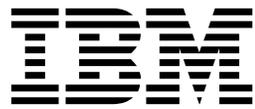


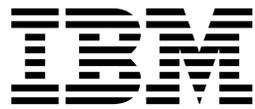
Elastic Storage Server
Version 5.3.2.1

Quick Deployment Guide



Elastic Storage Server
Version 5.3.2.1

Quick Deployment Guide



Note

Before using this information and the product it supports, read the information in “Notices” on page 99.

- | This edition applies to version 5.3.2.1 of the Elastic Storage Server (ESS) for Power, and to all subsequent releases and modifications until otherwise indicated in new editions.

IBM welcomes your comments; see the topic “How to submit your comments” on page ix. When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright IBM Corporation 2018, 2019.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Tables	v	Appendix H. Upgrading the Hardware Management Console (HMC)	67
About this information	vii	Appendix I. Obtaining kernel for system upgrades	69
Who should read this information.	vii	About the ESS Red Hat Linux Errata Kernel Update	70
Related information	vii	Appendix J. Obtaining systemd update for system upgrades.	73
Conventions used in this information	viii	About the ESS Red Hat Linux systemd update	74
How to submit your comments	ix	Appendix K. Obtaining Network Manager updates for system upgrades.	77
Chapter 1. Installation and upgrade related information and checklists	1	About the ESS Red Hat Linux Network Manager update	78
Chapter 2. Installing Elastic Storage Server	9	Appendix L. MES upgrade flow	81
Optional: Plug-N-Play	15	Appendix M. Running gssinstallcheck in parallel	85
Post Plug-N-Play cleanup steps.	17	Appendix N. NTP setup	87
Chapter 3. Upgrading Elastic Storage Server	19	Appendix O. Legacy deployment instructions	89
Chapter 4. gssutils - ESS Installation and Deployment Toolkit	29	Appendix P. Considerations for adding PPC64LE building blocks to ESS PPC64BE building blocks	95
Chapter 5. Adding building blocks to an existing ESS cluster	37	Appendix Q. Shutting down and powering up ESS	97
Appendix A. Known issues	39	Notices	99
Appendix B. Troubleshooting for ESS on PPC64LE	49	Trademarks	100
Appendix C. ESS networking considerations	51	Glossary	103
Appendix D. Support for hybrid enclosures	55		
Appendix E. Pre-installation tasks for ESS	57		
Appendix F. Installation: reference	63		
Appendix G. Updating the system firmware	65		

Tables

1.	Conventions	viii	4.	Pre-installation tasks	57
2.	Known issues in ESS 5.3.2.1	39			
3.	Recommended high speed switch firmware levels	54			

About this information

This information guides you in installing, or upgrading to, version 5.3.2.x of the Elastic Storage Server (ESS).

Who should read this information

This information is intended for experienced system installers and upgraders who are familiar with ESS systems.

Related information

ESS information

The ESS 5.3.2.x library consists of following information units. You can access these publications on IBM® Knowledge Center or IBM Publications Center.

- *Elastic Storage Server: Quick Deployment Guide*, SC27-9526
- *Elastic Storage Server: Protocol Nodes Quick Deployment Guide*, SC27-9527
- *Elastic Storage Server: Problem Determination Guide*, SC27-9529
- *Elastic Storage Server: Command Reference*, SC27-9530
- *IBM Spectrum Scale RAID: Administration*, SC27-9528
- *IBM ESS Expansion: Quick Installation Guide (Model 084)*, SC27-4627
- *IBM ESS Expansion: Installation and User Guide (Model 084)*, SC27-4628
- *IBM ESS Expansion: Hot Swap Side Card - Quick Installation Guide (Model 084)*, GC27-9210
- *IBM ESS Expansion: Hardware Installation and Maintenance Guide (Model 106)*, SC27-9211
- *IBM ESS Expansion: Overview of CMA and Rail Kit Hardware Fasteners (Model 106)*, SC27-9296
- *Installing the Model 024, ESLL, or ESLS storage enclosure*, GI11-9921
- *Removing and replacing parts in the 5147-024, ESLL, and ESLS storage enclosure*
- *Disk drives or solid-state drives for the 5147-024, ESLL, or ESLS storage enclosure*
- For information about the DCS3700 storage enclosure, see:
 - *System Storage® DCS3700 Quick Start Guide*, GA32-0960-04:
 - *IBM System Storage DCS3700 Storage Subsystem and DCS3700 Storage Subsystem with Performance Module Controllers: Installation, User's, and Maintenance Guide*, GA32-0959-07:
<http://www.ibm.com/support/docview.wss?uid=ssg1S7004920>
- For information about the IBM Power Systems™ EXP24S I/O Drawer (FC 5887), see IBM Knowledge Center :
http://www.ibm.com/support/knowledgecenter/8247-22L/p8ham/p8ham_5887_kickoff.htm

For more information, see IBM Knowledge Center:

http://www-01.ibm.com/support/knowledgecenter/SSYSP8_5.3.2/sts532_welcome.html

For the latest support information about IBM Spectrum Scale™ RAID, see the IBM Spectrum Scale RAID FAQ in IBM Knowledge Center:

<http://www.ibm.com/support/knowledgecenter/SSYSP8/gnrfaq.html>

Other related information

For information about:

- IBM Spectrum Scale, see:
http://www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale_welcome.html
- IBM Spectrum Scale call home, see Understanding call home.
- Installing IBM Spectrum Scale and CES protocols with the installation toolkit, see Installing IBM Spectrum Scale on Linux nodes with the installation toolkit.
- IBM POWER8[®] servers, see IBM Knowledge Center:
<http://www.ibm.com/support/knowledgecenter/POWER8/p8hdx/POWER8welcome.htm>
- Extreme Cluster/Cloud Administration Toolkit (xCAT), go to the xCAT website :
<http://xcat.org/>
 - xCAT 2.14.3 Release Notes
- Mellanox OFED Release Notes (4.4), go to https://www.mellanox.com/related-docs/prod_software/Mellanox_OFED_Linux_Release_Notes_4_4-1_0_0_0.pdf
- IBM Electronic Service Agent (ESA) documentation, go to <https://www-01.ibm.com/support/esa/>.
- Drive call home, go to Drive call home in 5146 and 5148 systems.

Conventions used in this information

Table 1 describes the typographic conventions used in this information. UNIX file name conventions are used throughout this information.

Table 1. Conventions

Convention	Usage
bold	Bold words or characters represent system elements that you must use literally, such as commands, flags, values, and selected menu options. Depending on the context, bold typeface sometimes represents path names, directories, or file names.
<u>bold underlined</u>	<u>bold underlined</u> keywords are defaults. These take effect if you do not specify a different keyword.
constant width	Examples and information that the system displays appear in constant-width typeface. Depending on the context, constant-width typeface sometimes represents path names, directories, or file names.
<i>italic</i>	<i>Italic</i> words or characters represent variable values that you must supply. <i>Italics</i> are also used for information unit titles, for the first use of a glossary term, and for general emphasis in text.
<key>	Angle brackets (less-than and greater-than) enclose the name of a key on the keyboard. For example, <Enter> refers to the key on your terminal or workstation that is labeled with the word <i>Enter</i> .
\	In command examples, a backslash indicates that the command or coding example continues on the next line. For example: <pre>mkcondition -r IBM.FileSystem -e "PercentTotUsed > 90" \ -E "PercentTotUsed < 85" -m p "FileSystem space used"</pre>
{item}	Braces enclose a list from which you must choose an item in format and syntax descriptions.
[item]	Brackets enclose optional items in format and syntax descriptions.
<Ctrl-x>	The notation <Ctrl-x> indicates a control character sequence. For example, <Ctrl-c> means that you hold down the control key while pressing <c>.

Table 1. Conventions (continued)

Convention	Usage
<i>item...</i>	Ellipses indicate that you can repeat the preceding item one or more times.
	In <i>synopsis</i> statements, vertical lines separate a list of choices. In other words, a vertical line means <i>Or</i> . In the left margin of the document, vertical lines indicate technical changes to the information.

How to submit your comments

Your feedback is important in helping us to produce accurate, high-quality information. You can add comments about this information in IBM Knowledge Center:

http://www.ibm.com/support/knowledgecenter/SSYSP8/sts_welcome.html

To contact the IBM Spectrum Scale development organization, send your comments to the following email address:

scale@us.ibm.com

Chapter 1. Installation and upgrade related information and checklists

Review the following installation and upgrade related information before starting with the installation or the upgrade of Elastic Storage Server (ESS).

- “New features and enhancements”
- “Component versions for this release” on page 2
- “Supported editions on each architecture” on page 2
- “ESS best practices and support statements” on page 2
- “Obtaining the required Red Hat Enterprise Linux and ESS code” on page 4
- “Supported upgrade paths” on page 5
- “Support for hardware call home” on page 5
- “Pre-installation or pre-upgrade checklist” on page 5
- “Post-installation or post-upgrade checklist” on page 6
- “Other topics” on page 7
- “Sample installation and upgrade flow” on page 7

New features and enhancements

Release	Changes
ESS 5.3.2.1	<ul style="list-style-type: none">• Support for IBM Spectrum Scale 5.0.2 PTF1 + efix• Support for new Network Manager version (1.12.0-8.el7_6)• Support for new kernel version (3.10.0-862.25.3)• Support for new MOFED (4.4-2.0.8.1)• Deployment bug fixes
ESS 5.3.2	<ul style="list-style-type: none">• Support for 4U106 storage enclosures<ul style="list-style-type: none">– Power® servers use 5 SAS adapters• Support for IBM Spectrum Scale 5.0.2 PTF1 + efix (5.0.2.1.0.1)• Support for new Power Firmware, ESA, xCAT, IPR• Support for new MOFED (4.4)• Support for Red Hat Enterprise Linux 7.5• Support for MES upgrades<ul style="list-style-type: none">– GLxS and GSxS models only• Security enhancements<ul style="list-style-type: none">– SSH shell timeout (5 minutes)– USB port disable– CD-ROM disable– xCAT SSL update (TLSv12) + Cipher support

Component versions for this release

Note: Your version might be slightly different from the version indicated in this document. Refer to the release notes document that comes with the installation package for the exact version of the installation package and the component version.

The respective versions for the core components in this release of ESS are as follows:

- | • Supported architectures: PPC64BE and PPC64LE
- | • IBM Spectrum Scale: 5.0.2.1 + efix
- | • xCAT: 2.14.3
- | • HMC: 860 SP3
- | • System firmware: SV860_165 (FW860.51)
- | • Red Hat Enterprise Linux: 7.5 (PPC64BE and PPC64LE)
- | • Kernel: 3.10.0-862.25.3
- | • Systemd: 219-57.el7_5.3
- | • Network Manager: 1.12.0-8.el7_6
- | • OFED: MLNX_OFED_LINUX-4.4-2.0.8.1
- | • IPR: 19511000
- | • ESA: 4.4.2-3

Supported editions on each architecture

The following are the ESS editions supported on the available architectures.

PPC64BE

- | • Standard Edition (Data Access edition)
- Advanced Edition
- Data Management Edition

PPC64LE

- | • Standard Edition (Data Access edition)
- Data Management Edition

ESS best practices and support statements

- | • It is advised that you set `autoload` to `on` to enable GPFS™ to recover automatically in case of a daemon problem. Deployment automatically enables this on new installations but you should disable `autoload` for an upgrade and re-enable it after an upgrade.
- | To disable, issue the following command:
 - | `mmchconfig autoload=no`
- | Once the maintenance operation or upgrade is complete, re-enable `autoload`.
 - | `mmchconfig autoload=yes`
- By default, file systems must only be mounted on the management server node (EMS). Do not mount the file system on any other ESS nodes besides the EMS (where the primary GUI runs) which is mandatory for the GUI to function correctly.
- It is advised that you disable `automount` for file systems when performing an upgrade to ESS 5.3.1 or later.
 - `mmchfs Device -A no`
 - Device* is the device name of the file system.
- | Automount should automatically be disabled when creating new file systems with **gssgenvdisks**.

Remember: Mount the file system only on the EMS node where the GUI and the PM collector run.

- Do not configure more than 5 failure groups in a single file system.
- Consider moving all supported Infiniband devices to the Datagram mode (CONNECTED_MODE=no). For more information, see Appendix C, “ESS networking considerations,” on page 51.
- If you have 40Gb adapters, enable flow control on your switch. Consider doing the same for 100Gb adapters.
- RDMA over Ethernet (RoCE) is not supported.
- Sudo on the ESS nodes is not supported.
- Enabling the firewall on any ESS node is not supported.
- Enabling SELinux on any ESS node is not supported.
- Running any additional service or protocols on any ESS node is not supported.
- Consider moving quorum, cluster, and file system management responsibilities from the ESS nodes to other server license nodes within the cluster.
- It is not required, though highly recommended, that the code levels match during a building block addition. Be mindful of changing the release and file system format in mixed IBM Spectrum Scale environments.
- You must take down the GPFS cluster to run firmware updates in parallel.
- Do not independently update IBM Spectrum Scale (or any component) on any ESS node unless specifically advised from the L2 service. Normally this is only needed to resolve an issue. Under normal scenarios it is advised to only upgrade in our tested bundles.
- It is acceptable for LBS or customers to update any security errata available from Red Hat Network (RHN). Only components checked and protected by ESS (kernel, network manager, systemd) must not be modified unless advised by the IBM service. For more information on applying security erratas see <https://access.redhat.com/solutions/10021>
- Client node deployment is not supported from the ESS management node.
- You must deploy or add building blocks from an EMS with the same architecture. There must be a dedicated EMS for each architecture (PPC64BE or PPC64LE).
- If running in a mixed architecture environment, the GUI and PM collector are recommended to run on the PPC64LE EMS node.
- Modifying any ESS nodes as a proxy server is not supported.
- Offline upgrades from any prior ESS version are supported. (No level hop required; building blocks updated in parallel).
- PPC64LE to PPC64BE conversions and vice versa are not supported.
- Multiple building blocks are ideal as ESS now by default uses file system level metadata replication. If a single building block is used, by default **gssgenvdisks** uses one failure group and only IBM Spectrum Scale RAID level metadata replication.
- It is recommended to use the highest available block size when creating vdisks or NSDs. This includes situations where the customer is only generating tiny files. The default block size is 16M (current maximum).
- Broadcom high-speed adapters (EL3Z PCIe2 LP 2-port 10/1GbE BaseT RJ45 Adapter) are supported by ESS but not currently supported by the deployment scripts. Contact the L2 service for the procedure to enable this adapter type.
- It is recommended that all nodes in a cluster run the same version of Mellanox OFED.
- Automatic EMS failover is not supported. For help in setting up a redundant, standby EMS, contact the L2 service.
- 4K MTU (InfiniBand) and 9000 MTU (Ethernet) are recommended. Changing to these MTU values requires associated switch-side changes.

Obtaining the required Red Hat Enterprise Linux and ESS code

If you are a member of IBM, you must contact ESS development or L2 service to obtain the code directly.

The required Red Hat components and SHA256 are:

- Red Hat Enterprise Linux 7.5 ISO

```
| f39deefb462ef927b4cc0db7b9bfb45fd144fa7e947947648e8edb80dcb7089 RHEL-7.5-20180322.0-Server-ppc64-dvd1.iso
| 6a3dbaa9e77b270f90ceedd93e75026bd2575e89103f4eeb81bb11137d7283bd RHEL-7.5-20180322.0-Server-ppc64le-dvd1.iso
```

- Network manager version : 1.12.0-8.el7_6

```
| be5b4bce1a60ea1e140924b551bbd9c78e81ebda81851aff0161d082c695ebb8 netmgr-RHSA-2018-3665-BE-ESS5321.tar.gz
| f6c3470a8da6c1b582ddcf1384a9512f4715a74118c70f23df77b9ab8779bb23 netmgr-RHSA-2018-3665-LE-ESS5321.tar.gz
```

- Systemd version: 219-57.el7_5.3

```
| 76eb76638ff3376767ebf4c5ec17df5f52f78f9e74ac798fe5831fa13f0fd3c1 systemd-532-RHBA-2018-BE.tar.gz
| d8b979c26ec72cbcf476e0d98f91d759e54fc84f523a0e09dcb9c54d1d1893f9 systemd-532-RHBA-2018-LE.tar.gz
```

- Kernel version: 3.10.0-862.25.3

```
| da42b57c6d55e3927647ba3ef4062eb8969905560b107f41c7eb19830f63d79f kernel-RHSA-2018-3843-BE-ESS5321.tar.gz
| 424fbb155b0631f787537209514b873cdd97c816f13d6ef220ac4db50c0c551c kernel-RHSA-2018-3843-LE-ESS5321.tar.gz
```

On ESS 5.3.2.1 systems shipped from manufacturing, these items can be found on the management server node in the `/home/deploy` directory.

Customers or business partners can download the required Red Hat components from Red Hat Network using the customer license. For more information, see:

- Appendix I, “Obtaining kernel for system upgrades,” on page 69
- Appendix J, “Obtaining systemd update for system upgrades,” on page 73
- Appendix K, “Obtaining Network Manager updates for system upgrades,” on page 77

The ESS software archive that is available in different versions for both PPC64BE and PPC64LE architectures.

Available PPC64BE packages:

```
ESS_STD_BASEIMAGE-5.3.2.1-ppc64-Linux.tgz
ESS_ADV_BASEIMAGE-5.3.2.1-ppc64-Linux.tgz
ESS_DM_BASEIMAGE-5.3.2.1-ppc64-Linux.tgz
```

Available PPC64LE packages:

```
ESS_STD_BASEIMAGE-5.3.2.1-ppc64le-Linux.tgz
ESS_DM_BASEIMAGE-5.3.2.1-ppc64le-Linux.tgz
```

ESS 5.3.2.1 can be downloaded from IBM FixCentral.

Once downloaded and placed in `/home/deploy`, `untar` and `uncompress` the package to view the contents. For example, for the standard edition PPC64LE package, use the following command:

```
tar -xvf ESS_STD_BASEIMAGE-5.3.2.1-ppc64le-Linux.tgz
```

For example, from the BASEIMAGE tar file, files such as the following get extracted with the preceding command:

- `ESS_5.3.2.1_ppc64le_Release_note_Standard.txt`: This file contains the release notes for the latest code.
- `gss_install-5.3.2.1_ppc64le_standard_20190121T212642Z.tgz`: This .tgz file contains the ESS code.
- `gss_install-5.3.2.1_ppc64le_standard_20190121T212642Z.md5`: This .md5 file to check the integrity of the tgz file.

Supported upgrade paths

The following upgrade paths are supported:

- | • ESS version 5.1.x, 5.2.x, and 5.3.x.y to version 5.3.2.1 on PPC64BE.
- | • ESS version 5.1.x, 5.2.x, and 5.3.x.y to version 5.3.2.1 on PPC64LE.
- | **Note:** For upgrading to ESS 5.3.2.1 from version 5.0.x or earlier (Support for PPC64LE began in 5.1.x) and performing an online (rolling) update, you must follow the hops below. Offline upgrade from any version is supported though.
- | • **3.5.5 (or earlier) > 4.5.2 > 5.1.x > 5.3.2.1**
- | • **4.0.x > 5.0.x > 5.1.x (or 5.2.x) > 5.3.2.1**
- | • **4.5.x (or 4.6.x) > 5.1.x > 5.3.2.1**

Offline upgrades to ESS 5.3.2.1 from any prior ESS version are supported.

Support for hardware call home

	PPC64BE	PPC64LE
Call home when disk needs to be replaced	X	X
Enclosure call home	Unsupported	Unsupported
Server call home	Through HMC	Unsupported

For more information, see Drive call home in 5146 and 5148 systems.

Note:

Software call home is supported on PPC64BE and PPC64LE architectures.

Pre-installation or pre-upgrade checklist

Before you arrive at a customer site or before upgrade, it is advised that you perform the following tasks:

	Obtain the kernel, systemd, networkmanager, RHEL ISO (Provided by ESS development or L2 Service), and ESS tarball (FixCentral). Verify that the checksum match with what is listed in this document. Also ensure that you have the correct architecture packages (PPC64LE or PPC64BE).
	Ensure that you read all the information in the ESS Quick Deployment Guide. Make sure that you have the latest copy from the IBM Knowledge Center and the version matches accordingly. You should also refer to the related ESS 5.3.2 documentation in IBM Knowledge Center.
	Obtain the customer RHEL license.
	Contact the local SSR and ensure that all hardware checks have been completed. Make sure all hardware found to have any issues has been replaced.
	If the 1Gb switch is not included in the order, contact the local network administrator to ensure isolated xCAT and FSP VLANs are in place.
	Develop an inventory and plan for how to upgrade, install, or tune the client nodes.
	Upgrade the HMC to SP3 if doing a PPC64BE installation. This can be done concurrently. The SSR or the customer might be able to do this ahead of time.
	Consider talking to the local network administrator regarding ESS switch best practices, especially the prospect of upgrading the high-speed switch firmware at some point prior to moving the system into production, or before an upgrade is complete. For more information, see “Customer networking considerations” on page 53.

	Review <i>Elastic Storage Server: Command Reference</i> .
	Review ESS FAQ and ESS best practices.
	Review the ESS 5.3.2 known issues.
	Ensure that all client node levels are compatible with the ESS version. If needed, prepare to update the client node software on site and possibly other items such as the kernel and the network firmware or driver.
	Power down the storage enclosures, or remove the SAS cables, until the gssdeploy -x operation is completed. Note: You would only use gssdeploy -x if the legacy installation sequence is used.
	If adding an PPC64LE building block to an existing PPC64BE building block, carefully review Appendix P, "Considerations for adding PPC64LE building blocks to ESS PPC64BE building blocks," on page 95.
	If installing or upgrading protocol nodes, carefully review <i>Elastic Storage Server: Protocol Nodes Quick Deployment Guide</i> .
	Carefully study the network diagram for the architecture used. For more information, see Appendix C, "ESS networking considerations," on page 51 and <i>5148-22L protocol node diagrams</i> in <i>Elastic Storage Server: Protocol Nodes Quick Deployment Guide</i> .
	It is recommended to use a larger block size with IBM Spectrum Scale 5.0.0 or later, even for small I/O tasks. Consult the documentation carefully.
	Ensure that the correct edition of ESS is to be deployed. For example, do not install the Data Management Edition if Standard Edition is on the order. This must be verified even before Plug-N-Play is attempted.
	If this is a non-Mellanox order (Broadcom 10Gb adapter) being used, consult L2 for procedure as the deployment scripts are unsupported.
	Determine the supported high-speed switch bonding mode and, if Infiniband, determine which MTU will be used. The default MTU is now 2048 but it can be changed to 4092.

Post-installation or post-upgrade checklist

After the installation or upgrade is completed, it is advised that you verify the following:

	Hardware and software call home have been set up and tested. If applicable, consider postponing the call home setup until the protocol nodes are deployed. <ul style="list-style-type: none"> • For more information, see Drive call home in 5146 and 5148 systems. • For information about HMC call home (Server PPC64BE Only), see Configuring HMC Version 8.8.3 and Later for Call Home. • For call home support information, see "Support for hardware call home" on page 5. • For software call home information, see Software call home.
	GUI has been set up and demonstrated to the customer. If applicable, consider postponing the GUI setup until the protocol nodes are deployed.
	GUI SNMP and SMTP alerts have been set up, if desired.
	The customer RHEL license is registered and active.
	No issues have been found with mmhealth , GUI, gnrhealthcheck , gssinstallcheck , serviceable events.
	No SAS width or speed issues have been found. Consult development for more information.
	Client nodes are properly tuned. For more information, see "Adding IBM Spectrum Scale nodes to an ESS cluster" on page 63.

	It is advised that you turn on autoload to enable GPFS to recover automatically in case of a daemon problem. mmchconfig autoload=yes
	Connect all nodes to Red Hat Network (RHN).
	Update any security related erratas from RHN if the customer desires (yum -y security). Do not update any kernel, systemd, or network manager erratas.
	Ensure that you have saved a copy of the xCAT database off to a secure location.
	Install or upgrade the protocols. For more information, see <i>Elastic Storage Server: Protocol Nodes Quick Deployment Guide</i> .
	Ensure (if possible) that all network switches have had the firmware updated.
	IBM Spectrum Scale release level and file system format have been updated, if applicable.
	If there are more than one building block, make sure multiple failure groups are used in a file system and that metadata replication is turned on (-m 2). Do not exceed 5 failure groups or more than 2 metadata replicas.
	Consider offering your clients the following trainings to better enable the proper administration of an IBM Spectrum Scale cluster and ESS storage.
	<ul style="list-style-type: none"> • IBM Spectrum Scale Redbook • IBM Spectrum Scale basic training classes • IBM Spectrum Scale advanced training classes

Other topics

For help with the following topics, and many others that are unlisted, contact L2 Service.

- Restoring a management server
- Part upgrades or replacements
- VLAN reconfiguration on the 1Gb switch
- | • Extending the 1Gb Cumulus management switch

Sample installation and upgrade flow

New installations go through manufacturing CSC. The system is fully installed with ESS 5.3.2.1, tested, malfunctioning parts replaced, and required RHEL pieces shipped in /home/deploy.

Installation

Note: It is recommended that all operations be completed using **gssutils**. Protocol node deployment are not supported using **gssutils**.

- SSR checkout complete
- LBS arrival on site
- Plug-n-Play mode demonstrated
- | • Decisions made on file system names and sizes, block size, host names, IP addresses, and so on
- Check high-speed switch settings and firmware
- | • Deploy EMS and building block
- Network bonds created
- Cluster created
- Recovery groups, NSDs, file system created
- Stress test performed
- Final checks performed

Proceed to install the protocol nodes, if applicable.

