

Introduction to Spectrum Scale Data Access Services (DAS) S3

Spectrum Scale UK User Group Meeting
London – June 30th, 2022

Kira Dickey (IBM)
Ulf Troppens (IBM)



The world is changing ...



2005

Luca Bruno/AP

The world is changing ...



2005

Michael Sohn/AP



Luca Bruno/AP

2013



Object Storage is evolving to support changing workload requirements

AWS S3

- Object Storage for AWS public cloud

AWS S3 API

- De-facto standard for object access protocols
 - Used by many vendors and applications
-

AWS S3

- Launch Date: March 14, 2006
- S3 API: <https://docs.aws.amazon.com/AmazonS3/latest/API>

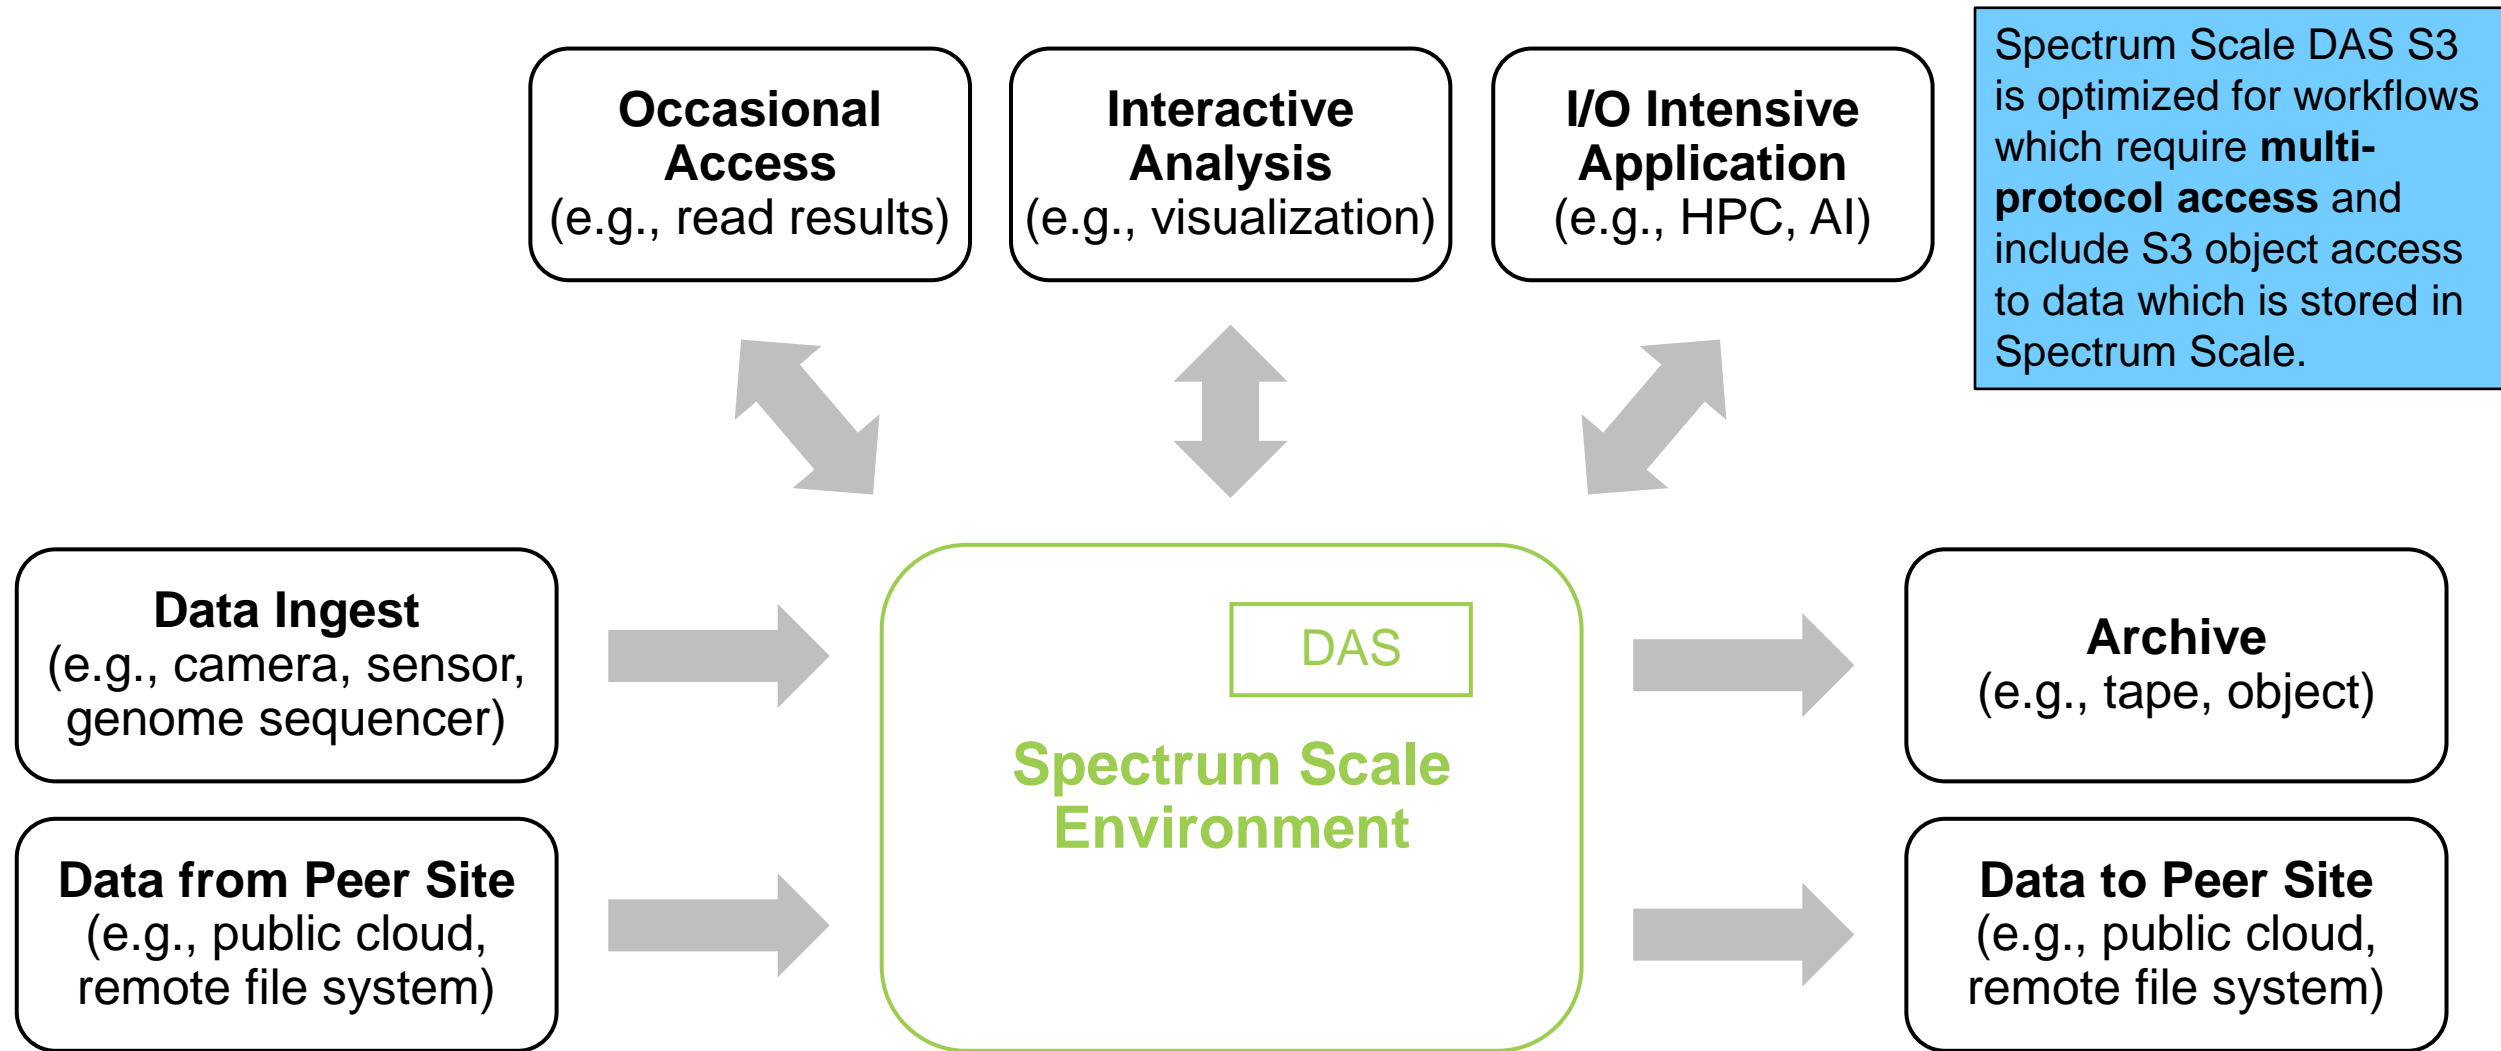
History of Object Storage

- The History Boys: Object storage ... from the beginning
https://www.theregister.com/2016/07/15/the_history_boys_cas_and_object_storage_map/

Spectrum Scale DAS S3

- Launch Date: May 27, 2022
<https://www.ibm.com/docs/en/scalecontainernative?topic=513-spectrum-scale-data-access-services>

