



Install VMware

StorageGRID

NetApp
October 03, 2025

This PDF was generated from <https://docs.netapp.com/us-en/storagegrid-116/vmware/index.html> on October 03, 2025. Always check docs.netapp.com for the latest.

Table of Contents

Install VMware	1
Install VMware: Overview	1
Plan and prepare for VMware installation	1
Before you install (VMware)	1
Required materials	2
Download and extract the StorageGRID installation files	3
Software requirements	5
CPU and RAM requirements	6
Storage and performance requirements	7
Deploy virtual machine grid nodes (VMware)	9
Collect information about your deployment environment	9
How grid nodes discover the primary Admin Node	11
Deploy a StorageGRID node as a virtual machine	12
Configure the grid and complete installation (VMware)	18
Navigate to the Grid Manager	18
Specify the StorageGRID license information	18
Add sites	19
Specify Grid Network subnets	20
Approve pending grid nodes	21
Specify Network Time Protocol server information	26
Specify Domain Name System server information	27
Specify the StorageGRID system passwords	28
Review your configuration and complete installation	30
Post-installation guidelines	32
Automate the installation (VMware)	32
Automate grid node deployment	32
Run the Bash script	42
Automate the configuration of StorageGRID	43
Overview of the installation REST API	45
StorageGRID Installation API	45
Where to go next	46
Required tasks	46
Optional tasks	47
Troubleshoot installation issues	47
Virtual machine resource reservation requires adjustment	48

Install VMware

Install VMware: Overview

Installing a StorageGRID system in a VMware environment includes three primary steps.

1. Preparation: During planning and preparation, you perform the following tasks:

- Learn about the hardware, software, virtual machine, storage, and performance requirements for StorageGRID.
- Learn about the specifics of [StorageGRID networking](#) so you can configure your network appropriately.
- Identify and prepare the physical servers you plan to use to host your StorageGRID grid nodes.
- On the servers you have prepared:
 - Install VMware vSphere Hypervisor
 - Configure the ESX hosts
 - Install and configure VMware vSphere and vCenter

2. Deployment: Deploy grid nodes using the VMware vSphere Web Client. When you deploy grid nodes, they are created as part of the StorageGRID system and connected to one or more networks.

- a. Use the VMware vSphere Web Client, a .vmdk file, and a set of .ovf file templates to deploy the software-based nodes as virtual machines (VMs) on the servers you prepared in step 1.
- b. Use the StorageGRID Appliance Installer to deploy StorageGRID appliance nodes.



Hardware-specific installation and integration instructions are not included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the installation and maintenance instructions for your appliance.

3. Configuration: When all nodes have been deployed, use the Grid Manager to configure the grid and complete the installation.

These instructions recommend a standard approach for deploying and configuring a StorageGRID system in a VMware environment. See also the information about the following alternative approaches:

- Use the deploy-vsphere-ovftool.sh Bash script (available from the installation archive) to deploy grid nodes in VMware vSphere.
- Automate the deployment and configuration of the StorageGRID system using a Python configuration script (provided in the installation archive).
- Automate the deployment and configuration of appliance grid nodes with a Python configuration script (available from the installation archive or from the StorageGRID Appliance Installer).
- If you are an advanced developer of StorageGRID deployments, use the installation REST APIs to automate the installation of StorageGRID grid nodes.

Plan and prepare for VMware installation

Before you install (VMware)

Before deploying grid nodes and configuring the StorageGRID grid, you must be familiar

with the steps and requirements for completing the procedure.

The StorageGRID deployment and configuration procedures assume that you are familiar with the architecture and operational functionality of the StorageGRID system.

You can deploy a single site or multiple sites at one time; however, all sites must meet the minimum requirement of having at least three Storage Nodes.



StorageGRID does not support the use of virtual storage area networks (vSANs), because the underlying disk protection is not hardware RAID.

Before starting the node deployment and grid configuration procedure, you must:

- Plan the StorageGRID deployment.
- Install, connect, and configure all required hardware, including any StorageGRID appliances, to specifications.



Hardware-specific installation and integration instructions are not included in the StorageGRID installation procedure. To learn how to install StorageGRID appliances, see the installation and maintenance instructions for your appliance.

- Understand the [available network options and how each network option should be implemented on grid nodes](#).
- 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.
- Decide which of the available deployment and configuration tools you want to use.

Related information

[SG100 and SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

Required materials

Before you install StorageGRID, you must gather and prepare required materials.

Item	Notes
NetApp StorageGRID license	You must have a valid, digitally signed NetApp license. Note: The StorageGRID installation archive includes a free license that does not provide any support entitlement for the product.
StorageGRID installation archive	You must download the StorageGRID installation archive and extract the files .

Item	Notes
VMware software and documentation	During installation, you use VMware vSphere Web Client to deploy virtual grid nodes on virtual machines. For supported versions, see the Interoperability Matrix.
Service laptop	<p>The StorageGRID system is installed through a service laptop. The service laptop must have:</p> <ul style="list-style-type: none"> • Network port • SSH client (for example, PuTTY) • Supported web browser
StorageGRID documentation	<ul style="list-style-type: none"> • Release notes • Instructions for administering StorageGRID

Related information

[NetApp Interoperability Matrix Tool](#)

Download and extract the StorageGRID installation files

You must download the StorageGRID installation archives and extract the files..

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.



You must apply any required hotfixes after you install the StorageGRID release. For more information, see the [hotfix procedure in the recovery and maintenance instructions](#)

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 VMware.



Use the .zip file if you are running Windows on the service laptop.

7. Save and extract the archive file.
8. Choose the files you need from the following list.

The files you need depend on your planned grid topology and how you will deploy your StorageGRID system.



The paths listed in the table are relative to the top-level directory installed by the extracted installation archive.

Path and file name	Description
./vsphere/README	A text file that describes all of the files contained in the StorageGRID download file.
./vsphere/NLF000000.txt	A free license that does not provide any support entitlement for the product.
./vsphere/NetApp-SG-version-SHA.vmdk	The virtual machine disk file that is used as a template for creating grid node virtual machines.
./vsphere/vsphere-primary-admin.ovf ./vsphere/vsphere-primary-admin.mf	The Open Virtualization Format template file (.ovf) and manifest file (.mf) for deploying the primary Admin Node.
./vsphere/vsphere-non-primary-admin.ovf ./vsphere/vsphere-non-primary-admin.mf	The template file (.ovf) and manifest file (.mf) for deploying non-primary Admin Nodes.
./vsphere/vsphere-archive.ovf ./vsphere/vsphere-archive.mf	The template file (.ovf) and manifest file (.mf) for deploying Archive Nodes.
./vsphere/vsphere-gateway.ovf ./vsphere/vsphere-gateway.mf	The template file (.ovf) and manifest file (.mf) for deploying Gateway Nodes.
./vsphere/vsphere-storage.ovf ./vsphere/vsphere-storage.mf	The template file (.ovf) and manifest file (.mf) for deploying virtual machine-based Storage Nodes.
Deployment scripting tool	Description
./vsphere/deploy-vsphere-ovftool.sh	A Bash shell script used to automate the deployment of virtual grid nodes.
./vsphere/deploy-vsphere-ovftool-sample.ini	An example configuration file for use with the deploy-vsphere-ovftool.sh script.
./vsphere/configure-storagegrid.py	A Python script used to automate the configuration of a StorageGRID system.
./vsphere/configure-sga.py	A Python script used to automate the configuration of StorageGRID appliances.
./vsphere/storagegrid-ssoauth.py	An example Python script that you can use to sign in to the Grid Management API when single sign-on is enabled.
./vsphere/configure-storagegrid.sample.json	An example configuration file for use with the configure-storagegrid.py script.

Path and file name	Description
./vsphere/configure-storagegrid.blank.json	A blank configuration file for use with the <code>configure-storagegrid.py</code> script.
./vsphere/extras/api-schemas	<p>API schemas for StorageGRID.</p> <p>Note: Before you perform an upgrade, you can use these schemas to confirm that any code you have written to use StorageGRID management APIs will be compatible with the new StorageGRID release if you do not have a non-production StorageGRID environment for upgrade compatibility testing.</p>

Related information

[Recover and maintain](#)

Software requirements

You can use a virtual machine to host any type of StorageGRID grid node. One virtual machine is required for each grid node installed on the VMware server.

VMware vSphere Hypervisor

You must install VMware vSphere Hypervisor on a prepared physical server. The hardware must be configured correctly (including firmware versions and BIOS settings) before you install VMware software.

- Configure networking in the hypervisor as required to support networking for the StorageGRID system you are installing.

Networking guidelines

- Ensure that the datastore is large enough for the virtual machines and virtual disks that are required to host the grid nodes.
- If you create more than one datastore, name each datastore so that you can easily identify which datastore to use for each grid node when you create virtual machines.

ESX host configuration requirements



You must properly configure the network time protocol (NTP) on each ESX host. If the host time is incorrect, negative effects, including data loss, could occur.

VMware configuration requirements

You must install and configure VMware vSphere and vCenter before deploying StorageGRID grid nodes.

For supported versions of VMware vSphere Hypervisor and VMware vCenter Server software, see the Interoperability Matrix.

For the steps required to install these VMware products, see the VMware documentation.

Related information

[NetApp Interoperability Matrix Tool](#)

CPU and RAM requirements

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

For information about supported servers, see the Interoperability Matrix.

Each StorageGRID node requires the following minimum resources:

- CPU cores: 8 per node
- RAM: At least 24 GB per node, and 2 to 16 GB 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 (not recommended), be sure to consider the resource requirements of the other applications.