- GUI setup (w/SNMP alerts if desired)
- Call home setup
- Nodes attached to RHN and security updates applied

Upgrade

To upgrade to an ESS 5.3.2.1 system at the customer site, it is recommended that you use **gssutils**. For more information, see Chapter 4, “**gssutils** - ESS Installation and Deployment Toolkit,” on page 29.

Note: It is recommended that all operations be completed using **gssutils**. Protocol node deployment are not supported using **gssutils**.

- Check high speed switch settings and firmware
- Ensure that there are no hardware issues
- Ensure client / protocol node compatibility
- Ensure no heavy IO operations are being performed
- Upgrade ESS (rolling upgrade or with cluster down)
 - Always ensure you have quorum (if rolling upgrade)
 - Always carefully balance the recovery groups and scale management functions as you upgrade each node (if rolling upgrade)
- Move the release level and the file system format, if applicable.
- Final checks are performed
- | • If applicable, upgrade the ESS protocol nodes
- | • Ensure that call home and GUI are still working as expected
- Use yum to upgrade any security related errata (**yum -y security**). Do not update any kernel, systemd, or network manager erratas.

Chapter 2. Installing Elastic Storage Server

| **Note:** Any versions mentioned in the following steps are just examples and they might vary from the
| actual product versions.

New installations from manufacturing provide all the necessary pieces to start the deployment.

Note: Though manufacturing supplies the relevant pieces needed to deploy the system, it is always a good practice to consult the latest release notes and *Elastic Storage Server: Quick Deployment Guide* for any changes before starting the deployment.

Inside the `/home/deploy` directory on the management server node (EMS), the following items are available for the architecture you are deploying:

- Kernel
- Systemd
- Network Manager
- RHEL ISO
- README containing critical xCAT information such as MAC addresses
- ESS tarball

| The xCAT database is intact using the default host name and IP addresses. All nodes are installed and do
| not need to be re-installed. Manufacturing has kept intact `/etc/hosts` on each node and a copy of
| `gssdeploy.cfg` in `/var/tmp` on the EMS node.

Note: It is recommended to use `gssutils` for all aspects of deployment or upgrade.

Assuming that the SSR has completed the full check of the system (no bad hardware, device paths, basic networking verified), you have the option to start by using the Plug-N-Play to demonstrate to the customer how fast and easy creating a file system can be and to provide an overview of the GUI. For more information, see “Optional: Plug-N-Play” on page 15.

| **Note:** If Plug-N-Play was used, clean up the system before proceeding with full installation. For more
| information, see “Post Plug-N-Play cleanup steps” on page 17.

| Before proceeding further, ensure that you have reviewed and completed the pre-installation and
| pre-upgrade checklist. For more information, see “Pre-installation or pre-upgrade checklist” on page 5.

ESS installation flow

| **Note:** All installation (or upgrade) steps can be done by using `gssutils`. Depending on your comfort
| level and preference, you may choose to use that interface. For more information, see Chapter 4,
| “`gssutils` - ESS Installation and Deployment Toolkit,” on page 29.

Extract the Elastic Storage Server software

Note: The package name depends on the platform and the edition on which you are installing the software.

1. Determine the current ESS version shipped from manufacturing.
| `xdsh ems1,gss_ppc64 "/opt/ibm/gss/tools/bin/gssinstallcheck -N localhost --get-version" | xcoll -n`
| Save this information because it is needed later in the installation flow.

2. Unpack the ESS software archive (This is contained in the ESS_STD_BASEIMAGE-5.3.2.1-ppc64le-Linux.tgz file.

```
tar -zxvf gss_install-5.3.2.1_ppc64_standard_20190121T212642Z.tgz
```

3. Check the MD5 checksum:

```
md5sum -c gss_install-5.3.2.1_ppc64_standard_20190121T212642Z.md5
```

4. Make sure the /opt/ibm/gss/install/rhel7/<ARCH> directory is clean:

```
/bin/sh gss_install-5.3.2.1_ppc64_standard_20190121T212642Z --remove
```

Depending on the architecture, replace <ARCH> with ppc64 or ppc64le.

5. Extract the ESS packages and accept the license as follows. By default, it is extracted to the /opt/ibm/gss/install directory:

```
/bin/sh gss_install-5.3.2.1_ppc64_standard_20190121T212642Z --text-only
```

Deploy the ESS

Before proceeding with the following steps, ensure that you have completed all the steps in “Extract the Elastic Storage Server software ” on page 9.

Follow these steps to perform a new installation of the ESS software on a management server node and I/O server nodes. Node host names **ems1**, **gssio1**, and **gssio2** are examples. Each environment could have its own unique naming conventions. For an xCAT command such as **updatenode**, use an xCAT host name. For the IBM Spectrum Scale commands (those start with mm), use an IBM Spectrum Scale host name. For example, **ems1** is an xCAT host name (typically a hostname associated with the management interface) and **ems1-hs** is the corresponding IBM Spectrum Scale host name (typically a host name associated with the high speed interface).

1. Make the **gssdeploy** script executable, if it is not yet executable:

```
chmod +x /opt/ibm/gss/install/rhel7/<ARCH>/samples/gssdeploy
```

2. Run one of the following commands depending on the architecture:

For PPC64BE:

```
cd /var/tmp ; ./gssinstall_ppc64 -u
```

For PPC64LE:

```
cd /var/tmp ; ./gssinstall_ppc64le -u
```

3. Change xCAT IPs, host names, domain, and netmasks.

- a. Copy the gsschenv.cfg from /opt/ibm/gss/tools/conf to /opt/ibm/gss/tools/bin.

- b. Modify the **gsschenv.cfg**.

```
# cat gsschenv.cfg
# Modify the following
# HOSTNAME_ORIG = Original hostnames in your xCAT ESS environment
# IP_ADDR_ORIG = Original IPs in your xCAT ESS environment want (1 to 1 with HOSTNAME_ORIG)
# HOSTNAME_NEW = The new hostname (1 to 1 with the HOSTNAME_ORIG)
# IP_ADDR_NEW = The new IPs you want (1 to 1 with HOSTNAME_NEW/ORIG)
# NETMASK = The new netmask associated with the IPs
# DOMAIN = The new domain associated with the IPs

HOSTNAME_ORIG=(ems1 gssio1 gssio2)
IP_ADDR_ORIG=(192.168.45.20 192.168.45.21 192.168.45.22)
HOSTNAME_NEW=(modems1 modgssio1 modgssio2)
IP_ADDR_NEW=(192.168.45.40 192.168.45.41 192.168.45.42)
NETMASK="255.255.255.0"
DOMAIN="gpfs.net"
```

4. Run **gsschenv** to modify your environment.

```
cd /opt/ibm/gss/tools/bin; ./gsschenv --modify /opt/ibm/gss/tools/conf/gsschenv.cfg --reboot
```

5. Run **systemctl reboot** to reboot the management server node.
6. After the environment is updated, a default `/etc/hosts` file is created on EMS. If you have the high-speed host name and IPs, add them to this file. After updating, copy `/etc/hosts` to all the I/O nodes.

```
xdcp gss_ppc64 /etc/hosts /etc/hosts
```

7. Run the **gssprecheck** script in the full install mode and address any errors.

```
/opt/ibm/gss/tools/samples/gssprecheck -N ems1 --install --file /var/tmp/gssdeploy.cfg
```

Note: First verify that `/var/tmp/gssdeploy.cfg` exists and that it has the default entries from manufacturing. If it does not exist, you need to create it.

8. Compare the installed ESS version and levels.
 Compare that the ESS version installed on each node matches the level downloaded and used from FixCentral. Also, compare that each core component (non-IBM Spectrum Scale configuration) is up-to-date.

```
xdsh ems1,gss_ppc64 "/opt/ibm/gss/tools/bin/gssinstallcheck -N localhost --get-version" | xcoll -n
```

Compare against the version checked at the beginning of the installation flow.

- **If the version matches with the version that was saved in step 1 of “Extract the Elastic Storage Server software ” on page 9**, skip to “Check the system hardware.”
- **If the version does not match with the version that was saved in step 1 of “Extract the Elastic Storage Server software ” on page 9**, run the following command sequence in the specified order:
 - a. Step 6 to step 10 of “Update the management server node” on page 21.
 - b. Step 5 to Step 10 of “Update the I/O server nodes” on page 23.

Note: Replace *CurrentIoServer* with *gss_ppc64* in these steps.

Check the system hardware

Ensure that this step is fully completed before beginning this procedure.

Power up the storage enclosures and then wait for at least 10 minutes from power on for discovery to complete before moving on to the next step. Here is the list of key log files that should be reviewed for possible problem resolution during deployment.

- By default, the `/var/log/message` logs from all I/O server nodes are directed to the xCAT computes log on the EMS (`/var/log/xcat/computes.log`).
- The **gssdeploy** log is located at `/var/log/gss`
- The xCAT log is located at `/var/log/xcat`
- Console outputs from the I/O server node during deployment are located at `/var/log/consol`es on the EMS.

1. Once rebooted, verify the installation by running **gssinstallcheck**:

```
gssinstallcheck -G ems1,gss_ppc64
```

By default, **gssinstallcheck** runs on all nodes sequentially. You can run **gssinstallcheck** in parallel from the management server node as follows.

```
# xdsh gss_ppc64 "/opt/ibm/gss/tools/bin/gssinstallcheck -N localhost" | xcoll -n
```

For more information, see Appendix M, “Running gssinstallcheck in parallel,” on page 85.

Check for any error with the following:

- a. Installed packages
- b. Linux kernel release
- c. OFED level

- | d. IPR SAS FW
 - | e. IPR SAS queue depth
 - | f. System firmware
 - | g. System profile setting
 - | h. Host adapter driver
2. Run **gssstoragequickcheck**:

```
gssstoragequickcheck -G gss_ppc64
```

3. Run **gssfindmissingdisks**:

```
gssfindmissingdisks -G gss_ppc64
```

If **gssfindmissingdisks** displays an error, run **mmgetpdisktopology** and save the output. Run **topsummary** using the saved output on each I/O server node to obtain more information about the error:

```
mmgetpdisktopology > /var/tmp/NODE_top.out
topsummary /var/tmp/NODE_top.out
```

4. Run **gsscheckdisks**:

```
GSENV=INSTALL gsscheckdisks -G gss_ppc64 --encl all --iotest a --write-enable
```

Attention: When run with **--iotest w** (write) or **--iotest a** (all), **gsscheckdisks** will perform write I/O to the disks attached through the JBOD. This will overwrite the disks and will result in the loss of any configuration or user data stored on the attached disks. **gsscheckdisks** should be run only during the installation of a building block to validate that read and write operations can be performed to the attached drives without any error. The **GSENV** environment variable must be set to **INSTALL** to indicate that **gsscheckdisks** is being run during installation.

5. Check for any hardware serviceable events and address them as needed. To view the serviceable events, issue the following command:

```
gssinstallcheck -N ems1,gss_ppc64 --srv-events
```

If any serviceable events are displayed, you can obtain more information by using the **--platform-events EVENTLIST** flag.

Set up the high-speed network

Customer networking requirements are site-specific. The use of bonding to increase fault-tolerance and performance is advised but guidelines for doing this have not been provided in this document. Consult with your local network administrator before proceeding further. Before creating network bonds, carefully read Appendix C, “ESS networking considerations,” on page 51.

- To set up bond over IB, run the following command.

```
gssgennetworks -G ems,gss_ppc64 --create --ipoib --suffix=-hs --mtu 4092
```

In this example, MTU is set to 4092. The default MTU is 2048 (2K) and the **gssgennetworks** command supports 2048 (2K) and 4092 (4K) MTU. Consult your network administrator for the proper MTU setting.

- To set up bond over Ethernet, run the following command.

```
gssgennetworks -N ems1,gss_ppc64 --suffix=-hs --create-bond
```

Note: For information on Infiniband issue with multiple fabrics, see *Infiniband with multiple fabric* in “Customer networking considerations” on page 53.

Create the cluster, recovery groups, and file system

1. Create the GPFS cluster:

```
gssgencluster -C test01 -G gss_ppc64 --suffix=-hs --accept-license
```

In this example, test01 is used as the cluster name and -hs is used as the suffix of the host name.

2. Verify healthy network connectivity:

```
xdsh gss_ppc64 /usr/lpp/mmfs/bin/mmnetverify
```

3. Create the recovery groups:

```
gssgenclusterrgs -G gss_ppc64 --suffix=-hs
```

4. Create the vdisks, NSDs, and file system:

```
gssgenvdisks --create-vdisk --create-nsds --create-filesystem --contact-node gssio1
```

Note: gssgenvdisk, by default, creates data vdisk with 8+2p RAID code and 16 MB block size, and metadata vdisk with 3WayReplication and 1 MB block size. These default values can be changed to suitable values for the customer environment.

5. Add the management server node to the cluster:

```
gssaddnode -N ems1 --cluster-node gssio1 --suffix=-hs --accept-license --no-fw-update
```

In this example, the management server hostname is ems1 with a suffix of -hs (ems1-hs) in the high-speed network. The **--no-fw-update** option is used because the management server node does not contain a SAS adapter or attached drives.

Check the installed software and system health

1. Run **gssinstallcheck** in parallel from the management server node.

```
# xdsh ems1,gss_ppc64 "/opt/ibm/gss/tools/bin/gssinstallcheck -N localhost" | xcoll -n
```

By default, **gssinstallcheck** runs on all nodes sequentially. For more information, see Appendix M, "Running gssinstallcheck in parallel," on page 85.

Note: When **gssinstallcheck** is run in parallel, you might get an error for the system firmware.

2. Shut down GPFS in all nodes and reboot all nodes.

- a. Shut down GPFS all nodes:

```
mmshutdown -a
```

- b. Reboot all server nodes:

```
xdsh gss_ppc64 "systemctl reboot"
```

- c. Reboot the management server node:

```
systemctl reboot
```

3. After reboots, run the following command (**Not applicable for PPC64LE**):

```
gssinstallcheck -G gss_ppc64 --phy-mapping
```

Ensure that the phy mapping check is OK.

4. Restart GPFS in all nodes and wait for all nodes to become active:

```
mmstartup -a
```

5. Mount the filesystem and perform a stress test. For example, run:

```
mmmount gpfs0 -a  
gssstress /gpfs/gpfs0 gssio1 gssio2
```

In this example, **gssstress** is invoked on the management server node. It is run on I/O server nodes gssio1 and gssio2 with /gpfs/gpfs0 as the target path. By default gssstress runs for 20 iterations and can be adjusted using the -i option (type **gssstress** and press Enter to see the available options). During the I/O stress test, check for network error by running from another console:

```
gssinstallcheck -N ems1,gss_ppc64 --net-errors
```

6. Perform a health check. Run:

```
gnrhealthcheck  
/usr/lpp/mmfs/bin/mmhealth node show -N all --verbose
```

Address any issues that are identified.

7. Check for any open hardware serviceable events and address them as needed. The serviceable events can be viewed as follows:

```
gssinstallcheck -N ems1,gss_ppc64 --srv-events
```

If any serviceable events are displayed, you can obtain more information by using the `--platform-events EVENTLIST` flag.

Note:

- On PPC64BE systems, investigate, manage, and close serviceable events from HMC.
- On PPC64LE systems, ASMI can be used to investigate issues.

8. Verify that NTP is set up and enabled. For more information, see Appendix N, “NTP setup,” on page 87.

Install the ESS GUI

Important:

- Complete all of the following steps carefully including the steps for configuring `mmperfmon` and restricting certain sensors to the management server node (EMS) only.
 - It is recommended to delay the GUI setup, if protocol nodes will be immediately deployed. Once the ESS and protocol nodes are deployed, run the wizard to properly discover and slot the nodes into the rack.
1. Generate performance collector on the management server node by running the following command. The management server node must be part of the ESS cluster and the node name must be the node name used in the cluster (e.g., ems1-hs).

```
mmperfmon config generate --collectors ems1-hs
```

2. Set up the nodes in the `ems nodeclass` and `gss_ppc64 nodeclass` for performance monitoring by running the following command.

```
mmchnode --perfmon -N ems,gss_ppc64
```

3. Start the performance monitoring sensors by running the following command.

```
xdsh ems1,gss_ppc64 "systemctl start pmsensors"
```

4. Capacity and fileset quota monitoring is not enabled in the GUI by default. You must correctly update the values and restrict collection to the management server node only.

- a. To modify the GPFS Disk Capacity collection interval, run the following command:

```
mmperfmon config update GPFSDiskCap.restrict=EMSNodeName  
GPFSDiskCap.period=PeriodInSeconds
```

The recommended period is 86400 so that the collection is done once per day.

- b. To restrict GPFS Fileset Quota to run on the management server node only, run the following command:

```
mmperfmon config update GPFSFilesetQuota.period=600 GPFSFilesetQuota.restrict=EMSNodeName
```

Here the *EMSNodeName* must be the name shown in the `mm1scluster` output.

Note: To enable quota, the filesystem quota checking must be enabled. Refer `mmchfs -Q` and `mmcheckquota` commands in the *IBM Spectrum Scale: Command and Programming Reference*.

5. Verify that the values are set correctly in the performance monitoring configuration by running the **mmpfmon config show** command on the management server node. Make sure that `GPFSDiskCap.period` is properly set, and `GPFFilesetQuota` and `GPFSDiskCap` are both restricted to the management server node only.

Note: If you are moving from manual configuration to auto configuration then all sensors are set to default. Make the necessary changes using the **mmpfmon** command to customize your environment accordingly. For information on how to configure various sensors using **mmpfmon**, see Manually installing IBM Spectrum Scale GUI.

6. Start the performance collector on the management server node:

```
systemctl start pmcollector
```

7. Enable and start `gpfsgui`:

```
systemctl enable gpfsgui.service
systemctl start gpfsgui
```

Note: If your ESS system came with 5148-22L protocol nodes, wait until the protocol nodes are installed before setting up the GUI.

8. To launch the ESS GUI in a browser, go to: `https://EssGuiNode` where `EssGuiNode` is the host name or IP address of the management server node for GUI access. To log in, type `admin` in the User Name field and your password in the Password field on the login page. The default password for `admin` is `admin001`. Walk through each panel and complete the GUI Setup Wizard.

This completes the installation task of the ESS system. For information on action items to be done after installation, see “Post-installation or post-upgrade checklist” on page 6.

Optional: Plug-N-Play

Using Plug-N-Play with **gssutils** is an option. For more information, see Chapter 4, “**gssutils** - ESS Installation and Deployment Toolkit,” on page 29.

The goal of Plug-N-Play is to allow customers to build a cluster, file system and begin sampling the GUI as soon as possible. The stated goal is for this to be achieved in under an hour after lab-based services (LBS) starts working on the system. Manufacturing now ships EMS with xCAT preconfigured with default settings.

Prerequisites

- Unpacking and basic power connectivity are completed.
- FSP and xCAT networks are set up in documented ports and they are connected to proper VLANs.
- SSRs have done validation checks using **gssutils** to ensure correct disk placement, cabling, networking, and server health.
- Access to EMS is available over SSH for LBS.

Basic assumptions:

- EMS has xCAT connection in T3 (1Gb card).
- All nodes have FSP connections in the HMC 1 port.
- On PPC64BE, HMC is properly configured with connections to the FSP and xCAT networks.
- On PPC64LE, EMS has an extra FSP connection in the T2 port (1Gb card).
- All standard VLANS (xCAT, FSP) are set up properly.
- The high-speed network may or may not be ready at the time of Plug-N-Play. For the purposes of this demonstration, the cluster is set up over the 1Gb network.

Note: The code level shipped from manufacturing might be lower than the code level currently available on FixCentral. This can be ignored in the context of Plug-N-Play. Use the version that is available on the system.

Overview

The primary objective is to build a very generic environment to allow the customer to preview their working Elastic Storage Server (ESS) system as fast as possible with the assumption that the final customizations are coming later. This gives the customers an opportunity to see their storage subsystem working right away. They start to get familiar with the installation process, the available file system space, start deciding on file system, and block sizes, and become familiar with the GUI.

Some basic health checks are also run in this mode that give LBS confidence that the actual installation will go smoothly:

- Default manufacturing host name, IPs, user IDs, passwords
- Networking over the 1Gb (provisioning) only.
- Basic hardware checking:
 - **gssstoragequickcheck**
 - **gssfindmissingdisks**
 - **gsscheckdisks**
- Basic file system creation (use of entire space, single pool, 16M block size, 8+2p RAID code)
- GUI and performance monitoring setup

Work flow

1. System delivered to customer site. SSR arrives and performs basic unpacking and cabling. All nodes are powered on to the operating system. SSR does a full hardware check using **gssutils**. They replace any bad components prior to LBS arrival.
2. LBS logs in to EMS through SSH.
3. Run the **gssprecheck** script in the full install mode and address any errors.
`/opt/ibm/gss/tools/samples/gssprecheck -N ems1 --install --file /var/tmp/gssdeploy.cfg`

Note: First verify that `/var/tmp/gssdeploy.cfg` exists and it has the default entries from manufacturing. If it does not exist, you need to create it.

4. Proceed to running the standard set of ESS verification checks.
 - **gssstoragequickcheck**
 - **gssfindmissingdisks**
 - **gsscheckdisks**

For more information, see the usage in “Check the system hardware” on page 11.

5. Create your cluster using **gssgencluster**.
`gssgencluster -C test01 -G gss_ppc64 --accept-license --no-fw-update`

6. Create your recovery groups using **gssgenclusterrgs**.

```
gssgenclusterrgs -G gss_ppc64
```

7. Create your file system using **gssgenvdisks**.

```
gssgenvdisks --create-vdisk --create-nsds --create-filesystem --contact-node gssi01
```

8. Add EMS using **gssaddnode**.

```
gssaddnode -N ems1 --cluster-node gssi01 --accept-license --no-fw-update
```

9. Set up the performance monitoring collector and sensors. For more information, see this section.

Note: Ensure that you use the low-speed (1Gb) host names.

- | 10. Start the GUI on the EMS and walk through the setup wizard.
- | `systemctl start gpfsGUI`

Conclusion

At this point, the customer must be able to do several tasks with their new ESS system. At a minimum, they should be able to mount the file system, view free space, and use the GUI. This mode shows how fast an ESS system can be brought up and used at a customer site.

Post Plug-N-Play cleanup steps

| Use these steps to clean up an ESS system if Plug-N-Play was used before doing full installation.

- | 1. Stop the GUI on EMS using `systemctl stop gpfsGUI`.
- | 2. Wipe the GUI database clean.
 - | `su -l postgres -c 'psql -d postgres -c "drop schema fssc cascade"'`
- | 3. Unmount the file systems.
 - | `mmumount all -a`
- | 4. SSH to one of the I/O nodes and delete the data and metadata vdisks.
 - | `/opt/ibm/gss/tools/samples/gssdelvdisks`
- | 5. Delete the log vdisks using `mmdelvdisk`.
 - | You can query the log vdisks with `mm1svdisk`.

| **Note:** Consult the man page for this command for usage help.

- | 6. Delete the recovery groups using `mmdelrecoverygroup`.
 - | You can query the recovery groups using `mmdelrecoverygroup`.

| **Note:** Consult the man page for this command for usage help.

- | 7. Shut down GPFS.
 - | `mmshutdown -a`
- | 8. Delete the cluster.
 - | `mmdelnode -a`
- | 9. Break the network bonds on each node.
 - | `cd /etc/sysconfig/network-scripts ; rm -f *bond*`
 - | `nmcli c reload`

Chapter 3. Upgrading Elastic Storage Server

These are the legacy steps required to complete a new ESS upgrade. It is highly recommended that these steps be issued from within **gssutils**.

Note:

- Protocol nodes must be upgraded after the upgrade on ESS nodes is completed.
- Offline upgrades from any prior ESS version are supported. An offline upgrade involves taking GPFS down cluster wide before doing the upgrade. Carefully follow the documented steps and look for actions specific to offline upgrades.
- Offline upgrades can be done to multiple building blocks at once.
- Power 8 firmware is not automatically updated. In the following procedure, the best time to upgrade the Power 8 firmware is advised. For more information, see Appendix G, “Updating the system firmware,” on page 65.
- The HMC (PPC64BE) can be upgraded at any time without disrupting a production system.
- Package names mentioned in the following steps are examples and they might not match with the package names that you are upgrading to. This is also applicable to the architecture mentioned in the example package names. The default architecture is PPC64LE.
- If possible, perform the upgrade while the system is offline. This is faster and it is less prone to issues.

During the upgrade process if a step fails, it must be addressed before moving to the next step. Follow these steps to perform an upgrade of the ESS system.

Note: For considerations and instructions to upgrade a cluster that contains ESS and protocol nodes, see *Upgrading a cluster containing ESS and protocol nodes in Elastic Storage Server: Protocol Nodes Quick Deployment Guide*.

Prerequisites

Before you begin the upgrade procedure, do the following:

- Obtain the ESS tarball, kernel, systemd, RHEL ISO, and network manager packages for the architecture being used.
- Archive the current contents of /home/deploy and move the 5.3.2.1 packages there.
- Make sure that the RHEL ISO is moved to /opt/ibm/gss/iso or the location specified in the gssdeploy.cfg file).
- Disable the subscription manager and any external repositories by issuing the following commands on each node that you want to upgrade:

```
subscription-manager config --rhsm.manage_repos=0
yum clean all
```
- All health checks must be clean before attempting to upgrade a system. This applies to offline and online upgrades.
- Ensure that the HMC is upgraded to the recommended level. For more information, see Appendix H, “Upgrading the Hardware Management Console (HMC),” on page 67.
- Understand the implications of upgrading the release level to LATEST and upgrading the file system format version. After you complete the upgrade to the latest code level, you cannot revert to the previous code level. For more information, see Completing the migration to a new level of IBM Spectrum Scale.
- Ensure that you have reviewed and completed the pre-installation and pre-upgrade checklist. For more information, see “Pre-installation or pre-upgrade checklist” on page 5.

