

IBM Spectrum Scale Install Toolkit Overview

May 17th, 2018 GPFS User's Group (Boston)

Aaron Palazzolo (aspalazz@us.ibm.com)

Muthu Muthiah (mutmuthi@in.ibm.com)



Background of the Install Toolkit

Background of Install Toolkit – where?

Where can I get it?

Spectrum Scale Protocols Standard package

Spectrum Scale Protocols Advanced package

Spectrum Scale Protocols Data Management package

Where is the Install Toolkit extracted?

`/usr/lpp/mmfs/5.0.1.0/installer/`

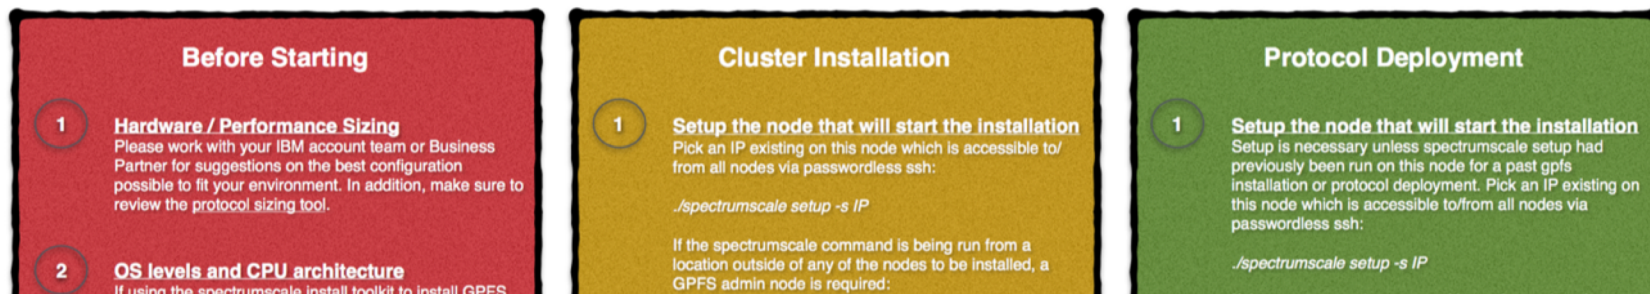
How do I first use it?

Type: `/usr/lpp/mmfs/5.0.1.0/installer/spectrumscale -h`

Background of Install Toolkit – finding help

- Is there a quick overview showing how to use the Install Toolkit?

[Spectrum Scale Protocols Quick Overview on the GPFS Wiki](#)



the self extracting package will give a URL link to this guide as well

- Is the Install Toolkit documented in the Knowledge Center?

See the Knowledge Center sections for installing GPFS on Linux nodes

Background of Install Toolkit – popular questions

Why would I want to use the Install Toolkit?

1. *Simplify GPFS cluster creation
2. Automatically create NSD and file system stanza files
3. Create new NSDs, file systems, and add new nodes
4. Single package manager for all components of a release
5. Deploy Object, SMB, NFS by automatically pulling in all pre-reqs
6. *Install and Deploy File Audit Logging
7. *Consistently set and sync time on all nodes
8. *Deploy the Spectrum Scale GUI
9. Configure Performance Monitoring consistently across all nodes
10. *Configure callhome across all nodes
11. Simplify upgrade with a single command to upgrade all components on all nodes

Background of Install Toolkit – popular questions

- **I heard there are pre-requisites?**

x86: RHEL7.x, SLES12 SPx, Ubuntu 16.04.x (*see FAQ for exact versions*)

PPC64 / PPC64le: RHEL7.x

Repos + same CPU architecture on all nodes to be acted upon

SSH access with no password prompts nor Y/N questions to and among all nodes

Firewall ports for the toolkit itself: 8889, 10080, 22

Firewall ports for the Install GUI: 9443 (https) 9080 (http)

Firewall ports for GPFS: 1191 + ephemeral port range

- **What if there are unsupported OSs or a mix of CPU architectures among nodes?**

Pick a grouping of nodes that are similar and only configure the Install Toolkit to be aware of these nodes. Manually add other nodes to the cluster. Manually add functions to the other nodes. Manually upgrade the other nodes.

Background of Install Toolkit – phases of the toolkit

What happens and when?

The Install Toolkit operation can be broken down into a few phases:

- 1) User input via 'spectrumscale' commands
- 2) A 'spectrumscale install' phase
- 3) A 'spectrumscale deploy' phase
- 4) A 'spectrumscale upgrade' phase

******Very Important******

Each phase can be run again at later points in time to introduce new NSDs, file systems, nodes. Note that we only support additions at this point.

Background of Install Toolkit – phases of the toolkit

What happens and when?

User Input Details:

- 1) User input via 'spectrumscale' commands:
 - a) All user input is recorded into a clusterdefinition.txt file in `/usr/lpp/mmfs/5.0.1.0/installer/configuration/`
 - b) As you input your cluster config, remember that you can have the toolkit act on parts of the cluster by simply not telling it about nodes that may have incompatible OS levels, architectures, etc...
 - c) Optionally the, `./spectrumscale config populate -N <node in cluster>`, command can be used to traverse an existing cluster and automatically populate the clusterdefinition.txt file.

Background of Install Toolkit – phases of the toolkit

What happens and when?

Install Details:

2) A 'spectrumscale install' phase:

- a) Install will act upon all nodes inputted into the clusterdefinition.txt file
- b) GPFS and perfmon rpms will be installed
- c) GPFS portability layer will be created
- d) GPFS will be started
- e) A cluster will be created
- f) Server and Client licenses will be applied
- g) GUI nodes may be created and the GUI may be started on these nodes
- h) NTP, perfmon, FAL, callhome, ephemeral ports, cluster profile... may be configured
- i) NSDs may be created - *note* filesystems are not created during install

Background of Install Toolkit – phases of the toolkit

What happens and when?

Deploy Details:

2) A 'spectrumscale deploy' phase:

- a) Deploy will act upon all nodes inputted into the clusterdefinition.txt file
- b) File systems will be configured - **note** it is possible to only configure file systems during deploy if you do not want protocols
- c) File Audit Logging (*FAL*) may be enabled
- d) SMB, NFS, Object protocol rpms will be copied to all protocol nodes
- e) SMB, NFS, Object services may be started
- f) Authentication may be configured

Background of Install Toolkit – phases of the toolkit

What happens and when?

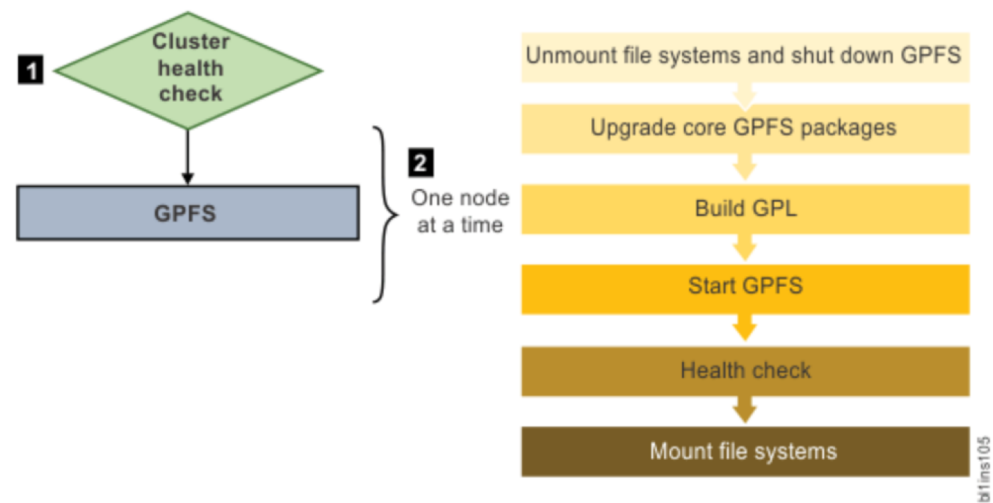
[See this KC link for details on the upgrade flow](#)

Upgrade Details:

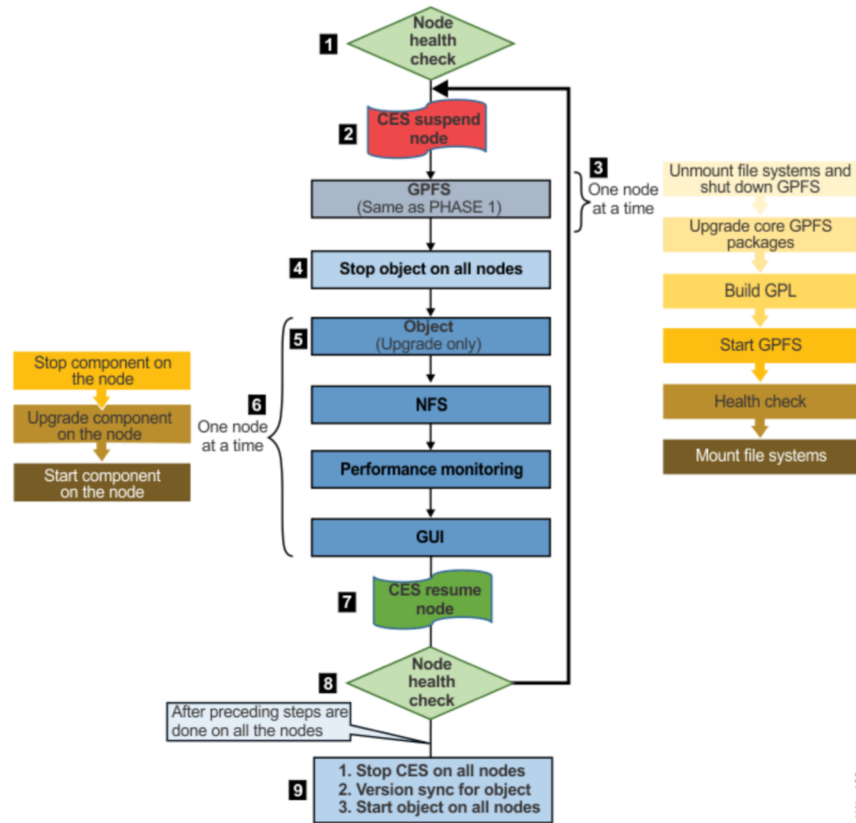
3) A 'spectrumscale upgrade' phase:

- a) Upgrade will act upon all nodes inputted into the clusterdefinition.txt file
- b) All installed/deployed components will be upgraded
- c) Upgrades are sequential with multiple passes
- d) Pass 1 of all nodes will upgrade GPFS sequentially
- e) Pass 2 of all nodes will upgrade Object sequentially
- f) Pass 3 of all nodes will upgrade NFS sequentially
- g) Pass 4 of all nodes will upgrade SMB sequentially

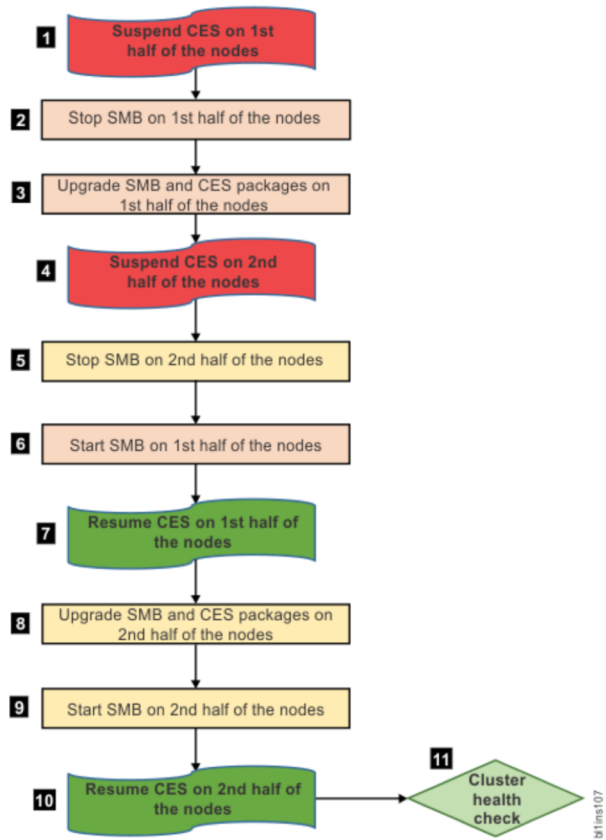
Phase 1 - upgrade non-protocol nodes



Phase 2 - upgrade protocol nodes (stage 1)



Phase 2 - upgrade protocol nodes (stage 2)



Logging and Debug of the Install Toolkit

Logging and Debug (General)

- The Install Toolkit reports progress to the console. Log files kept in the *logs* subdirectory of the `/usr/lpp/mmfs/5.0.1.x/installer` location will contain verbose output
- **Log file naming convention = action the toolkit was running followed by the timestamp.** Example: a log for './spectrumscale deploy' would look like: `DEPLOY-13-05-2018_17:56:53.log`
- **If a step fails, the Install Toolkit will provide information on the failure and suggested actions**
- **The `installer.snap.py` script can be run to collect all debug information. It will gather:**
 - Cluster configuration file
 - Install Toolkit logs
 - Information about protocol services
 - A GPFS snap from every node
 - Basic environment information

Logging and Debug – Knowledge Center Troubleshooting Updates

Example of troubleshooting an Install Toolkit failure during deploy

- 1) The failure will always indicate the log file to look at. Open up this log file
- 2) Go to the very end but ignore the last errors as they merely indicate that Chef encountered an error.

Don't spend much time analyzing this - we know a Deploy failed and the Log name

```
2016-01-15 15:31:14,912 [ TRACE ] Stopping chef zero
2016-01-15 15:31:14,913 [ ERROR ] The following error was encountered:
Traceback (most recent call last):
  File "/usr/lpp/mmfs/4.2.0.0/installer/espypylib/reporting.py", line 222, in log_to_file
    yield handler
  File "/usr/lpp/mmfs/4.2.0.0/installer/espypylib/install.py", line 167, in _install
    setup.install(config)
  File "/usr/lpp/mmfs/4.2.0.0/installer/espypylib/setup/ces.py", line 325, in install
    self.deploy(config.protocol_nodes, options_fn)
  File "/usr/lpp/mmfs/4.2.0.0/installer/espypylib/deploy.py", line 133, in deploy_nodes
    raise DeployError()
DeployError: Installation failed on one or more nodes. Check the log for more details.
2016-01-15 15:31:14,957 [ INFO ] Detailed error log: /usr/lpp/mmfs/4.2.0.0/installer/logs/DEPLOY-15-01-2016_15:29:59.log
```

Logging and Debug – Knowledge Center Troubleshooting Updates

Example of troubleshooting an Install Toolkit failure

- 3) Search back in time for the topmost occurrence of the word: FATAL
- 4) The word FATAL may occur multiple times, so make sure you find the first occurrence time-wise.

```
2016-01-15 15:31:09,447 [ TRACE ] objnode4 [2016-01-15T15:31:09+05:30] ERROR: Running exception handlers
2016-01-15 15:31:09,447 [ TRACE ] objnode4 Running handlers complete
2016-01-15 15:31:09,447 [ TRACE ] objnode4 [0m[2016-01-15T15:31:09+05:30] ERROR: Exception handlers complete
2016-01-15 15:31:09,448 [ TRACE ] objnode4 [2016-01-15T15:31:09+05:30] FATAL: Stacktrace dumped to /var/chef/cache/chef-stacktrace.out
2016-01-15 15:31:09,448 [ TRACE ] objnode4 Chef Client failed. 32 resources updated in 46.185382251 seconds[0m
2016-01-15 15:31:09,474 [ TRACE ] objnode4 [2016-01-15T15:31:09+05:30] ERROR: yum_package[spectrum-scale-object] (swift_on_gpfs::swift_node_install line 14) had an error: Ch
```

This example shows that Chef recipe `swift_node_install` failed

Logging and Debug – Knowledge Center Troubleshooting Updates

Example of troubleshooting an Install Toolkit failure

- 5) Manually scroll a bit upwards (backwards in time) and visually inspect the log
- 6) You're looking for an Error message encompassed in “=====” signs

We've now found the exact failure. Note the yum command and the text indicating a specific version of libcap-ng which is necessary to continue the deploy.

To recover: remove the conflicting version and install the required version

```
[ INFO ] [objnode3 15-01-2016 15:30:51] IBM SPECTRUM SCALE: Installing Object packages (SS50)
[ TRACE ] objnode3 * log[IBM SPECTRUM SCALE: Installing Object packages (SS50).] action write
[ TRACE ] objnode3
[ TRACE ] objnode4 [0m * yum_package[spectrum-scale-object] action install
[ TRACE ] objnode4 [0m
[ TRACE ] objnode4 =====[0m
[ TRACE ] objnode4 [31mError executing action `install` on resource `yum_package[spectrum-scale-object]`[0m
[ TRACE ] objnode4 =====[0m
[ TRACE ] objnode4
[ TRACE ] objnode4 [0m Chef::Exceptions::Exec[0m
[ TRACE ] objnode4 -----[0m
[ TRACE ] objnode4 yum -d0 -e0 -y install spectrum-scale-object-4.2.0-0 returned 1:
[ TRACE ] objnode4 [0m STDOUT: You could try using --skip-broken to work around the problem
[ TRACE ] objnode4 [0m You could try running: rpm -Va --nofiles --nodigest
[ TRACE ] objnode4 [0m
[ TRACE ] objnode4 [0m STDERR: Error: Package: libcap-ng-python-0.7.3-5.el7.x86_64 (ces_object)
[ TRACE ] objnode4 [0m Requires: libcap-ng = 0.7.3-5.el7
[ TRACE ] objnode4 [0m Installed: libcap-ng-0.7.5-4.el7.x86_64 (@RHEL7.1)
[ TRACE ] objnode4 [0m libcap-ng = 0.7.5-4.el7
[ TRACE ] objnode4 [0m
[ TRACE ] objnode4 [0m Resource Declaration:[0m
[ TRACE ] objnode4 [0m -----[0m
[ TRACE ] objnode4 # In /var/chef/cache/cookbooks/swift_on_gpfs/recipes/swift_node_install.rb
[ TRACE ] objnode4 [0m
[ TRACE ] objnode4 [0m 14: package pkg do
[ TRACE ] objnode4 [0m 15:   retries 3
[ TRACE ] objnode4 [0m 16:   retry_delay 3
[ TRACE ] objnode4 [0m 17: end
[ TRACE ] objnode4 [0m 18: end
```

Additional Content:

Step by Step Guides

Additional Content:

Setting up the Install Toolkit

Added Content – setting up the Install Toolkit

Extract the Spectrum Scale package

```
# ./Spectrum_Scale_Protocols_Data_Management-5.0.1.0-x86_64-Linux-install
```

Agree to the license

LICENSE INFORMATION

The Programs listed below are licensed under the following License Information terms and conditions in addition to the Program license terms previously agreed to by Client and IBM. If Client does not have previously agreed to license terms in effect for the Program, the International Program License Agreement (Z125-3301-14) applies.

Program Name (Program Number):

IBM Spectrum Scale Data Management Edition V5.0.1 (5737-F34)
IBM Spectrum Scale Data Management Edition V5.0.1 (5641-DM1)
IBM Spectrum Scale Data Management Edition V5.0.1 (5641-DM3)
IBM Spectrum Scale Data Management Edition V5.0.1 (5641-DM5)

Press Enter to continue viewing the license agreement, or **enter "1" to accept the agreement**, "2" to decline it, "3" to print it, "4" to read non-IBM terms, or "99" to go back to the previous screen.

1

License Agreement Terms accepted.

Added Content – setting up the Install Toolkit

=====

Product rpms successfully extracted to /usr/lpp/mmfs/5.0.1.0

Cluster installation and protocol deployment

To install a cluster or deploy protocols with the Spectrum Scale Install Toolkit: /usr/lpp/mmfs/5.0.1.0/installer/spectrumscale -h

To install a cluster manually: Use the gpfs rpms located within /usr/lpp/mmfs/5.0.1.0/gpfs_rpms

To upgrade an existing cluster using the Spectrum Scale Install Toolkit:

- 1) Copy your old clusterdefinition.txt file to the new /usr/lpp/mmfs/5.0.1.0/installer/configuration/ location
- 2) Review and update the config: /usr/lpp/mmfs/5.0.1.0/installer/spectrumscale config update
- 3) (Optional) Update the toolkit to reflect the current cluster config:
/usr/lpp/mmfs/5.0.1.0/installer/spectrumscale config populate -N <node>
- 4) Run the upgrade: /usr/lpp/mmfs/5.0.1.0/installer/spectrumscale upgrade -h

To add nodes to an existing cluster using the Spectrum Scale Install Toolkit:

- 1) Add nodes to the clusterdefinition.txt file: /usr/lpp/mmfs/5.0.1.0/installer/spectrumscale node add -h
- 2) Install GPFS on the new nodes: /usr/lpp/mmfs/5.0.1.0/installer/spectrumscale install -h
- 3) Deploy protocols on the new nodes: /usr/lpp/mmfs/5.0.1.0/installer/spectrumscale deploy -h

To add NSDs or file systems to an existing cluster using the Spectrum Scale Install Toolkit:

- 1) Add nsds and/or filesystems with: /usr/lpp/mmfs/5.0.1.0/installer/spectrumscale nsd add -h
- 2) Install the NSDs: /usr/lpp/mmfs/5.0.1.0/installer/spectrumscale install -h
- 3) Deploy the new file system: /usr/lpp/mmfs/5.0.1.0/installer/spectrumscale deploy -h

To update the toolkit to reflect the current cluster config examples:

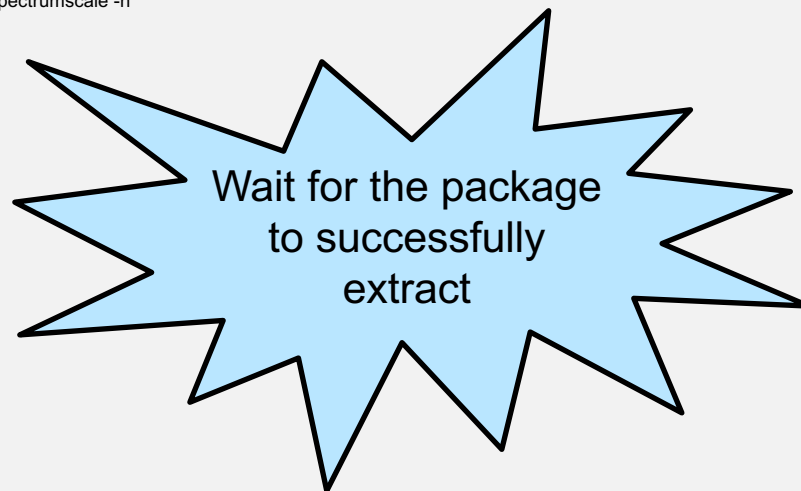
- /usr/lpp/mmfs/5.0.1.0/installer/spectrumscale config populate -N <node>
- 1) Manual updates outside of the install toolkit
 - 2) Sync the current cluster state to the install toolkit prior to upgrade
 - 3) Switching from a manually managed cluster to the install toolkit

=====

To get up and running quickly, visit our wiki for an IBM Spectrum Scale Protocols Quick Overview:

<https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/General%20Parallel%20File%20System%20%28GPFS%29/page/Protocols%20Quick%20Overview%20for%20IBM%20Spectrum%20Scale>

=====



Added Content – setting up the Install Toolkit

```
# cd /usr/lpp/mmfs/5.0.1.0/installer
```

Change to the Installer directory and check out the Install Toolkit (spectrumscale) help

```
## ./spectrumscale -h
```

```
usage: spectrumscale [-h] [-v] [--version]
```

```
{setup,node,config,nsd,filesystem,auth,installgui,callhome,fileauditlogging,enable,disable,install,deploy,upgrade}
```

This script will configure and install Spectrum Scale, setting up your nodes as specified in the configuration files.

positional arguments: {setup,node,config,nsd,filesystem,auth,installgui,callhome,fileauditlogging,enable,disable,install,deploy,upgrade}

setup	Install Chef components and activate control server
node	List, add or remove nodes from the install
config	Set properties in the clusterdefinition settings file
nsd	List, add or remove NSDs from the install
filesystem	List or modify filesystems in the install
auth	Configure authentication for Object and File protocols
installgui	Configures and starts install GUI
callhome	Configures and manages the callhome
fileauditlogging	list the file audit logging configuration
enable	Enable object, smb or nfs in the protocol configuration
disable	Disable object, smb or nfs in the protocol configuration
install	Install GPFS on the configured nodes
deploy	Deploy file systems and protocols on the configured nodes
upgrade	Upgrade GPFS, Object, NFS, SMB and Performance Monitoring components on the cluster

optional arguments:

-h, --help	show this help message and exit
-v, --verbose	Turn on verbose logging for traceback on error
--version	show program's version number and exit

For usage of sub-commands, specify a command followed by -h or --help

Added Content – setting up the Install Toolkit

Verify pre-reqs are met – Check the quick overview guide and the Knowledge Center

Setup the node that will perform the installation

In this case we've selected the node we are running the command from.