 Monitor your CPU and memory usage regularly to ensure that these resources continue to accommodate your workload. For example, doubling the RAM and CPU allocation for virtual Storage Nodes would provide similar resources to those provided for StorageGRID appliance nodes. Additionally, if the amount of metadata per node exceeds 500 GB, consider increasing the RAM per node to 48 GB or more. For information about managing object metadata storage, increasing the Metadata Reserved Space setting, and monitoring CPU and memory usage, see the instructions for administering, monitoring, and upgrading StorageGRID.

If hyperthreading is enabled on the underlying physical hosts, you can provide 8 virtual cores (4 physical cores) per node. If hyperthreading is not enabled on the underlying physical hosts, you must provide 8 physical cores per node.

If you are using virtual machines as hosts and have control over the size and number of VMs, you should use a single VM for each StorageGRID node and size the VM accordingly.

For production deployments, you should not run multiple Storage Nodes on the same physical storage hardware or virtual host. Each Storage Node in a single StorageGRID deployment should be in its own isolated failure domain. You can maximize the durability and availability of object data if you ensure that a single hardware failure can only impact a single Storage Node.

See also the information about storage requirements.

Related information

[NetApp Interoperability Matrix Tool](#)

[Storage and performance requirements](#)

[Administer StorageGRID](#)

[Monitor and troubleshoot](#)

[Upgrade software](#)

Storage and performance requirements

You must understand the storage and performance requirements for StorageGRID nodes hosted by virtual machines, so you can provide enough space to support the initial configuration and future storage expansion.

Performance requirements

The performance of the OS volume and of the first storage volume significantly impacts the overall performance of the system. Ensure that these provide adequate disk performance in terms of latency, input/output operations per second (IOPS), and throughput.

All StorageGRID nodes require that the OS drive and all storage volumes have write-back caching enabled. The cache must be on a protected or persistent media.

Requirements for virtual machines that use NetApp ONTAP storage

If you are deploying a StorageGRID node as a virtual machine with storage assigned from a NetApp ONTAP system, you have confirmed that the volume does not have a FabricPool tiering policy enabled. For example, if a StorageGRID node is running as a virtual machine on a VMWare host, ensure the volume backing the datastore for the node does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.

 Never use FabricPool to tier any data related to StorageGRID back to StorageGRID itself. Tiering StorageGRID data back to StorageGRID increases troubleshooting and operational complexity.

Number of virtual machines required

Each StorageGRID site requires a minimum of three Storage Nodes.

 In a production deployment, do not run more than one Storage Node on a single virtual machine server. Using a dedicated virtual machine host for each Storage Node provides an isolated failure domain.

Other types of nodes, such as Admin Nodes or Gateway Nodes, can be deployed on the same virtual machine host, or they can be deployed on their own dedicated virtual machine hosts as required. However, if you have multiple nodes of the same type (two Gateway Nodes, for example), do not install all instances on the same virtual machine host.

Storage requirements by node type

In a production environment, the virtual machines for StorageGRID grid nodes must meet different requirements, depending on the types of nodes.

 Disk snapshots cannot be used to restore grid nodes. Instead, refer to the recovery and maintenance procedures for each type of node.

Node Type	Storage
Admin Node	100 GB LUN for OS 200 GB LUN for Admin Node tables 200 GB LUN for Admin Node audit log
Storage Node	100 GB LUN for OS 3 LUNs for each Storage Node on this host Note: A Storage Node can have 1 to 16 storage LUNs; at least 3 storage LUNs are recommended. Minimum size per LUN: 4 TB Maximum tested LUN size: 39 TB.
Gateway Node	100 GB LUN for OS
Archive Node	100 GB LUN for OS

 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. As a general rule, 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.

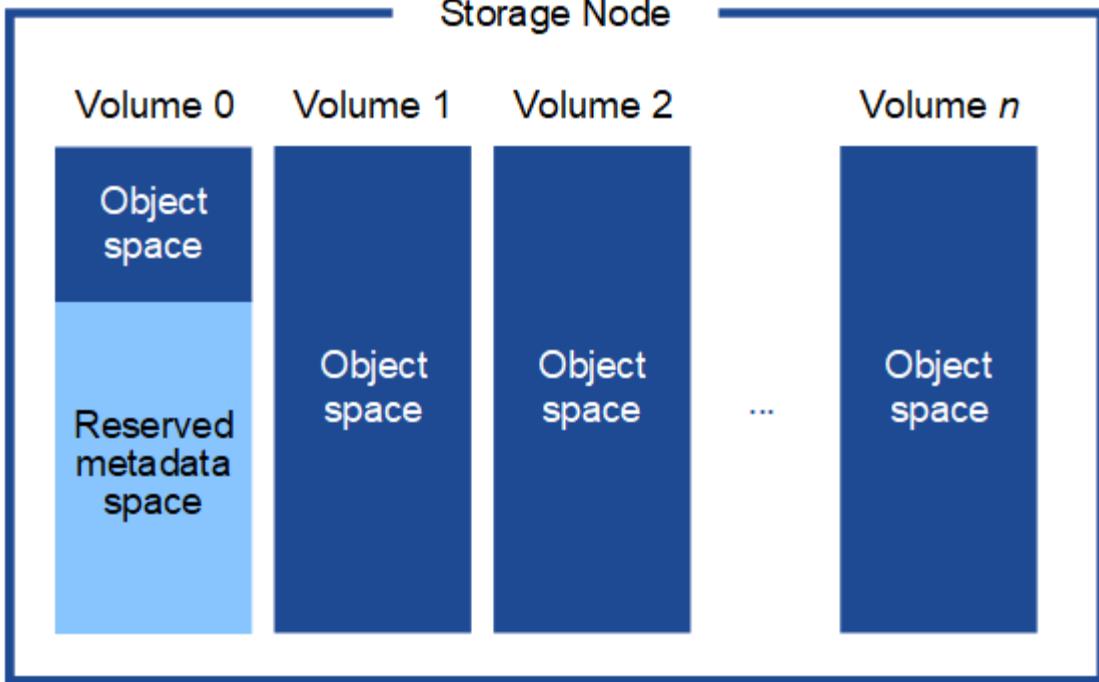
Storage requirements for Storage Nodes

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.



An appliance Storage Node can have up to 48 storage volumes.

As shown in the figure, StorageGRID reserves space for object metadata on storage volume 0 of each Storage Node. Any remaining space on storage volume 0 and any other storage volumes in the Storage Node are used exclusively for object data.



To provide redundancy and to protect object metadata from loss, StorageGRID stores three copies of the metadata for all objects in the system at each site. The three copies of object metadata are evenly distributed across all Storage Nodes at each site.

When you assign space to volume 0 of a new Storage Node, you must ensure there is adequate space for that node's portion of all object metadata.

- At a minimum, you must assign at least 4 TB to volume 0.



If you use only one storage volume for a Storage Node and you assign 4 TB or less to the volume, the Storage Node might enter the Storage Read-Only state on startup and store object metadata only.

- If you are installing a new StorageGRID 11.6 system and each Storage Node has 128 GB or more of RAM, you should assign 8 TB or more to volume 0. Using a larger value for volume 0 can increase the space allowed for metadata on each Storage Node.
- When configuring different Storage Nodes for a site, use the same setting for volume 0 if possible. If a site contains Storage Nodes of different sizes, the Storage Node with the smallest volume 0 will determine the metadata capacity of that site.

For details, go to [Manage object metadata storage](#).

Related information

[Recover and maintain](#)

Deploy virtual machine grid nodes (VMware)

Collect information about your deployment environment

Before deploying grid nodes, you must collect information about your network configuration and VMware environment.



It is more efficient to perform one single installation of all the nodes, rather than installing some nodes now and some nodes later.

VMware information

You must access the deployment environment and collect information about the VMware environment; the networks that were created for the Grid, Admin, and Client Networks; and the storage volume types you plan to use for Storage Nodes.

You must collect information about your VMware environment, including the following:

- The username and password for a VMware vSphere account that has appropriate permissions to complete the deployment.
- Host, datastore, and network configuration information for each StorageGRID grid node virtual machine.



VMware live vMotion causes the virtual machine clock time to jump and is not supported for grid nodes of any type. Though rare, incorrect clock times can result in loss of data or configuration updates.

Grid Network information

You must collect information about the VMware network created for the StorageGRID Grid Network (required), including:

- The network name.
- If you are not using DHCP, the required networking details for each grid node (IP address, gateway, and network mask).
- If you are not using DHCP, the IP address of the primary Admin Node on the Grid Network. See “How grid nodes discover the primary Admin Node” for more information.

Admin Network information

For nodes that will be connected to the optional StorageGRID Admin Network, you must collect information about the VMware network created for this network, including:

- The network name.
- The method used to assign IP addresses, either static or DHCP.
- If you are using static IP addresses, the required networking details for each grid node (IP address, gateway, network mask).
- The external subnet list (ESL) for the Admin Network.

Client Network information

For nodes that will be connected to the optional StorageGRID Client Network, you must collect information about the VMware network created for this network, including:

- The network name.
- The method used to assign IP addresses, either static or DHCP.
- If you are using static IP addresses, the required networking details for each grid node (IP address, gateway, network mask).

Information about additional interfaces

You can optionally add trunk or access interfaces to the VM in vCenter after you install the node. For example, you might want to add a trunk interface to an Admin or Gateway Node, so you can use VLAN interfaces to segregate the traffic belonging to different applications or tenants. Or, you might want to add an access interface to use in a high availability (HA) group.

The interfaces you add are displayed on the VLAN interfaces page and on the HA groups page in the Grid Manager.

