



StorageGRID® 11.2

Installation Guide

For VMware® Deployments

August 2019 | 215-13581_2019-08_en-us
doccomments@netapp.com

 **NetApp®**

Contents

Installation overview	5
Planning and preparation	6
Required materials	6
Downloading and extracting the StorageGRID installation files	7
Hardware requirements	9
Software requirements	9
Virtual machine requirements	10
Networking requirements	10
Network model	10
Networking guidelines	12
Network installation and provisioning	14
Internal grid node communications	15
External communications	17
Networking and ports for platform services and Cloud Storage Pools	20
Storage and performance requirements	21
Web browser requirements	22
Deploying grid nodes	23
Deploying virtual machine grid nodes in VMware vSphere Web Client	23
Collecting information about your deployment environment	23
Deploying the primary Admin Node virtual machine	25
Deploying additional grid node virtual machines	28
Deploying appliance Storage Nodes	31
Starting StorageGRID appliance installation	31
Monitoring StorageGRID appliance installation	34
Configuring the grid and completing installation	36
Navigating to the Grid Manager	36
Specifying the StorageGRID license information	37
Adding sites	37
Specifying Grid Network subnets	38
Approving pending grid nodes	39
Specifying Network Time Protocol server information	43
Specifying Domain Name System server information	44
Specifying the StorageGRID system passwords	45
Reviewing your configuration and completing installation	46
Automating the installation	49
Automating grid node deployment in VMware vSphere	49
Defining the configuration file for your deployment	50
Running the Bash script	57
Automating the configuration of StorageGRID	58
Automating the configuration and installation of appliance Storage Nodes	59
Overview of installation REST APIs	62

Where to go next	64
Troubleshooting	66
Virtual machine resource reservation requires adjustment	66
Copyright	67
Trademark	68
How to send comments about documentation and receive update notifications	69

Installation overview

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

1. **Preparation:** You must make sure you have all of the required materials, installation files, hardware, and networks.
2. **Deployment:** When you deploy grid nodes, the individual grid nodes are created and connected to one or more networks. When deploying grid nodes:
 - a. You deploy the virtual nodes as virtual machines (VMs) in VMware vSphere Web Client using a `.vmdk` file and a set of `.ovf` file templates.
 - b. You deploy any StorageGRID appliance Storage Nodes, using the StorageGRID Appliance Installer.
3. **Configuration:** When all nodes have been deployed, you use the StorageGRID Grid Manager to configure the grid and complete the installation.

This document recommends a standard approach for deploying and configuring a StorageGRID system in a VMware environment, but it also provides information about these alternative approaches:

- Using the `deploy-vmware-ovftool.sh` Bash script (available from the installation archive) to deploy grid nodes in VMware vSphere.
- Configuring the StorageGRID system with a Python configuration script (available from the installation archive).
- Deploying and configuring appliance Storage Nodes with a second Python configuration script (available from the installation archive or from the StorageGRID Appliance Installer).
- Using the installation REST APIs to automate the installation of StorageGRID grid nodes and appliance Storage Nodes.

Related concepts

[Planning and preparation](#) on page 6

[Overview of installation REST APIs](#) on page 62

Related tasks

[Deploying virtual machine grid nodes in VMware vSphere Web Client](#) on page 23

[Deploying appliance Storage Nodes](#) on page 31

[Configuring the grid and completing installation](#) on page 36

[Automating grid node deployment in VMware vSphere](#) on page 49

[Automating the configuration of StorageGRID](#) on page 58

[Automating the configuration and installation of appliance Storage Nodes](#) on page 59

Planning and preparation

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. If the StorageGRID system includes StorageGRID appliance Storage Nodes, you must be familiar with all steps and requirements necessary to install the appliance hardware.

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.

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.
 - Note:** 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. See “Networking requirements” for more information.
- 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 concepts

[Networking guidelines](#) on page 12

Related information

[SG6000 appliance installation and maintenance](#)

[SG5700 appliance installation and maintenance](#)

[SG5600 appliance installation and maintenance](#)

Required materials

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

Item	Notes
NetApp StorageGRID license	<p>You must have a valid, digitally signed NetApp license.</p> <p>Note: The StorageGRID installation archive includes a free license that does not provide any support entitlement for the product.</p>

Item	Notes
StorageGRID installation archive for VMware	<p>You must download one of the following StorageGRID installation archives and extract the files to your service laptop.</p> <ul style="list-style-type: none"> StorageGRID-Webscale-version-VMware-uniqueID.tgz StorageGRID-Webscale-version-VMware-uniqueID.zip
VMware software and documentation	<p>During installation, you deploy virtual grid nodes on virtual machines in VMware vSphere Web Client.</p> <p>For supported versions, see the Interoperability Matrix.</p>
Service laptop	<p>The StorageGRID system is installed through a service laptop</p> <p>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 tasks

[Downloading and extracting the StorageGRID installation files](#) on page 7

Related references

[Web browser requirements](#) on page 22

Related information

[NetApp Interoperability Matrix Tool](#)

[VMware: Featured Documentation Sets](#)

[Administering StorageGRID](#)

Downloading and extracting the StorageGRID installation files

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

Steps

1. Go to the Software Download page on the NetApp Support Site.
[NetApp Downloads: Software](#)
2. Sign in using the username and password for your NetApp account.
3. Scroll to **StorageGRID (formerly StorageGRID Webscale)**, select **All Platforms**, and click **Go**.
4. Select the StorageGRID release, and click **View & Download**.

5. From the Software Download section of the page, click **CONTINUE**, and accept the End User License Agreement.
6. Download the .tgz or .zip file for your platform.

- `StorageGRID-Webscale-version-VMware-uniqueID.tgz`

- `StorageGRID-Webscale-version-VMware-uniqueID.zip`

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

7. Extract the installation archive.
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 grid.

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

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

Filename	Description
/vsphere/configure-storagegrid.py	A Python script used to automate the configuration of a StorageGRID system.
/vsphere/configure-storagegrid.sample.json	A sample configuration file for use with the configure-storagegrid.py script.
/vsphere/configure-storagegrid.blank.json	A blank configuration file for use with the configure-storagegrid.py script.

Hardware requirements

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

For information about supported servers, see the Interoperability Matrix.

Related information

[NetApp Interoperability Matrix Tool](#)

Software requirements

All StorageGRID grid nodes, except for StorageGRID appliance grid nodes, are hosted by virtual machines. One virtual machine is required for each virtual 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.
- 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

Caution: 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](#)
[VMware: Featured Documentation Sets](#)

Virtual machine requirements

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

The maximum tested LUN size is 39 TB.

Node Type	vCPU	RAM	Storage
Admin Node	8	24 GB	100 GB LUN for OS 200 GB LUN for Admin Node tables 200 GB LUN for Admin Node audit log
Storage Node	8	24 GB	100 GB LUN for OS 1 to 16 Storage LUNs, as required
API Gateway Node	8	24 GB	100 GB LUN for OS
Archive Node	8	24 GB	100 GB LUN for OS

Note: You must provide at least 8 vCPUs per node, whether or not hyperthreading is enabled on the underlying physical ESX hosts.

See “Storage and performance requirements” for more information.

Related concepts

[Storage and performance requirements](#) on page 21

Networking requirements

You must verify that the networking infrastructure and configuration is in place to support your StorageGRID system.

For more information on networking configuration and supported network topologies, see the *Grid primer*.

Related information

[Grid primer](#)

Network model

You can configure three networks for use with the StorageGRID system. Each network type must be on a separate subnet with no overlap.

To understand how these three networks are used, consider the three types of network traffic that are processed by nodes in a StorageGRID system:

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

To allow you more precise control and security, you can configure one, two, or three networks to manage these three types of traffic.

Grid Network

The Grid Network is required. It is used for all internal StorageGRID traffic. The Grid Network provides connectivity between all nodes in the grid, across all sites and subnets. All hosts on the Grid Network must be able to talk to all other hosts. The Grid Network can consist of multiple subnets. Networks containing critical grid services, such as NTP, can also be added as Grid subnets.

When the Grid Network is the only StorageGRID network, it is also used for all admin traffic and all client traffic. The Grid Network gateway is the node default gateway unless the node has the Client Network configured.

