StorageGRID® 9.0
Installation Guide

NetApp, Inc.
495 East Java Drive
Sunnyvale, CA 94089 U.S.A.
Telephone: +1 (408) 822-6000
Fax: +1 (408) 822-4501
Support telephone: +1 (888) 463-8277
Web: http://www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-06991_C0
July 2015

Copyright and trademark information

Copyright information

Copyright © 1994-2013 NetApp, Inc. All rights reserved. Printed in the U.S.A.

No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

Trademark information

NetApp, the NetApp logo, Network Appliance, the Network Appliance logo, Akorri, ApplianceWatch, ASUP, AutoSupport, BalancePoint, BalancePoint Predictor, Bycast, Campaign Express, ComplianceClock, Cryptainer, CryptoShred, Data ONTAP, DataFabric, DataFort, Decru, Decru DataFort, DenseStak, Engenio, Engenio logo, E-Stack, FAServer, FastStak, FilerView, FlexCache, FlexClone, FlexPod, FlexScale, FlexShare, FlexSuite, FlexVol, FPolicy, GetSuccessful, gFiler, Go further, faster, Imagine Virtually Anything, Lifetime Key Management, LockVault, Manage ONTAP, MetroCluster, MultiStore, NearStore, NetCache, NOW (NetApp on the Web), Onaro, OnCommand, ONTAPI, OpenKey, PerformanceStak, RAID-DP, ReplicatorX, SANscreen, SANshare, SANtricity, SecureAdmin, SecureShare, Select, Service Builder, Shadow Tape, Simplicity, Simulate ONTAP, SnapCopy, SnapDirector, SnapDrive, SnapFilter, SnapLock, SnapManager, SnapMigrator, SnapMirror, SnapMover, SnapProtect, SnapRestore, Snapshot, SnapSuite, SnapValidator, SnapVault, StorageGRID, StoreVault, the StoreVault logo, SyncMirror, Tech OnTap, The evolution of storage, Topio, vFiler, VFM, Virtual File Manager, VPolicy, WAFL, Web Filer, and XBB are trademarks or registered trademarks of NetApp, Inc. in the United States, other countries, or both.

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. A complete and current list of other IBM trademarks is available on the Web at www.ibm.com/legal/copytrade.shtml.

Apple is a registered trademark and QuickTime is a trademark of Apple, Inc. in the U.S.A. and/or other countries. Microsoft is a registered trademark and Windows Media is a trademark of Microsoft Corporation in the U.S.A. and/or other countries. RealAudio, RealNetworks, RealPlayer, RealSystem, RealText, and RealVideo are registered trademarks and RealMedia, RealProxy, and SureStream are trademarks of RealNetworks, Inc. in the U.S.A. and/or other countries.

All other brands or products are trademarks or registered trademarks of their respective holders and should be treated as such.

NetApp, Inc. is a licensee of the CompactFlash and CF Logo trademarks. NetApp, Inc. NetCache is certified RealSystem compatible.

Contents

1

Overview

A summary of the installation process

Introduction

It is assumed that all hardware is installed and connected before you start the installation procedure. It is also assumed that the grid topology has been designed and specified in conjunction with a NetApp Solutions Engineer to meet deployment requirements.

Intended Audience

This guide is intended for technical personnel trained to install and support the StorageGRID system. An advanced level of computer literacy is required, including knowledge of Linux/Unix command shells, networking, and server hardware setup and configuration.

This guide assumes that you are familiar with the StorageGRID system's general design, configurations and options. It also assumes that you received training in grid provisioning, installation, and integration.

Installation Overview and Checklists

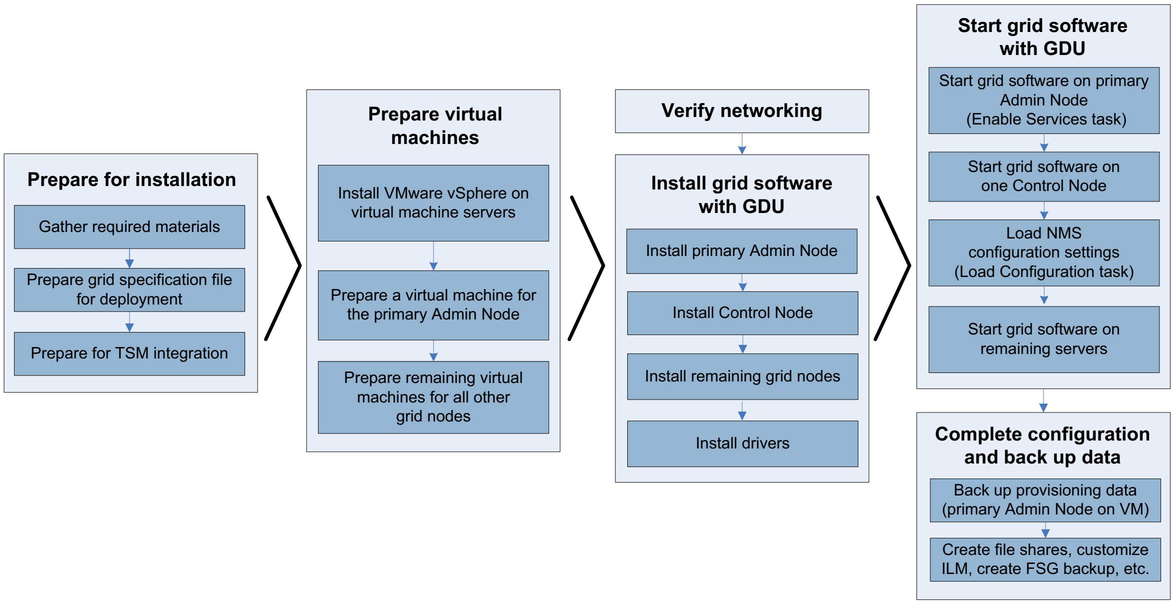
Prepare for installation
Gather required materials
Prepare grid specification file for deployment
Prepare for TSM integration
Prepare virtual machines
Install VMware vSphere on virtual machine servers
Prepare a virtual machine for the primary Admin Node
Prepare remaining virtual machines for all other grid nodes
Verify networking
Install grid software with GDU
Install primary Admin Node
Install Control Node
Install remaining grid nodes
Install drivers
Start grid software with GDU
Start grid software on primary Admin Node (Enable Services task)
Start grid software on one Control Node
Load NMS configuration settings (Load Configuration task)
Start grid software on remaining servers
Complete configuration and back up data
Back up provisioning data (primary Admin Node on VM)
Create file shares, customize ILM, create FSG backup, etc.

Figure 1: Software Installation Overview

The following checklists provide an overview of the software installation process. Follow and complete the following checklists to install StorageGRID:

1. Prepare for installation. See "Prepare for Installation Checklist" on page 10.
2. Complete the "Prepare Virtual Machines Checklist" on page 11.
3. Install, start, and configure grid software. See "Install and Configure Grid Software Checklist" on page 12.

Prepare for Installation Checklist

Complete the following checklist to prepare for the installation of the StorageGRID system.

Table 1: Prepare for Installation Checklist

✓	Step	Task	See
Prepare for Software Installation			
	1.	Read the *Installation Guide* and become familiar with all of the steps and requirements needed to successfully install a grid.	

Table 1: Prepare for Installation Checklist (cont.)

✓	Step	Task	See
	2.	Gather all materials required to perform the installation.	page 15
	3.	If the deployment includes an Archive Node that uses Tivoli® Storage Manager (TSM) middleware to write to archival media, read about TSM configuration, and perform any necessary steps.	page 18
	4.	Prepare the grid specification file for deployment. The grid specification file is an XML file that encapsulates the grid design. It specifies grid topology, grid configuration, and network information. The grid specification file specifies site-specific information such as network IP addresses.	page 22

Prepare Virtual Machines Checklist

Complete the following checklist to prepare virtual machines for the installation of the StorageGRID system.

Table 2: Prepare Virtual Machine Checklist

✓	Step	Task	See
Prepare virtual machine servers			
	1.	Set up virtual machine servers: • Install VMware ESX/ESXi. • Create and start virtual machines. • Configure automatic restart for virtual machines.	page 23
Prepare a virtual machine for the primary Admin Node			
	2.	Create the Provisioning media (floppy image) for installation of the primary Admin Node on a virtual machine. The Provisioning media contains the deployment grid specification file and an activation file required for the installation of Linux on the primary Admin Node.	page 29
	3.	Install Linux on the virtual machine that will host the primary Admin Node.	page 30
	4.	Install VM tools and configure settings on the virtual machine that will host the primary Admin Node.	page 33
	5.	Load installation ISOs onto the primary Admin Node using the `load_cds.py` script. The Grid Deployment Utility (GDU) uses these ISO images to install the grid software.	page 34
	6.	Install provisioning software on the primary Admin Node.	page 37

Table 2: Prepare Virtual Machine Checklist (cont.)

✓	Step	Task	See
	7.	Provision the grid and create the Server Activation floppy image.	page 38
Prepare virtual machines for all other grid nodes			
	8.	Install Linux on the virtual machines that will host the remaining grid nodes. You can perform this operation in parallel on multiple servers.	page 43
	9.	Install VM tools and configure settings on virtual machines that will host grid nodes.	page 46
	10.	Configure Storage Nodes for NFS storage.	page 47

Install and Configure Grid Software Checklist

Complete the following checklist to install StorageGRID software on prepared virtual machines.

Table 3: Install Grid Software Checklist

✓	Step	Task	See
Install grid software			
	1.	Verify networking to confirm that the primary Admin Node can communicate with the other virtual machines or servers in the grid.	
	2.	Use GDU to install grid software on the virtual machine that will hosts the primary Admin Node. Do not start grid software yet.	page 51
	3.	Use GDU to install grid software on the virtual machine that will hosts the Control Node. Do not start grid software yet.	page 51
	4.	Use GDU to install grid software on the virtual machines that will host the remaining grid nodes. You can install grid software in parallel on the remaining virtual machines.	page 51
Start grid software			
	5.	Use GDU to start grid software on the primary Admin Node.	page 54
	6.	Use GDU to start grid software on one Control Node.	page 54
	7.	Use GDU to load NMS configuration settings on the primary Admin Node. Then log in to the NMS management interface (MI) to start monitoring the grid.	page 55
	8.	Use GDU to start grid software on the other virtual machines.	page 57

Table 3: Install Grid Software Checklist (cont.)

✓	Step	Task	See
Complete integration and configuration			
	9.	Back up the provisioning data to two safe, secure locations. **WARNING This step is vital. Do not skip this step.**	Appendix A
	10.	Configure the grid.	page 61
	11.	Configure file shares. You must create file system shares for the clients that will access the grid. Both CIFS (Windows Workgroup or Active Directory) and NFS are supported.	*Administrator Guide*
	12.	Customize Gateway Node behavior. Consult the deployment's configuration information to determine if you need to customize the behavior of a Gateway Node, for example configure content protection or cache parameters. Customization is achieved by configuring File System Gateway (FSG) settings for replication groups. You may also have to create and edit FSG profiles.	*Administrator Guide*
	13.	Verify client and grid integration.	page 75

About Licensing Agreements

To complete aStorageGRID software installation, you must accept a license agreement. After installation is complete, the text of the license agreement is available on each server at /var/local/.license-accepted.txt.

About VMware vSphere

For the current supported versions of VMware software, see the Interoperability Matrix Tool (IMT).

Configuring Virtual Disks

StorageGRID software installation is optimized to align partitions to 4K boundaries. No special action is required when creating virtual disks in VMware for use with StorageGRID software. If using Fiber Channel for your VMware datastores, select "LUN type" for all attached drives.

2

Prepare for Installation

Gather materials, prepare the grid specification file for deployment, and prepare for TSM integration

Gather Materials

Gather the materials listed in Table 4 to prepare for the installation.

Table 4: Materials Checklist

✓	Item	Notes
	Default grid specification file	Use Grid Designer to deploy the default grid specification file, and update it with data specific to the grid. For more information on deployment, see the *Grid Designer User Guide*.
	provisioning-autoinst-<*dn*>.xml file	Used during Linux installation to customize the primary Admin Node for its role in the grid. <*dn*> is the device name of the system drive of the server. Located in the provisioning directory of the StorageGRID Software CD or the StorageGRID Software Service Pack CD. If a service pack is available, use the copy on the service pack CD.
	SUSE Linux Enterprise Server (SLES) DVD	Use only the supported version of SLES for the StorageGRID 9.0 system. For supported versions, see the Interoperability Matrix Tool (IMT). **NOTE** Use of any version of SLES other than the correct version will result in an installation failure. Since you can load Linux on multiple servers in parallel, it is helpful to have more than one copy of SLES DVD.

Table 4: Materials Checklist (cont.)

✓	Item	Notes
	StorageGRID Software CD	Version 9.0.0 Confirm that the version matches the grid activation information. Used to install the base version of the grid software on all grid servers. Works in conjunction with the Server Activation USB media that customizes each server to prepare it for its assigned role in the grid.
	Enablement Layer for StorageGRID Software CD	Version 9.0.0 Used to customize the Linux operating system installed on each grid server. Only the packages needed to support the services hosted on the server are retained to minimize the overall footprint occupied by the operating system and maximize the security of each server.
	If available, Service Pack CDs	The StorageGRID Service Pack consists of two CDs: • StorageGRID Software Service Pack CD version 9.0.x • Enablement Layer for StorageGRID Software Service Pack CD version 9.0.x where x is the service pack number. The service pack number is identical for both CDs. A service pack is the collection of fixes and enhancements since the release of 9.0. Service packs are cumulative. For example, Service Pack 9.0.2 includes the contents of Service Pack 9.0.1.
	Custom files	If applicable, custom files noted in the <custom-files> sections of the grid specification file.
	If the deployment includes an Archive Node that uses Tivoli® Storage Manager (TSM) middleware to write to archival media, the CD containing the required TSM Client packages.	The following Tivoli Storage Manager (TSM) Client packages are required: • Backup Archive RPM (TIVsm-BA.i386.rpm) • Tivoli API RPM (TIVsm-API.i386.rpm) • API 64-bit RPM (TIVsm-API64.i386.rpm) — for 64-bit Archive Nodes only For more information, see "TSM Client Packages" on page 21. Information on the supported version of the TSM packages is available in the Interoperability Matrix Tool (IMT).
	VMware® software and documentation	For the current supported versions of VMware software, see the Interoperability Matrix Tool (IMT).

Table 4: Materials Checklist (cont.)

✓	Item	Notes
	Service laptop	Laptop must have: • Microsoft® Windows® operating system • Network port • Supported browser for StorageGRID 9.0 • Telnet/ssh client (PuTTY version 0.57 or higher) • WinImage, used to create the floppy disk images required to install grid nodes in a virtual machine. WinImage is available at http://www.winimage.com) • Grid Designer **NOTE** Each version of Grid Designer supports a specific version of StorageGRID software. Ensure you have the correct version of Grid Designer. • WinSCP: if installing a primary Admin Node on a virtual machine, use a tool such as WinSCP to transfer files to and from the Admin Node. WinSCP can be downloaded from: http://winscp.net/eng/download.php
	StorageGRID documentation	*Administrator Guide* *Grid Designer User Guide* *Release Notes*

Prepare Grid Specification File for Deployment

Prepare the grid specification file for deployment. That is, use Grid Designer to customize the grid specification file by entering networking and other information specific to the installation.

Prerequisites

The following materials and information are available:

- Grid specification file
- provisioning-autoinst-<*dn*>.xml file

- Deployment-specific data such as networking IP addresses, NTP server IP addresses, and server names
- Grid Designer software
- Service laptop

Procedure

1. Update the grid documentation.

   Before deploying the grid, note all deployment-specific data such as networking IP addresses, NTP server IP addresses, and server names. This information is used for many integration, maintenance, and expansion procedures.
2. Use Grid Designer to update the grid specification file with site-specific information. For more information, see the *Grid Designer User Guide*.

   Grid Designer outputs OVF files used to automate the creation of virtual machines using VMware vCenter Server.

Prepare for TSM Integration

Tivoli Storage Manager (TSM) writes to archival media. You can integrate Archive Nodes that use TSM with either a new or an existing TSM server.

For background information on the Archive Node, see the “Archive Storage” chapter in the *Administrator Guide*.

NOTE You cannot co-host an Archive Node with a TSM server.

This section includes best practices for integrating the Archive Node with a TSM server before beginning the integration. This section documents details of Archive Node operation that impact the configuration of the TSM Server. It also provides required and recommended TSM Server settings.

Subsequent sections of this guide (“Configure the TSM Server” on page 63) include detailed sample instructions for preparing a TSM Server that follow these recommendations.

Best Practices for TSM Integration

Overview of Archive Node Configuration and Operation

After an object is ingested into the grid (via a Gateway Node, the StorageGRID API (SGAPI), or CDMI), copies are made in the required locations, including Archive Nodes, based on the grid's ILM rules. The grid manages the Archive Node as a location where objects are stored indefinitely and are always accessible.

The TSM Archive Node acts as a client to a TSM server, and the TSM client libraries are installed on the Archive Node by the StorageGRID software installation process. Objects directed to the Archive Node for storage are saved directly to the TSM as they are received. The Archive Node does not stage objects before saving them to the TSM, nor does it perform object aggregation. However, the Archive Node may submit multiple objects to the TSM in a single transaction when data rates warrant.

After the Archive Node saves an object to the TSM, the object is managed by the TSM using its lifecycle/retention policies. These retention policies must be defined to be compatible with the operation of the Archive Node. (That is, the objects saved by the Archive Node must be stored indefinitely and must always be accessible by the Archive Node, unless they are purged by the Archive Node.)

There is no connection between the grid's ILM rules and the TSM server's lifecycle/retention policies. Each operate independently of one another. However, as each object is ingested into the grid, you can assign it a TSM management class. This management class is passed to the TSM along with the object. Assigning different management classes to different types of objects permits you to configure the TSM to place objects in different storage pools, or to apply different migration or retention policies as required. For example, objects identified as database backups (temporary content than can be overwritten with newer data) could be treated differently than application data (fixed content that must be retained indefinitely).

The Archive Node can be integrated with a new or an existing TSM server.

The Archive Node does not require a dedicated TSM server. TSM servers may be shared with other clients, provided that the TSM server is sized appropriately for the maximum expected load.

It is possible to configure more than one Archive Node to write to the same TSM server; however, this configuration is *only* recommended if the Archive Nodes write different sets of data to the TSM. Configuring more than one Archive Node to write to the same TSM is *not* recommended when each Archive Node writes copies of the same objects to the archive. In the latter scenario, both copies are subject to a single point of failure (the TSM server) for what are supposed to be independent, redundant copies of data.

Archive Nodes do not make use of the Hierarchical Storage Management (HSM) component of TSM.

TSM Configuration: Best Practices

When sizing and configuring the TSM server, you should be aware of the following:

- Because the Archive Node does not aggregate objects before saving them to the TSM, the TSM database must be sized to hold references to all objects that will be written to the Archive Node.
- Archive Node software cannot tolerate the latency involved in writing objects directly to tape or other removable media. Therefore, the TSM server must be configured with a disk storage pool for the initial storage of data saved by the Archive Node whenever removable media are used.
- You must configure TSM retention policies to use event-based retention. The Archive Node does not support creation-based TSM retention policies. We recommend setting `retmin=0` and `retver=0` in the retention policy (which indicates that retention begins when the Archive Node triggers a retention event, and is retained for 0 days after that). However, these values of `retmin` and `retver` are optional.

The disk pool must be configured to migrate data to the tape pool (that is, the tape pool must be the `NXTSTGPOOL` of the disk pool). The tape pool must *not* be configured as a copy pool of the disk pool with simultaneous write to both pools (that is, the tape pool cannot be a `COPYSTGPOOL` for the disk pool). To create offline copies of the tapes containing Archive Node data, configure the TSM with a second tape pool that is a copy pool of the tape pool used for Archive Node data.

Settings for the Archive Node are made in the NMS MI, as directed during the installation (as described on “Configure the TSM Archive Node” on page 69).

TSM Client Packages

When installing software on the Archive Node server, you are required to insert the TSM Client packages CD.

Download the TSM client packages from IBM. These packages are included in a .tar file available at:

ftp://service.software.ibm.com/storage/tivoli-storage-management/maintenance/client/

Information on the supported version of the TSM packages is available in the Interoperability Matrix Tool (IMT).

Create CD

1. Download the LinuxX86 version of the .tar file corresponding to the version of the TSM Library in use at the deployment. The .tar file is found at <*version*> ▶ Linux ▶ LinuxX86 ▶ <*version*>.
2. Unpack the .tar file.
3. Copy the .rpm files to a CD using standard CD burning software.

This CD is specific to a particular release of TSM Library Software. For information on the correct version to use, see the site's configuration information.

Reference Materials for TSM Servers

You may find the following reference materials useful in helping you prepare your TSM for integration with the Archive Node in aStorageGRID system:

- *IBM Tape Device Drivers Installation and User's Guide*

  http://www-01.ibm.com/support/docview.wss?rs=577&uid=ssg1S7002972
- *IBM Tape Device Drivers Programming Reference*

  http://www-01.ibm.com/support/docview.wss?rs=577&uid=ssg1S7003032
- *Administrator's Reference for Tivoli Storage Manager for Linux Server:* Storage Manager for Linux Server ▶ Administrator's Reference ▶ Chapter 2: Administrative Commands

  http://publib.boulder.ibm.com/infocenter/tivihelp/v1r1/index.jsp

Next Steps

Your next step is to prepare virtual machines. See Chapter 3: “Prepare Virtual Machines”.

3

Prepare Virtual Machines

How to prepare virtual machines for the installation of StorageGRID software

Introduction

To prepare virtual machines for the installation of StorageGRID software, perform the following:

1. "Install VMware vSphere on VM Servers" on page 23
2. "Prepare a Virtual Machine for the Primary Admin Node" on page 29
3. "Prepare Virtual Machines for all Other Grid Nodes" on page 43

Install VMware vSphere on VM Servers

Install and configure VMware vSphere software on all servers that will host virtual machines. These virtual machines will host grid nodes. Perform the steps outlined in Table 5.

Table 5: Install and Configure VMware Software

✓	Step	Action	See
	1.	Install VMware vSphere software.	page 24
	2.	Create the virtual machines.	page 25
	3.	Configure ESX/ESXi for automatic restart.	page 26
	4.	Start the virtual machines.	page 28

NOTE For supported versions of VMware software, see the Interoperability Matrix Tool (IMT).

Install VMware vSphere Software

To install and configure a virtual machine, install and configure VMware vSphere software. In particular, you require VMware ESX/ESXi and VMware vCenter Server software. For the steps required to install these vSphere products see VMware documentation available at: http://www.vmware.com/support/pubs

VMware vCenter Server

Create and configure virtual machines using the Open Virtualization Format (OVF) files produced by Grid Designer:

- one OVF file per VM host — When deployed using vCenter Server, this OVF file creates and configures all virtual machines hosted on a physical server.

  To use these OVF files, each destination host must have a single datastore with enough free space to hold all virtual disks required on this VM host. For information on the size of the virtual disks required on a VM host, see the VM BOM created by Grid Designer.

  After the initial deployment, you can use vCenter to migrate virtual disks to multiple datastores.
- one OVF file per virtual machine — These OVF files create and configure a single virtual machine. For information on the size of the virtual disks required for each virtual machine, consult the VM BOM created by Grid Designer.

VMware ESX/ESXi

You must install VMware ESX/ESXi on a prepared physical server with correctly configured hardware. Grid hardware must be correctly configured (including firmware versions and BIOS settings) before you install VMware software.

Configure networking in the hypervisor as required to support networking for the grid. Note that an HAGC hosted on a virtual machine does not require the use of a crossover cable for heartbeat.

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.

VMware vSphere Client

Install the VMware vSphere client on your service laptop. The VMware vSphere client allows you to connect to vCenter Server, monitor ESX/ESXi servers, and create and configure virtual machines.

Create the Virtual Machines

After you install VMware ESX/ESXi, create one virtual machine for each grid node installed on the server.

NOTE For more information on creating virtual machines, see VMware vSphere documentation.

Create Virtual Machines Using OVF Files

Use this procedure when you are using VMware vCenter Server.

Prerequisites

- OVF files
- VM Bill of Materials (the GID<*grid_ID*>_REV<*revision_number*>_GSPEC_VMBOM.html file) includes information on the virtual machines to be installed on each physical server, including the type of grid node hosted in each one and the resources that each VM requires.
- vSphere client software
- when deploying one OVF file per VM Host:
  - vCenter Server software
  - a datastore large enough for all virtual disks for all virtual machines hosted on the server
- when deploying one OVF file per virtual machine
  - a datastore large enough for the virtual machine’s virtual disks

Procedure

- Connect to vCenter Server using vSphere Client software, and deploy each OVF file to a host system.

  This creates all of the required virtual machines, as described in the VM bill of materials.

Create Virtual Machines Manually

If required, use the *NetApp StorageGRID Deployment Guide* and the information in the VM BOM generated by Grid Designer to manually create virtual machines.

After you create all virtual machines hosted on the server, and you adjust the resource allocations if required, configure the VMs to automatically restart when the ESX/ESXi server restarts. See "Configure ESX/ESXi for Automatic Restart" below.

Configure ESX/ESXi for Automatic Restart

Configure the ESX/ESXi server to automatically restart the virtual machines when the server restarts. Without an automatic restart, the virtual machines and grid nodes remain shut down after VMware ESX/ESXi server restarts.

Prerequisites

- All virtual machines have been created and configured

Procedure

1. In the VMware vSphere client tree, click the root element (that is, the ESX/ESXi server), and then select the **Configuration** tab.

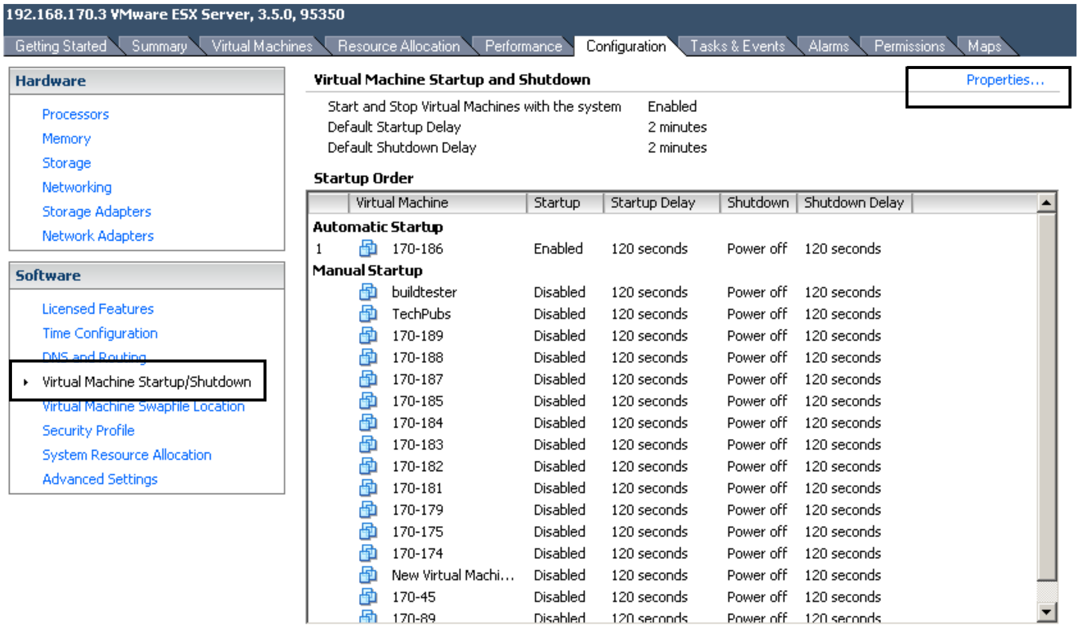

192.168.170.3 VMware ESX Server, 3.5.0, 95350

Virtual Machine Startup and Shutdown

Start and Stop Virtual Machines with the system	Enabled
Default Startup Delay	2 minutes
Default Shutdown Delay	2 minutes

Startup Order

	Virtual Machine	Startup	Startup Delay	Shutdown	Shutdown Delay
Automatic Startup					
1	170-186	Enabled	120 seconds	Power off	120 seconds
Manual Startup					
	buildtester	Disabled	120 seconds	Power off	120 seconds
	TechPubs	Disabled	120 seconds	Power off	120 seconds
	170-189	Disabled	120 seconds	Power off	120 seconds
	170-188	Disabled	120 seconds	Power off	120 seconds
	170-187	Disabled	120 seconds	Power off	120 seconds
	170-185	Disabled	120 seconds	Power off	120 seconds
	170-184	Disabled	120 seconds	Power off	120 seconds
	170-183	Disabled	120 seconds	Power off	120 seconds
	170-182	Disabled	120 seconds	Power off	120 seconds
	170-181	Disabled	120 seconds	Power off	120 seconds
	170-179	Disabled	120 seconds	Power off	120 seconds
	170-175	Disabled	120 seconds	Power off	120 seconds
	170-174	Disabled	120 seconds	Power off	120 seconds
	New Virtual Machi...	Disabled	120 seconds	Power off	120 seconds
	170-45	Disabled	120 seconds	Power off	120 seconds
	170-89	Disabled	120 seconds	Power off	120 seconds

Figure 2: Configuration Pane

2. Under Software, click **Virtual Machine Startup/Shutdown**, and click **Properties**.

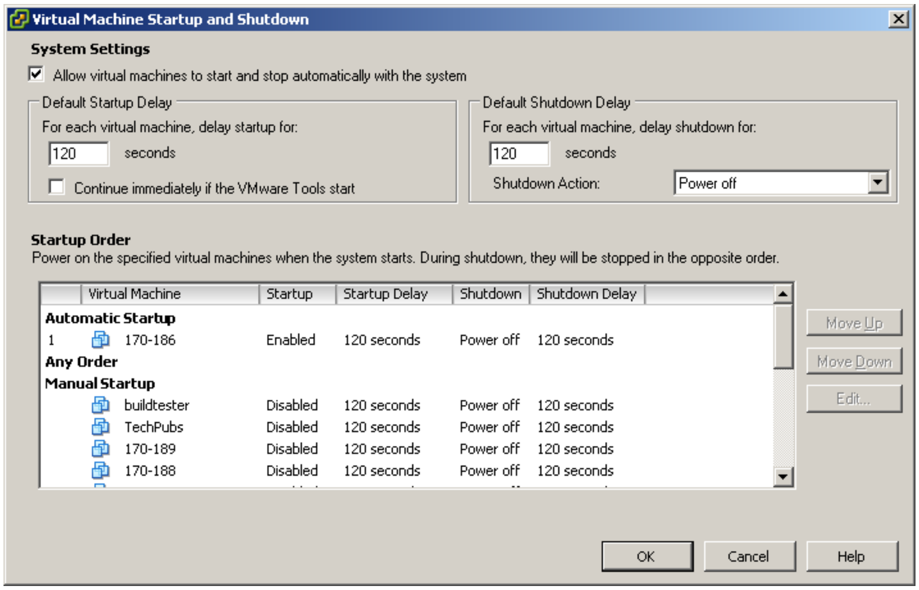

System Settings

☑ Allow virtual machines to start and stop automatically with the system

Default Startup Delay

For each virtual machine, delay startup for: 120 seconds

☐ Continue immediately if the VMware Tools start

Default Shutdown Delay

For each virtual machine, delay shutdown for: 120 seconds

Shutdown Action: Power off

Startup Order

Power on the specified virtual machines when the system starts. During shutdown, they will be stopped in the opposite order.

