Elastic Storage Server
Version 3.5

# *Quick Deployment Guide*

SC27-8580-00

Elastic Storage Server
Version 3.5

# *Quick Deployment Guide*

SC27-8580-00

**Note**
Before using this information and the product it supports, read the information in "Notices" on page 13.

This edition applies to version 3.5 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 viii. 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.

# Contents

# Tables

# About this information

This information guides you in quickly installing, or upgrading to, version 3.5 of the Elastic Storage Server (ESS).

For detailed ESS installation and upgrade information, see *Deploying the Elastic Storage Server*.

## Who should read this information

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

## Prerequisite and related information

### ESS information

The ESS 3.5 library consists of these information units:

- *Deploying the Elastic Storage Server*, SC27-6659
- *Elastic Storage Server: Quick Deployment Guide*, SC27-8580
- *IBM Spectrum Scale RAID: Administration*, SC27-6658

For more information, see IBM Knowledge Center:

**http://www.ibm.com/support/knowledgecenter/SSYSP8_3.5.0/sts35_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/sts_welcome.html**

### Related information

For information about:

- IBM Spectrum Scale, see IBM Knowledge Center:
  **http://www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale_welcome.html**
- IBM POWER8 servers, see IBM Knowledge Center:
  **http://www.ibm.com/support/knowledgecenter/POWER8/p8hdx/POWER8welcome.htm**
- The DCS3700 storage enclosure, see:
  - *System Storage DCS3700 Quick Start Guide*, GA32-0960-03:
    **http://www.ibm.com/support/docview.wss?uid=ssg1S7004915**
  - *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**
- 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**
- Extreme Cluster/Cloud Administration Toolkit (xCAT), go to the xCAT website:
  **http://sourceforge.net/p/xcat/wiki/Main_Page/**

## 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.<br><br>Depending on the context, **bold** typeface sometimes represents path names, directories, or file names. |
| **bold underlined** | **bold underlined** 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.<br><br>Depending on the context, `constant-width` typeface sometimes represents path names, directories, or file names. |
| *italic* | *Italic* words or characters represent variable values that you must supply.<br><br>*Italics* 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 *Enter*. |
| \ | In command examples, a backslash indicates that the command or coding example continues on the next line. For example:<br>`mkcondition -r IBM.FileSystem -e "PercentTotUsed > 90" \`<br>`-E "PercentTotUsed < 85" -m p "FileSystem space used"` |
| {*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**>. |
| *item*... | Ellipses indicate that you can repeat the preceding item one or more times. |
| \| | In *synopsis* statements, vertical lines separate a list of choices. In other words, a vertical line means *Or*.<br><br>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**

# Deploying the Elastic Storage Server - for experienced users

This topic includes a concise set of deployment instructions for those who are familiar with Elastic Storage Server (ESS) systems.

In these instructions:

- All version numbers shown are examples. The version depends on the release and edition that is being deployed.
- Node names `ems1`, `gssio1`, and `gssio2` are examples. Each environment could have its own unique naming conventions.

Follow these high-level steps:

1. Complete the prerequisite tasks.
2. Install the management server software.
3. Complete one of the following tasks:
   a. Install the ESS system.
   b. Upgrade the ESS system.

### Complete the prerequisite tasks

Complete these tasks before proceeding:

1. Obtain a Red Hat Enterprise Linux 7.1 ISO image file or DVD for 64-bit IBM Power Systems architecture.
2. Obtain the ESS software archive.
3. Review the list of known issues for the ESS version you are installing.
   See "Known issues," on page 9 for more information.

### Install the management server software

1. Unpack the ESS software archive:

   ```
   tar -zxvf gss_install-3.5.0_ppc64_advanced_20151009T141843Z.tgz
   ```

2. Check the MD5 checksum:

   ```
   md5sum -c gss_install-3.5.0_ppc64_advanced_20151009T141843Z.md5
   ```

3. Make sure the `/opt/ibm/gss/install` directory is clean:

   ```
   /bin/sh gss_install-3.5.0_ppc64_advanced_20151009T141843Z --remove
   ```

4. Extract the ESS packages and accept the license:

   ```
   /bin/sh gss_install-3.5.0_ppc64_advanced_20151009T141843Z --text-only
   ```

## Install the ESS system

Follow these steps to perform a new installation of the ESS software on a management server node and I/O server nodes.