- l • If you do not have an active cluster, ignore any commands related to starting and stopping IBM Spectrum Scale (**mmstartup** or **mmshutdown**). If you are performing an online upgrade, it is assumed that you have an active cluster.

Prepare the system for upgrade

1. Perform a health check by issuing the following command:

```
gnrhealthcheck
```

Address any issues that are identified.

2. Verify network connectivity and node health by issuing the following commands:

```
xdsh ems1,gss_ppc64 /usr/lpp/mmfs/bin/mmnetverify
/usr/lpp/mmfs/bin/mmhealth node show -N all
```

3. Wait for any of these commands that are performing file system maintenance tasks to complete:

```
mmadddisk
mmapplypolicy
mmcheckquota
mmdeldisk
mmfsck
mmlsnapshot
mmrestorefs
mmrestripefile
mmrestripefs
mmrpldisk
```

4. Stop the creation and deletion of snapshots using **mmcrsnapshot** and **mmdelsnapshot** during the upgrade window.
5. Run the following checks.
 - a. Run **gssinstallcheck** in parallel to verify that the current system is up to date.


```
xdsh ems1,gss_ppc64 "/opt/ibm/gss/tools/bin/gssinstallcheck -N localhost" | xcoll -n
```
 - b. Run the serviceable events check and rule out any bad hardware components.


```
gssinstallcheck -G ems1,gss_ppc64 --srv-events
```
 - c. Check for any open GUI events or tips.

Upgrading from ESS 5.1.x, 5.2.x, or 5.3.x (offline updates can be performed from any ESS version)

1. Unpack the ESS software archive (This is contained in the ESS_STD_BASEIMAGE-5.3.2.1-ppc64-Linux.tgz file.

```
tar -zxvf gss_install-5.3.2.1_ppc64le_standard_20190121T212642Z.tgz
```

2. Check the MD5 checksum:

```
md5sum -c gss_install-5.3.2.1_ppc64le_standard_20190121T212642Z.md5
```

3. Make sure the `/opt/ibm/gss/install/rhel7/<ARCH>` directory is clean:

```
/bin/sh gss_install-5.3.2.1_ppc64le_standard_20190121T212642Z --remove
```

Depending on the architecture, replace `<ARCH>` with `ppc64` or `ppc64le`.

Note: If you are upgrading to 5.3.2.1 from an earlier release, you might need to clean up the directory structure used in earlier releases. To do so, issue the following command:

```
/bin/sh gss_install-5.3.2.1_ppc64le_standard_20190121T212642Z --remove --dir /opt/ibm/gss/install
```

4. Extract the ESS packages and accept the license as follows. By default, it is extracted to the `/opt/ibm/gss/install` directory:

```
/bin/sh gss_install-5.3.2.1_ppc64le_standard_20190121T212642Z --text-only
```

5. Make the **gssdeploy** script executable:

```
chmod +x /opt/ibm/gss/install/rhel7/<arch>/samples/gssdeploy
```

6. Perform cleanup and save a backup copy of the xCAT database:

```
/opt/ibm/gss/install/rhel7/<arch>/samples/gssdeploy -c -r /var/tmp/xcatdb
```

7. Run one of the following commands depending on the architecture.

For PPC64BE:

```
cd /var/tmp ; ./gssinstall_ppc64 -u
```

For PPC64LE:

```
cd /var/tmp ; ./gssinstall_ppc64le -u
```

Note: Before copying `gssdeploy.cfg.default` to `gssdeploy.cfg`, it is advised that you save any current copy of `gssdeploy.cfg`. You can then use that as a reference.

8. Run the following command to copy the `gssdeploy.cfg.default` and customize it for your environment by editing it:

```
cp /var/tmp/gssdeploy.cfg.default /var/tmp/gssdeploy.cfg
```

Note: The directory from which you execute the **gssinstall** script determines where the `gssdeploy.cfg.default` is stored. It is recommended that you run **gssinstall** script from `/var/tmp`, but not mandatory.

Do not copy the `gssdeploy.cfg` configuration file to the `/tmp` directory because the **gssdeploy** script uses the `/tmp/gssdeploy` directory and the `/tmp` directory might get cleaned up in case of a system reboot.

9. Customize the `gssdeploy.cfg` configuration file according to your environment.

Update the management server node

1. On the management server node, stop GUI services:

```
systemctl stop gpfsGUI
```

2. Copy the RHEL 7.5 ISO file to the directory specified in the `gssdeploy.cfg` file. The default is `/opt/ibm/gss/iso`.

3. Install tools and xCAT and restore the xCAT database:

```
/var/tmp/gssdeploy -x -r /var/tmp/xcatdb
```

4. Perform precheck to detect any errors and address them before proceeding further:

```
/opt/ibm/gss/tools/samples/gssprecheck -N ems1 --upgrade --file /var/tmp/gssdeploy.cfg
```

Note: **gssprecheck** gives hints on ways to fix any discovered issues. It is recommended to review each found issue carefully though resolution of all might not be mandatory.

5. Shut down GPFS using one of the following steps depending on whether online or offline upgrade is being performed.

- **Online upgrade only**

- a. Review the quorum information prior to shutting down GPFS:

```
mmgetstate -a
```

- b. Shut down GPFS.

```
mmshutdown
```

- **Offline upgrade only**

Shut down GPFS on the entire cluster.

Note: Ensure that this is the step that you want to perform.

```
mmshutdown -a
```

6. Set up the kernel, systemd, and Network Manager errata repositories. For example, use the following command on PPC64BE systems:

```
| /var/tmp/gssdeploy -k /home/deploy/kernel-RHSA-2018-3843-BE-ESS5321.tar.gz -p \  
| /home/deploy/systemd-532-RHBA-2018-2753-BE.tar.gz,/home/deploy/netmgr-RHSA-2018-3665-BE-ESS5321.tar.gz --silent
```

Note: This command extracts the supplied tar zip files and builds the associated repository.

- -k option: Set up the kernel repository
- -p option: Set up the patch repository (For example: systemd, network manager). One or more patches might be specified at the same time separated by comma.
- Directory structure:

Kernel repository

```
/install/gss/otherpkgs/rhels7/<arch>/kernel
```

Patch repository

```
/install/gss/otherpkgs/rhels7/<arch>/patch
```

Important: Make sure that all RPMs in the /install directory including the extracted files in the kernel directory (/install/gss/otherpkgs/rhels7/<arch>/kernel), the patch directory (/install/gss/otherpkgs/rhels7/<arch>/patch), and xCAT RPMs, etc. have the correct read permission for user, group, and others (chmod 644 files). For example:

```
/install/gss/otherpkgs/rhels7/<arch>/kernel  
-rw-r--r-- 1 nobody nobody 44023568 Jan 11 10:17 kernel-3.10.0-862.25.3.e17.ppc64.rpm  
/install/gss/otherpkgs/rhels7/<arch>/patch  
-rw-r--r-- 1 nobody nobody 5412240 Oct 05 12:02 systemd-219-57.e17_5.3.ppc64.rpm  
-rw-r--r-- 1 nobody nobody 1899284 Jan 11 10:49 NetworkManager-1.12.0-8.e17_6.ppc64.rpm
```

Wrong file permission will lead to node deployment failure.

7. Update the management server node:

```
updatenode ems1 -P gss_updatenode
```

Use **systemctl reboot** to reboot the management server node and complete this step again as follows:

```
updatenode ems1 -P gss_updatenode
```

This additional step rebuilds OFED for the new kernel and builds GPFS Portability Layer (GPL) for IBM Spectrum Scale, if required.

Note: You can use the -V option with the **updatenode** command for a more verbose output on the screen for a better understanding of failures, if any.

8. Update OFED on the management server node:

```
updatenode ems1 -P gss_ofed
```

9. Update IP RAID Adapter firmware on the management server node:

```
updatenode ems1 -P gss_ipraid
```

10. Use **systemctl reboot** to reboot the management server node.

11. Perform the following steps to upgrade IBM Spectrum Scale RAID configuration parameters.

```
| /opt/ibm/gss/tools/samples/gssupg532.sh -b ems1-hs,gss_ppc64  
| /opt/ibm/gss/tools/samples/gssupg532.sh -c
```

12. Update the system firmware.

Before starting up GPFS, it is a good time to upgrade the system firmware. For more information, see Appendix G, "Updating the system firmware," on page 65.

- **Online upgrade only**

Update the EMS node only.

- | • **Offline upgrade only**
- | System firmware of all nodes can be updated at this time in case of an offline upgrade.
- | 13. Start up GPFS.
- | • **Online upgrade only**
- | Run this command on the management server node.
- | `mmstartup`
- | • **Offline upgrade only**
- | Skip this step
- | 14. Verify that GPFS is in the active state before upgrading the I/O server nodes.
- | • **Online upgrade only**
- | Run this command.
- | `mmgetstate`
- | • **Offline upgrade only**
- | Skip this step
- | Do not proceed if the system is not active.
- | 15. Ensure that the management server node is fully updated and active.
- | `gssinstallcheck -N ems1`
- | **Note: Offline upgrade only** Some checks cannot be performed at this time because the cluster is offline. It is good to verify that the base upgrade has completed on this node (OFED, IPR, Kernel, etc.).

| Update the I/O server nodes

| **Note:** This procedure is for online upgrade thus many of the steps can be skipped or done in parallel in case of offline upgrades. Look for **Offline upgrade only** tags in each step for guidance.

| Repeat the following steps for each I/O server node, one node at a time.

- | 1. Before shutting down GPFS on any I/O server node, run precheck from the management server node:

| `/opt/ibm/gss/tools/samples/gssprecheck -N IO_NODE --upgrade --file /var/tmp/gssdeploy.cfg`

| **Note: gssprecheck** gives hints on ways to fix any discovered issues. It is recommended to review each found issue carefully though resolution of all might not be mandatory.

- | 2. Move the cluster and file system manager role to another node if the current node is a cluster manager or file system manager.

- | • **Online upgrade only**
- | a. To find the cluster and file system managers, run:

| `mm1smgr`

- | b. To change the file system manager, run:

| `mmchmgr gpfs0 gssio2-hs`

| In this example, `gssio2-hs` is the new file system manager of file system `gpfs0`.

- | c. To change the cluster manager, run:

| `mmchmgr -c gssio2-hs`

| In this example, `gssio2-hs` is the new cluster manager.

- | • **Offline upgrade only**
- | Skip this step

3. Move the recovery group in the current I/O server node to the peer I/O server node in the same building block.
- **Online upgrade only**
 - a. To list the recovery groups, run:


```
mmisrecoverygroup
```
 - b. To list the active server, primary server, and secondary server, run:


```
mmisrecoverygroup rg_gssio1-hs -L | grep active -A2
```
 - c. To move the recovery group from the current active I/O server node (rg_gssio1-hs) to the peer I/O server node (gssio2-hs) in the same building block, run the following commands in the shown order:


```
mmchrecoverygroup rg_gssio1-hs --active gssio2-hs
```

```
mmchrecoverygroup rg_gssio1-hs --servers gssio2-hs,gssio1-hs
```
 - **Offline upgrade only**
Skip this step
4. After confirming that the recovery group has been successfully moved to the peer I/O server node, unmount all GPFS file systems if mounted, and shut down GPFS on the current I/O server node while maintaining quorum.
- **Online upgrade only**

```
mmumount all -N CurrentIoServer-hs
```

```
mmshutdown -N CurrentIoServer-hs
```
 - **Offline upgrade only**
Skip this step
5. Run **updatenode**.
- **Online upgrade only**

```
updatenode CurrentIoServer -P gss_updatenode
```
 - **Offline upgrade only**
Run in parallel.
Instead of *CurrentIoServer* use the xCAT group gss_ppc64. For example:


```
updatenode gss_ppc64 -P gss_updatenode
```
6. Reboot the I/O server node and complete this step again if you are instructed to do so in the **updatenode** output. Reboot the I/O server node as follows.
- **Online upgrade only**

```
xdsh CurrentIoServer "systemctl reboot"
```
 - **Offline upgrade only**
Run in parallel.
Instead of *CurrentIoServer* use the xCAT group gss_ppc64. For example:


```
xdsh gss_ppc64 "systemctl reboot"
```
7. Run **updatenode** again (if instructed to do so).
- **Online upgrade only**

```
updatenode CurrentIoServer -P gss_updatenode
```
 - **Offline upgrade only**
Run in parallel.
Instead of *CurrentIoServer* use the xCAT group gss_ppc64. For example:


```
updatenode gss_ppc64 -P gss_updatenode
```
8. Update OFED.
- **Online upgrade only**

```
updatenode CurrentIoServer -P gss_ofed
```

- | • **Offline upgrade only**
- | Run in parallel.
- | Instead of *CurrentIoServer* use the xCAT group *gss_ppc64*. For example:
- | `updatenode gss_ppc64 -P gss_ofed`
- | 9. Update IP RAID FW in the I/O Server node that is being upgraded.
- | • **Online upgrade only**
- | `updatenode CurrentIoServer -P gss_ipraid`
- | • **Offline upgrade only**
- | Run in parallel.
- | Instead of *CurrentIoServer* use the xCAT group *gss_ppc64*. For example:
- | `updatenode gss_ppc64 -P gss_ipraid`
- | 10. Reboot the I/O server node.
- | • **Online upgrade only**
- | `xdsh CurrentIoServer "systemctl reboot"`
- | • **Offline upgrade only**
- | Run in parallel.
- | Instead of *CurrentIoServer* use the xCAT group *gss_ppc64*. For example:
- | `xdsh gss_ppc64 "systemctl reboot"`
- | 11. Update the SAS host adapter firmware on *CurrentIoServer*.
- | • **Online upgrade only**
- | `CurrentIoServer$ mmchfirmware --type host-adapter`
- | Here *CurrentIoServer* is an I/O server node and the command is run on the I/O server node.
- | • **Offline upgrade only**
- | Run this command on both I/O server nodes.
- | 12. Update the node configuration.
- | • **Online upgrade only**
- | `/opt/ibm/gss/tools/samples/gssupg532.sh -s CurrentIoServer-hs`
- | This command is run from the EMS node.
- | • **Offline upgrade only**
- | Run the following command.
- | `/opt/ibm/gss/tools/samples/gssupg532.sh -s`
- | 13. On PPC64BE systems, run phy check and ensure that the phy mapping is OK.
- | • **Online upgrade only**
- | `gssinstallcheck -N CurrentIoServer --phy-mapping`
- | • **Offline upgrade only**
- | Run the following command.
- | `gssinstallcheck -G gss_ppc64 --phy-mapping`
- | 14. **Online upgrade only** Perform the system firmware update. For more information, see Appendix G, "Updating the system firmware," on page 65.
- | 15. Start GPFS on the I/O server node.
- | • **Online upgrade only**
- | `mmstartup -N CurrentIoServer-hs`
- | Once the GPFS daemon is successfully started, move back the recovery group that was moved to the peer I/O server node of the same building block in Step 3c above. Move back the cluster manager and the file system manager if required that was moved to the other nodes in step 2.
- | • **Offline upgrade only**

- | Run the following command to start GPFS cluster wide.
- | `mmstartup -a`
- | Moving recovery groups, or cluster or file system manager is not needed at this point.
- | 16. Wait until the I/O server can be seen active from the management server node, using the following command.
- | • **Online upgrade only**
- | `mmgetstate -N CurrentIoServer-hs`
- | The management server must be already running for issuing this command.
- | • **Offline upgrade only**
- | Run the following command to check cluster state in parallel.
- | `mmgetstate -a`
- | 17. Run **gssinstallcheck**.
- | • **Online upgrade only**
- | Run this command from the management server node.
- | `gssinstallcheck -N IO_NODE`
- | • **Offline upgrade only**
- | Run the following command to perform final installcheck in parallel.
- | `xdsh ems1,gss_ppc64 "/opt/ibm/gss/tools/bin/gssinstallcheck -N localhost" | xcoll -n`
- | 18. **Online upgrade only** Repeat preceding steps for the peer I/O server node of the same building block.
- | **Offline upgrade only** Skip this step.
- | 19. **Online upgrade only** Repeat all steps in this section for each additional building block.
- | **Offline upgrade only** Skip this step.

Update the enclosure and drive firmware

1. To update the storage enclosure firmware, run one of the following commands from one I/O server node of each building block.

- When upgrade is being performed concurrently:

```
CurrentIoServer$ mmchfirmware --type storage-enclosure
```

- When upgrade is being performed non-concurrently, all attached enclosures can be upgraded in parallel.

```
mmchfirmware --type storage-enclosure -N gss_ppc64
```

Note: The IBM Spectrum Scale daemon must be down on all nodes of the node class `gss_ppc64` for parallel upgrade.

2. To update the drive firmware, run the following command from **each** I/O Server node of each building block:

```
CurrentIoServer$ mmchfirmware --type drive
```

The drive update can take some time to complete. You can update the drives more quickly by taking the system offline (shutting down IBM Spectrum Scale) and using the **--fast-offline** option.

Check the installed software and system health

1. Perform a health check:

```
gnrhealthcheck
/usr/lpp/mmfs/bin/mmhealth node show -N all --verbose
```

2. Check for any hardware serviceable events and address them as needed. To view the serviceable events, issue the following command:

```
gssinstallcheck -N ems1,gss_ppc64 --srv-events
```

If any serviceable events are displayed, you can obtain more information by using the `--platform-events EVENTLIST` flag.

Note:

- On PPC64BE systems, investigate, manage, and close serviceable events from HMC.
- On PPC64LE systems, ASMI can be used to investigate issues.
- During initial deployment of the nodes, SRC BA15D001 may be logged as serviceable event by Partition Firmware. This is normal and should be cleared after the initial deployment. For more information, see Appendix A, “Known issues,” on page 39.

Note: Some of these steps might fail if they are already implemented in previous versions of ESS. If you see any failures indicating `mmperfmon` has already been configured, ignore these failure messages and continue with the remaining steps.

Upgrading GUI

Perform the following steps to upgrade the GUI:

Note: Some of these steps might fail if the GUI is already set up. However, it is important to rerun the upgrade steps using the latest changes.

1. Generate performance collector on the management server node by running the following command. The management server node must be part of the ESS cluster and the node name must be the node name used in the cluster (e.g., `ems1-hs`).

```
mmperfmon config generate --collectors ems1-hs
```

2. Set up the nodes in the `ems_nodeclass` and `gss_ppc64_nodeclass` for performance monitoring by running the following command.

```
mmchnode --perfmon -N ems,gss_ppc64
```

3. Start the performance monitoring sensors by running the following command.

```
xdsh ems1,gss_ppc64 "systemctl start pmsensors"
```

4. Capacity and fileset quota monitoring is not enabled in the GUI by default. You must correctly update the values and restrict collection to the management server node only.

- a. To modify the GPFS Disk Capacity collection interval, run the following command:

```
mmperfmon config update GPFSDiskCap.restrict=EMSNodeName  
GPFSDiskCap.period=PeriodInSeconds
```

The recommended period is 86400 so that the collection is done once per day.

- b. To restrict GPFS Fileset Quota to run on the management server node only, run the following command:

```
mmperfmon config update GPFSFilesetQuota.period=600 GPFSFilesetQuota.restrict=EMSNodeName
```

Here the *EMSNodeName* must be the name shown in the `mm1scluster` output.

Note: To enable quota, the filesystem quota checking must be enabled. Refer `mmchfs -Q` and `mmcheckquota` commands in the *IBM Spectrum Scale: Command and Programming Reference*.

5. Verify that the values are set correctly in the performance monitoring configuration by running the `mmperfmon config show` command on the management server node. Make sure that `GPFSDiskCap.period` is properly set, and `GPFSFilesetQuota` and `GPFSDiskCap` are both restricted to the management server node only.

Note: If you are moving from manual configuration to auto configuration then all sensors are set to default. Make the necessary changes using the `mmperfmon` command to customize your environment accordingly. For information on how to configure various sensors using `mmperfmon`, see Manually installing IBM Spectrum Scale GUI.

6. Start the performance collector on the management server node:
`systemctl start pmcollector`
7. Enable and start gpfs GUI:
`systemctl enable gpfs.service`
`systemctl start gpfs`
8. To launch the ESS GUI in a browser, go to: `https://EssGuiNode` where `EssGuiNode` is the hostname or IP address of the management server node for GUI access. To log in, type `admin` in the User Name field and your password in the Password field on the login page. The default password for `admin` is `admin001`. Walk through each panel and complete the GUI Setup Wizard.

After the GUI is up and running, do the following:

1. Enable the subscription manager by issuing the following commands on the upgraded nodes:
`subscription-manager config --rhsm.manage_repos=1`
`yum clean all`

This completes the upgrade task of the ESS system. For information on action items to be done after installation, see “Post-installation or post-upgrade checklist” on page 6.

Chapter 4. `gssutils` - ESS Installation and Deployment Toolkit

LBS team members or customers use ESS command line tools to perform SSR prechecks, installation, deployment, or upgrade tasks on the ESS server by following the ESS: Quick Deployment Guide (QDG).

The ESS Installation and Deployment Toolkit (`gssutils`) is designed to facilitate these tasks by using a menu driven tool.

The ESS Installation and Deployment Toolkit (`gssutils`) is designed to eliminate these problems.

Note: `gssutils` does not support protocol node deployments or upgrades at this time.

- “`gssutils` introduction”
- “`gssutils` prerequisites” on page 31
- “`gssutils` defaults” on page 32
- “`gssutils` usage” on page 32
- “`gssutils` customization” on page 34
- “`gssutils` restrictions” on page 36

`gssutils` introduction

`gssutils` is a text-based, menu-driven utility that facilitates SSR prechecks, installation, deployment, and upgrade tasks. It can be started with or without any optional arguments. It provides a set of task menus that are related to installation and deployment activities. When a task is selected from the menu, a command is issued to the system for that task. `gssutils` requires a minimum of 80 x 24 character window to operate.

When running `gssutils`, one can pass the EMS server, I/O server nodes, prefix, and suffix while initiating the tool and all the commands are automatically adjusted to use the provided value while running the actual commands. For example, the SSR check, Plug N Play mode, Fusion mode, installation, deployment, or upgrade commands are customized to run for user provided EMS server, I/O server nodes, prefix, and suffix.

`gssutils` can be customized depending on environment and it can be run by just pressing the Enter key without doing any command line changes such as changing EMS host name, I/O server node names, prefix or suffix for CLI and there is no need to remember the CLI options.

`gssutils` also guides users to run the command in the correct order as it is a menu-driven, step wise user interface. This allows the user to not having to remember the next command in the sequence.

`gssutils` also has a built-in help system available for each menu. All menus in `gssutils` have an option called **Help**. Help menu brings up a Linux manual page which explains the use of the menu and also explains the commands that are invoked from that menu. Users must read the help from each menu before running any tasks.

The `gssutils` console is as follows.

```
ESS INSTALLATION AND DEPLOYMENT TOOLKIT

1. Help
2. SSR Tools >
3. Plug n Play and Hybrid >
4. Install >
5. Upgrade >
6. Validation checks >
7. View/Collect service data (snaps) >
8. Exit

man gssutils_panel_1
Help
```

The **gssutils** command line help is as follows:

```
# gssutils -h
usage: gssutils [-h] [-N NODE-LIST | -G NODE-GROUP] [--prefix PREFIX][--suffix SUFFIX]
[--config CONFIG_FILE] [--customize][--ems-node-name EMS_NODE]
[--io-node-one-name IO_NODE1][--io-node-two-name IO_NODE2]
```

Optional arguments:

- N *NODE-LIST***
Provides a list of nodes. If node list or group name is not provided, -N localhost is assumed.
- G *NODE-GROUP***
Provides the name of node group. Nodes in the *NODE-LIST* are members of the *NODE-GROUP*.
- prefix *PREFIX***
Provides the host name prefix. Use = between --prefix and value if the value starts with -.
- suffix *SUFFIX***
Provides the host name suffix. Use = between --suffix and value if the value starts with -.
- config *CONFIG_FILE***
Provides the configuration file for **gssutils** for a specific environment.
- customize**
Customizes the EMS host name and I/O node name and generates the **gssutils** configuration file.
This file can be used with --config to run **gssutils** specific to an environment.
- ems-node-name *EMS_NODE***
Specifies EMS host name to populate the **gssutils** configuration file.
- io-node-one-name *IO_NODE1***
Specifies I/O node 1 host name to populate the **gssutils** configuration file.
- io-node-two-name *IO_NODE2***
Specifies I/O node 2 host name to populate the **gssutils** configuration file.

-h | --help

Displays usage information about this script and exits.

This tool has six major sections:

- **SSR Tools**
- **Plug n Play and Fusion**
- **Install/Deploy**
- **Upgrade**
- **Validation checks**
- **View/Collect service data (snaps)**

The **SSR Tools** menu option can be used to validate the system after it has arrived to the customer location and LBS teams have plugged in all servers and components together and the GPFS cluster is ready to be deployed. SSR tools help you to validate the system interconnect including whether the network is correct or not, the enclosures are connected to the I/O server nodes or not, etc. It is the primary menu option that must be used by LBS or the customer before creating the GPFS cluster.