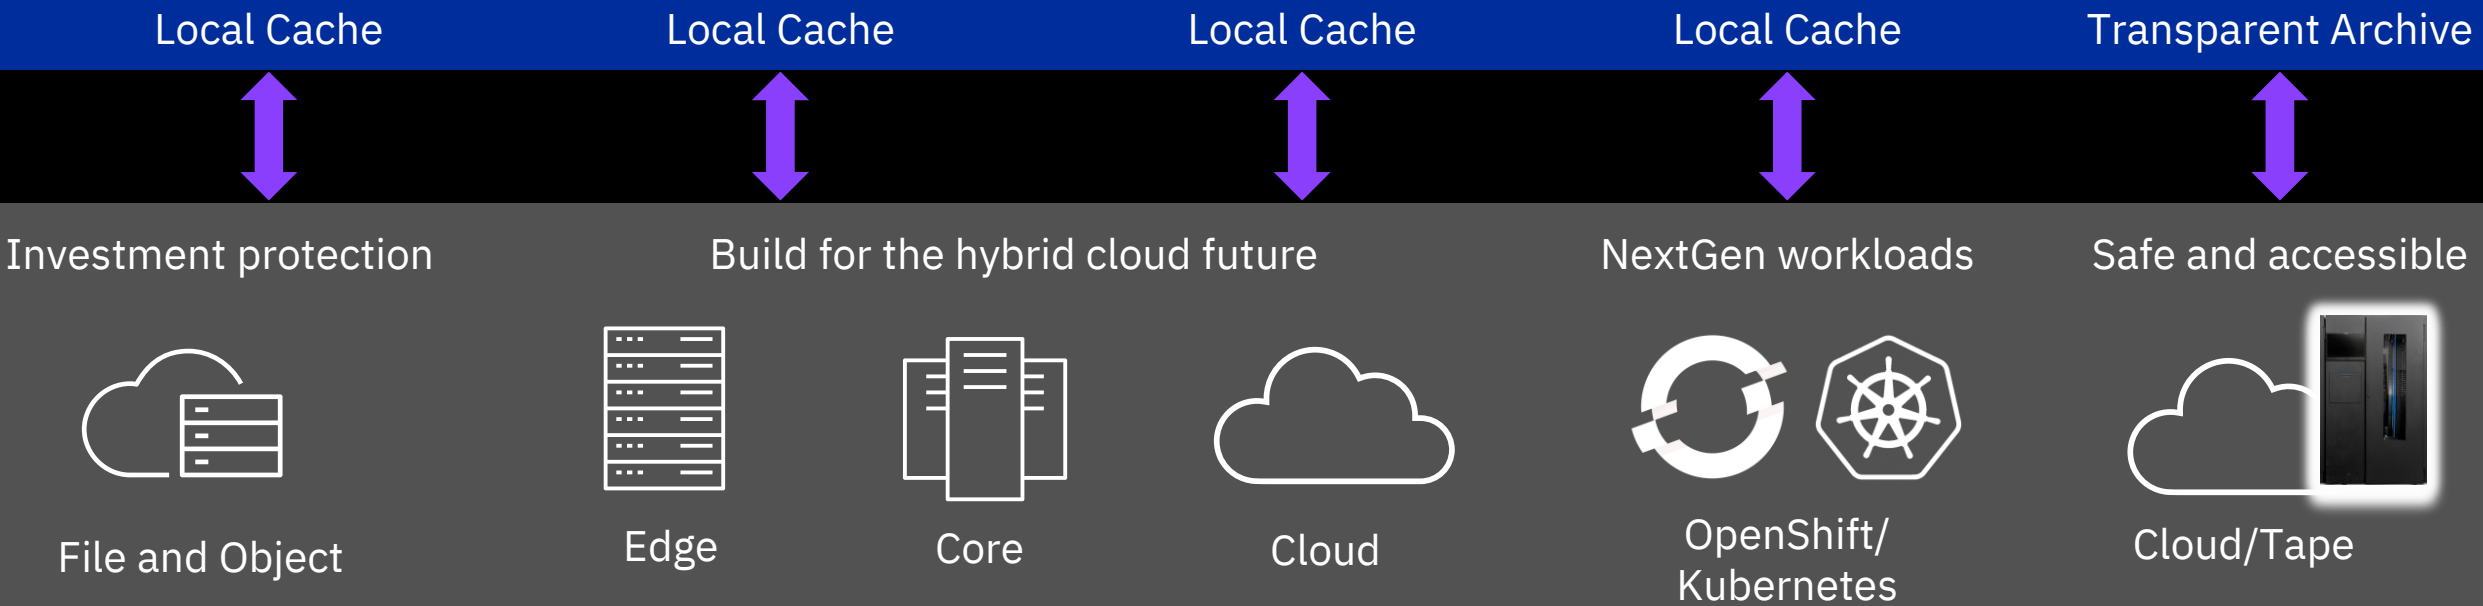
Spectrum Scale DAS S3 – System Context



Now your applications can simultaneously connect multiple data sources using multiple data interfaces from edge to core to cloud



A global data platform provides a
single source of truth



IBM Global Data Platform for Unstructured File & Object Data

Unstructured Data Services Framework




Applications and Workloads



Data Access Services



Data Caching Services



Data Management Services



Data
Security
Services


IBM's Global Data Platform for File & Object Data



4 Data Security Services

- Identify
- Protect
- Detect
- Respond
- Recover

1 Data Access Services



2 Data Caching Services

Global Data Platform (powered by Spectrum Scale)



Investment protection



File & Object
Storage

(NetApp, PowerScale, etc)

Object Storage



IBM COS

File Storage



Spectrum Scale

NextGen workloads



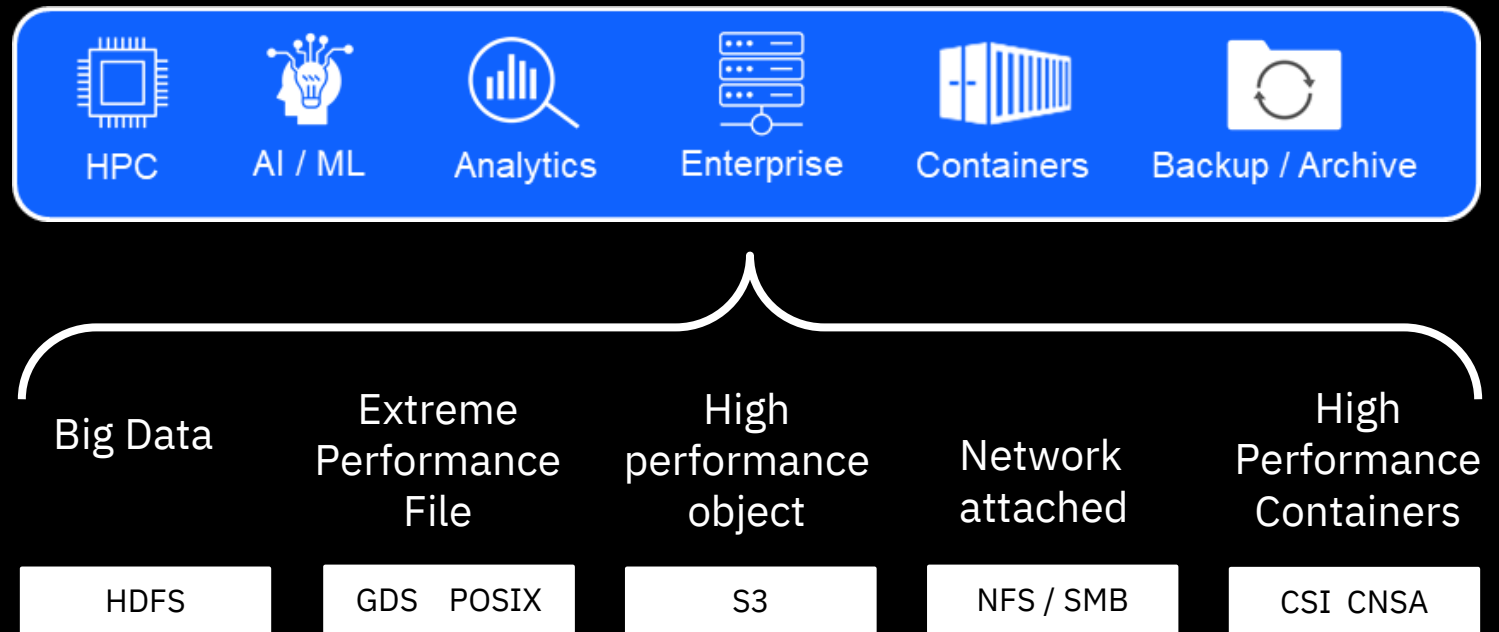
Spectrum Fusion

3 Data Management Services

Data Access Services - Let your application requirements define the interface to use

- **A single storage system** for multiple application requirements
- **Developers can choose** the interface they need to provide the results they require
- **Storage resources are optimized** and provide faster results with multiple high performance data access options

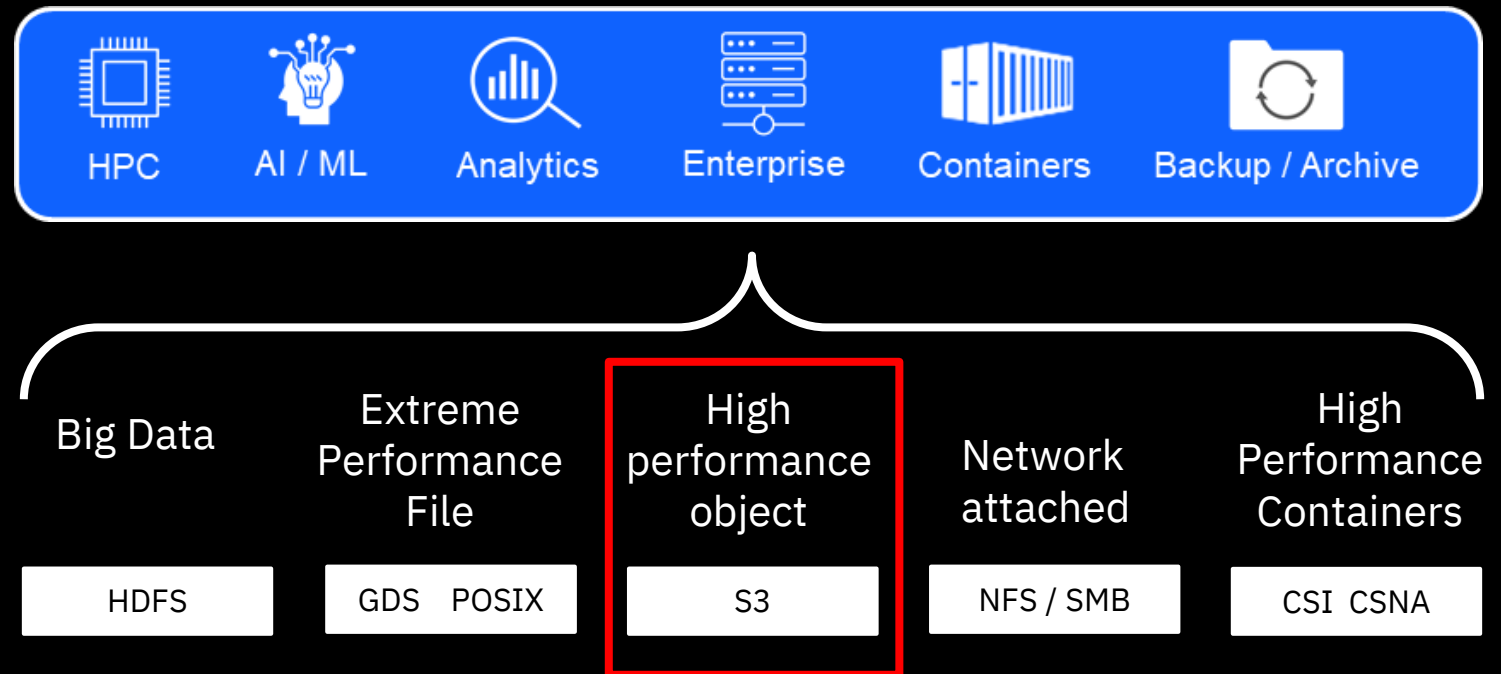
Multiple interfaces access the same data