- This node has access to all other nodes via SSH with no Y/N prompts
- All required firewall ports were opened on all nodes, prior to this

```
# ./spectrumscale setup -s 9.11.102.38 --setuptype ss
[ INFO ] Installing prerequisites for install node
[ INFO ] Chef successfully installed and configured
[ INFO ] Your control node has been configured to use the IP 9.11.102.38 to communicate with other nodes.
[ INFO ] Port 8889 will be used for chef communication.
[ INFO ] Port 10080 will be used for package distribution.
[ INFO ] Install Toolkit setup type is set to Spectrum Scale (default). If an ESS is in the cluster, run this command to set ESS mode: ./spectrumscale setup -s server_ip -st ess
[ INFO ] SUCCESS
[ INFO ] Tip : Designate protocol, nsd and admin nodes in your environment to use during install: ./spectrumscale -v node add <node> -p -a -n
```

- If an ESS is in the cluster, see the next page

Added Content – setting up the Install Toolkit

If an ESS is in the cluster, use 'ess' as the setuptype

```
# ./spectrumscale setup -s 9.11.102.38 --setuptype ess
```

```
[ INFO ] Installing prerequisites for install node
```

```
[ INFO ] Existing Chef installation detected. Ensure chef-zero is included and the PATH is configured.
```

```
[ INFO ] Your control node has been configured to use the IP 9.11.102.38 to communicate with other nodes.
```

```
[ INFO ] Port 8889 will be used for chef communication.
```

```
[ INFO ] Port 10080 will be used for package distribution.
```

```
[ INFO ] Install Toolkit setup type is set to ESS. This mode will allow the EMS node to execute Install Toolkit commands.
```

```
[ INFO ] SUCCESS
```

```
[ INFO ] Tip : Designate an EMS node as admin node: ./spectrumscale node add <node> -a
```

```
[ INFO ] Tip : After designating an EMS node, add nodes for the toolkit to act upon: ./spectrumscale node add <node> -p -n
```

```
[ INFO ] Tip : After designating the EMS node, if you want to populate the cluster definition file with the current configuration, you can run: ./spectrumscale config populate -N <ems_node>
```

Additional Content:

Installing GPFS (new cluster)

Added Content – Installing GPFS (new cluster)

```
[root@protocol-node1] # ./spectrumscale node list
```

```
[ INFO ] List of nodes in current configuration:
```

```
[ INFO ] [Installer Node]
```

```
[ INFO ] 9.11.102.38
```

```
[ INFO ]
```

```
[ INFO ] Setup Type: SpectrumScale
```

```
[ INFO ]
```

```
[ INFO ] [Cluster Name]
```

```
[ INFO ] No cluster name configured
```

```
[ INFO ]
```

```
[ INFO ] File Audit logging : Disabled
```

```
[ INFO ]
```

```
[ INFO ] No nodes configured. Use 'spectrumscale node add' to add nodes.
```

```
# ./spectrumscale node add protocol-node1 -a
```

```
[ INFO ] Adding node protocol-node1.tuc.stglabs.ibm.com as a GPFS node.
```

```
[ INFO ] Setting protocol-node1.tuc.stglabs.ibm.com as an admin node.
```

```
[ INFO ] Configuration updated.
```

```
[ INFO ] Tip : Designate protocol or nsd nodes in your environment to use during install:./spectrumscale node add <node> -p -n
```

```
# ./spectrumscale node add protocol-node2
```

```
[ INFO ] Adding node protocol-node2.tuc.stglabs.ibm.com as a GPFS node.
```

```
# ./spectrumscale node add client-node2
```

```
[ INFO ] Adding node client-node2.tuc.stglabs.ibm.com as a GPFS node.
```

The toolkit configuration is empty except for the installer node

- You'll be adding nodes and configuration details.
- The Installer node info is from the prior `./spectrumscale setup` command

Let's add a node

- First is the node that will administer the installation
- The admin node will need promptless ssh to and from all other nodes.
- We'll designate this as a protocol node later, for now it will be a client node

Now we add 2 more client nodes

Added Content – Installing GPFS (new cluster)

Let's add a GUI node

- GUI nodes must be added as admin nodes, GUIs must have promptless ssh connectivity to/from all other nodes
- Skip forward a few slides and you'll see what happens if you forget the -a

```
# ./spectrumscale node add client-node1 -g -a
[ INFO ] Adding node client-node1.tuc.stglabs.ibm.com as a GPFS node.
[ INFO ] Setting client-node1.tuc.stglabs.ibm.com as an admin node.
[ INFO ] Configuration updated.
[ INFO ] Tip : Designate protocol or nsd nodes in your environment to use during install:./spectrumscale node add <node> -p -n
[ INFO ] Setting client-node1.tuc.stglabs.ibm.com as a GUI server.
```

Let's add 2 NSD nodes

- If you don't want to add NSDs or FSs at this time, leave off the -n flag so they are marked as client nodes

```
# ./spectrumscale node add nsd-node1 -n
[ INFO ] Adding node nsd-node1.tuc.stglabs.ibm.com as a GPFS node.
[ INFO ] Adding node nsd-node1.tuc.stglabs.ibm.com as an NSD server.
[ INFO ] Configuration updated.
[ INFO ] Tip :If all node designations are complete, add NSDs to your cluster definition and
define required filesystems:./spectrumscale nsd add <device> -p <primary node> -s <secondary node> -fs <file system>
```

```
# ./spectrumscale node add nsd-node2 -n
[ INFO ] Adding node nsd-node2.tuc.stglabs.ibm.com as a GPFS node.
[ INFO ] Adding node nsd-node2.tuc.stglabs.ibm.com as an NSD server.
[ INFO ] Configuration updated.
[ INFO ] Tip :If all node designations are complete, add NSDs to your cluster definition and
define required filesystems:./spectrumscale nsd add <device> -p <primary node> -s <secondary node> -fs <file system>
```


Added Content – Installing GPFS (new cluster)

The Install Toolkit should now show all the nodes we've added

```
# ./spectrumscale node list
```

```
[ INFO ] List of nodes in current configuration:
[ INFO ] [Installer Node]
[ INFO ] 9.11.102.38
[ INFO ]
[ INFO ] Setup Type: SpectrumScale
[ INFO ]
[ INFO ] [Cluster Name]
[ INFO ] No cluster name configured
[ INFO ]
[ INFO ] File Audit logging : Disabled
[ INFO ]
[ INFO ] GPFS Node                               Admin  Quorum  Manager  NSD Server  Protocol  GUI Server  Callhome Node  OS      Arch
[ INFO ] client-node1.tuc.stglabs.ibm.com          X
[ INFO ] client-node2.tuc.stglabs.ibm.com
[ INFO ] nsd-node1.tuc.stglabs.ibm.com                X
[ INFO ] nsd-node2.tuc.stglabs.ibm.com                X
[ INFO ] protocol-node1.tuc.stglabs.ibm.com          X
[ INFO ] protocol-node2.tuc.stglabs.ibm.com
[ INFO ] client-node3.tuc.stglabs.ibm.com          X
[ INFO ]
[ INFO ] [Export IP address]
[ INFO ] No export IP addresses configured
```

Notice no Quorum/Manager nodes are defined yet

If this is a new cluster create

- The toolkit will automatically define these for you during precheck
- OR you can define exact nodes for quorum/manager as you desire

If this is an existing cluster that you're adding nodes to

- The toolkit will not automatically define any.

Added Content – Installing GPFS (new cluster)

Since we already added NSD nodes, let's add NSDs

- In this example, nsd-node1 & nsd-node2 share the same physical disks, we'll alternate primary/secondary server to keep things balanced
- All NSDs in this example will be set for both dataAndMetadata
- Half of the NSDs assigned to each FS will be in failure group 1 and the other half, failure group 2
- We're designating 3 FSs for these NSDs: cesSharedRoot, ObjectFS, and fs1. The Install Toolkit will create the NSDs during install, if they don't already exist. And the FSs during deploy if they don't already exist

```
# ./spectrumscale nsd add -p nsd-node1 -s nsd-node2 -u dataAndMetadata -fs cesSharedRoot -fg 1 "/dev/sdb"
```

```
[ INFO ] Connecting to nsd-node1.tuc.stglabs.ibm.com to check devices and expand wildcards.
```

```
[ INFO ] Looking up details of /dev/sdb.
```

```
[ INFO ] The installer will create the new file system cesSharedRoot if it does not exist.
```

```
[ INFO ] Adding NSD None on nsd-node1.tuc.stglabs.ibm.com using device /dev/sdb.
```

```
[ INFO ] Configuration updated
```

```
[ INFO ] Tip : If all node designations and any required protocol configurations are complete, proceed to check the installation configuration: ./spectrumscale install --precheck
```

```
# ./spectrumscale nsd add -p nsd-node2 -s nsd-node1 -u dataAndMetadata -fs cesSharedRoot -fg 2 "/dev/sdc"
```

```
[ INFO ] Connecting to nsd-node2.tuc.stglabs.ibm.com to check devices and expand wildcards.
```

```
[ INFO ] Looking up details of /dev/sdc.
```

```
[ INFO ] Adding NSD None on nsd-node2.tuc.stglabs.ibm.com using device /dev/sdc.
```

```
[ INFO ] Configuration updated
```

```
[ INFO ] Tip : If all node designations and any required protocol configurations are complete, proceed to check the installation configuration: ./spectrumscale install --precheck
```

```
# ./spectrumscale nsd add -p nsd-node1 -s nsd-node2 -u dataAndMetadata -fs ObjectFS -fg 1 "/dev/sdd"
```

```
# ./spectrumscale nsd add -p nsd-node2 -s nsd-node1 -u dataAndMetadata -fs ObjectFS -fg 2 "/dev/sde"
```

```
# ./spectrumscale nsd add -p nsd-node1 -s nsd-node2 -u dataAndMetadata -fs fs1 -fg 1 "/dev/sdf"
```

```
# ./spectrumscale nsd add -p nsd-node1 -s nsd-node2 -u dataAndMetadata -fs fs1 -fg 1 "/dev/sdg"
```

```
# ./spectrumscale nsd add -p nsd-node2 -s nsd-node1 -u dataAndMetadata -fs fs1 -fg 2 "/dev/sdh"
```

```
# ./spectrumscale nsd add -p nsd-node2 -s nsd-node1 -u dataAndMetadata -fs fs1 -fg 2 "/dev/sdi"
```

Added Content – Installing GPFS (new cluster)

Here's how the NSDs will be setup

- NSDs are created during the install phase if they do not already exist
- If performing an install, deploy, or upgrade and NSDs already exist – the toolkit does not need to be told of them

```
# ./spectrumscale nsd list
```

	INFO	Name	FS	Size (GB)	Usage	FG	Pool	Device	Servers
[INFO]	nsd1	cesSharedRoot	10	dataAndMetadata	1	Default	/dev/sdb	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com	
[INFO]	nsd3	ObjectFS	250	dataAndMetadata	1	Default	/dev/sdd	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com	
[INFO]	nsd5	fs1	750	dataAndMetadata	1	Default	/dev/sdf	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com	
[INFO]	nsd6	fs1	750	dataAndMetadata	1	Default	/dev/sdg	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com	
[INFO]	nsd2	cesSharedRoot	10	dataAndMetadata	2	Default	/dev/sdc	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com	
[INFO]	nsd4	ObjectFS	250	dataAndMetadata	2	Default	/dev/sde	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com	
[INFO]	nsd7	fs1	750	dataAndMetadata	2	Default	/dev/sdh	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com	
[INFO]	nsd8	fs1	750	dataAndMetadata	2	Default	/dev/sdi	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com	

Here's how the file systems will be setup

- File systems are created during the deploy phase if their NSDs already exist
- If performing an install, deploy, or upgrade and FSs already exist – the toolkit does not need to be told of them

```
# ./spectrumscale filesystem list
```

	INFO	Name	BlockSize	Mountpoint	NSDs Assigned	Default Data Replicas	Max Data Replicas	Default Metadata Replicas	Max
Metadata Replicas									
[INFO]	cesSharedRoot	Default	(4M)/ibm/cesSharedRoot	2	1	2	2	2	2
[INFO]	ObjectFS	Default	(4M)/ibm/ObjectFS	2	1	2	2	2	2
[INFO]	fs1	Default	(4M)/ibm/fs1	4	1	2	2	2	

Added Content – Installing GPFS (new cluster)

Verify the GPFS settings are as expected

```
# ./spectrumscale config gpfs
[ INFO ] No changes made to defaults. Current settings are as follows:
[ WARN ] No value for GPFS cluster name in clusterdefinition file.
[ INFO ] GPFS profile is default.
[ INFO ] Remote shell command is /usr/bin/ssh.
[ INFO ] Remote file copy command is /usr/bin/scp.
[ WARN ] No value for GPFS Daemon communication port range in clusterdefinition file.
```

Let's set the cluster name

```
# ./spectrumscale config gpfs -c democloud.tuc.stglabs.ibm.com
[ INFO ] Setting GPFS cluster name to democloud.tuc.stglabs.ibm.com
```

The warning for cluster name is now gone

```
# ./spectrumscale config gpfs
[ INFO ] No changes made to defaults. Current settings are as follows:
[ INFO ] GPFS cluster name is democloud.tuc.stglabs.ibm.com
[ INFO ] GPFS profile is default.
[ INFO ] Remote shell command is /usr/bin/ssh.
[ INFO ] Remote file copy command is /usr/bin/scp.
[ WARN ] No value for GPFS Daemon communication port range in clusterdefinition file.
```

Two warnings are displayed:

1) No cluster name is set

- If this is a new cluster, the toolkit will name the cluster after the admin node unless you set the name yourself.
- If you are adding nodes to an existing cluster, you must set the cluster name to match the existing cluster.

2) There are no ports set for the GPFS Daemon network to use

- This is OK as long as all firewalls are down and all ports opened
- You must set a port range if you intend to keep firewalls up. See Knowledge Center topic on Firewalls for more information.

Added Content – Installing GPFS (new cluster)

Let's check the Callhome settings

- With 5.0.0.0 and higher, the toolkit will enable callhome by default
- A precheck failure will be encountered unless configured or disabled

```
# ./spectrumscale callhome config
[ INFO ] No changes made to defaults. Current settings are as follows:
[ INFO ] Callhome is enabled.
[ INFO ] Schedule is daily,weekly.
[ WARN ] No value for Customer name in clusterdefinition file.
[ WARN ] No value for Customer id in clusterdefinition file.
[ WARN ] No value for Customer email in clusterdefinition file.
[ WARN ] No value for Customer country in clusterdefinition file.
[ WARN ] No value for Proxy IP in clusterdefinition file.
[ WARN ] No value for Proxy port in clusterdefinition file.
[ WARN ] No value for Proxy user in clusterdefinition file.
[ WARN ] No value for Proxy password in clusterdefinition file.
```

Here's an example of a basic callhome setup

```
# ./spectrumscale callhome config -n MyUserName -i IBM -e myemail@company.com -cn US -a
[ INFO ] Setting Customer name to MyUserName
[ INFO ] Setting Customer id to IBM
[ INFO ] Setting Customer email to myemail@company.com
[ INFO ] Setting Customer country to US
```

Added Content – Installing GPFS (new cluster)

We're about ready to install – let's run a pre-check first

```
# ./spectrumscale install --precheck
```

```
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/INSTALL-PRECHECK-12-05-2018_13:27:42.log
```

```
[ WARN ] Only one GUI server specified. The Graphical User Interface will not be highly available.
```

```
[ FATAL ] Invalid configuration for setting up GUI.
```

```
[ INFO ] Detailed error log: /usr/lpp/mmfs/5.0.1.0/installer/logs/INSTALL-PRECHECK-12-05-2018_13:27:42.log
```

```
[ FATAL ] A GUI server must also be a GPFS admin node. Ensure the following nodes have passwordless SSH to all other nodes and are specified as admin nodes using  
'./spectrumscale node add <node hostname or IP> -a': client-node1.tuc.stglabs.ibm.com
```

This example shows that we forgot to mark the GUI node as an admin node.

Here's how:

```
# ./spectrumscale node add client-node1 -g -a
```

```
[ INFO ] Setting client-node1.tuc.stglabs.ibm.com as an admin node.
```

```
[ INFO ] Configuration updated.
```

```
[ INFO ] Tip : Designate protocol or nsd nodes in your environment to use during install:./spectrumscale node add <node> -p -n
```

```
[ INFO ] Setting client-node1.tuc.stglabs.ibm.com as a GUI server.
```

Added Content – Installing GPFS (new cluster)

Time for another pre-check – this time it's a success

```
# ./spectrumscale install --precheck
```

```
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/INSTALL-PRECHECK-12-05-2018_13:38:17.log
[ WARN ] Only one GUI server specified. The Graphical User Interface will not be highly available.
[ WARN ] Install toolkit will not configure file audit logging as it has been disabled.
[ INFO ] Checking for knife bootstrap configuration...
[ INFO ] Performing Chef (deploy tool) checks.
[ INFO ] Performing GPFS checks.
[ INFO ] Running environment checks
[ INFO ] Skipping license validation as no existing GPFS cluster detected.
[ WARN ] No quorum nodes are configured. The Install Toolkit will assign quorum nodes.
[ WARN ] No manager nodes are configured. The Install Toolkit will assign manager nodes.
[ INFO ] Checking pre-requisites for portability layer.
[ INFO ] GPFS precheck OK
[ INFO ] Performing NSDs checks.
[ INFO ] Performing Performance Monitoring checks.
[ INFO ] Running environment checks for Performance Monitoring
[ INFO ] Performing GUI checks.
[ INFO ] Performing callhome checks.
[ INFO ] Checking input for callhome
[ INFO ] Checking network connectivity for callhome configuration.
[ WARN ] No call home node is specified. The installation toolkit will assign one of the nodes as the call home node.
[ INFO ] callhome precheck OK
[ INFO ] Network check from admin node protocol-node1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from admin node client-node3.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ WARN ] Ephemeral port range is not set. Please set valid ephemeral port range using the command ./spectrumscale config gpfs --ephemeral_port_range . You may set the default values as 60000-61000
[ INFO ] The install toolkit will configure call home as it is enabled.
[ INFO ] Pre-check successful for install.
[ INFO ] Tip : ./spectrumscale install
```

Added Content – Installing GPFS (new cluster)

One final look at the node list

- Notice that after running an install precheck, the quorum/manager and callhome nodes were automatically added
- If there are no GUI nodes, the toolkit will automatically designate the installer node as an admin node as well

```
# ./spectrumscale node list
[ INFO ] List of nodes in current configuration:
[ INFO ] [Installer Node]
[ INFO ] 9.11.102.38
[ INFO ]
[ INFO ] Setup Type: SpectrumScale
[ INFO ]
[ INFO ] [Cluster Name]
[ INFO ] democluster.tuc.stglabs.ibm.com
[ INFO ]
[ INFO ] File Audit logging : Disabled
[ INFO ]
[ INFO ]
```

GPFS Node	Admin	Quorum	Manager	NSD Server	Protocol	GUI Server	Callhome Node	OS	Arch
client-node1.tuc.stglabs.ibm.com								ubuntu16	x86_64
client-node2.tuc.stglabs.ibm.com							X	sles12	x86_64
nsd-node1.tuc.stglabs.ibm.com		X	X	X				rhel7	x86_64
nsd-node2.tuc.stglabs.ibm.com		X	X	X				rhel7	x86_64
protocol-node1.tuc.stglabs.ibm.com	X	X						rhel7	x86_64
protocol-node2.tuc.stglabs.ibm.com								rhel7	x86_64

```
[ INFO ]
[ INFO ] [Export IP address]
[ INFO ] No export IP addresses configured
```


Added Content – Installing GPFS (new cluster)

Start an install

- Although this example shows an install that will create a new cluster, keep in mind that install can be run again in the future, to add nodes, roles, NSDs, or functions.

```
# ./spectrumscale install
```

```
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/INSTALL-09-05-2018_12:52:29.log
```

```
[ INFO ] Validating configuration
```

Added Content – Installing GPFS (new cluster)

A successful installation

- This gave us a GPFS cluster, NSDs, GUI nodes, and Performance Monitoring

```
[ INFO ] Checking for a successful install
[ INFO ] Checking state of Chef (deploy tool)
[ INFO ] Chef (deploy tool) ACTIVE
[ INFO ] Checking state of GPFS
[ INFO ] GPFS callhome has been successfully installed. To configure callhome run 'mmcallhome -h' on one of your nodes.
[ INFO ] 3 server licenses and 1 client license enabled on the cluster. Ensure enough licenses are owned for this configuration.
[ INFO ] Checking state of GPFS on all nodes
[ INFO ] GPFS active on all nodes
[ INFO ] GPFS ACTIVE
[ INFO ] Checking state of NSDs
[ INFO ] NSDs ACTIVE
[ INFO ] Checking state of Performance Monitoring
[ INFO ] Running Performance Monitoring post-install checks
[ INFO ] pmcollector running on all nodes
[ INFO ] pmsensors running on all nodes
[ INFO ] Performance Monitoring ACTIVE
[ INFO ] Checking state of GUI
[ INFO ] Running Graphical User Interface post-install checks
[ INFO ] Graphical User Interface running on all GUI servers
[ INFO ] Enter one of the following addresses into a web browser to access the Graphical User Interface: client-node1.tuc.stglabs.ibm.com
[ INFO ] GUI ACTIVE
[ INFO ] Checking state of callhome
[ INFO ] callhome ACTIVE
[ INFO ] SUCCESS
[ INFO ] All services running
[ INFO ] StanzaFile and NodeDesc file for NSD, filesystem, and cluster setup have been saved to /usr/lpp/mmfs folder on node: protocol-node1.tuc.stglabs.ibm.com
[ INFO ] Installation successful. 6 GPFS nodes active in cluster democluster.tuc.stglabs.ibm.com. Completed in 5 minutes 36 seconds.
[ INFO ] Tip :If all node designations and any required protocol configurations are complete, proceed to check the deploy configuration:./spectrumscale deploy -precheck
```

Added Content – Installing GPFS (new cluster)

Verify cluster state with the following commands

```
# mmlscluster
```

```
GPFS cluster information
```

```
=====
```

```

GPFS cluster name:      democuster.tuc.stglabs.ibm.com
GPFS cluster id:        2254960516004008510
GPFS UID domain:        democuster.tuc.stglabs.ibm.com
Remote shell command:   /usr/bin/ssh
Remote file copy command: /usr/bin/scp
Repository type:        CCR

```

Node	Daemon node name	IP address	Admin node name	Designation
1	protocol-node1.tuc.stglabs.ibm.com	9.11.102.38	protocol-node1.tuc.stglabs.ibm.com	quorum-perfmon
2	nsd-node1.tuc.stglabs.ibm.com	9.11.102.62	nsd-node1.tuc.stglabs.ibm.com	quorum-manager-perfmon
3	nsd-node2.tuc.stglabs.ibm.com	9.11.102.66	nsd-node2.tuc.stglabs.ibm.com	quorum-manager-perfmon
4	protocol-node2.tuc.stglabs.ibm.com	9.11.102.55	protocol-node2.tuc.stglabs.ibm.com	perfmon
6	client-node2.tuc.stglabs.ibm.com	9.11.102.69	client-node2.tuc.stglabs.ibm.com	perfmon
7	client-node1.tuc.stglabs.ibm.com	9.11.102.67	client-node1.tuc.stglabs.ibm.com	perfmon

```
# mmgetstate -a -L
```

Node number	Node name	Quorum	Nodes up	Total nodes	GPFS state	Remarks
1	protocol-node1	2	3	6	active	quorum node
2	nsd-node1	2	3	6	active	quorum node
3	nsd-node2	2	3	6	active	quorum node
4	protocol-node2	2	3	6	active	
6	client-node2	2	3	6	active	
7	client-node1	2	3	6	active	

Added Content – Installing GPFS (new cluster)

Verify cluster health

```
# mmhealth cluster show
```

Component	Total	Failed	Degraded	Healthy	Other
-----	-----	-----	-----	-----	-----
NODE	6	0	0	0	6
GPFS	6	0	0	0	6
NETWORK	6	0	0	6	0
DISK	8	0	0	8	0
PERFMON	6	0	0	6	0
THRESHOLD	6	0	0	6	0

```
# mmhealth cluster show NODE
```

