

StorageGRID[®] 9.0.x

Software Upgrade 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: www.netapp.com
Feedback: doccomments@netapp.com

Part number: 215-06840_D0
October 2013

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

Copyright and trademark information	2
1 What's New	7
Introduction	7
Intended Audience	7
New Features in StorageGRID 9.0.0	8
API Gateway Nodes	8
CMS Database Expansion	8
Metadata Update	8
Audit Messages	8
HTTP Metadata Attributes	8
Enable Object Location Indexing	9
FSG and LDR Health Check Timeout Defaults Updated	9
Grid Designer Updates	9
Grid Size Increase	9
HTTP API Renamed StorageGRID API	9
Last Access Time Metadata for ILM Rules	10
Support for CDMI	10
CDMI Audit Messages	10
CDMI Component	10
Support for IPv6	11
Samba 3.6.3	11
Refresh Grid Nodes onto Virtual Machine	11
Updates to High Availability Gateway Nodes	11
Unicast Heartbeat	11
Updates to Grid Management Tree	12
New Features in StorageGRID 9.0.1	13
Control Node Hardware Refresh	13
SLES 11 SP2 64-Bit	13
Deprecated Features	13
Obsoleted Features	14
2 Prepare for Software Upgrade	15
Introduction	15
Upgrade Versus Update	15
Upgrade of the StorageGRID System	15
Update of the StorageGRID System	15
Prepare for Upgrade	16
Prepare for Upgrade Checklist	16
Verify Installed Version of StorageGRID Software	16
Verify Installed Version of SLES	17

Procedure	17
Upgrade and Update Paths for StorageGRID and SLES Software ...	17
Gather Required Materials	19
Schedule Downtime	20
Check the Condition of the Grid	20
Check Free Space Required for the Upgrade	21
Disable, Reschedule, or Cancel FSG Backups.....	22
Disable FSG Backups	22
Reschedule FSG Backups.....	22
Cancel an Active FSG Backup.....	23
Understand How the Upgrade Affects the Grid	23
Client Access	23
Operational Gateway Node Upgrade	23
NMS MI Access During the Upgrade of an HCAC	24
NMS Alarms	24
Hot Fixes	24
Configuration Restrictions	24
Gateway Node Replication Groups and Storage Nodes	25
Last Access Time Metadata	26
Custom Metadata	26
Expansion During Upgrade	26
3 Upgrade Software	27
Introduction	27
Load Software Distribution on the Primary Admin Node	28
Primary Admin Node is on a Virtual Machine	29
Primary Admin Node is on a Physical Server (With CDs)	30
A Copy of the ISOs is on the Primary Admin Node	30
Provision the Grid	31
Admin Node is on a Virtual Machine	32
Admin Node is on a Physical Server	34
Start the Software Upgrade Grid Task	37
Upgrade the Grid Software with GDU	38
After Grid Software Upgraded on all Grid Nodes	42
4 Gateway Node Upgrades	43
Introduction	43
Upgrade Basic Gateway Replication Groups	43
Fail Over From Primary to Secondary	44
Fail Back to Original Primary	46
Upgrade High Availability Gateway Replication Groups	47
Convert Heartbeat for High Availability Gateway Cluster	48
Gateway Node is Hosted by a Virtual Machine.....	48
Gateway Node is on a Physical Server	49
Fail Over Within the HAGC	49

5	Expansion During Upgrade	51
	Scope and Limitations	51
	Service Packs	51
	Gateway Nodes	51
	Storage Nodes	52
	Adding Grid Nodes During an Upgrade	53
6	Manually Upgrading	55
	Introduction	55
	Upgrading the Grid Software Manually	55
7	Complete the Upgrade	59
	Final Tasks	59
	Confirm Upgrade and Make Configuration Changes	59
	Enable Object Location Indexing	60
	Upgrade Operating System	61
8	Apply Service Packs	63
	Introduction	63
	Update Sequence	63
	About Updating Software on Gateway Nodes	64
	Apply the Service Pack	65
9	Upgrade the Operating System	69
	Introduction	69
	Server or Virtual Machine Update Sequence	70
	Update to SLES 10 SP3	71
	Prerequisites	71
	Procedure	71
	Update to SLES 11 SP2 (64-Bit)	73
	Prerequisites	74
	Procedure	74
10	Troubleshooting	77
	Service Hangs	77
	Terminate Service	77
	Provisioning Failures	78
	Server Crashes or Fails to Start	78
	Grid Node Fails	78
	Grid Task Pauses With Error	78
	Ingest or Data Retrieval is Interrupted	79
	Upgrading the Grid Software Manually	79
	Script Exceptions	79
	NMS MI is Unavailable on Grid with HCAC	80

A	How to Use GDU	81
	Start GDU	81
	GDU User Interface	83
	Entering Commands in GDU	85
	Install Drivers with GDU	85
	Close GDU	86
	GDU Troubleshooting	87
	GDU Display Problems	87
	Problems with Server Status	87
	GDU Log Files	87
	Missing GDU Task	87
	Troubleshooting with screen in Multi Display Mode	88
	About load_cds.py	89
	Copy ISO Files in Multi-Site Environment	89
	Glossary	93
	Index	105

What's New

Introduction

This guide describes how to upgrade from StorageGRID 8.5.2 or later software to StorageGRID 9.0.x, how to upgrade the operating system (OS) to SLES 10 service pack 3 64-bit and SLES 11 SP2 (64-bit), and how to apply 9.0.x service packs.

Read this guide in full before beginning an upgrade and become familiar with all of the steps and requirements needed to upgrade a grid.

Intended Audience

The content of this guide is intended for technical personnel trained to install and support the StorageGRID system. This guide assumes that you are familiar with the StorageGRID system's general design, configurations and options, and have received training in grid provisioning, installation, and integration.

An advanced level of computer literacy is assumed, including knowledge of Linux/Unix command shells, networking, and server hardware setup and configuration. This guide also assumes your familiarity with terms related to computer operations and programming, network communications, and operating system file operations. Acronyms are widely used.

New Features in StorageGRID 9.0.0

API Gateway Nodes

Grid Designer now lets you design grids with API Gateway Nodes. API Gateway Nodes provide read-write access to the grid for external SGAPI and CDMI applications, but do not support NFS and CIFS applications.

CMS Database Expansion

A CMS database can be expanded to a maximum size of 800 GiB. In order to access the new tablespace, a CMS must use metadata replication and have Object Location Indexing enabled.

Metadata Update

Metadata can be added, updated, or deleted from an object via the StorageGRID API (SGAPI). For more information, see the *StorageGRID API Reference*.

NOTE Metadata update is disabled until the software upgrade completes. If you use the metadata update feature prior to completing the upgrade to StorageGRID 9.0, an HTTP 503 Service Unavailable error appears.

Audit Messages

The following new audit messages may be generated when adding, updating, or deleting metadata:

- HGMD
- HPMD
- OMDU

For more information, see the *Audit Message Reference*.

HTTP Metadata Attributes

<Storage Node> ► LDR ► HTTP ► Overview ► Main now displays HTTP Metadata PUT and HTTP Metadata GET attributes to count attempts, suc-

cesses, and failures when updating, retrieving, adding or deleting metadata via the SGAPI.

<Control Node> ► CMS ► Metadata ► Overview ► Main now displays via the Total Object Metadata Updates attribute the total number of object metadata updates for objects managed by the CMS.

Enable Object Location Indexing

For grids upgraded to Release 9.0, Object Location Indexing must be enabled for each CMS. The result is improved performance when running an object lost task, performing foreground verification, or decommissioning Storage Nodes.

NOTE Object location indexing is only available for CMSs that use metadata replication.

FSG and LDR Health Check Timeout Defaults Updated

The default setting for both the FSG Health Check Timeout and the LDR Health Check Timeout has been updated to 200 seconds.

Grid Designer Updates

For information on new features and updates to Grid Designer, see the *Grid Designer User Guide*.

Grid Size Increase

The maximum number of bindings and thus grid services that an StorageGRID deployment can support has been increased to 700. The maximum number of bindings that a StorageGRID deployment can support depends on the grid's topology and selected hardware. Grid Designer displays a warning if the proposed grid topology exceeds the maximum number of bindings — depending on hardware selected, this may be less than the maximum of 700 bindings.

HTTP API Renamed StorageGRID API

To avoid confusion with CDMI, the HTTP API is now known as the StorageGRID API (SGAPI) in StorageGRID documentation. For

example, the *HTTP API Reference* is now called the *StorageGRID API Reference*.

Last Access Time Metadata for ILM Rules

Last Access Time metadata is now available when you create rules for Information Lifecycle Management (ILM). Last Access Time metadata identifies the date that the content was last retrieved by an HTTP client with Last Access Time enabled. (See [Grid Management](#) ► [Grid Configuration](#) ► [HTTP Management](#) ► [Permissions](#) ► [Overview](#) ► [Main](#) ► [HTTP /CDMI and /UUID Namespace](#).) You can use the Last Access Time metadata as a reference time and as a filter when you create ILM rules. Last Access Time metadata helps you identify and move old content to an appropriate storage pool.

NOTE The Last Access Time feature is available after you complete the software upgrade. The Last Access Time feature is not available while the software upgrade is in progress.

Support for CDMI

You can now use the Cloud Data Management Interface (CDMI) to access the grid. For more information, see the *CDMI Reference*.

CDMI Audit Messages

The following audit messages have been added for CDMI:

- CDMD — CDMI DELETE transaction
- CDMG — CDMI GET transaction
- CDMP — CDMI POST transaction
- CDMU — CDMI PUT transaction

For more information, see the *Audit Message Reference*.

CDMI Component

There is a new CDMI component at [<Storage Node>](#) ► [LDR](#) ► [CDMI](#) that tracks CDMI transactions; for example, total operations, operation rate, and operations failed.

Support for IPv6

IPv6 is only supported for grids running SLES 11 SP2 64-bit.

IPv6 is supported for external interfaces such as the NMS MI, the SG API, CDMI, and external NTP sources. For more information, see the *Administrator Guide*.

Samba 3.6.3

For new and upgraded installations of the StorageGRID 9.0 system, Samba has been upgraded to 3.6.3. By default, SMB 2.0 is listed as disabled in the `/etc/samba/smb.conf` file.

Refresh Grid Nodes onto Virtual Machine

During a hardware refresh, you can refresh grid nodes to virtual machines. You can refresh Admin Nodes, Control Nodes, Gateway Nodes, and Storage Nodes onto virtual machines. However, you cannot refresh combined grid nodes, such as Admin/Gateway Nodes, onto virtual machines. To refresh combined grid nodes to virtual machines, you must split them (with Grid Designer) into individual grid nodes. For details, see the *Expansion Guide*.

Updates to High Availability Gateway Nodes

Unicast Heartbeat

New and upgraded High Availability Gateway Node Clusters (HAGC) now use unicast rather than broadcast heartbeat. When virtual machines host HAGCs, one or more HAGCs can share a dedicated unicast heartbeat network for intra-cluster communications. A crossover cable is no longer needed to connect the virtual machine hosts. However, HAGCs hosted on physical servers still require crossover cables for heartbeat to maintain the highest levels of reliability.

Multiple HAGCs hosted on virtual machines can share a unicast heartbeat network. Therefore, the main and supplementary heartbeat IP addresses must be unique. Grid Designer 9.0.x now warns you when it finds duplicate heartbeat IP addresses.

Updates to Grid Management Tree

To accommodate the introduction of CDMI, the Grid Management tree has been updated to include a new HTTP Management branch. As well, several components from the Grid Configuration branch have been moved to the new HTTP Management branch and pages updated. For example, UUID Namespace on the Grid Management ► HTTP Management ► Permissions (was Grid Management ► Grid Configuration ► HTTP Advanced) has been updated to HTTP /CDMI and / UUID Namespaces.

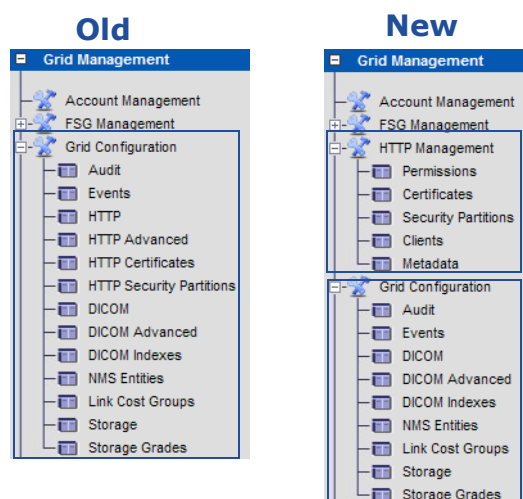


Figure 1: Grid Management Tree Updates

Table 1: Changes to Grid Management ►

Was: Grid Configuration ►	Now: HTTP Management ►
HTTP Advanced	Permissions
HTTP Certificates	Certificates
HTTP Security Partitions	Security Partitions
HTTP	Clients
None. Was section of HTTP Advanced.	Metadata

For more information, see the *Grid Primer* and the *Administrator Guide*.

New Features in StorageGRID 9.0.1

Control Node Hardware Refresh

When performing a hardware refresh of multiple Control Nodes, you now only need to update the grid specification file and provision the grid once. For more information on hardware refresh, see the *Expansion Guide*.

SLES 11 SP2 64-Bit

New installations of the StorageGRID 9.0.1 system only support the SLES 11 SP2 64-bit operating system.

- For StorageGRID 9.0.0, you cannot upgrade from SLES 10 to SLES 11.
- For StorageGRID 9.0.1 (or later service packs), you can upgrade from SLES 10 to SLES 11 SP2 64-bit.

Deprecated Features

The following features are no longer supported for newly installed grids, but are not removed from the NetApp StorageGRID product:

- Installation of the StorageGRID system on physical servers

NOTE A StorageGRID system installed on physical servers before 9.0.x, can be upgraded to release 9.0.x and remain on physical servers. However, it is recommended that all grid nodes eventually be refreshed to virtual machines.

- Deduplication
- File recovery
- Metadata synchronization
- SUSE Linux 10 service pack 3 32-bit for expansion or maintenance
- SUSE Linux 10 service pack 3 64-bit for installation

For information on the currently supported version of SLES, see the see the Interoperability Matrix Tool (IMT) at support.netapp.com/matrix

- VMware ESX/ESXi 4.0.