	Virtual Machine	Startup	Startup Delay	Shutdown	Shutdown Delay
Automatic Startup					
1	170-186	Enabled	120 seconds	Power off	120 seconds
Any Order					
Manual Startup					
	buildtester	Disabled	120 seconds	Power off	120 seconds
	TechPubs	Disabled	120 seconds	Power off	120 seconds
	170-189	Disabled	120 seconds	Power off	120 seconds
	170-188	Disabled	120 seconds	Power off	120 seconds

Figure 3: Virtual Machine Startup and Shutdown Window

3. Under System Settings, select **Allow virtual machines to start and stop automatically with the system**.
4. Under Default Startup Delay, leave the start-up delay time at the default of 120 seconds.

   A delay of 120 seconds gives each virtual machine time to start before the next virtual machine begins the process of starting. This delay allows for a smoother start-up process that will not overload the system.
5. Under Startup Order, in the **Manual Startup** list select any virtual machine that hosts an Admin Node or a Gateway Node, and click **Move up** to move it to the **Automatic Startup** list.

   Place any Admin Node at the top of the list and any Gateway Nodes below it in any order. Admin Nodes and Gateway Nodes are generally NTP primaries. Starting them first helps prevent timing alarms within the grid when the ESX/ESXi server restarts.
6. Move virtual machines that host other grid nodes to the **Any Order** list by clicking **Move up**.

   These virtual machines restart in any order.

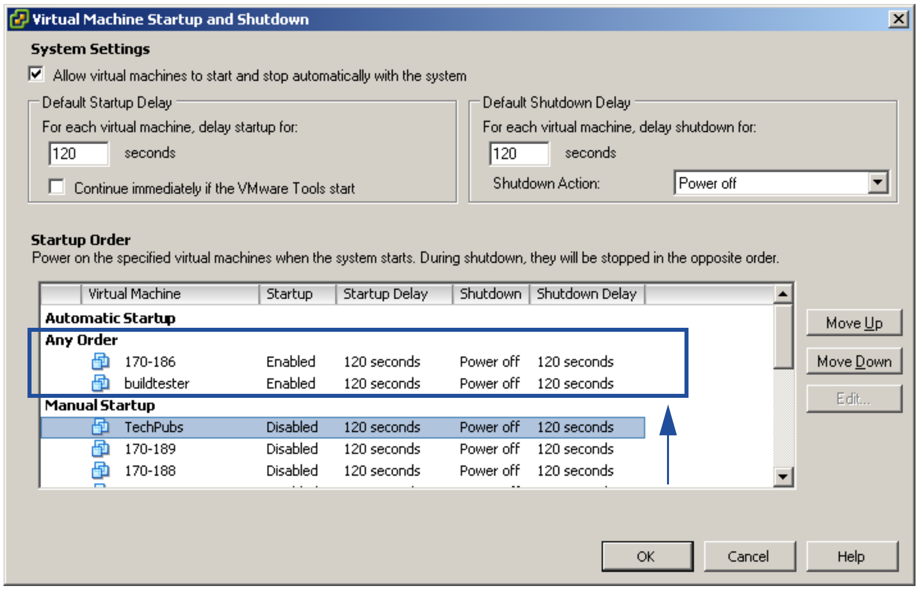

Virtual Machine Startup and Shutdown

System Settings

☑ Allow virtual machines to start and stop automatically with the system

Default Startup Delay

For each virtual machine, delay startup for: 120 seconds

☐ Continue immediately if the VMware Tools start

Default Shutdown Delay

For each virtual machine, delay shutdown for: 120 seconds

Shutdown Action: Power off

Startup Order

Power on the specified virtual machines when the system starts. During shutdown, they will be stopped in the opposite order.

Virtual Machine	Startup	Startup Delay	Shutdown	Shutdown Delay
Automatic Startup				
Any Order				
170-186	Enabled	120 seconds	Power off	120 seconds
buildtester	Enabled	120 seconds	Power off	120 seconds
Manual Startup				
TechPubs	Disabled	120 seconds	Power off	120 seconds
170-189	Disabled	120 seconds	Power off	120 seconds
170-188	Disabled	120 seconds	Power off	120 seconds

Figure 4: Move Virtual Machines to Control Startup Order

7. Click **OK**.

Start the Virtual Machines

Prerequisites

- All virtual machines have been configured for automatic restart.

Procedure

1. In the vSphere client, select the virtual machine.
2. Click the **Console** tab.
3. Click the green **Power On** button ▶.

No operating system is available to the virtual machine at this point, so it attempts to PXE boot and fails with the message “Operating System not found”.

Prepare a Virtual Machine for the Primary Admin Node

Complete the following tasks to prepare a virtual machine for the primary Admin Node.

Table 6: Prepare a Virtual Machine for the Primary Admin Node

✓	Step	Action	See
	1.	Create the provisioning floppy image.	page 29
	2.	Install Linux on the virtual machine for the primary Admin Node.	page 30
	3.	Install VMware Tools and configure VM settings.	page 33
	4.	Load the software distribution.	page 34
	5.	Install the provisioning software.	page 37
	6.	Provision the grid and create a Server Activation USB flash drive.	page 38

Create the Provisioning Floppy Image

A Provisioning floppy image contains site-specific information required to install, integrate, and maintain the StorageGRID system.

Prerequisites

- Utility such as WinImage (available at http://www.winimage.com), that permits you to create a floppy disk image
- Grid specification file GID<*Grid_ID*>_REV1_GSPEC.xml file. See "Prepare Grid Specification File for Deployment" on page 17.
- provisioning-autoinst-<*dn*>.xml file

  where <*dn*> is the device name of the system drive, for example sda. Located in the provisioning directory of the StorageGRID Software CD or the StorageGRID Software Service Pack CD. Use the version on the service pack CD if a service pack is available.
- Service laptop

Procedure

1. Start the WinImage software. From the **File** menu, select **New**.
   - In the Format selection dialog, select a standard format **1.44 MB** floppy, and click **OK**.

2. Create a floppy image that contains the deployment grid specification file and the provisioning-autoinst-<dn>.xml file:
    a. From the **Image** menu, select **Inject**.
    b. Browse for the grid specification file, and select **Open**.
    c. When prompted, confirm that you want to inject the file. Select **Yes**.
    d. Repeat from step **a** for the provisioning-autoinst-<*dn*>.xml file.

       where <*dn*> is the device name of the system drive, for example sda.
3. Save the floppy image:
    a. From the **File** menu, select **Save**.
    b. In the **Save** dialog, browse to the destination folder.
    c. Select Save as type: **Virtual floppy Image (*.vfd,*.flp)**
    d. Enter a file name ending in .flp, such as ***<servername>.flp***.

       You must enter the extension, or the vSphere client cannot use the image during installation.
    e. Click **Save**.

Install Linux on a Virtual Machine for the Primary Admin Node

Use this procedure to install Linux on the virtual machine that will host the primary Admin Node. The installation process completely erases the server drives and installs its own OS, applications, and support files.

Prerequisite

- The virtual machine for the primary Admin Node has been started. See "Start the Virtual Machines" on page 28.
- SLES DVD. For information on the supported version of SLES, see the Interoperability Matrix Tool (IMT).
- Provisioning floppy image. See "Create the Provisioning Floppy Image" on page 29.

Procedure

1. Insert the SLES DVD into the machine from which you are running the vSphere client. Skip this step if you are using an ISO image of Linux.

2. In the VMware vSphere client navigation tree, select the virtual machine.
3. Click the Connect/Disconnect CD/DVD drive to the virtual machine icon, then select **Connect CD/DVD 1 ▶ Connect to <*CD_drive_letter*>**

   — or —

   If you are using an ISO image of the Linux installation, select **Connect CD/DVD 1 ▶ Connect to ISO image on local disk**.
4. Click the **Console** tab.
5. Click anywhere inside the Console pane to enter the Console pane.

   Your mouse pointer disappears.

TIP Press <Ctrl>+<Alt> to release your mouse pointer from the VM console.

---

6. Press **<Ctrl>+<Alt>+<Insert>** to reset the virtual machine. The server performs the following steps:
   - The BIOS runs a hardware verification.
   - By default the system boots from the DVD, and loads the SUSE Linux Enterprise Server Boot Screen in the VMware vSphere client Console pane.

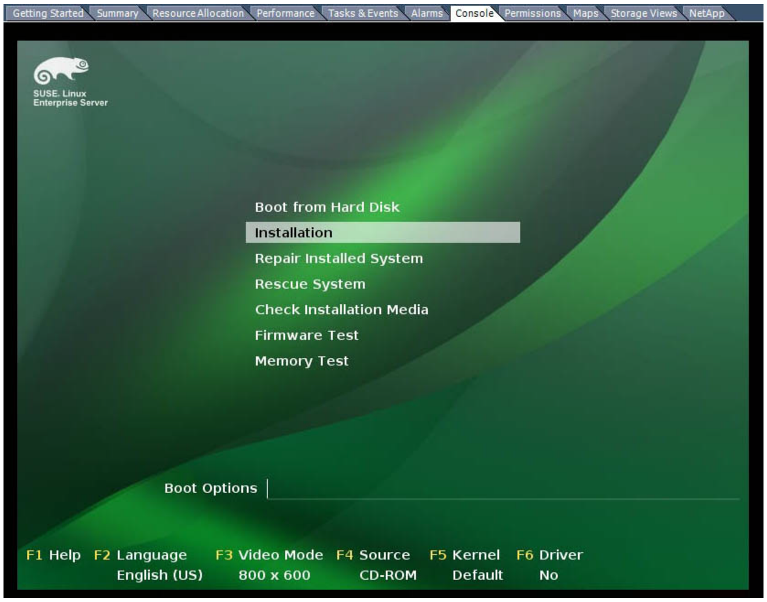

F2 Language English (US)
F3 Video Mode 800 x 600
F4 Source CD-ROM
F5 Kernel Default
F6 Driver No

Figure 5: SLES Boot Screen

7. From the SLES Boot Screen:
   a. Press the keyboard's down arrow and highlight **Installation**. (Do not press **<Enter>**.)

NOTE You must move the cursor to the Installation option within eight seconds. If you do not, SLES will automatically attempt to install from the hard drive and the installation process will fail. If this happens, you must begin the installation process again from the beginning.

b. Press **<Ctrl>+<Alt>** to leave the Console pane.

   The mouse icon disappears.

c. Click **Connect Floppy 1** and select Connect to Floppy Image on local disk.

d. On your service laptop, select the floppy image that contains the activation file for this server.

e. Click anywhere inside the Console pane to return to the Console pane.

f. Press **<Tab>**. At the bottom of the screen, adjacent to the **Boot Options** prompt, enter:

   ```
   autoyast=device://fd0/provisioning-autoinst-<dn>.xml
   ```

   where <*dn*> is the device name of the system drive, for example sda.

NOTE If you do not enter the path to the activation file when required, AutoYaST does not customize the installation for the server. Always enter the path to the activation file.

   If you enter an incorrect value, and are prompted to re-enter the device name and path, check the floppy device name.

g. Press **<Enter>**.

   Wait about one minute while the installer processes the information.

   The base SLES installation completes without further intervention.

8. During this installation process, disconnect the floppy image after it is no longer required. Click the Connect/Disconnect the floppy devices of the virtual machine icon and select **Disconnect Floppy 1 ▶ Disconnect from <*drive*>**.

NOTE If the floppy image is connected when the virtual machine reboots, the message “This is not a bootable disk is displayed.” To continue, disconnect the floppy image and hit any key to continue.

When SLES installation is complete, the server completes its configuration and starts the operating system. Installation is complete when the login prompt appears.

Install VMware Tools and Configure VM Settings

Install VMware Tools on each virtual machine that will host a grid node. For information on the supported versions of VMware software, see the Interoperability Matrix Tool (IMT).

Install VMware Tools from an ISO image that is packaged with vSphere client software.

The vmware_setup.py script configures the video resolution settings that the virtual machine uses on startup. This is necessary because the default resolution is not supported by the VMware video adapter, which results in an error on startup.

Prerequisites

- Linux has been installed on the virtual machine

Procedure

1. In the Virtual Machine tree, right-click the virtual machine, then select **Guest ▶ Install/ Upgrade VMware Tools**.

   The Install VMware Tools dialog appears.
2. Select **Interactive Tool Upgrade** and click **OK**.

   The VMware Tools package is made available to the virtual machine as an ISO at /cdrom.

   Wait until VMware disconnects the Linux CD/DVD.
3. In the VMware vSphere Client, click in the **Console** pane of the virtual machine and log in to the host that will be the new grid node.
4. Copy the VMware Tools packages to the virtual machine:
   a. Mount the ISO image of the VMware Tools CD. Enter:
      ```
      mount /cdrom
      ```
   b. Copy the gzip package from the CD to the virtual machine, and unpack it. Enter:
      ```
      mkdir /tmp/vmtools
      cd /tmp/vmtools
      tar -zxvf /cdrom/VMwareTools-*.tar.gz
      ```

5. Install VMware Tools, accepting the default installation options. Enter:

   ```
   cd /tmp/vmtools/vmware-tools-distrib/
   ./vmware-install.pl --default
   ```

   Wait for the installation to complete. This takes about a minute.

6. Check to make sure that VMware Tools is running. Enter:

   ```
   /etc/init.d/vmware-tools status
   ```

   You will see the message “vmtoolsd is running”.

7. Remove the installation files from the virtual machine. Enter:

   ```
   cd /tmp
   rm -rf vmtools
   ```

8. Run the vmware_setup.py script. Enter: `vmware_setup.py`

   The script completes silently.

9. Reboot to ensure the changes take effect. Enter: `reboot`

Load Installation ISOs

Use the load_cds.py script to load installation ISOs onto the primary Admin Node. For background on the load_cds.py script, see “About load_cds.py” on page 117.

This section contains two procedures:

- “Use load_cds.py with CDs on a Virtual Machine” on page 34
- “Use load_cds.py with ISOs” on page 36

Use load_cds.py with CDs on a Virtual Machine

Follow this procedure if installing the primary Admin Node on a virtual machine.

Prerequisites

- Linux has been installed on the primary Admin Node. See “Install Linux on a Virtual Machine for the Primary Admin Node” on page 30.
- Service laptop running vSphere Client
- The following StorageGRID CDs:
  - StorageGRID 9.0.0 Software CD
  - Enablement Layer for StorageGRID 9.0.0 Software CD
  - If available, StorageGRID 9.0.x Software Service Pack CD

- If available, Enablement Layer for StorageGRID 9.0.x Software Service Pack CD
- Any additional CDs required for the grid, such as TSM client packages CD

Procedure

NOTE The order in which CDs are loaded does not matter.

1. In vSphere Client, click in the console window of the primary Admin Node's virtual machine. Log in as root. When prompted for a password, press **<Enter>**.
2. Insert the latest StorageGRID Software CD in the service laptop.

   If available, use the StorageGRID 9.0.x Software Service Pack CD. Otherwise, use the StorageGRID 9.0.0 Software CD.

TIP Press <Ctrl>+<Alt> to release the mouse pointer from the VM console.

3. Click the Connect/Disconnect CD/DVD drive to the virtual machine icon, then select **Connect CD/DVD 1 ▶ Connect to <*CD_drive_letter*>**
4. Install the load_cds.py script:
   a. Mount the latest CD.

      Enter: `mount /cdrom`
   b. Install the updated load_cds.py script from the CD.

      Enter: `/cdrom/install-load-cds`
   c. Unmount the CD. Enter: `umount /cdrom`
5. Load the first CD. Enter: `load_cds.py`

   Wait while the ISO image of the CD is written to the correct directory.
6. When asked "Would you like to read another CD?" [Y/N]:
   a. Insert the next CD in the service laptop.
   b. In the vSphere client, click the Connect/Disconnect CD/DVD icon and then select **Disconnect CD/DVD 1**.
   c. Connect the next CD. Select **Connect CD/DVD 1 ▶ Connect to <*CD_drive_letter*>**.
   d. Click in the vSphere console window.
   e. Type **y**, and press **<Enter>**.

7. Repeat step **6** for all CDs. The order in which CDs are loaded does not matter.
8. To exit, type **n** and press **<Enter>** when prompted.
9. Log out. Enter: **`exit`**

The next step is to install the provisioning software. See "Install the Provisioning Software" on page 37.

Use load_cds.py with ISOs

Use this procedure to load the ISOs if you have already copied ISO images of the installation CDs to the primary Admin Node. For example, you can use this procedure when the server is at a remote site, and you used **scp** to copy the files.

NOTE Do not put the ISO images in the /var/local/install directory of the primary Admin Node. Use any other directory instead, for example, /var/local/tmp. The load_cds.py script copies files from the directory that you specify to the /var/local/install directory.

Prerequisites

- Linux has been installed on the primary Admin Node. See "Install Linux on a Virtual Machine for the Primary Admin Node" on page 30.
- ISO images of the following CDs are on the primary Admin Node server:
  - StorageGRID 9.0.0 Software
  - Enablement Layer for StorageGRID 9.0.0 Software
  - If available, StorageGRID 9.0.x Software Service Pack
  - If available, Enablement Layer for StorageGRID 9.0.x Software Service Pack
  - Any additional CDs required for the grid, such as the TSM client packages CD

  Ensure that the ISO files are correct. The load_cds.py command does not perform any validation.

Procedure

NOTE The order in which ISO images are loaded does not matter.

1. Log in the primary Admin Node as root. When prompted for a password, press **<Enter>**.
2. Remove any USB flash drives from the server.
3. Install the load_cds.py script from the latest StorageGRID software CD. If available, use the StorageGRID 9.0.x Software Service Pack CD. Otherwise, use the StorageGRID 9.0 Software CD.
   a. Mount the latest CD. Enter:

   ```
   mount -o loop,ro <iso_service_pack_CD_including_path> /cdrom
   ```

   b. Install the updated load_cds.py script from the CD.

   Enter: `/cdrom/install-load-cds`

   c. Unmount the CD. Enter: `umount /cdrom`
4. Load the ISO images. Enter on one line:

   ```
   load_cds.py <iso_software_CD_including_path>
   <iso_enablement_layer_CD_including_path>
   <iso_enablement_layer__service_pack_CD_including_path>
   <iso_software_service_pack_CD_including_path>
   <iso_other_CD_including_path>
   ```

   Separate ISO file names with a space. The order does not matter.

   Wait until the ISO images are written to the correct directory.

The next step is to install the provisioning software. See "Install the Provisioning Software" below.

Install the Provisioning Software

Prerequisites

- ISO images of the software CDs have been loaded onto the primary Admin Node using load_cds.py. See "Load Installation ISOs" on page 34.

Procedure

1. At the primary Admin Node server, log in as root. When prompted for a password, press **<Enter>**.

2. If there is no service pack, mount the StorageGRID 9.0.0 Software CD image. Enter:

   ```
   mount -o loop,ro /var/local/install/Bycast_\
   StorageGRID_9.0.0_Software_<buildnumber>.iso /cdrom
   ```

3. If there is a service pack, mount the StorageGRID 9.0.x Software Service Pack image. Enter:

   ```
   mount -o loop,ro /var/local/install/Bycast_\
   StorageGRID_9.0.<servicepacknumber>_\
   Software_Service_Pack_<buildnumber>.iso /cdrom
   ```

4. Load the provisioning software. Enter:

   ```
   /cdrom/load-provisioning-software
   ```

5. When prompted, read and accept the NetApp StorageGRID Licensing Agreement.

6. When prompted, confirm that the current time is within ten minutes of the time displayed. If not, set the system date and time in UTC time using the format YYYY-MM-DD-hh-mm.

WARNING **Make sure that you verify the time displayed. Provisioning may fail if the displayed time is not within 10 minutes of the current time.**

   To determine the difference between local time and UTC time, use UTC Time conversion tools available on the internet, for instance thetimeNOW at http://www.thetimenow.com/index.cgi.

   The official current UTC time is available from the International Bureau of Weights and Measures (BIPM) at http://www.bipm.org/en/scientific/tai/time_server.html.

The next step is to provision the grid. See "Provision the Grid and Create Server Activation Floppy Image" below.

Provision the Grid and Create Server Activation Floppy Image

Provisioning is the process of turning a grid design into the collection of files needed to create, expand, and upgrade the grid. This collection of files is known as the gpt (grid provisioning tool) repository and includes the SAID (Software Activation and Integration Data) package. For more information on provisioning and the SAID package, see Appendix A, "Grid Specification Files and Provisioning".

Prerequisites

- Provisioning software has been installed. See "Install the Provisioning Software" on page 37.
- The following materials are available:
  - Provisioning floppy image. See "Create the Provisioning Floppy Image" on page 29.
  - a utility such as WinImage (available at http://www.winimage.com), that permits you to create a floppy disk image
  - a tool such as WinSCP (available at http://winscp.net/eng/download.php) to transfer files to and from the primary Admin Node
  - Service laptop

Procedure

1. At the primary Admin Node, log in as root. When prompted for a password, press **<Enter>**.
2. In vSphere Client, connect the Provisioning floppy image by clicking the Connect/Disconnect the floppy devices of the virtual machine icon and selecting **Connect Floppy 1 ▶ Connect to <*drive*>.**
3. Click in the vSphere console window to return to the command line.
4. Copy the GID<*Grid_ID*>_REV1_GSPEC.xml file from the Provisioning floppy image to the primary Admin Node:
   a. Mount the floppy image. Enter: `mount /media/floppy`
   b. Copy its contents to the Admin Node. Enter:
      ```
      mkdir /root/usb
      cp /media/floppy/* /root/usb
      ```
   c. Unmount the floppy image. It is no longer needed. Enter:
      ```
      umount /media/floppy
      ```
5. Provision the grid:
   a. Enter: `provision /root/usb`
   b. When prompted, enter a secure provisioning passphrase.
   c. Write down the provisioning passphrase for future reference. This passphrase is required for most installation and maintenance procedures.

WARNING **Make sure that the passphrase is stored in a secure place and is available for future reference.**

   d. When prompted, re-enter the provisioning passphrase.

When the process is complete, “Provisioning complete” is displayed. For example:

```
linux:/var/local/tmp # provision /root/usb
Please set a new passphrase for the GPT encrypted repository.
It must be at least 6 characters.
Enter passphrase:
Enter passphrase again:
Provisioning
Grid 400052 has been provisioned for software version: 9.0.0
Revision number 1 for Grid 400052
Created Sat Jan 07 18:09:42 UTC 2012, by the user: provision
Ensuring all certificates exist.
 creating grid CA cert
 creating grid ldr recovery keys
 creating grid task keys
 creating new node certs on server: agcs-190-39
 creating new node certs on server: gcs-190-40
 generating host ssh keys for: agcs-190-39
 generating the root user's ssh keys for ssh-access point: agcs-190-39
 generating host ssh keys for: gcs-190-40
 generating the root user's ssh keys for ssh-access point: gcs-190-40
 creating new nms cert on server: agcs-190-39
Creating server staging area and placing some files in it.
Creating Gator config file.
Running Gator.
Creating SAID staging area and copying files into it.
Running SPT to create Autoyast files.
 running spt for agcs-190-39
 running spt for gcs-190-40
Creating SAID file.
Finished creating revision Sat Jan 07 18:10:58 UTC 2010.
***
*** Please log out and log back in because the hostname has changed.
***
*** Warning: Do not log out of the Admin Node before extracting the
*** Passwords.txt file from the SAID package. You require the Admin
*** Node password from Passwords.txt to log back into the server.
***
No grid tasks to run
Saving GPT Repository...done
Copying GID400052_REV1_SAID.zip to USB key...done
Copying repository backup to USB key...done
Provisioning complete.
```

If provisioning ends with an error message, see “Provisioning Troubleshooting” on page 105.

The SAID package was written to the /root/usb directory.

WARNING **Do not log out of the Admin Node until you have transferred the SAID package to the service laptop. You need the Passwords.txt file from the SAID package to log back in.**

6. Unzip the GID<*grid_ID*>_REV1_SAID.zip file from the /root/usb directory to retrieve the Passwords.txt file. Then use the Admin Node password and WinSCP to copy the zip file from the Admin Node to your service laptop.

   Alternatively, unzip the SAID on the Admin Node server and transfer the Passwords.txt file, the Configuration.txt file, and the contents of the Doc directory from the Admin Node to your service laptop using a floppy image. (The entire SAID file is usually too large to fit in a floppy image.) Create the floppy image on your service laptop as /media/floppy, connect to it in the vSphere client, and mount it under Linux. Copy the files to the image, unmount it, and then extract the images to your laptop.

7. Unzip the GID<*grid_ID*>_REV1_SAID.zip file on your service laptop, and review the contents of the Doc/Index.html file to confirm the grid is configured correctly. For more information, see "About the SAID Package" on page 87.

WARNING **You must confirm that the SAID package is correct before proceeding. If errors are discovered later, you will have to start the provisioning process from the beginning. This includes re-installing any installed grid nodes and may take several days.**

   If you discover any errors, you must reinstall the primary Admin Node from the beginning. Go back to "Next Steps" on page 22.

8. Create the Server Activation media that contains the server activation files that are used to install Linux on each remaining server and customize the server for its assigned grid node role.

   For grid nodes installed on a VM, create Server Activation floppy images:

   a. Start the WinImage software. From the **File** menu, select **New**.

   b. In the Format selection dialog, select a standard format **1.44 MB** floppy. Click **OK**.

   c. Copy one or more server activation files to the floppy image:

      - From the **Image** menu, select **Inject**.
      - Browse to the Grid_Activation directory of the unzipped SAID package and select a server activation file.
      - Add additional activation files if desired. Three or four activation files should fit on a single floppy image.

d. Save the floppy image(s) with a descriptive name, and a file name ending in .flp.
    - From the **File** menu, select **Save**.
    - Select Save as type: **Virtual floppy Image (*.vfd,*.flp)**
    - Enter a file name ending in .flp, such as ***<servername>.flp***.

      You must enter the extension, or the vSphere client cannot use the image during installation.
    - Click **Save**.

e. Repeat until all server activation files are on floppy images.

9. Return to the Admin Node server, and log out.

   When you log back in to the Admin Node, the server has been assigned a hostname, and uses the password recorded in the Passwords.txt file.

