Storage Scale

What's new in IBM Storage Scale (5.2.*) and Scale System (6.2.*)

The Global Platform For Storage *(GPFS) Services of unstructured data

Chris Maestas
CTO, Data and AI Storage Solutions
Chief Troublemaking Officer

IBM

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.

The IBM Storage Scale Journey and **Support reference guide**

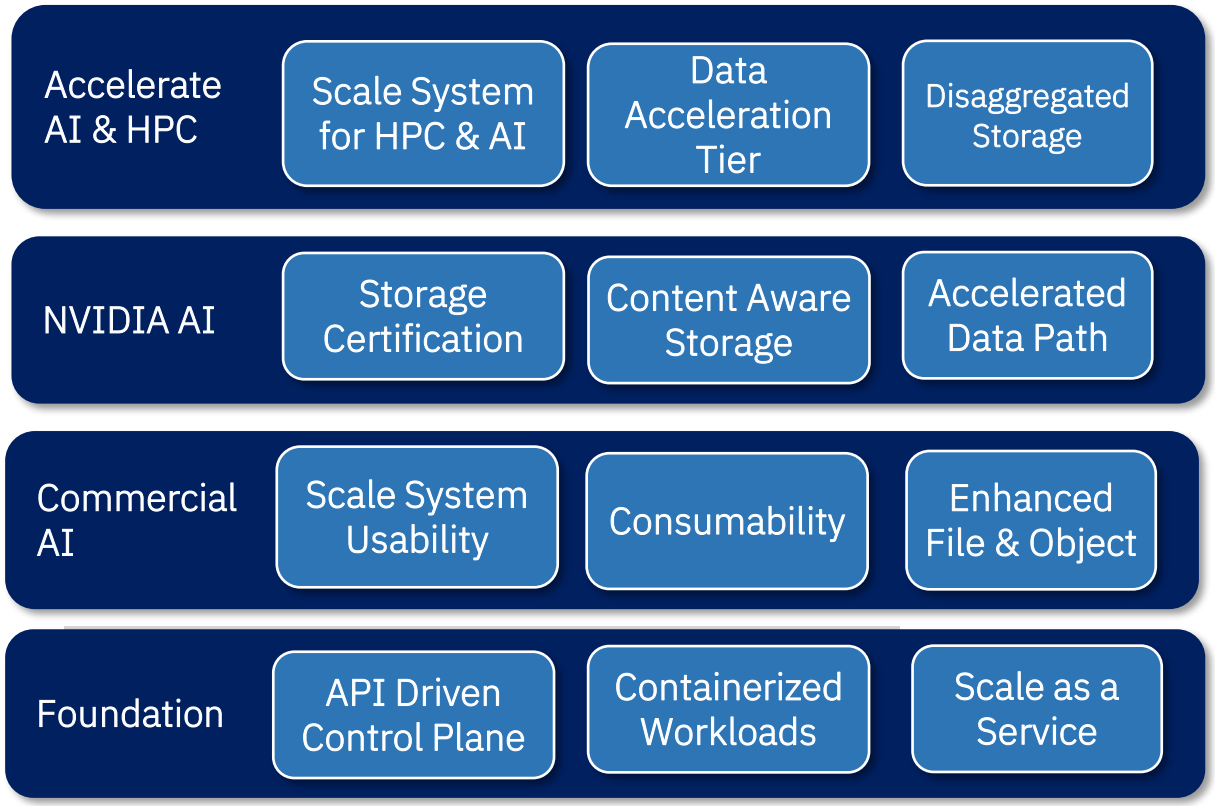

Support reference guide for Storage Scale:
https://www.ibm.com/support/pages/node/6252403

Support reference guide for Storage Scale System:
https://www.ibm.com/support/pages/node/6252477

Note: In a Severity 1 case, please describe your business impact.

IBM Storage Scale Software Version Recommendation Preventive Service Planning:
https://www.ibm.com/support/pages/ibm-storage-scale-software-version-recommendation-preventive-service-planning

IBM’s Storage Scale System works with NVIDIA solutions!

Storage Scale

https://www.nvidia.com/en-us/data-center/dgx-superpod/

NVIDIA Enterprise Reference Architecture (ERA)

NVIDIA DGX BasePOD

NVIDIA DGX SuperPOD

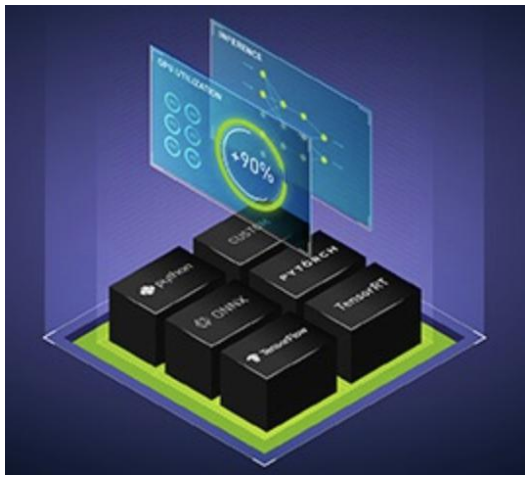

NVIDIA Cloud Partner (NCP)

And wait there’s more! NVIDIA certified system vendor based on HGX

https://www.nvidia.com/en-us/data-center/products/certified-systems/

IBM and NVIDIA Redbook

https://redbooks.ibm.com/abstracts/redp5746.html

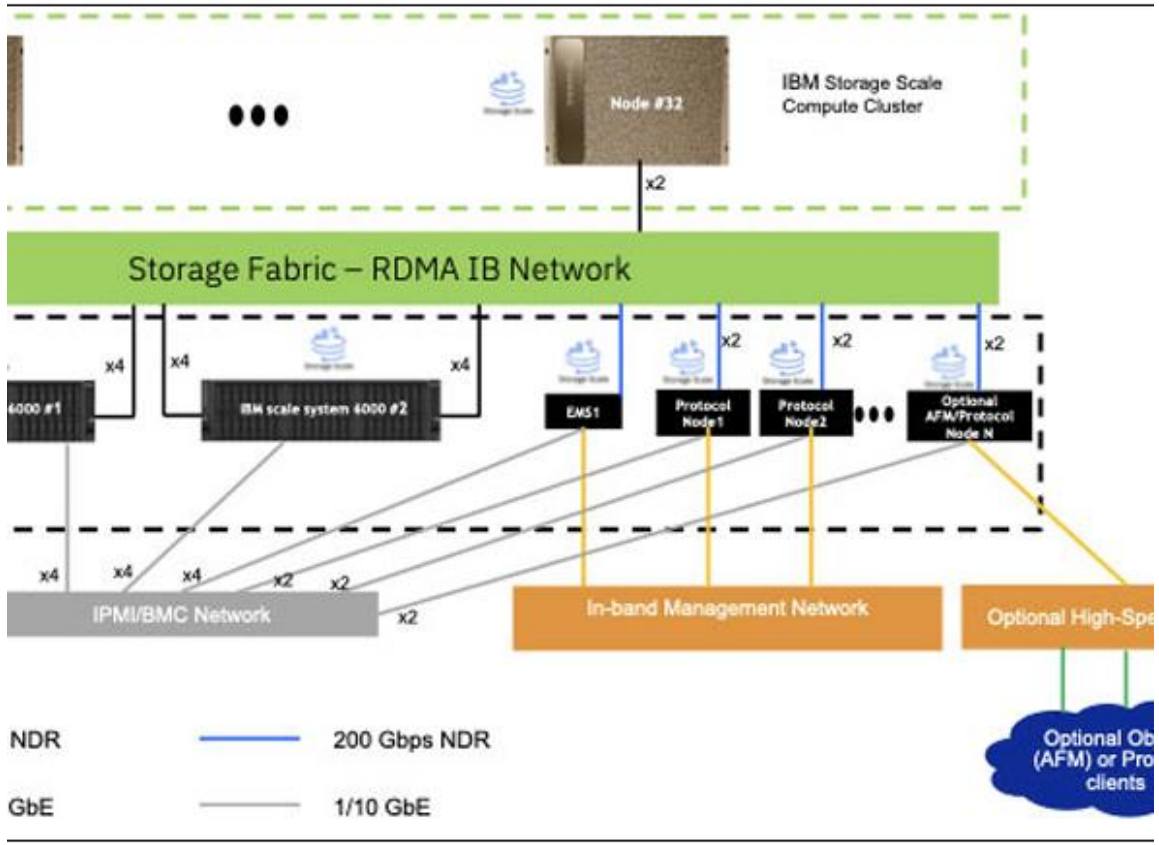

e network diagram with optional AFM connectivity

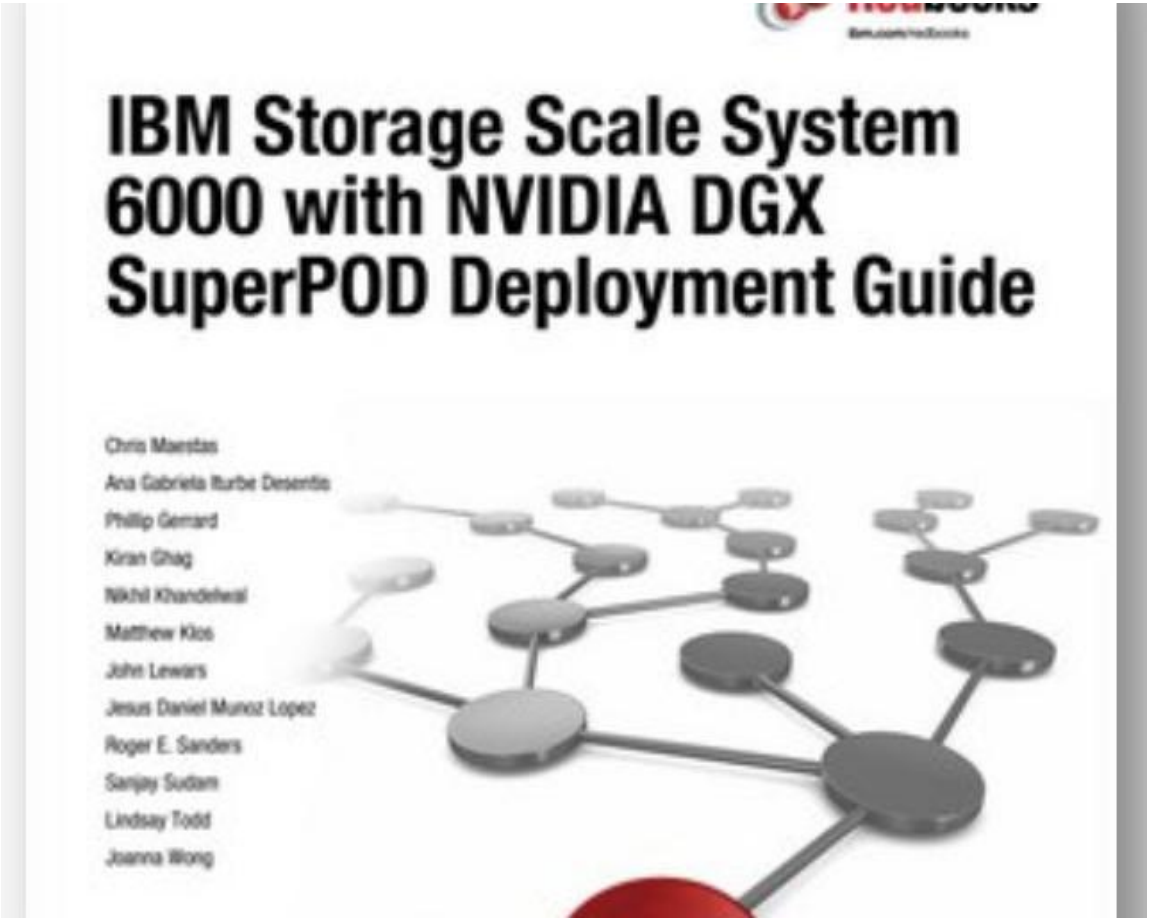

IBM Storage Scale - global platform for storage (gpfs) services of unstructured data

1) Accessibility, 2) Storage Acceleration and Abstraction 3) Data Aware Services, and 4) Assurance

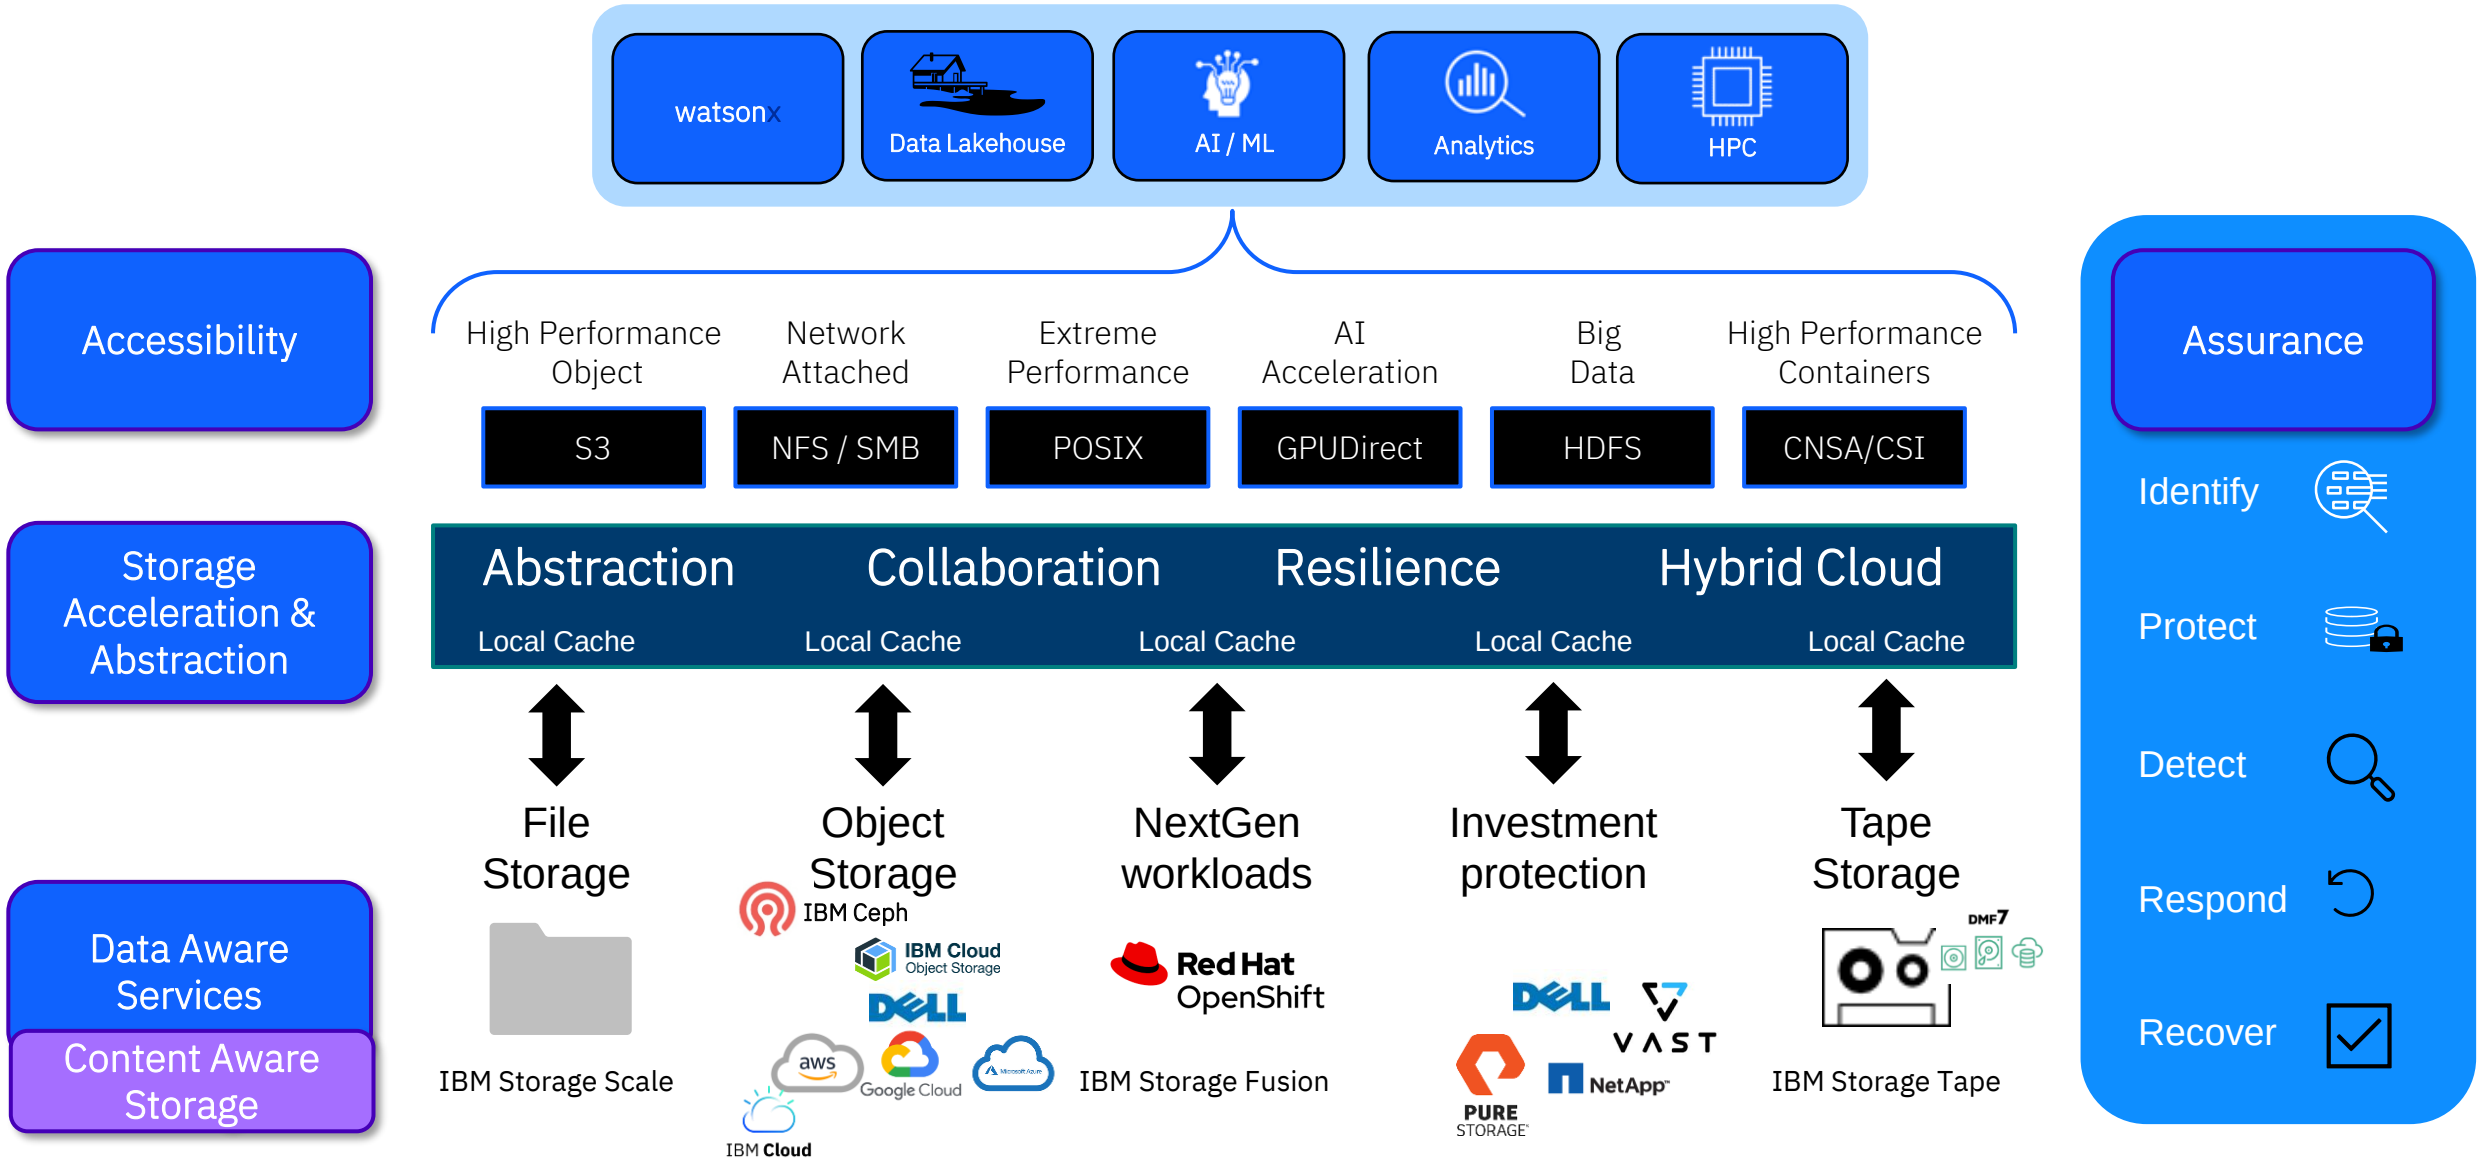

IBM Storage Scale - global platform for storage (gpfs) services of unstructured data

