IBM Storage Scale 5.2.3 Protocols Quick Overview

May, 2025.

Before starting

Always start here to understand:
• Common prerequisites
• Basic installation toolkit operation
• Requirements when an existing cluster exists, both with or without an ESS

1. How does the installation toolkit work?

IBM Storage Scale installation toolkit operation can be summarized in 4 phases:

I. User input via **spectrumscale** commands
II. A **spectrumscale install** phase
III. A **spectrumscale deploy** phase
IV. A **spectrumscale upgrade** phase

Each phase can be run again at later points in time to introduce new nodes, protocols, NSDs, file systems, or updates.

All user input via **spectrumscale** commands is recorded into a cluster definition file in: **/usr/lpp/mmfs/5.2.3/ansible-toolkit/ansible/vars**

Each phase acts upon all nodes inputted into the cluster definition file. For example, if you only want to deploy protocols in a cluster containing a mix of unsupported and supported operating systems, input only the supported protocol nodes and leave all the other nodes out of the cluster definition.

Note: To deploy a cluster in the cloud (AWS or GCP), use the **cloud-kit** command, which is available at: **/usr/lpp/mmfs/your_s-cale_version/cloudkit**. This quick guide does not cover the cloudkit usage. For more information, see **cloudkit** command.

2. Hardware or performance sizing

Work with your IBM account team or Business Partner for suggestions on the best configuration possible to fit your environment. Also, see IBM Storage Scale FAQ for related information.

3. Operating system and CPU architecture

The installation toolkit supports these operating systems:

- x86:
RHEL 8.x / 9.x, SLES 15, Ubuntu 22.04 / 24.04
- ppc64 LE:
RHEL 8.x / 9.x
- s390x:
RHEL 8.x / 9.x , SLES 15, Ubuntu 22.04
- aarch64:
RHEL 9.x, Ubuntu 22.04

All cluster nodes that the installation toolkit acts upon must be part of the same CPU architecture and endianness.

All protocol nodes must have the same operating system and architecture.

Refer to the support matrices available at IBM Storage Scale FAQ.

4. Repositories

A base repository must be set up on every node. For RHEL 8 and 9, also set up the AppStream repository.

- RHEL check: yum repolist, dnf repolist
- SLES check: zypper repos
- Ubuntu check: apt edit-sources

5. Firewall and networking, and SSH

Make sure that:

-All nodes are networked together and pingable via IP, FQDN, and hostname.
-Reverse DNS lookup are in place.
-If **/etc/hosts** is used for name resolution, ordering within is: IP FQDN hostname.
-Promptless SSH is setup between all nodes and themselves by using IP, FQDN, and hostname.

Firewalls should be turned off on all nodes. Else, specific ports must be opened both internally (for GPFS and the installer) and externally (for the protocols). See IBM Documentation for more details before proceeding.

6. Time synchronization among nodes is required

A consistent time must be established on all nodes of the cluster.

7. Cleanup prior SMB and NFS

Prior implementations of SMB and NFS must be completely removed before proceeding with a new protocol deployment. Refer to the cleanup guidance in IBM Documentation.

8. If a GPFS cluster preexists

Proceed to the "Protocol deployment" section if:

a. File systems are created and mounted ahead of time, and **nfs4** ACLs is in place
b. SSH promptless access among all nodes is set
c. Firewall ports are open
d. CCR is enabled
e. You configured **mmchconfig release=LATEST**
f. The installed GPFS **rpms** match the exact build dates of those included within the protocols package

9. If an IBM Storage Scale System is part of the cluster

Proceed to the "Cluster installation" section to use the installation toolkit to install GPFS and add new nodes to the existing ESS cluster. Proceed to the "Protocol deployment" section to deploy protocols.

a. CCR must be enabled
b. EMS nodes must be in the **ems** nodeclass. I/O nodes must be in their own nodeclass: **gss** or **gss_ppc64**
c. GPFS on the ESS nodes must be at minimum 5.0.5.x
d. All quorum and quorum-manager nodes are recommended to be at the latest levels possible
e. A CES shared root file system has been created and mounted on the EMS

10. Protocols in a stretch cluster

Refer to the stretch cluster use case within the IBM Documentation.

11. Extract IBM Storage Scale package

There is no protocols-specific package. Any standard, advanced, data access, or data management package is now sufficient for protocol deployment. Extracting the package presents a license agreement.

```
.Spectrum_Scale_Data_Management-5.2.3.x-<arch>-Linux-install
```

12. Explore the spectrumscale help

From **/usr/lpp/mmfs/5.2.3.x/ansible-toolkit**, use the **-h** flag.

```
./spectrumscale -h
./spectrumscale setup -h
./spectrumscale node add -h
./spectrumscale config -h
./spectrumscale config protocols -h
```

13. FAQ and quick reference

See the IBM Storage Scale Quick Reference.

Refer to the IBM Storage Scale FAQ page.

14. Installation toolkit scope

Everything that is directly related to IBM Storage Scale, like adding nodes to a cluster or setting up protocols, can be done manually by using the installation toolkit within the protocol virtual machine.

ESS is responsible for updating any non-IBM Storage Scale item, like the operating system, MOFED, or tuning profiles.