10. Back up the provisioning data to a directory on the Admin Node. This backup copy can be used to restore the grid in the case of an emergency or during an upgrade or grid expansion.

    a. Log in to the primary Admin Node as root with the password listed in the Passwords.txt file.

    b. Create a directory for the backup provisioning data. Enter:

    ```
    mkdir -p /var/local/backup
    ```

    c. Back up the provisioning data. Enter:

    ```
    backup-to-usb-key /var/local/backup
    ```

    d. When prompted, enter the provisioning passphrase.

11. Store the Provisioning directory and the Backup Provisioning directories separately in a safe place. For example, use WinSCP to copy these directories to your service laptop, and then store them to two separate USB flash drives that are stored in two separate, secure physical locations. For more information, see "Preserving Copies of the Provisioning Data" on page 104.

    The contents of the Provisioning directory is used during expansion and maintenance of the grid when a new SAID package must be generated.

WARNING **Store copies of the Provisioning directory in two separate and secure locations. The Provisioning directories contain encryption keys and passwords that can be used to obtain data from the grid. The Provisioning directory is also required to recover from a primary Admin Node failure.**

Prepare Virtual Machines for all Other Grid Nodes

Complete the following tasks to prepare virtual machines for all other grid nodes.

Table 7: Prepare Virtual Machines for All Other Grid Nodes

✓	Step	Action	See
	1.	Install Linux on virtual machines.	page 43
	2.	Install VMware Tools and configure VM settings.	page 46
	3.	Configure Storage Nodes for NFS storage volumes.	page 47

Install Linux on Virtual Machines

Use this procedure to install Linux on all virtual machines in the grid except, the virtual machine for the primary Admin Node.

Prerequisites

- The virtual machine for the grid node has been started. See “Start the Virtual Machines” on page 28.
- SLES DVD. For information on the supported version of SLES, see the Interoperability Matrix Tool (IMT).
- Server Activation floppy image. See “Provision the Grid and Create Server Activation Floppy Image” on page 38.

Procedure

1. Insert the SLES DVD into the machine from which you are running the vSphere client. Skip this step if you are using an ISO image of Linux.
2. In the VMware vSphere client navigation tree, select the virtual machine.
3. Click the Connect/Disconnect CD/DVD drive to the virtual machine icon, then select **Connect CD/DVD 1 ▶ Connect to <*CD_drive_letter*>**

   — or —

   If you are using an ISO image of the Linux installation, select **Connect CD/DVD 1 ▶ Connect to ISO image on local disk**.
4. Click the **Console** tab.

5. Click anywhere inside the Console pane to enter the Console pane.

   Your mouse pointer disappears.

TIP Press <Ctrl>+<Alt> to release your mouse pointer from the VM console.

---

6. Press **<Ctrl>+<Alt>+<Insert>** to reset the virtual machine. The server performs the following steps:
   - The BIOS runs a hardware verification.
   - By default the system boots from the DVD, and loads the SUSE Linux Enterprise Server Boot Screen in the VMware vSphere client Console pane.

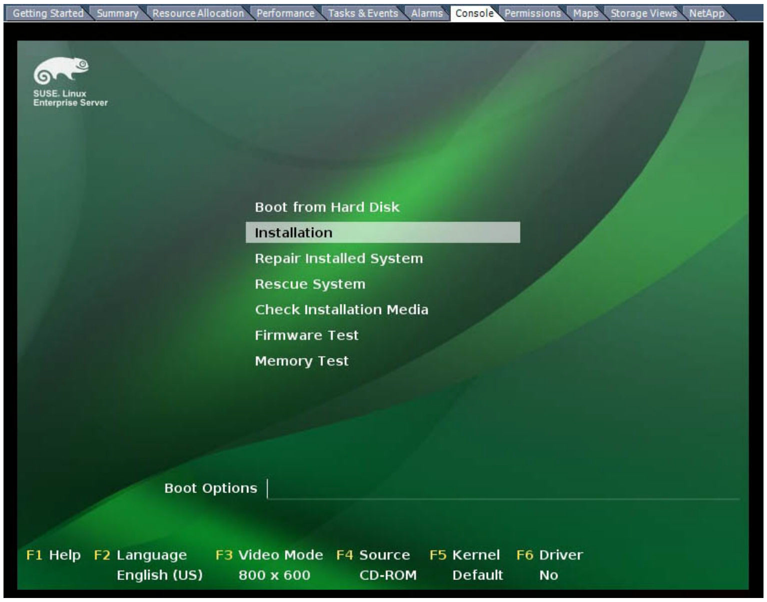

Figure 6: SLES Boot Screen

7. From the SLES Boot Screen:
   a. Press the keyboard's down arrow and highlight **Installation**. (Do not press **<Enter>**.)

NOTE You must move the cursor to the Installation option within eight seconds. If you do not, SLES will automatically attempt to install from the hard drive and the installation process will fail. If this happens, you must begin the installation process again from the beginning.

---

   b. Press **<Ctrl>+<Alt>** to leave the Console pane.

   The mouse icon disappears.

c. Click **Connect Floppy 1** and select Connect to Floppy Image on local disk.

d. On your service laptop, select the floppy image that contains the activation file for this server.

e. Click anywhere inside the Console pane to return to the Console pane.

f. Press **<Tab>**. At the bottom of the screen, adjacent to the **Boot Options** prompt, enter:

```
autoyast=device://fd0/<servername>-autoinst.xml
```

where <*servername*> is the server name used to name the activation file.

NOTE If you do not enter the path to the activation file when required, AutoYaST does not customize the installation for the server. Always enter the path to the activation file.

If you enter an incorrect value, and are prompted to re-enter the device name and path, check the floppy device name.

g. Press **<Enter>**.

Wait about one minute while the installer processes the information.

The base SLES installation completes without further intervention.

8. During this installation process, disconnect the floppy image after it is no longer required. Click the Connect/Disconnect the floppy devices of the virtual machine icon and select **Disconnect Floppy 1 ▶ Disconnect from <*drive*>**.

NOTE If the floppy image is connected when the virtual machine reboots, the message “This is not a bootable disk is displayed.” To continue, disconnect the floppy image and hit any key to continue.

When SLES installation is complete, the server completes its configuration and starts the operating system. Installation is complete when the login prompt appears.

For the next step, see “Install VMware Tools and Configure VM Settings” on page 46.

Install VMware Tools and Configure VM Settings

Install VMware Tools on each virtual machine that will host a grid node. For information on the supported versions of VMware software, see the Interoperability Matrix Tool (IMT).

Install VMware Tools from an ISO image that is packaged with vSphere client software.

The vmware_setup.py script configures the video resolution settings that the virtual machine uses on startup. This is necessary because the default resolution is not supported by the VMware video adapter, which results in an error on startup.

Prerequisites

- Linux has been installed on the virtual machine

Procedure

1. In the Virtual Machine tree, right-click the virtual machine, then select **Guest ▶ Install/ Upgrade VMware Tools**.

   The Install VMware Tools dialog appears.
2. Select **Interactive Tool Upgrade** and click **OK**.

   The VMware Tools package is made available to the virtual machine as an ISO at /cdrom.

   Wait until VMware disconnects the Linux CD/DVD.
3. In the VMware vSphere Client, click in the **Console** pane of the virtual machine and log in to the host that will be the new grid node.
4. Copy the VMware Tools packages to the virtual machine:
   a. Mount the ISO image of the VMware Tools CD. Enter:
      ```
      mount /cdrom
      ```
   b. Copy the gzip package from the CD to the virtual machine, and unpack it. Enter:
      ```
      mkdir /tmp/vmtools
      cd /tmp/vmtools
      tar -zxvf /cdrom/VMwareTools-*.tar.gz
      ```
5. Install VMware Tools, accepting the default installation options. Enter:
   ```
   cd /tmp/vmtools/vmware-tools-distrib/
   ./vmware-install.pl --default
   ```
   Wait for the installation to complete. This takes about a minute.

6. Check to make sure that VMware Tools is running. Enter:

   ```
   /etc/init.d/vmware-tools status
   ```

   You will see the message "vmtoolsd is running".

7. Remove the installation files from the virtual machine. Enter:

   ```
   cd /tmp
   rm -rf vmtools
   ```

8. Run the vmware_setup.py script. Enter: `vmware_setup.py`

   The script completes silently.

9. Reboot to ensure the changes take effect. Enter: `reboot`

Configure Storage Nodes for NFS Storage Volumes

Use this procedure to configure the server for the installation of a Storage Node integrated with NFS mounted storage volumes.

Prerequisites

- The NetApp storage system has been set up and exported via NFS.

NOTE Configuration of the NFS server is beyond the scope of this guide.

- NFS mounted storage volumes are integrated with the Storage Node as described in the the *Administrator Guide*.
- Ensure that you have the IP address of the NFS server.

Procedure

1. Log in to the Storage Node server as root, using the password provided in the Passwords.txt file.
2. Verify connectivity to the NFS server. Enter: `ping <NFS_Server_IP>`
3. Verify that the Linux NFS client package is present (it should be installed by default). Enter:
   - For SLES 11: `rpm -q nfs-client`
   - For SLES 10: `rpm -q nfs-utils`
4. Mount the NFS exports:
   a. At the Storage Node server, using a text editor such as vi, add this line to the /etc/fstab file for each storage volume (on one line):

```
<NFS_Server_IP>:<volume_path> /var/local/rangedb/<next_index>
nfs rw,rsize=65536,wsize=65536,nfsvers=3,tcp
```

where:

- *<NFS_Server_IP>* is the IP address of the NFS server
- *<volume_path>* is the path of the storage volume exported from the NFS server
- *<next_index>* is the Storage Node rangedb, a number between 0 and 15 in hexadecimal notation (0 to F, case-specific). For the first storage volume, the index number is 0.

For example:

```
192.168.130.16:/vol/vol1 /var/local/rangedb/0 nfs
rw,rsize=65536,wsize=65536,nfsvers=3,tcp
```

Repeat step **a** for each storage volume.

b. Create a mount point for each NFS storage volume, using the same index numbers used in the etc/fstab file in step **a**. Enter:

```
mkdir -p /var/local/rangedb/<next_available_index>
```

For example:

```
mkdir -p /var/local/rangedb/0
```

Do not create mount points outside of /var/local/rangedb.

Repeat step **b** for each storage volume.

c. Mount the storage volumes to the mount points. Enter:

```
mount -a
```

If mounting fails, make sure that the storage volumes are configured on the NFS server for read-write access to the Storage Node and that the IP address of the Storage Node supplied to the NFS server is correct.

Troubleshooting

This section includes troubleshooting topics to help you identify and solve problems that may occur while preparing virtual machines. See also “Provisioning Troubleshooting” on page 105 and “GDU Troubleshooting” on page 115.

If problems persist, contact Support. You may be asked to supply the following installation log files:

- /var/local/log/install.log (found on the server being installed)
- /var/local/log/gdu-console.log (found on the primary Admin Node)

Virtual Machine Not Started

If a virtual machine does not start after you create it, see "VM Resource Reservation Requires Adjustment" on page 49.

If a virtual machine does not restart after VMware ESX/ESXi is restarted, see "VM is Not Configured for Automatic Restart" on page 49.

VM Resource Reservation Requires Adjustment

The OVF files created by Grid Designer 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 ESX/ESXi and the predefined number of resources are not available, the virtual machines will not start.

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.

1. In the VMware vSphere 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**, then use the **Reservation** slider to adjust the MHz reserved for this virtual machine.
   b. Select **Memory**, 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.

VM is Not Configured for Automatic Restart

If the virtual machine does not restart after VMware ESX/ESXi is restarted, it is most likely that the virtual machine has not been configured for automatic restart.

1. In the VMware vSphere client tree, select the virtual machine that is not started.

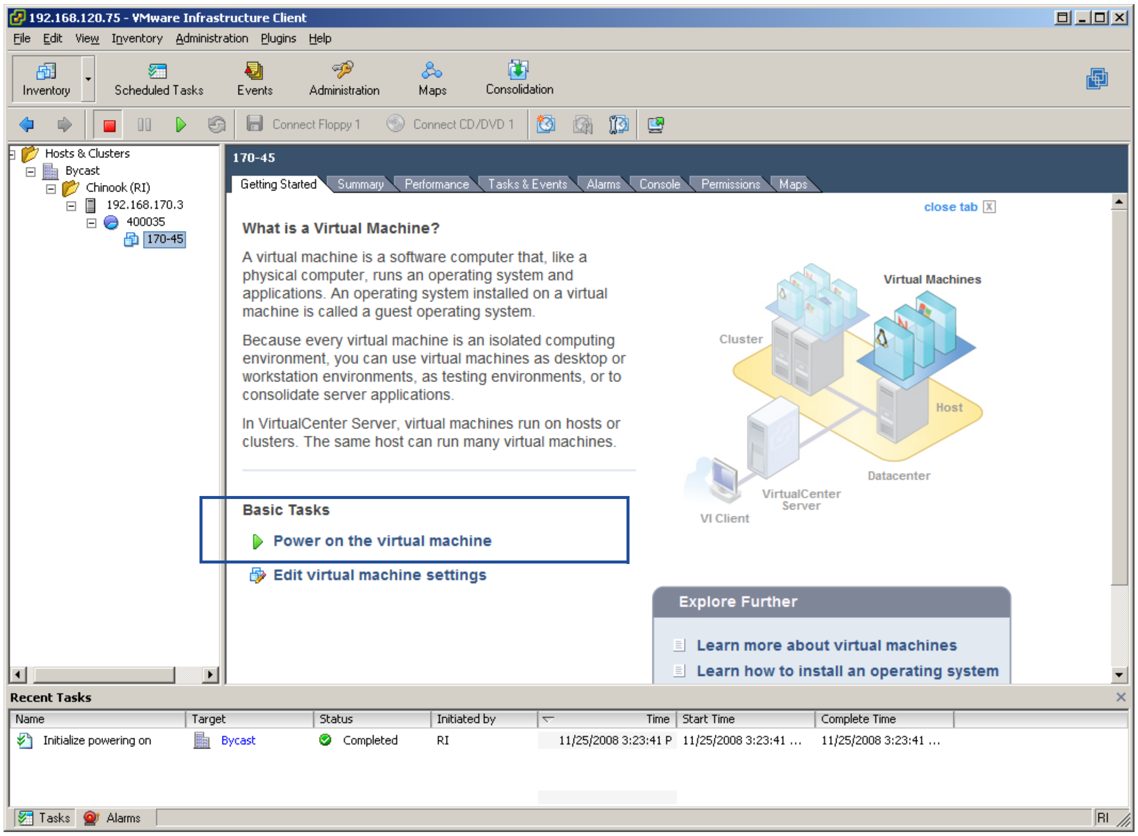

What is a Virtual Machine?

A virtual machine is a software computer that, like a physical computer, runs an operating system and applications. An operating system installed on a virtual machine is called a guest operating system.

Because every virtual machine is an isolated computing environment, you can use virtual machines as desktop or workstation environments, as testing environments, or to consolidate server applications.

In VirtualCenter Server, virtual machines run on hosts or clusters. The same host can run many virtual machines.

Name	Target	Status	Initiated by	Time	Start Time	Complete Time
Initialize powering on	Bycast	Completed	RI	11/25/2008 3:23:41 P	11/25/2008 3:23:41 ...	11/25/2008 3:23:41 ...

Figure 7: Virtual Machine Manual Restart

2. In the Getting Started pane, under Basic Tasks, click **Power on the virtual machine**.
3. Configure the virtual machine to restart automatically. See "Configure ESX/ESXi for Automatic Restart" on page 26.

Next Steps

Your next step is to install grid software on the prepared virtual machines. See Chapter 4: "Install and Start Grid Software".

4

Install and Start Grid Software

How to install and then start StorageGRID software in the correct order

Verify Networking

Before you install grid software, verify networking to confirm that the primary Admin Node can communicate with the other servers in the grid.

Install Grid Software

Use the following procedure to install grid software on all virtual machines in the following order: primary Admin Node, one Control Node, and then the remaining grid nodes.

You must install the primary Admin Node first because the primary Admin Node provides the provisioning environment for the grid. Provisioning is the process of turning a grid design into the collection of files needed to create, expand, or upgrade that grid. This collection of files is referred to as the `gpt` (grid provisioning tool) repository.

Control Nodes require extensive database initialization that can take approximately three hours to complete, which is much longer than other servers.

You can install grid software on multiple servers in parallel to reduce the total time required to install the grid.

Prerequisites

- There is connectivity between the primary Admin Node and the other servers. If there is no connectivity to the Admin Node, use

GDU to install software on the primary Admin Node and then go to Appendix C: "Install Grid Software Manually".

- Linux has been installed
- VM Tools have been installed and are running on the server. See "Prepare Virtual Machines" on page 23.
- If the server is not at the same site as the primary Admin Node, the ISO images of the StorageGRID Software CD and the Enablement Layer for StorageGRID Software CD have been copied to the server in order to reduce WAN traffic. See "Copy ISO Files in Multi-Site Environment" on page 117.

Procedure

1. Start GDU on the primary Admin Node. See "How to Use GDU" on page 109.
2. Install the StorageGRID software on the server:
    a. Select the server where you want to install the software in the **Servers** panel and confirm that its current state is Available.
    b. Select **Install Software** in the **Tasks** panel, and then select **Start Task** in the **Actions** panel and press **<Enter>**.

       Installation times vary depending on the size of the database being set up. It can take approximately 45 minutes on Admin Nodes, three hours on Control Nodes, and up to three hours on Storage Nodes which have storage installed. The script completes in less than 10 minutes on other servers.

       GDU fails with an error message if software is being installed in a virtual machine and VMTools have not been installed.

       The script uses configuration data from autoinst.xml files loaded earlier to identify the role of the server in the grid and the hardware configuration required. The script next adds or removes Linux packages as required to customize the operating system for that node. Finally, if required on this server, the script sets up the MySQL database.

       If the server hosts an LDR service and the provisioning hardware profile has not specified object store names, the installation script detects the unallocated drives and formats the disks.

       If the server hosts a Gateway Node, the server reboots automatically as part of the installation process.
3. If necessary, install drivers for the physical servers onto which the virtual machines are installed. See "Install Drivers" on page 53.

Install Drivers

When installing StorageGRID software on supported servers, you may also need to install drivers. The Enablement Layer CD for StorageGRID software may include drivers for supported servers, and you can use GDU to install drivers included on the Enablement Layer CD. However, when the Enablement Layer CD excludes required drivers, you cannot use GDU to install the drivers. You must locate and manually install the drivers. For information about supported servers, see the Interoperability Matrix posted on the NetApp Support Site (http://support.netapp.com/).

It is your responsibility to confirm that the drivers are the most recent qualified version. For the latest version, see your hardware vendor.

Start Grid Software

NOTE Do not start grid software until instructed to do so. Grid software can be started after you install the primary Admin Node, install one Control Node, and load NMS configuration settings.

After the grid software has been installed, start the grid software using GDU. Begin with the primary Admin Node followed by one Control Node. After the primary Admin Node and one Control Node are running, you must load the NMS configuration settings and connect to the NMS MI (management interface) via a web browser to start monitoring the grid. Finally, start the grid software on the other servers in the grid.

Start Software on Primary Admin Node

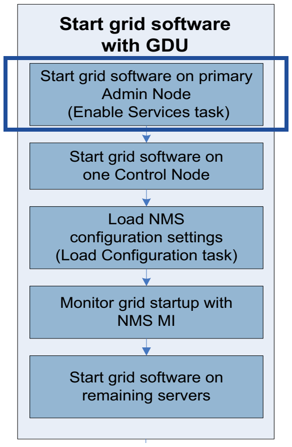
Start grid software with GDU
Start grid software on primary Admin Node (Enable Services task)
Start grid software on one Control Node
Load NMS configuration settings (Load Configuration task)
Monitor grid startup with NMS MI
Start grid software on remaining servers

Start the grid software on the primary Admin Node first.

Prerequisites and required materials

- If the grid has an HCAC, there is connectivity between the two Admin Node servers
- Grid software has been installed on the server

Procedure

1. Start GDU on the primary Admin Node if not already running. See "How to Use GDU" on page 109.
2. Start the StorageGRID software on the primary Admin Node:
   a. In the **Servers** panel, select the primary Admin Node and confirm that its current state is Available.
   b. In the **Tasks** panel, select **Enable Services** and then in the **Actions** panel select **Start Task** and press **<Enter>**. Wait for the task to complete.
3. If the primary Admin Node is part of a High Capacity Admin Cluster (HCAC), start the software on the processing Admin Node:
   a. In the **Servers** panel, select the processing Admin Node and confirm that its current state is Available.
   b. In the **Tasks** panel, select **Enable Services**, and then in the **Actions** panel select **Start Task** and press **<Enter>**.

      Wait for the task to complete.

The next step is to start one of the Control Nodes. See "Start Software on One Control Node" below.

Start Software on One Control Node

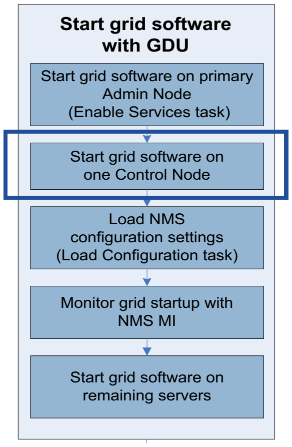
Start grid software with GDU
Start grid software on primary Admin Node (Enable Services task)
Start grid software on one Control Node
Load NMS configuration settings (Load Configuration task)
Monitor grid startup with NMS MI
Start grid software on remaining servers

After you have started the software on the primary Admin Node, start the software on one Control Node. This is required to bring an ADC service online.

Prerequisites and required materials

- There is connectivity between the primary Admin Node and the Control Node
- Grid software has been installed on the server

Procedure

1. Start GDU on the primary Admin Node if not already running. See "How to Use GDU" on page 109.
2. Start the StorageGRID software on the Control Node:
   a. In the **Servers** panel, select the server and confirm that its current state is Available.
   b. In the **Tasks** panel, select **Enable Services**, and then in the **Actions** panel select **Start Task** and press **<Enter>**. Wait for the task to complete.

The next step is to load the NMS configuration settings. See "Load NMS Configuration Settings" below.

Load NMS Configuration Settings

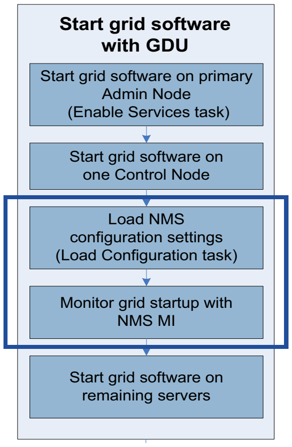
Start grid software with GDU
Start grid software on primary Admin Node (Enable Services task)
Start grid software on one Control Node
Load NMS configuration settings (Load Configuration task)
Monitor grid startup with NMS MI
Start grid software on remaining servers

Loading the NMS configuration settings for the grid can be done any time after at least one ADC service (the ADC is a service that runs on Control Nodes) is running. The primary Admin Node uses the ADC to distribute the required settings to all other grid nodes.

The NMS MI (management interface) allows you to verify that services are recognized and have joined the grid. Without the NMS MI, you may not be made aware of network connection or authentication problems.

Prerequisites

- Grid software has been started on the primary Admin Node
- Grid software has been started on one Control Node
- Service laptop is available to connect to the NMS MI

Procedure

1. Start GDU on the primary Admin Node if not already running. See "How to Use GDU" on page 109.
2. Select the primary Admin Node server in the **Servers** panel and confirm that its current state is Available.
3. Select **Load Configuration** in the **Tasks** panel, and then select **Start Task** in the **Actions** panel and press **<Enter>**. Wait for the task to complete.
4. Connect to the customer network from the service laptop.

   Work with the customer system administrator to establish the physical network connection to the service laptop. Using the customer's network rather than a direct connection within the rack verifies that the interface is accessible using the same infrastructure the customer uses.
5. Launch the web browser.
6. Ensure that the browser is configured to permit pop-ups.
7. Go to **https://<*IP_address*>**

   where <*IP_address*> is the client-side IP address of the primary Admin Node. In a grid with an HCAC, use the IP address of the reporting Admin Node.

   Wait for the NMS Log In page to appear. It may take several minutes after the configuration settings have been loaded.
8. Log in to the NMS MI using the username **Vendor** and the password listed in the Passwords.txt file.

   The Grid Overview page appears. All grid nodes and their services should appear in the Grid Topology tree. Services that have not yet been started are gray, indicating that they are "Administratively Down". As services are started, they assume the color appropriate to their status.

   Some alarms appear during the start-up as new services connect and locate required components. The alarms clear as the process continues. Archive Nodes show a major alarm (orange) when they are first started. This is normal. You must complete the TSM Archive Node configuration steps described in "Complete Setup of the TSM Archive Node" on page 62 to clear the alarm.

For detailed information on how to use the NMS MI, see the *Grid Primer*.

The next step is to start the grid software on the remaining servers. See "Start Software on Other Grid Nodes" below.

Start Software on Other Grid Nodes

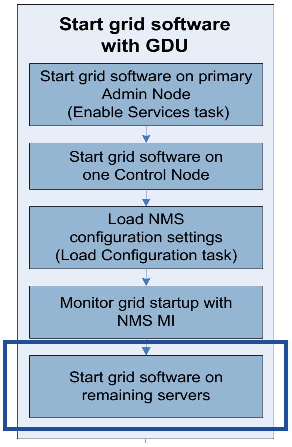
Start grid software with GDU
Start grid software on primary Admin Node (Enable Services task)
Start grid software on one Control Node
Load NMS configuration settings (Load Configuration task)
Monitor grid startup with NMS MI
Start grid software on remaining servers

You can now start the software on the other grid nodes in the grid.

If a site includes a secondary HCAC, start it before starting the other grid nodes at that site. This allows the Admin Nodes to start processing data at the earliest possible time, thereby making the databases of the HCACs more consistent.

Prerequisites

- There is connectivity between the primary Admin Node and the server
- Grid software has been installed on the server

Procedure

1. Start GDU on the primary Admin Node if not already running. See "How to Use GDU" on page 109.
2. Start the StorageGRID software on the server:
    a. In the **Servers** panel, select the server and confirm that its current state is Available.
    b. In the **Tasks** panel, select **Enable Services**, and then select **Start Task** in the **Actions** panel and press **<Enter>**. Wait for the task to complete.

Close GDU when you no longer need it. See "Close GDU" on page 114.

Next Steps

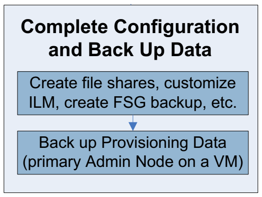
Complete Configuration and Back Up Data
Create file shares, customize ILM, create FSG backup, etc.
Back up Provisioning Data (primary Admin Node on a VM)

After you have installed StorageGRID software on all virtual machines:

- Ensure that you store the provisioning data in two separate, secure locations. See "Preserving Copies of the Provisioning Data" on page 104.
- Complete the configuration and customization steps described in Chapter 5: "Configure the Grid" on page 61.

Troubleshooting

This section includes troubleshooting topics to help you identify and solve problems that may occur while installing the StorageGRID software. See also "Provisioning Troubleshooting" on page 105 and "GDU Troubleshooting" on page 115.

If problems persist, contact Support. You may be asked to supply the following installation log files:

- /var/local/log/install.log (found on the server being installed)
- /var/local/log/gdu-console.log (found on the primary Admin Node)

Corrupt ISO Message When Using load_cds.py Script

If you get an error message that a corrupt ISO has been detected (for instance, a failed md5sum check), check the integrity of the CD and load the CD again.

GDU Fails When Re-running "Install Software" Task

If your initial attempt to install software on a grid node fails, GDU may transition into an error state that will not allow you to successfully install software.

When you encounter this issue, GDU displays only the Update Status task for the grid node. After refreshing the display by selecting Update Status, the Install Software task becomes available again, but subsequent attempts to install software fail.

To work around this issue, unmount volumes and remove them from the /etc/fstab file before attempting to install the grid node again.

1. Log in to the grid node as root using the password listed in the Passwords.txt file.
2. Check what is mounted on the server. Enter: **`mount`**
3. Unmount all storage volumes, FSG mount points, MySQL partitions and/or audit partitions from the server. Enter (as required):

```
umount /var/local/rangedb/*
umount /fsg
umount /var/local/mysq_ibdata
umount /var/local/audit
```

4. Edit /etc/fstab.

   Keep entries for NFS mounted storage and system drives.

   Remove the lines for any volumes whose device name includes “by-uuid” or “fsgvg-fsglv”. For example, remove any lines similar to the following:

```
/dev/disk/by-uuid/2a1204d2-db57-4358-b742-142e20d90ec6 /var/local/mysql_ibdata
ext3 errors=remount-ro,noatime,barrier=0 0 1
/dev/mapper/fsgvg-fsglv /fsg xfs dmapi,mtpt=/fsg,noalign,nobarrier,ikeep 0 2
/dev/disk/by-uuid/e2ibre68f0-6ed0-41c3-9094-c29e327173ef /var/local/rangedb/0
ext3 errors=remount-ro,dirsync,barrier=0,data=writeback
/dev/disk/by-uuid/e23e444d0-2cc1-85c3-1111-c56f324563cf /var/local/rangedb/1
ext3 errors=remount-ro,dirsync,barrier=0,data=writeback
```

5. Save /etc/fstab.
6. Return to GDU, and select **Update Status** for that grid node.
7. Retry the **Install Software** task.