Performance and scalability optimized

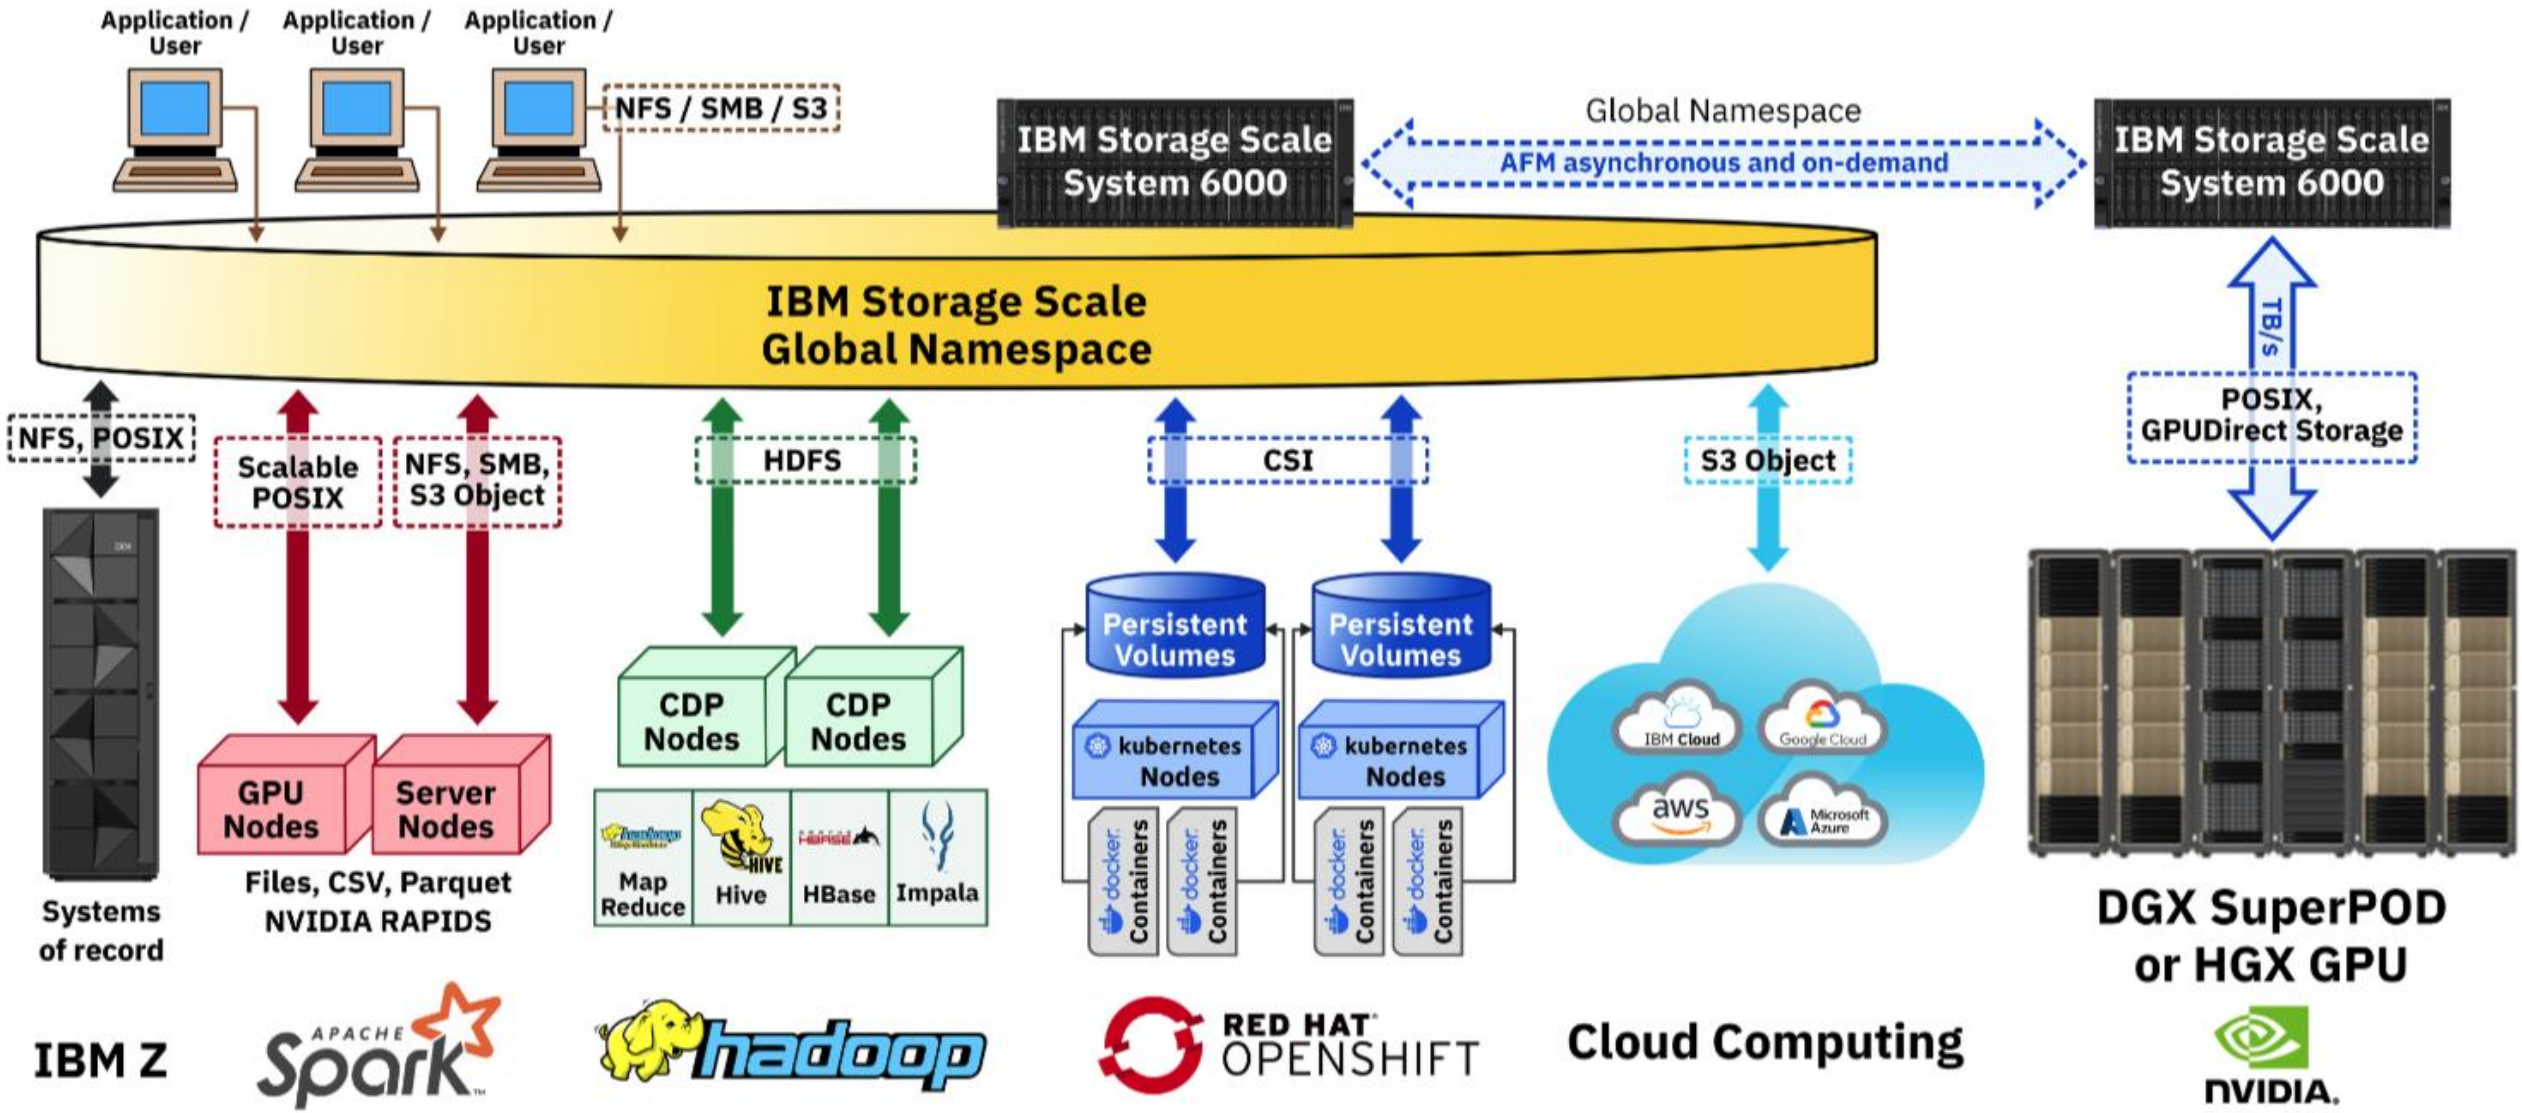

Julich Lab Jupiter Exascale AI: IBM Storage, NVIDIA GPU and ARM

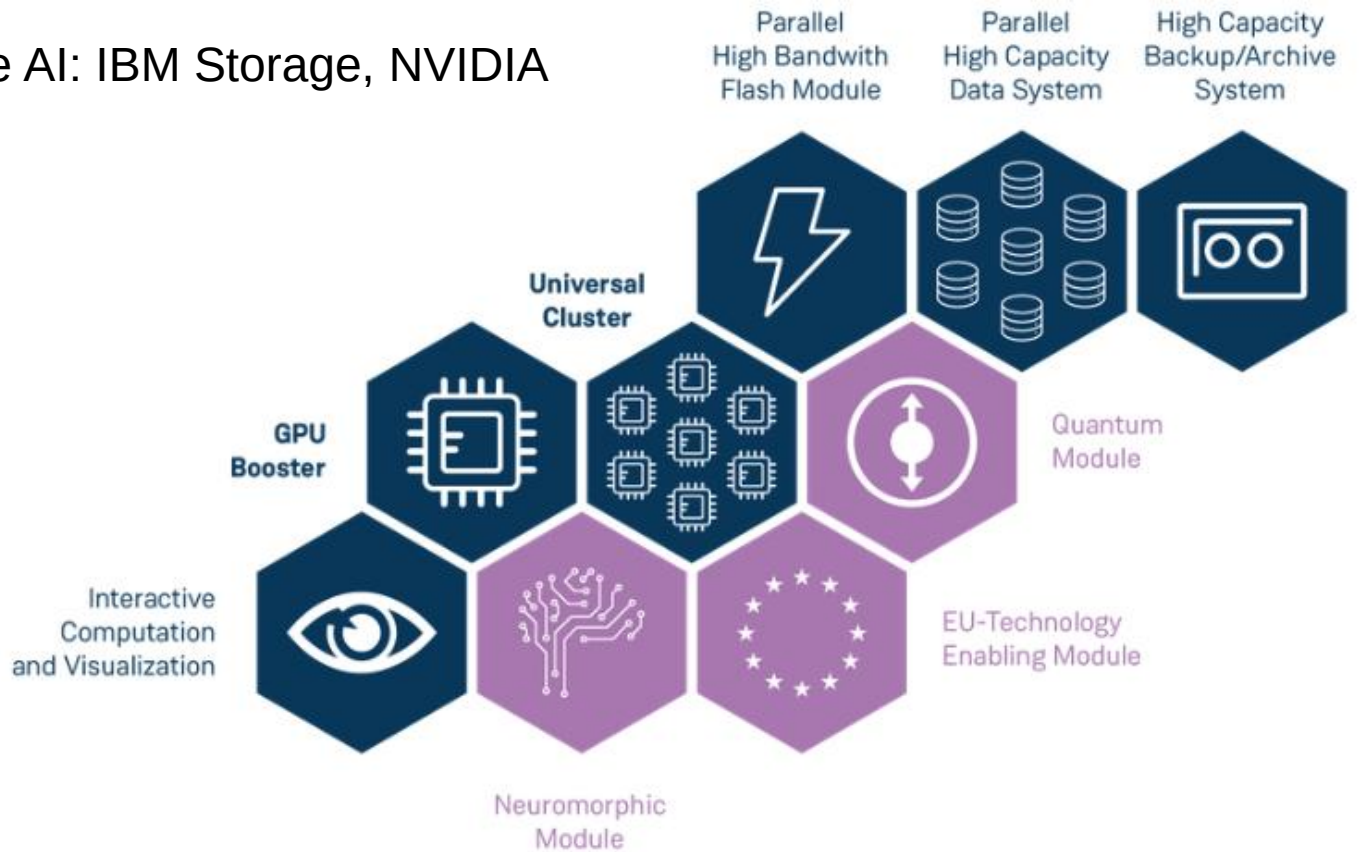

https://www.fz-juelich.de/en/ias/jsc/jupiter/tech

- Scratch storage is 20 IBM Storage Scale 6000 systems utilizing NVMe disk technology, based on the IBM Storage Scale solution.
  - With 29 PB of raw and 21 PB of useable capacity, it targets to provide more than 2 TB/s write and 3TB/s read performance.
- In addition, a high-capacity Storage Module with more than 300 PB of raw capacity as well as a tape infrastructure for backup and archive purposes with over 700 PB capacity will be provided.

IBM Storage Scale deployment model comparison

Storage Area Network (SAN) (NVMeoF, Fiber Channel, iSCSI)

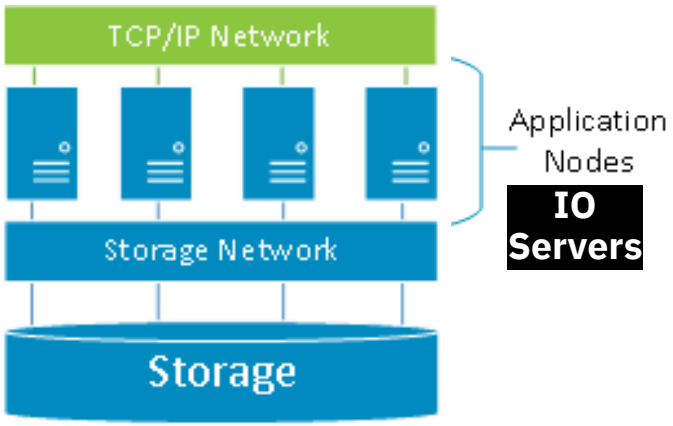

Block based

SAN based filesystem

PureScale

Sailfish

DS8k SDS Scale

Unify and parallelize storage silos

Twin tailed storage with erasure coding

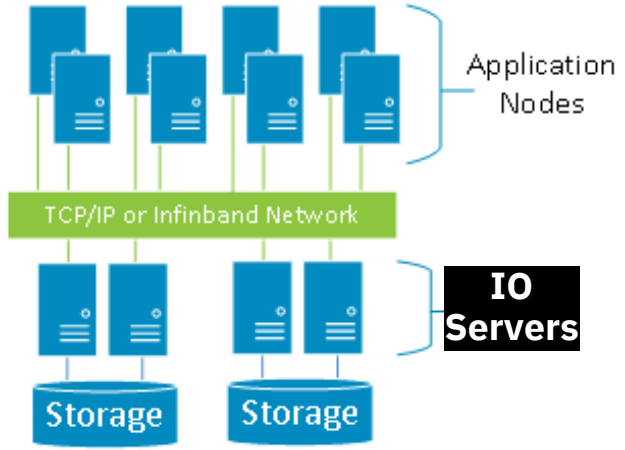

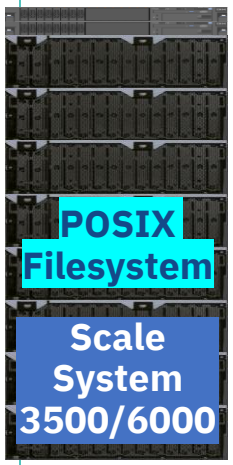

Modular High-Performance Scaling

Shared Nothing Cluster (SNC) Model (Storage Rich Servers (replication, erasure code))

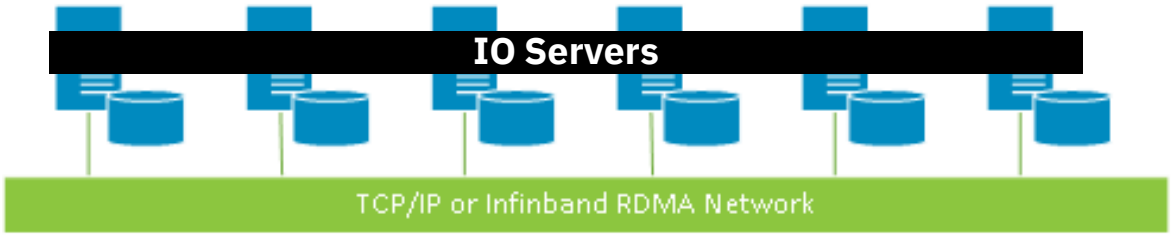

Scale (FPO)

Scale Erasure Code Edition (ECE)

Fusion HCI or SDS

Span storage rich servers for converged architecture or HDFS deployment

IBM Storage Ceph or OSS Ceph

IBM Cloud Object Storage (COS)

IBM Storage Scale Developer Edition

https://www.ibm.com/products/storage-scale

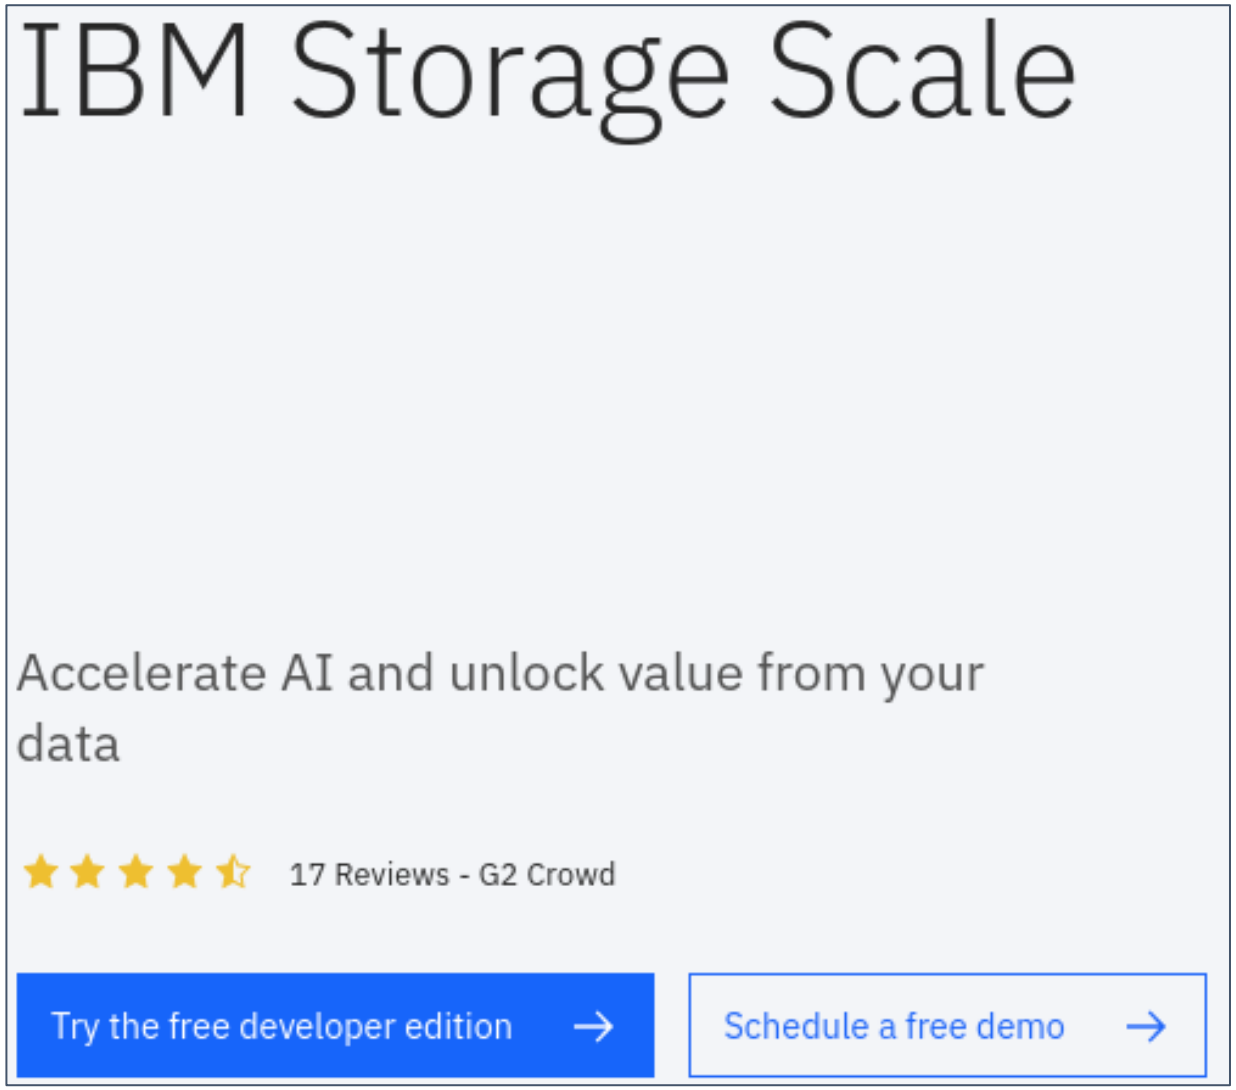

Scale User Group

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

Techzone – https://techzone.ibm.com/collection/ibm-storage-scale-developer-edition-labs