1. Copy the **gssdeploy** script and customize it for your environment by editing it:

   ```
   cp /opt/ibm/gss/install/samples/gssdeploy /var/tmp
   chmod +x /var/tmp/gssdeploy
   ```

`/var/tmp` is a sample directory name. You can specify a different directory name. The **gssdeploy** script uses a directory called `/tmp/gssdeploy`, so do *not* copy the script to `/tmp`.

2. Clean the current xCAT installation and associated configuration:

   ```
   ./gssdeploy -c
   ```

3. Update the ESS repositories on the management server node:

   ```
   cd /opt/ibm/gss/install
   installer/gssinstall -m manifest -u
   ```

4. Run the **gssdeploy** script:

   ```
   ./gssdeploy -x
   ```

5. Log out and then log back in to acquire the environment updates.
6. Update the management server node:

   ```
   updatenode ems1 -P gss_updatenode
   ```

   Reboot the management server node and run this step again if you are instructed to do so.
7. Update OFED on the management server node:

   ```
   updatenode ems1 -P gss_ofed
   ```

8. Reboot the management server node.

## Deploy the I/O server nodes

1. Before initiating the deployment of the I/O server nodes, verify that the JBODs are powered off.
2. Deploy on the I/O server nodes using the customized deploy script:

   ```
   ./gssdeploy -d
   ```

3. Run:

   ```
   nodestat gss_ppc64
   ```

   The installation is complete when **nodestat** displays `sshd` as the status for all I/O server nodes.
4. At the end of the deployment, reboot the node:

   ```
   xdsh gss_ppc64 systemctl reboot
   ```

## Check the system hardware

After the I/O server nodes have been installed successfully, power on the JBODs. Wait approximately five to 10 minutes from power on for discovery to complete before moving on to the next step.

1. Run **gssstoragequickcheck**:

   ```
   gssstoragequickcheck -G gss_ppc64
   ```

2. Run lsiFixNV:

   ```
   xdsh gss_ppc64 "/xcatpost/gss_sashba"
   ```

3. Run **gssfindmissingdisks**:

   ```
   gssfindmissingdisks -G gss_ppc64
   ```

   If **gssfindmissingdisks** display an error, run **mmgetpdisktopology** and pipe it to **topsummary** on each I/O server node to obtain more information about the error:

   ```
   mmgetpdisktopology | topsummary
   ```

4. Run **gsscheckdisks**:

   ```
   GSSENV=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 **GSSENV** environment variable must be set to **INSTALL** to indicate that **gsscheckdisks** is being run during installation.

## Set up the high-speed network

Set up the high-speed network that will be used for the cluster data communication. Copy the `/etc/hosts` to all nodes in the building block. After the `/etc/hosts` file is properly set with high-speed IP addresses and corresponding hostnames, you can use the **gssgennetworks** script, which is located in the `/opt/ibm/gss/tools/samples` directory, to create a bonded Ethernet network.

1. To see the current set of active (up) interfaces, run:

   ```
   /opt/ibm/gss/tools/samples/gssgennetworks -G gss_ppc64
   ```

2. To create a bonded interface, run:

   ```
   /opt/ibm/gss/tools/samples/gssgennetworks -G gss_ppc64 --create_bond
   ```

   The script sets `miimon` to `100`, the bonding mode to `802.3ad` (LACP), and `xmit_hash_policy` to `layer3+4`. The other bond options are left with the default values, including `lacp_rate` (the default is `slow`). For proper network operation, the Ethernet switch settings in the networking infrastructure must match the I/O server node interface bond settings.

See *Deploying the Elastic Storage Server* for:

- more information about **gssgennetworks**
- information about creating a bonded InfiniBand network.

## 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 hostname.

2. Create the recovery groups:

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

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

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

4. 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** on the management server:
   ```
   gssinstallcheck -N ems1
   ```
2. Run **gssinstallcheck** on the I/O server nodes:
   ```
   gssinstallcheck -G gss_ppc64
   ```
3. 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 EMS node:
      ```
      systemctl reboot
      ```
4. After reboots, run:
   ```
   gssinstallcheck -G gss_ppc64 | grep phy
   ```
   Ensure that the phy mapping check is OK.
5. Restart GPFS in all nodes and wait for all nodes to become active:
   ```
   mmstartup -a
   ```
6. 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.
7. Perform a health check. Run:

   ```
   gnrhealthcheck
   ```

   Address any issues that are identified.

### Install the ESS GUI

1. Start the ESS GUI:

   ```
   systemctl start gpfsgui
   ```
2. Complete the system setup wizard.

---

## Upgrade the ESS system

Follow these steps to perform an upgrade of the ESS system.

### Prepare the system for upgrade

1. Perform a health check. Run:

   ```
   gnrhealthcheck
   ```

   Address any issues that are identified.
2. Wait for any of these commands that are performing file system maintenance tasks to complete:
   **mmadddisk**
   **mmapplypolicy**
   **mmcheckquota**
   **mmdeldisk**

**mmfsck**
**mmlssnapshot**
**mmrestorefs**
**mmrestripefile**
**mmrestripefs**
**mmrpldisk**

3. It is recommended that you stop the creation and deletion of snapshots using **mmcrsnapshot** and **mmdelsnapshot** during the upgrade window.

## Upgrading from ESS 2.5.n

Perform the following steps if you are upgrading from ESS 2.5.*n*:

1. Copy the **gssdeploy** script and customize it for your environment by editing it. For example, run these commands:

   ```
   cp /opt/ibm/gss/install/samples/gssdeploy /var/tmp
   chmod +x /var/tmp/gssdeploy
   ```

   The `/var/tmp` directory is an example. You can specify a different directory name. Do *not* copy the script to `/tmp`, because **gssdeploy** uses a directory named `/tmp/gssdeploy`.
2. Clean the current xCAT installation and associated configuration:

   ```
   ./gssdeploy -c -r Directory
   ```

   *Directory* is the xCAT database dump directory that will be used for the restore operation.
3. Update ESS repositories on the management server node:

   ```
   cd /opt/ibm/gss/install
   installer/gssinstall -m manifest -u
   ```

   This step was run previously, but because a clean operation was needed for an ESS 2.5.*n* update, it needs to be run again.
4. Run the **gssdeploy** script:

   ```
   ./gssdeploy -x -r Directory
   ```
5. Go to the section titled Update the management server node.

## Upgrading from ESS 3.x.y

Perform the following steps if you are upgrading from ESS 3.x.y:

1. Update ESS repositories on the management server node:

   ```
   cd /opt/ibm/gss/install
   installer/gssinstall -m manifest -u
   ```
2. Go to the next section titled Update the management server node.

## Update the management server node

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

   ```
   systemctl stop gpfsgui
   ```
2. Save `ZIMonSensors.cfg` for later reference:

   ```
   cp /opt/IBM/zimon/bin/ZIMonSensors.cfg /tmp
   ```
3. Shut down IBM Spectrum Scale on the management server node while making sure quorum is still maintained. Run:

```
mmshutdown
```

4. Update the management server node:
   ```
   updatenode ems1 -P gss_updatenode
   ```

   Reboot the management server node and complete this step again if you are instructed to do so in the **updatenode** output.
5. Update OFED on the management server node:

   ```
   updatenode ems1 -P gss_ofed
   ```
6. Reboot the management server node.
7. Update the IBM Spectrum Scale configuration parameters for the management server. Use the cluster node name (e.g., ems1-hs) as opposed to the xCAT object name (e.g., ems1). Run:

   ```
   /opt/ibm/gss/tools/samples/gssupg350.sh -c ems1-hs
   ```
8. Start IBM Spectrum Scale on the management server node:
   ```
   mmstartup
   ```

## Update the I/O server nodes

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

1. Move the cluster and file system manager role to another node if the current node is a cluster manager or file system manager.
   a. To find the cluster and file system managers, run:

      ```
      mmlsmgr
      ```
   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.
2. 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, run:
      ```
      mmlsrecoverygroup
      ```
   b. To list the active server, primary server, and secondary server, run:

      ```
      mmlsrecoverygroup 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:

      ```
      mmchrecoverygroup rg_gssio1_hs --servers gssio2-hs,gssio1-hs
      ```
3. After confirming that the recovery group has been successfully moved to the peer I/O server node, shut down IBM Spectrum Scale on the current I/O server node while maintaining quorum:

   ```
   mmshutdown -N CurrentIoServer
   ```
4. Run **updatenode**:

   ```
   updatenode CurrentIoServer -P gss_updatenode
   ```

Reboot the I/O server node and complete this step again if you are instructed to do so in the **updatenode** output.

5. Update OFED.

   ```
   updatenode CurrentIoServer -P gss_ofed
   ```

6. Reboot the I/O server node.
7. Update the SAS host adapter firmware on *CurrentIoServer*:

   ```
   CurrentIoServer$ mmchfirmware --type host-adapter
   ```

8. Update the IBM Spectrum Scale configuration parameters on *CurrentIoServer*. Use the cluster node name (e.g., gssio1-hs) as opposed to the xCAT object name (e.g., gssio1). Run:

   ```
   /opt/ibm/gss/tools/samples/gssupg350.sh -s CurrentIoServer
   ```

9. Run phy check and ensure that the phy mapping is OK:

   ```
   gssinstallcheck -N CurrentIo server | grep phy
   ```

10. Start IBM Spectrum Scale on the I/O server node. Move back the cluster manager and the file system manager if required. Move the failback recovery group back to the current I/O server node. Run:

    ```
    mmstartup -N CurrentIoServer
    ```

## Update the enclosure and drive firmware

1. To update the storage enclosure firmware, run the following command from one of the nodes of the building block for each building block, for example:

   ```
   IOServer$ mmchfirmware --type storage-enclosure
   ```

2. The update of the drive firmware can be performed from one of the nodes in the building block. Run:

   ```
   IOServer$ mmchfirmware --type drive -N gss_ppc64
   ```

   The `-N gss_ppc64` option (or a list of I/O server nodes) must be specified to run the firmware upgrade for drives in all of the building blocks.

   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.

## Perform the ESS 3.5 configuration update

1. To enable the IBM Spectrum Scale callbacks, run:

   ```
   /opt/ibm/gss/tools/samples/gssupg350.sh -b
   ```

2. To copy the new disk preparation script, run:

   ```
   /opt/ibm/gss/tools/samples/gssupg350.sh -p
   ```

## Check the installed software and system health

1. Run **gssinstallcheck** on the management server:

   ```
   gssinstallcheck -N ems1
   ```

2. Run **gssinstallcheck** on the I/O server nodes:

   ```
   gssinstallcheck -G gss_ppc64
   ```

3. Perform a health check. Run:

   ```
   gnrhealthcheck
   ```

   Address any issues that are identified.

## Start the ESS GUI

1. Verify `/opt/IBM/zimon/bin/ZIMonSensors.cfg`. Refer to the saved copy of `ZIMonSensors.cfg` if needed.
2. Start the ESS GUI on the management server node:

   ```
   systemctl start gpfsgui
   ```

# Appendix. Known issues

This topic includes known issues for ESS.

## ESS 3.5.2 issues

Table 2 includes information about known issues in ESS 3.5 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 3.5.2.*

| Issue | Environment affected | Description | Resolution or action |
|---|---|---|---|
| **1.** After upgrading to to ESS 3.5.2, the Java 1.6 rpm still exists which may fail some security scans. | Cluster software upgrade<br><br>Type: Upgrade<br><br>IBM Spectrum ScaleVersion: Advanced or Standard<br><br>Affected nodes: I/O + EMS | ESS<br><br>ESS 3.5.2 no longer requires a separate java rpm (java is packaged within the ESS GUI rpm). The java rpm remains on the system and may cause some security scans to fail. The workaround is to simply remove the rpm via yum.<br><br>The rpm in question is:<br><br>java-1.6.0-ibm-1.6.0.16.7-1jpp.ppc64 | To work around this issue, do the following:<br>1. Run:<br>`xdsh ems1,gss_ppc64 "yum -y remove java-1.6*"`<br>2. Confirm the removal. Run:<br>`xdsh ems1,gss_ppc64 "rpm -qa \| grep -i java-1.6"`<br>No output should be returned |