- If you add a trunk interface, configure one or more VLAN interfaces for each new parent interface. See [configure VLAN interfaces](#).
- If you add an access interface, you must add it directly to HA groups. See [configure high availability groups](#).

Storage volumes for virtual Storage Nodes

You must collect the following information for virtual machine-based Storage Nodes:

- The number and size of storage volumes (storage LUNs) you plan to add. See “Storage and performance requirements.”

Grid configuration information

You must collect information to configure your grid:

- Grid license
- Network Time Protocol (NTP) server IP addresses
- Domain Name System (DNS) server IP addresses

Related information

[How grid nodes discover the primary Admin Node](#)

[Storage and performance requirements](#)

How grid nodes discover the primary Admin Node

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:

- You can use the ADMIN_IP parameter to enter the primary Admin Node’s IP address manually.
- 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.

Automatic discovery of the primary Admin Node is accomplished using a multicast Domain Name System (mDNS). When the primary Admin Node first starts up, it publishes its IP address using mDNS. Other nodes on the same subnet can then query for the IP address and acquire it automatically. However, because multicast IP

traffic is not normally routable across subnets, nodes on other subnets cannot acquire the primary Admin Node's IP address directly.

If you use automatic discovery:



- You must include the ADMIN_IP setting for at least one grid node on any subnets that the primary Admin Node is not directly attached to. This grid node will then publish the primary Admin Node's IP address for other nodes on the subnet to discover with mDNS.
- Ensure that your network infrastructure supports passing multi-cast IP traffic within a subnet.

Deploy a StorageGRID node as a virtual machine

You use VMware vSphere Web Client to deploy each grid node as a virtual machine. During deployment, each grid node is created and connected to one or more StorageGRID networks.

If you need to deploy any StorageGRID appliance Storage Nodes, see the installation and maintenance instructions for the appliance.

Optionally, you can remap node ports or increase CPU or memory settings for the node before powering it on.

What you'll need

- You have reviewed how to [plan and prepare for installation](#), and you understand the requirements for software, CPU and RAM, and storage and performance.
- You are familiar with VMware vSphere Hypervisor and have experience deploying virtual machines in this environment.



The open-vm-tools package, an open-source implementation similar to VMware Tools, is included with the StorageGRID virtual machine. You do not need to install VMware Tools manually.

- You have downloaded and extracted the correct version of the StorageGRID installation archive for VMware.



If you are deploying the new node as part of an expansion or recovery operation, you must use the version of StorageGRID that is currently running on the grid.

- You have the StorageGRID Virtual Machine Disk (.vmdk) file:

NetApp-SG-version-SHA.vmdk

- You have the .ovf and .mf files for each type of grid node you are deploying:

Filename	Description
vsphere-primary-admin.ovf	The template file and manifest file for the primary Admin Node.
vsphere-primary-admin.mf	

Filename	Description
vsphere-non-primary-admin.ovf	The template file and manifest file for a non-primary Admin Node.
vsphere-non-primary-admin.mf	
vsphere-archive.ovf	The template file and manifest file for an Archive Node.
vsphere-archive.mf	
vsphere-gateway.ovf	The template file and manifest file for a Gateway Node.
vsphere-gateway.mf	
vsphere-storage.ovf	The template file and manifest file for a Storage Node.
vsphere-storage.mf	

- The .vdmk, .ovf, and .mf files are all in the same directory.
- You have a plan to minimize failure domains. For example, you should not deploy all Gateway Nodes on a single virtual machine server.



In a production deployment, do not run more than one Storage Node on a single virtual machine server. Using a dedicated virtual machine host for each Storage Node provides an isolated failure domain.

- If you are deploying a node as part of an expansion or recovery operation, you have the [instructions for expanding a StorageGRID system](#) or the [recovery and maintenance instructions](#).
- If you are deploying a StorageGRID node as a virtual machine with storage assigned from a NetApp ONTAP system, you have confirmed that the volume does not have a FabricPool tiering policy enabled. For example, if a StorageGRID node is running as a virtual machine on a VMWare host, ensure the volume backing the datastore for the node does not have a FabricPool tiering policy enabled. Disabling FabricPool tiering for volumes used with StorageGRID nodes simplifies troubleshooting and storage operations.



Never use FabricPool to tier any data related to StorageGRID back to StorageGRID itself. Tiering StorageGRID data back to StorageGRID increases troubleshooting and operational complexity.

About this task

Follow these instructions to initially deploy VMware nodes, add a new VMware node in an expansion, or replace a VMware node as part of a recovery operation. Except as noted in the steps, the node deployment procedure is the same for all node types, including Admin Nodes, Storage Nodes, Gateway Nodes, and Archive Nodes.

If you are installing a new StorageGRID system:

- You must deploy the primary Admin Node before you deploy any other grid node.
- You must ensure that each virtual machine can connect to the primary Admin Node over the Grid Network.
- You must deploy all grid nodes before configuring the grid.

If you are performing an expansion or recovery operation:

- You must ensure that the new virtual machine can connect to the primary Admin Node over the Grid Network.

If you need to remap any of the node's ports, do not power on the new node until the port remap configuration is complete.

Steps

1. Using VCenter, deploy an OVF template.

If you specify a URL, point to a folder containing the following files. Otherwise, select each of these files from a local directory.

```
NetApp-SG-version-SHA.vmdk  
vsphere-node.ovf  
vsphere-node.mf
```

For example, if this is the first node you are deploying, use these files to deploy the primary Admin Node for your StorageGRID system:

```
NetApp-SG-version-SHA.vmdk  
sphere-primary-admin.ovf  
sphere-primary-admin.mf
```

2. Provide a name for the virtual machine.

The standard practice is to use the same name for both the virtual machine and the grid node.

3. Place the virtual machine in the appropriate vApp or resource pool.
4. If you are deploying the primary Admin Node, read and accept the End User License Agreement.

Depending on your version of vCenter, the order of the steps will vary for accepting the End User License Agreement, specifying the name of the virtual machine, and selecting a datastore.

5. Select storage for the virtual machine.

If you are deploying a node as part of recovery operation, perform the instructions in the [storage recovery step](#) to add new virtual disks, reattach virtual hard disks from the failed grid node, or both.

When deploying a Storage Node, use 3 or more storage volumes, with each storage volume being 4 TB or larger. You must assign at least 4 TB to volume 0.



The Storage Node .ovf file defines several VMDKs for storage. Unless these VMDKs meet your storage requirements, you should remove them and assign appropriate VMDKs or RDMs for storage before powering up the node. VMDKs are more commonly used in VMware environments and are easier to manage, while RDMs might provide better performance for workloads that use larger object sizes (for example, greater than 100 MB).



Some StorageGRID installations might use larger, more active storage volumes than typical virtualized workloads. You might need to tune some hypervisor parameters, such as MaxAddressableSpaceTB, to achieve optimal performance. If you encounter poor performance, contact your virtualization support resource to determine whether your environment could benefit from workload-specific configuration tuning.

6. Select networks.

Determine which StorageGRID networks the node will use by selecting a destination network for each source network.

- The Grid Network is required. You must select a destination network in the vSphere environment.
- If you use the Admin Network, select a different destination network in the vSphere environment. If you do not use the Admin Network, select the same destination you selected for the Grid Network.
- If you use the Client Network, select a different destination network in the vSphere environment. If you do not use the Client Network, select the same destination you selected for the Grid Network.

7. Under **Customize Template**, configure the required StorageGRID node properties.

a. Enter the **Node name**.



If you are recovering a grid node, you must enter the name of the node you are recovering.

b. In the **Grid Network (eth0)** section, select STATIC or DHCP for the **Grid network IP configuration**.

- If you select STATIC, enter the **Grid network IP**, **Grid network mask**, **Grid network gateway**, and **Grid network MTU**.
- If you select DHCP, the **Grid network IP**, **Grid network mask**, and **Grid network gateway** are automatically assigned.

c. In the **Primary Admin IP** field, enter the IP address of the primary Admin Node for the Grid Network.



This step does not apply if the node you are deploying is the primary Admin Node.

If you omit the primary Admin Node IP address, the IP address will be automatically discovered if the primary Admin Node, or at least one other grid node with ADMIN_IP configured, is present on the same subnet. However, it is recommended to set the primary Admin Node IP address here.

d. In the **Admin Network (eth1)** section, select STATIC, DHCP, or DISABLED for the **Admin network IP configuration**.

- If you do not want to use the Admin Network, select DISABLED and enter **0.0.0.0** for the Admin Network IP. You can leave the other fields blank.
- If you select STATIC, enter the **Admin network IP**, **Admin network mask**, **Admin network gateway**, and **Admin network MTU**.
- If you select STATIC, enter the **Admin network external subnet list**. You must also configure a gateway.
- If you select DHCP, the **Admin network IP**, **Admin network mask**, and **Admin network gateway** are automatically assigned.

e. In the **Client Network (eth2)** section, select STATIC, DHCP, or DISABLED for the **Client network IP configuration**.

- If you do not want to use the Client Network, select DISABLED and enter **0.0.0.0** for the Client network IP. You can leave the other fields blank.
- If you select STATIC, enter the **Client network IP**, **Client network mask**, **Client network gateway**, and **Client network MTU**.
- If you select DHCP, the **Client network IP**, **Client network mask**, and **Client network gateway** are automatically assigned.

8. Review the virtual machine configuration and make any changes necessary.

9. When you are ready to complete, select **Finish** to start the upload of the virtual machine.

10. If you deployed this node as part of recovery operation and this is not a full-node recovery, perform these steps after deployment is complete:

- a. Right-click the virtual machine, and select **Edit Settings**.
- b. Select each default virtual hard disk that has been designated for storage, and select **Remove**.
- c. Depending on your data recovery circumstances, add new virtual disks according to your storage requirements, reattach any virtual hard disks preserved from the previously removed failed grid node, or both.