Storage Scale

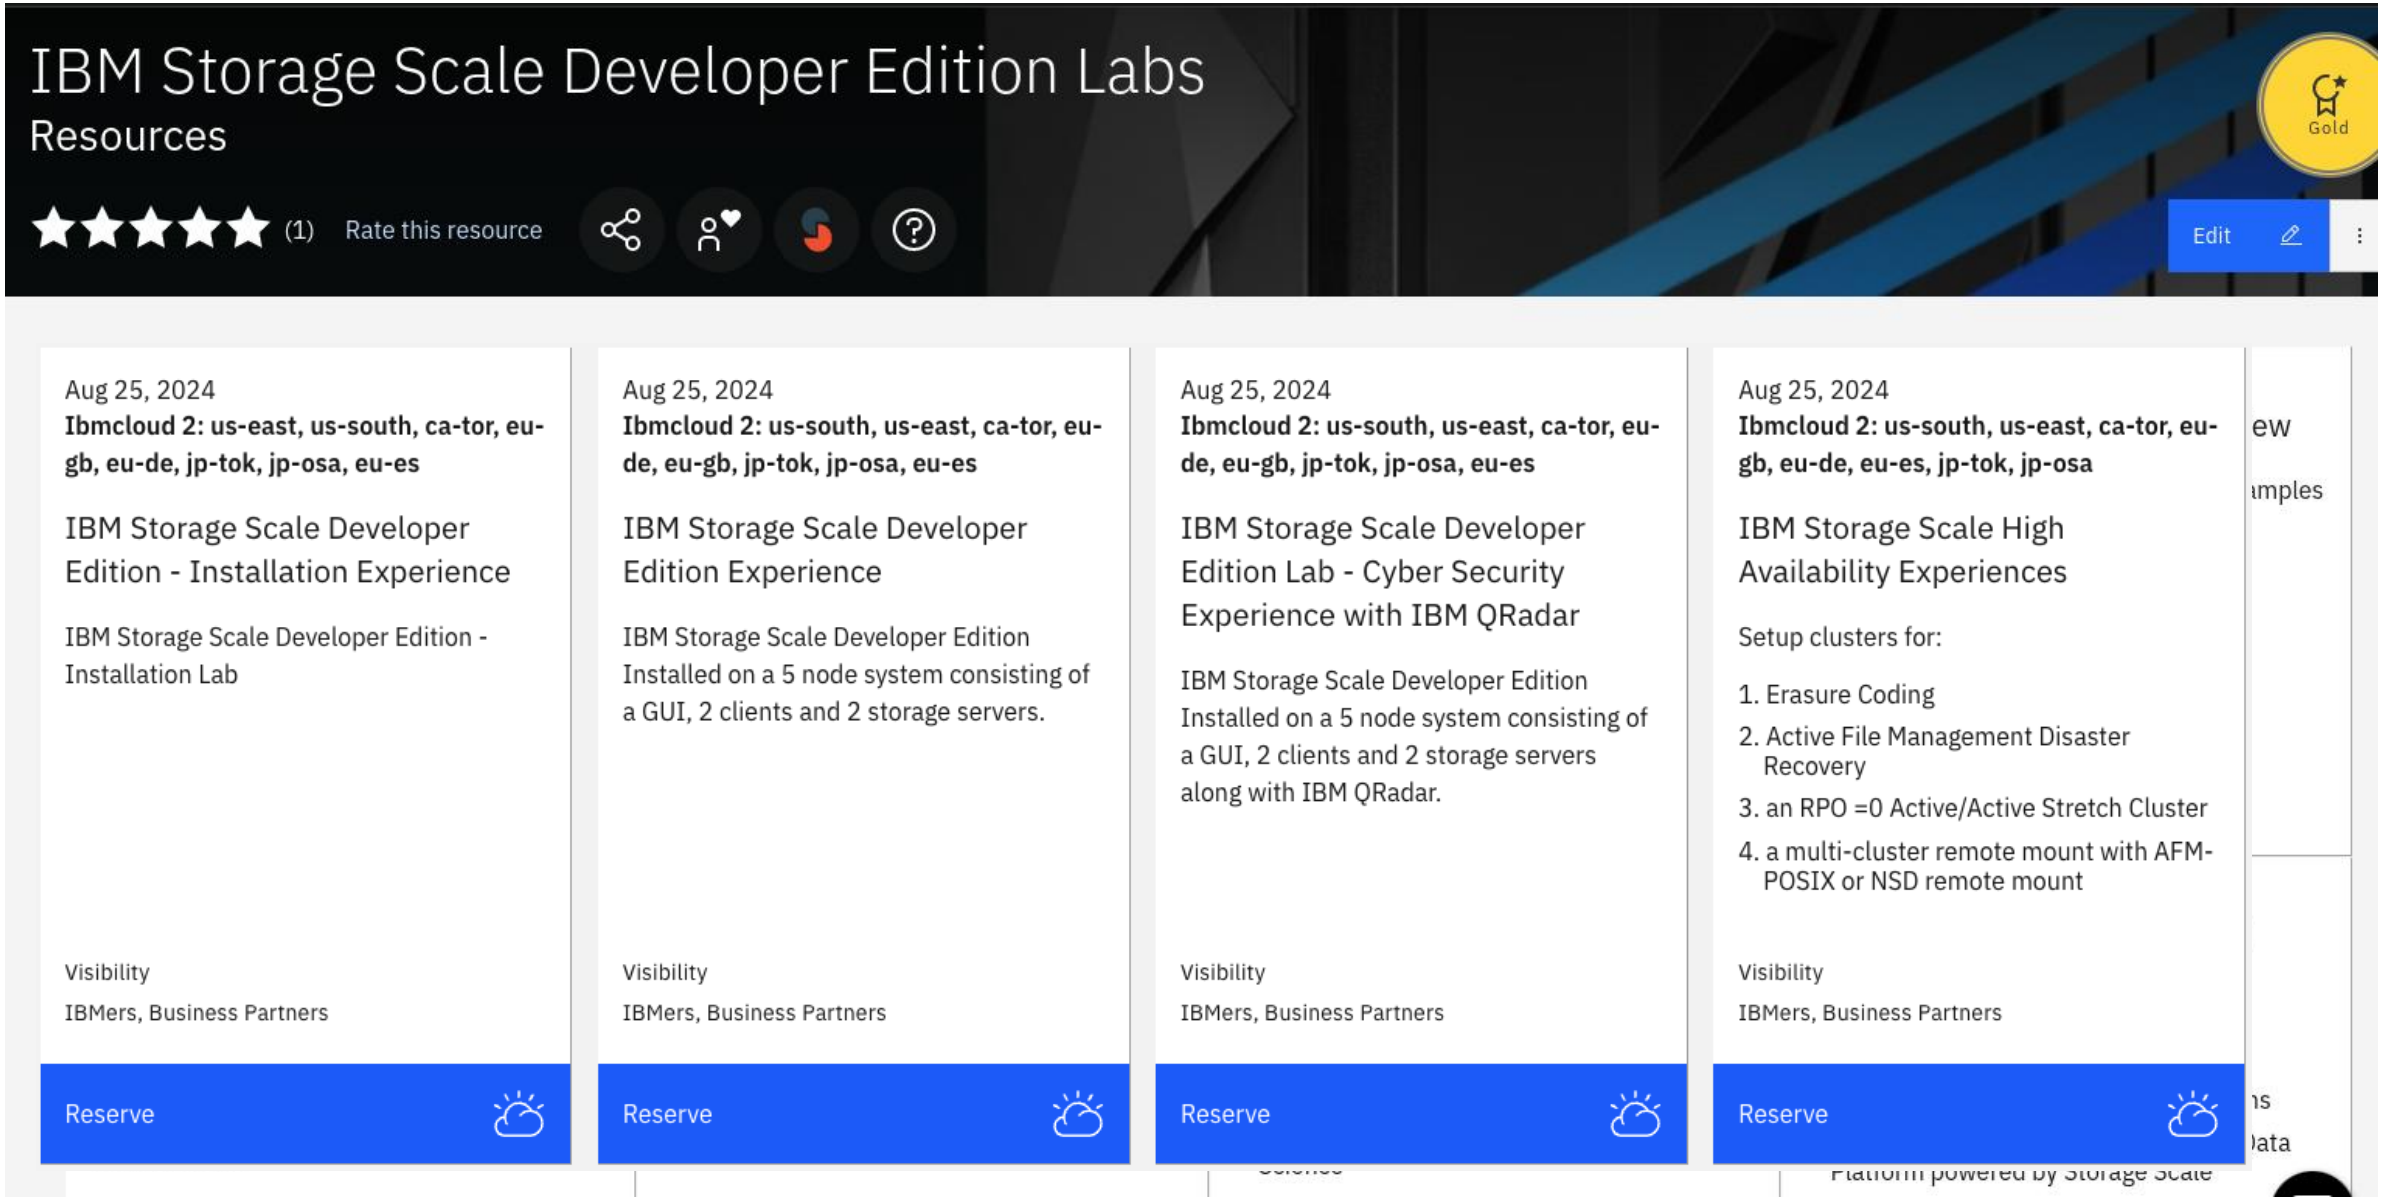

Scale Release Cadence

Extended Update Support goals:

- EUS with PTFs every 18 months
- Extended support on last EUS within a release
- Increase the number of Modification levels with new function
- Scale's Extended Update Support (EUS) approach is outlined in product FAQ

- **EUS release approach applies to non-containerized scale**
- CNSA currently doesn't have an EUS

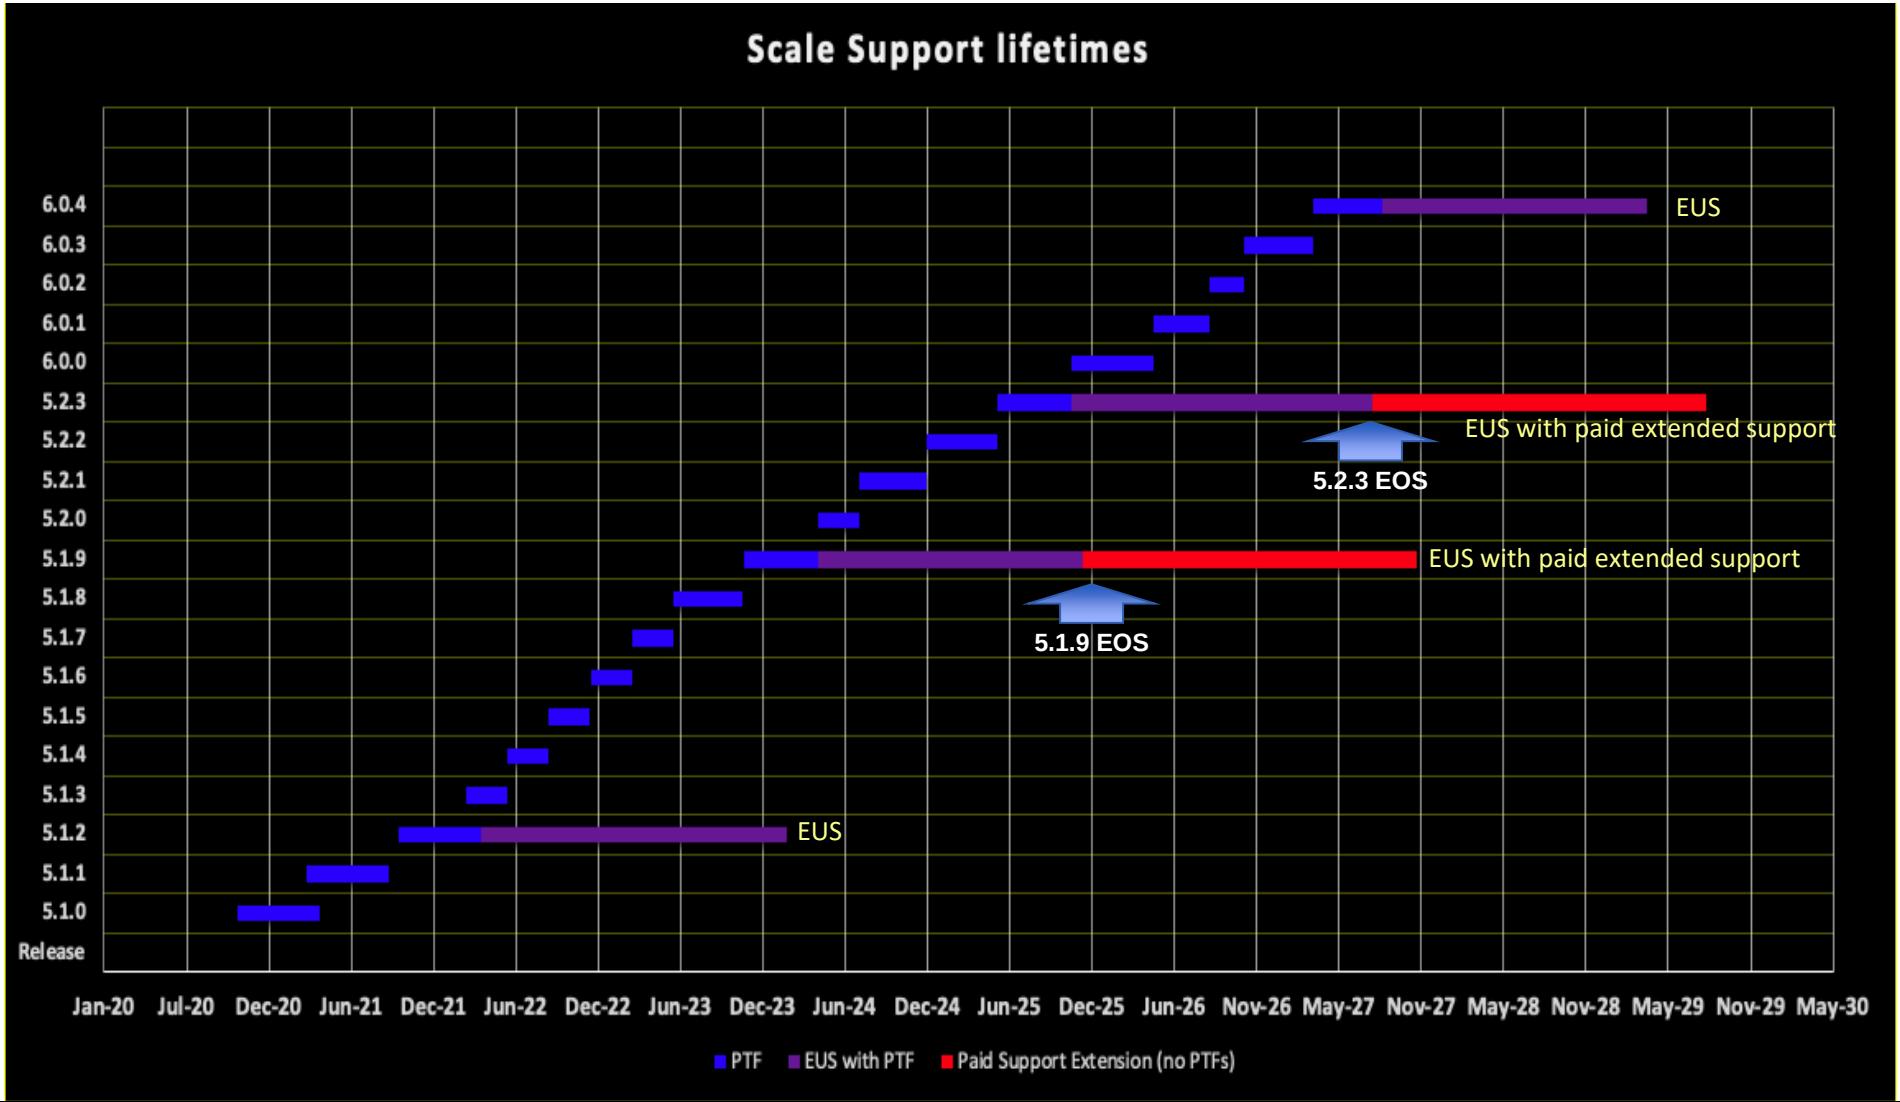

Note: Version numbers and release timing are for example purposes to demonstrate the goal of EUS ~ every 18 months and do **not** represent a commitment to deliver a specific version or on a specific timeline

Release Cadence Goals

Can different IBM Storage Scale maintenance levels coexist?

A2.8:

Different releases of IBM Storage Scale can coexist, that is, be active in the same cluster and simultaneously access the same file system. For release co-existence, IBM Storage Scale follows the N-1 rule. According to this rule, a particular IBM Storage Scale release (N) can co-exist with the prior release of IBM Storage Scale (N-1). This allows IBM Storage Scale to support an online (rolling) upgrade, that is a node by node upgrade. As expected, any given release of IBM Storage Scale can coexist with the same release. To clarify, the term release here refers to an IBM Storage Scale release stream and the release streams are currently defined as 4.2.x > 5.0.x > 5.1.x > 5.2.x.

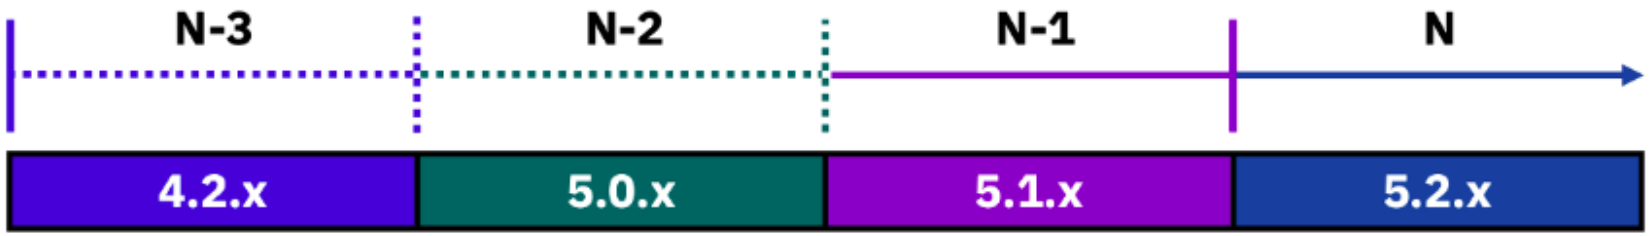

These coexistence rules also apply for remote cluster access (multi-cluster remote mount). A node running release N-2 cannot perform a remote mount from a cluster which has nodes running release N, and vice versa.

Accessibility

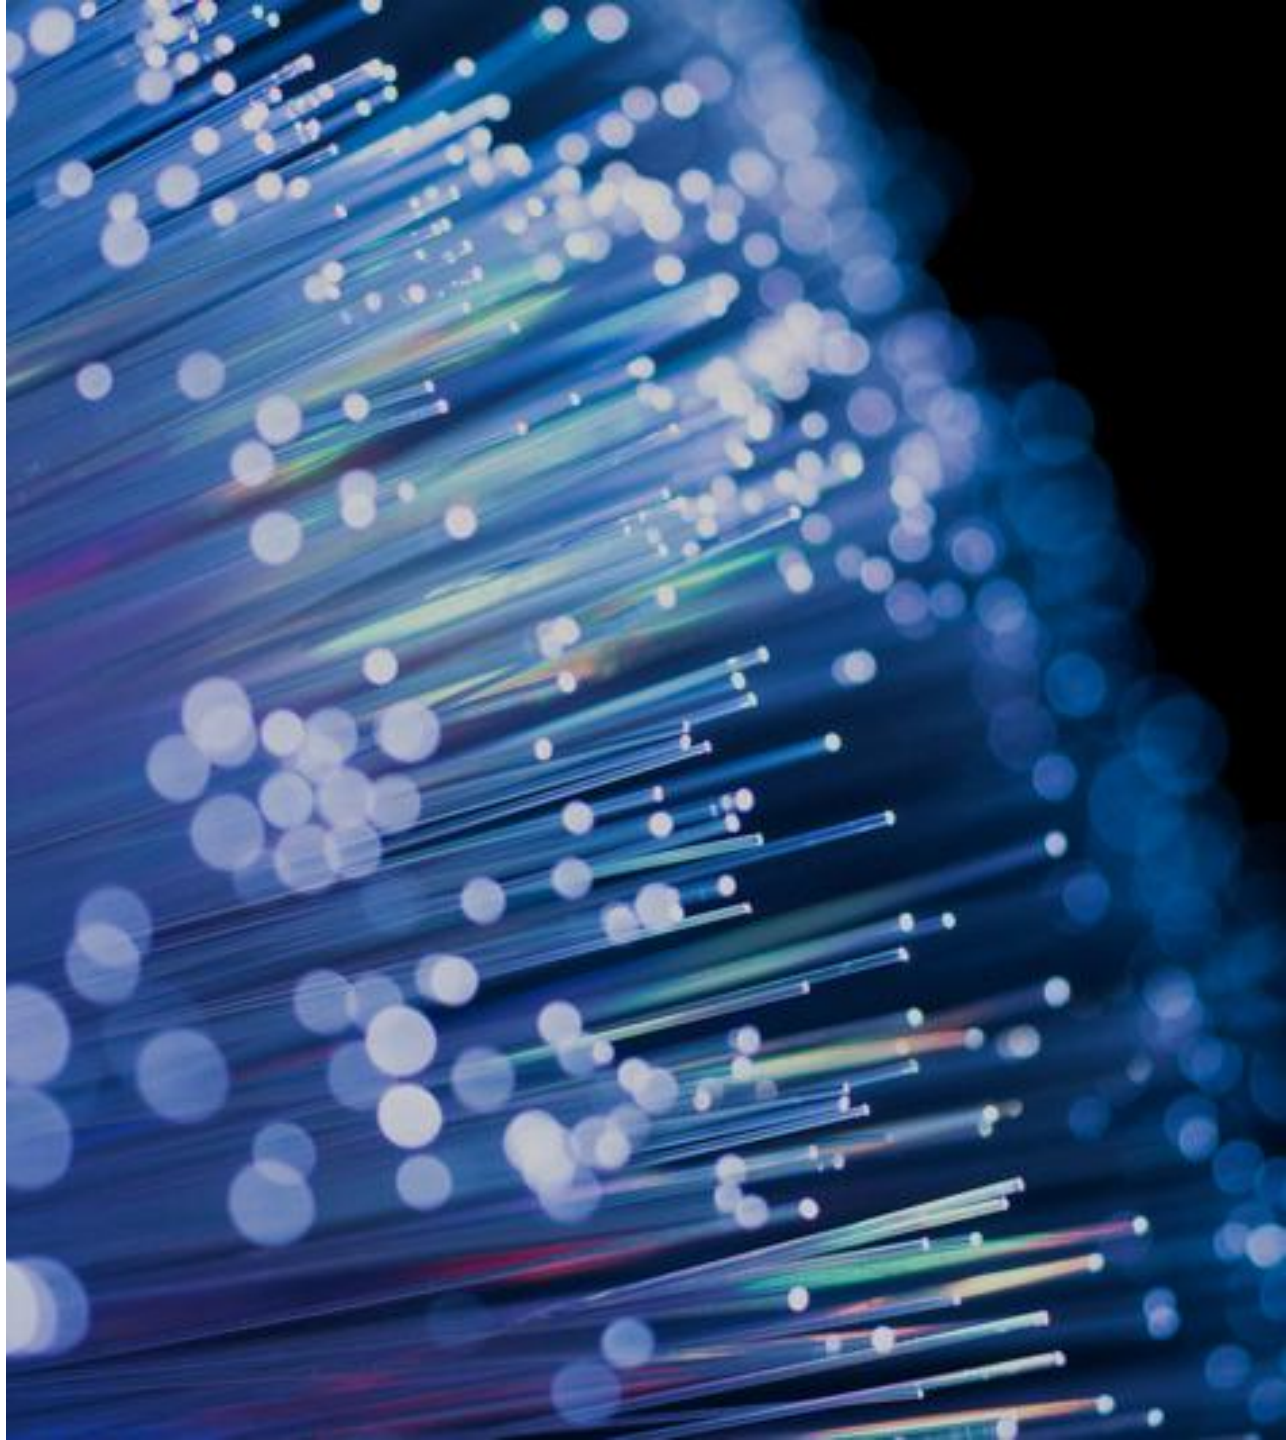

GPUDirect Storage

enables an explicit direct memory access (DMA) between GPU memory and storage when used in the application code

