

# Spectrum Scale Expert Talks

Episode 15:

## **IBM Spectrum Scale Container Native Storage Access**




**Show notes:**

[www.spectrumscaleug.org/experttalks](http://www.spectrumscaleug.org/experttalks)

**Join our conversation:**

[www.spectrumscaleug.org/join](http://www.spectrumscaleug.org/join)




# SSUG::Digital

**Welcome to digital events!**



**Show notes:**  
[www.spectrumscaleug.org/experttalks](http://www.spectrumscaleug.org/experttalks)


**Join our conversation:**  
[www.spectrumscaleug.org/join](http://www.spectrumscaleug.org/join)




# About the user group

- Independent, work with IBM to develop events
- Not a replacement for PMR!
- Email and Slack community
- <https://www.spectrumscaleug.org/join>

#SSUG






# We are ...

- Simon Thompson (UK)
- Kristy Kallback-Rose (USA)
- Bob Oesterlin (USA)
- Bill Anderson (USA)
- Chris Schipalius (Australia)




#SSUG






Check <https://www.spectrumscaleug.org/experttalks>  
for charts, show notes and upcoming talks


- Past talks:
  - 001: What is new in Spectrum Scale 5.0.5?
  - 002: Best practices for building a stretched cluster
  - 003: Strategy update
  - 004: Update on performance enhancements in Spectrum Scale (file create, MMAP, direct IO, ESS 5000)
  - 005: Update on functional enhancements in Spectrum Scale (inode management, vCPU scaling, NUMA considerations)
  - 006: Persistent Storage for Kubernetes and OpenShift environments
  - 007: Manage the lifecycle of your files using the policy engine
  - 008: Multi-node scaling of AI workloads using Nvidia DGX, OpenShift and Spectrum Scale
  - 009: Continental: Deep Thought – An AI Project for Autonomous Driving Development
  - 010: Data Accelerator for Analytics and AI (DAAA)
  - 011: What is new in Spectrum Scale 5.1.0?
  - 012: Lenovo - Spectrum Scale and NVMe Storage
  - 013: Event driven data management and security using Spectrum Scale Clustered Watch Folder and File Audit Logging
  - 014: What is new in Spectrum Scale 5.1.1?
- Today:
  - **June 29: IBM Spectrum Scale Container Native Storage Access**





# Speakers

- Laura O. Braverman (IBM)
- Harald Seipp (IBM)



# IBM Spectrum Scale

## Container Native Storage Access (CNSA)

Laura O. Braverman  
Senior Software Engineer  
IBM Spectrum Scale Development



# Disclaimer

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code, or functionality. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

IBM reserves the right to change product specifications and offerings at any time without notice. This publication could include technical inaccuracies or typographical errors. References herein to IBM products and services do not imply that IBM intends to make them available in all countries.

# Previous Expert Talks related to containerization

## Why to containerize Spectrum Scale?


- Episode 3: Spectrum Scale Strategy Update
- <https://www.spectrumscaleug.org/event/ssugdigital-spectrum-scale-expert-talk-strategy-update/>

## How to access Spectrum Scale from your containerized applications?

- Episode 6: Persistent Storage for Kubernetes and OpenShift environments
- <https://www.spectrumscaleug.org/event/ssugdigital-persistent-storage-for-containers-with-spectrum-scale/>

## Use case

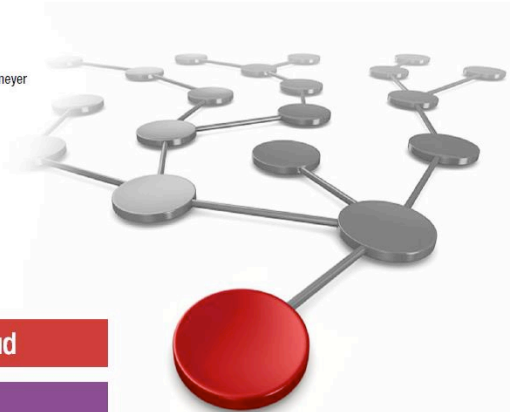
- Episode 8: Scalable multi-node training for AI workloads on NVIDIA DGX, Red Hat OpenShift and IBM Spectrum Scale
- <https://www.spectrumscaleug.org/event/ssugdigital-multi-node-training-for-ai-workloads/>



Draft Document for Review September 21, 2020 10:23 am REDP-5610-00

### Deployment and Usage Guide for Running AI Workloads on Red Hat OpenShift and NVIDIA DGX Systems with IBM Spectrum Scale

Simon Lorenz  
Gero Schmidt  
Thomas Schoenemeyer



Cloud

Storage

IBM. Redpaper

The image shows the cover of a Redbooks document. At the top right is the Redbooks logo with the URL ibm.com/redbooks. Below it is the document title: 'Deployment and Usage Guide for Running AI Workloads on Red Hat OpenShift and NVIDIA DGX Systems with IBM Spectrum Scale'. The authors listed are Simon Lorenz, Gero Schmidt, and Thomas Schoenemeyer. The central graphic is a network diagram with several grey nodes connected by lines, and one prominent red node at the bottom right. Below the diagram are two horizontal bars: a red one with a cloud icon and the word 'Cloud', and a purple one with the word 'Storage'. At the bottom of the cover are the IBM logo and the word 'Redpaper'.

# Contents

## **Overview**

What is IBM Spectrum Scale Container Native Storage access?

Why would I want CNSA?

Spectrum Scale Containerization Journey

Cluster Overview

CNSA/CSI Compatibility Matrix

How does Operator simplify things?

## **What's new in CNSA v5.1.1.1**

Notable changes in 5.1.1.1

Goodbye, Fix Central

IBM Cloud Container Registry – Get Entitlement Key

Public CNSA GitHub Repository

**A short demo...**

**Looking forward**

# What is IBM Spectrum Scale Container Native Storage access?

IBM Spectrum Scale Container Native Storage Access (CNSA) supports the Red Hat OpenShift Container Platform with a fully containerized deployment.

This allows deployment of IBM Spectrum Scale on Red Hat CoreOS worker nodes where classic packages cannot be installed.