Data Access Services - Let your application requirements define the interface to use

- **A single storage system** for multiple application requirements
- **Developers can choose** the interface they need to provide the results they require
- **Storage resources are optimized** and provide faster results with multiple high performance data access options

Multiple interfaces access the same data



Spectrum Scale DAS S3 – Reference Architecture


External Applications

Shared Data Network

DAS Cluster

Private Data Network

Storage Cluster



S3 clients run outside OpenShift cluster for DAS Cluster

IBM recommends well controlled high-speed network to connect S3 clients.

Spectrum Scale DAS 5.1.3 requires dedicate compact OpenShift cluster running on bare-metal x86_64

IBM recommends to configure dedicated high-speed network to connect all Spectrum Scall components

Spectrum Scale DAS 5.1.3 supports all ESS models

Spectrum Scale DAS S3 – Performance

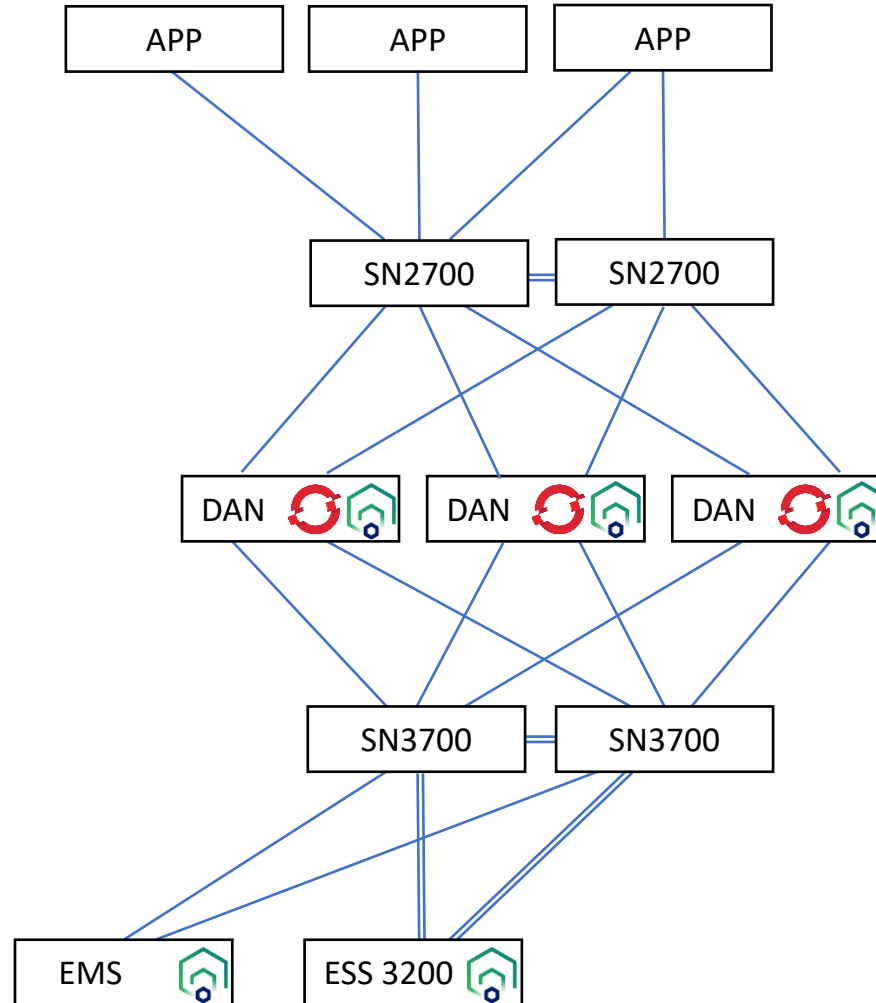
External Applications

Shared Data Network

DAS Cluster

Private Data Network

Storage Cluster



61 GB/s read for 1GB objects
23 GB/s write for 1GB objects
Performance was limited by single ESS 3200
using 24 NVMe drives

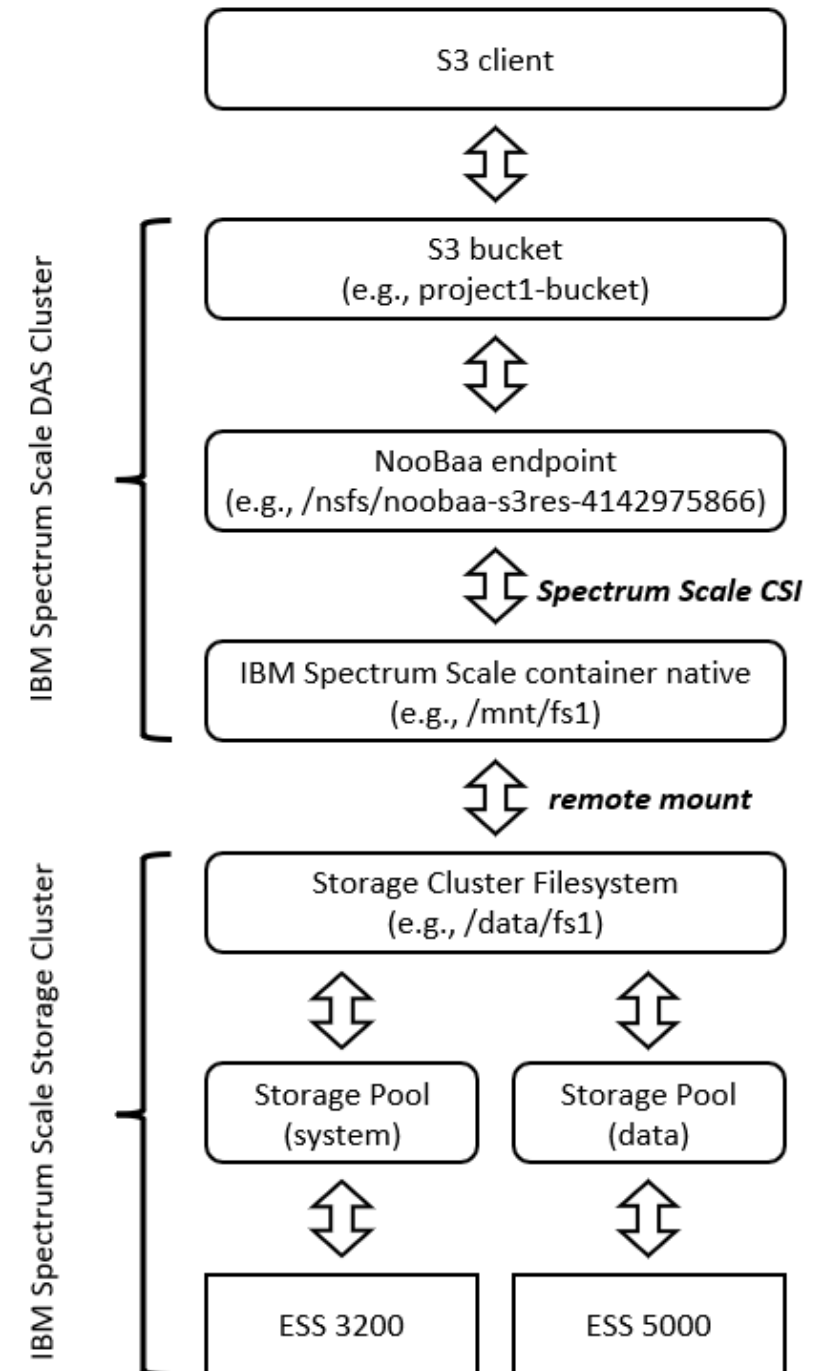
Performance is measured
for DAS Cluster
(aggregated throughput)

Blog post with details:

<https://community.ibm.com/community/user/storage/blogs/silvana-de-gyves-avila1/2022/05/20/ibm-data-access-services-performance-evaluation>

Spectrum Scale DAS S3

– Data Path



<https://www.ibm.com/docs/en/scalecontainernative?topic=architecture-data-path>

Spectrum Scale DAS S3 – Use Cases

Containerized S3 object access integrated within Spectrum Scale delivering high performance object for AI and analytics workloads


Data Access Services (DAS): Modernized, containerized protocol stack

- **HPO:** Containerized S3 object access component of DAS
- **Performance:** Spectrum Scale DAS S3 5.1.3 is optimized for large objects

Emerging Object Use Cases:

May 2022

- **High-Performance Object Storage** for AI & Analytics



Spectrum Scale DAS S3 – Use Cases

Containerized S3 object access integrated within Spectrum Scale delivering high performance object for AI and analytics workloads