Component	Node	Status	Reasons
-----	-----	-----	-----
NODE	protocol-node1.tuc.stglabs.ibm.com	TIPS	gpfs_maxfilestocache_small, gpfs_maxstatcache_high, gpfs_pagepool_small
NODE	nsd-node2.tuc.stglabs.ibm.com	TIPS	gpfs_maxstatcache_high, gpfs_pagepool_small
NODE	nsd-node1.tuc.stglabs.ibm.com	TIPS	gpfs_maxstatcache_high, gpfs_pagepool_small
NODE	protocol-node2.tuc.stglabs.ibm.com	TIPS	gpfs_maxfilestocache_small, gpfs_maxstatcache_high, gpfs_pagepool_small
NODE	client-node1.tuc.stglabs.ibm.com	TIPS	gpfs_maxstatcache_high, gpfs_pagepool_small
NODE	client-node2.tuc.stglabs.ibm.com	TIPS	gpfs_maxstatcache_high, gpfs_pagepool_small

Added Content – Installing GPFS (new cluster)

Verify NSDs were created

```
# mmlsnsd
```

File system	Disk name	NSD servers
(free disk)	nsd1	nsd-node1.tuc.stglabs.ibm.com,nsd-node2.tuc.stglabs.ibm.com
(free disk)	nsd2	nsd-node2.tuc.stglabs.ibm.com,nsd-node1.tuc.stglabs.ibm.com
(free disk)	nsd3	nsd-node1.tuc.stglabs.ibm.com,nsd-node2.tuc.stglabs.ibm.com
(free disk)	nsd4	nsd-node2.tuc.stglabs.ibm.com,nsd-node1.tuc.stglabs.ibm.com
(free disk)	nsd5	nsd-node1.tuc.stglabs.ibm.com,nsd-node2.tuc.stglabs.ibm.com
(free disk)	nsd6	nsd-node1.tuc.stglabs.ibm.com,nsd-node2.tuc.stglabs.ibm.com
(free disk)	nsd7	nsd-node2.tuc.stglabs.ibm.com,nsd-node1.tuc.stglabs.ibm.com
(free disk)	nsd8	nsd-node2.tuc.stglabs.ibm.com,nsd-node1.tuc.stglabs.ibm.com

..but remember that file systems are not created until the deploy phase

```
# mmlsmount all
mmremote: No file systems were found.
mmlsmount: Command failed. Examine previous error messages to determine cause.
```

Added Content – Installing GPFS (new cluster)

Check callhome

```
# mmcallhome info list
group      customerName  customerID  email                      countryCode
-----
global     MyUserName      IBM         myemail@company.com       US

# mmcallhome group list
callHomeGroup  callHomeNode  callHomeChildNodes
-----
autoGroup_1    client-node2  client-node1,client-node2,nsd-node1,nsd-node2,protocol-node1,protocol-node2

# mmcallhome capability list
group      callHomeNode  status
-----
global     ---          enabled
[root@protocol-node1 installer]# mmcallhome status list
Group      Task          Start time      Status      Package file name
autoGroup_1  daily        20180510020202.371  minor error  ...roup_1.gat_daily.g_daily.20180510020202371.c10.DC
```

Added Content – Installing GPFS (new cluster)

Check the perfmon config

Sensors available after initial install are shown on the next page

```
# mmperfmon config show
# This file has been generated automatically and SHOULD NOT
# be edited manually. It may be overwritten at any point
# in time.

cephMon = "/opt/IBM/zimon/CephMonProxy"
cephRados = "/opt/IBM/zimon/CephRadosProxy"
colCandidates = "client-nodel.tuc.stglabs.ibm.com", "nsd-
nodel.tuc.stglabs.ibm.com"
colRedundancy = 2
collectors = {
host = ""
port = "4739"
}
config = "/opt/IBM/zimon/ZIMonSensors.cfg"
ctdbstat = ""
daemonize = T
hostname = ""
ipfixinterface = "0.0.0.0"
logfile = "/var/log/zimon/ZIMonSensors.log"
loglevel = "info"
mmcmd = "/opt/IBM/zimon/MMCmdProxy"
mmdfcmd = "/opt/IBM/zimon/MMDFProxy"
mmpmon = "/opt/IBM/zimon/MmpmonSockProxy"
piddir = "/var/run"
release = "5.0.1-0"
```

You'll see these sensors in the mmperfmon config show output

43 © 2018 IBM Corporation

Added Content – Installing GPFS (new cluster)

The GUI and performance collectors should be running on their respective nodes

```
root@client-node1:~# systemctl status gpfsGUI
gpfsGUI.service - IBM_Spectrum_Scale Administration GUI
   Loaded: loaded (/usr/lib/systemd/system/gpfsGUI.service; enabled; vendor preset: enabled)
   Active: active (running) since Wed 2018-05-09 15:04:38 MST; 22h ago
   Main PID: 27120 (java)
   Status: "GSS/GPFS GUI started"
   CPU: 5min 51.030s
   CGroup: /system.slice/gpfsGUI.service
           └─27120 /usr/lpp/mmfs/java/jre/bin/java -XX:+HeapDumpOnOutOfMemoryError -Dcom.ibm.gpfs.platform=GPFS -Dcom.ibm.gpfs.vendor
```

```
root@client-node1:~# systemctl status pmcollector
pmcollector.service - zimon collector daemon
   Loaded: loaded (/lib/systemd/system/pmcollector.service; enabled;
   vendor preset: enabled)
   Active: active (running) since Wed 2018-05-09 15:02:56 MST; 22h ago
   Main PID: 18729 (pmcollector)
   CGroup: /system.slice/pmcollector.service
           └─18729 /opt/IBM/zimon/sbin/pmcollector -C
/opt/IBM/zimon/ZIMonCollector.cfg -R /var/run/perfmon
```

```
[root@nsd-node1 ~]# systemctl status pmcollector
pmcollector.service - zimon collector daemon
   Loaded: loaded (/usr/lib/systemd/system/pmcollector.service; enabled;
   vendor preset: disabled)
   Active: active (running) since Wed 2018-05-09 15:02:57 MST; 22h ago
   Main PID: 14591 (pmcollector)
   CGroup: /system.slice/pmcollector.service
           └─14591 /opt/IBM/zimon/sbin/pmcollector -C
/opt/IBM/zimon/ZIMonCollector.cfg -R /var/run/perfmon
```

Client-node1 runs a collector because:

- It's a GUI node

NSD-node1 runs a collector because:

- When perfmon config is set to on in the Toolkit (default), it will automatically create 2 collectors and configure them in federated mode for redundancy
- GUI nodes are chosen first, NSD nodes are second priority for a collector, client nodes are chosen last
- If in ESS mode, the Toolkit will always assume the EMS node is a collector and will not add more collectors

Additional Content:

Deploying Protocols

Added Content – deploying protocols

Check the Install Toolkit to see if any protocol nodes or CES-IPs are listed

```
# ./spectrumscale node list
[ INFO ] List of nodes in current configuration:
[ INFO ] [Installer Node]
[ INFO ] 9.11.102.38
[ INFO ]
[ INFO ] Setup Type: SpectrumScale
[ INFO ]
[ INFO ] [Cluster Name]
[ INFO ] democlustertuc.stglabs.ibm.com
[ INFO ]
[ INFO ] File Audit logging : Disabled
[ INFO ]
[ INFO ] GPFS Node
[ INFO ] client-node1.tuc.stglabs.ibm.com
[ INFO ] client-node2.tuc.stglabs.ibm.com
[ INFO ] nsd-node1.tuc.stglabs.ibm.com
[ INFO ] nsd-node2.tuc.stglabs.ibm.com
[ INFO ] protocol-node1.tuc.stglabs.ibm.com
[ INFO ] protocol-node2.tuc.stglabs.ibm.com
[ INFO ]
[ INFO ] [Export IP address]
[ INFO ] No export IP addresses configured
```

	Admin	Quorum	Manager	NSD Server	Protocol	GUI Server	Callhome Node	OS	Arch
client-node1.tuc.stglabs.ibm.com	X					X		ubuntu16	x86_64
client-node2.tuc.stglabs.ibm.com							X	sles12	x86_64
nsd-node1.tuc.stglabs.ibm.com		X	X	X				rhel7	x86_64
nsd-node2.tuc.stglabs.ibm.com		X	X	X				rhel7	x86_64
protocol-node1.tuc.stglabs.ibm.com	X	X						rhel7	x86_64
protocol-node2.tuc.stglabs.ibm.com								rhel7	x86_64

In this case, there are none

Added Content – deploying protocols

CES-IPs can be contained with a single hostname set for DNS-RR

```
# nslookup ces-exporttips
Server: 9.11.227.25
Address: 9.11.227.25#53

Name: ces-exporttips.tuc.stglabs.ibm.com
Address: 9.11.102.107
Name: ces-exporttips.tuc.stglabs.ibm.com
Address: 9.11.102.108
Name: ces-exporttips.tuc.stglabs.ibm.com
Address: 9.11.102.138
Name: ces-exporttips.tuc.stglabs.ibm.com
Address: 9.11.102.80
```

Here's how to set CES-IPs with the Install Toolkit

```
# ./spectrumscale config protocols -e 9.11.102.107,9.11.102.108,9.11.102.138,9.11.102.80
[ INFO ] Setting Export IP Pool to 9.11.102.107,9.11.102.108,9.11.102.138,9.11.102.80
[ INFO ] Tip :Enable NFS, Object or SMB protocols as appropriate:./spectrumscale enable nfs|object|smb
```

****Note that the Install Toolkit does not currently support CES-groups****

- If all protocol nodes cannot see the same CES-IP subnet, you'll need to use CES groups

Added Content – deploying protocols

This is what CES-IPs will look like AFTER you finish the deploy

- Spectrum Scale code will take the CES IPs you provide and alias them to an appropriate adapter
- The adapter for aliasing of CES IPs is chosen based upon its route. Make sure you have a preset route for this to work.

```
[root@protocol-node1 installer]# ip addr
1: lo: <LOOPBACK,UP,LOWER_UP> mtu 65536 qdisc noqueue state UNKNOWN
    link/loopback 00:00:00:00:00:00 brd 00:00:00:00:00:00
    inet 127.0.0.1/8 scope host lo
        valid_lft forever preferred_lft forever
    inet6 ::1/128 scope host
        valid_lft forever preferred_lft forever
2: eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qdisc mq state UP qlen 1000
    link/ether 00:50:56:83:72:a2 brd ff:ff:ff:ff:ff:ff
    inet 9.11.102.38/24 brd 9.11.102.255 scope global eth0
        valid_lft forever preferred_lft forever
    inet 9.11.102.107/24 brd 9.11.102.255 scope global secondary eth0:0
        valid_lft forever preferred_lft forever
    inet 9.11.102.108/24 brd 9.11.102.255 scope global secondary eth0:1
        valid_lft forever preferred_lft forever
    inet6 2002:90b:e006:102:250:56ff:fe83:72a2/64 scope global noprefixroute dynamic
        valid_lft 2591994sec preferred_lft 604794sec
    inet6 fe80::250:56ff:fe83:72a2/64 scope link
        valid_lft forever preferred_lft forever
3: virbr0: <NO-CARRIER,BROADCAST,MULTICAST,UP> mtu 1500 qdisc noqueue state DOWN
    link/ether 52:54:00:40:b8:7d brd ff:ff:ff:ff:ff:ff
    inet 192.168.122.1/24 brd 192.168.122.255 scope global virbr0
        valid_lft forever preferred_lft forever
4: virbr0-nic: <BROADCAST,MULTICAST> mtu 1500 qdisc pfifo_fast master virbr0 state DOWN qlen 500
    link/ether 52:54:00:40:b8:7d brd ff:ff:ff:ff:ff:ff
```

```
[root@protocol-node1 installer]# route
```

```
Kernel IP routing table
```

Destination	Gateway	Genmask	Flags	Metric	Ref	Use	Iface
default	ip9-11-102-1.tu	0.0.0.0	UG	100	0	0	eth0
9.11.102.0	0.0.0.0	255.255.255.0	U	100	0	0	eth0
192.168.122.0	0.0.0.0	255.255.255.0	U	0	0	0	virbr0

Added Content – deploying protocols

Before pointing to file systems, see if the Install Toolkit already knows about any

```
[root@protocol-node1 installer]# ./spectrumscale filesystem list
```

[INFO]	Name	BlockSize	Mountpoint	NSDs Assigned	Default Data Replicas	Max Data Replicas	Default Metadata Replicas	Max Metadata Replicas
[INFO]	cesSharedRoot	Default	(4M)/ibm/cesSharedRoot	2	1	2	2	2
[INFO]	ObjectFS	Default	(4M)/ibm/ObjectFS	2	1	2	2	2
[INFO]	fs1	Default	(4M)/ibm/fs1	4	1	2	2	2

In this case, the Install Toolkit knows about 3 file systems.

This is because we added them with the 'spectrumscale nsd add' commands

```
# ./spectrumscale nsd list
```

[INFO]	Name	FS	Size (GB)	Usage	FG	Pool	Device	Servers
[INFO]	nsd1	cesSharedRoot	10	dataAndMetadata	1	Default	/dev/sdb	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com
[INFO]	nsd3	ObjectFS	250	dataAndMetadata	1	Default	/dev/sdd	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com
[INFO]	nsd5	fs1	750	dataAndMetadata	1	Default	/dev/sdf	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com
[INFO]	nsd6	fs1	750	dataAndMetadata	1	Default	/dev/sdg	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com
[INFO]	nsd2	cesSharedRoot	10	dataAndMetadata	2	Default	/dev/sdc	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com
[INFO]	nsd4	ObjectFS	250	dataAndMetadata	2	Default	/dev/sde	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com
[INFO]	nsd7	fs1	750	dataAndMetadata	2	Default	/dev/sdh	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com
[INFO]	nsd8	fs1	750	dataAndMetadata	2	Default	/dev/sdi	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com

If you have no file systems listed:

- you'll have to create them manually and allow the Install Toolkit to remain ignorant of them
- Or you'll have to add them via `spectrumscale nsd add / spectrumscale filesystem modify` and then follow this up with a `spectrumscale install` before continuing
- Or you can run `./spectrumscale config populate -N <node in cluster>` to traverse the cluster and pull the existing file system configuration into the Install Toolkit

Added Content – deploying protocols

Assuming you now have file systems, let's point the Install Toolkit to them

```
# ./spectrumscale config protocols -f cesSharedRoot -m /ibm/cesSharedRoot
[ INFO ] Setting Shared File System Name to cesSharedRoot
[ INFO ] Setting Shared File System Mountpoint to /ibm/cesSharedRoot
[ INFO ] Tip :Enable NFS, Object or SMB protocols as appropriate:./spectrumscale enable nfs|object|smb
```

Here we configure a shared root file system specifically for the protocols to use.

- You can point this command to an existing FS
- You can also point it to a FS that the toolkit has yet to create. Remember, they won't be created until deploy commences
- Shared Root is where all protocols will store state information so it's important to make this highly redundant.

Added Content - deploying protocols

Now we can enable the file protocols: nfs and smb

```
# ./spectrumscale enable nfs
[ INFO ] Enabling NFS on all protocol nodes.

# ./spectrumscale enable smb
[ INFO ] Enabling SMB on all protocol nodes.
```

This is how to enable Object

```
# ./spectrumscale enable object
[ INFO ] Enabling OBJECT on all protocol nodes.
# ./spectrumscale config object -e ces-exporttips
[ INFO ] Setting Endpoint Hostname to ces-exporttips
# ./spectrumscale config object -o Object_Fileset
[ INFO ] Setting GPFS Object Base to Object_Fileset
# ./spectrumscale config object -f ObjectFS -m /ibm/ObjectFS
[ INFO ] Setting Object File System Name to ObjectFS
[ INFO ] Setting Object File System Mountpoint to /ibm/ObjectFS
# ./spectrumscale config object -au admin -ap -dp
[ INFO ] Setting Admin User to admin
```

At this point you will be asked to type the admin and database passwords. They will be encrypted with the password you type as a secret key

Added Content – deploying protocols

Verify that protocol nodes, protocols, and CES-IPs have been added

- Remember, all of this information is stored within the clusterdefinition.txt file

```
# ./spectrumscale node list
[ INFO ] List of nodes in current configuration:
[ INFO ] [Installer Node]
[ INFO ] 9.11.102.38
[ INFO ]
[ INFO ] Setup Type: SpectrumScale
[ INFO ]
[ INFO ] [Cluster Name]
[ INFO ] democluster.tuc.stglabs.ibm.com
[ INFO ]
[ INFO ] [Protocols]
[ INFO ] Object : Enabled
[ INFO ] SMB : Enabled
[ INFO ] NFS : Enabled
[ INFO ]
[ INFO ] File Audit logging : Disabled
[ INFO ]
[ INFO ] GPFS Node
[ INFO ] client-node1.tuc.stglabs.ibm.com      Admin  Quorum  Manager  NSD Server  Protocol  GUI Server  Callhome Node  OS      Arch
[ INFO ] client-node2.tuc.stglabs.ibm.com      X
[ INFO ] client-node2.tuc.stglabs.ibm.com      X
[ INFO ] nsd-node1.tuc.stglabs.ibm.com          X      X      X
[ INFO ] nsd-node2.tuc.stglabs.ibm.com          X      X      X
[ INFO ] protocol-node1.tuc.stglabs.ibm.com      X      X      X
[ INFO ] protocol-node2.tuc.stglabs.ibm.com      X      X
[ INFO ]
[ INFO ] [Export IP address]
[ INFO ] 9.11.102.80 (pool)
[ INFO ] 9.11.102.107 (pool)
[ INFO ] 9.11.102.108 (pool)
[ INFO ] 9.11.102.138 (pool)
```

Added Content – deploying protocols

Check the Toolkit config for GPFS, Object, Protocols, and Perfmon

```
# ./spectrumscale config gpfs
[ INFO ] No changes made to defaults. Current settings are as follows:
[ INFO ] GPFS cluster name is democloud.tuc.stglabs.ibm.com.
[ INFO ] GPFS profile is default.
[ INFO ] Remote shell command is /usr/bin/ssh.
[ INFO ] Remote file copy command is /usr/bin/scp.
[ WARN ] No value for GPFS Daemon communication port range in clusterdefinition file.

# ./spectrumscale config object
[ INFO ] No changes made to defaults. Current settings are as follows:
[ INFO ] Object File System Name is ObjectFS.
[ INFO ] Object File System Mountpoint is /ibm/ObjectFS.
[ INFO ] Endpoint Hostname is ces-exporttips.
[ INFO ] GPFS Object Base is Object_Fileset.
[ WARN ] No value for GPFS Fileset inode allocation in clusterdefinition file.
[ WARN ] No value for Admin Token in clusterdefinition file.
[ INFO ] Admin User is admin.
[ INFO ] Admin Password is stored in the clusterdefinition file
[ WARN ] No value for Swift User in clusterdefinition file.
[ WARN ] No value for Swift Password in clusterdefinition file.
[ INFO ] Database Password is stored in the clusterdefinition file
[ WARN ] No value for Multi-region Data File Path in clusterdefinition file.
[ WARN ] No value for Region Number for Multi-region in clusterdefinition file.
[ INFO ] Tip :If all node designations and any required protocol configurations are complete, proceed to check the installation configuration:
[ INFO ] ./spectrumscale deploy --precheck

# ./spectrumscale config protocols
[ INFO ] No changes made to defaults. Current settings are as follows:
[ INFO ] Shared File System Name is cesSharedRoot.
[ INFO ] Shared File System Mountpoint is /ibm/cesSharedRoot.
[ INFO ] Export IP Pool is 9.11.102.80,9.11.102.107,9.11.102.108,9.11.102.138.
[ INFO ] Tip :Enable NFS, Object or SMB protocols as appropriate:./spectrumscale enable nfs|object|smb

# ./spectrumscale config perfmon
[ INFO ] No changes made to defaults. Current settings are as follows:
[ INFO ] Performance Monitoring reconfiguration is on (default). Collectors may be moved to different nodes, sensors may be added to nodes, and sensors may be reset to defaults.
```

Added Content – deploying protocols

Let's start

```
# ./spectrumscale deploy --precheck
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/DEPLOY-PRECHECK-10-05-2018_13:40:16.log
Enter the secret encryption key:
Repeat the secret encryption key:
[ WARN ] File Authentication will not be configured on your nodes. Run ./spectrumscale auth --help for information on auth.
[ WARN ] Object Authentication will be configured with a self-signed certificate. It is recommended to use stronger authentication schemes for a production environment. Run ./spectrumscale auth --help for information on auth.
[ WARN ] Only one GUI server specified. The Graphical User Interface will not be highly available.
[ WARN ] Install toolkit will not configure file audit logging as it has been disabled.
[ INFO ] Checking for knife bootstrap configuration...
[ INFO ] Performing Chef (deploy tool) checks.
[ INFO ] Performing Filesystem checks.
[ INFO ] NSDs are in a valid state
[ INFO ] Performing Cluster Export Services checks.
[ INFO ] Running environment checks for protocols
[ INFO ] Checking state of GPFS on all nodes
[ INFO ] GPFS active on all nodes
[ INFO ] Checking state of GPFS on all nodes
[ INFO ] GPFS active on all nodes
[ INFO ] protocol precheck OK
[ INFO ] Performing Object Store checks.
[ INFO ] Running environment checks for Object Storage
[ INFO ] Object Storage ready for install
[ INFO ] Performing SMB checks.
[ INFO ] Running environment checks for SMB
[ INFO ] SMB precheck OK
[ INFO ] Performing NFS checks.
[ INFO ] Running environment checks for NFS
[ INFO ] NFS precheck OK
[ INFO ] Performing Performance Monitoring checks.
[ INFO ] Running environment checks for Performance Monitoring
[ INFO ] Performing GUI checks.
[ INFO ] Performing callhome checks.
[ INFO ] Checking input for callhome
[ INFO ] Checking network connectivity for callhome configuration.
[ INFO ] callhome precheck OK
[ INFO ] Network check from admin node protocol-node1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from admin node client-node1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node2.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node1.tuc.stglabs.ibm.com to all other protocol nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node2.tuc.stglabs.ibm.com to all other protocol nodes in the cluster passed
[ WARN ] Ephemeral port range is not set. Please set valid ephemeral port range using the command ./spectrumscale config gpfs --ephemeral_port_range . You may set the default values as 60000-61000
[ INFO ] The install toolkit will reconfigure the call home group [autoGroup_1] with current configuration.
[ INFO ] Pre-check successful for deploy.
[ INFO ] Tip : ./spectrumscale deploy
```

Added Content – deploying protocols

Pre-check was successful. Now we can start the deploy

```
# ./spectrumscale deploy  
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/DEPLOY-10-05-2018_13:59:22.log  
[ INFO ] Validating configuration
```

Added Content – deploying protocols

Here's an example of a successful deployment

```
[ INFO ] Checking for a successful install
[ INFO ] Checking state of Chef (deploy tool)
[ INFO ] Chef (deploy tool) ACTIVE
[ INFO ] Checking state of Filesystem
[ INFO ] File systems have been created successfully
[ INFO ] Filesystem ACTIVE
[ INFO ] Checking state of Cluster Export Services
[ INFO ] Checking state of CES on all nodes
[ INFO ] CES healthy on all nodes
[ INFO ] 4 server licenses and 2 client licenses enabled on the cluster. Ensure enough licenses are owned for this configuration.
[ INFO ] Cluster Export Services ACTIVE
[ INFO ] Checking state of Object Store
[ INFO ] Running Object post-install checks
[ INFO ] Checking state of OBJ on all nodes
[ INFO ] OBJ healthy on all nodes
[ INFO ] Object Store ACTIVE
[ INFO ] Checking state of SMB
[ INFO ] Running SMB post-install checks
[ INFO ] Checking state of SMB on all nodes
[ INFO ] SMB healthy on all nodes
[ INFO ] SMB ACTIVE
[ INFO ] Checking state of NFS
[ INFO ] Running NFS post-install checks
[ INFO ] Checking state of NFS on all nodes
[ INFO ] NFS healthy on all nodes
[ INFO ] NFS ACTIVE
[ INFO ] Checking state of Performance Monitoring
[ INFO ] Running Performance Monitoring post-install checks
[ INFO ] pmcollector running on all nodes
[ INFO ] pmsensors running on all nodes
[ INFO ] Performance Monitoring ACTIVE
[ INFO ] Checking state of GUI
[ INFO ] Running Graphical User Interface post-install checks
[ INFO ] Graphical User Interface running on all GUI servers
[ INFO ] Enter one of the following addresses into a web browser to access the Graphical User Interface: client-node1.tuc.stglabs.ibm.com
[ INFO ] GUI ACTIVE
[ INFO ] Checking state of callhome
[ INFO ] callhome ACTIVE
[ INFO ] SUCCESS
[ INFO ] All services running
[ INFO ] StanzaFile and NodeDesc file for NSD, filesystem, and cluster setup have been saved to /usr/lpp/mmfs folder on node: protocol-node1.tuc.stglabs.ibm.com
[ INFO ] Successfully installed and configured protocols. 2 protocol nodes were enabled. Components installed: Chef (deploy tool), Filesystem, Cluster Export Services, Object Store, SMB, NFS, Performance Monitoring, GUI, callhome.
It took 19 minutes 21 seconds.
```

