

IBM Elastic Storage System

6.1.4.1



ESS Quick Deployment Guide

This guide is for system administrators, installers, and programmers of IBM Spectrum Scale clusters who are experienced with the operating systems on which each IBM Spectrum Scale cluster is based.

ESS quick deployment sheet

This quick sheet lists concise sets of steps for the ESS deployment procedures. For more information, see respective procedures.

<p style="text-align: center;">COMMON INSTRUCTIONS</p> <ol style="list-style-type: none"> Download the package from IBM Fix Central. Populate the <code>/etc/hosts</code> file on the BMS node. Extract the package. <code>ess??00_6.1.x.x_???-??_dme_ppc64le.sh --start-container</code> Enter the following information: <ul style="list-style-type: none"> Is the BMS hostname correct? Type container hostname. Type the container FSP IP address. Run <code>essrun config load</code> with the password set by IBM SSR on your setup. <code>essrun -N ems1,io1,io2,io3,io4 config load -p ibmesscluster</code> Run <code>essrun config check</code> with the password set by IBM SSR on your setup. <code>essrun -N ems1,io1,io2,io3,io4 config check -p ibmesscluster</code> <p style="text-align: center;">AFTER SPECIFIC BLOCK HAS BEEN RUN</p> <ol style="list-style-type: none"> Check health of the system, run the command. <code>essrun -N io1,io2,io3,io4 healthcheck</code> 	<p style="text-align: center;">NEW DEPLOYMENT</p> <ol style="list-style-type: none"> If the BMS node needs to be updated, run this command (You might be prompted to run the command again, if the kernel is changed). <code>essrun -N ems1 update --offline</code> If canisters or I/O nodes need to be updated, run this command. <code>essrun -N io1,io2,io3,io4 update --offline --no-fw-update</code> Create network bonds. <code>essrun -N ems1,io1,io2,io3,io4 network --suffix=hs</code> Set up the NTP server and clients. <code>essrun -N io1,io2,io3,io4 timeset client --server-ip ntp.server.ip</code> Create the cluster. <code>essrun -N io1,io2,io3,io4 cluster --suffix=hs</code> Add the BMS node to the cluster. <code>essrun -N io1 cluster --add-ems ems1 --suffix=hs</code> Create the file system. <code>essrun -N io1,io2,io3,io4 filesystem --suffix=hs</code> If needed, update the firmware; on each canister/I/O node, run these commands. <code>mmchfirmware --type storage-enclosure</code> <code>mmchfirmware --type drive</code> Configure GUI hardware monitoring <code>essrun -N ems1,io1,io2,io3,io4 gui enable --configure</code> Set up call home. For more information about the call home configuration, see Configuring call home on ESS systems.
<p style="text-align: center;">OFFLINE UPDATE</p> <ol style="list-style-type: none"> Update the BMS node; from the BMS node, not from the container, run this command. <code>mmshutdwn -a</code> <code>systemctl stop gpfsogui</code> Update the BMS node; from the container, run this command [You may be prompted to run a second time if the kernel is changed]. <code>essrun -N ems1 update --offline</code> If the POWER firmware needs to be updated, refer to the online documentation. After the BMS node is updated, from the container, run this command. <code>essrun -N io1,io2,io3,io4 update --offline --no-fw-update</code> From each canister or I/O node, run this command. <code>mmchfirmware --type storage-enclosure</code> <code>mmchfirmware --type drive</code> Start IBM Spectrum Scale; from the BMS node, run this command. <code>mmstartup -a</code> Enable <code>gpfsui</code> start on reboot. <code>systemctl enable gpfsui</code> Start the GUI; from the BMS node, not from the container, run this command. <code>systemctl start gpfsui</code> 	<p style="text-align: center;">ONLINE UPDATE</p> <ol style="list-style-type: none"> Update the BMS node; from the BMS node, not from the container, run this command. <code>mmshutdwn</code> <code>systemctl stop gpfsogui</code> Update the BMS node; from the container, run this command [You may be prompted to run a second time if the kernel is changed]. <code>essrun -N ems1 update --offline</code> Start IBM Spectrum Scale on the BMS node. <code>mmstartup</code> If the POWER firmware needs to be updated, refer to the documentation. After the BMS node is online, from the container, run this command. <code>essrun -N io1,io2,io3,io4 update</code> Enable <code>gpfsui</code> start on reboot. <code>systemctl enable gpfsui</code> Start the GUI; from the BMS node, not from the container, run this command. <code>systemctl start gpfsui</code> Rerun the <code>esscallhomeconf</code> command with the same arguments that you used when you configured the system. For more information see, Configuring call home on ESS systems.

