



IBM Spectrum Scale 5.0.1 Protocols Quick Overview

May 11, 2018

Before Starting

Always start here to understand:

Common pre-requisites

Basic Install Toolkit operation

Requirements when an existing cluster exists, both with or without an ESS

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

Protocol & File System Deployment

Start here if you already have a cluster and would like to:

Add/Enable protocols on existing cluster nodes

Create a file system on existing NSDs

Configure File or Object Protocol Authentication

Configuration

Start here if you already have a cluster with protocols enabled and would like to:

Check cluster state and health

Configure a basic SMB or NFS export or Object

Configure and Enable File Audit Logging

Upgrade & Product Lifecycle

Start here to gain a basic of understanding of:

Upgrade guidance

How to add nodes, NSDs, FSSs, protocols, auth to an existing cluster

Basic Logging, Debugging, Data Capture

- 1 **How does the Install Toolkit work?**
IBM Spectrum Scale Install Toolkit operation can be summarized by 4 phases:
- 1) User input via 'spectrumscale' commands
 - 2) A 'spectrumscale install' phase
 - 3) A 'spectrumscale deploy' phase
 - 4) A 'spectrumscale upgrade' phase

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

All user input via 'spectrumscale' commands is recorded into a clusterdefinition.txt file in /usr/lpp/mmfs/5.0.1.x/installer/configuration/

Each phase will act 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 OSs, input only the supported protocol nodes and leave all other nodes out of the cluster definition.

- 2 **Hardware / Performance Sizing**
Please work with your IBM account team or Business Partner for suggestions on the best configuration possible to fit your environment. In addition, make sure to review the protocol sizing guide.

- 3 **OS levels and CPU architecture**
The Install Toolkit supports the following OSs:

x86: RHEL7.x, SLES12, Ubuntu16.04
ppc64 BE: RHEL7.x
ppc64 LE: RHEL7.x, SLES12

All cluster nodes the Install Toolkit acts upon must be of the same CPU architecture and endianness.

All protocol nodes must be of the same OS, architecture, and endianness.

- 4 **Repositories**
A base repository must be setup on every node.
RHEL check: yum repolist
SLES12 check: zypper repos
Ubuntu check: apt edit-sources

- 5 **Firewall & Networking & SSH**
All nodes must be networked together and pingable via IP, FQDN, and hostname

Reverse DNS lookup must be in place

If /etc/hosts is used for name resolution, ordering within must be: IP FQDN hostname

Promptless ssh must be setup between all nodes and themselves 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 the IBM Knowledge Center for more details before proceeding.

- 6 **Time sync among nodes is required**
A consistent time must be established on all nodes of the cluster. NTP can be automatically configured during spectrumscale install. See step 9 of the installation stage.

- 7 **Cleanup prior SMB, NFS, Object**
Prior implementations of SMB, NFS, and Object must be completely removed before proceeding with a new protocol deployment. Refer to the cleanup guide within the IBM Knowledge Center.

- 8 **If a GPFS cluster pre-exists**
Proceed to the Protocol Deployment section as long as you have:

- a) file system(s) created and mounted ahead of time & nfs4 ACLs in place
- b) ssh promptless access among all nodes
- c) firewall ports open
- d) CCR enabled
- e) set mmhconfig release=LATEST
- f) installed GPFS rpms should match the exact build dates of those included within the protocols package

- 9 **If an ESS is part of the cluster**
Proceed to the Cluster Installation section to use the Install 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 node(s) must be in the ems nodeclass. IO nodes must be in their own nodeclass: gss or gss_ppc64.
- c) GPFS on the ESS nodes must be at minimum 4.2.0.0
- 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 Knowledge Center.

- 11 **Extract Spectrum Scale package**
The standard, advanced, or data management protocol packages are necessary for protocol deployment. Extracting the package will present a license agreement.