Attention: When configuring the Grid Network, you must ensure that the network is secured from untrusted clients, such as those on the open internet.

The Grid Network is mapped using `GRID_NETWORK_TARGET` (eth0 inside the container).

IP/ Mask	Gateway	Static routes	Default route (0.0.0.0/0)
CIDR for static IP	The Grid Network gateway must be configured if there are multiple grid subnets. The Grid Network gateway is the node default gateway until grid configuration is complete.	Static routes are generated automatically for all nodes to all subnets configured in the global Grid Network Subnet List.	The Grid Network Gateway IP is the default gateway. If a Client Network is added, the default gateway switches from the Grid Network gateway to the Client Network gateway when grid configuration is complete.

Admin Network

The Admin Network is optional. It is a closed network used for system administration and maintenance. The Admin Network is typically a private network and does not need to be routable between sites.

Using the Admin Network for administrative access allows the Grid Network to be isolated and secure. Typical uses of the Admin Network include access to the Grid Manager, access to critical services, such as NTP and DNS, access to audit logs on Admin Nodes, and SSH access to all nodes for maintenance and support. The Admin Network is never used for internal grid traffic. An Admin Network gateway is provided and allows the Admin Network to span multiple subnets. However, the Admin Network gateway is never used as the node default gateway.

The Admin Network is mapped using `ADMIN_NETWORK_TARGET` (eth1 inside the container).

IP/Mask	Gateway	Static routes	Default route (0.0.0.0/0)
CIDR for static IP	The Admin Network gateway is required if multiple admin subnets are defined.	Static routes are generated automatically to each subnet configured in the node's Admin Network Subnet List.	N/A

Client Network

The Client Network is also optional. It is an open network used to provide access to grid services for client applications such as S3 and Swift. The Client Network enables grid nodes to communicate

with any subnet reachable through the Client Network gateway. The Client Network does not become operational until you complete the StorageGRID configuration steps.

You can use the Client Network to provide client access to the grid, so you can isolate and secure the Grid Network. The following nodes are often configured with a Client Network:

- API Gateway Nodes and Storage Nodes, because these nodes provide S3 and Swift protocol access to the grid.
- Admin Nodes, because these nodes provide access to the Tenant Manager.

When a Client Network is configured, the Client Network gateway is required and becomes the node default gateway after the grid has been configured.

The Client Network is mapped using `CLIENT_NETWORK_TARGET` (eth2 inside the container).

IP/Mask	Gateway	Static routes	Default route (0.0.0.0/0)
CIDR for static IP	The Client Network gateway is required if the Client Network is configured. The Client Network gateway becomes the default route for the grid node when grid configuration is complete.	N/A	Added if a Client Network Gateway IP is configured

Networking guidelines

Follow these guidelines when configuring StorageGRID networks.

Grid Network guidelines

- You configure the Grid Network Subnet List using the Grid Manager to enable static route generation between subnets on the Grid Network.
- Each node must be attached to the Grid Network and must be able to communicate with the primary Admin Node using the networking configuration you specify when deploying the node.

Subnet guidelines

Note: The following restrictions are enforced by the Grid Manager during deployment. They are provided here to assist in pre-deployment network planning.

- The subnets defined by a given node's IP/mask (CIDR) values for all networks it uses cannot overlap each other.
- The subnet defined by a given node's Admin Network IP/mask (CIDR) value cannot overlap any subnet in the Grid Network Subnet List.
- The subnet defined by a given node's Client Network IP/mask (CIDR) value cannot overlap any subnet in the Grid Network Subnet List, nor any subnet present in that node's Admin Network External Subnet List (A-ESL).
- The subnet defined by the Grid Network IP/mask (CIDR) value of any given node in the grid must be contained within at least one subnet present in the Grid Network Subnet List.
- No subnet in the Grid Network Subnet List can overlap with any subnet in any node's Admin Network External Subnet List (A-ESL).

Gateway guidelines

- If set, the gateway for a given network must lie within the subnet defined by the node's IP/mask (CIDR) value for that network.
- If you configure an interface using static addressing, you must specify a gateway address other than 0.0.0.0.
- If you do not have a gateway, best practice is to set the gateway address to be equal to the grid network IP address or to the .1 address on the subnet.

General guidelines

- At least one NTP server must be reachable by the primary Admin Node, using the networking configuration you specified when deploying the primary Admin Node.
- If you are not ready to configure the optional Admin and Client Networks during deployment, you can configure these networks when you approve grid nodes during the configuration steps. See “Approving pending grid nodes” for more information.
- Admin Nodes must always be secured from untrusted clients, such as those on the open internet. You must ensure that no untrusted client can access any Admin Node on the Grid Network, the Admin Network, or the Client Network.

Post-configuration guidelines

After completing configuration:

- If DHCP was used to assign IP addresses, you should configure a DHCP reservation for each IP address on these networks. You can only set up DHCP during the deployment phase. You cannot set up DHCP during configuration.
 - Attention:** Because nodes reboot when their IP addresses change, outages can occur if a DHCP address change affects multiple nodes at once.
- You must use the IP address change procedures if you want to change IP addresses, subnet masks, and default gateways for a grid node. See information about configuring IP addresses in the recovery and maintenance instructions.
- 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.

For more information on the StorageGRID network model and various ways to use it, review the networking topics in the *Grid primer*.

Related tasks

[Approving pending grid nodes](#) on page 39

Related information

[Recovery and maintenance](#)
[Grid primer](#)

Network installation and provisioning

You must understand how the Grid Network and the optional Admin and Client Networks are used during node deployment and grid configuration.

When you first deploy a node, you must attach the node to the Grid Network and ensure it has access to the primary Admin Node. If the Grid Network is isolated, you can configure the Admin Network on the primary Admin Node for configuration and installation access from outside the Grid Network. If the Grid Network has a gateway configured, it is the default gateway for the node during deployment. This allows grid nodes on separate subnets to communicate with the primary Admin Node before the grid has been configured.

Once the nodes have been deployed, the nodes register themselves with the primary Admin Node using the Grid Network. You can then use the Grid Manager, the `configure-storagegrid.py` Python script, or the Installation API to configure the grid and approve the registered nodes. During grid configuration, you can configure multiple grid subnets. Static routes to these subnets through the Grid Network gateway will be created on each node when you complete grid configuration. If necessary, subnets containing NTP servers or requiring access to the Grid Manager or API can also be configured as grid subnets.

If you want to disable the Admin Network or Client Network, you can remove the configuration from them during the node approval process.

Note: When using the Client Network, keep in mind that a node's default gateway will switch from the Grid Network to the Client Network when you complete the grid configuration steps. For all nodes, you must ensure that the node does not lose access to external NTP servers when the gateway switches. For Admin Nodes, you must also ensure that browsers or API clients do not lose access to the Grid Manager. To maintain access, perform one of the following steps:

- When configuring the node, route Grid Manager traffic (Admin Nodes only) and NTP traffic through the Admin Network.
- Add subnets to the Grid Network Subnet List (GNSL) that include the IPs of remote clients and servers that should communicate with the grid over the Grid Network.
- Ensure that both the Grid and Client Network gateways can route traffic to and from the external NTP servers and browsers or other Grid Manager API clients.

If you are creating...	Behavior	Recommended configuration
Grid Network only	All Grid, Admin, and Client traffic flows over the Grid Network. The Grid Network gateway is the node default gateway.	
Grid Network and Admin Network	Grid and Client traffic flows over the Grid Network. Administrative traffic flows over the Admin Network. The Grid Network gateway is the node default gateway.	

If you are creating...	Behavior	Recommended configuration
Grid Network and Client Network (no Admin Network)	When a node is deployed, the Grid Network gateway is the node default gateway. Subnets providing access to the Grid Manager and NTP servers should be included as Grid Network subnets during configuration. When you complete the grid configuration steps, the Client Network gateway becomes the node default gateway.	Allow NTP and installer client access through both the Grid and Client Network gateways. <i>or</i> Add the NTP or installer Client subnets, or both, as Grid Networks.
All three networks (Grid, Admin, and Client)	When a node is deployed, the Grid Network gateway is the node default gateway. Subnets providing access to the Grid Manager and NTP servers should be included on the Grid Network subnets or as Admin Network subnets during configuration. When you complete the grid configuration steps, the Client Network gateway becomes the node default gateway.	Allow NTP and installer client access through both the Grid and Client Network gateways. <i>or</i> Add the NTP or installer client subnets, or both, as Grid Networks (so explicit routes will be created). <i>or</i> Add NTP and installer client subnets to the Admin Network External Subnet List (AESL).
Client Network, but at a later time	Subnets providing access to the Grid Manager and NTP servers should be included as Grid Networks or as Admin subnets. The Client Network gateway will become the node default gateway	Allow NTP and installer client access through both the Grid and Client Network gateways. <i>or</i> Add the NTP or installer client subnets, or both, as Grid Networks (so explicit routes will be created). <i>or</i> Add NTP and installer client subnets to the AESL.