Support matrix

Release	OS	Runs on	Can upgrade or deploy
ESS 3500 6.1.4	Red Hat® Enterprise Linux® 8.6 (x86_64)	POWER9™ EMS	<ul style="list-style-type: none"> ESS 3500 nodes POWER9 EMS POWER9 protocol nodes
ESS 3200 6.1.4	Red Hat Enterprise Linux 8.6 (x86_64)	<ul style="list-style-type: none"> POWER9 EMS 	<ul style="list-style-type: none"> ESS 3200 nodes POWER9 EMS POWER9 protocol nodes

Release	OS	Runs on	Can upgrade or deploy
ESS 3000 6.1.4	Red Hat Enterprise Linux 8.6 (x86_64)	<ul style="list-style-type: none"> POWER8® EMS POWER9 EMS 	<ul style="list-style-type: none"> ESS 3000 nodes POWER8 EMS POWER9 EMS POWER8 protocol nodes POWER9 protocol nodes
ESS 5000 6.1.4	Red Hat Enterprise Linux 8.6 (x86_64)	<ul style="list-style-type: none"> POWER9 EMS 	<ul style="list-style-type: none"> ESS 5000 nodes POWER9 EMS POWER9 protocol nodes
ESS Legacy 6.1.4	<ul style="list-style-type: none"> Red Hat Enterprise Linux 8.6 (PPC64LE) Red Hat Enterprise Linux 7.9 (PPC64LE) 	<ul style="list-style-type: none"> POWER8 EMS POWER9 EMS 	<ul style="list-style-type: none"> ESS POWER8 I/O nodes (PPC64LE) ESS POWER8 protocol nodes (PPC64LE) ESS POWER9 protocol nodes (PPC64LE)* POWER8 EMS POWER9 EMS

