Storage Scale CES S3

-High-Performant Object Access-

IBM Storage Scale User Meeting @ ISC2024

May 13, 2024 | Hotel Reichshof Hamburg

Ulf Troppens

Disclaimer

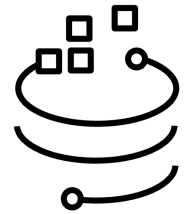

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision. 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.

Agenda

- Storage Scale CES S3
- Evolution of Multi-Protocol Data Access in Storage Scale
- CES S3 Examples
- CES S3 CLI Reference
- CES S3 Installation
- Migrating to CES S3

Storage Scale

Cluster Export Services for S3 (CES S3)

For data-intensive workflows which need S3 object access to data which is stored in Storage Scale filesystems

- Optimized for multi-protocol data access to enable workflows which access the same instance of data using S3 object and other access protocols
  - High-performant and scalable S3 object access to data which is stored in Storage Scale filesystems
  - Support for S3 API calls which are required to process data which is stored in Storage Scale filesystems
  - Files and directories in Storage Scale filesystems are represented 1:1 as S3 objects and S3 buckets
- Use cases include
  - AI and Analytics
  - watsonx
- Replaces CES Swift Object and DAS S3 (HPO)
- Tech Preview with Storage Scale 5.2.0
- GA targeted for 2H24

Multi-Protocol Data Access

Simultaneous multi-protocol access including GPUDirect support

Outcome: Enable globally dispersed teams to collaborate on data regardless of protocol, location or format

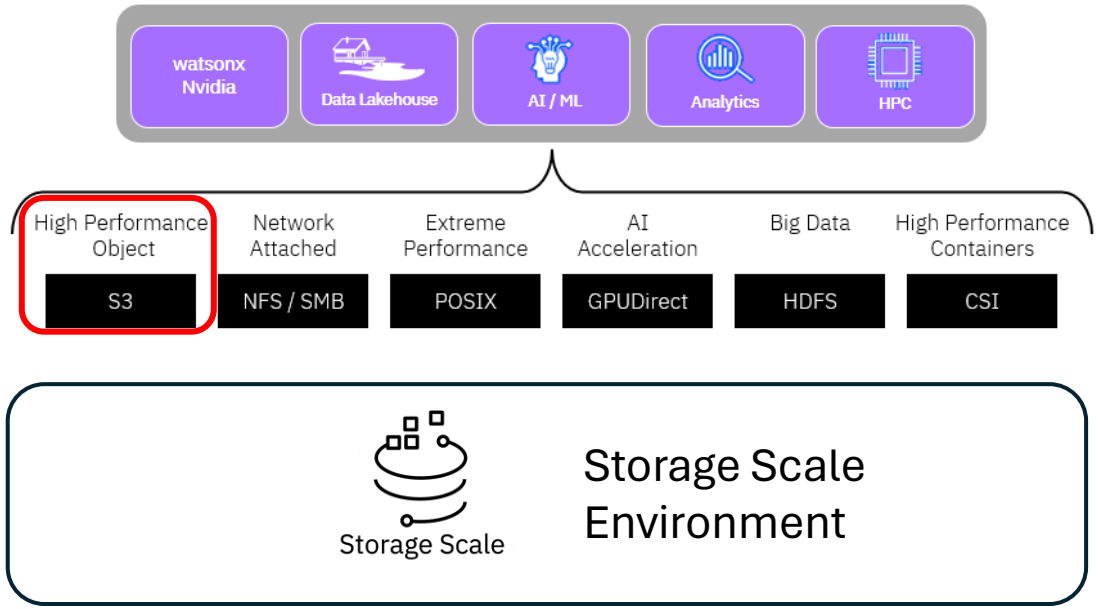

CES S3 Tech Preview

Deployment Requirements:	
Spectrum Scale Cluster:	Spectrum Scale 5.2.0
Operating System:	RHEL8.x or RHEL9.x
Architecture:	x86_64
Spectrum Scale CES Cluster Size:	Up to 4-node CES cluster (test limit)

Scaling limits for CES S3 Tech Preview:

- Up to 10TB single object size
- Up to 5000 S3 accounts
- Up to 5000 S3 buckets
- Up to 4,000,000 objects/bucket

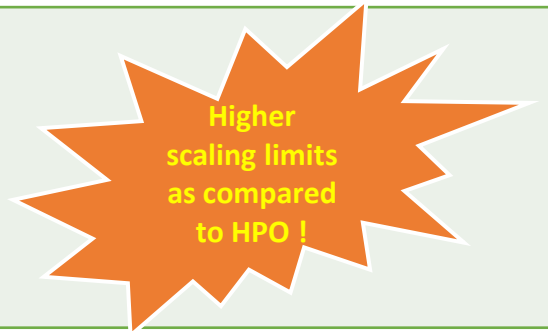

➔ Tech Preview included in Storage Scale 5.2.1
➔ To participate in Tech Preview: https://www.ibm.com/support/pages/node/7145681

Data-Intensive Workflows for Unstructured Data

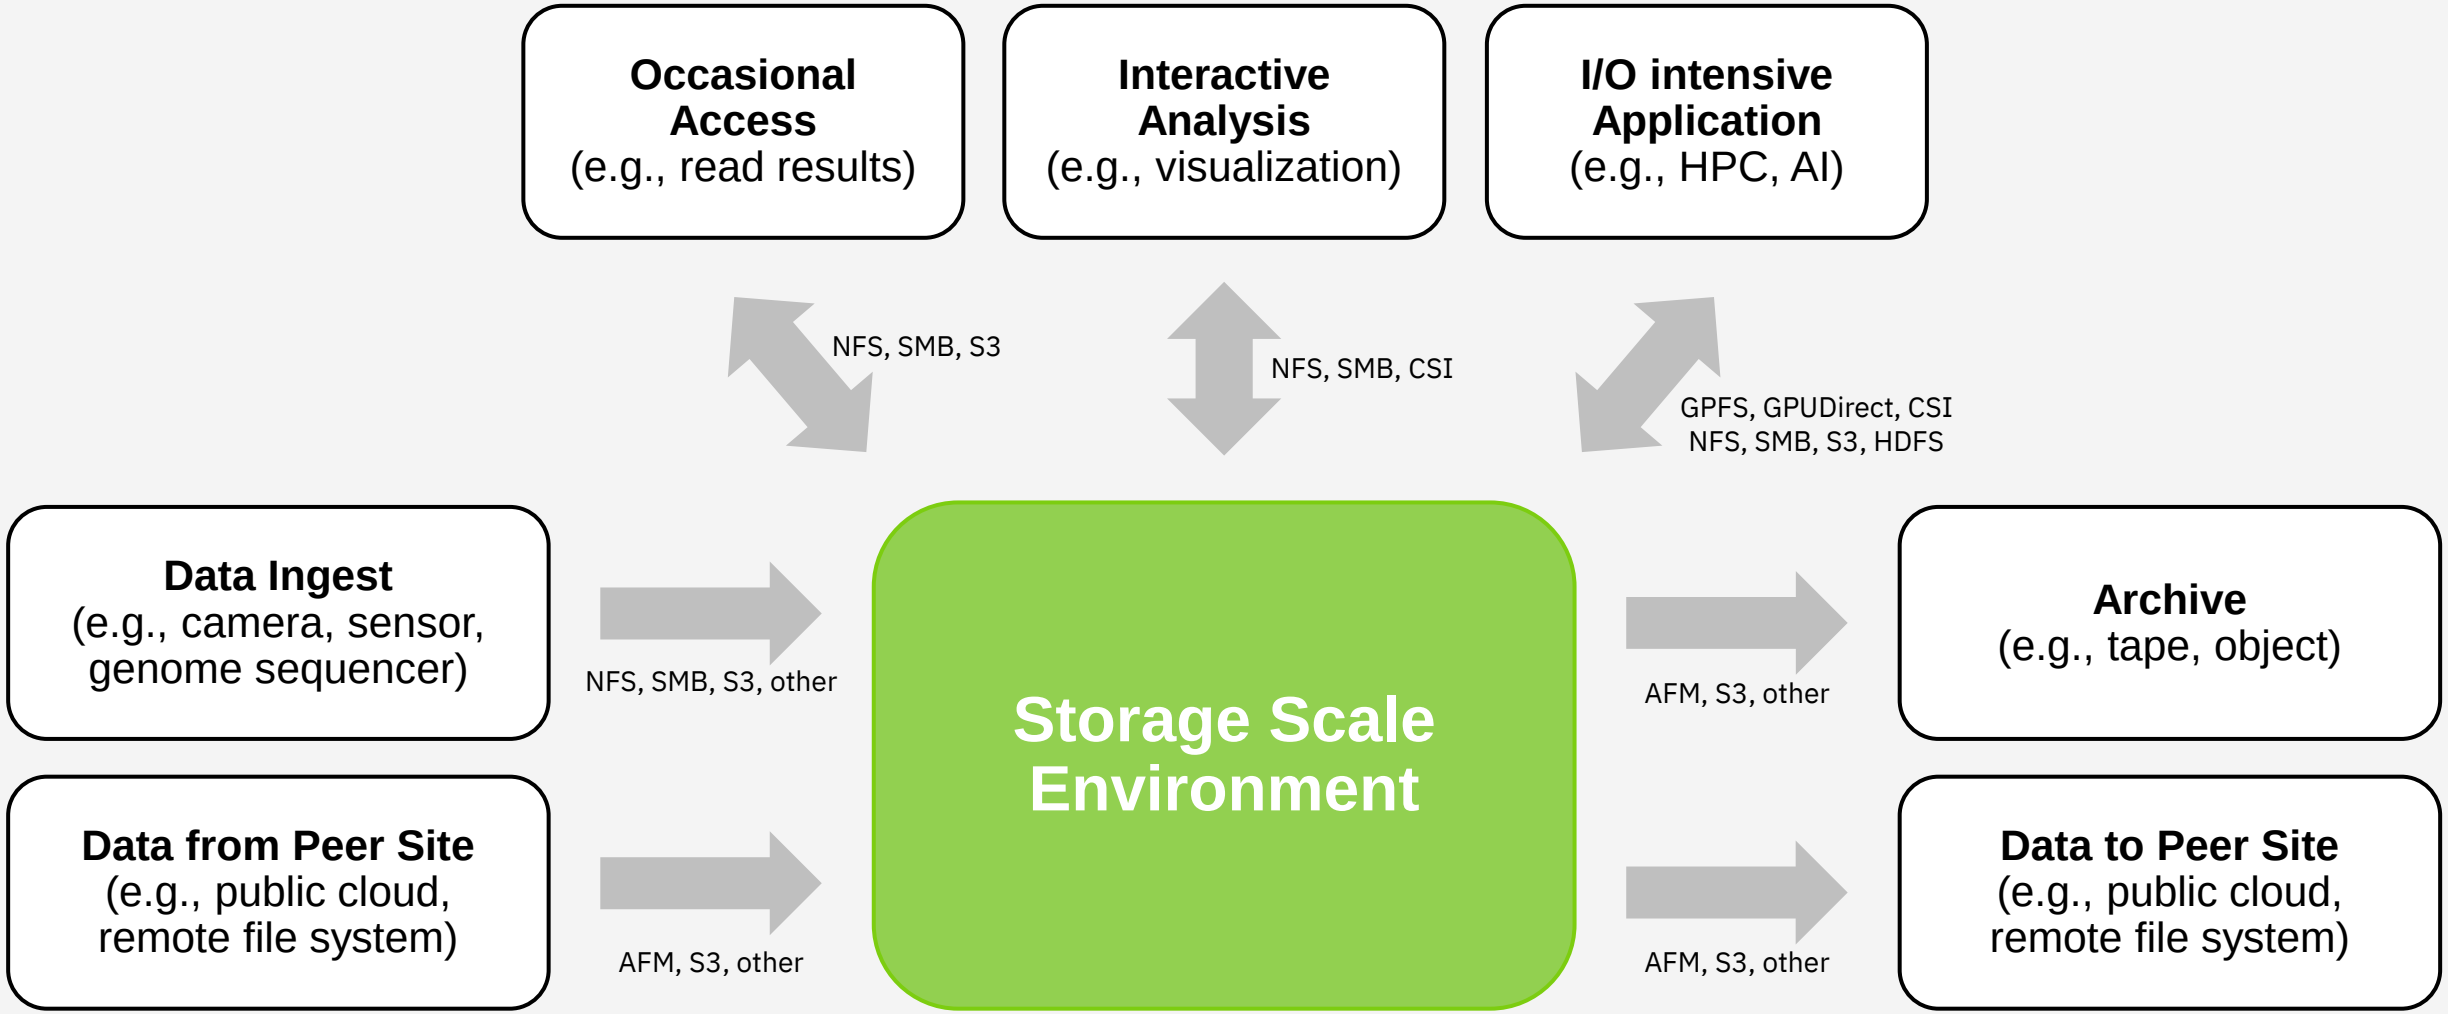

- Many unstructured data is generated and processed outside the data center.
- File and object protocols allows devices and applications to access and process data on remote servers and systems.
- Many of the remote applications and devices stick to one of the many file and object protocols (e.g., genome sequencers).

IBM watsonx.data + Storage Scale (with S3)