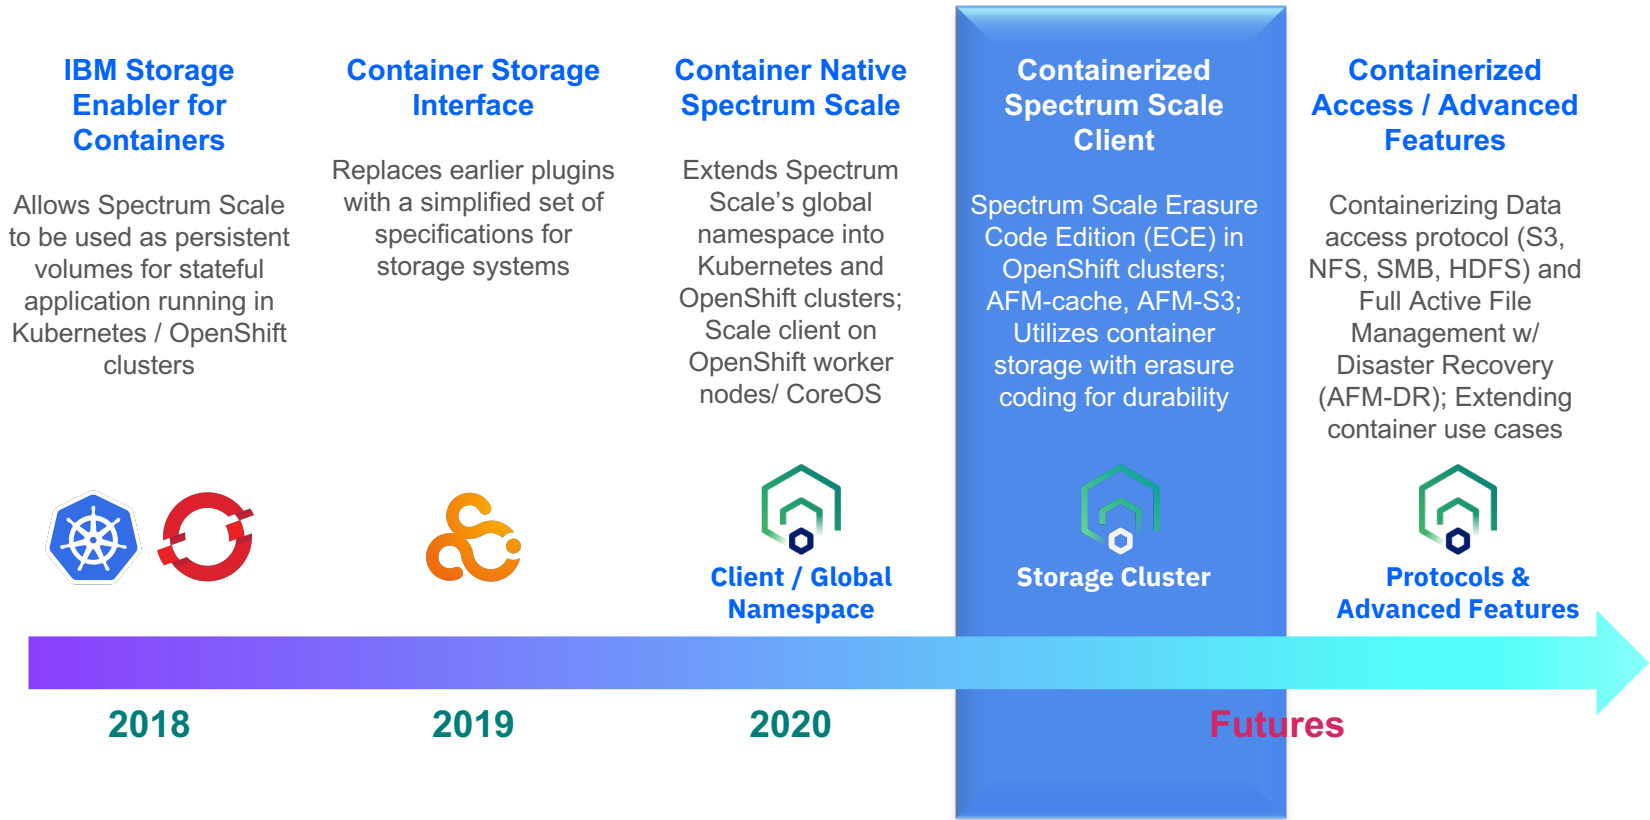


# Why would I want CNSA?

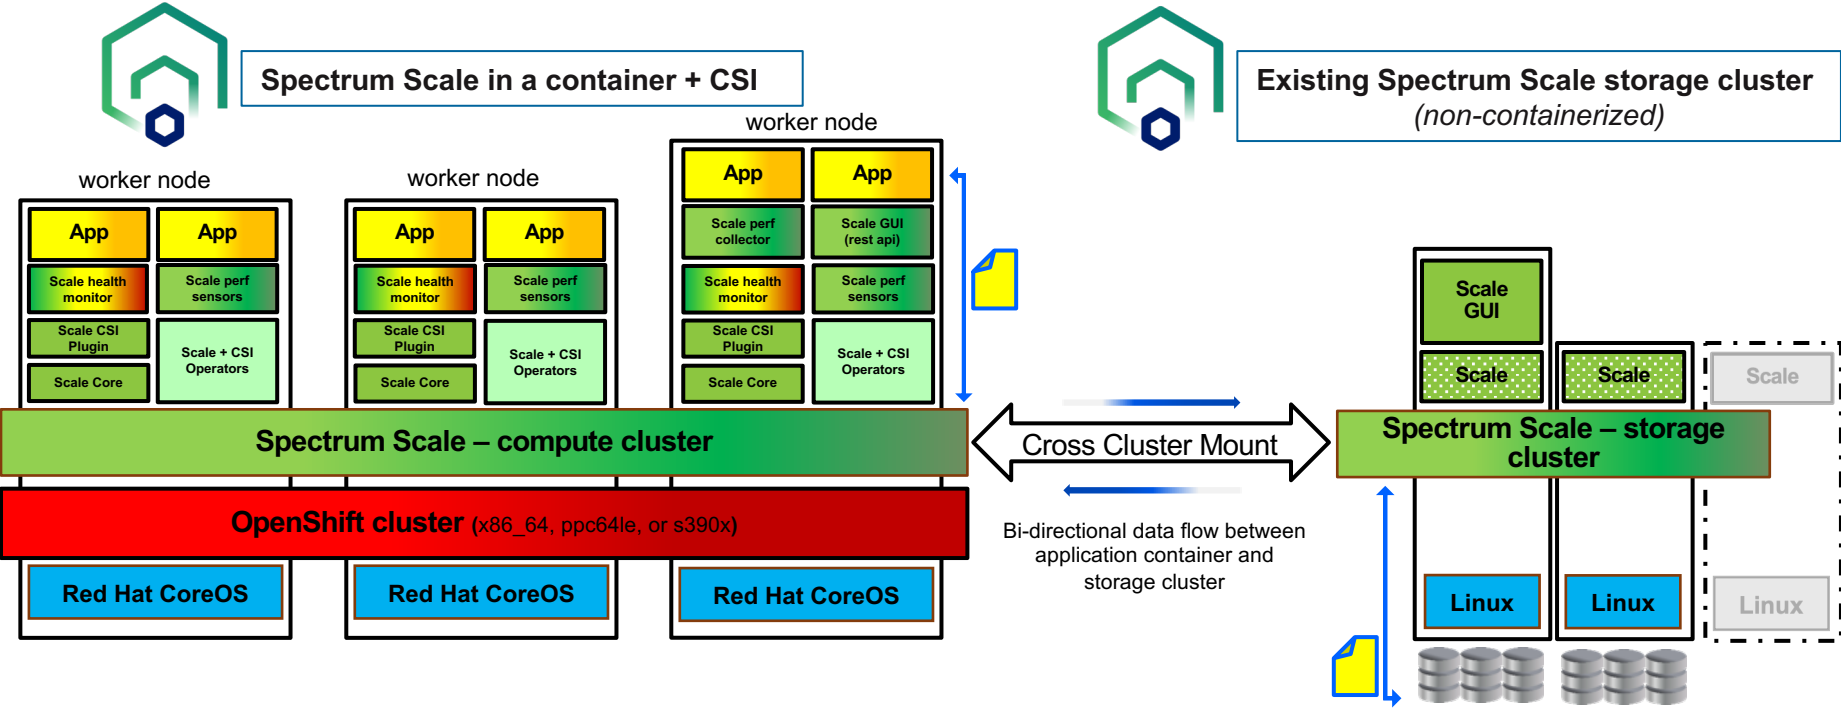
- Brings scalability and performance of Spectrum Scale to your containerized workloads
- Integrates OpenShift in your data intensive workflows
- Greatly simplifies the end management of the Spectrum Scale cluster by use of our operator and its controllers
- Supports architectures X/P/Z
- **Note:** Try it out! The DAE or DME version of Spectrum Scale Container Native storage access is available to current users (free of charge)



# Spectrum Scale Containerization Journey



# Cluster Overview





# CNSA/CSI Compatibility Matrix



GA	Version	Architecture	Scale level	OCP level	K8s level	RHEL	Ubuntu
12/10/2019	<b>CSI 1.0.0</b>	x86	5.0.4.1+	4.2	1.13, 1.14 1.15. 1.16. 1.17	7.5, 7.6, 7.7	N/A
3/25/2020	<b>CSI 1.1.0</b>	x86	5.0.4.1+	4.2, 4.3,	1.14, 1.15, 1.16, 1.17	7.5, 7.6, 7.7	N/A
6/30/2020	<b>CSI 2.0.0</b>	x86	5.0.4.1+	4.3, 4.4, 4.5	1.16, 1.17, 1.18	7.5, 7.6, 7.7	18.04 via RPQ
11/13/2020	<b>CSI 2.1.0</b>	x86, ppc64le	5.0.4.1+	4.5, 4.6	1.18, 1.19	7.7, 7.8	20.04 (req 5.1.0.x)
12/10/2020	<b>CNSA 5.1.0-1 + CSI 2.1.0</b>	x86, ppc64le, s390	5.1.0-1	4.5, 4.6	NA	N/A	N/A
3/18/2021	<b>CNSA 5.1.0-3 + CSI 2.1.0</b>	x86, ppc64le, s390	5.1.0-3	4.5, 4.6	NA	N/A	N/A
4/23/2021	<b>CSI 2.2.0</b>	x86, ppc64le	5.1.0-1+	4.6, 4.7	1.19, 1.20	7.7, 7.8, 7.9	20.04 (x86)
4/23/2021	<b>CNSA 5.1.0-3 + CSI 2.2.0</b>	x86, ppc64le, s390	5.1.0-3	4.6, 4.7	NA	N/A	N/A
6/21/2021	<b>CNSA 5.1.1.1 + CSI 2.2.0</b>	x86, ppc64le, s390	5.1.1.1	4.6, 4.7	NA	N/A	N/A

*\*CSI and CNSA customers are expected to keep up to date for the latest fixes*

[Refer to section 17 of the Spectrum Scale FAQ for current interop tables](#)

# How does Operator simplify things?

- Operator ensures state via its controllers and their control (reconcile) loops.
- Each controller has a declarative state, the *desired* state it aims to achieve

## Remote Mount Controller

Uses the information defined by the user in the ScaleCluster Custom Resource file to mount an IBM Spectrum Scale remote filesystem to the CNSA cluster

This is done using the GUI API on both the local CNSA GUI and Remote GUI

## Call Home Controller

Responsible for configuring call home

Monitors the CNSA cluster and sends alerts and events to IBM support

## ScaleCluster Controller

Creates all the resources necessary to deploy the Scale, GUI, and pmcollector pods

# What's new in CNSA v5.1.1.1

Support for upgrading IBM Spectrum Scale Container Native Storage Access (CNSA) to v5.1.1.1.

IBM Spectrum Scale Container Native Storage Access (DAE, DME) images are now hosted on the entitled IBM Cloud Container Registry.

GitHub will be used to host the YAML manifest files required for IBM Spectrum Scale Container Native Storage Access installation, replacing FixCentral.

Support for a Spectrum Scale specific must-gather image that facilitates the collection of a gpfs.snap and relevant Kubernetes diagnostic logs.

## Notable Changes in 5.1.1.1

### Where are all the images in the ScaleCluster Custom Resource?

That's the operator's job now.

Based on the edition (DME, DAE) selected in the license field of the ScaleCluster Custom Resource, the respective IBM Cloud Container Registry images will be used for deployment.

### What happened to the OpenShift Internal Image Registry steps?

Now that our images are hosted by IBM Cloud Container Registry, there is no need to configure an OpenShift Internal Image Registry.

### What happened to all the yaml files for deployment?

We've condensed them into a single yaml file.





Support Downloads ▾ Documentation ▾ Forums Cases ▾ Monitoring ▾ Manage support account ▾

### Container Native

Filter fix details:

	▲	🔍 Description	📅 Release date
<input type="checkbox"/>	1	<p>🔑 fix pack: → <a href="#">Spectrum Scale Container Native Storage Access-5.1.1.1-x86_64_has_moved_to_the_IBM_Image_Registry</a></p> <p>⚠️ <b>Notice:</b> Spectrum Scale Container Native Storage Access (CNSA) is moving from Fix Central to the IBM Image Registry located at: <a href="https://myibm.ibm.com/products-services/containerlibrary">https://myibm.ibm.com/products-services/containerlibrary</a> Please refer to the Readme below for more information.</p> <p>📄 <a href="#">Link to IBM Image Registry</a>      📄 <a href="#">Readme</a></p>	2021/06/22



# IBM Cloud Container Registry

## Get entitlement key

<https://myibm.ibm.com/products-services/containerlibrary>



Log into the IBM Container software library with the IBMid and password that is associated with the entitled software

Use your IBMid to login to MyIBM. However, your username to access cp.icr.io images will be **cp**



Click  
*Get entitlement key*



Save the key to a secure location. This key will be used to configure your IBM Spectrum Scale Container Native storage access deployment

After receiving an entitlement key, follow the IBM Spectrum Scale Container Native documentation for the release you are wanting to install or upgrade

<https://www.ibm.com/docs/en/scalecontainernative>



# Public CNSA GitHub Repository

## *Location and Contents*

- <https://github.com/IBM/ibm-spectrum-scale-container-native>
- Contains manifests needed to deploy CNSA and MCO files needed to configure OpenShift for CNSA.
- Fix Central no longer will host CNSA tar archives containing these manifests and MCO files.
- End users need not directly know about this repo. CNSA documentation utilizes links to files in this repo to facilitate installation.



```
oc create -f https://raw.githubusercontent.com/IBM/ibm-spectrum-scale-container-native/v5.1.1/generated/installer/ibm-spectrum-scale-operator.yaml -n ibm-spectrum-scale-ns
```










*Example inline GitHub link for operator deployment*

# Public CNSA GitHub Repository

## Location and Contents

v5.1.1.1 1 branch 1 tag Code

 **lostasi** Delete MCO readme files 2fccccf 7 days ago  17 commits

 base	Operator digest update part 3	9 days ago
 generated	Operator digest update part 3	9 days ago
 machineconfig	Delete MCO readme files	7 days ago
 CODE_OF_CONDUCT.md	Create CODE_OF_CONDUCT.md	4 months ago
 CONTRIBUTING.md	Create CONTRIBUTING.md	4 months ago
 LICENSE	Initial commit	4 months ago
 README.md	Flesh out README.md	9 days ago
 ibm_v1_scalecluster_cr.yaml	Updates to sample ScaleCluster CR	7 days ago
 kustomization.yaml	Initial Spectrum Scale container native 5.1.1.1 manifest population	17 days ago

# A short demo...

- Get our entitlement key from MyIBM for IBM Cloud Container Registry
- Update global pull secret
- Create CA Cert
- Create Storage Cluster GUI secret
- Deploy operator using IBM Cloud Container Registry image
- Deploy CNSA using IBM Cloud Container Registry images



# Looking forward

## Containerized Storage

- AFM
- ECE

## Containerization of existing Spectrum Scale functionality

- Encryption
- Multiple filesystem, multiple remote cluster support
- Protocols (S3, SMB, NFS, HDFS)


## Vanilla k8s support

## Support for Public Clouds

## Support for Edge Computing

## Better User Experience

- Integrated deployment of CNSA and CSI
- Operator Lifecycle Manager
  - Simplified deployment of CNSA operator
- Rolling Upgrade



**Tell us  
what  
you'd  
like!**

# IBM Spectrum Scale

## CNSA Field Lessons Learned

Harald Seipp  
Senior Technical Staff Member  
IBM Systems EMEA Storage Competence Center