Added Content – deploying protocols

Verify all nodes are active

```
# mmgetstate -a -L
```

Node number	Node name	Quorum	Nodes up	Total nodes	GPFS state	Remarks
1	protocol-node1	2	3	6	active	quorum node
2	nsd-node1	2	3	6	active	quorum node
3	nsd-node2	2	3	6	active	quorum node
4	protocol-node2	2	3	6	active	
6	client-node2	2	3	6	active	
7	client-node1	2	3	6	active	

Verify all protocol services are running

```
# mmces service list -a
```

```
Enabled services: OBJ SMB NFS
```

```
protocol-node2.tuc.stglabs.ibm.com: OBJ is running, SMB is running, NFS is running
```

```
protocol-node1.tuc.stglabs.ibm.com: OBJ is running, SMB is running, NFS is running
```

Added Content – deploying protocols

Verify CES-IPs are assigned

```
# mmlscluster --ces
```

```
GPFS cluster information
```

```
=====
```

```
GPFS cluster name:      democloud.tuc.stglabs.ibm.com
GPFS cluster id:        2254960516004008510
```

```
Cluster Export Services global parameters
```

```
-----
```

```
Shared root directory:  /ibm/cesSharedRoot/ces
Enabled Services:       OBJ SMB NFS
Log level:              0
Address distribution policy: even-coverage
```

Node	Daemon node name	IP address	CES IP address list
1	protocol-node1.tuc.stglabs.ibm.com	9.11.102.38	9.11.102.107 9.11.102.108
4	protocol-node2.tuc.stglabs.ibm.com	9.11.102.55	9.11.102.138 9.11.102.80

Added Content – deploying protocols

Check the overall cluster health

```
# mmhealth cluster show
```

Component	Total	Failed	Degraded	Healthy	Other
-----	-----	-----	-----	-----	-----
NODE	6	0	0	0	6
GPFS	6	0	0	0	6
NETWORK	6	0	0	6	0
FILESYSTEM	3	0	0	3	0
DISK	8	0	0	8	0
CES	2	0	0	0	2
CESIP	1	0	0	1	0
PERFMON	6	0	0	6	0
THRESHOLD	6	0	0	6	0

```
# mmhealth cluster show CES
```

Component	Node	Status	Reasons
-----	-----	-----	-----
CES	protocol-node1.tuc.stglabs.ibm.com	TIPS	nfs_sensors_not_configured, smb_sensors_not_configured, smb_sensors_not_configured
CES	protocol-node2.tuc.stglabs.ibm.com	TIPS	nfs_sensors_not_configured, smb_sensors_not_configured, smb_sensors_not_configured
[root@protocol-node1 installer]# mmhealth cluster show NODE			
Component	Node	Status	Reasons
-----	-----	-----	-----
NODE	protocol-node1.tuc.stglabs.ibm.com	TIPS	gpfs_maxfilestocache_small, gpfs_maxstatcache_high, gpfs_pagepool_small
NODE	nsd-node2.tuc.stglabs.ibm.com	TIPS	gpfs_maxstatcache_high, gpfs_pagepool_small
NODE	nsd-node1.tuc.stglabs.ibm.com	TIPS	gpfs_maxstatcache_high, gpfs_pagepool_small
NODE	protocol-node2.tuc.stglabs.ibm.com	TIPS	gpfs_maxfilestocache_small, gpfs_maxstatcache_high, gpfs_pagepool_small
NODE	client-node1.tuc.stglabs.ibm.com	TIPS	gpfs_maxstatcache_high, gpfs_pagepool_small
NODE	client-node2.tuc.stglabs.ibm.com	TIPS	gpfs_maxstatcache_high, gpfs_pagepool_small

Added Content – deploying protocols

Let's check performance monitoring

```
# mmperfmon config show
# This file has been generated automatically and SHOULD NOT
# be edited manually. It may be overwritten at any point
# in time.

cephMon = "/opt/IBM/zimon/CephMonProxy"
cephRados = "/opt/IBM/zimon/CephRadosProxy"
colCandidates = "client-nodel.tuc.stglabs.ibm.com", "nsd-nodel.tuc.stglabs.ibm.com"
colRedundancy = 2
collectors = {
host = ""
port = "4739"
}
config = "/opt/IBM/zimon/ZIMonSensors.cfg"
ctdbstat = ""
daemonize = T
hostname = ""
ipfixinterface = "0.0.0.0"
logfile = "/var/log/zimon/ZIMonSensors.log"
loglevel = "info"
mmcmd = "/opt/IBM/zimon/MMCmdProxy"
mmdfcmd = "/opt/IBM/zimon/MMDFProxy"
mmpmon = "/opt/IBM/zimon/MmpmonSockProxy"
piddir = "/var/run"
release = "5.0.1-0"
```

Added Content – deploying protocols

These protocol sensors have been added

NFS

```
{
  name = "NFSIO"
  period = 1
  restrict = "cesNodes"
  type = "Generic"
},
```

SMB

```
{
  name = "SMBStats"
  period = 1
  restrict = "cesNodes"
  type = "Generic"
},
{
  name = "SMBGlobalStats"
  period = 1
  restrict = "cesNodes"
  type = "Generic"
},
{
  name = "CTDBStats"
  period = 1
  restrict = "cesNodes"
  type = "Generic"
},
{
  name = "CTDBDBStats"
  period = 1
  restrict = "cesNodes"
  type = "Generic"
},
```

Object

```
{
  name = "SwiftAccount"
  period = 1
  restrict = "cesNodes"
  type = "generic"
},
{
  name = "SwiftContainer"
  period = 1
  restrict = "cesNodes"
  type = "generic"
},
{
  name = "SwiftObject"
  period = 1
  restrict = "cesNodes"
  type = "generic"
},
{
  name = "SwiftProxy"
  period = 1
  restrict = "cesNodes"
  type = "generic"
}
```

Additional Content:

Setting up Authentication

Added Content – setting up File authentication (without toolkit)

We highly recommend NOT using the Install Toolkit for setting up Authentication

- Why? It's easier to run a single mmuserauth command and get instant feedback regarding success or failure

Pre-reqs

- All protocol nodes must be able to resolve and ping the Authentication server
- openldap-clients, sssd, ypbind, krb5-workstation must be installed on all protocol nodes

Double check that File Authentication is not yet configured

```
# mmuserauth service list
FILE access not configured
PARAMETERS          VALUES
-----

OBJECT access configuration : LOCAL
PARAMETERS          VALUES
-----

ENABLE_KS_SSL        false
ENABLE_KS_CASIGNING   false
KS_ADMIN_USER         admin
```

Added Content – setting up authentication (without toolkit)

Here's an example of running mmuserauth to configure AD + SUA

```
# /usr/lpp/mmfs/bin/mmuserauth service create --data-access-method file --type ad --servers '9.11.139.2' --netbios-name 'democluster' --idmap-role 'master' --user-name 'Administrator' --unixmap-domains 'storage4test(10000000-299999999)' --idmap-range '10000-1000000' --idmap-range-size '10000'
```

Enter Active Directory User 'Administrator' password:
File authentication configuration completed successfully.

Verify authentication is now configured

```
# mmuserauth service list
```

FILE access configuration : AD	
PARAMETERS	VALUES

ENABLE_NFS_KERBEROS	false
SERVERS	9.11.139.2
USER_NAME	Administrator
NETBIOS_NAME	democluster
IDMAP_ROLE	master
IDMAP_RANGE	10000-1000000
IDMAP_RANGE_SIZE	10000
UNIXMAP_DOMAINS	none
LDAPMAP_DOMAINS	none

OBJECT access configuration : LOCAL	
PARAMETERS	VALUES

ENABLE_KS_SSL	false
ENABLE_KS_CASIGNING	false
KS_ADMIN_USER	admin

Added Content – setting up authentication (not recommended toolkit method)

If you skipped mmuserauth and want to use the Install Toolkit for setting up Authentication.....

Start with enabling AD for the file protocols to use

```
# ./spectrumscale auth file ad
[ INFO ] A configuration template has been created at configuration/authconfig.txt. Open this file in the text editor of your choice and
complete the template.
Would you like to open this file now? [Y/n]:
```

Added Content – setting up authentication (not recommended toolkit method)

```
# For detailed information on the authentication settings refer to the
# mmuserauth page in the GPFS commands section of the GPFS Knowledge
# Center.
```

```
[file_auth]
enable_file_auth = True
backend_server = ad
```

```
[file_ad]
# mandatory settings for file authentication:
```

```
# Specifies the host name or IP address of the authentication server.
servers = 9.11.139.2
# Specifies the name used to advertise the cluster in the Windows network
# neighborhood. It is limited to 15 ASCII characters and must not
# contain any white space or one of the following characters:
# / : * ? . - ; |
netbios_name = democluster
```

```
# Specifies the ID map role of the Spectrum Scale system. The ID map role
# of a stand-alone or singular system deployment must be "master".
idmap_role = master
```

```
# Specifies the user name to perform operations against the
# authentication server.
bind_username = Administrator
```

```
# Specifies the password of the user name that is specified with the
# bind_username option.
# Leave as [prompt] to be prompted for the password in
# a secure manner.
bind_password = [prompt]
```

```
# additional settings for file authentication:
```

```
# Specifies the AD domains for which user ID and group ID should be fetched
# from the AD server. Input value format is
# DOMAIN1[(L1+H1)][:Domain2[(L2+H2)][:Domain3[(L3+H3)]...]
# Multiple domains can be specified separate by semi-colon(:)
unixmap_domains =
```

```
# Specifies the AD domains for which user ID and group ID should be fetched
# from the LDAP server.
# The value can be split onto multiple lines providing that continuing lines
# begin with a space.
ldapmap_domains =
```

```
# Enables kerberized authentication for NFSv4 based access
enable_nfs_kerberos = False
```


```
# optional settings for file authentication: (initialized with default values)
# If required, un-comment and specify your environment configuration
```

```
# Specifies the range of values from which the Spectrum Scale UIDs and GIDs
# are assigned by the system to the AD users and groups.
idmap_range = 10000-1000000
```

```
# Specifies the total number of UIDs and GIDs that are assignable per domain.
idmap_range_size = 10000
```

A sample template will pop up

- Fill out the template for your specific settings
- Save the template by typing pressing ctrl-x
- Make sure to save and write the file
- Fill in all of the passwords, including your bind_password used by the AD user who has permission to setup this acct
- You can delete the template afterwards since it will be written to the clusterdefinition.txt file



Are you ready to save your changes now? [Y/n]: y
 [WARN] Object Authentication will be configured with a self-signed certificate. It is recommended to use stronger authentication schemes for a production environment. Run
 ./spectrumscale auth --help for information on auth.
 Enter the secret encryption key:
 Repeat the secret encryption key:
 Enter (file_ad) bind_password:
 Repeat (file_ad) bind_password:
 [INFO] Your authentication settings have been merged into the main cluster definition file (configuration/clusterdefinition.txt). The completed template can now be deleted.
 Would you like to delete the template? [Y/n]: y

Added Content – setting up authentication (not recommended toolkit method)

Let's try a deploy pre-check now to see if we're ready

./spectrumscale deploy --precheck

```
[INFO ] Logging to file: /usr/tp/mmfs/5.0.1.0/installer/logs/DEPLOY-PRECHECK-10-05-2018_14:30:58.log
Enter the secret encryption key:
Repeat the secret encryption key:
[WARN ] Object Authentication will be configured with a self-signed certificate. It is recommended to use stronger authentication schemes for a production environment. Run ./spectrumscale auth --help for information on auth.
[WARN ] Only one GUI server specified. The Graphical User Interface will not be highly available.
[WARN ] Install toolkit will not configure file audit logging as it has been disabled.
[INFO ] Checking for knife bootstrap configuration...
[INFO ] Performing Chef (deploy tool) checks.
[INFO ] Performing Filesystem checks.
[INFO ] NSDs are in a valid state
[INFO ] Performing Cluster Export Services checks.
[INFO ] Running environment checks for protocols
[INFO ] Checking state of GPFS on all nodes
[INFO ] GPFS active on all nodes
[INFO ] Checking state of GPFS on all nodes
[INFO ] GPFS active on all nodes
[INFO ] protocol precheck OK
[INFO ] Performing Object Store checks.
[INFO ] Running environment checks for Object Storage
[INFO ] Object Storage ready for install
[INFO ] Performing SMB checks.
[INFO ] Running environment checks for SMB
[INFO ] SMB precheck OK
[INFO ] Performing NFS checks.
[INFO ] Running environment checks for NFS
[INFO ] NFS precheck OK
[INFO ] Performing Authentication checks.
[INFO ] Running environment checks for Authentication setup
[INFO ] Checking promptless ssh for authentication server
[WARN ] Pre requisite packages not found on protocol-node2.tuc.stglabs.ibm.com (Pre requisites: sssd, ybind, openldap-clients, krb5-workstation). Ensure this is resolved before proceeding with the install toolkit. For more information about file authentication prerequisites for supported AUTH methods, see 'Configuring authentication and ID mapping for file access' subject in the Knowledge Center.
[WARN ] Pre requisite packages not found on protocol-node1.tuc.stglabs.ibm.com (Pre requisites: sssd, ybind, openldap-clients, krb5-workstation). Ensure this is resolved before proceeding with the install toolkit. For more information about file authentication prerequisites for supported AUTH methods, see 'Configuring authentication and ID mapping for file access' subject in the Knowledge Center.
[INFO ] Cluster ready for authentication setup
[INFO ] Performing Performance Monitoring checks.
[INFO ] Running environment checks for Performance Monitoring
[INFO ] Performing GUI checks.
[INFO ] Performing callhome checks.
[INFO ] Checking input for callhome
[INFO ] Checking network connectivity for callhome configuration.
[INFO ] callhome precheck OK
[INFO ] Network check from admin node protocol-node1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[INFO ] Network check from admin node client-node1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[INFO ] Network check from protocol node protocol-node2.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[INFO ] Network check from protocol node protocol-node1.tuc.stglabs.ibm.com to all other protocol nodes in the cluster passed
[INFO ] Network check from protocol node protocol-node2.tuc.stglabs.ibm.com to all other protocol nodes in the cluster passed
[WARN ] Ephemeral port range is not set. Please set valid ephemeral port range using the command ./spectrumscale config gpfs --ephemeral_port_range . You may set the default values as 60000-61000
[INFO ] The install toolkit will reconfigure the call home group [autoGroup_1] with current configuration.
[INFO ] Pre-check successful for deploy.
[INFO ] Tip : ./spectrumscale deploy
```

Notice the warnings? We forgot to install 4 critical pre-requisites needed by authentication on all protocol nodes.

Added Content – setting up authentication (not recommended toolkit method)

Install these packages on all protocol nodes

```
# yum install sssd ypbind openldap-clients krb5-workstation -y
```

Re-run the pre-check

```
# ./spectrumscale deploy --precheck
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/DEPLOY-10-05-2018_14:32:52.log
[ WARN ] Object Authentication will be configured with a self-signed certificate. It is recommended to use stronger authentication schemes for a production environment. Run ./spectrumscale auth --help for information on auth.
[ WARN ] Only one GUI server specified. The Graphical User Interface will not be highly available.
[ WARN ] Install toolkit will not configure file audit logging as it has been disabled.
[ INFO ] Checking for knife bootstrap configuration...
[ INFO ] Performing Chef (deploy tool) checks.
[ INFO ] Performing Filesystem checks.
[ INFO ] NSDs are in a valid state
[ INFO ] Performing Cluster Export Services checks.
[ INFO ] Running environment checks for protocols
[ INFO ] Checking state of GPFS on all nodes
[ INFO ] GPFS active on all nodes
[ INFO ] Checking state of GPFS on all nodes
[ INFO ] GPFS active on all nodes
[ INFO ] protocol precheck OK
[ INFO ] Performing Object Store checks.
[ INFO ] Running environment checks for Object Storage
[ INFO ] Object Storage ready for install
[ INFO ] Performing SMB checks.
[ INFO ] Running environment checks for SMB
[ INFO ] SMB precheck OK
[ INFO ] Performing NFS checks.
[ INFO ] Running environment checks for NFS
[ INFO ] NFS precheck OK
[ INFO ] Performing Authentication checks.
[ INFO ] Running environment checks for Authentication setup
[ INFO ] Checking promptless ssh for authentication server
[ INFO ] Cluster ready for authentication setup
[ INFO ] Performing Performance Monitoring checks.
[ INFO ] Running environment checks for Performance Monitoring
[ INFO ] Performing GUI checks.
[ INFO ] Performing callhome checks.
[ INFO ] Checking input for callhome
[ INFO ] Checking network connectivity for callhome configuration.
[ INFO ] callhome precheck OK
[ INFO ] Network check from admin node protocol-node1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from admin node client-node1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node2.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node1.tuc.stglabs.ibm.com to all other protocol nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node2.tuc.stglabs.ibm.com to all other protocol nodes in the cluster passed
[ WARN ] Ephemeral port range is not set. Please set valid ephemeral port range using the command ./spectrumscale config gpfs --ephemeral_port_range. You may set the default values as 60000-61000
[ INFO ] The install toolkit will reconfigure the call home group [autoGroup_1] with current configuration.
[ INFO ] Pre-check successful for deploy.
[ INFO ] Tip : ./spectrumscale deploy
```

You're now ready for a deploy

```
# ./spectrumscale deploy
```

```
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/DEPLOY-10-05-2018_14:35:52.log
```

```
[ INFO ] Validating configuration
```

Added Content – setting up authentication (not recommended toolkit method)

Example of a successful deploy with authentication post checks

```
[INFO] Checking for a successful install
[INFO] Checking state of Chef (deploy tool)
[INFO] Chef (deploy tool) ACTIVE
[INFO] Checking state of Filesystem
[INFO] File systems have been created successfully
[INFO] Filesystem ACTIVE
[INFO] Checking state of Cluster Export Services
[INFO] Checking state of CES on all nodes
[INFO] CES healthy on all nodes
[INFO] 4 server licenses and 2 client licenses enabled on the cluster. Ensure enough licenses are owned for this configuration.
[INFO] Cluster Export Services ACTIVE
[INFO] Checking state of Object Store
[INFO] Running Object post-install checks
[INFO] Checking state of OBJ on all nodes
[INFO] OBJ healthy on all nodes
[INFO] Object Store ACTIVE
[INFO] Checking state of SMB
[INFO] Running SMB post-install checks
[INFO] Checking state of SMB on all nodes
[INFO] SMB healthy on all nodes
[INFO] SMB ACTIVE
[INFO] Checking state of NFS
[INFO] Running NFS post-install checks
[INFO] Checking state of NFS on all nodes
[INFO] NFS healthy on all nodes
[INFO] NFS ACTIVE
[INFO] Checking state of Authentication
[INFO] Running Authentication post-install checks
[INFO] Checking state of AUTH on all nodes
[INFO] AUTH healthy on all nodes
[INFO] Authentication ACTIVE
[INFO] Checking state of Performance Monitoring
[INFO] Running Performance Monitoring post-install checks
[INFO] pmcollector running on all nodes
[INFO] pmsensors running on all nodes
[INFO] Performance Monitoring ACTIVE
[INFO] Checking state of GUI
[INFO] Running Graphical User Interface post-install checks
[INFO] Graphical User Interface running on all GUI servers
[INFO] Enter one of the following addresses into a web browser to access the Graphical User Interface: client-node1.tuc.stglabs.ibm.com
[INFO] GUI ACTIVE
[INFO] Checking state of callhome
[INFO] callhome ACTIVE
[INFO] SUCCESS
[INFO] All services running
[INFO] StanzaFile and NodeDesc file for NSD, filesystem, and cluster setup have been saved to /usr/lpp/mmfs folder on node: protocol-node1.tuc.stglabs.ibm.com
[INFO] Successfully installed and configured protocols. 2 protocol nodes were enabled. Components installed: Chef (deploy tool), Filesystem, Cluster Export Services, Object Store, SMB, NFS, Authentication, Performance Monitoring, GUI, callhome. It took 6 minutes 2 seconds.
```

Added Content – setting up authentication (not recommended toolkit method)

Double check to make sure the authentication services are healthy

mmuserauth service list

FILE access configuration : AD

PARAMETERS	VALUES
------------	--------

ENABLE_NFS_KERBEROS	false
SERVERS	9.11.139.2
USER_NAME	Administrator
NETBIOS_NAME	democluster
IDMAP_ROLE	master
IDMAP_RANGE	10000-1000000
IDMAP_RANGE_SIZE	10000
UNIXMAP_DOMAINS	none
LDAPMAP_DOMAINS	none

OBJECT access configuration : LOCAL

PARAMETERS	VALUES
------------	--------

ENABLE_KS_SSL	false
ENABLE_KS_CASIGNING	false
KS_ADMIN_USER	admin

mmhealth cluster show AUTH

Component	Node	Status	Reasons
AUTH	protocol-node1.tuc.stglabs.ibm.com	HEALTHY	-
AUTH	protocol-node2.tuc.stglabs.ibm.com	HEALTHY	-

Additional Content:

**Adding nodes, nsds, and file
systems to an existing cluster**

Added Content – adding to an existing cluster

We're going to start with a 6 node cluster and add:

- 1 GUI node, 1 protocol node, and 1 NSD node
- 2 non-shared NSDs belonging to the new NSD node

This procedure applies to:

- Clusters that have been installed without the toolkit
- Clusters that have been installed with the toolkit

Pre-reqs:

- OS pre-installed on the new nodes that will be added to the cluster
- Promptless ssh to/from new nodes to all other nodes including themselves
- Firewalls / SELinux settings configured to all for Toolkit and Spectrum Scale
- Download the Spectrum Scale protocol package
- Extract the protocol package on either an existing node or a new node

Added Content – adding to an existing cluster

First, setup the Install Toolkit

```
# ./spectrumscale setup -s 9.11.102.38 -st ss
```

```
[ INFO ] Installing prerequisites for install node
```

```
[ INFO ] Existing Chef installation detected. Ensure chef-zero is included and the PATH is configured.
```

```
[ INFO ] Your control node has been configured to use the IP 9.11.102.38 to communicate with other nodes.
```

```
[ INFO ] Port 8889 will be used for chef communication.
```

```
[ INFO ] Port 10080 will be used for package distribution.
```

```
[ INFO ] Install Toolkit setup type is set to Spectrum Scale (default). If an ESS is in the cluster, run this command to set ESS mode: ./spectrumscale setup -s server_ip -st ess
```

```
[ INFO ] SUCCESS
```

```
[ INFO ] Tip : Designate protocol, nsd and admin nodes in your environment to use during install: ./spectrumscale -v node add <node> -p -a -n
```

Next, try config populate

- Config populate needs to be pointed at a node in the cluster
- When run, it will ssh to that node and run various mm commands in order to traverse the cluster and pull its configuration into the Install Toolkit.
- Config populate doesn't support all configurations, but it doesn't hurt to try it. If it works, you won't have to manually manually tell the Install Toolkit about the cluster

Added Content – adding to an existing cluster

Run config populate

Run config populate from the node you extracted code on

- Here, it is run on protocol-node1

Point config populate to any node in the cluster

- Here, it points to itself, protocol-node1, which is already in the cluster.
- If you have an ESS, point to the EMS node
- If you've extracted the code to a node not in the cluster, point it to a node in the cluster. The node not in the cluster can still be the installer node.
- The node you point to will become the admin node for the toolkit. It will need promptless ssh access to/from all nodes. The toolkit will run mm commands on this node.

What if config populate doesn't work with my config?