5

Configure the Grid

Process Overview

After the grid software has been installed, you must complete a number of configuration steps to enable the grid to work its deployment environment. Some of the procedures are required to ensure the grid works correctly; others need only be performed to customize the grid in a particular way.

Required Grid Configuration

Perform the following procedures to enable the grid to work correctly:

- Complete the installation of the TSM Archive Node. See "Complete Setup of the TSM Archive Node" on page 62.
- Create the first backup of the FSG managed file system to create initial conditions for future maintenance. See "Create the Initial FSG File System Backup" on page 72.
- If the grid includes a High Availability Gateway Cluster (HAGC), ensure that it is configured with more than one ping node for increased stability. For more information, see the *Administrator Guide*.
- If the primary Admin Node is installed in a virtual machine, ensure that you have preserved two copies of the provisioning data in separate locations as described in "Preserving Copies of the Provisioning Data" on page 104.

WARNING **You must preserve copies of the provisioning data. Store copies of the Provisioning directory in two separate and secure locations. The Provisioning directories contain encryption keys and passwords that can be used to obtain data from the grid. The Provisioning directory is also required to recover from a primary Admin Node failure.**

Optional Grid Customization

After completing an installation, you can customize the grid as follows:

- Create NMS MI accounts for users and administrators.
- Customize the ssh access point for increased security
- Configure DNS
- Configure e-mail notification for Admin Nodes
- Configure SNMP monitoring
- Enable deduplication

NOTE Deduplication is deprecated and no longer supported.

- Reconfigure IP addresses
- Configure storage compression
- Configure the grid's ILM policy
- Create CIFS or NFS file shares to enable clients to store and retrieve objects to the grid

For more information and procedures, see the *Administrator Guide*.

Complete Setup of the TSM Archive Node

The TSM Archive Node is not functional after you complete the installation of the Archive Node. Before the grid can save objects to the TSM Archive Node, you must complete the installation and configuration of the TSM server and configure the Archive Node to communicate with the TSM server.

For information on optimizing TSM retrieval and store sessions, see the *Administrator Guide*.

Install a New TSM Server

You can integrate the Archive Node with either a new or an existing TSM server. In either case, use the information in "Best Practices for TSM Integration" on page 19 to help plan the integration of the Archive Node with the TSM server.

If you are integrating with an existing TSM, skip to "Configure the TSM Server" below.

If you are installing a new TSM on its own server, follow the instructions in your TSM documentation to complete the installation. After installation is complete, go on to "Configure the TSM Server" below.

NOTE An Archive Node cannot be co-hosted with a TSM server.

Configure the TSM Server

This section includes sample instructions for preparing a TSM Server that follow the recommendations outlined in "Best Practices for TSM Integration" on page 19. These instructions guide you through the process of:

- Defining a disk storage pool, and a tape storage pool (if required) on the TSM server.
- Defining a domain policy that uses the TSM management class for the data saved from the Archive Node, and registering a node to use this domain policy.

These instructions are provided for your guidance only; they are not intended to replace TSM Server documentation, or to provide complete and comprehensive instructions suitable for any customer configuration. Instructions suitable for your site should be provided by a TSM administrator who is familiar both with your detailed requirements, and with the complete set of TSM Server documentation.

Define Storage Pools

The Archive Node writes to a disk storage pool. To archive content to tape, the disk storage pool must be configured to move content to a tape storage pool.

- For a TSM server, you must define a tape storage pool and a disk storage pool within Tivoli Storage Manager. After the disk pool is defined, create a disk volume and assign it to the disk pool.

A tape pool is not required if your TSM server uses disk-only storage.

There are a number of steps that must be completed on your TSM server before you can create a tape storage pool. (Create a tape library and at least one drive in the tape library. Define a path from the server to the library and from the server to the drives, and then define a device class for the drives.) The details of these steps may vary depending upon the hardware configuration and storage requirements of the site. For more information, see the TSM documentation.

A sample set of instructions are included here: be aware that the requirements of your site may vary. For configuration details and for instructions, see the TSM documentation for your server.

You must log onto the server with administrative privileges and use the dsmadmc tool to execute the following commands.

1. Create a tape library. Enter:

   ```
   define library <tapelibrary> libtype=scsi
   ```

   where *<tapelibrary>* is an arbitrary name chosen for the tape library, and the value of libtype may vary depending upon the type of tape library.
2. Define a path from the server to the tape library. Enter (on one line):

   ```
   define path servername tapelibrary srctype=server
   desttype=library device=lib-devicename
   ```

   where:

   - servername is the name of the TSM server
   - lib-devicename is the device name for the tape library
3. Define a drive for the library. Enter:

   ```
   define drive tapelibrary <drivename>
   ```

   where you may choose any *<drivename>* you require.

   You may want to configure an additional drive or drives, depending upon your hardware configuration. (For example, if the TSM server is connected to a fibre channel switch that has two inputs from a tape library, you may wish to define a drive for each input.)
4. Define a path from the server to the drive(s) you defined. Enter (on one line):

   ```
   define path servername drivename srctype=server
   desttype=drive library=tapelibrary device=drive-dname
   ```

   where *drive-dname* is the device name for the drive, and tapelibrary is the name of the tape library as defined in step **1**.

   Repeat for each drive that you have defined for the tape library, using a separate drivename and drive-dname for each drive.

5. Define a device class for the drives. Enter (on one line):

   ```
   define devclass DeviceClassName devtype=lto
   library=tapelibrary format=ultrium3
   ```

   where:

   - *DeviceClassName* is the name of the device class
   - *lto* describes the type of drive connected to the server
   - *tapelibrary* is the tape library name defined in step **1**
   - substitute the appropriate value for *ultrium3* in the format= *parameter* to match your tape type

6. Add tape volumes to the inventory for the library. Enter:

   ```
   checkin libvolume tapelibrary
   ```

   where *tapelibrary* is the tape library name defined in step **1** of "Create a tape library. Enter:" on page 64.

7. Create the primary tape storage pool. On the TSM server, enter (on one line):

   ```
   define stgpool BycastTapePool DeviceClassName
   description=description collocate=filespace maxscratch=XX
   ```

   where:

   - *BycastTapePool* is the name of the Archive Node's tape storage pool. You can select any name for the tape storage pool (as long as the name uses the syntax conventions expected by the TSM).
   - *DeviceClassName* is the name of the device class name for the tape library.
   - *description* is a description of the storage pool that can be displayed on the TSM server using the 'query stgpool' command. For example: "Tape storage pool for the Archive Node".
   - Setting collocate to "filespace" specifies that the TSM server should write objects from the same filespace into a single tape.
   - *XX* is:
     - the number of empty tapes in the tape library (in the case that the Archive Node is the only application using the library).

       — or —
     - the number of tapes allocated for use by the StorageGRID system (in instances where the tape library is shared).

8. On a TSM server, create a disk storage pool. At the TSM server's administrative console, enter (on one line):

   ```
   define stgpool BycastDiskPool disk description=description
   maxsize=maximum_file_size nextstgpool=BycastTapePool
   highmig=percent_high lowmig=percent_low
   ```

   where:

   - *BycastDiskPool* is the name of the Archive Node's disk pool. You can select any name for the disk storage pool (as long as the name uses the syntax conventions expected by the TSM).
   - *description* is a description of the storage pool that can be displayed on the TSM server using the 'query stgpool' command. For example, "Disk storage pool for the Archive Node".
   - *maximum_file_size* forces objects larger than this size to be written directly to tape, rather than being cached in the disk pool. It is recommended to set *maximum_file_size* to 10 GB.
   - nextstgpool=*BycastTapePool* refers the disk storage pool to the tape storage pool defined for the Archive Node.
   - *percent_high* sets the value at which the disk pool begins to migrate its contents to the tape pool. It is recommended to set *percent_high* to 0 so that data migration begins immediately.
   - *percent_low* sets the value at which migration to the tape pool stops. It is recommended to set *percent_low* to 0 to clear out the disk pool.

9. On a TSM Server, create a disk volume (or volumes) and assign it to the disk pool. Enter:

   ```
   define volume BycastDiskPool volume_name formatsize=size
   ```

   where:

   - *BycastDiskPool* is the disk pool name defined above
   - *volume_name* is the full path to the location of the volume (for example /var/local/arc/stage6.dsm) on the TSM server where it writes the contents of the disk pool in preparation for transfer to tape.
   - *size* is the size, in MB, of the disk volume.

   For example, to create a single disk volume such that the contents of a disk pool fill a single tape, set the value of *size* to 200000 when the tape volume has a capacity of 200 GB.

   However, it may be desirable to create multiple disk volumes of a smaller size, as the TSM server can write to each volume in the disk pool. For example, if the tape size is 250 GB, create 25 disk volumes with a *size* of 10 GB (10000) each.

The TSM server preallocates space in the directory for the disk volume. This can take some time to complete (more than three hours for a 200 GB disk volume).

Create a Domain Policy and Register a Node

Next, define a domain policy that uses the TSM management class for the data saved from the Archive Node, and then register a node to use this domain policy.

NOTE Archive Node processes may leak memory if the client password for the Archive Node in Tivoli Storage Manager (TSM) expires. Ensure that the TSM is configured such that the client username/password for the Archive Node never expires.

When registering a node on the TSM server for the use of the Archive Node (or updating an existing node), you must specify the number of mount points that the node can use for write operations by specifying the MAXNUMMP parameter to the REGISTER NODE command. The number of mount points is typically equivalent to the number of tape drive heads allocated to the Archive Node. The number specified for MAXNUMMP on the TSM server must be at least as large as the value set for the ARC ▶ Middleware ▶ Configuration ▶ Main ▶ Maximum Store Sessions for the Archive Node. Which is set to a value of 0 or 1, as concurrent store sessions are not supported by the Archive Node. For more information, see Table 8 on page 69.

The value of MAXSESSIONS set for the TSM Server controls the maximum number of sessions that can be opened to the TSM server by all client applications.The value of MAXSESSIONS specified on the TSM must be at least as large as the value specified for ARC ▶ Middleware ▶ Configuration ▶ Main ▶ Number of Sessions in the NMS for the Archive Node. The Archive Node concurrently creates at most one session per mount point plus a small number (< 5) of additional sessions.

TSM Server

The TSM node assigned to the Archive Node uses a custom domain policy *tsm-domain*. The *tsm-domain* domain policy is a modified version of the “standard” domain policy, configured to write to tape and with the archive destination set to be the grid’s storage pool (*BycastDiskPool*).

You must log in to the TSM server with administrative privileges and use the dsmadmc tool to create and activate the domain policy.

Create and Activate the Domain Policy

1. Create a Domain Policy. Enter: `copy domain standard tsm-domain`
2. Determine the name of the Management Class to be used with the Archive Node:
    a. If you are using an existing management class go to step **3**.
    b. If necessary, create the management class you need. Enter:

    ```
    define policyset tsm-domain standard
    define mgmtclass tsm-domain standard default
    ```

    where *default* is the default management class for the grid deployment (as specified in "Configure the TSM Archive Node" on page 69) and we first create a policy set called standard.
3. Create a copygroup to the appropriate storage pool. Enter (on one line):

    ```
    define copygroup tsm-domain standard default type=archive
    destination=BycastDiskPool retinit=event retmin=0 retver=0
    ```

    where *default* is the default Management Class for the Archive Node. The values of retinit, retmin, and retver have been chosen to reflect the retention behavior currently used by the Archive Node.

    NOTE Do not set `retinit` to `retinit=create`. Setting `retinit=create` blocks the Archive Node from deleting content since retention events are used to remove content from TSM

4. Assign the management class to be the default. Enter:

    ```
    assign defmgmtclass tsm-domain standard default
    ```

5. Set the new policy set as active. Enter:

    ```
    activate policyset tsm-domain standard
    ```

    Ignore the "no backup copy group" warning that appears when you enter the activate command.
6. Register a node to use the new policy set on the TSM. On the TSM server, enter (on one line):

    ```
    register node arc-user arc-password passexp=0
    domain=tsm-domain MAXNUMMP=number-of-sessions
    ```

    where *arc-user* and *arc-password* are same client node name and password as you define on the Archive Node, as described in "Configure the TSM Archive Node" on page 69, and the value of MAXNUMMP is set to the number of tape drives reserved for Archive Node store sessions (as described on page 71).

    Note that by default, registering a node creates an administrative user ID with client owner authority, with the password defined for the node).

Configure the TSM Archive Node

Before the Archive Node can communicate with the TSM middleware running on the TSM server, you must configure a number of settings in the NMS MI that permit the Archive Node to recognize it. Until these settings are made, the ARC service remains in a Major alarm state (as it is unable to communicate with the TSM).

1. Access the NMS MI and go to **Archive Node ▶ ARC ▶ Middleware ▶ Configuration ▶ Main**.
2. Use the information in Table 8 to make the appropriate settings.

   As a minimum, you need to customize Server IP, Node Name, User Name, Password, and, depending on the situation, Server Port and Management Class.

   The Number of Sessions, Maximum Retrieval Sessions, and Maximum Archive Sessions have been selected for the case where you have a tape library with 2 drives allocated to the Archive Node.
3. Click **Apply Changes**.

   The Archive Node can now connect to the middleware, and the alarm on Middleware Connectivity clears within a few minutes.

Table 8: ARC ▶ Middleware Component Configuration Settings

Prompt	Type	Description
Middleware Account		
Maximum Retrieve Sessions	Text	The maximum number of concurrent sessions that the ARC can open to the middleware server to retrieve objects. In most cases, set Maximum Retrieve Sessions to the Number of Sessions – Maximum Store Sessions. For example, if the Number of Sessions is 5, and Maximum Store Sessions has its default value of 1, then Maximum Retrieve Sessions should be 4. To share one tape drive for both storage and retrieval, set the Maximum Retrieve Sessions equal to the Number of Sessions. When there are multiple tape drives, making this setting optimizes performance when objects on a Storage Node are being restored from copies on the Archive Node, or if the archive is full and operating in a read-only mode.

Table 8: ARC ▶ Middleware Component Configuration Settings (cont.)

Prompt	Type	Description
Management Class	Text	The name of the default Tivoli Storage Manager (TSM) management class (*default*) assigned to objects by the ARC when it saves them to the TSM. Middleware management classes outline how the middleware's backup and archive operations function, and may be used to specify business rules that are applied by the middleware server. (Such business rules operate independently of the grid's business rules, and must be consistent with the grid's requirement that objects are stored permanently and are always available for retrieval by the Archive Node.) The default management class is used if a management class is not specified for an object when it is saved to the grid (via an FSG Profile as described in the *Administrator Guide*, the StorageGRID API (SGAPI), or CDMI), or if the management class that is specified is not defined on the TSM middleware server. If the default management class you specify here does not exist on the TSM server, then an object saved using this management class is not stored to the middleware archive. Instead the grid retains the object in a queue, and increments the value of CMS ▶ Content ▶ Overview ▶ Objects with ILM Evaluation Pending. Recall that the name of a TSM management class can include no more than 30 characters, is not case-sensitive, and can include only the following characters: • alphabetic characters: A – Z • numerals: 0 – 9 • the following “special” characters: . (period), - (hyphen), + (plus sign), & (ampersand), _ (underscore)
Maximum Store Sessions	Text	The maximum number of concurrent sessions that the ARC can open to the middleware server to store objects. Concurrent store sessions are not supported by the Archive Node. Set Maximum Store Sessions to 1 when the Archive Node is able to store or retrieve objects. Set the value to 0 when the archive managed by the middleware server is full, and the Archive Node can only retrieve objects.
Node Name	Text	Sets the name of the Archive Node, as it appears to the middleware. The name that you enter here must be the same as the node name (*arc-user*) that you registered on the TSM server in “Create a Domain Policy and Register a Node” on page 67.

Table 8: ARC ▸ Middleware Component Configuration Settings (cont.)

Prompt	Type	Description
Number of Sessions	Text	The number of tape drives on the middleware server that are dedicated to the Archive Node. Set this value to be the same as the value for MAXNUMMP (maximum number of mount points) set when the TSM node was registered or updated. (In the register command, the default value of MAXNUMMP used is 1, if no value is set.) The value of MAXSESSIONS for the TSM server must be set to be at least as large as Number of Sessions set here for the ARC. (The default value of MAXSESSIONS on the TSM server is 25). The Archive Node concurrently creates at most one session per mount point plus a small number (<5) of additional sessions.
Password	Text	The value of the password used by the ARC to log in to the TSM middleware. This is the password for the node (*arc-user*) that you registered for the client (*arc-password*) in “Create a Domain Policy and Register a Node” on page 67, or the password of the administrative user you selected for the node. You are prompted to confirm the password after you enter it.
Server IP	Text	Sets the IP address of the middleware server used by the ARC. The default value is 127.0.0.1.
Server Port	Text	Sets the value of the port number on the middleware server that the ARC uses for communications. The default port used clients (such as the Archive Node) to communicate with the TSM server is 1500.
User Name	Text	The user name of the account that the ARC uses to log in to the middleware. By default, when you register a node, the TSM server creates an administrative user ID with client owner authority. If you chose to use this default administrative user, enter the client node name (*arc-user*) as the user name. If you chose to define (or use) a different administrative user for the node, enter its name here.

Set Custom Alarms for the Archive Node

You should establish custom alarms for the following two attributes that are used to monitor the speed and efficiency of data retrieval from the Archive Node:

- ARQL – Average Queue Length. The average time, in microseconds, that an object is queued for retrieval from the middleware.

- ARRL – Average Request Latency. The average time, in microseconds, needed by the Archive Node to retrieve objects via the middleware.

The normal values for these attributes (found in the NMS MI under ARC ▶ Retrieve) depend heavily on how the middleware and storage is configured and used by the customer. For the TSM Archive Node, the values set in the middleware for request timeouts and the number of sessions made available for retrieve requests are particularly influential.

After integration is complete, monitor retrieval from the Archive Node to establish values for normal retrieval times and queue lengths. Then use the NMS MI to create custom alarms for ARQL and ARRL that will notify an administrator of abnormal operating conditions that may require investigation. See the *Administrator Guide* for the procedure on how to create custom alarms.

Create the Initial FSG File System Backup

To create initial conditions for future maintenance of the managed file system, the grid must contain a backup of each replication group. Force an initial file system backup as a safety precaution to ensure you can recover from any problems that might arise prior to the first scheduled backup (the file system is automatically backed up every day during normal operations).

Perform this procedure for each Gateway Node replication group.

1. Find the FSG designated to perform the backup.
    a. In the NMS MI, go to **Grid Management** ▶ **FSG Management** ▶ **<*Replication_Group*>** ▶ **Overview** ▶ **Main**.
    b. Click the link for Backup FSG to go to **FSG** ▶ **Overview** ▶ **Main**.
2. Ensure the backup FSG is running normally:
    a. Verify the FSG State attribute reports Online.
    b. Verify the FSG Status attribute reports No Errors.
    c. Go to **FSG** ▶ **Backup** ▶ **Overview** ▶ **Main**.
    d. Verify that Current Status is Idle.
    e. Verify the Backup Schedule: Next Scheduled Backup is not imminent.

3. Initiate the manual backup:
    a. Go to **FSG ▶ Backup ▶ Configuration**.
    b. Select **Force Manual Backup**.
    c. Click **Apply Changes**.

       The FSG now performs the backup, which stores a file into the grid.
4. Verify the file system backup:
    a. Go to **FSG ▶ Backup**.
    b. Verify the following:
        - The attribute Successful Backups reports 1 or more. If this value is zero (0), check the Current Status attribute of the Current Backup group on this page. If the value is Active, wait for the backup to complete.
        - The Previous Backup section for the Backup Result reports "Successful".
        - Start Time and End Time report a reasonably current time indicating the manual backup.

    If the backup fails, repeat the process to force a backup and re-validate.
5. Check that the backup was stored to the grid:
    a. Go to the grid's **Overview** page.
    b. Verify that the value of the summary attribute Total Managed Objects reports 1 or more objects stored to the grid, representing the backup object(s). (There should be one successful backup for each replication group. If the automatic backup process began while you were completing grid integration, the value may be larger than the number of replication groups.)

       Alternatively, get the value of the backup object identifier from FSG ▶ Backup and look up where the backup object has been stored using <*primary Admin Node*> ▶ CMN ▶ Object Lookup ▶ Configuration.
6. If there is more than one replication group in the deployment, repeat this procedure for each group. Verify that the number of backups continues to increment.
7. Log out of the NMS MI. Click **Logout**.

Next Steps

Continue with the configuration steps listed below.

1. Configure the grid's ILM policy. For more information, see the *Administrator Guide*.
2. Create file system shares. For more information, see the *Administrator Guide*.
3. Customize Gateway Node behavior. This includes configuring content protection, configuring the cache, configuring FSG profiles, and so on. For more information, see the *Administrator Guide*.
4. Configure client-side entities.
5. Verify gateway access. For more information, see Chapter 6: "Verify Client and Grid Integration".

6

Verify Client and Grid Integration

Introduction

To complete the integration of a StorageGRID system, you must verify that the customer's client workstations can access the grid, store content, retrieve content, and delete content (content deletion depends on the grid's business rules).

Verify Client Access

1. Create NFS and/or CIFS fileshares as described in the *Administrator Guide*.
2. Have the customer system administrator perform these tests with your guidance, using a client (and user account) integrated with one of the file shares.
3. The system administrator may want to perform this process for all clients to verify each is correctly integrated.

The test uses source test files that are located on the StorageGRID Software CD. The client system needs a CD drive to access these test files.

Test cases must be executed sequentially in order to preserve dependencies between test cases.

Set up Test

CIFS Integration

Use the following procedure to prepare the integration test if the client application connects to the StorageGRID system via CIFS shares:

1. If the client computer is using a client-side firewall (such as Zone-Alarm), ensure that access is enabled to the gateway IP addresses of the StorageGRID system.
2. From the client workstation, ping the primary FSG to verify connectivity:
    a. From the Windows **start** button, select **Run**.
    b. In the Run dialog box, enter: **`cmd.exe`**
    c. Enter: **`ping <IP_address>`**

    where <*IP_address*> is the client-side IP address of the FSG. Verify that the server responds, indicating connectivity.

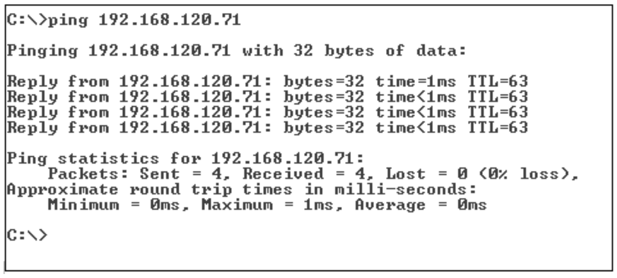

```
C:\>ping 192.168.120.71

Pinging 192.168.120.71 with 32 bytes of data:

Reply from 192.168.120.71: bytes=32 time=1ms TTL=63
Reply from 192.168.120.71: bytes=32 time<1ms TTL=63
Reply from 192.168.120.71: bytes=32 time<1ms TTL=63
Reply from 192.168.120.71: bytes=32 time<1ms TTL=63

Ping statistics for 192.168.120.71:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 0ms, Maximum = 1ms, Average = 0ms

C:\>
```

Figure 8: Sample PING Verification

    d. Close the Command Prompt window. Enter: **exit**
3. Map a network drive to the share:
    a. From the Windows desktop, click **My Computer**.
    b. From the **Tools** menu, select **Map Network Drive**.
    c. In the **Drive** box, select an unused letter to assign to the StorageGRID system gateway folder.
    d. In the **Folder** box, type the CIFS share name in the form:

    `\\<IP_address>\<share_name>`

    where <*IP_address*> is the address used in step **2c** above and <*share_name*> is the share name used for CIFS integration. The example below shows \\192.168.180.71\myDirectory.

The share name does not use the \fsg path.

You may not want to reconnect a secondary FSG share at logon.

e. Select **Reconnect at logon**.
f. Click **Finish**.
g. If prompted, enter a User name and Password and click **OK**.

4. Repeat the process for each secondary FSG.

   You can now perform the “Test File Ingest” on page 77.

NFS Integration

Use the following procedure to prepare the integration test if the client application connects to the StorageGRID system via NFS shares. This test assumes you are at a command shell of a Unix/Linux system.

1. Verify connectivity to the primary FSG. Enter: `ping <IP_address>`

   (or variant for the client system)

   where <*IP_address*> is the client-side IP address of the Gateway Node. Verify that the server responds, indicating connectivity.
2. Mount the primary FSG writable share using a command appropriate to the client OS. A sample Linux command is:

   ```
   mount -t nfs -o hard,intr <IP_address>:/fsg/<share> <myGrid>
   ```

   Use the IP address of the primary FSG and the share name for the client directory. The mount point can be any name selected by the client (*myGrid* in the command above).
3. Repeat the process for each secondary FSG. The mount point must be distinct from the first primary.

You can now perform the “Test File Ingest” below.

Test File Ingest

To verify that the StorageGRID system allows files to be stored through a primary FSG file share:

1. At the test workstation, insert the StorageGRID Software CD and open the testdata directory which contains the test files Sample_X-Ray.jpg and Sample-MRI.jpg.
2. Log in to the NMS MI.

3. Record the number of objects stored in the system on the Overview Summary page:
   - Total Managed Objects
   - Total Files Stored to Grid – Successful

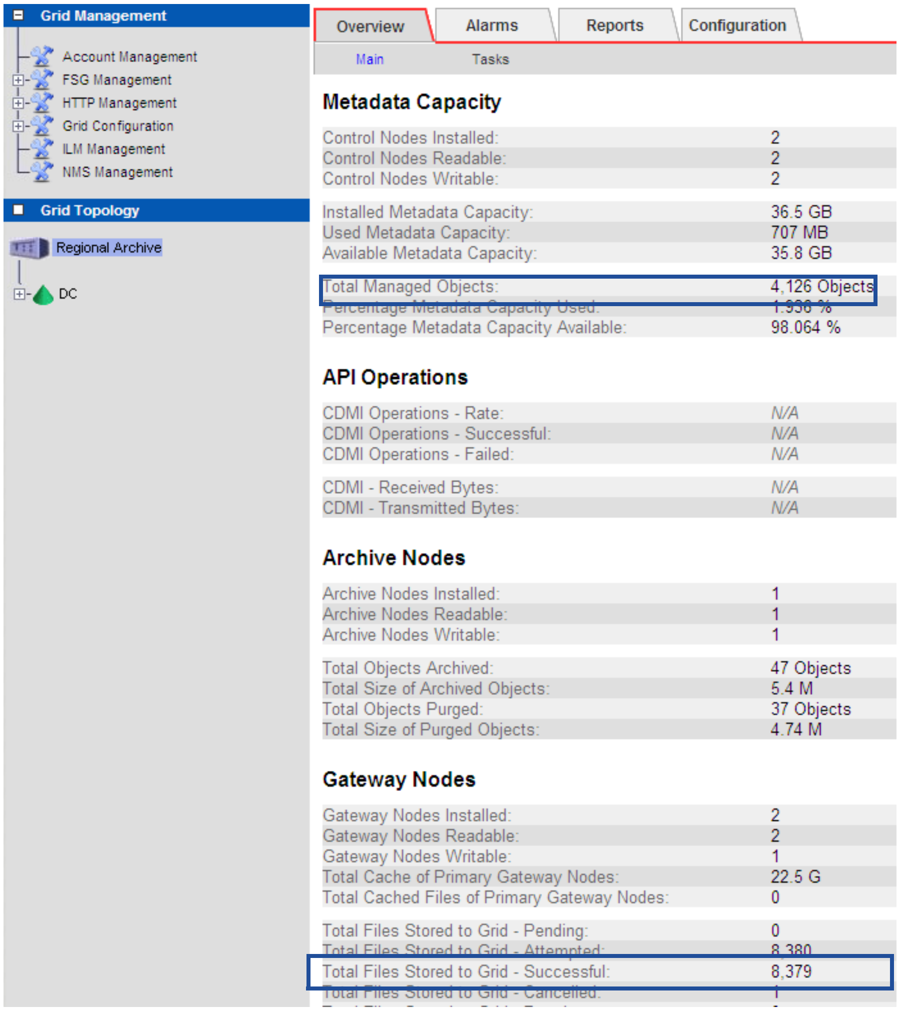

Metadata Capacity

Control Nodes Installed:	2
Control Nodes Readable:	2
Control Nodes Writable:	2
Installed Metadata Capacity:	36.5 GB
Used Metadata Capacity:	707 MB
Available Metadata Capacity:	35.8 GB
Total Managed Objects:	4,126 Objects
Percentage Metadata Capacity Used:	1.936 %
Percentage Metadata Capacity Available:	98.064 %

API Operations

CDMI Operations - Rate:	N/A
CDMI Operations - Successful:	N/A
CDMI Operations - Failed:	N/A
CDMI - Received Bytes:	N/A
CDMI - Transmitted Bytes:	N/A

Archive Nodes

Archive Nodes Installed:	1
Archive Nodes Readable:	1
Archive Nodes Writable:	1
Total Objects Archived:	47 Objects
Total Size of Archived Objects:	5.4 M
Total Objects Purged:	37 Objects
Total Size of Purged Objects:	4.74 M

Gateway Nodes

Gateway Nodes Installed:	2
Gateway Nodes Readable:	2
Gateway Nodes Writable:	1
Total Cache of Primary Gateway Nodes:	22.5 G
Total Cached Files of Primary Gateway Nodes:	0
Total Files Stored to Grid - Pending:	0
Total Files Stored to Grid - Attempted:	8,380
Total Files Stored to Grid - Successful:	8,379
Total Files Stored to Grid - Cancelled:	1

Figure 9: Grid Summary Page: Verifying Ingest

NOTE If the integration test is being performed for an existing grid with client activity on other gateway replication groups, use the value of “Files Stored to Grid – Successful” for the primary FSG being tested (on ▶FSG ▶Storage ▶Overview) instead.