Network requirements

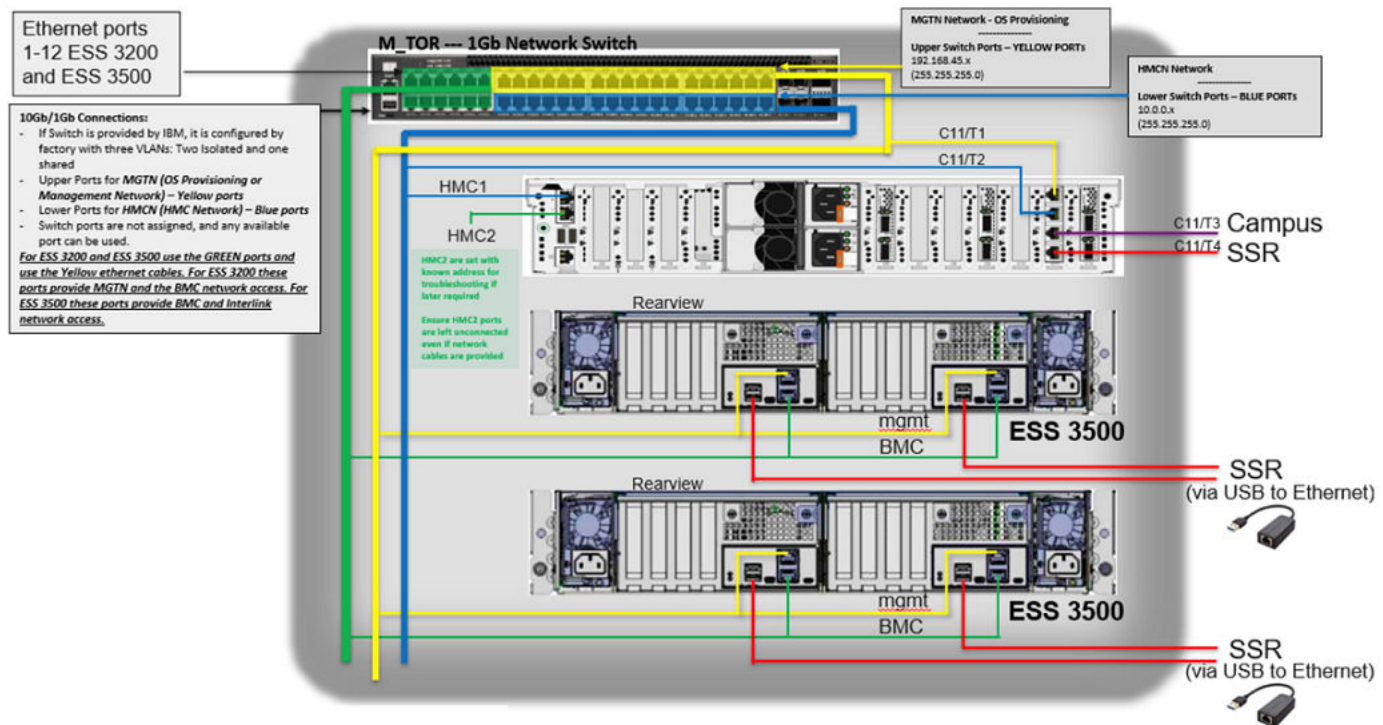


Figure 1: 1Gb network switch



Figure 2: Logical view of two switches

The orange cable that is shown in this figure must be connected between port 22 of the upper switch and 21 of the lower switch as part of the configuration. This cable works as inter-switch link (ISL) between the two switches.

Code version

ESS Legacy, ESS 3000, ESS 3200, ESS 5000, and ESS 3500 releases are included in ESS 6.1.4.x with two editions: Data Management Edition and Data Access Edition. An example of package names is as follows:

```
ess_6.1.4.1_0919-18_dme_ppc64le.tar.xz
ess_6.1.4.1_0919-18_dae_ppc64le.tar.xz
```

Note:

- The versions shown here might not be the GA version available on IBM® Fix Central. It is recommended to go to IBM Fix Central and download the latest code.
- ppc64le in the package name implies that each container runs on a POWER®-based EMS. For details about functions supported by respective containers, see [Support matrix](#) on page 1.

You can download the latest 6.1.x.x code (6.1.4.1 is the latest) from IBM Fix Central by using the following link.

- [IBM Fix Central download link](#)

A unified container is offered with two versions (Data management + Data access). Example package names for each container are as follows:

```
// Unified Container (Data Access and Data Management versions)
ESS_DAE_BASEIMAGE-6.1.4.1-ppc64LE-Linux.tgz
ESS_DME_BASEIMAGE-6.1.4.1-ppc64LE-Linux.tgz
```

Note: The container installs and runs on the EMS only. The EMS supported is Power-based only. Running container on a x86-based node is not supported as of now.

Upgrade guidance

From \ To	6.0.2.0	6.0.2.1	6.0.2.2	6.1.0.0	6.1.0.1	6.1.1.0	6.1.1.1	6.1.1.2	6.1.2.0	6.1.2.1	6.1.2.2	6.1.2.3	6.1.3.0	6.1.3.1	6.1.4.0	6.1.4.1
6.0.2.0		Green	Green	Yellow	Yellow	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red
6.0.2.1			Green	Yellow	Yellow	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red
6.0.2.2				Green	Yellow	Yellow	Red	Red	Red	Red	Red	Red	Red	Red	Red	Red
6.1.0.0					Green	Yellow	Yellow	Red	Red	Red	Red	Red	Red	Red	Red	Red
6.1.0.1						Green	Yellow	Yellow	Red	Red	Red	Red	Red	Red	Red	Red
6.1.1.0							Green	Yellow	Yellow	Red	Red	Red	Red	Red	Red	Red
6.1.1.1								Green	Yellow	Yellow	Red	Red	Red	Red	Red	Red
6.1.1.2									Green	Yellow	Yellow	Red	Red	Red	Red	Red
6.1.2.0										Green	Yellow	Yellow	Red	Red	Red	Red
6.1.2.1											Green	Yellow	Yellow	Red	Red	Red
6.1.2.2												Green	Yellow	Yellow	Red	Red
6.1.2.3													Green	Yellow	Yellow	Red
6.1.3.0														Green	Yellow	Yellow
6.1.3.1															Green	Yellow
6.1.4.0																Green
6.1.4.1																

RED - NOT SUPPORTED
 YELLOW - SUPPORTED NOT TESTED
 GREEN - SUPPORTED

Note: Upgrades to ESS 6.1.2.x follow the N-2 rule. You can upgrade from ESS 6.1.2.x, 6.1.1.x (that is, 6.1.1.2) or 6.1.0.x.

