



Western Digital

An Intelligent Cloud-Enabled Data Storage Solution Powered by Western Digital's Ultrastar[®] Data60 Storage Platform and NetApp[®] StorageGRID[®]



Our Vision	Our Mission	Our Purpose
To create breakthrough innovation—inspired by the convergence of human potential and digital transformation—that enables the world to actualize its aspirations	To unlock the potential of data by harnessing the possibility to use it	To be the world's iconic data storage company
		Our Values
		We Think Big. We Create Possibility. We Make It Happen. We Do It Together.

Contents

Executive Summary	2
Problem Statement	2
Solution Highlights	2
Technology Overview	3
Ultrastar Data60 60-Bay Hybrid Storage Platform Overview	3
NetApp StorageGRID Overview	6
NetApp StorageGRID Solution Design on Ultrastar Data60	8
NetApp StorageGRID Deployment Topology on Ultrastar Data60	8
StorageGRID Setup Configuration Details	9
StorageGRID Configuration Requirements	9
StorageGRID Software Installation and Configuration Details	11
Performance Result	19
Features and Benefits.....	20
Use Cases.....	20
Conclusion.....	20
Contributors	21
Version History	21
References.....	21

Executive Summary

In our digital age, the ability to collect and exploit data has become a decisive factor in every organization's ability to compete and thrive. Terabytes, or even petabytes, of data are generated every day by many organisations. Transaction data was neatly stored in structured block storage devices, but now we need to deal with a wild and fast-growing mix of unstructured data from social media, video, audio, log files, sensor data, emails, and more. Putting all of this on traditional block or file storage is expensive and hard to manage due to their inflexibility and lack of scalability. A scalable storage solution is needed to store and manage all this unstructured data effectively.

Managing petabytes of data is tedious and requires a different approach. Object storage was designed to scale and address the needs of unstructured data. Software-Defined Storage (SDS) is an architecture that abstracts software from the underlying industry standard hardware to keep total cost of ownership (TCO) low without compromising performance. Many organisations have started implementing on-premises object storage solutions based on a SDS architecture, which provides a resilient, low-cost, highly-available, manageable, and scalable storage solution.

This document demonstrates NetApp StorageGRID, a fully compliant S3 API hybrid cloud storage software deployed on Western Digital's Ultrastar Data60, which is a reliable, high-performant, high density storage platform for software-defined storage.

Problem Statement

IT leaders have long struggled to economically keep pace with growing data storage needs, but today's exploding data volumes are pushing this challenge to the breaking point. Legacy implementations using traditional SAN/NAS and DAS infrastructure can no longer provide the scale and flexibility required in modern environments. In today's world, where year-over-year generated data growth can be measured in zettabytes, overprovisioning in the traditional approach is a recipe for wasting resources and constantly scrambling to keep pace with business needs. IDC reports that year-on-year data creation growth is now measured in zettabytes and is expected to grow up to 175 ZB by 2025.

As always with growing data, the storage volume pools continue to grow, leading to scaling issues in the current deployments. Traditional SAN and NAS vendors are not able to keep pace with the volume and velocity of unstructured data. Future data centers need a storage solution that is both linearly scalable and low-TCO. Fortunately, cloud object storage architectures have emerged, which can help support this massive data growth. Object storage architectures can provide nearly infinite scalability and manage vast amounts of unstructured data under a single pane of glass. It provides a broader view of your storage resources.

In short, what enterprises need is a low-TCO, scalable, high-availability, and easy-to-manage solution to meet their growing data storage needs. In this document, we demonstrate a solution for the growing demand of unstructured data by deploying NetApp StorageGRID software on a Western Digital Ultrastar Data60 storage platform.

Solution Highlights

Western Digital, a pioneer in reliable, high-density industry-standard hardware for software-defined storage projects, worked with NetApp, a file and object storage company specializing in hybrid cloud object storage solutions on such a solution. The following sections of this paper provide an overview of this software-defined object-based storage solution - built on Western Digital's Ultrastar Data60 and powered by NetApp StorageGRID.

The solution combines:

- **Western Digital Ultrastar Data60** is a key element of next-generation disaggregated storage and software-defined storage (SDS) systems. It addresses the demanding storage needs of large enterprise customers, storage OEMs, cloud service providers, as well as resellers/integrators requiring dense, shared HDDs.

- **NetApp StorageGRID software:** NetApp StorageGRID is a software-defined, object-based storage solution that supports the Amazon Simple Storage Service (S3) API. StorageGRID provides secure, durable storage for unstructured data at scale. Integrated, metadata-driven lifecycle management policies optimize where your data lives throughout its life.

By combining NetApp StorageGRID software with Western Digital Ultrastar Data60, organizations can realize:

- Massive scalability
- Flexible deployment
- Resilience and high availability
- Low TCO
- 100% native Amazon S3 API support
- Automated easy data management
- Management and monitoring
- Cloud platform services

Technology Overview

Ultrastar Data60 60-Bay Hybrid Storage Platform Overview

The Western Digital's **Ultrastar Data60 60-Bay Hybrid Storage Platform** provides several core capabilities and advantages in on-premises cloud storage environments.



Designed for High Density and Flexibility

The Ultrastar Data60 is a key element of next-generation disaggregated storage and software-defined storage (SDS) systems, delivering high density and the flexibility to balance performance with cost. The Ultrastar Data60 provides up to 1.32PB¹ of raw storage using our 22TB HDDs in a compact and efficient form factor. Western Digital HelioSeal® drives ensure cool running, quiet operation, and high reliability. A high-performance data tier can be set up for demanding applications by using SSDs in up to 24 of the drive slots, enabling the ability to serve both fast and big data from a single platform²

¹ One terabyte (TB) is equal to one trillion bytes and one petabyte (PB) is equal to one quadrillion bytes. Actual user capacity may be less due to operating environment.


² Ultrastar Data60 Hybrid Storage Platform supports third-party device features that are within SAS specification as mandatory. Any third-party drive features that are vendor specific are not guaranteed to function.



Ultrastar Data60 Storage Platform

Building on 50+ Years of Storage Design Experience

Conventional dense disk shelves frequently suffer from performance degradation due to induced vibration from adjacent drives. Traditional platforms also have cooling challenges as the cooling air passes over successive rows of drives, losing effectiveness as it gets heated up along the airflow path. Developing storage devices and platforms side-by-side, we address these challenges through Silicon to Systems Design, a set of technologies developed based on a holistic view of devices, platforms, and their interactions. The first two of these innovative technologies are IsoVibe™ and ArcticFlow™. IsoVibe reduces vibration-induced performance degradation, while ArcticFlow overcomes the cooling issues by introducing cool air into the middle of the platform. Both technologies contribute to long-term reliability, enabling our five-year limited warranty on the entire platform.



IsoVibe Patented Vibration Isolation Technology

Precise cuts in the baseboard provide a suspension for the drives in the chassis, isolating them from transmitted vibration. The result is that consistent performance is maintained, even when all the drives are working hard.




ArcticFlow Innovative Thermal Zone Cooling Technology

By introducing cool air into the center of the chassis, drives operate at lower and more consistent temperatures than conventional systems. This results in lower fan speeds, reduced vibration, lower power consumption, quieter operation, and ultimately higher reliability.