# Agenda

- Architecture/Planning

- Accessing vs. owning cluster
- Resilience
- Versions
- Firewalls
- Infra nodes / Control Plane nodes
- Image Registry
- DNS

- Installation

- Firewalls
- File System Naming
- Accessing vs Owning cluster in k8s resources

- Operation

- Namespaces
- Default Storage Class
- Small PVs

- Troubleshooting

- CNSA Operator and ScaleCluster resource
- CSI Operator and scale-csi pod
- Log collection

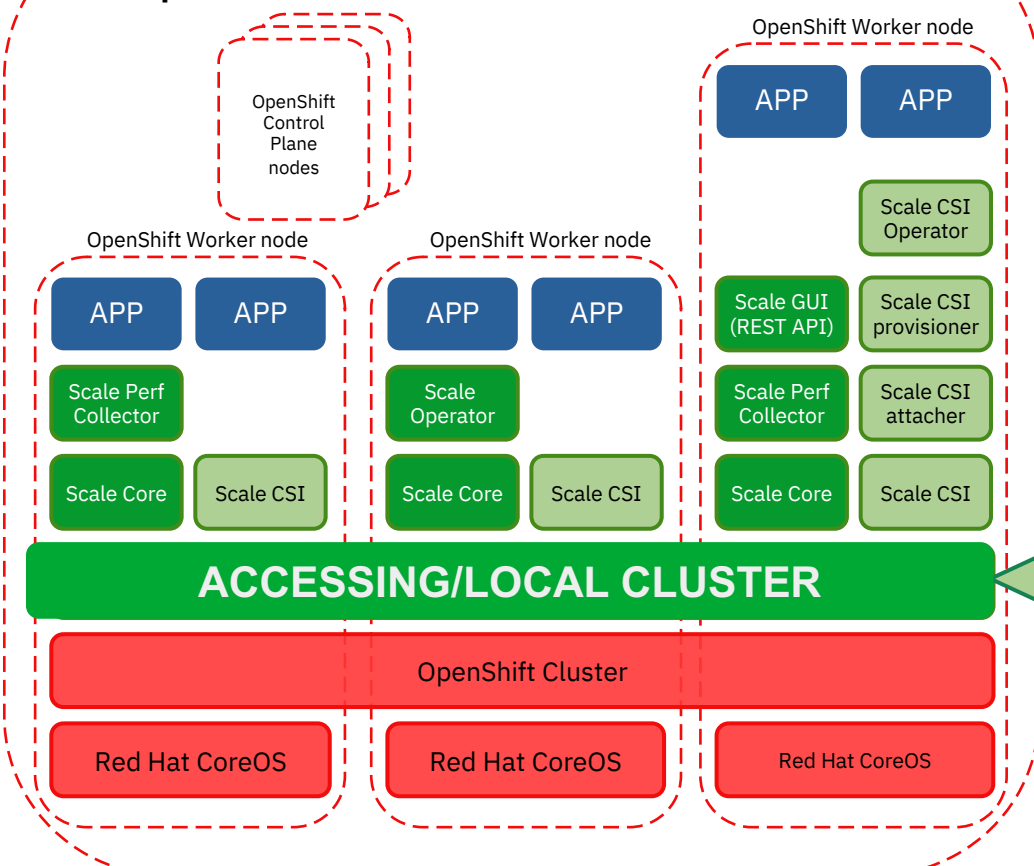


# Architecture Planning


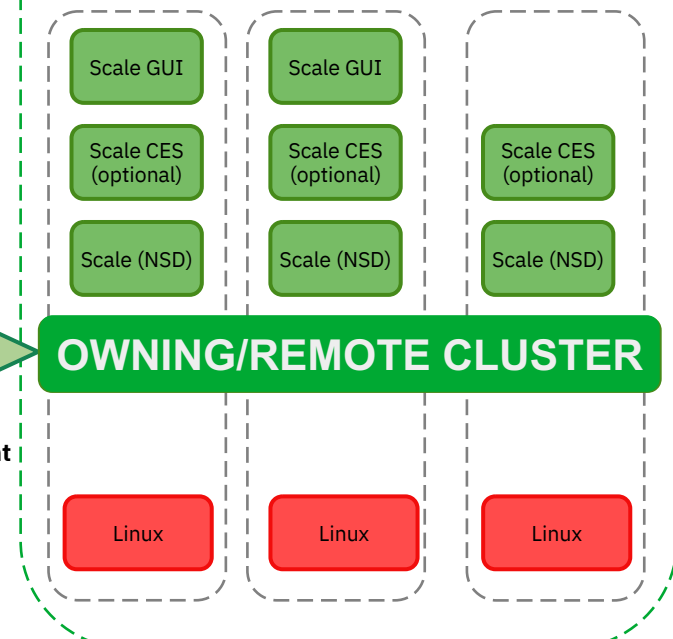


# Spectrum Scale Cluster Nomenclature

## Spectrum Scale Container Native Storage Access (CNSA) on OpenShift

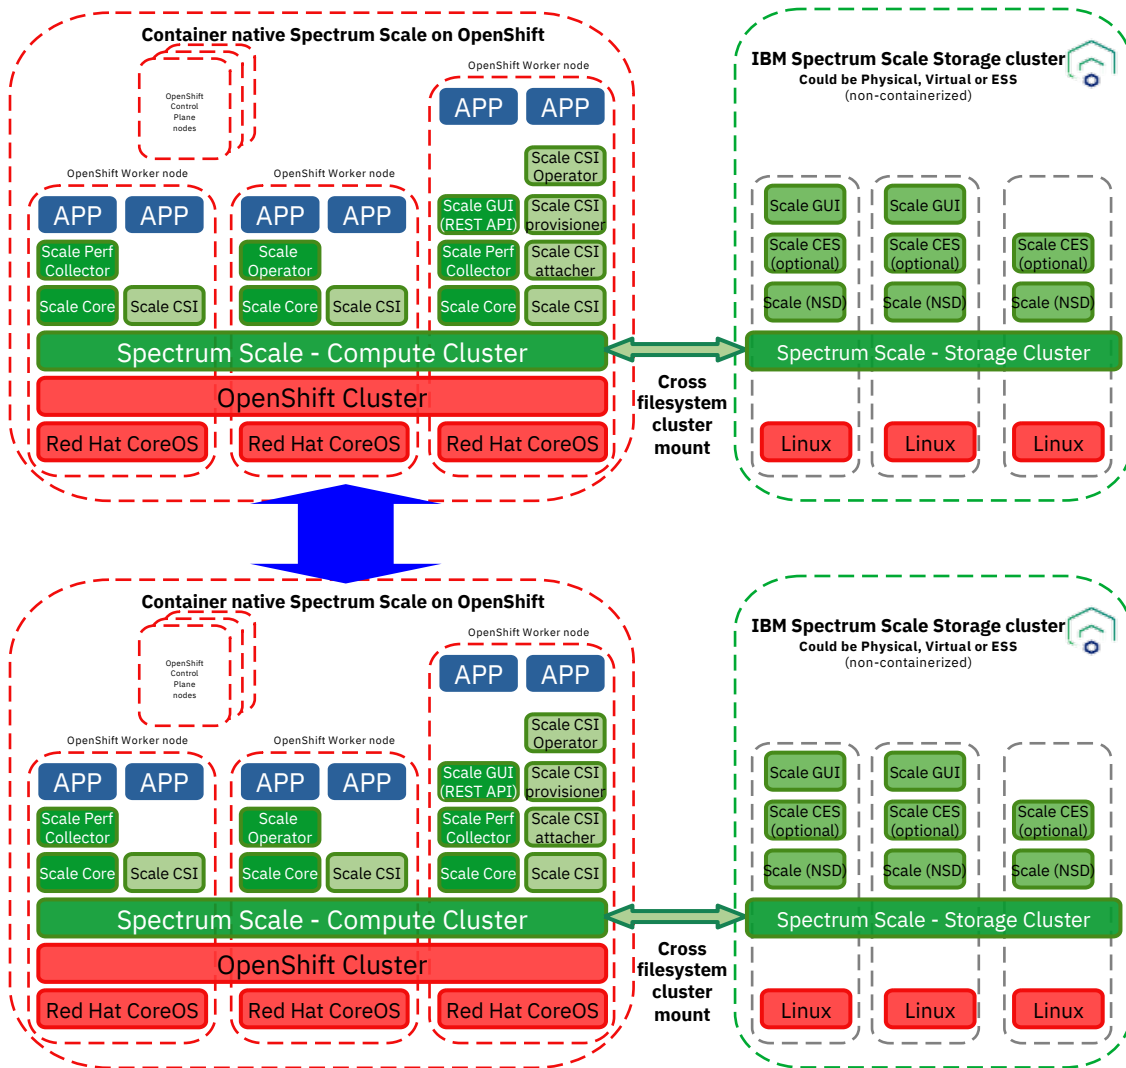


## IBM Spectrum Scale Storage cluster Could be Physical, Virtual or ESS (non-containerized)



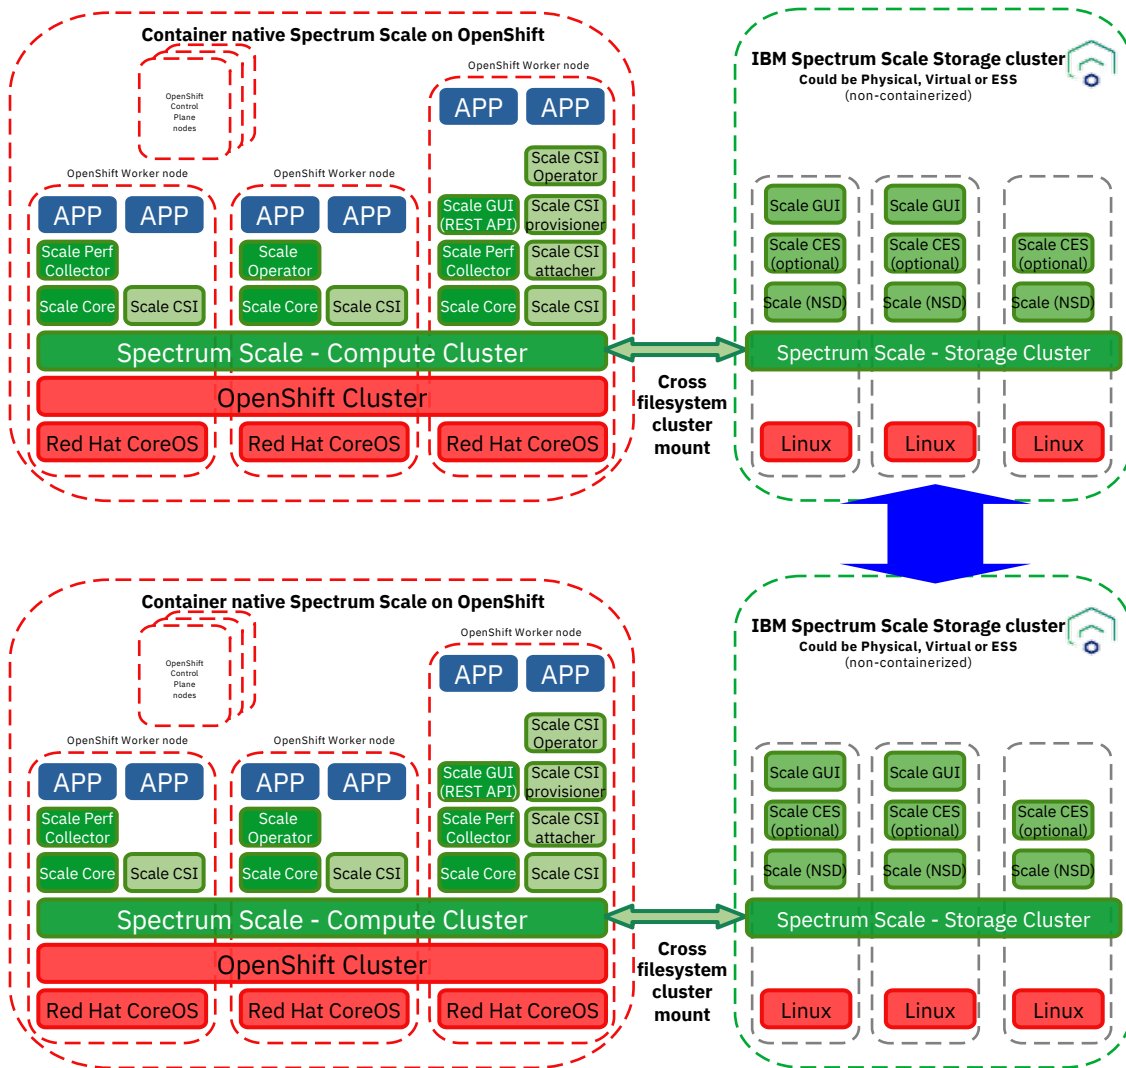
# Resilience considerations – Application-level replication