Note the following important guidelines:

- If you are adding new disks you should use the same type of storage device that was in use before node recovery.
- The Storage Node .ovf file defines several VMDKs for storage. Unless these VMDKs meet your storage requirements, you should remove them and assign appropriate VMDKs or RDMs for storage before powering up the node. VMDKs are more commonly used in VMware environments and are easier to manage, while RDMs may provide better performance for workloads that use larger object sizes (for example, greater than 100 MB).

11. If you need to remap the ports used by this node, follow these steps.

You might need to remap a port if your enterprise networking policies restrict access to one or more ports that are used by StorageGRID. See the [networking guidelines](#) for the ports used by StorageGRID.



Do not remap the ports used in load balancer endpoints.

- a. Select the new VM.
- b. From the Configure tab, select **Settings > vApp Options**. The location of **vApp Options** depends on the version of vCenter.
- c. In the **Properties** table, locate **PORT_REMAP_INBOUND** and **PORT_REMAP**.
- d. To symmetrically map both inbound and outbound communications for a port, select **PORT_REMAP**.



If only **PORT_REMAP** is set, the mapping that you specify applies to both inbound and outbound communications. If **PORT_REMAP_INBOUND** is also specified, **PORT_REMAP** applies only to outbound communications.

- i. Scroll back to the top of the table, and select **Edit**.
- ii. On the Type tab, select **User configurable**, and select **Save**.
- iii. Select **Set Value**.
- iv. Enter the port mapping:

```
<network type>/<protocol>/<default port used by grid node>/<new port>
```

<network type> is grid, admin, or client, and <protocol> is tcp or udp.

For example, to remap ssh traffic from port 22 to port 3022, enter:

```
client/tcp/22/3022
```

v. Select **OK**.

e. To specify the port used for inbound communications to the node, select **PORT_REMAP_INBOUND**.



If you specify PORT_REMAP_INBOUND and do not specify a value for PORT_REMAP, outbound communications for the port are unchanged.

i. Scroll back to the top of the table, and select **Edit**.

ii. On the Type tab, select **User configurable**, and select **Save**.

iii. Select **Set Value**.

iv. Enter the port mapping:

```
<network type>/<protocol>/<remapped inbound port>/<default inbound port used by grid node>
```

<network type> is grid, admin, or client, and <protocol> is tcp or udp.

For example, to remap inbound SSH traffic that is sent to port 3022 so that it is received at port 22 by the grid node, enter the following:

```
client/tcp/3022/22
```

v. Select **OK**.

12. If you want to increase the CPU or memory for the node from the default settings:

a. Right-click the virtual machine, and select **Edit Settings**.

b. Change the number of CPUs or the amount of memory as required.

Set the **Memory Reservation** to the same size as the **Memory** allocated to the virtual machine.

c. Select **OK**.

13. Power on the virtual machine.

After you finish

If you deployed this node as part of an expansion or recovery procedure, return to those instructions to complete the procedure.

Configure the grid and complete installation (VMware)

Navigate to the Grid Manager

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

What you'll need

The primary Admin Node must be deployed and have completed the initial startup sequence.

Steps

1. Open your web browser and navigate to one of the following addresses:

`https://primary_admin_node_ip`

`client_network_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.

2. Click **Install** **Install a StorageGRID system**.

The page used to configure a StorageGRID grid appears.

NetApp® StorageGRID®

Help ▾

Install

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

License

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

Grid Name

License File

Specify the StorageGRID license information

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

Steps

1. 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.

2. 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.



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

NetApp® StorageGRID®

Help ▾

Install

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

License

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

Grid Name	Grid1
New License File	Browse
License Serial Number	950719
Storage Capacity (TB)	240

3. Click **Next**.

Add sites

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.

Steps

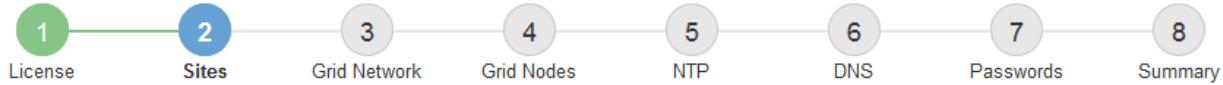
1. On the Sites page, enter the **Site Name**.
2. To add additional sites, click the plus sign next to the last site entry and enter the name in the new **Site Name** text box.

Add as many additional sites as required for your grid topology. You can add up to 16 sites.

NetApp® StorageGRID®

Help ▾

Install



Sites

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

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

Site Name 1	<input type="text" value="Raleigh"/>	
Site Name 2	<input type="text" value="Atlanta"/>	

3. Click **Next**.

Specify Grid Network subnets

You must specify the subnets that are used on the Grid Network.

About this task

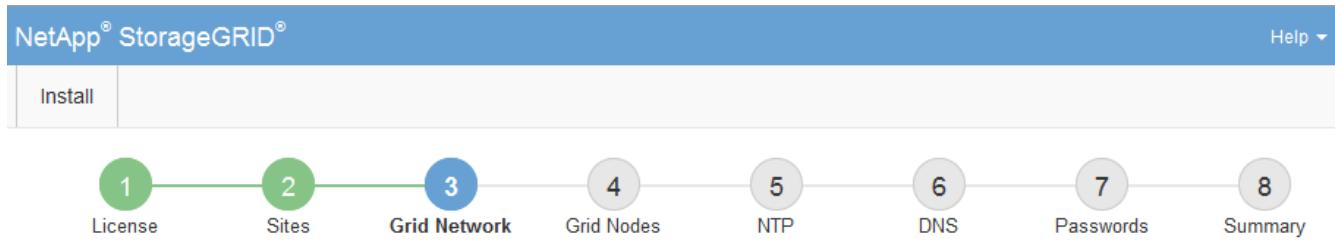
The subnet entries include the subnets for the Grid Network for each site in your StorageGRID system, along with any subnets that need to be reachable via the Grid Network.

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

Steps

1. Specify the CIDR network address for at least one Grid Network in the **Subnet 1** text box.
2. Click the plus sign next to the last entry to add an additional network entry.

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



Grid Network

You must specify the subnets that are used on the Grid Network. These entries typically include the subnets for the Grid Network for each site in your StorageGRID system. Select Discover Grid Networks to automatically add subnets based on the network configuration of all registered nodes.

Note: You must manually add any subnets for NTP, DNS, LDAP, or other external servers accessed through the Grid Network gateway.

Subnet 1	172.16.0.0/21	+
Discover Grid Network subnets		

3. Click **Next**.

Approve pending grid nodes

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

What you'll need

You have deployed all virtual and StorageGRID appliance grid nodes.



It is more efficient to perform one single installation of all the nodes, rather than installing some nodes now and some nodes later.

Steps

1. Review the Pending Nodes list, and confirm that it shows all of the grid nodes you deployed.
 - If a grid node is missing, confirm that it was deployed successfully.
2. Select the radio button next to a pending node you want to approve.



Grid Nodes

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

Pending Nodes

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

						Search
	Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address	
<input checked="" type="radio"/>	50:6b:4b:42:d7:00	NetApp-SGA	Storage Node	StorageGRID Appliance	172.16.5.20/21	

Approved Nodes

Grid nodes that have been approved and have been configured for installation. An approved grid node's configuration can be edited if errors are identified.

						Search
	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address
<input checked="" type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21
<input checked="" type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21
<input checked="" type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21
<input checked="" type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21
<input checked="" type="radio"/>	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21

3. Click **Approve**.

4. In General Settings, modify settings for the following properties, as necessary:

Storage Node Configuration

General Settings

Site	<input type="text" value="Raleigh"/>
Name	<input type="text" value="NetApp-SGA"/>
NTP Role	<input type="text" value="Automatic"/>
ADC Service	<input type="text" value="Automatic"/>

Grid Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="172.16.5.20/21"/>
Gateway	<input type="text" value="172.16.5.20"/>

Admin Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="10.224.5.20/21"/>
Gateway	<input type="text" value="10.224.0.1"/>
Subnets (CIDR)	<input type="text" value="10.0.0.0/8"/> ×
	<input type="text" value="172.19.0.0/16"/> ×
	<input type="text" value="172.21.0.0/16"/> + ×

Client Network

Configuration	STATIC
IPv4 Address (CIDR)	<input type="text" value="47.47.5.20/21"/>
Gateway	<input type="text" value="47.47.0.1"/>

- **Site:** The name of the site with which this grid node will be associated.
- **Name:** The name that will be assigned to the node, and the name that will be displayed in the Grid Manager. The name defaults to the name you specified when you configured the node. During this step of the installation process, you can change the name as required.



After you complete the installation, you cannot change the name of the node.



For a VMware node, you can change the name here, but this action will not change the name of the virtual machine in vSphere.

- **NTP Role:** The Network Time Protocol (NTP) role of the grid node. The options are **Automatic**, **Primary**, and **Client**. Selecting **Automatic** assigns the Primary role to Admin Nodes, Storage Nodes with ADC services, Gateway Nodes, and any grid nodes that have non-static IP addresses. All other grid nodes are assigned the Client role.



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

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

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

- **IPv4 Address (CIDR):** The CIDR network address for the Grid Network interface (eth0 inside the container). For example: 192.168.1.1234/21
- **Gateway:** The Grid Network gateway. For example: 192.168.0.1



The gateway is required if there are multiple grid subnets.