Data Access Services (DAS): Modernized, containerized protocol stack

- **HPO:** Containerized S3 object access component of DAS
- **Performance:** Spectrum Scale DAS S3 5.1.3 is optimized for large objects

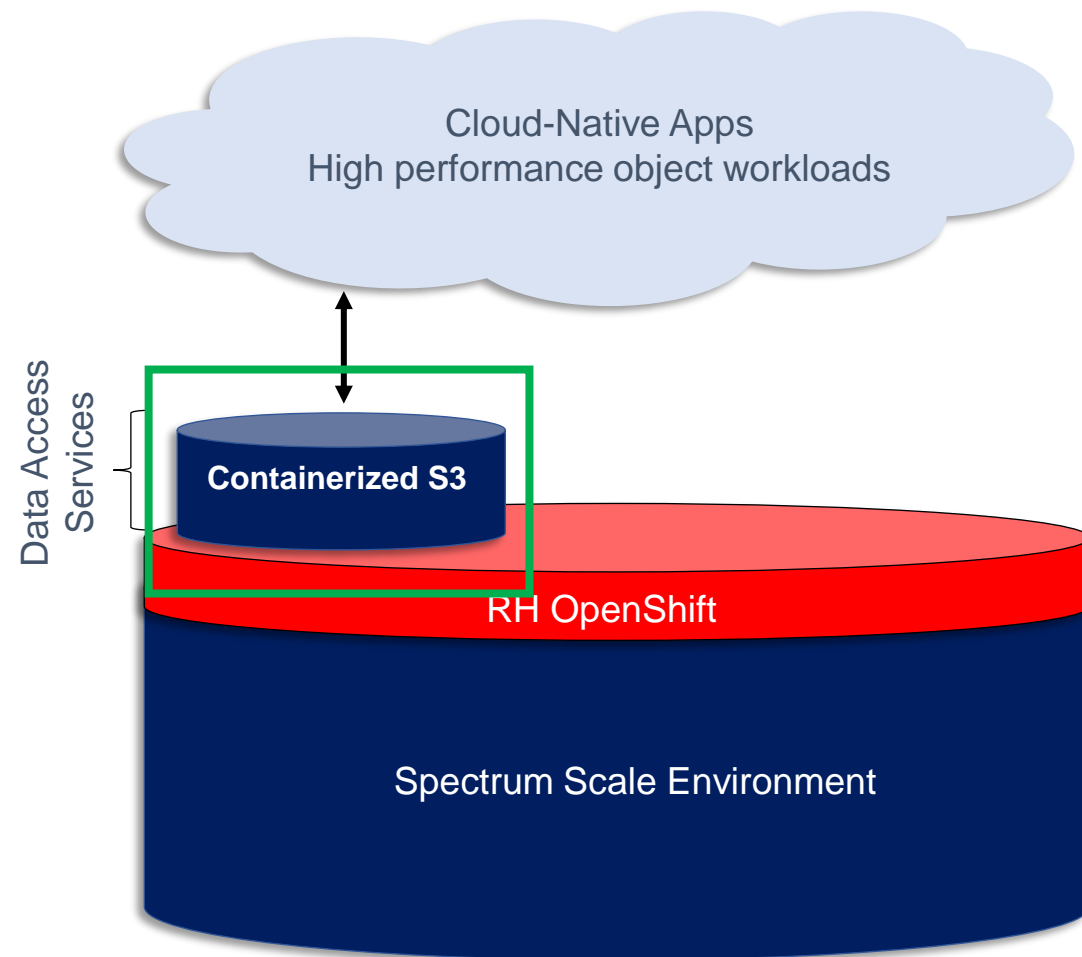
Emerging Object Use Cases:

May 2022

- **High-Performance Object Storage** for AI & Analytics

Considered future Use Cases

- **Multi-protocol data access** for AI & Analytics using NFS, S3 and CSI without copying data
- **S3 access to AFM cached data** for accelerated S3 access to data stored on remote object stores and legacy storage systems
- **Backup and Restore** using high-performance object storage.
- **Tape aware object storage** for cost-effective Active Archive

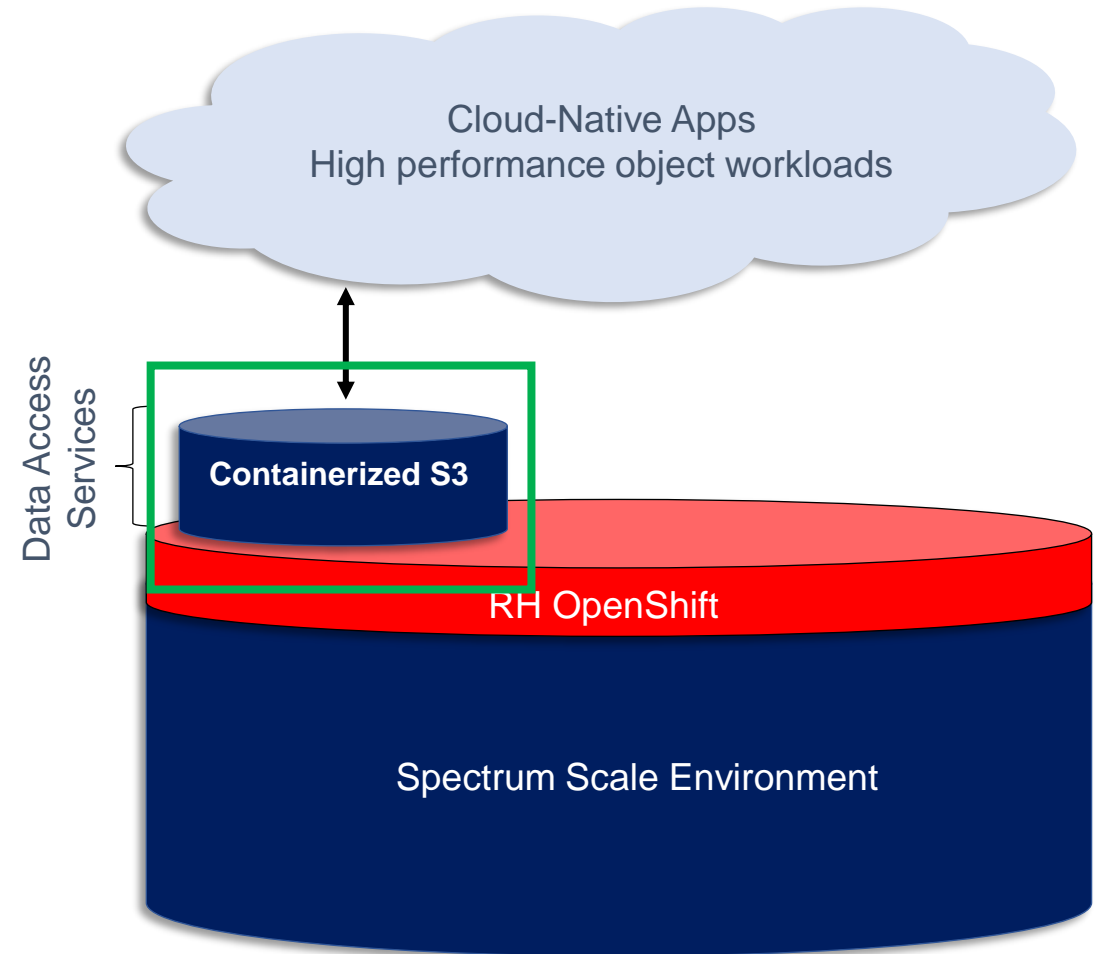


Spectrum Scale DAS S3 – Software Prerequisites

Containerized S3 object access integrated within Spectrum Scale delivering high performance object for AI and analytics workloads

Customer Requirements & HPO Dependencies:

- Spectrum Scale 5.1.3.1: DAE, DME, ESS for DAE, ESS for DME, ECE (future)
- OpenShift 4.9.31
- CNSA 5.1.3.1 / CSI 2.5.1
- All ESS models at GA



<https://www.ibm.com/docs/en/scalecontainernative?topic=services-planning>

Packaging, Pricing, and Support

Packaging

- Included in Spectrum Scale DAE, DME, ESS for DAE, ESS for DME
- Installed as part of a Spectrum Scale implementation.