For information on the current supported version of VM ware software, see the Interoperability Matrix Tool (IMT) at support.netapp.com/matrix.

Obsoleted Features

The following features are not supported and are removed from the NetApp StorageGRID system:

- DICOM
- Distributed File System Gateway (DFSG)
- SUSE Linux 10 service packs 1 and 2 32-bit for installation, expansion, or maintenance
- SUSE Linux 10 service pack 3 32-bit for installation

Note that for expansion and maintenance procedures, SUSE Linux 10 service pack 3 32-bit is deprecated.

- VMware ESX 3.5

For information on the current supported version of VM ware software, see the Interoperability Matrix Tool (IMT) at support.netapp.com/matrix.

Prepare for Software Upgrade

Introduction

Before beginning the upgrade process, determine whether you need to perform an upgrade or an update of the current StorageGRID system. As well determine whether you need to upgrade the SLES operating system (OS). Depending on the version of the StorageGRID system that you are currently running and the version of the SLES operating system that the StorageGRID system is running on, there are different procedural tasks that you must perform.

Upgrade Versus Update

Whether or not you need to perform an upgrade or an update of the StorageGRID system depends on the currently installed version of the StorageGRID system.

Upgrade of the StorageGRID System

If you are currently running a version of the StorageGRID system that is earlier than 9.0 then you must perform an upgrade.

Note that the only supported upgrade path is from StorageGRID 8.5.2 or later to StorageGRID 9.0.x. You cannot apply the upgrade to releases earlier than StorageGRID 8.5.2. If the grid is running a release earlier than StorageGRID 8.5.2, first upgrade the grid to StorageGRID 8.5.3 or later, and then upgrade the grid to StorageGRID 9.0.x. For more information, see [Chapter 3: “Upgrade Software”](#).

Update of the StorageGRID System

If you are currently running StorageGRID 9.0.x software, you do not upgrade grid software, but rather update it by applying the latest

service pack. For more information and procedures on the update process, see [Chapter 8: “Apply Service Packs”](#).

Prepare for Upgrade

Prepare for Upgrade Checklist

Complete the following checklist to prepare for a software upgrade of the StorageGRID system.

Table 2: Prepare for Upgrade Checklist

✓	Step	Task	See
Prepare for Software Upgrade			
	1.	Read the <i>Upgrade Guide</i> in its entirety and become familiar with all of the steps and requirements needed to successfully complete an upgrade of the StorageGRID system. Plan your upgrade before you begin.	This guide.
	2.	Verify installed version of StorageGRID software and procedural path (upgrade or update).	page 17
	3.	Verify installed version of SLES.	page 19
	4.	Gather all materials required to perform the upgrade.	page 19
	5.	Schedule downtime.	page 20
	6.	Check the condition of the grid.	page 20
	7.	Become familiar with how the upgrade process affects a running grid.	page 23

Verify Installed Version of StorageGRID Software

The current version of StorageGRID software must be 8.5.2

Before beginning the upgrade process, determine the currently installed version of StorageGRID software and the upgrade path needed to successfully complete the upgrade process.

1. In the NMS MI, go to **SSM ► Services ► Overview ► Main**.
 - a. For the installed version of StorageGRID, check the value of storage-grid-release in the Packages table.
 - b. Determine if you need to perform an upgrade or an update of the StorageGRID software.

Verify Installed Version of SLES

Running StorageGRID software on the correct version of SLES is critical to the successful operation of the system. Before upgrading to StorageGRID 9.0.x, confirm the version of SLES that is currently running. If the StorageGRID system is not running on at least SLES 10 SP3 (32 or 64-bit), perform OS upgrades before upgrading to StorageGRID 9.0.x.

After completing the upgrade of all grid nodes to StorageGRID 9.0.x, it is recommended that you upgrade the operating system from SLES 10 SP3 (64-bit) to SLES 11 SP2 (64-bit). Before upgrading the operating system to SLES 11 SP2 (64-bit) all grid nodes must be upgraded to StorageGRID 9.0.x. Failure to complete the upgrade of all grid nodes to StorageGRID 9.0.x before upgrading the SLES operating system may result in the failure of your StorageGRID system.

Note that you cannot upgrade StorageGRID software and the underlining operating system in parallel on different grid nodes. The upgrade to StorageGRID 9.0.x must finish in its entirety before upgrading the SLES operating system. For more information on upgrade paths and procedures see, [“Upgrade and Update Paths for StorageGRID and SLES Software”](#) below and [Chapter 9: “Upgrade the Operating System”](#).



WARNING Do not upgrade to SLES 11 SP2 (64-bit) until all grid nodes are upgraded to StorageGRID 9.0.x.

Procedure

You cannot upgrade from SLES 10 SP3 (32-bit) to SLES 11 SP2 (64-bit).

1. In the NMS MI, go to **SSM ► Services ► Overview ► Main**.
 - a. For the installed version of SLES, check the value of Operating System.
 - b. If it is not SLES 10 SP3 (32 or 64-bit), you must upgrade SLES before upgrading or updating to StorageGRID 9.0.x.

Upgrade and Update Paths for StorageGRID and SLES Software

The following diagram illustrates the supported paths to upgrade a StorageGRID system to StorageGRID 9.0.x running on SLES 11 SP2 64-bit.

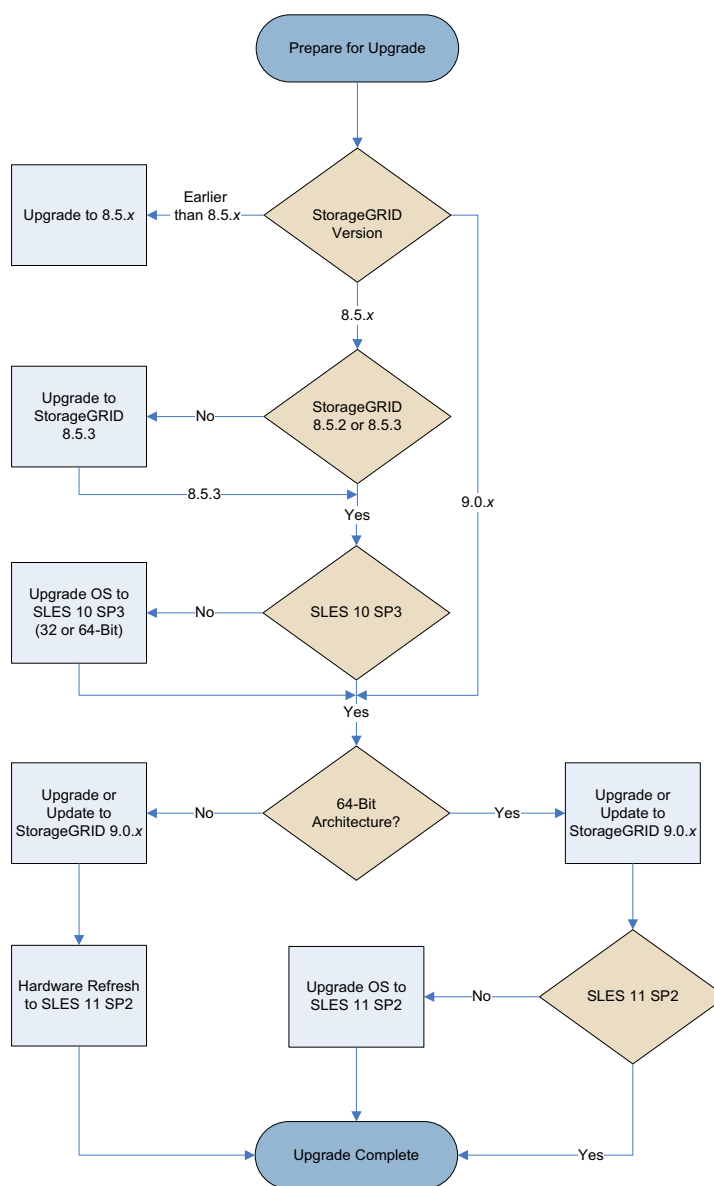


Figure 2: Supported Upgrade and Update Paths

Gather Required Materials

Before you begin, ensure that you have all required materials:

Table 3: Materials Checklist for Software Upgrade

✓	Item	Notes
	StorageGRID Software CD for the base version	Version 9.0
	Enablement Layer for StorageGRID Software CD	Version 9.0
	If any service packs have been released, the latest Service Pack CDs	<p>There are two Service Pack CDs:</p> <ul style="list-style-type: none"> StorageGRID 9.0.x. Software Service Pack Enablement Layer for StorageGRID 9.0.x. Software Service Pack <p>where x is the service pack number. The service pack number is identical for both CDs.</p>
	Documentation	<ul style="list-style-type: none"> <i>Administrator Guide</i> <i>Expansion Guide</i> <i>Release Notes</i> <hr/> <p>NOTE Carefully read the <i>Release Notes</i> prior to starting the upgrade.</p> <hr/>
	Passwords.txt file for the grid	
	Provisioning passphrase	
	Provisioning USB flash drive and Backup Provisioning USB flash drive	Get a copy of the most recent provisioning USB flash drives. These flash drives are updated each time the grid is modified.
	Keyboard and monitor	Used to upgrade unconnected Gateway Nodes and Storage Nodes.
	Service laptop computer	<p>Laptop must have:</p> <ul style="list-style-type: none"> Microsoft® Windows® operating system Network port Supported browser for StorageGRID 9.0.x telnet/ssh client to run GDU. For example, PuTTY, available from: http://www.chiark.greenend.org.uk/~sgtatham/putty/

Schedule Downtime

While the upgrade is designed as a rolling upgrade that enables the grid to remain continuously available, clients are unable to access the grid for a brief period of time when the software on operational Gateway Nodes is upgraded. For more information, see [“Client Access”](#) on page 23.

Check the Condition of the Grid

Before upgrading the StorageGRID software, perform the checks listed in [Table 4](#) to ensure that the grid is running normally, that all connected servers are operational, that the grid is running a supported version of SLES, and that the grid can accommodate the upgrade.

Table 4: Condition of the Grid Checklist

✓	Item	Notes	See
Grid Condition			
	Free Space	Check the amount of free space on the root partition of every server.	page 21
	Alarms	Address any active alarms in the grid.	<i>Troubleshooting Guide</i>
	FSGs	Confirm that no FSG backup is scheduled to take place during the upgrade.	page 22 <i>Maintenance Guide</i>
		If an FSG has recently been restored from backup, ensure that restoration has completed.	
		If the secondary FSG is acting as primary, correct the problem that led to the failover and restore the FSGs to their original roles before proceeding.	

Table 4: Condition of the Grid Checklist (cont.)

✓	Item	Notes	See
	Grid tasks	<p>Go to CMN ► Grid Tasks ► Configuration and confirm that no grid tasks are running or pending. The only exceptions are:</p> <ul style="list-style-type: none"> • LDR content rebalancing grid task (LBAL) • ILM evaluation (ILME) <p>These grid tasks can run concurrently with the Software Upgrade (SWUP) grid task. If any other grid tasks are running or pending, wait for them to complete or release their lock, run the grid tasks, or abort them as appropriate.</p> <p>After you start the Software Upgrade (SWUP) grid task and finish the upgrade of the primary Admin Node and all Control Nodes, you can start the Grid Expansion (GEXP) grid task and run it concurrently with the Software Upgrade (SWUP) grid task.</p>	<i>Administrator Guide</i>
	Decommissioning and refreshing servers	If a server is being decommissioned or refreshed, complete all decommissioning/hardware refresh procedures before proceeding with the upgrade.	<i>Expansion Guide</i> <i>Maintenance Guide</i>

Check Free Space Required for the Upgrade

Before beginning the upgrade, ensure that there is a minimum of 500 MB space available on the root partition of each server in the grid.

If necessary, you can create free space by:

- Removing excess NTP log files from `/var/lib/ntp/var/log/ntpstats`. Only the log files from the last seven days are needed.
- Uninstalling IBM Director Agent
- Manually moving kernel source files, if installed, to `/var/local/src` and creating a symlink. At the command line of each server, enter:

```
if [ -d /usr/src ]; then
    mv /usr/src /var/local/src
    ln -s /var/local/src /usr/src
fi
```

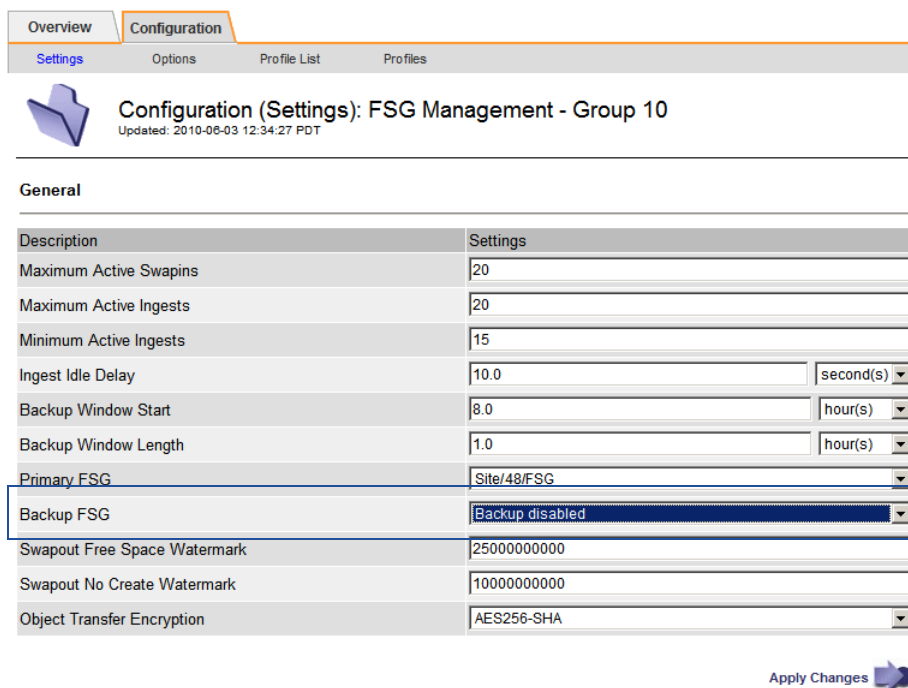
After the upgrade is complete, reinstall IBM Director Agent if required.