Cluster installation

Start here if you would like to:
• Create a new cluster from scratch
• Add and install new GPFS nodes to an existing cluster (client, NSD, GUI)
• Create new NSDs on an existing cluster

1. Setup the node that will start the installation

Pick an IP that exists on this node and is accessible to or from all nodes via promptless SSH:

```
./spectrumscale setup -s IP
```

- ESS: If the **spectrumscale** command is being run on nodes in a cluster with an ESS, make sure to switch to ESS mode (see page 2 for ESS examples):
```
./spectrumscale setup -s IP -st ess
```
- ECE: If the **spectrumscale** command is being run on an Erasure Code Edition cluster, make sure to switch to ECE mode:
```
./spectrumscale setup -s IP -st ece
```

Ansible is a prerequisite for using the **spectrumscale** command.

2. Populate the cluster

If a cluster preexists, the installation toolkit can automatically traverse the existing cluster and populate its **scale_clusterdefinition.json** file with the current cluster's configuration details.

Point it at a node within the cluster with promptless SSH access to all other cluster nodes:

```
./spectrumscale config populate -N hostname
```

If you are in ESS mode, point config populate to the EMS:

```
./spectrumscale config populate -N ems1
```

Note: Consult the limitations of the **config populate** command.
Note: To remotely mount a file system using the installation toolkit, use **remote_mount**. This quick guide does not cover the **remote_mount** command. For more information, see **spectrumscale** command.

3. Add NSD server nodes (non-ESS nodes)

Adding NSD nodes is necessary if you would like the installation toolkit to configure new NSDs and file systems.

```
./spectrumscale node add hostname -n
./spectrumscale node add hostname –n
...
```

4. Add NSDs (non-ESS devices)

NSDs can be added as non-shared disks seen by a primary NSD server. NSDs can also be added as shared disks seen by a primary and multiple secondary NSD servers.

In this example, we add 4 /dev/dm disks seen by both primary and secondary NSD servers:

```
./spectrumscale nsd add –p primary_nsdnode_hostname
-s secondary_nsdnode_hostname /dev/dm-1 /dev/dm-2
/dev/dm-3 /dev/dm-4
```

5. Define file systems (non-ESS)

File systems are defined by assigning a file system name to one or more NSDs. File systems are defined but not created until this install is followed by a deployment.

In this example, we assign all 4 NSDs to the **fs1** file system:

```
./spectrumscale nsd list
./spectrumscale filesystem list
./spectrumscale nsd modify nsd1 -fs fs1
./spectrumscale nsd modify nsd2 -fs fs1
./spectrumscale nsd modify nsd3 -fs fs1
./spectrumscale nsd modify nsd4 -fs fs1
```

If desired, multiple file systems can be assigned at this point. See IBM Documentation for details on "spectrumscale nsd modify". We recommend a separate file system for shared root to be used with protocols.

6. Add GPFS client nodes

```
./spectrumscale node add hostname
```

The installer assigns quorum and manager nodes by default. Refer to IBM Documentation if a specific configuration is desired.

7. Add Storage Scale GUI nodes

```
./spectrumscale node add hostname -g -a
…
```

The management GUI automatically starts after installation and allows for further cluster configuration and monitoring.

8. Configure performance monitoring

Configure performance monitoring consistently across nodes:

```
./spectrumscale config perfmon -r on
```

9. Configure call home

Call home is enabled by default within the installation toolkit. Refer to the call home settings and configure mandatory options for call home:

```
./spectrumscale callhome config -h
```

Alternatively, disable the call home:

```
./spectrumscale callhome disable
```

10. Name your cluster

```
./spectrumscale config gpfs -c my_cluster_name
```

11. Review your configuration

```
./spectrumscale node list
./spectrumscale nsd list
./spectrumscale filesystem list
./spectrumscale config gpfs --list
./spectrumscale install --precheck
```

12. (Optional) Enable installation of IBM Storage Scale native REST API

```
./spectrumscale scaleadmd enable
./spectrumscale nodeid define
```

Note: With **ansible-toolkit**, root and passwordless SSH access are necessary to deploy a native REST API cluster.

13. Start the installation

```
./spectrumscale install
```

After completion, you have an active GPFS cluster with available NSDs, file systems, performance monitoring, time synchronization, call home, and a GUI. File systems are fully created and protocols installed in the next stage, deployment.

Install can be re-run in the future to:
- Add GUI nodes
- Add NSD server nodes
- Add GPFS client node
- Add NSDs
- Add file systems
- Enable and configure or update **callhome** settings

Protocol and file system deployment

Start here if you already have a cluster and would like to:
• Add or enable protocols on existing cluster nodes
• Create a file system on existing NSDs

1. Setup the node that will start the installation

Setup is necessary unless **spectrumscale** setup was previously run on this node for a past GPFS installation or protocol deployment. Pick an IP that exists on this node and is accessible to or from all nodes via promptless SSH:

```
./spectrumscale setup -s IP
```

-ESS: If the **spectrumscale** command is being run on nodes in a cluster with an ESS, make sure to switch to ESS mode (see page 2 for ESS examples):
```
./spectrumscale setup -s IP -st ess
```
-ECE: If the **spectrumscale** command is being run on an Erasure Code Edition cluster, make sure to switch to ECE mode:
```
./spectrumscale setup -s IP -st ece
```

Ansible is a prerequisite for using the **spectrumscale** command.