Internal grid node communications

The following ports must be accessible to grid nodes on the Grid Network. Ensure that the required ports for the grid node type are open on the server.

Configuration of the following ports is only required if you need to define firewall restrictions that are external to VMware networking.

If enterprise networking policies restrict the availability of any of these ports, you can remap ports when you deploy nodes using the VMware Vsphere Web Client, or by using a configuration file setting when automating grid node deployment.

Note: NetApp recommends that you enable ICMP traffic between grid nodes. Allowing ICMP traffic can improve failover performance when a grid node cannot be reached.

Port	TCP or UDP	Protocol	From	To	Details
22	TCP	SSH	Primary Admin Node	All nodes	For maintenance procedures, the primary Admin Node must be able to communicate with all other nodes using SSH on port 22. Allowing SSH traffic from other nodes is optional.
80	TCP	HTTP	Appliances	Primary Admin Node	Used by StorageGRID appliance Storage Nodes to communicate with the primary Admin Node to start the installation.
123	UDP	NTP	All nodes	All nodes	Network time protocol service. Every node synchronizes the time with every other node using NTP.
443	TCP	HTTPS	All nodes	Primary Admin Node	Used for communicating status to the primary Admin Node during installation and other maintenance procedures.
1139	TCP	HTTPS	Storage Nodes	Storage Nodes	ILM-related internal traffic between Storage Nodes.
1501	TCP	HTTPS	All nodes	Storage Nodes with ADC	Reporting, auditing, and configuration internal traffic.
1502	TCP	HTTPS	All nodes	Storage Nodes	S3- and Swift-related internal traffic.
1503	TCP	NTTPS	All nodes	Storage Nodes	ILM-related internal traffic for the deprecated CMS service.
1504	TCP	HTTPS	All nodes	Admin Nodes	NMS service reporting and configuration internal traffic.
1505	TCP	HTTPS	All nodes	Admin Nodes	AMS service internal traffic.
1506	TCP	HTTPS	All nodes	All nodes	Server status internal traffic.
1507	TCP	HTTPS	All nodes	API Gateway Nodes	Load balancer internal traffic.
1508	TCP	HTTPS	All nodes	Primary Admin Node	Configuration management internal traffic.
1509	TCP	HTTPS	All nodes	Archive Nodes	Archive Node internal traffic.
1511	TCP	HTTPS	All nodes	Storage Nodes	Metadata internal traffic.
5353	UDP	mDNS	All nodes	All nodes	Optionally used for full-grid IP changes and for primary Admin Node discovery during installation, expansion, and recovery.

Port	TCP or UDP	Protocol	From	To	Details
7001	TCP	Cassandra	Storage Nodes	Storage Nodes	Cassandra TLS inter-node cluster communication.
9042	TCP	Cassandra	Storage Nodes	Storage Nodes	Cassandra client port.
9999	TCP	HTTPS	All nodes	All nodes	Internal traffic for multiple services. Includes maintenance procedures, metrics, and networking updates.
11139	TCP	HTTPS	Archive/Storage Nodes	Archive/Storage Nodes	ILM-related internal traffic between Storage Nodes and Archive Nodes.
18000	TCP	HTTPS	Admin/Storage Nodes	Storage Nodes with ADC	Account service internal traffic.
18001	TCP	HTTPS	Admin/Storage Nodes	Storage Nodes with ADC	Identity Federation internal traffic.
18002	TCP	HTTPS	Admin/Storage Nodes	Storage Nodes	Internal API traffic related to object protocols.
18003	TCP	HTTPS	Admin/Storage Nodes	Storage Nodes with ADC	Platform services internal traffic.
18017	TCP	HTTPS	Admin/Storage Nodes	Storage Nodes	Data Mover service internal traffic for Cloud Storage Pools.
18082	TCP	HTTPS	Admin/Storage Nodes	Storage Nodes	S3-related internal traffic.
18083	TCP	HTTPS	All nodes	Storage Nodes	Swift-related internal traffic.
18200	TCP	HTTPS	Admin/Storage Nodes	Storage Nodes	Additional statistics about client requests.
19000	TCP	HTTPS	Admin/Storage Nodes	Storage Nodes with ADC	Keystone service internal traffic.

External communications

Clients need to communicate with grid nodes and, by extension, the servers that host them in order to ingest and retrieve content. The ports used depends on the protocols chosen to ingest and retrieve content.

If enterprise networking policies restrict the availability of any of the ports used for traffic into or out of the nodes, you can remap ports when deploying nodes.

The following table shows the ports used for traffic into the nodes.

Port	TCP or UDP	Protocol	From	To	Details
22	TCP	SSH	Service laptop	All nodes	SSH or console access is required for procedures with console steps. Optionally, you can use port 2022 instead of 22.
25	TCP	SMTP	Admin Nodes	Email server	Used for alerts and email-based AutoSupport. You can override the default port setting of 25 using the Email Servers page.
53	TCP/UDP	DNS	All nodes	DNS servers	Used for domain name system.
67	UDP	DHCP	All nodes	DHCP service	Optionally used to support DHCP-based network configuration. The dhclient service does not run for statically-configured grids.
68	UDP	DHCP	DHCP service	All nodes	Optionally used to support DHCP-based network configuration. The dhclient service does not run for grids that use static IP addresses.
80	TCP	HTTP	Browser	Admin Nodes	Port 80 redirects to port 443 for the Admin Node user interface.
			Browser	Appliances	Port 80 redirects to port 8443 for the StorageGRID appliance's install/maintenance user interface.
			Storage Nodes with ADC	AWS	Used for platform services messages sent to AWS or other external services that use HTTP. Tenants can override the default HTTP port setting of 80 when creating an endpoint.
			Storage Nodes	AWS	Cloud Storage Pools requests sent to AWS targets that use HTTP. Grid administrators can override the default HTTP port setting of 80 when configuring a Cloud Storage Pool.
111	TCP/UDP	RPCBind	NFS client	Admin Nodes	Used by NFS-based audit export (portmap).
123	UDP	NTP	Primary NTP nodes	External NTP	Network time protocol service. Nodes selected as primary NTP sources also synchronize clock times with the external NTP time sources.
137	UDP	NetBIOS	SMB client	Admin Nodes	Used by SMB-based audit export for clients that require NetBIOS support.
138	UDP	NetBIOS	SMB client	Admin Nodes	Used by SMB-based audit export for clients that require NetBIOS support.
139	TCP	SMB	SMB client	Admin Nodes	Used by SMB-based audit export for clients that require NetBIOS support.

Port	TCP or UDP	Protocol	From	To	Details
161	TCP/UDP	SNMP	SNMP client	All nodes	Used for SNMP polling. All nodes provide basic information; Admin Nodes additionally provide alarm data. Note: For information about using SNMP with StorageGRID, contact your NetApp sales representative.
162	TCP/UDP	SNMP Traps	All nodes	Trap destinations	Outbound SNMP traps default to port 162. Note: For information about using SNMP with StorageGRID, contact your NetApp sales representative.
389	TCP/UDP	LDAP	Storage Nodes with ADC	Active Directory/ LDAP	Used for connecting to an Active Directory or LDAP server for Identity Federation.
443	TCP	HTTPS	Browser	Admin Nodes	Used by web browsers and management API clients for accessing the Grid Manager and Tenant Manager.
			Admin Nodes	Active Directory	Used by Admin Nodes connecting to Active Directory if single sign-on (SSO) is enabled.
			Archive Nodes	AWS S3	Used for accessing AWS S3 from Archive Nodes.
			Storage Nodes with ADC	AWS	Used for platform services messages sent to AWS or other external services that use HTTPS. Tenants can override the default HTTP port setting of 443 when creating an endpoint.
			Storage Nodes	AWS	Cloud Storage Pools requests sent to AWS targets that use HTTPS. Grid administrators can override the default HTTPS port setting of 443 when configuring a Cloud Storage Pool.
445	TCP	SMB	SMB client	Admin Nodes	Used by SMB-based audit export.
903	TCP/UDP	NFS	NFS client	Admin Nodes	Used by NFS-based audit export (<code>rpc.mountd</code>).
2022	TCP	SSH	Service laptop	All nodes	SSH or console access is required for procedures with console steps. Optionally, you can use port 2022 can instead of 22.
2049	TCP/UDP	NFS	NFS client	Admin Nodes	Used by NFS-based audit export (<code>nfs</code>).