./Spectrum_Scale_Protocols_Data_Management-5.0.1.x-<arch>-Linux-install

- 12 **Explore the spectrumscale help**
From location /usr/lpp/mmfs/5.0.1.x/installer
Use the -h flag for base spectrumscale as well as any subcommand. Examples:
./spectrumscale -h
./spectrumscale setup -h
./spectrumscale node add -h
./spectrumscale config -h
./spectrumscale config protocols -h

- 13 **FAQ and Quick Reference**
Refer to the Knowledge Center Quick Reference
Refer to the Spectrum Scale FAQ

- 1 **Setup the node that will start the installation**
Pick an IP existing on this node which is accessible to/ from all nodes via promptless ssh:

./spectrumscale setup -s IP

Setup in an ESS environment
If the spectrumscale command is being run on a node(s) 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

- 2 **Populate the cluster**
If a cluster pre-exists, the Install Toolkit can automatically traverse the existing cluster and populate its clusterdefinition.txt file with current cluster 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 in ESS mode, point config populate to the EMS:

./spectrumscale config populate -N ems1

* Note the limitations of the config populate command

- 3 **Add NSD server nodes (non-ESS nodes)**
Adding NSD nodes is necessary if you would like the install 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 FSSs)**
File systems are defined by assigning a file system name to one or more NSDs. Filesystems will be defined but not created until this install is followed by a deploy.

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 the IBM Knowledge Center 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 will assign quorum and manager nodes by default. Refer to the IBM Knowledge Center if a specific configuration is desired.

- 7 **Add Spectrum Scale GUI nodes**
./spectrumscale node add hostname -g -a
...

The management GUI will automatically start after installation and allow for further cluster configuration and monitoring.

- 8 **Configure performance monitoring**
Configure performance monitoring consistently across nodes.

./spectrumscale config perfmon -r on

- 9 **Configure network time protocol (NTP)**
The network time protocol can be automatically configured and started on all nodes provided the NTP package has been pre-installed on all nodes:

./spectrumscale config ntp -e on -s ntp_server1,
ntp_server2, ntp_server3, ...

- 10 **Configure Callhome**
Starting with 5.0.0.0, callhome is enabled by default within the Install Toolkit. Refer to the callhome settings and configure mandatory options for callhome:

./spectrumscale callhome config -h

Alternatively, disable callhome:

./spectrumscale callhome disable

- 11 **Name your cluster**
./spectrumscale config gpfs -c my_cluster_name

- 12 **Review your config**
./spectrumscale node list
./spectrumscale nsd list
./spectrumscale filesystem list
./spectrumscale config gpfs --list
./spectrumscale install --precheck

- 13 **Start the installation**
./spectrumscale install

Upon completion you will have an active GPFS cluster with available NSDs, performance monitoring, time sync, callhome, and a GUI. File systems will be 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 nodes
- add NSDs
- enable and configure or update callhome settings

- 1 **Setup the node that will start the installation**
Optionally, the Install Toolkit can automatically traverse the existing cluster and populate its clusterdefinition.txt file with current cluster configuration details. Point it at a node within the cluster with promptless ssh access to all other cluster nodes:

./spectrumscale setup -s IP

Setup in an ESS environment
If the spectrumscale command is being run on a node(s) 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

- 2 **Populate the cluster**
Optionally, the Install Toolkit can automatically traverse the existing cluster and populate its clusterdefinition.txt 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 in ESS mode, point config populate to the EMS:

./spectrumscale config populate -N ems1

*Note the limitations of the config populate command

- 3 **Add protocol nodes**
./spectrumscale node add hostname -p
./spectrumscale node add hostname -p
....

- 4 **Assign protocol IPs (CES-IPs)**
Add a comma separated list of IPs to be used specifically for cluster export services such as NFS, SMB, Object. Reverse DNS lookup must be in place for all IPs. CES-IPs must be unique and different than cluster node IPs.

./spectrumscale config protocols -e EXPORT_IP_POOL

*All protocol nodes must see the same CES-IP network(s). 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

*Skip this step if you setup file systems / NSDs manually and not through the install toolkit.

- 6 **Configure protocols to point to a shared root file system location**
A ces directory will be automatically created at root of the specified file system mount point. This is used for protocol admin/config and needs >=4GB free. Upon completion of protocol deployment, GPFS configuration will point to this as cesSharedRoot. It is recommended that cesSharedRoot be a separate file system.

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

*If you setup file systems / NSDs manually, perform a manual check of <mmnsd> and <mmfs all -L> to make sure all NSDs and file systems required by the deploy are active and mounted before continuing.