- Data replication is handled on application level
  - The stateful application replicates the data between sites
- Storage is separate on each of the sites



# Resilience considerations – Storage replication

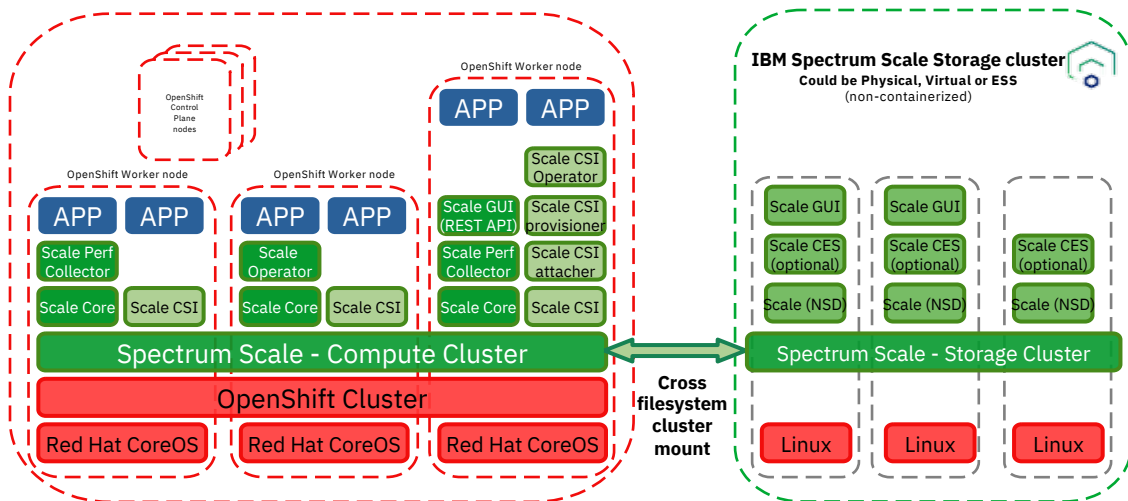
- Storage is replicated between the sites
  - Stretched cluster with near-zero RPO
  - AFM for long distance
- Application might handle failover
  - using a shared lock on data
- Alternatively manual failover
  - Re-create PV definitions and application context on secondary site
  - Scale up application at failover
- Future: automation



# Versions

- CNSA is at 5.1.0.3, CSI at 2.2.0, the Storage Cluster is at 5.1.1.x
- Volume Creation from Snapshot does not work?
  - Do not forget `mmchconfig release=LATEST` on the storage cluster

IBM Spectrum Scale container native	OpenShift Container Platform	IBM Spectrum Scale Container Storage Interface	IBM Spectrum Scale remote cluster
5.1.0.1	4.5, 4.6	2.1.0	5.1.0.1+
5.1.0.3	4.5, 4.6	2.1.0	5.1.0.1+
5.1.0.3	4.6, 4.7	2.2.0	5.1.1.x
5.1.1.1	4.6, 4.7	2.2.0	5.1.1.x

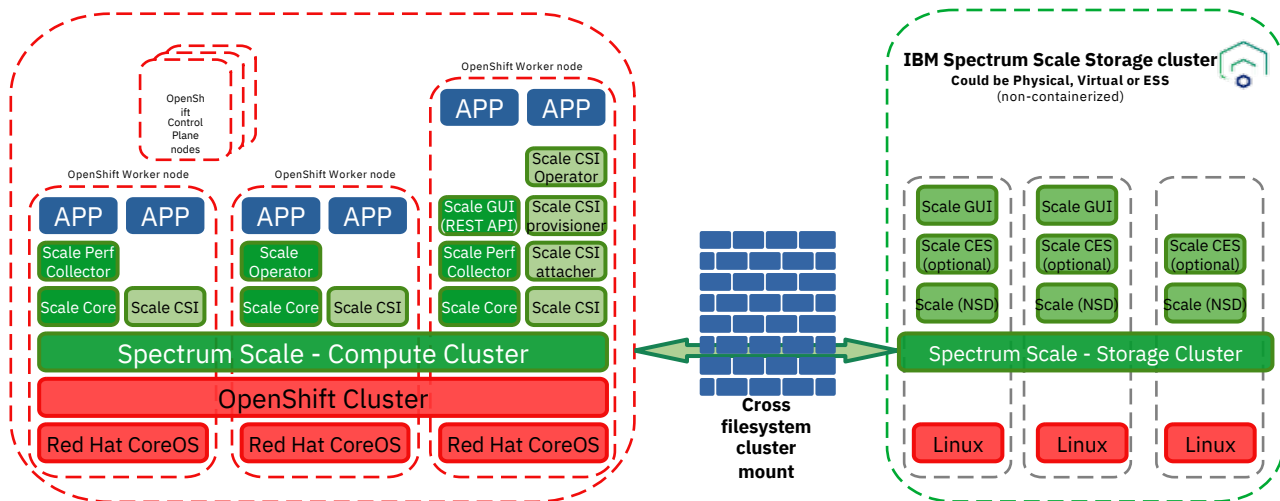


# Firewalls 1/2

- Firewall in-between  
Accessing and Owning cluster needs thorough planning
  - Have all required ports opened in both directions
  - Set the ephemeral port range on both clusters
  - Set ephemeral ports to the same range
- Symptom when not configured correctly:  
`mmismount`  
command on the CNSA cluster fails (timeout) although mount is present

Table 1. Recommended port numbers that can be used for internal communication

Port Number	Protocol	Service Name	Components that are involved in communication
1191	TCP	GPFS	Intra-cluster
22	TCP	Remote shell command, such as SSH.	Commands
22	TCP	Remote file copy command, such as SCP.	Commands
---	ICMP	ICMP ECHO (ping).	Intra-cluster
User-selected range	TCP	GPFS ephemeral port range	Intra-cluster



# Firewalls 2/2

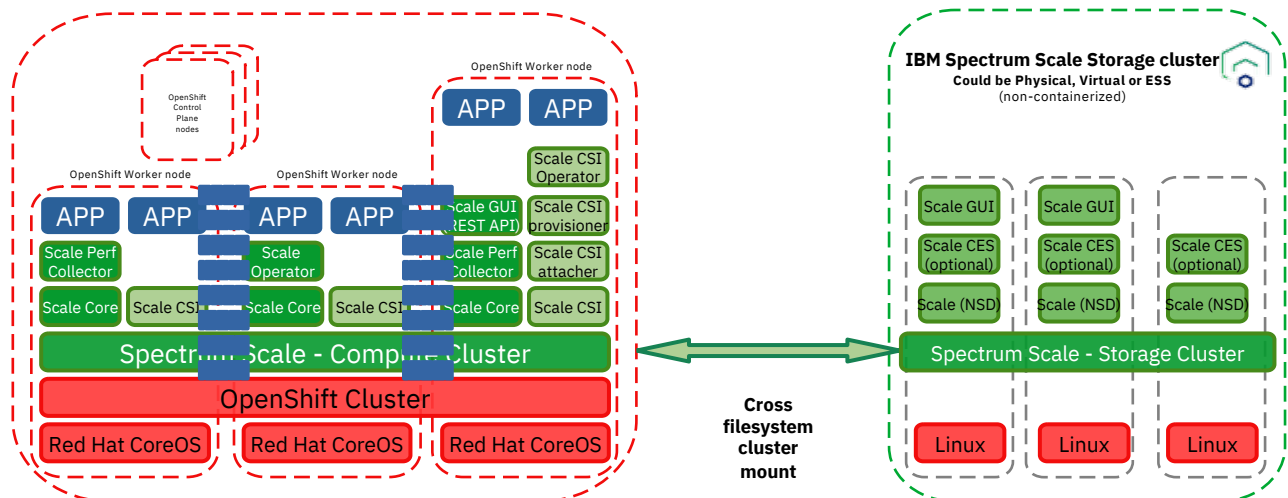
- Firewall in-between cluster nodes
  - Seen on an RHOSP-based OpenShift installation
- Watch out for port 12345!
- For additional ports see [documentation](#)

Table 1. Firewall recommendations for GUI

Port Number	Functions
47080	Management GUI
47443	Management GUI
80	Management GUI IBM Spectrum Scale ma
443	Management GUI IBM Spectrum Scale ma
4444	Management GUI
4739	Performance monitorin
9085 and 9084	Performance monitorin

Table 1. Recommended port numbers that can be used for Performance Monitoring tool

Port Number	Protocol	Service Name	Components that are involved in communication
4739	TCP and UDP	Performance Monitoring tool	Intra-cluster. This port needs to be accessible by all sensor nodes that are sending performance monitoring data.
8123	TCP	Object Metric collection	Intra-cluster
8124	TCP	Object Metric collection	Intra-cluster
8125	UDP	Object Metric collection	Intra-cluster
8126	TCP	Object Metric collection	Intra-cluster
8127	TCP	Object Metric collection	Intra-cluster
9085	TCP	Performance Monitoring Tool	Intra-cluster. This port needs to be open to all nodes where the performance collector is running. This port is used for internal communication between the collectors.
9980	TCP	Performance Monitoring Tool	Intra-cluster



# Infra Nodes

- Infra nodes in OpenShift are workers with a special label (and license)
- It is allowed to deploy Spectrum Scale CNSA and CSI to these nodes!

Where applicable, end users can use infrastructure nodes without disqualifying the node for infrastructure licensing to house these software components.

Examples may include:

- CNI and CSI drivers and controllers (also known as plugins).