Port	TCP or UDP	Protocol	From	To	Details
8022	TCP	SSH	Service laptop	All nodes	SSH on port 8022 grants access to the base operating system for support and troubleshooting. This port is not required to be accessible between grid nodes or during normal operations.
8082	TCP	HTTPS	S3 clients	API Gateway Nodes	S3-related external traffic to API Gateway Nodes (HTTPS).
8083	TCP	HTTPS	Swift clients	API Gateway Nodes	Swift-related external traffic to API Gateway Nodes (HTTPS).
8084	TCP	HTTP	S3 clients	API Gateway Nodes	S3-related external traffic to API Gateway Nodes (HTTP).
8085	TCP	HTTP	Swift clients	API Gateway Nodes	Swift-related external traffic to API Gateway Nodes (HTTP).
9022	TCP	SSH	Service laptop	Appliances	Grants access to StorageGRID appliances in pre-configuration mode for support and troubleshooting. This port is not required to be accessible between grid nodes or during normal operations.
18082	TCP	HTTPS	S3 clients	Storage Nodes	S3-related external traffic to Storage Nodes (HTTPS).
18083	TCP	HTTPS	Swift clients	Storage Nodes	Swift-related external traffic to Storage Nodes (HTTPS).
18084	TCP	HTTP	S3 clients	Storage Nodes	S3-related external traffic to Storage Nodes (HTTP).
18085	TCP	HTTP	Swift clients	Storage Nodes	Swift-related external traffic to Storage Nodes (HTTP).

Networking and ports for platform services and Cloud Storage Pools

If you plan to use platform services or Cloud Storage Pools, you must configure grid networking and ports to ensure that the destination endpoints can be reached.

To ensure that platform services messages can be delivered to the configured endpoints, you must configure the network or networks containing ADC Storage Nodes. Similarly, to ensure that ILM rules can move objects to and from a specified Cloud Storage Pool, you must configure the network or networks that contain your grid's Storage Nodes.

By default, platform services and Cloud Storage Pool communications use the following ports:

- **80:** For endpoint URIs that begin with `http`
- **443:** For endpoint URIs that begin with `https`
- **8082:** For external HTTPS communications received by an API Gateway Node

- **8084:** For external HTTP communications received by an API Gateway Node

A different port can be specified when the endpoint is created or edited.

If you use a non-transparent proxy server, you must also configure proxy settings to allow messages to be sent to external endpoints, such as an endpoint on the internet. See *administering StorageGRID* to learn how to configure proxy settings.

Related information

[Administering StorageGRID](#)

[Using tenant accounts](#)

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 AFF storage

If a StorageGRID node is deployed in a virtual machine with storage assigned from a NetApp All Flash FAS (AFF) system, confirm that the FlexVol does not have a tiering policy enabled.

Attention: Never assign storage for a StorageGRID node from a FlexVol with an active tiering policy. Node outages might occur if storage used by the StorageGRID node is tiered to a capacity tier.

Number of virtual machines required

Each StorageGRID site requires a minimum of three Storage Nodes.

Attention: 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 for Storage Nodes

A Storage Node can have 1 to 16 storage volumes; however, 3 or more storage volumes are recommended. Each storage volume should be 4 TB or larger.

You must assign at least 3 TB to volume 0 of each Storage Node. StorageGRID uses volume 0 to reserve 3 TB of space (the Metadata Reserved Space, or CAWM) for object metadata and for essential database operations, such as compaction and repair.

If you assign 4 TB or more to volume 0, as recommended, the additional space is used for object storage.

Note: If you use only one storage volume for a Storage Node and you assign 3 TB or less to the volume, the Storage Node immediately enters the Storage Read-Only state on startup and stores object metadata only.

Note: If you assign less than 500 GB to volume 0 (non-production use only), 10% of the storage volume's capacity is reserved for metadata.

Web browser requirements

You must use a supported web browser.

Web browser	Minimum supported version
Google Chrome	70
Microsoft Internet Explorer	11 (Native Mode)
Mozilla Firefox	63

You should set the browser window to a recommended width.

Browser width	Pixels
Minimum	1024
Optimum	1280

Deploying grid nodes

When you deploy grid nodes in a VMware environment, the individual grid nodes are created and connected to one or more networks.

Steps

1. [Deploying virtual machine grid nodes in VMware vSphere Web Client](#) on page 23
2. [Deploying appliance Storage Nodes](#) on page 31

Related tasks

[Automating the installation](#) on page 49

Deploying virtual machine grid nodes in VMware vSphere Web Client

You can use VMware vSphere Web Client to deploy each grid node that will run on a virtual machine.

Collecting information about your deployment environment

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

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.

VMware information

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.

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).

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 concepts

[Storage and performance requirements](#) on page 21

Related references

[How grid nodes discover the primary Admin Node](#) on page 24

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.

Attention: 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.

Deploying the primary Admin Node virtual machine

You deploy the primary Admin Node before deploying other grid nodes.

Before you begin

- You have access to the `.ovf` and `.mf` files for the primary Admin Node:
 - `vsphere-primary-admin.ovf`
 - `vsphere-primary-admin.mf`
- You have the StorageGRID Virtual Machine Disk (`.vmdk`) file, and it is in the same folder as the `.ovf` and `.mf` files:

`NetApp-SG-version-SHA.vmdk`

Note: The same `.vmdk` file is used for all types of nodes.

- You must be familiar with VMware vSphere Hypervisor and have experience deploying virtual machines in this environment.

Note: 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.

Steps

1. Open the VMware vSphere Web Client, and sign in.
2. Navigate to the appropriate location in the **vCenter** hierarchy. From there, right-click, select **Deploy OVF Template**, and select `vsphere-primary-admin.ovf` and `NetApp-SG-version-SHA.vmdk`.
Normally you can navigate to **Hosts and Clusters**, and deploy the template as a specific user or within a specific resource pool.
3. Read the information on the **End User License Agreement** page, and click **Accept**.
4. Specify the name of the virtual machine, and select a datastore for the disks.
The best practice is to use the same name for the virtual machine as you used for the grid node.
5. In the **Network Mapping** page, select the networks to use by associating a network port to each network.
The Grid Network is required. The Admin and Client Networks are optional.
 - If you are going to use the Admin Network, assign the Admin Network adapter to a network in the vSphere environment.

- If you are going to use the Client Network, assign the Client Network adapter to a network in the vSphere environment
- If you are not going to use an Admin Network or Client Network, assign their network adapters to the same network as the Grid Network.

Note: The following warning is displayed, but you can ignore it because you will disable the Admin and Client Networks on the next page.

```
Warning: Multiple source networks are mapped to the host network:
<Destination Network Name>
```

6. Provide the required StorageGRID information in the **Properties** page:

Note: This step is for setting Network IP information. After deploying grid nodes, you use the Grid Manager to configure networks.

a. Enter the **Node Name**.

Attention: If you are performing this installation task because you are recovering a grid node, you must use the same name for the replacement node that was used for the node you are recovering.

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

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

c. In the **Admin Network (eth1)** section, under **Admin Network IP configuration**, select STATIC, DHCP, or DISABLED.

- If you select STATIC, enter the **Admin Network IP**, **Admin Network mask**, and **Admin Network gateway**.
- If you select DHCP, the **Admin Network IP**, **Admin Network mask**, and **Admin Network gateway** are automatically assigned.
- If you do not want to use the Admin Network (eth1), select DISABLED and enter **0.0.0.0** for the Admin Network IP. You can leave the other fields blank.

d. Under the **Admin Network (eth1)** section, enter the **Admin Network external subnet list**.