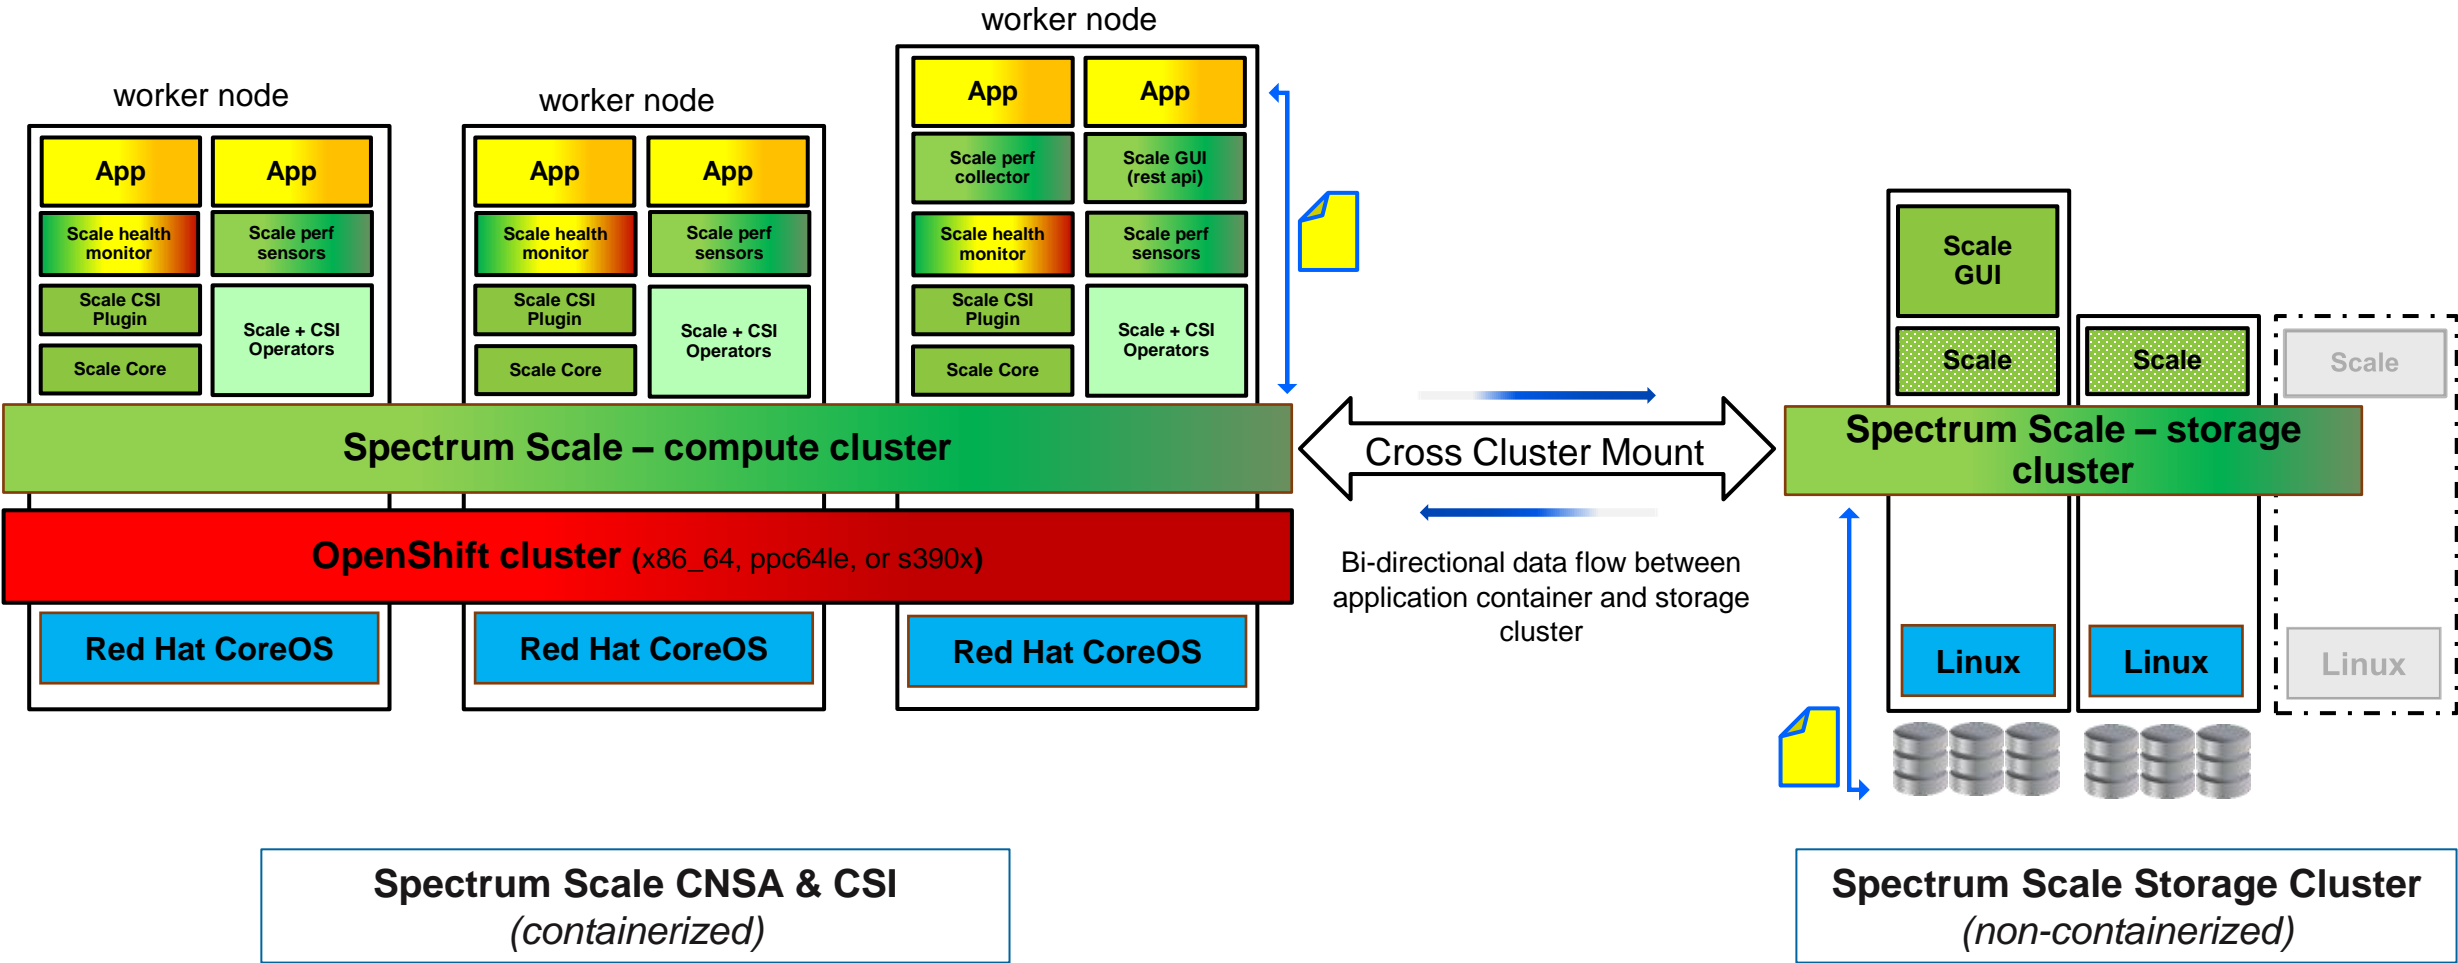
Pricing

- Included in Spectrum Scale DAE, DME, ESS for DAE, ESS for DME
- No separate pricing, included in Spectrum Scale

Support

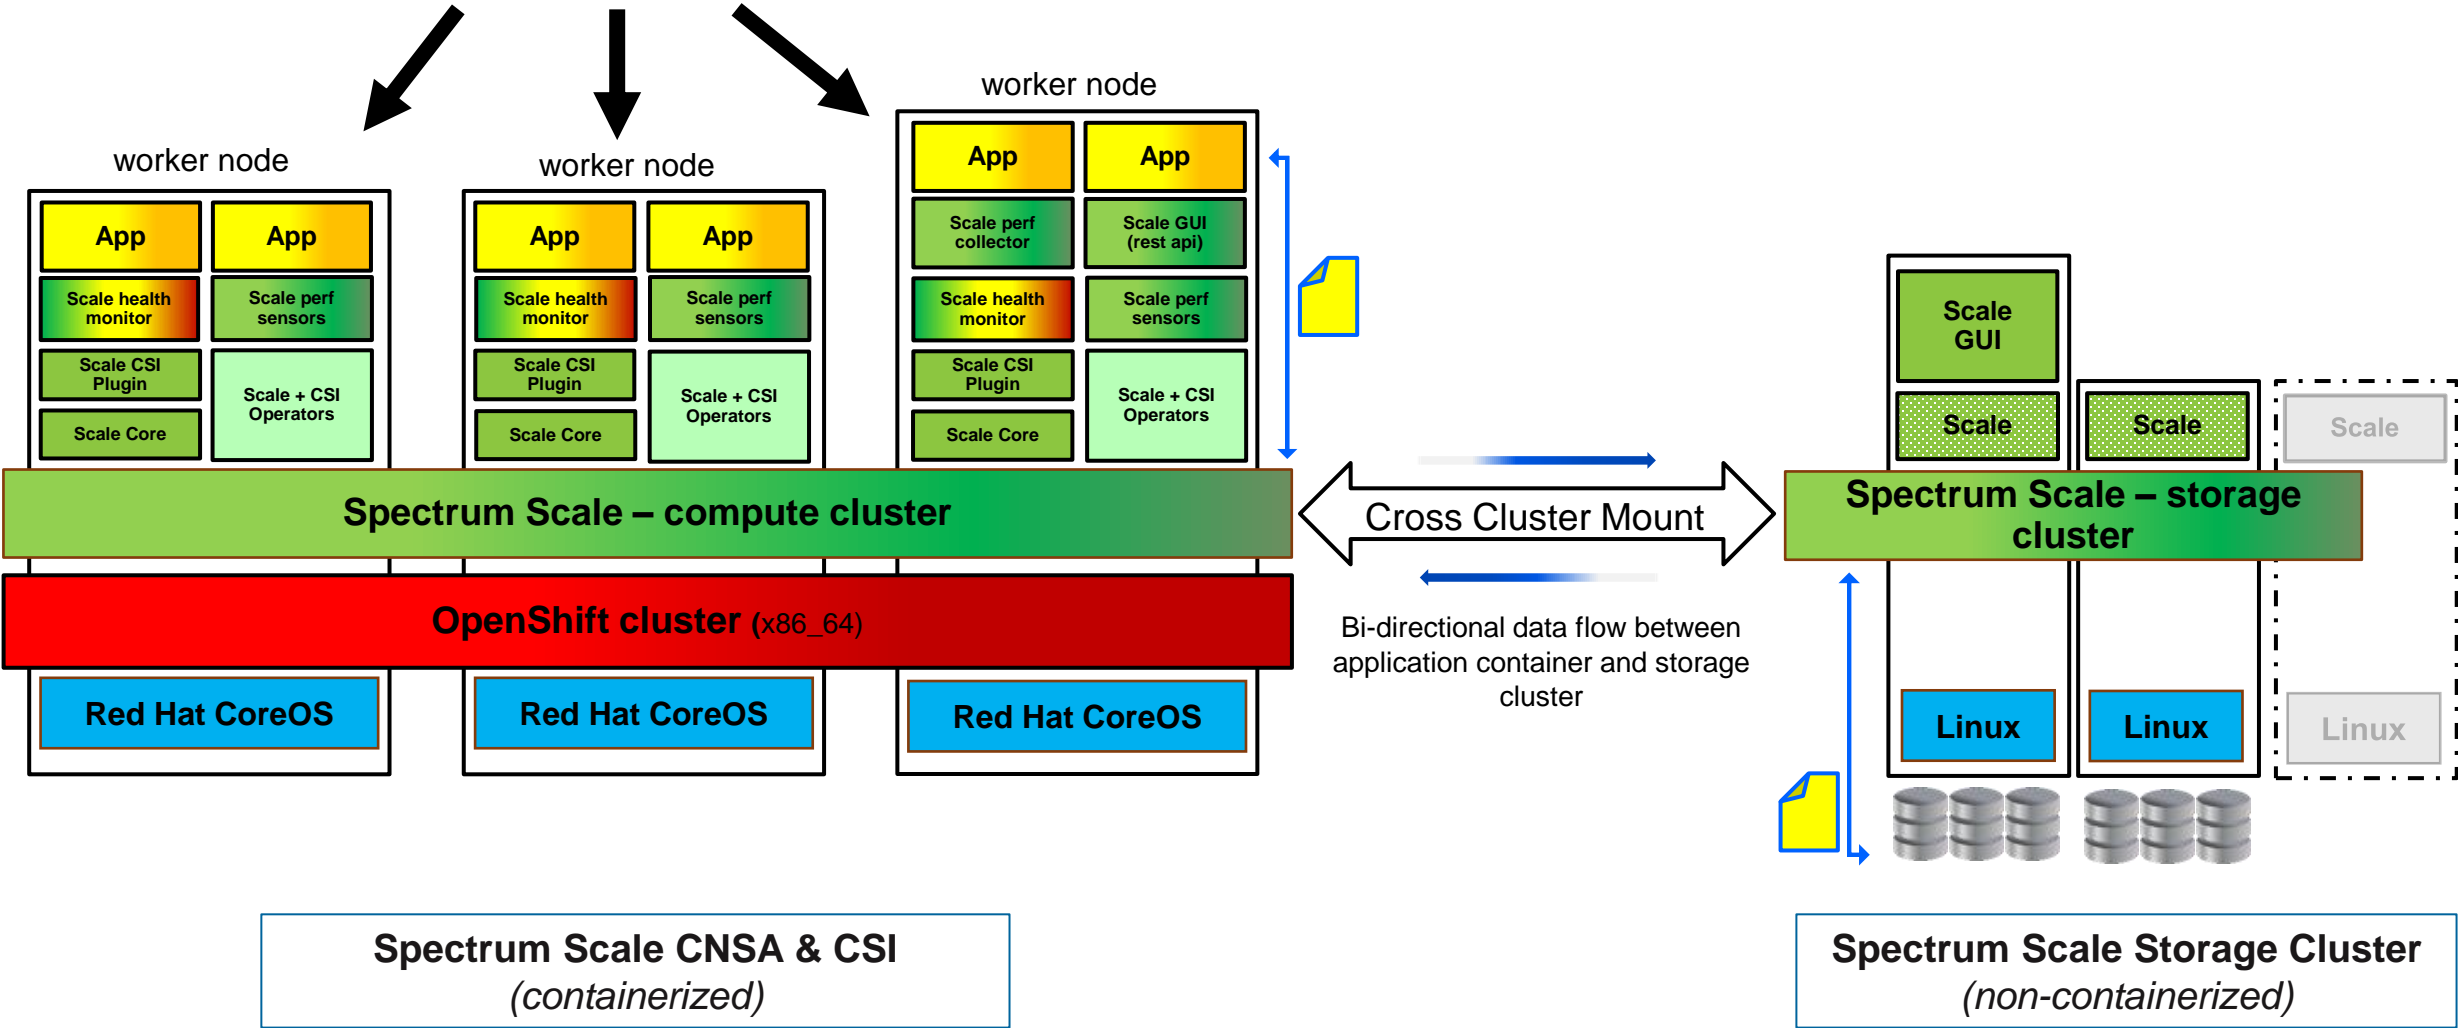
- **IBM Support is the single point of contact for case management**
 - IBM: L1 and L2
 - Escalation with Red Hat if necessary
- Customer receives fixes through IBM Cloud Registry