The **Plug n Play and Fusion** menu option provides two modes: **Plug N Play** and **Fusion**. The Plug N Play mode allows customers to build a cluster, file system and begin sampling the GUI as soon as possible. The stated goal is for this to be achieved in under an hour after LBS starts working on the system. Manufacturing now ships EMS with xCAT preconfigured with default settings. The Fusion mode no longer requires that Elastic Storage Server (ESS) systems be rediscovered or re-deployed at a customer site. The end goal of this mode is to greatly reduce the time and the complexity in bringing up an ESS system.

The **Install/Deploy** menu option allows you to re-deploy the EMS along with I/O server nodes in case the customer has cleaned up (re-imaged) the ESS nodes and they want to restart the EMS and I/O server nodes deployment from scratch. This menu option is a menu-driven representation of the installation section of the ESS: Quick Deployment Guide.

The **Upgrade** menu option allows you to upgrade the existing ESS server to the latest version of the ESS software. This menu item is a menu-driven representation of the upgrade section of the ESS: Quick Deployment Guide.

The **Validation checks** menu option can be used once the ESS system has been deployed and the GPFS cluster is up and running. Using this option, users can collect firmware level, run stress test on network and file system, etc. It can be used as a post deployment utility to validate the ESS server.

The **View/Collect service data** option can be used to collect the **gpfs.snap** SOS reports.

gssutils prerequisites

gssutils is a part of `gpfs.gss.tools-<version>.el7.noarch.rpm`. The `gpfs.gss.tools` RPM must be installed either manually (using `yum` or `rpm -ivh`) or using the `gssinstall_<arch> -u` command. Once the `gpfs.gss.tools` RPM is installed, users can type **gssutils** on the command prompt to use this utility, in case the ESS systems were wiped and re-imaged for fresh deployment.

If the system has been shipped from manufacturing, the ESS server is already deployed and the `gpfs.gss.tools` RPM is already installed. Users can type **gssutils** on the command prompt to use it for activities such as doing SSR checks or creating GPFS cluster as GPFS cluster needs to be created on newly shipped ESS systems.

If you want to use **gssutils** for upgrading to ESS 5.3.x, you must do these steps before you can use **gssutils** for upgrade.

1. Make the **gssdeploy** script executable:
`chmod +x /opt/ibm/gss/install/rhel7/<ARCH>/samples/gssdeploy`
2. Clean the current xCAT installation and associated configuration to remove any preexisting xCAT configuration, and then address any errors before proceeding:
`/opt/ibm/gss/install/rhel7/<ARCH>/samples/gssdeploy -c`

Note: If you are doing an upgrade, use the following command to clean up the current xCAT configuration and to save a backup copy of the xCAT database.

```
/opt/ibm/gss/install/rhel7/<ARCH>/samples/gssdeploy -c -r /var/tmp/xcatdb
```

3. Run one of the following commands depending on the architecture:

For PPC64BE:

```
cd /var/tmp ; ./gssinstall_ppc64 -u
```

For PPC64LE:

```
cd /var/tmp ; ./gssinstall_ppc64le -u
```

gssutils defaults

gssutils comes with some defaults such as the default EMS host name, I/O server node names, prefix, and suffix. By default, all **gssutils** commands assume that the default EMS host name is `ems1`, default I/O server node name group `gss_ppc64`, I/O server node names are `gssi01` and `gssi02`, and prefix and suffix are empty.

```

CHECK SYSTEM HARDWARE AND SOFTWARE

1. Help
2. Show node details
3. Check and validate various install parameters
4. Quick storage configuration check
5. Check enclosure cabling and paths to disks
6. Check disks for IO operations
7. Ping all nodes
8. Check ssh to all nodes
9. Run lsscsi from all nodes
10. Check for open serviceable events
11. Back

/opt/ibm/gss/tools/bin/gssnodedetails -N ems1,gss_ppc64
Shows miscellaneous node information.

```

gssutils usage

The **-N** and **-G** options:

gssutils has **-N** and **-G** options which can be used to replace the default node list or node group in ESS commands. The default is node list is `ems1,gss_ppc64`. If you want to run the **gssutils** with a different node list option, do it as follows.

```
$ gssutils -N ems2
```

```
                CHECK SYSTEM HARDWARE AND SOFTWARE

1. Help
2. Show node details
3. Check and validate various install parameters
4. Quick storage configuration check
5. Check enclosure cabling and paths to disks
6. Check disks for IO operations
7. Ping all nodes
8. Check ssh to all nodes
9. Run lsscsi from all nodes
10. Check for open serviceable events
11. Back

/opt/ibm/gss/tools/bin/gssnodedetails -N ems2
Shows miscellaneous node information.
```

With this command, the default node list is changed from ems1, gss_ppc64 (in default section) to ems2. The changes specified using the -N and -G options are only applicable until the **gssutils** instance is running.

Restriction with -N and -G options:

- If -N and -G options have been used with **gssutils**, none of the other **gssutils** options can be used along with them.
- Some of the ESS server deployment commands do not require -N or -G option, however they still need EMS or I/O server node names. Those commands are not a part of core ESS deployment however they are a part of other binaries in the ESS deployment toolkit. For example, the **updatenode** command belongs to xCAT however this command used to update the node and needs EMS or I/O server node name without -N and -G. For example:

```
$ updatenode ems1 -P gss_updatenode
```

```
updatenode ems1 -P gss_updatenode
Update the ESS Management Server (EMS) node and installs required rpms.
This must be run twice with a reboot in between. See QDG for details.
Press c to change nodename and options of the command.
```

In this example, ems1 is an argument to **updatenode** (xCAT command) without -N. **gssutils** -N and -G do not change the values for such command. To change these default values, user must use the **gssutils** customization option, which is described in the **gssutils customization** section.

--prefix and --suffix options:

These options can be used in conjunction with any switches with **gssutils** to provide prefix and suffix names to host name for high speed network. For example:

```
$ gssutils --prefix=hr --suffix=-10g
```

```

CHECK SYSTEM HARDWARE AND SOFTWARE

1. Help
2. Show node details
3. Check and validate various install parameters
4. Quick storage configuration check
5. Check enclosure cabling and paths to disks
6. Check disks for IO operations
7. Ping all nodes
8. Check ssh to all nodes
9. Run lsscsi from all nodes
10. Check for open serviceable events
11. Back

/opt/ibm/gss/tools/bin/gssnodedetails -N ems2 --prefix hr --suffix=-10g
Shows miscellaneous node information.
```

When user provides --prefix and --suffix option then now all deployment commands also get prefix and suffix.

gssutils customization

gssutils allows customization for specific targeted environment. As a result, you do not need to use the -N or -G or --prefix or --suffix option every time while using **gssutils**. One can generate the **gssutils** configuration file specific to their environment using the --customize option and the generate a configuration file that can be used subsequently. For example, in an environment where the EMS node is ems2 and I/O server node names are io3 and io4, and suffix -40g is to be specified, the following command needs to be issued.

```
$ gssutils --customize --config /var/tmp/env2 --ems-node-name ems2 \
--io-node-one-name io3 --io-node-two-name io4 --suffix=-40g
```

Successfully generated the customization configuration file.

This command generates a **gssutils** customization file at /var/tmp/env2. Thereafter, **gssutils** can be invoked using the customization file as follows.

```
$ gssutils --config /var/tmp/env2
```

Here, running **gssutils** with a customized configuration file replaced **gssutils** defaults to user provided defaults used while generating the configuration file.

The custom configuration file is capable of replacing the defaults of those commands also which are not a part of the ESS core deployment.

```

updatenode ems2 -P gss_updatenode
Update the ESS Management Server (EMS) node and installs required rpms. This must be run
twice with a reboot in between. See QDG for details. Press c to change nodename and options
of the command.
```

Note: Running **gssutils** with the custom configuration file is the recommended way of using **gssutils**.

The “c” button customization in **gssutils**

gssutils supports a special kind of customization called the “c” button customization. At any point of time in **gssutils**, users can press the “c” button and change the command inline. This customization allows users to modify the command inline without persisting the customization forever. For example:

\$ **gssutils**

```

CHECK SYSTEM HARDWARE AND SOFTWARE

1. Help
2. Show node details

/opt/ibm/gss/tools/bin/gssnodedetails -N ems2 hr -10g

Change command option(s) below. Press Ctrl+G to return.
Nodelist/Group:-N ems2
Prefix:hr
Suffix:-10g
cmd options:█

/opt/ibm/gss/tools/bin/gssnodedetails -N ems2 --prefix hr --suffix=-10g
Shows miscellaneous node information.
```

When user presses the “c” button, a dialog opens which allows user to do changes and commit the changes of command using Ctrl+g. Another example is as follows:

```
UPDATE EMS NODE TO LATEST FIXES

1. Help
2. Update Kernel Errata Repo (Optional)

updatenode ems1 -P gss_updatenode

Change command option(s) below. Press Ctrl+G to return.
cmd options:ems1 -P gss_updatenode

updatenode ems1 -P gss_updatenode
Update the ESS Management Server (EMS) node and installs required rpms.
This must be run twice with a reboot in between. See QDG for details.
```

gssutils restrictions

- **gssutils** must be invoked from the EMS node and not from I/O server nodes.
- Make sure that you use a console of adequate height and width to specify a large file name while extracting the tar ball. A minimum 90 x 32 console size or higher is required. If you experience an error, this task can be performed from the shell prompt also.
- **gssutils** must only be invoked from a PuTTY session or an equivalent standard Linux terminal. **gssutils** must not be invoked from an IPMI SOL connection or from an xCAT Remote Console (rcons) session. It might not work because of the different terminal type used by IPMI SOL and xCAT Remote Console connections.

Chapter 5. Adding building blocks to an existing ESS cluster

| When adding building blocks in an ESS cluster, perform the following steps. You may add more than one building blocks at a time.

Attention: Be careful when running this procedure on a production system. Failure to follow the listed procedure might result in re-deployment of a running I/O server node, loss of quorum, and/or data loss. Be especially carefully before running **gssdeploy -d** or **gsscheckdisks**. It is always a good practice to back up the xCAT database prior to starting this procedure and for the customer to back up their data, if possible.

- | 1. Power on and boot the new building blocks. Do not power up the enclosures.
2. Add the new building block to /etc/hosts file.
3. Find the new building block serial numbers.
`/var/tmp/gssdeploy -f <subnet></mask>`
4. Find rack positions that are used for GUI.
`/var/tmp/gssdeploy -i`
5. Update gssdeploy.cfg to make following changes.
 - Change DEPLOYMENT_TYPE to ADD_BB.
 - Change GSS_GROUP to something other than gss_ppc64 or ces_ppc64.
 - Add the new serial numbers and node names.
6. Run **gssdeploy**.
`/var/tmp/gssdeploy -o`
7. Run **gssprecheck** on the new node group.
`/opt/ibm/gss/tools/samples/gssprecheck -G new_group --install --file /var/tmp/gssdeploy.cfg`
- | 8. Run **gssdeploy -d**.
`/var/tmp/gssdeploy -d`
- | 9. Perform the steps in this procedure: “Set up the high-speed network ” on page 12.
- | 10. Run **gssinstallcheck** on the new group. Ignore any GPFS related settings for now.
`xdsh new_group "/opt/ibm/gss/tools/bin/gssinstallcheck -N localhost" | xcoll -n`
- | 11. Add the new nodes to the cluster.
`gssaddnode -G new_group --cluster-node gssi01 --suffix=-hs --accept-license --nodetype gss`
12. Create recovery groups.
`gssgenclusterrgs -G new_group --suffix=-hs`
13. Create vdisks and NSDs. Make sure that you match the block size and the RAID code.
`gssgenvdisks --create-vdisk --create-nsds -contact-node gssi01 --recovery-group new_recovery_group`
14. Add NSDs to an existing file system by using the **mmadddisk** command. For more information, see mmadddisk command.
15. Run restripe if needed.
`mmresripefs -b`
16. Update the performance monitoring list using the **mmchnode** command. For more information, see mmchnode command.
17. Put the new nodes back in the gss_ppc64 node group and delete the temporary group, and then comment out the GSS_GROUP line in gssdeploy.cfg.
For example, add nodes to gss_ppc64 as follows.
`chdef -t group gss_ppc64 -o gssi03,gssi04 -p`

| 18. Update call home for new building blocks.

```
| gsscallhomeconf -E ems1 -N ems1,gss_ppc64 --suffix=-hs --register=all
```

| **Note:** Now that you have additional building blocks, ensure that file system metadata replication is enabled (-m 2). The new pair of NSDs should be in a new failure group. The maximum number of failure groups is 5.

Appendix A. Known issues

This topic describes known issues for ESS.

ESS 5.3.2.1 issues

The following table describes known issues in ESS 5.3.2.1 and how to resolve these issues. Depending on which fix level you are installing, these might or might not apply to you.

Table 2. Known issues in ESS 5.3.2.1

Issue	Environment affected	Description	Resolution or action
The gssgennetworks script requires high-speed host names to be derived from I/O server (xCAT) host names using suffix, prefix, or both.	High-speed network generation Type: Install Version: All Arch: All Affected nodes: I/O server and EMS nodes	gssgennetworks requires that the target host name provided in -N or -G option are reachable to create the high-speed network on the target node. If the xCAT node name does not contain the same base name as the high-speed name you might be affected by this issue. A typical deployment scenario is: gssio1 // xCAT name gssio1-hs // high-speed An Issue scenario is: gssio1 // xCAT name foolabc-hs // high-speed name	Create entries in the /etc/hosts with node names that are reachable over the management network such that the high-speed host names can be derived from it using some combination of suffix and/or prefix. For example, if the high-speed host names are foolabc-hs, goolabc-hs: 1. Add fool1 and gool1 to the /etc/hosts using management network address (reachable) in the EMS node only. 2. Use: gssgennetworks -N fool1,gool1 - suffix abc-hs --create-bond 3. Remove the entries fool1 and gool1 from the /etc/hosts file on the EMS node once the high-speed networks are created. Example of how to fix (/etc/hosts): // Before <IP><Long Name><Short Name> 192.168.40.21 gssio1.gpfs.net gssio1 192.168.40.22 gssio2.gpfs.net gssio2 X.X.X.X foolabc-hs.gpfs.net foolabc-hs X.X.X.Y goolabc-hs.gpfs.net goolabc-hs // Fix 192.168.40.21 gssio1.gpfs.net gssio1 fool 192.168.40.22 gssio2.gpfs.net gssio2 gool X.X.X.X foolabc-hs.gpfs.net foolabc-hs X.X.X.Y goolabc-hs.gpfs.net goolabc-hs gssgennetworks -N fool1, gool1 --suffix=abc-hs --create-bond
Running gssutils over PuTTY might show horizontal lines as "qqq" and vertical lines as "xxx".	ESS Install and Deployment Toolkit Type: Install or Upgrade Version: All Arch: All Affected Nodes: EMS and I/O server nodes	PuTTY translation default Remote Character set UTF-8 might not translate horizontal line and vertical character sets correctly.	1. On the PuTTY terminal Window > Translation, change Remote character set from UTF-8 to ISO-8859-1:1998 (Latin-1, West Europe) (this should be the first option after UTF-8). 2. Open session.

Table 2. Known issues in ESS 5.3.2.1 (continued)

Issue	Environment affected	Description	Resolution or action
<p>gssinstallcheck might flag an error regarding page pool size in multi-building block situations if the physical memory sizes differ.</p>	<p>Software Validation Type: Install or Upgrade Arch: Big Endian or Little Endian Version: All Affected nodes: I/O server nodes</p>	<p>gssinstallcheck is a tool introduced in ESS 3.5, that helps validate software, firmware, and configuration settings. If adding (or installing) building blocks of a different memory footprint installcheck will flag this as an error. Best practice states that your I/O servers must all have the same memory footprint, thus pagepool value. Page pool is currently set at ~60% of physical memory per I/O server node. Example from gssinstallcheck: [ERROR] pagepool: found 142807662592 expected range 147028338278 - 179529339371</p>	<p>1. Confirm each I/O server node's individual memory footprint. From the EMS, run the following command against your I/O xCAT group: xdsh gss_ppc64 "cat/ proc/meminfo grep MemTotal" Note: This value is in KB. If the physical memory varies between servers and/or building blocks, consider adding memory and re-calculating pagepool to ensure consistency. 2. Validate the pagepool settings in IBM Spectrum Scale: mm1sconfig grep -A 1 pagepool Note: This value is in MB. If the pagepool value setting is not roughly ~60% of physical memory, then you must consider recalculating and setting an updated value. For information about how to update the pagepool value, see IBM Spectrum Scale documentation on IBM Knowledge Center.</p>
<p>Creating small file systems in the GUI (below 16G) will result in incorrect sizes</p>	<p>GUI Type: Install or Upgrade Arch: Big Endian or Little Endian Version: All Affected nodes: All</p>	<p>When creating file systems in the GUI smaller than 16GB (usually done to create CES_ROOT for protocol nodes) the size will come out larger than expected.</p>	<p>There is currently no resolution. The smallest size you might be able to create is 16GB. Experienced users might consider creating a customer <code>vdisk.stanza</code> file for specific sizes you require. You can try one of the following workarounds:</p> <ul style="list-style-type: none"> • Use three-way replication on the GUI when creating small file systems. • Use gssgenvdisks which supports the creation of small file systems especially for CES_ROOT purposes (Refer to the <code>--crcefs</code> flag).
<p>Creating file systems in the GUI might immediately result in lack of capacity data</p>	<p>GUI Type: Install or Upgrade Arch: Big Endian or Little Endian Version: All Affected nodes: All</p>	<p>When creating file systems in the GUI you might not immediately see the capacity data show up.</p>	<p>You may wait up to 24 hours for the capacity data to display or simply use the command line which should accurately show the file system size.</p>

Table 2. Known issues in ESS 5.3.2.1 (continued)

Issue	Environment affected	Description	Resolution or action
Canceling disk replacement through GUI leaves original disk in unusable state	GUI Type: Install or Upgrade Arch: Big Endian or Little Endian Version: All Affected nodes: I/O server nodes	Canceling a disk replacement can lead to an unstable system state and must not be performed. However, if you did this operation, use the provided workaround.	Do not cancel disk replacement from the GUI. However, if you did, then use the following command to recover the disk to state: mmchpdisk <RG> --pdisk <pdisk> --resume
Under Monitoring > Hardware details , you might see enclosures missing location information.	GUI Type: Install or Upgrade Arch: Big Endian or Little Endian Version: All Affected nodes: N/A	After install or upgrade to ESS 5.3.2.1, you might see missing location information for the enclosures in your system. This does not reflect the true frame U location which can be observed in the Monitoring > Hardware details panel.	The current workaround is to wait up to 24 hours for the GUI services to refresh. After this period you will see the enclosure location information fill in.
The GUI wizard might start again after completing the initial setup.	GUI Type: Install Arch: Big Endian Version: All Affected nodes: N/A	After completing the GUI wizard setup on ESS 5.3.2.1 PPC64BE, you might see the start screen again.	If you see the GUI wizard start screen a second time, type the address of the EMS into the browser and press enter. https://<ip of EMS over management network> You will then be taken to the GUI home screen.
Upon upgrades to ESS 5.3.2.1, you might notice missing groups and users in the Monitoring > Capacity GUI panel	GUI Type: Upgrade Arch: All Version: All Affected nodes: N/A	You might notice one or more missing pools or users after upgrading to ESS 5.3.2.1 in the Monitoring > Capacity GUI panel. You may also see missing capacity and throughput data under the Monitoring > Nodes panel.	There is currently no resolution or workaround. Try waiting 24 hours for the GUI to refresh. You can also try clicking Refresh .
Upon upgrades to ESS 5.3.2.1, you might see several Mellanox OFED weak-updates and unknown symbols messages on the console during gss_updatenode .	OFED Type: Upgrade Arch: Big Endian and Little Endian Version: All Affected nodes: N/A	When building the new OFED driver against the kernel, you might see many messages such as weak-updates and unknown symbols.	There is currently no resolution or workaround. These messages can be ignored.

Table 2. Known issues in ESS 5.3.2.1 (continued)

Issue	Environment affected	Description	Resolution or action
<p>During firmware upgrades on PPC64LE, update_flash might show the following warning:</p> <p>Unit kexec.service could not be found.</p>	<p>Firmware Type: Installation or Upgrade Arch: Little Endian Version: All Affected nodes: N/A</p>		<p>This warning can be ignored.</p>
<p>Setting target node names within gssutils might not persist for all panels. The default host names, such as ems1 might still show.</p>	<p>Deployment Type: Install or Upgrade Arch: Big Endian or Little Endian Version: All Affected nodes: All</p>	<p>gssutils allows users to conveniently deploy, upgrade, or manage systems within a GUI-like interface. If you run gssutils -N NODE, it must store that node name and use it throughout the menu system. There is a bug that might prevent this from working as designed.</p>	<p>Use on one of the following resolutions:</p> <ul style="list-style-type: none"> • Change the JSON file directly as follows. <ol style="list-style-type: none"> 1. Issue this command. <pre data-bbox="954 695 1360 772">/opt/ibm/gss/tools/bin/gssutils \ --customize --ems-node-name ems2 \ --config /tmp/LE.json</pre> 2. Issue this command. <pre data-bbox="954 821 1360 877">/opt/ibm/gss/tools/bin/gssutils \ --config /tmp/LE.json</pre> • For any given command within gssutils, you can press “c” to customize. At that point, you can change the node name(s) accordingly.
<p>The GUI does not display the firmware levels for drives.</p>	<p>GUI Type: Upgrade Arch: Big Endian Version: All</p>	<p>This behavior is seen during upgrade.</p>	<p>Use the mmlsfirmware command to view this information.</p>
<p>The 1Gb links show as unknown or unhealthy.</p>	<p>GUI Type: Install and Upgrade Arch: Big Endian or Little Endian Version: All</p>	<p>This behavior is seen during installation or upgrade.</p>	<p>mmhealth does not monitor the health state of IP interfaces that are not used by IBM Spectrum Scale. These are the IP interfaces that have the value None in the grid column Networks.</p>

Table 2. Known issues in ESS 5.3.2.1 (continued)

Issue	Environment affected	Description	Resolution or action
<p>The mmhealth command shows the status as degraded for an empty slot. (DCS3700 only – 5U84)</p>	<p>GUI / mmhealth Type: Install and Upgrade Arch: Big Endian or Little Endian Version: All</p>	<p>The handling of mmlsfirmware now marks empty slots. (DCS3700 only – 5U84) For example: Running mmlsfirmware --serial-number enclosure_serial results in: drive EMPTY SLOT <enclosure serial> not_available not_available is marked for slots that have no drive inserted, by design.</p>	<p>Currently there is no workaround for this issue. It is limited to DCS3700 – 5U84 enclosures.</p>
<p>The md5sum command works only under a folder where binaries are available.</p>	<p>Type: Install and Upgrade Arch: Big Endian or Little Endian Version: All</p>	<p>Upon running this command: md5sum -c /home/deploy/gss_install-5.3.2.1_ppc64le_datamanagement_20190121T232849Z.md5 The following error occurs: md5sum: gss_install-5.3.2.1_ppc64le_datamanagement_20190121T232849Z: No such file or directory gss_install-5.3.2.1_ppc64le_datamanagement_20190121T232849Z: FAILED open or read md5sum: WARNING: 1 listed file could not be read Press Enter to continue...</p>	<p>The md5sum -c command must be ran from CLI mode and from the folder in which the binary resides. For example:md5sum -c /home/deploy/gss_install-5.3.2.1_ppc64le_datamanagement_20190121T232849Z.md5</p>

Table 2. Known issues in ESS 5.3.2.1 (continued)

Issue	Environment affected	Description	Resolution or action
<p>Infiniband with multiple fabric is not supported.</p>	<p>Type: Install and Upgrade</p> <p>Arch: Big Endian or Little Endian</p> <p>Version:</p> <p>All</p>	<p>In a multiple fabric network, the Infiniband Fabric ID might not be properly appended in the verbsPorts configuration statement during the cluster creation. Incorrect verbsPort setting might cause the outage of the IB network.</p>	<p>It is advised to do the following to ensure that the verbsPorts setting is accurate:</p> <ol style="list-style-type: none"> 1. Use gssgennetworks to properly set up IB or Ethernet bonds on the ESS system. 2. Create a cluster. During cluster creation, the verbsPorts setting is applied and there is a probability that the IB network becomes unreachable, if multiple fabric are set up during the cluster deployment. 3. Ensure that the GPFS daemon is running and then run the mmfsadm test verbs config grep verbsPorts command. <p>These steps show the Fabric ID found for each link.</p> <p>For example:</p> <pre># mmfsadm test verbs config grep verbsPorts mmfs verbsPorts: mlx5_0/1/4 mlx5_1/1/7</pre> <p>In this example, the adapter <code>mlx5_0</code>, port 1 is connected to fabric 4 and the adapter <code>mlx5_1</code> port 1 is connected to fabric 7. Now, run the following command and ensure that verbsPorts settings are correctly configured to the GPFS cluster.</p> <pre># mmlsconfig grep verbsPorts verbsPorts mlx5_0/1 mlx5_1/1</pre> <p>Here, it can be seen that the fabric has not been configured even though IB was configured with multiple fabric. This is a known issue.</p> <p>Now using mmchconfig, modify the verbsPorts setting for each node or node class to take the subnet into account.</p> <pre>[root@gssi01 ~]# verbsPorts="\$(echo \$(mmfsadm test verbs config grep verbsPorts awk '{ \$1="" ; \$2="" ; \$3="" ; print \$0 } '))" # echo \$verbsPorts mlx5_0/1/4 mlx5_1/1/7 # mmchconfig verbsPorts="\$verbsPorts" -N gssi01 mmchconfig: Command successfully completed mmchconfig: Propagating the cluster configuration data to all affected nodes. This is an asynchronous process.</pre> <p>Here, the node can be any GPFS node or node class. Once the verbsPorts setting is changed, make sure that the new, correct verbsPorts setting is listed in the output of the mmlsconfig command.</p> <pre># mmlsconfig grep verbsPorts verbsPorts mlx5_0/1/4 mlx5_1/1/7</pre>

Table 2. Known issues in ESS 5.3.2.1 (continued)

Issue	Environment affected	Description	Resolution or action
<p>During an ESS upgrade, part information and firmware levels under the Hardware Details might be missing.</p>	<p>GUI Type: Upgrade Arch: Big Endian or Little Endian Version: All Affected nodes: N/A</p>	<p>The ESS GUI might be missing information under Part Info and Firmware version within the Hardware Details panel.</p>	<p>There are two workarounds:</p> <ol style="list-style-type: none"> 1. Wait for up to 24 hours for the GUI refresh tasks to run. 2. Try running a series of manual tasks to speed up the process of refreshing the GUI. <ol style="list-style-type: none"> a. Log in to the EMS node. b. Change directory to <code>/usr/1pp/mmfs/gui/cli</code>. c. Run the following tasks in order. <pre>runtask -c CLUSTER RECOVERY_GROUP runtask -c CLUSTER DISK_ENCLOSURES runtask -c CLUSTER ENCLOSURE_FW runtask -c CLUSTER CHECK_FIRMWARE</pre> <p>Where <i>CLUSTER</i> is either the cluster name or the cluster ID that can be determined by using the <code>mm1scluster</code> command.</p> <p>After running these tasks, the GUI should refresh with the issues resolved.</p>
<p>During file system creation in the ESS GUI, several inputs are ignored under Configure Properties.</p>	<p>GUI Arch: Big Endian or Little Endian Version: 5.3.1.x Affected nodes: N/A</p>	<p>When creating file systems in the ESS GUI, there are several properties that can be set under Configure Properties. Some of those values are:</p> <ul style="list-style-type: none"> • Enable quota • Quota scope • Inode access time update • Enable DMAPI • Enable file system features compatible with release <p>The GUI ignores these input fields and instead just passes only default values to the <code>mmcrfs</code> command.</p>	<p>You can use the following workarounds:</p> <ul style="list-style-type: none"> • Create a file system with the required values from the command line only. • Create the file system from the GUI and modify the values using <code>mmchfs</code> on the command line afterward.