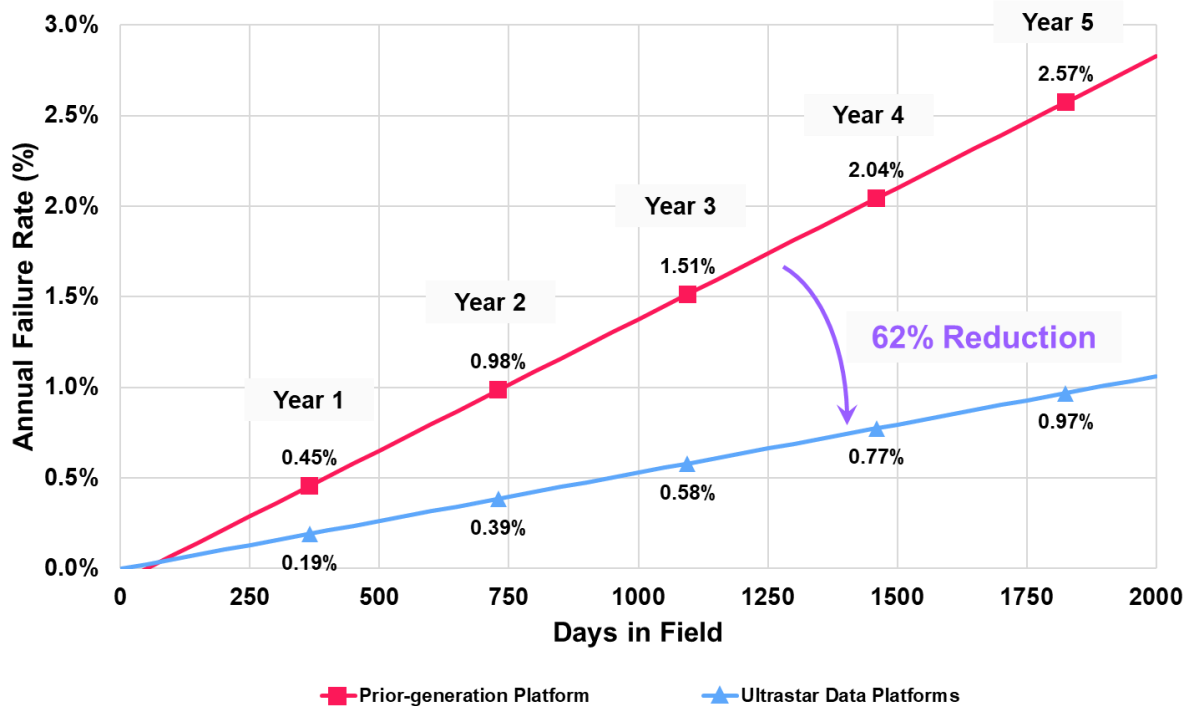


Western Digital Resource Manager

A GUI-based tool that enables real-time monitoring and management of the platform and provides a consolidated dashboard displaying the most critical information. Other views allow platform configuration, health monitoring and maintenance.



After implementing both ArcticFlow and IsoVibe in the Ultrastar Data60 and Data102 enclosures, Western Digital compared the field-return rates for these two product sets to a previous generation enclosure without ArcticFlow and IsoVibe. Both enclosure generations have been sold into the same types of customer environments and were supporting the same types of workloads with the same kind of disk drives. In the span of about three years, field-return rates have dropped by 62% (see Figure below). Those findings are based on a very large sample size – in the range of hundreds of thousands of drives.



Designed for the Enterprise and the Cloud

The Ultrastar Data60 addresses the demanding storage needs of large enterprise customers, storage OEMs, cloud service providers and resellers/integrators that require dense, shared HDD or hybrid storage. It provides the flexibility to specify the HDD and SSD combinations to balance capacity, performance, and cost.

Ultrastar Data60 60-Bay Storage Platform Features

Max. Drives	<ul style="list-style-type: none"> • 60 x 3.5" drive bays •
Drive Interface	<ul style="list-style-type: none"> • 12Gb/s SAS • 6Gb/s SATA
Available Drive Capacities	<ul style="list-style-type: none"> • HDD up to 22TB¹ CMR or up to 26TB¹ SMR
Host Interface	<ul style="list-style-type: none"> • Dual redundant I/O Modules (IOM), • 6 Mini-SAS HD ports per IOM
Cooling	<ul style="list-style-type: none"> • 4 main enclosure fans, front-to-rear system cooling with zero-loss backflow prevention • 1 IO module fan • Dual PSUs with built-in fans
Power	<ul style="list-style-type: none"> • Dual 1600W, 80+ Platinum • 200-240V AC input, auto ranging, 50-60Hz
Serviceability	Cable-free hot-swappable IOM, power supply, fans and drives

NetApp StorageGRID Overview

NetApp StorageGRID is a software-defined object storage suite that supports a wide range of use cases across public, private, and hybrid multi-cloud 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 is composed of globally distributed, redundant, heterogenous nodes that collectively support file and object protocols, allowing integration with both existing and next-generation client applications.

Grid nodes and services

The basic building block of a StorageGRID system is the grid node. Nodes contain services, which are software modules that provide a set of capabilities to a grid node. The values and statuses for all the functions of the StorageGRID system are reported through attributes.

The StorageGRID system uses three types of grid nodes.

Admin Nodes

Admin nodes provide management services such as system configuration, monitoring, and logging. Each grid must have one primary Admin node and might have additional non-primary Admin nodes for redundancy. Admin nodes can also be used to load balance S3 client traffic.

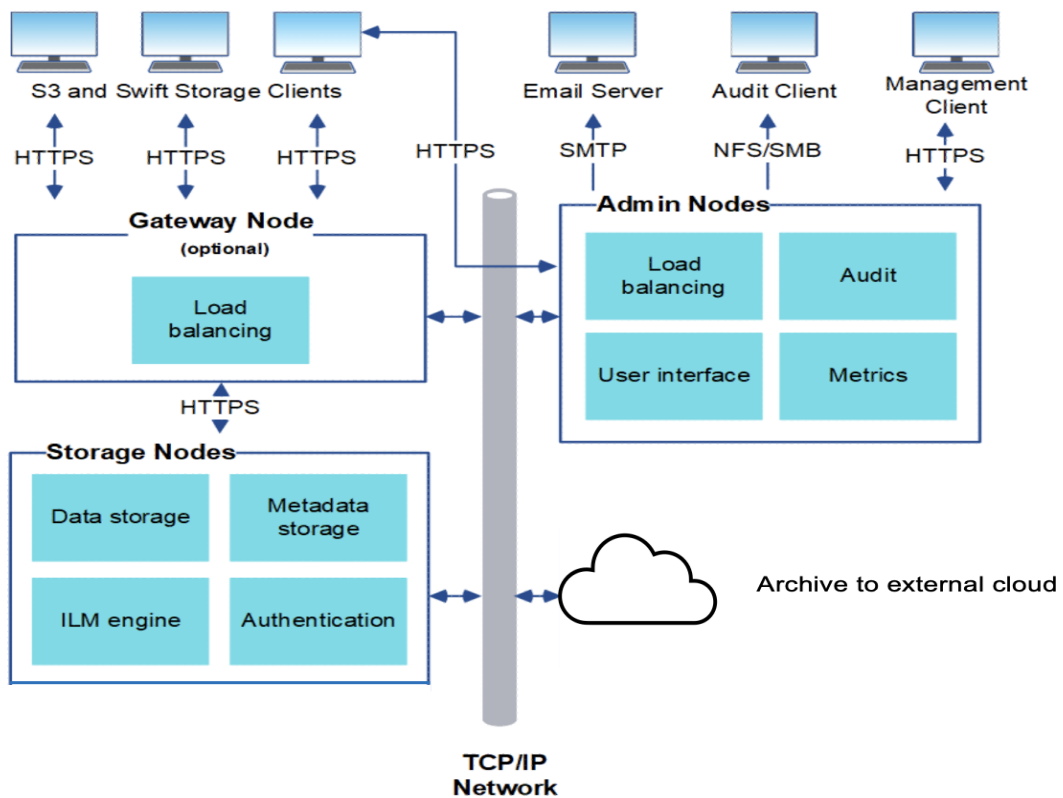
Storage Nodes

Storage nodes manage and store object data and metadata. Each StorageGRID system must have at least three Storage nodes. If you have multiple sites, each site within your StorageGRID system must also have three Storage nodes.

Gateway Nodes

Gateway nodes provide a load-balancing interface that client applications can use to connect to StorageGRID. A load balancer seamlessly directs clients to an optimal Storage node, so that the failure of nodes or even an entire site is transparent. You can use a combination of Gateway nodes and Admin nodes for load balancing, or you can implement a third-party HTTP load balancer.

¹ One terabyte (TB) is equal to one trillion bytes and one petabyte (PB) is equal to one quadrillion bytes. Actual user capacity may be less due to operating environment.



Primary services for Admin Nodes

The following list shows the primary services for Admin nodes; however, this list does not list all node services.

- Audit Management System (AMS) service tracks system activity.
- Configuration Management Node (CMN) service manages system-wide configuration. This service runs on the primary admin node only.
- Management Application Program Interface (mgmt-api) service processes requests from the Grid Management API and the Tenant Management API.
- High Availability service manages high availability virtual IP addresses for groups of Admin Nodes and Gateway Nodes.
- Load Balancer service provides load balancing of S3 traffic from clients to Storage Nodes.
- Network Management System (NMS) service provides functionality for the Grid Manager.
- Prometheus service collects and stores metrics.
- Server Status Monitor (SSM) service monitors the operating system and underlying hardware.