4. Copy the two test files from the StorageGRID Software CD to the share. Copying the files to the share automatically stores the files into the StorageGRID system.
   - For CIFS, drag the files using Windows Explorer from the CD to the mapped share.

- For NFS, enter: **`cp <cd_dev/testdata/>*.jpg <myGrid>`**

  where <*myGrid*> is the mount point.

  - Verify the files are visible on the primary FSG share. Enter: **`ls <myGrid>/*`**

5. In the NMS MI, verify that these two values increased by 2:
   - Total Files Stored to Grid – Successful
   - Total Managed Objects

   It may take a couple of minutes for the value to change. Stored files may appear as Total Files Stored to Grid – Pending briefly, before the store is complete.

6. Verify the files are “mirrored” in each secondary FSG:
   - For CIFS, look at the secondary FSG Explorer window. You may need to refresh the window by pressing **<F5>**.
   - For NFS, enter: **`ls <mySecondary>/Sample*.jpg`**

     where <*mySecondary*> is the mount point of the secondary FSG.

Test File Retrieval

Verify that the StorageGRID system allows files to be retrieved through the primary and secondary FSG file shares.

1. Confirm that the test files have been successfully stored to the grid. See “Test File Ingest” on page 77.
2. In the NMS MI, go to the Overview Summary page and note the value of Total Files Retrieved from Grid – Successful.

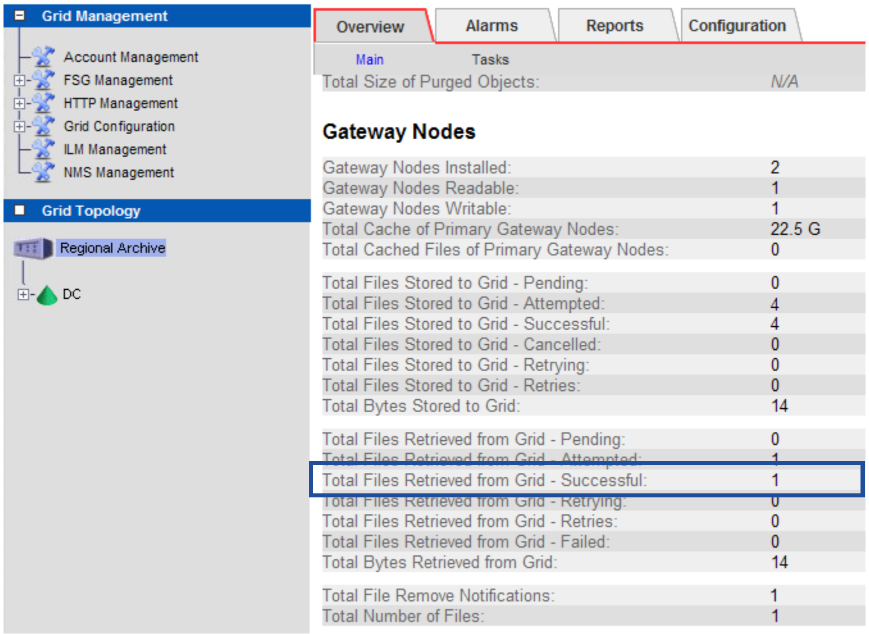

Total Size of Purged Objects:	N/A

Gateway Nodes

Gateway Nodes Installed:	2
Gateway Nodes Readable:	1
Gateway Nodes Writable:	1
Total Cache of Primary Gateway Nodes:	22.5 G
Total Cached Files of Primary Gateway Nodes:	0
Total Files Stored to Grid - Pending:	0
Total Files Stored to Grid - Attempted:	4
Total Files Stored to Grid - Successful:	4
Total Files Stored to Grid - Cancelled:	0
Total Files Stored to Grid - Retrying:	0
Total Files Stored to Grid - Retries:	0
Total Bytes Stored to Grid:	14
Total Files Retrieved from Grid - Pending:	0
Total Files Retrieved from Grid - Attempted:	1
Total Files Retrieved from Grid - Successful:	1
Total Files Retrieved from Grid - Retrying:	0
Total Files Retrieved from Grid - Retries:	0
Total Files Retrieved from Grid - Failed:	0
Total Bytes Retrieved from Grid:	14
Total File Remove Notifications:	1
Total Number of Files:	1

Figure 10: Grid Summary Page: Verifying Retrieval

NOTE If the integration test is being performed for an existing grid with client activity on other gateway replication groups, use the value of “Files Retrieved from Grid – Successful” for the primary FSG being tested (on ▶FSG ▶Storage ▶Overview) instead.

3. Copy the Sample_X-Ray.jpg test file from the primary FSG share to a temporary location:
   - For CIFS, use Windows Explorer to drag the test file from the primary FSG share to the desktop.
   - For NFS, enter: **`cp <myGrid>/Sample_X-Ray.jpg <tmp>/`**

     where <*myGrid*> is the mount point and <*tmp*> is an available temporary directory on the client system.

     Verify that the files appear in the local temporary file system. Enter: **`ls <tmp>/Sample*.jpg`**
4. Compare the source file on StorageGRID Software CD to the file in the temporary location to confirm that both copies are identical. The file retrieved from the grid should be an exact duplicate of the file that was stored in the first test case.
5. If the replication group includes a secondary FSG, repeat step **3** and step **4** for the Sample-MRI.jpg test file.

6. In the NMS MI, go to the Overview Summary page and confirm that the value of Total Files Retrieved from Grid – Successful has increased by 2.
7. Delete the temporary copies from the client's system:
   - For CIFS, delete the test files from the Windows desktop.
   - For NFS, enter:

     ```
     rm <tmp>/Sample_X-Ray.jpg
     rm <tmp>/Sample-MRI.jpg
     ```

Test File Deletion

What happens during the file deletion test depends on how content protection is configured on the grid, that is, whether:

- Deletion is permitted
- WORM is enabled
- File recovery is enabled

NOTE File recovery is deprecated and no longer supported.

Deletion Is Permitted

Follow this procedure to verify that files can be deleted from a primary FSG file share when file deletion is permitted, that is, when WORM is *not* enabled.

1. Confirm that the test files have been successfully stored to the grid. See "Test File Ingest" on page 77.
2. In the NMS MI, go to the grid's Overview Summary page and note the value of Total File Remove Notifications.

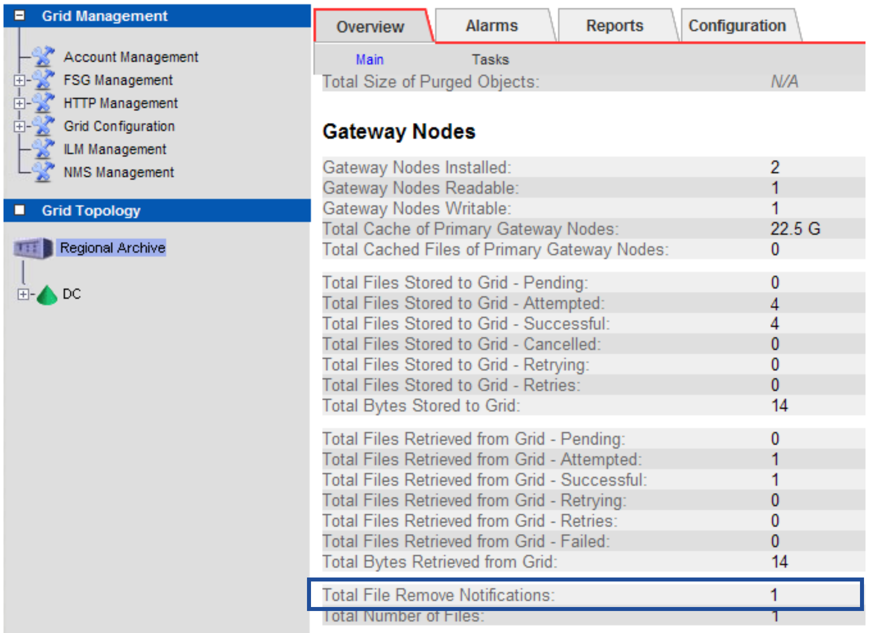

Total Size of Purged Objects:	N/A

Gateway Nodes

Gateway Nodes Installed:	2
Gateway Nodes Readable:	1
Gateway Nodes Writable:	1
Total Cache of Primary Gateway Nodes:	22.5 G
Total Cached Files of Primary Gateway Nodes:	0
Total Files Stored to Grid - Pending:	0
Total Files Stored to Grid - Attempted:	4
Total Files Stored to Grid - Successful:	4
Total Files Stored to Grid - Cancelled:	0
Total Files Stored to Grid - Retrying:	0
Total Files Stored to Grid - Retries:	0
Total Bytes Stored to Grid:	14
Total Files Retrieved from Grid - Pending:	0
Total Files Retrieved from Grid - Attempted:	1
Total Files Retrieved from Grid - Successful:	1
Total Files Retrieved from Grid - Retrying:	0
Total Files Retrieved from Grid - Retries:	0
Total Files Retrieved from Grid - Failed:	0
Total Bytes Retrieved from Grid:	14
Total File Remove Notifications:	1
Total Number of Files:	1

Figure 11: Total File Remove Notifications on the Grid Summary

---

NOTE If the integration test is being performed for an existing grid with client activity on other gateway replication groups, use the value of “File Remove Notifications” for the primary FSG being tested (on ▶FSG ▶Storage ▶Overview) instead.

---

3. Delete the test file Sample_X-Ray.jpg from the primary FSG.
   - For CIFS, delete the file from the Windows Explorer share window.
   - For NFS, enter: **`rm <myGrid>/Sample_X-Ray.jpg`**

     where <*myGrid*> is the mount point.
4. In the NMS MI, verify that the value of Total File Remove Notifications has increased by 1. It may take a couple of minutes for the value to change.
5. Verify that the file has disappeared from the FSG share:
   - For CIFS, use Windows Explorer to confirm that the file has disappeared from the primary FSG and secondary FSG shares. Press **<F5>** to refresh the windows as needed.
   - For NFS, enter: **`ls <myGrid>/Sample*.jpg`**

     The Sample-MRI.jpg test file should not be listed.

     Verify the file deleted from the grid has also disappeared from the secondary FSG shares. Confirm that the listing on the primary and secondary FSGs is the same.

WORM Is Enabled

Follow this procedure to confirm that files cannot be deleted from the primary FSG when WORM is enabled.

1. Confirm that the test files have been successfully stored to the grid. See "Test File Ingest" on page 77.
2. Go to **Grid Management ▶ FSG Management ▶ <*Replication Group*> ▶ Overview ▶ Main** and confirm that WORM is Enabled and that the protection period is not set so small as to effectively disable WORM for that profile.
3. Delete the test file Sample_X-Ray.jpg from the primary FSG. See step **3** on page 82.

The value of File Remove Notifications in the NMS MI does not change.

4. Confirm that the Permission Denied message appears.

   For CIFS, press **F5** to refresh the window and confirm the file is still present in the share. In Windows XP or later, the Denied message does not appear even though the file is preserved.

File Recovery Is Enabled

Follow this procedure to confirm that files deleted from the FSG are transferred to the share's recovery directory when file recovery is enabled.

1. Confirm that the test files have been successfully stored to the grid. See "Test File Ingest" on page 77.
2. Go to **Grid Management ▶ FSG Management ▶ <*Replication Group*> ▶ Configuration ▶ Profiles** and confirm that File Recovery is Enabled for the FSG profile.
3. Delete the test file Sample_X-Ray.jpg from the primary FSG. See step **3** on page 82.

The value of File Remove Notifications in the NMS MI does not change.

4. Verify that the file has disappeared from the FSG share. See step **5** on page 82.
5. Verify that the file has been transferred to the recovery directory. The recovery directory is named <*share*>-recovery, where <*share*> is the name of the parent directory. For example, the recovery directory of the share fsg/myshare is /fsg/myshare-recovery.
   - If a share was created for the <*share*>*-recovery* directory using either the CIFS or NFS configuration utility, mount the recovery share and look for the file.
   - If a share was not created for the <*share*>*-recovery* directory, log in to a command shell on the FSG and look for the file under the /fsg/<*share*>-recovery directory.

Test Secondary FSG Read-Only Permissions

This test verifies that the secondary FSG is read-only and does not permit the writing or deleting of files. An HAGC may not include a separate secondary FSG. In this case, skip this test.

1. Attempt to delete the file Sample-MRI.jpg from the secondary FSG. A Permission Denied message should result.
   - For CIFS, use Windows Explorer to delete the file from the secondary FSG share.

   NOTE In Windows XP or later, this action does not result in an Access Denied message. Press **<F5>** to refresh Explorer and confirm that the file is still present.

   - For NFS, enter: `rm <mySecondary>/Sample-MRI.jpg`
2. Attempt to copy the file Sample_X-Ray.jpg to the secondary FSG. A Permission Denied message should result.
   - For CIFS, use Windows Explorer to copy the file to the secondary FSG share.

   NOTE In Windows XP or later, this action does not result in an Access Denied message.

   - For NFS, enter:

     ```
     cp <cd_dev/testdata/>Sample_X-Ray.jpg <mySecondary>
     ```

Repeat this test for each secondary FSG.

Test High Availability Gateway Clusters

If desired, repeat the tests from the supplementary FSG for the HAGC.

1. Fail over from the main FSG to the supplementary FSG. Follow the instructions in the *Maintenance Guide*.
2. Perform the entire verification test.
3. Fail back from the supplementary FSG to the main FSG, making it the active primary.

Conclude the Test

To conclude the test, leave a clean file system for the end user client.

1. If deletion is permitted, delete the remaining test files from the primary FSG share:
    - For CIFS, use Windows Explorer to delete the file.
    - For NFS, enter: `rm <myGrid>/sample*.jpg`
2. If file recovery is enabled, remove the file from the recovery directory.
3. If deletion is not permitted (WORM mode is enabled), ask the customer for their preference. If they prefer that you remove the test files, temporarily disable WORM mode and remove the test files:
    a. Follow the procedures in the *Administrator Guide* to disable WORM mode on the file share.
    b. Delete the remaining test file from the primary FSG share.
    c. Re-enable WORM mode on the share, as described in the *Administrator Guide*.

WARNING **Ensure that you re-enable WORM mode. Customer data is at risk until deletion protection is restored.**

4. For CIFS, close all Explorer windows.
5. Log out of the NMS MI and close the browser window.
6. Remove the StorageGRID Software CD from the client computer.

Test the Audit Client

If the Audit option is deployed, you must set up and verify the audit client.

To perform the verification, create an NFS or CIFS audit fileshare as described in the *Administrator Guide*. Have the customer system administrator perform these tests with your guidance, using a client.

Perform the test for each grid node (Admin Node or Audit Node) that hosts an AMS service.

CIFS Integration

The connection setup of the audit client is the same as for FSG shares, substituting the IP address of the Admin Node or Audit Node and using **audit-export** as <*Share_Name*>.

The audit share is read-only. Log files are intended to be read by computer applications; verification does not include opening a file. It is considered sufficient verification that the audit log file(s) appear in a Windows Explorer window. Following connection verification, close all windows.

NFS Integration

Mount and Verify Audit Client Share

1. Verify connectivity. Enter: `ping <IP_address>`

   (or variant for the client system) using the client-side IP address of the grid node (Admin Node or Audit Node) hosting the AMS service. Verify that the server responds, indicating connectivity.
2. Mount the audit read-only share using a command appropriate to the client OS. A sample Linux command is (enter on one line):

   ```
   mount -t nfs -o hard,intr <Admin_Node_IP_address>:/var/
   local/audit/export <myAudit>
   ```

   Use the IP address of the grid node (Admin Node or Audit Node) hosting the AMS service and the predefined share name for the audit system. The mount point can be any name selected by the client (for example, *myAudit* in the command above).
3. Verify the files are available from the audit share. Enter:

   ```
   ls <myAudit>/*
   ```

   where <*myAudit*> is the mount point of the audit share. There should be at least one log file listed.

A

Grid Specification Files and Provisioning

What is Provisioning

Provisioning is the process of turning a grid design into the collection of files needed to create, expand, maintain, or upgrade the grid. That collection of files, referred to as the GPT (grid provisioning tool) repository, includes the SAID package. The key input for provisioning is the grid specification file.

About the SAID Package

The Software Activation and Integration Data (SAID) package contains site-specific files for the grid. It is generated during the provisioning process as a zip file and is named using the following naming convention:

GID<*grid_ID*>_REV<*revision_number*>_SAID.zip

The SAID package contains the following items:

Table 9:

Item	Description
Doc directory	Contains html files used to confirm provisioning.
Escrow_Keys directory	Encryption keys used by the Data Recovery Tool.
Grid_Activation directory	Contains activation files, one for each server. Activation files are named <*servername*>-autoinst.xml. Activation files are keyed to work with the hardware used for the grid deployment and the version of StorageGRID software.

Table 9:

Item	Description
Configuration.txt	Lists grid-wide configuration and integration data generated during the provisioning process.
Grid_Tasks directory	Contains files created by some types of changes to the grid specification file, such as adding a server or converting the grid to use metadata replication. Grid tasks are used to trigger various actions within the grid that are required to implement the specified changes to the grid.
Grid specification file	XML file that encapsulates the grid design. File name is: GID<*grid_ID*>_REV<*revision_number*>_GSPEC.xml
Passwords.txt	Passwords used to access the grid.

Grid Configuration Files

The Doc directory of the SAID package contains html files documenting the specifications of the grid's configuration. Use these pages to confirm that the grid configuration is correct and complete.

The index.html can only be opened in a Windows Internet Explorer browser.

- Click the <*SAID_package*>/doc/index.html file. For an example, see Figure 12 below.

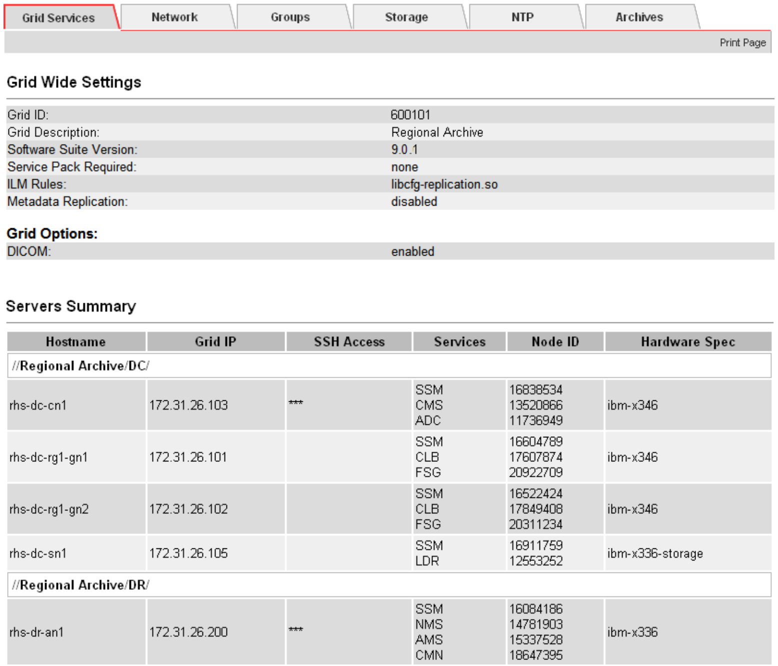

Grid Wide Settings

Grid ID:	600101
Grid Description:	Regional Archive
Software Suite Version:	9.0.1
Service Pack Required:	none
ILM Rules:	libcfg-replication.so
Metadata Replication:	disabled

Grid Options:

DICOM:	enabled

Servers Summary

Hostname	Grid IP	SSH Access	Services	Node ID	Hardware Spec
//Regional Archive/DC/					
rhs-dc-cn1	172.31.26.103	***	SSM CMS ADC	16838534 13520866 11736949	ibm-x346
rhs-dc-rg1-gn1	172.31.26.101		SSM CLB FSG	16604789 17607874 20922709	ibm-x346
rhs-dc-rg1-gn2	172.31.26.102		SSM CLB FSG	16522424 17849408 20311234	ibm-x346
rhs-dc-sn1	172.31.26.105		SSM LDR	16911759 12553252	ibm-x336-storage
//Regional Archive/DR/					
rhs-dr-an1	172.31.26.200	***	SSM NMS AMS CMN	16084186 14781903 15337528 18647395	ibm-x336

Figure 12: Index.html File

About Grid Specification Files

The grid specification file is an XML file that encapsulates the configuration information needed to install, expand, and maintain a grid. The file includes topology, servers, options, and networking details for the grid.

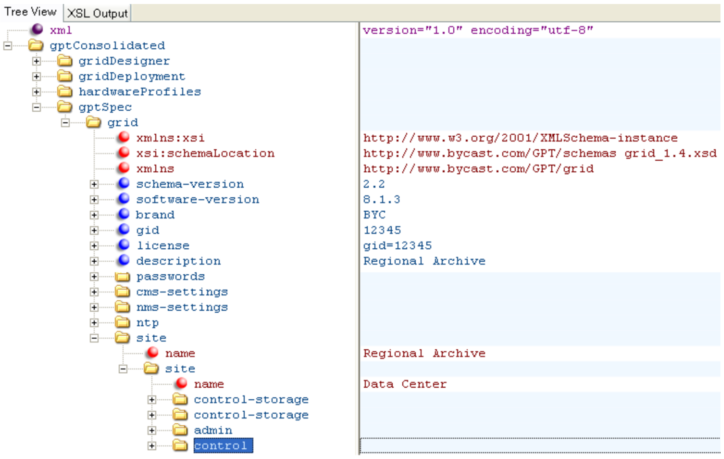

```
xml                                         version="1.0" encoding="utf-8"
gptConsolidated
    gridDesigner
    gridDeployment
    hardwareProfiles
    gptSpec
        grid
            xmlns:xsi                       http://www.w3.org/2001/XMLSchema-instance
            xsi:schemaLocation              http://www.bycast.com/GPT/schemas grid_1.4.xsd
            xmlns                           http://www.bycast.com/GPT/grid
            schema-version                  2.2
            software-version                8.1.3
            brand                           BYC
            gid                             12345
            license                         gid=12345
            description                     Regional Archive
            passwords
            cms-settings
            nms-settings
            ntp
            site
                name                        Regional Archive
                site
                    name                    Data Center
                    control-storage
                    control-storage
                    admin
                    control
```

Figure 13: Grid Specification File in XML Notepad 2007

All new grid specification files are created and deployed using Grid Designer. As well, all grid specification files updated to StorageGRID 9.0 are edited and deployed using Grid Designer. For more information, see the *Grid Designer User Guide*.

Grid Specification File Stages

The grid specification file goes through a number of stages as the grid is designed and then installed:

- Default grid specification file — The default grid specification file describes the basic grid topology and grid configuration.
- Deployment grid specification file — The deployment grid specification file is created from the default grid specification file by updating the grid specification file with customer-specific data, for example IP addresses
- Provisioned grid specification file — The provisioned grid specification file is created when the provision command is run.

These stages are summarized in Figure 14 below.

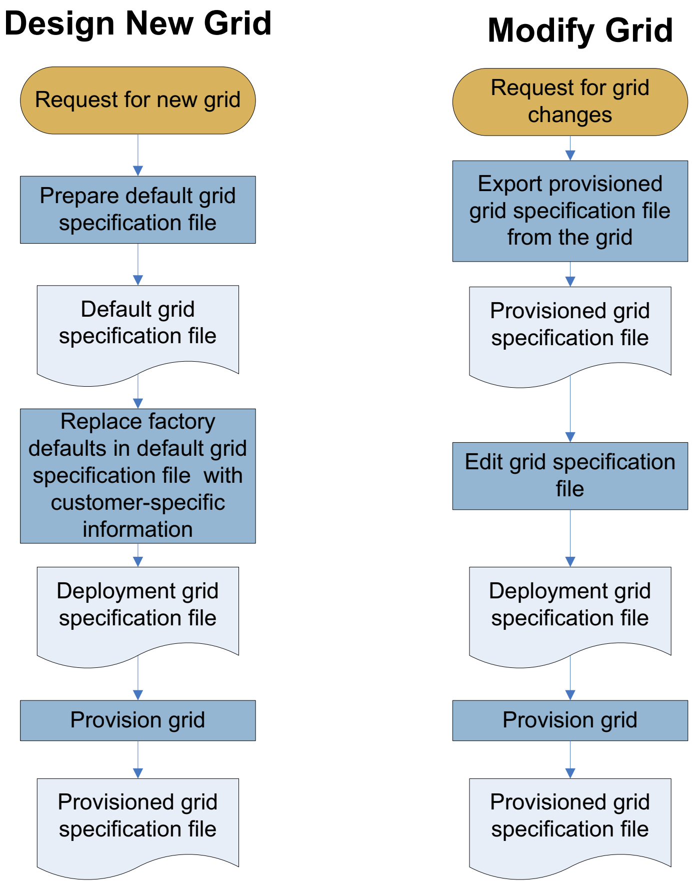
Design New Grid
Request for new grid
Prepare default grid specification file
Default grid specification file
Replace factory defaults in default grid specification file with customer-specific information
Deployment grid specification file
Provision grid
Provisioned grid specification file
Modify Grid
Request for grid changes
Export provisioned grid specification file from the grid
Provisioned grid specification file
Edit grid specification file
Deployment grid specification file
Provision grid
Provisioned grid specification file

Figure 14: Editing the Grid Specification File

Naming Convention

Grid specification files use the naming convention GID<*grid_ID*>_REV<*revision_number*>_GSPEC.xml, where <*grid_ID*> refers to the grid's unique identifier and <*revision_number*> refers to the revision number of the grid specification file, for example, GID1234_REV1_GSPEC.xml.

The default grid specification file has a <*revision_number*> of zero (REV0). The revision number is increased by 1 each time the grid specification file is modified, for example to add servers, change IP addresses, or refresh hardware. For the initial installation of the grid, the revision number must be 1 (REV1) — that is, the default grid specification file has been modified once for the installation of the StorageGRID system. Any other revision number will cause provisioning to fail.

Grid Specification File Structure

Grid specification files are edited and deployed using Grid Designer. The following section describes the xml structure of the grid specification file.

Server Names

To review grid information using the grid specification file, you must ensure that you select the correct server. Table 10 below lists the server tags. In addition, check the server name (gptSpec>grid>site>site>*server*>name) to confirm that you are using the correct server.

Table 10: Server Tags

Server	XML Tag
Admin Node	admin
API Gateway Node	gateway
Archive Node	archive
Audit Node	custom
Control Node	control
Gateway Node	gateway
Storage Node	storage

Tags

Table 11 below lists the XML tags of the attributes most likely to be reviewed or updated.

Table 11: Common Changes to Grid Specification Files

Setting	XML Tag	Notes
External NTP time sources 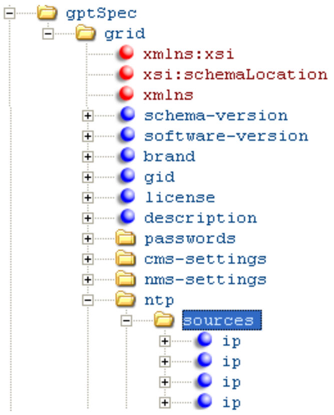 gptSpec grid xmlns:xsi xsi:schemaLocation xmlns schema-version software-version brand gid license description passwords cms-settings nms-settings ntp sources ip ip ip ip	gptSpec>grid>ntp>sources>ip	To provide a stable time source, it is recommended that four NTP time servers be used. External time sources must use the NTP protocol and not the SNTP protocol. In particular, do not use the Windows Time Service: it does not provide enough synchronization accuracy because it uses SNTP.
Networking information in a grid that does not have a private network 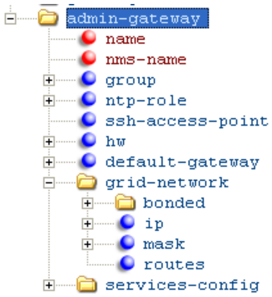 admin-gateway name nms-name group ntp-role ssh-access-point hw default-gateway grid-network bonded ip mask routes services-config	gptSpec>grid>site>site>*server*>default-gateway	
	gptSpec>grid>site>site>*server*>grid-network>ip	
	gptSpec>grid>site>site>*server*>grid-network>mask	
	gptSpec>grid>site>site>*server*>grid-network>routes	
NMS Entity Name 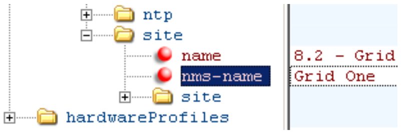 ntp site name nms-name site hardwareProfiles 8.2 - Grid Grid One	gptSpec>grid>site>site>nms-name	Adds the attribute nms-name="<name>" to the element tags for grid, site, and server. For example: **`<site name="grid name" nms-name= "grid one">`**

Table 11: Common Changes to Grid Specification Files (cont.)