Disable, Reschedule, or Cancel FSG Backups

An upgrade completes more slowly if an FSG backup is in progress. Either disable backups temporarily or reschedule the backup. You can also cancel an active backup.

Disable FSG Backups

1. Log in to the NMS MI using the Vendor account.
2. Go to **Grid Management ► FSG Management ► <Replication_Group> ► Configuration ► Settings**.
3. Change the value of **Backup FSG** to **Backup disabled**.



Configuration (Settings): FSG Management - Group 10
Updated: 2010-06-03 12:34:27 PDT

General

Description	Settings
Maximum Active Swapins	20
Maximum Active Ingests	20
Minimum Active Ingests	15
Ingest Idle Delay	10.0 second(s)
Backup Window Start	8.0 hour(s)
Backup Window Length	1.0 hour(s)
Primary FSG	Site/48/FSG
Backup FSG	Backup disabled
Swapout Free Space Watermark	25000000000
Swapout No Create Watermark	10000000000
Object Transfer Encryption	AES256-SHA

Apply Changes

Figure 3: Disabling FSG Backups

4. Click **Apply Changes**.

Reschedule FSG Backups

1. Log in to the NMS MI using the Vendor account.
2. Go to **Grid Management ► FSG Management ► <Replication_Group> ► Configuration ► Settings**.
3. Change the value of **Backup Window Start**.
4. Click **Apply Changes**.

Cancel an Active FSG Backup

1. Log in to the NMS MI using the Vendor account.
2. Go to **<Backup Gateway Node> ► FSG ► Backup ► Configuration ► Main**.
3. Select **Cancel Active Backup**.

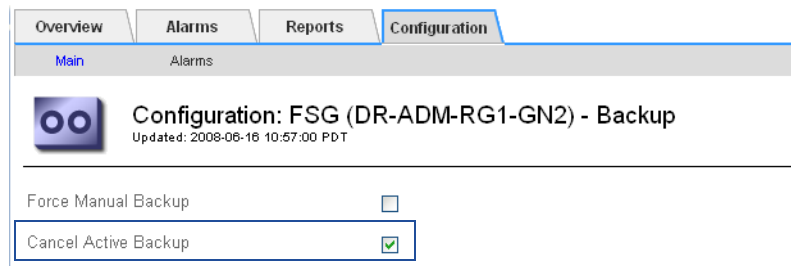


Figure 4: Canceling the Active FSG Backup

4. Click **Apply Changes**.

Understand How the Upgrade Affects the Grid

Client Access

During upgrade, the grid remains continuously available, providing a maximum number of grid services and operating transparently in a mixed-version environment. Clients can ingest and retrieve data throughout the upgrade process except for a period of time when operational Gateway Nodes are unavailable.

NFS shares may have to be remounted. The specific behavior varies depending on the NFS client implementation. For example, the NFS mount may become stale after a failover, depending on the client system's automount settings and any active transactions at the time of failover.

Operational Gateway Node Upgrade

Client services are unavailable for a time while Gateway Nodes are being upgraded. Client service is interrupted for a short period of time while an HAGC automatically fails over or while a Gateway Node is manually failed over. Plan client downtime accordingly.

NOTE If the grid does not support business continuity failover, that is, if the grid cannot be configured to fail over to a writable FSG (for example, due to lack of network connectivity from the client to the writable FSG), client access is unavailable for the duration of the primary upgrade.

For failover procedures and the upgrade order for operational Gateway Nodes, see [Chapter 4: “Gateway Node Upgrades”](#).

NMS MI Access During the Upgrade of an HCAC

When upgrading a High Capacity Admin Node Cluster (HCAC), access to the NMS MI (via that HCAC’s reporting) is lost until both the reporting and processing Admin Nodes have been upgraded.

NMS Alarms

Software upgrades can cause NMS alarms while the grid is operating in a mixed-version environment. In general, these alarms clear when the upgrade completes.

Hot Fixes

Any StorageGRID 8.5.x hotfixes up to, but not including 9.0.x, should be applied prior to starting the upgrade.

Unless otherwise directed, StorageGRID 9.0.x hotfixes should be applied after the upgrade.



WARNING Do not apply hot fixes in the middle of a software upgrade. Results are unpredictable and the grid may not function as intended.

Configuration Restrictions

During an upgrade to release StorageGRID 9.0.x, the following restrictions apply:

- Do not make configuration changes to Gateway Nodes until the entire replication group has been upgraded.

- Do not edit any Grid Options (Grid Management ► Grid Configuration ► Configuration ► Main) until the upgrade completes.
- Do not enable any new features (for example, network transfer compression or object segmentation) until the upgrade completes.

Gateway Node Replication Groups and Storage Nodes

Gateway Node replication groups and Storage Nodes can be upgraded or added at any time after the upgrade has started. Note, however, that HAGCs must be converted from broadcast to unicast heartbeat before the Software Upgrade grid task is run. For more information on upgrading Gateway Node, see [Chapter 4: “Gateway Node Upgrades”](#).

An upgrade may proceed even if all Gateway Node replication groups or Storage Nodes are not currently connected to the grid. Gateway Nodes and Storage Nodes may be connected to the grid at a later date and then upgraded via GDU or Gateway Nodes and Storage Nodes can be manually updated when not connected to the grid. The upgrade completes when these unconnected Gateway Nodes and Storage Nodes are connected and grid software is started.

Connecting to the grid means establishing network connectivity between the upgraded grid node and the grid and then starting grid software on the grid node. Grid software can be started either via GDU (see [Chapter 3: “Upgrade Software”](#)) or manually (see [Chapter 6: “Manually Upgrading”](#)).

As well, the grid can be expanded with new Gateway Node replication groups and Storage Nodes at any time after the upgrade process begins and the primary Admin Node and all Control Nodes are upgraded. When adding Gateway Node replication groups or Storage Nodes during an upgrade, the expansion grid nodes must be installed with the software to which the grid is being upgraded.

Upgrade must be completed for all grid nodes before beginning a new upgrade. This means that any unconnected grid nodes must be upgraded, connected to the grid, and grid software started before a new upgrade can begin.

Last Access Time Metadata

Last access time metadata is only available after you complete a software upgrade. The last access time feature is not available while a software upgrade is in progress.

Custom Metadata

You cannot update custom metadata during a software upgrade. During an upgrade the custom metadata feature is disabled. Wait until the software upgrade completes before you update custom metadata. If you attempt to update custom metadata while a software upgrade is in progress, error 503 appears.

Expansion During Upgrade

You can add new Gateway Node replication groups and Storage Nodes to the grid while performing an upgrade. If you want to expand and upgrade the grid at the same time, for grid node types other than Gateway Nodes and Storage Nodes, complete the upgrade to Storage-GRID 9.0.x, and then expand the grid. For more information see, [Chapter 5: “Expansion During Upgrade”](#).

Upgrade Software

Procedures to provision the grid and upgrade the grid software with GDU

Introduction

This chapter describes how to upgrade the grid software from StorageGRID 8.5.2 or later to StorageGRID 9.0.x using GDU. Complete the checklist in [Table 5](#) to upgrade the software.

Before upgrading to StorageGRID 9.0.x, confirm that you are running a supported version of SLES and a supported version of the StorageGRID software. For more information, see [Chapter 2: “Prepare for Software Upgrade”](#) on page 15.



WARNING If the value of storage-grid-release (go to: **SSM ► Services ► Overview**) is not at least 8.5.2, do not use this procedure. First, upgrade to the latest StorageGRID 8.5.x software (StorageGRID 8.5.2 or later).

Table 5: Upgrade Software Checklist

✓	Step	Task	See
Confirm StorageGRID Version and Operation System			
	1.	Confirm that the version of StorageGRID software is at least 8.5.2.	page 16
	2.	Confirm that SLES has been upgraded to a supported version.	page 17
Upgrade Software			
	1.	Load the software distribution on the primary Admin Node.	page 28
	2.	Provision the grid.	page 31
	3.	Start the Software Upgrade grid task.	page 37

Table 5: Upgrade Software Checklist (cont.)

✓	Step	Task	See
	4.	<p>Upgrade the grid software:</p> <ol style="list-style-type: none"> 1. Upgrade the primary Admin Node or HCAC. For a primary HCAC, upgrade the reporting Admin Node and then the processing Admin Node. 2. Use GDU to upgrade the grid software on the remaining servers. <hr/> <p>NOTE After you upgrade the primary Admin Node or HCAC and all Control Nodes, you can add Gateway Nodes and Storage Nodes at any time while you upgrade the rest of the grid. See Chapter 5: “Expansion During Upgrade”.</p> <hr/>	page 38
	3.	Complete the upgrade	page 42

If problems unrelated to the upgrade occur during the upgrade process (such as a failed server or a network problem that leaves some servers without connectivity to the rest of the grid), correct the problem as quickly as possible and then resume the upgrade.

Load Software Distribution on the Primary Admin Node

Use the `load_cds.py` script to load a copy of the software distribution onto the primary Admin Node. For background on the `load_cds.py` script, see [“About load_cds.py” on page 89](#).

NOTE The primary Admin Node is the Admin Node hosting the CMN service. For an HCAC, the primary reporting Admin Node hosts the CMN service. There is one CMN service per grid. For more information, see the *Administrator Guide*.

This section contains the following procedures:

- [“Primary Admin Node is on a Virtual Machine” on page 29](#)
- [“Primary Admin Node is on a Physical Server \(With CDs\)” on page 30](#)
- [“A Copy of the ISOs is on the Primary Admin Node” on page 30](#)

Primary Admin Node is on a Virtual Machine

Use this procedure to load the software distribution CDs for a service pack when the primary Admin Node is installed on a virtual machine.

Prerequisites

- Ensure you have:
 - StorageGRID 9.0 Software
 - Enablement Layer for StorageGRID 9.0 Software
 - If available, StorageGRID 9.0.x Software Service Pack
 - If available, Enablement Layer for StorageGRID 9.0.x Software Service Pack
- Service laptop running vSphere Client

Procedure

1. In vSphere Client, click in the console window of the primary Admin Node's virtual machine. Access a command shell and log in as root using the password listed in the `Passwords.txt` file.
2. Insert a CD in the service laptop.
If available, use the StorageGRID 9.0.x Software Service Pack CD. Otherwise, use the StorageGRID 9.0 Software CD. The order in which you insert the CDs does not matter.



TIP Press **<Ctrl>+<Alt>** to release your 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. Enter: `load_cds.py`
Wait while the ISO image is written to the correct directory.
5. When prompted:
 - 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>**.
6. Repeat step 5 for all CDs.

7. To exit, type **n** and press **<Enter>** when prompted.
8. Log out. Enter: **exit**

Primary Admin Node is on a Physical Server (With CDs)

Prerequisites

- Ensure you have:
 - StorageGRID 9.0 Software
 - Enablement Layer for StorageGRID 9.0 Software
 - If available, StorageGRID 9.0.x Software Service Pack
 - If available, Enablement Layer for StorageGRID 9.0.x Software Service Pack

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. Remove any USB flash drives from the server.
3. Insert a CD.
If available, use the StorageGRID 9.0.x Software Service Pack CD. Otherwise, use the StorageGRID 9.0 Software CD. The order in which you insert the CDs does not matter.
4. Load the ISO image. Enter: **load_cds.py**
5. When prompted, insert the next CD, type **y**, and press **<Enter>**.
6. Repeat step 5 for all CDs.
7. To exit, type **n** and press **<Enter>** when prompted.
8. Log out of the command shell. Enter: **exit**

The order in which you load the CDs does not matter.

The next step is to provision the grid. See [“Provision the Grid”](#) below.

A Copy of the ISOs is on the Primary Admin Node

Use this procedure to load ISOs if you have already copied ISO images of the installation CDs to the primary Admin Node; for example, if the server is at a remote site and you have used `scp` to copy 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 instance `/var/local/tmp`. The `load_cds.py` script will copy the files from the directory you specify to the `/var/local/install` directory.

Prerequisites

- ISO images of the following CDs are on the primary Admin Node server:
 - StorageGRID 9.0 Software
 - Enablement Layer for StorageGRID 9.0 Software
 - If available, StorageGRID 9.0.x Software Service Pack
 - If available, Enablement Layer for StorageGRID 9.0.x Software Service Pack

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. Remove any USB flash drives from the server.
3. Load the ISO images. Enter (on one line):

```
load_cds.py <iso_software_CD_including_path>  
<iso_enablement_layer_CD_including_path>  
<iso_software_service_pack_CD_including_path>  
<iso_enablement_layer_service_pack_CD_including_path>
```

Enter multiple ISO filenames separated by a space. The order does not matter.
Wait until the ISO images are written to the correct directory.
4. Log out of the command shell. Enter: **exit**

The next step is to provision the grid. See [“Provision the Grid”](#) below.

Provision the Grid

You must provision the grid to create the files needed for the upgrade to 9.0.x. The following upgrade procedure assumes that Provisioning data is backed up onto USB flash drives. For alternative commands if this is not the case, see the “Grid Specification Files and Provisioning” appendix of the *Installation Guide*.

This section contains the following procedures:

- [“Admin Node is on a Virtual Machine” on page 32](#)
- [“Admin Node is on a Physical Server” on page 34](#)

Admin Node is on a Virtual Machine

Prerequisites

- Grid is being upgraded from release 8.5.2 or later.
- ISO images of the software CDs have been loaded onto the primary Admin Node server with `load_cds.py`. See [“Load Software Distribution on the Primary Admin Node” on page 28](#).
- The condition of the grid has been reviewed. See [“Check the Condition of the Grid” on page 20](#).
- A tool like WinSCP to copy files to and from the Admin Node.
- Ensure you have:
 - Passwords.txt file
 - Provisioning passphrase
 - Provisioning media

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. If there is no service pack, mount the StorageGRID 9.0 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 CD image. Enter:

```
mount -o loop,ro /var/local/install/  
Bycast_StorageGRID_9.0.<servicepacknumber>_Software_  
Service_Pack_<buildnumber>.iso /cdrom
```
4. Create the /root/usb directory for provisioning data if it does not already exist. Enter: `mkdir -p /root/usb`
5. Copy the contents of the provisioning media to the /root/usb directory. For example, use WinSCP to copy the files from your service laptop.
6. Load the provisioning software. Enter (on one line):

```
/cdrom/load-provisioning-software --alternate-usb-dir=/\  
root/usb
```

NOTE The script outputs “Converting spec file for use with software version: 9.0”. The version number in the message does not change when the service pack number changes.