- 7 **Enable the desired file protocols**
./spectrumscale enable nfs
./spectrumscale enable smb

- 8 **Enable the Object protocol if desired**
./spectrumscale enable object

Configure an admin user, password, and database password to be used for Object operations:

./spectrumscale config object -au admin -ap -dp

Configure the Object endpoint using a single hostname with a round robin DNS entry mapping to all CES IPs:

./spectrumscale config object -e hostname

Specify a file system and filesset name where your Object data will go:

./spectrumscale config object -f fs1 -m /ibm/fs1
./spectrumscale config object -o Object_Fileset

*The Object filesset must not pre-exist. If an existing filesset is detected at the same location, deployment will fail so that existing data is preserved.

- 9 **Setup Authentication**
Authentication must be setup prior to using any protocols. If you are unsure of the appropriate authentication config you may skip this step and revisit by re-running the deployment at a later time or manually using the mmuserauth commands. Refer to the IBM Knowledge Center for the many supported authentication configurations.

Install Toolkit AD example for File and/or Object
./spectrumscale auth file ad
./spectrumscale auth object ad

- 10 **Configure Callhome**
Starting with 5.0.0.0, callhome is enabled by default within the Install Toolkit. Refer to the callhome settings and configure mandatory options for callhome:

./spectrumscale callhome config -h

Alternatively, disable callhome:

./spectrumscale callhome disable

- 11 **Review your config**
./spectrumscale node list
./spectrumscale deploy --precheck

- 12 **Start the deployment**
./spectrumscale deploy

Upon completion you will have protocol nodes with active cluster export services and IPs. File systems will have been created and Authentication will be configured and ready to use. Performance Monitoring tools will also be ready at this time.

Deploy can be re-run in the future to:

- enable additional protocols
- enable authentication for file or Object
- create additional file systems (run install first to add more NSDs)
- add additional protocol nodes (run install first to add more nodes)
- enable and configure or update callhome settings

- 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

Basic GPFS Health

mmgetstate -aL
mmiscluster
mmiscluster --ces
mmnetverify

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

Authentication

mmuserauth service list
mmuserauth service check

Callhome

mmcalhome info list
mmcalhome group list
mmcalhome status list

File protocols (NFS & SMB)

Verify all file systems to be used with protocols have nfs4 ACLs and locking in place. Protocols will not work correctly without this setting in effect. Check with: mmisfs 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
,Protocols=3:4)"

mmnfs export list

Example SMB export creation:
mkdir /ibm/fs1/smb_export1

chown "DOMAINUSER" /ibm/fs1/smb_export1

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

mmmsmb export list

Object protocol

Verify the Object protocol by listing users and uploading an object to a container:

source \$HOME/openrc
openstack user list
openstack project list
swift stat
date > test_object1.txt
swift upload test_container test_object1.txt
swift list test_container

Performance Monitoring

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

File Audit Logging

File audit logging functionality is available with Advanced and Data Management Editions of Spectrum Scale.

a) Enable and configure using the Install Toolkit as follows:

./spectrumscale fileauditinglogging enable
./spectrumscale filesystem modify --fileauditingloggingenable gpfs1
./spectrumscale fileauditinglogging 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
mmmsgqueue status
mmaudit all consumeStatus -N <node list>

Config Populate of Install Toolkit

Cluster configuration changes may be made outside of the Install Toolkit. If planning to use the Install Toolkit after configuration changes, it's recommended to update the Install Toolkit to match the current cluster configuration.

a) Automatically update the Install Toolkit to match the current config by pointing it to any node within the cluster that has ssh access to/from all other nodes and run config populate.

./spectrumscale config populate -N hostname

If an ESS is in the cluster, point config populate at the EMS node:

./spectrumscale config populate -N ems1

OR

b) If config populate hits any of the known limitations, you may have to update the Install Toolkit config manually in order to match your current cluster configuration.

Upgrading 4.1.1.x to 5.0.1.x:

A direct path from 4.1.1.x to 5.0.1.x is not possible unless all nodes of the cluster are offline. However, it is possible to upgrade first, from 4.1.1.x to 4.2.x.x, and second, from 4.2.x.x to 5.0.1.x, while the cluster is online.