Existing Deployment watsonx.data
With Storage Scale SDS

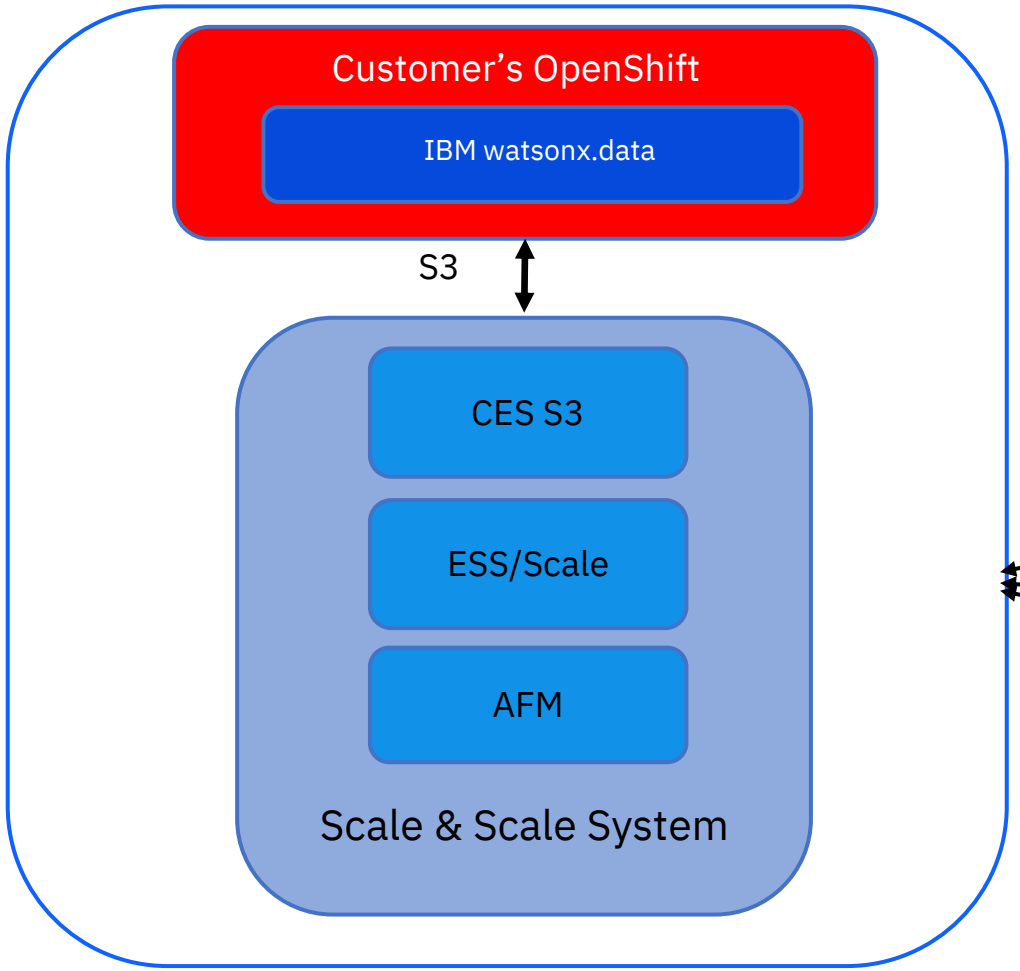

A single source to
access global data

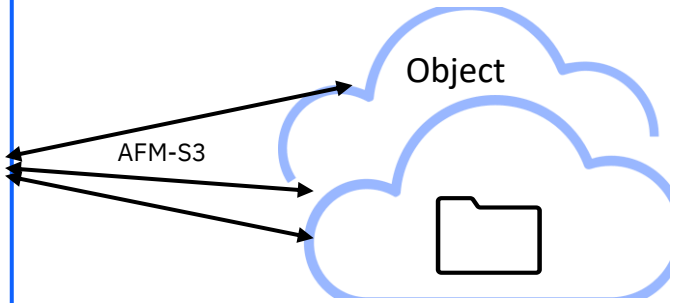

Remote S3 bucket
Ceph and/or another object store
Provided by Scale AFM-> S3

Value Proposition: Storage Scale Provides a Virtualized Caching Layer

to Deliver High Performance Tier to Unstructured Data

IBM Storage Scale

Storage Abstraction and Acceleration

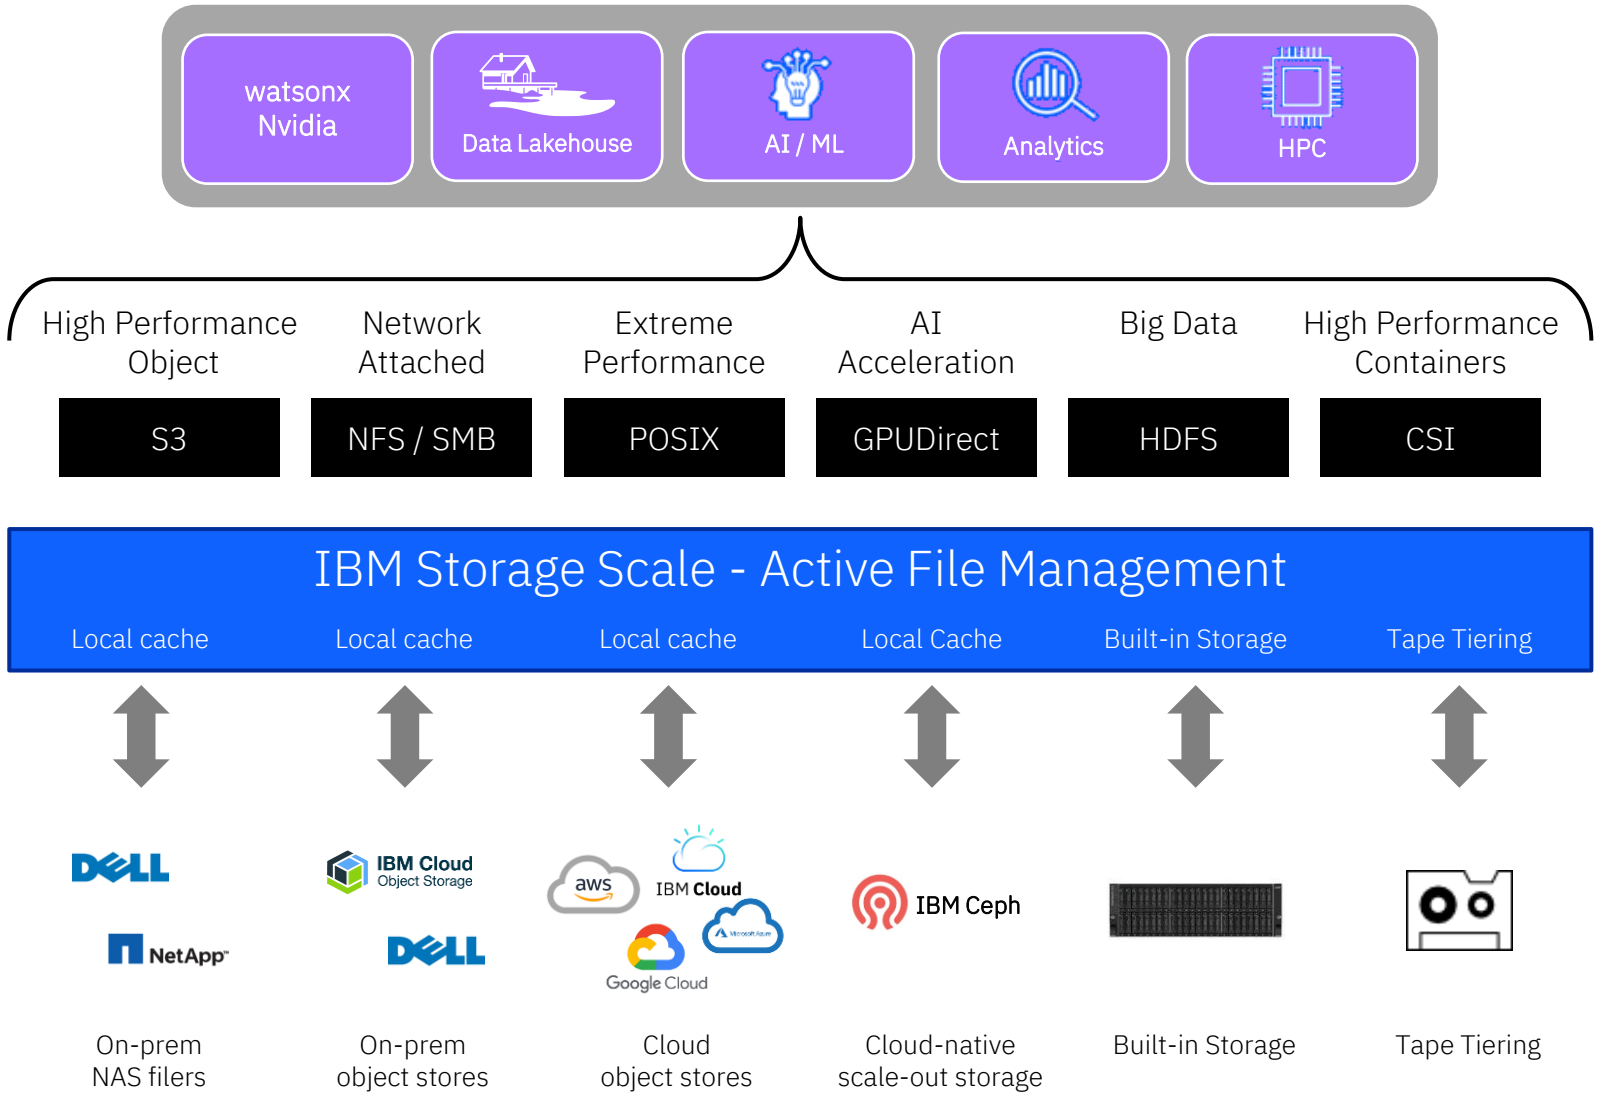

Multi-Protocol Data Access

Simultaneous multi-protocol access including GPUDirect Storage support

Outcome: Enable globally dispersed teams to collaborate on data regardless of protocol, location or format

Storage Acceleration

Automatic, transparent caching of back-end storage systems

Outcome: Accelerates data queries and improves economics by fronting lower performance storage

Storage Abstraction

Single global namespace delivers a consistent, seamless experience for new or existing storage

Outcome: Reduce unnecessary data copies and improve efficiency, security and governance

Storage Scale CES S3

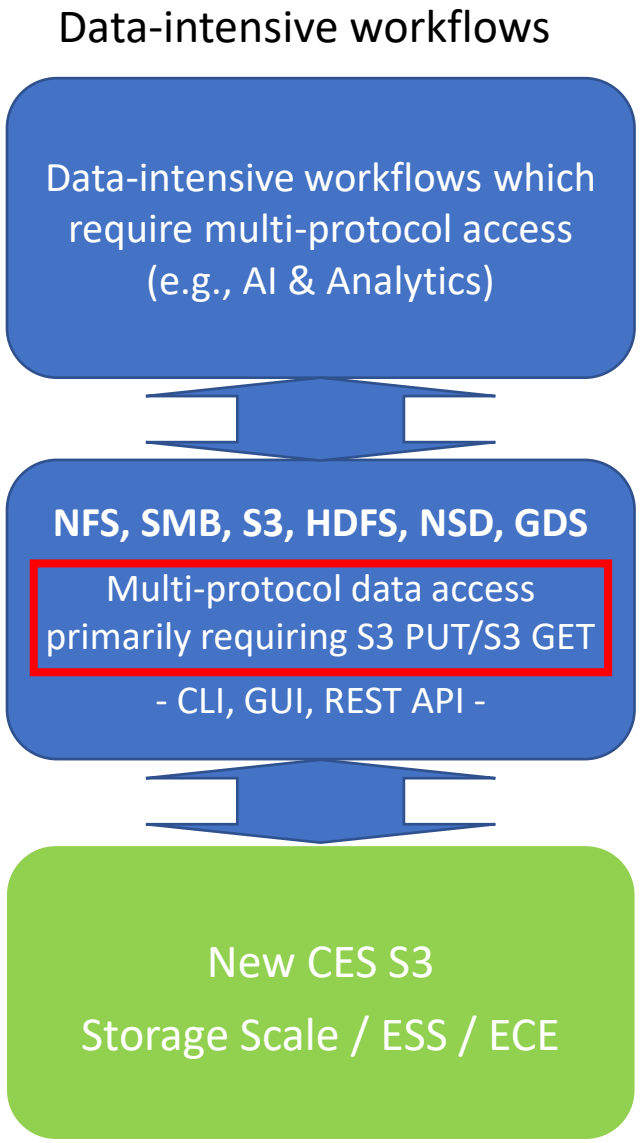

New CES S3

- Strategic approach to provide S3 access to data in Storage Scale
- Will replace CES Swift Object and DAS S3 (HPO)
  ➔ Guidance for migration will be provided
- Integrated with Cluster Export Services (CES)
- Included in DAE, DAE for ESS, DME, DME for ESS, ECE, Standard Edition, Advanced Edition
- Optimized for multi-protocol data access to enable workflows which access the same instance of data using S3 and other access protocols
- Support for S3 API calls which are required to process data which is stored in Storage Scale filesystems
- No support for S3 Rest API calls which are primarily used for S3 only workloads such as S3 Versioning and S3 Locking
- Tech Preview with Storage Scale 5.2.0
- GA targeted for 2H24