If you selected DHCP for the Grid Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the resulting IP address is not within a DHCP address pool.

6. If you want to configure the Admin Network for the grid node, add or update the settings in the Admin Network section as necessary.

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



If you selected DHCP for the Admin Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the resulting IP address is not within a DHCP address pool.

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

- a. Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.
Rebooting can take several minutes.
- b. Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- c. Select **Configure Networking > IP Configuration** and configure the enabled networks.
- d. Return to the Home page and click **Start Installation**.
- e. In the Grid Manager: If the node is listed in the Approved Nodes table, reset the node.
- f. Remove the node from the Pending Nodes table.
- g. Wait for the node to reappear in the Pending Nodes list.

- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page.

For additional information, see the installation and maintenance instructions for your appliance model.

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



If you selected DHCP for the Client Network configuration and you change the value here, the new value will be configured as a static address on the node. You must make sure the resulting IP address is not within a DHCP address pool.

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

- a. Reboot the appliance: In the Appliance Installer, select **Advanced > Reboot**.

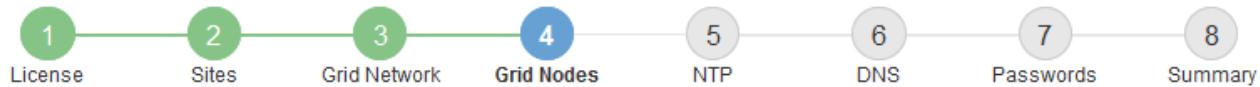
Rebooting can take several minutes.

- b. Select **Configure Networking > Link Configuration** and enable the appropriate networks.
- c. Select **Configure Networking > IP Configuration** and configure the enabled networks.
- d. Return to the Home page and click **Start Installation**.
- e. In the Grid Manager: If the node is listed in the Approved Nodes table, reset the node.
- f. Remove the node from the Pending Nodes table.
- g. Wait for the node to reappear in the Pending Nodes list.
- h. Confirm that you can configure the appropriate networks. They should already be populated with the information you provided on the IP Configuration page.

For additional information, see the installation and maintenance instructions for your appliance.

- 8. Click **Save**.

The grid node entry moves to the Approved Nodes list.



Grid Nodes

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

Pending Nodes

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

							Search	🔍
	Grid Network MAC Address	Name	Type	Platform	Grid Network IPv4 Address			
No results found.								

Approved Nodes

Grid nodes that have been approved and have been configured for installation. An approved grid node's configuration can be edited if errors are identified.

							Search	🔍
	Grid Network MAC Address	Name	Site	Type	Platform	Grid Network IPv4 Address		
<input type="radio"/>	00:50:56:87:42:ff	dc1-adm1	Raleigh	Admin Node	VMware VM	172.16.4.210/21		
<input type="radio"/>	00:50:56:87:c0:16	dc1-s1	Raleigh	Storage Node	VMware VM	172.16.4.211/21		
<input type="radio"/>	00:50:56:87:79:ee	dc1-s2	Raleigh	Storage Node	VMware VM	172.16.4.212/21		
<input type="radio"/>	00:50:56:87:db:9c	dc1-s3	Raleigh	Storage Node	VMware VM	172.16.4.213/21		
<input type="radio"/>	00:50:56:87:62:38	dc1-g1	Raleigh	API Gateway Node	VMware VM	172.16.4.214/21		
<input type="radio"/>	50:6b:4b:42:d7:00	NetApp-SGA	Raleigh	Storage Node	StorageGRID Appliance	172.16.5.20/21		

9. Repeat these steps for each pending grid node you want to approve.

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

10. When you are done approving grid nodes, click **Next**.

Specify Network Time Protocol server information

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.

About this task

You must specify IPv4 addresses for the NTP servers.

You must specify external NTP servers. The specified NTP servers must use the NTP protocol.

You must specify four NTP server references of Stratum 3 or better to prevent issues with time drift.



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

[Support boundary to configure the Windows Time service for high-accuracy environments](#)

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



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

Perform additional checks for VMware, such as ensuring that the hypervisor uses the same NTP source as the virtual machine, and using VMTools to disable the time sync between the hypervisor and StorageGRID virtual machines.

Steps

1. Specify the IPv4 addresses for at least four NTP servers in the **Server 1 to Server 4** text boxes.
2. If necessary, select the plus sign next to the last entry to add additional server entries.

NetApp® StorageGRID®

Help ▾

Install

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

Network Time Protocol

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

Server 1	10.60.248.183
Server 2	10.227.204.142
Server 3	10.235.48.111
Server 4	0.0.0.0

+

3. Select **Next**.

Specify Domain Name System server information

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.

About this task

Specifying DNS server information allows you to use Fully Qualified Domain Name (FQDN) hostnames rather than IP addresses for email notifications and AutoSupport. Specifying at least two DNS servers is recommended.

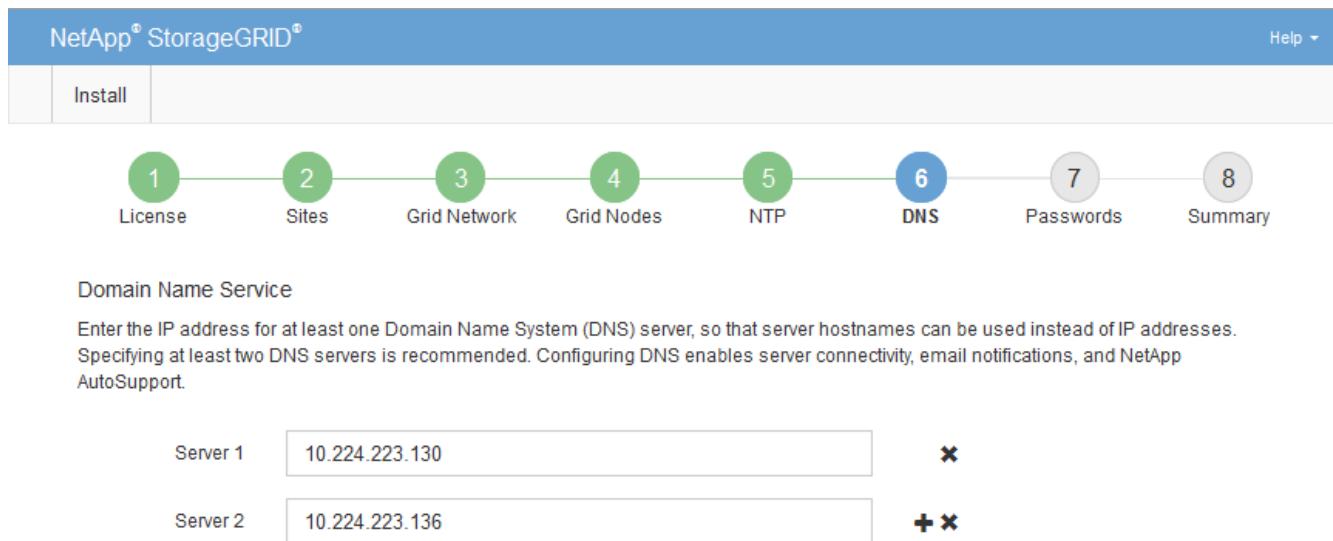


Provide two to six IPv4 addresses for DNS servers. You should select DNS servers that each site can access locally in the event of network islanding. This is to ensure an islanded site continues to have access to the DNS service. After configuring the grid-wide DNS server list, you can further customize the DNS server list for each node. For details, see the information about modifying the DNS configuration in the recovery and maintenance instructions.

If the DNS server information is omitted or incorrectly configured, a DNST alarm is triggered on each grid node's SSM service. The alarm clears when DNS is configured correctly and the new server information has reached all grid nodes.

Steps

1. Specify the IPv4 address for at least one DNS server in the **Server 1** text box.
2. If necessary, select the plus sign next to the last entry to add additional server entries.



The best practice is to specify at least two DNS servers. You can specify up to six DNS servers.

3. Select **Next**.

Related information

[Recover and maintain](#)

Specify the StorageGRID system passwords

As part of installing your StorageGRID system, you need to enter the passwords to use to secure your system and perform maintenance tasks.

About this task

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.

Steps

1. In **Provisioning Passphrase**, enter the provisioning passphrase that will be required to make changes to the grid topology of your StorageGRID system.

Store the provisioning passphrase in a secure place.



If after the installation completes and you want to change the provisioning passphrase later, you can use the Grid Manager. Select **CONFIGURATION > Access control > Grid passwords**.

2. In **Confirm Provisioning Passphrase**, reenter the provisioning passphrase to confirm it.
3. In **Grid Management Root User Password**, enter the password to use to access the Grid Manager as the “root” user.

Store the password in a secure place.

4. In **Confirm Root User Password**, reenter the Grid Manager password to confirm it.

NetApp® StorageGRID®

Help ▾

Install

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

Passwords

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

Provisioning Passphrase
Confirm Provisioning Passphrase
Grid Management Root User Password
Confirm Root User Password

Create random command line passwords.

5. If you are installing a grid for proof of concept or demo purposes, optionally deselect the **Create random command line passwords** check box.

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



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

6. Click **Next**.

Review your configuration and complete installation

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

Steps

1. View the **Summary** page.

NetApp® StorageGRID®

Help ▾

Install

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

Summary

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

General Settings

Grid Name	Grid1	Modify License
Passwords	Auto-generated random command line passwords	Modify Passwords

Networking

NTP	10.60.248.183 10.227.204.142 10.235.48.111	Modify NTP
DNS	10.224.223.130 10.224.223.136	Modify DNS
Grid Network	172.16.0.0/21	Modify Grid Network

Topology

Topology	Atlanta	Modify Sites	Modify Grid Nodes				
	Raleigh	dc1-adm1	dc1-g1	dc1-s1	dc1-s2	dc1-s3	NetApp-SGA

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

3. Click **Install**.



If a node is configured to use the Client Network, the default gateway for that node switches from the Grid Network to the Client Network when you click **Install**. If you lose connectivity, you must ensure that you are accessing the primary Admin Node through an accessible subnet. See [Networking guidelines](#) for details.