2. Populate the cluster

Optionally, the installation toolkit can automatically traverse the existing cluster and populate its **scale_clusterdefinition.json** file with current cluster details. Point it at a node within the cluster with promptless SSH access to all other cluster nodes:

```
./spectrumscale config populate -N hostname
```

If you are in ESS mode, point config populate to the EMS:

```
./spectrumscale config populate -N ems1
```

Note: Consult the limitations of the **config populate** command.

3. Add protocol nodes

```
./spectrumscale node add hostname -p
./spectrumscale node add hostname -p
...
```

Note: In an IBM Storage Scale System, the protocol node and the manager server run on the same VM. You cannot run multiple VMs or run items (protocol node, manager server) directly on the host itself.

4. Assign protocol IPs (CES-IPs)

Add a comma-separated list of IPs to be used specifically for cluster export services such as NFS and SMB. Reverse DNS lookup must be in place for all IPs. CES IPs must be unique and different from cluster node IPs.

```
./spectrumscale config protocols -e EXPORT_IP_POOL
```

Note: All protocol nodes must see the same CES-IP networks. If CES-groups are to be used, apply them after the deployment is successful.

5. Verify file system mount points are as expected

```
./spectrumscale filesystem list
```

Note: Skip this step if you set up file systems or NSDs manually and not through the installation toolkit.

6. Configure protocols to point to a shared root file system location

A CES directory gets automatically created at the root of the specified file system mount point. This is used for protocol admin/config and needs >= 4 GB free.

Upon completion of protocol deployment, GPFS configuration points to this as **cesSharedRoot**. It is recommended that **cesSharedRoot** be a separate file system.

```
./spectrumscale config protocols -f fs1 -m /ibm/fs1
```

Note: If you setup file systems or NSDs manually, perform a manual check of **<mmlsnsd>** and **<mmlsfs all -L>** to make sure all NSDs and file systems required by the deployment are active and mounted before continuing.

7. Enable the desired protocols

```
./spectrumscale enable nfs
./spectrumscale enable smb
./spectrumscale enable s3
```

Note: In an IBM Storage Scale System, the only protocols available are NFS and SMB.

8. Configure call home

Call home is enabled by default within the installation toolkit. Refer to the call home settings and configure mandatory options for call home:

```
./spectrumscale callhome config -h
```

Alternatively, disable call home:

```
./spectrumscale callhome disable
```

9. Review your configuration

```
./spectrumscale node list
./spectrumscale deploy –precheck
```

10. Start the deployment

```
./spectrumscale deploy
```

Upon completion, you have protocol nodes with active cluster export services and IPs. Performance monitoring tools are also ready for use at this time.

Deploy can be re-run in the future to:
-Enable additional protocols
-Add additional protocol nodes (run install first to add more nodes)
-Enable and configure or update call home settings

Configuration

Start here if you already have a cluster with protocols enabled and would like to:
• Check cluster state and health, basic logging or debugging
• Configure a basic SMB or NFS export
• Configure and enable file audit logging or watch folders

1. Path to binaries

Add the following **PATH** variable to your shell profile to allow convenient access of GPFS **mm** commands:

```
export PATH=$PATH:/usr/lpp/mmfs/bin
```

2. Basic GPFS health

```
mmgetstate -aL
mmlscluster
mmlscluster --ces
mmnetverify
```

3. CES service and IP check

```
mmces address list
mmces service list -a
mmhealth cluster show
mmhealth node show -N all -v
mmhealth node show <component> -v
mmces events list –a
```

4. Authentication

```
mmuserauth service list
mmuserauth service check
```

5. Call home

```
mmcallhome info list
mmcallhome group list
mmcallhome status list
```

6. File protocols (NFS and SMB)

Make sure that all file systems to be used with protocols have **nfs4** ACLs and locking in effect. Protocols do not work correctly without this setting in place.

Check with:

```
mmlsfs all -D -k
```

Example NFS export creation:

```
mkdir /ibm/fs1/nfs_export1

mmnfs export add /ibm/fs1/nfs_export1 -c "*(Access_-
Type=RW,Squash=no_root_squash,SecType=sys,Proto-cols=3:4)"

mmnfs export list
```

Example SMB export creation:

```
mkdir /ibm/fs1/smb_export1

chown "DOMAIN\USER" /ibm/fs1/smb_export1

mmsmb export add smb_export1 /ibm/fs1/smb_export1
--option "browseable=yes"

mmsmb export list
```

7. S3 protocol

Example for S3 account creation:
```
# mms3 account create s3user9001 --uid 9001 --gid 9000 --
newBucketsPath "/ibm/fs/s3user9001-dir"
```

List the created accounts:
```
# mms3 account list
```

Example for creating a bucket for an existing S3 user account:
```
# mms3 bucket create newbucket-9k --accountName s3user9001 --
filesystemPath /ibm/fs/s3user9001-dir/newbucket-9k-dir
```