Spectrum Scale CNSA & CSI



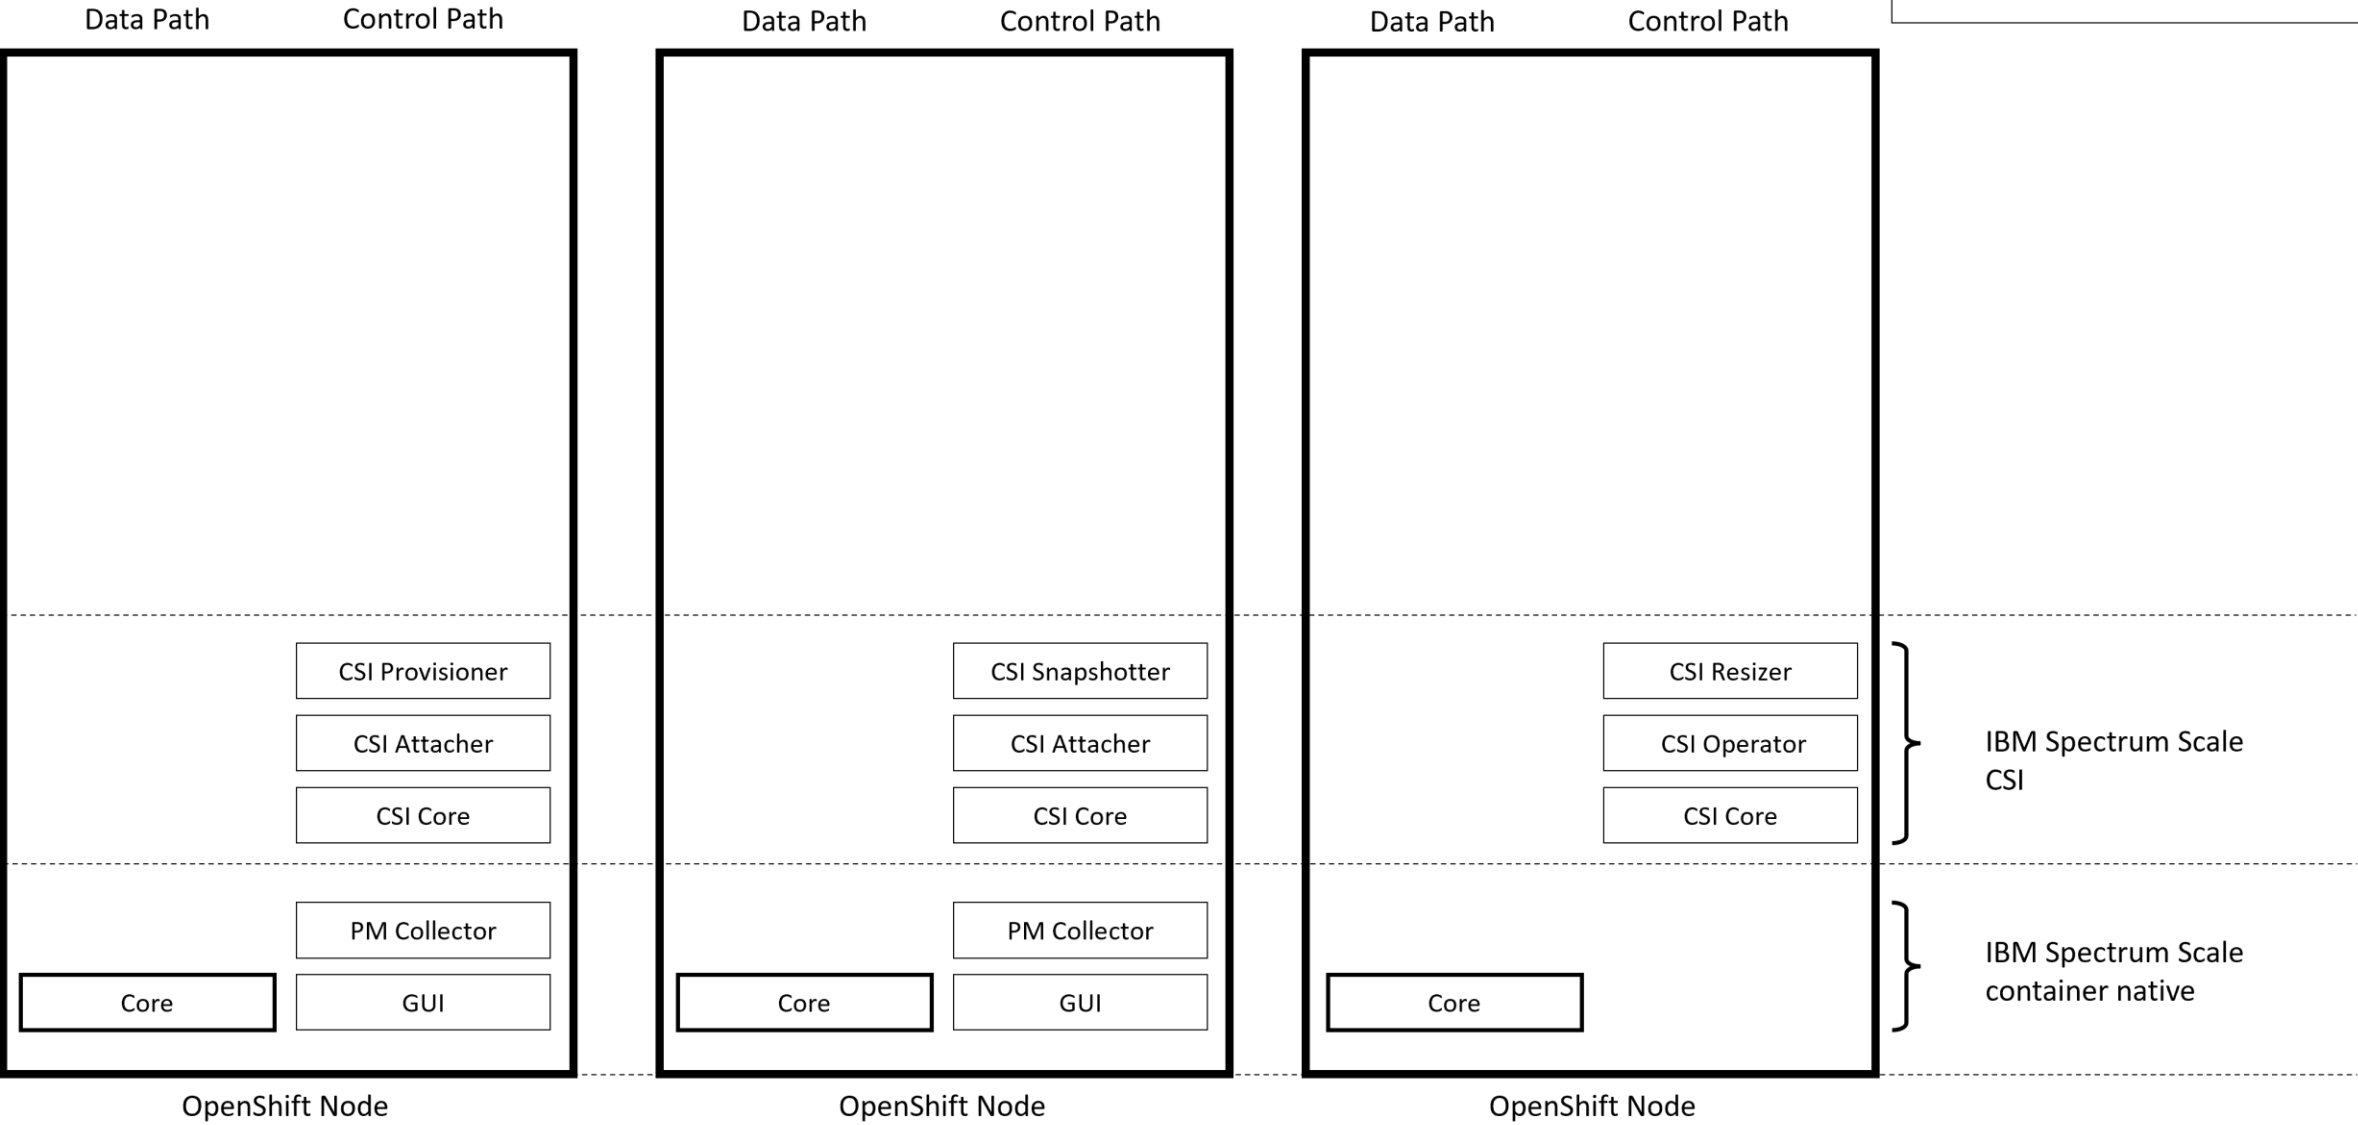
Spectrum Scale CNSA & CSI & DAS S3

The new **Data Access Services (DAS) S3** adds pods