Strategic Directions for Storage Scale S3 Access

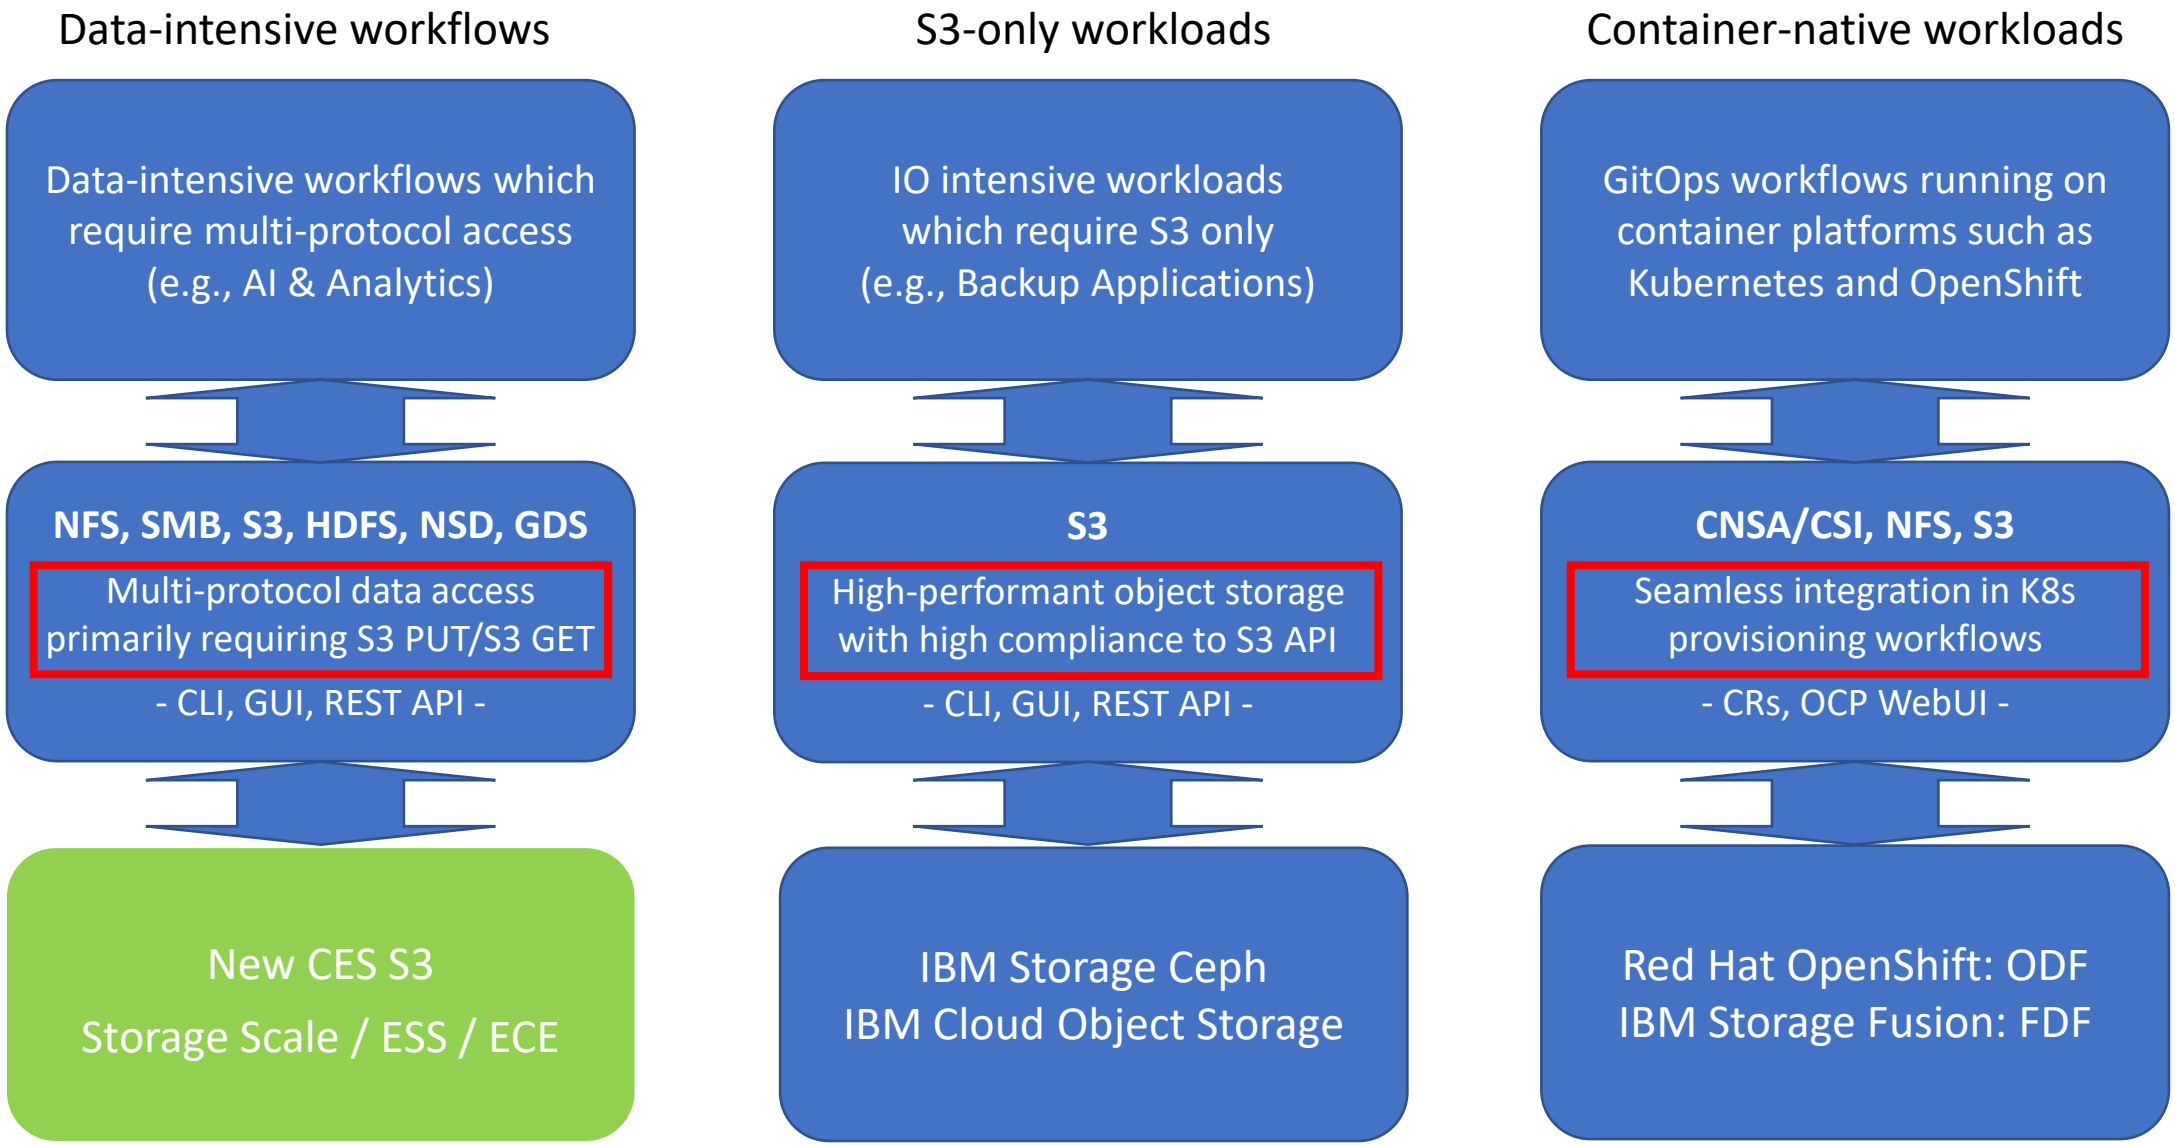

Storage Scale CES S3 Architecture

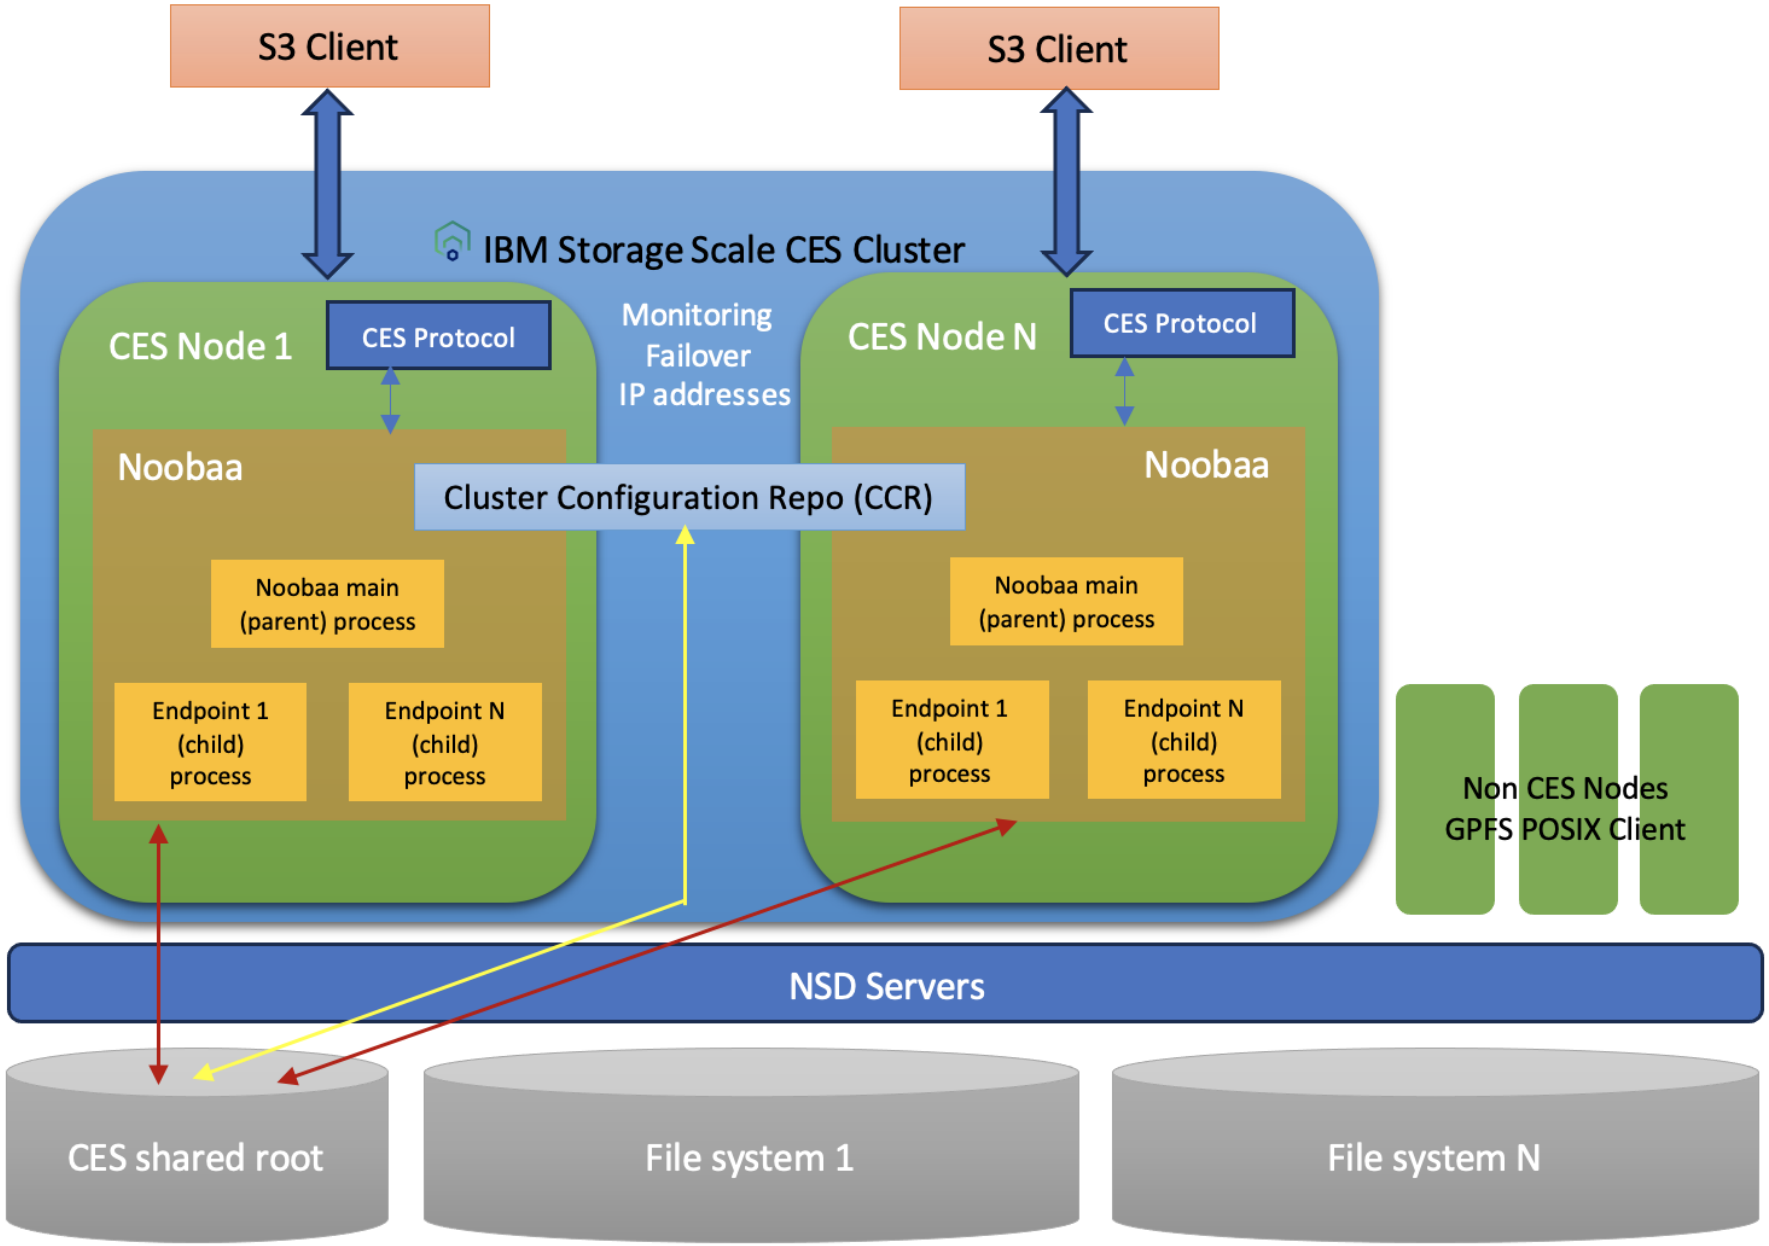

Control Path Architecture

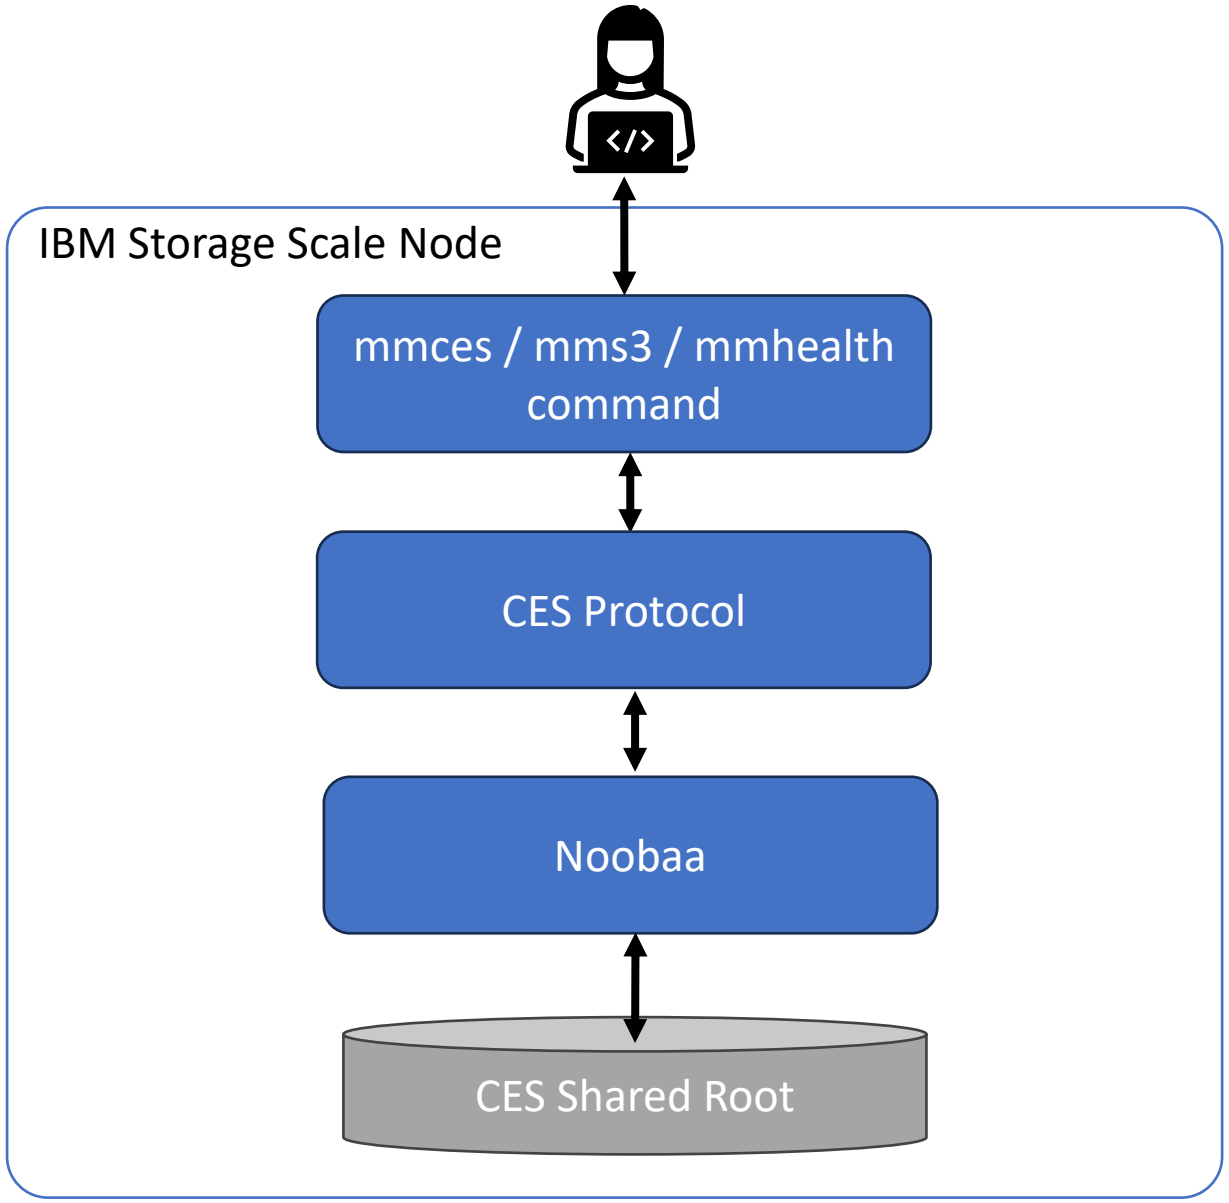

Data Path Architecture

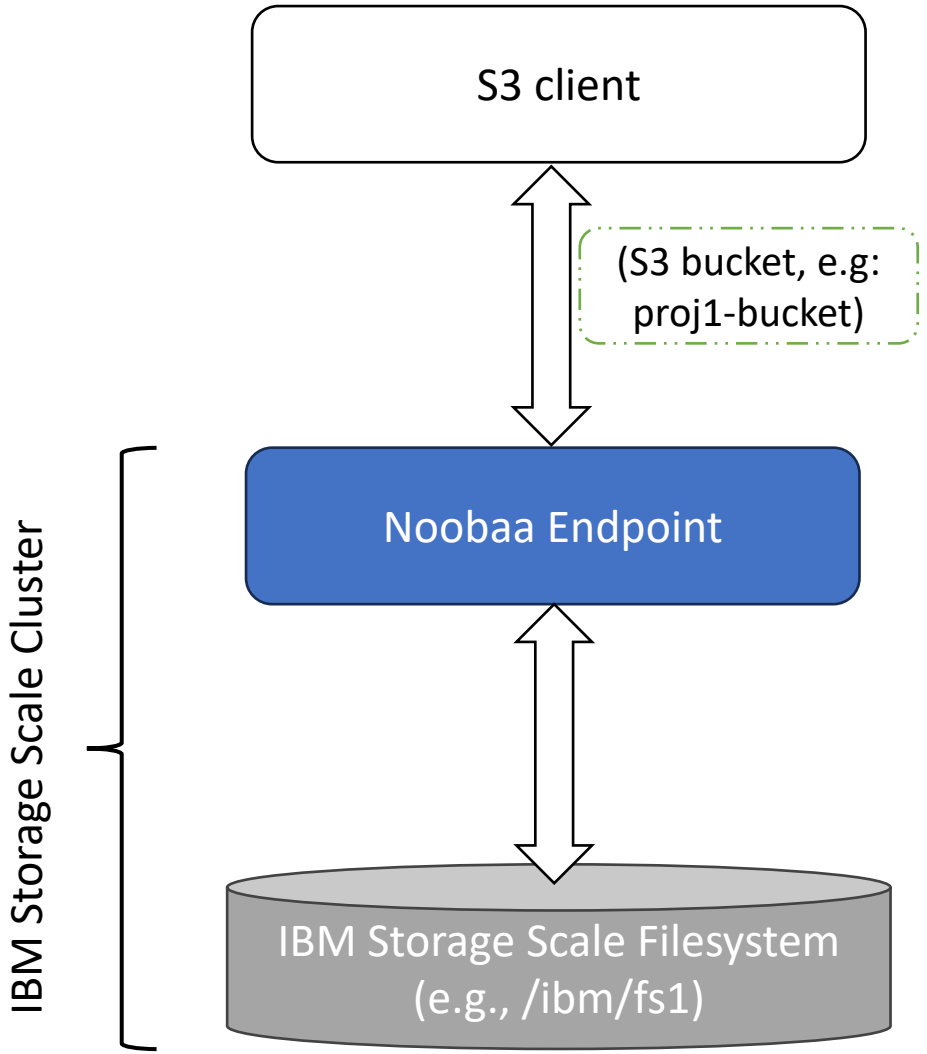

Benchmark Environment

Application Nodes: **6x x86_64 Servers each server**
- **2x CPU AMD EPYC 7F72 24-Core**
- **512GB Memory**
- **Bond – 2x 100GbE 1 port Ethernet**
- RHEL 9.3
- COSBench v0.4.2

CES Nodes: **3x x86_64 servers each server**
- **2x CPU Intel(R) Xeon(R) Gold 6346 CPU 32-Core**
- **256GB Memory**
- **Bond1 – 2x 100GbE 1 port Ethernet (to APP nodes)**
- **Bond2 – 2x 200GbE 1 port Ethernet (to ESS)**
- RHEL 9.3
- GPFS 5.2.0
- Ganesha v5.7
- Noobaa v5.15.0

GPFS Storage cluster – Dedicated ESS3200
- **Bond – 2x 200GbE 1 port from 2 dual port Ethernet card (each canister)**

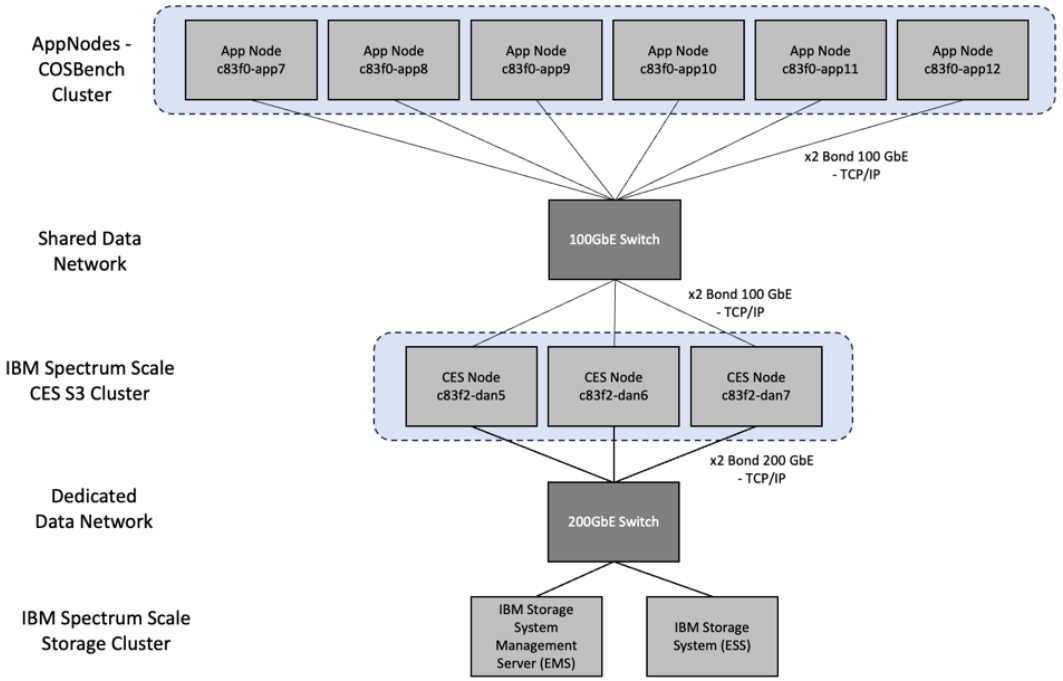

Fig. 1. Benchmark environment.

COSBench application, 12 drivers were created to have in total 2 drivers each APP node.

```
[root@c83f2-dan5 ~]# mmlscluster --ces

GPFS cluster information
========================
  GPFS cluster name:         danrf.gpfs
  GPFS cluster id:           15299468826331659279

Cluster Export Services global parameters
-----------------------------------------
  Shared root directory:                /gpfs/rf-cesshared
  Enabled Services:                     NFS SMB S3
  Log level:                            0
  Address distribution policy:          even-coverage

Node   Daemon node name       IP address       CES IP address list
-----------------------------------------------------------------
   1   dan5ib.gpfs.net        172.20.200.136   172.20.100.50
   2   dan6ib.gpfs.net        172.20.200.137   172.20.100.52
   3   dan7ib.gpfs.net        172.20.200.138   172.20.100.51
```

```
[root@c83f2-dan5 ~]# mms3 config list

 S3 NOOBAA Configuration:
 ========================
 ALLOW_HTTP : true
 DEBUGLEVEL : default
 ENABLEMD5 : false
 ENDPOINT_FORKS : 12
 ENDPOINT_PORT : 6001
 ENDPOINT_SSL_PORT : 6443
 GPFSDLPATH : /usr/lpp/mmfs/lib/libgpfs.so
 NSFS_DIR_CACHE_MAX_DIR_SIZE : 268435456
 NSFS_DIR_CACHE_MAX_TOTAL_SIZE : 805306368
 NSFS_NC_CONFIG_DIR_BACKEND : GPFS
 NSFS_NC_STORAGE_BACKEND : GPFS
 UVTHREADPOOLSIZE : 16
 ========================
```

Blog post with details:
https://community.ibm.com/community/user/storage/blogs/rogelio-rivera-gutierrez/2024/04/25/ibm-storage-scale-performance-ces-s3-tech-preview

Benchmark Workload

The performance benchmark tool used for this evaluation is COSBench. Workload files exercise the whole environment between APP and CES nodes with the least amount of overhead.

Benchmark 1 (Large Objects):

- Buckets: 10
- Objects: 100 (evenly distributed in the 10 buckets)
- Run duration: 5 minutes per Workstage.
- Workers: 1, 8, 32, 64, 128, 256, 512
- Object size: 1GB.
- Operation: 100% read, 100% write

Benchmark 2 (Small Objects):

- Buckets: 10
- Objects: 1000 (per bucket / per object size)
- Run duration: 5 minutes per Workstage.
- Workers: 1, 8, 32, 64, 128, 256, 512
- Object size: 4KB, 32KB, 64KB, 128KB, 256KB 1MB, 4MB, 8MB.
- Operation: 100% read, 100% write

Blog post with details:
https://community.ibm.com/community/user/storage/blogs/rogelio-rivera-gutierrez/2024/04/25/ibm-storage-scale-performance-ces-s3-tech-preview

PERFORMANCE WRITE

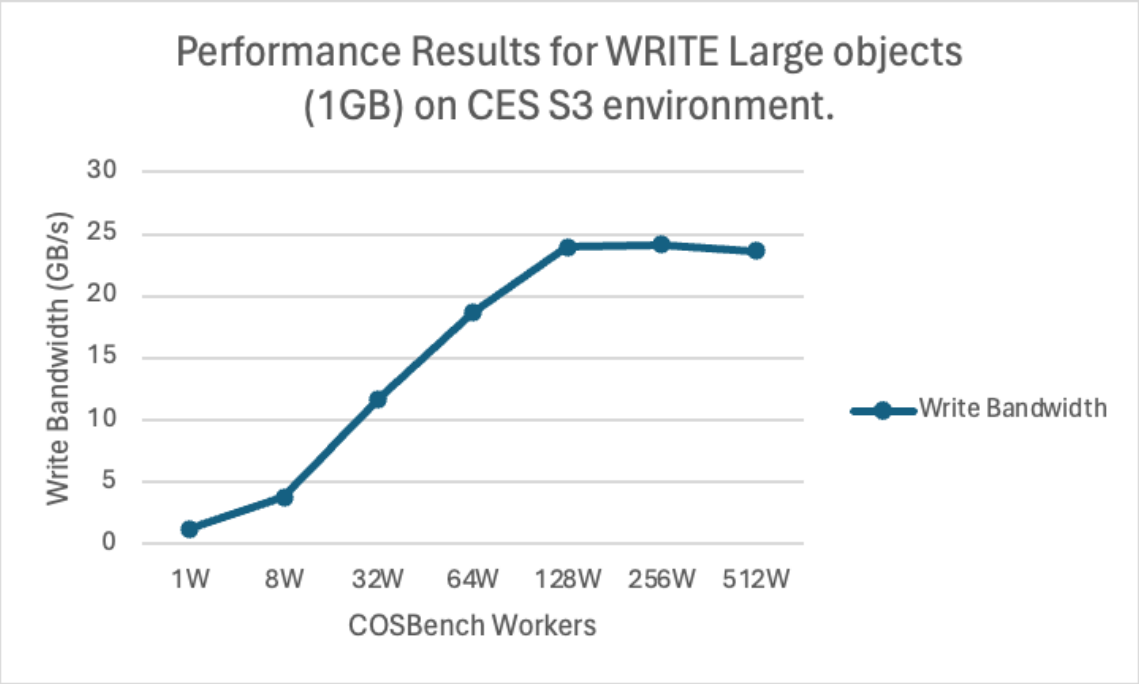

Op-Type	Obj Size	Workers	Op-Count	Byte-Count	Avg-ResTime	Avg-ProcTime	Throughput	Bandwidth	Succ-Ratio
WRITE	1GB	1	357 ops	365.57 GB	839.51 ms	23.96 ms	1.19 op/s	1.22 GB/S	100%
		8	1.11 kops	1.14 TB	2147.99 ms	25.47 ms	3.72 op/s	3.81 GB/S	100%
		32	3.41 kops	3.49 TB	2796.12 ms	46.83 ms	11.43 op/s	11.7 GB/S	100%
		64	5.44 kops	5.57 TB	3508.95 ms	104.65 ms	18.23 op/s	18.66 GB/S	100%
		128	6.96 kops	7.13 TB	5470.89 ms	622.74 ms	23.39 op/s	23.95 GB/S	100%
		256	6.96 kops	7.12 TB	10843.78 ms	1830.04 ms	23.6 op/s	24.17 GB/S	100%
		512	6.67 kops	6.83 TB	22211.44 ms	1973.41 ms	23.05 op/s	23.61 GB/S	100%

Table 1. Performance Results for WRITE Large objects (1GB) on CES S3 envronment.