List the bucket:
```
# mms3 bucket list newbucket-9k
```

8. Performance monitoring

```
systemctl status pmsensors
systemctl status pmcollector
mmperfmon config show
mmperfmon query -h
```

9. File audit logging

File audit logging (FAL) is available only with Data Management and Advanced editions of IBM Storage Scale.

a. Enable and configure using the installation toolkit as follows:
```
./spectrumscale fileauditlogging enable
./spectrumscale filesystem modify –fileauditlog-gingenable
gpfs1
./spectrumscale fileauditlogging list
./spectrumscale filesystem modify –logfileset
<LOGFILESET> retention <days> gpfs1
```

b. Install the file audit logging rpms on all nodes:
```
./spectrumscale install --precheck
./spectrumscale install
```

c. Deploy the file audit logging configuration *gpfs.adv.* or gpfs.dm.* rpms must be installed on all nodes:
```
./spectrumscale deploy --precheck
./spectrumscale deploy
```

d. Check the status:
```
mmhealth node show FILEAUDITLOG -v
mmhealth node show MSGQUEUE -v
mmaudit all list
mmaudit all consumeStatus -N <node list>
mmwatch all list
```

10. Logging and debugging

Installation or deployment:

```
/usr/lpp/mmfs/5.2.3.x/ansible-toolkit/logs
```

Verbose logging for all **spectrumscale** commands by adding a '-v' immediately after **./spectrumscale**:

```
/usr/lpp/mmfs/5.2.3.x/ansible-toolkit/spectrumscale
-v <cmd> GPFS default log location:/var/adm/ras/
```

Enabling Linux system log or journal is recommended.

11. Data capture for support

A consistent time must be established on all nodes of the cluster.

System-wide data capture:

```
/usr/lpp/mmfs/bin/gpfs.snap
```

Installation/Deploy/Upgrade specific:

```
/usr/lpp/mmfs/5.2.3.x/ansible-toolkit/installer.snap.py
```

Upgrade and cluster additions

Always start here to understand:
• Upgrade guidance
• How to add nodes, NSDs, FSs, protocols, to an existing cluster

Important note: CES Swift Object feature is not supported from IBM Storage Scale 5.1.9 onwards. For more information, see *Stabilized, deprecated, and discontinued features in IBM Storage Scale*.

1. Upgrading 5.1.x.x to 5.2.3.x

a. Extract the 5.2.3.x IBM Storage Scale package:
```
./Spectrum_Scale_Data_Management-5.2.3.x-Linux
```

b. Setup and configure the installation toolkit:
```
. spectrumscale setup -s <IP of installer node>
. spectrumscale config populate -N <any cluster  node>
```

If **config populate** is incompatible with your cluster configuration, you must manually add the nodes and configure to the installation toolkit.

If desired, enable the prompt to users to shut down their work-loads before starting the upgrade: **./spectrumscale upgrade config workload -p** on:

```
./spectrumscale node list
./spectrumscale nsd list
./spectrumscale filesystem list
./spectrumscale config gpfs
./spectrumscale config protocol
./spectrumscale upgrade precheck
./spectrumscale upgrade run
```

2. Upgrading 5.2.3.x to future PTFs

Follow the same procedure as indicated in the previous item.

3. Upgrade compatibility with LTFS-EE

a. **ltfsee stop** (on all LTFSEE nodes)
b. **umount /ltfs** (on all LTFSEE nodes)
c. **dsmmigfs disablefailover** (on all LTFSEE nodes)
d. **dsmmigfs stop** (on all LTFSEE nodes)
e. **systemctl stop hsm.service** (on all LTFSEE nodes)
f. Upgrade using the installation toolkit
g. Upgrade LTFS-EE if desired
h. Reverse steps e through a and restart/enable

4. Upgrade compatibility with TCT

a. Stop TCT on all nodes prior to the upgrade
```
mmcloudgateway service stop -N Node | Nodeclass
```
b. Upgrade using the installation toolkit
c. Upgrade the TCT **rpms** manually, then restart the TCT

5. Offline upgrade using the installation toolkit

The installation toolkit supports offline upgrade of all nodes in the cluster or a subset of nodes in the cluster. This is useful for 4.2.3.x - 5.2.x.x upgrades. It is also useful when nodes are unhealthy and cannot be brought into a healthy or active state for upgrade. Parallel offline upgrade of all nodes in the cluster is also supported.

a. Check the upgrade configuration:
```
./spectrumscale upgrade config list
```

b. Add nodes that are already shutdown:
```
./spectrumscale upgrade config offline -N <node1,node2>
./spectrumscale upgrade config offline -N all
./spectrumscale upgrade config list
```

Note: If all the nodes are included in the offline upgrade, the installation toolkit performs a parallel offline upgrade on them.

c. Start the upgrade:
```
./spectrumscale upgrade precheck
./spectrumscale upgrade run
```

6. Upgrading subsets of nodes (excluding nodes)

The installation toolkit supports excluding groups of nodes from the upgrade. This allows for staging cluster upgrades across multiple windows. For example, upgrading only NSD nodes and then, at a later time, upgrading only protocol nodes. This is also useful if specific nodes are down and unreachable. See IBM Documentation to learn about limitations.