to modernize Spectrum Scale's built-in multi-protocol data access for S3.



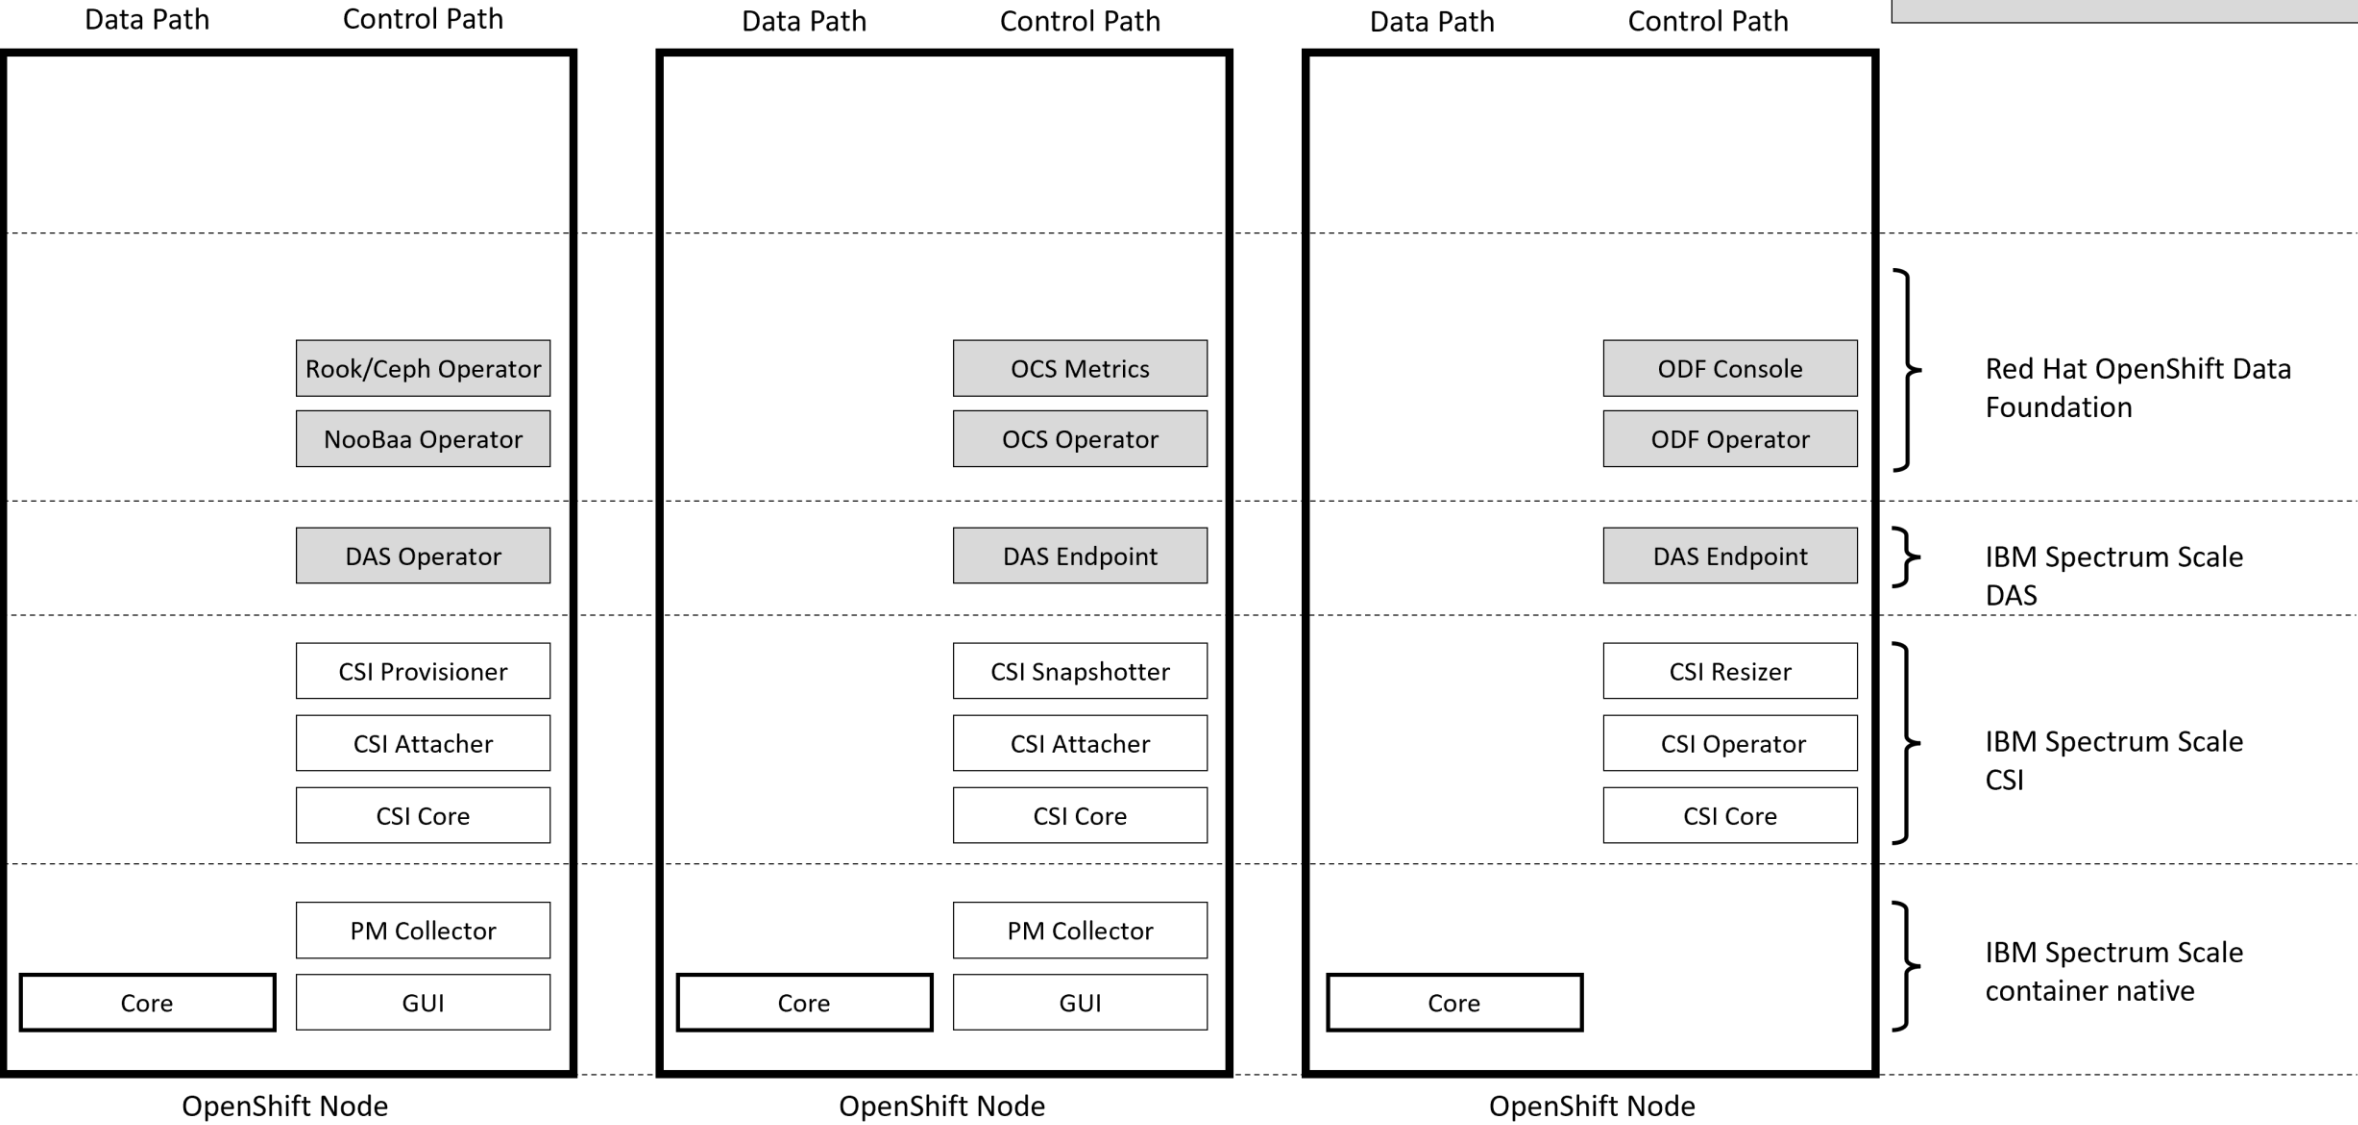
Spectrum Scale CNSA & CSI (prerequisite)