7. When prompted, enter the provisioning passphrase.
If provisioning fails, see [“Provisioning Failures” on page 78](#).
8. Go to **CMN ► Grid Task ► Configuration ► Main** and confirm that the Software Upgrade grid task is displayed in the Pending table.
9. Back up the provisioning data:
 - 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.
 - d. Log out of the command shell. Enter: `exit`
10. Optionally, you can archive firewall settings to protect against possible loss of custom settings. The upgrade process may overwrite custom firewall settings with default settings.
 - a. Setup passwordless access from the Admin Node to the other servers in the grid. Enter: `ssh-add`
When prompted, enter the ssh-access password from the Passwords.txt file.
 - b. Copy the Provisioned grid specification file. Enter:
`copy-grid-spec <directory>`
where `<directory>` is the directory where a copy of the Provisioned grid specification file is to be saved. This copy of the Provisioned grid specification file will be used in the next step to determine the location of servers from which firewall settings must be obtained.
 - c. Create an archive of firewall settings for each server in the grid. Enter:
`archive-fw-config --target-dir=<target_directory> --grid-spec=<directory>/<grid_spec_file_name>`
where `<target_directory>` is the directory to which archive files of firewall settings will be saved, and `<directory>/<grid_spec_file_name>` is the location and file name of the Provisioned grid specification file copied in step **10 b**.
Firewall settings for each server in the grid are archived to a `<target_dir>/<hostname>.SuSEfirewall2` file. One file for each

server in the grid. Firewall settings can be reviewed later and updated if necessary at the end of the upgrade process. For more information, see [“After Grid Software Upgraded on all Grid Nodes”](#) on page 42.

11. Store the Provisioning directory (/root/usb) and the Backup Provisioning directories (/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, secure physical locations.

For more advice on where to store provisioning data, see “Preserving Copies of Provisioning Data” in the *Installation Guide*.



WARNING Store two copies of the Provisioning directory separately in safe 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.

The next step is to convert all HAGCs to unicast from broadcast heartbeat before the Software Upgrade grid task is run. For more information on upgrading Gateway Nodes, see [Chapter 4: “Gateway Node Upgrades”](#).

After converting all HAGCs to unicast heartbeat, the next step is to run the Software Upgrade grid task. See [“Start the Software Upgrade Grid Task”](#) below.

NOTE If you store a copy of the provisioning data 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 the grid in the event of a catastrophic failure. If the only copy of these keys are in the grid, it may not be possible to recover data.

Admin Node is on a Physical Server

Prerequisites

- Grid is being upgraded from release 8.5.2 or later
- ISO images of the software CDs have been loaded onto the primary Admin Node server with load_cds.py. See [“Load Software Distribution on the Primary Admin Node”](#) on page 28.

- The condition of the grid has been reviewed. See [“Check the Condition of the Grid” on page 20](#).
- Ensure you have:
 - Provisioning USB flash drive
 - Backup Provisioning USB flash drive
 - Passwords.txt file
 - 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. If there is no service pack, mount the StorageGRID 9.0 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 CD 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
```

NOTE The script outputs “Converting spec file for use with software version: 9.0”. The version number in the message does not change when the service pack number changes.

5. When prompted, enter the provisioning passphrase.
6. When prompted, insert the Provisioning USB flash drive.
7. When provisioning is complete, remove the Provisioning USB flash drive.

If provisioning fails, see [“Provisioning Failures” on page 78](#).

8. Go to **CMN ► Grid Task ► Configuration ► Main** and confirm that the Software Upgrade grid task is displayed in the Pending table.
9. 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.
 - d. Remove the Backup Provisioning USB flash drive.
 - e. Log out of the command shell. Enter: **exit**

10. Optionally, you can archive firewall settings to protect against possible loss of custom settings. The upgrade process may overwrite custom firewall settings with default settings.
 - a. Setup passwordless access from the Admin Node to the other servers in the grid. Enter: `ssh-add`
When prompted, enter the ssh-access password from the `Passwords.txt` file.
 - b. Copy the provisioned grid specification file. Enter:
`copy-grid-spec <directory>`
where `<directory>` is the directory where a copy of the Provisioned grid specification file is to be saved. This copy of the Provisioned grid specification file will be used in the next step to determine the location of servers from which firewall settings must be obtained.
 - c. Create an archive of firewall settings for each server in the grid. Enter:
`archive-fw-config --target-dir=<target_directory> --grid-spec=<directory>/<grid_spec_file_name>`
where `<target_directory>` is the directory to which archive files of firewall settings will be saved, and `<directory>/<grid_spec_file_name>` is the location and file name of the Provisioned grid specification file copied in step 10 b.
Firewall settings for each server in the grid are archived to a `<target_dir>/<hostname>.SuSEfirewall2` file. One file for each server in the grid. Firewall settings can be reviewed later and updated if necessary at the end of the upgrade process. For more information, see [“After Grid Software Upgraded on all Grid Nodes”](#) on page 42.
11. Store both USB flash drives separately in safe locations.



WARNING Store the Provisioning USB flash drive and the Backup Provisioning USB flash drive separately in safe 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. The Provisioning USB is also required to recover from a primary Admin Node failure.

The next step is to convert all HAGCs to unicast from broadcast heartbeat before the Software Upgrade grid task is run. For more information on upgrading Gateway Nodes, see [Chapter 4: “Gateway Node Upgrades”](#).

After converting all HAGCs to unicast heartbeat, the next step is to run the Software Upgrade grid task. See [“Start the Software Upgrade Grid Task”](#) below.

Start the Software Upgrade Grid Task

For detailed information on how to work with grid tasks, see the “Grid Tasks” chapter of the *Administrator Guide*.

The software upgrade starts with the Software Upgrade grid task.

NOTE HAGCs must be converted from broadcast to unicast heartbeat before the Software Upgrade grid task is run. For more information on upgrading Gateway Nodes, see [Chapter 4: “Gateway Node Upgrades”](#).

Prerequisites

- The grid has been provisioned. See [“Provision the Grid”](#) on page 31.
- The condition of the grid has been reviewed. See [“Check the Condition of the Grid”](#) on page 20.
- All HAGCs have been converted to unicast heartbeat
- No grid tasks are running. The only exceptions are:
 - LDR Content Rebalancing grid task (LBAL)
 - ILM Evaluation (ILME)

These grid tasks can run concurrently with the Software Upgrade (SWUP) grid task. If any other grid tasks are running, wait for them to complete or release their lock, or abort them as appropriate.

- Grid Expansion (GEXP)

The Grid Expansion (GEXP) grid task can be started and run concurrently with the Software Upgrade (SWUP) grid task after the primary Admin Node and all Control Nodes are upgraded, but only for Gateway Nodes and Storage Nodes. See [Chapter 5: “Expansion During Upgrade”](#).

For details on how to work with grid tasks, see the “Grid Tasks” chapter of the *Administrator Guide*.

Procedure

1. Log in to the NMS MI using the Vendor account.
2. Go to **<Grid_Root> ► Configuration ► Tasks**.
3. In the Pending table, locate the grid task Software Upgrade. Under Actions, select **Start**.

NOTE There could be more than one pending grid task. Make sure that you select the correct one.

4. Click **Apply Changes**.

The grid task moves to the Active table.

NOTE You must wait for the page to auto-refresh before the change is visible. Do not click **Apply Changes** again.

5. Wait until the Software Upgrade grid task moves to the stage Wait for Software Version. This may take up to 20 minutes in a large or busy grid.

NOTE Do not wait for the grid task to complete. Wait only until the grid task moves to the stage “Wait for Software Version”.

The next step is to upgrade the software on the primary Admin Node. See “[Upgrade the Grid Software with GDU](#)” below.

Upgrade the Grid Software with GDU

Upgrade the primary Admin Node or the HCAC that includes the primary reporting Admin Node first. If the Admin Node hosts a secondary FSG, the secondary FSG is unavailable until this procedure is complete. If you are upgrading an HCAC, upgrade its primary reporting Admin Node and then its processing Admin Node before any other grid node. The procedure to upgrade the grid software uses GDU (Grid Deployment Utility).

After the primary Admin Node or both grid nodes in an HCAC are upgraded, you can start the NMS MI and use it to monitor both old and upgraded servers.

Once you have upgraded the primary Admin Node or both grid nodes in an HCAC, the order in which you upgrade the servers within a site does not matter except for operational Gateway Nodes or when an expansion of Gateway Nodes or Storage Nodes is planned. Before an expansion of Gateway Nodes or Storage Nodes can occur, all Control Nodes must be upgraded. Detailed procedures to upgrade operational Gateway Nodes are covered in [Chapter 4: “Gateway Node Upgrades”](#). For more information on expansion during upgrade, see [Chapter 5: “Expansion During Upgrade”](#).

You can upgrade the software on servers hosting different types of grid nodes in parallel, but do not upgrade all servers of the same type simultaneously. For example, do not upgrade all Storage Nodes at the same time: you must ensure that enough Storage Nodes remain available to handle ingest and retrieval. Similarly, do not upgrade all Control Nodes at once: where possible, ensure that at least two Control Nodes are running at all times.

If the TSM middleware is co-hosted on the Archive Node server, you may wish to shut down the TSM before the update. Depending on configuration, you may need to restart the TSM manually after the update.

Once you have upgraded the primary Admin Node or both grid nodes in an HCAC and all Control Nodes, you can expand the grid by adding Gateway Node replication groups and Storage Nodes. This grid expansion is permitted before the upgrade of all grid nodes completes. Grid node types other than Gateway Nodes and Storage Nodes can only be added after the upgrade process completes. When adding Gateway Nodes and Storage Nodes, you must install the software version to which the grid is being upgraded.

Gateway Nodes and Storage Nodes not currently connected to the grid can be upgraded at any time after the primary Admin Node or HCAC and all Control Nodes are upgraded. The upgrade of an unconnected Gateway Node or Storage Node can be performed either manually or via GDU. To upgrade via GDU, network connections between the primary Admin Node and the Gateway Node or Storage Node must be established. When this occurs, the Gateway Node or Storage Node becomes available in the Servers panel of GDU. Note that even when manually upgrading a Gateway Node or Storage Node, the final step of the upgrade process (starting grid services) is usually performed via GDU. If upgrading manually, see [Chapter 6: “Manually Upgrading”](#).

The Software Upgrade (SWUP) grid task remains active until all grid nodes are upgraded and connected to the grid (network connections are established and grid services are started).

Prerequisites

- The software CDs have been copied to the primary Admin Node. See [“Load Software Distribution on the Primary Admin Node”](#) on page 28.
- If the server is not at the same site as the primary Admin Node, the StorageGRID Software CD and the Enablement Layer for StorageGRID Software CD have been copied to the remote site in order to reduce WAN traffic. See [“Copy ISO Files in Multi-Site Environment”](#) on page 89.
- The grid has been provisioned. See [“Provision the Grid”](#) on page 31.
- The condition of the grid has been reviewed. See [“Check the Condition of the Grid”](#) on page 20.
- All HAGCs have been converted to unicast heartbeat
- There is network connectivity between the primary Admin Node and the server you are upgrading.

NOTE Network connectivity requirements are not necessarily applicable to Gateway Nodes and Storage Nodes not currently connected to the grid. Unconnected Gateway Nodes and Storage Nodes can be upgraded manually without being connected to the grid.

- The ssh key is installed: `/root/.ssh/id_rsa` is on the server
If the ssh key has been removed, reinstall it. For more information and the procedure, see the *Administrator Guide*.
- The SSH Access Password is available (consult the `Passwords.txt` file)

NOTE If a grid node fails while an upgrade of the grid is in progress, contact Support.

Procedure

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

1. Start GDU. See [“How to Use GDU” on page 81](#).
2. Upgrade the StorageGRID software on each server:
 - First, upgrade the primary Admin Node or both Admin Nodes in an HCAC.
 - If you are upgrading the HCAC that includes the primary reporting Admin Node, upgrade its primary reporting Admin Node and then its processing Admin Node.
 - After upgrading the primary Admin Node or the HCAC that includes the primary reporting Admin Node, the order in which you upgrade the servers within a site does not matter except for operational Gateway Nodes or when an expansion of Gateway Nodes or Storage Nodes is planned (see [Chapter 4: “Gateway Node Upgrades”](#) and [Chapter 5: “Expansion During Upgrade”](#)), and provided that at least one server of each type is running while you upgrade the other servers of that type.



WARNING To ensure that the grid remains operational, do not upgrade all servers of the same type simultaneously.

- The following procedures also apply to Gateway Nodes and Storage Nodes not currently connected to the grid and that are not upgraded manually. Follow these procedures when the Gateway Node or Storage Node is connected to the grid.

When an unconnected Gateway Node or Storage Node is connected to the grid it becomes available in GDU for upgrade.

To manually upgrade unconnected Gateway Nodes and Storage Nodes, see [Chapter 6: “Manually Upgrading”](#).

- a. In the **Servers** panel, select the server and confirm that the server state is Available.
- b. In the **Tasks** panel, select **Upgrade Software**, and then in the **Actions** panel, select **Start Task** and press **<Enter>**.
 - If you are upgrading the HCAC that includes the primary reporting Admin Node, perform step **2 b** for the primary reporting Admin Node first and then the processing Admin Node before continuing to step **2 c**.

Wait until the task completes.

The server may reboot automatically if it hosts an FSG.

The Upgrade Software task upgrades the software to the base version 9.0 and, if applicable, applies the service pack 9.0.x.

- c. Log in to the NMS MI and verify that the services hosted on that server have rejoined the grid, that is, the state is not Unknown (shown in blue) or Administratively Down (shown in gray).
3. After you have finished using the GDU console, quit GDU and remove the ability to access servers without a server password. See [“Close GDU”](#) on page 86.

After Grid Software Upgraded on all Grid Nodes

After all unconnected Gateway Nodes and Storage Node are connected (network connections established and grid services started) to the grid and the software has been upgraded on all servers or virtual machines, complete post-upgrade configuration steps. For more information, see [Chapter 7: “Complete the Upgrade”](#).

NOTE The software upgrade is not complete until all unconnected Gateway Nodes and Storage Nodes are upgraded and connected to the grid.

After completing the upgrade of all grid nodes to StorageGRID 9.0.x, it is recommended that you upgrade the operating system from SLES 10 SP3 (64-bit) to SLES 11 SP2 (64-bit). Before upgrading the operating system to SLES 11 SP2 (64-bit) all grid nodes must be upgraded to StorageGRID 9.0.x. Failure to complete the upgrade of all grid nodes to StorageGRID 9.0.x before upgrading the SLES operating system may result in the failure of your StorageGRID system.

Note that you cannot upgrade StorageGRID software and the underlining operating system in parallel on different grid nodes. The upgrade to StorageGRID 9.0.x must finish in its entirety before upgrading the SLES operating system. See [Chapter 9: “Upgrade the Operating System”](#).



WARNING Do not upgrade to SLES 11 SP2 (64-bit) until all grid nodes are upgraded to StorageGRID 9.0.x.

Gateway Node Upgrades

How to upgrade the StorageGRID software on connected Gateway Nodes to minimize service interruptions

Introduction

This chapter contains information on how to upgrade the StorageGRID software on operational Gateway Nodes. The procedure depends on the type of replication group. Except for the procedure to convert HAGCs to unicast heartbeat, this chapter is not applicable to Gateway Nodes not connected to the grid or Gateway Nodes added during upgrade. Non-connected Gateway Nodes do not follow these procedures when the replication group is connected and software is upgraded with GDU. Because non-connected Gateway Nodes are not yet operational, they can be upgraded in any order.



WARNING Do not shut down the next FSG in a replication group until the previous FSG has been upgraded and is back online, the replication status is normal, and replication queues are clear.

Upgrade Basic Gateway Replication Groups

In a grid that supports business continuity failover, that is, if the grid can be configured to fail over to a writable secondary, follow these steps to minimize service interruptions while the operational Gateway Nodes are upgraded.

1. Upgrade the software on the secondary Gateway Node. See [“Upgrade the Grid Software with GDU” on page 38](#).

2. Fail over from the primary Gateway Node to the secondary Gateway Node. See [“Fail Over From Primary to Secondary”](#) on page 44.

The secondary Gateway Node becomes the acting primary Gateway Node.

The original primary Gateway Node becomes a secondary Gateway Node.

3. Upgrade the software on the original primary Gateway Node that is now a secondary Gateway Node. See [“Upgrade the Grid Software with GDU”](#) on page 38.
4. Fail back to the original primary Gateway Node. See [“Fail Back to Original Primary”](#) on page 46.
5. Upgrade any other Gateway Nodes in the replication group.

Fail Over From Primary to Secondary

NOTE There is a brief interruption of service to clients during the failover.

1. In the NMS MI, verify that the secondary FSG is operating normally:
 - a. Go to **<Configured_Secondary> ► FSG ► Overview ► Main** and verify that FSG State is Online and FSG Status is No Errors.
 - b. Go to **<Configured_Secondary> ► FSG ► Replication ► Main** and verify that Current Role is Active Secondary and Replication Status is Normal.
2. Stop client ingest.

If you fail over without stopping client ingest, any client operations that are in progress when the failover occurs will be interrupted.
3. Verify that all pending content has been stored to the grid:
 - a. Go to **<Configured_Primary> ► FSG ► Storage ► Overview ► Main**.
 - b. Wait until the value of Files Stored to Grid – Pending decreases to zero.

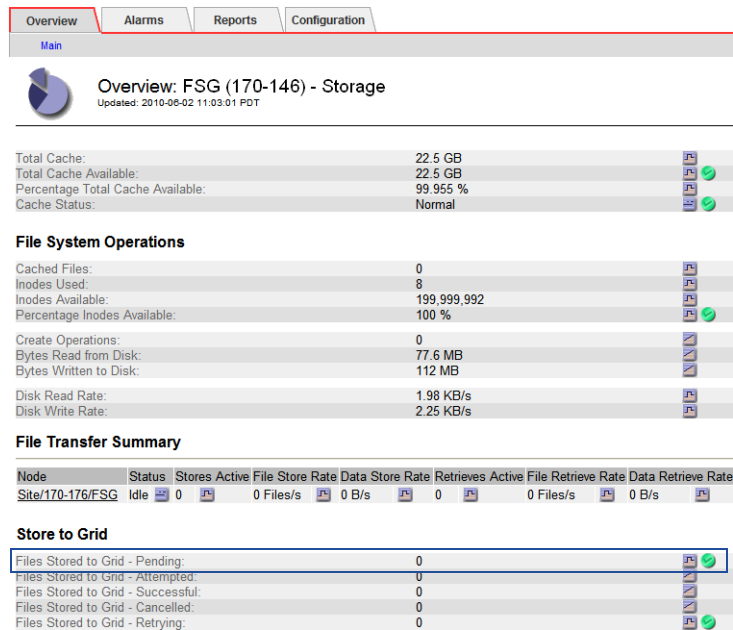


Figure 5: Files Stored to Grid – Pending

Files Stored to Grid – Pending is greater than zero when files are cached on the FSG faster than they can be written to a Storage Node. If the value is not zero, ensure that all Storage Nodes and Control Nodes are operating normally. You may need to temporarily stop ingests to the grid before you can proceed with the upgrade.

The value may not decrease to zero if you have not stopped client ingests. If the FSG is stopped with a pending queue, the objects will not be ingested until the FSG is restarted and therefore will not be available on the failover FSG.

4. Confirm that a backup is not in progress. See “Disable, Reschedule, or Cancel FSG Backups” on page 22.
5. Fail over:
 - a. Go to **Grid Management ► FSG Management ► <Replication_Group> ► Configuration ► Settings**.
 - b. Select **Primary FSG ► <current_backup_FSG>**
where **<current_backup_FSG>** is the FSG selected for Backup FSG.

Description	Settings
Maximum Active Swapins	20
Maximum Active Ingests	20
Minimum Active Ingests	15
Ingest Idle Delay	10.0 second(s)
Backup Window Start	8.0 hour(s)
Backup Window Length	1.0 hour(s)
Primary FSG	Site/aurora-49/FSG
Backup FSG	Site/aurora-49/FSG
Swapout Free Space Watermark	25000000000
Swapout No Create Watermark	10000000000
Object Transfer Encryption	AES256-SHA

Apply Changes

Figure 6: Configuring the Backup FSG to be Active Primary

c. Click **Apply Changes**.

6. Go to **<Configured_Secondary> ► FSG ► Replication ► Overview ► Main** and verify that Current Role is Active Primary
7. Go to **<Configured_Primary> ► FSG ► Replication ► Overview ► Main** and verify that Current Role is Active Secondary and that the Active Session ID value is the same as for the active primary.
8. Redirect client applications if applicable and resume client ingest.
The details of this redirection, which includes configuring IP addresses on the client, depend upon the nature of the client application and are not covered in this guide.

Fail Back to Original Primary

NOTE There is a brief interruption of service to clients during the failover.

1. In the NMS MI, verify that the original primary FSG is operating normally:
 - a. Go to **<Configured_Primary> ► FSG ► Overview ► Main** and verify that FSG State is Online and FSG Status is No Errors.

- b. Go to **<Configured_Primary> ► FSG ► Replication ► Main** and verify that Current Role is Active Secondary and Replication Status is Normal.
2. If the grid was ingesting content while the primary was being upgraded:
 - a. Stop client ingest.
 - b. Verify that all pending content has been stored to the grid. Go to **<Configured_Secondary> ► FSG ► Storage ► Overview ► Main** and wait until the value of Files Stored to Grid - Pending decreases to zero.
3. Fail over:
 - a. Go to **Grid Management ► FSG Management ► <Replication Group> ► Configuration ► Settings**.
 - b. Change the **Primary FSG** back to the original primary FSG.
 - c. Click **Apply Changes**.
4. Verify that the FSG services are operating in their original roles by reviewing the FSG ► Replication component of each:
 - The Current Role attribute reports the role the service is performing (Primary or Secondary).
 - The Active Session ID numbers are the same for the primary and the secondary.
5. Redirect client applications if applicable and resume client ingest.

Upgrade High Availability Gateway Replication Groups

Follow these steps to minimize service interruptions while operational Gateway Nodes in High Availability Gateway (HAGC) replication groups are upgraded.

1. Convert heartbeat to unicast. See [“Convert Heartbeat for High Availability Gateway Cluster”](#) on page 48.
2. Start the Software Upgrade grid task. See [“Start the Software Upgrade Grid Task”](#) on page 37.
3. Upgrade the software on the secondary Gateway Node. See [“Upgrade the Grid Software with GDU”](#) on page 38.
4. Upgrade the software on the FSG whose Current Role is Standby Primary. See [“Upgrade the Grid Software with GDU”](#) on page 38.

5. Force the cluster to fail over. See [“Fail Over Within the HAGC”](#) below.
6. Upgrade the software on the other FSG, that is, the one whose Current Role is now Standby Primary. See [“Upgrade the Grid Software with GDU”](#) on page 38.

Convert Heartbeat for High Availability Gateway Cluster

HAGCs originally installed before StorageGRID 9.0 use broadcast rather than unicast heartbeat. Before starting the Software Upgrade grid task, convert all HAGCs to unicast from broadcast heartbeat.

When virtual machines host HAGCs, one or more HAGCs can share a dedicated unicast heartbeat network for intra-cluster communications and a crossover cable is no longer needed to connect the virtual machine hosts. Note, however, that HAGCs hosted on physical servers still require crossover cables for heartbeat to maintain the highest levels of reliability.

Gateway Node is Hosted by a Virtual Machine

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. Extract updated heartbeat configuration files (ha.cf) from the primary Admin Node and save then to a directory on the server.

Enter: **get-ha-config /root/usb**

A sub-directory is created within the /root/usb directory for each server hosting a Gateway Node where the network interface is setup to use heartbeat and the correct ha.cf file is copied to this directory.

3. Copy the correct ha.cf from the Admin Node to the Gateway Node's /etc/ha.d/ directory. Enter (on one line):