- You can still use the toolkit, you'll just need to input at least one node in the cluster via `./spectrumscale node add` and you'll need to make this node an admin node.

```
[root@protocol-node1 installer]# ./spectrumscale config populate -N protocol-node1
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/config-populate-10-05-2018_15:14:49.log
[ INFO ] Detected clusterdefinition file present in configuration /usr/lpp/mmfs/5.0.1.0/installer/configuration.
Installer will keep backup of existing clusterdefinition.txt file in /usr/lpp/mmfs/5.0.1.0/installer/configuration path and populate a new one. Do you want to
continue [Y/n]: y
[ INFO ] Updating existing configuration. It may take few moments ....
[ INFO ] Getting installer path
[ INFO ] Populating protocols detail
[ INFO ] Populating ntp details
[ INFO ] Populating enabled protocols detail
[ INFO ] Object is enabled. so populating object details.
[ INFO ] The passwords entered for object can be those that were initially used for setting up object or optionally enter new passwords to change to new
value. Enter the encryption passwords that will be used to encrypt the sensitive data in the cluster configuration file.
[ INFO ] Populating export ip of ces nodes
[ INFO ] Adding admin node into configuration
[ INFO ] Populating cluster details
[ INFO ] Total 6 node(s) found in the cluster
[ INFO ] Populating the details of protocol-node2.tuc.stglabs.ibm.com
[ INFO ] Populating the details of protocol-node1.tuc.stglabs.ibm.com
[ INFO ] Populating the details of client-node2.tuc.stglabs.ibm.com
[ INFO ] Populating the details of nsd-node1.tuc.stglabs.ibm.com
[ INFO ] Populating the details of client-node1.tuc.stglabs.ibm.com
[ INFO ] Populating the details of nsd-node2.tuc.stglabs.ibm.com
[ INFO ] Defining client-node1.tuc.stglabs.ibm.com as a GUI node
[ INFO ] Populating the GUI addresses details
[ INFO ] Populating GPFS details
[ INFO ] Populating the zimmon_collector details
[ INFO ] Populating the secondary_zimmon_collector details
[ INFO ] Authentication file _ad is set. so populating the details
[ INFO ] Populating NSD's details
[ INFO ] Total 8 NSDs found in the cluster
[ INFO ] Populating the details of nsd3
[ INFO ] Populating the details of nsd4
[ INFO ] Populating the details of nsd1
[ INFO ] Populating the details of nsd2
[ INFO ] Populating the details of nsd5
[ INFO ] Populating the details of nsd6
[ INFO ] Populating the details of nsd7
[ INFO ] Populating the details of nsd8
[ INFO ] Populating filesystem details
[ INFO ] Total 3 filesystem found in the cluster.
[ INFO ] Populating filesystem details for ObjectFS
[ INFO ] Populating filesystem details for cesSharedRoot
[ INFO ] Populating filesystem details for fs1
[ INFO ] Populating Callhome details
[ INFO ] Populating file audit logging details
[ INFO ] Configuration successfully updated.
```

Added Content – adding to an existing cluster

Check the Toolkit config after config populate finishes

```
# ./spectrumscale node list
[ INFO ] List of nodes in current configuration:
[ INFO ] [Installer Node]
[ INFO ] 9.11.102.38
[ INFO ]
[ INFO ] Setup Type: SpectrumScale
[ INFO ]
[ INFO ] [Cluster Name]
[ INFO ] democluster.tuc.stglabs.ibm.com
[ INFO ]
[ INFO ] [Protocols]
[ INFO ] Object : Enabled
[ INFO ] SMB : Enabled
[ INFO ] NFS : Enabled
[ INFO ]
[ INFO ] File Audit logging : Disabled
[ INFO ]
[ INFO ] GPFS Node
[ INFO ] client-node1.tuc.stglabs.ibm.com      Admin  Quorum  Manager  NSD Server  Protocol  GUI Server  Callhome Node  OS      Arch
[ INFO ] client-node2.tuc.stglabs.ibm.com      X
[ INFO ] client-node2.tuc.stglabs.ibm.com      X
[ INFO ] nsd-node1.tuc.stglabs.ibm.com          X      X      X
[ INFO ] nsd-node2.tuc.stglabs.ibm.com          X      X      X
[ INFO ] protocol-node1.tuc.stglabs.ibm.com     X      X
[ INFO ] protocol-node2.tuc.stglabs.ibm.com     X
[ INFO ]
[ INFO ] [Export IP address]
[ INFO ] 9.11.102.80 (pool)
[ INFO ] 9.11.102.107 (pool)
[ INFO ] 9.11.102.108 (pool)
[ INFO ] 9.11.102.138 (pool)
```


Added Content – adding to an existing cluster

Add 1 new protocol node

```
# ./spectrumscale node add protocol-node3 -p
[ INFO ] Adding node protocol-node3.tuc.stglabs.ibm.com as a GPFS node.
[ INFO ] Setting protocol-node3.tuc.stglabs.ibm.com as a protocol node.
[ INFO ] Configuration updated.
[ INFO ] Tip : If all node designations are complete, configure the protocol environment as needed:
./spectrumscale config protocols -f cesSharedRoot -m /ibm/cesSharedRoot
```

Add a 1 new GUI node

```
# ./spectrumscale node add client-node3 -g -a
[ INFO ] Adding node client-node3.tuc.stglabs.ibm.com as a GPFS node.
[ INFO ] Setting client-node3.tuc.stglabs.ibm.com as an admin node.
[ INFO ] Configuration updated.
[ INFO ] Tip : Designate protocol or nsd nodes in your environment to use during install:./spectrumscale
node add <node> -p -n
[ INFO ] Setting client-node3.tuc.stglabs.ibm.com as a GUI server.
```

Add a 1 new NSD node

```
# ./spectrumscale node add nsd-node3 -n
[ INFO ] Adding node nsd-node3.tuc.stglabs.ibm.com as a GPFS node.
[ INFO ] Adding node nsd-node3.tuc.stglabs.ibm.com as an NSD server.
[ INFO ] Configuration updated.
[ INFO ] Tip :If all node designations are complete, add NSDs to your cluster definition and
define required filesystems:./spectrumscale nsd add <device> -p <primary node> -s <secondary node> -fs <file system>
```

Added Content – adding to an existing cluster

Re-check the Toolkit config to see those 3 new nodes that we'll be adding

```
# ./spectrumscale node list
[ INFO ] List of nodes in current configuration:
[ INFO ] [Installer Node]
[ INFO ] 9.11.102.38
[ INFO ]
[ INFO ] Setup Type: SpectrumScale
[ INFO ]
[ INFO ] [Cluster Name]
[ INFO ] democluster.tuc.stglabs.ibm.com
[ INFO ]
[ INFO ] [Protocols]
[ INFO ] Object : Enabled
[ INFO ] SMB : Enabled
[ INFO ] NFS : Enabled
[ INFO ]
[ INFO ] File Audit logging : Disabled
[ INFO ]
[ INFO ] GPFS Node
[ INFO ] client-node1.tuc.stglabs.ibm.com      Admin Quorum Manager NSD Server Protocol GUI Server Callhome Node OS Arch
[ INFO ] client-node2.tuc.stglabs.ibm.com      X                                     X                                     ubuntu16 x86_64
[ INFO ] nsd-node1.tuc.stglabs.ibm.com           X      X      X                                     rhel7    x86_64
[ INFO ] nsd-node2.tuc.stglabs.ibm.com           X      X      X                                     rhel7    x86_64
[ INFO ] protocol-node1.tuc.stglabs.ibm.com      X      X                                     rhel7    x86_64
[ INFO ] protocol-node2.tuc.stglabs.ibm.com      X      X                                     rhel7    x86_64
[ INFO ] protocol-node3.tuc.stglabs.ibm.com      X      X                                     rhel7    x86_64
[ INFO ] client-node3.tuc.stglabs.ibm.com      X                                     X                                     rhel7    x86_64
[ INFO ] nsd-node3.tuc.stglabs.ibm.com           X                                     X                                     rhel7    x86_64
[ INFO ]
[ INFO ] [Export IP address]
[ INFO ] 9.11.102.80 (pool)
[ INFO ] 9.11.102.107 (pool)
[ INFO ] 9.11.102.108 (pool)
[ INFO ] 9.11.102.138 (pool)
```

Added Content – Installing GPFS (new cluster)

Since we already added NSD nodes, let's add NSDs

- In this example, nsd-node1 & nsd-node2 share the same physical disks, we'll alternate primary/secondary server to keep things balanced
- All NSDs in this example will be set for both dataAndMetadata
- Half of the NSDs assigned to each FS will be in failure group 1 and the other half, failure group 2
- We're designating 3 FSs for these NSDs: cesSharedRoot, ObjectFS, and fs1. The Install Toolkit will create the NSDs during install, if they don't already exist. And the FSs during deploy if they don't already exist

```
# ./spectrumscale nsd add -p nsd-node1 -s nsd-node2 -u dataAndMetadata -fs cesSharedRoot -fg 1 "/dev/sdb"
```

```
[ INFO ] Connecting to nsd-node1.tuc.stglabs.ibm.com to check devices and expand wildcards.
```

```
[ INFO ] Looking up details of /dev/sdb.
```

```
[ INFO ] The installer will create the new file system cesSharedRoot if it does not exist.
```

```
[ INFO ] Adding NSD None on nsd-node1.tuc.stglabs.ibm.com using device /dev/sdb.
```

```
[ INFO ] Configuration updated
```

```
[ INFO ] Tip : If all node designations and any required protocol configurations are complete, proceed to check the installation configuration: ./spectrumscale install --precheck
```

```
# ./spectrumscale nsd add -p nsd-node2 -s nsd-node1 -u dataAndMetadata -fs cesSharedRoot -fg 2 "/dev/sdc"
```

```
[ INFO ] Connecting to nsd-node2.tuc.stglabs.ibm.com to check devices and expand wildcards.
```

```
[ INFO ] Looking up details of /dev/sdc.
```

```
[ INFO ] Adding NSD None on nsd-node2.tuc.stglabs.ibm.com using device /dev/sdc.
```

```
[ INFO ] Configuration updated
```

```
[ INFO ] Tip : If all node designations and any required protocol configurations are complete, proceed to check the installation configuration: ./spectrumscale install --precheck
```

```
# ./spectrumscale nsd add -p nsd-node1 -s nsd-node2 -u dataAndMetadata -fs ObjectFS -fg 1 "/dev/sdd"
```

```
# ./spectrumscale nsd add -p nsd-node2 -s nsd-node1 -u dataAndMetadata -fs ObjectFS -fg 2 "/dev/sde"
```

```
# ./spectrumscale nsd add -p nsd-node1 -s nsd-node2 -u dataAndMetadata -fs fs1 -fg 1 "/dev/sdf"
```

```
# ./spectrumscale nsd add -p nsd-node1 -s nsd-node2 -u dataAndMetadata -fs fs1 -fg 1 "/dev/sdg"
```

```
# ./spectrumscale nsd add -p nsd-node2 -s nsd-node1 -u dataAndMetadata -fs fs1 -fg 2 "/dev/sdh"
```

```
# ./spectrumscale nsd add -p nsd-node2 -s nsd-node1 -u dataAndMetadata -fs fs1 -fg 2 "/dev/sdi"
```

Added Content – Installing GPFS (new cluster)

Here's how the NSDs will be setup

- NSDs are created during the install phase if they do not already exist
- If performing an install, deploy, or upgrade and NSDs already exist – the toolkit does not need to be told of them

```
# ./spectrumscale nsd list
```

	INFO	Name	FS	Size (GB)	Usage	FG	Pool	Device	Servers
[INFO]	nsd1	cesSharedRoot	10	dataAndMetadata	1	Default	/dev/sdb	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com	
[INFO]	nsd3	ObjectFS	250	dataAndMetadata	1	Default	/dev/sdd	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com	
[INFO]	nsd5	fs1	750	dataAndMetadata	1	Default	/dev/sdf	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com	
[INFO]	nsd6	fs1	750	dataAndMetadata	1	Default	/dev/sdg	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com	
[INFO]	nsd2	cesSharedRoot	10	dataAndMetadata	2	Default	/dev/sdc	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com	
[INFO]	nsd4	ObjectFS	250	dataAndMetadata	2	Default	/dev/sde	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com	
[INFO]	nsd7	fs1	750	dataAndMetadata	2	Default	/dev/sdh	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com	
[INFO]	nsd8	fs1	750	dataAndMetadata	2	Default	/dev/sdi	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com	

Here's how the file systems will be setup

- File systems are created during the deploy phase if their NSDs already exist
- If performing an install, deploy, or upgrade and FSs already exist – the toolkit does not need to be told of them

```
# ./spectrumscale filesystem list
```

	INFO	Name	BlockSize	Mountpoint	NSDs Assigned	Default Data Replicas	Max Data Replicas	Default Metadata Replicas	Max
Metadata Replicas									
[INFO]	cesSharedRoot	Default	(4M)/ibm/cesSharedRoot	2	1	2	2	2	2
[INFO]	ObjectFS	Default	(4M)/ibm/ObjectFS	2	1	2	2	2	2
[INFO]	fs1	Default	(4M)/ibm/fs1	4	1	2	2	2	

Added Content – adding to an existing cluster

Add 2 new NSDs and associate them with a new file system

```
# ./spectrumscale nsd add -p nsd-node3 -u dataAndMetadata -fs fs2 -fg 1 "/dev/sdb"
```

```
# ./spectrumscale nsd add -p nsd-node3 -u dataAndMetadata -fs fs2 -fg 2 "/dev/sdc"
```

Verify the new NSDs and file systems show up in the listing

```
# ./spectrumscale nsd list
```

	INFO	Name	FS	Size(GB)	Usage	FG	Pool	Device	Servers
[INFO]	nsd1	cesSharedRoot		10	dataAndMetadata	1	Default	/dev/sdb	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com]
[INFO]	nsd3	ObjectFS		250	dataAndMetadata	1	Default	/dev/sdd	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com]
[INFO]	nsd5	fs1		750	dataAndMetadata	1	Default	/dev/sdf	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com]
[INFO]	nsd6	fs1		750	dataAndMetadata	1	Default	/dev/sdg	[nsd-node1.tuc.stglabs.ibm.com],nsd-node2.tuc.stglabs.ibm.com]
[INFO]	nsd2	cesSharedRoot		10	dataAndMetadata	2	Default	/dev/sdc	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com]
[INFO]	nsd4	ObjectFS		250	dataAndMetadata	2	Default	/dev/sde	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com]
[INFO]	nsd7	fs1		750	dataAndMetadata	2	Default	/dev/sdh	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com]
[INFO]	nsd8	fs1		750	dataAndMetadata	2	Default	/dev/sdi	[nsd-node2.tuc.stglabs.ibm.com],nsd-node1.tuc.stglabs.ibm.com]
[INFO]	nsd9	fs1		750	dataAndMetadata	1	Default	/dev/sdb	[nsd-node3.tuc.stglabs.ibm.com]
[INFO]	nsd10	fs1		750	dataAndMetadata	2	Default	/dev/sdc	[nsd-node3.tuc.stglabs.ibm.com]

```
# ./spectrumscale filesystem list
```

	INFO	Name	BlockSize	Mountpoint	NSDs Assigned	Default Data Replicas	Max Data Replicas	Default Metadata Replicas	Max
Metadata Replicas									
[INFO]	cesSharedRoot	Default	(4M)/ibm/cesSharedRoot	2	1	2	2	2	2
[INFO]	ObjectFS	Default	(4M)/ibm/ObjectFS	2	1	2	2	2	2
[INFO]	fs1	Default	(4M)/ibm/fs1	4	1	2	2	2	2
[INFO]	fs2	Default	(4M)/ibm/fs2	2	1	2	2	2	2

Added Content – adding to an existing cluster

First - run an install precheck

```
# ./spectrumscale install --precheck
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/INSTALL-PRECHECK-10-05-2018_16:12:03.log
[ WARN ] Install toolkit will not configure file audit logging as it has been disabled.
[ INFO ] Checking for knife bootstrap configuration...
[ INFO ] Performing Chef (deploy tool) checks.
[ INFO ] Performing GPFS checks.
[ INFO ] Running environment checks
[ INFO ] Installed license validation passed. The installed license on GPFS cluster matching with the installer license version.
[ INFO ] The following nodes will be added to cluster democloud.tuc.stglabs.ibm.com: client-node3.tuc.stglabs.ibm.com, protocol-node3.tuc.stglabs.ibm.com, protocol-node4.tuc.stglabs.ibm.com
[ INFO ] Checking pre-requisites for portability layer.
[ INFO ] GPFS precheck OK
[ INFO ] Performing NSDs checks.
[ INFO ] Performing Performance Monitoring checks.
[ INFO ] Running environment checks for Performance Monitoring
[ INFO ] Performing GUI checks.
[ INFO ] Performing callhome checks.
[ INFO ] Checking input for callhome
[ INFO ] Checking network connectivity for callhome configuration.
[ INFO ] callhome precheck OK
[ INFO ] Network check from admin node protocol-node1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from admin node client-node3.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from admin node client-node1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node3.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node2.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node1.tuc.stglabs.ibm.com to all other protocol nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node3.tuc.stglabs.ibm.com to all other protocol nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node2.tuc.stglabs.ibm.com to all other protocol nodes in the cluster passed
[ WARN ] Ephemeral port range is not set. Please set valid ephemeral port range using the command ./spectrumscale config gpfs --ephemeral_port_range . You may set the default values as 60000-61000
[ INFO ] The install toolkit will reconfigure the call home group [autoGroup_1] with current configuration.
[ INFO ] Pre-check successful for install.
[ INFO ] Tip : ./spectrumscale install
```

Added Content – adding to an existing cluster

Next - run the Install:

```
# ./spectrumscale install
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/INSTALL-10-05-2018_16:14:03.log
[ INFO ] Validating configuration
```

When the Install finishes, you'll have:

- **1 new GUI node** (log into the GUI to activate it)
- **1 new client node** (remember that protocols will not be installed/enabled on the new node until the deploy phase)
- **1 new nsd node**
- **2 new NSDs** (remember that the new file system will not be built upon these NSDs until the deploy phase)

Now, it's time to run the Deploy Precheck:

```
# ./spectrumscale deploy --precheck
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/DEPLOY-PRECHECK-10-05-2018_16:30:16.log
```

If the Deploy Precheck is successful, run a deploy

Added Content – adding to an existing cluster

If the Deploy Precheck is successful, run a deploy

```
# ./spectrumscale deploy  
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/DEPLOY-10-05-2018_16:46:03.log
```

When the Deploy finishes, you'll have:

- **1 new protocol node** (CES-IPs may auto-distribute to this node dependent upon your CES IP address distribution policy)
- **1 new file system**

Adding protocol nodes to an ESS

Sample of what's to come

we don't currently support this so use at your own risk

SAMPLE of what's to come - CES osimage definition

ESS 5.3.0 now contains an osimage definition for CES nodes

```
[root@ems1 ~]# lsdef -t osimage
rhels7.0-ppc64-install-gss (osimage)
rhels7.1-ppc64-install-gss (osimage)
rhels7.2-ppc64-install-gss (osimage)
rhels7.2-ppc64le-install-gss (osimage)
rhels7.3-ppc64-install-ces (osimage)
rhels7.3-ppc64-install-gss (osimage)
rhels7.3-ppc64le-install-ces (osimage)
rhels7.3-ppc64le-install-gss (osimage)
```

Note

Using the CES definitions is not supported with 5.3.0, but if you're already 'hacking' the ESS to deploy protocol nodes, please give them a try. These definitions will be fully utilized in later releases of ESS. We're calling this a 'sample' for now.

SAMPLE of what's to come - CES osimage definition

GPFS rpms are excluded from the pkglist so the Install Toolkit can install GPFS.

```
[root@ems1 ~]# lsdef -t osimage rhels7.3-ppc64le-install-ces
```

```
Object name: rhels7.3-ppc64le-install-ces
```

```
  addcmdline=modprobe.blacklist=mpt3sas
```

```
  groups=all
```

```
  imagetype=linux
```

```
  osarch=ppc64le
```

```
  osname=Linux
```

```
  osvers=rhels7.3
```

```
  otherpkgdir=/install/gss/otherpkgs/rhels7/ppc64le
```

```
  otherpkglist=/opt/ibm/gss/xcat/install/rh/ces.rhels7.ppc64le.otherpkgs.pkglist
```

```
  pkgdir=/install/rhels7.3/ppc64le
```

```
  pkglist=/opt/ibm/gss/xcat/install/rh/gss.rhels7.ppc64le.pkglist
```

```
  postbootscripts=setupntp,gss_postboot,gss_ofed,gss_sashba
```

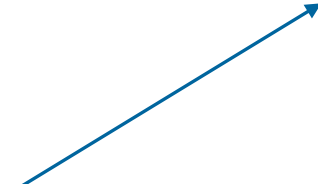
```
  postscripts=otherpkgs,gss_instncic,gss_post
```

```
  profile=ces
```

```
  provmethod=install
```

```
  synclists=/opt/ibm/gss/xcat/install/rh/gss.rhels7.ppc64le.synclist
```

```
  template=/opt/ibm/gss/xcat/install/rh/gss.rhels7.ppc64le.tmpl
```



```
kernel/kernel  
kernel/kernel-devel  
kernel/kernel-abi-whitelists  
kernel/kernel-bootwrapper  
kernel/kernel-doc  
kernel/kernel-headers  
kernel/kernel-tools  
kernel/kernel-tools-libs  
kernel/perf  
kernel/kmod-mpt3sas  
patch/systemd*  
patch/libgudev1*  
patch/NetworkManager*
```

Adding protocol nodes to an ESS

Recommendations

- Spectrum Scale 5.0.0.2 or higher is recommended
- If GPFS on the ESS will differ from GPFS on the Protocol nodes, make sure the versions are within the same release for the best compatibility.
- Hand install the OS/drivers or use the EMS XCAT (*at your own risk*), on all protocol nodes
- If deploying protocol nodes via XCAT (*at your own risk*) on the EMS:
 - Do not use the `ems1` or `gss_ppc64` groups. Make your own group.
 - Remove the protocol nodes and groups afterwards, to prevent an accidental future upgrade via XCAT.
 - Do not use the GPFS rpms from the ESS install. Instead, download a protocol node package, extract it on a protocol node, and use this.
 - Use the Install Toolkit, in ESS mode, running on a protocol node, to install GPFS rpms and add the protocol node(s) to the cluster
 - Check `resolv.conf` and `/etc/hosts` on each protocol node and configure as necessary
- A SEPARATE, small (4->10GB) filesystem is recommended for `cesSharedRoot`.

Adding protocol nodes to an ESS

Mandatory pre-reqs

- All protocol nodes need a base OS repo to be accessible
- /etc/hosts on all nodes must be formatted correctly: IP FQDN alias
- Promptless SSH must be setup for all nodes to all nodes
- Do not extract or run the Install Toolkit on the EMS node
- A filesystem to be used for cesSharedRoot (protocol state data) must be pre-created
- CES-IPs must be available and unused. A pre-existing adapter must be available on each protocol node with a placeholder IP. The placeholder IP must establish a route on a subnet compatible with the CES-IPs that will be added later. All CES-IPs must be in a DNS and setup for both forward and reverse DNS lookup. A single hostname set for DNS-RR is recommended in order to group all CES-IPs together.
- The ESS GPFS level must be within N-1 of the desired protocol node levels
- Protocol nodes must all be of the same OS and CPU architecture (*this can differ from ESS*)
- Firewalls must be down completely OR all nodes must have the correct ports opened (*see the KC for details*)
- SELinux settings should be checked. In some Scale levels, NFS needs a workaround to function with SELinux set at enforcing levels.

Adding protocol nodes to an ESS

Suggested flow

1) Start with an active ESS cluster (*GUI and perfmon collector should be configured and running*)

```
[root@ems1 ~]# mmgetstate -a -L
```

Node number	Node name	Quorum	Nodes up	Total nodes	GPFS state	Remarks
1	ems1	2	3	9	active	quorum node
2	gssio1	2	3	9	active	quorum node
3	gssio2	2	3	9	active	quorum node

```
[root@ems1 ~]# mmlscluster
```

GPFS cluster information

=====

```
GPFS cluster name:      democlustertuc.stglabs.ibm.com
GPFS cluster id:        16461288785736326788
GPFS UID domain:        democlustertuc.stglabs.ibm.com
Remote shell command:   /usr/bin/ssh
Remote file copy command: /usr/bin/scp
Repository type:        CCR
```

Node	Daemon node name	IP address	Admin node name	Designation
1	ems1.tuc.stglabs.ibm.com	9.11.84.9	ems1.tuc.stglabs.ibm.com	quorum-perfmon
2	gssio1.tuc.stglabs.ibm.com	9.11.84.12	gssio1.tuc.stglabs.ibm.com	quorum-manager-
perfmon				
3	gssio2.tuc.stglabs.ibm.com	9.11.84.13	gssio2.tuc.stglabs.ibm.com	quorum-manager-
perfmon				

Adding protocol nodes to an ESS

2) Choose a method for installing the OS on the protocol nodes

a) *Manually install the OS on each protocol node. Afterwards:*

- *Setup promptless-ssh to/from every node in the cluster for hostname, FQDN*
- *Configure an OS base repo on each protocol node. <yum repolist> must run clean*
- *If the OS matches the EMS, copy over the OFED drivers and install if applicable*
- *If the OS matches the EMS, copy over any kernel errata/system/drivers and install if applicable*

OR

b) *Use XCAT on the EMS node to install the OS/kernel/drivers/OFED on each protocol node*

- **this is not yet supported – we are working on the code + safeguards**
- *See the recommendations slide*

Adding protocol nodes to an ESS

3) Pick a protocol node and Extract the Spectrum Scale package

- In this example, we've picked a node that is not yet in the cluster. It will become a protocol node when we're finished

```
[root@protocol-node1 ~]# ./Spectrum_Scale_Protocols_Data_Management-5.0.1.0-x86_64-Linux-install
```

4) Agree to the license

LICENSE INFORMATION

The Programs listed below are licensed under the following License Information terms and conditions in addition to the Program license terms previously agreed to by Client and IBM. If Client does not have previously agreed to license terms in effect for the Program, the International Program License Agreement (Z125-3301-14) applies.