Setting	XML Tag	Notes
Networking information for grid communication in grid with a private network 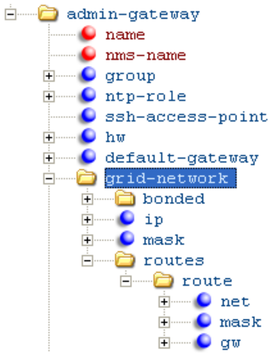 admin-gateway name nms-name group ntp-role ssh-access-point hw default-gateway grid-network bonded ip mask routes route net mask gw	gptSpec>grid>site>site> *server*>default-gateway	
	gptSpec>grid>site>site> *server*>grid-network>ip	
	gptSpec>grid>site>site> *server*>grid-network>mask	
	gptSpec>grid>site>site> *server*>grid-network>routes	
Networking information for client access in a grid with a private network 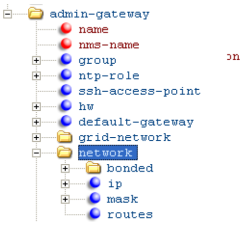 admin-gateway name nms-name group ntp-role ssh-access-point hw default-gateway grid-network network bonded ip mask routes	gptSpec>grid>site>site> *server*>default-gateway	Servers that have client-side IP addresses are Admin Nodes, Gateway Nodes, and Archive Nodes.
	gptSpec>grid>site>site> *server*>network>ip	
	gptSpec>grid>site>site> *server*>network>mask	
	gptSpec>grid>site>site> *server*>network>routes	
Virtual IP address of Gateway Node cluster 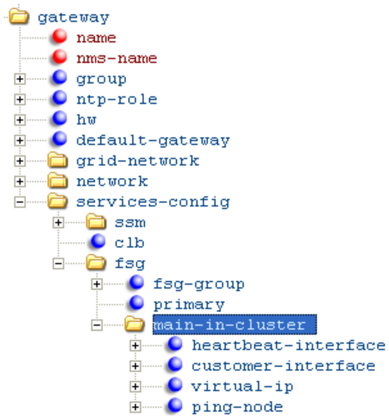 gateway name nms-name group ntp-role hw default-gateway grid-network network services-config ssm clb fsg fsg-group primary main-in-cluster heartbeat-interface customer-interface virtual-ip ping-node	gptSpec>grid>site>site> gateway>services-config> fsg>main-in-cluster>virtual-ip	Virtual IP addresses are used with high availability clusters.

Table 11: Common Changes to Grid Specification Files (cont.)

Setting	XML Tag	Notes
Heartbeat IP addresses 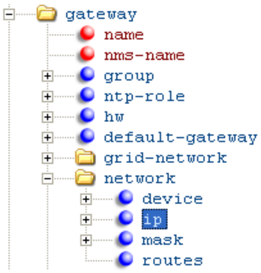 gateway name nms-name group ntp-role hw default-gateway grid-network network device ip mask routes	gptSpec>grid>site>site>gateway>network>ip	Heartbeat IP addresses are only used with High Availability Gateway Clusters. They do not need to be modified unless they conflict with another network. If necessary, substitute the 10.1.1.x network with another unused non-routeable network.

View a Copy of the Grid Specification File

Follow this procedure if you need to quickly check the grid specification file. If you need to edit the file, use the procedure in "Export the Latest Grid Specification File" on page 95 to obtain a copy of the file.

1. In the NMS MI, go to **Grid Management ▶ Grid Configuration ▶ Configuration ▶ Main**.

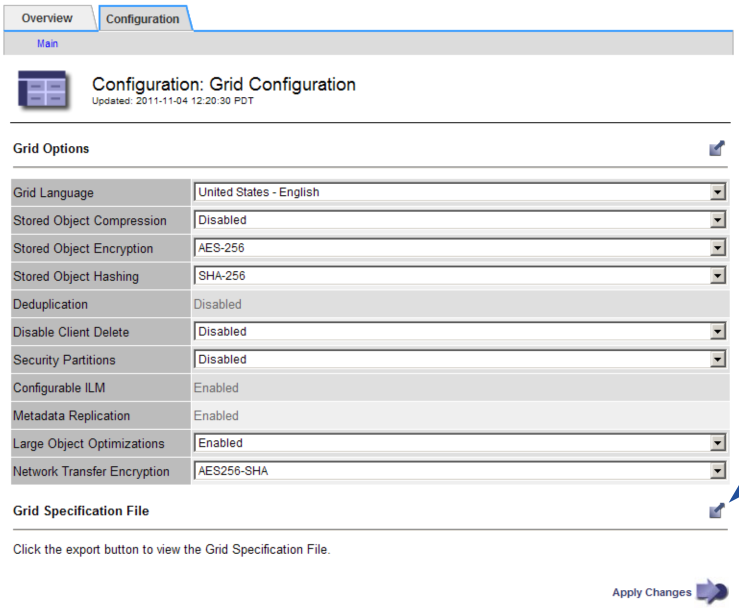

Configuration: Grid Configuration
Updated: 2011-11-04 12:20:30 PDT

Grid Options

Grid Language	United States - English
Stored Object Compression	Disabled
Stored Object Encryption	AES-256
Stored Object Hashing	SHA-256
Deduplication	Disabled
Disable Client Delete	Disabled
Security Partitions	Disabled
Configurable ILM	Enabled
Metadata Replication	Enabled
Large Object Optimizations	Enabled
Network Transfer Encryption	AES256-SHA

Grid Specification File

Click the export button to view the Grid Specification File.

View Grid Specification File

Figure 15: View the Grid Specification File

2. Click **Export** at the bottom of the page, next to Grid Specification File. A new browser window opens, showing the grid specification file in raw XML.

Export the Latest Grid Specification File

Admin Node Hosted on a Virtual Machine

Prerequisites and Required Materials

- Passwords.txt file
- a utility such as WinImage (available at http://www.winimage.com), that permits you to create a floppy disk image
- a tool such as WinSCP (available at http://winscp.net/eng/download.php) to transfer files to and from the Admin Node
- Service laptop

Procedure

1. At the primary Admin Node server, access a command shell and log in as root using the password listed in the Passwords.txt file.
2. Create a directory to hold the provisioned grid specification file. Enter: **`mkdir -p /root/usb`**
3. Copy the provisioned grid specification file to the directory. Enter: **`copy-grid-spec /root/usb`**
4. Use WinSCP to copy the GID<*Grid_ID*>_REV<*revision_number*>_GSPEC.xml file from the Admin Node to your service laptop.

   Alternatively, copy the file from the Admin Node to your service laptop using a floppy image. Create the floppy image on your service laptop, connect to it in the vSphere client, and then mount it under Linux as /media/floppy. Copy the grid specification file to the floppy image from the Admin Node virtual machine, unmount it, and then extract the files to your laptop.
5. Log out. Enter: **`exit`**

Admin Node Hosted on a Physical Server

NOTE New installations of the StorageGRID 9.0 system are not supported on physical servers. Virtual machines must be used.

Prerequisites and Required Materials

- USB flash drive
- Passwords.txt file

Procedure

1. At the primary Admin Node server, access a command shell and log in as root using the password listed in the Passwords.txt file.
2. Insert a USB flash drive.
3. Copy the provisioned grid specification file to the USB flash drive. Enter: **`copy-grid-spec`**
4. Log out. Enter: **`exit`**

Provision the Grid

Use this procedure to implement the changes made to the grid specification file. The provisioning script imports the updated grid specification file into the grid and generates any grid tasks required to complete the implementation of the changes.

On a Primary Admin Node on a Virtual Machine

Prerequisites and Required Materials

- Deployment grid specification file. See "Export the Latest Grid Specification File" on page 95.
- A utility such as WinImage (available at http://www.winimage.com), that permits you to create a floppy disk image
- Passwords.txt file
- Provisioning passphrase
- A tool such as WinSCP (available at http://winscp.net/eng/download.php) to transfer files to and from the Admin Node
- Service laptop

Procedure

1. Create a floppy image that contains the deployment grid specification file:
   a. Start the WinImage software on your service laptop.
   b. From the **File** menu, select **New**. In the Format selection dialog, select a standard format **1.44 MB** floppy. Click **OK**.

   c. From the **Image** menu, select **Inject**. Browse for the grid specification file, and select **Open**. When prompted, confirm that you want to inject the file. Select **Yes**.

2. Save the floppy image:

   a. From the **File** menu, select **Save**.

   b. In the **Save** dialog, browse to the destination folder.

   c. Select **Save as** type: **Virtual floppy Image (*.vfd,*.flp)**

   d. Enter a filename ending in .flp. For example, `<`*`servername`*`>.flp`

      You must enter the extension, or the vSphere client cannot use the image during installation.

   e. Click **Save**.

3. At the primary Admin Node, log in as root. When prompted for a password, press **<Enter>**.

4. In vSphere Client, connect the Provisioning floppy image by clicking the Connect/Disconnect the floppy devices of the virtual machine icon and selecting **Connect Floppy 1 ▶ Connect to floppy image on local disk**.

5. Click in the vSphere console window to return to the command line.

6. Copy the GID<*Grid_ID*>_REV<*revision_number*>_GSPEC.xml file from the deployment floppy image to the primary Admin Node:

   a. Mount the floppy image. Enter: **`mount /media/floppy`**

   b. Copy its contents to the Admin Node. Enter:

```
mkdir /root/usb
cp /media/floppy/* /root/usb
```

   c. Unmount the floppy image. It is no longer needed. Enter:

```
umount /media/floppy
```

7. Remove any old grid specification files from /root/usb.

   Ensure that there is only one file named GID<*grid_ID*>_REV<*revision_number*>_GSPEC.xml.

NOTE The /root/usb directory must contain only one grid specification file. Otherwise, provisioning will fail.

---

8. Run the provisioning script:

   a. At the primary Admin Node server, access a command shell and log in as root using the password listed in the Passwords.txt file.

   b. Run the provisioning script. Enter: **`provision /root/usb`**

c. When prompted, enter the provisioning passphrase.

When the process is complete, “Provisioning complete” is displayed.

NOTE If provisioning ends with an error message, see “Provisioning Troubleshooting” on page 105.

9. Back up the provisioning data to another directory on the Admin Node.
   a. Create a directory for the backup provisioning data. Enter:
      `mkdir -p /var/local/backup`
   b. Back up the provisioning data. Enter:
      `backup-to-usb-key /var/local/backup`
   c. When prompted, enter the provisioning passphrase.
10. Store the Provisioning directory (found at /root/usb) and the Backup Provisioning directory (found at /var/local/backup) separately in a safe place. For example, use WinSCP to copy these directories to your service laptop, and then store them to two separate USB flash drives that are stored in two separate and secure locations. For more information, see “Preserving Copies of the Provisioning Data” on page 104.

    The contents of the Provisioning directory are used during expansion and maintenance of the grid when a new SAID package must be generated.

WARNING **Store copies of the Provisioning directory in two separate and secure locations. The Provisioning directories contain encryption keys and passwords that can be used to obtain data from the grid. The Provisioning directory is also required to recover from a primary Admin Node failure.**

On a Primary Admin Node on a Physical Server

NOTE New installations of the StorageGRID 9.0 system are not supported on physical servers. Virtual machines must be used.

Prerequisites and Required Materials

- Deployment grid specification file. See “Export the Latest Grid Specification File” on page 95.
- Provisioning USB flash drive

- Backup Provisioning USB flash drive
- Passwords.txt file
- Provisioning passphrase

Procedure

1. Copy the edited grid specification file to the root level of the Provisioning USB flash drive.
2. Remove the old grid specification file from the root level of the Provisioning USB flash drive.

The Provisioning USB flash drive must contain only one grid specification file at the root level. Otherwise, provisioning will fail.

3. Verify that the Provisioning USB flash drive contains only one grid specification file at the root level, that is, there is only one file named GID<*grid_ID*>_REV<*revision_number*>_GSPEC.xml.
4. Run the provisioning script:
   a. At the primary Admin Node server, access a command shell and log in as root using the password listed in the Passwords.txt file.
   b. Run the provisioning script. Enter: **`provision`**
   c. When prompted, insert the Provisioning USB flash drive.
   d. When prompted, enter the provisioning passphrase.
   e. When provisioning is complete, remove the Provisioning USB flash drive.

NOTE If provisioning ends with an error message, see "Provisioning Troubleshooting" on page 105.

5. Back up the provisioning data:
   a. Insert the Backup Provisioning USB flash drive.
   b. Enter: **`backup-to-usb-key`**
   c. When prompted, enter the provisioning passphrase.
6. When backup is complete, remove the Backup Provisioning USB flash drive.
7. Review the current configuration to confirm all settings are correct:
   a. Copy the file GID<*grid_ID*>_REV<*revision_number*>_SAID.zip on the USB Provisioning flash drive to the service laptop and extract the contents.
   b. Inspect the file Doc\Index.html to make sure that the settings are correct. If there is an error, you need to provision the grid again. For more information, see "Errors in Grid Specification File" on page 106.

8. Store the Provisioning USB flash drive and the Backup Provisioning USB flash drive separately in safe locations.

WARNING **Store copies of the Provisioning USB flash drive and the Backup Provisioning USB flash drive in two separate and secure locations. The USB flash drives contain encryption keys and passwords that can be used to obtain data from the grid. The Provisioning USB flash drive is also required to recover from a primary Admin Node failure.**

Change the Provisioning Passphrase

Use this procedure to update the provisioning passphrase. The provisioning passphrase is used to encrypt the GPT repository. It is created when the grid is first installed and is required for software upgrades, grid expansions, and many maintenance procedures.

WARNING **The provisioning passphrase is required for many installation and maintenance procedures. The provisioning passphrase is not listed in the Passwords.txt file. Make sure that it is documented and kept in a safe location.**

On a Primary Admin Node on a Virtual Machine

Prerequisites and Required Materials

- Passwords.txt file
- Current provisioning passphrase
- New provisioning passphrase

Procedure

1. At the primary Admin Node server, access a command shell and log in as root using the password listed in the Passwords.txt file.

2. Change the passphrase:
   a. Enter: `change-repository-passphrase <path>`

      where *<path>* is the location on the server where you want to store a copy of the updated GPT repository.
   b. When prompted, enter the old passphrase.
   c. When prompted, enter the new passphrase. It must be at least six characters.
   d. When prompted, enter the passphrase again.

      The passphrase of the GPT repository is changed to the new value, and an updated copy of the repository that uses this password is saved to *<path>*.

WARNING **The provisioning passphrase is required for many installation and maintenance procedures. The provisioning passphrase is not listed in the Passwords.txt file. Make sure that it is documented and kept in a safe location.**

3. Back up the provisioning data to another directory on the Admin Node. This backup copy can be used to restore the grid in the case of an emergency or during an upgrade or grid expansion.
   a. Create a directory for the backup provisioning data. Enter:

      `mkdir -p /var/local/backup`
   b. Back up the provisioning data. Enter:

      `backup-to-usb-key /var/local/backup`
   c. When prompted, enter the provisioning passphrase.
4. Store the contents of the Provisioning directory (found at </var/local/gpt-data/>) and the Backup Provisioning directories (/var/local/backup) separately in a safe place. For more information, see “Preserving Copies of the Provisioning Data” on page 104.

WARNING **Protect the contents of the Provisioning directory. The Provisioning directory contains encryption keys and passwords that can be used to obtain data from the grid. The Provisioning directory is also required to recover from a primary Admin Node failure.**

5. Close the command shell. Enter: `exit`
6. Write down the provisioning passphrase for future reference.

On a Primary Admin Node on a Physical Server

NOTE New installations of the StorageGRID 9.0 system are not supported on physical servers. Virtual machines must be used.

Prerequisites and Required Materials

- Provisioning USB flash drive
- Backup Provisioning USB flash drive
- Passwords.txt file
- Current provisioning passphrase
- New provisioning passphrase

Procedure

1. At the primary Admin Node server, access a command shell and log in as root using the password listed in the Passwords.txt file.
2. Change the passphrase:
    a. Enter: **`change-repository-password`**
    b. When prompted, enter the old passphrase.
    c. When prompted, enter the new passphrase. It must be at least six characters.
    d. When prompted, enter the passphrase again.
    e. When prompted, insert the Provisioning USB flash drive.
3. Remove the Provisioning USB flash drive and store in a safe place.
4. When prompted, insert the Backup Provisioning USB flash drive.
5. When backup is complete, remove the Backup Provisioning USB flash drive and store it in a safe place.

WARNING **Store copies of the Provisioning USB flash drive and the Backup Provisioning USB flash drive in two separate and secure locations. The USB flash drives contain encryption keys and passwords that can be used to obtain data from the grid. The Provisioning USB flash drive is also required to recover from a primary Admin Node failure.**

6. Close the command shell. Enter: **`exit`**

7. Write down the provisioning passphrase for future reference.

WARNING **The provisioning passphrase is required for many installation and maintenance procedures. The provisioning passphrase is not listed in the Passwords.txt file. Make sure that it is documented and kept in a safe location.**

Provisioning without a USB Flash Drive

The following commands are run at the command shell interface of the primary Admin Node to either update or copy grid provisioning data:

- provision
- change-repository-passphrase
- copy-grid-spec
- backup-to-usb-key
- restore-from-usb-key
- load-provisioning-software

By default, provisioning data is assumed to be stored on the Provisioning USB flash drive and the Backup USB flash drive. You are prompted to insert the appropriate device so that updated data can be written to these locations.

However, you can store provisioning information to another location by optionally entering a file path as an argument to each of these commands, as follows:

- provision <*path*>
- change-repository-passphrase <*path*>
- copy-grid-spec <*path*>
- backup-to-usb-key <*path*>
- restore-from-usb-key <*path*>

When loading provisioning software during upgrade, use the following flag to specify an alternate location for the updated provisioning information:

- load-provisioning-software --alternate-usb-dir=<*path*>

If you choose to store provisioning data to another location, be aware of the following:

- the size of the GPT repository increases every time the provision command is run because a new revision is created and is added to the GPT repository
- preserving copies of the GPT repository is *critical* to the continued operation of the grid, as described in "Preserving Copies of the Provisioning Data" below.

For information on supported versions of VMware software, see the Interoperability Matrix Tool (IMT)

Because VMware vSphere software does not support the use of USB flash drives, it is required that you store provisioning data to an alternate location when the primary Admin Node is installed in a virtual machine. Floppy disk images are sometimes used to transfer data to or from software running in a virtual machine, but you should be aware of the following:

- except for very small grids, the SAID package is too large to place on a floppy disk image
- except for very small grids that have only a few revisions, the GPT repository is too large to place on a floppy disk image

Therefore you will generally store provisioning data to a location on the primary Admin Node, and immediately make additional copies in alternate locations as described in "Preserving Copies of the Provisioning Data" below.

Preserving Copies of the Provisioning Data

Preserving the grid's GPT repository is critical to the continued operation of StorageGRID software. The contents of the GPT repository are required to upgrade, maintain, or expand the grid. The GPT repository is also required to restore the primary Admin Node, should it fail and require replacement.

VMware vSphere does not permit you to store data from the grid node directly to a USB flash drive. Each time you run one of the commands that update or copy grid provisioning data, you *must* back up the provisioning data to two secure locations, preferably in two distinct physical locations. For example, use a tool such as WinSCP to copy the provisioning data to your service laptop, and then store it to two USB flash drives. Store these USB flash drives separately in two geographically distinct secure locations such as a locked cabinet or safe. The USB flash drives contain encryption keys and passwords that can be used to obtain data from the grid.

WARNING **Always store two copies of the GPT repository in two separate and secure locations. The GPT repository is essential to the continued operation of the grid, and to recover from a failed primary Admin Node.**

It is possible to back up provisioning data directly into the grid by saving a copy to an FSG file share. If you store a copy to the grid, it is recommended that you *always* store a second copy in another location outside of the grid. The SAID package includes a two-part encryption key that permits you to recover data from grid nodes in the event of a catastrophic grid failure. If the only copies of these keys are in the grid itself, it is not possible to recover data after such a failure.

In a grid where the primary Admin Node was upgraded to StorageGRID 9.0 and is installed directly on a physical server, it is recommended that you store provisioning data on two USB flash drives (the Provisioning USB flash drive and the Backup USB flash drive). Store the Provisioning USB flash drive and the Backup Provisioning USB flash drive separately in two geographically distinct secure locations such as a locked cabinet or safe. The USB flash drives contain encryption keys and passwords that can be used to obtain data from the grid.

NOTE New installations of the StorageGRID 9.0 system are not supported on physical servers. Virtual machines must be used.

Provisioning Troubleshooting

In case of provisioning errors, follow the guidelines below.

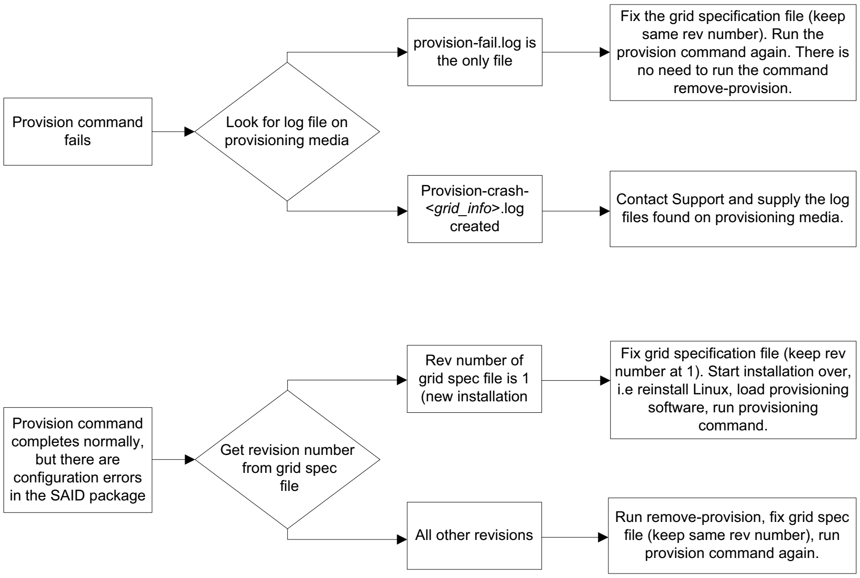
Provision command fails
Look for log file on provisioning media
provision-fail.log is the only file
Fix the grid specification file (keep same rev number). Run the provision command again. There is no need to run the command remove-provision.
Provision-crash-<grid_info>.log created
Contact Support and supply the log files found on provisioning media.
Provision command completes normally, but there are configuration errors in the SAID package
Get revision number from grid spec file
Rev number of grid spec file is 1 (new installation
Fix grid specification file (keep rev number at 1). Start installation over, i.e reinstall Linux, load provisioning software, run provisioning command.
All other revisions
Run remove-provision, fix grid spec file (keep same rev number), run provision command again.

Figure 16: Troubleshooting Provisioning Errors

Provision Command Fails

If provisioning fails because the grid specification file is incorrect, the file provision-fail.log is created on the Provisioning Media. This file contains the error message that the provisioning software displayed before terminating.

If the provisioning program terminates abnormally (crash), two identical log files are saved to the Provisioning Media:

- provision-fail.log
- provision-crash-<*grid_info*>.log

  where <*grid_info*> includes the grid ID, the grid revision being created and a timestamp.

If the provision-fail.log file is the only file created, fix the grid specification file by updating it with Grid Designer and run provisioning again. If the provision-crash-<*grid_info*>.log file is created, contact Support.

If provisioning ends with an error, no information is saved and the remove-revision command does not need to be run.

Errors in Grid Specification File

If the provision command completes normally, but you discover an error in the provisioning data after examining the configuration pages in the SAID package, fix the grid specification file by updating it with Grid Designer and then reprovision the grid.

Initial Installation

NOTE Follow this procedure if the revision number of the grid specification file is 1.

If during the initial installation you discover errors in the SAID package, you must fix the grid specification file and reinstall the primary Admin Node from the beginning, that is, you must reinstall Linux, load provisioning software, and provision the grid.

Upgrades, Expansion, and Maintenance Procedures

NOTE Follow this procedure if the revision number of the grid specification file is greater than 1. This procedure cannot be used for a new installation.

1. Confirm that no scripts or grid tasks generated by provisioning have been started.
2. Remove the provisioning data from the grid. Enter:

   ```
   remove-revision
   ```

   The remove-revision command does not remove grid tasks generated by provisioning nor does it roll back grid tasks that have already been run.

WARNING **Do not use the remove-revision command if you have started any scripts or grid tasks that were generated by provisioning. Contact Support for assistance.**

3. When prompted, enter the provisioning passphrase.
4. Cancel any pending grid tasks created by the provisioning.
5. Fix the deployment grid specification file with Grid Designer and save it to the root directory of the Provisioning Media. Do not change the REV<*revision_number*>.

   NOTE The Provisioning Media must contain only one grid specification file at the root level. Otherwise, provisioning will fail.

6. Run provisioning again and generate a new SAID package. For more information, see "Provision the Grid" on page 96. The old SAID package is overwritten and a new one is generated that uses the same naming convention.
7. Review the contents of the SAID package to confirm that the provisioning information is correct.

B

How to Use GDU

Start GDU

NOTE GDU is always run from the primary Admin Node or the HCAC's primary reporting Admin Node.

1. At the primary Admin Node server or the HCAC's primary reporting Admin Node, access a command shell.

   — or —

   If using GDU remotely:

   a. Start a Telnet/ssh client such as PuTTY.

   b. Select **Window ▶ Translation ▶ Remote character set ▶ UTF-8**.

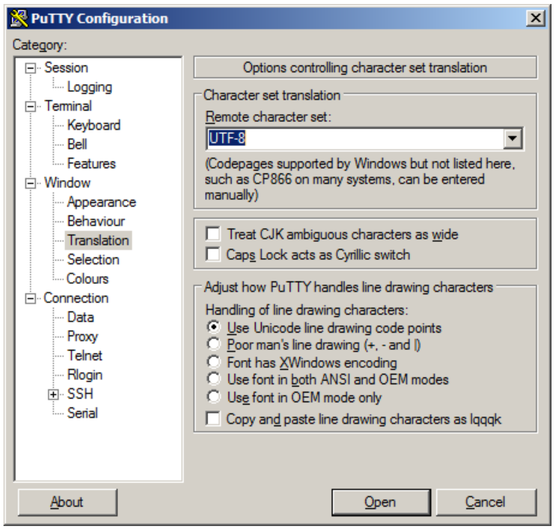
PuTTY Configuration
Category:
Session
Logging
Terminal
Keyboard
Bell
Features
Window
Appearance
Behaviour
Translation
Selection
Colours
Connection
Data
Proxy
Telnet
Rlogin
SSH
Serial
Options controlling character set translation
Character set translation
(Codepages supported by Windows but not listed here, such as CP866 on many systems, can be entered manually)
☐ Treat CJK ambiguous characters as wide
☐ Caps Lock acts as Cyrillic switch
Adjust how PuTTY handles line drawing characters
Handling of line drawing characters:
● Use Unicode line drawing code points
○ Poor man's line drawing (+, - and |)
○ Font has XWindows encoding
○ Use font in both ANSI and OEM modes
○ Use font in OEM mode only
☐ Copy and paste line drawing characters as lqqqk
About
Open
Cancel

Remote character set: UTF-8

Figure 17: PuTTY Settings for GDU

2. Log in as root using the password listed in the Passwords.txt file.
3. If you are using GDU for an upgrade, enter: **`exec bash`**

4. Enter: `ssh-add`

   You need to run ssh-add, which adds the ssh private key to the ssh agent, each time you start a new shell session.

   For more information on ssh access points, see the *Administrator Guide*.

5. If prompted, enter the SSH Access Password listed in the Passwords.txt file.

6. If using PuTTY, start screen. For example, enter: `screen -S GDU`

   NOTE Do not use screen if running GDU locally because the GDU console characters will not display properly.

   The name of the session (for example GDU in the command above) is optional, but recommended since it is useful for managing screen sessions.