Table 2. Known issues in ESS 5.3.2.1 (continued)

Issue	Environment affected	Description	Resolution or action
<p>A failed disk's state wouldn't be changed to drained or replace in new enclosures that are added by MES procedure and never be used for any file system.</p>	<p>Type: IBM Spectrum Scale RAID Arch: Little Endian Version: ESS 5.3.2.1 Affected Nodes: N/A</p>	<p>If a user runs mmvdisk pdisk change --simulate-failing to fail two pdisks in the new enclosure(s) that are added by using the MES procedure and never be used for any file system, the state of the second pdisk stays at <code>simulate-failing</code>. Then, GUI cannot detect that the second failed disk is replaceable, same as command line which fails because of the state of the disk.</p>	<p>Run the mmchpdisk --diagnose on the failing disk. Or, run mmshutdown or mmstartup on the I/O server node that serves the recovery group that the <code>simulate-failing</code> pdisk belongs to. Important: This workaround causes a failover.</p>
<p>Cable pulls might result in I/O hang and application failure.</p>	<p>Type: IBM Spectrum Scale RAID Arch: Big Endian and Little Endian Version: ESS 5.3.2.1 Affected Nodes: I/O server nodes</p>	<p>If SAS cables are pulled, I/O might hang for an extended period of time during RG or path recovery which could lead to application failures.</p>	<p>Change <code>nsdRAIDEventLogShortTermDelay</code> to 30ms (The default is 3000ms):</p> <ol style="list-style-type: none"> 1. Run mmchconfig nsdRAIDEventLogShortTermDelay=30. 2. Restart GPFS.
<p>gssinstall_<arch> and gssinstallcheck report <code>NOT_INST</code> for a few GPFS group RPMs.</p>	<p>Type: Deployment Arch: Little Endian Version: ESS 5.3.2.1 Affected Nodes: ALL</p>	<p>By default, deployment does not install RPMs for file audit logging support.</p> <p>Also, you may see <code>NOT_INST</code> for the Power 8 firmware RPM.</p>	<p>This is the expected behavior.</p> <p>If the file audit logging feature is required, you can manually install these packages from the EMS GPFS repository.</p> <ul style="list-style-type: none"> • <code>gpfs.kafka</code> • <code>gpfs.libkafka</code> <p>For information on how to install Power 8 firmware package (01SV860_165_165-1.1-1), see Appendix G, "Updating the system firmware," on page 65.</p>
<p>gssinstallcheck flags the MT4115 adapter with the incorrect firmware.</p>	<p>Type: Deployment Arch: Big Endian and Little Endian Version: ESS 5.3.2.1 Affected Nodes: ALL</p>	<p>There is a bug that causes gssinstallcheck to flag that the MT4115 adapter has incorrect firmware after deploying or upgrading of ESS 5.3.2.1.</p>	<p>If the firmware is showing 12.23.1020 for the MT4115 adapter you can safely ignore this message.</p>

Table 2. Known issues in ESS 5.3.2.1 (continued)

Issue	Environment affected	Description	Resolution or action
ESS GUI System Setup wizard fails on the Verify Installation screen in the IBM Spectrum Scale active check.	Type: GUI Arch: Big Endian or Little Endian Version: 5.3.2.1 Affected nodes: N/A	ESS System Setup wizard fails on the Verify Installation screen. The displayed error message is: Health monitoring is not active on 'X' nodes. Run 'mmhealth node show GPFS -N all' command to check why mmhealth does not provide health information for those nodes.	Click Verify again. The error should clear after that.

Appendix B. Troubleshooting for ESS on PPC64LE

Note: Most issues on ESS (PPC64LE) are not applicable if the Fusion mode is used.

Here are some tips on how to avoid common issues on ESS (PPC64LE).

- Always use /24 for the FSP network. It is advised to use 10.0.0.0/24.
- If possible use /24 for the xCAT network. It is advised to use 192.168.X.X/24.
- Do not overlap subnets on the EMS node. For instance, do not use 192.168.X.X on both networks.
- Do not use non-traditional subnets such as /26.
- Always verify that all nodes are visible on the FSP network by using **gssdeploy -f**.
- If you get a timeout (8 min/16min) during Genesis discovery look into the following:
 - Is the DHCP server started and without issue (**systemctl status dhcpd**)?
 - Is your subnetting correct?
 - Is your /etc/hosts file correct?
 - Look into using a genesis IP range.

Genesis IP range example:

Add the following to gssdeploy.cfg

```
EMS_GENESIS_IP_RANGE=192.168.202.13-192.168.202.14
```

In this case, 202.13 and 202.14 are the nodes that are being tried for deployment. There cannot be any nodes up with IPs in the given range. After setting the range, use **gssdeploy -x** or **gssdeploy -o** again. If all else fails, you can power on the node and boot into petitboot to obtain the deployment MAC address. Once obtained, you can add into the node xCAT definition and complete the steps manually to start deployment.

Each node ships with an extra cable in HMC port2 intended to be used to troubleshoot issues and access the FSP. It is advised you plug each cable into the FSP VLAN post deployment and set a static IP on the same subnet. Once this is done you can access ASMI remotely from the EMS.

Alternatively, you can use this cable to hook a laptop to in the lab to access each node via the default manufacturing static IP.

Another workaround to the Genesis timeout issue is to manually retrieve and insert the MAC addresses into the xCAT node definitions. Use the following steps if the Genesis discovery fails (**gssdeploy -x** times out):

1. Exit **gssdeploy**.
2. Use **rpower** to power off the node(s).
`rpower NodeName off`
3. Power on the node.
`rsetboot NodeName hd ; rpower NodeName on`
4. Bring up a console immediately.
`rcons NodeName -f`
5. When Petitboot comes up on the node, select **Exit to shell**.
6. Use the Linux **ifconfig** command to determine the interface that is holding the IP address.
7. Copy the MAC address and return to the EMS node.
8. Insert the MAC address.

```
chdef NodeName mac=MacAddress
```

```
nodeset NodeName osimage=install-gss_osimage_you_are_deploying
```

```
makedhcp -n ; makedns -n
```

At this point you can skip **gssdeploy -x** and move on to the next step in the Quick Deployment Guide.

- Ensure that the storage enclosures are powered off or SAS cables are disconnected before running the **gssdeploy -x** command. If you are unable to power off the storage enclosures or remove the SAS connections before running **gssdeploy -x**, genesis discovery might fail. In that case, exit **gssdeploy** and log in to the I/O server nodes by using the temporary dynamic IP address.
- In most cases, the node IP is different from the one in the `/etc/hosts` file. You can find it from the **dhcp status** or the **systemctl status dhcpd** commands, or from the journal or from `/var/log/messages`. It is also displayed in the **rcons** output. Log in and remove the `mpt3sas` driver (**modprobe -r mpt3sas**) and the nodes finish discovering. Confirm with the command **nodediscover1s** from the EMS node.

Appendix C. ESS networking considerations

This topic describes the networking requirements for installing ESS.

Note: The references to HMC are not applicable for the PPC64LE platform.

Networking requirements

The following networks are required:

- **Service network**

This network connects the flexible service processor (FSP) on the management server and I/O server nodes (with or without the HMC, depending on the platform) as shown in blue in Figure 1 and 2 on the following pages.

- **Management and provisioning network**

This network connects the management server to the I/O server nodes (and HMCs, if available) as shown in yellow in in Figure 1 and 2 on the following pages. The management server runs DHCP on the management and provisioning network. If a management server is not included in the solution order, a customer-supplied management server is used.

- **Clustering network**

This high-speed network is used for clustering and client node access. It can be a 10 Gigabit Ethernet (GbE), 40 GbE, or InfiniBand network. It might not be included in the solution order.

- **External and campus management network**

This public network is used for external and campus management of the management server, the HMC (if available), or both.

Figure 1, Network Topology, is a high-level logical view of the management and provisioning network and the service network for an ESS building block (on **PPC64BE**).

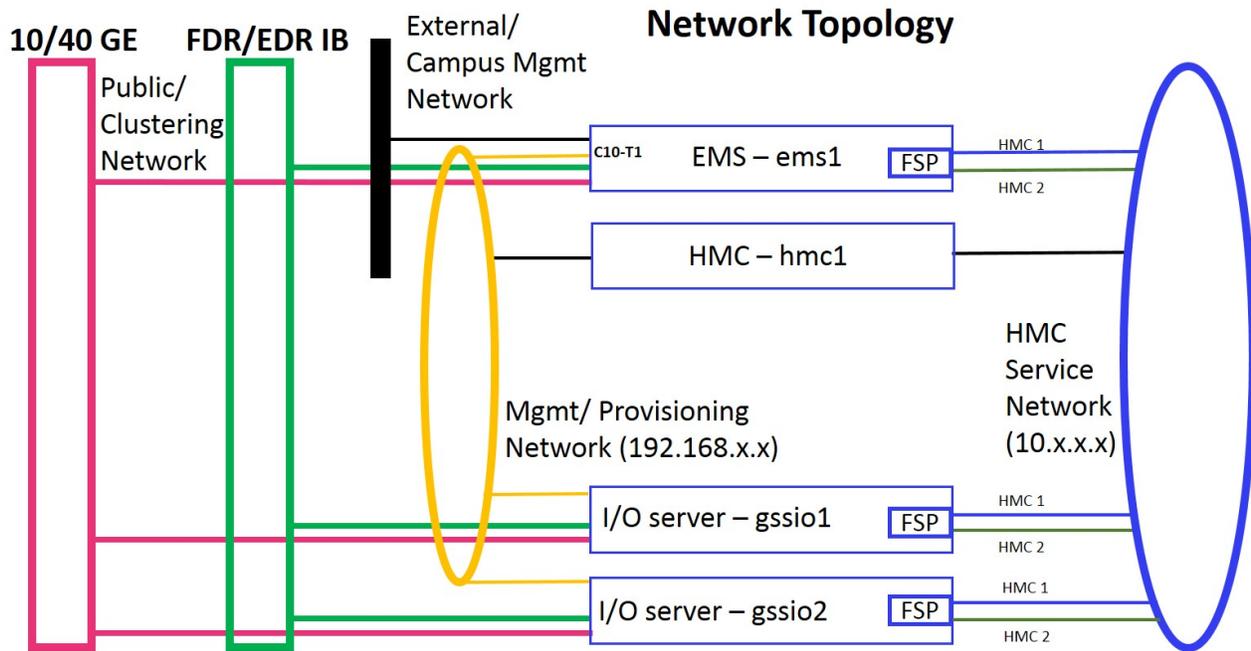


Figure 1. The management and provisioning network and the service network: a logical view (on PPC64BE)

Figure 2, Network Topology, is a high-level logical view of the management and provisioning network and the service network for an ESS building block (on PPC64LE).

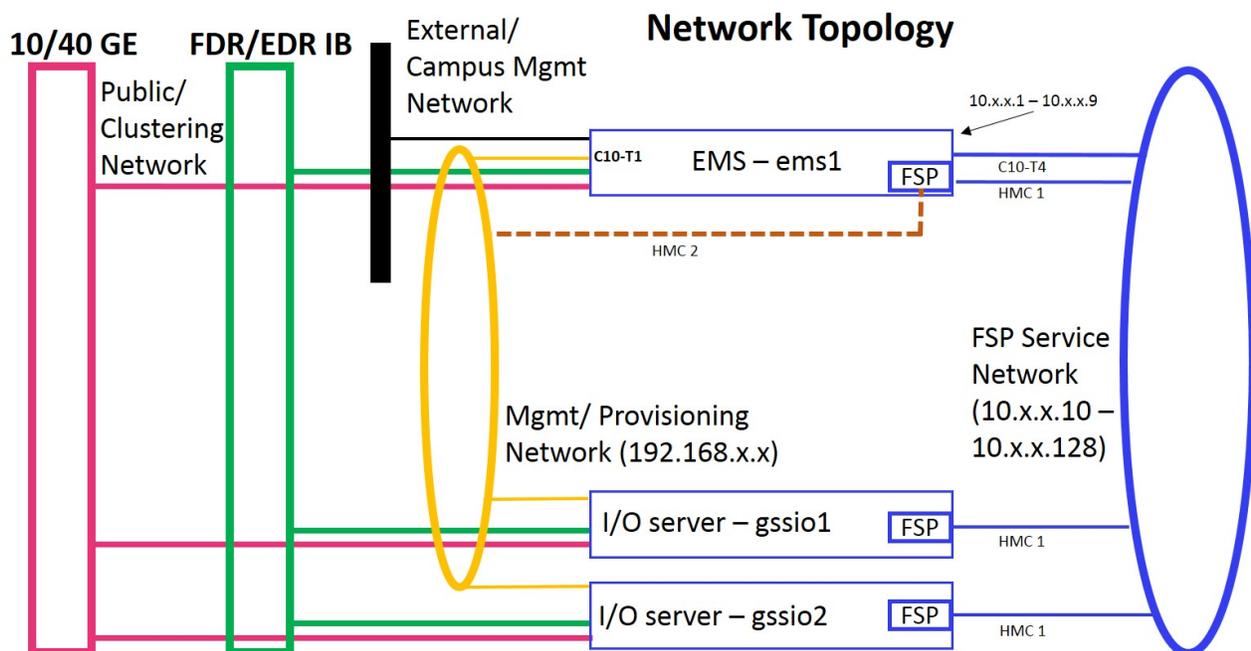


Figure 2. The management and provisioning network and the service network: a logical view (on PPC64LE)

The management and provisioning network and the service network must run as two non-overlapping networks implemented as two separate physical networks or two separate virtual local-area networks (VLANs).

Tip: HMC 2 is an optional third cable on the management server node that can be connected either to the management network or any other external network provided by the customer. This connection can be added in case the ability to service or control the management server node remotely is required.

The HMC, the management server, and the switches (1 GbE switches and high-speed switches) might not be included in a solution order in which an existing or customer-supplied HMC or management server is used. Perform any advance planning tasks that might be needed to access and use these solution components.

Customer networking considerations

Review the information about switches and switch firmware that were used to validate this ESS release. For information about available IBM networking switches, see the IBM networking switches page on IBM Knowledge Center.

It is recommended that if two switches are used in a high availability (HA) configuration, both switches be at the same firmware level.

To check the firmware version, do the following:

1. SSH to the switch.
2. Issue the following commands.

```
# en
# show version
```

For example:

```
login as: admin
Mellanox MLNX-OS Switch Management
Using keyboard-interactive authentication.
Password:
Last login: Mon Mar 5 12:03:14 2018 from 9.3.17.119
Mellanox Switch
io232 [master] >
io232 [master] > en
io232 [master] # show version
```

Example output:

```
Product name: MLNX-OS
Product release: 3.4.3002
Build ID: #1-dev
Build date: 2015-07-30 20:13:19
Target arch: x86_64
Target hw: x86_64
Built by: jenkins@fit74
Version summary: X86_64 3.4.3002 2015-07-30 20:13:19 x86_64
Product model: x86
Host ID: E41D2D52A040
System serial num: Defined in system VPD
System UUID: 03000200-0400-0500-0006-000700080009
```

Infiniband with multiple fabric

In a multiple fabric network, the Infiniband Fabric ID might not be properly appended in the verbsPorts configuration statement during the cluster creation. Incorrect verbsPort setting might cause the outage of the IB network. It is advised to do the following to ensure that the verbsPorts setting is accurate:

1. Use **gssgennetworks** to properly set up IB or Ethernet bonds on the ESS system.
2. Create a cluster. During cluster creation, the verbsPorts setting is applied and there is a probability that the IB network becomes unreachable, if multiple fabric are set up during the cluster deployment.

3. Ensure that the GPFS daemon is running and then run the `mmfsadm test verbs config | grep verbsPorts` command.

These steps show the Fabric ID found for each link.

For example:

```
# mmfsadm test verbs config | grep verbsPorts
mmfs verbsPorts: mlx5_0/1/4 mlx5_1/1/7
```

In this example, the adapter `mlx5_0`, port 1 is connected to fabric 4 and the adapter `mlx5_1` port 1 is connected to fabric 7. Now, run the following command and ensure that `verbsPorts` settings are correctly configured to the GPFS cluster.

```
# mmlsconfig | grep verbsPorts
verbsPorts mlx5_0/1 mlx5_1/1
```

Here, it can be seen that the fabric has not been configured even though IB was configured with multiple fabric. This is a known issue.

Now using `mmchconfig`, modify the `verbsPorts` setting for each node or node class to take the subnet into account.

```
[root@gssio1 ~]# verbsPorts="$(echo $(mmfsadm test verbs config | \
grep verbsPorts | awk '{ $1=""; $2=""; $3=""; print $0} '))"
# echo $verbsPorts
mlx5_0/1/4 mlx5_1/1/7
# mmchconfig verbsPorts="$verbsPorts" -N gssio1
mmchconfig: Command successfully completed
mmchconfig: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
```

Here, the node can be any GPFS node or node class. Once the `verbsPorts` setting is changed, make sure that the new, correct `verbsPorts` setting is listed in the output of the `mmlsconfig` command.

```
# mmlsconfig | grep verbsPorts
verbsPorts mlx5_0/1/4 mlx5_1/1/7
```

Switch information

ESS release updates are independent of switch updates. Therefore, it is recommended that Ethernet and Infiniband switches used with the ESS cluster be at their latest switch firmware levels. Customers are responsible for upgrading their switches to the latest switch firmware.

Table 3. Recommended high speed switch firmware levels

Switch model	Recommended firmware level
sb7700_sb7800	3.6.6164
sn2410_sn2700	3.6.8100
sx1410_sx1700	3.6.8008
sx6036	3.6.6106

Appendix D. Support for hybrid enclosures

ESS supports hybrid enclosures that comprise four or two enclosures containing only hard disk drives (HDDs) and one or two enclosures of only solid state drives (SSDs).

There are three hybrid enclosure models that are supported. The support for GH14S and GH24S is added in ESS 5.3.1. The support for GH12S is added in ESS 5.3.1.1:

GH14S

1 2U24 (SSD) and 4 5U84 (HDD) enclosures

GH14S recovery groups (RGs)

Each of the two GH14S RGs have a 12-disk SSD user declustered array (DA) and a 167-disk HDD user DA.

GH24S

2 2U24 (SSD) and 4 5U84 (HDD) enclosures

GH24S recovery groups

Each of the two GH24S RGs have a 24-disk SSD user DA and a 167-disk HDD user DA.

GH12S

1 2U24 (SSD) and 2 5U84 (HDD) enclosures

GH12S recovery groups

Each of the two GH12S RGs have a 12-disk SSD user DA and an 83-disk HDD user DA.

For information about 2U24 and 5U84 enclosures, see IBM ESS Expansion documentation.

Hybrid enclosure support in gssgenvdisk

The **gssgenvdisk** command can detect hybrid enclosures. In case of hybrid enclosures, the **gssdgenvdisk** command requires two declustered arrays (DAs). One DA comprises HDDs only and the other one comprises SSDs only.

gssgenvdisk provides a default placement policy in case hybrid enclosures are used. According to this policy, any data vdisk is placed on the DA which is composed of HDDs only and any metadata vdisk is placed on the DA which is composed of SSDs only.

In the following example, DA1 consists of HDDs and DA2 consists of SSDs. This example shows that the default data vdisk is placed in DA1 and the metadata vdisk is placed in DA2.

```
# gssgenvdisk --create-vdisk --create-nsds --create-filesystem --verbose \  
--contact-node essio41 --filesystem-name gpfs0 --reserved-space-percent 1
```

```
...  
...
```

```
# mmlsvdisk
```

vdisk name	RAID code	recovery group	declustered array	block size in KiB	remarks
rg_io41_ce_Data_16M_2p_1	8+2p	rg_io41-ce	DA1	16384	
rg_io41_ce_Data_16M_2p_2	8+2p	rg_io41-ce	DA1	16384	
rg_io41_ce_MetaData_1M_3W_1	3WayReplication	rg_io41-ce	DA2	1024	
rg_io41_ce_loghome	4WayReplication	rg_io41-ce	DA1	2048	log
rg_io41_ce_logtip	2WayReplication	rg_io41-ce	NVR	2048	logTip
rg_io41_ce_logtipbackup	Unreplicated	rg_io41-ce	SSD	2048	logTipBackup
rg_io42_ce_Data_16M_2p_1	8+2p	rg_io42-ce	DA1	16384	
rg_io42_ce_Data_16M_2p_2	8+2p	rg_io42-ce	DA1	16384	

rg_io42_ce_MetaData_1M_3W_1	3WayReplication	rg_io42-ce	DA2	1024	
rg_io42_ce_loghome	4WayReplication	rg_io42-ce	DA1	2048	log
rg_io42_ce_logtip	2WayReplication	rg_io42-ce	NVR	2048	logTip
rg_io42_ce_logtipbackup	Unreplicated	rg_io42-ce	SSD	2048	logTipBackup

You can override the default vdisk placement policy used in case of a hybrid enclosure system, by using the `--use-only-da` option. If the `--use-only-da` option is used, only the specified DA is considered for the vdisk creation. If the system is a hybrid enclosure and there are multiple DAs, the DAs available in the recovery group are not considered except for the one specified with the `--use-only-da` option.

You can create vdisks on other DAs using the BM Spectrum Scale RAID command.

Appendix E. Pre-installation tasks for ESS

This topic provides the pre-installation tasks required for ESS.

Note: The references to HMC are not applicable for the PPC64LE platform.

Table 4. Pre-installation tasks

ESS component	Description	Required actions	System settings
1. Service network Note: This network varies depending on the platform (PPC64BE or PPC64LE).	<p>HMC service network: This private network connects the HMC with the management server's FSP and the I/O server nodes. The service network must not be seen by the OS running on the node being managed (that is, the management server or the I/O server node).</p> <p>The HMC uses this network to discover the management server and the I/O server nodes and perform such hardware management tasks as creating and managing logical partitions, allocating resources, controlling power, and rebooting. Note: HMC is not applicable for the PPC64LE platform.</p> <p>FSP service network: This private network connects the FSP interface on EMS and the I/O server nodes. The service network must be seen by the OS running on the EMS node but not by the I/O server nodes being managed.</p>	<p>Perform any advance planning tasks that might be needed to access and use the HMC if it is not part of the solution order and a customer-supplied HMC will be used.</p> <p>Set up this network if it has not been set up already.</p>	Set the HMC to be the DHCP server for the service network.
2. Management and provisioning network	<p>This network connects the management server node with the HMC (when present) and the I/O server nodes. It typically runs over 1Gb.</p> <ul style="list-style-type: none"> This network is visible to the OS that is running on the nodes. The management server uses this network to communicate with the HMC (when present) and to discover the I/O server nodes. The management server will be the DHCP server on this network. There cannot be any other DHCP server on this network. This network is also used to provision the node and therefore deploy and install the OS on the I/O server nodes. 	<p>Perform any advance planning tasks that might be needed to access and use the management server if it is not part of the solution order and a customer-supplied management server will be used.</p> <p>Set up this network if it has not been set up already.</p>	
3. Clustering network	<p>This network is for high-performance data access. In most cases, this network is also part of the clustering network. It is typically composed of 10GbE, 40GbE, or InfiniBand networking components.</p>	Set up this network if it has not been set up already.	