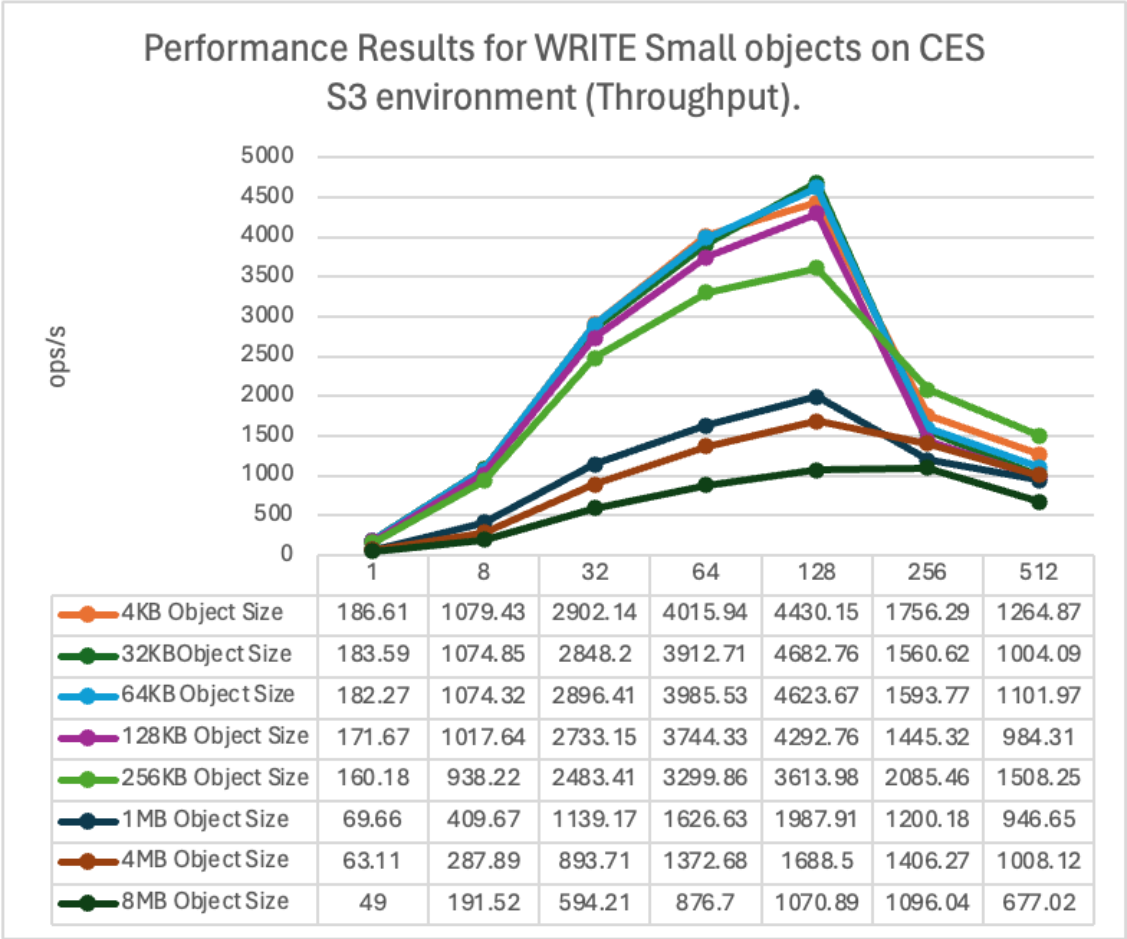

	1	8	32	64	128	256	512
4KB Object Size	186.61	1079.43	2902.14	4015.94	4430.15	1756.29	1264.87
32KBObject Size	183.59	1074.85	2848.2	3912.71	4682.76	1560.62	1004.09
64KB Object Size	182.27	1074.32	2896.41	3985.53	4623.67	1593.77	1101.97
128KB Object Size	171.67	1017.64	2733.15	3744.33	4292.76	1445.32	984.31
256KB Object Size	160.18	938.22	2483.41	3299.86	3613.98	2085.46	1508.25
1MB Object Size	69.66	409.67	1139.17	1626.63	1987.91	1200.18	946.65
4MB Object Size	63.11	287.89	893.71	1372.68	1688.5	1406.27	1008.12
8MB Object Size	49	191.52	594.21	876.7	1070.89	1096.04	677.02

Blog post with details:
https://community.ibm.com/community/user/storage/blogs/rogelio-rivera-gutierrez/2024/04/25/ibm-storage-scale-performance-ces-s3-tech-preview

PERFORMANCE READ

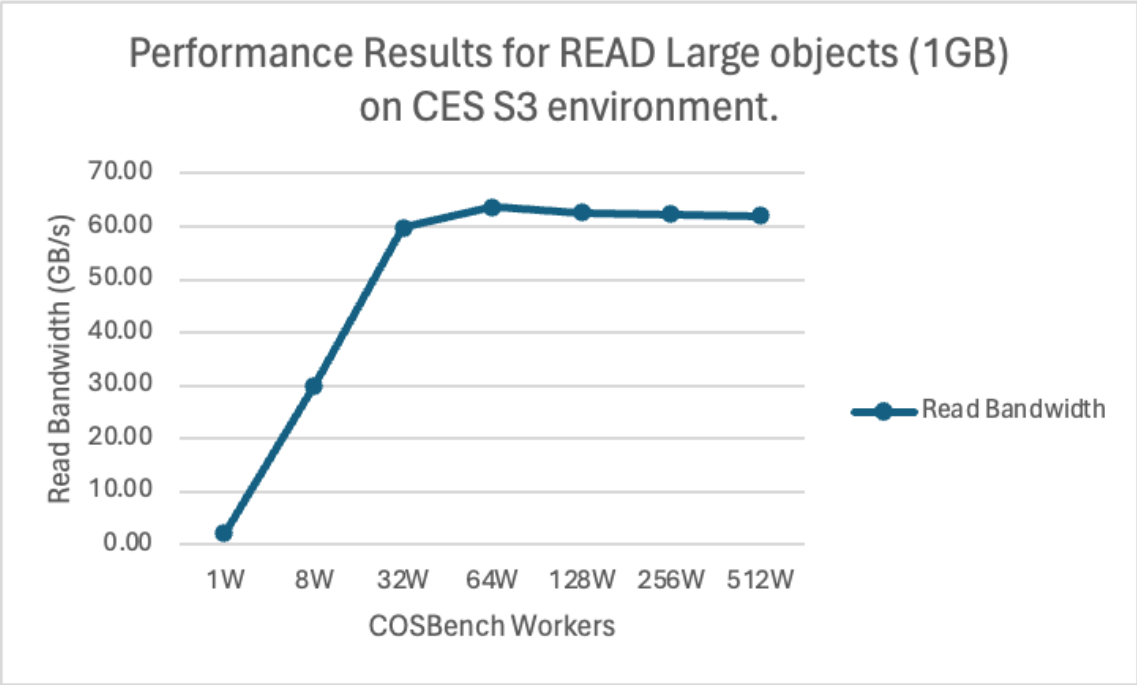

Op-Type	Obj Size	Workers	Op-Count	Byte-Count	Avg-ResTime	Avg-ProcTime	Throughput	Bandwidth	Succ-Ratio
READ	1GB	1	611 ops	625.66 GB	490.86 ms	4.83 ms	2.04 op/s	2.09 GB/S	100%
		8	8.78 kops	8.99 TB	273.26 ms	5.13 ms	29.27 op/s	29.98 GB/S	100%
		32	17.49 kops	17.91 TB	548.53 ms	7.67 ms	58.33 op/s	59.73 GB/S	100%
		64	18.61 kops	19.06 TB	1029.76 ms	15.79 ms	62.15 op/s	63.64 GB/S	100%
		128	18.28 kops	18.72 TB	2093.59 ms	28.6 ms	61.13 op/s	62.6 GB/S	100%
		256	18.12 kops	18.55 TB	4210.39 ms	60.39 ms	60.79 op/s	62.25 GB/S	100%
		512	17.91 kops	18.34 TB	8453.23 ms	106.39 ms	60.55 op/s	62.01 GB/S	100%

Table 2. Performance Results for READ Large objects (1GB) on CES S3 envronment.

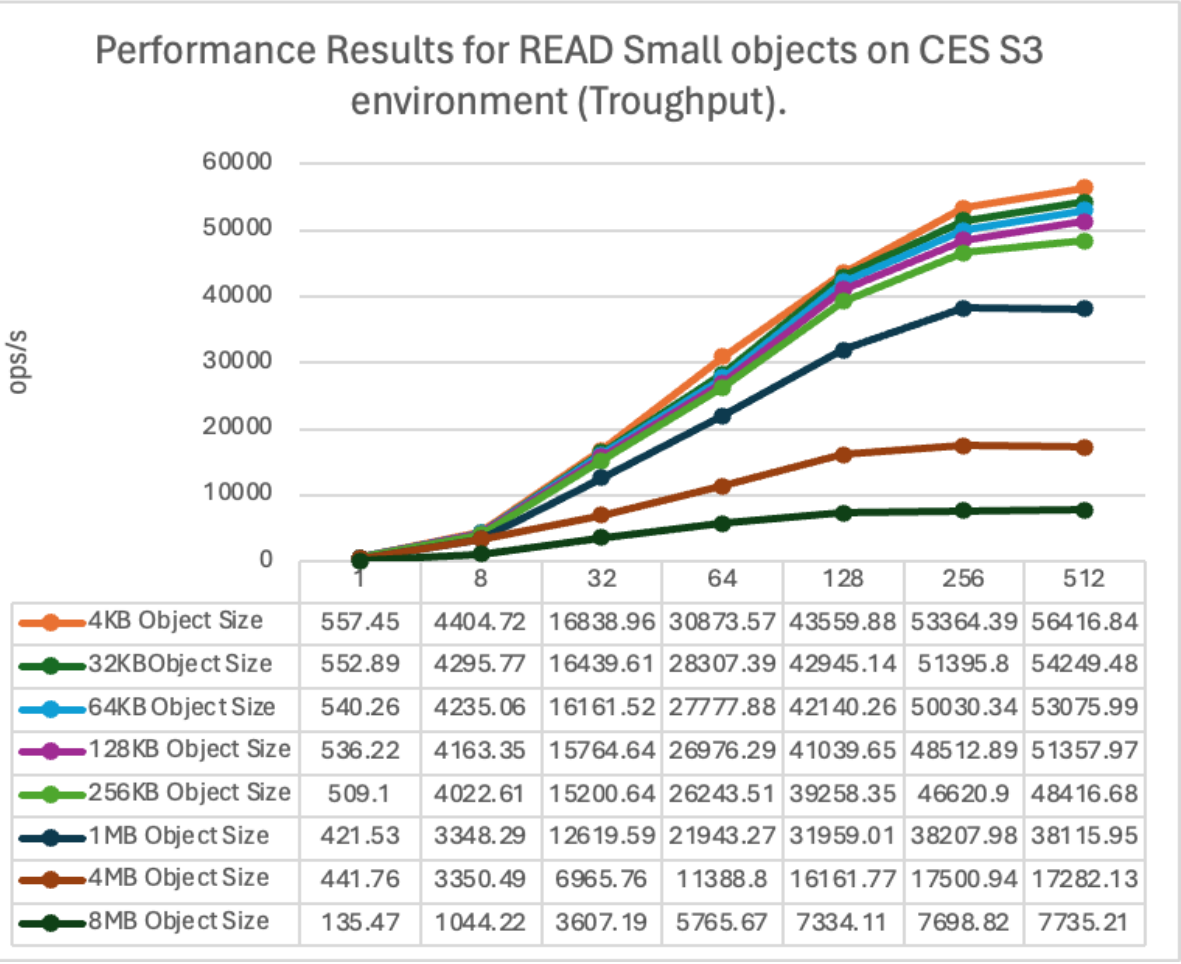

Blog post with details:
https://community.ibm.com/community/user/storage/blogs/rogelio-rivera-gutierrez/2024/04/25/ibm-storage-scale-performance-ces-s3-tech-preview

CES S3 Roadmap

Now (Tech Preview with Storage Scale 5.2.0; GA targeted for 2H24)

Workload
- S3 access to files and directories in Storage Scale filesystems
- S3 PUT, S3 GET, and mandatory additional S3 API Calls
- Multi-protocol data access including CES S3, CES NFS, CES SMB, CSI and POSIX
- S3 access to AFM managed data
- AFM S3 to CES S3
- AI and Analytics
- watsonx.data

Data Management
- Filesets
- Storage Pools
- ILM using NVMe, SSD, NL-SAS
- Backup and restore (mmbackup)

Storage
- Any NVMe, SSD and NL-SAS based storage supported by Storage Scale

Performance Objectives (using 24 NVMe drives)
- Large objects (sequential)
  - 60GB/s read throughput
  - 20GB/s write throughput
- Small objects (random)
  - 30K/s read operations
  - 4K/s write operations

Scaling
- Up to 10TB single object size
- Up to 5000 S3 accounts
- Up to 5000 S3 buckets
- Up to 4,000,000 objects/bucket
- Up to 4-node CES cluster

Deployment
- RHEL 8, RHEL 9
- x86_64
- Linux on Z (RHEL 9 only, no Tech Preview)
- Filesystem with NFSv4 ACLs

Security
- S3 accounts and credentials stored in local repository
- S3 object access protected by SSL certificates
- All access by S3 applications will be authenticated
- ACLs in the filesystem will be enforced
- Object Etags based on mtime and inode, or MD5 checksums

License
- Included in DAE, DAE for ESS, DME, DME for ESS, ECE, Standard Edition, Advanced Edition, Developer Edition

Under consideration for future

Workload
- S3 Select

Data Management
- Tiering to tape

Deployment
- Rolling Update
- Linux on Power

Scaling
- Improve scaling pending customer feedback

Performance
- Improve performance

Security
- AD/LDAP

New CES S3 – Details

Envisioned product capabilities for CES S3 GA. This is work in progress and might change. In particular the claims on supported operating systems and versions, and the support for CPU architectures might change.

Integrated in Storage Scale Cluster Export Services (CES)
- Based on NooBaa, using Namespace Filesystem (NS-FS)
- S3 objects and S3 buckets are mapped 1:1 to files and directories in Spectrum Scale filesystems and vice versa

Workload
- AI and Analytics, watsonx.data
- Multi-protocol data access including CES S3, CES NFS, CES SMB, CSI and POSIX
- S3 access to AFM managed data

Scaling
- Up to 10TB single object size
- Up to 5000 locally managed S3 accounts
- Up to 5000 S3 buckets
- Up to 100,000,000 objects per S3 bucket

Performance
- Basic load distribution (e.g., DNS Round Robin)
- Current benchmark environment
  - 1x ESS3200 with 24x NVMe drives
  - High-Speed Ethernet
- Sequential access using COSBench along with 1GB object
  - 60GB/s read throughput
  - 20GB/s write throughput
- Random accusing using COSBench along with small objects
  - 30K/s reads operations
  - 4K/s write operations
- Chart deck with benchmark environment and benchmark results
  - Architecture shall show different networks for Storage Scale (dedicated private network) and S3 access (shared data center network)

Deployment
- Up to 4-node CES Cluster
- No dependency to OpenShift
- Remote mounted to Storage Cluster or part of Storage Cluster
- OS: RHEL8.x or RHEL 9.3 or later (latest RHEL 9 version which is supported at CES S3 GA)
- CPU Architecture: x86_64
- Underlying filesystem configured with NFSv4 ACLs

Storage
- Any NVMe, SSD and NL-SAS based storage supported by Storage Scale

Resilience
- Backup- and Restore procedure in case of CES Cluster loss
- HA against CES node failures

Security
- Audit logging of configuration changes
- Check summing of data in-flight
- Replication and erasure encoding for data at-rest.
- Sensitive data such as the S3 access credentials are transferred and stored encrypted
- S3 object access is protected by SSL certificates.
- All access by S3 applications will be authenticated and ACLs in the filesystem will be enforced in subsequent data access
- Object Etags based on mtime and inode (default) or MD5 checksums

Installation
- Supported by installation toolkit

New CES S3 – Details (continued)

Envisioned product capabilities for CES S3 GA. This is work in progress and might change. In particular the claims on supported operating systems and versions, and the support for CPU architectures might change.