Further legacy container migration guidance

You must migrate first to ESS 5.3.7.x before you upgrade to ESS 6.1.x.x (container version).

ESS 5.3.x.x upgrade guidance

- You can upgrade to 5.3.7.x from 5.3.5.x (online) or 5.3.6.x (online).
- For online upgrade you can jump one OS version and for offline upgrade you can jump two OS versions.
 - Only exception is RHEL 7.7 to RHEL 7.9 upgrade. Because there is no RHEL 7.8.
 - Online upgrade to RHEL 7.7 from RHEL 7.6 can be done.
 - Upgrade to RHEL 7.7 from RHEL 7.5 must be done online.

ESS 6.1.x.x upgrade guidance

- It is recommended to convert from ESS 5.3.7.x to ESS 6.1.2.x and follow the normal N-X rules. To convert to ESS 6.1.2.x, use the following table (based on the RHEL 7.9 kernel):

Table 1: RHEL kernels	
ESS	Kernel
6.1.2.4	3.10.0-1160.71.1.el7
6.1.2.3	3.10.0-1160.62.1.el7
6.1.2.2	3.10.0-1160.49.1.el7
5.3.7.6	3.10.0-1160.62.1.el7
5.3.7.5	3.10.0-1160.59.1.el7
5.3.7.4	3.10.0-1160.49.1.el7
5.3.7.3	3.10.0-1160.45.1.el7
5.3.7.2	3.10.0-1160.31.1
5.3.7.1	3.10.0-1160.24.1
5.3.7.0	3.10.0-1160.11.1.el7

An example of upgrade jump is as follows:

- To upgrade to ESS 6.1.2.2, you can only upgrade from 5.3.7.4 or lower versions (that is, less than equal to 5.3.7.4).
- To upgrade to ESS 6.1.2.3, you can only upgrade from 5.3.7.6 or lower versions.
- It is not recommended to upgrade from ESS 5.3.7.x to ESS 6.1.1.2 anymore. Upgrade directly to ESS 6.1.2.3 or ESS 6.1.2.4. If you are updating from ESS 6.1.1.2, upgrade to 6.1.2.3 or higher (do not upgrade to 6.1.2.2).
- For ESS 5.3.7.3, consider downgrading MOFED to MLNX_OFED_LINUX-4.9-3.1.5.3, and then convert to 6.1.2.3 or 6.1.2.4. This is to obtain full support for online upgrade when converting to RDMA core libs.
- When upgrading to 5.3.x.x, first upgrade to ESS 5.3.7.2 or ESS 5.3.7.3, and then upgrade to 6.1.2.3 or 6.1.2.4. This upgrade is to obtain full support for online upgrade when converting to RDMA core libs.
- You may need to modify the container unblock jumps from a specific 5.3.7.x level. Issue to the following command to upgrade the ESS level in the container:

```
vim /opt/ibm/ess/deploy/ansible/vars.yml
```

- Change (an example if you want to convert from ESS 5.3.7.1 or higher) LEGACY_SUPPORTED_VERSION: "5.3.7.3" to LEGACY_SUPPORTED_VERSION: "5.3.7.1".

For more information about the ESS 6.1.x.x upgrade, see [IBM Spectrum Scale Alert: Mellanox OFED 5.x considerations in IBM ESS V6.1.2.x+](#).

Example of the /etc/hosts file

```
127.0.0.1 localhost localhost.localdomain.local localhost4 localhost4.localdomain4

### Management IPs 192.168.45.0/24
192.168.45.20 ems1.localdomain.local ems1
192.168.45.21 essio1.localdomain.local essio1
192.168.45.22 essio2.localdomain.local essio2
192.168.45.23 prt1.localdomain.local prt1
192.168.45.24 prt2.localdomain.local prt2

### High-speed IPs 10.0.11.0/24
10.0.11.1 ems1-hs.localdomain.local ems1-hs
10.0.11.2 essio1-hs.localdomain.local essio1-hs
10.0.11.3 essio2-hs.localdomain.local essio2-hs
10.0.11.4 prt1-hs.localdomain.local prt1-hs
10.0.11.5 prt2-hs.localdomain.local prt2-hs

### Container info 192.168.45.0/24
192.168.45.80 cems0.localdomain.local cems0

### Protocol CES IPs
10.0.11.100 prt_ces1.localdomain.local prt_ces1
```

```
10.0.11.101 prt_ces1.localdomain.local prt_ces1
10.0.11.102 prt_ces2.localdomain.local prt_ces2
10.0.11.103 prt_ces2.localdomain.local prt_ces2
```