Table 4. Pre-installation tasks (continued)

ESS component	Description	Required actions	System settings
4. Management network domain	The management server uses this domain for the proper resolution of hostnames.	Set the domain name using <i>lowercase</i> characters. Do <i>not</i> use any uppercase characters.	Example: gpfs.net
5. HMC node (IP address and hostname) Note: HMC is not applicable for the PPC64LE platform.	The IP address of the HMC node on the management network has a console name, which is the hostname and a domain name. <ul style="list-style-type: none"> This IP address must be configured and the link to the network interface must be up. The management server must be able to reach the HMC using this address. 	Set the fully-qualified domain name (FQDN) and the hostname using <i>lowercase</i> characters. Do <i>not</i> use any uppercase characters. Do <i>not</i> use a suffix of <i>-enx</i> , where <i>x</i> is any character. Do <i>not</i> use an <i>_</i> (underscore) in the hostname.	Example: IP address: 192.168.45.9 Hostname: hmc1 FQDN: hmc1.gpfs.net
6. Management server node (IP address)	The IP address of the management server node has an FQDN and a hostname. <ul style="list-style-type: none"> This IP address must be configured and the link to the network interface must be up. The management network must be reachable from this IP address. 	Set the FQDN and hostname using <i>lowercase</i> characters. Do <i>not</i> use any uppercase characters. Do <i>not</i> use a suffix of <i>-enx</i> , where <i>x</i> is any character. Do <i>not</i> use an <i>_</i> (underscore) in the hostname.	Example: IP address: 192.168.45.10 Hostname: ems1 FQDN: ems1.gpfs.net
7. I/O server nodes (IP addresses)	The IP addresses of the I/O server nodes have FQDNs and hostnames. <ul style="list-style-type: none"> These addresses are assigned to the I/O server nodes during node deployment. The I/O server nodes must be able to reach the management network using this address. 	Set the FQDN and hostname using <i>lowercase</i> characters. These names must match the name of the partition created for these nodes using the HMC. Do <i>not</i> use any uppercase characters. Do <i>not</i> use a suffix of <i>-enx</i> , where <i>x</i> is any character. Do <i>not</i> use an <i>_</i> (underscore) in the host name.	Example: I/O server 1: IP address: 192.168.45.11 Hostname: gssio1 FQDN: gssio1.gpfs.net I/O server 2: IP address: 192.168.45.12 Hostname: gssio2 FQDN: gssio2.gpfs.net
8. Management server node management network interface (PPC64BE) Management server node FSP network interface (PPC64LE)	The management network interface of the management server node must have the IP address that you set in item 6 assigned to it. This interface must have only one IP address assigned. For the PPC64LE system, one additional interface is assigned to FSP network. This interface must have only one IP address assigned.	To obtain this address, run: ip addr	Example: enP7p128s0f0

Table 4. Pre-installation tasks (continued)

ESS component	Description	Required actions	System settings
9. HMC (hscroot password) Note: HMC is not applicable for the PPC64LE platform.		Set the password for the hscroot user ID.	Example: abc123 This is the default password.
10. Kernel	Updating the kernel is required for all ESS nodes and it is verified by using gssinstallcheck .		Example: kernel-RHSA-2018-3843-LE-ESS5321.tar.gz
11. systemd	Updating the systemd service is required for all ESS nodes and it is verified by using gssinstallcheck .		Example: systemd-532-RHBA-2018-2753-BE.tar.gz
12. Network Manager	Updating the Network Manager service is required for all ESS nodes and it is verified by using gssinstallcheck .		Example: netmgr-RHSA-2018-3665-LE-ESS5321.tar.gz
13. Customer Red Hat Network (RHN) license keys	If possible, retrieve the RHN license keys for the customer in advance. This allows you to download the kernel, ISO, systemd, and Network Manager ahead of time. This also allows you to register and connect the newly deployed ESS system to RHN to apply security updates prior to leaving the site.		The keys must be available from the customer order. Contact offering management if help is required. Note: The customer must have an EU license.
14. I/O servers (user IDs and passwords)	The user IDs and passwords of the I/O servers are assigned during deployment.		Example: User ID: root Password: cluster (this is the default password)
15. FSP IPMI password	The IPMI password of the FSP. FSP IPMI of all the nodes assumed to be identical.		Example: PASSWORD
16. Clustering network (hostname prefix or suffix)	This high-speed network is implemented on a 10Gb Ethernet, 40Gb Ethernet or InfiniBand network.	Set a hostname for this network. It is customary to use hostnames for the high-speed network that use the prefix and suffix of the actual hostname. Do <i>not</i> use a suffix of -enx, where <i>x</i> is any character.	Examples: Suffixes: -bond0, -ib, -10G, -40G Hostnames with a suffix: gssi01-ib, gssi02-ib

Table 4. Pre-installation tasks (continued)

ESS component	Description	Required actions	System settings
17. High-speed cluster network (IP address)	<p>The IP addresses of the management server nodes and I/O server nodes on the high-speed cluster network have FQDNs and hostnames.</p> <p>In the example, 172.10.0.11 is the IP address that the GPFS daemon uses for clustering. The corresponding FQDN and hostname are gssio1-ib and gssio1-ib.data.net, respectively.</p>	<p>Set the FQDNs and hostnames.</p> <p>Do <i>not</i> make changes in the <code>/etc/hosts</code> file for the high-speed network until the deployment is complete. Do <i>not</i> create or enable the high-speed network interface until the deployment is complete.</p>	<p>Example:</p> <p>Management server: IP address: 172.10.0.10 Hostname: ems1-ib FQDN: ems1-ib.gpfs.net</p> <p>I/O server 1: IP address: 172.10.0.11 Hostname: gssio1-ib FQDN: gssio1-ib.data.net</p> <p>I/O server 2: IP address: 172.10.0.12 Hostname: gssio2-ib FQDN: gssio2-ib.data.net</p>
18. Red Hat Enterprise Linux 7.5	<p>The Red Hat Enterprise Linux 7.5 DVD or ISO file is used to create a temporary repository for the xCAT installation. xCAT uses it to create a Red Hat Enterprise Linux repository on the management server node.</p>	<p>Obtain this DVD or ISO file and download.</p> <p>For more information, see the Red Hat Enterprise Linux website:</p> <p>http://access.redhat.com/products/red-hat-enterprise-linux/</p>	<p>Example: RHEL-7.5-20180919.1-Server-ppc64-dvd1.iso</p> <p>Note: The Red Hat Enterprise Linux 7.5 ISO name depends on the architecture (PPC64BE or PPC64LE).</p>
19. Management network switch	<p>The switch that implements the management network must allow the Bootstrap Protocol (BOOTP) to go through.</p>	<p>Obtain the IP address and access credentials (user ID and password) of this switch.</p> <p>Some switches generate many Spanning Tree Protocol (STP) messages, which interfere with the network boot process. You need to disable STP to mitigate this.</p>	

Table 4. Pre-installation tasks (continued)

ESS component	Description	Required actions	System settings
20. Target file system	You need to provide information about the target file system that is created using storage in the ESS building blocks. This information includes name, block size, file system size, RAID code, etc. This information you is passed on to gssgenvdisks to create the customer file system.	Set the target file system name, the mount point, the block size, the number of data NSDs, and the number of metadata NSDs.	Example: Block size = 16M, #datansd=4, #metadatansd=2

Appendix F. Installation: reference

This topic provides information on adding IBM Spectrum Scale nodes to an ESS cluster and node name considerations.

Adding IBM Spectrum Scale nodes to an ESS cluster

IBM Spectrum Scale node configuration is optimized for running IBM Spectrum Scale RAID functions.

1. ESS cluster node configuration is optimized for running IBM Spectrum Scale RAID functions. Protocols, other gateways, or any other non-ESS services must not be run on ESS management server nodes or I/O server nodes. In a cluster with high IO load, avoid using ESS nodes as cluster manager or filesystem manager. For optimal performance the NSD client nodes accessing ESS nodes should be properly configured. ESS ships with **gssClientConfig.sh script** located in `/usr/lpp/mmfs/samples/gss/` directory. This script can be used to configure the client as follows:

```
/usr/lpp/mmfs/samples/gss/gssClientConfig.sh <Comma Separated list of  
client nodes or nodeclass>
```

You can run the following to see configuration parameter settings without setting them:

```
/usr/lpp/mmfs/samples/gss/gssClientConfig.sh -D
```

After running this script, restart GPFS on the affected nodes for the optimized configuration settings to take effect.

Important: Do not run **gssClientConfig.sh** unless you fully understand the impact of each setting on the customer environment. Make use of the `-D` option to decide if all or some of the settings might be applied. Then, individually update each client node settings as required.

2. When IBM Spectrum Scale nodes deployed with protocols are added to the ESS cluster, quorum, cluster manager, and filesystem manager functions should be moved from the ESS to the protocol nodes after adding protocol nodes to the cluster.

For information about adding an IBM Spectrum Scale protocol node to an ESS cluster, see:

- Overview of the IBM Spectrum Scale installation toolkit
- Preparing a cluster that contains ESS for adding protocols
- Spectrum Scale Protocols Quick Overview

Node name considerations

Carefully select the hostname, suffix, and prefix of the management server and I/O server so that the hostname used in the high-speed network and by the ESS cluster can be generated from the suffix or prefix.

High-speed hostnames

Example 1:

```
a-bcd-edf-1  
a-bcd-edf-2  
a-bcd-edf-3  
a-bcd-edf-4
```

Here, `a-bcd-` is the prefix and `edf-1`, `edf-2`, `edf-3`, and `edf-4` are the xCAT names of the nodes.

Example 2:

```
1-a-bcd-edf
2-b-bcd-edf
3-c-bcd-edf
4-d_bcd_edf
```

Here, `-edf` is the suffix and `1-a-bcd`, `2-a-bcd`, `3-a-bcd`, and `4-a-bcd` are the xCAT names of the nodes.

If possible, avoid using high-speed node names with variations at the beginning and the end, such as:

```
A-a-bcd-edf-1
B-b-bdc-edf-2
C-c-bcd-edf-3
D-d-bcd-edf-4
```

In such cases, use the `-N` option and specify the node list with the `gssgencluster` and `gssgenclusterrgs` commands. The node names must be reachable from the management server node. xCAT requires that the target nodes be part of a node group and a warning might be issued if the hostname is not defined as an xCAT object.

Example:

1. The xCAT hostnames are `gssio1`, `gssio2`, `gssio3`, and `gssio4`.
2. The high-speed hostnames are `A-test1`, `B-test2`, `C-test3`, `D-test4`. These hostnames are reachable from the management server node. They are not defined in xCAT.

Run:

```
gssgencluster -C test01 -N A-test1,B-test2,C-test3,D-test4
```

Appendix G. Updating the system firmware

Use this information to obtain and apply the system firmware updates.

The system firmware packages are available in one of the following directories depending on the architecture of the management server node in newly shipped systems:

- **PPC64BE:** /opt/ibm/gss/install/rhel7/ppc64/firmware
- **PPC64LE:** /opt/ibm/gss/install/rhel7/ppc64le/firmware
- System firmware update files for PPC64BE for updating using HMC:
 - | 01SV860_165_165.rpm
 - | 01SV860_165_165.xml
- System firmware update file for PPC64LE for updating using the command line:
 - | 01SV860_165_165.img

Depending on your platform, use one of the following sets of steps for updating system firmware.

- Update the system firmware on PPC64LE systems as follows.
 1. Unpack the *img file in the /tmp/fwupdate directory.
 - | cd /opt/ibm/gss/install/rhel7/ppc64le/firmware
 - | rpm -ivh 01SV860_165_165.rpm
 2. Shutdown IBM Spectrum Scale and stop any ongoing I/O on the node.
 3. Verify the firmware level.
 - | update_flash -v -f /tmp/fwupdate/01SV860_165_165.img
 4. Update the system firmware.
 - | update_flash -f /tmp/fwupdate/01SV860_165_165.img

After issuing this command, the node reboots and updates the firmware. It could take up to 30 minutes for the node to reboot with the new firmware level. You can then run **gssinstallcheck** on the node to verify if the firmware is successfully updated.

To update system firmware on PPC64BE systems, you must use HMC and you must upgrade HMC to 860 SP3 before updating system firmware. For information about upgrading HMC, see HMC V8 Upgrade Procedure.

- Update the system firmware on PPC64BE systems as follows.
 1. From the HMC navigation area, click **Resources > All Systems > Server > Updates**.
 2. From the **Updates** menu, click **Change Licensed Internal Code > for the Current Release...**
 3. Using SFTP, point to the /opt/ibm/gss/install/rhel7/ppc64/firmware directory on the EMS node. The following files should be present:
 - | 01SV860_165_165.rpm
 - | 01SV860_165_165.xml

Note: For updating the system firmware using HMC, if SFTP to the EMS node does not work, move the *rpm and the *xml files to a server which is accessible using FTP or SFTP.

4. Select the update file and update the system firmware. It could take up to 30 minutes to update the firmware using HMC.

Appendix H. Upgrading the Hardware Management Console (HMC)

For PPC64BE deployments, ensure that HMC is properly configured for the management server node and I/O server nodes and partition names are correctly set.

- To apply the HMC V8 update, use the following resources:
 - HMC V8 upgrade procedure: <https://www-01.ibm.com/support/docview.wss?uid=nas8N1020108>
 - HMC V8 860 files: <ftp://public.dhe.ibm.com/software/server/hmc/network/v8860/>
 - HMC V8 860 SP3 ISO: ftp://public.dhe.ibm.com/software/server/hmc/updates/HMC_Update_V8R860_SP3.iso

After upgrading, the HMC configuration should be similar to:
Release: 8.6.0 Base Version: V8R8.6.0 Service Pack: 3

Note: This is not applicable for the PPC64LE platform.

Appendix I. Obtaining kernel for system upgrades

For new system installation, the kernel is shipped with the system. However, for upgrades, you need to obtain and package the kernel update, and then follow the kernel update installation procedure.

You must have a EUS license to download the kernel from Red Hat Network.

Use the following steps during an upgrade to obtain and package the kernel update.

1. Clear the version locks.

```
yum versionlock clear
```

2. Connect the management server node to the Red Hat Network.

```
subscription-manager register --username=<X> --password=<Y>
subscription-manager list --available // list pools
subscription-manager attach --pool=<X>
```

Or

```
subscription-manager attach --auto
```

3. Create a directory for the kernel update package.

For PPC64BE, issue:

```
| mkdir -p /tmp/kernel/RHSA-2018-3843-BE/
```

For PPC64LE, issue:

```
| mkdir -p /tmp/kernel/RHSA-2018-3843-LE/
```

4. List all repositories and enable the repositories that are disabled, as required.

```
yum repolist all
yum-config-manager --enable rhel*
```

Or

```
subscription-manager config --rhsm.manage_repos=1
```

5. Download the kernel update package.

For PPC64BE, issue:

```
| yum update *862.25.3* --downloadonly --downloadaddir=/tmp/kernel/RHSA-2018-3843-BE
| yum update perf-3.10.0-862.25.3.e17.ppc64.rpm --downloadonly --downloadaddir=/tmp/kernel/RHSA-2018-3843-BE
| yum update python-perf-3.10.0-862.25.3.e17.ppc64.rpm --downloadonly \
| --downloadaddir=/tmp/kernel/RHSA-2018-3843-BE
```

For PPC64LE, issue:

```
yum update *862.25.3* --downloadonly --downloadaddir=/tmp/kernel/RHSA-2018-3843-LE
yum update perf-3.10.0-862.25.3.e17.ppc64le.rpm --downloadonly --downloadaddir=/tmp/kernel/RHSA-2018-3843-LE
yum update python-perf-3.10.0-862.25.3.e17.ppc64le.rpm --downloadonly \
--downloadaddir=/tmp/kernel/RHSA-2018-3843-LE
```

The command-line kernel download method might fail if a newer kernel is available. In that case, use these steps.

- a. Use one of the following steps depending on your platform:

- For PPC64BE, go to the following URL: https://access.redhat.com/search/#/?q=kernel*862*25*.3*ppc64.rpm&p=1&srch=any&documentKind=
- For PPC64LE, go to the following URL: https://access.redhat.com/search/#/?q=kernel*862*25*.3*ppc64le.rpm&p=1&srch=any&documentKind=

- b. Search for the required or any additional RPMs listed in “About the ESS Red Hat Linux Errata Kernel Update” on page 70 and download them.

6. Package the directory.

For PPC64BE, issue:

```
cd /tmp/kernel ; tar -zcvf kernel-RHSA-2018-3843-BE.tar.gz RHSA-2018-3843-BE
```

For PPC64LE, issue:

```
cd /tmp/kernel ; tar -zcvf kernel-RHSA-2018-3843-LE.tar.gz RHSA-2018-3843-LE
```

Note: Make sure that the RPM files are in the RHSA-2018-3843-BE or the RHSA-2018-3843-LE folder. Do not create any nested folder inside the RHSA-2018-3843-BE or the RHSA-2018-3843-LE folder and try to place the RPM file in that nested folder. Doing so results in failure of the kernel patch installation during the cluster deployment.

7. Disable the Red Hat Network connection on the management server node.

```
subscription-manager config --rhsm.manage_repos=0
yum clean all
```

Continue with the kernel update installation steps for `kernel-RHSA-2018-3843-BE.tar.gz` or `kernel-RHSA-2018-3843-LE.tar.gz` using `gssdeploy -k`. For example, use one of the following commands depending on the architecture to place the kernel updates in the kernel repository:

For PPC64BE, issue:

```
/var/tmp/gssdeploy -k kernel-RHSA-2018-3843-BE.tar.gz --silent
```

This command places the kernel update in `/install/gss/otherpkgs/rhels7/ppc64/kernel`

For PPC64LE, issue:

```
/var/tmp/gssdeploy -k kernel-RHSA-2018-3843-LE.tar.gz --silent
```

This command places the kernel update in `/install/gss/otherpkgs/rhels7/ppc64le/kernel`

For more information about the kernel update, see “About the ESS Red Hat Linux Errata Kernel Update.”

About the ESS Red Hat Linux Errata Kernel Update

This topic provides information about the Red Hat Linux Errata Kernel Update for ESS.

At the time of shipping from factory, most current recommended kernel errata and associated RPMs are provided in the `/home/deploy` directory.

Kernel errata updates can be obtained from Red Hat network (RHN) using the supplied license:
<https://access.redhat.com/errata/#/>.

For information about the kernel update for the current release, see <https://access.redhat.com/errata/RHSA-2018:3843>.

This example shows errata update (RHSA-2018-3843) provided in the `/home/deploy` directory of the EMS node when shipped from factory.

The following packages are provided in `kernel-RHSA-2018-3843-BE.tar.gz`:

```
kernel-3.10.0-862.25.3.el7.ppc64.rpm
kernel-abi-whitelists-3.10.0-862.25.3.el7.noarch.rpm
kernel-bootwrapper-3.10.0-862.25.3.el7.ppc64.rpm
kernel-devel-3.10.0-862.25.3.el7.ppc64.rpm
kernel-doc-3.10.0-862.25.3.el7.noarch.rpm
kernel-headers-3.10.0-862.25.3.el7.ppc64.rpm
```

```
| kernel-tools-3.10.0-862.25.3.el7.ppc64.rpm
| kernel-tools-libs-3.10.0-862.25.3.el7.ppc64.rpm
| kernel-tools-libs-devel-3.10.0-862.25.3.el7.ppc64.rpm
| perf-3.10.0-862.25.3.el7.ppc64.rpm
| python-perf-3.10.0-862.25.3.el7.ppc64.rpm

| The following packages are provided in kernel-RHSA-2018-3843-LE.tar.gz:
| kernel-3.10.0-862.25.3.el7.ppc64le.rpm
| kernel-abi-whitelists-3.10.0-862.25.3.el7.noarch.rpm
| kernel-bootwrapper-3.10.0-862.25.3.el7.ppc64le.rpm
| kernel-devel-3.10.0-862.25.3.el7.ppc64le.rpm
| kernel-doc-3.10.0-862.25.3.el7.noarch.rpm
| kernel-headers-3.10.0-862.25.3.el7.ppc64le.rpm
| kernel-tools-3.10.0-862.25.3.el7.ppc64le.rpm
| kernel-tools-libs-3.10.0-862.25.3.el7.ppc64le.rpm
| kernel-tools-libs-devel-3.10.0-862.25.3.el7.ppc64le.rpm
| perf-3.10.0-862.25.3.el7.ppc64le.rpm
| python-perf-3.10.0-862.25.3.el7.ppc64le.rpm
```

Appendix J. Obtaining systemd update for system upgrades

For new system installation, the systemd update is shipped with the system and it is available in the /home/deploy directory. However, for upgrades, you need to obtain and package the systemd update, and then install the systemd update.

You must have a EUS license to download the systemd update from Red Hat Network.

Use the following steps during an upgrade to obtain and package the systemd update.

1. Clear the version locks.

```
yum versionlock clear
```

2. Connect the management server node to the Red Hat Network.

```
subscription-manager register --username=<X> --password=<Y>
subscription-manager list --available // list pools
subscription-manager attach --pool=<X>
```

Or

```
subscription-manager attach --auto
```

3. Create a directory for the systemd update package.

For PPC64BE, issue:

```
mkdir -p /tmp/systemd/RHBA-2018-2753-BE/
```

For PPC64LE, issue:

```
mkdir -p /tmp/systemd/RHBA-2018-2753-LE/
```

4. List all repositories and enable the repositories that are disabled, as required.

```
yum repolist all
yum-config-manager --enable rhel*
```

Or

```
subscription-manager config --rhsm.manage_repos=1
```

5. Download the systemd update package.

For PPC64BE, issue:

```
| yum update systemd*219-57.e17_5.3* --downloadonly --downloadaddir=/tmp/systemd/RHBA-2018-2753-BE
| yum update libgudev1-219-57.e17_5.3.ppc64.rpm --downloadonly --downloadaddir=/tmp/systemd/RHBA-2018-2753-BE
| yum update libgudev1-devel-219-57.e17_5.3.ppc64.rpm --downloadonly --downloadaddir=/tmp/systemd/RHBA-2018-2753-BE
| yum update dracut*033-535*.ppc64.rpm --downloadonly --downloadaddir=/tmp/systemd/RHBA-2018-2753-BE
```

For PPC64LE, issue:

```
| yum update systemd*219-57.e17_5.3* --downloadonly --downloadaddir=/tmp/systemd/RHBA-2018-2753-LE
| yum update libgudev1-219-57.e17_5.3.ppc64le.rpm --downloadonly --downloadaddir=/tmp/systemd/RHBA-2018-2753-LE
| yum update libgudev1-devel-219-57.e17_5.3.ppc64le.rpm --downloadonly --downloadaddir=/tmp/systemd/RHBA-2018-2753-LE
| yum update dracut*033-535*.ppc64le.rpm --downloadonly --downloadaddir=/tmp/systemd/RHBA-2018-2753-LE
```

The command-line kernel download method might fail if a newer kernel is available. In that case, use these steps.

- a. Use one of the following steps depending on your platform:

- For PPC64BE, go to the following URL: https://access.redhat.com/search/#/?3Fq=systemd*219*57*e17*5*3*ppc64.rpm%26p=1%26sort=relevant%26scoped=false%26language=en
- For PPC64LE, go to the following URL: https://access.redhat.com/search/#/?3Fq=systemd*219*57*e17*5*3*ppc64le.rpm%26p=1%26sort=relevant%26scoped=false%26language=en

b. Search for the required or any additional RPMs listed in “About the ESS Red Hat Linux systemd update” and download them.

6. Package the directory.

For PPC64BE, issue:

```
| cd /tmp/systemd ; tar -zcvf systemd-RHBA-2018-2753-BE.tar.gz RHBA-2018-2753-BE
```

For PPC64LE, issue:

```
| cd /tmp/systemd ; tar -zcvf systemd-RHBA-2018-2753-LE.tar.gz RHBA-2018-2753-LE
```

Note: Make sure that the RPM files are in the RHBA-2018-2753-BE or the RHBA-2018-2753-LE folder. Do not create any nested folder inside the RHBA-2018-2753-BE or the RHBA-2018-2753-LE folder and try to place the RPM file in that nested folder. Doing so results in failure of the systemd patch installation during the cluster deployment.

7. Disable the Red Hat Network connection on the management server node.

```
subscription-manager config --rhsm.manage_repos=0  
yum clean all
```

| Continue with the systemd update installation steps for systemd-RHBA-2018-2753-BE.tar.gz or using **gssdeploy -p**. For example, use one of the following commands depending on the architecture to place the systemd update in the patch repository:

For PPC64BE, issue:

```
| /var/tmp/gssdeploy -p RHBA-2018-2753-BE.tar.gz --silent
```

This command places the systemd updates in /install/gss/otherpkgs/rhels7/ppc64/patch

For PPC64LE, issue:

```
| /var/tmp/gssdeploy -p systemd-RHBA-2018-2753-LE.tar.gz --silent
```

This command places the systemd updates in /install/gss/otherpkgs/rhels7/ppc64le/patch

For more information, see “About the ESS Red Hat Linux systemd update.”

About the ESS Red Hat Linux systemd update

This topic provides information about the Red Hat Linux systemd update for ESS.

This example shows systemd update (RHBA-2018-2753) provided in the /home/deploy directory of the EMS node when shipped from factory.