Management
- Basic integration in Storage Scale management framework
- New CLI commands and REST API to manage CES S3
- mmhealth, gpfs.snap and Call Home must include CES S3 related configuration and status
- Note: With MVP GA mmperfmon will not support metrics for S3

S3 Service
- CLI and REST API to create, delete, enable, disable and report status of S3 service
- CLI and REST API to configure IP addresses for S3 access
- CLI and REST API to enable (default) and disable automatic fail-over and fail-back of IP addresses
- CLI and REST API to enable and disable (default) MD5 based ETags

S3 Accounts
- Stored in local repository
- Map S3 access keys (account key, account identifier) to UIDs and GIDs
- CLI and REST API to create, delete and list S3 Accounts
- CLI and REST API to configure default path for new S3 buckets which are created via the S3 CreateBucket request

S3 Buckets
- Each S3 bucket is mapped to one directory and vice versa
- CLI and REST API to create, delete and list S3 Buckets
- S3 Rest API requests: S3 CreateBucket, S3 ListObjects, S3 ListObjectsV2, S3 DeleteBucket, S3 HeadBucket, S3 ListBuckets, S3 ListMultipartUploads

S3 Objects
- Each S3 object is mapped to one file and vice versa
- S3 Rest API Requests: S3 PutObject, S3 GetObject, S3 HeadObject, S3 CopyObject, S3 DeleteObject, S3 DeleteObjects, S3 CreateMultipartUpload, S3 CompleteMultipartUpload, S3 AbortMultipartUpload, S3 UploadPart, S3 UploadPartCopy, S3 ListParts
- Allows S3 applications to store user defined object metadata in addition to the object data itself.

Data Management
- Filesets
- Storage Pools
- ILM using NVMe, SSD and NL-SAS
- Backup and restore using mmbackup

Migration from CES Swift Object and DAS S3 (HPO) to new CES S3
- Expertly engineered migration procedure to DAS S3 (HPO)
- In-place data migration (=no copying of data), in case -swiftonfile is used
- AFM or NooBaa based migration for customers who can’t do in-place migration
- Expert Lab offering
- Chart deck for sellers

S3 Service
- watsonx.data – S3 interface for Scale

IBM Storage
Quick Reference Guide

Perform Migration from IBM Storage Scale SWIFT/S3 to Storage Scale CES S3

Overview

This service will perform the migration from IBM Storage Scale SWIFT/S3 to Storage Scale CES S3, including authentication and data migration and reconfiguration

Perform Storage Scale Software upgrade is a pre-requisite requirement to a minimum version of 5.2.x, orderable as a separate service offering

Custom SOW
Implement
Data and AI

IBM Technology
Expert Labs

Target Audience

- Existing Storage Scale SWIFT/S3 Clients
- Clients who need to prepare and plan for the deprecation of Storage Scale SWIFT/S3
- Clients who are currently on other S3 services and want to migrate to Storage Scale CES S3

Why Use This Service?

- Reduce the risks of the migration from SWIFT/S3 to Storage Scale CES S3
- Leverage best practices and knowledge from IBM Expert Labs services

Benefits

- This service delivers file system design and implementation with Storage Scale CES S3
- It helps with the seamless integration of Storage Scale CES S3 into your infrastructure.
- It provides knowledge and understanding of advanced functions and use cases for Storage Scale CES S3

Activities

- Assess the current infrastructure requirements and pre-requisites for new Storage Scale CES S3 environment
- Consult on the deployment methods and mode of operation of Storage Scale CES S3
- Advise on Storage Scale CES S3 protocols and gateways including external load balancing
- Implement Storage Scale CES S3 into Client infrastructure
- Document configuration As-Built

Deliverables

- Documentation of current architecture and cluster configuration
- Recommended actions based on known best practices
- Implement Storage Scale CES S3
- Migrate from Storage Scale SWIFT/S3 to Storage Scale CES S3
- Skills enablement and guidance on best practices for Storage Scale CES S3

Scope

- Migration of IBM Storage Scale SWIFT/S3 to Storage Scale CES S3
- Duration: Variable based on environment assessment i.e. data volume and deployment mode

Contacts

- Contact us at systems-expert-labs@ibm.com or your local Technology Expert Labs team

IBM Storage
Quick Reference Guide

Perform Upgrade IBM Storage Scale Software

Overview

The Upgrade Storage Scale Software service is designed to help you upgrade up to 10 Storage Scale nodes in a single cluster, based on a jointly developed plan and targeted use cases

Benefit from software advances since first implementation. Use the latest efficiency, performance and automation features.

Standard Offering
Implement
Data and AI

IBM Technology Expert Labs

Target Audience

- Clients that would like to upgrade their Storage Scale software
- Clients reaching end-of-support on current software levels
- Clients wanting to migrate from Storage Scale SWIFT/S3 to Storage Scale CES S3

Why Use This Service?

- To maintain currency of the Storage Scale Software
- To mitigate end-of-support or end of life on the current software version you are running
- To eliminate code interoperability issues between Storage Scale and application software levels

Benefits

- Expertise to upgrade your Storage Scale software
- Ensure all upgrade options are understood, including code dependencies and the effects that these upgrades will have on other Storage Scale cluster components
- Skills and best practices knowledge transfer

Activities

- Review current Storage Scale software state to prepare for the upgrade
- Plan upgrade per client operational requirements
- Upgrade Storage Scale for up to ten(10) Storage Scale nodes in a single cluster
- Verify upgrade of Storage Scale software and ensure that normal operation successfully resumes
- Document post upgrade environment

Deliverables

- Confirmation of successful upgrade of Storage Scale software and an end of engagement document

Scope

- Limited to up to 10 Storage Scale nodes in a single Storage Scale cluster
- Duration: ~1 week

Contacts

- Contact us at systems-expert-labs@ibm.com or your local Technology Expert Labs team

Agenda

- Storage Scale CES S3
- Evolution of Multi-Protocol Data Access in Storage Scale
- CES S3 Examples
- CES S3 CLI Reference
- CES S3 Installation
- Migrating to CES S3

The world is changing …

2005

Luca Bruno/AP

The world is changing …

2005

Michael Sohn/AP

Luca Bruno/AP

2013

Worldwide File- and Object-Based Storage Forecast, 2021-2027

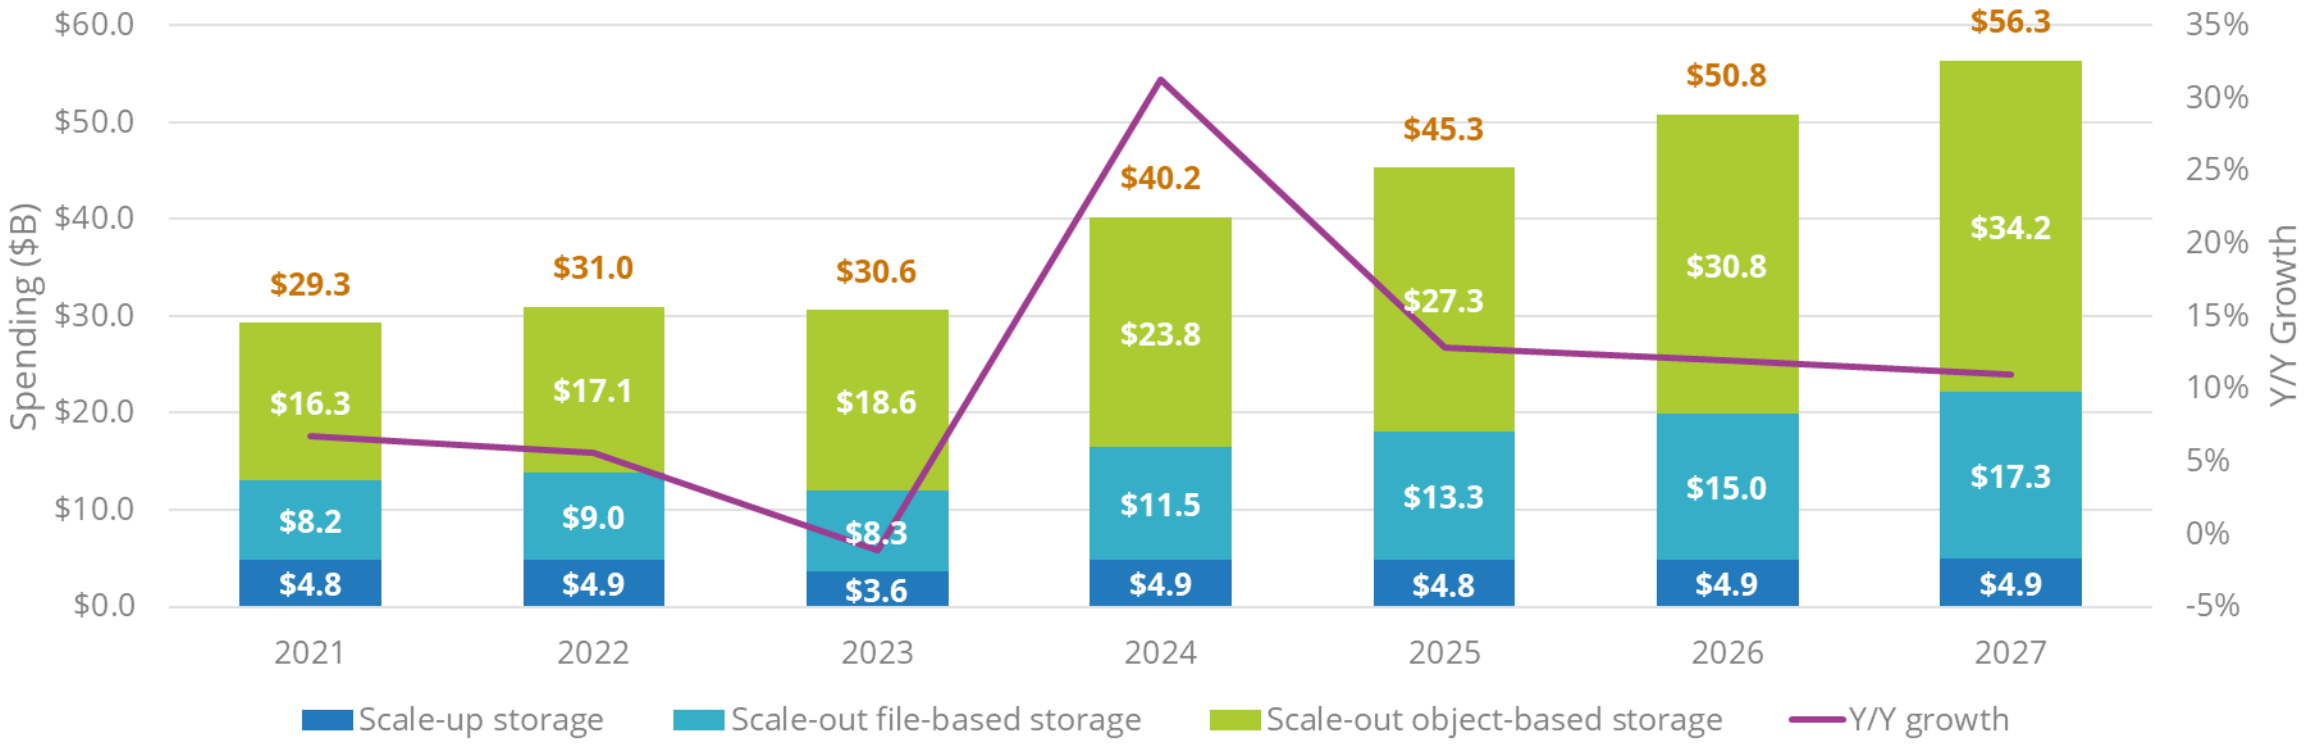

Comments:

- Scale-out file storage (13.9% CAGR) and object storage (14.9% CAGR) will increase in line with unstructured data growth. Drivers include GenAI/AI/ML, analytics, data lakes, IoT, cloud-native applications, and backup and archive.
- Scale-up file storage (0.2% CAGR) will remain roughly flat as demand continues to shift to scale-out FOBS that can better accommodate the data-intensive apps that enterprises deploy as part of digital transformation initiatives.

Source: IDC Worldwide File- and Object-Based Storage Market Update and Forecast, December 2023
https://www.idc.com/getdoc.jsp?containerId=US49940822

Multi-Protocol Data Access for Data-Intensive Workflows

File and object protocols allow applications, services and devices to access and process data which is stored on remote systems.

- Sun Microsystems released the **Network File System (NFS)** protocol in 1984. NFS is the predominant file protocol in Linux/UNIX environments.
- IBM developed the **Server Message Block (SMB)** protocol in 1983. SMB is the predominant file protocol in Windows environments.
- Amazon AWS launched the **Simple Storage Service (S3)** in 2006. The S3 API defines the predominant object protocol for modern modular applications – on-cloud, on-prem and hybrid.
- There are more file and object protocols such as the **Hadoop Distributed File System (HDFS)**.

Multi-Protocol Data Access simplifies Data-Intensive Workflows for unstructured data.

- There is a **rich ecosystem** of old and new applications, services and devices, each **sticking** to one of those protocols.
- **AI & Analytics** is based on **data-intensive workflows** which acquire and process massive amounts of unstructured data. These workflows require the integration of many of those applications, services and devices.
- **Simplification and Cost Savings:** Multi-Protocol Data Access to the same instance of data without copying eliminates the need to create, store and manage multiple copies of the same instance of data.
- **Acceleration:** Eliminating the need to create copies accelerates data-intensive workflows, because it just takes time to create copies of large data sets.

Storage for Data and AI therefore must support various file and object protocols along with multi-protocol data access to the same instance of data without copying of data.

- The **Data Access Services (DAS)** of **IBM's Global Data Platform (GDP)** addresses this need.
- Other market segments of file and object storage such as **General-Purpose File Sharing** and **Enterprise NAS** (e.g., NetApp, Isilon) require a different level of NFS, SMB and S3 and a different level of product packaging.