a. Check the upgrade configuration:
```
./spectrumscale upgrade config list
```

a. Check the upgrade configuration:
```
./spectrumscale upgrade config list
```

b. Add nodes that are NOT to be upgraded:
```
./spectrumscale upgrade config exclude –N <node1,node2>
./spectrumscale upgrade config list
```

c. Start the upgrade:
```
./spectrumscale upgrade precheck
./spectrumscale upgrade run
```

d. Prepare to upgrade the previously excluded nodes:
```
./spectrumscale upgrade config list
./spectrumscale upgrade config exclude --clear
./spectrumscale upgrade exclude -N <already_up-graded_nodes>
```

e. Start the upgrade:
```
./spectrumscale upgrade precheck
./spectrumscale upgrade run
```

7. Resume of a failed upgrade

If an installation toolkit upgrade fails, it is possible to correct the failure and resume the upgrade without needing to recover all nodes or services. Resume with: **./spectrumscale upgrade run**.

8. Handling Linux kernel updates

The GPFS portability layer must be rebuilt on every node that under- goes a Linux kernel update.

Apply the kernel, reboot, rebuild the GPFS portability layer on each node with this command before starting GPFS: **/usr/lpp/mmfs/bin/mmbuildgpl**

Or: **mmchconfig autoBuildGPL=yes and mmstartup**

9. Adding to the installation

The following procedures can be combined to reduce the number of installations and deployments necessary.

To add a node:
a. Choose one or more node types to add:
Client node: **./spectrumscale node add hostname**
NSD node: **./spectrumscale node add hostname -n**
Protocol node: **./spectrumscale node add hostname -p**
GUI node: **./spectrumscale node add hostname -g -a**
… Repeat for as many nodes as you'd like to add.
b. Install GPFS on the new nodes:
```
./spectrumscale install -pr
./spectrumscale install
```
c. If a protocol node is being added, also run deployment:
```
./spectrumscale deploy -pr
./spectrumscale deploy
```

To add an NSD:
a. Verify that the NSD server that connects this new disk exists within the cluster.
b. Add the NSD(s) to the installation toolkit:
```
./spectrumscale nsd add -h
```
… Repeat for as many NSDs as you'd like to add.
c. Run an install:
```
./spectrumscale install -pr
./spectrumscale install
```

To add a file system:
a. Verify free NSDs exist and are known to the installation toolkit.
b. Define the file system:
```
./spectrumscale nsd list
./spectrumscale nsd modify nsdX -fs file_sys tem_name
```
c.Deploy the new file system:
```
./spectrumscale deploy -pr
./spectrumscale deploy
```

To enable another protocol:
See the "Protocol and file system deployment" section. Proceed with steps 7, 8, 9, 10. Note that some protocols require removal of the authentication configuration prior to enablement.

** URLs are subject to change **

May, 2025.

Examples

Example of readying Red Hat Linux nodes for Storage Scale installation and deployment of protocols

Configure promptless SSH (promptless ssh is required):
```
# ssh-keygen
# ssh-copy-id <FQDN of node>
# ssh-copy-id <IP of node>
# ssh-copy-id <non-FQDN hostname of node>
```

Repeat on all nodes to all nodes, including current node.

Turn off firewalls (alternative is to open ports specific to each Storage Scale functionality):
```
#systemctl stop firewalld
#systemctl disable firewalld
```

Repeat on all nodes.

How to check if a yum repository is configured correctly:
`# yum repolist` or **`# dnf repolist`** > Should return no errors. It must also list RHEL base operating system and AppStream repositories.

Other repository possibilities include a satellite site, a custom yum repository, an RHELx.x DVD .iso, an RHELx.x physical DVD.

What if I don't want to use the installation toolkit? How do I get a repository for all the IBM Storage Scale **`rpms`**?
```
# cd /usr/lpp/mmfs/5.2.3.x/tools/repo
# ./local-repo --repo
# yum repolist
```

Preinstall prerequisite rpms to make installation and deployment easier:
```
# yum install kernel-devel cpp gcc gcc-c++ glibc sssd ypbind openldap-clients krb5-workstation
```

Turn off SELinux (or set to permissive mode):
```
# sestatus
# vi /etc/selinux/config
```

Change **`SELINUX=xxxxxx`** to **`SELINUX=disabled`**.

Save and reboot. Repeat on all nodes.

Setup a default path to IBM Storage Scale commands (not required):
```
#vi /root/.bash_profile
——add this line——
export PATH=$PATH:/usr/lpp/mmfs/bin
——save/exit——
```

Log out and back in for these changes to take effect.