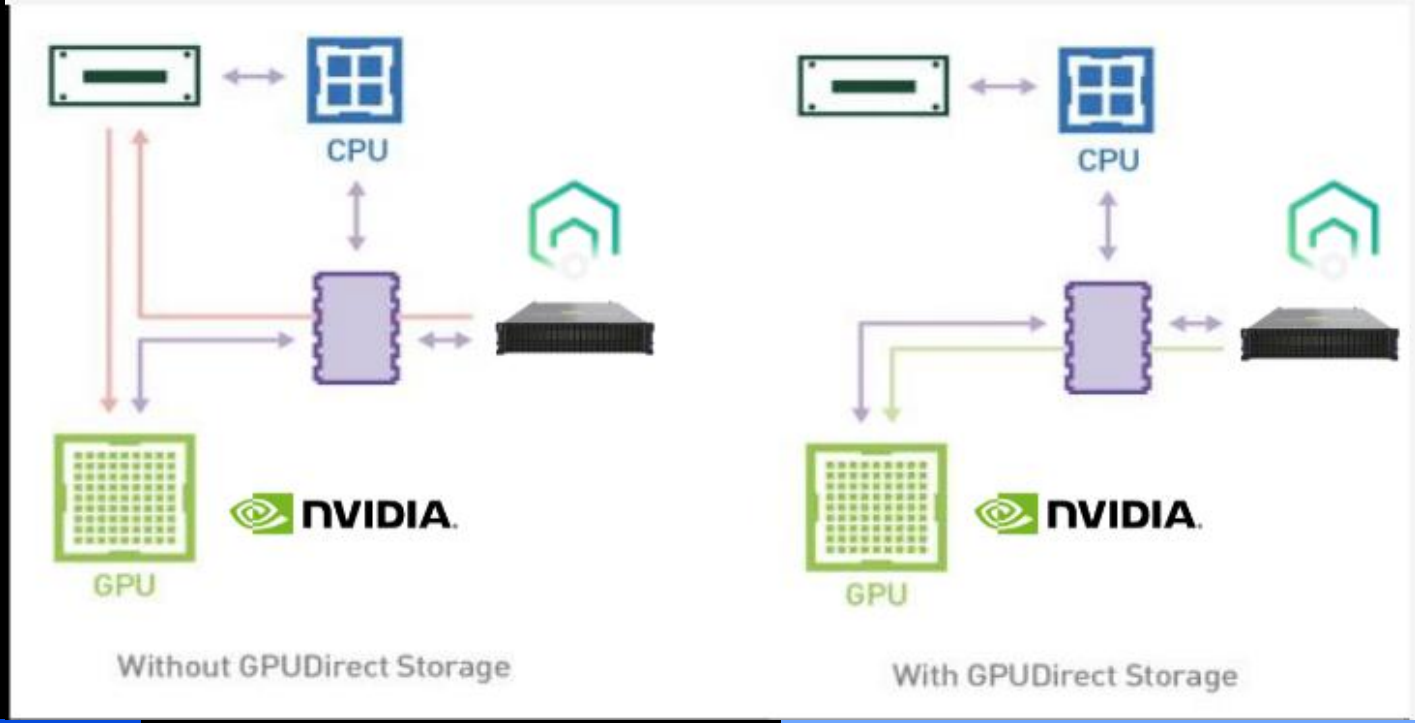

NVIDIA Magnum IO

Family of I/O Optimizations for GPU accelerated data centers

GPU Direct RDMA: Access peer node's memory without copying to host memory

GPU Direct Storage: Transfer data to/from GPU directly from storage without involving CPU and CPU memory

CUDA Toolkit

GDS will be in the CUDA toolkit

Development environment for GPU accelerated applications

Libraries, compilers, debuggers, optimizers, and tools

Leading GPU compute platform since 2006

GDS for Applications

Invoked using the CUDA Toolkit (cuFile) API

APIs must be explicitly called by the applications

Storage must be GDS enabled. If not, GDS call falls back to regular data movement.

Why it matters

AI, HPC, Analytics are data hungry and require a very high data throughput.

GPUs are starved by slow I/O (and NFS is particularly slow)

NFS, SMB, HDFS

Support and Currency:

- Added mmces service restart option – (5.2.3.1)
  - For NFS will attempt to restart when nfs_hung event occurs
    - If restart fails will move CES IPS
- Updated Samba on IBM Storage Scale to the upstream Samba 4.20 version.
  - Fix slowness of SMB sharesec (**mmsmb exportacl**) command when the number of Samba shares increases.
  - Fix issue in smbd deadlock on two locking DBs.
- IBM Storage Scale supports the NFS 4.2
- The **mmnfs export load** command can process configuration files without restarting NFS CES.
- Added support for CES HDFS Transparency 3.3.6-0 with Open Source Apache BigTop 3.2.1 on RHEL 8.x and RHEL 9.x

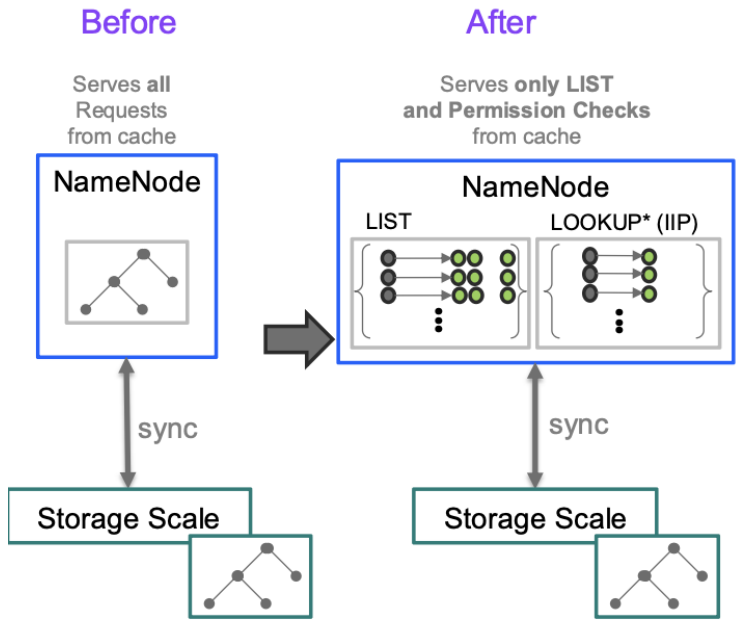

Improved performance:

- HDFS transparency metadata design update
- **High Performance SMB - Continued partnership with Tuxera and investigating others**

High Performance Object 2.0 (CES-S3)!

Storage Scale

Multi-protocol data access support with POSIX, S3, NFS, SMB and CSI

Support and Currency:

- Supports S3 bucket notifications feature for events such as object creation and object removal.
  - Configure webhooks to get these notifications.
  - Swift is discontinued - Use 5.1.8 Swift code in CES of 5.1.9 but migrate soon!

ILM support including Tiering to Tape support via RPQ

S3 over RDMA - initial implementation - https://github.com/noobaa/noobaa-core/pull/8817 - work to be done for CES-S3

Client Experience:

2 billion objects in a single bucket ! - https://github.com/ghcoelhopsa/scale_s3_benchmark

Client upgrade to 5.2.3 while uploading 20 million objects

Targeting 5 Billion Object upload test now

IBM Technology Expert Labs can provide billable migration services (Swift to CES S3 and DAS S3 (HPO 1.0) to CES S3 (HPO 2.0))

Improved performance:

- IBM Storage Scale CES S3 Performance evaluation of large and small objects using COSBench: https://community.ibm.com/community/user/storage/blogs/rogelio-rivera-gutierrez/2024/04/25/ibm-storage-scale-performance-ces-s3-tech-preview

Container Native Storage Access (CNSA)

Improvements introduced in CNSA 5.2.3.0

OpenShift and regular Kubernetes!

	Supported Versions
Kubernetes	1.29, 1.30, and 1.31
Architectures	x86_64
RHEL	8.10, 9.4 and 9.5
Ubuntu	22.04 and 24.04

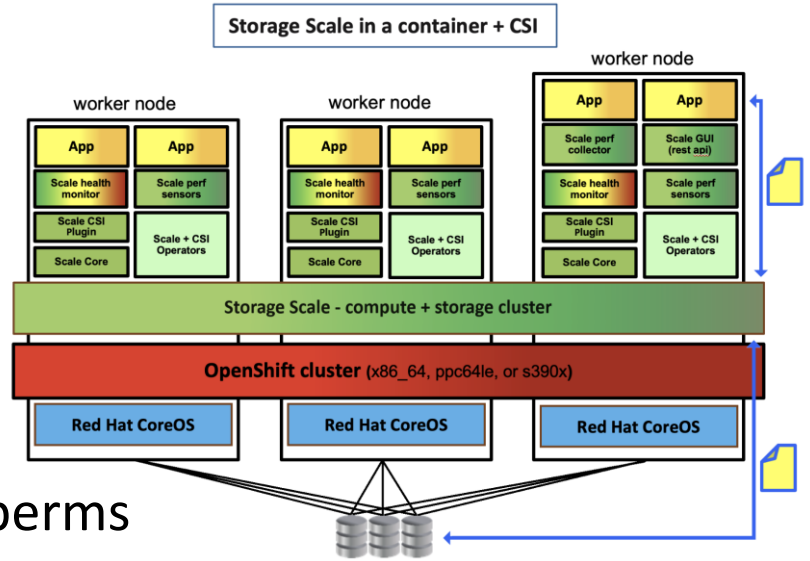

- 5.2.3 GUI: **ContainerOperator** user role includes **CsiAdmin** user role perms
- Support for Google Kubernetes Engine (GKE) is technology preview

- AFM Cache Volumes Support

- RDMA support with IBM Storage Scale container native
  - RDMA over Converged Ethernet in Scale 5.2.3 (InfiniBand in 5.2.2)

- Shared disks can be used (have connection to at least two K8s nodes)
  - SNC disks are not supported (have connection to one K8s node)
- Red Hat OpenShift Virtualization (Fusion Access)
  - Goal is to gain VM use cases and customers
  - Uses CNSA local file system shared disk functionality: *SAN attach FC / iSCSI*
    - Announced @ Red Hat Summit

Demos are available to show!

Works with OpenShift Virtualization Engine!

Container Storage Interface (CSI)

Improvements introduced in CSI 2.14.0

Storage Scale

Upgrades for OpenShift, Kubernetes and Ansible as well as improved functionality that support simpler administration and configuration.

- RedHat OpenShift 4.[16-18] (x86_64 architecture)
- Vanilla Kubernetes 1.[29,30-32]
  - (x86_64 and ppc64le)
- Ubuntu 22.04/24.04 (x86_64 architecture) with Vanilla Kubernetes.
- Support for linking fileset at customized path
- Support for enabling dynamic provisioning workflow for pre-created volumes by using a custom storage class parameter

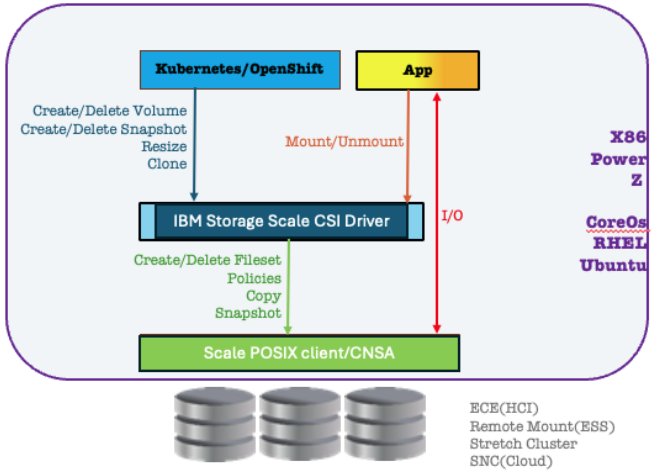

Dynamic Provisioning - Create/Delete Volume
Static Provisioning
Volume Snapshot
Volume Expansion
Shallow Copy
Volume Cloning
Compression
Tiering
ConsistencyG
Remote Mour
Multiple Filesys
FsGroup
GUI HA
Lightweight Volume/Fileset Based Volu

Works with NVIDIA Grace systems running RedHat or Ubuntu!

Acceleration and Abstraction

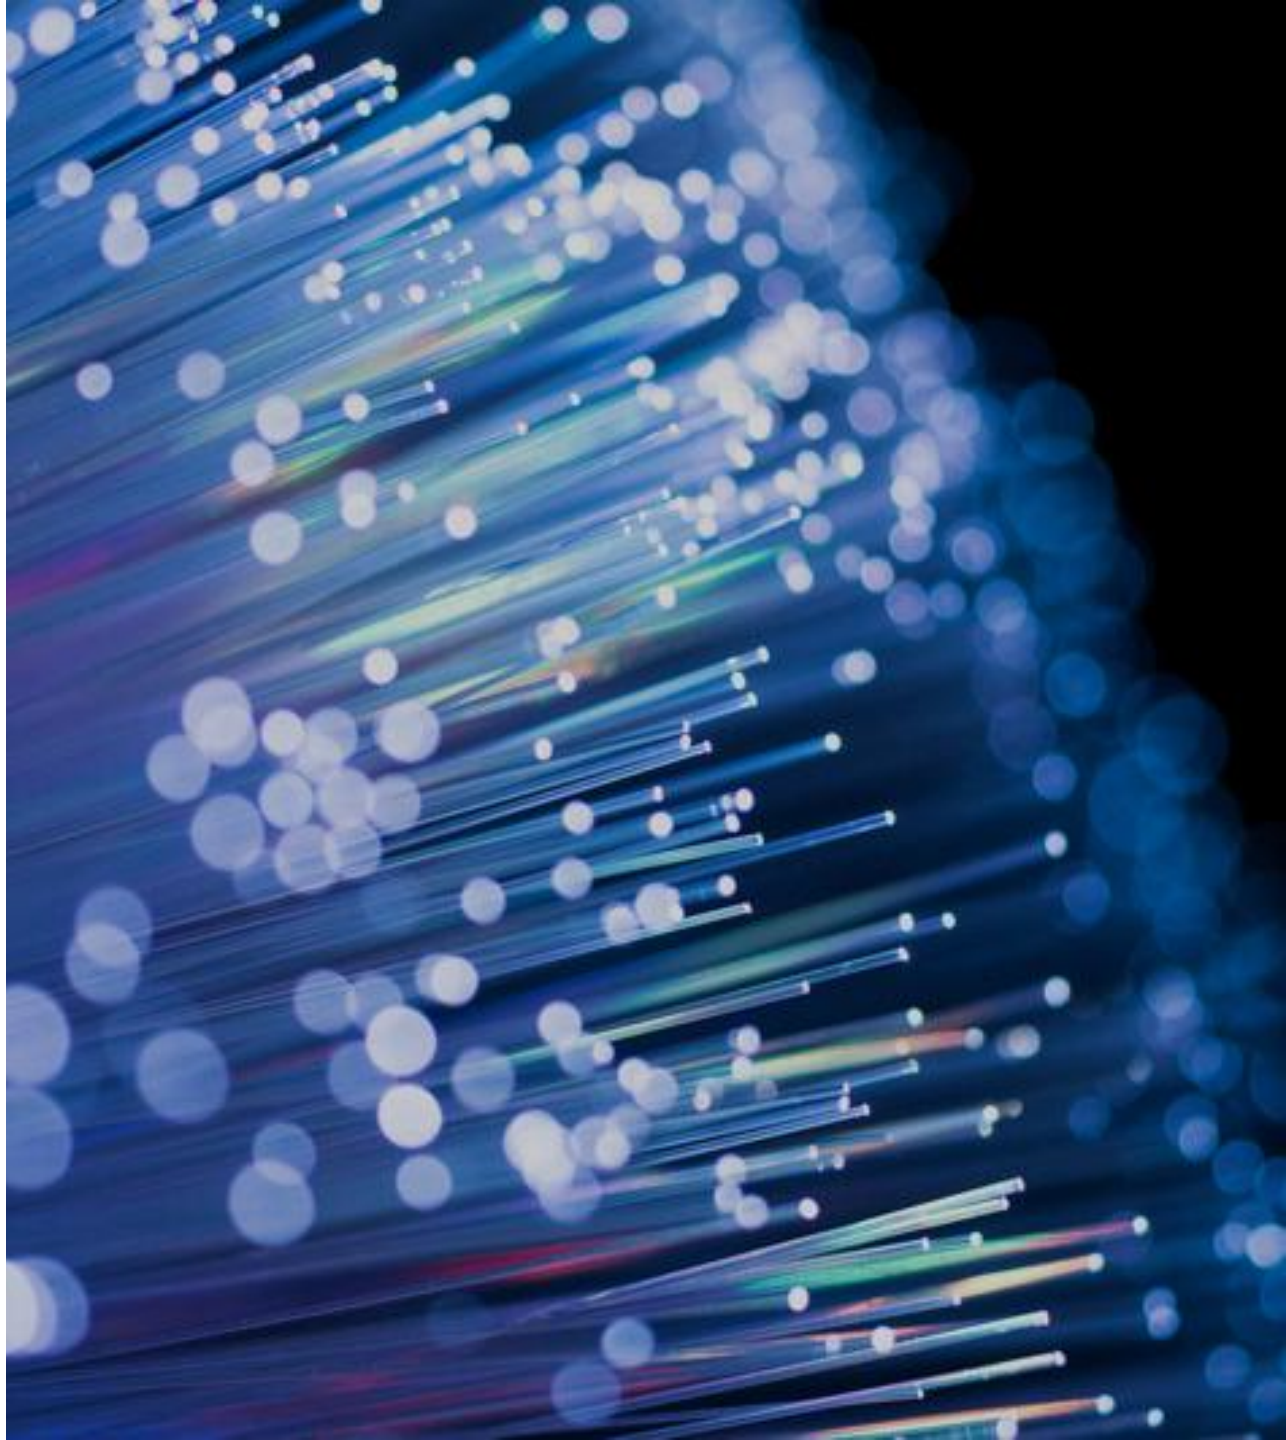

Default configuration changes with 5.2.N

Storage Scale

Provide better out-of-the-box performance for a wide variety of workloads.

Apply only for new 5.2.0 clusters. Do not apply for existing clusters, even with a 5.2.0 upgrade.