Current Data Access Services for Storage Scale

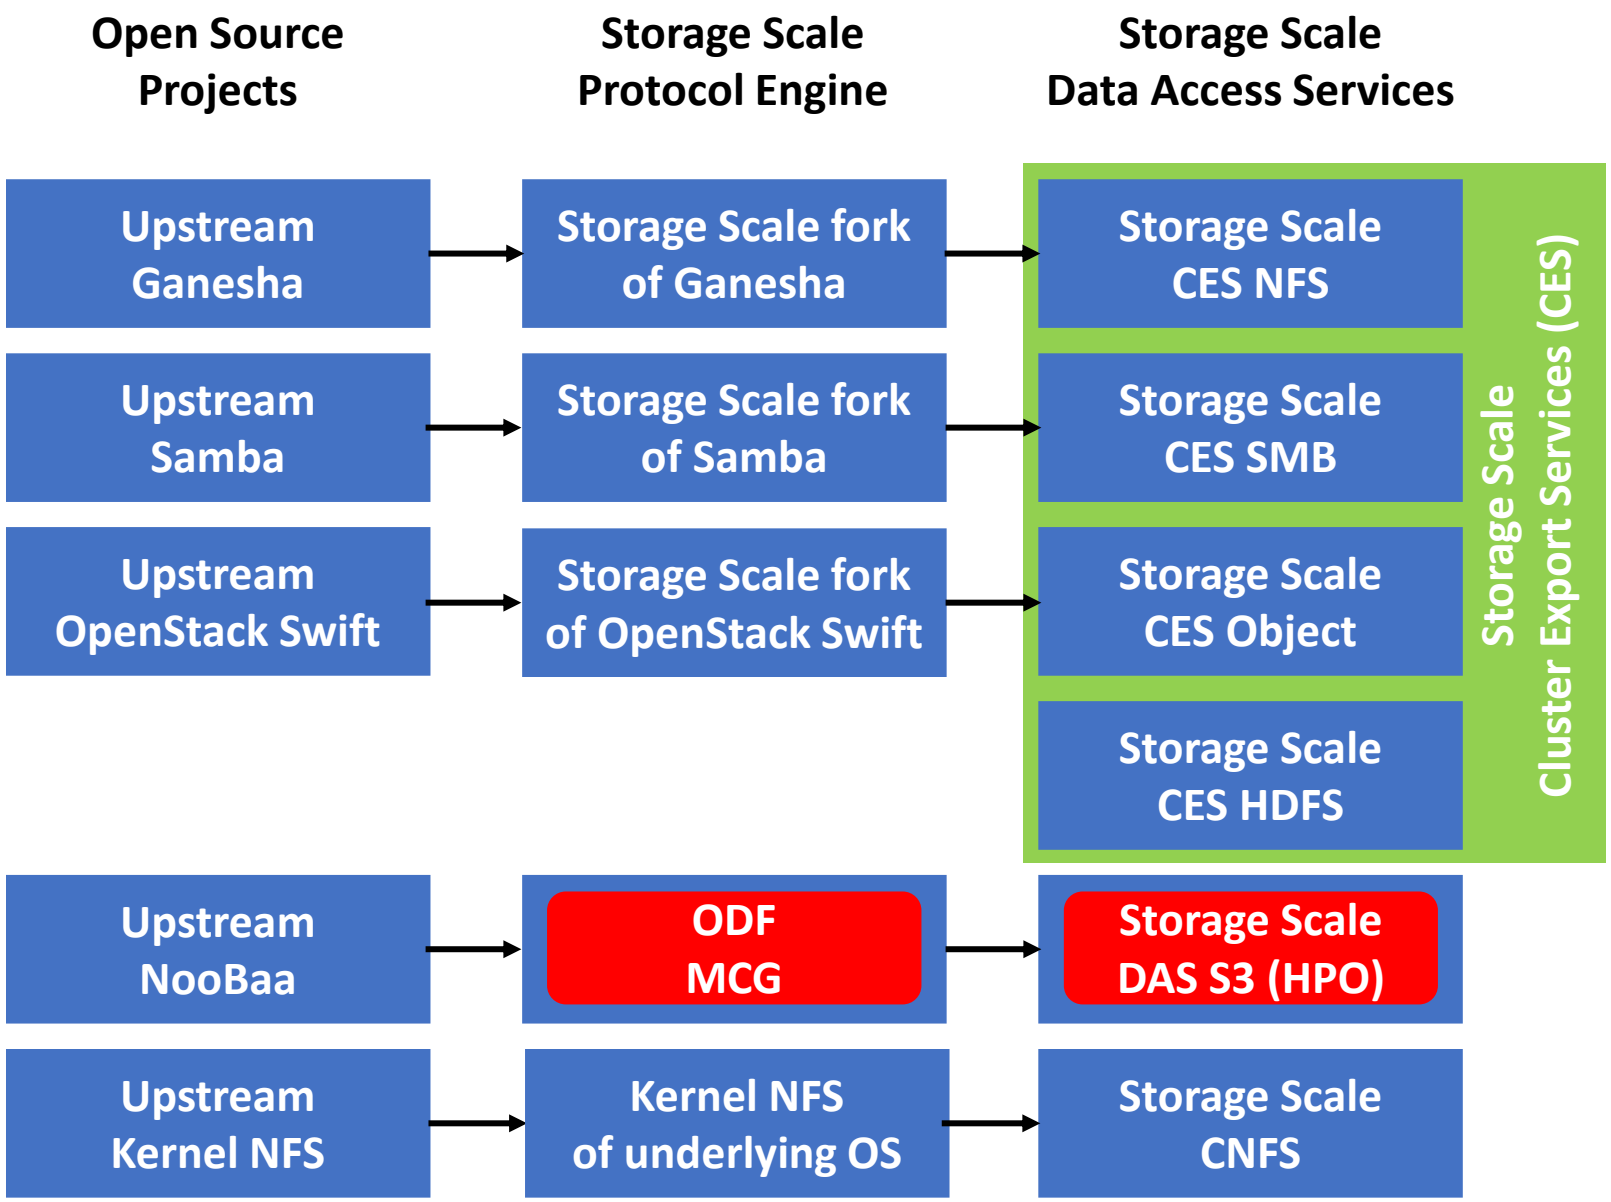

Desired Data Access Services for Storage Scale

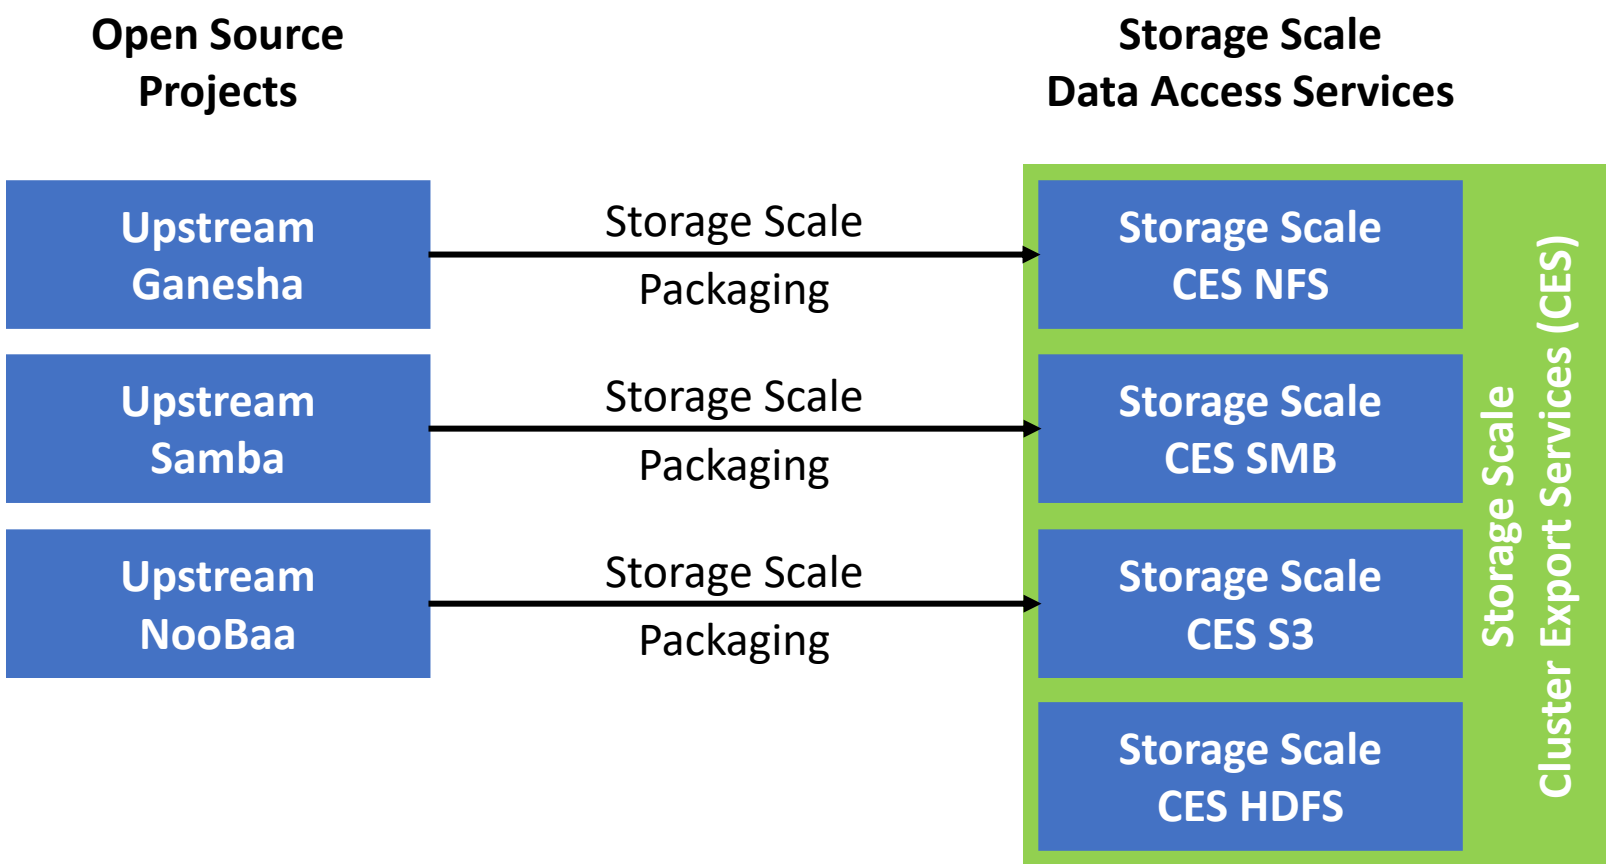

Multi-Protocol Data Access for Data Intensive Workloads

The Red Hat Storage team joined IBM in 2023. This group included software engineers for Samba, Ganesha and NooBaa.
IBM now streamlines our engagement in the upstream projects, by having IBM teams for Samba, Ganesha and NooBaa, and each team being responsible for multiple IBM Storage products.

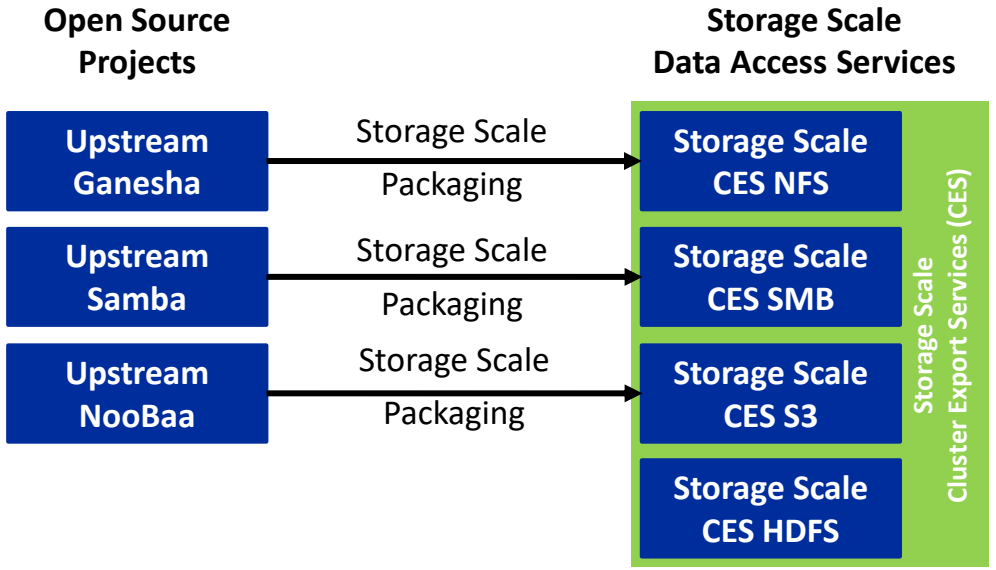

Strategic approach to Storage Scale protocols will be to leverage upstream contributions

- Faster integration of upstream enhancements for fixes and improvements
- Storage Scale fixes and improvements will be contributed upstream
- Leverage leadership in open-source community
- Integrated with Storage Scale Cluster Export Services (CES)

Optimized for multi-protocol data access to enable data-intensive workflows which access the same instance of data using S3 and other access protocols

Timeline:

- CES S3 Object protocol (HPO 2.0) will replace CES Swift Object protocol and refactor DAS S3 (HPO) Noobaa stack removing OpenShift prerequisite.
  - Tech Preview with Storage Scale 5.2.0
  - GA targeted for 2H24

Agenda

- Storage Scale CES S3
- Evolution of Multi-Protocol Data Access in Storage Scale
- CES S3 Examples
- CES S3 CLI Reference
- CES S3 Installation
- Migrating to CES S3

Early preview: CES S3 deployment

- Starting point:
  - Storage Scale cluster with 5 nodes (stretched)
  - Storage Scale version 5.1.9, RHEL 9.2
  - Two CES nodes with SMB and NFS enabled

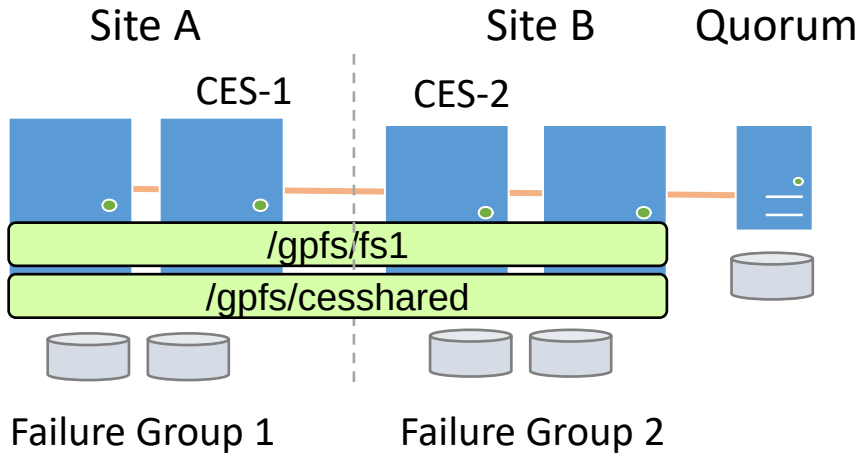

- Deployment in three simple steps
  - Upgrade to 5.2.0 using installation toolkit: `spectrumscale upgrade run`
  - Enable S3 using installation toolkit: `spectrumscale enable S3`
  - Deploy the configuration: `spectrumscale deploy`

CES S3 configuration (accounts and buckets)

- S3 service is enabled and started after deployment on two CES nodes

```
# mmces service list
Enabled services: NFS SMB S3
NFS is running, SMB is running, S3 is running
```

- Create account name *user1* with UID 1001 and GID 100 and bucket-path */gpfs/fs1/buckets*

```
# mms3 account create --uid 1001 --gid 100 --newBucketsPath /gpfs/fs1/buckets user1
The secret and access keys are as follows.
Access Key                Secret Key
fCZCVhcdL9YKrexcUSFU      hT7jZP1Hd8kDARAvspO4HKa0HHswd2UrpVGBmZPL
```

  - If *newBucketPath* does not exist, it is created as directory with proper permissions (770 for UID and GID)

- Create bucket *test1* for *user1* in subdirectory *test1*

```
# mms3 bucket create --accountName user1 --filesystemPath /gpfs/fs1/buckets/test1 test1
INFO Bucket test1 created successfully.
```

  - Buckets can also be created using the S3 API

CES S3 list accounts and buckets

- List accounts:

```
# mms3 account list
Name   New Buckets Path        Uid     Gid
-----  -----------------       ---     ---
user1   /gpfs/fs1/buckets       1001    100
```

- List account detail (keys)

```
# mms3 account list user1
Name   New Buckets Path        Uid     Gid     Access Key              Secret Key
----   -----------------       ---     ---     -----------             -----------
user1  /gpfs/fs1/buckets       1001    100     fCZCVhcdL9YKrexcUSFU    hT7jZP1Hd8k….
```

- List buckets:

```
# mms3 bucket list
Name
------
test1
```

Using CES S3 service

- Created a profile for user1 with access_key and secret_key with AWS CLI
  - Created a command alias for AWS cli (s3u1) wrapping essential aws-command parameters (AWS_ACCESS_KEY_ID, AWS_SECRET_ACCESS_KEY, AWS_ENDPOINT_URL --no-verify-ssl)

- List buckets using s3u1 alias

```
$ s3u1 ls s3://
2024-02-21 21:37:16 test1
```

- PUT an object in bucket *test1*

```
$ s3u1 cp files/file_0.pdf s3://test1
upload: files/file_0.pdf to s3://test1/file_0.pdf
```

- List content of bucket *test1*

```
$ s3u1 ls s3://test1
2024-02-21 21:22:15     413696 file_0.pdf
```

Create new S3 account and store object

- Create S3 *user2* with UID 1000 and GID 100 in bucket-path */gpfs/fs1/buckets*

```
# mms3 account create --uid 1000 --gid 100 --newBucketsPath /gpfs/fs1/buckets user2
  Access Key              Secret Key
  o2xc26I25IJNe5r33ph4    rH4s3cXBVfcej/6eyDre4sV07BUU29jOKZt3nPN7
```

- Create bucket via S3 API (using alias s3u2 for *user2*)

```
$ s3u2 mb s3://test2
make_bucket: test2
```

- List buckets

```
$ s3u2 ls
2024-02-23 12:23:32 test2
2024-02-23 12:23:32 test1
```

- PUT object in bucket test2

```
$ s3u2 cp files/file_2.pdf s3://test2
upload: files/file_2.pdf to s3://test2/file_2.pdf
```

Adjusting buckets access policy

- List bucket *test1* created by *user1* as *user2* – Permission denied

```
$ s3u2 ls s3://test1
An error occurred when calling the ListObjectsV2 operation: Access Denied
```

- Adjust bucket policy for bucket *test1* allowing access for *user2*:

```
$ s3u1api put-bucket-policy --bucket test1 --policy file:///home/user/b-test1.policy
```

- List buckets *test1* as *user2* - works

```
$ s3u2 ls s3://test1
2024-02-21 21:22:15     413696 file_0.pdf
```

- List bucket *test2* as *user2*

```
$ s3u2 ls s3://test2
2024-02-23 12:29:32     248832 file_2.pdf
```

```
{
"Version":"2024-02-23",
"Statement":[
   {  "Sid":"AllowUser2",
      "Effect":"Allow",
      "Principal": { "AWS": "user2" },
      "Action":["s3:*"],
      "Resource":["arn:aws:s3:::*"]
    }
  ]
}
```

Access object through SMB

- Let’s look at the directory structure in the buckets path */gpfs/fs/buckets*:

```
# tree  /gpfs/fs1/buckets/
/gpfs/fs1/buckets/
├── test1
│   └── file_0.pdf
└── test2
    └── file_2.pdf
```

- Export bucket path

```
# mmsmb export add buckets /gpfs/fs1/buckets/
```

- Look into the SMB export (smb user UID = 1000, GID = 100)

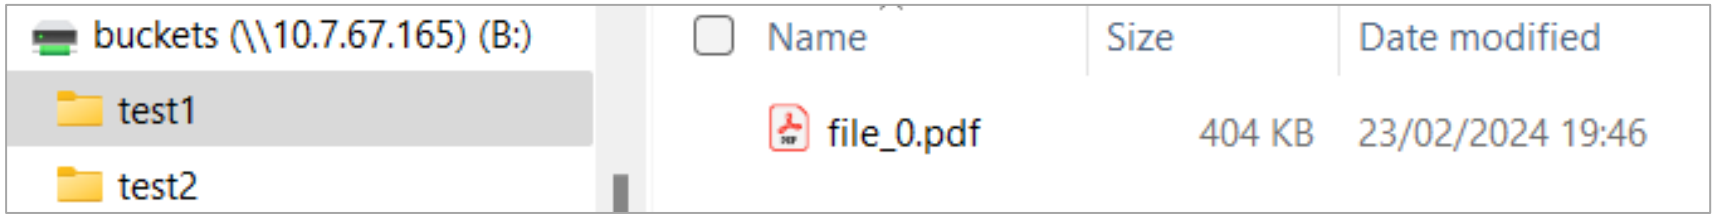

Multi-protocol sharing (S3 and SMB)

- Store files in buckets via SMB in directory *test1*

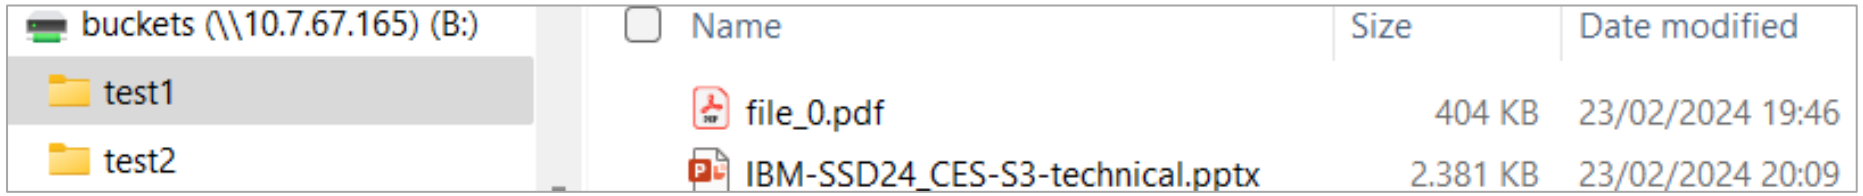

- List bucket *test1*

```
$ s3u1 ls s3://test1
2024-02-23 20:09:57    2438094 IBM-SSD24_CES-S3-technical.pptx
2024-02-23 19:46:03     413696 file_0.pdf
```

Manage object tags

- Add tag *project=s3* to object *test1/file_0.pdf* using S3 API

```
s3u1api put-object-tagging --bucket test1 --key file_0.pdf \
        --tagging TagSet='[{Key=project,Value=s3}]'
```

- Show tag through S3 API using *get-object-tagging* function:

```
$ s3u1api get-object-tagging --bucket test1 --key file_0.pdf
{  "TagSet": [ {"Key": "project","Value": "s3"}]}
```

- Show extended attributes of files

```
# mmlsattr -L -d /gpfs/fs1/buckets/test1/file_0.pdf
file name:            /gpfs/fs1/buckets/test1/file_0.pdf
creation time:        Fri Feb 23 19:46:03 2024
user.noobaa.content_type: "application/pdf"
user.storage_class:   "STANDARD"
user.noobaa.tag.project: "s3"
```

Use tags in policy

- Policy to list tag *user.noobaa.tag.project*

```
RULE 'extlist' EXTERNAL LIST 'tag' EXEC ''
RULE 'listtag' LIST 'tag' WHERE XATTR('user.noobaa.tag.project') IS NOT NULL
```

- Execute policy and create output file *my.list.tag*

```
# mmapplypolicy fs1 -P list-tag.policy -I defer -f ./my
```

- Result of the policy shows file_0.pdf that was tagged before:

```
# cat my.list.tag
6336 571644372 0   -- /gpfs/fs1/buckets/test1/file_0.pdf
```

Observations

- Some S3 related events popped up

```
# mmhealth cluster show ces
  Component     Node          Status            Reasons
  ------------------------------------------------------------------------
  CES           scaleA-1i     DEGRADED      s3_config_backup_notok
  CES           scaleA-2i     DEGRADED      s3_access_denied
  CES           scaleB-2i     HEALTHY       -
```