```
scp /root/usb/<Gateway Node>/etc/hac.d/ha.cf\  
<Gateway_Node_IP_address>:/etc/ha.d/
```

NOTE For unconnected Gateway Nodes you may need to connect a keyboard and monitor to the server and manually copy files.

4. Repeat step 3, for each Gateway Node.
5. Log out of the primary Admin Node. Enter: **exit**
6. Proceed to the next step in the upgrade or update process.

Gateway Node is on a Physical Server

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. Extract updated heartbeat configuration files (`ha.cf`) from the primary Admin Node and copy them to a USB flash drive. Enter:
get-ha-config
You are prompted to insert a USB flash drive. A directory is created for each server hosting a Gateway Node where the network interface is setup to using heartbeat and the correct `ha.cf` file is copied to this directory.
3. Log out of the primary Admin Node. Enter: **exit**
4. At the server hosting a Gateway Node where the network interface is setup to use heartbeat, access a command shell and log in as root using the password listed in the `Passwords.txt` file.

NOTE For unconnected Gateway Nodes you may need to connect a keyboard and monitor to the server.

5. Insert the USB flash drive containing `ha.cf` files obtained in step 2 and copy the correct `ha.cf` to `/etc/ha.d/`.
Each `ha.cf` file is stored in a subdirectory named for the server onto which the `ha.cf` file must be copied.
6. Remove the USB flash drive and log out of the Gateway Node. Enter: **exit**
7. Repeat steps 4 to 6, for each Gateway Node.
8. Proceed to the next step in the upgrade or update process.

Fail Over Within the HAGC

NOTE There is a brief interruption of service to clients during the failover.

1. In the NMS MI, verify that the Active Primary is operating normally:
 - a. Go to **<Active_Primary> ► FSG ► Replication ► Overview ► Main**.
 - b. Verify that Current Role is Active Primary.
 - c. Verify that Cluster Status is Normal.
The status is Normal when the main FSG is the active primary and the supplementary FSG is the standby primary.

Changing the status of the active primary in a cluster with a status of Transitional or Vulnerable may render gateway services unavailable to the grid.

d. Verify that Replication Status is Normal.

2. Stop client ingest.

If you execute the failover without stopping clients, any client operations that are in progress when the failover occurs will be interrupted.

3. Verify that all pending content has been stored to the grid. Go to **<Active_Primary> ► FSG ► Storage ► Overview ► Main** and wait until the value of Files Stored to Grid – Pending decreases to zero.

This value may not decrease to zero if you have not stopped client ingests. If the FSG is stopped with a pending queue, the objects will not be ingested until the FSG is restarted and therefore will not be available on the failover FSG.

4. Stop client services on the active primary:

a. Go to **<Active_Primary> ► FSG ► Client Services ► Configuration ► Main**.

b. Change the Client Services State to **Stopped**.

c. Click **Apply Changes**.

d. Go to **FSG ► Client Services ► Overview ► Main** and monitor the status of Client Services. Wait until all services are Stopped.

e. Go to **FSG ► Replication ► Overview ► Main** and verify that Current Role has changed to Standby Primary and Cluster Status to Vulnerable.

If the cluster status is Transitional, monitor the cluster for a few minutes until the status changes.

5. Resume client services on the FSG that is now the standby primary:

a. Go to **<Standby_Primary> ► FSG ► Client Services ► Configuration ► Main**.

b. Change the Client Services State to **Running**.

c. Click **Apply Changes**.

6. Remount file shares for NFS clients if required (that is, if client operations were interrupted).

7. Resume client ingest.

Expansion During Upgrade

Adding Gateway Node replication groups and Storage Nodes during an upgrade

Scope and Limitations

This chapter provides an overview of how to add a Gateway Node replication group or Storage Node while the grid is undergoing a software upgrade. Expansion can take place any time after the **Software Upgrade (SWUP)** grid task has started, and the primary Admin Node and all Control Nodes have been upgraded. Details of the expansion process are described at the referenced locations. Once expansion begins, it follows the normal expansion process as documented in the *Expansion Guide*.

When adding Gateway Nodes and Storage Nodes during an upgrade, the expansion Gateway Nodes and Storage Nodes must be installed with the software release version to which the grid is being upgraded.

Read through this section before beginning an expansion during an upgrade for important information required to successfully expand your grid during an upgrade.

Service Packs

Expansion can also occur during the application of a service pack. For more information on the application of service packs, see [Chapter 8 “Apply Service Packs”](#).

Gateway Nodes

Only new Gateway Node replication groups can be added to the grid during an upgrade. This includes a High Availability Gateway Cluster (HAGC) replication group. An existing replication group cannot be

expanded during an upgrade. Note that HAGCs must be converted to unicast heartbeat before expansion begins. For more information, see [“Convert Heartbeat for High Availability Gateway Cluster”](#) on page 48.

When a Gateway Node replication group is added to the grid during an upgrade, configuration of the Gateway Nodes can be performed when the addition of the replication group to the grid is complete. The upgrade runs independent of the expansion and does not have to complete before configuration takes place. However, features new to the installed software version cannot be enabled until the upgrade completes.

The expansion Gateway Node replication group can begin ingesting objects before the upgrade completes.

An existing Gateway Node replication group cannot be converted to an HAGC during an upgrade.

Storage Nodes

When Storage Nodes are added to the grid during an upgrade, configuration of the Storage Node including ILM rules can be performed when the expansion of the Storage Node is complete. The upgrade runs independent of the expansion and does not have to complete before configuration takes place. However, features new to the installed software version cannot be enabled until the upgrade completes.

The expansion Storage Node can begin storing objects before the upgrade completes.

Existing Storage Nodes cannot be expanded during an upgrade.

Adding Grid Nodes During an Upgrade

The following outlines when and how to perform an expansion while a grid is being upgraded.

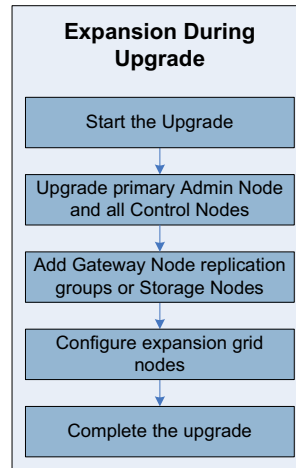


Figure 7: Expansion During Upgrade

Table 6: Expansion During Upgrade

✓	Step	Action	See
	1.	<p>Start the upgrade.</p> <p>The upgrade from StorageGRID 8.5.2 or later software to StorageGRID 9.0.x begins as per normal. The software distribution is loaded onto the primary Admin Node. The grid is provisioned. HAGCs are converted to unicast heartbeat. The Software Upgrade grid task is started. The primary Admin Node is upgraded, and upgrade is started on the grid's other grid nodes.</p> <p>All Control Nodes must be upgraded before starting an expansion.</p>	Chapter 3: "Upgrade Software"

Table 6: Expansion During Upgrade (cont.)

✓	Step	Action	See
	2.	<p>Add Gateway Nodes or Storage Nodes.</p> <hr/> <p>NOTE Before starting expansion, confirm that the grid has the bindings capacity necessary to support the additional grid nodes.</p> <hr/> <p>Once expansion begins, it follows the normal expansion process as documented in the <i>Expansion Guide</i>.</p> <p>Begin the expansion process by updating the provisioned grid specification file. Add the expansion grid nodes to the grid specification file and then reprovision the grid.</p> <p>Next, update networking, load Linux, prepare hardware, and install StorageGRID software on the new server.</p> <p>Then, add, customize, and start the expansion Gateway Nodes and Storage Nodes.</p> <hr/> <p>NOTE The version of the StorageGRID software installed on the expansion server must be that to which the grid is being upgraded.</p> <hr/>	“Expanding the Grid” chapter of the <i>Expansion Guide</i> .
	3.	<p>Make configuration changes to expansion Gateway Node replication group and Storage Nodes.</p> <p>All aspects of Gateway Node and Storage Node configuration can be performed on the expansion grid node before upgrade completes. For example, configuring FSG profiles for Gateway Nodes and ILM rules for Storage Nodes.</p> <hr/> <p>NOTE Features new to the installed software version cannot be enabled until the upgrade completes.</p> <hr/>	“Gateways (FSG and CLB)”, “Disk Storage (LDR)”, and “Information Lifecycle Management (ILM)” chapters of the <i>Administrator Guide</i> .
	4.	Complete the upgrade.	Chapter 7: “Complete the Upgrade”

Manually Upgrading

Manually upgrade a server including Gateway Nodes and Storage Nodes not connected to the grid

Introduction

The preferred method to upgrade grid software is with GDU. However, there are scenarios when the manual process may be required. For instance, if there is no connectivity between the server and the primary Admin Node or when upgrading unconnected Gateway Nodes and Storage Nodes.

Upgrading the Grid Software Manually

The following procedure can be used when either troubleshooting the upgrade process or upgrading Gateway Nodes and Storage Nodes not connected to the grid. Depending on the grid node being upgraded, procedures differ slightly. Follow the procedure exactly as presented below.

NOTE The following procedure does not use ISO images.

Prerequisites and required materials

- Ensure you have:
 - StorageGRID 9.0 Software
 - Enablement Layer for StorageGRID 9.0 Software
 - If available, StorageGRID 9.0.x Software Service Pack
 - If available, Enablement Layer for StorageGRID 9.0.x Software Service Pack
- The grid software has been upgraded on the primary Admin Node or the HCAC that includes the primary reporting Admin Node server. See [“Upgrade the Grid Software with GDU”](#) on page 38.

If an expansion during upgrade of Gateway Nodes or Storage Nodes is planned, all Control Nodes must be upgraded before beginning the expansion process.

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

- If you are upgrading a High Capacity Gateway Cluster (HAGC), a USB flash drive to copy updated heartbeat configuration files

NOTE If a grid node fails while an upgrade of the grid is in progress, contact Support.

Procedure

1. For all servers hosting Gateway Nodes where the network interface is setup to use heartbeat (for HAGC or otherwise), extract updated heartbeat configuration files (ha.cf) from the primary Admin Node and copy them to a USB flash drive:
 - a. At the primary Admin Node, access a command shell and log in as root using the password listed in the `Passwords.txt` file.
 - b. Enter: **get-ha-config**

You are prompted to insert a USB flash drive. A directory is created for each server hosting Gateway Nodes where the network interface is setup to using heartbeat and the correct ha.cf file is copied to this directory.
 - c. Log out of the primary Admin Node. Enter: **exit**
2. At the server to be upgraded, access a command shell and log in as root using the password listed in the `Passwords.txt` file.

NOTE For unconnected Gateway Nodes and Storage Nodes you may need to connect a keyboard and monitor to the server.

3. For all servers hosting Gateway Nodes where the network interface is setup to using heartbeat (for HAGC or otherwise), insert the USB flash drive containing ha.cf files obtained in step 1 and copy the correct ha.cf to `/etc/ha.d/`.

Each ha.cf file is stored in a subdirectory named for the server onto which the ha.cf file must be copied.
4. Remove any USB flash drive from the server.

The order in which you load the CDs does not matter.

5. Load software distribution onto the grid Node:
 - a. Insert the StorageGRID 9.0 Software CD.
 - b. Load the ISO image. Enter: `load_cds.py`
 - c. When prompted, insert the next CD, type **y**, and press **<Enter>**.
 - d. Repeat step 5 c for all CDs.
 - e. To exit, type **n** and press **<Enter>** when prompted.
6. Mount the StorageGRID 9.0 Software CD image. Enter:


```
mount -o loop,ro /var/local/install/Bycast_StorageGRID_\
9.0_Software_<buildnumber>.iso /cdrom
```
7. Run the update script to upgrade to the base version. Enter:


```
/cdrom/swupdate/updategrid.rb --iso=/var/local/install/\
Enablement_Layer_for_StorageGRID_9.0_Software_\
<buildnumber>.iso
```
8. If there is a service pack available, apply the service pack:
 - a. Mount the StorageGRID 9.0.x Software Service Pack CD image.
Enter:


```
mount -o loop,ro /var/local/install/Bycast_StorageGRID_\
9.0.<servicepacknumber>_Software_Service_Pack_\
<buildnumber>.iso /cdrom
```
 - b. Run the update script again to apply the service pack. Enter:


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

NOTE For new unconnected grid nodes, to continue, you must establish network connections to the grid. This can be done at any time; however, upgrade cannot complete until network connections are established and the rest of this procedure completed.

9. If this is a Gateway Node or Storage Node that has never before been connected to the grid, start the grid software:
 - In GDU, select **Enable Services** for the grid node.
 - or —
 - Enter: `postinstall.rb start`

Grid services are started on the grid node.
10. If the grid node you are upgrading is connected to the grid and is not a new grid node, when the update script has completed, restart the grid services. Enter: `/etc/init.d/servermanager start`
 - or —

If prompted to reboot, enter: `reboot`

NOTE On servers hosting a Gateway Node, you may be prompted to reboot if the Linux kernel has been upgraded. Grid services will start automatically when the server reboots. In the NMS, the grid services will turn blue (Unknown state) as you reboot the server but eventually will become green (Connected state).

11. Log out. Enter: **exit**
12. If the grid node you are updating is connected to the grid, confirm in the NMS MI that the software has been upgraded:
 - a. Go to **SSM ► Services ► Overview ► Main**.
 - b. Confirm that the value for storage-grid-release has been updated.

Complete the Upgrade

Final Tasks

Confirm Upgrade and Make Configuration Changes

Check the NMS MI to confirm the upgrade completed successfully and make any required configuration changes.

1. Log in to the NMS MI.
2. Confirm that the Software Upgrade grid task completed successfully:
 - a. Go to **<Grid_Root> ► Configuration ► Grid Tasks**.
 - b. Confirm that the Software Upgrade grid task is in the Historical table with a Status of Successful.

The Software Upgrade grid task will not move to the Historical table with a Status of Successful until all unconnected Gateway Nodes and Storage Nodes have been upgraded and operational on the grid.

If you do not see the grid task or it does not have the status Successful after all grid nodes have been upgraded and all unconnected Gateway Nodes and Storage Nodes have been upgraded and connected to the grid, contact Support.

If an expansion occurs during upgrade, the expansion does not have to complete before the upgrade. Upgrade can complete before expansion.

3. Verify that the upgraded services have joined the grid. For each server:
 - a. Check that the services recover from the Unknown state and that they are displayed in green.
 - b. Go to **SSM ► Services ► Main** and check that the Packages section reports the storage-grid-release version 9.0.x on all servers.

- c. Go to **SSM ► Overview ► Main** and verify that each service reports the version number noted in the *Release Notes*.
4. Check that the services are operating normally and that there are no alarms on the grid. Diagnose and address any issues.
5. If you disabled or rescheduled FSG backups, change back to the original setting.
6. Review all custom alarms to verify that they will work as expected. The software may have changed such that these alarms are no longer required or have to be modified.
7. Because the upgrade process may overwrite custom firewall settings with default settings, if you archived firewall settings during the upgrade process, it is recommended that you compare these archived settings with current firewall settings to determine if any have been reset. For more information and the procedure to archive firewall settings, see [“Provision the Grid” on page 31](#).

If created earlier, an archived version of the firewall settings for each server is available in a `<target-dir>/<hostname>.SuSEfirewall2` file. (The location of these files was determined at creation; see step 10 of the procedure [“Provision the Grid” on page 31](#).)

Compare these archive files with the current `SuSEfirewall2` file on each server. The grid stores the current version of the firewall settings on each server in the `/etc/sysconfig/SuSEfirewall2` file.

If necessary, use the `open-port.rb` and `close-port.rb` scripts to recover custom firewall settings. For example, to open port 1024, enter:

```
open-port.rb -tcp=1024
```

The `close-port.rb` script only closes ports listed as open in the `/etc/sysconfig/SuSEfirewall2` file. See the line:

```
FW_SERVICES_EXT_TCP="ssh snmp snmptrap <open_port_number>"
```

Ports opened as a range (for example, `open-port.rb -tcp=100:125`) must be closed as a range. Individual ports within the range cannot be closed. Note that if a request to close a port is not successful, no feedback is given.

Enable Object Location Indexing

After your upgrade to StorageGRID 9.0.x is complete, if your grid's Control Nodes still use metadata synchronization, you must convert your grid to use metadata replication. After the conversion to metadata replication is complete, you must then enable Object

Location Indexing for each CMS in the grid. For more information and procedures, see the *Maintenance Guide*.

NOTE If you do not convert your grid to metadata replication and enable Object Location Indexing, you will not be able to upgrade your grid beyond StorageGRID 9.0.

Upgrade Operating System

After completing the upgrade of all grid nodes to StorageGRID 9.0.x, it is recommended that you upgrade the operating system from SLES 10 SP3 (64-bit) to SLES 11 SP2 (64-bit). Before upgrading the operating system to SLES 11 SP2 (64-bit) all grid nodes must be upgraded to StorageGRID 9.0.x. Failure to complete the upgrade of all grid nodes to StorageGRID 9.0.x before upgrading the SLES operating system may result in the failure of your StorageGRID system.

Note that you cannot upgrade StorageGRID software and the underlining operating system in parallel on different grid nodes. The upgrade to StorageGRID 9.0.x must finish in its entirety before upgrading the SLES operating system. For more information, see [Chapter 9: “Upgrade the Operating System”](#).

Apply Service Packs

Introduction

This chapter describes how to update software, that is, how to apply service packs.

A service pack is the collection of fixes and enhancements since the release of the base version. A service pack cannot be installed unless the base version 9.0 is installed.

Service packs are cumulative. For example, service pack 8.5.2 includes the contents of service pack 8.5.1.

Each 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 the same for both CDs.

Follow this procedure each time you need to apply a service pack to an 9.0.x grid, for instance to go from 9.0.0 to 9.0.2.

Update Sequence

There is no prescribed order in which to apply service packs except:

- Apply the service pack to the primary Admin Node first.
- Avoid applying the service pack to all servers of the same type simultaneously.
- Follow the recommended sequence for operational Gateway Nodes described below.

- For a High Capacity Admin Cluster (HCAC), first apply the service pack to the reporting Admin Node and then apply the service pack to the processing Admin Node. Complete the application of the service pack to both servers in an HCAC before applying the service pack to another server type.

About Updating Software on Gateway Nodes

The grid can continue to ingest data throughout the update process except for a period of time when operational Gateway Nodes are unavailable:

- If the grid does not support business continuity, operational Gateway Nodes are unavailable for the duration of the Gateway Node update.

NOTE If the grid does not support business continuity failover, Gateway Nodes are unavailable while they are being updated. Plan client downtime accordingly.

- If the grid does support business continuity, operational Gateway Nodes are unavailable during the failover procedure. The length of the service interruption depends on the grid configuration and grid activity at the time of the update.

NOTE Client service is interrupted for a short period of time during the failover procedure. Plan client downtime accordingly.

See [Table 7](#) below for a summary of how to update the software on Gateway Nodes to minimize interruptions to grid operations.

Table 7: Update Sequence for Operational Gateway Nodes

Replication Group	Upgrade Sequence to Minimize Interruptions to Grid Operation
Basic Gateway – No Business Continuity Failover	<ol style="list-style-type: none"> 1. Update all Gateway Nodes in the same replication group except for the primary. 2. Update the primary.

Table 7: Update Sequence for Operational Gateway Nodes (cont.)

Replication Group	Upgrade Sequence to Minimize Interruptions to Grid Operation
Basic Gateway – Business Continuity Failover	<ol style="list-style-type: none"> 1. Update the secondary. 2. Manually fail over to the secondary. 3. Update the primary. 4. Fail back to the primary. <p>For guidance on the failover procedures, see “Fail Over From Primary to Secondary” on page 44 and “Fail Back to Original Primary” on page 46.</p>
High Availability Gateway	<ol style="list-style-type: none"> 1. Update the secondary. 2. Update the standby primary. 3. Fail the cluster over. 4. Update the server that is now the standby primary. <p>For guidance on the failover procedures, see “Fail Over Within the HAGC” on page 49.</p>

Apply the Service Pack

Prerequisites

- The base version release 9.0 or a 9.0.x service pack is installed and running on all servers
- There is network connectivity between the primary Admin Node and the server you are upgrading
- The ssh key is installed: /root/.ssh/id_rsa is on the server
- The following materials and information are available:
 - StorageGRID Software 9.0.x Service Pack CD
 - Enablement Layer for StorageGRID 9.0.x Software Service Pack CD
 - SSH Access Password (consult the 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. Copy the ISO images of the two service pack CDs to the primary Admin Node using the `load_cds.py` command.

See “Load Software Distribution on the Primary Admin Node” on page 28 for more information, including instructions for when the primary Admin Node is installed on a virtual machine (“Primary Admin Node is on a Virtual Machine” on page 29).

3. Quit GDU if it is running. Select **Quit** from the **Actions** panel, and press **<Enter>**.

This is required in case the service pack includes changes that affect GDU. GDU must not be running while the provisioning software is loaded.

To determine if GDU is running, enter: `pgrep gdu-console`

The result should not display any output.