Program Name (Program Number):

IBM Spectrum Scale Data Management Edition V5.0.1 (5737-F34)
IBM Spectrum Scale Data Management Edition V5.0.1 (5641-DM1)
IBM Spectrum Scale Data Management Edition V5.0.1 (5641-DM3)
IBM Spectrum Scale Data Management Edition V5.0.1 (5641-DM5)

Press Enter to continue viewing the license agreement, or enter **"1" to accept the agreement**, "2" to decline it, "3" to print it, "4" to read non-IBM terms, or "99" to go back to the previous screen.

1

License Agreement Terms accepted.

Adding protocol nodes to an ESS

=====
Product rpms successfully extracted to /usr/lpp/mmfs/5.0.1.0

Cluster installation and protocol deployment

To install a cluster or deploy protocols with the Spectrum Scale Install Toolkit: `/usr/lpp/mmfs/5.0.1.0/installer/spectrumscale -h`

To install a cluster manually: Use the gpfs rpms located within `/usr/lpp/mmfs/5.0.1.0/gpfs_rpms`

To upgrade an existing cluster using the Spectrum Scale Install Toolkit:

- 1) Copy your old clusterdefinition.txt file to the new `/usr/lpp/mmfs/5.0.1.0/installer/configuration/` location
- 2) Review and update the config: `/usr/lpp/mmfs/5.0.1.0/installer/spectrumscale config update`
- 3) (Optional) Update the toolkit to reflect the current cluster config:
`/usr/lpp/mmfs/5.0.1.0/installer/spectrumscale config populate -N <node>`
- 4) Run the upgrade: `/usr/lpp/mmfs/5.0.1.0/installer/spectrumscale upgrade -h`

To add nodes to an existing cluster using the Spectrum Scale Install Toolkit:

- 1) Add nodes to the clusterdefinition.txt file: `/usr/lpp/mmfs/5.0.1.0/installer/spectrumscale node add -h`
- 2) Install GPFS on the new nodes: `/usr/lpp/mmfs/5.0.1.0/installer/spectrumscale install -h`
- 3) Deploy protocols on the new nodes: `/usr/lpp/mmfs/5.0.1.0/installer/spectrumscale deploy -h`

To add NSDs or file systems to an existing cluster using the Spectrum Scale Install Toolkit:

- 1) Add nsds and/or filesystems with: `/usr/lpp/mmfs/5.0.1.0/installer/spectrumscale nsd add -h`
- 2) Install the NSDs: `/usr/lpp/mmfs/5.0.1.0/installer/spectrumscale install -h`
- 3) Deploy the new file system: `/usr/lpp/mmfs/5.0.1.0/installer/spectrumscale deploy -h`

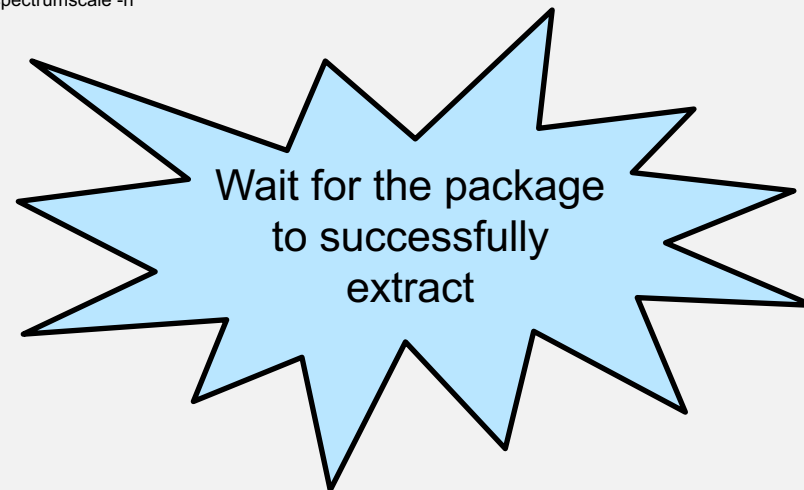
To update the toolkit to reflect the current cluster config examples:

- `/usr/lpp/mmfs/5.0.1.0/installer/spectrumscale config populate -N <node>`
- 1) Manual updates outside of the install toolkit
 - 2) Sync the current cluster state to the install toolkit prior to upgrade
 - 3) Switching from a manually managed cluster to the install toolkit

=====
To get up and running quickly, visit our wiki for an IBM Spectrum Scale Protocols Quick Overview:

<https://www.ibm.com/developerworks/community/wikis/home?lang=en#/wiki/General%20Parallel%20File%20System%20%28GPFS%29/page/Protocols%20Quick%20Overview%20for%20IBM%20Spectrum%20Scale>

=====



Adding protocol nodes to an ESS

6) Verify yum repolist is clean and has a base repo on each protocol node

```
# yum repolist
```

```
Loaded plugins: langpacks, product-id, search-disabled-repos, subscription-manager
repo id          repo name          status
ss.InstallMedia   ss.RHEL-7.2 Server.x86_64  4,620
repolist: 4,620
```

**Your repo config is contained here: /etc/yum.repos.d/*

**If you're unsure if the repo works, try a yum install on an rpm that doesn't exist on the node*

7) Change directories to the installer dir

```
# cd /usr/lpp/mmfs/5.0.1.0/installer
```

Adding protocol nodes to an ESS

8) Setup the Install Toolkit for ESS mode

```
# ./spectrumscale setup -s 9.11.102.38 -st ess
[ INFO ] Installing prerequisites for install node
[ INFO ] Chef successfully installed and configured
[ INFO ] Your control node has been configured to use the IP 9.11.102.38 to communicate with other nodes.
[ INFO ] Port 8889 will be used for chef communication.
[ INFO ] Port 10080 will be used for package distribution.
[ INFO ] Install Toolkit setup type is set to ESS. This mode will allow the EMS node to execute Install Toolkit commands.
[ INFO ] SUCCESS
[ INFO ] Tip : Designate an EMS node as admin node: ./spectrumscale node add <node> -a
[ INFO ] Tip : After designating an EMS node, add nodes for the toolkit to act upon: ./spectrumscale node add <node> -p -n
[ INFO ] Tip : After designating the EMS node, if you want to populate the cluster definition file with the current configuration,
you can run: ./spectrumscale config populate -N <ems_node>
```

- *9.11.102.38 is an IP on the installer node that can talk to the EMS and all other protocol nodes via promptless ssh*
- *-st ess is the new flag in Spectrum Scale 5.0.0.0 and onwards, allowing ESS mode*

9) Verify (from the EMS) that the ems and gss_ppc64 nodeclasses contain the right nodes

```
[root@ems1 ~]# mmlsnodeclass
Node Class Name      Members
-----
ems                  ems1.tuc.stglabs.ibm.com
gss_ppc64            gssio2.tuc.stglabs.ibm.com,gssio1.tuc.stglabs.ibm.com
GUI_MGMT_SERVERS    ems1.tuc.stglabs.ibm.com
```

- *The Install Toolkit will key off these nodeclasses to find the EMS node and protect the IO nodes*
- *No non-ESS nodes should exist within these nodeclasses*

Adding protocol nodes to an ESS

10) Populate the Install Toolkit with the current cluster config by pointing it to the EMS

```
[root@protocol-node1 installer]# ./spectrumscale config populate -N ems1.tuc.stglabs.ibm.com
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/config-populate-13-05-2018_19:54:06.log
[ INFO ] Detected clusterdefinition file present in configuration /usr/lpp/mmfs/5.0.1.0/installer/configuration.
Installer will keep backup of existing clusterdefinition.txt file in /usr/lpp/mmfs/5.0.1.0/installer/configuration path and populate a new
one. Do you want to continue [Y/n]: y
[ INFO ] Updating existing configuration. It may take few moments ....
[ INFO ] Getting installer path
[ INFO ] Populating protocols detail
[ INFO ] Populating ntp details
[ INFO ] Populating enabled protocols detail
[ INFO ] Cluster Export Services are not enabled in this cluster
[ INFO ] Populating export ip of ces nodes
[ INFO ] Cluster Export Services are not enabled in this cluster
[ INFO ] Adding admin node into configuration
[ INFO ] Populating cluster details
[ INFO ] Total 3 node(s) found in the cluster including 1 EMS and 2 IO node(s). EMS node details will be populated so that it may act as
an admin node for Install Toolkit functions. IO nodes will not be added and must be managed outside of the Install Toolkit
[ INFO ] Populating the details of ems1.tuc.stglabs.ibm.com
[ INFO ] Populating GPFS details
[ INFO ] Populating the zimon_collector details
[ INFO ] Populating NSD's details
[ INFO ] Total 5 NSDs detected that are hosted on the ESS IO nodes. These NSDs will not be added and must be managed outside of the
Install Toolkit
[ INFO ] Total 3 filesystems detected that are hosted on the ESS IO nodes. These filesystems will not be added and must be managed outside
of the Install Toolkit
[ INFO ] Populating Callhome details
[ INFO ] Populating file audit logging details
[ INFO ] Configuration successfully updated.
```

Tip: Not all cluster configs are compatible with config populate.

- If config populate fails, add the EMS node and set the cluster name in the toolkit to match, using ./spectrumscale commands.

***Notice the Toolkit will ignore the IO nodes, existing NSDs and FSs served by them**

Adding protocol nodes to an ESS

11) List the nodes/config that the Install Toolkit now knows about

```
# ./spectrumscale node list
[ INFO ] List of nodes in current configuration:
[ INFO ] [Installer Node]
[ INFO ] 9.11.102.38
[ INFO ]
[ INFO ] Setup Type: ESS
[ INFO ]
[ INFO ] [Cluster Name]
[ INFO ] democluster.tuc.stglabs.ibm.com
[ INFO ]
[ INFO ] File Audit logging : Disabled
[ INFO ]
[ INFO ] GPFS Node                               Admin  Quorum  Manager  NSD Server  Protocol  GUI Server  EMS  OS  Arch
[ INFO ] ems1.tuc.stglabs.ibm.com                X      X                               X      X      rhel7 ppc64le
```

The toolkit has picked up:

- *Cluster name*
- *EMS node*
- *GUI, Quorum roles*
- *Admin assignment simply means the Install Toolkit expects this node to be able to issue all mm commands*

Adding protocol nodes to an ESS

12) Let's add 2 client nodes to the Install Toolkit

```
[root@protocol-node1 installer]# ./spectrumscale node add protocol-node1
[ INFO ] Adding node protocol-node1.tuc.stglabs.ibm.com as a GPFS node.

[root@protocol-node1 installer]# ./spectrumscale node add protocol-node2
[ INFO ] Adding node protocol-node2.tuc.stglabs.ibm.com as a GPFS node.

[root@protocol-node1 installer]# ./spectrumscale node list
[ INFO ] List of nodes in current configuration:
[ INFO ] [Installer Node]
[ INFO ] 9.11.102.38
[ INFO ]
[ INFO ] Setup Type: ESS
[ INFO ]
[ INFO ] [Cluster Name]
[ INFO ] democloud.tuc.stglabs.ibm.com
[ INFO ]
[ INFO ] File Audit logging : Disabled
[ INFO ]
[ INFO ] GPFS Node                               Admin  Quorum  Manager  NSD Server  Protocol  GUI Server  EMS  OS  Arch
[ INFO ] emsl.tuc.stglabs.ibm.com                 X      X
[ INFO ] protocol-node1.tuc.stglabs.ibm.com          rhel17 x86_64
[ INFO ] protocol-node2.tuc.stglabs.ibm.com          rhel17 x86_64
```

Now that the Install Toolkit knows about these nodes, a subsequent install will add them to the cluster as client nodes

Adding protocol nodes to an ESS

13) Check perfmon and gpfs settings in the toolkit

```
# ./spectrumscale config gpfs
[ INFO ] No changes made to defaults. Current settings are as follows:
[ INFO ] GPFS cluster name is democluster.tuc.stglabs.ibm.com.
[ INFO ] GPFS profile is default.
[ INFO ] Remote shell command is /usr/bin/ssh.
[ INFO ] Remote file copy command is /usr/bin/scp.
[ WARN ] No value for GPFS Daemon communication port range in clusterdefinition file.

# ./spectrumscale config perfmon
[ INFO ] No changes made to defaults. Current settings are as follows:
[ INFO ] Performance Monitoring reconfiguration is on (default). Collectors may be moved to different nodes, sensors may be added to nodes, and sensors may be reset to defaults.
```

- *It's important to make sure the cluster name matches the existing cluster (it should after config populate)*
- *perfmon reconfig set to on (default) means the Install Toolkit will install sensors on each non-ESS node and activate them. The ESS perfmon config will remain in place and all sensors will point back to the EMS node so long as it was configured as a collector.*
- *The warning for no GPFS Daemon communication port range is not an issue if firewalls are down.*

Adding protocol nodes to an ESS

14) Run an install precheck

```
[root@protocol-node1 installer]# ./spectrumscale install --precheck
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/INSTALL-PRECHECK-13-05-2018_20:09:34.log
[ WARN ] NTP is already running on: ems1.tuc.stglabs.ibm.com. The install toolkit will no longer setup NTP.
[ WARN ] No NSD servers specified. The install toolkit will continue without creating any NSDs. If you still want to continue, please
ignore this warning. Otherwise, for information on adding a node as an NSD server,
see: 'http://www.ibm.com/support/knowledgecenter/STXKQY_5.0.1/com.ibm.spectrum.scale.v5r01.doc/bllins_configuringgpfs.htm'
[ INFO ] Checking for knife bootstrap configuration...
[ INFO ] Performing Chef (deploy tool) checks.
[ INFO ] Performing GPFS checks.
[ INFO ] Running environment checks
[ INFO ] The following nodes will be added to cluster democluster.tuc.stglabs.ibm.com: protocol-node1.tuc.stglabs.ibm.com, protocol-
node2.tuc.stglabs.ibm.com
[ WARN ] No manager nodes are configured. The Install Toolkit will assign manager nodes.
[ INFO ] Checking pre-requisites for portability layer.
[ INFO ] GPFS precheck OK
[ INFO ] Performing Performance Monitoring checks.
[ INFO ] Running environment checks for Performance Monitoring
[ INFO ] Network check from admin node ems1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ WARN ] Ephemeral port range is not set. Please set valid ephemeral port range using the command ./spectrumscale config gpfs --
ephemeral_port_range . You may set the default values as 60000-61000
[ INFO ] The install toolkit will not configure call home as it is disabled. To enable call home, use the following CLI command:
./spectrumscale callhome enable
[ INFO ] Pre-check successful for install.
[ INFO ] Tip : ./spectrumscale install
```

- Check the warnings and make sure they are expected before proceeding
- Note that callhome (mmcallhome) is not touched in an ESS environment by the Install Toolkit. Feel free to set this up on your own

Adding protocol nodes to an ESS

15) Run an install if the precheck was successful

```
[root@protocol-node1 installer]# ./spectrumscale install
```

You'll see messages indicating:

- *Repositories being created/removed*
- *Components being installed / started / restarted*
- *Nodes being added to the cluster*
- *Licenses being enabled*

The end result should be a successful install

```
[ INFO ] Checking for a successful install
[ INFO ] Checking state of Chef (deploy tool)
[ INFO ] Chef (deploy tool) ACTIVE
[ INFO ] Checking state of GPFS
[ INFO ] GPFS callhome has been successfully installed. To configure callhome run 'mmcallhome -h' on one of your nodes.
[ INFO ] 3 server licenses and 2 client license enabled on the cluster. Ensure enough licenses are owned for this configuration.
[ INFO ] Checking state of GPFS on all nodes
[ INFO ] GPFS active on all nodes
[ INFO ] GPFS ACTIVE
[ INFO ] Checking state of Performance Monitoring
[ INFO ] Running Performance Monitoring post-install checks
[ INFO ] pmcollector running on all nodes
[ INFO ] pmsensors running on all nodes
[ INFO ] Performance Monitoring ACTIVE
[ INFO ] SUCCESS
[ INFO ] Installation successful. 3 GPFS nodes active in cluster democuster.tuc.stglabs.ibm.com. Completed in 5 minutes 36 seconds.
[ INFO ] Tip :If all node designations and any required protocol configurations are complete, proceed to check the deploy
configuration:./spectrumscale deploy --precheck
```

Adding protocol nodes to an ESS

16) Check the cluster afterwards

```
[root@protocol-node1 installer]# mmlscluster
```

```
GPFS cluster information
```

```
=====
```

```
GPFS cluster name:      democuster.tuc.stglabs.ibm.com
GPFS cluster id:        16461288785736326788
GPFS UID domain:        democuster.tuc.stglabs.ibm.com
Remote shell command:   /usr/bin/ssh
Remote file copy command: /usr/bin/scp
Repository type:        CCR
```

Node	Daemon node name	IP address	Admin node name	Designation
1	ems1.tuc.stglabs.ibm.com	9.11.84.9	ems1.tuc.stglabs.ibm.com	quorum-perfmon
2	gssio1.tuc.stglabs.ibm.com	9.11.84.12	gssio1.tuc.stglabs.ibm.com	quorum-manager-perfmon
3	gssio2.tuc.stglabs.ibm.com	9.11.84.13	gssio2.tuc.stglabs.ibm.com	quorum-manager-perfmon
4	protocol-node1.tuc.stglabs.ibm.com	9.11.102.38	protocol-node1.tuc.stglabs.ibm.com	perfmon
5	protocol-node2.tuc.stglabs.ibm.com	9.11.102.55	protocol-node2.tuc.stglabs.ibm.com	perfmon

Our new nodes have been added

- *Check `mmgetstate -a` to verify the new nodes are active*
- *Check `'mmlsmount all'` to verify FSs are mounted on these nodes*
- *Check `systemctl status pmsensors` to verify sensors are running*
- *Check `mmperfmon config show` to verify collector is still the EMS node*

Adding protocol nodes to an ESS

17) We've just added client nodes, let's make them protocol nodes now

```
[root@protocol-node1 installer]# ./spectrumscale node add protocol-node1 -p
[ INFO ] Setting protocol-node1.tuc.stglabs.ibm.com as a protocol node.
[ INFO ] Configuration updated.
[ INFO ] Tip : If all node designations are complete, configure the protocol environment as needed: ./spectrumscale config protocols -f cesSharedRoot -m /ibm/cesSharedRoot

[root@protocol-node1 installer]# ./spectrumscale node add protocol-node2 -p
[ INFO ] Setting protocol-node2.tuc.stglabs.ibm.com as a protocol node.
[ INFO ] Configuration updated.
[ INFO ] Tip : If all node designations are complete, configure the protocol environment as needed: ./spectrumscale config protocols -f cesSharedRoot -m /ibm/cesSharedRoot
```

Before proceeding, identify (or create) an FS that will be used for cesSharedRoot

- *Create this FS from the EMS if it does not already exist*
- *The FS should be >=4GB*
- *An existing FS of any size can be used, but it's always recommended to have a separate FS just for cesSharedRoot. This will allow for easier FS service in the future.*

Adding protocol nodes to an ESS

18) Verify cesSharedRoot is mounted on all to-be protocol nodes

```
[root@protocol-node1 installer]# mmlsmount all -L
```

```
File system cesSharedRoot is mounted on 5 nodes:
```

9.11.84.13	gssio2
9.11.84.9	ems1
9.11.84.12	gssio1
9.11.102.55	protocol-node2
9.11.102.38	protocol-node1

Adding protocol nodes to an ESS

19) Add CES-IPs to the Install Toolkit

```
# ./spectrumscale config protocols -e 9.11.102.107,9.11.102.108,9.11.102.138,9.11.102.80
[ INFO ] Setting Export IP Pool to 9.11.102.107,9.11.102.108,9.11.102.138,9.11.102.80
[ INFO ] Tip :Enable NFS, Object or SMB protocols as appropriate:./spectrumscale enable nfs|object|smb
```

20) Point the Install Toolkit to the FS name and mount point of cesSharedRoot

```
# ./spectrumscale config protocols -f cesSharedRoot -m /ibm/cesSharedRoot
[ INFO ] Setting Shared File System Name to cesSharedRoot
[ INFO ] Setting Shared File System Mountpoint to /ibm/cesSharedRoot
[ INFO ] Tip :Enable NFS, Object or SMB protocols as appropriate:./spectrumscale enable nfs|object|smb
```

This is mandatory – make sure the FS name and path is correct

21) Enable the File Protocols if desired

```
# ./spectrumscale enable nfs
[ INFO ] Enabling NFS on all protocol nodes.

# ./spectrumscale enable smb
[ INFO ] Enabling SMB on all protocol nodes.
```

Adding protocol nodes to an ESS

22) Configure Object if desired

```
# ./spectrumscale enable object
[ INFO ] Enabling OBJECT on all protocol nodes.
# ./spectrumscale config object -e ces-exportips
[ INFO ] Setting Endpoint Hostname to ces-exportips
# ./spectrumscale config object -o Object_Fileset
[ INFO ] Setting GPFS Object Base to Object_Fileset
# ./spectrumscale config object -f ObjectFS -m /ibm/ObjectFS
[ INFO ] Setting Object File System Name to ObjectFS
[ INFO ] Setting Object File System Mountpoint to /ibm/ObjectFS
# ./spectrumscale config object -au admin -ap -dp
[ INFO ] Setting Admin User to admin
```

At this point you will be asked to type the admin and database passwords. They will be encrypted with the password you type as a secret key

- *Make sure the FS to be used for Object is mounted on all protocol nodes.*
- *Make sure the Fileset to be used for Object does not pre-exist*
- *The Object endpoint can be a hostname that corresponds to the DNS-RR hostname containing all CES-IPs*

Adding protocol nodes to an ESS

23) Check the Install Toolkit config

```
[root@protocol-node1 installer]# ./spectrumscale node list
[ INFO ] List of nodes in current configuration:
[ INFO ] [Installer Node]
[ INFO ] 9.11.102.38
[ INFO ]
[ INFO ] Setup Type: ESS
[ INFO ]
[ INFO ] [Cluster Name]
[ INFO ] democluster.tuc.stglabs.ibm.com
[ INFO ]
[ INFO ] [Protocols]
[ INFO ] Object : Enabled
[ INFO ] SMB : Enabled
[ INFO ] NFS : Enabled
[ INFO ]
[ INFO ] File Audit logging : Disabled
[ INFO ]
[ INFO ] GPFS Node
[ INFO ] ems1.tuc.stglabs.ibm.com      Admin  Quorum  Manager  NSD Server  Protocol  GUI Server  EMS  OS  Arch
[ INFO ] protocol-node1.tuc.stglabs.ibm.com      X      X      X      X      X      X      X  rhel7  ppc64le
[ INFO ] protocol-node2.tuc.stglabs.ibm.com      X      X      X      X      X      X      X  rhel7  x86_64
[ INFO ]
[ INFO ] [Export IP address]
[ INFO ] 9.11.102.80 (pool)
[ INFO ] 9.11.102.107 (pool)
[ INFO ] 9.11.102.108 (pool)
[ INFO ] 9.11.102.138 (pool)
```

If deploying NFS – RHEL 7.3 / RHEL7.4 sometimes start NFS by default on each node after install/upgrade.