Installation
Prerequisite



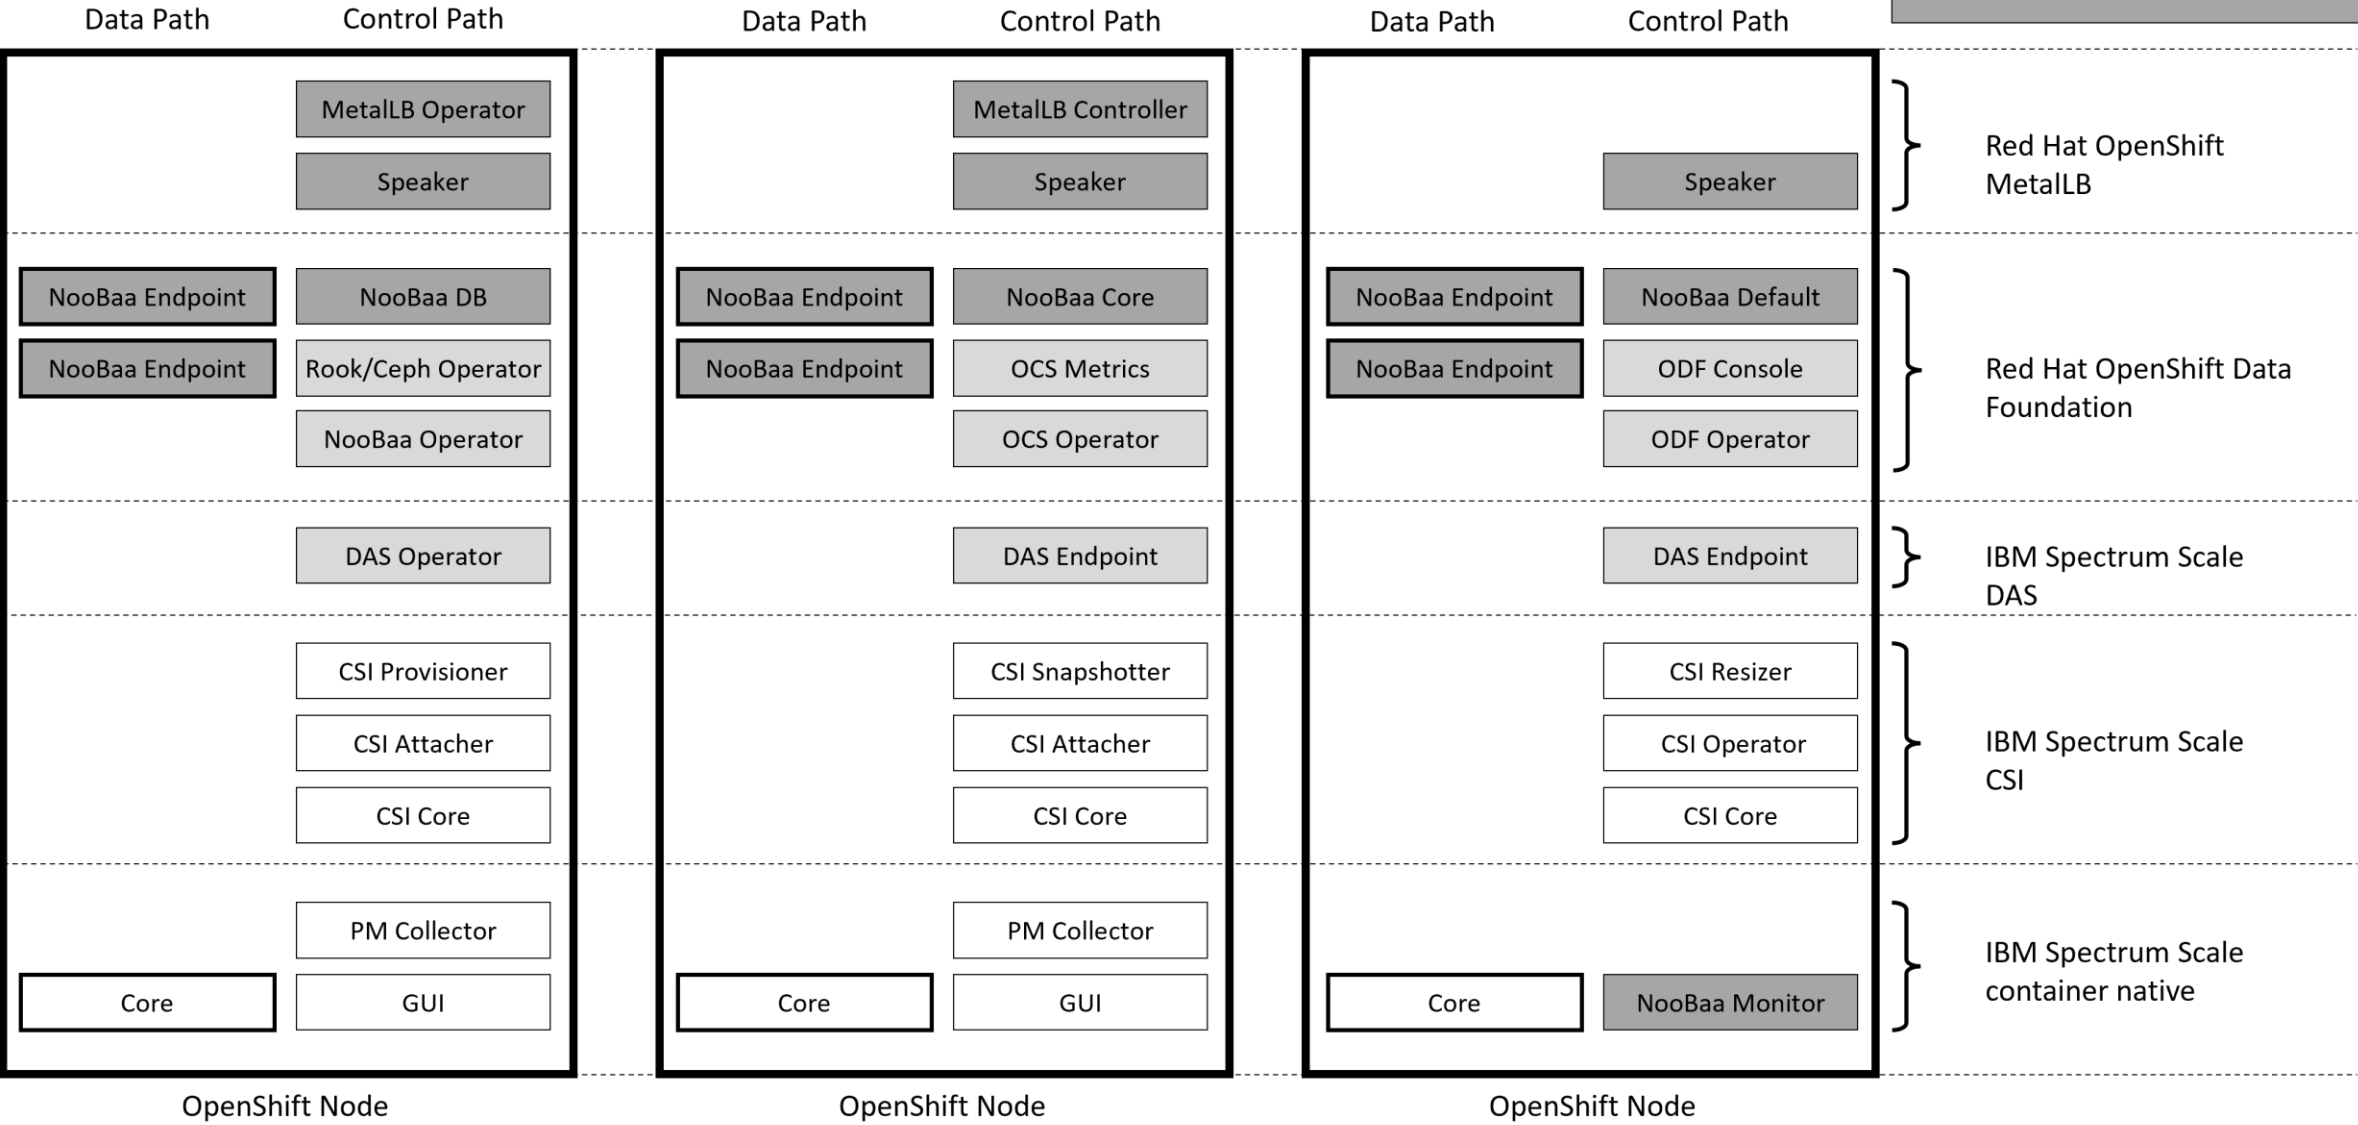
Spectrum Scale DAS S3 (deployed)

```
oc apply -f
ibm-spectrum-scale-
das-operator.yaml
```



Spectrum Scale DAS S3 (configured)

mmdas service create



Thank you for using
IBM Spectrum Scale!