For information about the systemd update for the current release, see <https://access.redhat.com/errata/RHBA-2018:2753>.

| The following packages are provided in systemd-RHBA-2018-2753-BE.tar.gz:

```
| systemd-219-57.e17_5.3.ppc64.rpm  
| systemd-devel-219-57.e17_5.3.ppc64.rpm  
| systemd-devel-219-57.e17_5.3.ppc.rpm  
| systemd-journal-gateway-219-57.e17_5.3.ppc64.rpm  
| systemd-libs-219-57.e17_5.3.ppc64.rpm  
| systemd-libs-219-57.e17_5.3.ppc.rpm  
| systemd-networkd-219-57.e17_5.3.ppc64.rpm
```

```
| systemd-python-219-57.e17_5.3.ppc64.rpm
| systemd-resolved-219-57.e17_5.3.ppc64.rpm
| systemd-resolved-219-57.e17_5.3.ppc.rpm
| systemd-sysv-219-57.e17_5.3.ppc64.rpm
| libgudev1-219-57.e17_5.3.ppc64.rpm
| libgudev1-219-57.e17_5.3.ppc.rpm
| libgudev1-devel-219-57.e17_5.3.ppc64.rpm
| libgudev1-devel-219-57.e17_5.3.ppc.rpm
| dracut-033-535.e17_5.1.ppc64.rpm
| dracut-network-033-535.e17_5.1.ppc64.rpm
| dracut-config-rescue-033-535.e17_5.1.ppc64.rpm

| The following packages are provided in the systemd-RHBA-2018-2753-LE.tar.gz:
| systemd-219-57.e17_5.3.ppc64le.rpm
| systemd-devel-219-57.e17_5.3.ppc64le.rpm
| systemd-journal-gateway-219-57.e17_5.3.ppc64le.rpm
| systemd-libs-219-57.e17_5.3.ppc64le.rpm
| systemd-networkd-219-57.e17_5.3.ppc64le.rpm
| systemd-python-219-57.e17_5.3.ppc64le.rpm
| systemd-resolved-219-57.e17_5.3.ppc64le.rpm
| systemd-sysv-219-57.e17_5.3.ppc64le.rpm
| libgudev1-219-57.e17_5.3.ppc64le.rpm
| libgudev1-devel-219-57.e17_5.3.ppc64le.rpm
| dracut-033-535.e17_5.1.ppc64le.rpm
| dracut-network-033-535.e17_5.1.ppc64le.rpm
| dracut-config-rescue-033-535.e17_5.1.ppc64le.rpm
```

Appendix K. Obtaining Network Manager updates for system upgrades

For new system installation, the Network Manager update is shipped with the system and it is available in the `/home/deploy` directory. However, for upgrades, you need to obtain and package the Network Manager update, and then install the Network Manager update.

You must have a EUS license to download the Network Manager update from Red Hat Network.

Use the following steps during an upgrade to obtain and package the Network Manager update.

1. Clear the version locks.

```
yum versionlock clear
```

2. Connect the management server node to the Red Hat Network.

```
subscription-manager register --username=<X> --password=<Y>
subscription-manager list --available // list pools
subscription-manager attach --pool=<X>
```

Or

```
subscription-manager attach --auto
```

3. Create a directory for the Network Manager update package.

For PPC64BE, issue:

```
mkdir -p /tmp/netmgr/RHSA-2018-3665-BE
```

For PPC64LE, issue:

```
mkdir -p /tmp/netmgr/RHSA-2018-3665-LE
```

4. List all repositories and enable the repositories that are disabled, as required.

```
yum repolist all
yum-config-manager --enable rhel*
```

Or

```
subscription-manager config --rhsm.manage_repos=1
```

5. Download the Network Manager update package.

For PPC64BE, issue:

```
yum update NetworkManager*1.12.0*8* --downloadonly --downloadaddir=/tmp/netmgr/RHSA-2018-3665-BE
yum update glib2-2.56.1-1.e17.ppc64.rpm --downloadonly --downloadaddir=/tmp/netmgr/RHSA-2018-3665-BE
yum update glib2-devel-2.56.1-1.e17.ppc64.rpm --downloadonly --downloadaddir=/tmp/netmgr/RHSA-2018-3665-BE
```

For PPC64LE, issue:

```
yum update NetworkManager*1.12.0*8* --downloadonly --downloadaddir=/tmp/netmgr/RHSA-2018-3665-LE
yum update glib2-2.56.1-1.e17.ppc64le.rpm --downloadonly --downloadaddir=/tmp/netmgr/RHSA-2018-3665-LE
yum update glib2-devel-2.56.1-1.e17.ppc64le.rpm --downloadonly --downloadaddir=/tmp/netmgr/RHSA-2018-3665-LE
```

The command-line kernel download method might fail if a newer kernel is available. In that case, use these steps.

- a. Use one of the following steps depending on your platform:

- For PPC64BE, go to the following URL: https://access.redhat.com/search/#/?3Fq=NetworkManager*1.12.0*8*ppc64.rpm*%26p=1%26sort=relevant%26scoped=false
- For PPC64LE, go to the following URL: https://access.redhat.com/search/#/?3Fq=NetworkManager*1.12.0*8*ppc64le*%26p=1%26sort=relevant%26scoped=false

- b. Search for the required or any additional RPMs listed in “About the ESS Red Hat Linux Network Manager update” on page 78 and download them.

6. Package the directory.

For PPC64BE, issue:

```
| cd /tmp/systemd ; tar -zcvf netmgr-RHSA-2018-3665-BE.tar.gz RHSA-2018-3665-BE
```

For PPC64LE, issue:

```
| cd /tmp/systemd ; tar -zcvf netmgr-RHSA-2018-3665-LE.tar.gz RHSA-2018-3665-LE
```

Note: Make sure that the RPM files are in the RHSA-2018-3665-BE or the RHSA-2018-3665-LE folder. Do not create any nested folder inside the RHSA-2018-3665-BE or the RHSA-2018-3665-LE folder and try to place the RPM file in that nested folder. Doing so will result in failure of the network manager patch installation during the cluster deployment.

7. Disable the Red Hat Network connection on the management server node.

```
subscription-manager config --rhsm.manage_repos=0
yum clean all
```

8. Place the Network Manager updates in the patch repository.

For PPC64BE, issue:

```
| /var/tmp/gssdeploy -p netmgr-RHSA-2018-3665-BE.tar.gz --silent
```

This command places the Network Manager updates in `/install/gss/otherpkgs/rhels7/ppc64/patch`

For PPC64LE, issue:

```
| /var/tmp/gssdeploy -p netmgr-RHSA-2018-3665-LE.tar.gz --silent
```

This command places the Network Manager updates in `/install/gss/otherpkgs/rhels7/ppc64le/patch` For more information, see “About the ESS Red Hat Linux Network Manager update.”

About the ESS Red Hat Linux Network Manager update

This topic provides information about the Red Hat Linux Network Manager update for ESS.

| This example shows Network Manager update (RHSA-2018-3665) provided in the `/home/deploy` directory of the EMS node when shipped from factory.

| For information about the Network Manager update for the current release, see <https://access.redhat.com/errata/RHSA-2018:3665>.

| The following packages are provided in `netmgr-RHSA-2018-3665-BE.tar.gz`:

```
| NetworkManager-1.12.0-8.e17_6.ppc64.rpm
| NetworkManager-ads1-1.12.0-8.e17_6.ppc64.rpm
| NetworkManager-bluetooth-1.12.0-8.e17_6.ppc64.rpm
| NetworkManager-glib-1.12.0-8.e17_6.ppc64.rpm
| NetworkManager-glib-devel-1.12.0-8.e17_6.ppc64.rpm
| NetworkManager-libnm-1.12.0-8.e17_6.ppc64.rpm
| NetworkManager-libnm-devel-1.12.0-8.e17_6.ppc64.rpm
| NetworkManager-ppp-1.12.0-8.e17_6.ppc64.rpm
| NetworkManager-team-1.12.0-8.e17_6.ppc64.rpm
| NetworkManager-tui-1.12.0-8.e17_6.ppc64.rpm
| NetworkManager-wifi-1.12.0-8.e17_6.ppc64.rpm
| NetworkManager-wwan-1.12.0-8.e17_6.ppc64.rpm
```

```
| NetworkManager-dispatcher-routing-rules-1.12.0-8.e17_6.noarch.rpm
| NetworkManager-config-server-1.12.0-8.e17_6.noarch.rpm
| glib2-2.56.1-1.e17.ppc64.rpm
| glib2-devel-2.56.1-1.e17.ppc64.rpm
```

| The following packages are provided in the netmgr-RHSA-2018-3665-LE.tar.gz:

```
| NetworkManager-1.12.0-8.e17_6.ppc64le.rpm
| NetworkManager-adsl-1.12.0-8.e17_6.ppc64le.rpm
| NetworkManager-bluetooth-1.12.0-8.e17_6.ppc64le.rpm
| NetworkManager-glib-1.12.0-8.e17_6.ppc64le.rpm
| NetworkManager-glib-devel-1.12.0-8.e17_6.ppc64le.rpm
| NetworkManager-libnm-1.12.0-8.e17_6.ppc64le.rpm
| NetworkManager-libnm-devel-1.12.0-8.e17_6.ppc64le.rpm
| NetworkManager-ppp-1.12.0-8.e17_6.ppc64le.rpm
| NetworkManager-team-1.12.0-8.e17_6.ppc64le.rpm
| NetworkManager-tui-1.12.0-8.e17_6.ppc64le.rpm
| NetworkManager-wifi-1.12.0-8.e17_6.ppc64le.rpm
| NetworkManager-wwan-1.12.0-8.e17_6.ppc64le.rpm
| NetworkManager-config-server-1.12.0-8.e17_6.noarch.rpm
| NetworkManager-dispatcher-routing-rules-1.12.0-8.e17_6.noarch.rpm
| glib2-2.56.1-1.e17.ppc64le.rpm
| glib2-devel-2.56.1-1.e17.ppc64le.rpm
```

Appendix L. MES upgrade flow

For customers who are looking to add storage to an existing building block, this option is now supported. The goal of MES is to expand a customer's available storage, either to existing file system or a new one, without the need to buy an entirely new building block. MES is designed to be an online operation and to not interrupt customer workloads.

Supported paths

- GS1S -> GS2S
- GS2S -> GS4S
- GL1S -> GL2S
- GL2S -> GL4S
- GL4S -> GL6S

You may not hop multiple building block types at a time in the same MES upgrade session. For example:

GL1S -> GL4S is not supported.

You must first go from GL1S -> GL2S and then from GL2S -> GL4S.

Prerequisites

1. All new or existing building blocks must be at ESS 5.3.2. If there are protocol nodes in the setup they must also be upgraded to ESS 5.3.2 levels (Underlying code + IBM Spectrum Scale 5.0.2.1 verified by **gssinstallcheck**).
2. A system must be completely healthy prior to the **mmvdisk** conversion and the MES migration.
3. The recovery groups must be converted to **mmvdisk** before MES can begin. See the Flow section that follows.
4. Heavy IBM Spectrum Scale and I/O operations must be suspended prior to the **mmvdisk** conversion and the MES operation.
5. Additional enclosures cannot span frames.
6. If space needs to be made, for example for moving of the EMS, this has to be planned for accordingly.
7. Legacy enclosures (PPC64BE) and the new 4U106 (LE) model are not supported.
8. LBS must wear an ESD wrist band when physically working on the hardware (like plugging in SAS cables).

MES upgrade considerations

- Do not try to configure call home before MES is complete, that is until resizing is done.
- You can perform additional MES upgrades while the DA's are rebalancing.
- You can restripe the file system while the DA's are rebalancing.
- Although it is recommended, you do not have to rebalance the file system if NSDs are added during MES.

SAS cable plug-in tips

- Unlatch the cable arm from the I/O server node into which you will be plugging in the SAS cable.
- Remove the blue cap.
- Make sure the location code label from the cable matches the port location code and the port number.
- Remove the cap from the SAS cable connector and plug it into the port.

- You should hear a click when the cable is inserted correctly.

SSR tasks

SSR is responsible for the following tasks.

1. Code 20 of the new enclosures - replacing parts as needed.
2. Running or labeling the new SAS cable connections.
3. Potentially making space in the frame - Moving the EMS.

SSR is not responsible for checking system health using **gssutils** like in a rackless or a rackful solution.

LBS tasks

LBS is responsible for the following tasks.

1. Upgrade of ESS 5.3.2 - prior to MES engagement.
2. **mmvdisk** conversion
3. Pre and post **mmvdisk** conversion and MES health checks.
4. Plugging the SAS cables into the adapters and enclosures.
5. Performing MES software functions such as conversion and resizing.
6. New storage management functions such as adding new space to existing file system and creating a new file system.
7. Restriping the file system.
8. Replacing any bad parts such as disks or cables.
9. Pre and post engagement operations

Flow

TDA process ensures that the customer is prepared for the MES upgrade. Considerations such as if there is enough room in the rack or usage of the file system space are planned out.

LBS

1. LBS performs normal ESS software upgrade. Customer must be at ESS 5.3.2 for MES. This upgrade is treated as a separate engagement than the future MES operation. No **mmvdisk** conversion is done at this time.

=== MES operation begins ===

SSR

1. The SSR arrives at the customer site. If the EMS needs to be moved, the SSR shuts down GPFS and powers down the server to move. For more information, see Appendix Q, "Shutting down and powering up ESS," on page 97.
2. The SSR places the new enclosures in the rack and establishes power connection. Based on the lights, the SSR performs a code 20 operation. If lights indicate any problem, they might need to take a service action.
3. The SSR runs the new SAS cable connections and labels in a bundle and hooks them to the frame. Later when LBS comes they simply plug in the connections when required in the flow.
4. The SSR places the EMS (if required) back into the existing frame or a new frame. Network connections and power are restored. Once the server is powered on, the SSR (or customer) can start GPFS to return the EMS back into the cluster.

LBS

1. Upon arrival LBS should first perform the normal upgrade-related system verification steps. Run the following from the EMS:

- a. **gnrhealthcheck** - This command determines if there are any issues in various areas of ESS. Any problems that show up must be addressed before MES starts.
 - b. **gssinstallcheck -G gss_ppc64** - This command checks the system to ensure all components match ESS 5.3.2 levels. If there are protocol nodes in the setup, check them using **gssinstallcheck -G ces_ppc64**
 - c. **/opt/ibm/gss/tools/samples/gssprecheck -N ems1 --upgrade --file /var/tmp/gssdeploy.cfg** - This command checks for common issues prior to doing any upgrade.
 - d. **mmhealth node show -N all --verbose** - This command shows any system health related issues to address.
 - e. Look for any events or tips that are open in the GUI. These also show up when you issue **mmhealth** but it is good to check in the GUI as well.
 - f. **gssinstallcheck -N ems1,gss_ppc64 --srv-events** - This command checks for any serviceable events reported from the Power 8 servers.
2. Convert to **mmvdisk**:
 - a. **gssgenclusterrgs -G gss_ppc64 --suffix=-hs --convert --legacy=no**
 - b. Wait for 5 minutes for daemons to recycle; file system stays up.
 3. Verification steps:
 - a. **mmvdisk nodeclass list** - This command shows if the **mmvdisk** node class exists.
 - b. **mmgetstate -a** - Issue this command to ensure that all daemons are active.
 - c. **mmlsmount all -L** - Issue this command to ensure that all mount points are still up.
 4. After these issues are resolved, MES upgrade can begin. Start by moving both recovery groups to **gssio2-hs**.
Move the recovery group in the current I/O server node to the peer I/O server node in the same building block.
 - a. To list the recovery groups and the current master server, run:
`mmvdisk recoverygroup list`
 - b. To move the recovery group from the current active I/O server node (**rg_gssio1-hs**) to the peer I/O server node (**gssio2-hs**) in the same building block, run the following commands in the shown order:
`mmvdisk recoverygroup change --recovery-group rg_gssio1-hs --active gssio2-hs`
Running **mmvdisk recoverygroup list** should show both RGs actively managed by **gssio2-hs**.
 5. Plug in the SAS cables for **gssio1** on the server and enclosure ends. Reboot the I/O node. Wait for 5 minutes for the node to reboot and paths to be rediscovered. Run the following commands to ensure that **gssio1** has discovered the new enclosures.
 - a. **gssstoragequickcheck -N gssio1**
 - b. **gssfindmissingdisks -N gssio1**
 Both commands should return with no issues and recognize the new enclosure and disk counts. The paths should also be without error. Once this is complete, start IBM Spectrum Scale on the node in question. Once IBM Spectrum Scale is active proceed to the next step.
 6. Move the recovery group ownership to **gssio1-hs**. Use the same commands as step 4. Make sure to use the correct RG names.
 7. Perform step 5 for **gssio2**.
 8. Rebalance both recovery groups
 - a. **mmvdisk recoverygroup change --recovery-group rg_gssio1-hs --active gssio1-hs**
 - b. **mmvdisk recoverygroup change --recovery-group rg_gssio2-hs --active gssio2-hs**
 - c. Check that the ownership has changed using the **mmvdisk recoverygroup list** command listed in step 4.

9. Perform step 5 again before starting the resize command. Use the group name to check all servers at once; instead of `-N gssioX` use `-G gss_ppc64`.
10. Update MES enclosure and drive firmware. If there are any issues, you should stop and replace any disks or enclosures that could not be updated for some reason.
 CurrentIoServer implies running the command from either of I/O server nodes in the MES building block.
 - a. `CurrentIoServer$ mmchfirmware --type storage-enclosure`
 - b. `CurrentIoServer$ mmchfirmware --type drive`
11. Add new storage into recovery groups.
`gssgenclusterrgs -N gssio3,gssio4 --suffix=-te0 --resize --legacy=no`
12. Verify that the new storage is available and the DA is rebalancing.
`mmvdisk recoverygroup list --recovery-group RG --all` - Run for both recovery groups.
 Note that the additional free space available in the DA and the background task for each DA is showing as rebalancing.
13. Start up the GUI and use **Edit rack components** to have the GUI discover the new topologies and make changes to the frame accordingly. Changes such as modify ESS model to consume more U space, move EMS, and so on.
14. Reconfigure call home.
`gsscallhomeconf -E ems1 -N ems1,gss_ppc64 --suffix=-hs --register all`

At this point, discussions with the customers need to be had on what to do with the free space (**gssgenvdisks**).

1. Add to the existing file system?
 - a. See the add building block flow for tips on creating new NSDs and adding to an existing file system.
 - b. Consider file system restripe at the end which might take time. (`mmrestripefs FileSystem -b`)
2. Create a new file system ?
 - See the installation section on how to use **gssgenvdisks** on creating a new file system.

Before LBS operation is complete, it is good to run the health checks in step #1 again besides **gssprecheck**. If everything is clean, the operation is complete.

Appendix M. Running gssinstallcheck in parallel

The **gssinstallcheck** command checks various aspects of the installation on all nodes. This command runs on each node sequentially. It has been enhanced such that you can run the **gssinstallcheck** command on all nodes in parallel.

It is advisable to run **gssinstallcheck** in parallel if the number of nodes in the cluster is more than 40 nodes. This is because running this command sequentially on such a large number of nodes takes a significant amount of time.

Note: Parallel **gssinstallcheck** can only be invoked from the management server node. Invoking **gssinstallcheck** parallelly from I/O server nodes will not work.

You can run **gssinstallcheck** in parallel as follows.

```
# xdsh ems1,gss_ppc64 "/opt/ibm/gss/tools/bin/gssinstallcheck -N localhost" | xcoll -n
```

In this command, **gssinstallcheck** is being run from the management server node and all I/O server nodes are a part of the gss_ppc64 xCAT group. Following is a sample output of this command. The output of all the nodes is grouped together if the **gssinstallcheck** output is same across nodes. In the following example, the output texts from gssio1 and gssio2 nodes are identical thus they have been grouped together in a single output. The ems1 output has been separately printed as the output of **gssinstallcheck** on the ems1 node is different. For more information, see the **xdsh** and **xcoll** command documentation.

```
=====
gssio1,gssio2
=====
Start of install check
xCAT objects not found for the nodelist localhost
nodelist:      localhost
...
...

=====
ems1
=====
Start of install check
xCAT objects not found for the nodelist localhost
nodelist:      localhost

Getting package information.
...
...
```

Appendix N. NTP setup

Ensure that NTP is configured on the management server node to act as an NTP server. The NTP service must be enabled and the chronyd service must be disabled.

Run the following from the management node.

1. Stop the NTP server on all ESS nodes.

```
xdsh ems,gss_ppc64 "systemctl stop ntpd"
xdsh ems,gss_ppc64 "systemctl enable ntpd"
```

2. Stop and disable the chronyd service.

```
xdsh ems,gss_ppc64 "systemctl stop chronyd"
xdsh ems,gss_ppc64 "systemctl disable chronyd"
```

3. Assign the management server node as the NTP server.

```
makentp
```

This command assigns the management server node as the NTP server. If there is any other NTP server already running then follow the Red Hat Enterprise Linux documentation to synchronize the time.

4. Edit the `/etc/ntp.conf` file accordingly.

```
# Use the local clock
server 127.127.1.0 prefer
fudge 127.127.1.0 stratum 10

restrict default kod nomodify notrap noquery nopeer
restrict 127.0.0.1

# Modify the line below to match your system
restrict 192.168.202.0 mask 255.255.255.0 nomodify notrap
driftfile /var/lib/ntp/ntp.drift
logfile /var/log/ntp.log
broadcastdelay 0.009
keys /etc/ntp/keys
```

5. Restart NTP on the management server.

```
systemctl restart ntpd
```

6. On the I/O server nodes, modify `/etc/ntp.conf` accordingly.

```
# Modify the line below to match your system
server 192.168.202.20
driftfile /var/lib/ntp/drift
disable auth
restrict 127.0.0.1
```

7. Synchronize time with the management server node.

```
ntpdate ems1
```

8. Start NTP.

```
systemctl start ntpd
```

Appendix O. Legacy deployment instructions

Use these steps for the legacy method of rebuilding the xCAT database from scratch and reinstalling the I/O server nodes.

Use these steps only if there are problems with the normal installation flow. These instructions are abbreviated. If you need the context, refer to the ESS 5.3.1.1 Quick Deployment Guide.

Rediscover and rebuild the xCAT database

1. If deploying on the PPC64LE platform, gather information for the `gssdeploy.cfg` configuration file using the following commands when you are in close proximity with the rack containing the nodes:

a. Scan the nodes in the FSP subnet range:

```
/var/tmp/gssdeploy -f FSP_Subnet_Range
```

FSP_Subnet_Range is the FSP management node interface subnet range. For example, 10.0.0.0/24.

Note:

- It is recommended to use the IP address 10.0.0.1 for the management interface, if possible.
- It is highly recommended that you use the /24 netmask because scanning of the subnet takes a considerable duration of time if a wider network range is used.
- The **gssdeploy -f** command first determines if a DHCP server is running on the network. If the DHCP server is not running, it prompts you to start one so that the I/O server nodes can obtain addresses. Select Y to start the DHCP server when prompted.
- This command scans the specified subnet range to ensure that only the nodes on which you want to deploy are available. These include I/O server nodes and management server node (EMS).
- This command returns the chassis serial numbers and FSP IP addresses of the EMS and I/O server nodes in the building block(s).
- There is a slight hang when **gssdeploy -f** attempts to query the FSP IP address configured on the EMS operating system. This operation eventually times out and fails which is the normal behavior. The only EMS FSP IP that should be discovered is the one assigned to HMC port 1.
- Do not proceed to the next step until FSP IP addresses and serial numbers of all known nodes are visible using the **gssdeploy -f** script.

b. Physically identify the nodes in the rack:

```
/var/tmp/gssdeploy -i
```

With the `-i` option, *Node_IP*, *Default_Password*, and *Duration* need to be provided as input, where:

- *Node_IP* is the returned FSP IPMI IP address of the node obtained by using the **gssdeploy -f** command.
- *Default_Password* is the FSP IPMI default password of the node, which is `PASSWORD`
- *Duration* is the time duration in seconds for which the LED on the node should blink.

After you issue this command, the LED blinks on the specified node for the specified duration. You can identify the node in the rack using the blinking LED.

Depending on the order of a node in the rack, its corresponding entry is made in the `gssdeploy.cfg` file. For example, for the bottommost node in the rack, its corresponding entry is put first in `gssdeploy.cfg`.