*Table 2. Known issues in ESS 3.5.2. (continued)*

| Issue | Environment affected | Description | Resolution or action |
|---|---|---|---|
| **2. gssinstallcheck** may flag an error regarding page pool size in multi-building block situations if the physical memory sizes differ. | Software validation<br><br>Type: Install or Upgrade<br><br>IBM Spectrum Scale Version: Advanced or Standard<br><br>Affected nodes: I/O | **gssinstallcheck** is a new tool introduced in ESS 3.5, that helps validate software, firmware, and configuration settings.<br><br>If adding (or installing) building blocks of a different memory footprint installcheck will flag this as an error. Best practice states that your IO servers should all have the same memory footprint thus pagepool value. Page pool is currently set at ~60% of physical memory per IO node.<br><br>Example from gssinstallcheck:<br>`[ERROR] pagepool: found 142807662592 expected range 147028338278 - 179529339371` | 1. Confirm each IO node's individual memory footprint. From the EMS, run the following command against your IO xCAT group:<br>`xdsh gss_ppc64 "cat/proc/meminfo \| grep MemTotal"`<br><br>**Note:** This value is in KB<br>If the physical memory varies between servers and/or building blocks, consider adding mrmory and re-calculating pagepool to ensure consistency.<br>2. Validate the pagepool settings in Spectrum Scale:<br>`mmlsconfig \| prep -A 1 pagepool`<br><br>**Note:** This value is in MB.<br>If the pagepool value setting is not roughly ~60% of physical memory, then you should consider recalculating and setting an updated value.<br>Please refer to the Spectrum Scale documentation on how to update the pagepool value.<br>http://www-01.ibm.com/support/knowledgecenter/SSFKCN/gpfs_welcome.html |

*Table 2. Known issues in ESS 3.5.2. (continued)*

| Issue | Environment affected | Description | Resolution or action |
|---|---|---|---|
| **3.** Spectrum Scale filesystems may fail to mount on Red Hat 7.1 nodes due to a systemd issue. | Cluster<br><br>filesystem<br><br>Type: Install or Upgrade<br><br>Version: Advanced or Standard<br><br>Affected Nodes: IO +EMS | ESS 3.0.X or higher contains Spectrum Scale 4.1.0.8 which has a problem mounting on Red Hat 7.1 due to an issue with systemd. Updating cluster nodes to the latest systemd packages will correct the issue.<br><br>As of ESS 3.5.X the systemd issue may no longer present itself. If you are still unable to mount Spectrum Scale filesystems proceed with the workaround to correct this issue. | ESS 3.0.X or higher contain Spectrum Scale 4.1.0.8 which has a problem mounting on Red Hat 7.1 due to an issue with systemd. Updating cluster nodes to the latest systemd packages will correct the issue.<br><br>As of ESS 3.5.X the systemd issue may no longer present itself. If you are still unable to mount Spectrum Scale filesystems proceed with the workaround to correct this issue.<br><br>To workaround this issue please do the following:<br>• Connect your cluster nodes to the Red Hat network (RHN) and apply the RHBA-2015-0738 errata prior to mounting your GPFS filesystems.<br>• If your unable to connect to the RHN directly you may download the required rpms and updated manually on each node or create a local yum repository.<br><br>The minimum required rpms and levels are:<br><br>systemd-sysv-208-20.el7_1.2.ppc64.rpm<br><br>systemd-libs-208-20.el7_1.2.ppc64.rpm<br><br>systemd-208-20.el7_1.2.ppc64.rpm<br><br>libgudev1-208-20.el7_1.2.ppc64.rpm<br><br>The advisory can be found here:<br><br>https://rhn.redhat.com/errata/RHBA-2015-0738.html |
| **4.** If upgrading from ESS 3.5.1 to ESS 3.5.2,gssinstall on the EMS node will flag one of the "Group gui RPMs" as 'Old". | Cluster software upgrade<br><br>Type: Upgrade<br><br>Version: Standard or Advanced<br><br>Affected Nodes: EMS | Users are required to run gssinstall to configure the repositories required for installation or upgrade. Group RPMs marked 'Old' normally indicate that there is a higher software version already installed, thus one should stop to investigate. One of the GUI RPMs is called out as a false-positive. | The workaround is to proceed as normal. During the upgrade flow the correct GUI rpm will be installed. Gssinstallcheck should be used at the end of the upgrade to validate that all required rpms are in place. You should not see any 'OLD' or 'NEW' RPMS displayed in the output for each node. |

# 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

## B

**building block**
A pair of servers with shared disk enclosures attached.

**BOOTP**
See *Bootstrap Protocol (BOOTP)*.

**Bootstrap Protocol (BOOTP)**
A computer networking protocol thst 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.

**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 provices 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 can 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

SC27-8580-00