key-value pairs to an object. Grid admins can create ILM rules on object tags to define where specific objects are placed and how they are protected.

Another feature added in 11.0 is support for location constraint. Grid admins can define S3 regions and create ILM rules based on these regions, restricting object placement to specific data centers or nodes. S3 users can then optionally specify a bucket's region by using the location constraint request element.

Finally, the policy editor UI has been revamped with a new integrated policy creation and simulation workflow. NetApp strongly advises using the simulator to validate a change in the ILM policy against a sample bucket before activating the new policy, to avoid unintended consequences.

How many ILM rules can a StorageGRID Webscale deployment have?

Answer: You can have as many ILM rules as you want. StorageGRID Webscale doesn't limit the number of ILM rules you can create. However, it can become very difficult to manage a large set of rules, and this could increase the probability of conflicting or overlapping rules, eventually leading to unintended data loss.

Can ILM policies be set for each tenant in a multitenant deployment?

Answer: No. ILM policies are not yet available at a tenant level. Policies are currently applied at a global level.

Can I retrieve an object that has been deleted, but not purged from the grid?

Answer: There is no supported method to retrieve a deleted object from the grid. Objects deleted are purged from the grid.

StorageGRID does support versioning, and deleted objects in a versioning-enabled bucket could be restored. For details, go to the [versioning section](#) of the FAQ.

Can I manage ILM policies through a RESTful API?

Answer: Yes. RESTful API functions for ILM policy management include:

- Evaluating a proposed ILM policy
- Listing the metadata values that can be used to create an ILM rule
- Listing ILM rules
- Creating ILM rules
- Deleting ILM rules
- Retrieving and replacing ILM rules

(New in 11.0)

- Creating proposed ILM policies
- Deleting proposed ILM policies
- Retrieving proposed ILM policies

- Activating proposed ILM policies

11 Versioning

What is versioning?

Answer: Versioning enables customers to keep multiple variants of an object in a bucket. This protects the objects from accidental modification or deletion. It can also be used to archive objects and to retrieve previous versions of the object later. Versioning is enabled at a bucket level granularity.

Starting in 10.4, ILM rules can be created to manage versioned object. Grid administrators can manage noncurrent versioned objects by specifying how long they want to retain these objects, emulating a recycle bin for versioned objects.

Currently, up to 1,000 versions of a single object are supported.

The S3 APIs supported for versioning are:

- Object operations:
 - PUT Object Version
 - GET Object Version
 - HEAD Object Version
- Bucket operations:
 - GET Bucket Object Versions
 - GET Bucket Versioning
 - PUT Bucket Versioning

Note:

- 1) At present, S3 versioning can be enabled only on buckets created in 10.3 and later.
- 2) Do NOT enable versioning on buckets serving the NAS Bridge or AltaVault.

For details, see the [StorageGRID Webscale 11.0 SE presentation](#).

NetApp field team: For details, go to this Internal-only link:

https://wikid.netapp.com/w/NGS_NPI/TOI/Content_Repository/StorageGRID_10.3/Webscale#S3_Object_Versioning

12 Performance

What questions do I ask my customers to gather the correct information about the workload that is intended for a NetApp StorageGRID Webscale environment?

Answer: To understand the customer's workload for a StorageGRID Webscale environment, collect the following information:

- Size of data to be ingested
- API interface (S3 or Swift)
- Ratio of read operations to write operations
- Average write rate
- Peak write rate
- Duration of peak write rate

- Average read rate
- Peak read rate
- Duration of peak read rate
- Copies of data required
- ILM policy required

What performance should I expect from StorageGRID Webscale?

Answer: For the latest performance enhancements, see the [StorageGRID Webscale performance technical report](#).

13 Partner integration

Where can I find information about partner solutions with NetApp StorageGRID Webscale?

Answer: Go to the Partner Solutions page on netapp.com: <https://www.netapp.com/us/products/object-storage/storagegrid/storagegrid-partner-solutions.aspx>.

Where can I find more information about AltaVault and StorageGRID Webscale integration?

Answer: See the StorageGRID Webscale solution brief at <https://fieldportal.netapp.com/content/370731?assetComponentId=371355>.

14 StorageGRID Webscale Appliance

14.1 Overview

What is the NetApp StorageGRID Webscale appliance?