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

- If you select STATIC, enter the **Client Network IP**, **Client Network mask**, and **Client Network gateway**.
- If you select DHCP, the **Client Network IP**, **Client Network mask**, and **Client Network gateway** are automatically assigned.
- If you do not want to use the Client Network (eth2), select DISABLED and enter **0.0.0.0** for the Client Network IP. You can leave the other fields blank.

7. Click **Finish**.

8. If you want to remap ports used by the primary Admin Node:

- a. Right-click on the VM, and select **Edit Settings**.
- b. Select **vApp Options**.
- c. In the **Authoring** section, expand **Properties** and scroll down until you see **PORT_REMAP_INBOUND** and **PORT_REMAP**.

You might need to remap a port if your enterprise networking policies restrict access to one or more ports that is used by StorageGRID.

- d. To symmetrically map both inbound and outbound communications for a port, select **PORT_REMAP** and click **Edit**.

Enter the port mapping as *<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.

Example

To remap ssh traffic from port 22 to port 3022, enter the following:

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

Click **OK**.

Note: 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.

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

Enter the port mapping as *<network type>/<protocol>/<remapped inbound port>/<default inbound port used by grid node>*, where network type is grid, admin, or client, and protocol is tcp or udp.

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
```

Click **OK**.

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

9. Power on the virtual machine if it is not active.

Related concepts

[Internal grid node communications](#) on page 15

[External communications](#) on page 17

Deploying additional grid node virtual machines

After deploying the primary Admin Node, you can deploy the other virtual machine-based grid nodes in your StorageGRID system.

Before you begin

- You have the `.ovf` and `.mf` files for the grid nodes you are deploying:

Filename	Description
<code>vsphere-non-primary-admin.ovf</code> <code>vsphere-non-primary-admin.mf</code>	The template file and manifest file for deploying non-primary Admin Nodes.
<code>vsphere-archive.ovf</code> <code>vsphere-archive.mf</code>	The template file and manifest file for deploying Archive Nodes.
<code>vsphere-gateway.ovf</code> <code>vsphere-gateway.mf</code>	The template file and manifest file for deploying API Gateway Nodes.
<code>vsphere-storage.ovf</code> <code>vsphere-storage.mf</code>	The template file and manifest file for deploying Storage Nodes.

- You have placed all of these files in the same directory.
- You have the StorageGRID Virtual Machine Disk (`.vmdk`) file, and it is in the same folder as the `.ovf` and `.mf` files:

`NetApp-SG-version-SHA.vmdk`

Note: The same `.vmdk` file is used for all types of nodes.

About this task

You must deploy all grid nodes before configuring the grid. When deploying a StorageGRID virtual machine, ensure that it can connect to the primary Admin Node over the Grid Network.

Steps

- Open VMware vSphere Web Client, and sign in.
- Navigate to the vApp or resource pool where you want to deploy the StorageGRID grid, and select **Actions > All vCenter Actions > Deploy OVF Template**.
- Select the `vsphere-node.ovf` and `NetApp-SG-version-SHA.vmdk` files.
- Specify the name of the virtual machine.
The best practice is to use the same name for the virtual machine as you used for the grid node.
- In the **Network Mapping** page, select the networks to use by associating a network port to each network. The Grid Network is required. The Admin and Client Networks are optional. Select the Grid Network to use, and then choose the following as applicable:
 - If you are planning to use the Admin Network, assign the Admin Network adapter to a network in the vSphere environment.
 - If you are planning to use the Client Network, assign the Client Network adapter to a network in the vSphere environment.
 - If you do not plan to use an Admin Network or Client Network, assign their network adapters to the same network as the Grid Network.

6. Provide the required StorageGRID information in the **Properties** page, and click **Finish**.
 - a. Enter the **Node Name**.

Attention: If you are performing this task because you are recovering a grid node, you must use the same name for the replacement node that was used for the node you are recovering.
 - b. Enter the **Primary Admin IP**.

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.
 - c. In the **Grid Network (eth0)** section, under **Grid Network IP configuration**, select STATIC or DHCP.
 - If you select STATIC, enter the **Grid Network IP**, **Grid Network mask**, and **Grid Network gateway**.
 - If you select DHCP, the **Grid Network IP**, **Grid Network mask**, and **Grid Network gateway** are automatically assigned.
 - d. In the **Admin Network (eth1)** section, under **Admin Network IP configuration**, select STATIC, DHCP, or DISABLED.
 - If you select STATIC, enter the **Admin Network IP**, **Admin Network mask**, and **Admin Network gateway**.
 - 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.
 - If you do not want to use the Admin Network (eth1), select DISABLED and enter **0.0.0.0** for the Admin Network IP. You can leave the other fields blank.
 - e. In the **Client Network (eth2)** section, under **Client Network IP configuration**, select STATIC, DHCP, or DISABLED.
 - If you select STATIC, enter the **Client Network IP**, **Client Network mask**, and **Client Network gateway**.
 - If you select DHCP, the **Client Network IP**, **Client Network mask**, and **Client Network gateway** are automatically assigned.
 - If you do not want to use the Client Network (eth2), select DISABLED and enter **0.0.0.0** for the Client Network IP. You can leave the other fields blank.
7. Click **Next** and then **Finish** to start the upload of the virtual machine.
8. Assign storage to the virtual machine.

If you are using only one storage volume for each Storage Node, you must assign it at least 4 TB of space. If you are using more than one storage volume for each Storage Node, you must assign at least 3 TB to object store 0; however, 4 TB is recommended. See “Storage and performance requirements” for more information.

Attention: The Storage Node OVF provided 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).

9. If you want to remap ports used by a node:
 - a. If you specified **DISABLED** for the Client network IP configuration, you must enter **0.0.0.0** for the Client Network IP under the **Client Network (eth2)** section. Completing this field is required.
 - b. Right-click on the VM, and select **Edit Settings**.
 - c. Select **vApp Options**.
 - d. In the **Authoring** section, expand **Properties** and scroll down until you see **PORT_REMAP_INBOUND** and **PORT_REMAP**.

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 “Internal grid node communications” or “External client communications” for the list of ports used by StorageGRID.

- e. To symmetrically map both inbound and outbound communications for a port, select **PORT_REMAP** and click **Edit**.

Enter the port mapping as *<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.

Example

To remap ssh traffic from port 22 to port 3022, enter the following:

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

Click **OK**.

Note: 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.

- f. To specify the port used for inbound communications to the node, select **PORT_REMAP_INBOUND** and click **Edit**.

Enter the port mapping as *<network type>/<protocol>/<remapped inbound port>/<default inbound port used by grid node>*, where network type is grid, admin, or client, and protocol is tcp or udp.

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
```

Click **OK**.

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

10. Power on the virtual machine.

Related concepts

[Storage and performance requirements](#) on page 21

[Internal grid node communications](#) on page 15

[External communications](#) on page 17

Deploying appliance Storage Nodes

You can deploy any StorageGRID appliances into your StorageGRID grid. Each appliance functions as a single Storage Node and can connect to the Grid Network, the Admin Network, and the Client Network.

Steps

1. [Starting StorageGRID appliance installation](#) on page 31
2. [Monitoring StorageGRID appliance installation](#) on page 34

Starting StorageGRID appliance installation

To install StorageGRID on an appliance Storage Node, you use the StorageGRID Appliance Installer, which is included on the appliance.

Before you begin

- The appliance has been installed in a rack or cabinet, connected to your networks, and powered on.
- Network links, IP addresses, and port remapping (if necessary) have been configured for the appliance using the StorageGRID Appliance Installer.
- The primary Admin Node for the StorageGRID grid has been deployed.
- All Grid Network subnets listed on the IP Configuration page of the StorageGRID Appliance Installer have been defined in the Grid Network Subnet List on the primary Admin Node.

For instructions for completing these prerequisite tasks, see the installation and maintenance instructions for your appliance.

- You have a service laptop with a supported web browser.
- You know one of the IP addresses assigned to the appliance's compute controller. You can use the IP address for any attached StorageGRID network.