Run the following to stop/disable the OS version of NFS on each protocol node so it doesn't conflict with Spectrum Scale NFS: a) `systemctl disable nfs` b) `systemctl stop nfs`

Adding protocol nodes to an ESS

24) Start a deploy precheck

```
[root@protocol-node1 installer]# ./spectrumscale deploy --precheck
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/DEPLOY-PRECHECK-13-05-2018_20:27:02.log
[ WARN ] File Authentication will not be configured on your nodes. Run ./spectrumscale auth --help for information on auth.
[ WARN ] Object Authentication will be configured with a self-signed certificate. It is recommended to use stronger authentication schemes for a production environment. Run ./spectrumscale auth --help for information on auth.
[ INFO ] Checking for knife bootstrap configuration...
[ INFO ] Performing Chef (deploy tool) checks.
[ INFO ] Performing Filesystem checks.
[ INFO ] NSDs are in a valid state
[ INFO ] Performing Cluster Export Services checks.
[ INFO ] Running environment checks for protocols
[ INFO ] Checking state of GPFS on all nodes
[ INFO ] GPFS active on all nodes
[ INFO ] Checking state of GPFS on all nodes
[ INFO ] GPFS active on all nodes
[ INFO ] Specified GPFS device cesSharedRoot was detected on node protocol-node1.tuc.stglabs.ibm.com. Will be used for Protocols shared root.
[ INFO ] Specified GPFS device cesSharedRoot was detected on node protocol-node2.tuc.stglabs.ibm.com. Will be used for Protocols shared root.
[ INFO ] Protocols shared root file system is correctly mounted at /gpfs/cesSharedRoot on node protocol-node2.tuc.stglabs.ibm.com.
[ INFO ] Protocols shared root file system is correctly mounted at /gpfs/cesSharedRoot on node protocol-node1.tuc.stglabs.ibm.com.
[ INFO ] protocol precheck OK
[ INFO ] Performing Object Store checks.
[ INFO ] Running environment checks for Object Storage
[ INFO ] Specified GPFS device ObjectFS was detected on node protocol-node2.tuc.stglabs.ibm.com. Will be used for Object base.
[ INFO ] Specified GPFS device ObjectFS was detected on node protocol-node1.tuc.stglabs.ibm.com. Will be used for Object base.
[ INFO ] Object base file system is correctly mounted at /gpfs/ObjectFS on node protocol-node2.tuc.stglabs.ibm.com.
[ INFO ] Object base file system is correctly mounted at /gpfs/ObjectFS on node protocol-node1.tuc.stglabs.ibm.com.
[ INFO ] Object Storage ready for install
[ INFO ] Performing SMB checks.
[ INFO ] Running environment checks for SMB
[ INFO ] SMB precheck OK
[ INFO ] Performing NFS checks.
[ INFO ] Running environment checks for NFS
[ INFO ] NFS precheck OK
[ INFO ] Performing Performance Monitoring checks.
[ INFO ] Running environment checks for Performance Monitoring
[ INFO ] Network check from admin node ems1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node2.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node2.tuc.stglabs.ibm.com to all other protocol nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node1.tuc.stglabs.ibm.com to all other protocol nodes in the cluster passed
[ WARN ] Ephemeral port range is not set. Please set valid ephemeral port range using the command ./spectrumscale config gpfs --ephemeral_port_range . You may set the default values as 60000-61000
[ INFO ] The install toolkit will not configure call home as it is disabled. To enable call home, use the following CLI command: ./spectrumscale callhome enable
[ INFO ] Pre-check successful for deploy.
[ INFO ] Tip : ./spectrumscale deploy
```

Adding protocol nodes to an ESS

25) If the deploy precheck was successful and all warnings are understood, proceed with a deploy:

```
[root@protocol-node1 installer]# ./spectrumscale deploy
```

You'll see messages indicating:

- *Repositories being created/removed*
- *Components being installed / started / restarted*
- *Protocols being configured and enabled*
- *Licenses being enabled*

The end result should be a successful deploy with all desired protocols active

Adding protocol nodes to an ESS

26) Check the cluster afterwards

```
[root@protocol-node1 installer]# mmhealth cluster show
```

Component	Total	Failed	Degraded	Healthy	Other
-----	-----	-----	-----	-----	-----
NODE	5	0	0	5	0
GPFS	5	0	0	5	0
NETWORK	5	0	0	5	0
FILESYSTEM	2	0	0	2	0
DISK	8	0	0	8	0
CES	2	0	0	2	0
NATIVE_RAID	2	0	0	2	0
GUI	1	0	0	1	0
PERFMON	3	0	0	3	0
THRESHOLD	3	0	0	3	0

- All nodes are healthy
- All protocols are active
- Neither protocol node show any failure node flags
- All CES IPs are assigned to the nodes

```
[root@protocol-node1 installer]# mmces service list -a
```

Enabled services: OBJ SMB NFS

protocol-node2.tuc.stglabs.ibm.com: OBJ is running, SMB is running, NFS is running

protocol-node1.tuc.stglabs.ibm.com: OBJ is running, SMB is running, NFS is running

```
[root@protocol-node1 installer]# mmces address list
```

Address	Node	Group	Attribute
-----	-----	-----	-----
9.11.102.80	protocol-node1.tuc.stglabs.ibm.com	none	object_database_node
9.11.102.107	protocol-node2.tuc.stglabs.ibm.com	none	none
9.11.102.108	protocol-node2.tuc.stglabs.ibm.com	none	none
9.11.102.138	protocol-node1.tuc.stglabs.ibm.com	none	object_singleton_node

```
[root@protocol-node1 installer]# mmces node list
```

Node	Name	Node Flags	Node Groups
-----	-----	-----	-----
4	protocol-node1.tuc.stglabs.ibm.com	none	
5	protocol-node2.tuc.stglabs.ibm.com	none	

Adding protocol nodes to an ESS

27) What can we do now?

The Install Toolkit can be used to add more nodes or features to the cluster. For example:

- Add nodes and give them the `-n` flag, add disks to these nodes, follow with an install, and you'll have non-ESS NSD nodes with NSDs
- Modify the NSDs within the Install Toolkit and associate new or existing file systems, follow with a deploy, and you'll have either new file systems built upon these NSDs or you'll have added NSDs to existing file systems.
- Add more protocol nodes if desired, by re-running install and deploy
- Enable more protocols if desired, by re-running deploy
- For Authentication, we recommend using `mmuserauth` instead of the Install Toolkit

Using the Install Toolkit with LTFS-EE nodes in the same cluster

Using the Install Toolkit with LTFS-EE nodes in the same cluster

What if I want a Spectrum Archive (LTFS-EE) node within my cluster?

- Build a cluster using the Install Toolkit
- Before installing LTFS-EE, you must enable DMAPI on the file system to be used with LTFS-EE. This requires an unmount of the file system. If this file system is also used for protocols then it will be impossible to unmount unless you first issue an mmshutdown on all protocol nodes first. Once dmapi is set (*mmchfs Device -z yes*), re-mount the file system and issue an mmstartup on all protocol nodes.
- Install / configure LTFS-EE
- To upgrade, perform steps a through e on all LTFS-EE nodes:

- | | |
|-------------------------------|--|
| a) ltfsee stop | f) upgrade using the toolkit |
| b) unmount /ltfs | g) upgrade LTFS-EE |
| c) dsmmigfs disablefailover | h) reverse the steps an through e and restart/enable |
| d) dsmmigfs stop | |
| e) systemctl stop hsm.service | |

NTP Configuration

(only for non-ESS environments)

Added Content – NTP configuration *(only for non-ESS environments)*

NTP pre-requisites

- All nodes must be able to reach the NTP server IPs

- The NTP package must already be installed on all nodes

- Full bidirectional access to the UDP port 123 must be allowed on all nodes

When does NTP get configured?

- The Install GUI will automatically configure NTP if IPs are added

- The Install toolkit will configure NTP upon 'spectrumscale install' provided it is set to on

When does NTP not get configured?

- If it is spectrumscale config ntp -e off is set

- If a cluster pre-exists

- If NTP is already running

- During 'spectrumscale deploy' or 'spectrumscale upgrade'

Added Content – NTP configuration *(only for non-ESS environments)*

New NTP help showing all options

```
# ./spectrumscale config ntp -h
usage: spectrumscale config ntp [-h] [-l] [-e {on,off}] [-s UPSTREAM_SERVERS]

optional arguments:
  -h, --help            show this help message and exit
  -l, --list            List the current settings in the configuration
  -e {on,off}, --enabled {on,off}
                        Specify whether NTP should be enabled
  -s UPSTREAM_SERVERS, --upstream-servers UPSTREAM_SERVERS
                        Specify the Upstream NTP Servers(comma separated IP's
                        with NO space between multiple IPs)
```

Configuring the cluster nodes to point to two NTP servers

```
# ./spectrumscale config ntp -e on -s 9.11.107.11,9.11.107.12
[ WARN  ] The NTP package must already be installed and full bidirectional access to the UDP port 123 must be allowed.
[ WARN  ] If NTP is already running on any of your nodes, NTP setup will be skipped. To stop NTP run 'service ntpd stop'.
[ INFO  ] Setting Upstream NTP Servers(comma separated IP's with NO space between multiple IPs) to 9.11.107.11,9.11.107.12
[ WARN  ] Your SLES12 nodes may be running AppArmor as it is a default install option for SLES.  If so, enabling NTP will cause the
installation to fail unless AppArmor profiles are updated to include the NTP services.  See the following Knowledge Center link for details:
http://www.ibm.com/support/knowledgecenter/STXKQY_5.0.1/com.ibm.spectrum.scale.v5r01.doc/bllins_protocolsprerequisites.htm
```

Performance Monitor re-configuration

Added Content – Perfmon re-configuration

What does perfmon re-configuration mean?

- Performance monitoring configuration may need changing in situations such as when adding a GUI node as this requires performance statistics to be available locally on the node, and thus a perfmon collector must be installed. In other cases, it may be necessary to keep the toolkit from either upgrading perfmon rpms or reconfiguring the collectors.
- The install toolkit can be set to not touch perfmon or it can be set to reconfigure when necessary.

When does perfmon get re-configured?

- When GUI nodes are added.
- If only one collector exists in a non-ESS environment, the toolkit will add a 2nd . It will always choose GUI nodes first. NSD nodes will be preferred next. Client nodes will be preferred last.

When does NTP not get re-configured?

- In an ESS environment (*unless the toolkit is adding non-ESS NSD nodes*)
- When the flag is set to no. If this is selected, the Knowledge Center should be consulted for manual configuration.

Added Content – Perfmon re-configuration

New perfmon help showing all options

```
# ./spectrumscale config perfmon -h
```

```
usage: spectrumscale config perfmon [-h] [-l] [-r {on,off}]
```

optional arguments:

- h, --help show this help message and exit
- l, --list List the current settings in the configuration
- r {on,off}, --enable-reconfig {on,off}
Specify if the install toolkit can reconfigure
Performance Monitoring. When set to on,
reconfiguration may move the collector to different
nodes and may reset sensor data. Custom sensors and
data may be erased.

Setting perfmon re-configuration to on for the next 'spectrumscale install'

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

```
[ INFO ] Setting Performance Monitoring reconfiguration to on
```

Cautions:

Depending upon your Scale version, some sensor values may be reset to defaults when the Install Toolkit is run. Sensors to watch for include:

- GPFSSetQuota, GPFSDiskCap
- NFSIO *(5.0.1.0 and higher will not reset this sensor)*

Added Content: Cluster Upgrade

Added Content – cluster Upgrade

Extract the 5.0.1.0 package

- The Install Toolkit is only included within protocol packages
- Ideally, use the same licensed package as is currently installed on the cluster. However, if desired, the Install Toolkit can migrate the user from a lower license to a higher one. For example:

Standard to Advanced

Standard to Advanced

Advanced to Data Management

```
# ./Spectrum_Scale_Protocols_Data_Management-5.0.1.0-x86_64-Linux-install
```

Setup the Toolkit

- The node you extract the package and setup the toolkit on, will become the installer node
- **If you have an ESS in the cluster:** use the --setuptype ess. If not, use --setuptype ss. The default is 'ss'

```
# ./spectrumscale setup -s 9.11.102.38 --setuptype ss
```

```
[ INFO ] Installing prerequisites for install node
```

```
[ INFO ] Chef successfully installed and configured
```

```
[ INFO ] Your control node has been configured to use the IP 9.11.102.38 to communicate with other nodes.
```

```
[ INFO ] Port 8889 will be used for chef communication.
```

```
[ INFO ] Port 10080 will be used for package distribution.
```

```
[ INFO ] Install Toolkit setup type is set to Spectrum Scale (default). If an ESS is in the cluster, run this command to set ESS mode: ./spectrumscale setup -s server_ip -st ess
```

```
[ INFO ] SUCCESS
```

```
[ INFO ] Tip : Designate protocol, nsd and admin nodes in your environment to use during install:./spectrumscale -v node add <node> -p -a -n
```

Added Content – cluster Upgrade

You have a few options to tell the Toolkit about the existing cluster prior to upgrade

#1: Config Populate

- Point it at a node and it will traverse the cluster and add all of nodes of the same architecture to the toolkit.
- If you'd rather choose the exact nodes to upgrade, simply delete the nodes from the toolkit afterwards.
- Config populate has a few limitations that may cause it not to work with your config, though there is no harm in trying.

#2: Copy an old clusterdefinition.txt file

- If you've used the Toolkit before, you'll have a clusterdefinition.txt in your prior code's extraction point
- Copy this old clusterdefinition.txt file to the new location
- You will have to run a `./spectrumsale config update` for the toolkit to update the clusterdefinition.txt version

#3: Input cluster information into the toolkit by hand

- You can always use `./spectrumscale` commands to tell the toolkit about your cluster.
- The Toolkit will not need to know about NSDs or File Systems for an upgrade to proceed
- The Toolkit does not need to know if a node is an NSD node.
- The Toolkit does need to know if a node is a protocol node, or GUI node, and it needs to know the protocol configuration

Added Content – cluster Upgrade

Run config populate

Run config populate from the node you extracted code on

- Here, it is run on protocol-node1

Point config populate to any node in the cluster

- Here, it points to itself, protocol-node1, which is already in the cluster.
- If you have an ESS, point to the EMS node
- If you've extracted the code to a node not in the cluster, point it to a node in the cluster. The node not in the cluster can still be the installer node.
- The node you point to will become the admin node for the toolkit. It will need promptless ssh access to/from all nodes. The toolkit will run mm commands on this node.

What if config populate doesn't work with my config?

- You can still use the toolkit, you'll just need to input at least one node in the cluster via `./spectrumscale node add` and you'll need to make this node an admin node.

```
[root@protocol-node1 installer]# ./spectrumscale config populate -N protocol-node1
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/config-populate-10-05-2018_15:14:49.log
[ INFO ] Detected clusterdefinition file present in configuration /usr/lpp/mmfs/5.0.1.0/installer/configuration.
Installer will keep backup of existing clusterdefinition.txt file in /usr/lpp/mmfs/5.0.1.0/installer/configuration path and populate a new one. Do you want to
continue [Y/n]: y
[ INFO ] Updating existing configuration. It may take few moments ....
[ INFO ] Getting installer path
[ INFO ] Populating protocols detail
[ INFO ] Populating ntp details
[ INFO ] Populating enabled protocols detail
[ INFO ] Object is enabled. so populating object details.
[ INFO ] The passwords entered for object can be those that were initially used for setting up object or optionally enter new passwords to change to new
value. Enter the encryption passwords that will be used to encrypt the sensitive data in the cluster configuration file.
[ INFO ] Populating export ip of ces nodes
[ INFO ] Adding admin node into configuration
[ INFO ] Populating cluster details
[ INFO ] Total 6 node(s) found in the cluster
[ INFO ] Populating the details of protocol-node2.tuc.stglabs.ibm.com
[ INFO ] Populating the details of protocol-node1.tuc.stglabs.ibm.com
[ INFO ] Populating the details of client-node2.tuc.stglabs.ibm.com
[ INFO ] Populating the details of nsd-node1.tuc.stglabs.ibm.com
[ INFO ] Populating the details of client-node1.tuc.stglabs.ibm.com
[ INFO ] Populating the details of nsd-node2.tuc.stglabs.ibm.com
[ INFO ] Defining client-node1.tuc.stglabs.ibm.com as a GUI node
[ INFO ] Populating the GUI addresses details
[ INFO ] Populating GPFS details
[ INFO ] Populating the zimmon_collector details
[ INFO ] Populating the secondary_zimmon_collector details
[ INFO ] Authentication file _ad is set. so populating the details
[ INFO ] Populating NSD's details
[ INFO ] Total 8 NSDs found in the cluster
[ INFO ] Populating the details of nsd3
[ INFO ] Populating the details of nsd4
[ INFO ] Populating the details of nsd1
[ INFO ] Populating the details of nsd2
[ INFO ] Populating the details of nsd5
[ INFO ] Populating the details of nsd6
[ INFO ] Populating the details of nsd7
[ INFO ] Populating the details of nsd8
[ INFO ] Populating filesystem details
[ INFO ] Total 3 filesystem found in the cluster.
[ INFO ] Populating filesystem details for ObjectFS
[ INFO ] Populating filesystem details for cesSharedRoot
[ INFO ] Populating filesystem details for fs1
[ INFO ] Populating Callhome details
[ INFO ] Populating file audit logging details
[ INFO ] Configuration successfully updated.
```


Added Content – cluster Upgrade

Check the Toolkit config after config populate finishes

```
# ./spectrumscale node list
[ INFO ] List of nodes in current configuration:
[ INFO ] [Installer Node]
[ INFO ] 9.11.102.38
[ INFO ]
[ INFO ] Setup Type: SpectrumScale
[ INFO ]
[ INFO ] [Cluster Name]
[ INFO ] democluster.tuc.stglabs.ibm.com
[ INFO ]
[ INFO ] [Protocols]
[ INFO ] Object : Enabled
[ INFO ] SMB : Enabled
[ INFO ] NFS : Enabled
[ INFO ]
[ INFO ] File Audit logging : Disabled
[ INFO ]
[ INFO ]
[ INFO ] GPFS Node      Admin  Quorum  Manager  NSD Server  Protocol  GUI  Server  Callhome Node  OS      Arch
[ INFO ] client-node1.tuc.stglabs.ibm.com  X                               X                               ubuntu16 x86_64
[ INFO ] client-node2.tuc.stglabs.ibm.com                               X                               X                               sles12   x86_64
[ INFO ] nsd-node1.tuc.stglabs.ibm.com      X      X      X      X
[ INFO ] nsd-node2.tuc.stglabs.ibm.com      X      X      X      X
[ INFO ] protocol-node1.tuc.stglabs.ibm.com X      X                               X      X      rhel7    x86_64
[ INFO ] protocol-node2.tuc.stglabs.ibm.com                               X      X      rhel7    x86_64
[ INFO ]
[ INFO ] [Export IP address]
[ INFO ] 9.11.102.80 (pool)
[ INFO ] 9.11.102.107 (pool)
[ INFO ] 9.11.102.108 (pool)
[ INFO ] 9.11.102.138 (pool)
```

What happens when you upgrade?

- All nodes listed in the Install Toolkit (except for an EMS) will be upgraded
- Upgrade is sequential, node at a time.

What if I don't want to upgrade all nodes now?

- Remove the nodes you don't want to upgrade from the Install Toolkit
- All CES nodes must be upgraded together. Please don't attempt to split these up or the toolkit's upgrade logic will assume you have less CES nodes than you do, potentially causing outages.

Added Content – cluster Upgrade

Run the Upgrade Precheck

```
# ./spectrumscale upgrade --precheck
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/UPGRADE-PRECHECK-13-05-2018_10:16:07.log
[ INFO ] Checking state of OBJ on all nodes
[ INFO ] OBJ healthy on all nodes
[ INFO ] Object precheck OK
[ INFO ] Checking state of NFS on all nodes
[ INFO ] NFS healthy on all nodes
[ INFO ] NFS precheck OK
[ INFO ] pmsensors running on all nodes
[ INFO ] pmcollector running on all nodes
[ INFO ] Performance monitoring precheck OK
[ INFO ] Checking state of GPFS
[ INFO ] Checking state of GPFS
[ INFO ] GPFS active on nodes
[ INFO ] GPFS health check OK
[ INFO ] Checking for packages dependent on GPFS
[ INFO ] GPFS dependent check OK
[ INFO ] Checking pre-requisites for portability layer.
[ INFO ] Checking state of SMB on all nodes
[ INFO ] SMB healthy on all nodes
[ INFO ] SMB precheck OK
```

Continued...

Added Content – cluster Upgrade

During precheck: pay attention to the detected code levels

- The Install Toolkit can upgrade a cluster containing mixed code levels as long as all levels fit in the N-1 -> N rules
- This example shows nodes at levels 5.0.0-1, 4.2.3-8, 5.0.0-2 will all be upgraded to GPFS 5.0.1-0

```
[ INFO ] Checking installed and available GPFS and protocol package versions...
[ INFO ] GPFS
[ INFO ] GPFS version available in repo: 5.0.1-0
[ INFO ] GPFS version(s) installed on nodes: 5.0.0-1, 4.2.3-8, 5.0.0-2, 5.0.0-2
[ INFO ] New version of GPFS available in repo
```

During precheck: pay attention to the success of the network checks

- If any network checks fail, analyze promptless ssh (FQDN and hostname) and firewall, before continuing

```
[ INFO ] Network check from admin node protocol-node1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from admin node client-node1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node2.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node2.tuc.stglabs.ibm.com to all other protocol nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node1.tuc.stglabs.ibm.com to all other protocol nodes in the cluster passed
```

You're ready to upgrade if the upgrade precheck is successful

```
[ INFO ] SUCCESS: Ready for upgrade
```

Added Content – cluster Upgrade

Start the upgrade

```
# ./spectrumscale upgrade
```

Do you really want to begin upgrading? This may cause a brief Object, NFS, SMB and Performance Monitoring outage. [y/N]: y

What does this outage warning mean?

- **SMB**

Mixed versions of SMB cannot be active in the same cluster at the same time. If SMB is enabled, Install Toolkit logic will divide the CES nodes into 2 groups. It will suspend and stop SMB on the first group and then upgrade it. This suspend will cause CES-IPs on these nodes to failover to the remaining active CES nodes. Any SMB I/O running to the nodes being suspend will need to be re-driven by the client. Next in the upgrade, the second group of CES nodes will be suspended and SMB will be stopped. At this moment, SMB will experience a cluster-wide outage. The moment the second group of CES nodes have SMB stopped, the first group of CES nodes will be resumed and SMB started. At this point, the SMB outage is over. Now, the second group of CES nodes will undergo an SMB upgrade. When finished, the second group of CES nodes will be resumed and SMB started. It's recommended to quiesce SMB I/O for the duration of the upgrade, if possible.

- **NFS**

NFS is concurrent and active mixed versions within the same cluster are compatible. Depending upon the settings of your clients driving I/O to NFS, they may see brief pauses in I/O during each suspend that occurs during upgrade. It's recommended to quiesce NFS I/O for the duration of the upgrade, if possible.

- **Object**

Mixed versions of Object are not compatible when active at the same time. During upgrade, Object will be stopped on all nodes and upgraded on all nodes at the same time. During this time period, I/O via the object protocol will pause and depending upon your client settings, may need to be restarted once the upgrade has finished. It's recommended to quiesce Object I/O for the duration of the upgrade, if possible.

- **Performance monitoring**

Performance statistics will continue to be gathered on all sensor nodes during the course of the upgrade, however there may be small gaps in the statistics of a node at the moment it is undergoing upgrade. In addition, the collector node will not receive statistics from sensors during the short time it takes to stop the collector, upgrade, and restart the collector.

- **Callhome**

In progress callhomes may be interrupted by an upgrade.

Upgrading Protocol nodes in an ESS cluster


Spectrum Scale and ESS releases

5.0.1 release upgrade paths



Online upgrade supported from N-1 to N
Offline upgrade supported from any level to N

5.3.1 ESS upgrade paths



Online/Offline upgrade supported from N-2 or N-1 to N

Upgrading Protocol nodes in an ESS cluster

Recommendations

- Spectrum Scale 5.0.0.2 or later is recommended (*as of 4-11-2018*)
- If GPFS on the ESS will differ from GPFS on the Protocol nodes, make sure the versions are within the same release for the best compatibility.
- If at all possible, suggest performing the upgrade during a period of downtime with quiesced protocol data access.

Upgrading Protocol nodes in an ESS cluster

Mandatory pre-reqs

- All protocol nodes need a base OS repo to be accessible
- /etc/hosts on all nodes must be formatted correctly: IP FQDN alias
- Promptless SSH must be setup for all nodes to all nodes
- Do not extract or run the Install Toolkit on the EMS node
- The ESS GPFS level must be within N-1 of the desired protocol node levels
- Firewalls must be down completely OR all nodes must have the correct ports opened (*see the KC for details*)
- SELinux should settings should be checked. In some Scale levels, NFS needs a workaround to function with SELinux set at enforcing levels.
- Do not use XCAT / ESS scripts for upgrading the protocol nodes. They are unaware of the logic required to upgrade the nodes safely.
- Prepare the customer for an outage of Object and SMB. Expect NFS traffic to pause multiple times.

Upgrading Protocol nodes in an ESS cluster

Suggested flow using the Install Toolkit

- 1) Start with an active ESS cluster with healthy and active protocols
- 2) Extract the Protocols package to a Protocol node of your choice
- 3) `./spectrumscale setup -s <IP of this node> -st ess`
- 4) `./spectrumscale config populate -N <EMS node>`
 - If this fails, populate the toolkit with the protocol nodes, CES-IPs, cesSharedRoot location, manually
 - NSDs/Filesystems are not necessary
- 5) `./spectrumscale upgrade --precheck`
- 6) `./spectrumscale upgrade`
- 7) If OS/kernels/drivers/FW/OFED need to be upgraded, proceed with one protocol node at a time by first suspending the node, then stopping all services on the node, then shutting down GPFS on the node. Once done, make sure a reboot has occurred after any kernel change, run `mmbuildgpl`, and bring GPFS back online. Resume all protocols on the node. Repeat on each node. This can be done before or after the toolkit upgrade.