Example of a new IBM Storage Scale cluster installation followed by a protocol deployment

Installation toolkit commands for installation
- Toolkit is running from **`cluster-node1`** with an internal cluster network IP of 10.11.10.11, which all nodes can reach: **`cd /usr/lpp/mmfs/5.2.3.x/ansible-toolkit/`**

```
./spectrumscale setup -s 10.11.10.11
./spectrumscale node add cluster-node1 -a -g
./spectrumscale node add cluster-node2 -a -g
./spectrumscale node add cluster-node3
./spectrumscale node add cluster-node4
./spectrumscale node add cluster-node5 -n
./spectrumscale node add cluster-node6 -n
./spectrumscale nsd add -p node5.tuc.stglabs.ibm.com -s node6.tuc.stglabs.ibm.com -u dataAndMetadata -fs cesSharedRoot -fg 1 "/dev/sdb"
./spectrumscale nsd add -p node6.tuc.stglabs.ibm.com -s node5.tuc.stglabs.ibm.com -u dataAndMetadata -fs cesSharedRoot -fg 2 "/dev/sdc"
./spectrumscale nsd add -p node5.tuc.stglabs.ibm.com -s node6.tuc.stglabs.ibm.com -u dataAndMetadata -fs fs1 -fg 1 "/dev/sdh"
./spectrumscale nsd add -p node5.tuc.stglabs.ibm.com -s node6.tuc.stglabs.ibm.com -u dataAndMetadata -fs fs1 -fg 1 "/dev/sdi"
./spectrumscale nsd add -p node5.tuc.stglabs.ibm.com -s node6.tuc.stglabs.ibm.com -u dataAndMetadata -fs fs1 -fg 2 "/dev/sdj"
./spectrumscale nsd add -p node5.tuc.stglabs.ibm.com -s node6.tuc.stglabs.ibm.com -u dataAndMetadata -fs fs1 -fg 2 “/dev/sdk"
./spectrumscale config perfmon -r on
./spectrumscale callhome enable                                  ← This is the list of nodes the installation toolkit will upgrade. Remove any non-CES nodes you would rather do manually.
./spectrumscale callhome config -n COMPANY_NAME -i COMPANY_ID -cn MY_COUNTRY_CODE -e MY_EMAIL_ADDRESS
./spectrumscale config gpfs -c mycluster
./spectrumscale node list
./spectrumscale install --precheck
./spectrumscale install
```

Installation Outcome: A 6-node IBM Storage Scale cluster with active NSDs
- 2 GUI nodes
- 2 NSD nodes
- 2 client nodes
- 10 NSDs
- Configured performance monitoring
- Configured **`callhome`**
- 3 file systems created, each with 2 failure groups

Example of readying Red Hat Linux nodes for Storage Scale installation and deployment of protocols

Starting point
- If you have a 5148-22L protocol node, stop following these directions and refer to the ESS 5.3.7 (or higher) Quick Deployment Guide.
- The cluster containing ESS is active and online.
- RHEL 8.x/9.x, SLES 15, or Ubuntu 20.04 is installed on all nodes that are going to serve as protocol nodes.
- RHEL 8.x/9.x, SLES 15, or Ubuntu 20.04 base repository is set up on nodes that are going to serve as protocol nodes.
- The nodes that will serve as protocol nodes have connectivity to the GPFS cluster network.

Limitations
- If you have a 5148-22L protocol node, stop following these directions and refer to the ESS 5.3.7 (or higher) Quick Deployment Guide.
- The cluster containing ESS is active and online.
- RHEL 8.x/9.x, SLES 15, or Ubuntu 20.04 is installed on all nodes that are going to serve as protocol nodes.
- RHEL 8.x/9.x, SLES 15, or Ubuntu 20.04 base repository is set up on nodes that are going to serve as protocol nodes.

Getting started
1. Create a **`cesSharedRoot`** from the EMS: **`gssgenvdisks --create-vdisk --create-nsds --create-filesystem --contact-node gssio1-hs --crcesfs`**
2. Mount the CES shared root file system on the EMS node and set it to automount. When done with this full procedure, make sure the protocol nodes are set to automount the CES shared root file system as well.
3. Use the ESS GUI or CLI to create additional file systems for protocols if desired. Configure each file system for **`nfsv4`** ACLs.
4. Pick a protocol node to run the installation toolkit from.
5. Locate the installation toolkit, which is contained within these packages: Storage Scale Standard or Data Access or Advanced or Data Management Edition.
6. Download and extract one of the Storage Scale packages to the protocol node that will run the installation toolkit. Once extracted, the installation toolkit is in the **`/usr/lpp/mmfs/5.2.3.x/ansible-toolkit`** directory.

Inputting the configuration into the installation toolkit with the commands detailed below, involves pointing the installation toolkit to the EMS node, telling the installation toolkit about the mount points and paths to the CES shared root, and designating the protocol nodes and protocol config to be installed/deployed. Refer to the IBM Storage Scale FAQ.

Installation Toolkit commands
```
./spectrumscale setup -s 10.11.10.11 -st ess                    ← Internal GPFS network IP on the current installer node that can see all protocol nodes.
./spectrumscale config populate -N ems-node                     ← OPTIONAL. Have the installation toolkit traverse the existing cluster and autopopulate its configuration.
./spectrumscale node list                                       ← OPTIONAL. Check the node configuration discovered by config populate.
./spectrumscale node add ems-node -a –e                         ← Designate the EMS node for the installation toolkit to use for coordination of the install/deploy.
./spectrumscale node add cluster-node1 -p
./spectrumscale node add cluster-node2 -p
./spectrumscale node add cluster-node3 -p
./spectrumscale node add cluster-node4 –p
./spectrumscale config protocols -e 172.31.1.10,172.31.1.11,172.31.1.12,172.31.1.13,172.31.1.14
./spectrumscale config protocols -f cesSharedRoot -m /ibm/cesSharedRoot
./spectrumscale enable nfs
./spectrumscale enable smb
./spectrumscale enable s3
./spectrumscale node list                                       ← It is normal for ESS IO nodes to not be listed in the installation toolkit. Do not add them.
./spectrumscale install --precheck
./spectrumscale install                                         ← The install will install GPFS on the new protocol nodes and add them to the existing ESS cluster.
./spectrumscale deploy --precheck                               ← It is important to make sure CES shared root is mounted on all protocol nodes before continuing.
./spectrumscale deploy                                          ← The deploy will install / configure protocols on the new protocol nodes.
```