About this task

To install StorageGRID on an appliance Storage Node:

- You specify or confirm the IP address of the primary Admin Node and the name of the Storage Node.
- You start the installation and wait as volumes are configured and the software is installed.
- Partway through the appliance installation tasks, the installation pauses. To resume the installation, you sign into the Grid Manager, approve all grid nodes, and complete the StorageGRID installation process.

Note: If you need to deploy multiple StorageGRID appliance Storage Nodes at one time, you can automate the installation process by using the `configure-sga.py` Appliance Installation Script.

Steps

1. Open a browser, and enter one of the IP addresses for the compute controller.

https://Controller_IP:8443

The StorageGRID Appliance Installer Home page appears.

2. In the **Primary Admin Node connection** section, determine whether you need to specify the IP address for the primary Admin Node.

If you have previously installed other nodes in this data center, the StorageGRID Appliance Installer can discover this IP address automatically, assuming the primary Admin Node, or at least one other grid node with ADMIN_IP configured, is present on the same subnet.

3. If this IP address is not shown or you need to change it, specify the address:

Option	Description
Manual IP entry	<ol style="list-style-type: none"> a. Unselect the Enable Admin Node discovery check box. b. Enter the IP address manually. c. Click Save. d. Wait for the connection state for the new IP address to become ready.
Automatic discovery of all connected primary Admin Nodes	<ol style="list-style-type: none"> a. Select the Enable Admin Node discovery check box. b. Wait for the list of discovered IP addresses to be displayed. c. Select the primary Admin Node for the grid where this appliance Storage Node will be deployed. d. Click Save. e. Wait for the connection state for the new IP address to become ready.

4. In the **Node name** field, enter the name you want to use for this Storage Node, and click **Save**.

The node name is assigned to this Storage Node in the StorageGRID system. It is shown on the Grid Nodes page in the Grid Manager. If required, you can change the name when you approve the node.

5. In the **Installation** section, confirm that the current state is “Ready to start installation of *node name* into grid with Primary Admin Node *admin_ip*” and that the **Start Installation** button is enabled.

If the **Start Installation** button is not enabled, you might need to change the network configuration or port settings. For instructions, see the installation and maintenance instructions for your appliance.

6. From the StorageGRID Appliance Installer home page, click **Start Installation**.

NetApp® StorageGRID® Appliance Installer

Home
Configure Networking ▾
Configure Hardware ▾
Monitor Installation
Advanced ▾

Home

i The installation is ready to be started. Review the settings below, and then click Start Installation.

Primary Admin Node connection

Enable Admin Node discovery

Primary Admin Node IP

Connection state Connection to 172.16.4.210 ready

Cancel
Save

Node name

Node name

Cancel
Save

Installation

Current state Ready to start installation of NetApp-SGA into grid with Admin Node 172.16.4.210.

Start Installation

The Current state changes to “Installation is in progress,” and the Monitor Installation page is displayed.

Note: If you need to access the Monitor Installation page manually, click **Monitor Installation** from the menu bar.

- If your grid includes multiple StorageGRID appliance Storage Nodes, repeat these steps for each appliance.

Note: If you need to deploy multiple StorageGRID appliance Storage Nodes at one time, you can automate the installation process by using the `configure-sga.py` appliance installation script.

Related tasks

[Automating the configuration and installation of appliance Storage Nodes](#) on page 59

Related references

[How grid nodes discover the primary Admin Node](#) on page 24

Related information

[SG6000 appliance installation and maintenance](#)

[SG5700 appliance installation and maintenance](#)

*SG5600 appliance installation and maintenance***Monitoring StorageGRID appliance installation**

The StorageGRID Appliance Installer provides status until installation is complete. When the software installation is complete, the appliance is rebooted.

Steps

1. To monitor the installation progress, click **Monitor Installation** from the menu bar.

The Monitor Installation page shows the installation progress.

Monitor Installation

1. Configure storage		Running
Step	Progress	Status
Connect to storage controller	<div style="width: 100%; height: 10px; background-color: green;"></div>	Complete
Clear existing configuration	<div style="width: 100%; height: 10px; background-color: green;"></div>	Complete
Configure volumes	<div style="width: 50%; height: 10px; background-color: blue;"></div>	Creating volume StorageGRID-obj-00
Configure host settings	<div style="width: 0%; height: 10px; background-color: blue;"></div>	Pending

2. Install OS	Pending
3. Install StorageGRID	Pending
4. Finalize installation	Pending

The blue status bar indicates which task is currently in progress. Green status bars indicate tasks that have completed successfully.

Note: The installer ensures that tasks completed in a previous install are not re-run. If you are re-running an installation, any tasks that do not need to be re-run are shown with a green status bar and a status of “Skipped.”

2. Review the progress of first two installation stages.


1. Configure storage

During this stage, the installer connects to the storage controller, clears any existing configuration, communicates with SANtricity software to configure volumes, and configures host settings.

2. Install OS

During this stage, the installer copies the base operating system image for StorageGRID from the primary Admin Node to the appliance.

3. Continue monitoring the installation progress until the **Install StorageGRID** stage pauses and a message appears on the embedded console, prompting you to approve this node on the Admin Node using the Grid Manager.



Monitor Installation

1. Configure storage	Complete
2. Install OS	Complete
3. Install StorageGRID	Running
4. Finalize installation	Pending

```

Connected (unencrypted) to: QEMU
/platform.type: Device or resource busy
[2017-07-31T22:09:12.362566] INFO -- [INSG] NOTICE: seeding /var/local with c
ontainer data
[2017-07-31T22:09:12.366205] INFO -- [INSG] Fixing permissions
[2017-07-31T22:09:12.369633] INFO -- [INSG] Enabling syslog
[2017-07-31T22:09:12.511533] INFO -- [INSG] Stopping system logging: syslog-n
g.
[2017-07-31T22:09:12.570096] INFO -- [INSG] Starting system logging: syslog-n
g.
[2017-07-31T22:09:12.576360] INFO -- [INSG] Beginning negotiation for downloa
d of node configuration
[2017-07-31T22:09:12.581363] INFO -- [INSG]
[2017-07-31T22:09:12.585066] INFO -- [INSG]
[2017-07-31T22:09:12.588314] INFO -- [INSG]
[2017-07-31T22:09:12.591851] INFO -- [INSG]
[2017-07-31T22:09:12.594886] INFO -- [INSG]
[2017-07-31T22:09:12.598360] INFO -- [INSG]
[2017-07-31T22:09:12.601324] INFO -- [INSG]
[2017-07-31T22:09:12.604759] INFO -- [INSG]
[2017-07-31T22:09:12.607800] INFO -- [INSG]
[2017-07-31T22:09:12.610985] INFO -- [INSG]
[2017-07-31T22:09:12.614597] INFO -- [INSG]
[2017-07-31T22:09:12.618282] INFO -- [INSG] Please approve this node on the A
dmin Node GMI to proceed...

```

4. Go to the Grid Manager, approve all pending grid nodes, and complete the StorageGRID installation process.

When you click **Install** from the Grid Manager, stage 3 completes and stage 4, Finalize Installation, begins. When stage 4 completes, the controller is rebooted.

Related tasks

[Configuring the grid and completing installation](#) on page 36

Configuring the grid and completing installation

You complete installation by configuring the StorageGRID grid from the Grid Manager on the primary Admin Node.

Steps

1. [Navigating to the Grid Manager](#) on page 36
2. [Specifying the StorageGRID license information](#) on page 37
3. [Adding sites](#) on page 37
4. [Specifying Grid Network subnets](#) on page 38
5. [Approving pending grid nodes](#) on page 39
6. [Specifying Network Time Protocol server information](#) on page 43
7. [Specifying Domain Name System server information](#) on page 44
8. [Specifying the StorageGRID system passwords](#) on page 45
9. [Reviewing your configuration and completing installation](#) on page 46

Navigating to the Grid Manager

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

Before you begin

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

Steps

1. Open your web browser and navigate to the following address:

`https://primary_admin_node_ip`

Note: 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 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

Specifying 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**. The name is displayed as the top level in the grid topology tree after installation.

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.

Note: 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.