Answer: The StorageGRID Webscale appliance (SGA) is a hardware-based storage node that is based on the E-series DE212 or DE460 enclosure. Compute and storage are both provided in a single chassis and are specifically configured for StorageGRID Webscale. The current family of SGA is the SG5700, consisting of the following components:

- StorageGRID compute server (E5700SG):
 - Modified E-Series E5700 controller running custom StorageGRID firmware
 - Runs the StorageGRID Webscale software as a bare-metal Debian host using Docker containers
 - StorageGRID Pre-Grid Environment (PGE) firmware, version 2.0
 - CPU: 8-core; Memory: 64GB RAM
- Storage controller (E-Series E2800):
 - Runs NetApp SANtricity OS 8.30
 - Includes SANtricity System Manager 11.30 for hardware management and monitoring
- Available in two form factors: 12 drives, 2U enclosure; 60 drives, 4U enclosure.
- Available with 4TB, 8TB, or 10TB non-FDE drives and 10TB FIPS drive configurations.
- Support for up to four 10-GbE or 25-GbE connections to the StorageGRID Webscale Grid Network and Client Network.
- External 16Gb/s fiber cable interconnect to connect E5700SG to the E2800.

- Include the StorageGRID Webscale appliance installer to simplify storage node deployment and configuration.
- Support Full Disk Encryption (FDE) drives or Federal Information Processing Standard (FIPS) drives. When these drives are used with the Drive Security feature in SANtricity System Manager, unauthorized access to data is prevented.

What is the value proposition of a StorageGRID Webscale appliance?

Answer:

- Simplified installation and management
- Reduction in hypervisor licensing and administration
- No hypervisor server infrastructure is required because it's built into the appliance

14.2 Configuration

What are the available models in the SG5700 family?

Answer: Two models are available:

- SG5760: 4U 60-drive appliance
- SG5712: 2U 12-drive appliance

The models are identical in terms of compute, performance, and operation; only the capacity differs.

What are the available configurations?

Answer: The SG5760 is available in the following configurations:

- SG5760 (4TB) = 240TB raw, 180TB usable
- SG5760 (8TB) = 480TB raw, 355TB usable
- SG5760 (10TB) = 600TB raw, 443TB usable (Also available with 10TB FIPS drives.)

The SG5712 is available in the following configurations:

- SG5712 (4TB) = 48TB raw, 34TB usable
- SG5712 (8TB) = 96TB raw, 67TB usable
- SG5712 (10TB) = 120TB raw, 84TB usable (Also available with 10TB FIPS drives.)

How is the appliance configured?

Answer: Both models are configured as follows:

- 16 x data volumes (that is, LUNs) (Size increases or decreases with drive size and quantity.)
- 1 OS volume (512GB)
- 1 PGE backup volume (32GB)

DDP configuration is different for each model depending on the disk quantity:

- SG5760 = 3-drive preservation capacity
- SG5712 = 1-drive preservation capacity

14.3 Software

Which versions of StorageGRID Webscale are supported with the appliance?

Answer: StorageGRID Webscale 10.3 and later are supported.

What is the pre-grid environment (PGE)?

Answer: PGE is the factory-installed firmware image of the E5700SG. On boot, the PGE checks to see whether the StorageGRID Webscale software has been installed and either boots the OS or runs setup as appropriate.

What are the enhancements in the new PGE 2.0?

Answer: In StorageGRID Webscale 11.0, the PGE gets an upgrade and adds the following enhancements:

- Fast persistence to handle unclean node shutdown. In 11.0, node containers across all platforms (appliance, virtual machines, bare metal), are truly ephemeral. Persistent data required by the container is written out to disks, and containers are never restarted. Containers are always removed and re-created, even when the node did not shut down cleanly.
- New appliance installer UI and management APIs. PGE 2.0 delivers a new UI and a new set of REST management APIs, in a style similar to the Grid Management Interface UI and APIs. The new UI delivers the same functionality as the old installer and adds new features, such as installation monitor UI, remap ports UI, and more.

For more details, see the *Hardware Installation and Maintenance Guide for StorageGRID Webscale SG5700 series appliances* in the [NetApp StorageGRID Webscale Resources](#) page.

How are software upgrades performed on the appliance?

Answer: Software upgrades are centrally managed by the GMI. The process is identical to the process for the VM storage node.

For details, see the *StorageGRID Webscale Upgrade Guide* in the [NetApp StorageGRID Webscale Resources](#) page.

How do I upgrade the firmware on the appliance?

Contact NetApp Support for assistance. Upgrading the firmware on the appliance should be performed only if directed by Support.

14.4 Hardware Infrastructure

Are any physical servers required?

Answer: No. The storage nodes can exist in virtual (VMware ESX or KVM virtual machine) and appliance form. The admin and gateway nodes must be installed as virtual machines.

Do I still need a hypervisor?