Label the CNSA nodes with

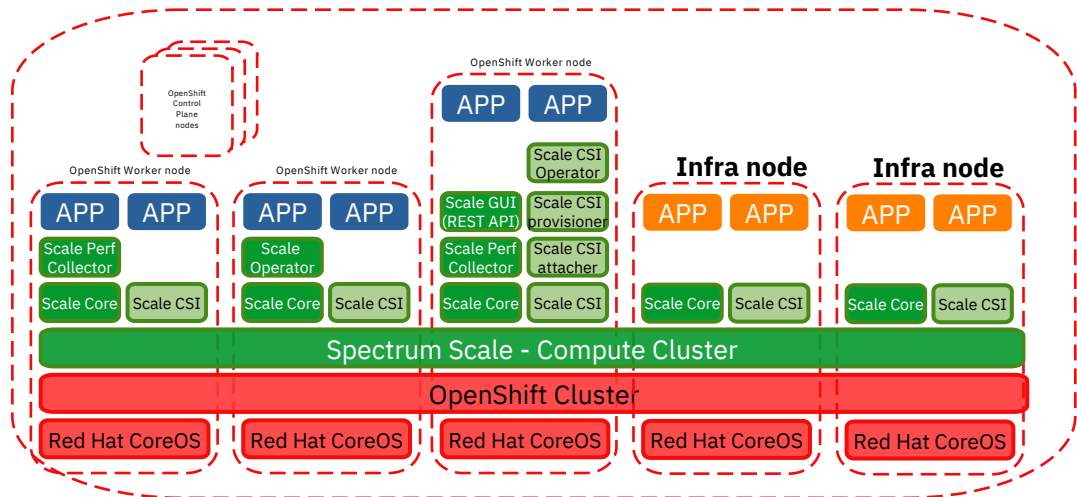
```
app.kubernetes.io/component: "scale"
```

and then use exactly that node selector in the ScaleCluster Custom Resource (CR):

```
spec:
```

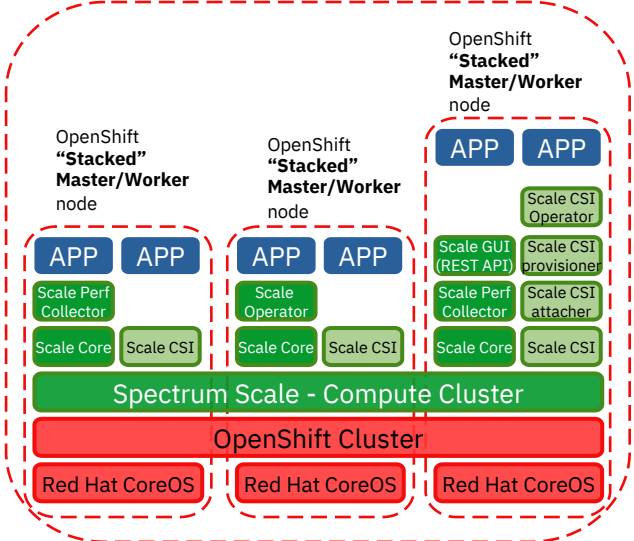
```
  nodeSelector:
```

```
    app.kubernetes.io/component: "scale"
```



# Control Plane Nodes?

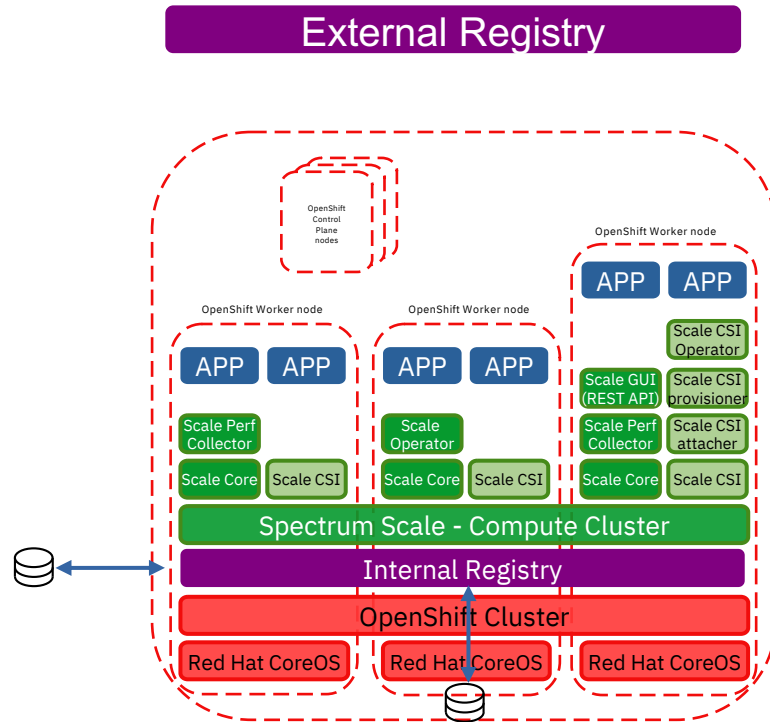
- Example: Minimum 3-node OpenShift deployment with “Stacked master/workers”
- Not yet supported by Scale CNSA/CSI



# Image Registry

- External registry
  - Internet public registries
    - Docker.io pull limit!
    - “Internet-only” with all images through IBM Container Registry new in 5.1.1.1
  - Private registry on internal network (Quay Enterprise , JFrog Artifactory, ...)
    - for air-gapped installations
    - Ensure all containers (including postgres, ubi-minimal) are accessible
- Internal registry
  - Using external Object, NFS or CSI storage
  - Using emptyDir (don't do this!)
    - Ephemeral temporary directory on a single machine → no HA

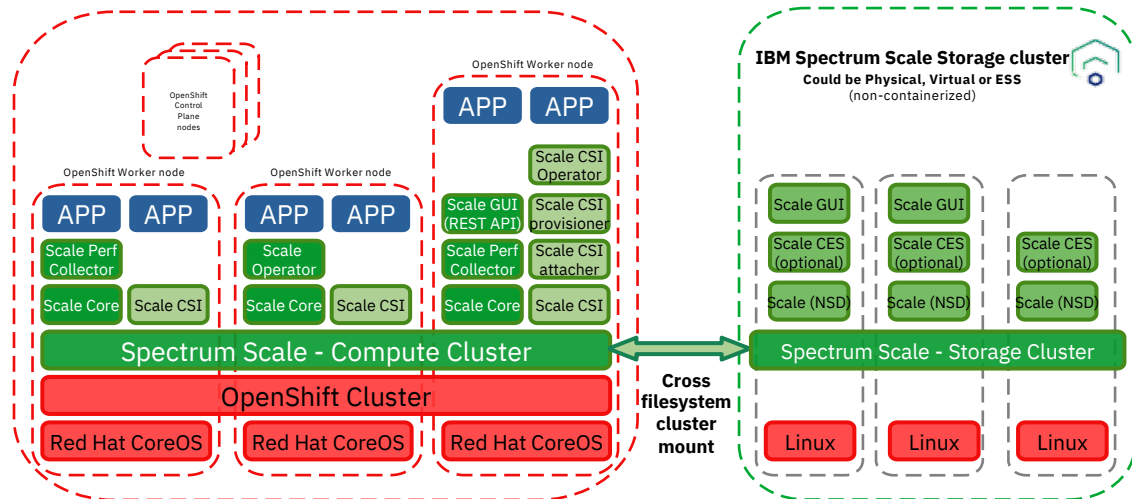
```
postgres: "docker.io/library/postgres@sha256:a21"
```



# DNS

- Ensure your Storage Cluster nodes can be forward- and backward DNS resolved
  - Work-around is using `hostAliases` within ScaleCluster CR
- Ensure consistency between Kubernetes node names and DNS entries
  - Work-around is using node mapping in CSI driver configuration

```
[root@ocp4-helpernode ~]# oc get nodes
NAME                                STATUS    ROLES    AGE   VERSION
master1.ocp.4ibm.de                Ready    master   53d   v1.19.0+a5a0987
master2.ocp.4ibm.de                Ready    master   53d   v1.19.0+a5a0987
master3.ocp.4ibm.de                Ready    master   53d   v1.19.0+a5a0987
worker1.ocp.4ibm.de                Ready    worker   104d  v1.19.0+a5a0987
worker2.ocp.4ibm.de                Ready    worker   104d  v1.19.0+a5a0987
worker3.ocp.4ibm.de                Ready    worker   104d  v1.19.0+a5a0987
worker4.ocp.4ibm.de                Ready    worker   104d  v1.19.0+a5a0987
worker5.ocp.4ibm.de                Ready    worker   104d  v1.19.0+a5a0987
worker6.ocp.4ibm.de                Ready    worker   104d  v1.19.0+a5a0987
```



Installation

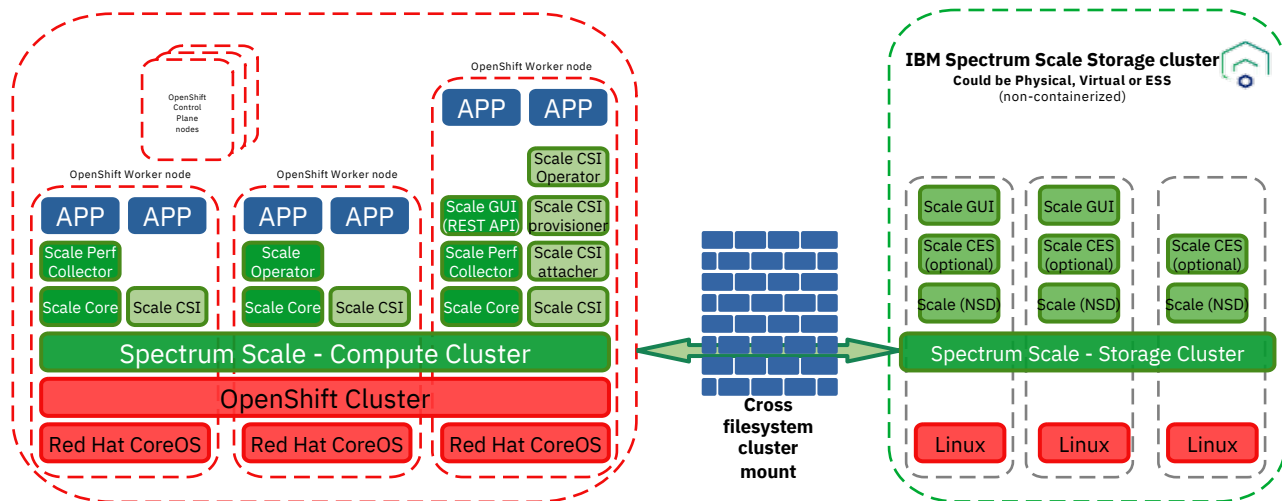


# Firewalls

- How to set the ephemeral port range on the CNSA cluster?
  - Add `tscCmdPortRange=60000-61000` to `scale-profile.yaml` at CNSA installation time

```
--
apiVersion: v1
kind: ConfigMap
metadata:
  name: ibm-spectrum-scale-core-profile
data:
  profile: |
    %cluster:
      pagepool=1G
      ignorePrefetchLUNCount=yes
      workerThreads=1024
      maxFilesToCache=128K
      maxStatCache=128K
      maxblocksize=16M

      # Prefetch Tunings
      initPrefetchBuffers=128
      prefetchTimeout=30
      prefetchPct=25
      prefetchMaxBuffersPerInstance=1024
      controlSetxattrImmutableSELinux=yes
      enforceFilesetQuotaOnPost=yes
      tscCmdPortRange=60000-61000
```