The main purpose of **gssdeploy -i** is to properly identify the slot of the ESS components within the IBM Spectrum Scale GUI. This is important for disk and other hardware replacements in the future. If using the default naming conventions, the bottom most server found in a frame is `gssio1`, then `gssio2`, and so on.

| **Note:** Upgrading to HMC SP3 might affect support for hardware call home.

| 2. Update the `gssdeploy.cfg` file according to your requirements and the gathered information.

| The options that you can specify in the `gssdeploy.cfg` file include:

- | • Whether use DVD for installation: `RHEL_USE_DVD`
- | The default option is to use ISO.
- | • If DVD, then device location: `RHEL_DVD`
- | • Mount point to use for RHEL media: `RHEL_MNT`
- | • ISO location: `RHEL_ISODIR`
- | The default location is `/opt/ibm/gss/iso`.
- | • ISO file name: `RHEL_ISO`
- | • EMS host name: `EMS_HOSTNAME`
- | • Network interface for xCAT management network: `EMS_MGTNETINTERFACE`
- | • Network interface for FSP network: `FSP_MGTNETINTERFACE` [**Not applicable for PPC64BE**]
- | • FSP default IPMI password: `FSP_PASSWD` [**Not applicable for PPC64BE**]
- | • HMC host name: `HMC_HOSTNAME` [**Not applicable for PPC64LE**]
- | • HMC default user ID: `HMC_ROOTUID` [**Not applicable for PPC64LE**]
- | • HMC default password: `HMC_PASSWD` [**Not applicable for PPC64LE**]
- | • Type of deployment: `DEPLOYMENT_TYPE`
- | The default type of deployment is ESS. It can also be `ADD_BB`.
- | ESS: Deploys I/O server nodes.
- | ADD_BB: Adds new building block of I/O server nodes.
- | PPC64LE protocol nodes can be deployed using the `CES` or `ADD_CES` deployment types.
- | CES: Deploys protocol nodes.
- | ADD_CES: Adds new protocol nodes.
- | • I/O server user ID: `IOSERVERS_UID`
- | • I/O server default password: `IOSERVERS_PASSWD`
- | • I/O server serial numbers: `IOSERVERS_SERIAL` [**Not applicable for PPC64BE**]
- | • I/O server node names: `IOSERVERS_NODES`
- | For example, `gssio1 gssio2`
- | • Deployment OS image: `DEPLOY_OSIMAGE`
- | • xCAT Group: `GSS_GROUP`
- | For example, `gss_ppc64`, `ces_ppc64`

| **Note:** Modification is only required when adding a protocol node to existing setup or adding an ESS I/O server node. A temporary group name is used for that operation.

| **Note:** For PPC64LE, there must be a one-to-one relationship between serial number and node in `gssdeploy.cfg` and for every node specified in `gssdeploy.cfg`, there must be a matching entry in `/etc/hosts`.

| 3. Copy the RHEL 7.5 ISO file to the directory specified in the `gssdeploy.cfg` file.

| 4. Perform precheck to detect any errors and address them before proceeding further:

```
| /opt/ibm/gss/tools/samples/gssprecheck -N ems1 --pre --install --file /var/tmp/gssdeploy.cfg
```

| **Note:** `gssprecheck` gives hints on ways to fix any discovered issues. It is recommended to review each found issue carefully though resolution of all might not be mandatory.

| **Attention:** Do the following steps before running `gssdeploy -x`.

- | • Power down the storage enclosures, or remove the SAS cables.
- | • Make sure that you update the `/etc/hosts` file with the xCAT node names and IP addresses that match with values defined in `gssdeploy.cfg`.
- | 5. Verify that the ISO is placed in the location specified in the `gssdeploy.cfg` configuration file and then run the **gssdeploy** script:
- | `/var/tmp/gssdeploy -x`
- | **Note:** To perform I/O server discovery task this step will power cycle the I/O server nodes specified in the `gssdeploy.cfg` file.
- | 6. Log out and then log back in to acquire the environment updates.
- | 7. Back up the xCAT database and save it to a location not on the management server node:
- | `dumpxCATdb -p /var/tmp/db`
- | `tar -zcvf xCATDB-backup.tar.gz /var/tmp/db`
- | 8. Set up the kernel, systemd, and Network Manager errata repositories. For example, use the following command on PPC64BE systems:
- | `/var/tmp/gssdeploy -k /home/deploy/kernel-RHSA-2018-3843-BE-ESS5321.tar.gz -p \`
- | `/home/deploy/systemd-532-RHBA-2018-2753-BE.tar.gz,/home/deploy/netmgr-RHSA-2018-3665-BE-ESS5321.tar.gz --silent`

| **Note:** This command extracts the supplied tar zip files and builds the associated repository.

- | • `-k` option: Set up the kernel repository
- | • `-p` option: Set up the patch repository (For example: systemd, network manager). One or more patches might be specified at the same time separated by comma.
- | • Directory structure:

```
| Kernel repository
| /install/gss/otherpkgs/rhels7/<arch>/kernel
| Patch repository
| /install/gss/otherpkgs/rhels7/<arch>/patch
```

| **Important:** Make sure that all RPMs in the `/install` directory including the extracted files in the kernel directory (`/install/gss/otherpkgs/rhels7/<arch>/kernel`), the patch directory (`/install/gss/otherpkgs/rhels7/<arch>/patch`), and xCAT RPMs, etc. have the correct read permission for user, group, and others (`chmod 644` files). For example:

```
| /install/gss/otherpkgs/rhels7/<arch>/kernel
| -rw-r--r-- 1 nobody nobody 44023568 Jan 11 10:17 kernel-3.10.0-862.25.3.e17.ppc64.rpm
| /install/gss/otherpkgs/rhels7/<arch>/patch
| -rw-r--r-- 1 nobody nobody 5412240 Oct 05 12:02 systemd-219-57.e17_5.3.ppc64.rpm
| -rw-r--r-- 1 nobody nobody 1899284 Jan 11 10:49 NetworkManager-1.12.0-8.e17_6.ppc64.rpm
```

| Wrong file permission will lead to node deployment failure.

| Deploy the I/O server nodes

- | 1. Before initiating the deployment of the I/O server nodes, do the following on the management server node:
 - | a. Verify that the running kernel level is at the desired level (for example, 862.25.3) using the **uname -a** command.
 - | b. Verify that there are no repository errors and all repositories are in place (patch, kernel, etc) using the **yum repolist** command.
 - | c. Ensure that the attached storage enclosures are powered off.
- | 2. Run the **gssinstallcheck** script:
- | `gssinstallcheck -N ems1`

| By default, **gssinstallcheck** runs on all nodes sequentially. You can run **gssinstallcheck** in parallel from the management server node as follows.

```
| # xdsH ems1,gss_ppc64 "/opt/ibm/gss/tools/bin/gssinstallcheck -N localhost" | xcoll -n
```

| For more information, see Appendix M, "Running gssinstallcheck in parallel," on page 85.

| This script is used to verify IBM Spectrum Scale profile, OFED, and kernel, etc.

| a. Check for any error with the following:

- | 1) Installed packages
- | 2) Linux kernel release
- | 3) OFED level
- | 4) IPR SAS FW
- | 5) IPR SAS queue depth
- | 6) System firmware
- | 7) System profile setting
- | 8) Host adapter driver

| Ignore other errors that may be flagged by the **gssinstallcheck** script. They will go away after the remaining installation steps are completed.

| 3. Run the **gssprecheck** script in full install mode and address any errors:

```
| /opt/ibm/gss/tools/samples/gssprecheck -N ems1 --install --file /var/tmp/gssdeploy.cfg
```

| **Note:** **gssprecheck** gives hints on ways to fix any discovered issues. It is recommended to review each found issue carefully though resolution of all might not be mandatory.

| 4. Deploy on the I/O server nodes using the customized deploy script:

```
|  
| ./gssdeploy -d
```

| 5. After a duration of about five minutes, run the following command:

```
|  
| nodestat gss_ppc64
```

| After running the command, the output displays the OS image name or packages being installed. For example:

| **PPC64LE installations:**

```
| node: rhels7.5-ppc64le-install-gss  
| node: rhels7.5-ppc64le-install-gss
```

| **PPC64BE installations:**

```
| node: rhels7.5-ppc64-install-gss  
| node: rhels7.5-ppc64-install-gss
```

| After about 30 minutes, the following output displays:

```
| node: sshd  
| node: sshd
```

| The installation is complete when **nodestat** displays sshd for all I/O server nodes. Here **gss_ppc64** is the xCAT node group containing I/O server nodes. To follow the progress of a node installation, you can tail the console log by using the following command:

```
| tailf /var/log/console/NodeName
```

| where *NodeName* is the node name.

| You can also use the following command to view the progress of the node installation:

```
| rcons NodeName -f
```

| To exit an rcons session, press Ctrl+E followed by C and then the period key (.).

| **Note:** Make sure the xCAT post-installation script is complete before rebooting the nodes. You can
| check xCAT post process running on the I/O server nodes as follows:

| `xdsh gss_ppc64 "ps -eaf | grep -v grep | grep xcatpost"`

| If there are any processes still running, wait for them to complete.

| 6. At the end of the deployment, wait for approximately five minutes and reboot the node:

| `xdsh gss_ppc64 systemctl reboot`

| 7. Once rebooted, verify the installation by running **gssinstallcheck**:

| `gssinstallcheck -G ems1,gss_ppc64`

| By default, **gssinstallcheck** runs on all nodes sequentially. You can run **gssinstallcheck** in parallel
| from the management server node as follows.

| `# xdsh gss_ppc64 "/opt/ibm/gss/tools/bin/gssinstallcheck -N localhost" | xcoll -n`

| For more information, see Appendix M, "Running gssinstallcheck in parallel," on page 85.

| Check for any error with the following:

- | a. Installed packages
- | b. Linux kernel release
- | c. OFED level
- | d. IPR SAS FW
- | e. IPR SAS queue depth
- | f. System firmware
- | g. System profile setting
- | h. Host adapter driver

| Ignore other errors that may be flagged by the **gssinstallcheck** script. They will go away after the
| remaining installation steps are completed.

Appendix P. Considerations for adding PPC64LE building blocks to ESS PPC64BE building blocks

When adding PPC64LE nodes to ESS PPC64BE systems, following considerations apply.

- All nodes must be at IBM Spectrum Scale 4.2.3.6 or later, and ESS 5.0.2 or later.
- The tested flow that is recommended for this procedure is available in the Box.
- Each architecture must have its own EMS node.
- PPC64LE must be the primary EMS node where GUI and performance collector services run.
- Shut down the GUI on the PPC64BE system and then change the performance collector to the PPC64LE EMS.
- Add the PPC64LE nodes as sensor nodes.
- You must have a flat network for the 1Gb xCAT. All nodes must be reachable and resolvable. The same consideration applies for the high speed network.
- Update and copy `/etc/hosts` to all nodes. Run `makedns` on both EMS nodes.
- Be mindful of pools. Do not mix SSD and HDD for instance. You might need to also set up policy files.
- Before starting the GUI, update the component database.
`mmaddcompspec default --replace`
- After starting the GUI, add the xCAT PPC64LE IP to the hardware monitoring list.
- Run the **Edit Rack Components** wizard in the GUI.

Note:

- | • In the **Hardware Monitoring** page of the GUI, add the IPs of all EMS' over the xCAT network.
- | • The collector must run on the same node as the GUI. Make sure to shut down the PPC64BE GUI and move the collector to the PPC64LE EMS before starting the GUI there.
- | • Before starting the collector on the PPC64LE EMS and GUI, back up and restore the data from the PPC64BE EMS (`/opt/IBM/zimon/data`).

Appendix Q. Shutting down and powering up ESS

The ESS components and frame may need to be powered off in cases such as data center maintenance, relocation, or emergencies. Use the following information to shut down and power up ESS.

Shutting down ESS

1. Verify that the file systems are not needed by users during the time the system will be unavailable.
2. If you are using a remote cluster to mount the ESS file system, unmount the file system by issuing the **mmumount** command from the remote client nodes.
3. Shut down the nodes using the **mmshutdown -N** command. For example:

```
mmshutdown -N ems1,gssio1,gssio2
```
4. If other nodes are attached and ESS nodes are the only quorum and manager nodes, it is recommended that you use the **mmshutdown -a** command to shut down the entire cluster.
5. Verify that IBM Spectrum Scale is shut down on the I/O nodes by issuing the **mmgetstate -a** command.
6. Power off the EMS and I/O nodes by issuing the **mmshutdown -h now** command on each individual node.

If you are using the Big Endian (BE) platform:

- a. The EMC LPAR, I/O node1 LPAR, and I/O node 2 LPAR will be shut down after you issue the **shutdown -h now**.
- b. Use the HMC to shut down the physical servers.
- c. Verify that the power light on the front of the frame is blinking after the LPARs are shut down.

If you are using the Big Endian (BE) platform and the HMC resides within this frame:

- a. Power off the HMC. If the HMC controls servers that are outside of this frame, plan appropriately before shutting down.

If you are using the Little Endian (LE) platform:

- a. The EMC LPAR, I/O node1 LPAR, and I/O node 2 LPAR will be completely shut down after you issue the **shutdown -h now** command.
- b. Verify that the power light on the front of the frame is blinking.
7. Power off all storage by flipping the power switches to off.
8. Before shutting off power to the frame, verify there are no components within the frame that are relied on by external infrastructure such as IB or Ethernet switches. If any of these exist and hardware outside of the frame needs access, plan appropriately before shutting off power to the frame.

Powering up ESS

1. Verify that power is connected to the frame.
2. Turn on all PDUs within the ESS frame.
3. Power on the components in the following order:

If you are using the Big Endian (BE) platform:

- a. Power on the HMC.
- b. Power on the storage drawers by flipping the power switches on each storage module to on.
- c. Power on the EMS node, I/O node 1 and I/O node 2.
- d. Wait for the HMC to come online and log in.
- e. Wait for the EMS node, I/O node 1 and I/O node 2 to be accessible to the HMC.
- f. Once the EMS sees that node, I/O node 1 and I/O node 2 are powered on, move to the LPAR view for each and power on the associated LPARs:

EMS LPAR

I/O node 1 LPAR

I/O node 2 LPAR

- g. Once all LPARs are powered on, ssh to the EMS node and verify that IBM Spectrum Scale has come online by issuing **mmgetstate -N ems1,gssio1,gssio2**. If IBM Spectrum Scale does not automatically start, start it manually by issuing **mmstartup -N ems1,gssio1,gssio2**.
- h. Issue the **gnrhealthcheck** and the **mmhealth cluster show** commands, and check the GUI event logs.

If you are using the Little Endian (LE) platform:

- a. Power on the storage drawers by flipping the power switches on each storage module to on.
- b. Power on the EMS node, I/O node 1 and I/O node 2.
- c. Once all LPARs are powered on, ssh to the EMS node and verify that IBM Spectrum Scale has come online by issuing **mmgetstate -N ems1,gssio1,gssio2**. If IBM Spectrum Scale does not automatically start, start it manually by issuing **mmstartup -N ems1,gssio1,gssio2**.
- d. Issue the **gnrhealthcheck** and the **mmhealth cluster show** commands, and check the GUI event logs.

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation North Castle Drive Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing Legal and Intellectual Property Law IBM Japan Ltd. 19-21,
Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
Dept. 30ZA/Building 707
Mail Station P300

2455 South Road,
Poughkeepsie, NY 12601-5400
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment or a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Intel is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

Java™ and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, and Windows NT are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Glossary

This glossary provides terms and definitions for the ESS solution.

The following cross-references are used in this glossary:

- *See* refers you from a non-preferred term to the preferred term or from an abbreviation to the spelled-out form.
- *See also* refers you to a related or contrasting term.

For other terms and definitions, see the IBM Terminology website ([opens in new window](http://www.ibm.com/software/globalization/terminology)):

<http://www.ibm.com/software/globalization/terminology>

B

building block

A pair of servers with shared disk enclosures attached.

BOOTP

See Bootstrap Protocol (BOOTP).

Bootstrap Protocol (BOOTP)

A computer networking protocol that is used in IP networks to automatically assign an IP address to network devices from a configuration server.

C

CEC *See central processor complex (CPC).*

central electronic complex (CEC)

See central processor complex (CPC).

central processor complex (CPC)

A physical collection of hardware that consists of channels, timers, main storage, and one or more central processors.

cluster

A loosely-coupled collection of independent systems, or *nodes*, organized into a network for the purpose of sharing resources and communicating with each other. *See also GPFS cluster.*

cluster manager

The node that monitors node status using disk leases, detects failures, drives recovery, and selects file system

managers. The cluster manager is the node with the lowest node number among the quorum nodes that are operating at a particular time.

compute node

A node with a mounted GPFS file system that is used specifically to run a customer job. ESS disks are not directly visible from and are not managed by this type of node.

CPC *See central processor complex (CPC).*

D

DA *See declustered array (DA).*

datagram

A basic transfer unit associated with a packet-switched network.

DCM *See drawer control module (DCM).*

declustered array (DA)

A disjoint subset of the pdisks in a recovery group.

dependent fileset

A fileset that shares the inode space of an existing independent fileset.

DFM *See direct FSP management (DFM).*

DHCP *See Dynamic Host Configuration Protocol (DHCP).*

direct FSP management (DFM)

The ability of the xCAT software to communicate directly with the Power Systems server's service processor without the use of the HMC for management.

drawer control module (DCM)

Essentially, a SAS expander on a storage enclosure drawer.

Dynamic Host Configuration Protocol (DHCP)

A standardized network protocol that is used on IP networks to dynamically distribute such network configuration parameters as IP addresses for interfaces and services.

E

Elastic Storage Server (ESS)

A high-performance, GPFS NSD solution

made up of one or more building blocks that runs on IBM Power Systems servers. The ESS software runs on ESS nodes - management server nodes and I/O server nodes.

ESS Management Server (EMS)

An xCAT server is required to discover the I/O server nodes (working with the HMC), provision the operating system (OS) on the I/O server nodes, and deploy the ESS software on the management node and I/O server nodes. One management server is required for each ESS system composed of one or more building blocks.

encryption key

A mathematical value that allows components to verify that they are in communication with the expected server. Encryption keys are based on a public or private key pair that is created during the installation process. See also *file encryption key (FEK)*, *master encryption key (MEK)*.

ESS See *Elastic Storage Server (ESS)*.

environmental service module (ESM)

Essentially, a SAS expander that attaches to the storage enclosure drives. In the case of multiple drawers in a storage enclosure, the ESM attaches to drawer control modules.

ESM See *environmental service module (ESM)*.

Extreme Cluster/Cloud Administration Toolkit (xCAT)

Scalable, open-source cluster management software. The management infrastructure of ESS is deployed by xCAT.

F

failback

Cluster recovery from failover following repair. See also *failover*.

failover

(1) The assumption of file system duties by another node when a node fails. (2) The process of transferring all control of the ESS to a single cluster in the ESS when the other clusters in the ESS fails. See also *cluster*. (3) The routing of all transactions to a second controller when the first controller fails. See also *cluster*.

failure group

A collection of disks that share common access paths or adapter connection, and could all become unavailable through a single hardware failure.

FEK See *file encryption key (FEK)*.

file encryption key (FEK)

A key used to encrypt sectors of an individual file. See also *encryption key*.

file system

The methods and data structures used to control how data is stored and retrieved.

file system descriptor

A data structure containing key information about a file system. This information includes the disks assigned to the file system (*stripe group*), the current state of the file system, and pointers to key files such as quota files and log files.

file system descriptor quorum

The number of disks needed in order to write the file system descriptor correctly.

file system manager

The provider of services for all the nodes using a single file system. A file system manager processes changes to the state or description of the file system, controls the regions of disks that are allocated to each node, and controls token management and quota management.

fileset A hierarchical grouping of files managed as a unit for balancing workload across a cluster. See also *dependent fileset*, *independent fileset*.

fileset snapshot

A snapshot of an independent fileset plus all dependent filesets.

flexible service processor (FSP)

Firmware that provides diagnosis, initialization, configuration, runtime error detection, and correction. Connects to the HMC.

FQDN

See *fully-qualified domain name (FQDN)*.

FSP See *flexible service processor (FSP)*.

fully-qualified domain name (FQDN)

The complete domain name for a specific computer, or host, on the Internet. The FQDN consists of two parts: the hostname and the domain name.

G

GPFS cluster

A cluster of nodes defined as being available for use by GPFS file systems.

GPFS portability layer

The interface module that each installation must build for its specific hardware platform and Linux distribution.

GPFS Storage Server (GSS)

A high-performance, GPFS NSD solution made up of one or more building blocks that runs on System x servers.

GSS See *GPFS Storage Server (GSS)*.

H

Hardware Management Console (HMC)

Standard interface for configuring and operating partitioned (LPAR) and SMP systems.

HMC See *Hardware Management Console (HMC)*.

I

IBM Security Key Lifecycle Manager (ISKLM)

For GPFS encryption, the ISKLM is used as an RKM server to store MEKs.

independent fileset

A fileset that has its own inode space.

indirect block

A block that contains pointers to other blocks.

inode The internal structure that describes the individual files in the file system. There is one inode for each file.

inode space

A collection of inode number ranges reserved for an independent fileset, which enables more efficient per-fileset functions.

Internet Protocol (IP)

The primary communication protocol for relaying datagrams across network boundaries. Its routing function enables internetworking and essentially establishes the Internet.

I/O server node

An ESS node that is attached to the ESS storage enclosures. It is the NSD server for the GPFS cluster.

IP See *Internet Protocol (IP)*.

IP over InfiniBand (IPoIB)

Provides an IP network emulation layer on top of InfiniBand RDMA networks, which allows existing applications to run over InfiniBand networks unmodified.

IPoIB See *IP over InfiniBand (IPoIB)*.

ISKLM

See *IBM Security Key Lifecycle Manager (ISKLM)*.

J

JBOD array

The total collection of disks and enclosures over which a recovery group pair is defined.

K

kernel The part of an operating system that contains programs for such tasks as input/output, management and control of hardware, and the scheduling of user tasks.

L

LACP See *Link Aggregation Control Protocol (LACP)*.

Link Aggregation Control Protocol (LACP)

Provides a way to control the bundling of several physical ports together to form a single logical channel.

logical partition (LPAR)

A subset of a server's hardware resources virtualized as a separate computer, each with its own operating system. See also *node*.

LPAR See *logical partition (LPAR)*.

M

management network

A network that is primarily responsible for booting and installing the designated server and compute nodes from the management server.

management server (MS)

An ESS node that hosts the ESS GUI and xCAT and is not connected to storage. It must be part of a GPFS cluster. From a system management perspective, it is the

central coordinator of the cluster. It also serves as a client node in an ESS building block.

master encryption key (MEK)

A key that is used to encrypt other keys. See also *encryption key*.

maximum transmission unit (MTU)

The largest packet or frame, specified in octets (eight-bit bytes), that can be sent in a packet- or frame-based network, such as the Internet. The TCP uses the MTU to determine the maximum size of each packet in any transmission.

MEK See *master encryption key (MEK)*.

metadata

A data structure that contains access information about file data. Such structures include inodes, indirect blocks, and directories. These data structures are not accessible to user applications.

MS See *management server (MS)*.

MTU See *maximum transmission unit (MTU)*.

N

Network File System (NFS)

A protocol (developed by Sun Microsystems, Incorporated) that allows any host in a network to gain access to another host or netgroup and their file directories.

Network Shared Disk (NSD)

A component for cluster-wide disk naming and access.

NSD volume ID

A unique 16-digit hexadecimal number that is used to identify and access all NSDs.

node An individual operating-system image within a cluster. Depending on the way in which the computer system is partitioned, it can contain one or more nodes. In a Power Systems environment, synonymous with *logical partition*.

node descriptor

A definition that indicates how IBM Spectrum Scale uses a node. Possible functions include: manager node, client node, quorum node, and non-quorum node.

node number

A number that is generated and maintained by IBM Spectrum Scale as the cluster is created, and as nodes are added to or deleted from the cluster.

node quorum

The minimum number of nodes that must be running in order for the daemon to start.

node quorum with tiebreaker disks

A form of quorum that allows IBM Spectrum Scale to run with as little as one quorum node available, as long as there is access to a majority of the quorum disks.

non-quorum node

A node in a cluster that is not counted for the purposes of quorum determination.

O

OFED See *OpenFabrics Enterprise Distribution (OFED)*.

OpenFabrics Enterprise Distribution (OFED)

An open-source software stack includes software drivers, core kernel code, middleware, and user-level interfaces.

P

pdisk A physical disk.

PortFast

A Cisco network function that can be configured to resolve any problems that could be caused by the amount of time STP takes to transition ports to the Forwarding state.

R

RAID See *redundant array of independent disks (RAID)*.

RDMA

See *remote direct memory access (RDMA)*.

redundant array of independent disks (RAID)

A collection of two or more disk physical drives that present to the host an image of one or more logical disk drives. In the event of a single physical device failure, the data can be read or regenerated from the other disk drives in the array due to data redundancy.

recovery

The process of restoring access to file

system data when a failure has occurred. Recovery can involve reconstructing data or providing alternative routing through a different server.

recovery group (RG)

A collection of disks that is set up by IBM Spectrum Scale RAID, in which each disk is connected physically to two servers: a primary server and a backup server.

remote direct memory access (RDMA)

A direct memory access from the memory of one computer into that of another without involving either one's operating system. This permits high-throughput, low-latency networking, which is especially useful in massively-parallel computer clusters.

RGD See *recovery group data (RGD)*.

remote key management server (RKM server)

A server that is used to store master encryption keys.

RG See *recovery group (RG)*.

recovery group data (RGD)

Data that is associated with a recovery group.

RKM server

See *remote key management server (RKM server)*.

S

SAS See *Serial Attached SCSI (SAS)*.

secure shell (SSH)

A cryptographic (encrypted) network protocol for initiating text-based shell sessions securely on remote computers.

Serial Attached SCSI (SAS)

A point-to-point serial protocol that moves data to and from such computer storage devices as hard drives and tape drives.

service network

A private network that is dedicated to

managing POWER8 servers. Provides Ethernet-based connectivity among the FSP, CPC, HMC, and management server.

SMP See *symmetric multiprocessing (SMP)*.

Spanning Tree Protocol (STP)

A network protocol that ensures a loop-free topology for any bridged Ethernet local-area network. The basic function of STP is to prevent bridge loops and the broadcast radiation that results from them.

SSH See *secure shell (SSH)*.

STP See *Spanning Tree Protocol (STP)*.

symmetric multiprocessing (SMP)

A computer architecture that provides fast performance by making multiple processors available to complete individual processes simultaneously.

T

TCP See *Transmission Control Protocol (TCP)*.

Transmission Control Protocol (TCP)

A core protocol of the Internet Protocol Suite that provides reliable, ordered, and error-checked delivery of a stream of octets between applications running on hosts communicating over an IP network.

V

VCD See *vdisk configuration data (VCD)*.

vdisk A virtual disk.

vdisk configuration data (VCD)

Configuration data that is associated with a virtual disk.

X

xCAT See *Extreme Cluster/Cloud Administration Toolkit*.



Printed in USA

SC27-9526-01