Answer: If you are deploying in a VMware environment, you need VMs for the admin and gateway nodes. These nodes run as VMs on VMware ESX, with dependencies on VMware vCenter.

StorageGRID also supports the container-only deployment of all node types on RHEL, CentOS, Ubuntu, or Debian physical hosts or VMs.

The SG5660 8TB and 10TB configurations exceed the 256TB limit of a VM storage node. Is that OK?

Answer: Yes. The 256TB limit is a soft limit.

How does the appliance communicate with the grid?

Answer: StorageGRID Webscale implements a Grid, Admin, Client (GAC) network model. The GAC model provides a simple method for configuring three different networks to meet customer requirements for segregating StorageGRID network traffic:

- Grid network. Globally routable so that all nodes can talk to all other nodes. This network is required for internal StorageGRID traffic, such as replication.
- Admin network. A "closed" network that is used for maintenance and administration.
- Client network – An "open" network that is used for S3 or Swift API access.

The appliance communicates with the grid by using four 10/25GbE interfaces on the StorageGRID compute blade. When configured in aggregate mode, the Grid and Client network interfaces can share up to 100G of network bandwidth.

Is there a concern with saturating a single 10GbE connection on a node, given the network connection for the appliance?

Answer: If there are problems with network saturation, the nodes should be able to retry communications. The effect would be seen only as a slower transfer of data between nodes, but it should not result in a communication disruption.

Is the appliance highly available?

Answer: High availability is provided at the following levels:

- System level. The StorageGRID Webscale software that is running in the appliance functions as part of a grid.
- Disk level. The disks are protected from failure by NetApp SANtricity DDP.

Does the appliance have support for LACP and VLAN?

Answer: Yes. For details, go to the [networking section](#).

How can I use the 7-segment display on the appliance to find out the IP address of the appliance?

Answer: The IP addresses are shown with a repeating sequence of two digits. A "-" is added to separate each octet of the IP address. For example, the IP address 10.224.1.190 is displayed as 10, -2, 24, -1, -1, 90 in the 7-segment display.

What does the customer receive? Are there additional installation tasks?

Answer: The customer receives two boxes, one containing the disk enclosure with E2700 preinstalled. The StorageGRID controller (E5600SG) ships in a separate box. The customer or installer removes the controller blank and inserts the E5600SG. Firmware images are preinstalled.

Are SSD configurations available?

Answer: Only NL-SAS is available for the current release.

Are partially populated configurations supported for the SG5760?

Answer: No. The 60-drive enclosure must be full for setup to run. All 60 disks are required. The same applies to the SG5612.

14.5 Performance

What about performance?

Answer: See this [Performance TR](#) for details.

Would the appliance be sized differently than a VM storage node?

Answer: Similar rules apply. A minimum of three storage nodes is required.

14.6 Miscellaneous

Can a grid consist of a mix of appliances and VM storage nodes?

Answer: Yes.

How is it StorageGRID licensed?

Answer: StorageGRID Webscale software licensing must be purchased in addition to the appliance. See the [licensing section](#) of this document for details.

Are expansion shelves supported?

Answer: No. StorageGRID Webscale is a scale-out storage system. If you need more storage, add more storage nodes.

15 StorageGRID 9

15.1 StorageGRID 9 Comparison

What is the difference between NetApp StorageGRID 9 and StorageGRID Webscale?

Answer: StorageGRID Webscale is a brand-new object storage software release to meet customers' cloud storage needs, including but not limited to:

- Native support for S3 APIs
- 100 billion objects
- 70PiB capacity (soft limit)
- Greatly improved simplicity, including modular, scalable, resilient architecture and simplified deployment
- Erasure coding
- Amazon AWS archiving
- A dedicated storage appliance

And we're adding improvements with each release.

What is the future of StorageGRID 9?

Answer: NetApp continues to maintain StorageGRID 9 as a software product separate from StorageGRID Webscale. Adding new features to StorageGRID 9 will be slower compared with StorageGRID Webscale.

15.2 StorageGRID 9 Transition

Is it possible to upgrade existing StorageGRID 9 customers to StorageGRID Webscale?

Answer: No. StorageGRID 9 systems cannot currently be upgraded to StorageGRID Webscale.

What is the transition plan for StorageGRID 9 customers to StorageGRID Webscale?

Answer: StorageGRID 9 customers must install a new StorageGRID Webscale system, then migrate their data from the StorageGRID 9 system.

Version History

Version	Date	Document Version History
1.0	October 2017	Initial release

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2017 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.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.277-7103 (October 1988) and FAR 52-227-19 (June 1987).

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.