# Accessing vs. Owning Cluster in k8s resources

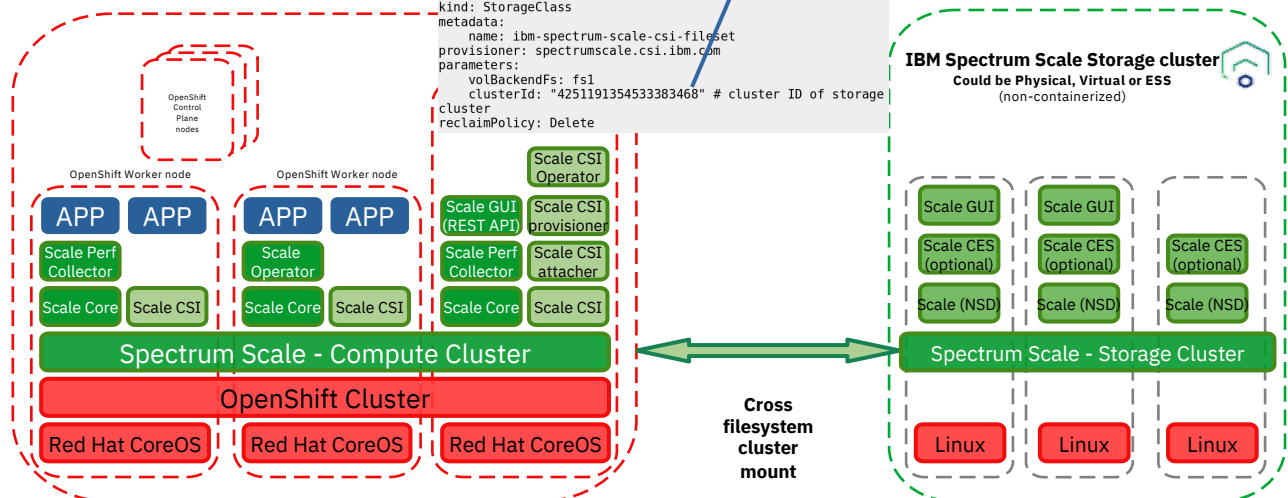
- CNSA – **ScaleCluster** CR holds a “descriptive name” for the Owning cluster within the **storageCluster** field, **mountPoint** is the Accessing cluster one
- CSI – **CSIScaleOperator** CR holds one entry for the Accessing cluster and another one for the Owning Cluster
- CSI Storage Class – needs to have **clusterId** of Owning cluster for fileset-based volumes

```
ScaleCluster CR:
fileSystems:
- name: "fs1"
  remoteMount:
    storageCluster: "storageCluster1"
    storageFs: "fs1"
    mountPoint: "/mnt/fs1"
remoteClusters:
- name: storageCluster1
  gui:
    cacert: "cacert-storage-cluster-1"
    host: ""
    secretName: "cnsa-remote-mount-storage-cluster-1"
    insecureSkipVerify: false
  contactNodes:
  - storageclusterInode1
  - storageclusterInode2
```

CSIScaleOperator CR:

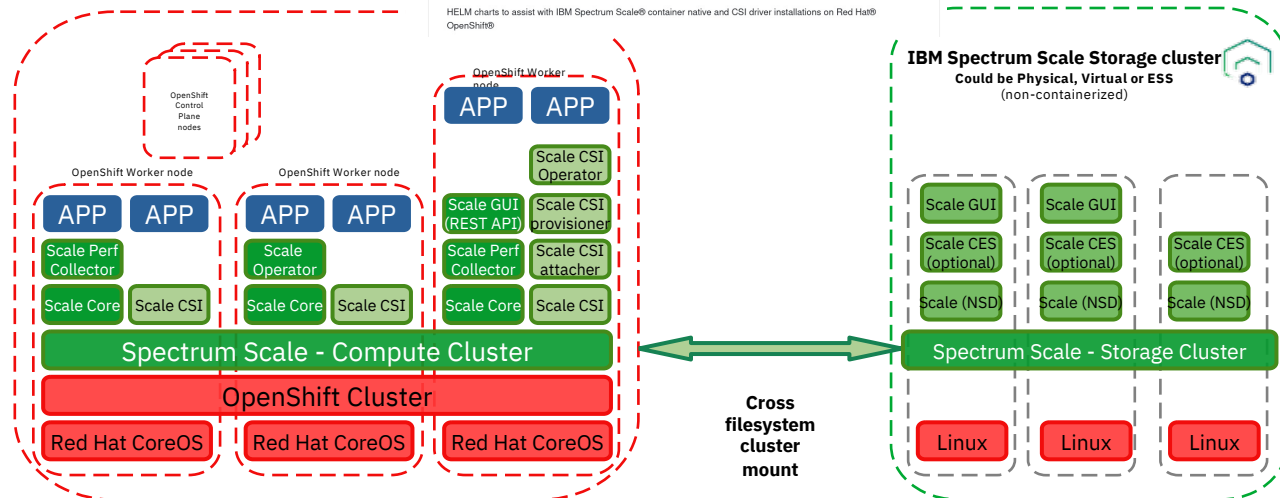
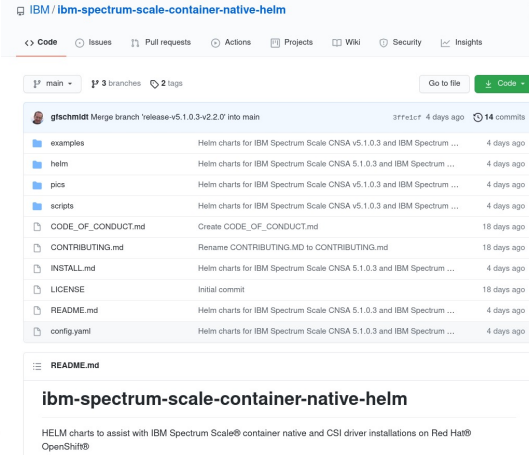
Key	Value	Description	Notes
clusters[0].id	"7061811038141825260"	IBM Spectrum Scale cluster ID of the CNSA cluster	oc exec <ibm-spectrum-scale-core-pod> -- nmlscuster   grep 'GPF5 cluster id:'
clusters[0].secrets	"secret-cnsa"	GUI secret with CNSA GUI credentials	
clusters[0].primary.primaryFs	"fs1"	File system on CNSA cluster	
clusters[0].primary.primaryFset	"scale-fset1"		Optional
clusters[0].primary.remoteCluster	"425119135453383468"	IBM Spectrum Scale cluster ID of the storage cluster	
clusters[0].restApi.guiHost	ibm-spectrum-scale-gui.ibm-spectrum-scale-ns	CNSA GUI Host	ibm-spectrum-scale-gui.ibm-spectrum-scale-ns
clusters[1].id	"425119135453383468"	IBM Spectrum Scale cluster ID of the storage cluster	Owning Cluster
clusters[1].secrets	"secret-storage"	GUI secret with storage cluster GUI credentials	
clusters[1].secureSslMode	false		
clusters[1].restApi.guiHost	"storagegui.example.com"	Storage cluster GUI Host	

```
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: ibm-spectrum-scale-csi-fileset
  provisioner: spectrumscale.csi.ibm.com
parameters:
  volBackendFs: fs1
  clusterId: "425119135453383468" # cluster ID of storage cluster
reclaimPolicy: Delete
```



# Installation simplification through Helm chart

- Note: These Helm charts are not supported by the IBM Spectrum Scale CNSA nor CSI offerings and are outside the scope of the IBM PMR process.
- All CNSA and CSI configuration in one place (`config.yaml`)
- Only need to specify
  - File system name
  - GUI hostnames
  - Owning cluster ID



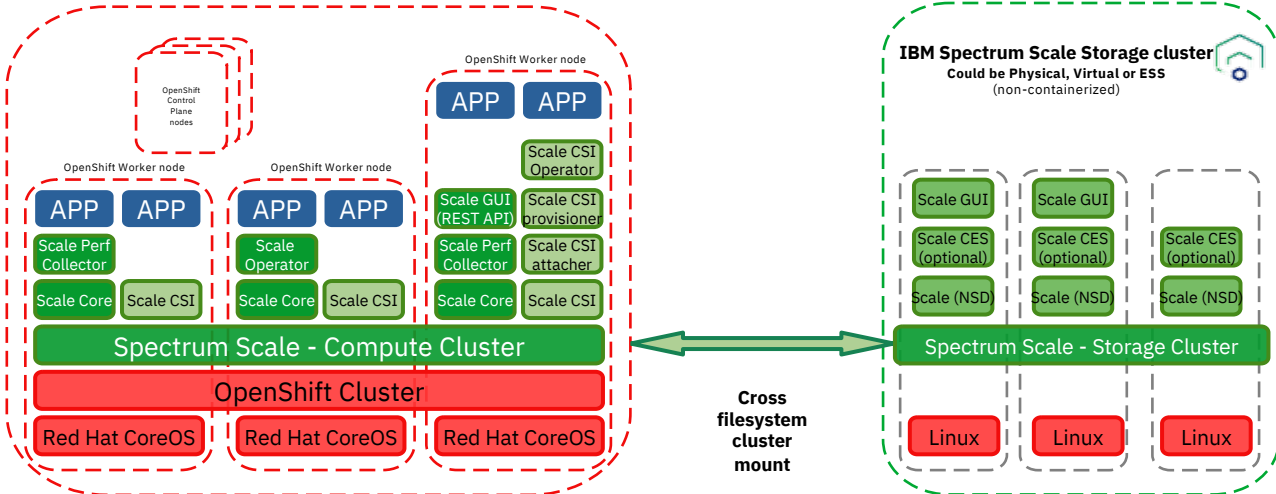


Operation



# Namespaces

- Avoid scheduling PVCs to
  - Default
  - openshift-
  - CNSA and CSI namespaces



# Default Storage Class

- Many workloads and templates do not have an option to specify the storage class
  - Set a default storage class to get started quickly

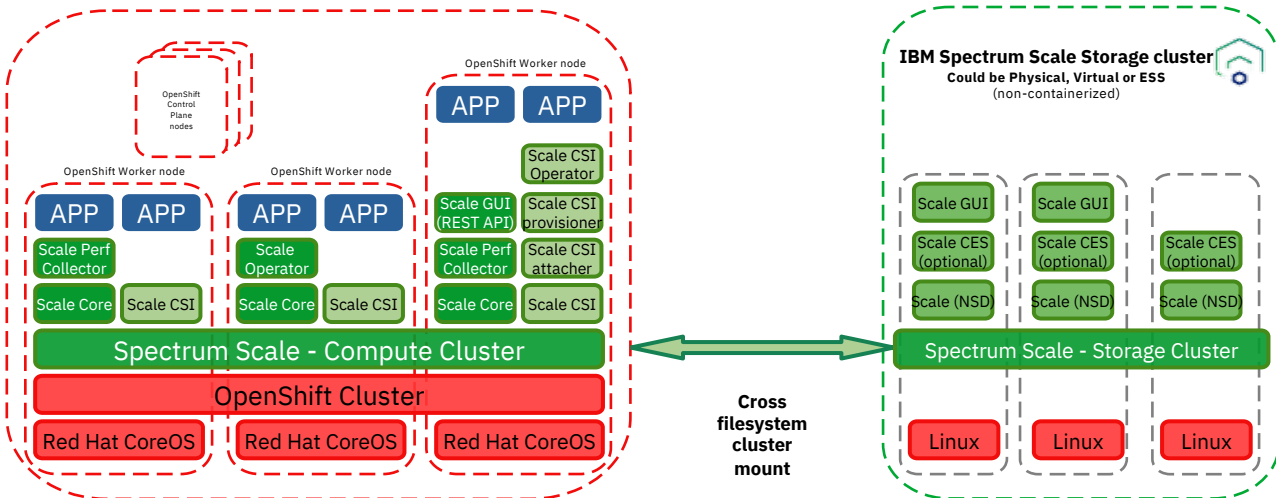
SC ibm-spectrum-scale-csi-fileset-dependent	spectrumscale.csi.ibm.com	Delete
SC ibm-spectrum-scale-csi-independent-fileset - Default	spectrumscale.csi.ibm.com	Delete
SC ibm-spectrum-scale-csi-lw	spectrumscale.csi.ibm.com	Delete