4. Click **Download Recovery Package**.

When the installation progresses to the point where the grid topology is defined, you are prompted to download the Recovery Package file (.zip), and confirm that you can 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.

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



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

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

Download Recovery Package

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

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

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

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

Download Recovery Package

I have successfully downloaded and verified the Recovery Package file.

If the installation is still in progress, the status page appears. This page indicates the progress of the installation for each grid node.

Installation Status						
If necessary, you may Download the Recovery Package file again.						
Name	Site	Grid Network IPv4 Address	Progress	Stage	Search	
dc1-adm1	Site1	172.16.4.215/21	<div style="width: 80%;"></div>	Starting services		
dc1-g1	Site1	172.16.4.216/21	<div style="width: 100%; background-color: green;"></div>	Complete		
dc1-s1	Site1	172.16.4.217/21	<div style="width: 70%;"></div>	Waiting for Dynamic IP Service peers		
dc1-s2	Site1	172.16.4.218/21	<div style="width: 10%;"></div>	Downloading hotfix from primary Admin if needed		
dc1-s3	Site1	172.16.4.219/21	<div style="width: 10%;"></div>	Downloading hotfix from primary Admin if needed		

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

7. Sign in to the Grid Manager using the “root” user and the password you specified during the installation.

Post-installation guidelines

After completing grid node deployment and configuration, follow these guidelines for DHCP addressing and network configuration changes.

- If DHCP was used to assign IP addresses, configure a DHCP reservation for each IP address on the networks being used.

You can only set up DHCP during the deployment phase. You cannot set up DHCP during configuration.



Nodes reboot when their IP addresses change, which can cause outages if a DHCP address change affects multiple nodes at the same time.

- You must use the Change IP procedures if you want to change IP addresses, subnet masks, and default gateways for a grid node. See [Configure IP addresses](#).
- If you make networking configuration changes, including routing and gateway changes, client connectivity to the primary Admin Node and other grid nodes might be lost. Depending on the networking changes applied, you might need to re-establish these connections.

Automate the installation (VMware)

You can use VMware vSphere to automate the deployment of grid nodes. You can also automate the configuration of StorageGRID.

Automate grid node deployment

Use VMware vSphere to automate the deployment of grid nodes.

What you'll need

- You have access to a Linux/Unix system with Bash 3.2 or later.
- You have VMware OVF Tool 4.1 installed and correctly configured.
- You know the username and password required to access VMware vSphere using the OVF Tool.
- You know the virtual infrastructure (VI) URL for the location in vSphere where you want to deploy the StorageGRID virtual machines. This URL will typically be a vApp, or Resource Pool. For example: `vi://vcenter.example.com/vi/sgws`



You can use the VMware `ovftool` utility to determine this value (see the `ovftool` documentation for details).



If you are deploying to a vApp, the virtual machines will not start automatically the first time, and you must power them on manually.

- You have collected all the required information for the configuration file. See [Collect information about your deployment environment](#) for information.
- You have access to the following files from the VMware installation archive for StorageGRID:

Filename	Description
<code>NetApp-SG-version-SHA.vmdk</code>	The virtual machine disk file that is used as a template for creating grid node virtual machines. Note: This file must be in the same folder as the <code>.ovf</code> and <code>.mf</code> files.
<code>vsphere-primary-admin.ovf</code> <code>vsphere-primary-admin.mf</code>	The Open Virtualization Format template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying the primary Admin Node.
<code>vsphere-non-primary-admin.ovf</code> <code>vsphere-non-primary-admin.mf</code>	The template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying non-primary Admin Nodes.
<code>vsphere-archive.ovf</code> <code>vsphere-archive.mf</code>	The template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying Archive Nodes.
<code>vsphere-gateway.ovf</code> <code>vsphere-gateway.mf</code>	The template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying Gateway Nodes.
<code>vsphere-storage.ovf</code> <code>vsphere-storage.mf</code>	The template file (<code>.ovf</code>) and manifest file (<code>.mf</code>) for deploying virtual machine-based Storage Nodes.
<code>deploy-vsphere-ovftool.sh</code>	The Bash shell script used to automate the deployment of virtual grid nodes.
<code>deploy-vsphere-ovftool-sample.ini</code>	The sample configuration file for use with the <code>deploy-vsphere-ovftool.sh</code> script.

Define the configuration file for your deployment

You specify the information needed to deploy virtual grid nodes for StorageGRID in a configuration file, which is used by the `deploy-vsphere-ovftool.sh` Bash script. You can modify a sample configuration file, so that you do not have to create the file from scratch.

Steps

1. Make a copy of the sample configuration file (`deploy-vsphere-ovftool.sample.ini`). Save the new

file as `deploy-vsphere-ovftool.ini` in the same directory as `deploy-vsphere-ovftool.sh`.

2. Open `deploy-vsphere-ovftool.ini`.
3. Enter all of the information required to deploy VMware virtual grid nodes.

See [Configuration file settings](#) for information.

4. When you have entered and verified all of the necessary information, save and close the file.

Configuration file settings

The `deploy-vsphere-ovftool.ini` configuration file contains the settings that are required to deploy virtual grid nodes.

The configuration file first lists global parameters, and then lists node-specific parameters in sections defined by node name. When the file is used:

- *Global parameters* are applied to all grid nodes.
- *Node-specific parameters* override global parameters.

Global parameters

Global parameters are applied to all grid nodes, unless they are overridden by settings in individual sections. Place the parameters that apply to multiple nodes in the global parameter section, and then override these settings as necessary in the sections for individual nodes.

- **OVFTOOL_ARGUMENTS:** You can specify OVFTOOL_ARGUMENTS as global settings, or you can apply arguments individually to specific nodes. For example:

```
OVFTOOL_ARGUMENTS = --powerOn --noSSLVerify --diskMode=eagerZeroedThick  
--datastore='datastore_name'
```

You can use the `--powerOffTarget` and `--overwrite` options to shut down and replace existing virtual machines.



You should deploy nodes to different datastores and specify OVFTOOL_ARGUMENTS for each node, instead of globally.

- **SOURCE:** The path to the StorageGRID virtual machine template (.vmdk) file and the .ovf and .mf files for individual grid nodes. This defaults to the current directory.

```
SOURCE = /downloads/StorageGRID-Webscale-version/vsphere
```

- **TARGET:** The VMware vSphere virtual infrastructure (vi) URL for the location where StorageGRID will be deployed. For example:

```
TARGET = vi://vcenter.example.com/vm/sgws
```

- **GRID_NETWORK_CONFIG**: The method used to acquire IP addresses, either STATIC or DHCP. The default is STATIC. If all or most of the nodes use the same method for acquiring IP addresses, you can specify that method here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_CONFIG = DHCP
```

- **GRID_NETWORK_TARGET**: The name of an existing VMware network to use for the Grid Network. If all or most of the nodes use the same network name, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_TARGET = SG-Admin-Network
```

- **GRID_NETWORK_MASK**: The network mask for the Grid Network. If all or most of the nodes use the same network mask, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_MASK = 255.255.255.0
```

- **GRID_NETWORK_GATEWAY**: The network gateway for the Grid Network. If all or most of the nodes use the same network gateway, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
GRID_NETWORK_GATEWAY = 10.1.0.1
```

- **GRID_NETWORK_MTU**: Optional. The maximum transmission unit (MTU) on the Grid Network. If specified, the value must be between 1280 and 9216. For example:

```
GRID_NETWORK_MTU = 8192
```

If omitted, 1400 is used.

If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value.



The MTU value of the network must match the value configured on the switch port the node is connected to. Otherwise, network performance issues or packet loss might occur.



For the best network performance, all nodes should be configured with similar MTU values on their Grid Network interfaces. The **Grid Network MTU mismatch** alert is triggered if there is a significant difference in MTU settings for the Grid Network on individual nodes. The MTU values do not have to be the same for all network types.