****This procedure can be used for non-protocol nodes as well**

Install GUI

(only for non-ESS environments)

Added Content – Install GUI *(only for non-ESS environments)*

Install GUI pre-requisites

Standard or Advanced Protocols package must be extracted on the node that will run the Install GUI

Port 9443 (https) or 9080 (http) should be opened on this node

When is the Install GUI available?

A user can start and use the Install GUI so long as no GPFS cluster already exists

Any limitations?

It can only be used if no GPFS cluster already exists

It will not configure NSDs, file systems, protocols, authentication

2 NSD nodes are required

1 GUI node is required

Added Content – Install GUI *(only for non-ESS environments)*

Install GUI flow

- User input of Cluster name, NTP server, # of nodes
- User input or import of node names
- User selection of Protocol nodes, NSD servers, GUI nodes
- Automatic or manual selection of Quorum/Manager nodes
- A pre-check, which is when the Install GUI creates the clusterdefinition.txt file and runs 'spectrumscale install –precheck'
- A pause after the pre-check to fix errors or continue with the Installation
- A full GPFS installation *(including addition of protocol related rpms for the future)*
- Necessitates user creates NSDs / file systems afterwards manually or with the Install Toolkit
- Necessitates user deploys protocols and authentication afterwards manually or with the Install Toolkit

Added Content – Install GUI *(only for non-ESS environments)*

1) Start the GUI

```
# ./spectrumscale installgui start
```

```
[ INFO ] Starting install toolkit Graphical User Interface on protocol-node1
```

```
[ INFO ] Shared GPFS JRE is already installed
```

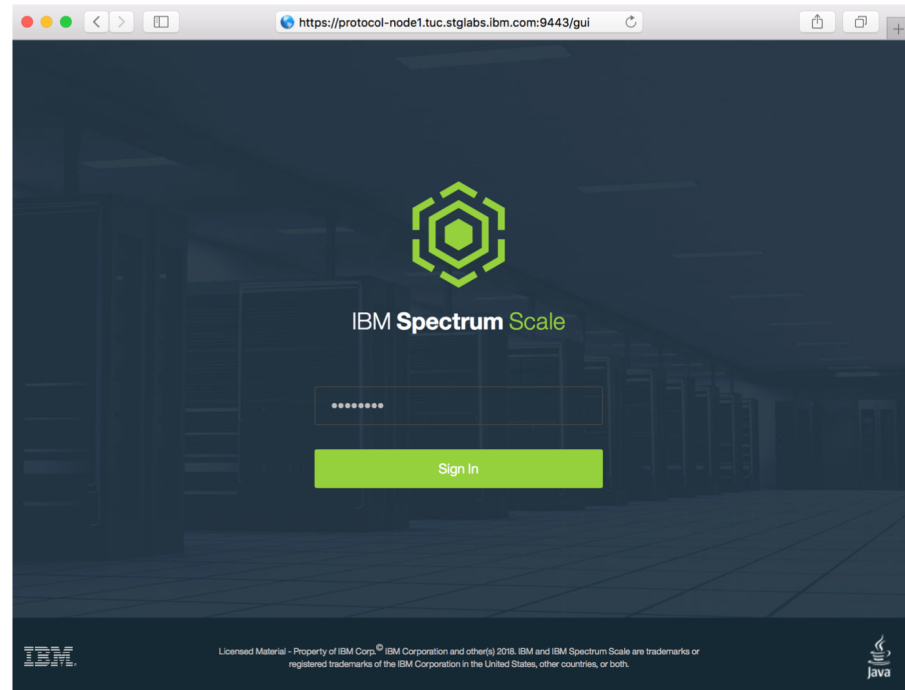
```
[ INFO ] Enter the following URL with the resolvable hostname or IP of protocol-node1 in your browser in order to launch the Install GUI:
```

```
[ INFO ] https://<Hostname or IP>:9443/gui
```

```
[ INFO ] Enter the following password:
```

```
[ INFO ] Passw0rd
```

2) Point your browser to the Install GUI node and login with the provided password



Added Content – Install GUI *(only for non-ESS environments)*



- Introduction
- Cluster Information
- Add Nodes
- Protocol Servers
- Storage Servers
- Graphical User Interface Servers
- Manager and Quorum Servers
- Summary
- Install
- Cluster Ready

Cluster Information

Cluster name:

democluster

NTP servers:

9.11.107.11,9.11.107.12

Number of nodes:

6

⬆
⬇
⬆

Added Content – Install GUI *(only for non-ESS environments)*

Introduction

Cluster Information

Add Nodes

Protocol Servers

Storage Servers

Graphical User Interface Servers

Manager and Quorum Servers

Summary

Install

Cluster Ready

Add Nodes

Specify the IP address or hostname of the nodes that should be part of the cluster. The IP address or hostname is used for communications within the cluster so the nodes must belong to the same network.

To specify IP address or hostname, you can upload a file that contains a comma separated list of IPs or hostnames,

Upload File

add a range of contiguous IPs,

From:

Hostname or IP address

To:

Hostname or IP address

or add IPs individually

protocol-node1

protocol-node2

client-node1

client-node2

nsd-node1

nsd-node2

Added Content – Install GUI *(only for non-ESS environments)*



- Introduction
- Cluster Information
- Add Nodes
- Protocol Servers
- Storage Servers
- Graphical User Interface Servers
- Manager and Quorum Servers
- Installation Network
- Summary
- Install
- Cluster Ready

Protocol Servers

If you are planning to use Object, NFS, or SMB protocols to host and share data, select protocol servers. Otherwise, skip this step.

Nodes	Protocol (2)
protocol-node1	<input checked="" type="checkbox"/>
protocol-node2	<input checked="" type="checkbox"/>
client-node1	<input type="checkbox"/>
client-node2	<input type="checkbox"/>
nsd-node1	<input type="checkbox"/>
nsd-node2	<input type="checkbox"/>

Use the Shift key to select a range of servers.
You need to assign different roles to the servers that are part of the IBM Spectrum Scale cluster to perform certain dedicated functions. [Learn more.](#)

Added Content – Install GUI *(only for non-ESS environments)*



- Introduction
- Cluster Information
- Add Nodes
- Protocol Servers
- Storage Servers
- Graphical User Interface Servers
- Manager and Quorum Servers
- Installation Network
- Summary
- Install
- Cluster Ready

Storage Servers

Physical storage devices are connected to and managed by the storage servers.

Nodes	Protocol (2)	Storage (2)
protocol-node1	✓	<input type="checkbox"/>
protocol-node2	✓	<input type="checkbox"/>
client-node1		<input type="checkbox"/>
client-node2		<input type="checkbox"/>
nsd-node1		<input checked="" type="checkbox"/>
nsd-node2		<input checked="" type="checkbox"/>

Use the Shift key to select a range of servers.
You need to assign different roles to the servers that are part of the IBM Spectrum Scale cluster to perform certain dedicated functions. [Learn more.](#)

Added Content – Install GUI *(only for non-ESS environments)*



- Introduction
- Cluster Information
- Add Nodes
- Protocol Servers
- Storage Servers
- Graphical User Interface Servers
- Manager and Quorum Servers
- Installation Network
- Summary
- Install
- Cluster Ready

Graphical User Interface Servers

Graphical User Interface (GUI) files will be installed on the GUI servers.

It is highly recommended to select at least two servers to maintain high availability. A maximum of three servers can be designated as the GUI servers in a cluster. A client node license is sufficient for a user interface server; a storage node license is not required.

Nodes	Protocol (2)	Storage (2)	GUI (2)
protocol-node1	✓		<input type="checkbox"/>
protocol-node2	✓		<input type="checkbox"/>
client-node1			<input checked="" type="checkbox"/>
client-node2			<input checked="" type="checkbox"/>
nsd-node1		✓	<input type="checkbox"/>
nsd-node2		✓	<input type="checkbox"/>

Use the Shift key to select a range of servers.
You need to assign different roles to the servers that are part of the IBM Spectrum Scale cluster to perform certain dedicated functions. [Learn more.](#)

Added Content – Install GUI *(only for non-ESS environments)*

Introduction

Cluster Information

Add Nodes

Protocol Servers

Storage Servers

Graphical User Interface Servers

Manager and Quorum Servers

Installation Network

Summary

Install

Cluster Ready

Manager and Quorum Servers

Manager servers manage file system configuration, disk space allocation, tokens that grant read and write access, and quotas. Quorum servers maintain data consistency in case of a node failure.

Manager and quorum servers are mandatory for the cluster. The installer automatically designates manager and quorum servers if you do not manually assign them.

Do you want to manually assign manager and quorum servers?

Yes


No

Select at least 1 Manager and 3 Quorum servers:

Nodes	Protocol (2)	Storage (2)	GUI (2)	Manager (2)	Quorum (3)
protocol-node1	✓			<input checked="" type="checkbox"/>	<input type="checkbox"/>
protocol-node2	✓			<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
client-node1			✓	<input type="checkbox"/>	<input type="checkbox"/>
client-node2			✓	<input type="checkbox"/>	<input type="checkbox"/>
nsd-node1		✓		<input type="checkbox"/>	<input checked="" type="checkbox"/>
nsd-node2		✓		<input type="checkbox"/>	<input checked="" type="checkbox"/>

Use the Shift key to select a range of servers.
You need to assign different roles to the servers that are part of the IBM Spectrum Scale cluster to perform certain dedicated functions. [Learn more.](#)

Added Content – Install GUI *(only for non-ESS environments)*



Introduction

Cluster Information

Add Nodes

Protocol Servers

Storage Servers

Graphical User Interface Servers

Manager and Quorum Servers

Installation Network

Summary

Install

Cluster Ready

Installation Network

Select the IP address of the installer node that can be accessed by all cluster nodes through passwordless SSH.

9.11.102.38

▼

Added Content – Install GUI *(only for non-ESS environments)*

Introduction

Cluster Information

Add Nodes

Protocol Servers

Storage Servers

Graphical User Interface Servers

Manager and Quorum Servers

Installation Network

Summary

Install

Cluster Ready

Summary

Verify the details before the appropriate files are installed on the specified nodes of the cluster.

Cluster name: democluster

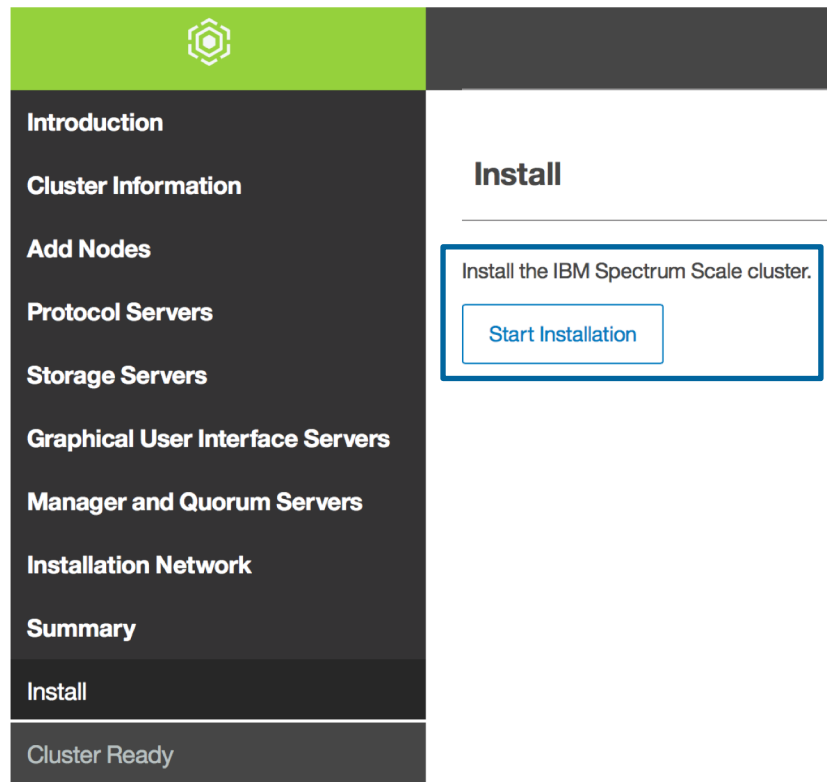
NTP servers: 9.11.107.11,9.11.107.12

Number of server node licenses required: 4

Number of client node licenses required: 2

Nodes	License Type	Protocol (2)	Storage (2)	GUI (2)	Manager (2)	Quorum (3)
protocol-node1	Server	✓			✓	
protocol-node2	Server	✓			✓	✓
client-node1	Client			✓		
client-node2	Client			✓		
nsd-node1	Server		✓			✓
nsd-node2	Server		✓			✓

Added Content – Install GUI *(only for non-ESS environments)*



What to expect:

- 1) 'Start Installation' will populate the `/usr/lpp/mmfs/5.0.1.0/installer/configuration/clusterdefinition.txt` file
- 2) A 'spectrumscale install –precheck' will be initiated
- 3) The precheck will either be successful or give a fatal error indicating issues to be resolved
- 4) 'Continue Installation' can then be clicked and the cluster will be installed
- 5) When finished, the Install GUI will stop and the user will be transitioned to the Spectrum Scale GUI
- 6) A basic GPFS cluster with no NSDs / FS will exist. Follow-up with another install & deploy

Added Content – Install GUI *(only for non-ESS environments)*

Yellow = warnings

Try to resolve if possible, but install can continue.

Red = fatal

Must be resolved before continuing

Pre-check Successful

Warnings have been found, but you can click Continue Installation as long as they are understood

The screenshot displays the IBM Spectrum Scale Install Toolkit GUI. On the left is a dark sidebar with a green header containing the IBM logo. The sidebar menu includes: Introduction, Cluster Information, Add Nodes, Protocol Servers, Storage Servers, Graphical User Interface Servers, Manager and Quorum Servers, Installation Network, Summary, Install, and Cluster Ready. The 'Install' option is currently selected. The main window has a dark header with a close button (X). Below the header, the title 'Warnings found' is displayed. A message states: 'Several warnings were found during the precheck. Please review them and if necessary, take the appropriate actions to fix them.' Below this message are three buttons: 'Continue Installation', 'Retry Installation', and 'Download log file'. The 'Continue Installation' button is highlighted. Below the buttons is a scrollable log window showing the pre-check output. The log includes several warning messages highlighted in yellow, such as 'NTP is not set to be configured with the install toolkit' and 'Ephemeral port range is not set'. At the bottom of the log window, there are buttons for 'Back' and 'Next'.

Warnings found

Several warnings were found during the precheck. Please review them and if necessary, take the appropriate actions to fix them. ⚠

[Continue Installation](#) [Retry Installation](#) [Download log file](#)

```

/usr/lpp/mmfs/5.0.1.0/installer/spectrumscale install --precheck
[ INFO ] Logging to file: /usr/lpp/mmfs/5.0.1.0/installer/logs/INSTALL-PRECHECK-12-05-2018 20:35:43.log
[ WARN ] NTP is not set to be configured with the install toolkit.See './spectrumscale config ntp -h' to setup.
[ WARN ] Install toolkit will not configure file audit logging as it has been disabled.
[ INFO ] Checking for knife bootstrap configuration...
[ INFO ] Performing Chef (deploy tool) checks.
[ INFO ] Performing GPFS checks.
[ INFO ] Running environment checks
[ INFO ] Installed license validation passed. The installed license on GPFS cluster matching with the installer license version.
[ INFO ] Checking pre-requisites for portability layer.
[ INFO ] GPFS precheck OK
[ INFO ] Performing Performance Monitoring checks.
[ INFO ] Running environment checks for Performance Monitoring
[ INFO ] Performing GUI checks.
[ INFO ] Network check from admin node client-node2.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from admin node client-node1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node1.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node2.tuc.stglabs.ibm.com to all other nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node2.tuc.stglabs.ibm.com to all other protocol nodes in the cluster passed
[ INFO ] Network check from protocol node protocol-node1.tuc.stglabs.ibm.com to all other protocol nodes in the cluster passed
[ WARN ] Ephemeral port range is not set. Please set valid ephemeral port range using the command ./spectrumscale config gpfs --ephemeral port range . You may set the default values as 60000-61000
[ INFO ] The install toolkit will not configure call home as it is disabled. To enable call home, use the following CLI command: ./spectrumscale callhome enable
[ INFO ] Pre-check successful for install.
[ INFO ] Tip : ./spectrumscale install
    
```

[Quit the Installer](#) [Back](#) [Next](#)

Added Content – Install GUI *(only for non-ESS environments)*

Install Success

- A cluster has now been created.
- The Install GUI will transition to the Spectrum Scale GUI
- Move back to the spectrumscale commands and re-run the install to add NSDs. Then follow-up with a deploy

The screenshot displays the IBM Spectrum Scale Install GUI. On the left is a dark sidebar with a green header containing the Spectrum Scale logo. The sidebar menu includes: Introduction, Cluster Information, Add Nodes, Protocol Servers, Storage Servers, Graphical User Interface Servers, Manager and Quorum Servers, Installation Network, Summary, Install, and Cluster Ready. The main content area has a dark header with a close button (X). Below the header, it says "6-Node Cluster Created Successfully". A message reads: "Congratulations! You have successfully installed IBM Spectrum Scale cluster democluster . [Download the installation log](#)".

Below the message is a section titled "GPFS cluster information" with the following details:

```

GPFS cluster name: democluster.tuc.stglabs.ibm.com
GPFS cluster id: 2254960516004008510
GPFS UID domain: democluster.tuc.stglabs.ibm.com
Remote shell command: /usr/bin/ssh
Remote file copy command: /usr/bin/scp
Repository type: CCR
    
```

At the bottom is a table with 5 columns: Node, Daemon node name, IP address, Admin node name, and Designation.

Node	Daemon node name	IP address	Admin node name	Designation
1	protocol-node1.tuc.stglabs.ibm.com	9.11.102.38	protocol-node1.tuc.stglabs.ibm.com	quorum-perfmon
2	nsd-node1.tuc.stglabs.ibm.com	9.11.102.62	nsd-node1.tuc.stglabs.ibm.com	quorum-manager-perfmon
3	nsd-node2.tuc.stglabs.ibm.com	9.11.102.66	nsd-node2.tuc.stglabs.ibm.com	quorum-manager-perfmon
4	protocol-node2.tuc.stglabs.ibm.com	9.11.102.55	protocol-node2.tuc.stglabs.ibm.com	perfmon
6	client-node2.tuc.stglabs.ibm.com	9.11.102.69	client-node2.tuc.stglabs.ibm.com	perfmon
7	client-node1.tuc.stglabs.ibm.com	9.11.102.67	client-node1.tuc.stglabs.ibm.com	perfmon

You are NOT done yet

Added Content – Install GUI *(only for non-ESS environments)*

```
# ./spectrumscale node list
[ INFO ] List of nodes in current configuration:
[ INFO ] [Installer Node]
[ INFO ] 9.11.102.61
[ INFO ]
[ INFO ] Setup Type: SpectrumScale
[ INFO ]
[ INFO ] [Cluster Name]
[ INFO ] democloud.tuc.stglabs.ibm.com
[ INFO ]
[ INFO ] [Protocols]
[ INFO ] Object : Disabled
[ INFO ] SMB : Disabled
[ INFO ] NFS : Disabled
[ INFO ]
[ INFO ] File Audit logging : Disabled
[ INFO ]
[ INFO ] GFS Node
```

	Admin	Quorum	Manager	NSD Server	Protocol	GUI Server	Callhome Node	OS	Arch
client-node1.tuc.stglabs.ibm.com	X					X		ubuntu16	x86_64
client-node2.tuc.stglabs.ibm.com	X					X	X	sles12	x86_64
nsd-node1.tuc.stglabs.ibm.com		X		X				rhel7	x86_64
nsd-node2.tuc.stglabs.ibm.com		X		X				rhel7	x86_64
protocol-node1.tuc.stglabs.ibm.com	X	X	X		X			rhel7	x86_64
protocol-node2.tuc.stglabs.ibm.com			X		X			rhel7	x86_64

```

[ INFO ]
[ INFO ] [Export IP address]
[ INFO ] No export IP addresses configured

```

Example showing how the Install GUI pre-populated the clusterdefinition.txt file

/usr/lpp/mmfs/5.0.1.0/installer/configuration/clusterdefinition.txt

Added Content – Install GUI *(only for non-ESS environments)*

```
# mmlscluster
```

```
GPFS cluster information
```

```
=====
```

```

GPFS cluster name:      democcluster.tuc.stglabs.ibm.com
GPFS cluster id:       2254960516004008510
GPFS UID domain:       democcluster.tuc.stglabs.ibm.com
Remote shell command:  /usr/bin/ssh
Remote file copy command: /usr/bin/scp
Repository type:       CCR

```

Node	Daemon node name	IP address	Admin node name	Designation
1	protocol-node1.tuc.stglabs.ibm.com	9.11.102.38	protocol-node1.tuc.stglabs.ibm.com	quorum-perfmon
2	nsd-node1.tuc.stglabs.ibm.com	9.11.102.62	nsd-node1.tuc.stglabs.ibm.com	quorum-manager-perfmon
3	nsd-node2.tuc.stglabs.ibm.com	9.11.102.66	nsd-node2.tuc.stglabs.ibm.com	quorum-manager-perfmon
4	protocol-node2.tuc.stglabs.ibm.com	9.11.102.55	protocol-node2.tuc.stglabs.ibm.com	perfmon
6	client-node2.tuc.stglabs.ibm.com	9.11.102.69	client-node2.tuc.stglabs.ibm.com	perfmon
7	client-node1.tuc.stglabs.ibm.com	9.11.102.67	client-node1.tuc.stglabs.ibm.com	perfmon

mmlscluster output after Install GUI has installed the cluster

```

# /usr/lpp/mmfs/bin/mmlsnd
mmlsnd: [I] No disks were found.

```

Notice that no NSDs have been created – this is a limitation of the Install GUI

Added Content – Install GUI *(only for non-ESS environments)*

How to add NSDs using the Install Toolkit after a successful Install GUI installation

```
# ./spectrumscale nsd add -p nsd-node1 -s nsd-node2 -u dataAndMetadata -fs cesSharedRoot -fg 1 "/dev/sdb"
```

```
[ INFO ] Connecting to nsd-node1.tuc.stglabs.ibm.com to check devices and expand wildcards.
```

```
[ INFO ] Looking up details of /dev/sdb.
```

```
[ INFO ] The installer will create the new file system cesSharedRoot if it does not exist.
```

```
[ INFO ] Adding NSD None on nsd-node1.tuc.stglabs.ibm.com using device /dev/sdb.
```

```
[ INFO ] Configuration updated
```

```
[ INFO ] Tip : If all node designations and any required protocol configurations are complete, proceed to check the installation configuration: ./spectrumscale install --precheck
```

```
# ./spectrumscale nsd add -p nsd-node2 -s nsd-node1 -u dataAndMetadata -fs cesSharedRoot -fg 2 "/dev/sdc"
```

```
[ INFO ] Connecting to nsd-node2.tuc.stglabs.ibm.com to check devices and expand wildcards.
```

```
[ INFO ] Looking up details of /dev/sdc.
```

```
[ INFO ] Adding NSD None on nsd-node2.tuc.stglabs.ibm.com using device /dev/sdc.
```

```
[ INFO ] Configuration updated
```

```
[ INFO ] Tip : If all node designations and any required protocol configurations are complete, proceed to check the installation configuration: ./spectrumscale install --precheck
```

```
# ./spectrumscale nsd add -p nsd-node1 -s nsd-node2 -u dataAndMetadata -fs ObjectFS -fg 1 "/dev/sdd"
```

```
# ./spectrumscale nsd add -p nsd-node2 -s nsd-node1 -u dataAndMetadata -fs ObjectFS -fg 2 "/dev/sde"
```

```
# ./spectrumscale nsd add -p nsd-node1 -s nsd-node2 -u dataAndMetadata -fs fs1 -fg 1 "/dev/sdf"
```

```
# ./spectrumscale nsd add -p nsd-node1 -s nsd-node2 -u dataAndMetadata -fs fs1 -fg 1 "/dev/sdg"
```

```
# ./spectrumscale nsd add -p nsd-node2 -s nsd-node1 -u dataAndMetadata -fs fs1 -fg 2 "/dev/sdh"
```

```
# ./spectrumscale nsd add -p nsd-node2 -s nsd-node1 -u dataAndMetadata -fs fs1 -fg 2 "/dev/sdi"
```

List the NSDs

```
# ./spectrumscale nsd list
```

List the File systems

```
# ./spectrumscale filesystem list
```

Start the install

```
# ./spectrumscale install
```

Added Content – Install GUI *(only for non-ESS environments)*

- You will have added NSDs to your prior install once the install phase completes successfully
- At this point you will have a complete GPFS cluster but no file systems nor protocols.
- Proceed to the deploy phase to create file systems, install/enable protocols, and configure authentication.

THE END