Primary services for Storage Nodes

The following list shows the primary services for Storage nodes; however, this list does not list all node services.

- Account (acct) service manages tenant accounts.
- Administrative Domain Controller (ADC) service maintains topology and grid-wide configuration.
- Cassandra service stores and protects object metadata.
- Cassandra reaper service performs automatic repairs of object metadata.
- Chunk service manages erasure-coded data and parity fragments.
- Data mover (dmv) service moves data to cloud storage pools.
- Distributed Data Store (DDS) service monitors object metadata storage.
- Identity (idnt) service federates user identities from LDAP and Active Directory.
- Local Distribution Router (LDR) service processes object storage protocol requests and manages object data on disk.
- Replicated State Machine (RSM) service ensures that S3 platform service requests are sent to their respective endpoints.
- Server Status Monitor (SSM) service monitors the operating system and underlying hardware.

NetApp StorageGRID Solution Design on Ultrastar Data60

Western Digital, a pioneer in reliable, high-density industry-standard hardware for software-defined storage projects, is partnering with NetApp to provide a verified, highly scalable object storage solution architecture for your data. It enables backup capabilities with high-performance computing and erasure coding and mirroring (HA) with replication for the on-premises and cloud systems architectures. The following sections of this paper provide an overview of the deployed solution.

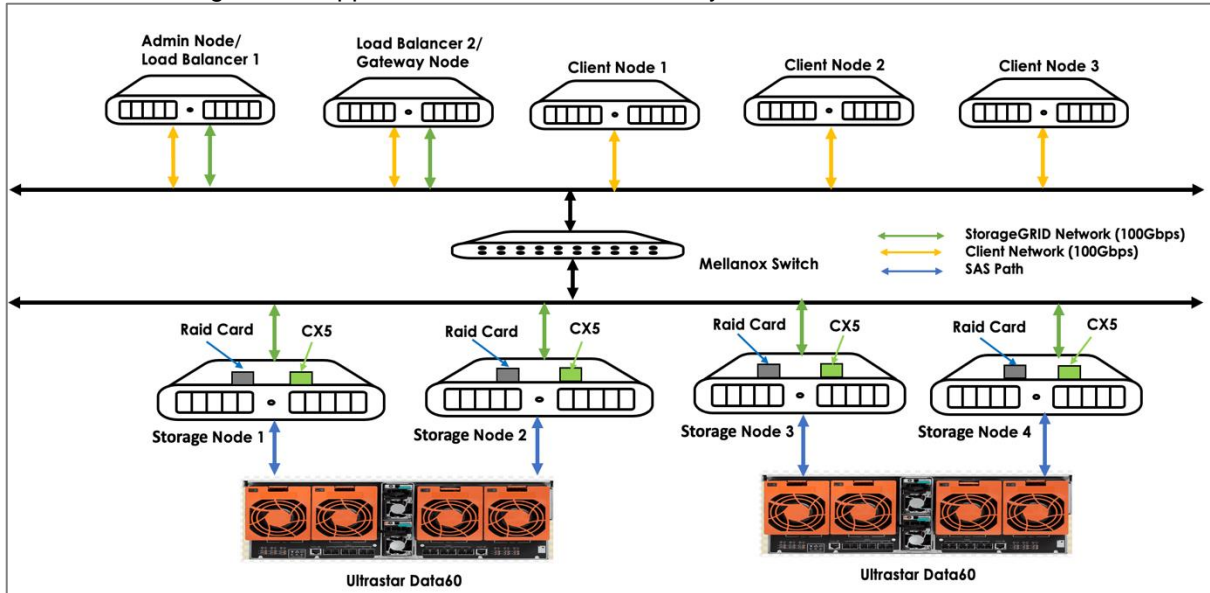
NetApp StorageGRID Deployment Topology on Ultrastar Data60

Ultrastar Data60 and NetApp StorageGRID have created a unique reference architecture with industry-standard servers and ethernet network switches. This end-to-end solution can be specifically designed to accelerate large-scale production while delivering backup, compute, replication, and storage performance as per need. The figure below illustrates a verified deployment topology with Ultrastar Data60.

NetApp StorageGRID is a software-defined, object-based storage solution that supports industry-standard object APIs such as Amazon S3. It allows you to build a single name space across many sites, with multiple service levels for metadata-driven object life-cycle policies. StorageGRID protects data via intelligent policy with options including replica, erasure coding and cloud tier.

As part of our current deployment to validate the object storage solution we have used below grid nodes:

- **Admin Nodes:** To provide management services such as system configuration, monitoring, and logging.
- **Storage Nodes:** To manage object data and metadata storage, including data protection.
- **API Gateway Nodes:** To provide a load balancing interface to the StorageGRID system through which applications can connect to the system.



StorageGRID Setup Configuration Details

Deployment Configuration Details	
Storage Product	2 x Ultrastar Data60
Software	NetApp StorageGRID 11.6
Storage Interface	8 x 12Gb/s SAS-3 host connections
Host OS	Ubuntu 22.04
Host NIC	1 x CX5/CX4 - (100Gbps)
CX5 OFED Package Version	5.8.1.2
CPU	Intel® Xeon® Gold 5318Y
CPU Core Details	Dual socket server with 24 core CPU each. 96 logical cores in total with HT enabled
Memory	128GB
Total Servers	9 (Storage-4, Admin/Load Balancer-1, Load Balancer-2, Client-3)
No of HDD on each Storage Node	30
HDD Capacity	18TB or 22TB (Total 60HDD per JBOD)
No of SSD on each Storage Node	2 SSD per server. Total 8 SSD.
No of SSD on each Load Balancer server	1 SSD per server. Total 2 SSD. This is needed for OS installation. 1TB SSD.
HBA	1 x 9480-8e/9580-8e per server. Total 4 RAID card.
SSD Capacity	7.68 TB (Each server with 2 SSDs. Total 6 SSDs on 3 servers)

For this deployment, we have used the below configuration:

- 60 HDDs are part of a single Ultrastar Data60 JBOD. Each Storage Node has 30 x 18TB or 22TB HDDs, which are part of the JBODs.
- 30 HDDs are exposed to each server using zoning configuration. Total of 120 HDDs on 4 Storage Nodes.
- There are two SSDs (7.68 TB each) connected per Storage Node, which are plugged into the backplane of each Storage Node to store the metadata information and have no connectivity to the JBODs (UD60).
- RAID6 (Hardware RAID) configured using (28 + 2) spare drives on each Storage Node.
- For production environment, deploying storage from one UltrastarData60 to multiple Storage Nodes will impact data availability in extreme situations. Please consult Western Digital or NetApp representatives for detail planning.

For Ultrastar Data60 installation and troubleshooting guide refer to [Installation-Guide-Ultrastar-Data60, UltrastarData60 TroubleshootingGuide](#). For more information on configuration & deployment contact [support portal](#).

StorageGRID Configuration Requirements

Before starting a StorageGRID installation, you must:

- Understand StorageGRID's compute requirements, including the minimum CPU and RAM requirements for each node.
- Understand the storage and performance requirements of each type of grid node.

- Understand how StorageGRID supports multiple networks for traffic separation, security, and administrative convenience, and have a plan for which networks you intend to attach to each StorageGRID node. Refer to [storagegrid-network-types](#) and [storagegrid-network-topology](#) for more details.
- Gather all networking information in advance. Unless you are using DHCP, gather the IP addresses to assign to each grid node, and the IP addresses of the domain name system (DNS) and network time protocol (NTP) servers that will be used.

Deploy grid nodes using the appropriate user interface:

When you deploy grid nodes, they are created as part of the StorageGRID system and connected to one or more networks.

These instructions recommend a standard approach for deploying and configuring a StorageGRID system in an Ubuntu or Debian environment.

Pre-requisite:

Before installing StorageGRID software, verify and configure the hardware so that it is ready to support the StorageGRID system.

- Each StorageGRID node requires the following minimum resources: CPU cores: 8 per node.
- RAM: At least 24 GiB per node, and 2 to 16 GiB less than the total system RAM, depending on the total RAM available and the amount of non-StorageGRID software running on the system.
- Ensure that the number of StorageGRID nodes you plan to run on each physical or virtual host does not exceed the number of CPU cores or the physical RAM available. If the hosts are not dedicated to running StorageGRID be sure to consider the resource requirements of the other applications.