Upgrading 4.2.x.x to 5.0.1.x

a) Extract the 5.0.1.x Spectrum Scale PTF package
./Spectrum_Scale_Protocols_Data_Management-5.0.1.x-Linux

b) Setup and Configure the Install Toolkit

./spectrumscale setup -s <IP of installer node>
./spectrumscale config populate -N <any cluster node>

"If config populate is incompatible with your cluster config, you will have to manually add the nodes and config to the Install Toolkit OR copy the last used clusterdefinition.txt file to the new 5.0.1.x Install Toolkit." cp -p /usr/lpp/mmfs/<4.2.x.x>your_last_level>installer/configuration/
clusterdefinition.txt /usr/lpp/mmfs/5.0.1.x/installer/configuration/

./spectrumscale node list
./spectrumscale nsd list
./spectrumscale filesystem list
./spectrumscale config gpfs
./spectrumscale config protocols
./spectrumscale upgrade --precheck
./spectrumscale upgrade

"The Install Toolkit can be used to upgrade a subset of nodes if desired. In this case, only input the nodes desired for an upgrade into the Install Toolkit. Likewise, remove any nodes added by config populate that are not desired for an upgrade at this time"

c) Run the upgrade precheck:

cd /usr/lpp/mmfs/4.2.3.x/installer/
./spectrumscale upgrade -pr

d) Run the upgrade:

./spectrumscale upgrade

Upgrading 5.0.1.x to future PTFs

Follow the same procedure as indicated above.

Upgrade compatibility with LTFS-EE

a) Itfsee stop (on all LTFSEE nodes)
b) umount /lfts (on all LTFSEE nodes)
c) dsmmigfs disableallover (on all LTFSEE nodes)
d) dsmmigfs stop (on all LTFSEE nodes)
e) systemctl stop hsm.service (on all LTFSEE nodes)
f) Upgrade using the Install Toolkit
g) Upgrade LTFS-EE if desired
h) Reverse steps e through a and restart/enable

Upgrade compatibility with TCT

a) Stop TCT on all nodes prior to the upgrade
mmioadgateway service stop -N Node1 Nodeclass
b) Upgrade using the Install Toolkit
c) Upgrade the TCT rpm(s) manually, then restart TCT

Handling Linux kernel updates

The GPFS portability layer must be rebuilt on every node that undergoes a Linux kernel update. Once the new kernel is active, typically after a reboot, rebuild the GPFS portability layer on each node with this command prior to starting GPFS: /usr/lpp/mmfs/bin/mmbuildgpl

Adding to the installation

The procedures below can be combined to reduce the number of installs and deploys 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 node(s):
./spectrumscale install -pr
./spectrumscale install
c) If a protocol node is being added, also run deploy
./spectrumscale deploy -pr
./spectrumscale deploy

To add an NSD:

To service the NSD server connecting this new disk runs an OS compatible with the install toolkit and that the NSD server exists within the cluster.
b) Add the NSD(s) to the install 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 about by the install toolkit
b) Define the file system
./spectrumscale nsd list
./spectrumscale nsd modify nsdX -fs file_system_name
c) Deploy the new file system
./spectrumscale deploy -pr
./spectrumscale deploy

To enable another protocol:

See the Protocol & File System deployment column.
Proceed with steps 7, 8, 9, 10, 11. Note that some protocols necessitate removal of the Authentication configuration prior to enablement.

To add authentication

See the Protocol & File System deployment column.
Proceed with steps 9, 10, 11

Service Procedures

To service a specific protocol node, simply suspend the node; the CES IPs will automatically move to another node

mmces node suspend -N Node1,Node2..

To service all protocol nodes or the shared root file system, proceed with an mmshutdown -N Node1, Node2..

<

Examples

Example of readying Red Hat 7.x nodes for Spectrum 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 Spectrum Scale functionality*)

```
systemctl stop firewalld
systemctl disable firewalld
- repeat on all nodes
```

How to check if a yum repository is configured correctly

yum repolist -> should return no errors. It must also show an RHEL7.x base repository. Other repository possibilities include a satellite site, a custom yum repository, an RHEL7.x DVD iso, an RHEL7.x physical DVD.