The new defaults are described in the mmchconfig man page!

config option	old default	new default
numaMemoryInterleave	no	yes
workerThreads	48	256
page pool	min(1G, 1/3 system mem)	min(4G, 1/3 system mem)
ignorePrefetchLUNCount	no	yes
dioRentryThreshold (undocumented)	0	1

Performance update from 6000

Storage Scale

- **SED ENABLED!**
- **5x iterations:**
- Mean Write:
  - 156.35 GB/s
- Mean Read:
  - 320.49 GB/s

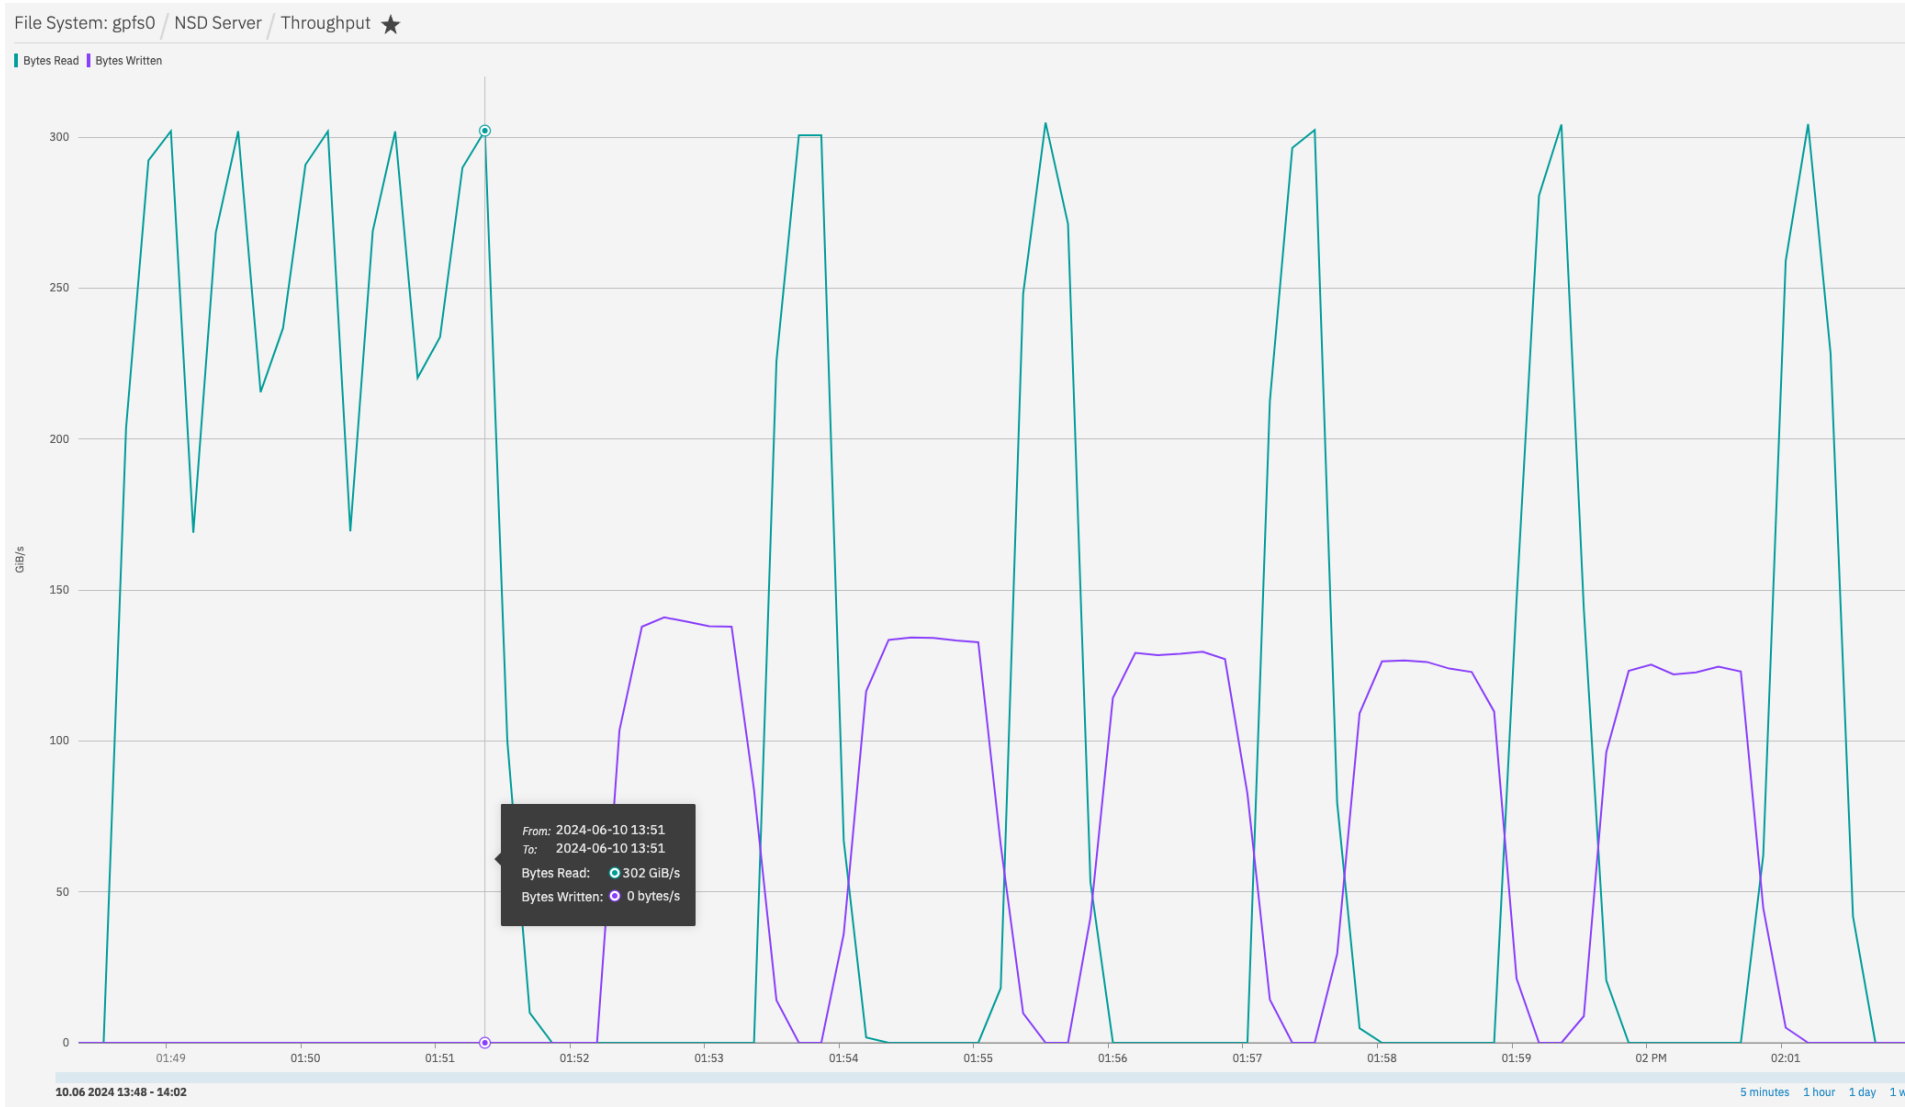

Filesystem Core improvements

Storage Scale

- The dynamic pagepool feature is now supported on NSD server nodes.
  - Note that this does not imply Scale System I/O nodes

- static pagepool stays STATIC when RDMA is enabled.
  - Use a dynamic pagepool when the size needs to change while running.

- The **`mmcrfileset`** command supports the **`-J`** parameter.
  - No longer needing 2 commands to create and then link a fileset!

- Use the LD_LIBRARY_PATH wisely! - https://access.redhat.com/solutions/49292
  - Observation that shared paths (e.g. /ibm/fs1/dir/) usually grant permissions to everyone.
    - Run: **chmod a+x** on shared directories as much as possible.
    - New feature to check directory ACL
      - this is cached with an in-memory flag to shortcut permission check
  - No other config changes or file system format changes necessary, this is a node-local optimization!

Active File Management (AFM) overview

Storage Scale

Active File Management enables caching data across other data sources.

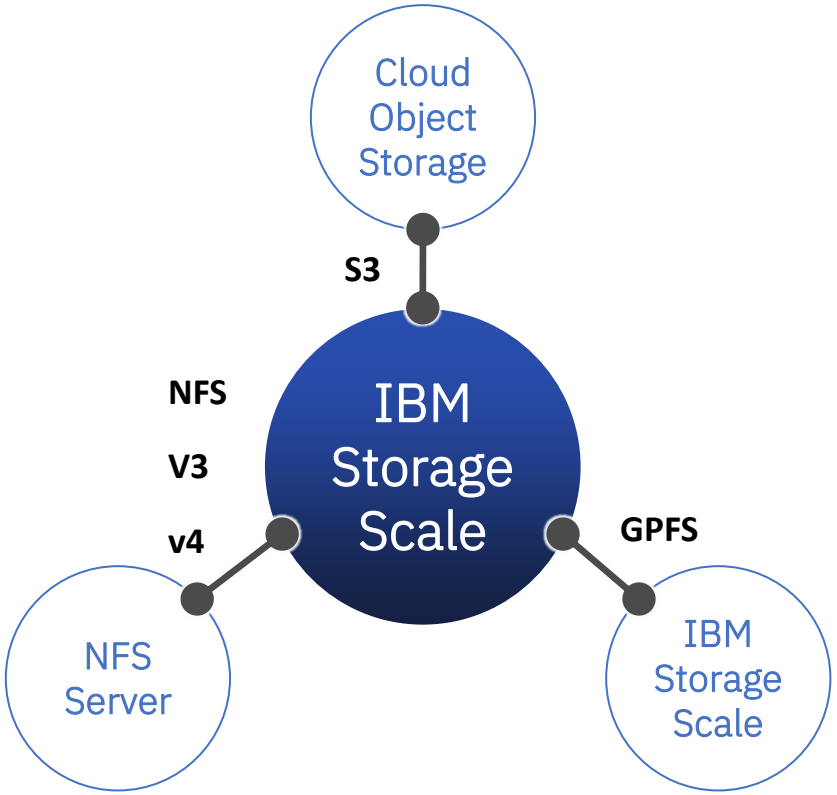

- Each AFM fileset has a distinct set of AFM attributes.
- An IBM Storage Scale cluster that contains AFM filesets is called a cache cluster.
- A cache cluster has a relationship with another remote site called the home, where either the cache or the home can be the data source or destination.
- A cache cluster must be an IBM Storage Scale Cluster.
- A home can be an IBM Storage Scale, NFS server and Object Storage.

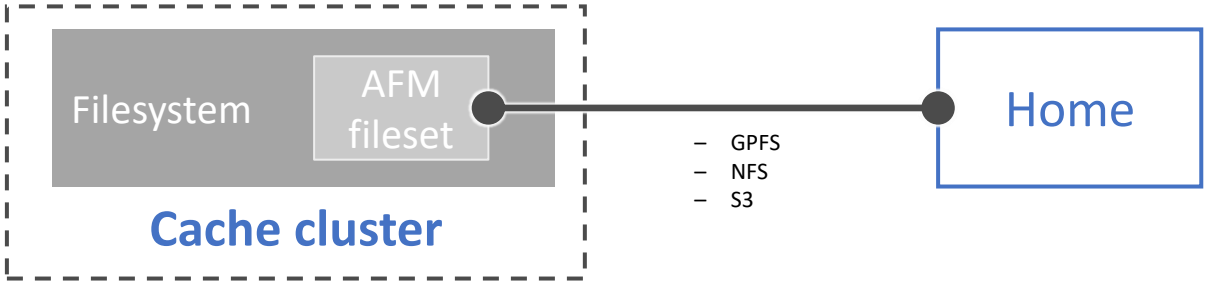

Abstraction: Active File Management (AFM)

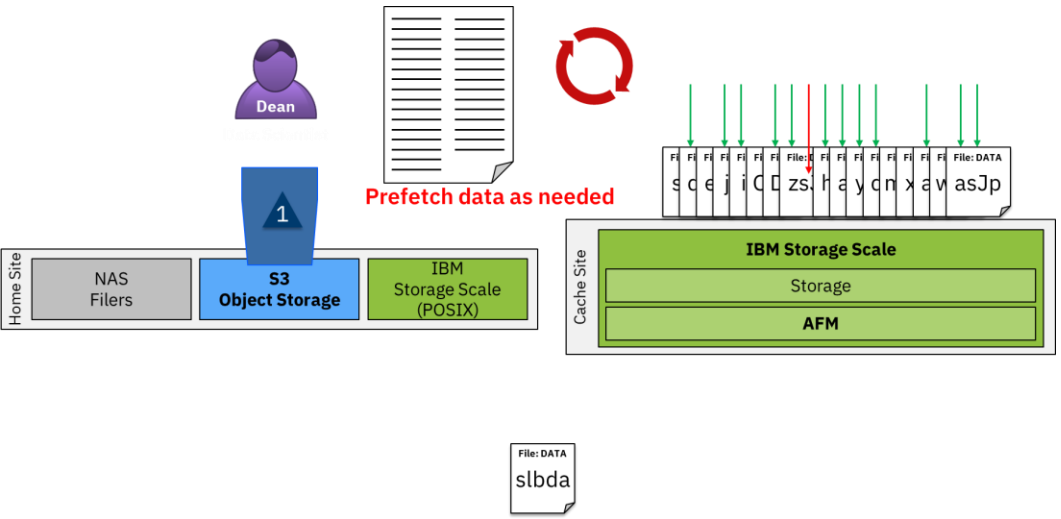

1. Support of adding failed file list under AFM Cloud Object Fileset
2. Allow numeric uid/gid for **mmafmcosctl** and **mmafmcosconfig**
3. Support for downloading object and directory with single list with COS backend

Support for Content Aware Storage (CAS)

- Watch folder support for AFM COS filesets
- Async Auto Prefetch
  - **afmAsyncPrefetchInterval=Xs**
- Fast ReadDir2
  - **'afmObjectFastReaddir2=yes|no**
    - **Sync deletion of data from S3**

Data Aware Services

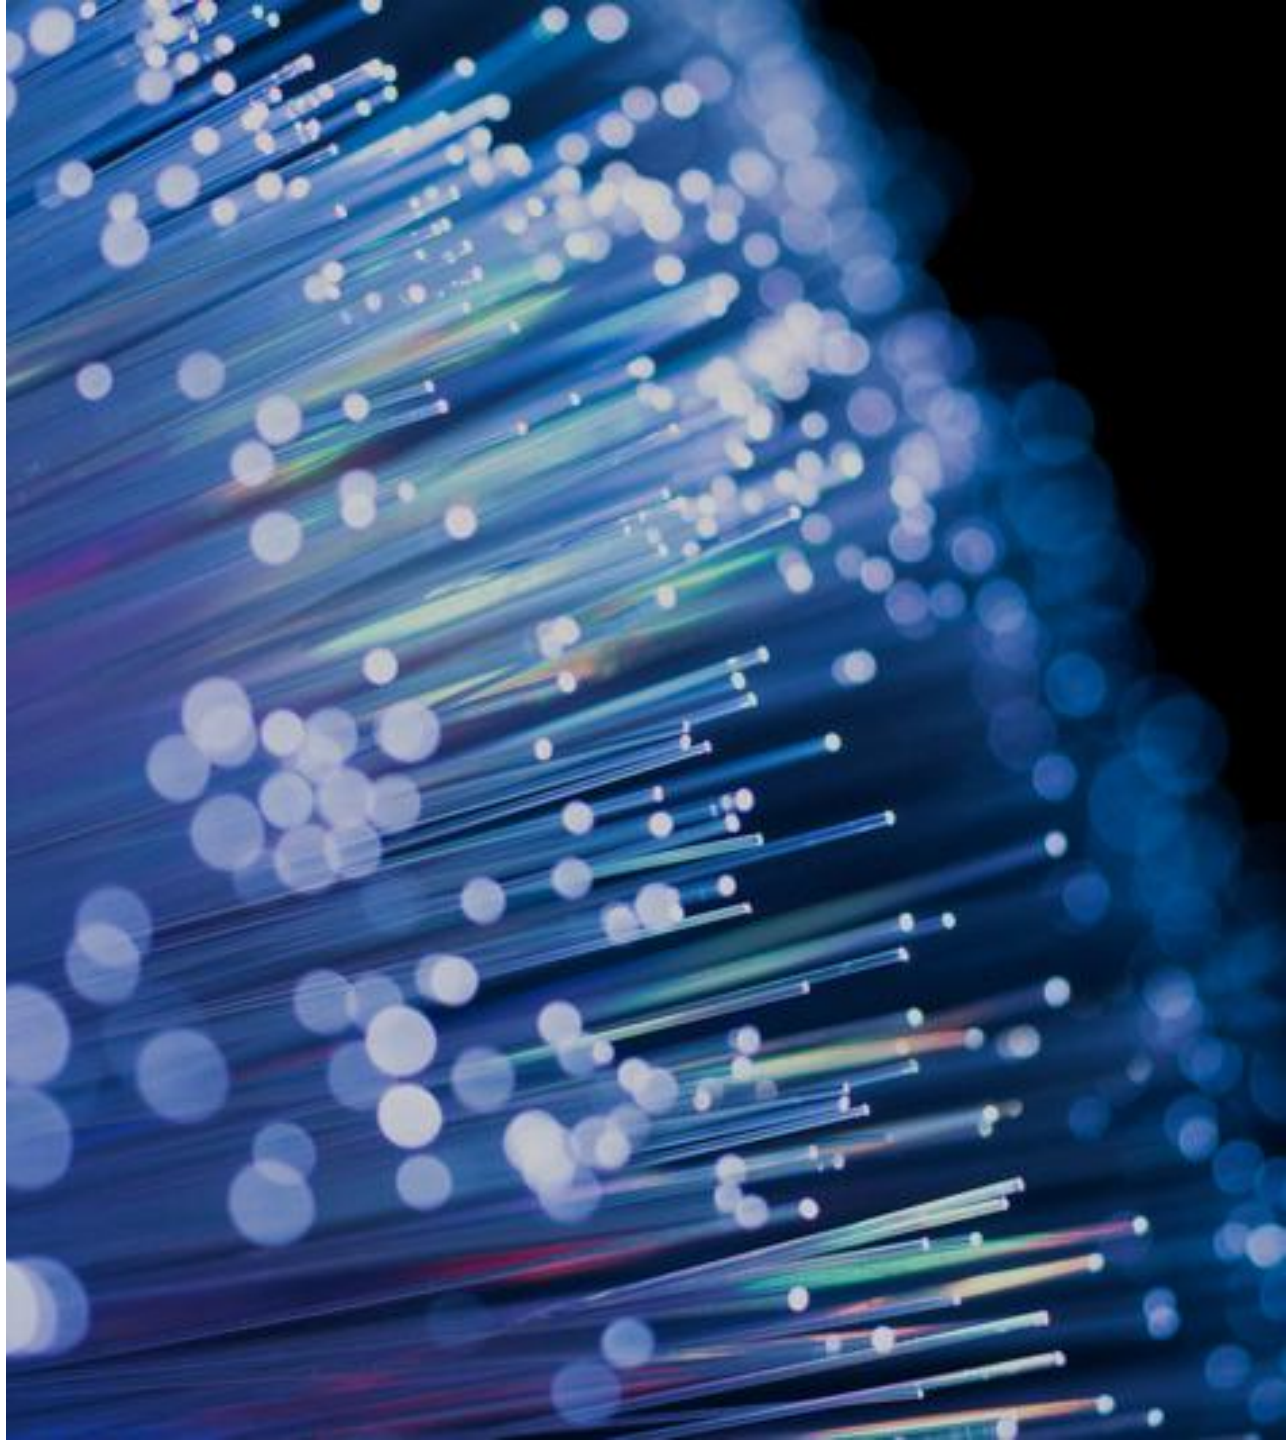

Data aware performance optimization and Quality of Service (QOS)

Storage Scale

Total IO Capability
B/W or IOPS

Management
Capped at 20%
(example)

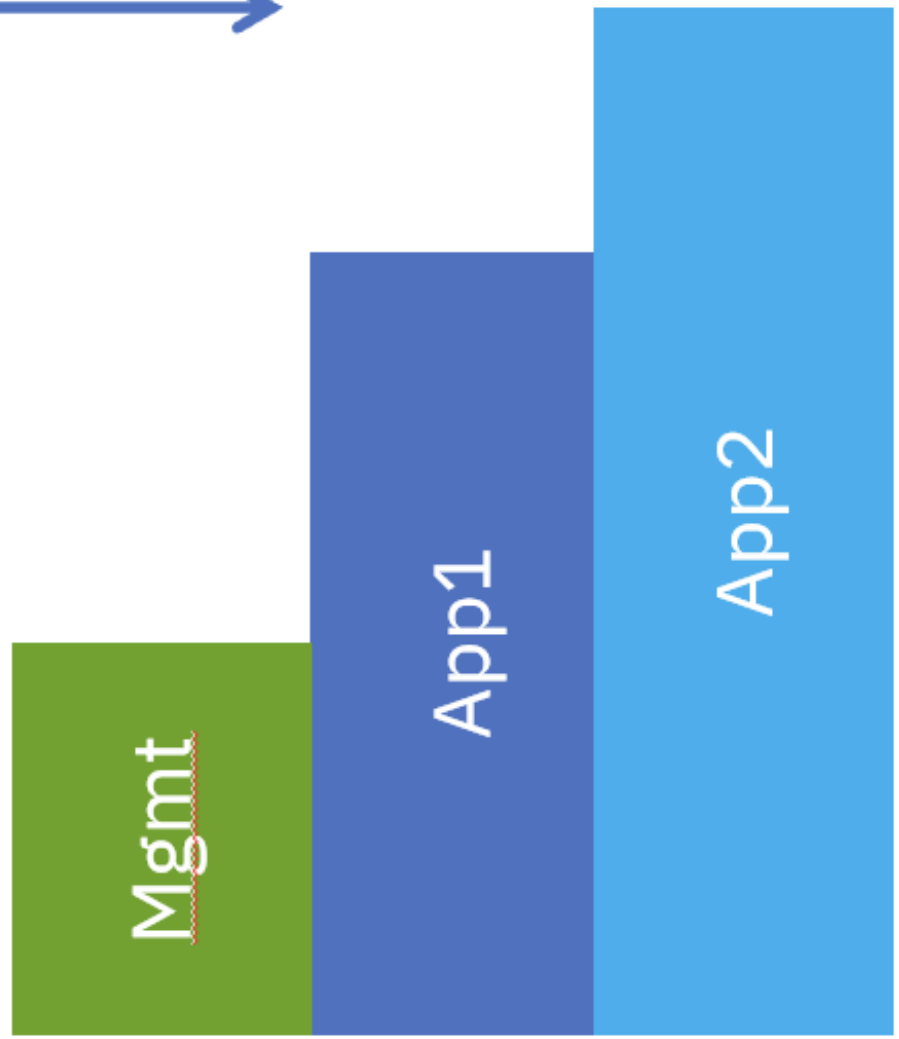

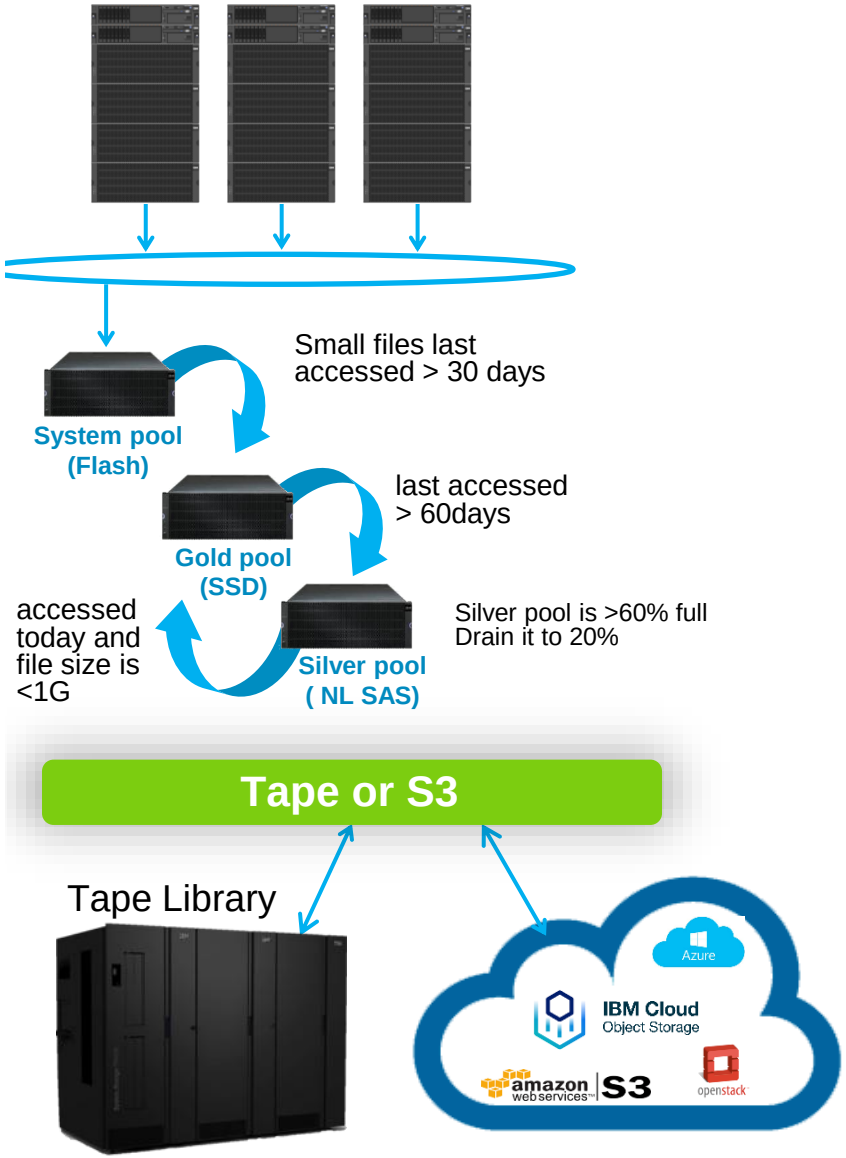

Content Aware Storage (CAS) optimized storage provides differentiation

Storage Scale

IBM AI optimized storage and AI runtime

- **Enterprise grade secure:** Consistent ACL and encryptions
- **Efficient:** Detect data change for incremental data processing
- **Support your legacy storage:** Connect to heterogeneous storage systems, including legacy unstructured data storage
- **Accelerate and scale:** GPU-optimized storage solution