Edit annotations

KEY	VALUE
ss.kubernetes.io/is-default-class	true

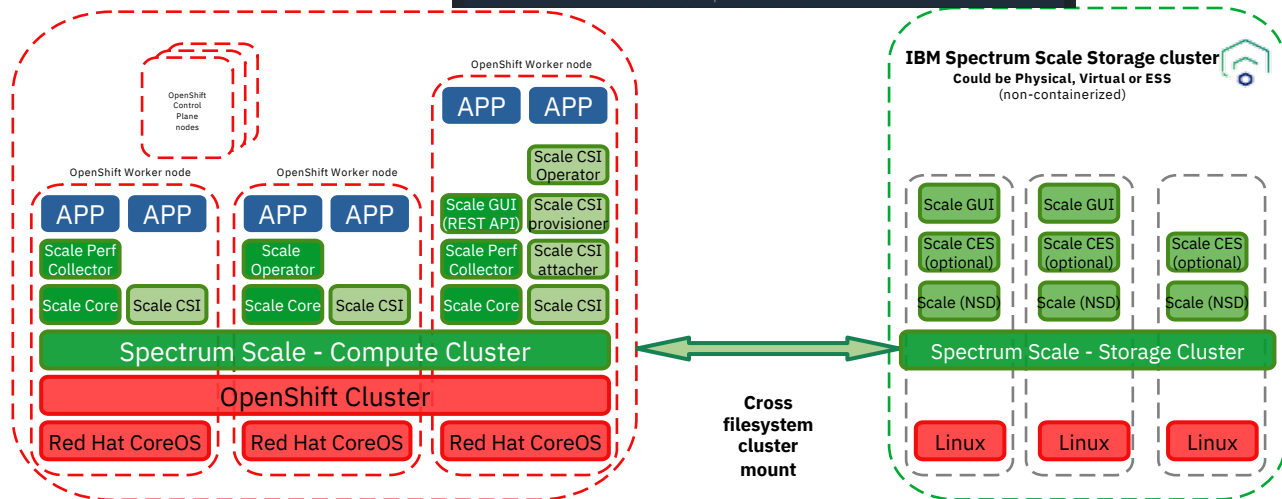
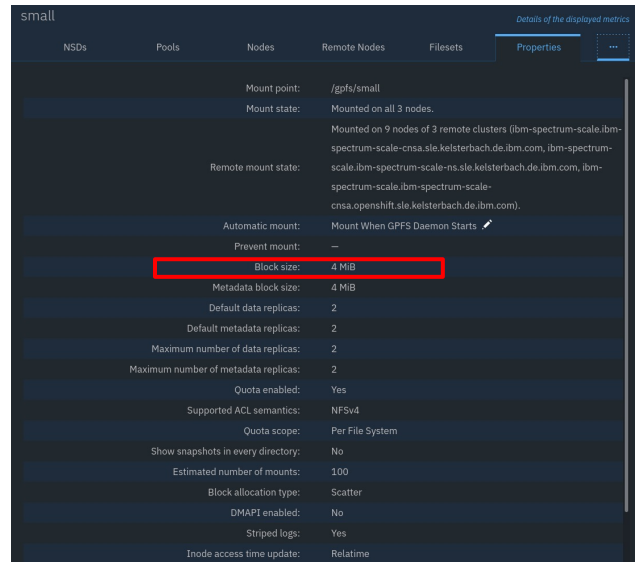
storageclass.kubernetes.io/is-default-class

Cancel Save



# Small PVs

- Some workloads might want to provision PVs as small as 1MiB
  - Might be below the FS block size
  - Quota will not be below the FS block size



# Troubleshooting



# CNSA Operator and ScaleCluster Resource

- The high-level logs are provided by the `ScaleCluster` CR
- Watch the Operator logs for the REST API commands
- In case the `scale-core` pods should not come up, have a look at the `config` init container!

CustomResourceDefinitions > CustomResourceDefinition details

CRD scaleclusters.scale.ibm.com

Details YAML Instances

Create ScaleCluster

Name Search by name...

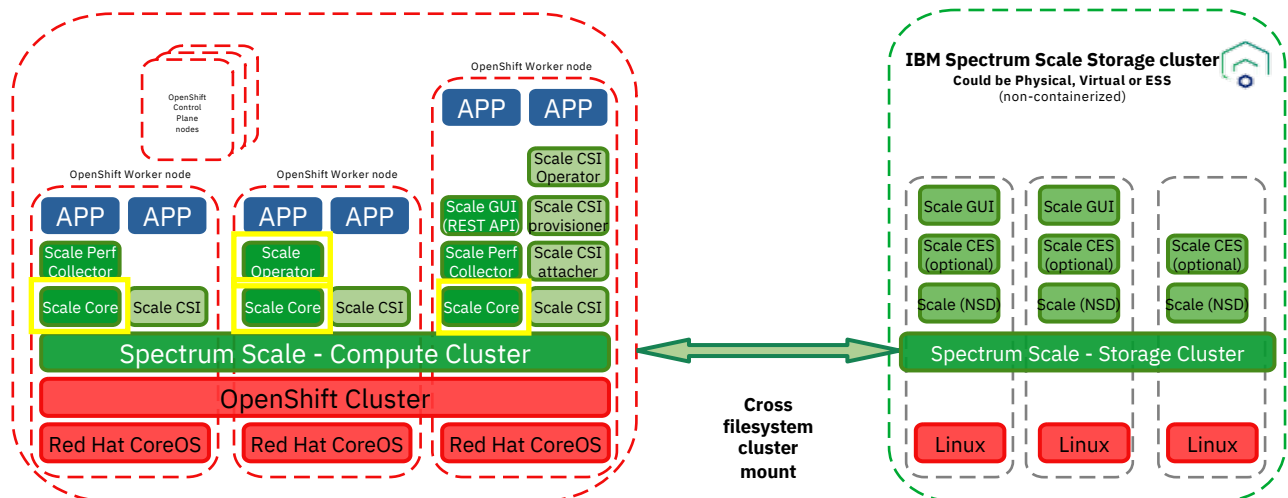
Name ↑

SC ibm-spectrum-scale

Conditions

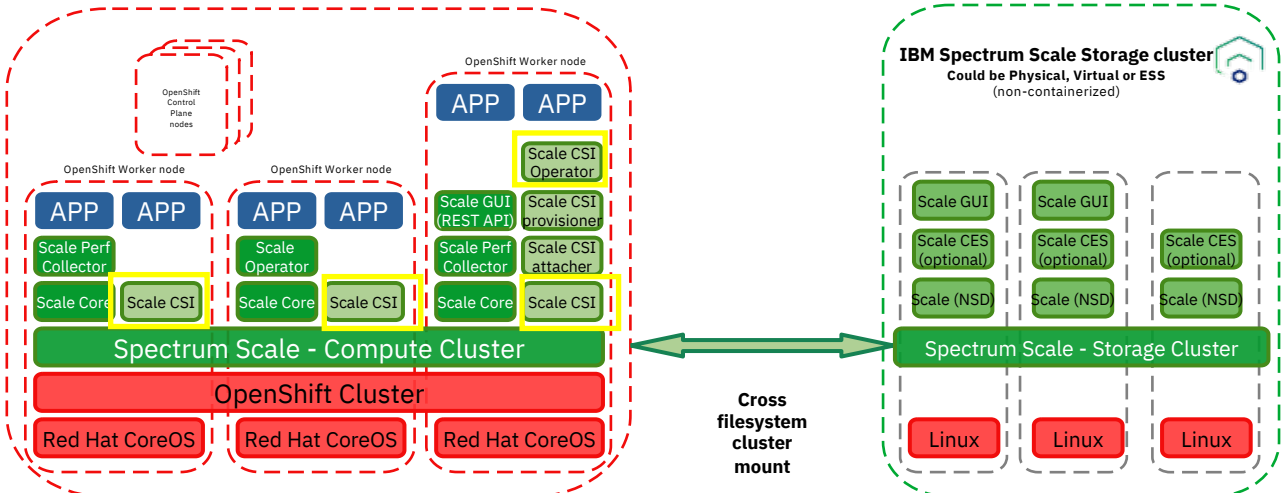
Type	Status	Updated	Reason	Message
CallhomeEnabled	False	Jun 2, 3:55 pm	CallhomeLicenseDeclined	Callhome disabled. The 'acceptLicense' parameter in the 'callhome' section of the configuration must be set to 'true' in order to enable callhome.
PodsConfigured	True	Jun 2, 4:41 pm	PodsConfigured	The pods have been configured successfully. Creation of pods and IBM Spectrum Scale cluster is ongoing or completed.
RemoteMountEnabled	True	Jun 2, 4:43 pm	RemoteMountCompleted	fsl, Configured successfully. File system mounted at /mnt/fsl.
GuiUserCreated	True	Jun 22, 11:31 am	GuiUserExists	GUI user ContainerOperator exists.

NS ibm-spectrum-scale-ns



# CSI Operator and scale-csi Pods

- The CSI Operator watches for Custom Resource changes and applies them via Ansible playbooks
- The `scale-csi` Pods (container `ibm-spectrum-scale-csi`) perform the “heavy lifting”



# Log Data Collection

- New data collection for CNSA 5.1.1.1 integrates with `oc adm must-gather`
- Script `spectrum-scale-driver-snap.sh` for CSI driver data collection

Download the tool from this location:

<https://raw.githubusercontent.com/IBM/ibm-spectrum-scale-csi/v2.2.0/tools/spectrum-scale-driver-snap.sh>

## Usage of the tool

```
spectrum-scale-driver-snap.sh [-n namespace] [-o output-dir] [-h]
```

```
-n: Debug data for CSI resources under this namespace will be collected. If not specified, default namespace is used. The tool returns  
-o: Output directory where debug data will be stored. If not specified, the debug data is stored in current directory.  
-h: Prints the usage
```

## Execute the `ibm-spectrum-scale-must-gather`

1. In the directory where your must-gather contents will be stored, run must-gather command using the `ibm-spectrum-scale-must-gather` image:

```
oc adm must-gather --image=image-registry.openshift-image-registry.svc:5000/ibm-spectrum-scale-ns/ibm-spectrum
```