How to configure a yum repository from an RHEL DVD iso image (*an iso repo is not required but a base OS repository is required*)

- *Alternatively, a satellite site or custom yum repository can be used so long as all RHEL7.x base rpms are available within*

```
mkdir /mnt/dvdiso
mount -o loop /path_to_dvd_iso/dvd.iso /mnt/dvdiso
cp /mnt/dvdiso/media.repo /etc/yum.repos.d/rhel7dvd.repo
chmod 644 /etc/yum.repos.d/rhel7dvd.repo
vi /etc/yum.repos.d/rhel7dvd.repo
--change--
gpgcheck=1
--add--
enabled=1
baseurl=file:///mnt/rhel7dvd
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-redhat-release
--save/exit--
yum clean all
subscription-manager clean
yum list --noplugin
```

Pre-install pre-req rpms to make installation and deployment easier

```
yum install kernel-devel gcc gcc-c++ glibc sssd ypbind openldap-clients
```

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 Spectrum Scale commands (*not required*)

```
vi /root/.bash_profile
--add this line--
export PATH=$PATH:/usr/lpp/mmfs/bin
--save/exit--
logout and back in for changes to take effect
```

Example of a new Spectrum Scale cluster installation followed by a protocol deployment

Install 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/5.0.1.x/installer/
./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 ObjectFS -fg 1 "/dev/sdd"
./spectrumscale nsd add -p node5.tuc.stglabs.ibm.com -s node6.tuc.stglabs.ibm.com -u dataAndMetadata -fs ObjectFS -fg 1 "/dev/sde"
./spectrumscale nsd add -p node6.tuc.stglabs.ibm.com -s node5.tuc.stglabs.ibm.com -u dataAndMetadata -fs ObjectFS -fg 2 "/dev/sdf"
./spectrumscale nsd add -p node5.tuc.stglabs.ibm.com -s node6.tuc.stglabs.ibm.com -u dataAndMetadata -fs ObjectFS -fg 2 "/dev/sdg"
./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 config ntp -e on -s ntp_server1,ntp_server2,ntp_server3
./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 config gpfs -c mycluster
./spectrumscale node list
./spectrumscale install --precheck
./spectrumscale install
```

Install Outcome: A 6node Spectrum Scale cluster with active NSDs

- 2 GUI nodes
- 2 NSD nodes
- 2 client nodes
- 10 NSDs
- configured performance monitoring
- callhome configured
- **3 file systems defined, each with 2 failure groups. File systems will not be created until a deployment**

Install Toolkit commands for Protocol Deployment (*assumes cluster created from above configuration./*

- Toolkit is running from the same node that performed the install above, 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 object
./spectrumscale config object -e mycluster-ces
./spectrumscale config object -o Object_Fileset
./spectrumscale config object -f ObjectFS -m /ibm/ObjectFS
./spectrumscale config object -au admin -ap -dp
./spectrumscale node list
./spectrumscale deploy --precheck
./spectrumscale deploy
```

Deploy Outcome:

- 2 Protocol nodes
- Active SMB and NFS file protocols
- Active Object protocol
- cesSharedRoot file system created and used for protocol configuration and state data
- ObjectFS file system created with an Object_Fileset created within
- fs1 file system created and ready

Next Steps:

- Configure Authentication with mmuserauth or by configuring authentication with the Install Toolkit and re-running the deployment

Example of adding protocols to an existing cluster

Pre-req 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-IPs 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 Install Toolkit).
- In this example, we will add the protocol functionality to nodes already within the cluster.