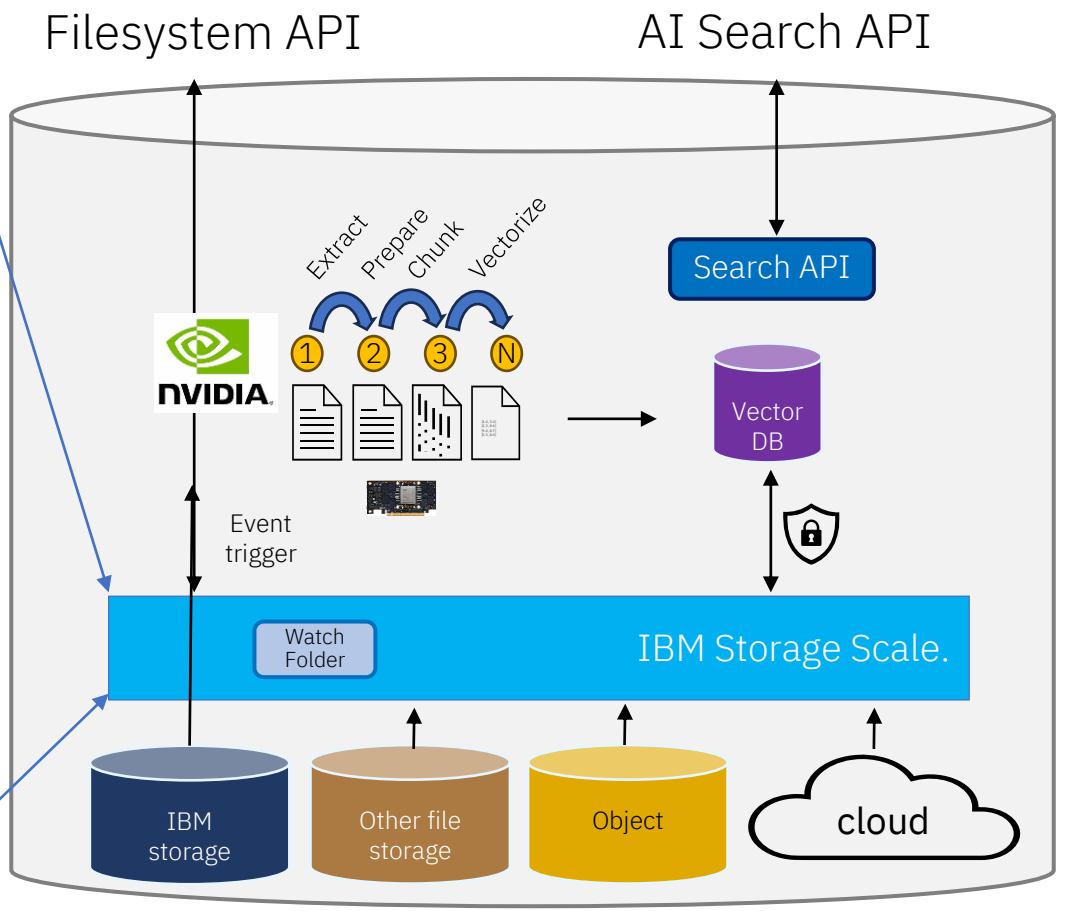

Monitoring and Health

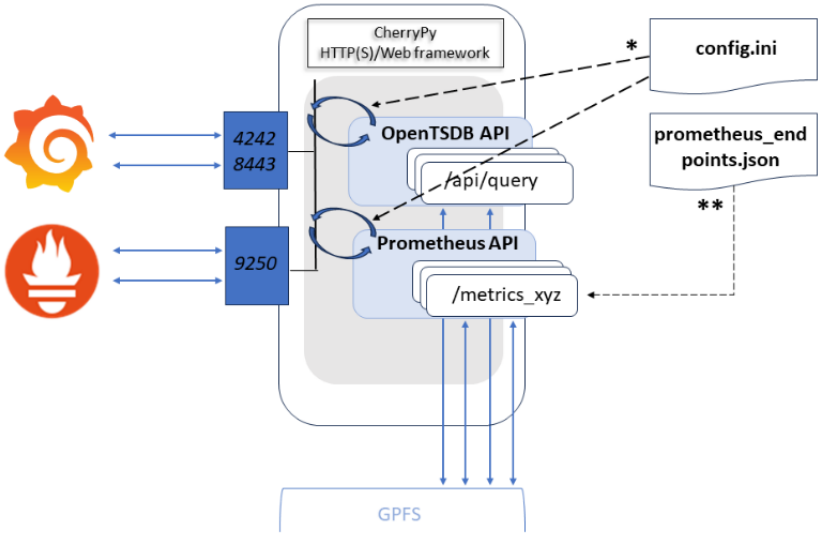

Storage Scale

Health monitoring changes

Added monitoring of **`mmauth`** certificate expiration.
Added swap usage monitoring.
Added SCALEMGMT monitoring through apihealth.

Performance monitoring changes

The following flags added to the **`mmperfmon`** command:

- **--interactive** | **-i** for colored and scrollable interactive mode, especially useful for inspecting multiple rows of output.
- **--statistics** to show statistics for the specified sensor.
- **--hour**, **--day**, **--week**, **--month**, **--group-by GROUP**, **--sort** for quality of life.

New sensor GPFSTSCOM: **`mmfsd`** socket statistics from **`mmpmon tscom`**

`mmptop/mmpstat:` Added **-N** option to support node classes and daemon or admin names.

Removed dependency on libboost-regex.

Support mmbackup for remoteFS with ILM

Storage Scale

RFE 307622

- By supporting backup on remotely mounted file system, mmbackup will use a backup network.
- Customer prefers to use dedicated networks to perform backups of data
- Added a new option: **--allow-backup-on-remote**
  - When specified, this parameter enables a backup on a remotely mounted file system. This option requires the IBM Storage Scale cluster level to be 5.1.8.0 or higher.
- This function relies on mmapplypolicy **–allow-scan-on-remote**
  - this option is not externalized yet, used internally (by **mmbackup, mmxcp**), but plan to externalize this in 6.0.0.
  - If this option is specified for the local file system, this option will be simply ignored.

expelnode mechanisms

Storage Scale

1. Explicit calls to mmexpelnode
2. Failure to respond to checks for disk leases
3. "RPC-related expels" happen when a cluster member detects that a partner node is not responsive to a comms check RPC (commMsgCheckMessages).
   - *any* member of a cluster (local or remote) can send a message to the cluster manager asking to expel its unresponsive partner.
   - The cluster to manager decides what to do! (complainer, complainee, nothing)
   - A non-cluster manager node does not EVER actually expel some other node.

Added a short "reason" in the mmfs.log entry that already gets made for RPC node expels

Storage Scale Deployment Toolkit

Storage Scale

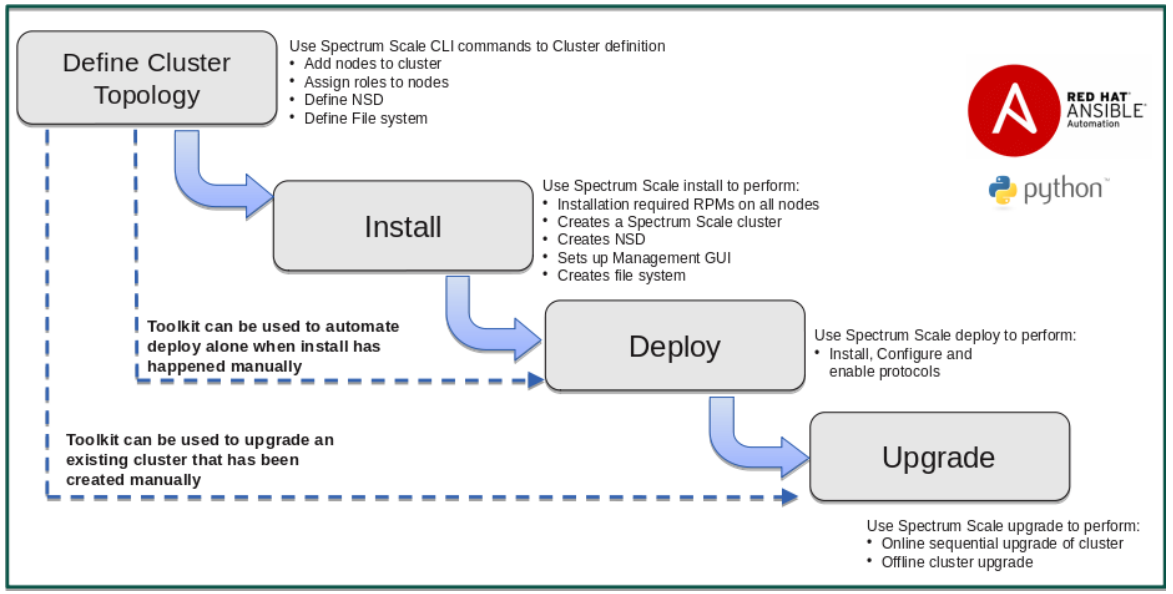

[CES S3] CES S3 Online upgrade support with noobaa changes

[CES S3] CES S3 offline upgrade support with noobaa changes

IBM Storage Scale native REST API Installation, Deployment and upgrade support.

Toolkit certification to work with the latest Ansible library.

[Python] Toolkit support with Python 3.12

[Currency] Extended OS currency

[ESS Protocol Node] Scale System (ESS) Protocol node certification with 5.2.3 Toolkit.

[ESS Protocol Node] CES S3 Protocol Deployment certification with 5.2.3 Toolkit.

[Open Source] Open Source Ansible Role certification with 5.2.3.

[CloudKit] Continued support and enhancements for toolkit integration.

[CloudKit] Ansible playbook to work with CES group for protocol deployment.

[Toolkit] [BDA]: Supporting HDFS Transparency 3.3.6 in Install Toolkit

[Toolkit]: Enabling FIPS environment certification support for Spectrum Scale deployments except GPFS GUI component.

Abstracting Cloud Service Deployment – Cloudkit!

Storage Scale

What is Storage Scale Cloudkit?

Create Storage Scale clusters on the cloud with

Bring Your Own License (BYOL) Model
Look in */usr/lpp/mmfs/VERSION/cloudkit*

Automates provisioning and deployment of Storage Scale on the cloud

Applies Storage Scale best practices for deploying on the cloud

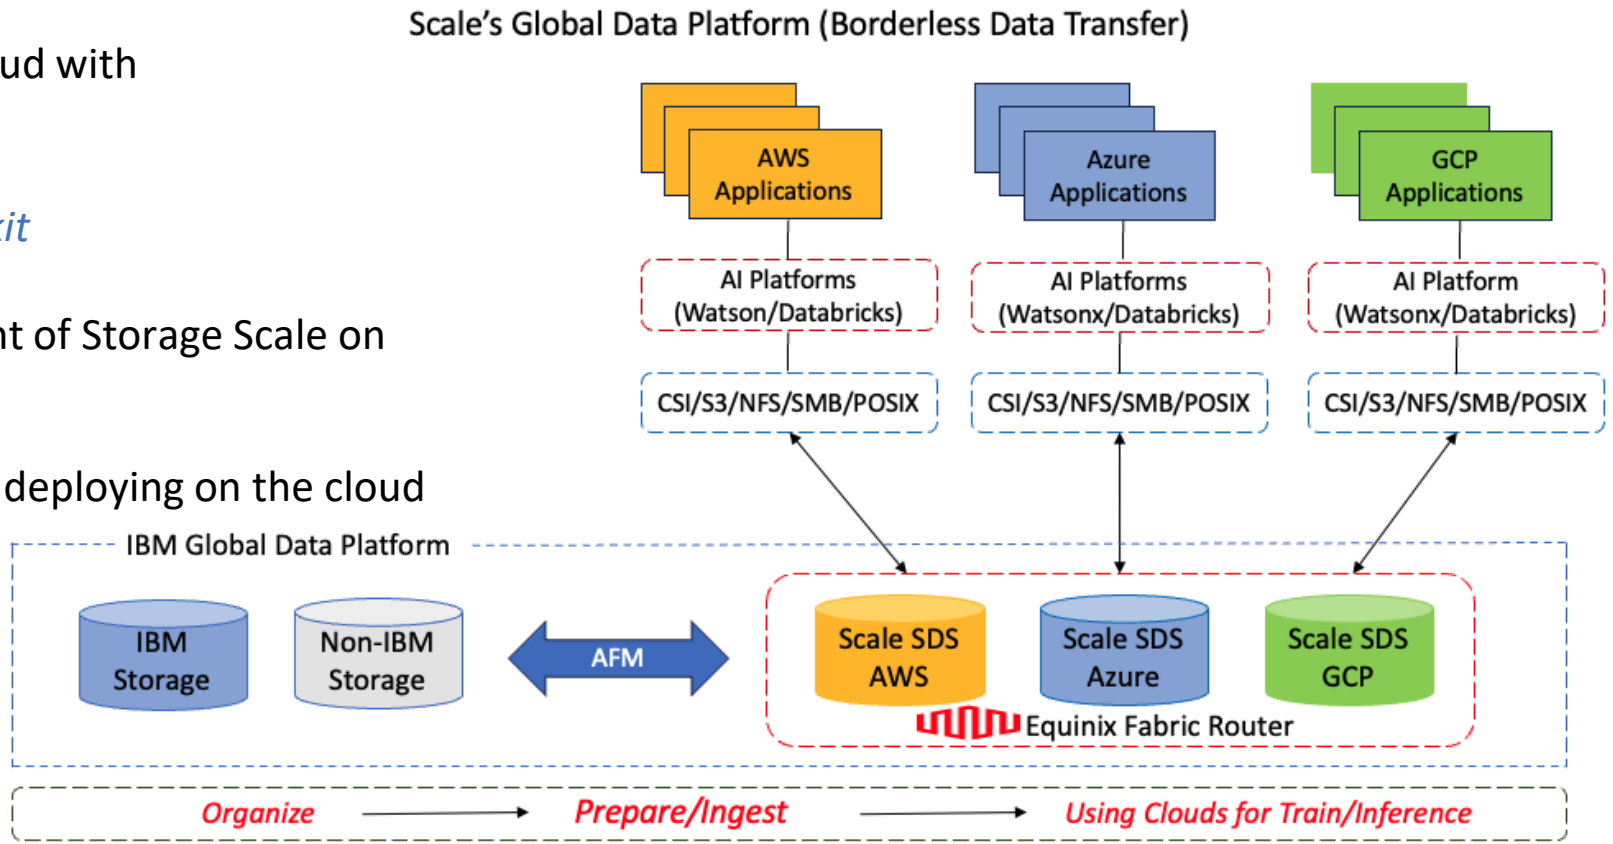

Updates

- On Microsoft Azure cloud, upgrade IBM Storage Scale clusters.
- On Microsoft Azure cloud, AFM to cloud object storage.
- On AWS, automated deployment and configuration of CES or protocols nodes.
- On GCP, the following disk types: pd-extreme, hyperdisk-extreme, hyperdisk-balanced.

Assurance Services

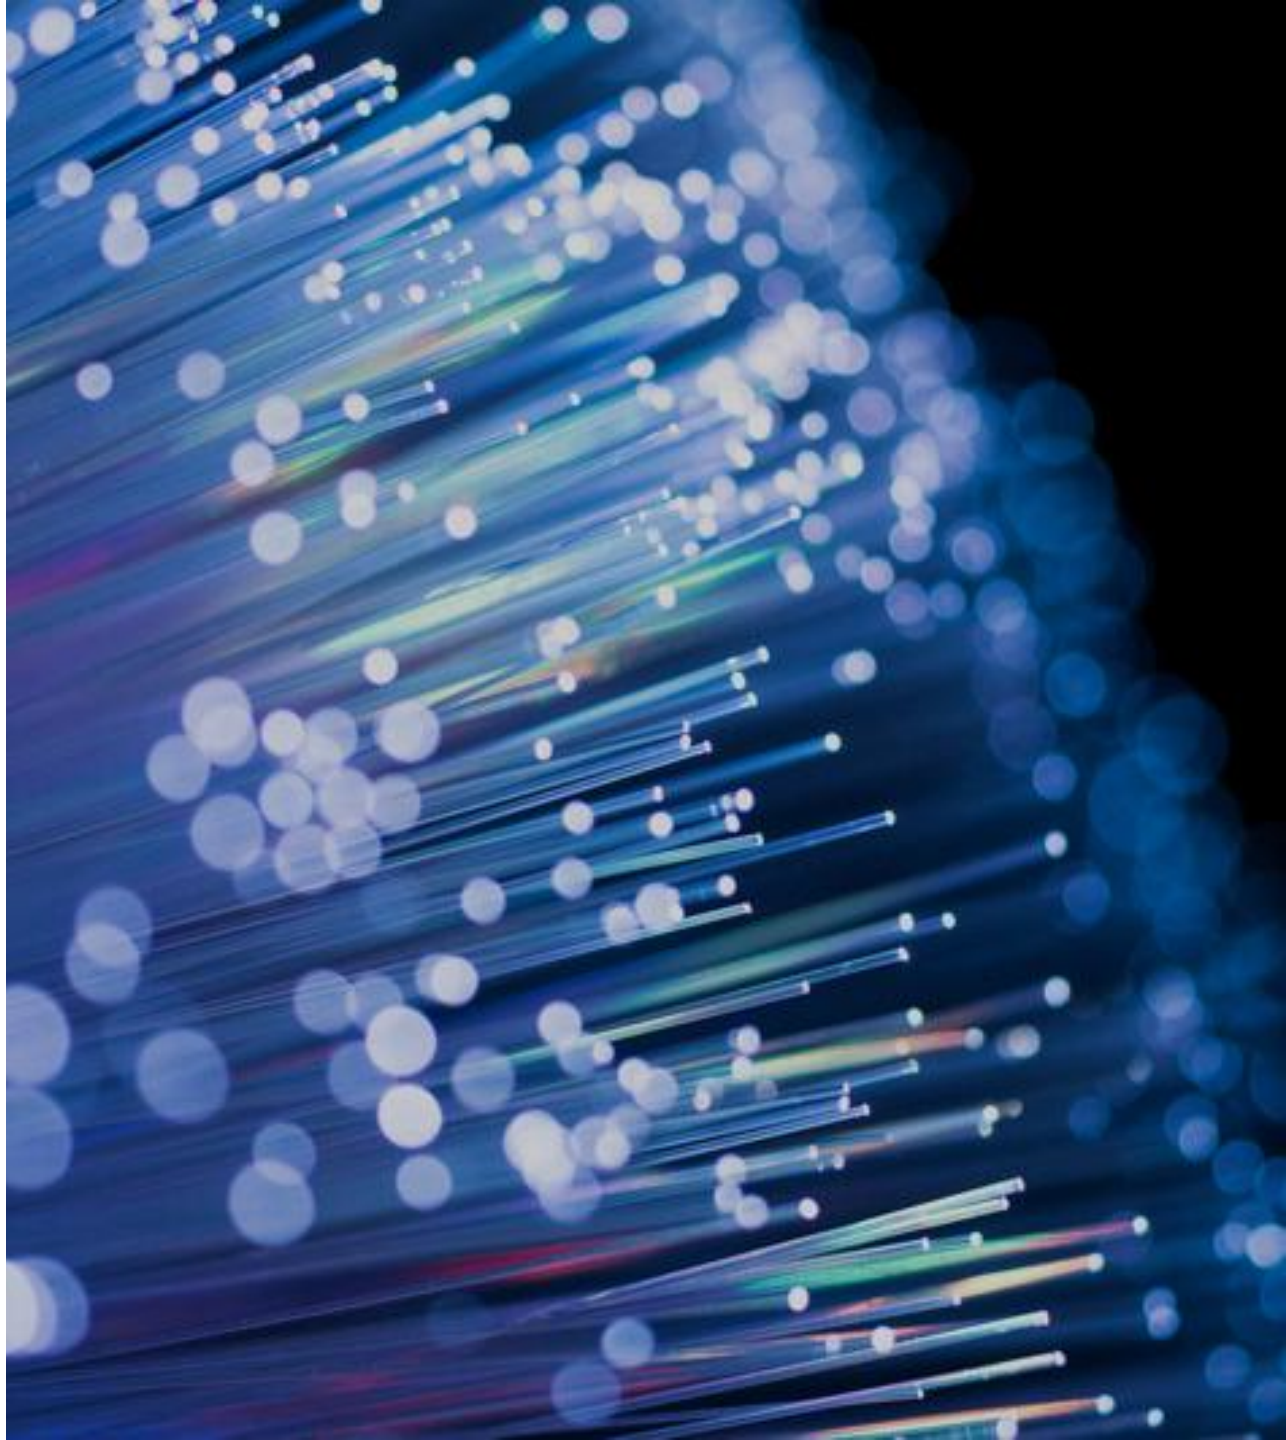

IBM Storage Scale native REST API

Security Improvements

Removal of SSH dependency

Removal of root requirement for control plane

Remote Administration

Fine-Grained Role Based Access Control
Declarative policy rules based on
Open Policy Agent