LUN purpose	Storage category	Number of LUNs	Minimum size/LUN
Container engine storage pool	Container pool	1	Total number of nodes × 100 GB
/var/local volume	System data	1 for each node on this host	90 GB
Storage Node	Object data	3 for each Storage Node on this host Note: A software-based Storage node can have 1 to 16 storage volumes; at least 3 storage volumes are recommended.	12 TB (4 TB/LUN) See storage requirements for Storage Nodes for more information.
Admin Node audit logs	System data	1 for each Admin node on this host	200 GB
Admin Node tables	System data	1 for each Admin node on this host	200 GB

Note: The maximum tested LUN size is 39 TB.

Depending on the audit level configured, the size of user inputs such as S3 object key name, and how much audit log data you need to preserve, you might need to increase the size of the audit log LUN on each Admin node. A grid generates approximately 1 KB of audit data per S3 operation, which would mean that a 200 GB LUN would support 70 million operations per day or 800 operations per second for two to three days.

Minimum storage space for a host

The following table shows the minimum storage space required for each type of node. You can use this table to determine the minimum amount of storage you must provide to the host in each storage category, based on which nodes will be deployed on that host.

Type of node	Container pool	System data	Object data
Storage Node	100 GB	90 GB	4,000 GB

Admin Node	100 GB	490 GB (3 LUNs)	not applicable
Gateway Node	100 GB	90 GB	not applicable

Note: A software-based Storage node can have 1 to 16 storage volumes - 3 or more storage volumes are recommended. Each storage volume should be 4 TB or larger. For more details regarding the configuration and requirement details, please follow @ [storagegrid-storage-and-performance-requirements.html#example-calculating-the-storage-requirements](#).

Note: For the current deployment, we have created 15*RAID6 volumes/LUNS, each of size 30TB (Hardware RAID configured on the Ultrastar Data60 HDDs) on all the storage nodes to store the storage node data. We have used the 1*RAID1 volume/LUN which is created from the attached SSD disks (size of 7.68TB each) to store metadata information on the storage nodes. For Ultrastar Data60 installation and troubleshooting guide refer @ [Installation-Guide-Ultrastar-Data60](#), [UltrastarData60_TroubleshootingGuide](#). For more information on configuration & deployment contact support @ [support portal](#).

StorageGRID Software Installation and Configuration Details

Install Linux®

You must install Linux on all grid hosts. Use the [NetApp Interoperability Matrix Tool](#) to get a list of supported versions.

1. Ensure that all hosts have access to Ubuntu or Debian package repositories.
2. If swap is enabled:

- a. Run the following command:

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

- b. Remove all swap entries from /etc/fstab to persist the settings.

Note: Failing to disable swap entirely can severely lower performance.

3. Disable individual profiles for the AppArmor packages installed on the base system that overlap with the packages in the StorageGRID system container. When you disable individual profiles, an entry appears in the StorageGRID log files indicating that AppArmor is enabled. Use the following commands:

```
$ sudo ln -s /etc/apparmor.d/<profile.name>
/etc/apparmor.d/disable/
$ sudo apparmor_parser -R /etc/apparmor.d/<profile.name>
```

4. The StorageGRID system runs on Linux as a collection of Docker containers. Before you can install StorageGRID, you must install Docker. Install docker and enable the service. If Docker is not included with your Linux distribution, you can download it from the Docker website.

```
$ sudo systemctl enable docker
$ sudo systemctl start docker
$ root@blr-r20-s9:~# sudo docker version
```

Client:

```
Version:          20.10.21
API version:      1.41
Go version:       go1.18.1
Git commit:       20.10.21-0ubuntu1~22.04.3
Built:            Thu Apr 27 05:57:17 2023
OS/Arch:          linux/amd64
Context:          default
Experimental:    true
```

Server:

```
Engine:
Version:          20.10.21
API version:      1.41 (minimum version 1.12)
Go version:       go1.18.1
Git commit:       20.10.21-0ubuntu1~22.04.3
Built:            Thu Apr 27 05:37:25 2023
OS/Arch:          linux/amd64
Experimental:    false
```

```

containerd:
  Version:          1.5.9-0ubuntu3.1
  GitCommit:
runc:
  Version:          1.1.4-0ubuntu1~22.04.3
  GitCommit:
docker-init:
  Version:          0.19.0
  GitCommit:
root@blr-r20-s9:~#

```

We have used the local storage for the Docker storage volume of the host partition (/var/lib).

5. After completing the Linux installation on your hosts, you might need to perform some additional configuration to prepare a set of network interfaces on each host that are suitable for mapping into the StorageGRID nodes you will deploy later. For more information refer @ [storagegrid-host-network-configuration-details](#).

6. You must allocate block storage volumes to each host. When allocating block storage volumes (LUNs) to hosts, follow the @ [Storage-requirement-guidelines](#).

Note: For the current deployment, we have created 15*RAID6 volumes/LUNS, each of size 30TB (Hardware RAID configured on the Ultrastar Data60 HDDs) on all the storage nodes to store the storage node data. We have used the 1*RAID1 volume/LUN which is created from the attached SSD disks (size of 7.68TB each) to store metadata information on the storage nodes. For Ultrastar Data60 installation and troubleshooting guide refer @ [installation-guide-ultrastar-data60](#), [UltrastarData60_TroubleshootingGuide](#). For more information on configuration & deployment contact support @ [support portal](#).

Note: You will use this information as well as the persistent name assigned by Linux to each physical volume when you deploy StorageGRID nodes on the host. Avoid using “raw” special device files (/dev/sdb, for example) as you compose your list of volume names. These files can change across reboots of the host, which will impact proper operation of the system.

Installation Steps

1. Go to the [NetApp Downloads page for StorageGRID](#).
2. Select the button for downloading the latest release, or select another version from the drop-down menu and select **Go**.
3. Sign in with the username and password for your NetApp account.
4. If a Caution/MustRead statement appears, read it and select the check box.
5. Read the End User License Agreement, select the check box, and then select **Accept & Continue**.
6. In the **Install StorageGRID** column, select the .tgz or .zip file for Ubuntu or Debian.
7. Save and extract the archive file.

Install StorageGRID host services

You use the StorageGRID DEB package to install the StorageGRID host services.

Steps

1. Copy the StorageGRID DEB packages to each of your hosts, or make them available on shared storage.

For example, place them in the /tmp directory, so you can use the example command in the next step.

2. Log in to each host as root or using an account with sudo permission, and run the following commands.

You must install the images package first, and the service package second. If you placed the packages in a directory other than /tmp, modify the command to reflect the path you used.

```

$ sudo dpkg --install /tmp/storagegrid-webscale-images-version-SHA.deb

$ sudo dpkg --install /tmp/storagegrid-webscale-service-version-
SHA.deb

```

Note: Python 2.7 must already be installed before the StorageGRID packages can be installed. The “`sudo dpkg --install /tmp/storagegrid-webscale-images-version-SHA.deb`” command will fail until you have done so.

3. You need to create node configuration files for Ubuntu or Debian deployments. Node configuration files are small text files that provide the information the StorageGRID host service needs to start a node and connect it to the appropriate network, and block storage resources. Node configuration files are used for virtual nodes and are not used for appliance nodes. You must place the configuration file for each StorageGRID node in the `/etc/storagegrid/nodes` directory on the host where the node will run. For example, if you plan to run one Admin node, one Gateway node, and one Storage node on Host1, you must place three node configuration files in `/etc/storagegrid/nodes` on Host. Refer to @ [node-configuration-files](#) for more information regarding all the details.

Note: Any files in `/etc/storagegrid/nodes` that do not follow these naming conventions will not be parsed by the host service.

4. Grid nodes communicate with the primary Admin node for configuration and management. Each grid node must know the IP address of the primary Admin node on the grid network. To ensure that a grid node can access the primary Admin node, you can do either of the following when deploying the node:

- a. You can use the `ADMIN_IP` parameter to manually enter the primary Admin Node’s IP address.
- b. You can omit the `ADMIN_IP` parameter to have the grid node discover the value automatically. Automatic discovery is especially useful when the grid network uses DHCP to assign the IP address to the primary Admin node.

5. **Example node configuration files:** You can use the example node configuration files to help set up the node configuration files for your StorageGRID system. The examples show node configuration files for all types of grid nodes. For more details on different configuration refer @ [storagegrid-node-configuration-example-files](#).