Installation outcome
- EMS node used as an administration node by the installation toolkit, to coordinate the installation
- 4 new nodes installed with GPFS and added to the existing ESS cluster
- Performance sensors automatically installed on the 4 new nodes and pointed back to existing collector / GUI on the EMS node
- ESS I/O nodes, NSDs/vdisks, left untouched by the installation toolkit

Deploy outcome
- CES Protocol stack added to 4 nodes, now designated as protocol nodes with server licenses
- 4 CES-IPs distributed among the protocol nodes
- Protocol configuration and state data will use the **`cesSharedRoot`** file system, which was pre-created on the ESS.

Installation toolkit commands for protocol deployment (assumes cluster created from the previous configuration)
- Toolkit is running from the same node that performed the previous installation, **`cluster-node1`**.

```
./spectrumscale node add cluster-node3 -p
./spectrumscale node add cluster-node4 -p
./spectrumscale config protocols -e 172.31.1.10,172.31.1.11,172.31.1.12,172.31.1.13,172.31.1.14
./spectrumscale config protocols -f cesSharedRoot -m /ibm/cesSharedRoot
./spectrumscale enable nfs
./spectrumscale enable smb
./spectrumscale enable s3
./spectrumscale node list
./spectrumscale deploy --precheck
./spectrumscale deploy
```

Deploy outcome
- 2 Protocol nodes
- Active SMB, NFS, and S3 file protocols

Next steps
- Configure authentication with **`mmuserauth`**

Example of adding protocols to an existing cluster

Prerequisites configuration
- Decide on a file system to use for **`cesSharedRoot`** (>=4GB). Preferably, a standalone file system solely for this purpose.
- Take note of the file system name and mount point. Verify the file system is mounted on all protocol nodes.
- Decide which nodes will be the Protocol nodes.
- Set aside CES-IPs that are unused in the current cluster and network. Do not attempt to assign the CES-IPs to any adapters.
- Verify each Protocol node has a pre-established network route and IP not only on the GPFS cluster network, but on the same network the CES-IP will belong to. When protocols are deployed, the CES-IPs will be aliased to the active network device matching their subnet. The CES-IPs must be free to move among nodes during failover cases.
- Decide which protocols to enable. The protocol deployment will install all protocols but will enable only the ones you choose.
- Add the new to-be protocol nodes to the existing cluster using **`mmaddnode`** (or use the installation toolkit).
- In this example, we will add the protocol functionality to nodes already within the cluster.

Installation toolkit commands (toolkit is running on a node that will become a protocol node)
```
./spectrumscale setup -s 10.11.10.15                            ← Internal GPFS network IP on the current installer node that can see all protocol nodes
./spectrumscale config populate -n cluster-node5                ← Pick a node in the cluster for the toolkit to use for automatic configuration
./spectrumscale node add cluster-node5 -a -p
./spectrumscale node add cluster-node6 -p
./spectrumscale node add cluster-node7 -p
./spectrumscale node add cluster-node8 -p
./spectrumscale config protocols -e 172.31.1.10,172.31.1.11,172.31.1.12,172.31.1.13,172.31.1.14
./spectrumscale config protocols -f cesSharedRoot -m /ibm/cesSharedRoot
./spectrumscale enable nfs
./spectrumscale enable smb
./spectrumscale enable s3
./spectrumscale callhome enable                                 ← If you prefer not to enable callhome, change the enable to a disable
./spectrumscale callhome config -n COMPANY_NAME -i COMPANY_ID -cn MY_COUNTRY_CODE -e MY_EMAIL_ADDRESS
./spectrumscale node list
./spectrumscale deploy --precheck
./spectrumscale deploy
```

Deploy outcome
- CES Protocol stack added to 4 nodes, now designated as protocol nodes with server licenses
- 4 CES-IPs distributed among the protocol nodes
- Protocol configuration and state data will use the `cesSharedRoot` file system
- The `callhome` will be configured

Example of upgrading protocol nodes or other nodes in the same cluster as an ESS

Pre-upgrade planning
- Refer to IBM Documentation for supported upgrade paths of Storage Scale nodes
- If you have a 5148-22L protocol node attached to an ESS, please refer to the ESS 5.3.7 (or higher) Quick Deployment Guide
- Consider whether OS, FW, or drivers on the protocol node(s) should be upgraded and plan this either before or after the installation toolkit upgrade
- SMB: requires quiescing all I/O for the duration of the upgrade. Due to the SMB clustering functionality, differing SMB levels cannot co-exist within a cluster at the same time. This requires a full outage of SMB during the upgrade.
- NFS: Recommended to quiesce all I/O for the duration of the upgrade. NFS experiences I/O pauses, and depending upon the client, mount may disconnect during the upgrade.
- Performance Monitoring: Collector(s) may experience small durations in which no performance data is logged, as the nodes upgrade.