- Event description

```
Event                     Event Message
-------------------------------------------------------------------------
s3_config_backup_notok  S3 config backup failed, last backup CCR version is None.
s3_access_denied         Account user1 does not have access to the storage path.
```

Agenda

- Storage Scale CES S3
- Evolution of Multi-Protocol Data Access in Storage Scale
- CES S3 Examples
- CES S3 CLI Reference
- CES S3 Installation
- Migrating to CES S3

mmces enhancement for S3

The **mmces service** command now supports management of S3 service.

```
# mmces service
Usage:
        mmces service enable  {NFS | OBJ | SMB | S3 | BLOCK | HDFS}
        mmces service disable {NFS | OBJ | SMB | S3 | BLOCK | HDFS} [--force]
        mmces service {start | stop} {NFS | OBJ | SMB | S3 | BLOCK | HDFS} [-N {Node[,Node...] |
NodeFile | NodeClass} | -a]
        mmces service list [-N {Node[,Node...] | NodeFile | NodeClass} | -a] [--verbose]
```

Example:

```
# mmces service list -a
Enabled services: NFS SMB S3
node1.vmlocal:  NFS is running, SMB is running, S3 is running
node2.vmlocal:  NFS is running, SMB is running, S3 is running
node3.vmlocal:  NFS is running, SMB is running, S3 is running
```

mms3 CLI – CES S3 management

The new **mms3** command supports management of S3 configuration, S3 accounts and S3 buckets. It also supports backup and restore of S3 configuration.

```
# mms3
mms3 is a CLI that empowers users to configure S3 service, buckets and accounts for IBM
Storage Scale system.
Usage:
  mms3 [command]
Available Commands:
  account     Manage IBM Storage Scale S3 accounts
  bucket      Manage IBM Storage Scale S3 buckets
  config      Manage IBM Storage Scale S3 configuration for a CES cluster.
Flags:
  -h, --help   help for mms3
Use "mms3 [command] --help" for more information about a command.
```

mms3 CLI – S3 Configuration management

```
# mms3 config
Usage:
  mms3 config [flags]
  mms3 config [command]


Available Commands:
  backup      Backup the IBM Storage Scale S3 configuration.
  change      Modifies the IBM Storage Scale S3 configuration parameters.
  list        List the IBM Storage Scale S3 configuration parameters and their values.
  restore     Restore the IBM Storage Scale S3 configuration.


Flags:
  -h, --help   help for config


Use "mms3 config [command] --help" for more information about a command.
```

Example:

```
# mms3 config change ENDPOINT_FORKS=6
S3 service will get restarted on all CES nodes where S3 service is running. Do you want to continue
with the configuration change? [y|n]: y
S3 Configuration successfully changed. S3 service restarted on the CES node: node1.vmlocal.
S3 Configuration successfully changed. S3 service restarted on the CES node: node2.vmlocal.
S3 Configuration successfully changed. S3 service restarted on the CES node: node3.vmlocal.
```

Default S3 Configuration

```
# mms3 config list


 S3 NOOBAA Configuration:
 ========================
 ALLOW_HTTP : false
 DEBUGLEVEL : default
 ENABLEMD5 : false
 ENDPOINT_FORKS : 2
 ENDPOINT_PORT : 6001
 ENDPOINT_SSL_PORT : 6443
 GPFSDLPATH : /usr/lpp/mmfs/lib/libgpfs.so
 NSFS_DIR_CACHE_MAX_DIR_SIZE : 536870912
 NSFS_DIR_CACHE_MAX_TOTAL_SIZE : 1073741824
 NSFS_NC_CONFIG_DIR_BACKEND : GPFS
 NSFS_NC_STORAGE_BACKEND : GPFS
 UVTHREADPOOLSIZE : 16
=======================
```

mms3 CLI - Accounts Management

```
# mms3 account

Usage:
  mms3 account [flags]
  mms3 account [command]


Available Commands:
  create      Create a IBM Storage Scale S3 account
  delete      Delete a IBM Storage Scale S3 account
  list        List IBM Storage Scale S3 account(s)


Flags:
  -h, --help   help for account


Use "mms3 account [command] --help" for more information about a command.

Example:
# mms3 account create account2 --uid 1001 --gid 1001 --newBucketsPath /mnt/gpfs0/account2
```

mms3 CLI - Buckets Management

```
# mms3 bucket

Usage:
  mms3 bucket [flags]
  mms3 bucket [command]


Available Commands:
  create      Create a IBM Storage Scale S3 bucket.
  delete      Delete a IBM Storage Scale S3 bucket
  list        List IBM Storage Scale S3 bucket(s)


Flags:
  -h, --help   help for bucket


Use "mms3 bucket [command] --help" for more information about a command.
```

Example:

```
# mms3 bucket create bucket2 --filesystemPath /mnt/gpfs0/account2/bucket2 --accountName account2
```

Sample script to update S3 Bucket Policy

With S3 bucket policies, you can secure access to objects in your buckets, so that only users with the appropriate permissions can access them. You can even prevent authenticated users without the appropriate permissions from accessing your S3 buckets/objects.

- Sample script to update S3 bucket policy is provided with CES S3 at: `/usr/lpp/mmfs/samples/s3SetBucketPolicy.sample.sh`

```
# /usr/lpp/mmfs/samples/s3SetBucketPolicy.sample.sh
Usage: ./s3SetBucketPolicy.sample.sh [-op <op-type>] [-b <bucket-name>] [-a <access-id>] [-s <secret-key>]

Options:
        -op:    Policy operation type (get/set)
        -b:  Bucket name
        -a:    S3 user AWS_ACCESS_KEY_ID
        -s:   S3 user AWS_SECRET_ACCESS_KEY
-h, --help: For help

Example:
# /usr/lpp/mmfs/samples/s3SetBucketPolicy.sample.sh -op get -b bucket-test28 -a Q7FjtLhmbXVOKXYOC -s
kVbB0yFMPgPAkASCSQ0U7t5ZbyXaazk7ihgQs
```

Failover and Failback Support

When a CES node fails or is suspended/unavailable, the CES addresses assigned to that node are redistributed among the remaining nodes. The S3 workload that was served by the node leaving the cluster, then gets handled by the active node that receives the associated CES addresses of the unavailable node.

Example, the **mmces** command used to suspend/stop services on a node.

```
# mmces node {suspend [--stop] | resume [--start]}
```

```
# mmlscluster --ces
GPFS cluster information
========================
  GPFS cluster name:         rkomandu-ip-cls-x.fyre.ibm.com
  GPFS cluster id:           2660669109853635319
Cluster Export Services global parameters
-----------------------------------------
  Shared root directory:                /mnt/ces-shared-root
  Enabled Services:                     S3
  Log level:                            0
  Address distribution policy:          even-coverage
Node   Daemon node name            IP address       CES IP address list
-----------------------------------------------------------------------
   2   rkomandu-ip-cls-x-worker1.fyre.ibm.com 10.11.77.120     10.11.77.139, 10.11.77.133
   3   rkomandu-ip-cls-x-worker2.fyre.ibm.com 10.11.77.128     Node suspended
```

Backing up and restoring S3 configuration

CES S3 supports back up and restore of S3 configuration files (stored in CES Shared Root file system).

- Periodic/automatic (every 10 minutes) backing up of S3 configuration files to CCR.

```
# mmccr flist
 48  mmsdrfs
  1  mmsysmon.json
  1  _callhomeconfig
597  Cesiplist
2033 _s3-config-backup.tar.bz2
```

- Commands to backup and restore S3 configuration files.

```
# mms3 config
Available Commands:
  backup      Backup the IBM Storage Scale S3 configuration.
  restore     Restore the IBM Storage Scale S3 configuration.
```

RAS / Monitoring

Monitoring capabilities are extended to check the health of S3 service on each node of the CES cluster. The **mmhealth** command displays the results of the background monitoring for the health of S3 service

```
# mmhealth node show CES

Node name:      rkomandu-ip-cls-x-worker1.fyre.ibm.com

Component                      Status        Status Change     Reasons & Notices
-------------------------------------------------------------------------------
CES                            HEALTHY       12 days ago       -
  AUTH                         DISABLED      19 days ago       -
  AUTH_OBJ                     DISABLED      19 days ago       -
  BLOCK                        DISABLED      19 days ago       -
  CESNETWORK                   HEALTHY       19 days ago       -
  HDFS_NAMENODE                DISABLED      19 days ago       -
  NFS                          HEALTHY       12 days ago       -
  OBJECT                       DISABLED      19 days ago       -
  S3                           HEALTHY       19 days ago       -
  SMB                          DISABLED      19 days ago       -

There are no active error events for the component CES on this node (rkomandu-ip-cls-x-
worker1.fyre.ibm.com).
```

mmhealth (Example)

Use cases shown:

- For a S3 bucket the bucket storage directory path does not exist.
- For a S3 user the buckets directory path does not exist.

```
# mmhealth node show CES S3 -v
rkomandu-ip-cls-x-worker2.fyre.ibm.com: Mon Apr 22 23:38:11 2024

Node name:      rkomandu-ip-cls-x-worker2.fyre.ibm.com

Component                      Status        Status Change            Reasons & Notices
-------------------------------------------------------------------------------------------
S3                             DEGRADED      2024-04-22 23:37:47      s3_newbucketspath_not_exist(s3user-16003@fvt.com),
s3_storage_not_exist(newbucket-16003-dir)
  newbucket-16003-dir          DEGRADED      2024-04-22 23:37:47      s3_storage_not_exist(newbucket-16003-dir)
  newbucket-backup-demo        HEALTHY       2024-04-09 21:03:08      -
  s3user-16003@fvt.com         DEGRADED      2024-04-22 23:37:47      s3_newbucketspath_not_exist(s3user-16003@fvt.com)
  s3user-16004@fvt.com         HEALTHY       2024-04-22 21:31:45      -

Event                          Parameter                     Severity    Active Since         Event Message
-------------------------------------------------------------------------------------------------------------
s3_active                      S3                            INFO        2024-04-10 09:03:46  The S3 service is now active.
s3_newbucketspath_not_exist    s3user-16003@fvt.com          WARNING     2024-04-22 23:37:47  The new buckets path for account
s3user-16003@fvt.com mentioned in the schema does not exist.
s3_storage_not_exist           newbucket-16003-dir           WARNING     2024-04-22 23:37:47  The storage path for bucket
newbucket-16003-dir mentioned in the schema does not exist.
s3_account_ok                  s3user-16004@fvt.com          INFO        2024-04-22 21:31:45  S3 account s3user-16004@fvt.com
is OK.
```

Debug Data Collection (gpfs.snap, mmcallhome)

- gpfs.snap

mmcallhome when issued which will capture similar data for CES S3

```
gpfs.snap -a
s3.snap.rkomandu-ip-cls-x-worker1.fyre.ibm.com.20240408_072434

├── CommandOutput
│   ├── mms3_account_list
│   ├── rsyslog_status
│   ├── s3_account_ls_l
│   ├── s3_noobaa_health
│   ├── s3_noobaa_journalctl_status
│   ├── s3_noobaa_nsfs_status
│   └── s3_noobaa_stats
├── etc
│   ├── logrotate.d
│   │   └── noobaa
│   │       └── logrotate_noobaa.conf
│   ├── noobaa.conf.d
│   │   └── config_dir_redirect
│   └── rsyslog.d
│       ├── noobaa_rsyslog.conf
│       └── noobaa_syslog.conf
├── mnt
│   └── ces-shared-root
│       └── ces
│           └── s3-config
│               ├── config.json
│               └── system.json
└── var
    └── log
        └── noobaa_events.log
        └── noobaa.log
```

Self-signed SSL/TLS certificates for secure communication between S3 client and S3 service (on CES node)

Steps:

- Generate TLS key, CSR and CRT file using openssl command
- Place tls.key and tls.crt files under {cesSharedRoot_path}/ces/s3-config/certificates.
- Restart the S3 service on all CES nodes
- At the client, when using S3 commands set AWS_CA_BUNDLE to point the CRT file. Example:

```
# alias s3_ssl='AWS_CA_BUNDLE=/path/to/tls.crt AWS_ACCESS_KEY_ID=add_your_access_key
AWS_SECRET_ACCESS_KEY=add_your_secret_key aws --endpoint https://<endpoint>:6443 s3'
```

Agenda

- Storage Scale CES S3
- Evolution of Multi-Protocol Data Access in Storage Scale
- CES S3 Examples
- CES S3 CLI Reference
- CES S3 Installation
- Migrating to CES S3

Ansible based Install Toolkit Overview

Install toolkit Workflow

Define Cluster Topology

Use Spectrum Scale CLI commands to Cluster definition
- Add nodes to cluster
- Assign roles to nodes
- Define NSD
- Define File system

Install

Use Spectrum Scale install to perform:
- Installation required RPMs on all nodes
- Creates a Spectrum Scale cluster
- Creates NSD
- Sets up Management GUI
- Creates file system

Deploy

Use Spectrum Scale deploy to perform:
- Install, Configure and enable protocols

Upgrade

Use Spectrum Scale upgrade to perform:
- Online sequential upgrade of cluster
- Offline cluster upgrade

Toolkit can be used to automate deploy alone when install has happened manually

Toolkit can be used to upgrade an existing cluster that has been created manually

Ansible Toolkit: Key Features enhancement

- Install Toolkit CES S3 Support
  - As a technology preview feature, toolkit supports CES S3 installation and configuration.

Install Toolkit CES S3 Deployment: **Workflow**

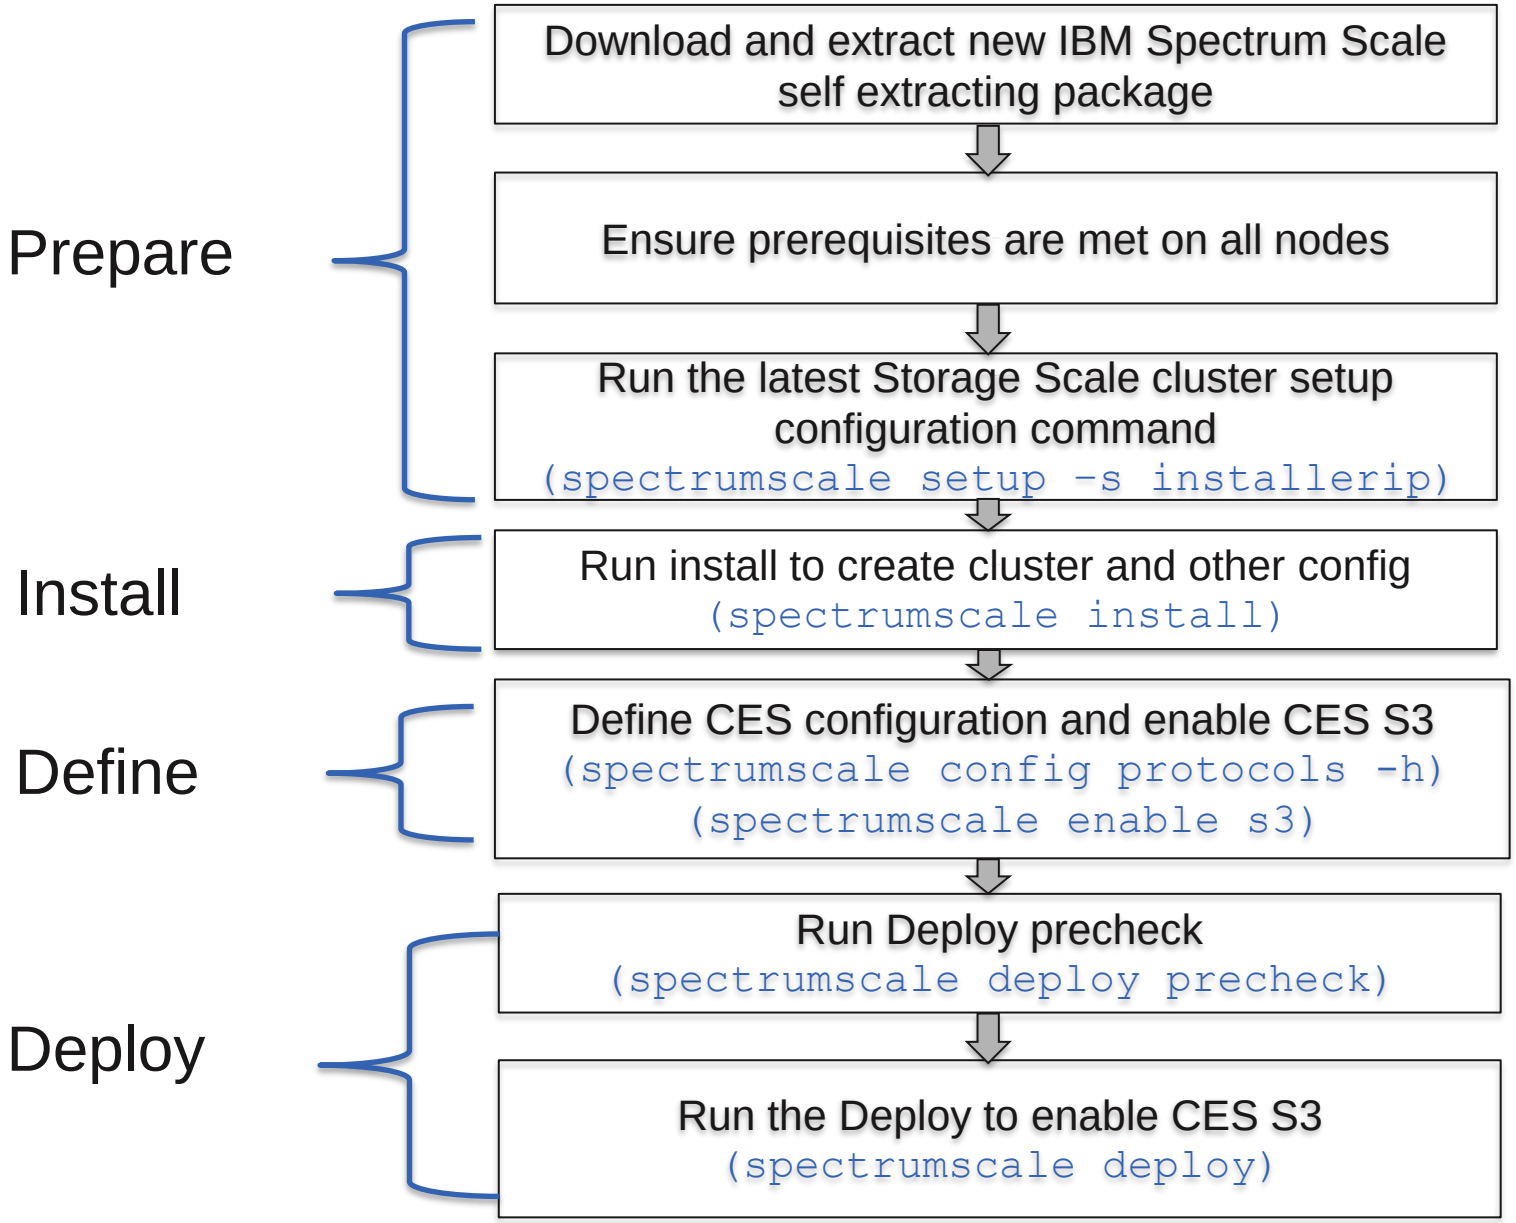

Ansible Toolkit: CES S3 steps

CES S3 deployment of IBM Storage Scale by using the installation toolkit.