Example for Primary Admin Node

Example file name: `/etc/storagegrid/nodes/dc1-adm1.conf`

Example file contents:

```

NODE_TYPE = VM_Admin_Node
ADMIN_ROLE = Primary
MAXIMUM_RAM = 120g

BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/ubuntu--vg-dc1--adm1--var--local
BLOCK_DEVICE_AUDIT_LOGS = /dev/mapper/ubuntu--vg-dc1--adm1--audit--logs
BLOCK_DEVICE_TABLES = /dev/mapper/ubuntu--vg-dc1--adm1--tables
GRID_NETWORK_TARGET = enp202s0f0np0
ADMIN_NETWORK_TARGET = eno8303
CLIENT_NETWORK_TARGET = enp202s0f1np1

GRID_NETWORK_CONFIG = STATIC
GRID_NETWORK_IP = 192.168.10.51
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 192.168.10.51
GRID_NETWORK_MTU = 9216

ADMIN_NETWORK_CONFIG = STATIC
ADMIN_NETWORK_IP = 100.100.100.101
ADMIN_NETWORK_MASK = 255.255.254.0
ADMIN_NETWORK_GATEWAY = 100.100.100.1
ADMIN_NETWORK_ESL = 100.100.100.0/23
ADMIN_NETWORK_MTU = 1500

```

Note: Here the “`BLOCK_DEVICE_VAR_LOCAL`”, “`BLOCK_DEVICE_AUDIT_LOGS`” and “`BLOCK_DEVICE_TABLES`” volumes are created on the system’s SSD disk.

Example for Storage Node

Example file name: `/etc/storagegrid/nodes/dc1-sn1.conf`

Example file contents:


```

NODE_TYPE = VM_Storage_Node
MAXIMUM_RAM = 120g

ADMIN_IP = 192.168.10.51

BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/ubuntu--vg-dc1--sn1--var--local
BLOCK_DEVICE_RANGEDB_000 = /dev/md/dc1-sn1-rangedb-0
BLOCK_DEVICE_RANGEDB_001 = /dev/disk/by-path/pci-0000:17:00.0-scsi-0:2:0:0
BLOCK_DEVICE_RANGEDB_002 = /dev/disk/by-path/pci-0000:17:00.0-scsi-0:2:1:0
BLOCK_DEVICE_RANGEDB_003 = /dev/disk/by-path/pci-0000:17:00.0-scsi-0:2:10:0
BLOCK_DEVICE_RANGEDB_004 = /dev/disk/by-path/pci-0000:17:00.0-scsi-0:2:11:0
BLOCK_DEVICE_RANGEDB_005 = /dev/disk/by-path/pci-0000:17:00.0-scsi-0:2:12:0
BLOCK_DEVICE_RANGEDB_006 = /dev/disk/by-path/pci-0000:17:00.0-scsi-0:2:13:0
BLOCK_DEVICE_RANGEDB_007 = /dev/disk/by-path/pci-0000:17:00.0-scsi-0:2:14:0
BLOCK_DEVICE_RANGEDB_008 = /dev/disk/by-path/pci-0000:17:00.0-scsi-0:2:2:0
BLOCK_DEVICE_RANGEDB_009 = /dev/disk/by-path/pci-0000:17:00.0-scsi-0:2:3:0
BLOCK_DEVICE_RANGEDB_010 = /dev/disk/by-path/pci-0000:17:00.0-scsi-0:2:4:0
BLOCK_DEVICE_RANGEDB_011 = /dev/disk/by-path/pci-0000:17:00.0-scsi-0:2:5:0
BLOCK_DEVICE_RANGEDB_012 = /dev/disk/by-path/pci-0000:17:00.0-scsi-0:2:6:0
BLOCK_DEVICE_RANGEDB_013 = /dev/disk/by-path/pci-0000:17:00.0-scsi-0:2:7:0
BLOCK_DEVICE_RANGEDB_014 = /dev/disk/by-path/pci-0000:17:00.0-scsi-0:2:8:0
BLOCK_DEVICE_RANGEDB_015 = /dev/disk/by-path/pci-0000:17:00.0-scsi-0:2:9:0
GRID_NETWORK_TARGET = enp152s0f0np0
ADMIN_NETWORK_TARGET = eno8303
#CLIENT_NETWORK_TARGET = bond0.1003

GRID_NETWORK_CONFIG = STATIC
GRID_NETWORK_IP = 192.168.10.61
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 192.168.10.51
GRID_NETWORK_MTU = 9216

ADMIN_NETWORK_CONFIG = STATIC
ADMIN_NETWORK_IP = 100.100.100.111
ADMIN_NETWORK_MASK = 255.255.254.0
ADMIN_NETWORK_GATEWAY = 100.100.100.1
ADMIN_NETWORK_MTU = 1500

```

Note: Here the “BLOCK_DEVICE_VAR_LOCAL” is created on the system’s SSD disk. And “BLOCK_DEVICE_RANGEDB_000” is RAID1 volume created using the locally attached SSD disks. And “BLOCK_DEVICE_RANGEDB_001-015” are RAID6 volumes created from the Ultrastar Data60 HDD disks.

Example for Gateway Node

Example file name: /etc/storagegrid/nodes/dc1-gw1.conf

Example file contents:

```

NODE_TYPE = VM_API_Gateway
MAXIMUM_RAM = 120g

ADMIN_IP = 192.168.10.51

BLOCK_DEVICE_VAR_LOCAL = /dev/mapper/ubuntu--vg-dc1--gw1--var--local
GRID_NETWORK_TARGET = enp202s0f0np0
ADMIN_NETWORK_TARGET = eno8303
CLIENT_NETWORK_TARGET = enp202s0f1np1

GRID_NETWORK_CONFIG = STATIC
GRID_NETWORK_IP = 192.168.10.53
GRID_NETWORK_MASK = 255.255.255.0
GRID_NETWORK_GATEWAY = 192.168.10.51
GRID_NETWORK_MTU = 9216

ADMIN_NETWORK_CONFIG = STATIC
ADMIN_NETWORK_IP = 100.100.100.112
ADMIN_NETWORK_MASK = 255.255.254.0

```



```
ADMIN_NETWORK_GATEWAY = 100.100.100.1
ADMIN_NETWORK_MTU = 1500
```

Note: Here the “BLOCK_DEVICE_VAR_LOCAL” volume is created on the system’s SSD disk.

6. After creating 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, as shown below.

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

7. You need to start the StorageGRID services at the hosts. To start your StorageGRID nodes, and ensure they restart after a host reboot, you must enable and start the StorageGRID host service.

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

Run the following command to ensure the deployment is proceeding:

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

8. Use the Grid Manager to define all of the information required to configure your StorageGRID system. Open your web browser and navigate to one of the following addresses:

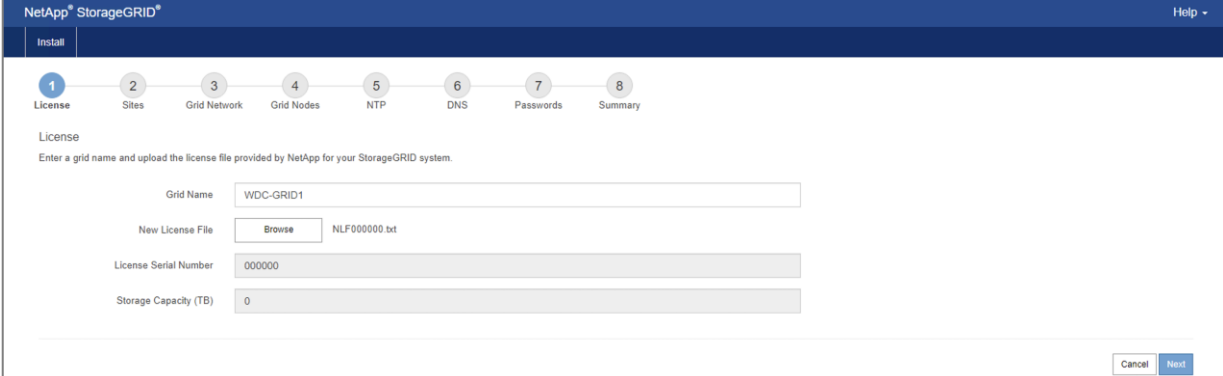
```
$ https://primary_admin_node_ip
```

Alternatively, you can access the Grid Manager on port 8443:

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

You can use the IP address for the primary Admin Node IP on the Grid Network or on the Admin Network, as appropriate for your network configuration.