Control Plane Designed For Applications / Operators

Retain CLI for human management

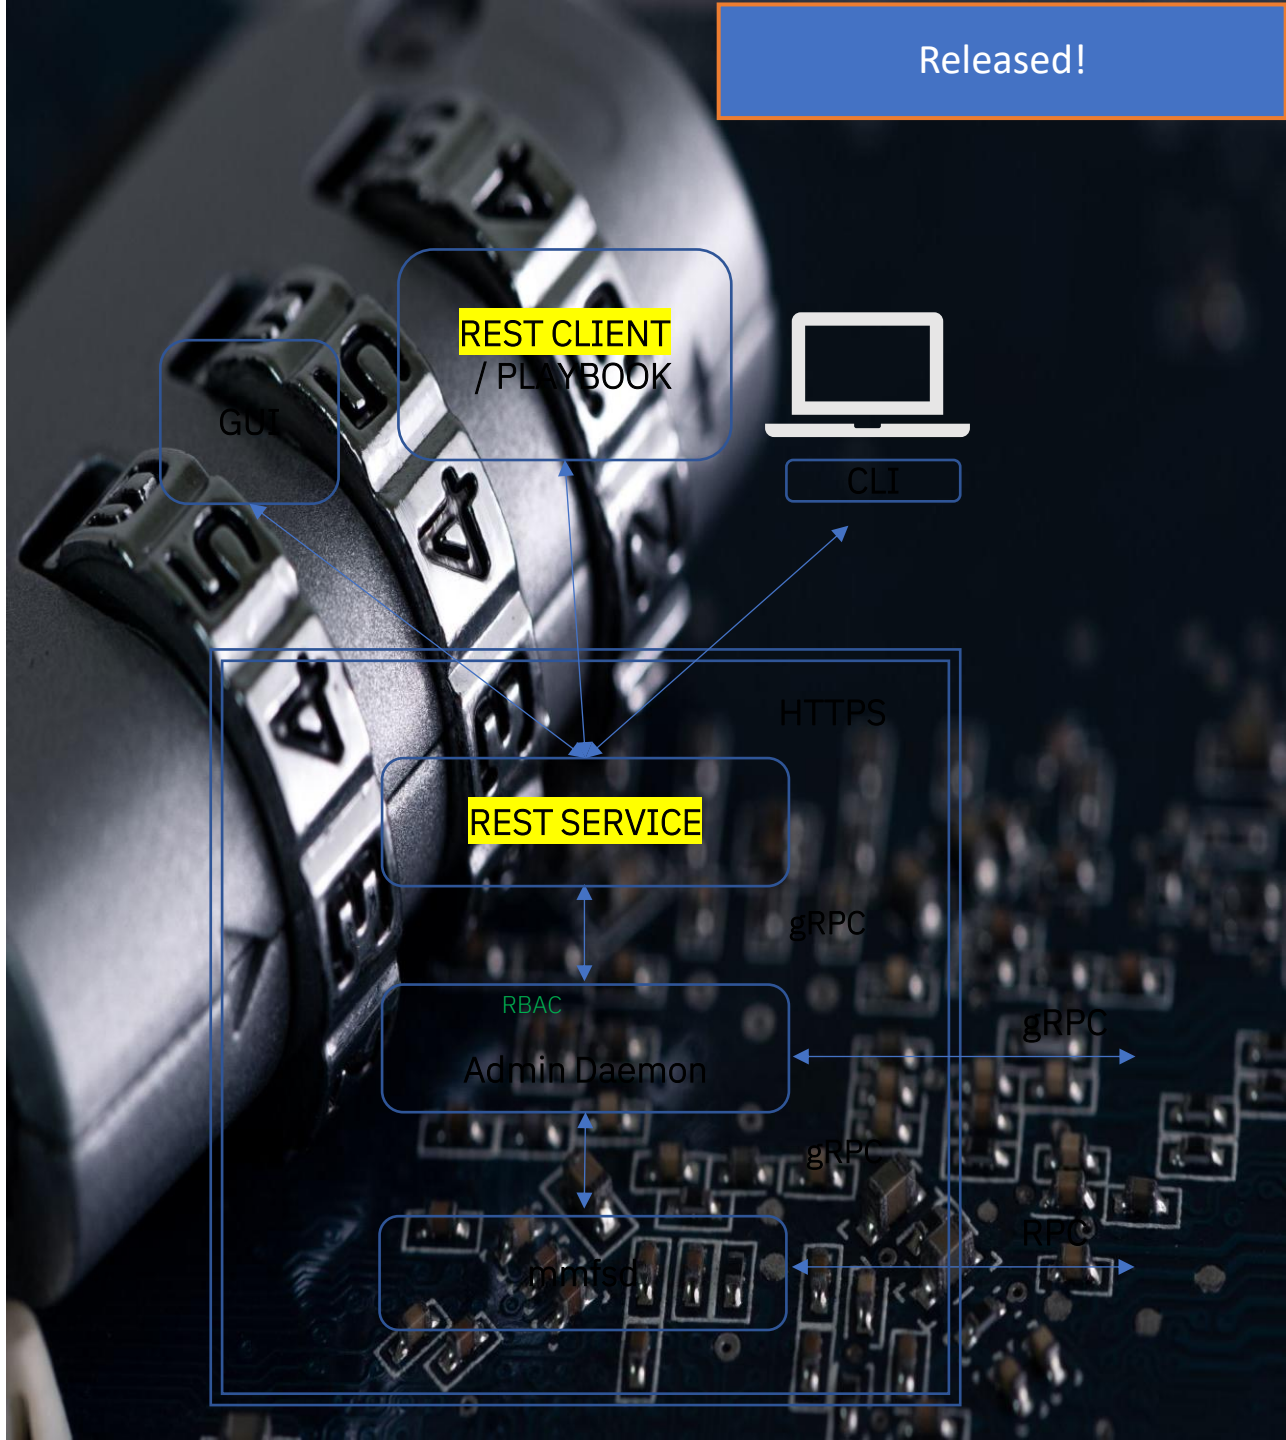

IBM Storage Scale native REST API

Storage Scale

Eliminates dependency on SSH, requirement for privileged users, and the need to issue commands locally on the node.

CLI and REST API of this feature are new interfaces to perform administrative operations. GUI is coming!

Architecture and OS

- x86_64 - RHEL 8/9, Ubuntu 22/24, SLES 15 SP6
- Power LE - RHEL 8/9, SLES 15 SP6
- Linux on Z - RHEL 8/9, SLES 15 SP6

ARM, Windows, AIX not supported.

- Do not intermix unsupported nodes

gpfs.scaleapi rpm/deb included

New user:group created scaleapiadmd:scaleapiadm as part of package install

New systemd service

scaleadmd.service

Migrate exsiting cluster: **scalectl cluster migrate**

Logs are located in: */var/mmfs/scaleadm/log/scaleadmd.log*

Installtoolkit supported

Ansible still requires ssh setup first

```
scalectl
scalectl apihealth
scalectl authorization
scalectl cluster
scalectl config
scalectl diagnostics
scalectl fileset
scalectl filesystem
scalectl node
scalectl nodeclass
scalectl operations
scalectl policy
scalectl troubleshooting
```

IBM Storage Scale native REST API

Storage Scale

Setting up a remote cluster –
https://ibmdocs-test.dcs.ibm.com/docs/en/IBM_Spectrum_Scale_test?topic=administering-remote-cluster-support-storage-scale-native-rest-api

Key exchange protocol - No need to explicitly provide public key file for remote cluster setup

Communication between clusters uses mutual TLS (mTLS) for secure authentication and data integrity.

Both clusters must trust each other's certificate authority.
Each node must have an X.509 certificate for communication across scaleadmd services.

Node identities are managed on a per node basis through scalectl node config command.

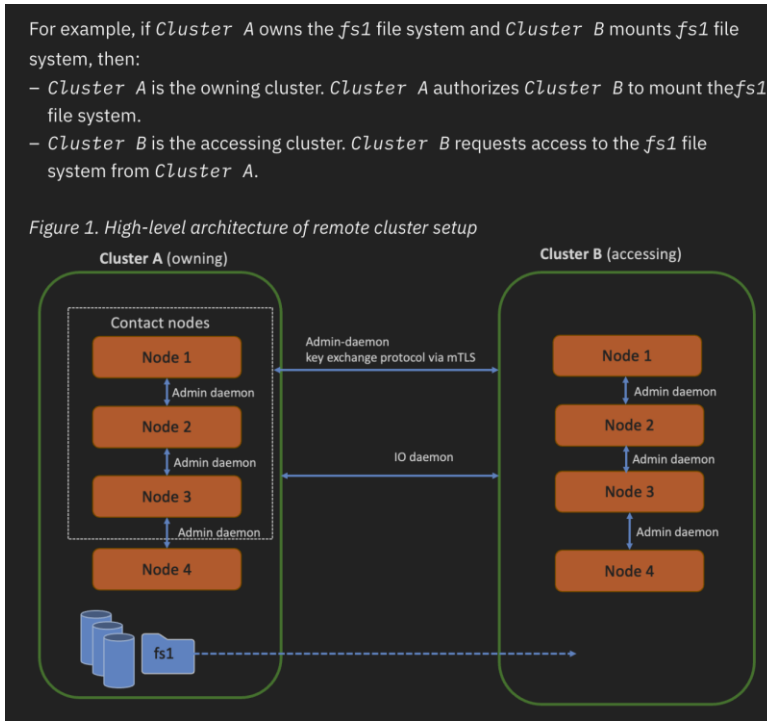


```
scalectl cluster create -N <NodeName>  --cluster-name "<ClusterName>" -A

scalectl cluster remote authorize --name <AccessingClusterName> \
--cluster-id <37939524262924283 82> --file resourcefile.json

scalectl cluster remote add --name <OwningingClusterName> --contact-nodes node1

scalectl filesystem mount <FilesystemName> -T <MountPoint>
```

Storage Scale

What’s new in

IBM Storage Scale System 6.2.*

Chris Maestas
IBM CTO, IBM Data and AI Storage Solutions
Chief Troublemaking Officer

IBM Storage Scale System

Integrated scale-out data management
for file and object

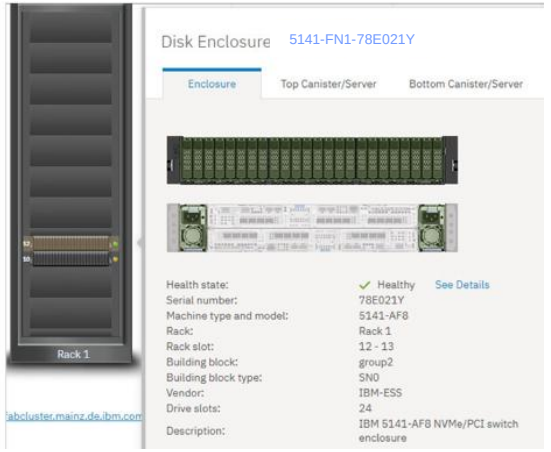

Simple GUI and wizards

Optimal building block for high-performance, scalable, reliable enterprise Storage Scale storage

- Faster data access with the choice to scale-up or scale-out
- Easy to deploy clusters with unified system GUI
- Simplified storage administration with IBM Storage Control integration

6000

3500

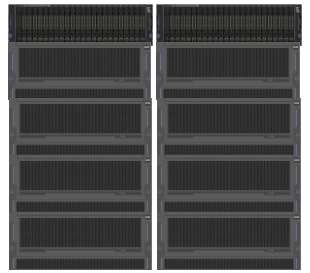

3500 or 6000
Hybrid/Capactity

One solution for all Storage Scale data needs

- Single repository of data with unified file and object support
- Anywhere access with multi-protocol support using protocol nodes: NFS 3/4.[0-2], SMB, Object, and HDFS
- Ideal for big data analytics including full Hadoop transparency

Ready for business-critical data

- Disaster recovery with synchronous or asynchronous replication
- Ensure reliability and fast rebuild times using Storage Scale RAID’s dispersed data and erasure code
- Six 9s (99.9999%) of availability and online scalability and upgrades

Scale System models are built for speed and capacity

Storage Scale

Storage Scale

Speed

Capacity

Hybrid

IBM Storage Scale System 6000

4U48 Enclosure
24 or 48 NVMe drives

Performance
Machine type 5149-F48

Up to 310+GB/s
IOR 100% read, InfiniBand

IBM Storage Scale System 3500

2U24 Enclosure
12 or 24 NVMe drives

Performance
Machine type 5141-FN2

Up to 125 GB/s
IOR 100% read, InfiniBand

IBM Storage Scale System 3500 Hybrid (NVMe + HDD)

Up to 8x 4U102 JBOD

3500 Hx
Machine type 5141-FN2

Up to 91 GB/s – NVMe
IOR 100% read InfiniBand
Up to 48 GB/s – HDD
IOR 100% read InfiniBand

IBM Storage Scale System 6000 Hybrid (NVMe + HDD)

Up to 9x 4U91 JBOD

6000 Hx
Machine type 5149-F48

Up to 280 GB/s – NVMe
IOR 100% read InfiniBand
Up to 100 GB/s – HDD
IOR 100% read InfiniBand

IBM Storage Scale System 3500 Capacity (HDD-Only)

Up to 8x 4U102 JBOD

3500 Cx
Machine types
5141-FN2

Up to 48 GB/s
IOR 100% read C4+ model, InfiniBand

IBM Storage Scale System 6000 Capacity (HDD-Only)

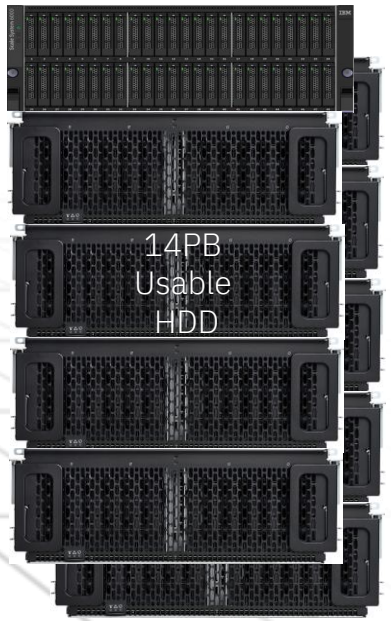

Up to 9x 4U91 JBOD

6000 Cx
Machine types
5141-FN2

Up to 90 GB/s
IOR 100% read C9 model, InfiniBand

SED Support with GKLM : Overview

GA in Dec 2022 for ESS 6.1.5.0 (ESS 3500)

Background:

- SED enabled by enrolling with MEK
- Auto lock on power off
- Data Security at Rest
- Need to unlock at Power ON using MEK
- Crypto erase by changing DEK

Challenges:

- External Key Managers are expensive
- Different Key Managers

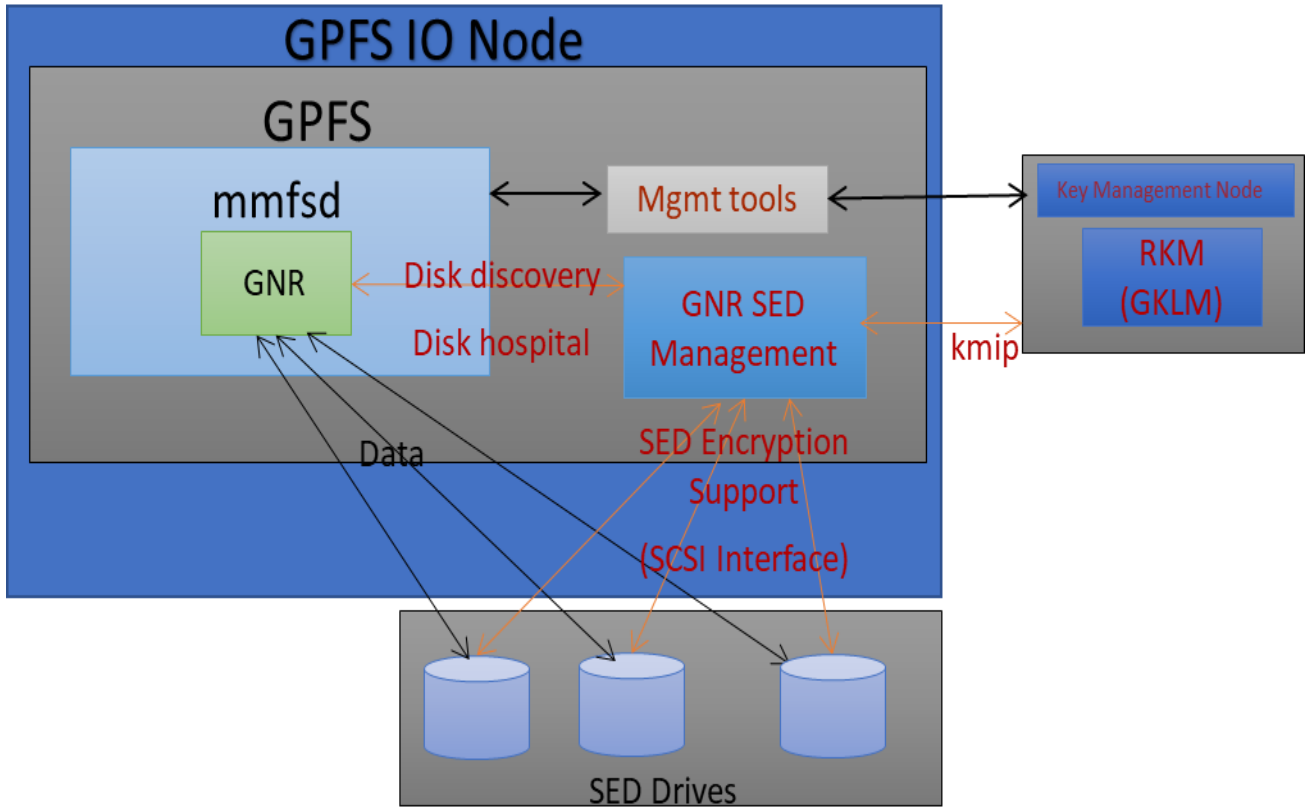

SED Support using TPM : Overview

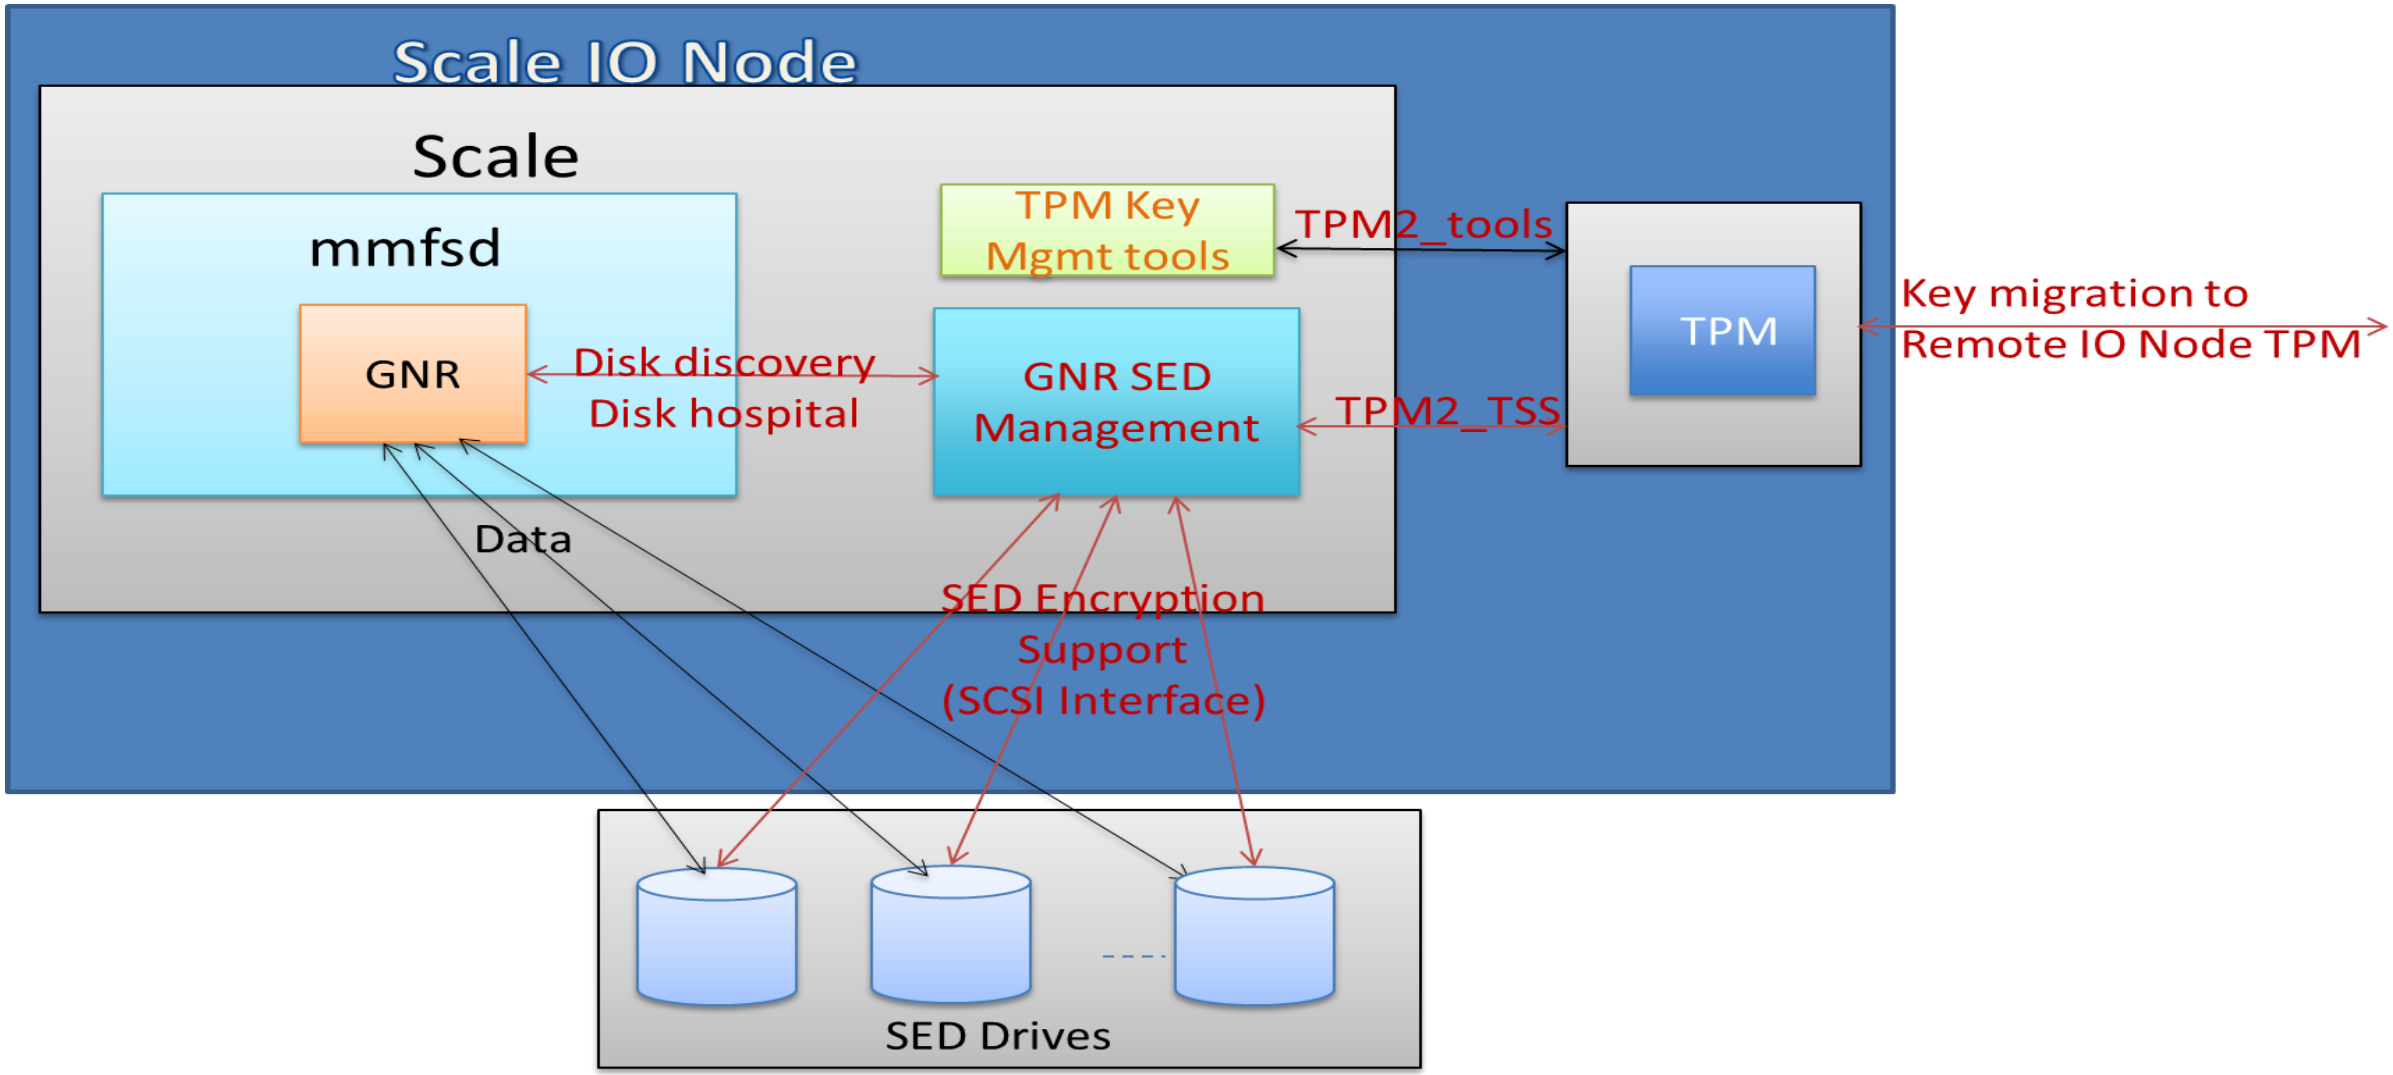

IBM Storage Scale 6000 with TLC, FCM4, and now QLC!

Storage Scale

https://www.ibm.com/docs/en/announcements/storage-scale-system-introduces-support-crypto-enabled-cx-7-adapter-second-source-hdd-drives

IBM Storage Scale System unveils next generation High-Capacity QLC-based NVMe SSDs for AI Applications

Published: 27 May 2025 AD25-0122 Hardware

Overview

IBM Storage Scale System announces availability of the next generation high-capacity quad-level cell (QLC) based NVMe SSDs optimized for AI workloads, with support for 30 TB and 60 TB drive options. The QLC flash tier is combined with TLC flash to achieve high endurance and high performance, enabling up to 2.2 PB of more cost-effective flash capacity within the Storage Scale System.

Planned availability date

13 June 2025

Description

The Storage Scale System QLC SSD is a 2.5" dual-ported U.2/U.3 PCIe Gen5 drive that offers incredible density, and performance and reliability, especially optimized for read-intensive workloads. The QLC also supports SED (Self-Encrypting Drive) providing enhanced data protection and security. SED leverages hardware-based encryption to transparently encrypt and decrypt the data written and read from the drives. IBM Storage Scale system 6000 offers both 30.72 TB and 61.44 TB drive options, delivering more than 2.2 PB of high-performance flash storage within a single building block.

Exactly 24 QLC drives of the same capacity (30.72 TB or 61.44 TB) are supported within a Storage Scale System 6000.

QLC support in the Storage Scale System 6000 requires 24x TLC NVMe drives of the same capacity (3.84 TB, 7.68 TB, 15.36 TB or 30.72 TB).

The ConnectX-7 400GbE/NDR Single-Port OSFP Adapter (NVIDIA PN MCX75310AAC-NEAT) is a crypto-enabled half-height, half length x16 PCIe Gen5 high-speed adapter, supporting both ethernet and Infiniband 400GbE speeds in a single port.

Grafana pdisks dashboard (example)

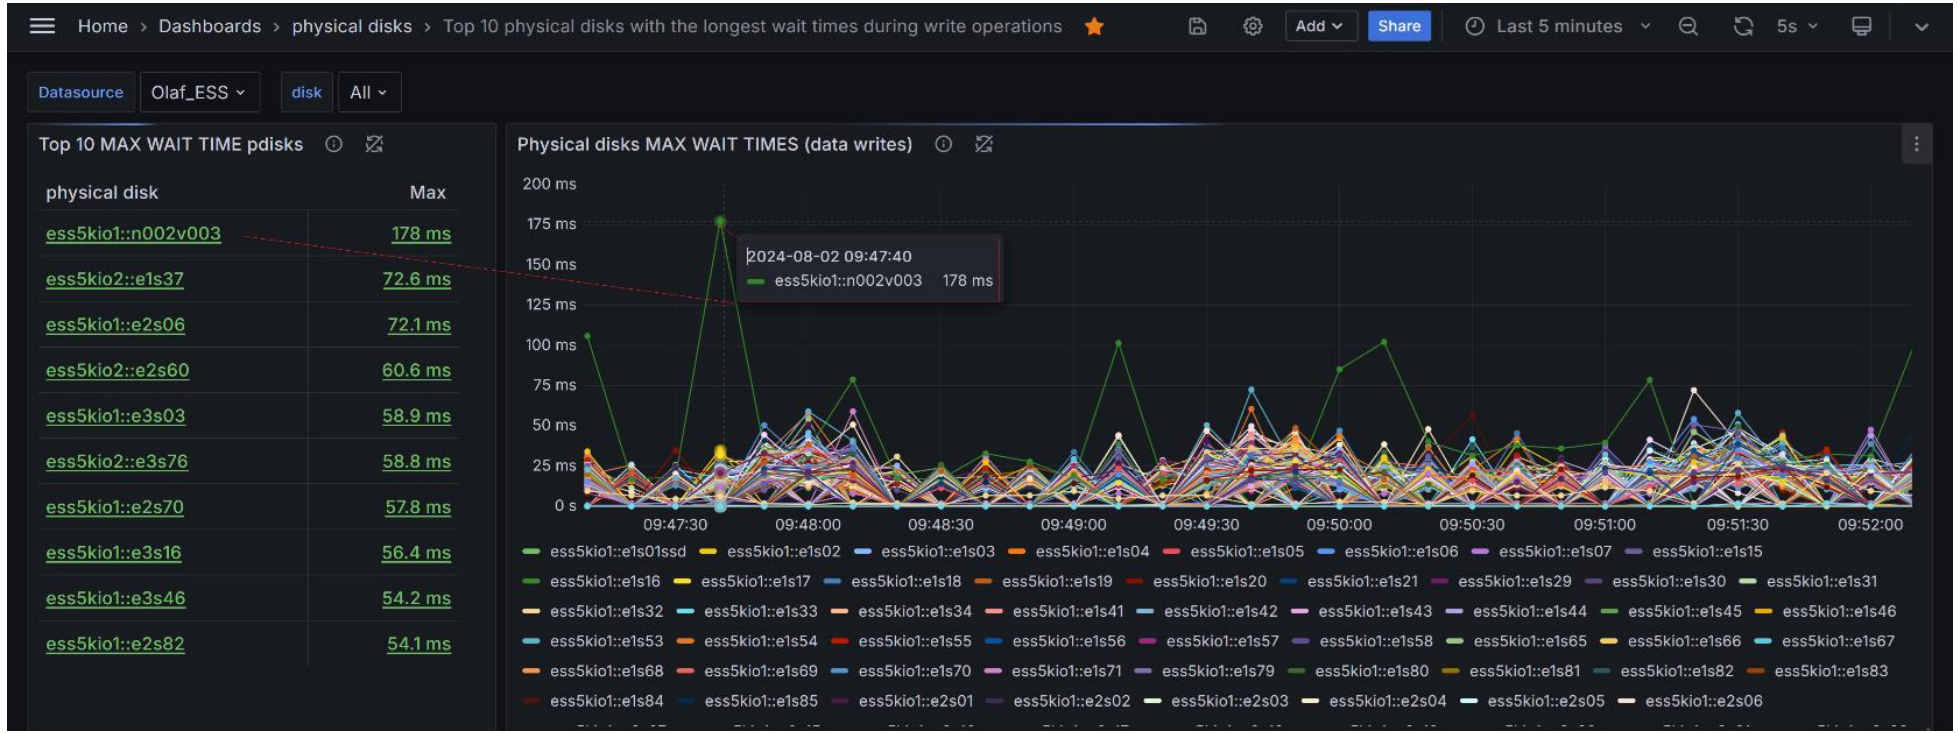

One of the most important measures of physical disk performance is the wait time for a disk write operation.

The **new** bundle of sample dashboards physical disks allows you to identify ***the top 10 physical disks with the longest wait time*** for a write operation for the selected time period.

For more details on a particular disk, you can ***drill down*** from the table to the individual disk view.

Watch DEMO video *on the IBM Storage Scale bridge for Grafana Wiki >>>*

New Scale System Software Features

IBM Storage & SDI
Storage Scale

- Scale 5.2.3.X
- RH9.4 3500/6000 RH8.10 5000/3000/3200
- DOCA 25.01
- Deploy from the GUI
- FCM space monitoring
- CX7 400G VPI card
- Firmware updates FCM4/4u102/4u91/6000
- MES for Falcon/HBA (1-9 enclosures) ESS 6000 4u91
- QLC support with TLC tier

IBM Storage Scale System 6000

mixed QLC Tier

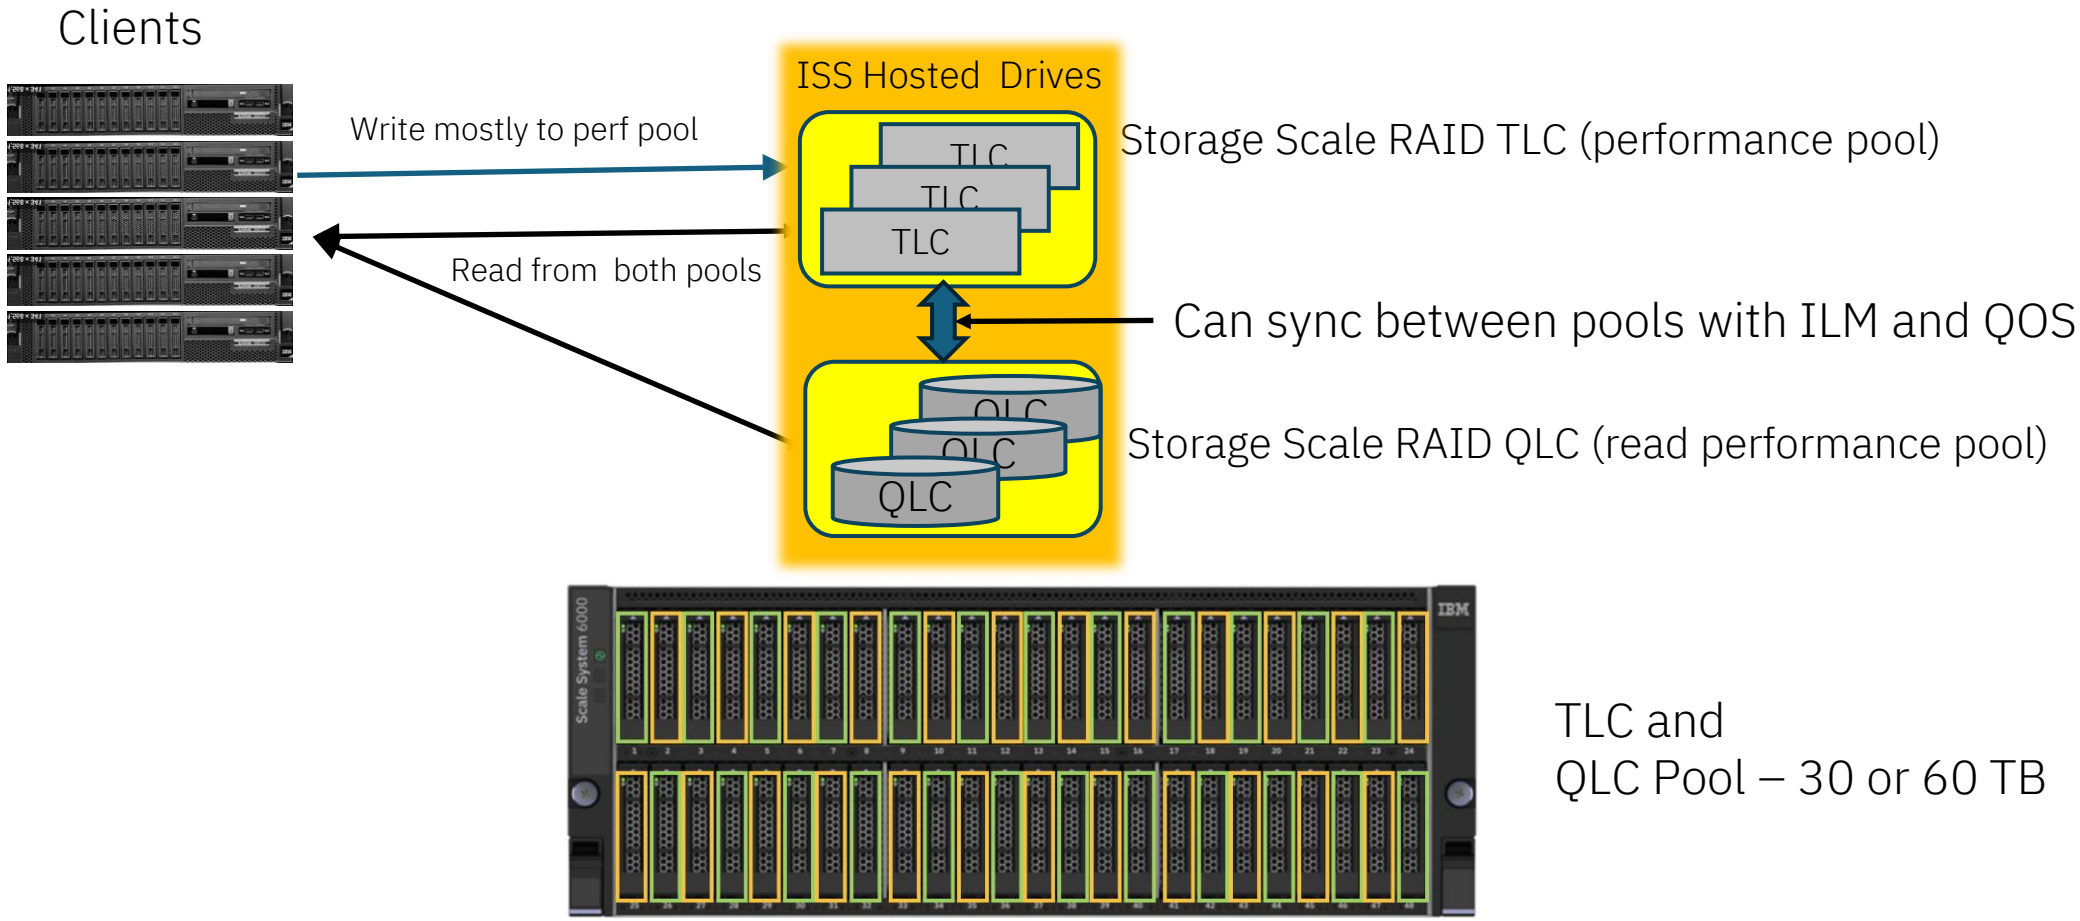

- Recommend data written to TLC pool first – use placement policies!
- Provides very high read to either pool (QLC read performance is like TLC)

IBM FlashCore™ Module 4

Capacity and Performance

2.5” dual ported U.2 NVMe Gen 4 PCIe
Industry leading density at 38.4 TB per drive
Inline hardware FIPS 140-3 encryption
Inline hardware 3:1 compression = 116 TB!

Internally tiered storage
-> MRAM -> SLC -> 3D QLC

Industry leading QLC endurance
15K Program/Erase cycles
Compared to 1500 for enterprise QLC

IBM Unique QLC management (100+ patents)
read calibration, heat binning, he[illegible]binning,
error correcting codes, optimized [illegible]tage

Continuous health [illegible]nitoring
keeps wear across all cells within 5%

FCM Space Management Best Practices in IBM Storage System 6000

Does your data avoid encryption or compression?

On average data if we achieve 1.2-1.3 it’s about 45-50 TB per drive

Estimate your data via gzip/zip/lz4 or scale software compression and get the best idea of actual compression.

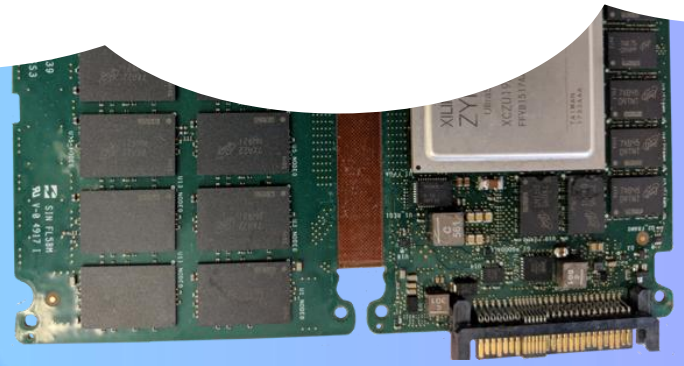

Thank you for using

Applied Data Systems
725 followers
1mo •

Applied Data Systems is excited to offer a fully integrated AI infrastructure solution that combines NVIDIA HGX BasePOD (H200 and B200 servers), DDC water-cooled cabinets, and the IBM Storage Scale System 6000 to address the most demanding AI cooling and data challenges.

Storage Scale

Now more than 1PB Usable Flash per 4U

Using TLC and QLC

IBM

Storage Scale System

*Performance based on tuned network for RDMA with TLC drives