

Before Starting

- 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/mmts/4.2.3.x/installer/configuration/
- 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.
- OS levels and CPU architecture**
 The Install Toolkit supports the following OSs:

 x86: RHEL6.x, RHEL7.x, SLES12, Ubuntu16.04 (4.2.3.1)
 ppc64 BE: RHEL6.x, 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.
- Repositories**
 A base repository must be setup on every node.
 RHEL check: `yum repolist`
 SLES12 check: `zypper repos`
 Ubuntu check: `apt edit-sources`
- 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.
- 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.
- 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.
- 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 `mmchconfig release=LATEST`
 f) installed GPFS rpms should match the exact build dates of those included within the protocols package
- If an ESS is part of the cluster**
 a) CCR must be enabled
 b) ESS nodes must be in their own node class: `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) Protocol nodes must be manually added to the ESS cluster prior to performing the protocol deployment step
- 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-4.2.3.x-<arch>-Linux-install`
- Explore the spectrumscale help**
 From location `/usr/lpp/mmts/4.2.3.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`
- FAQ and Quick Reference**
 Refer to the Knowledge Center Quick Reference
 Refer to the Spectrum Scale FAQ

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
- 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`

 If the spectrumscale command is being run from a location outside of any of the nodes to be installed, a GPFS admin node is required:

`./spectrumscale node add hostname -a`
 - Populate the cluster**
 *Skip this step if this is a new cluster creation

 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`
 * Note the limitations of the config populate command
 - Add NSD server 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`

 *Do not add ESS nodes to the Install Toolkit
 - Add NSDs**
 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`
 - Define file systems**
 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.
 - 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.
 - 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.
 - Configure performance monitoring**
 Configure performance monitoring consistently across nodes.

`./spectrumscale config perfmon -i on`
 - 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, ...`
 - Name your cluster**
`./spectrumscale config gpfs -c my_cluster_name`
 - Review your config**
`./spectrumscale node list`
`./spectrumscale nsd list`
`./spectrumscale filesystem list`
`./spectrumscale config gpfs --list`
`./spectrumscale install --precheck`
 - Start the installation**
`./spectrumscale install`

 Upon completion you will have an active GPFS cluster with available NSDs, performance monitoring, time sync, 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

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
- Setup the node that will start the installation**
 Setup is necessary unless spectrumscale setup had previously been run on this node for a past GPFS installation or protocol deployment. Pick an IP existing on this node which is accessible to/from all nodes via promptless ssh:

`./spectrumscale setup -s IP`

 If the spectrumscale command is run from a location outside of any of the nodes to be installed, a GPFS admin node is required.

`./spectrumscale node add hostname -a`
 - Populate the cluster**
 Optionally, starting with 4.2.3, 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`
 *Note the limitations of the config populate command
 - Add protocol nodes**
`./spectrumscale node add hostname -p`
`./spectrumscale node add hostname -p`

 - 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.
 - 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.
 - 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 <mmifs all -L> to make sure all NSDs and file systems required by the deploy are active and mounted before continuing.
 - Enable the desired file protocols**
`./spectrumscale enable nfs`
`./spectrumscale enable smb`
 - 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 fileset name where your Object data will go:

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

 *The Object fileset must not pre-exist. If an existing fileset is detected at the same location, deployment will fail so that existing data is preserved.
 - 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`
 - Review your config**
`./spectrumscale node list`
`./spectrumscale deploy --precheck`
 - 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 usable 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)

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
 - Test basic Object operation
- 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/mmts/bin`
 - Basic GPFS Health**
`mmgetstate -al`
`mmiscuster`
`mmiscuster -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`
 - File protocols (NFS & SMB)**
 Verify all file systems to be used with protocols have nfs4 ACLs and locking in effect. Protocols will not work correctly without this setting in place.
 Check with: `mmifs 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`

`mm smb export add smb_export1 /ibm/fs1/smb_export1 --option "browseable=yes"`

`mm smb 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`

Upgrade & Product Lifecycle

- Upgrading 4.1.1.x to 4.2.2.x:**
 CES Protocols require a hop from 4.1.1.x to 4.2.0.0 prior to upgrading to 4.2.3.x. Please see the 4.1.1.x to 4.2.0.0 upgrade instructions prior to returning here for the 4.2.3.x upgrade.
- Upgrading 4.2.x.x to 4.2.3.x**
 a) Extract the 4.2.3.x Spectrum Scale PTF package
`./Spectrum_Scale_Protocols_Data_Management-4.2.3.x-Linux`
 b) Copy the clusterdefinition.txt file from the last 4.2.x.x installer location to the 4.2.3.x installer:
`cp -p /usr/lpp/mmts/4.2.x.x/your_last_level/installer/configuration/clusterdefinition.txt /usr/lpp/mmts/4.2.3.x/installer/configuration/`
 c) Verify current cluster config is matched
`./spectrumscale node list`
`./spectrumscale nsd list`
`./spectrumscale filesystem list`
`./spectrumscale config gpfs`
`./spectrumscale config protocols`