9. Click **Install a StorageGRID** system. The page used to configure a StorageGRID appears as below.

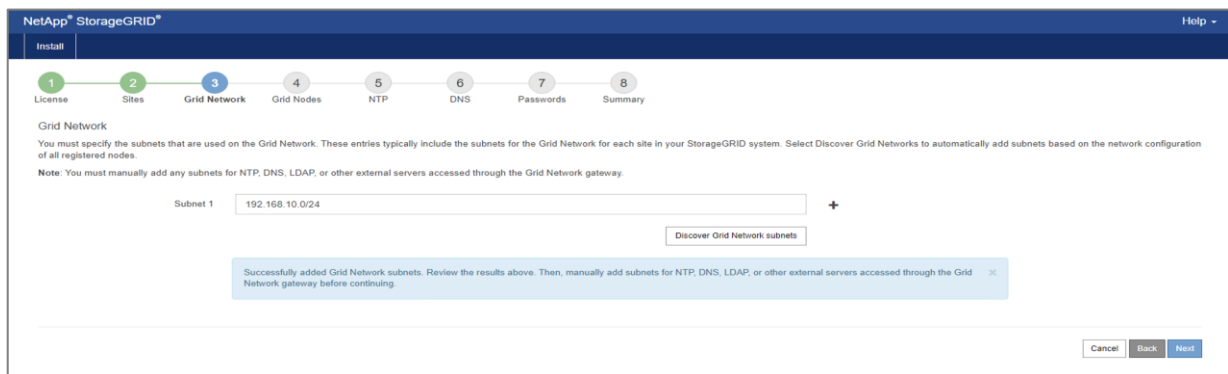


10. On the License page, enter a meaningful name for your StorageGRID system in **Grid Name**.

After installation, the name is displayed at the top of the Nodes menu. Click **Browse**, locate the NetApp License File (NLFunique_id.txt), and click **Open**. The license file is validated, and the serial number and licensed storage capacity are displayed. Click **Next**.

11. You must create at least one site when you are installing StorageGRID. You can create additional sites to increase the reliability and storage capacity of your StorageGRID system. On the Sites page, enter the **Site Name**. 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. Click **Next**.

12. 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, along with any subnets that need to be reachable via the Grid Network.



13. You must approve each grid node before it can join the StorageGRID system. Review the Pending Nodes list, and confirm that it shows all of the grid nodes you deployed. Select the radio button next to a pending node you want to approve.

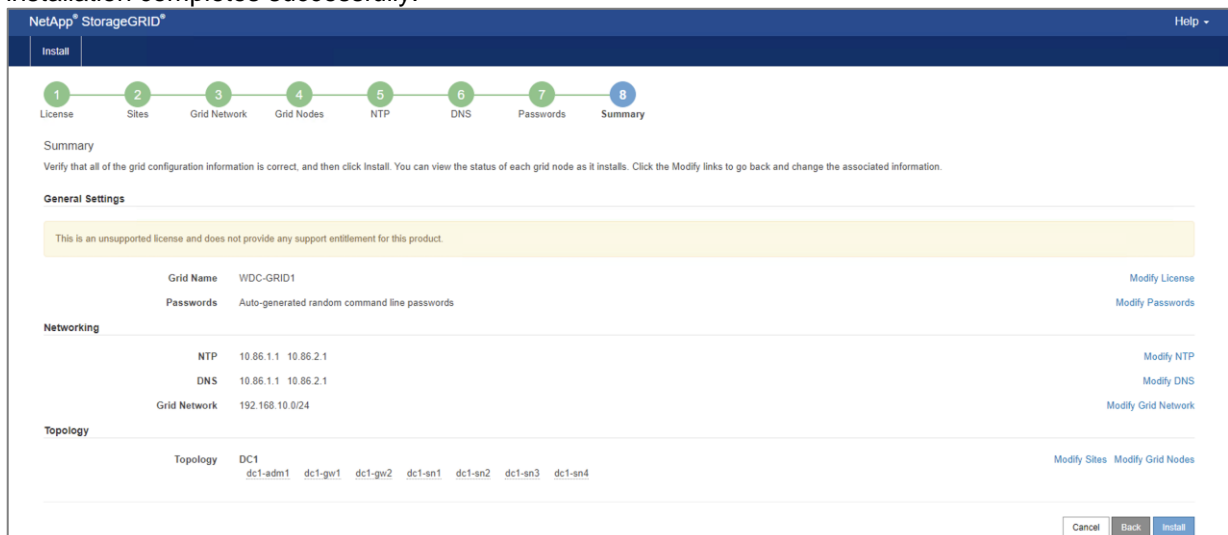
14. You must specify the **Network Time Protocol (NTP)** configuration information for the StorageGRID system, so that operations performed on separate servers can be kept synchronized.

15. You must specify **Domain Name System (DNS)** information for your StorageGRID system, so that you can access external servers using hostnames instead of IP addresses.

16. As part of installing your StorageGRID system, you need to enter the passwords to use to secure your system and perform maintenance tasks. Use the **Install Passwords** page to specify the provisioning passphrase and the grid management root user password.

- The provisioning passphrase is used as an encryption key and is not stored by the StorageGRID system.
- You must have the provisioning passphrase for installation, expansion, and maintenance procedures, including downloading the **Recovery Package**. Therefore, it is important that you store the provisioning passphrase in a secure location.
- You can change the provisioning passphrase from the **Grid Manager** if you have the current one.
- The grid management root user password may be changed using the Grid Manager.
- Randomly generated command line console and SSH passwords are stored in the Passwords.txt file in the Recovery Package.

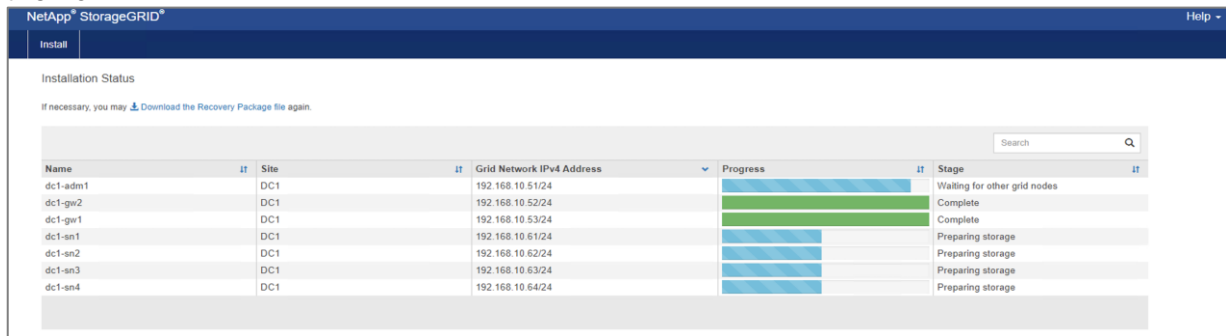
17. You must carefully review the configuration information you have entered to ensure that the installation completes successfully.



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. Click **Install**.

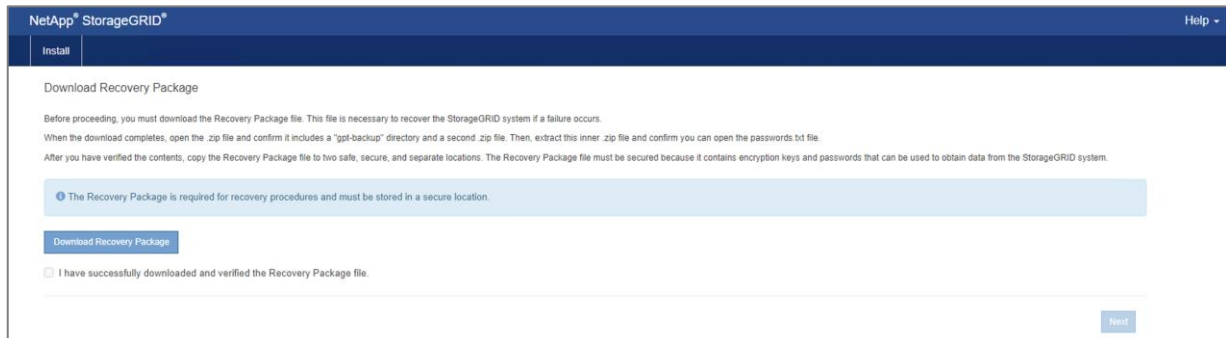
18. When the installation progresses to the point where the grid topology is defined, you are prompted to download the **Recovery Package** file (.zip), and confirm that you can successfully access the contents of this file. You must download the Recovery Package file so that you can recover the

StorageGRID system if one or more grid nodes fail. The installation continues in the background, but you cannot complete the installation and access the StorageGRID system until you download and verify this file.