The screenshot shows the NetApp StorageGRID installation wizard interface. At the top, there is a blue header with the text "NetApp® StorageGRID®" and a "Help" link. Below the header is a navigation bar with an "Install" button. A progress indicator shows eight steps: 1. License (highlighted), 2. Sites, 3. Grid Network, 4. Grid Nodes, 5. NTP, 6. DNS, 7. Passwords, and 8. Summary. The main content area is titled "License" and contains the instruction: "Enter a grid name and upload the license file provided by NetApp for your StorageGRID system." Below this instruction are four input fields: "Grid Name" with the value "Grid1", "New License File" with a "Browse" button, "License Serial Number" with the value "950719", and "Storage Capacity (TB)" with the value "240".

3. Click **Next**.

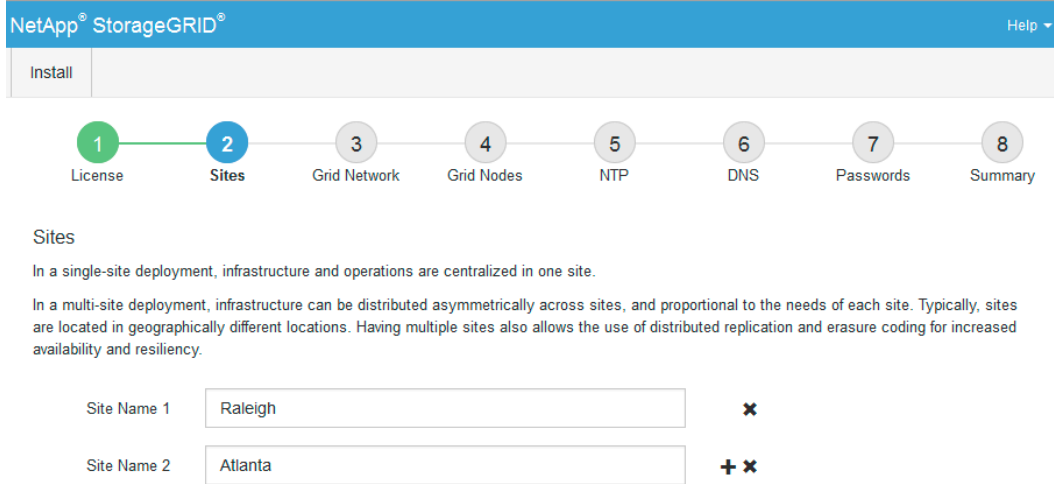
Adding sites

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

Steps

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

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



3. Click **Next**.

Specifying 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 (for example, the subnets hosting your NTP servers).

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

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.



3. Click **Next**.

Approving pending grid nodes

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

Before you begin

All virtual and StorageGRID appliance grid nodes must have been deployed.

Steps

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

Note: 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

<input type="radio"/>	Grid Network MAC Address <i>↑</i>	Name <i>↑</i>	Type <i>↑</i>	Platform <i>↑</i>	Grid Network IPv4 Address <i>↓</i>
<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

<input type="radio"/>	Grid Network MAC Address <i>↑</i>	Name <i>↑</i>	Site <i>↑</i>	Type <i>↑</i>	Platform <i>↑</i>	Grid Network IPv4 Address <i>↓</i>
<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

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"/> x
	<input type="text" value="172.19.0.0/16"/> x
	<input type="text" value="172.21.0.0/16"/> + x

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 host 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 during node deployment, but you can change the name here as required.
- **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, API Gateway Nodes, and any grid nodes that have non-static IP addresses. All other grid nodes are assigned the Client role.

Attention: 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:** For Storage Nodes, whether the selected node will run the Administrative Domain Controller service. Select **Automatic** to have this option applied automatically by the system as required, or select **Yes** or **No** to explicitly set this option for the grid node. For example, you might need to select **Yes** if you want to have more than three ADC services at a site.

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: 192.168.1.234/21
- **Gateway:** The Grid Network gateway. For example: 192.168.0.1

Note: The gateway is required if there are multiple grid subnets.

Note: 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.

Note: 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. Instead, you must follow these steps:

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

Note: 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. 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.

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.

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.

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

Specifying 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 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.

Note: 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.

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

The screenshot shows the NetApp StorageGRID installation wizard interface. At the top, there is a blue header with the text "NetApp StorageGRID" and a "Help" dropdown menu. Below the header is a navigation bar with the word "Install" and a progress indicator consisting of eight numbered circles (1-8) connected by a line. The circles are labeled: 1 License, 2 Sites, 3 Grid Network, 4 Grid Nodes, 5 NTP (highlighted in blue), 6 DNS, 7 Passwords, and 8 Summary. Below the progress bar, the section is titled "Network Time Protocol" and contains the instruction: "Enter the IP addresses for at least four Network Time Protocol (NTP) servers, so that operations performed on separate servers are kept in sync." There are four input fields labeled "Server 1" through "Server 4". Server 1 contains "10.60.248.183", Server 2 contains "10.227.204.142", Server 3 contains "10.235.48.111", and Server 4 contains "0.0.0.0". A plus sign (+) is located to the right of the Server 4 input field.

3. Click **Next**.

Specifying 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.

Attention: Provide two to six IP 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 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 IP address for at least one DNS server in the **Server 1** text box.
2. If necessary, click 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

Domain Name Service

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

Server 1 ✕

Server 2 + ✕

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

3. Click **Next**.

Related information

[Recovery and maintenance](#)

Specifying the StorageGRID system passwords

You need to enter the passwords to use to secure your StorageGRID system.

Steps

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

Attention: 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**.

Reviewing your configuration and completing installation

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

Steps

1. View the **Summary** page.



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**.

Note: 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 “Network installation and provisioning” 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.

Attention: 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.

i 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 is displayed indicating 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
dc1-adm1	Site1	172.16.4.215/21	<div style="width: 100%; background-color: #0070c0;"></div>	Starting services
dc1-g1	Site1	172.16.4.216/21	<div style="width: 100%; background-color: #0070c0;"></div>	Complete
dc1-s1	Site1	172.16.4.217/21	<div style="width: 80%; background-color: #0070c0;"></div>	Waiting for Dynamic IP Service peers
dc1-s2	Site1	172.16.4.218/21	<div style="width: 20%; background-color: #0070c0;"></div>	Downloading hotfix from primary Admin if needed
dc1-s3	Site1	172.16.4.219/21	<div style="width: 20%; background-color: #0070c0;"></div>	Downloading hotfix from primary Admin if needed

Once the installation status for all grid nodes reaches 100%, the installation is complete, and the sign-in page for the Grid Manager is displayed.

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

Automating the installation

You can automate the deployment of VMware virtual grid nodes, the configuration of grid nodes, and the configuration of StorageGRID appliances.

Steps

1. [Automating grid node deployment in VMware vSphere](#) on page 49
2. [Automating the configuration of StorageGRID](#) on page 58
3. [Automating the configuration and installation of appliance Storage Nodes](#) on page 59

Automating grid node deployment in VMware vSphere

You can automate the deployment of StorageGRID grid nodes in VMware vSphere.

Before you begin

- 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`

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

Note: 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 “Collecting 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.

Filename	Description
vsphere-gateway.ovf vsphere-gateway.mf	The template file (.ovf) and manifest file (.mf) for deploying API Gateway Nodes.
vsphere-storage.ovf vsphere-storage.mf	The template file (.ovf) and manifest file (.mf) for deploying virtual machine-based Storage Nodes.
deploy-vsphere-ovftool.sh	The Bash shell script used to automate the deployment of virtual grid nodes.
deploy-vsphere-ovftool-sample.ini	The sample configuration file for use with the deploy-vsphere-ovftool.sh script.

Steps

1. [Defining the configuration file for your deployment](#) on page 50
2. [Running the Bash script](#) on page 57

Related references

[Collecting information about your deployment environment](#) on page 23

Defining 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.

Related concepts

[Configuration file settings](#) on page 50

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=thin --
datastore='datastore_name'
```

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

Attention: 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 68 and 65535. If omitted, 1400 is used. If all or

most of the nodes use the same MTU for the Grid 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:

```
GRID_NETWORK_MTU = 8192
```

- **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 68 and 65535. If omitted, 1400 is used. 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 68 and 65535. If omitted, 1400 is used. 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 client communications. Remapping ports is necessary if enterprise networking policies restrict one or more ports used by StorageGRID. See “Internal grid node communications” or “External client communications” for the list of ports used by StorageGRID.

Note: 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.

The format used is: *<network type>/<protocol>/<new inbound port>/<default inbound port used by grid node>*, 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. See “How grid nodes discover the primary Admin Node” for more information.
 - Note:** 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
```

Note: 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
```

API Gateway Node

The following additional setting is required for API Gateway Nodes:

- **NODE_TYPE:** VM_API_Gateway

This example entry is for an example API 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
```

Note: For Admin nodes, INSTANCES must always equal 2.