 *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!
- Run the upgrade precheck:**
`cd /usr/lpp/mmts/4.2.3.x/installer/`
`./spectrumscale upgrade -pr`
- Run the upgrade:**
`./spectrumscale upgrade`
- Upgrading 4.2.3.x to future PTFs**
 Follow the same procedure as indicated above. Optionally, for Step b, the Install Toolkit can automatically traverse the existing cluster and populate its clusterdefinition.txt file with current cluster details:

`./spectrumscale config populate -n hostname`
- Upgrade compatibility with LTFS-EE**
 a) `ltsf stop` (on all LTFS-EE nodes)
 b) `umount /ltsf` (on all LTFS-EE nodes)
 c) `dsmmglfs disable/over` (on all LTFS-EE nodes)
 d) `systemctl stop hsm.service` (on all LTFS-EE nodes)
 e) use the toolkit to upgrade
 f) upgrade LTFS-EE if desired
- Upgrade compatibility with TCT**
 a) Suspend all TCT function prior to upgrade
 b) Use the toolkit to upgrade
 c) Upgrade the TCT rpms(s), then resume 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/mmts/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:
 a) Verify 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..`
- Logging & debugging**
 Installation / deployment:
`./usr/lpp/mmts/4.2.3.x/installer/logs`

 Verbose logging for all spectrumscale commands by adding a '-v' immediately after `./spectrumscale`:
`./usr/lpp/mmts/4.2.3.x/installer/spectrumscale -v <cmd>`

 GPFS default log location:
`/var/adm/ras/`

 Linux syslog or journal is recommended to be enabled
- Data Capture for Support**
System-wide data capture:
`./usr/lpp/mmts/bin/gpfs.snap`

Installation/Deploy/Upgrade specific:
`./usr/lpp/mmts/4.2.3.x/installer/installer.snap.py`
- Further IBM Spectrum Scale Education**
 Best Practices, hints, tips, videos, white papers, and up to date news regarding IBM Spectrum Scale can be found on the [IBM Spectrum Scale wiki](#).

Examples

Example of readying RHEL7 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 cpp gcc gcc-c++ glibc sssd ybind openldap-clients
```

Turn off selinux (*or set to permissive mode*)

```
setstatus
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 adding protocol nodes to an ESS

Starting point

- 1) The cluster containing ESS is active and online
- 2) RHEL7.x is installed on nodes that are going to serve as protocol nodes
- 3) RHEL7.x 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) Extract the Spectrum Scale Protocols Standard or Advanced package to each protocol node
- 7) /usr/lpp/mmfs/4.2.3.x/gpfs_rpms contains the GPFS levels you'll want to install on the protocol nodes
- 8) Install gpfs.base, gpfs.ext, gpfs.gpl, gpfs.gskit, gpfs.docs, gpfs.msg, gpfs.rpms from /usr/lpp/mmfs/4.2.3.0/gpfs_rpms on each to-be protocol node. Do NOT install performon, gui, callhome, java, nor protocols rpms at this time.
- 9) Run /usr/lpp/mmfs/bin/mmbuildgpl on each to-be protocol node
- 10) Add the protocol nodes to the existing ESS cluster using mmaddnode
- 11) Start the newly added nodes by running /usr/lpp/mmfs/bin/mmstartup -N nodeA, nodeB, nodeC, ...
- 12) Enable CCR on the existing GPFS cluster if it is not already enabled
- 13) Re-configure the existing cluster so that quorum/manager functions reside only upon nodes equal to the same GPFS version installed on the protocol nodes
- 14) Mount the CES shared root file system on the protocol nodes if it is not already mounted
- 15) Use the ESS GUI or CLI to create additional file systems for protocols if desired. Configure each file system for nfsv4 ACLs
- 16) Pick a protocol node to run the Install Toolkit from. The install toolkit is located in the Spectrum Scale Protocols Standard or Advanced package in /usr/lpp/mmfs/4.2.3.x/installer/
- 17) When inputting the configuration into the Install Toolkit, designate only the nodes that you have planned to use as protocol nodes. Do not designate ESS nodes such as EMS or I/O nodes as protocol nodes. Do not input EMS or I/O nodes into the Install Toolkit
- 18) Point the Install Toolkit to existing file systems and mount points only. Do not attempt to create new file systems or NSDs using the Install Toolkit.

Install Toolkit commands:

```
Do NOT use the new ./spectrumscale config populate 4.2.3 command if an ESS is in the cluster
./spectrumscale setup -s 10.11.10.11 <- internal gpfs network IP on the current installer node that can see all protocol nodes
./spectrumscale config perfmon -r off
./spectrumscale node add cluster-node1 -a -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 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, which was pre-created on the ESS
- Object protocol will use the ObjectFS filesystem, which was pre-created on the ESS

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/4.2.3.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 node6.tuc.stglabs.ibm.com -s node5.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 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
- **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.
- 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 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