```
an1-a-1# pgrep gdu-console
an1-a-1# _
```

If any output appears, GDU is running. You must quit GDU from the GDU console.

4. Mount the StorageGRID 9.0.x Software Service Pack CD image.
Enter:

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

5. Load the provisioning software from the service pack ISO.

If the Admin Node is hosted by a virtual machine:

- a. Create the `/root/usb` directory for provisioning data if it does not already exist. Enter: `mkdir -p /root/usb`

- b. Load the provisioning software, enter:

```
/cdrom/load-provisioning-software --alternate-usb-dir=/root/usb
```

— or —

If the Admin Node is hosted on a physical server:

- a. Load the provisioning software, enter:

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

NOTE The script outputs “Converting spec file for use with software version: 9.0”. The version number in the message does not change when the service pack number changes.

6. Start GDU. See “How to Use GDU” on page 81.

If you are updating the software on the Archive Node and the TSM middle-ware is co-hosted, you may wish to shut down the TSM before the update. Depending on configuration, you may need to restart the TSM manually after the update.

7. Apply the service pack to each server, following the order described in [“Update Sequence” on page 63](#):
 - For an HCAC, apply the service pack to the reporting Admin Node and then the processing Admin Node. Complete the application of the service pack to both servers in an HCAC before applying the service pack to any other server type.
 - a. Select the server in the **Servers** panel and confirm that the server state is Available.
 - b. Select **Update Software** in the **Tasks** panel, select **Start Task** in the **Actions** panel, and press **<Enter>**. Wait for the script to complete.
8. Verify that the updated services have joined the grid. For each server:
 - a. Check that the services recover from the Unknown state and that they are displayed in green.
 - b. Go to **SSM ► Services ► Overview ► Main** and check that the Packages section reports the storage-grid-release version 9.0.x on all servers.
 - c. Under Services, verify that each service reports the version number noted in the *Release Notes*.
9. After you have applied the service pack to all servers and have finished using the GDU console, quit GDU and remove the ability to access servers without a server password. See [“Close GDU” on page 86](#).

Upgrade the Operating System

Upgrade of the operating system to SLES 10 SP3 and SLES 11 SP2

Introduction

NOTE You must update the operating system (OS).

Before upgrading to StorageGRID 9.0.x, if not already running on at least SLES 10 SP3 (32 or 64-bit), you must upgrade SLES to at least this version of SLES. For information on qualified hardware, see the Interoperability Matrix Tool (IMT).

NOTE Both SLES 10 SP1 and SLES 10 SP2 are no longer supported.

After completing the upgrade of all grid nodes to StorageGRID 9.0.x, it is recommended that you upgrade the operating system from SLES 10 SP3 (64-bit) to SLES 11 SP2 (64-bit). Before upgrading the operating system to SLES 11 SP2 (64-bit) all grid nodes must be upgraded to StorageGRID 9.0.x. Failure to complete the upgrade of all grid nodes to StorageGRID 9.0.x before upgrading the SLES operating system may result in the failure of your StorageGRID system.

Note that you cannot upgrade StorageGRID software and the underlining operating system in parallel on different grid nodes. The upgrade to StorageGRID 9.0.x must finish in its entirety before upgrading the SLES operating system. For more information on upgrade paths and procedures see, [“Upgrade and Update Paths for StorageGRID and SLES Software”](#).



WARNING Do not upgrade to SLES 11 SP2 (64-bit) until all grid nodes are upgraded to StorageGRID 9.0.x.

Note that upgrading to SLES 11 SP2 (64-bit) may require a hardware refresh. You cannot upgrade from SLES 10 SP3 (32-bit) to SLES 11 SP2 (64-bit). For more information on how to perform a hardware refresh, see the *Expansion Guide*.

Server or Virtual Machine Update Sequence

Update the operating system one server or virtual machine at a time. Wait until the update has completed on one server or virtual machine before starting the next one.

For grid node type, you can update the OS in any sequence except for Gateway Nodes. For Gateway Nodes, update the operating system (OS) in the sequence listed below. For guidance on failover procedures, see [Chapter 4: “Gateway Node Upgrades”](#). If the grid does not support business continuity, client access is unavailable for the duration of the primary Gateway Node update. If the grid does support business continuity, client service is interrupted for a short period during the failover procedure. Plan client downtime accordingly. If the OS update and the grid software upgrade are scheduled for the same time, you may want to update the OS on the Gateway Node immediately following the software upgrade in order to minimize client disruptions.

Gateway Node OS Update Sequence

Table 8: Sequence for updating the OS on Gateway Nodes

Replication Group Type	Update Sequence
Gateway Node: No Business Continuity Failover	<ol style="list-style-type: none">1. Update all Gateway Nodes in the same replication group except for the primary Gateway Node.2. Update the primary Gateway Node.
Gateway Node: Business Continuity Failover	<ol style="list-style-type: none">1. Update the secondary Gateway Node.2. Manually fail over to the secondary Gateway Node.3. Update the primary Gateway Node.4. Fail back to the primary Gateway Node.

Table 8: Sequence for updating the OS on Gateway Nodes

Replication Group Type	Update Sequence
High Availability Gateway	<ol style="list-style-type: none"> 1. Update the secondary Gateway Node. 2. Update the standby primary Gateway Node. 3. Fail the cluster over 4. Update the server that is now the standby primary Gateway Node.

Update to SLES 10 SP3

Before upgrading to StorageGRID 9.0.x, perform the following procedure on each qualified server to update the OS to SLES 10 SP3 from either SLES 10 SP1 or SLES 10 SP2.

NOTE Both SLES 10 SP1 and SLES 10 SP2 are no longer supported.

For information on qualified hardware, see the Interoperability Matrix Tool (IMT) at support.netapp.com/matrix.

NOTE This is a manual procedure that does not use GDU.

Prerequisites

- Drivers for the hardware have been validated for the new service pack
- SLES 10 SP3 32 or 64-bit DVD

Procedure

1. Confirm that the server is running release 8.5.2 or later of the grid software. Go to **<grid_node> ► SSM ► Services ► Overview ► Main**. The following should be displayed:
 - The Operating System attribute displays SP1 or SP2.
 - The storage-grid-release attribute displays 8.5.2. or later.

2. Confirm that there is at least three GB of free disk space in `/var/local`:
 - a. At the server, access a command shell and log in as root using the password listed in the `Passwords.txt` file.
 - b. Enter: `df -h /var/local`
3. Confirm that the Enablement Layer for StorageGRID iso (`Enablement_Layer_for_StorageGRID_8.5.x_Software_<buildnumber>.iso`) is located at `/var/local/install`.
 If it is not, copy the iso to `/var/local/install` using the procedure listed in [Chapter 3: "Upgrade Software"](#).
4. At the server, insert the SLES 10 SP3 DVD.
5. Prepare the Smart package manager channel for SLES 10 SP3.
 Enter:


```
mount /cdrom
mkdir -p /var/local/install/sles-10-sp3
cd /var/local/install/sles-10-sp3
cp /cdrom/suse/*/*.rpm .
rm heartbeat-*.rpm
smart channel --add sles-10-sp3 type=rpm-dir path=/var/
local/install/sles-10-sp3
```

 When prompted to include this channel, press **Y**.
6. If the server hosts an FSG service, prepare the Smart channel for the enablement layer. Enter:


```
mkdir /tmp/el
mount -o loop,ro /var/local/install/Enablement_Layer_for_StorageGRID_8.5.x_Software_<buildnumber>.iso /tmp/el
smart channel --add el type=rpm-md baseurl=/tmp/el
```

 When prompted to include this channel, press **Y**.
7. Stop Server Manager and prevent services from starting. Enter:


```
/etc/init.d/servermanager stop
touch /etc/sv/ntp/DoNotStart
```
8. If Java is installed on the server, remove the current version. Enter:


```
smart remove java-1_4_2-sun
```
9. Upgrade the packages. Enter (on one line):


```
script -c "smart upgrade --update 2>&1" /var/local/log/os-upgrade.log
```
10. Inspect the screen output (or the content of `/var/local/log/os-upgrade.log`) for errors.

11. If the command fails, remove the cache and try to upgrade the packages again. Enter:


```
rm /var/lib/smart/cache
script -c "smart upgrade --update 2>&1" /var/local/log/os-upgrade.log
```
12. Remove the smart channel for SP3 and delete the SP3 rpm packages. Enter:


```
smart channel --remove sles-10-sp3
cd /var/local/install
rm -rf sles-10-sp3
```
13. If the server hosts an FSG, remove the smart channel for the enablement layer and unmount the enablement layer ISO. Enter:


```
smart channel --remove el
umount /tmp/el
```
14. Reboot the server. Enter: **reboot**
15. Update or recompile any other kernel drivers to match the new kernel version.
16. If you have modified any other drivers in step 15, reboot the server. Enter: **reboot**
17. Remove the DoNotStart file and start services. Enter:


```
rm /etc/sv/ntp/DoNotStart
/etc/init.d/servermanager restart
```
18. Log out of the command shell. Enter: **exit**
19. Go to <grid_node> ► **SSM** ► **Services** ► **Overview** ► **Main** and confirm that Operating System displays SP3.

Update to SLES 11 SP2 (64-Bit)

NOTE This is a manual procedure that does not use GDU.

After completing the upgrade of all grid nodes to StorageGRID 9.0.x, it is recommended that you upgrade the operating system from SLES 10 SP3 (64-bit) to SLES 11 SP2 (64-bit). Before upgrading the operating system to SLES 11 SP2 (64-bit) all grid nodes must be upgraded to StorageGRID 9.0.x. Failure to complete the upgrade of all grid nodes to StorageGRID 9.0.x before upgrading the SLES operating system may result in the failure of your StorageGRID system.

Note that you cannot upgrade StorageGRID software and the underlining operating system in parallel on different grid nodes. The upgrade to StorageGRID 9.0.x must finish in its entirety before upgrading the SLES operating system.

For information on qualified hardware, see the Interoperability Matrix Tool (IMT) at support.netapp.com/matrix.

Prerequisites

- Drivers for the hardware have been validated for the new service pack
- All grid nodes have been upgraded to StorageGRID 9.0.x
- SLES 11 SP2 (64-bit) DVD

Procedure



WARNING Do not upgrade to SLES 11 SP2 (64-bit) until all grid nodes are upgraded to StorageGRID 9.0.x.

1. Confirm that all grid nodes are running release 9.0.1 or later service pack of the grid software. Go to **<grid_node> ► SSM ► Services ► Overview ► Main**.

The following should be displayed:

- The Operating System attribute displays SLES 10 SP3 (x86_64).
- The storage-grid-release attribute displays 9.0.1 or later.

2. At the Admin Node, access a command shell and log in as root using the password listed in the Passwords.txt file.
3. Copy the autoupg.xml file to the server's root partition:

- a. For the Admin Node, enter:

```
cp /usr/local/gpt/versions/9.0.0/resources/autoupg.xml /root/
```

- b. For all other grid node types, enter:

```
scp /usr/local/gpt/versions/9.0.0/resources/autoupg.xml \
root@<IP_address>:/root
```

where **<IP_address>** is the IP address of the grid node to be upgraded.

4. Load the SLES 11 SP2 (64-bit) ISO onto the physical server or virtual machine.
5. Reboot the server. Enter: **reboot**

6. From the SLES Boot Screen:
 - a. Press the keyboard's down arrow and highlight Installation. (Do not press <Enter>.)
 - b. At **Boot Options**, enter: `autoupgrade=1`




Figure 8: SLES Boot Screen

- c. Press <Enter>.
 - d. Wait for upgrade to complete, log out of the server, and go to `<grid_node> ► SSM ► Services ► Overview ► Main` and confirm that Operating System displays SLES 11 SP2 (64-bit).

NOTE If you do not see the SLES install splash screen, ensure that you have Boot ► CD-ROM Drive listed above Boot ► From Hard drive in the bios settings.

7. After the OS is upgraded, upgrade VMware tools to the latest version.

Troubleshooting

Service Hangs

The GDU Upgrade Software task cannot complete if the server is prevented from shutting down because a service hangs. When this occurs, the GDU Upgrade Software task stops at the "Stopping services" stage and manual intervention is required to stop the service.

Terminate Service

1. At the affected server, access a command shell and log in as root using the password listed in the Passwords.txt file.
2. Display the status of services on the server. Enter:
`/usr/local/servermanager/reader.rb`
3. Examine the output to determine if any services are stuck in the "Stopping..." state.
4. If the service that is hung is Database Engine, shut down the MySQL service. Enter:
`mysqladmin --defaults-file=/etc/my.default.cnf shutdown`
5. If the service that is hung is one of adc, ams, arc, clb, cms, cmn, fsg, ldr, nms, or ssm, shut down the service forcibly. Enter:
`killall -9 <service-name>`

NOTE The service may appear as unknown (blue) in the NMS MI until it is restarted.

6. If the service is fsg, enter this command to stop any managed FSG services (NFS, CIFS, Heartbeat) that are still running:
`/usr/local/fsg/pre-start.sh`

Provisioning Failures

If provision fails, contact Support and provide the information contained in the log files written to the Provisioning USB flash drive or, if installed on a virtual machine, written to the location as set in the command `load-provisioning-software --alternate-usb-dir=<path>`.

If the grid specification file is incorrect, provisioning displays an error message and exits. This error message is written to the Provisioning USB flash drive or, if installed on a virtual machine, written to the location as set in the command `load-provisioning-software --alternate-usb-dir=<path>` and is named `provision-fail.log`.

If the provisioning process results in a program crash, two identical log files are saved to the Provisioning USB flash drive or, if installed on a virtual machine, written to the location as set in the command `load-provisioning-software --alternate-usb-dir=<path>`:

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

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

Server Crashes or Fails to Start

If a server crashes during the upgrade process or fails to start successfully after the upgrade completes, contact Support to investigate and to correct any underlying issues.

Grid Node Fails

If a grid node fails while an upgrade of the grid is in progress, contact Support.

Grid Task Pauses With Error

If a grid task pauses with error, contact Support for assistance.

Ingest or Data Retrieval is Interrupted

If data ingest or retrieval is unexpectedly interrupted (that is, other than when a Gateway Node is being upgraded), contact Support for assistance.

Upgrading the Grid Software Manually

If upgrading a grid node with GDU fails, it may be necessary to manually upgrade the software. For more information and the procedure, see [Chapter 6: “Manually Upgrading”](#).

Script Exceptions

The core upgrade process is scripted and should execute without problem. Should the script encounter an exception, the problem is reported with a message in this format:

```
updategrid.rb:<#>:<message> (RuntimeError)
```

The line number and an error message appear in the line.

If an error is encountered, use [Table 9](#) to identify and resolve the issue before running the script again. You can safely run the `updategrid.rb` script, that is, run the GDU Software Upgrade task, any number of times to retry an upgrade.

If the problem persists, contact Support for assistance. You can continue to upgrade other servers at the site while waiting for resolution of the issue unless the problem occurs on the first Admin Node or Control Node. It may be necessary run the grid without the failed server until such time as the problem can be resolved.

If the script fails without an error message, or with a message not listed in [Table 9](#), contact Support.

NOTE The file `/var/local/log/updategrid.log` is useful to diagnose problems. You may be asked to forward this file to Support if the upgrade does not complete successfully.

Table 9: updategrid.rb Error Messages

Error Message	Description and Troubleshooting
Wrong CD. Try again.	Check that the correct Enablement Layer for StorageGRID Software CD is inserted. Replace and continue.
This upgrade must include a new version of Enablement Layer for StorageGRID Software CD. Please include the ISO file for Enablement Layer for StorageGRID Software CD 9.0.	Check that the Enablement Layer for StorageGRID Software CD ISO was provided and that the version is correct. Try the upgrade again.
Please do not run this script from <code>/cdrom</code> . Change directory to somewhere else, and run this script again.	Make sure you did not run the upgrade script from <code>/cdrom</code> . Change the directory (enter <code>cd</code>) and try the upgrade script again.
Unknown argument: <code><argument></code>	Check that you typed the command line argument correctly. Try the command again.
Newest version already installed.	The upgrade has already been performed. Do not run the upgrade again.
Please verify you have ran the Software Upgrade grid task, which is part of the upgrade. Aborting upgrade.	Make sure you ran the Software Upgrade grid task and that it reached the stage Wait for Software Version before running the Software Upgrade grid task for the primary Admin Node. If you have not already run the grid task, run the task and try the upgrade again. If problems persist, contact Support.
Cannot find version file.	The script cannot verify that the server is running the correct version. Do not proceed with the upgrade.
Installed version: <code>x.x</code> not supported for upgrade.	The server is not running a supported version. The grid must be upgraded to a supported version before proceeding.

NMS MI is Unavailable on Grid with HCAC

During an upgrade or application of a service pack, if the NMS MI is unavailable on a grid with a High Capacity Admin Node Cluster (HCAC), confirm that both the reporting and processing Admin Nodes have completed the upgrade or service pack application process.

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**.

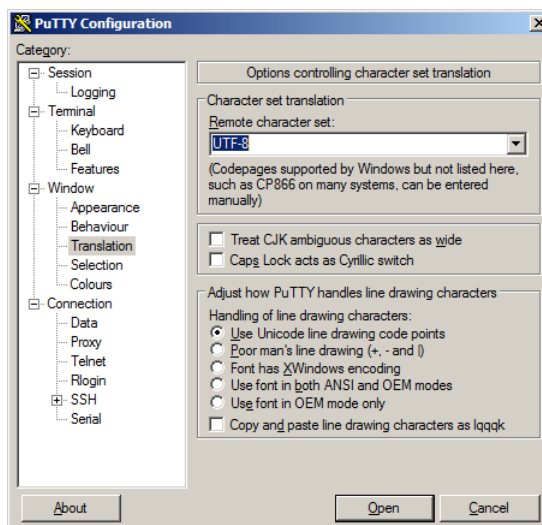


Figure 9: PuTTY Settings for GDU

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>**.




Figure 10: Entering the Provisioning Passphrase to Start GDU

If the characters do not display properly, see “GDU Display Problems” on page 87.

GDU User Interface




Figure 11: 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 <code>/etc/init.d/servermanager start</code>
Stop Services	Stop Server Manager and all services. This is equivalent to the command <code>/etc/init.d/servermanager stop</code>
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 <code>/var/local/install</code> 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.




Figure 12: 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 13](#) below.




Figure 13: 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.

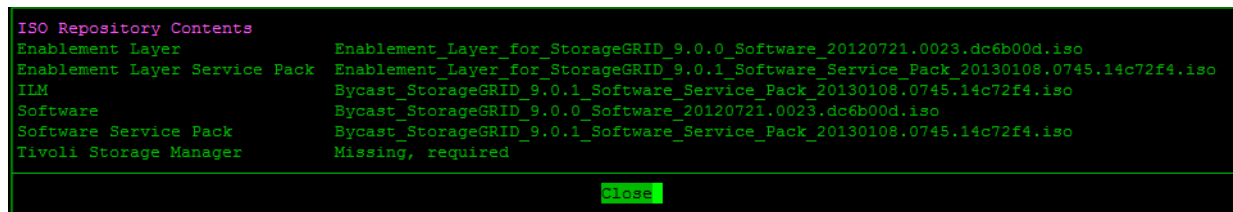
GDU Log Files

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 14](#) 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 14: 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/usbcreens/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 15: 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

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.

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.
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_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 primary Admin Node to the destination server. Enter:


```
scp /var/local/install/Enablement_Layer_for_StorageGRID_\
9.0_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_Software_<build>.iso \
Enablement_Layer_for_StorageGRID_9.0_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_Software_<build>.iso \
Enablement_Layer_for_StorageGRID_9.0_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`

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/cdm .
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 95.
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.
<hr/> NOTE Deduplication is deprecated and no longer supported. <hr/>	
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.
<hr/> 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

Symbols

/var/local/install directory 84, 88

A

ACLs
 defined 93
active primary FSG
 defined 93
ADC
 defined 93
ADE
 defined 93
Admin Node
 defined 93
AMS
 defined 93
API
 defined 93
API Gateway Node
 defined 93
ARC
 defined 93
Archive Node
 defined 94
atom
 defined 94
audit messages
 defined 94
Audit Node
 defined 94
AutoYaST
 defined 94
Available server state, in GDU 84

B

BASE64
 defined 94
basic gateway replication
 group 94
bundle
 defined 94
business continuity failover 94
Busy state, in GDU 84

C

CBID
 defined 95
checklist, software installation 16, 27
CIDR
 defined 95
CIFS
 defined 95
CLB
 defined 95

CMN

 defined 95

CMS

 defined 95
command, HTTP
 defined 95
container
 defined 95
content block ID
 defined 96
Continue Install GDU task 83
Control Node
 defined 96
CSTR
 defined 96

D

Data Center
 defined 96
DC site
 defined 96
deduplication
 defined 96
distributed CMS. See metadata replication
DR
 defined 96

E

EMR
 defined 96
Enable Services GDU task 83
Enablement Layer
 defined 96
Error server state, in GDU 84
exec bash, with GDU 81
expansion 51
 add Gateway Node 51
 add Storage Node 52
 Gateway Node 51
 grid nodes 51–54
 replication group 51

F

FCS
 defined 96
files
 gdu-console.log 87
FSG
 defined 97
 IPv6 support 11
FSG replication group
 defined 97

G

Gateway Node
 add during upgrade 51
 adding 51
 apply service pack 64
 defined 97
 expansion 51
 replication group
 adding 51
 defined 97
GDU 81–88
 Actions panel 84
 console 81
 defined 97
 entering commands 85
 ISOs 88
 Log Messages panel 84
 missing task 87
 navigation 85
 PuTTY, with 81, 87
 quit, how to 86
 screen, with 82, 88
 Server Info panel 84
 Servers panel 83
 start command 82
 Tasks panel 83
 troubleshooting 87
GDU tasks
 Continue Install 83
 Enable Services 83
 Install Driver 83, 86
 Install Software 83
 Load Configuration 83
 Reboot Server 83, 86
 Remount Storage 83
 Start Services 84
 Stop Services 84
 Update Software 84
 Update Status 84
 Upgrade Software 84
gdu-console command 82
gdu-console.log 87
GNU screen. See screen
grid
 defined 103
Grid ID
 signed text block 97
grid node
 defined 97
 expansion 51–54
grid specification file
 defined 98
grid tasks
 defined 98
 error 78

H

HAGC
definition 98

HCAC
definition 98

HTTP
defined 98

HTTPS
defined 98

I

ILM
defined 98

inode
defined 98

Install Driver GDU task 83, 86

Install Software GDU task 83

installation checklist 16, 27

ISO List action, in GDU 84

ISOs
copy to another site 89
missing, in GDU 88
required, missing, in GDU 88
service packs 89

K

KVM
defined 98

L

LAN
defined 99

latency
defined 99

LDR
defined 99

Load Configuration GDU task 83

load software 28

load_cds.py 28

M

main primary FSG
definition 99

manual upgrade procedure 55

metadata
defined 99

metadata replication
definition 99

metadata synchronization
definition 99

missing ISOs, in GDU 88

N

namespace
defined 99

nearline
defined 99

NFS

defined 100

NMS
defined 100

NMS MI unavailable,
troubleshoot 80

node ID
defined 100

NTP
defined 100

O

object identifier. See OID 100

object segmentation
defined 100

object stores
defined 100

OID
defined 100

operating system, upgrade 61, 69,
71, 73

P

Pingable server state, in GDU 84

primary Admin Node
defined 100

primary site, definition 63

processing Admin Node
defined 100

processing NMS
defined 100

provision the grid 31

provisioning 31
defined 101
failure, troubleshoot 78

PuTTY, with GDU 81, 87

Q

Quit action, in GDU 84

R

Reachable server state, in GDU 84

Reboot Server GDU task 83, 86

Remount Storage task, in GDU 83

replication group, adding 51

reporting Admin Node
defined 101

reporting NMS
defined 101

required
missing ISOs, in GDU 88

S

SAID package
defined 101

Samba
defined 101

SATA
defined 101

screen

GDU, with 82

multidisplay mode 88

SCSI
defined 101

security partitions
defined 101

segment container
defined 102

server
defined 102

server crash, troubleshoot 78

Server Manager
defined 102

service
defined 102

service pack 89
apply 65
ISOs 89

service, terminate 77

SLES
defined 102

SLES 10 SP1 69

SLES 10 SP2 69

SLES 10 SP3 69

SLES 11, IPv6 support 11

upgrade 61, 69

upgrade SLES 10 SP2 (64-bit) 73

upgrade SLES 10 SP3 71

SQL
defined 102

ssh
defined 102

SSH Access Password, with
GDU 82

ssh key, for upgrade 40

ssh-add -D 86

ssh-add, with GDU 82

SSL
defined 102

SSM
defined 102

standby primary FSG
definition 102

Start Services task, in GDU 84

Start Task action, in GDU 84

Stop Services task, in GDU 84

Storage Node
add during upgrade 52
defined 102

supplementary primary FSG
definition 103

SUSE
defined 103

synchronized CMS. See metadata
synchronization

T

task signed text block
defined 103

TCP/IP
defined 103

Telnet/ssh client, with GDU [81](#), [87](#)

throughput
defined [103](#)

TLS
defined [103](#)

transfer syntax
defined [103](#)

troubleshooting
GDU [87](#)

TSM
defined [103](#)

U

Update Software task, in GDU [84](#)

Update Status task, in GDU [84](#)

upgrade [31](#)
complete [59](#)
expansion during [51](#)
load software [28](#)
manually [55](#)
SLES [61](#), [69](#)
SLES 10 SP2 (64-bit) [73](#)
SLES 10SP3 [71](#)
troubleshoot [77](#)

Upgrade Software task, in
GDU [84](#)

URI
defined [103](#)

Use font encoding setting in
PuTTY, with GDU [87](#)

UTC
defined [103](#)

UTF-8 setting in PuTTY, with
GDU [81](#)

UUID
defined [103](#)

V

virtual machine
definition [104](#)

W

WAN
defined [104](#)

X

XFS
defined [104](#)

XML
defined [104](#)