Node classes

All building blocks have node classes except the EMS node. Instead of a node list, you can use these node classes in the `essrun` command when the inventory file contains these nodes. (The `essrun config load` command adds the nodes to the inventory file.)

Node classes are as follows:

- legacy: gss_ppc64le
- 5000: ess_ppc64le
- 3000: ess_x86_64
- 3200: ess3200_x86_64
- 3500: ess3500_x86_64
- Protocol Power9: ces_ppc64le

Deployment/upgrade instructions

1. Log in to the EMS (over management), set root password, set campus connection.
 - a. Campus connection interface is called 'campus'. Use `nmtui` to set an address.
 - b. Set the EMS hostname.
2. Log out and log in over the campus connection.
3. Set up the `/etc/hosts` file.
4. Copy the GA code into the `/etc/deploy` directory.
5. Extract outer tarball.
6. Decompress the inner xz file.

```
xz --decompress ess_6.1.4.1_0919-18_dae_ppc64le.tar.xz
```

7. Extract tar file.

```
tar xvf ess_6.1.4.1_0919-18_dae_ppc64le.tar
```

8. Run the installer.
 - a. Start the container.

```
sh ess_6.1.4.1_0919-18_dae_ppc64le --start-container
```

- b. Answer any user prompts.

Configuration in the container

1. Run Config load.

```
essrun -N ems1,essio1,essio2 config load -p ibmesscluster
```

2. Run Config check.

```
essrun -N ems1,essio1,essio2 config check -p ibmesscluster
```

3. Update EMS.

- a. Update EMS offline.

```
essrun -N ems1 update --offline
```

- b. If you are prompted to reboot then run the command again.

c. If you need to restart the container after reboot run the following command:

```
essmgr -r
```

d. Update the EMS again.

```
essrun -N ems1 update --offline
```

4. Update I/O nodes.

- Update I/O nodes online.

```
essrun -N <Node List/Node Name> update
```

- Update I/O nodes offline.

```
essrun -N <Node List/Node Name> update --offline
```

5. Update POWER firmware. For more information see, Upgrading POWER9 firmware appendix in ESS Deployment Guide.

6. Create network bonds.

```
essrun -N ems1,essio1,essio2 network --suffix=-hs
```

7. Run network test.

```
ssh essio1
```

```
ESSENV=TEST essnettest -N essio1,essio2 --suffix=-hs
```

8. Create a cluster.

```
essrun -N <Node List/Node Name> cluster --suffix=-hs
```

9. Add EMS.

```
essrun -N essio1 cluster --add-ems ems1 --suffix=-hs
```

10. Create a file system.

```
essrun -N ess3500_x86_64 filesystem --name fs1 --suffix=-hs
```

Note: For ESS 3500, you must keep 1.5 TB or more space free if future capacity MES is planned (performance to hybrid). Thus, it is recommended to not use all available space when you create a file system for the performance model. The default allocation is 80% of available space when you use the `essrun filesystem` command (for x86 nodes).

11. Final installation check per node.

```
essinstallcheck -N localhost
```

12. Check whether all nodes performance.

```
gnrhealthcheck
```

```
mmhealth node show -a
```

13. Set up Chrony/NTP.

14. Set up call home.

15. Enable added security (SELinux/firewall/sudo/admin central).

16. Configure the GUI/collectors.

```
essrun -N ems1,essio1,essio2 gui -configure
```

17. Set up conserver.

18. Set up protocol services (3500 or P8/P9 Protocol nodes).

19. Set up client nodes.

Note: If you want to add a building block, see Summary of ESS deployment scenarios appendix in ESS Deployment Guide.

November 2022

IBM and the IBM logo are trademarks of the IBM Corporation in the United States, other countries, or both.

Printed in USA

© **Copyright International Business Machines Corporation 2022.**

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

SC27-9889-01