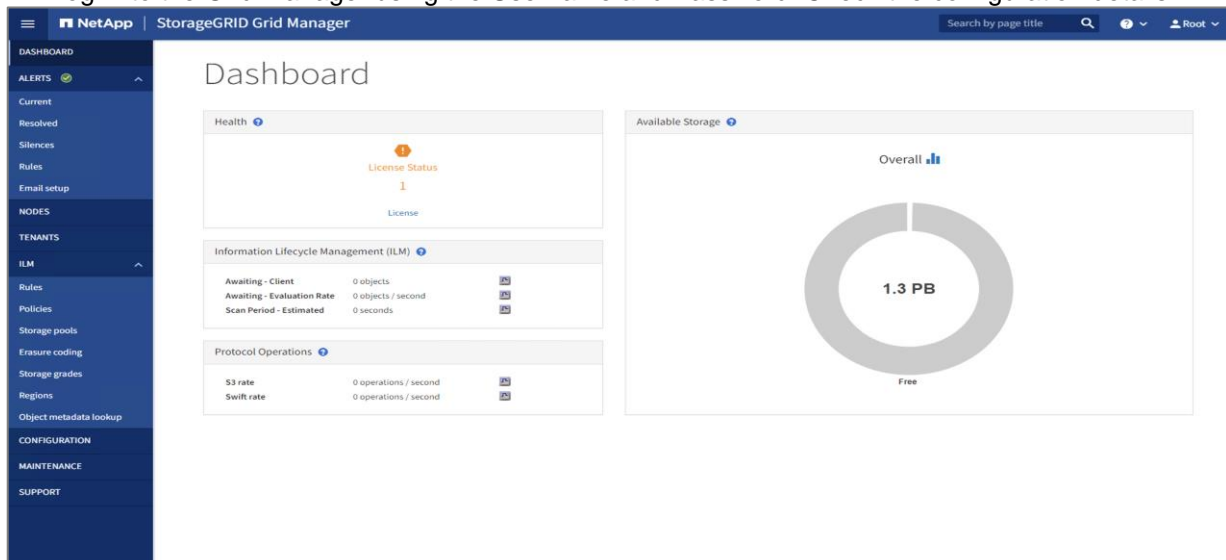
19. Verify that you can extract the contents of the .zip file, and then save it in two safe, secure, and separate locations.

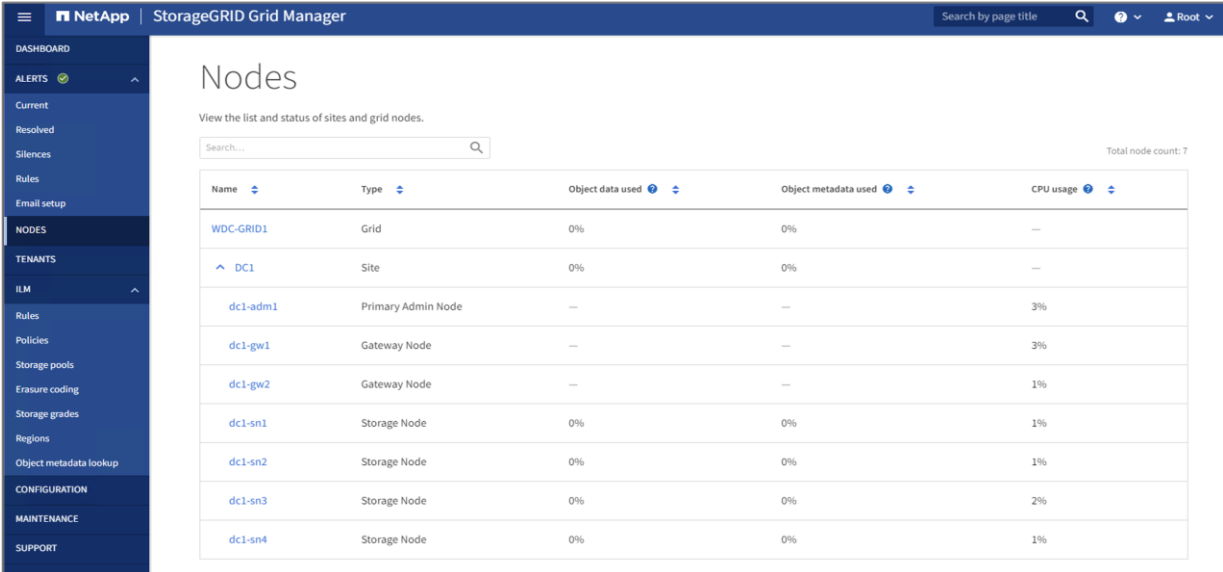
20. Select the **I have successfully downloaded and verified the Recovery Package file** check box, and click **Next**.



21. When the Complete stage is reached for all grid nodes, the **sign-in** page for the Grid Manager appears. Sign in to the Grid Manager using the "root" user and the password you specified during the installation.

22. Log in to the Grid Manager using the Username and Password. Check the configuration details.





Name	Type	Object data used	Object metadata used	CPU usage
WDC-GRID1	Grid	0%	0%	—
DC1	Site	0%	0%	—
dc1-adm1	Primary Admin Node	—	—	3%
dc1-gw1	Gateway Node	—	—	3%
dc1-gw2	Gateway Node	—	—	1%
dc1-sn1	Storage Node	0%	0%	1%
dc1-sn2	Storage Node	0%	0%	1%
dc1-sn3	Storage Node	0%	0%	2%
dc1-sn4	Storage Node	0%	0%	1%

23. You must create at least one tenant account to control access to the storage in your StorageGRID system. When you create a tenant account, you specify a name, a client protocol, and optionally a storage quota. If single sign-on (SSO) is enabled for StorageGRID, you also specify which federated group has root access permission to configure the tenant account. If StorageGRID is not using single sign-on, you must also specify whether the tenant account will use its own identity source and configure the initial password for the tenant's local root user. For more information refer to [tenant-account-creation](#).

24. Configuring StorageGRID to accept connections from client applications requires creating one or more tenant accounts and setting up the connections. An S3 tenant account is required before S3 API clients can store and retrieve objects on StorageGRID. Each tenant account has its own account ID, groups, users, containers, and objects. S3 tenant accounts are created by a StorageGRID grid administrator using the Grid Manager or the Grid Management API. When creating an S3 tenant account, the grid administrator specifies the following information:

- Display name for the tenant (the tenant's account ID is assigned automatically and cannot be changed).
- Whether the tenant account is allowed to use platform services. If the use of platform services is allowed, the grid must be configured to support their use.
- Optionally, a storage quota for the tenant account - the maximum number of gigabytes, terabytes, or petabytes available for the tenant's objects. A tenant's storage quota represents a logical amount (object size), not a physical amount (size on disk).
- If identity federation is enabled for the StorageGRID system, which federated group has root access permission to configure the tenant account.
- If single sign-on (SSO) is not in use for the StorageGRID system, whether the tenant account will use its own identity source or share the grid's identity source, and the initial password for the tenant's local root user.

25. After an S3 tenant account is created, tenant users can access the Tenant Manager to perform following tasks:

- Set up identity federation (unless the identity source is shared with the grid) and create local groups and users.
- Manage S3 access keys.
- Create and manage S3 buckets, including buckets that have S3 Object Lock enabled.
- Use platform services (if enabled).
- Monitor storage usage.

26. A grid administrator makes configuration choices that affect how S3 clients connect to StorageGRID to store and retrieve data. The specific information you need to make a connection depends on the configuration that was chosen. Client applications can store or retrieve objects by connecting to any of the following:

- The load balancer service on Admin nodes or Gateway nodes, or optionally, the virtual IP address of a high availability (HA) group of Admin nodes or Gateway nodes.

- The CLB service on Gateway nodes, or optionally, the virtual IP address of a high availability group of Gateway nodes.
 - Storage nodes, with or without an external load balancer.
27. When configuring StorageGRID, a grid administrator can use the Grid Manager or the Grid Management API to perform the following steps, all of which are optional:
- **Configure endpoints for the load balancer service:** You must configure endpoints to use the load balancer service. The load balancer service on Admin nodes or Gateway nodes distributes incoming network connections from client applications to Storage nodes. When creating a load balancer endpoint, the StorageGRID administrator specifies a port number, whether the endpoint accepts HTTP or HTTPS connections, the type of client (S3) that will use the endpoint, and the certificate to be used for HTTPS connections (if applicable).
 - **Configure untrusted client networks:** If a StorageGRID administrator configures a node's client network to be untrusted, the node only accepts inbound connections on the client network on ports that are explicitly configured as load balancer endpoints.
 - **Configure high availability groups:** If an administrator creates an HA group, the network interfaces of multiple Admin nodes or Gateway nodes are placed into an active-backup configuration. Client connections are made using the virtual IP address of the HA group.
28. Client applications connect to StorageGRID using the IP address of a grid node and the port number of a service on that node. If high availability (HA) groups are configured, client applications can connect using the virtual IP address of the HA group. For more information on client configuration refer to [client-connections](#).