   The screen program allows you to manage multiple shell instances concurrently, connect to the same session from different locations, detach from a session without stopping the program running within the session, and resume a session that was previously detached.

   To detach from a screen, press **<Ctrl>+<A>** and then **<Ctrl>+<D>**.

   To reattach to a screen, enter: `screen -r`

7. Start GDU. Enter: `gdu-console`

   NOTE If you get an error using GDU during an upgrade, it is likely because the session was already open. Either log out of the session and log back in, or enter: `exec bash`

8. When prompted, enter the provisioning passphrase. Type the passphrase, press **<Tab>** to select OK, and then press **<Enter>**.

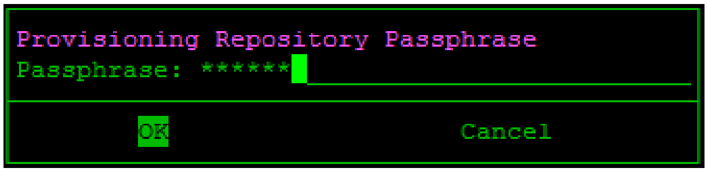

```
Provisioning Repository Passphrase
Passphrase: ******
```

Figure 18: Entering the Provisioning Passphrase to Start GDU

If the characters do not display properly, see "GDU Display Problems" on page 115.

GDU User Interface

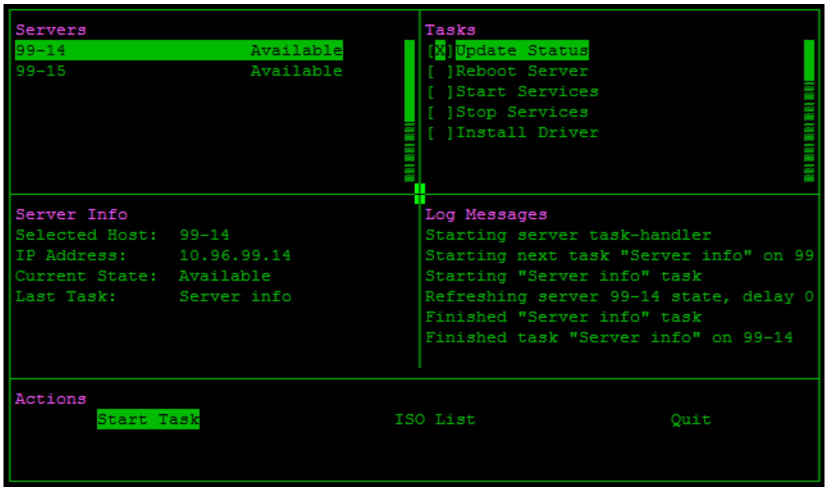

```
Servers                                     Tasks
99-14                       Available       [X]Update Status
99-15                       Available       [ ]Reboot Server
                                            [ ]Start Services
                                            [ ]Stop Services
                                            [ ]Install Driver

Server Info                                 Log Messages
Selected Host:   99-14                      Starting server task-handler
IP Address:      10.96.99.14                Starting next task "Server info" on 99
Current State:   Available                  Starting "Server info" task
Last Task:       Server info                Refreshing server 99-14 state, delay 0
                                            Finished "Server info" task
                                            Finished task "Server info" on 99-14

Actions
        Start Task                    ISO List                    Quit
```

Figure 19: GDU Console

The GDU console consists of five panels:

- Servers — Displays the servers in the grid.
- Tasks — Displays the procedures that can be performed on the server selected in the Servers panel. Only the tasks applicable to the current situation are displayed. It is possible to run GDU tasks in parallel on different servers.

  The list of tasks includes:

Task	Select To
Continue Install	Continue the software installation on the primary Admin Node or the HCAC's primary reporting Admin Node server if it has rebooted.
Enable Services	Start the grid software.
Install Driver	Install a driver from the Enablement Layer for StorageGRID Software CD.
Install Software	Install the grid software on a new server.
Load Configuration	Load NMS configuration settings.
Reboot Server	Reboot the server and start the services.
Remount Storage	Check for preserved storage volumes and remount them. Used for maintenance procedures on Storage Nodes.

Task	Select To
Start Services	Start Server Manager and all services. This is equivalent to the command **`/etc/init.d/servermanager start`**
Stop Services	Stop Server Manager and all services. This is equivalent to the command **`/etc/init.d/servermanager stop`**
Update Software	Apply a service pack.
Upgrade Software	Install a new base version of the software and apply a service pack.
Update Status	Display the current server status.

These tasks are described in detail in the procedures where they are used.

- Server Info — Displays the state of the server selected in the Servers panel. The status can be one of:

Current State	Notes
Available	The server is available for the tasks listed in the Tasks panel.
Busy	A GDU task is running on this server.
Error	A GDU task has failed.
Pingable	The server is pingable, but cannot be reached because there is a problem with the hostname.
Reachable	The server can be reached but is not available because the ssh host keys do not match.

- Log Messages — Displays the output of the GDU task executed for the server selected in the Servers panel. If you are running multiple GDU tasks in parallel, you can display the output of each task by selecting the appropriate server in the Servers panel.
- Actions — The actions are:

Action	Select to
Start Task	Start the procedure selected in the Tasks panel.
ISO List	List the ISO images that are in the /var/local/install directory of the primary Admin Node.
Quit	Quit GDU.

Entering Commands in GDU

Use the keyboard to enter commands:

To	Do
Go from panel to panel	Press **<Tab>**.
Go back from panel to panel	Press **<Shift> <Tab>**.
Go up and down within a panel	Press **<Up Arrow>** and **<Down Arrow>** Press **<Page Up>** and **<Page Down>** Press **<Home>** and **<End>**
Go right and left within a panel	Press **<Left Arrow>** and **<Right Arrow>**.
Select a task	Press the space bar. X appears next to the selected task.
Activate a command	Press **<Enter>**.

Install Drivers with GDU

You can use GDU to install drivers that are included on the Enablement Layer for StorageGRID Software CD in the `drivers` directory. It is your responsibility to confirm that the drivers included on the Enablement Layer for StorageGRID Software CD are the most recent qualified version. If they are not, get the latest version from the hardware vendor and install the drivers manually.

Prerequisites

- Connectivity to the primary Admin Node
- List of drivers required for this server
- Provisioning passphrase

Procedure

1. Start GDU.
2. Select the server in the **Servers** panel and confirm that its state is `Available`.

3. Install the driver:
    a. Select **Install Driver** in the **Tasks** panel. A panel listing the available drivers opens automatically. If the driver you need is not listed, you must install this driver manually.

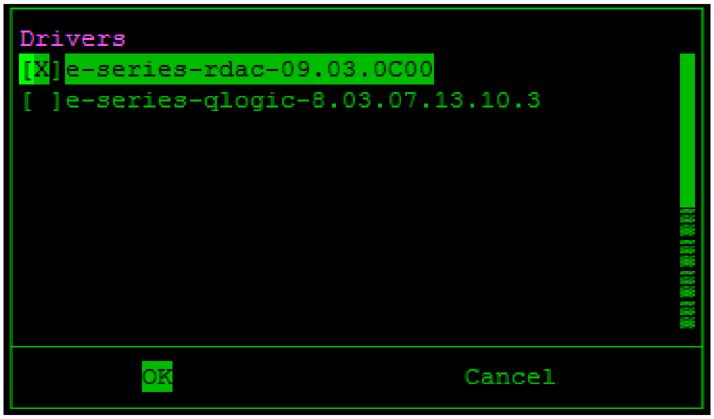

```
Drivers
[X]e-series-rdac-09.03.0C00
[ ]e-series-qlogic-8.03.07.13.10.3

          OK                      Cancel
```

Figure 20: Installing Drivers with GDU

    b. Select a driver from the list.
    c. Select **OK**. Wait for the driver installation script to complete.
    d. Reboot the server: Select **Reboot Server** in the **Tasks** panel, then select **Start Task** in the **Actions** panel and press **<Enter>**.
4. Repeat step **3** if you need to install any other driver on this server using GDU.
5. If you have finished using GDU, close it and remove passwordless access.

Close GDU

If you quit GDU while a task is in progress, GDU pauses until the task completes, and then closes. Some tasks, such as formatting storage volumes on a new Storage Node, can take hours to complete. Avoid quitting GDU while long-running tasks are in progress. Continue working in another terminal window.

1. Quit GDU. Select **Quit** in the **Actions** panel and then press **<Enter>**.

   When prompted, confirm that you want to quit GDU.
2. Remove the ability to access servers without a server password. Enter: `ssh-add -D`
3. Close the screen session. Enter: `exit`

GDU Troubleshooting

GDU Display Problems

Under certain circumstances, the GDU console may not display properly. For an example, see Figure 21 below.

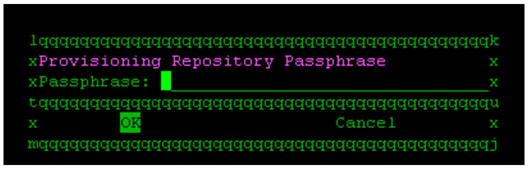

```
lqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqk
xProvisioning Repository Passphrase           x
xPassphrase: _________________________________x
tqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqu
x        OK                   Cancel          x
mqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqqj
```

Figure 21: GDU Display Problems

- If using PuTTY, change the Window Translation setting to Use font encoding.
- If running GDU locally, do not use screen.

Problems with Server Status

When starting GDU, the status update of all servers may hang, or take a long time to complete. After server status is updated, many appear as Unknown or Pingable when it is known that the servers are Available. This typically occurs in large grids with many servers.

To correct the problem, quit GDU, and restart with the -k option. Enter:

```
gdu-console -k
```

Starting GDU with the -k option bypasses its initial status update on startup. The state of all servers remains Unknown in GDU until you manually update them using the Update Status task.

GDU Log Files

The GDU logs are located on the primary Admin Node in /var/local/log/gdu-console.log.

Missing GDU Task

If a GDU task that you must execute is missing from the Tasks panel, check the GDU log for the reason. For instance, it is possible that a required ISO image is missing. To list the ISO images currently in the /var/local/install directory, use the ISO List GDU action. For a sample output, see Figure 22 below.

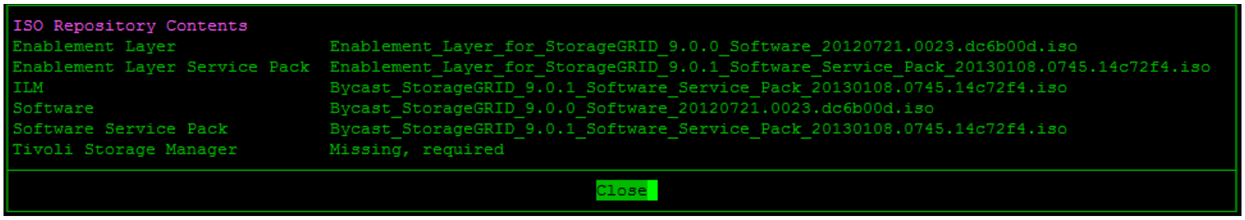

```
ISO Repository Contents
Enablement Layer             Enablement_Layer_for_StorageGRID_9.0.0_Software_20120721.0023.dc6b00d.iso
Enablement Layer Service Pack  Enablement_Layer_for_StorageGRID_9.0.1_Software_Service_Pack_20130108.0745.14c72f4.iso
ILM                          Bycast_StorageGRID_9.0.1_Software_Service_Pack_20130108.0745.14c72f4.iso
Software                     Bycast_StorageGRID_9.0.0_Software_20120721.0023.dc6b00d.iso
Software Service Pack        Bycast_StorageGRID_9.0.1_Software_Service_Pack_20130108.0745.14c72f4.iso
Tivoli Storage Manager       Missing, required
```

Figure 22: ISO Images in /var/local/install

The label Missing, required means that the ISO image of the CD required for the installation is not in the /var/local/install directory.

The label Not present means that an ISO image that GDU expected to find is not in the /var/local/install directory, but GDU does not know whether this ISO image is actually required.

Troubleshooting with screen in Multi Display Mode

The screen program is useful when two or more people need to interact with a shell session simultaneously for troubleshooting purposes. Below is an example of two users connecting to GDU at the same time.

User 1 creates a named screen session and starts GDU.

```
# screen -S GDU
# gdu-console
```

User 2 lists the screen sessions and connects without detaching User 1.

```
# screen -ls
There is a screen on:
        5361.GDU        (Attached)
1 Socket in /var/run/uscreens/S-root.
# screen -r -x GDU
```

Now both users are viewing GDU and inputs can come from either user.

About load_cds.py

The load_cds.py command accepts two different inputs: physical installation CDs or ISO images of the installation CDs stored in a directory on the primary Admin Node or the HCAC's primary reporting Admin Node.

You can run the load_cds.py script as many times as you need.

The script automatically deletes older service pack software when you load the latest service pack software.

If you insert the same CD twice, no new ISO is created. The existing ISO will not be overwritten.

If the load_cds.py script fails because you inserted a CD unrecognized by the script, eject the CD and continue with the correct CD (you do not have start over from the first CD you loaded).

Copy ISO Files in Multi-Site Environment

In a multi-site environment, copy ISO files to the servers in the remote location prior to installing or upgrading the software with GDU. This is an optional, but recommended, step to reduce the number of large files that would otherwise be transferred over a slow WAN link.

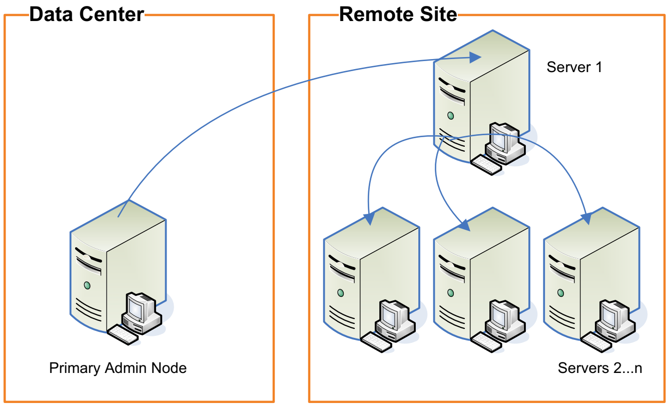

Figure 23: Copying Files to Remote Site

Prerequisites and Required Materials

- ISO images of the StorageGRID Software CD and the Enablement Layer for StorageGRID Software CD have been copied to the primary Admin Node or the HCAC's primary reporting Admin Node using the load_cds.py command
- ssh access between the Admin Node and the servers at the remote location

Procedure

1. At the primary Admin Node server or the HCAC's primary reporting Admin Node, access a command shell and log in as root using the password listed in the Passwords.txt file.

It is not usually necessary to copy the service pack ISO images since these files are small: there is no real gain over letting GDU copy the files automatically.

2. Copy the ISO image of the StorageGRID Software CD to a server at the remote location. If the site has an Admin Node, copy the ISO to it. Otherwise, use a Gateway Node, preferably a secondary. Enter:

   ```
   scp /var/local/install/Bycast_StorageGRID_9.0.0_Software_\
   <build>.iso <destination>:/var/local/tmp
   ```

   where *<destination>* is the hostname or IP address of the first server at the remote location.

3. Copy the ISO image of the Enablement Layer for StorageGRID Software CD from the primaryAdmin Node to the destination server. Enter:

   ```
   scp /var/local/install/Enablement_Layer_for_StorageGRID_\
   9.0.0_Software_<build>.iso <destination>:/var/local/tmp
   ```

4. If this is an upgrade, install the load_cds.py script. Enter:

```
scp /usr/local/sbin/load_cds.py <destination>:/usr/local/sbin/
```

5. Log in to the server at the remote site where you copied the ISO files. Enter: **`ssh <destination>`**

   When prompted, enter the password for the remote server listed in the Passwords.txt file.

6. Change to the /var/local/tmp directory. Enter: **`cd /var/local/tmp`**

7. Load the ISOs using the load_cds.py script. Enter:

   ```
   load_cds.py Bycast_StorageGRID_9.0.0_Software_<build>.iso \
   Enablement_Layer_for_StorageGRID_9.0.0_Software_<build>.iso
   ```

   Separate the ISO file names with a space.

8. Empty the temporary directory. Enter: **`rm -r /var/local/tmp/*`**

9. Copy the ISOs from the first server at the remote location to the remaining servers at the remote location. For each remaining server:
   a. Copy the ISO files needed for the update to the server. Enter:
      ```
      scp /var/local/install/* <next_server>:/var/local/tmp
      ```
      where *<next_server>* is the hostname or IP address of the next server at the remote site.
   b. If this is an upgrade, install the load_cds.py script. Enter:
      ```
      scp /usr/local/sbin/load_cds.py <next_server>:/usr/local/sbin/
      ```
   c. Log in to the next server at the remote site where you copied the ISO files. Enter: `ssh <next_server>`

      When prompted, enter the password for the remote server listed in the Passwords.txt file.
   d. Change to the /var/local/tmp directory. Enter: `cd /var/local/tmp`
   e. Load the ISOs using the load_cds.py script. Enter:
      ```
      load_cds.py Bycast_StorageGRID_9.0.0_Software_<build>.iso \
      Enablement_Layer_for_StorageGRID_9.0.0_Software_<build>.iso
      ```
      Separate the ISO file names with a space.
   f. Empty the temporary directory. Enter: `rm -r /var/local/tmp/*`
   g. End the ssh session. Enter: `exit`
   h. Repeat from step **a** for each server at the remote site.
10. End the ssh session on the remote server. Enter: `exit`
11. Log out of the Admin Node. Enter: `exit`

C

Install Grid Software Manually

Introduction

This chapter includes procedures to install StorageGRID software manually. While the preferred and recommended procedure to install grid software is with GDU, the manual installation of grid software is sometimes required.

NOTE Always use GDU to install grid software on the primary Admin Node.

Once you start installing grid software on a server manually, you cannot switch to GDU to complete the installation of that server.

The manual grid software installation procedure uses the glsetup and postinstall scripts. Table 12 below compares the GDU tasks to the equivalent manual installation scripts.

Table 12: GDU Tasks and Equivalent Installation Scripts

GDU Task	Equivalent Scripts
Install Software	glsetup.sh and postinstall.rb activate
Enable Services	postinstall.rb start
Update Software	updategrid.rb

Install Grid Software Manually

Prerequisites

- The grid software has been installed on the primary Admin Node.
- You have a copy of the ISOs from the /var/local/install/ directory of the primary Admin Node. Note that the Service Pack ISOs are only required if a service pack is available for the release:
  - Bycast_StorageGRID_9.0.0_Software_<*build*>.iso
  - Enablement_Layer_for_StorageGRID_9.0.0_Software_<*build*>.iso
  - Bycast_StorageGRID_9.0.<*servicepack*>_Software_Service_Pack_<*build*>.iso
  - Enablement_Layer_for_StorageGRID_9.0.<*servicepack*>_Software_Service_Pack_<*build*>.iso
  - For Archive Nodes, the ISO of the TSM client packages CD
- Linux has been installed on the server.
- The virtual machines have prepared. For more information, see "Prepare Virtual Machines" on page 23.

Procedure

In each step, omit the service pack CDs if a service pack is not available for the release.

1. Log in to the server as root using the password listed in the Passwords.txt file.
2. Copy these ISO images to the server's /var/local/tmp directory:
   - Bycast_StorageGRID_9.0.0_Software_<*build*>.iso
   - Enablement_Layer_for_StorageGRID_9.0.0_Software_<*build*>.iso
   - Bycast_StorageGRID_9.0.<*servicepack*>_Software_Service_Pack_<*build*>.iso
   - Enablement_Layer_for_StorageGRID_9.0.<*servicepack*>_Software_Service_Pack_<*build*>.iso
   - For an Archive Node, the ISO of the TSM client packages CD
3. Remove any USB flash drive from the server.
4. Change to the /var/local/tmp directory. Enter: `cd /var/local/tmp`
5. Mount the ISO of Software Service Pack CD. Enter:
   ```
   mount -o loop,ro <service_pack_CD_iso> /cdrom
   ```
6. Install the updated load_cds.py script from the Service Pack CD. Enter: `/cdrom/install-load-cds`

7. Unmount the Service Pack CD ISO. Enter: `umount /cdrom`

Do not copy the images to /var/local/install directly. Use the load_cds.py script instead because the script verifies the integrity of the ISOs.

8. Copy the ISOs to the /var/local/install/ directory using the load_cds.py script. Enter:

```
load_cds.py Bycast_StorageGRID_9.0.0_Software_<build>.iso \
Enablement_Layer_for_StorageGRID_9.0.0_Software_<build>.iso
\
Bycast_StorageGRID_9.0.<servicepack>_Software_Service_\
Pack_<build>.iso Enablement_Layer_for_StorageGRID_\
9.0.<servicepack>_Software_Service_Pack_<build>.iso
```

9. Remove the ISOs from /var/local/tmp. Enter: `rm -r /var/local/tmp`

10. Set up the server:

    a. Run the setup script. Enter:

    ```
    glsetup.sh --iso=/var/local/install/Enablement_Layer\
    _for_StorageGRID_9.0.0_Software_<build>.iso
    ```

    b. Follow the prompts to read and accept the software license agreement.

    c. If prompted to reformat partitions, for example, if the server storage was previously configured and you are re-running the glsetup script, enter `y` to reformat the partition, or `n` to preserve the existing partition.

    The script ends with the message “Completed Enablement Layer setup” when it completes normally.

11. If prompted to reboot the server:

    a. Enter: `y`

    b. When the server finishes rebooting, log in as root using the password listed in the Passwords.txt file.

12. Activate the software:

    a. If this is for an Archive Node that uses TSM for middleware, enter:

    ```
    postinstall.rb activate --iso=/var/local/install\
    /Bycast_StorageGRID_9.0.0_Software_<build>.iso \
    --iso=<TSM_ISO_name>
    ```

    b. For all other cases, enter:

    ```
    postinstall.rb activate --iso=/var/local/install\
    /Bycast_StorageGRID_9.0.0_Software_<build>.iso
    ```

    Upon completion of activation, the message “StorageGRID activation completed” appears.

13. Mount the StorageGRID Software Service Pack CD image. Enter:

    ```
    mount -o loop,ro /var/local/install/Bycast_StorageGRID_\
    9.0.<servicepack>_Software_Service_Pack_<build>.iso /cdrom
    ```

14. Run the update script to apply the service pack. Enter:

```
/cdrom/swupdate/updategrid.rb --iso=/var/local/install\
/Enablement_Layer_for_StorageGRID_9.0.<servicepack>_\
Software_Service_Pack_<build>.iso
```

15. If prompted to reboot:
    a. Enter: `reboot`
    b. When the server finishes rebooting, log in as root using the password listed in the Passwords.txt file.
16. Repeat this procedure from step **1** for any other grid nodes that require manual installation.

NOTE If you are preforming a manual software installation as part of a recovery procedure, do not start software until instructed.

17. If you are installing an entire StorageGRID system:
    a. Start the grid software. If servers have connectivity to the primary Admin Node, use GDU to start grid software as described in "Start Grid Software" on page 53.
    b. To complete grid installation, ensure that you load NMS configuration settings as described in "Load NMS Configuration Settings" on page 55.
    c. Ensure that you log out of any server that was manually installed. Enter: `exit`
18. If necessary, manually start grid software. Enter:

```
postinstall.rb start ; exit
```

Troubleshooting

glsetup Script Fails

If the script reports an error, run the script again with the verbose option (enter: `glsetup.sh --verbose`) to gather more information about the issue. After the issue has been resolved, you can safely rerun glsetup.sh.

Wrong Answer in glsetup When Formatting Partitions

The glsetup.sh script checks if the required disk partitions exist and are mounted. If they are (because you are re-running the glsetup.sh script), glsetup.sh identifies each partition, displays a warning message, and asks if you want to reformat the existing partition.

If you inadvertently enter No and then rerun glsetup.sh to correct the issue, the prompt to reformat the partition does not appear.

Force Prompt to Appear

1. Manually unmount the disk partitions, using the Linux umount command.
2. Re-run the setup utility. Enter: **`glsetup.sh --enable-reformat`**

postinstall Script Fails

If the postinstall script is halted or fails before it completes, restart the installation procedure on that server from the beginning. The postinstall script cannot be safely run twice, in whole or in part.

D

Required Ports

Inter-grid Communication

Servers that are part of a grid must communicate with one another using the listener ports listed in Table 14 below. These ports are required on the grid network.

Table 13: Inter-grid Communication Ports

Port	Description	Connecting To	Outgoing Connection from Gateway Node to Service
1501 (TCP)	ADC service connection	Servers hosting an ADC service.	Yes
1502 (TCP)	LDR service connection	Servers hosting an LDR service.	No
1503 (TCP)	CMS service connection	Servers hosting a CMS service.	No
1504 (TCP)	NMS service connection	Servers hosting an NMS service.	No
1505 (TCP)	AMS service connection	Servers hosting an AMS service.	No
1506 (TCP)	SSM service connection	Server hosting an SSM service (all servers).	No
1507 (TCP)	CLB service connection	Server hosting a CLB service.	Yes
1508 (TCP)	CMN service connection	Server hosting a CMN service.	No
1509 (TCP)	ARC service connection	Server hosting an ARC service.	No
1510 (TCP)	FSG service connection	Server hosting an FSG service.	Yes

Table 13: Inter-grid Communication Ports (cont.)

Port	Description	Connecting To	Outgoing Connection from Gateway Node to Service
1602 (TCP)	CMS maintenance connection	Server hosting a CMS service.	No
2013(TCP)	NMS service connection	Server hosting a High Capacity Admin Cluster.	No
18080, 18081 (TCP)	HTTP query/ retrieve and ingest	Servers hosting an LDR service in grids that also have FSG services.	Yes
1139 (TCP)	LDR replication	Servers hosting an LDR service.	No
11139 (TCP)	ARC replication	Servers hosting an ARC service.	No
22 (TCP)	ssh	All grid servers for centralization of installation, update, and other maintenance procedures.	Yes
123 (UDP)	ntp	All servers.	Yes
694 (UDP)	ha-client	Servers hosting a High Availability FSG	Yes

Client to Grid Communication

Clients communicate with grid servers to ingest and retrieve content. The connections differ depending on the protocols chosen to ingest and retrieve content. Consult Table 14 on page 129 to choose which ports will ingest and retrieve content. Third-party software (such as hardware monitoring) may require ports not listed here. Other client software may be used to connect with grid servers for monitoring and maintenance. In this case, additional ports are needed depending on the degree to which external monitoring and access is required.

Table 14: Client to Grid Communication Ports

Port	Description	Connecting To	Gateway Node Needs Port Open to Client Network
22 (TCP)	ssh	Servers being used as an access point into the grid for software installation, update, and maintenance.	No
111 (TCP/UDP)	portmap	Servers hosting an FSG service and offering an NFS share.	Yes
903 (TCP/UDP)	rpc.mountd	Servers hosting an FSG service and offering an NFS share.	Yes
904 (TCP/UDP)	nlm	Servers hosting an FSG service and offering an NFS share.	Yes
2049 (TCP/UDP)	nfsd	Servers hosting an FSG service and offering a NFS share.	Yes
137 (UDP)	netbios-ns	Servers hosting an FSG service and offering a CIFS share.	Yes
138 (TCP)	netbios-dgm	Servers hosting an FSG service and offering a CIFS share.	Yes
139 (TCP)	netbios.ssn	Servers hosting an FSG service and offering a CIFS share.	Yes
445 (TCP)	microsoft-ds	Servers hosting an FSG service and offering a CIFS share.	Yes
514 (TCP)	shell	Servers hosting an FSG service and offering an RSH connection.	Yes
161 (TCP/UDP)	snmp	Servers that are being monitored via SNMP.	Yes Dependant on where servers are located.

Table 14: Client to Grid Communication Ports (cont.)

Port	Description	Connecting To	Gateway Node Needs Port Open to Client Network
162 (TCP/UDP)	snmptrap	Servers that are being monitored via SNMP.	Yes Dependant on where servers are located.
80 (TCP)	HTTP	Servers hosting an NMS service.	No
443 (TCP)	HTTPS	Servers hosting an NMS service.	No
8080, 8081 (TCP)	HTTP query/ retrieve and ingest	The StorageGRID API (SGAPI) or CDMI on gateway servers hosting a CLB.	Yes
18080, 18081 (TCP)	HTTP query/ retrieve and ingest	The SGAPI or CDMI on servers hosting the LDR service (most grids access these services through a CLB).	No

E

Connectivity

Connection requirements for the NMS management interface and the server command shell

Browser Settings

Verify Internet Explorer Settings

If you are using Internet Explorer, verify that the settings for temporary internet files, security and privacy are correct.

1. Go to **Tools ▶ Internet Options ▶ General**
2. In the Browsing history box, click **Settings**.
3. For Check for newer versions of stored pages, select **Automatically**.