2. Once completed, a new directory with `must-gather` prefix will be present in your working directory.

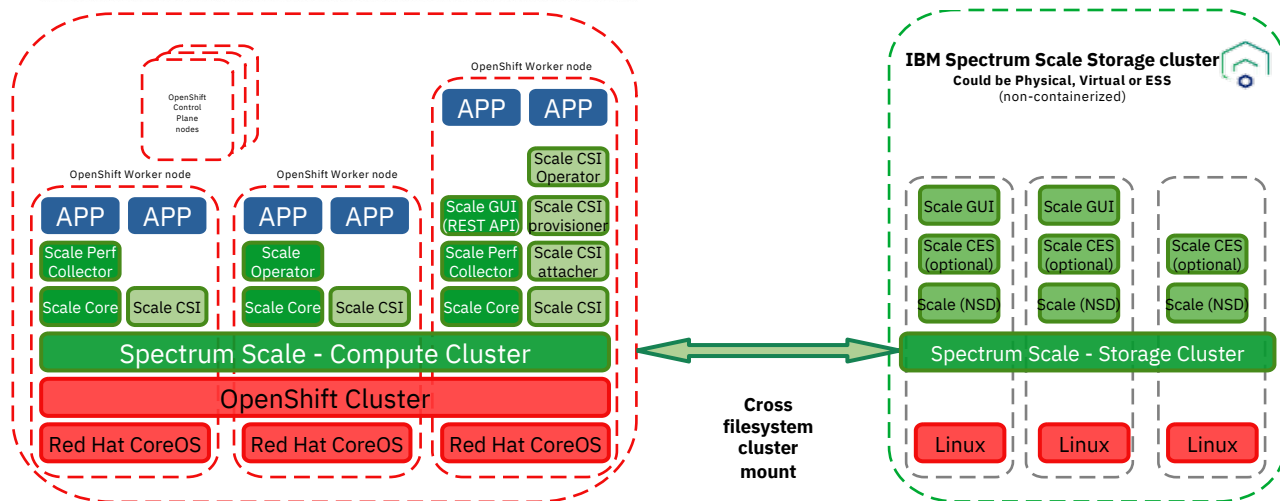
For example:

```
# ls -ltr  
drwxr-xr-x 3 root root 229 Jun 14 09:11 must-gather.local.681612165636007567
```

3. Create a compressed file from the must-gather directory that was just created in your working directory.

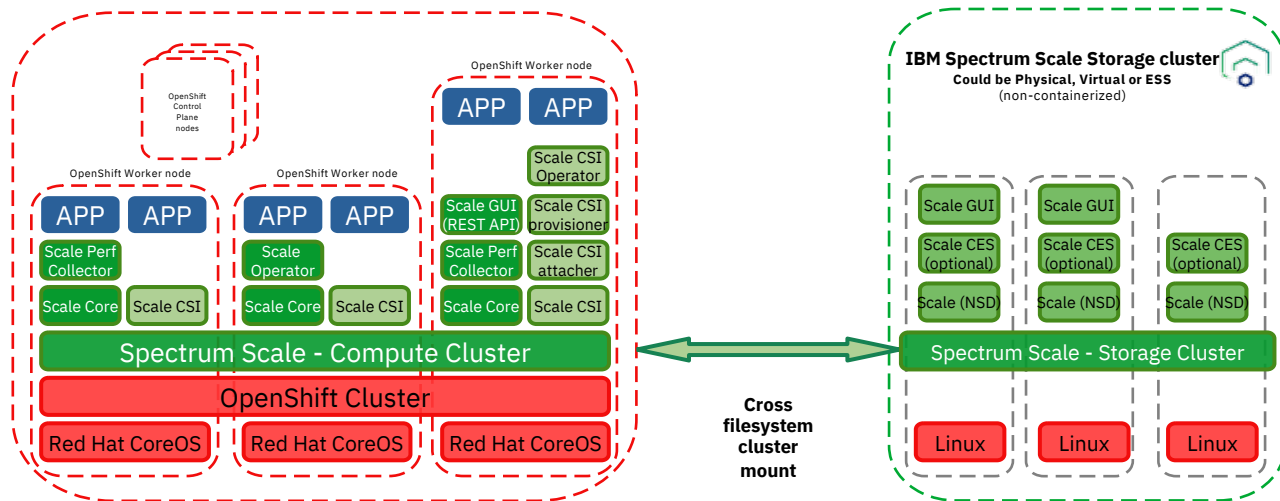
```
tar cvaf must-gather.tar.gz must-gather.local.681612165636007567/
```

**Note:** Replace the directory name used in this command with your respective must-gather directory.



# Call to action

- Spectrum Scale CNSA is a perfect match for OpenShift container environments featuring
  - Operator-driven automated installation, operation and storage provisioning
  - Virtually unlimited throughput
  - Exabyte scalability
  - Single namespace eliminating data silos
  - Flexible hardware- and software-based Storage Cluster options
  - Multicloud readiness
- Talk to your IBM sales contact for an RPQ or write to [scale@us.ibm.com](mailto:scale@us.ibm.com)



# Log your RFE!

[https://www.ibm.com/developerworks/rfe/execute?use\\_case=productsList](https://www.ibm.com/developerworks/rfe/execute?use_case=productsList)

- Spectrum Scale (formerly known as GPFS) - Private RFEs
- Spectrum Scale (formerly known as GPFS) - Public RFEs

*contact*

Filter the page content by brand and product

Servers and Systems So... ▾

Spectrum Scale (formerly known as GPFS) - Pu... ▾ ▶

[Hot](#)

[Top](#)

[New](#)

Search



35  
votes

### **eleminate lack of I/O on mmdelsnapshot start**

When deletion of bunch of snapshots starts we a lack of I/O for about three minutes. NFS Clients see a huge delay of I/O. Related applications hanging for this time and user connections and run into t...


**Under Consideration**

21  
votes

### **start services after gpfs (filesystems) is ready using systemd**


the current systemd units gpfs.service and <fs-mountpoint>.mount units can't be used to depend on (After/Required/.. systemd attributes) for other services. GPFS service is reporting itself as success...

**Under Consideration**



Check <https://www.spectrumscaleug.org/experttalks>  
for charts, show notes and upcoming talks

- Past talks:
  - 001: What is new in Spectrum Scale 5.0.5?
  - 002: Best practices for building a stretched cluster
  - 003: Strategy update
  - 004: Update on performance enhancements in Spectrum Scale (file create, MMAP, direct IO, ESS 5000)
  - 005: Update on functional enhancements in Spectrum Scale (inode management, vCPU scaling, NUMA considerations)
  - 006: Persistent Storage for Kubernetes and OpenShift environments
  - 007: Manage the lifecycle of your files using the policy engine
  - 008: Multi-node scaling of AI workloads using Nvidia DGX, OpenShift and Spectrum Scale
  - 009: Continental: Deep Thought – An AI Project for Autonomous Driving Development
  - 010: Data Accelerator for Analytics and AI (DAAA)
  - 011: What is new in Spectrum Scale 5.1.0?
  - 012: Lenovo - Spectrum Scale and NVMe Storage
  - 013: Event driven data management and security using Spectrum Scale Clustered Watch Folder and File Audit Logging
  - 014: What is new in Spectrum Scale 5.1.1?
- Today:
  - **June 29: IBM Spectrum Scale Container Native Storage Access**




# Thank you!

Please help us to improve Spectrum Scale with your feedback

- If you get a survey in email or a popup from the GUI, please respond
- We read every single reply

Provide Feedback ×

---



Tell IBM What You Think

Let us know what you think about IBM Spectrum Scale. It takes only a couple of minutes for you to help us improve our service. [IBM Privacy Policy](#)



## Spectrum Scale User Group

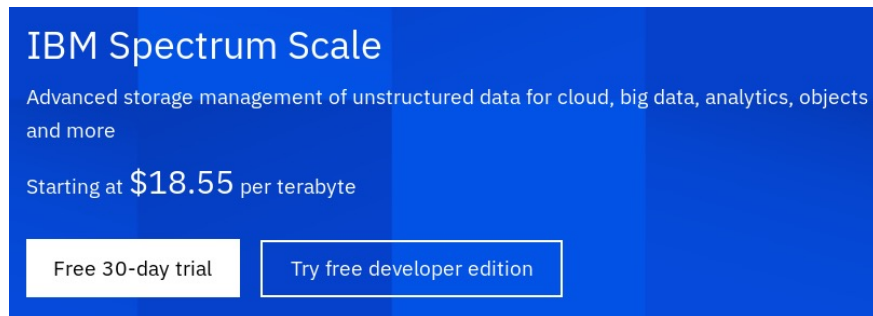
The Spectrum Scale (GPFS) User Group is free to join and open to all using, interested in using or integrating IBM Spectrum Scale.

The format of the group is as a web community with events held during the year, hosted by our members or by IBM.

See our web page for upcoming events and presentations of past events. Join our conversation via mail and Slack.

[www.spectrumscaleug.org](http://www.spectrumscaleug.org)

# Spectrum Scale Developer Edition!



IBM Spectrum Scale

Advanced storage management of unstructured data for cloud, big data, analytics, objects and more

Starting at **\$18.55** per terabyte

Free 30-day trial    Try free developer edition

## Fully functional!

- Based on first PTF of a release
- Derived from **Data Management Edition (DME)**
- Limited to 12 TBs:  
enough for a small test cluster
- Available from the Scale "try and buy" page on [ibm.com](https://ibm.com)

Free for non-production use, e.g. test, learning, upgrade prep...

- If you have to ask, it's probably not permitted

Not formally supported

# Spectrum Scale Early Programs

## Types of Programs:

### Alpha

Influence the development of new technology by gaining before market access to product code. Alpha programs are typically confidential and the first opportunity for you to interact with a feature or function.

### Beta

Try out a new offering with the team who owns the product and influence its usability and design. A Beta program gives you the ability to evaluate and provide feedback on IBM products before the products general availability. Beta programs are typically confidential and run prior to GA.

### Early Support Program (ESP)

Be one of the few selected participants to validate new Software or Hardware and potentially give your enterprise an edge over the competition. The IBM early support programs give you and IBM the opportunity to develop, evaluate, and gain experience with a product or a set of products in your enterprise environment.



### Customer Success

- Evaluate new IBM HW or SW in your environment.
- Validate procedures and interoperability with other products in your enterprise.
- Opportunity to Influence Product Design
- Early Enablement and education
- Strengthen Partnership with IBM

Talk to your IBM contact or Partner to be nominated!

# Spectrum Scale on GitHub!

<https://github.com/IBM/SpectrumScaleTools>

- IBM Spectrum Scale Bridge for Grafana
- IBM Spectrum Scale cloud install
- IBM Spectrum Scale Container Storage Interface driver
- IBM Spectrum Scale install infra
- IBM Spectrum Scale Security Posture
- Oracle Cloud Infrastructure IBM Spectrum Scale terraform template
- SpectrumScale\_ECE\_CAPACITY\_ESTIMATOR
- SpectrumScale\_ECE\_OS\_OVERVIEW
- SpectrumScale\_ECE\_OS\_READINESS
- SpectrumScale\_ECE\_STORAGE\_READINESS
- SpectrumScale\_ECE\_tuned\_profile
- SpectrumScale\_NETWORK\_READINESS

Find open source tools that are related with IBM Spectrum Scale.

Unless stated otherwise, the tools compiled in this list come with no warranty of any kind from IBM.

Check out the FAQ!

<https://www.ibm.com/support/knowledgecenter/en/STXKQY/gpfsclustersfaq.html>

<https://www.ibm.com/support/knowledgecenter/STXKQY/gpfsclustersfaq.pdf?view=kc>

<https://www.ibm.com/support/knowledgecenter/SSYSP8/gnrfaq.html>

HTML or PDF



Spectrum Scale version  
compatibility with OS or  
kernels

Updated regularly!