For example, if the virtual IP address of the HA group is 192.0.2.5 and the port number of an S3 load balancer endpoint is 10443, then an S3 client could use the following URL to connect to StorageGRID : <https://192.0.2.5:10443>.

Performance Result

We have used s3tester to run the performance benchmark on the StorageGrid system. The s3tester needs to be configured on all the clients before proceeding with the performance test. Refer @ [s3tester](#) for the configuration details.

Examples

Writing objects into a bucket

```
$ ./s3tester -concurrency=128 -size=20000000 -operation=put -requests=20000
-endpoint="https://10.96.105.5:18443" -prefix=3
```

- Starts writing objects into the default bucket test.
- The bucket needs to be created prior to running s3tester.
- The naming of the ingested objects will be 3-object

Here 3 is the prefix specified and object is a sequential number starting from zero and going to the number of requests.

- This command will perform a total of 20,000 PUT requests (or in this case slightly less because 20,000 does not divide by 128).
- The object size is 20,000,000 bytes.
- Replace the sample IP/port combination with the one you are using.

Reading objects from a bucket (and other operations)

```
$ ./s3tester -concurrency=128 -operation=get -requests=200000 -
endpoint="https://10.96.105.5:18443" -prefix=3
```

- Matches the request above and will read the same objects written in the same sequence.
- If you use the randget operation the objects will be read in random order simulating a random-access workload.
- If you use the head operation then the S3 HEAD operation will be performed against the objects in sequence.

- If you use the delete operation then the objects will be deleted.

The performance test is run using s3tester, simultaneously running tests from three clients. We have used different object sizes (1M, 2M, 4M, 8M, 10M, 32M, 64M, 128M and 265M etc.), different threads, and different policies to measure the performance. Refer the below details for the maximum performance achieved with the NetApp StorageGRID with Ultrastar Data60. Contact support @ [support portal](#) for the complete test reports.

Performance Results

StorageGRID Performance Results		
Storage Policy	Read BW	Write BW
Replication (Two-copy policy)	6.75 GBps	3.44 GBps
EC (2+1)	6.71 GBps	2.66 GBps

Features and Benefits

- Massive scalability and flexible infrastructure.
- Massively parallel transaction engine with integrated load balancing.
- Transaction multithread pipelining.
- Advanced security and encryption capabilities.
- Native support for Amazon S3.
- Metadata and content awareness.
- Cloud platform services.
- Service-level objective and performance monitoring.
- Management and monitoring.
- Reinforced data integrity with compliance-grade WORM.
- Flexible data protection through various replication and erasure coding schemes.
- Bucket-level granularity for all storage policies.
- Centralized and automatable installation and expansions.
- Automated monitoring and tenant management through an API.
- Reduced deployment risk, streamlined implementation, and the ability to migrate quickly with minimal disruption.

Use Cases

- Backup and disaster recovery
- File consolidation
- Media active archive
- Collaboration
- File lifecycle management
- AI / Machine Learning
- Data analytics
- Storage-as-a-Service (STaaS)

Conclusion

In a world full of generalists, Western Digital is a specialist. We're focused on one thing, helping your business get the most out of its data. Western Digital and NetApp together, are offering enterprise-grade data services you can rely on. Our industry-leading solutions work across diverse customer environments and the world's biggest public clouds.

As a globally leading IT provider, Western Digital consistently monitors the changing technological, strategic, and organizational requirements in the datacenter. Our observations show that as the complexities of IT environments have increased in general, so have the complexities of data protection environments in particular. Efficient data protection and data availability in all scenarios call for infrastructure and management consolidation, with the main goal of ensuring high data availability and quick and fast disaster recoveries. Together, Western Digital's Ultrastar Data-60 and NetApp

StorageGRID enable the tiering of file data to the cloud and lower the total cost of ownership for storage. To meet your requirements, we are offering a consistent data protection strategy with our solution.

By combining NetApp StorageGRID with Western Digital's Ultrastar Data60 platform, organizations can avail:

- Unlimited Scale
- Multidata-center storage
- Fully automated data tiering
- Resilient, high-availability storage
- Data Integrity and Protection
- Minimize the cost of the installation
- Analyse petabytes of data in hours, not weeks
- Access data natively in the cloud
- Cut cold data costs
- Migrate without headache
- Reduce Your Data center Footprint
- Lower Total Cost of Ownership

Contributors

Name	Company	Title
Puspanjali Panda	Western Digital	Principal Engineer, Test Engineering
Saravanakumar Pandian	Western Digital	Principal Engineer, Test Engineering
Pavan Gururaj	Western Digital	Senior Manager, Test Engineering
Raymond Lee	NetApp	Solution Architect
Narender Gupta	NetApp	Solution Engineer

Version History

Version	Revision Date	Notes
01	July 2023	Initial Release

Document Feedback

For feedback, questions, and suggestions for improvements to this document send an email to the Data Center Systems (DCS) Technical Marketing Engineering (TME) team distribution list at pd1-dcs-tm@wdc.com.

References

- https://documents.westerndigital.com/content/dam/doc-library/en_us/assets/public/western-digital/product/platforms/ultrastar-data60-hybrid-platform/product-brief-ultrastar-data60-hybrid-storage-platform.pdf
- https://documents.westerndigital.com/content/dam/doc-library/en_us/assets/public/western-digital/product/platforms/ultrastar-data60-hybrid-platform/installation-guide-ultrastar-data60.pdf
- https://support-en.westerndigital.com/ci/fattach/get/5172196/1676392546/redirect/1/filename/UltrastarData102_60TroubleshootingGuide_01.pdf
- https://portal.wdc.com/s/login/?!language=en_US&ec=302&startURL=%2Fs%2F
- <https://docs.netapp.com/us-en/storagegrid-116/ubuntu/storage-and-performance-requirements.html#example-calculating-the-storage-requirements-for-a-host>
- <https://docs.netapp.com/us-en/storagegrid-116/ubuntu/configuring-host-network.html>

- <https://docs.netapp.com/us-en/storagegrid-116/ubuntu/storage-and-performance-requirements.html#storage-requirements-for-storage-nodes>
- <https://mysupport.netapp.com/site/products/all/details/storagegrid/downloads-tab>
- <https://docs.netapp.com/us-en/storagegrid-116/ubuntu/creating-node-configuration-files.html#what-is-in-a-node-configuration-file>
- <https://docs.netapp.com/us-en/storagegrid-116/ubuntu/example-node-configuration-files.html#example-for-primary-admin-node>
- <https://docs.netapp.com/us-en/storagegrid-116/admin/creating-tenant-account.html>
- <https://docs.netapp.com/us-en/storagegrid-116/s3/configuring-tenant-accounts-and-connections.html#summary-ip-addresses-and-ports-for-client-connections>
- <https://github.com/s3tester/s3tester>
- <https://www.westerndigital.com/en-us/solutions/data-center#scale-efficiently>
- <https://docs.netapp.com/us-en/storagegrid-116/network/grid-network-topology.html>
- <https://docs.netapp.com/us-en/storagegrid-116/network/topology-for-all-three-networks.html>



5601 Great Oaks Parkway
San Jose, CA 95119, USA
www.westerndigital.com

© 2023 Western Digital Corporation or its affiliates. All rights reserved. Western Digital, the Western Digital design, the Western Digital logo, ArcticFlow, HelioSeal, IsoVibe, and Ultrastar are registered trademarks or trademarks of Western Digital Corporation or its affiliates in the US and/or other countries. Amazon Simple Storage Service is a trademark of Amazon.com, Inc. or its affiliates. Intel and Xeon are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. NetApp®, The NetAPP logo, and StorageGRID® are trademarks of NetApp, Inc. and are registered in the United States and/or other jurisdictions. All other marks are the property of their respective owners. Products may not be available in all regions. Images may differ from actual products.