- **ADMIN_NETWORK_CONFIG**: The method used to acquire IP addresses, either DISABLED, STATIC, or DHCP. The default is DISABLED. If all or most of the nodes use the same method for acquiring IP

addresses, you can specify that method here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_CONFIG = STATIC
```

- **ADMIN_NETWORK_TARGET**: The name of an existing VMware network to use for the Admin Network. This setting is required unless the Admin Network is disabled. If all or most of the nodes use the same network name, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_TARGET = SG-Admin-Network
```

- **ADMIN_NETWORK_MASK**: The network mask for the Admin Network. This setting is required if you are using static IP addressing. If all or most of the nodes use the same network mask, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_MASK = 255.255.255.0
```

- **ADMIN_NETWORK_GATEWAY**: The network gateway for the Admin Network. This setting is required if you are using static IP addressing and you specify external subnets in the **ADMIN_NETWORK_ESL** setting. (That is, it is not required if **ADMIN_NETWORK_ESL** is empty.) If all or most of the nodes use the same network gateway, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_GATEWAY = 10.3.0.1
```

- **ADMIN_NETWORK_ESL**: The external subnet list (routes) for the Admin Network, specified as a comma-separated list of CIDR route destinations. If all or most of the nodes use the same external subnet list, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_ESL = 172.16.0.0/21,172.17.0.0/21
```

- **ADMIN_NETWORK_MTU**: Optional. The maximum transmission unit (MTU) on the Admin Network. Do not specify if **ADMIN_NETWORK_CONFIG** = **DHCP**. If specified, the value must be between 1280 and 9216. If omitted, 1400 is used. If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value. If all or most of the nodes use the same MTU for the Admin Network, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
ADMIN_NETWORK_MTU = 8192
```

- **CLIENT_NETWORK_CONFIG**: The method used to acquire IP addresses, either **DISABLED**, **STATIC**, or

DHCP. The default is DISABLED. If all or most of the nodes use the same method for acquiring IP addresses, you can specify that method here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_CONFIG = STATIC
```

- **CLIENT_NETWORK_TARGET**: The name of an existing VMware network to use for the Client Network. This setting is required unless the Client Network is disabled. If all or most of the nodes use the same network name, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_TARGET = SG-Client-Network
```

- **CLIENT_NETWORK_MASK**: The network mask for the Client Network. This setting is required if you are using static IP addressing. If all or most of the nodes use the same network mask, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_MASK = 255.255.255.0
```

- **CLIENT_NETWORK_GATEWAY**: The network gateway for the Client Network. This setting is required if you are using static IP addressing. If all or most of the nodes use the same network gateway, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_GATEWAY = 10.4.0.1
```

- **CLIENT_NETWORK_MTU**: Optional. The maximum transmission unit (MTU) on the Client Network. Do not specify if CLIENT_NETWORK_CONFIG = DHCP. If specified, the value must be between 1280 and 9216. If omitted, 1400 is used. If you want to use jumbo frames, set the MTU to a value suitable for jumbo frames, such as 9000. Otherwise, keep the default value. If all or most of the nodes use the same MTU for the Client Network, you can specify it here. You can then override the global setting by specifying different settings for one or more individual nodes. For example:

```
CLIENT_NETWORK_MTU = 8192
```

- **PORT_REMAP**: Remaps any port used by a node for internal grid node communications or external communications. Remapping ports is necessary if enterprise networking policies restrict one or more ports used by StorageGRID. For the list of ports used by StorageGRID, see internal grid node communications and external communications in [Networking guidelines](#).



Do not remap the ports you are planning to use to configure load balancer endpoints.



If only PORT_REMAP is set, the mapping that you specify is used for both inbound and outbound communications. If PORT_REMAP_INBOUND is also specified, PORT_REMAP applies only to outbound communications.

The format used is: *network type/protocol/default port used by grid node/new port*, where network type is grid, admin, or client, and protocol is tcp or udp.

For example:

```
PORT_REMAP = client/tcp/18082/443
```

If used alone, this example setting symmetrically maps both inbound and outbound communications for the grid node from port 18082 to port 443. If used in conjunction with PORT_REMAP_INBOUND, this example setting maps outbound communications from port 18082 to port 443.

- **PORT_REMAP_INBOUND:** Remaps inbound communications for the specified port. If you specify PORT_REMAP_INBOUND but do not specify a value for PORT_REMAP, outbound communications for the port are unchanged.



Do not remap the ports you are planning to use to configure load balancer endpoints.

The format used is: *network type/protocol/_default port used by grid node/new port*, where network type is grid, admin, or client, and protocol is tcp or udp.

For example:

```
PORT_REMAP_INBOUND = client/tcp/443/18082
```

This example takes traffic that is sent to port 443 to pass an internal firewall and directs it to port 18082, where the grid node is listening for S3 requests.

Node-specific parameters

Each node is in its own section of the configuration file. Each node requires the following settings:

- The section head defines the node name that will be displayed in the Grid Manager. You can override that value by specifying the optional NODE_NAME parameter for the node.
- **NODE_TYPE:** VM_Admin_Node, VM_Storage_Node, VM_Archive_Node, or VM_API_Gateway_Node
- **GRID_NETWORK_IP:** The IP address for the node on the Grid Network.
- **ADMIN_NETWORK_IP:** The IP address for the node on the Admin Network. Required only if the node is attached to the Admin Network and ADMIN_NETWORK_CONFIG is set to STATIC.
- **CLIENT_NETWORK_IP:** The IP address for the node on the Client Network. Required only if the node is attached to the Client Network and CLIENT_NETWORK_CONFIG for this node is set to STATIC.
- **ADMIN_IP:** The IP address for the primary Admin node on the Grid Network. Use the value that you specify as the GRID_NETWORK_IP for the primary Admin Node. If you omit this parameter, the node attempts to discover the primary Admin Node IP using mDNS. For more information, see [How grid nodes discover the primary Admin Node](#).



The ADMIN_IP parameter is ignored for the primary Admin Node.

- Any parameters that were not set globally. For example, if a node is attached to the Admin Network and you did not specify ADMIN_NETWORK parameters globally, you must specify them for the node.

Primary Admin Node

The following additional settings are required for the primary Admin Node:

- **NODE_TYPE**: VM_Admin_Node
- **ADMIN_ROLE**: Primary

This example entry is for a primary Admin Node that is on all three networks:

```
[DC1-ADM1]
ADMIN_ROLE = Primary
NODE_TYPE = VM_Admin_Node

GRID_NETWORK_IP = 10.1.0.2
ADMIN_NETWORK_IP = 10.3.0.2
CLIENT_NETWORK_IP = 10.4.0.2
```

The following additional setting is optional for the primary Admin Node:

- **DISK**: By default, Admin Nodes are assigned two additional 200 GB hard disks for audit and database use. You can increase these settings using the DISK parameter. For example:

```
DISK = INSTANCES=2, CAPACITY=300
```



For Admin nodes, INSTANCES must always equal 2.

Storage Node

The following additional setting is required for Storage Nodes:

- **NODE_TYPE**: VM_Storage_Node

This example entry is for a Storage Node that is on the Grid and Admin Networks, but not on the Client Network. This node uses the ADMIN_IP setting to specify the primary Admin Node's IP address on the Grid Network.

```
[DC1-S1]
NODE_TYPE = VM_Storage_Node

GRID_NETWORK_IP = 10.1.0.3
ADMIN_NETWORK_IP = 10.3.0.3

ADMIN_IP = 10.1.0.2
```

This second example entry is for a Storage Node on a Client Network where the customer's enterprise networking policy states that an S3 client application is only permitted to access the Storage Node using either port 80 or 443. The example configuration file uses PORT_REMAP to enable the Storage Node to send and receive S3 messages on port 443.

```
[DC2-S1]
NODE_TYPE = VM_Storage_Node

GRID_NETWORK_IP = 10.1.1.3
CLIENT_NETWORK_IP = 10.4.1.3
PORT_REMAP = client/tcp/18082/443

ADMIN_IP = 10.1.0.2
```

The last example creates a symmetric remapping for ssh traffic from port 22 to port 3022, but explicitly sets the values for both inbound and outbound traffic.

```
[DC1-S3]
NODE_TYPE = VM_Storage_Node

GRID_NETWORK_IP = 10.1.1.3

PORT_REMAP = grid/tcp/22/3022
PORT_REMAP_INBOUND = grid/tcp/3022/22

ADMIN_IP = 10.1.0.2
```

The following additional setting is optional for Storage Nodes:

- **DISK:** By default, Storage Nodes are assigned three 4 TB disks for RangeDB use. You can increase these settings with the DISK parameter. For example:

```
DISK = INSTANCES=16, CAPACITY=4096
```

Archive Node

The following additional setting is required for Archive Nodes:

- **NODE_TYPE**: VM_Archive_Node

This example entry is for an Archive Node that is on the Grid and Admin Networks, but not on the Client Network.

```
[DC1-ARC1]
NODE_TYPE = VM_Archive_Node

GRID_NETWORK_IP = 10.1.0.4
ADMIN_NETWORK_IP = 10.3.0.4

ADMIN_IP = 10.1.0.2
```

Gateway Node

The following additional setting is required for Gateway Nodes:

- **NODE_TYPE**: VM_API_Gateway

This example entry is for an example Gateway Node on all three networks. In this example, no Client Network parameters were specified in the global section of the configuration file, so they must be specified for the node:

```
[DC1-G1]
NODE_TYPE = VM_API_Gateway

GRID_NETWORK_IP = 10.1.0.5
ADMIN_NETWORK_IP = 10.3.0.5

CLIENT_NETWORK_CONFIG = STATIC
CLIENT_NETWORK_TARGET = SG-Client-Network
CLIENT_NETWORK_MASK = 255.255.255.0
CLIENT_NETWORK_GATEWAY = 10.4.0.1
CLIENT_NETWORK_IP = 10.4.0.5

ADMIN_IP = 10.1.0.2
```

Non-primary Admin Node

The following additional settings are required for non-primary Admin Nodes:

- **NODE_TYPE**: VM_Admin_Node
- **ADMIN_ROLE**: Non-Primary

This example entry is for a non-primary Admin Node that is not on the Client Network:

```
[DC2-ADM1]
ADMIN_ROLE = Non-Primary
NODE_TYPE = VM_Admin_Node

GRID_NETWORK_TARGET = SG-Grid-Network
GRID_NETWORK_IP = 10.1.0.6
ADMIN_NETWORK_IP = 10.3.0.6

ADMIN_IP = 10.1.0.2
```

The following additional setting is optional for non-primary Admin Nodes:

- **DISK**: By default, Admin Nodes are assigned two additional 200 GB hard disks for audit and database use. You can increase these settings using the DISK parameter. For example:

```
DISK = INSTANCES=2, CAPACITY=300
```



For Admin nodes, INSTANCES must always equal 2.