Install 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 object
./spectrumscale config object -e mycluster-ces
./spectrumscale config object -o Object_Fileset
./spectrumscale config object -f ObjectFS -m /ibm/ObjectFS
./spectrumscale config object -au admin -ap -dp
./spectrumscale config object -e mycluster-ces
./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
- Object protocol will use the ObjectFS filesystem

Example of adding protocol nodes to an ESS

Starting point

- 1) The cluster containing ESS is active and online
- 2) RHEL7.x, SLES12, or Ubuntu16.04 is installed on all nodes that are going to serve as protocol nodes
- 3) RHEL7.x, SLES12, or Ubuntu 16.04 base repository is set up on nodes that are going to serve as protocol nodes
- 4) The nodes that will serve as protocol nodes have connectivity to the GPFS cluster network
- 5) Use the ESS GUI or CLI to create a CES shared root file system >=4GB.
- 6) 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.
- 7) Use the ESS GUI or CLI to create additional file systems for protocols if desired. Configure each file system for nfsv4 ACLs
- 8) Enable CCR on the existing GPFS cluster if it is not already enabled
- 9) Pick a protocol node to run the Install Toolkit from.
- 10) The Install Toolkit is contained within these packages: Spectrum Scale Protocols Standard or Advanced or Data Management Edition
- 11) Download and extract one of the Spectrum Scale Protocols packages to the protocol node that will run the Install Toolkit
- 12) Once extracted, the Install Toolkit is located in the /usr/lpp/mmfs/5.0.1.x/installer directory.
- 13) Inputting the configuration into the Install Toolkit with the commands detailed below, involves pointing the Install Toolkit to the EMS node, telling the Install Toolkit about the mount points and paths to the CES shared root and optionally, the Object file systems, and designating the protocol nodes and protocol config to be installed/deployed.

Install 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 Install Toolkit traverse the existing cluster and auto-populate its config.
./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 Install 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 object
./spectrumscale config object -e mycluster-ces
./spectrumscale config object -o Object_Fileset
./spectrumscale config object -f ObjectFS -m /ibm/ObjectFS
./spectrumscale config object -au admin -ap -dp
./spectrumscale node list
```

```
./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's 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 provided they are part of the GPFS cluster
```

Install Outcome:

- EMS node used as an admin node by the Install 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 Install 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
- Object protocol will use the ObjectIFS filesystem, which was pre-created on the ESS

Example of Upgrading protocol nodes / other nodes in the same cluster as an ESS

Pre-Upgrade planning:

- Refer to the Knowledge Center for supported upgrade paths of Spectrum Scale nodes
- *"Caution" if protocol nodes are included in the ESS xcat config: updatenode scripts are unaware of CES suspend/upgrade requirements*
- Consider whether OS, FW, or drivers on the protocol node(s) should be upgraded and plan this either before or after the install 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, mounts may disconnect during the upgrade.
- Object: Recommended to quiesce all I/O for the duration of the upgrade. Object service will be down or interrupted at multiple times during the upgrade process. Clients may experience errors or they might be unable to connect during this time. They should retry as appropriate.
- Performance Monitoring: Collector(s) may experience small durations in which no performance data is logged, as the nodes upgrade.

Install Toolkit commands for Scale 5.0.0.0 or higher ***If upgrading to a Scale version below 5.0.0.0, specifically refer to the ESS 5.2 QDG chapter 2***

```
./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
** If config populate is incompatible with your configuration, add the nodes and CES configuration to the install toolkit manually **
```

```
./spectrumscale node list <- This is the list of nodes the Install Toolkit will upgrade. Remove any non-CES nodes you would rather do manually
./spectrumscale upgrade --precheck
./spectrumscale upgrade
```

Example of Upgrading protocol nodes / other nodes (not in an ESS)

Pre-Upgrade planning:

- Refer to the Knowledge Center for supported upgrade paths of Spectrum Scale nodes
- Consider whether OS, FW, or drivers on the protocol node(s) should be upgraded and plan this either before or after the install 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, mounts may disconnect during the upgrade.
- Object: Recommended to quiesce all I/O for the duration of the upgrade. Object service will be down or interrupted at multiple times during the upgrade process. Clients may experience errors or they might be unable to connect during this time. They should retry as appropriate.
- Performance Monitoring: Collector(s) may experience small durations in which no performance data is logged, as the nodes upgrade.

Install 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>
** If config populate is incompatible with your configuration, add the nodes and CES configuration to the install toolkit manually **
```

```
./spectrumscale node list <- This is the list of nodes the Install Toolkit will upgrade. Remove any non-CES nodes you would rather do manually
./spectrumscale upgrade --precheck
./spectrumscale upgrade
```