Related concepts

[Internal grid node communications](#) on page 15

[External communications](#) on page 17

Related references

[How grid nodes discover the primary Admin Node](#) on page 24

Running 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.

Before you begin

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

About this task

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
```

Result

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

Automating the configuration of StorageGRID

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

Before you begin

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

Filename	Description
<code>configure-storagegrid.py</code>	Python script used to automate the configuration
<code>configure-storagegrid.sample.json</code>	Sample configuration file for use with the script
<code>configure-storagegrid.blank.json</code>	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`).

About this task

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

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

Steps

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

For example:

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

where *platform* is *debs*, *rpms*, 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 from which 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.

Attention: 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 concepts

[Overview of installation REST APIs](#) on page 62

Related tasks

[Navigating to the Grid Manager](#) on page 36

Automating the configuration and installation of appliance Storage Nodes

You can use the `configure-sga.py` script to configure and install StorageGRID appliance Storage Nodes.

Before you begin

- The appliance has been installed in a rack, connected to your networks, and powered on.
- Network links and IP addresses have been configured for the appliance using the StorageGRID Appliance Installer.
- The primary Admin Node for the StorageGRID grid has been deployed, and you know its IP address.

- All Grid Network subnets listed on the IP Configuration page of the StorageGRID Appliance Installer have been defined in the Grid Network Subnet List on the primary Admin Node.

See the installation and maintenance instructions for your appliance.

- You know the location of the `configure-sga.py` file. The file is included in the installation archive, or you can access it by clicking **Help > Appliance Installation Script** from the StorageGRID Appliance Installer.

Steps

1. Log in to the Linux machine you are using to run the Python script.
2. For help with the script syntax and to see a list of the available parameters, enter the following:

```
configure-sga.py --help
```

The `configure-sga.py` script uses four subcommands: `configure`, `install`, `monitor`, and `reboot`. To get help for each subcommand, enter the following:

```
configure-sga.py subcommand --help
```

3. To confirm the current configuration of the appliance Storage Node, enter the following:

```
configure-sga.py configure SGA-INSTALL-IP
```

The results show current IP information for the appliance, including the IP address of the primary Admin Node and information about the Admin, Grid, and Client Networks.

```
Connecting to https://10.224.2.32:8443 (Checking version and
connectivity)
2017/07/30 21:56:34: Performing GET on /api/versions... Received 200
2017/07/30 21:56:34: Performing GET on /api/v2/system-info...
Received 200
2017/07/30 21:56:34: Performing GET on /api/v2/admin-connection...
Received 200
2017/07/30 21:56:34: Performing GET on /api/v2/link-config...
Received 200
2017/07/30 21:56:34: Performing GET on /api/v2/networks... Received
200
2017/07/30 21:56:35: Performing GET on /api/v2/system-config...
Received 200
```

```
StorageGRID Webscale Appliance
Name: MS-SGA1
```

```
StorageGRID primary Admin Node
IP: 172.16.1.170
State: ready
Message: Connection validated at 2017-07-31T01:55:52.982161
```

```
Bond and VLAN configuration
Port bond mode: FIXED
Link speed: 10GBE

Admin Network: ENABLED
Bonding mode: no-bond
VLAN: novlan

Grid Network: ENABLED
Bonding mode: lacp
VLAN: novlan

Client Network: ENABLED
```

```

Bonding mode:   lacp
VLAN:          novlan

Admin Network
CIDR:          10.224.2.32/21 (Static)
MAC:           00:80:E5:43:AE:BC
Gateway:       10.224.0.1
Subnets:      192.168.8.0/24
                172.19.0.0/16
                10.0.0.0/8
                192.168.9.0/24
                172.21.0.0/16

Grid Network
CIDR:          172.16.2.32/21 (Static)
MAC:           00:A0:98:A1:85:AE
Gateway:       172.16.0.1
Subnets:      172.18.0.0/21
                172.17.0.0/21
                192.168.0.0/21

Client Network
CIDR:          47.47.2.32/21 (Static)
MAC:           00:A0:98:A1:85:AD
Gateway:       47.47.0.1

```

```

#####
#####   If you are satisfied with this configuration,   #####
##### execute the script with the "install" sub-command. #####
#####

```

4. If you need to change any of the values in the current configuration, use the `configure` subcommand to update them. For example, if you want to change the IP address for the primary Admin Node to 10.224.2.99, enter the following:

```
configure-sga.py configure --admin-ip 10.224.2.99 SGA-INSTALL-IP
```

5. When you are satisfied with the configuration, use the `install` and `monitor` subcommands to install the appliance:

```
configure-sga.py install --monitor SGA-INSTALL-IP
```

6. If you are experiencing difficulties trying to configure the appliance and want to try rebooting it, enter the following:

```
configure-sga.py reboot SGA-INSTALL-IP
```

Related information

[SG6000 appliance installation and maintenance](#)

[SG5700 appliance installation and maintenance](#)

[SG5600 appliance installation and maintenance](#)

Overview of installation REST APIs

StorageGRID provides two REST APIs for performing installation tasks: the StorageGRID Installation API and the StorageGRID Appliance Installer API.

Both APIs use 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.

Attention: Any API operations you perform using the API Docs (Swagger) user interface are live operations. Be careful not to create, update, or delete configuration 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 Docs** 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.
- **sites** – Site-level configuration operations. You can create, view, delete, and modify a site.

StorageGRID Appliance Installer API

The StorageGRID Appliance Installer API can be accessed over HTTPS from *Controller_IP*: 8443.

To access the API documentation, go to the StorageGRID Appliance Installer on the appliance and select **Help > API Docs** from the menu bar.

The StorageGRID Appliance Installer API includes the following sections:

- **installation** – Operations for starting the appliance installation and for monitoring installation status.
- **networking** – Operations related to the Grid, Admin, and Client Network configuration for a StorageGRID appliance and appliance port settings.

- **setup** – Operations to help with initial appliance installation setup including requests to get information about the system and update the primary Admin Node IP.
- **support** – Operations for rebooting the controller and getting logs.
- **upgrade** – Operations related to upgrading appliance firmware.

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

Optional tasks

- Configure mailing lists and contents of notification emails for system alarms and service status changes.
- Configure the settings for the Simple Network Management Protocol (SNMP) agent used to monitor system status.
- 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.

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

Related information

[Administering StorageGRID](#)

[Implementing S3 client applications](#)

[Implementing Swift client applications](#)

[SG6000 appliance installation and maintenance](#)

[SG5700 appliance installation and maintenance](#)

[SG5600 appliance installation and maintenance](#)

[Recovery and maintenance](#)

[Upgrading StorageGRID](#)

Troubleshooting

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)

For help with installation problems, see instructions for troubleshooting and appliance installation and maintenance, or contact technical support.

Related information

[*Troubleshooting StorageGRID*](#)

[*SG6000 appliance installation and maintenance*](#)

[*SG5700 appliance installation and maintenance*](#)

[*SG5600 appliance installation and maintenance*](#)

[*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

Copyright © 2019 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.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, 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).

Trademark

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

How to send comments about documentation and receive update notifications

You can help us to improve the quality of our documentation by sending us your feedback. You can receive automatic notification when production-level (GA/FCS) documentation is initially released or important changes are made to existing production-level documents.

If you have suggestions for improving this document, send us your comments by email.

[*doccomments@netapp.com*](mailto:doccomments@netapp.com)

To help us direct your comments to the correct division, include in the subject line the product name, version, and operating system.

If you want to be notified automatically when production-level documentation is released or important changes are made to existing production-level documents, follow Twitter account @NetAppDoc.

You can also contact us in the following ways:

- NetApp, Inc., 1395 Crossman Ave., Sunnyvale, CA 94089 U.S.
- Telephone: +1 (408) 822-6000
- Fax: +1 (408) 822-4501
- Support telephone: +1 (888) 463-8277