Download and extract the IBM Spectrum Scale self-extracting package on the installer node.

root@test-21:~# /root/Storage_Scale_Advanced-5.2.0.0-x86_64-Linux-install –silent

Go to the IBM Spectrum Scale installation toolkit directory.

root@test-21:~# **cd /usr/lpp/mmfs/5.2.0.0/ansible-toolkit/**

Issue the spectrumscale setup command to set up the new installer node.

root@test-21:/usr/lpp/mmfs/5.2.0.0/ansible-toolkit# **./spectrumscale setup -s 10.0.100.95**

```
root@test-21:/usr/lpp/mmfs/5.1.4.0/ansible-toolkit# ./spectrumscale setup -s 10.0.100.95
[ INFO  ] Installing prerequisites for install node
[ INFO  ] Found existing Ansible installation on system.
[ 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  ] Your ansible controller node has been configured to use the IP 10.0.100.95 to communicate with other nodes.
[ INFO  ] Port 10080 will be used for package distribution.
[ 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
```

Ansible Toolkit: CES S3 steps

- **# Use the Ansible based install toolkit Install to create a cluster or you can directly use deploy to install and configure CES S3 protocols if cluster is already exists.**

- **CES S3 configuration of IBM Storage Scale by using the installation toolkit.**

 # Use the Ansible based install toolkit cli config to define CES config to define CesShared root and export IP. so the toolkit knows what to act upon.

[root@test-21 ansible-toolkit**]# ./spectrumscale config protocols -f cesSharedRoot -m /mnt/cesSharedRoot**
[**INFO**] Setting filesystem to cesSharedRoot
[**INFO**] Setting mountpoint to /mnt/cesSharedRoot
[**INFO**] Tip :Enable NFS, S3, SMB or HDFS protocols as appropriate:./spectrumscale enable nfs|s3|smb|hdfs

[root@tes-21 ansible-toolkit]# **./spectrumscale config protocols -e 192.168.0.131,192.168.0.132**
[**INFO**] Setting export_ip_pool to ['192.168.0.131', '192.168.0.132']
[**INFO**] Tip :Enable NFS, S3, SMB or HDFS protocols as appropriate:./spectrumscale enable nfs|s3|smb|hdfs

Ansible Toolkit: CES S3 steps

- **# Use the Ansible based install toolkit Install to create a cluster or you can directly use deploy to install and configure CES S3 protocols if cluster is already exists.**

- **Enable CES S3 protocols in the installation toolkit configuration.**

 # Use the Ansible based install toolkit cli to enable s3 protocol. so the toolkit knows what to act upon.

[root@apr22x86-21 ansible-toolkit]# ./spectrumscale enable -h

usage: spectrumscale enable [-h] {s3,smb,nfs,hdfs} [{s3,smb,nfs,hdfs} ...]

positional arguments:

{s3,smb,nfs,hdfs} Specify the protocols to be enable or disable; either enable or disable one of the protocols, or enable or disable multiple protocols separated by space.

[root@test-21 ansible-toolkit]# **./spectrumscale enable s3**

[**INFO**] **IBM Storage Scale introduces CES S3 as a technology preview (non-production environments) feature. For instructions on enabling and running CES S3, please sign up for the Technical preview by contacting scale@us.ibm.com.**

[**INFO**] Enabling S3 on all protocol nodes.

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

Ansible Toolkit: CES S3 steps

 # Use the Ansible based install toolkit node list cli to check if S3 is defined

```
[root@apr22x86-21 ansible-toolkit]# ./spectrumscale node list
[ INFO  ] List of nodes in current configuration:
[ INFO  ] [Installer Node]
[ INFO  ] 10.0.100.22
[ INFO  ]
[ INFO  ] [Cluster Details]
[ INFO  ] Name: scale-cluster-2.openstacklocal
[ INFO  ] Setup Type: IBM Storage Scale
[ INFO  ]
[ INFO  ] [Protocols]
[ INFO  ] S3      : Enabled
[ INFO  ] SMB     : Enabled
[ INFO  ] NFS     : Enabled
[ INFO  ] HDFS    : Disabled
[ INFO  ]
[ INFO  ] [Extended Features]
[ INFO  ] File Audit logging     : Disabled
[ INFO  ] Management GUI         : Enabled
[ INFO  ] Performance Monitoring : Enabled
[ INFO  ] Callhome               : Enabled
[ INFO  ]
[ INFO  ] GPFS                    Admin  Quorum  Manager   NSD   Protocol   GUI   Callhome  Perf Mon    OS    Arch
[ INFO  ] Node                     Node   Node     Node   Server   Node    Server  Server  Collector
[ INFO  ] apr22x86-21.openstacklocal   X      X       X      X       X       X       X         X      rhel9  x86_64
[ INFO  ] apr22x86-22.openstacklocal           X       X      X       X                         X      rhel9  x86_64
[ INFO  ] apr22x86-23.openstacklocal           X       X      X                                 X      rhel9  x86_64
[ INFO  ]
[ INFO  ] [Export IP address]
[ INFO  ] 192.168.0.131 (pool)
[ INFO  ] 192.168.0.132 (pool)
[ INFO  ] 192.168.0.133 (pool)
[ INFO  ] 192.168.0.134 (pool)
[ INFO  ] 192.168.0.135 (pool)
[ INFO  ] 192.168.0.136 (pool)
```

CES S3

Ansible Toolkit: CES S3 steps

Use the Ansible based install toolkit to run deploy precheck and ensure the environment is ready for CES S3.

root@test-21:/usr/lpp/mmfs/5.2.0.0/ansible-toolkit# **./spectrumscale deploy precheck**

[root@test-21 ansible-toolkit]# ./spectrumscale deploy --pre
[INFO] Logging to file: /usr/lpp/mmfs/5.2.0.0/ansible-toolkit/logs/DEPLOY-PRECHECK-24-04-2024_02:14:57.log
[INFO] Validating configuration
[WARN] Ensure that base OS repositories are configured and enabled so that package dependencies can be satisfied during installation.
[WARN] test-21.openstacklocal is specified as both an NSD server and a protocol node. It is recommended that a node is not both a protocol and NSD server.
[WARN] test-22.openstacklocal is specified as both an NSD server and a protocol node. It is recommended that a node is not both a protocol and NSD server.
[WARN] Only one GUI server specified. The Graphical User Interface will not be highly available.
[INFO] Install toolkit will not configure file audit logging as it has been disabled.
[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 S3 checks.
[INFO] Running environment checks for S3
[FATAL] **test-22.openstacklocal: Pre requisite package not found on apr22x86-22.openstacklocal (Pre requisite: yum install bzip2). Ensure this is resolved before proceeding with the install toolkit.**
[INFO] Detailed error log: /usr/lpp/mmfs/5.2.0.0/ansible-toolkit/logs/DEPLOY-PRECHECK-24-04-2024_02:14:57.log
[FATAL] Pre requisite check failed on one or more nodes

Ansible Toolkit: CES S3

Run the Deploy process by using the ./spectrumscale deploy command.

root@test-21:/usr/lpp/mmfs/5.2.0.0/ansible-toolkit# **./spectrumscale deploy**

```
[root@test-21 ansible-toolkit]# ./spectrumscale deploy
[ INFO ] Logging to file: /usr/lpp/mmfs/5.2.0.0/ansible-toolkit/logs/DEPLOY-24-04-2024_02:45:30.log
[ INFO ] Validating configuration
[ WARN ] Ensure that base OS repositories are configured and enabled so that package dependencies can be satisfied during installation.
[ WARN ] test-21.openstacklocal is specified as both an NSD server and a protocol node. It is recommended that a node is not both a protocol and NSD server.
[ WARN ] test-22.openstacklocal is specified as both an NSD server and a protocol node. It is recommended that a node is not both a protocol and NSD server.
[ WARN ] Only one GUI server specified. The Graphical User Interface will not be highly available.
[ INFO ] Install toolkit will not configure file audit logging as it has been disabled.
[ INFO ] Running pre-install checks
[ INFO ] NSDs are in a valid state
[ 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 ] Running environment checks for S3
[ INFO ] S3 precheck OK
[ INFO ] TASK [ibm.spectrum_scale.s3_install : install | Add GPFS s3 packages to list] ***
[ INFO ] ok: [test-21.openstacklocal] => (item=noobaa-core)
[ INFO ] ok: [test-22.openstacklocal] => (item=noobaa-core)
[ INFO ] ok: [test-23.openstacklocal] => (item=noobaa-core)
[ INFO ] ok: [test-21.openstacklocal] => (item=gpfs.mms3)
[ INFO ] ok: [test-22.openstacklocal] => (item=gpfs.mms3)
[ INFO ] ok: [test-23.openstacklocal] => (item=gpfs.mms3)
[ INFO ] TASK [ibm.spectrum_scale.s3_configure : configure | Enable S3] *****************
[ INFO ] changed: [test-21.openstacklocal]
[ INFO ] Checking state of S3
[ INFO ] Running S3 post-install checks
[ INFO ] Checking state of S3 on all nodes
[ INFO ] Checking state of S3 on all nodes
[ INFO ] S3 ACTIVE
```

```
[root@test-21 ansible-toolkit]# mmdsh -N cesNodes "rpm -qa | grep noobaa-core"
test-22.openstacklocal: noobaa-core-5.15.0-20240320.el9.x86_64
test-21.openstacklocal: noobaa-core-5.15.0-20240320.el9.x86_64

[root@test-21 ansible-toolkit]# mmdsh -N cesNodes "rpm -qa | grep mms3"
test-22.openstacklocal: gpfs.mms3-5.2.0-0.el9.x86_64
test-21.openstacklocal: gpfs.mms3-5.2.0-0.el9.x86_64

[root@test-21 ansible-toolkit]# mmces state show S3 -a
NODE                         S3
---------------------------  ----
test-21.openstacklocal   HEALTHY
test-22.openstacklocal   HEALTHY
```

Agenda

- Storage Scale CES S3
- Evolution of Multi-Protocol Data Access in Storage Scale
- CES S3 Examples
- CES S3 CLI Reference
- CES S3 Installation
- Migrating to CES S3

Storage Scale S3 Convergence

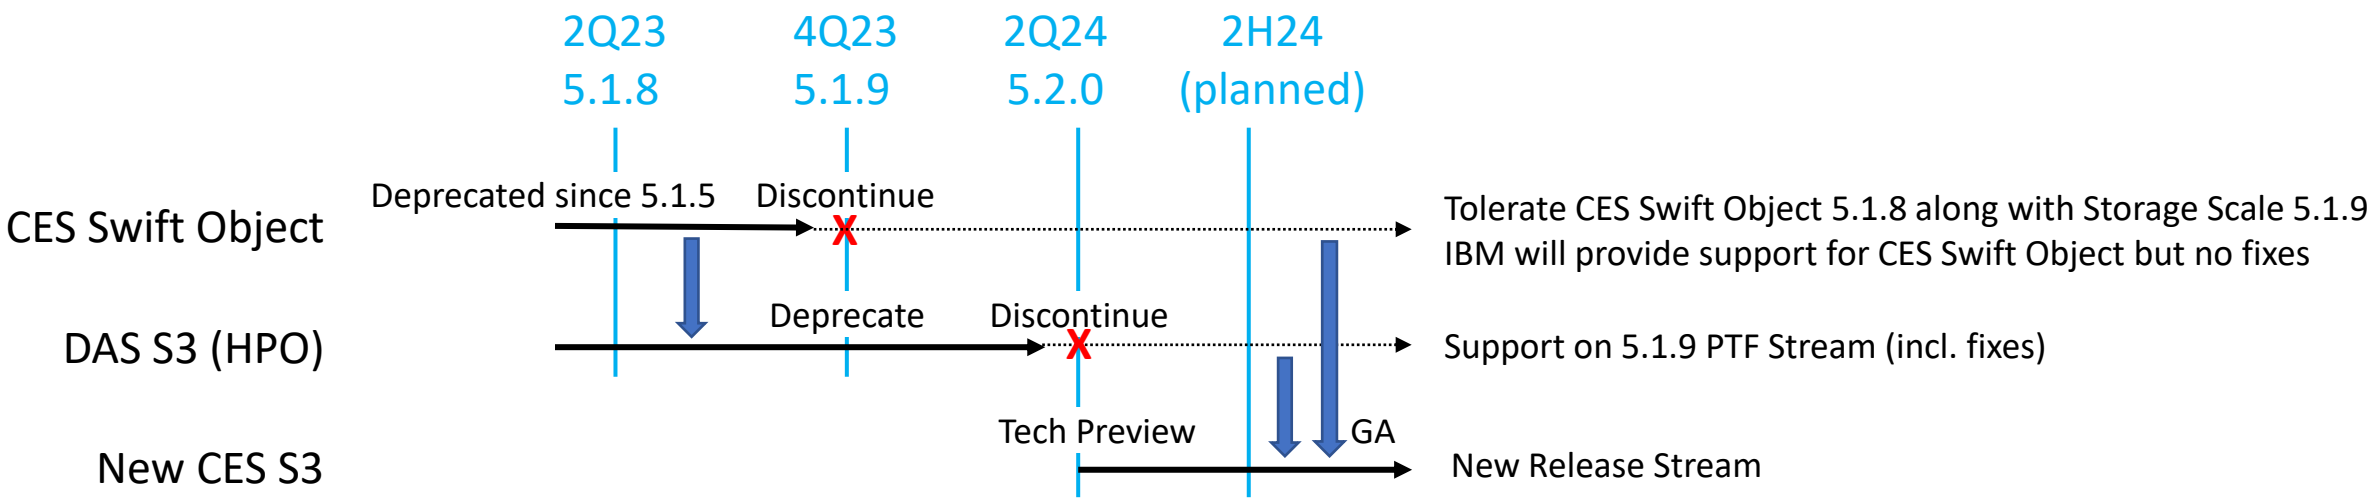

Discontinue CES Swift Object
- Discontinue CES Swift Object with Storage Scale 5.1.9
- Storage Scale 5.1.8 will be the last release with CES Swift Object
- Storage Scale 5.1.9 will tolerate upgrade from 5.1.8 and coexistence of CES Swift Object 5.1.8 with Storage Scale 5.1.9
- Storage Scale 5.1.9 is a long-term support release
- IBM will provide support but no fixes for CES Swift Object 5.1.8

Migration CES Swift Object to DAS S3 or new CES S3
- Migration procedure will be provided
- IBM Technology Expert Lab service offering will be provided
- Complexity of migration will depend on used object API (Swift API vs S3 API), disk layout (e.g., swift-on-file) and tiering to tape

Discontinue DAS S3 (HPO)
- Discontinue DAS S3 (HPO) with Storage Scale 5.2.0
- Storage Scale 5.2.0 is the last release with DAS S3 (HPO)
- Intend to support DAS S3 (HPO) till 2H24

Migration DAS S3 to new CES S3
- Migration procedure will be provided
- It is expected that migration will be in place (no copying of data)

Migration from CES Swift Object and DAS S3 to CES S3

1. Swift to CES S3 migration
   - Procedure to migrate data for Swift to CES S3 migration. Covers migration procedure with/without swiftonfile option enabled for Swift.
   - Reference migration time numbers
   - Procedure to measure migration timing at customer site.
   - User migration guidance.

2. DAS S3 to CES S3 migration
   - Procedure to migrate S3 accounts from DAS S3 to CES S3.
   - No need of data migration.

3. Migration Services
   - IBM Technology Expert Labs provides billable migration services.
   - Details covered in first part of this chart deck.

CES Swift Object Migration Considerations

What	Description	Recommendation
Discontinue Swift API	• When object storage technology emerged, many different REST APIs have been defined to access object storage systems of different vendors. Examples include AWS S3, Azure BLOB, OpenStack Swift and SNIA CDMI. • Meanwhile the S3 REST API as defined by AWS evolved to the de-facto standard for object-based access protocols of on-premise object storage. • CES Swift Object supported Swift and S3. • New CES S3 will support S3.	• Modernize your object-based application to use the AWS S3 API to access on-premise and hybrid cloud object storage.
Compliance to S3 API	• The AWS S3 API includes cloud services which are specific to AWS. • All other object storage systems support a subset of the S3 REST API, and each object storage system has a different level of compliance to the S3 REST API as defined by Amazon AWS. • The capabilities of the S3 REST API of CES Swift Object are documented here: https://docs.openstack.org/swift/train/s3_compat.html • The capabilities of the S3 REST API of CES S3 are yet to be documented. But since CES S3 will provide same REST APIs like DAS S3, check the documentation for DAS S3 here: https://www.ibm.com/docs/en/ssdas?topic=spectrum-scale-data-access-services-517 • The S3 REST API of new CES S3 is planned to be compatible with the S3 REST API of DAS S3 but different to CES Swift Object.	• Validate the interoperability of your S3 applications using new CES S3. • Report any interoperability issues to IBM.

CES Swift Object Migration Considerations

What	Description	Recommendation

Data Migration	• Per default, CES Swift Object stores object files on the file system in hashed paths calculated by the Swift ring files.
	  • To migrate to CES S3, all objects must be extracted from CES Swift Object using a tool such as rclone.
	  • There will be no in-place migration.
	• The optional swiftonfile policy maps objects 1:1 to files and vice versa.
	  • IBM plans to provide migration procedures to new CES S3 which do not require the copying of data. Container (bucket), project (account) and user configuration must be updated. It may be required to update file and object metadata, too.
		• The migration of large data sets can take a while and depends on the underlying infrastructure.
		• Plan respective sufficient time to migrate your data from CES Swift Object to new CES S3.
		• Consider in-place migration to minimize impact to your end users.
		• Contact IBM for available migration procedures and migration services.

Thank you for using

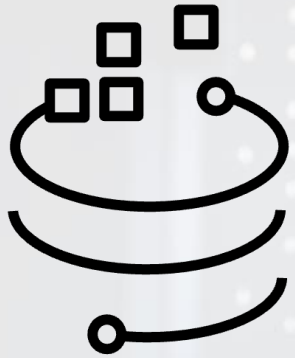

Storage Scale

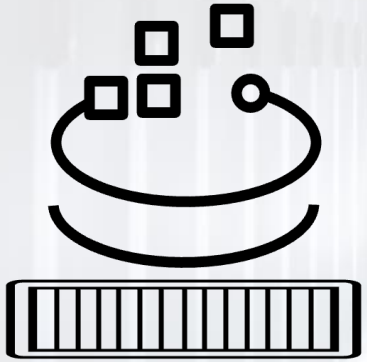

Storage Scale System