Installation toolkit commands for IBM Storage Scale 5.0.0.0 or higher
```
./spectrumscale setup -s 10.11.10.11 -st ess                    ← Internal GPFS network IP on the current Installer node that can see all protocol nodes.
./spectrumscale config populate -N ems1                         ← Always point config populate to the EMS node when an ESS is in the same cluster.
```
Note: If **`config populate`** is incompatible with your configuration, add the nodes and CES configuration to the installation toolkit manually.

```
./spectrumscale node list                                       ← This is the list of nodes the installation toolkit will upgrade. Remove any non-CES nodes you would rather do manually.
./spectrumscale upgrade precheck
./spectrumscale upgrade run
```

Example of upgrading protocol nodes or other nodes (not in an ESS)

Pre-upgrade planning
- Refer to IBM Documentation for supported upgrade paths of IBM Storage Scale nodes
- Consider whether OS, FW, or drivers on the protocol node(s) should be upgraded and plan this either before or after the installation toolkit upgrade
- SMB: Requires quiescing all I/O for the duration of the upgrade. Due to the SMB clustering functionality, differing SMB levels cannot co-exist within a cluster at the same time. This requires a full outage of SMB during the upgrade.
- NFS: Recommended to quiesce all I/O for the duration of the upgrade. NFS experiences I/O pauses, and depending upon the client, mount may disconnect during the upgrade.
- Performance Monitoring: Collector(s) may experience small durations in which no performance data is logged, as the nodes upgrade.

Installation toolkit commands
```
./spectrumscale setup -s 10.11.10.11 -st ss                     ← Internal GPFS network IP on the current Installer node that can see all protocol nodes.
./spectrumscale config populate -N <hostname_of_any_node_in_cluster>
```
Note: If **`config populate`** is incompatible with your configuration, add the nodes and CES configuration to the installation toolkit manually.

```
./spectrumscale upgrade config workload -p on                   ← Enable prompt to shut down workloads before upgrade.
./spectrumscale node list                                       ← This is the list of nodes the installation toolkit will upgrade. Remove any non-CES nodes you would rather do manually.
./spectrumscale upgrade precheck
./spectrumscale upgrade run
```

Example of a new IBM Storage Scale native REST API cluster installation followed by a protocol deployment

Installation toolkit commands for installation
The toolkit is running from **`cluster-node1`** with an internal cluster network IP: `10.11.10.11`. All nodes can reach the IP: **`cd /usr/lpp/mmfs/5.2.3.x/ansible-toolkit/`**
```
./spectrumscale setup -s 10.11.10.11
./spectrumscale node add cluster-node1 -a -g
./spectrumscale node add cluster-node2 -a -g
./spectrumscale node add cluster-node3
./spectrumscale node add cluster-node4
./spectrumscale node add cluster-node5 -n
./spectrumscale node add cluster-node6 -n
```

Enable the installation toolkit to install the administration daemon package
```
./spectrumscale scaleadmd enable
```

Import the node identity certificate
```
./spectrumscale nodeid define --cert <path to certificate> --key <path to private key> --chain <path to ca chain> -N all
```

For example:
```
./spectrumscale nodeid define --cert /root/tls/server.pem --key /root/tls/server.key --chain /root/tls/ca.crt -N all
```

For more information about generating a certificate, see *Node identities*.

```
./spectrumscale nsd add -p node5.tuc.stglabs.ibm.com -s node6.tuc.stglabs.ibm.com -u dataAndMetadata -fs cesSharedRoot -fg 1 "/dev/sdb"
./spectrumscale nsd add -p node6.tuc.stglabs.ibm.com -s node5.tuc.stglabs.ibm.com -u dataAndMetadata -fs cesSharedRoot -fg 2 "/dev/sdc"
./spectrumscale nsd add -p node5.tuc.stglabs.ibm.com -s node6.tuc.stglabs.ibm.com -u dataAndMetadata -fs fs1 -fg 1 "/dev/sdh"
./spectrumscale nsd add -p node5.tuc.stglabs.ibm.com -s node6.tuc.stglabs.ibm.com -u dataAndMetadata -fs fs1 -fg 1 "/dev/sdi"
./spectrumscale nsd add -p node5.tuc.stglabs.ibm.com -s node6.tuc.stglabs.ibm.com -u dataAndMetadata -fs fs1 -fg 2 "/dev/sdj"
./spectrumscale nsd add -p node5.tuc.stglabs.ibm.com -s node6.tuc.stglabs.ibm.com -u dataAndMetadata -fs fs1 -fg 2 “/dev/sdk"
./spectrumscale config perfmon -r on
./spectrumscale callhome enable
./spectrumscale callhome config -n COMPANY_NAME -i COMPANY_ID -cn MY_COUNTRY_CODE -e MY_EMAIL_ADDRESS
./spectrumscale config gpfs -c mycluster
./spectrumscale node list
./spectrumscale install --precheck
./spectrumscale install
```

** URLs are subject to change **