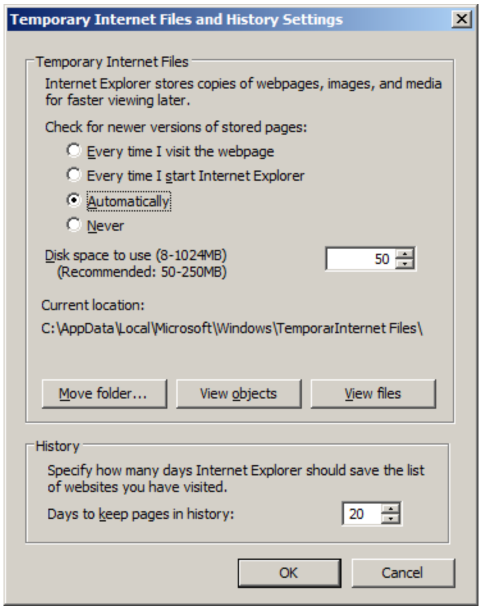

Temporary Internet Files

Internet Explorer stores copies of webpages, images, and media for faster viewing later.

Check for newer versions of stored pages:

○ Every time I visit the webpage
○ Every time I start Internet Explorer
● Automatically
○ Never

Disk space to use (8-1024MB) (Recommended: 50-250MB): 50

Current location:
C:\AppData\Local\Microsoft\Windows\TemporarInternet Files\

History

Specify how many days Internet Explorer should save the list of websites you have visited.

Days to keep pages in history: 20

Figure 24: Temporary Files Setting

4. Go to **Tools ▶ Internet Options ▶ Security ▶ Custom Level** and ensure that the Active Scripting setting is Enable.

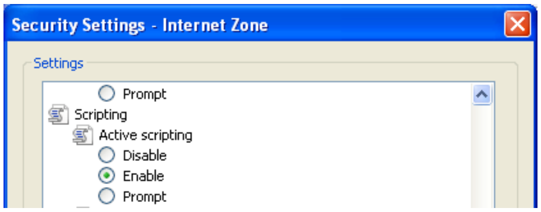

Security Settings - Internet Zone

Settings

○ Prompt

Scripting

Active scripting

○ Disable

● Enable

○ Prompt

Figure 25: Active Scripting Setting

5. Go to **Tools ▶ Internet Options ▶ Privacy** and ensure that the privacy setting is Medium or lower (cookies must be enabled).

Enable Pop-ups

To make any changes to passwords, you must ensure that your browser allows pop-up windows. For more information on allowing pop-up windows, see your browser's documentation.

NMS Connection Procedure

Connecting to the NMS MI at the customer site requires access to the customer's network.

If the grid is configured with a High Capacity Admin Cluster (HCAC), you can only connect to the reporting Admin Node. You cannot connect to the processing Admin Node. In a grid with two Admin Nodes or HCACs, you can connect to either Admin Node or HCACs' reporting Admin Node. Each Admin Node or HCAC displays a similar view of the grid; however, alarm acknowledgments made at one Admin Node or HCAC are not copied to the other Admin Node or HCAC. It is therefore possible that the Grid Topology tree will not look the same for each Admin Node or HCAC.

1. Work with the customer system administrator to establish the physical network connection to the service laptop. Using the customer's network rather than a direct connection within the rack verifies that the interface is accessible using the same infrastructure the customer uses.
2. From the Configuration.txt file, note the IP address of the Admin Node (reporting Admin Node in an HCAC) on the customer network. This is needed to access the NMS MI.

3. From the Passwords.txt file, note the NMS MI password for the Vendor account or the Admin account.
4. Launch the web browser.
5. Open the address **https://<*IP_address*>**

   where <*IP_address*> is the address of the Admin Node (reporting Admin Node in an HCAC) on the customer network specified in the Configuration.txt file.

Security Certificate

Depending on your version of Windows and web browser, you may be warned of a problem with the security certificate when you access the NMS MI URL.

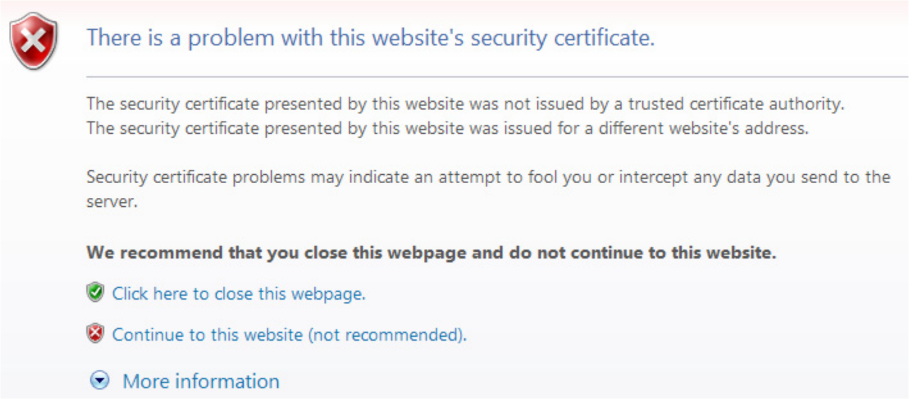

There is a problem with this website's security certificate.

The security certificate presented by this website was not issued by a trusted certificate authority.
The security certificate presented by this website was issued for a different website's address.

Security certificate problems may indicate an attempt to fool you or intercept any data you send to the server.

We recommend that you close this webpage and do not continue to this website.

Click here to close this webpage.

Continue to this website (not recommended).

More information

Figure 26: Example of a Security Alert Window

If this appears, you can either:

- Proceed with this session. The alert will appear again the next time you access this URL.
- Install the certificate. Follow the instructions of your browser.

The NMS MI uses a self-signed certificate. For information on importing this certificate into a browser see the browser's documentation. Note that the self-signed certificate used by the NMS MI is based on the grid's IP address. The expected URL to the interface is this IP address and not a domain name. In cases where the domain name is used, browsers may not be able to match the self-signed certificate to the identity of the NMS server. For more information see the browser's documentation.

Log In

1. Enter the username **Vendor** for full access to the NMS MI. If you are not making grid-wide configuration changes, you can also use the Admin account.
2. Enter the password for the NMS MI specified in the Passwords.txt file.

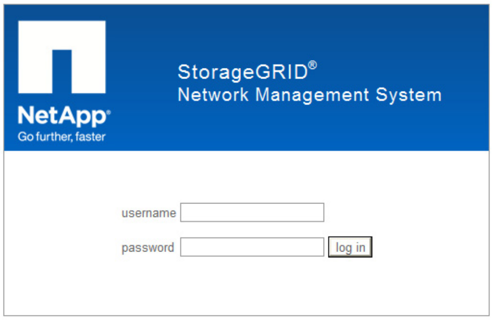

Figure 27: NMS MI Login Window

Log Out

When you finish your NMS MI session, log out to keep the system secure.

1. Click **Logout** Logout ➡ located at the top-right corner of the screen. The logging out message appears.
2. You may safely close the browser or use other applications.

NOTE Failure to log out may give unauthorized users access to your NMS session. Simply closing your browser is *not* sufficient to log out of the session.

Command Shell Access Procedures

Log In

- At the server, access a command shell and log in as root using the password listed in the Passwords.txt file.

Log Out

1. Enter **`exit`** to close the current command shell session.
2. Press **<Alt>+<F7>** to return to the Server Manager GUI.

Accessing a Server Remotely

There are three ways to connect to a server remotely using ssh:

- From any server, using the remote server password
- From the primary Admin Node, using the ssh private key password
- From the primary Admin Node, without entering any password except the ssh private key password once

The primary Admin Node acts as an ssh access point for other grid servers. The procedures to change the ssh private key password and to enable passwordless access from the primary Admin Node to other servers are described in the "Network Configuration" chapter of the *Administrator Guide*.

Connect Using the Remote Server Password

1. Log in to any local server.
2. Enter: **`ssh <IP_address>`**

   where <*IP_address*> is the IP address of the remote server.
3. When prompted, enter the password for the remote server listed in the Passwords.txt file.

Connect Using the ssh Private Key Password

1. Log in to the primary Admin Node.
2. Enter: **`ssh <hostname>`**

   where <*hostname*> is the name of the remote server.

   — or —

   Enter: **`ssh <IP_address>`**

   where <*IP_address*> is the IP address of the remote server.
3. When prompted, enter the SSH Access Password listed in the Passwords.txt file.

Connect to a Server without Using a Password

1. Log in to the primary Admin Node.
2. Add the ssh private key to the ssh agent to allow the primary Admin Node passwordless access to the other servers in the grid. Enter: **`ssh-add`**

   You need to add the ssh private key to the ssh agent each time you start a new shell session on the primary Admin Node.
3. When prompted, enter the SSH Access Password.

   You can now access any grid server from the primary Admin Node via ssh without entering additional passwords.
4. When you no longer require passwordless access to other servers, remove the private key from the ssh agent. Enter: **`ssh-add -D`**
5. Log out of the primary Admin Node command shell. Enter: **`exit`**

Using GNU screen

The GNU screen program, which is installed by default, allows you to manage many shell instances concurrently, to connect to the same session from different locations, to detach from a session without stopping the program running within the session, and to resume a session that was previously detached.

The screen program decouples the terminal emulator from the running program. This means that the program keeps running even if you detach from the session or close the terminal emulator, or lose the connection.

Consider using screen when you execute maintenance procedures remotely and there is a possibility of losing the connection or where it would be useful to 'hang up' and connect back later. For example, you may want to use screen when cloning a CMS database since this procedure can take a few hours to complete.

1. Log in to a server remotely. For more information, see "Accessing a Server Remotely" on page 135.
2. Start screen. Enter: **`screen`**
3. Enter the command or script you need to execute in the new window.
4. To quit the screen session, enter: **`exit`**

Screen has a number of command-line options. For example:

`-d`	To detach a session. The running program disappears from the terminal. However, it continues to run behind the scenes.
`-r`	To resume a detached screen session. This brings the program back to your terminal.
`-ls`	To list existing sessions.
`-S <name>`	To name a session.
`-x <name>`	To attach to a screen session that is already attached by another user (multi display mode). Either user can interact with the screen session or detach from it.

See below for an example of how to use screen.

Create named session

Detach screen session

List screen sessions

Reattach screen session

```
# screen -S CMScloneproc
# <commands to start cloning process>
# <CTRL+A> <CTRL+D>
# screen -ls
There is a screen on:
        20849.CMScloneproc      (Detached)
1 Socket in /var/run/uscreens/S-root.
# screen -d -r CMScloneproc
```

For more information, display the man page for screen, and consult the GNU official web site http://www.gnu.org/software/screen/.

Glossary

ACL	Access control list—Specifies what users or groups of users are allowed to access an object and what operations are permitted, for example read, write, and execute.
active primary FSG	In an HAGC, the FSG that is currently providing read-write service to clients. See also "FSG replication group".
ADC	Administrative Domain Controller—A software component of the StorageGRID system. The ADC service maintains topology information, provides authentication services, and responds to queries from the LDR, CMS, CMN, and CLB. The ADC service is found on the Control Node.
ADE	Asynchronous Distributed Environment—Proprietary development environment used as a framework for grid services within the NetApp StorageGRID Software.
Admin Node	A building block of the StorageGRID system. The Admin Node provides services for the web interface, grid configuration, and audit logs. See also "reporting Admin Node", "processing Admin Node", "primary Admin Node", "Audit Node", and "HCAC".
AMS	Audit Management System—A software component of the StorageGRID system. The AMS service monitors and logs all audited system events and transactions to a text log file. The AMS service is found on the Admin Node—reporting Admin Node in a High Capacity Admin Cluster (HCAC) and the Audit Node.
API	Application Programming Interface—A set of commands and functions, and their related syntax, that enable software to use the functions provided by another piece of software.
API Gateway Node	Application Programming Interface Gateway Node provides read-write access for HTTP clients (via StorageGRID API or CDMI). API Gateway Nodes are configured to include a "CLB" service, but not an "FSG" service. As a result, API Gateway Nodes do not support NFS/CIFS file systems and are not configured as part of a replication group.
ARC	Archive—A software component of the StorageGRID system. The ARC service manages interactions with archiving middleware that controls nearline archival media devices such as tape libraries. The ARC service is found on the Archive Node.

Archive Node	A building block of the StorageGRID system. The Archive Node manages storage of data to nearline data storage devices such as such as tape libraries (via IBM Tivoli® Storage Manager).
Audit Node	A building block of the StorageGRID system. The Audit Node logs all audit system events. It is an optional grid node that is generally reserved for larger grid deployment.
audit message	Information about an event occurring in the StorageGRID system that is captured and logged to a file.
atom	Atoms are the lowest-level component of the container data structure, and generally encode a single piece of information. (Containers are sometimes used when interacting with the grid via the StorageGRID API).
AutoYaST	An automated version of the Linux installation and configuration tool YaST ("Yet another Setup Tool"), which is included as part of the SUSE Linux distribution.
BASE64	A standardized data encoding algorithm that enables 8-bit data to be converted into a format that uses a smaller character set, enabling it to safely pass through legacy systems that can only process basic (low order) ASCII text excluding control characters. See RFC 2045 for more details.
Basic Gateway replication group	A Basic Gateway replication group contains a primary FSG and one or more secondary FSGs.
Binding	The persistent assignment of a grid service (for example, an FSG or SSM) to the consolidated NMS service or processing NMS service. This assignment is based on grid topology (consolidated Admin Node or HCAC). See also "Admin Node".
bundle	A structured collection of configuration information used internally by various components of the grid. Bundles are structured in container format.
business continuity failover	A business continuity failover within a Gateway Node replication group is one where a secondary Gateway Node is manually configured to act as a primary after the primary Gateway Node fails. Clients can continue to read and write to the grid after they are manually redirected to the acting primary. This is a temporary measure to maintain service while the primary Gateway Node is repaired.

CBID	Content Block Identifier — A unique internal identifier of a piece of content within the StorageGRID system.
CDMI	Cloud Data Management Interface — An industry standard defined by SNIA that includes a RESTful interface for object storage. For more information, see http://www.snia.org/cdmi.
CIDR	Classless Inter-Domain Routing—A notation used to compactly describe a subnet mask used to define a range of IP addresses. In CIDR notation, the subnet mask is expressed as an IP address in dotted decimal notation, followed by a slash and the number of bits in the subnet. For example, 192.0.2.0/24.
CIFS	Common Internet File System—A file system protocol based on SMB (Server Message Block, developed by Microsoft) which coexists with protocols such as HTTP, FTP, and NFS.
CLB	Connection Load Balancer—A software component of the StorageGRID system. The CLB service provides a gateway into the grid for clients connecting via the HTTP protocol. The CLB service is part of the Gateway Node.
Cloud Data Management Interface	See "CDMI" on page 141.
CMN	Configuration Management Node— A software component of the StorageGRID system. The CMN service manages system-wide configuration and grid tasks. The CMN service is found on the primary Admin Node.
CMS	Content Management System—A software component of the StorageGRID system. The CMS service manages content metadata and content replication according to the rules specified by the ILM policy. The CMS service is found on the Control Node.
command	In HTTP, an instruction in the request header such as GET, HEAD, DELETE, OPTIONS, POST, or PUT. Also known as an HTTP method.
container	A container is a data structure used by the internals of grid software. In the StorageGRID API, an XML representation of a container is used to define queries or audit messages submitted using the POST command. Containers are used for information that has hierarchical relationships between components. The lowest-level component of a container is an atom. Containers may contain 0 to N atoms, and 0 to N other containers.

content block ID	See "CBID".
content handle	See "UUID".
consolidated Admin Node	Admin Node hosting the consolidated NMS service. Can be the primary Admin Node.
consolidated NMS	Hosted by the consolidated Admin Node. It is the equivalent of a combined reporting NMS and processing NMS service. See also "NMS".
Control Node	A building block of the StorageGRID system. The Control Node provides services for managing content metadata and content replication.
CSTR	Null-terminated, variable length string.
DC	Data Center site.
deduplication	If enabled, when the grid identifies two files as being identical, it "deduplicates" them by redirecting all content handles to point to a single stored instance of the file. The end result is that only the number of copies required by the ILM policy are stored in the grid. The feature was designed for use with applications that save two identical copies of a file to the grid via different Gateway Nodes. **NOTE** Deduplication is deprecated and no longer supported.
distributed CMS	A CMS that uses metadata replication. See also "metadata replication".
DR	Disaster Recovery site.
EMR	Electronic Medical Records—A computerized system for managing medical data that may be interfaced to the grid.
Enablement Layer	The Enablement Layer for StorageGRID Software CD is used during installation to customize the Linux operating system installed on each grid server. Only the packages needed to support the services hosted on the server are retained, which minimizes the overall footprint occupied by the operating system and maximize the security of each grid node.
FCS	Fixed Content Storage—a class of stored data where the data, once captured, is rarely changed and must be retained for long periods of time in its original form. Typically this includes images, documents,

and other data where alterations would reduce the value of the stored information.

FSG	File System Gateway—A software component of the StorageGRID system. The FSG service enables standard network file systems to interface with the grid. The FSG service is found on the Gateway Node.
FSG replication group	A replication group is a group of FSGs that provide grid access to a specified set of clients. Within each replication group, there is a primary FSG (or a primary FSG cluster) and one or more secondary FSGs. The primary FSG allows clients read and write access to the grid, while storing file system information (file pointers) for all files saved to the grid. The secondary FSG "replicates" file system information, and backs up this information to the grid on a regular schedule.
Gateway Node	A building block of the StorageGRID system. The Gateway Node provides connectivity services for NFS/CIFS file systems and the HTTP protocol.
Gateway Node replication group	See "FSG replication group".
GDU	Grid Deployment Utility—A StorageGRID software utility used to facilitate the installation and update of software on all grid nodes. GDU is installed and available on the primary Admin Node.
GPT	Grid Provisioning Tool— a software tool included with StorageGRID software that permits you to provision a grid for installation, upgrade, maintenance, or expansion. GPT creates and maintains an encrypted "repository" of information about the grid that is required to maintain the grid and recover failed grid nodes.
Grid Designer	A Microsoft Windows based application used to create the configuration information needed to install, expand, and maintain a grid. It produces a grid specification file containing the grid's configuration details (in an XML format) required for the successful deployment of a StorageGRID system.
Grid ID signed text block	A BASE64 encoded block of cryptographically signed data that contains the grid ID which must match the grid ID (`gid`) element in the grid specification file. See also "provisioning".
grid node	The name of the StorageGRID system building blocks, for example Admin Node or Control Node. Each type of grid node consists of a set of services running on a server.

Grid Specification File	An XML file that provides a complete technical description of a specific grid deployment. It describes the grid topology, and specifies the hardware, grid options, server names, network settings, time synchronization, and gateway clusters included in the grid deployment. The Deployment Grid Specification file is used to generate the files needed to install the grid.
Grid Task	A managed sequence of actions that are coordinated across a grid to perform a specific function (such as adding new node certificates). Grid Tasks are typically long-term operations that span many entities within the grid. See also "Task Signed Text Block".
HAGC	High Availability Gateway Cluster—An HAGC is a primary gateway cluster that consists of a main FSG and a supplementary FSG. A high availability gateway replication group optionally includes one or more secondary FSGs.
HCAC	High Capacity Admin Cluster—An HCAC is the clustering of a reporting Admin Node and processing Admin Node. The result is an increase to a grid's capacity for grid services and thus grid nodes. See also "reporting Admin Node", "processing Admin Node", and "Admin Node".
HTTP	Hyper-Text Transfer Protocol—A simple, text based client/server protocol for requesting hypertext documents from a server. This protocol has evolved into the primary protocol for delivery of information on the World Wide Web.
HTTPS	Hyper-Text Transfer Protocol, Secure—URIs that include HTTPS indicate that the transaction must use HTTP with an additional encryption/authentication layer and often, a different default port number. The encryption layer is usually provided by SSL or TLS. HTTPS is widely used on the internet for secure communications.
ILM	Information Lifecycle Management—A process of managing content storage location and duration based on content value, cost of storage, performance access, regulatory compliance and other such factors.
inode	On Unix/Linux systems, a data structure that contains information about each file, for example, permissions, owner, file size, access time, change time, and modification time. Each inode has a unique inode number.
KVM	Keyboard, Video, Mouse—A hardware device consisting of a keyboard, LCD screen (video monitor), and mouse that permits a user to control all servers in a rack.

LAN	Local Area Network—A network of interconnected computers that is restricted to a small area, such as a building or campus. A LAN may be considered a node to the Internet or other wide area network. Contrast with WAN.
latency	Time duration for processing a transaction or transmitting a unit of data from end to end. When evaluating system performance, both throughput and latency need to be considered. See also "throughput".
LDR	Local Distribution Router—A software component of the StorageGRID system. The LDR service manages the storage and transfer of content within the grid. The LDR service is found on the Storage Node.
LUN	See "object store".
main primary FSG	In an HAGC, the FSG that is configured to be the active primary FSG by default.
metadata	Information related to or describing an object stored in the grid, for example file ingest path or ingest time.
metadata replication	In a grid that uses metadata replication, a CMS makes copies of metadata on the subset of CMSs that are in its CMS replication group, and then applies the grid's ILM policy to content metadata. In the NMS MI, CMSs that use metadata replication display the Metadata component. Called "distributed CMS" in a previous release.
metadata synchronization	In a grid that uses metadata synchronization, a CMS synchronizes metadata with all other read-write CMSs in the grid. Called "synchronized CMS" in a previous release.

NOTE Metadata synchronization is deprecated.

MI	Management Interface—The web-based interface for managing and monitoring the StorageGRID system provided by the NMS software component. See also "NMS".
namespace	A set whose elements are unique names. There is no guarantee that a name in one namespace is not repeated in a different namespace.
nearline	A term describing data storage that is neither "online" (implying that it is instantly available like spinning disk) nor "offline" (which could include offsite storage media). An example of a nearline data storage location is a tape that is loaded in a tape library, but is not necessarily mounted.

NFS	Network File System—A protocol (developed by SUN Microsystems) that enables access to network files as if they were on local disks.
NMS	Network Management System—A software component of the StorageGRID system. The NMS service provides a web-based interface for managing and monitoring the StorageGRID system. The NMS service is found on the Admin Node (both the reporting and processing Admin Nodes in an HCAC). There are three types of NMS service: consolidated, reporting, and processing. See also "MI" and "Admin Node".
node ID	An identification number assigned to a grid service within the StorageGRID system. Each service (such as an CMS or ADC) in a single grid must have a unique node ID. The number is set during system configuration and tied to authentication certificates.
NTP	Network Time Protocol—A protocol used to synchronize distributed clocks over a variable latency network such as the internet.
object store	A configured file system on a disk volume. The configuration includes a specific directory structure and resources initialized at system installation.
object segmentation	A StorageGRID process that splits a large object into a collection of small objects (segments) and creates a segment container to track the collection. The segment container contains the UUID for the collection of small objects as well as the header information for each small object in the collection. All of the small objects in the collection are the same size. See also "segment container".
OID	Object Identifier—The unique identifier of an object.
primary Admin Node	Admin Node that hosts the CMN service. There is one per grid. In an HCAC, the CMN service is hosted by the primary reporting Admin Node. See also "Admin Node" and "HCAC".
primary FSG	In an FSG replication group, the FSG that provides read-write services to clients. See also "FSG replication group".
processing Admin Node	Performs attribute and configuration processing that is passed on to the reporting Admin Node as part of a High Capacity Admin Cluster. See also "reporting Admin Node" and "HCAC".
processing NMS	Hosted by the processing Admin Node. Provides attribute and data processing functionality. Only operates in conjunction with a reporting Admin Node and the reporting NMS. See also "NMS".

provisioning	The process of editing the Grid Specification File (if required) and generating a new or updated SAID package and GPT repository. This is done on the primary Admin Node using the provision command. The new or updated SAID package is saved to the Provisioning Media. See also "Grid Specification File" and "SAID".
quorum	A simple majority: 50% + 1 of the total number in the grid. In StorageGRID software, some functionality may require a quorum of the total number of some types of service to be available.
reporting Admin Node	Reports attribute and configuration information to web clients as part of a High Capacity Admin Cluster. See also "processing Admin Node" and "HCAC".
reporting NMS	Hosted by the reporting Admin Node. Reports status information about the grid and provides a browser-based interface. Only operates in conjunction with a processing Admin Node and the processing NMS. See also "NMS".
SAID	Software Activation and Integration Data—Generated during provisioning, the SAID package contains site-specific files and software needed to install a grid.
Samba	A free suite of programs which implement the Server Message Block (SMB) protocol. Allows files and printers on the host operating system to be shared with other clients. For example, instead of using telnet to log in to a Unix machine to edit a file there, a Windows user might connect a drive in Windows Explorer to a Samba server on the Unix machine and edit the file in a Windows editor. A Unix client called "smbclient", built from the same source code, allows FTP-like access to SMB resources.
SATA	Serial Advanced Technology Attachment—A connection technology used to connect servers and storage devices.
SCSI	Small Computer System Interface— A connection technology used to connect servers and peripheral devices such as storage systems.
secondary FSG	A read-only FSG that may also perform backups of the FSG replication group. See also "FSG replication group".
security partition	If enabled, access to content ingested into the grid is restricted to the application, HTTP client, or FSG replication group that ingested the object.

segment container	An object created by StorageGRID during the segmentation process. Object segmentation splits a large object into a collection of small objects (segments) and creates a segment container to track the collection. A segment container contains the UUID for the collection of segmented objects as well as the header information for each segment in the collection. When assembled, the collection of segments creates the original object. See also "object segmentation".
server	Used when referring specifically to hardware.
Server Manager	Application that runs on all grid servers, supervises the starting and stopping of grid services, and monitors all grid services on the server.
service	A unit of the StorageGRID software such as the ADC, CMS or SSM.
SGAPI	StorageGRID Application Programming Interface—A set of commands and functions, and their related syntax, that provides HTTP clients with the ability to connect directly to the StorageGRID system (to store and retrieve objects) without the need for a Gateway Node.
SLES	SUSE Linux Enterprise Server—A commercial distribution of the SUSE Linux operating system, used with the StorageGRID system.
SQL	Structured Query Language— An industry standard interface language for managing relational databases. An SQL database is one that supports the SQL interface.
ssh	Secure Shell— A Unix shell program and supporting protocols used to log in to a remote computer and execute commands over an authenticated and encrypted channel.
SSM	Server Status Monitor—A unit of the StorageGRID software that monitors hardware conditions and reports to the NMS. Every server in the grid runs an instance of the SSM. The SSMS service is present on all grid nodes.
SSL	Secure Socket Layer—The original cryptographic protocol used to enable secure communications over the internet. See also "TLS".
standby primary FSG	In an HAGC, the FSG that is available to take over and provide read-write services to clients in event of the failure of the active primary FSG.
Storage Node	A building block of the StorageGRID system. The Storage Node provides storage capacity and services to store, move, verify, and retrieve objects stored on disks.

StorageGRID®	A registered trademark of NetApp Inc. for their fixed-content storage grid architecture and software system.
StorageGRID API	See “SGAPI”.
storage volume	See “object store”.
supplementary primary FSG	In an HAGC, the FSG that is configured to be the standby primary FSG by default.
SUSE	See “SLES”—SUSE Linux Enterprise Server.
synchronized CMS	See “metadata synchronization”.
Task Signed Text Block	A BASE64 encoded block of cryptographically signed data that provides the set of instructions that define a grid task.
TCP/IP	Transmission Control Protocol / Internet Protocol—A process of encapsulating and transmitting packet data over a network. It includes positive acknowledgment of transmissions.
throughput	The amount of data that can be transmitted or the number of transactions that can be processed by a system or subsystem in a given period of time. See also “latency”.
TLS	Transport Layer Security—A cryptographic protocol used to enable secure communications over the internet. See RFC 2246 for more details.
transfer syntax	The parameters, such as the byte order and compression method, needed to exchange data between systems.
TSM	Tivoli® Storage Manager—IBM storage middleware product that manages storage and retrieval of data from removable storage resources.
URI	Universal Resource Identifier—A generic set of all names or addresses used to refer to resources that can be served from a computer system. These addresses are represented as short text strings.
UTC	A language-independent international abbreviation, UTC is neither English nor French. It means both “Coordinated Universal Time” and “Temps Universel Coordonné”. UTC refers to the standard time common to every place in the world.
UUID	Universally Unique Identifier—Unique identifier for each piece of content in the StorageGRID. UUIDs provide client applications with a

	content handle that permits them to access grid content in a way that does not interfere with the grid's management of that same content. A 128-bit number which is guaranteed to be unique. See RFC 4122 for more details.
VM	Virtual Machine—A software platform that enables the installation of an operating system and software, substituting for a physical server and permitting the sharing of physical server resources amongst several virtual "servers".
XFS	A scalable, high performance journaled file system originally developed by Silicon Graphics.
WAN	Wide Area Network—A network of interconnected computers that covers a large geographic area such as a country. Contrast with "LAN".
XML	eXtensible Markup Language—A text format for the extensible representation of structured information; classified by type and managed like a database. XML has the advantages of being verifiable, human readable, and easily interchangeable between different systems.

Index