Run the Bash script

You can use the `deploy-vsphere-ovftool.sh` Bash script and the `deploy-vsphere-ovftool.ini` configuration file you modified to automate the deployment of StorageGRID grid nodes in VMware vSphere.

What you'll need

- You have created a `deploy-vsphere-ovftool.ini` configuration file for your environment.

You can use the help available with the Bash script by entering the help commands (`-h`/`--help`). For example:

```
./deploy-vsphere-ovftool.sh -h
```

or

```
./deploy-vsphere-ovftool.sh --help
```

Steps

1. Log in to the Linux machine you are using to run the Bash script.
2. Change to the directory where you extracted the installation archive.

For example:

```
cd StorageGRID-Webscale-version/vsphere
```

3. To deploy all grid nodes, run the Bash script with the appropriate options for your environment.

For example:

```
./deploy-vsphere-ovftool.sh --username=user --password=pwd ./deploy-vsphere-ovftool.ini
```

4. If a grid node failed to deploy because of an error, resolve the error and rerun the Bash script for only that node.

For example:

```
./deploy-vsphere-ovftool.sh --username=user --password=pwd --single -node="DC1-S3" ./deploy-vsphere-ovftool.ini
```

The deployment is complete when the status for each node is “Passed.”

Deployment Summary

node	attempts	status
DC1-ADM1	1	Passed
DC1-G1	1	Passed
DC1-S1	1	Passed
DC1-S2	1	Passed
DC1-S3	1	Passed

Automate the configuration of StorageGRID

After deploying the grid nodes, you can automate the configuration of the StorageGRID system.

What you'll need

- You know the location of the following files from the installation archive.

Filename	Description
configure-storagegrid.py	Python script used to automate the configuration
configure-storagegrid.sample.json	Sample configuration file for use with the script

Filename	Description
configure-storagegrid.blank.json	Blank configuration file for use with the script

- You have created a `configure-storagegrid.json` configuration file. To create this file, you can modify the sample configuration file (`configure-storagegrid.sample.json`) or the blank configuration file (`configure-storagegrid.blank.json`).

You can use the `configure-storagegrid.py` Python script and the `configure-storagegrid.json` configuration file to automate the configuration of your StorageGRID system.



You can also configure the system using the Grid Manager or the Installation API.

Steps

1. Log in to the Linux machine you are using to run the Python script.
2. Change to the directory where you extracted the installation archive.

For example:

```
cd StorageGRID-Webscale-version/platform
```

where `platform` is `debs`, `rpm`s, or `vsphere`.

3. Run the Python script and use the configuration file you created.

For example:

```
./configure-storagegrid.py ./configure-storagegrid.json --start-install
```

Result

A Recovery Package `.zip` file is generated during the configuration process, and it is downloaded to the directory where you are running the installation and configuration process. You must back up the Recovery Package file so that you can recover the StorageGRID system if one or more grid nodes fails. For example, copy it to a secure, backed up network location and to a secure cloud storage location.



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

If you specified that random passwords should be generated, you need to extract the `Passwords.txt` file and look for the passwords required to access your StorageGRID system.

```
#####
##### The StorageGRID "recovery package" has been downloaded as: #####
#####           ./sgws-recovery-package-994078-rev1.zip           #####
#####   Safeguard this file as it will be needed in case of a   #####
#####           StorageGRID node recovery.                      #####
#####
```

Your StorageGRID system is installed and configured when a confirmation message is displayed.

StorageGRID has been configured and installed.

Related information

[Navigate to the Grid Manager](#)

[Overview of the installation REST API](#)

Overview of the installation REST API

StorageGRID provides the StorageGRID Installation API for performing installation tasks.

The API uses the Swagger open source API platform to provide the API documentation. Swagger allows both developers and non-developers to interact with the API in a user interface that illustrates how the API responds to parameters and options. This documentation assumes that you are familiar with standard web technologies and the JSON (JavaScript Object Notation) data format.



Any API operations you perform using the API Docs webpage are live operations. Be careful not to create, update, or delete configuration data or other data by mistake.

Each REST API command includes the API's URL, an HTTP action, any required or optional URL parameters, and an expected API response.

StorageGRID Installation API

The StorageGRID Installation API is only available when you are initially configuring your StorageGRID system, and in the event that you need to perform a primary Admin Node recovery. The Installation API can be accessed over HTTPS from the Grid Manager.

To access the API documentation, go to the installation web page on the primary Admin Node and select **Help > API Documentation** from the menu bar.

The StorageGRID Installation API includes the following sections:

- **config** — Operations related to the product release and versions of the API. You can list the product release version and the major versions of the API supported by that release.
- **grid** — Grid-level configuration operations. You can get and update grid settings, including grid details, Grid Network subnets, grid passwords, and NTP and DNS server IP addresses.
- **nodes** — Node-level configuration operations. You can retrieve a list of grid nodes, delete a grid node,

configure a grid node, view a grid node, and reset a grid node's configuration.

- **provision**— Provisioning operations. You can start the provisioning operation and view the status of the provisioning operation.
- **recovery**— Primary Admin Node recovery operations. You can reset information, upload the Recover Package, start the recovery, and view the status of the recovery operation.
- **recovery-package**— Operations to download the Recovery Package.
- **schemas**— API schemas for advanced deployments
- **sites**— Site-level configuration operations. You can create, view, delete, and modify a site.

Where to go next

After completing an installation, you must perform a series of integration and configuration steps. Some steps are required; others are optional.

Required tasks

- Configure VMware vSphere Hypervisor for automatic restart.

You must configure the hypervisor to restart the virtual machines when the server restarts. Without an automatic restart, the virtual machines and grid nodes remain shut down after the server restarts. For details, see the VMware vSphere Hypervisor documentation.

- Create a tenant account for each client protocol (Swift or S3) that will be used to store objects on your StorageGRID system.
- Control system access by configuring groups and user accounts. Optionally, you can configure a federated identity source (such as Active Directory or OpenLDAP), so you can import administration groups and users. Or, you can create local groups and users.
- Integrate and test the S3 or Swift API client applications you will use to upload objects to your StorageGRID system.
- When you are ready, configure the information lifecycle management (ILM) rules and ILM policy you want to use to protect object data.



When you install StorageGRID, the default ILM policy, Baseline 2 Copies Policy, is active. This policy includes the stock ILM rule (Make 2 Copies), and it applies if no other policy has been activated.

- If your installation includes appliance Storage Nodes, use SANtricity software to complete the following tasks:
 - Connect to each StorageGRID appliance.
 - Verify receipt of AutoSupport data.
- If your StorageGRID system includes any Archive Nodes, configure the Archive Node's connection to the target external archival storage system.



If any Archive Nodes will use Tivoli Storage Manager as the external archival storage system, you must also configure Tivoli Storage Manager.

- Review and follow the StorageGRID system hardening guidelines to eliminate security risks.

- Configure email notifications for system alerts.

Optional tasks

- If you want to receive notifications from the (legacy) alarm system, configure mailing lists and email notifications for alarms.
- Update grid node IP addresses if they have changed since you planned your deployment and generated the Recovery Package. See information about changing IP addresses in the recovery and maintenance instructions.
- Configure storage encryption, if required.
- Configure storage compression to reduce the size of stored objects, if required.
- Configure audit client access. You can configure access to the system for auditing purposes through an NFS or a CIFS file share. See the instructions for administering StorageGRID.



Audit export through CIFS/Samba has been deprecated and will be removed in a future StorageGRID release.

Related information

[Administer StorageGRID](#)

[Use S3](#)

[Use Swift](#)

[Manage objects with ILM](#)

[Monitor and troubleshoot](#)

[Recover and maintain](#)

[SG100 and SG1000 services appliances](#)

[SG5600 storage appliances](#)

[SG5700 storage appliances](#)

[SG6000 storage appliances](#)

[Release notes](#)

[System hardening](#)

[Review audit logs](#)

[Upgrade software](#)

Troubleshoot installation issues

If any problems occur while installing your StorageGRID system, you can access the installation log files.

The following are the main installation log files, which technical support might need to resolve issues.

- `/var/local/log/install.log` (found on all grid nodes)
- `/var/local/log/gdu-server.log` (found on the primary Admin Node)

To learn how to access the log files, see the [instructions for monitoring and troubleshooting StorageGRID](#). For help troubleshooting appliance installation issues, see the installation and maintenance instructions for your appliances. If you need additional help, contact technical support.

[SG100 and SG1000 services appliances](#)

[SG6000 storage appliances](#)

[SG5700 storage appliances](#)

[SG5600 storage appliances](#)

[NetApp Support](#)

Virtual machine resource reservation requires adjustment

OVF files include a resource reservation designed to ensure that each grid node has sufficient RAM and CPU to operate efficiently. If you create virtual machines by deploying these OVF files on VMware and the predefined number of resources are not available, the virtual machines will not start.

About this task

If you are certain that the VM host has sufficient resources for each grid node, manually adjust the resources allocated for each virtual machine, and then try starting the virtual machines.

Steps

1. In the VMware vSphere Hypervisor client tree, select the virtual machine that is not started.
2. Right-click the virtual machine, and select **Edit Settings**.
3. From the Virtual Machines Properties window, select the **Resources** tab.
4. Adjust the resources allocated to the virtual machine:
 - a. Select **CPU**, and then use the Reservation slider to adjust the MHz reserved for this virtual machine.
 - b. Select **Memory**, and then use the Reservation slider to adjust the MB reserved for this virtual machine.
5. Click **OK**.
6. Repeat as required for other virtual machines hosted on the same VM host.

Copyright information

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

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

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

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

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

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

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

Trademark information

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