

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


IBM Storage Scale Days 2025 DE

March 19th – 20th, 2025 | Heidelberg, Germany

Olaf Weiser / Norbert Schuld / (Chris Maestas)

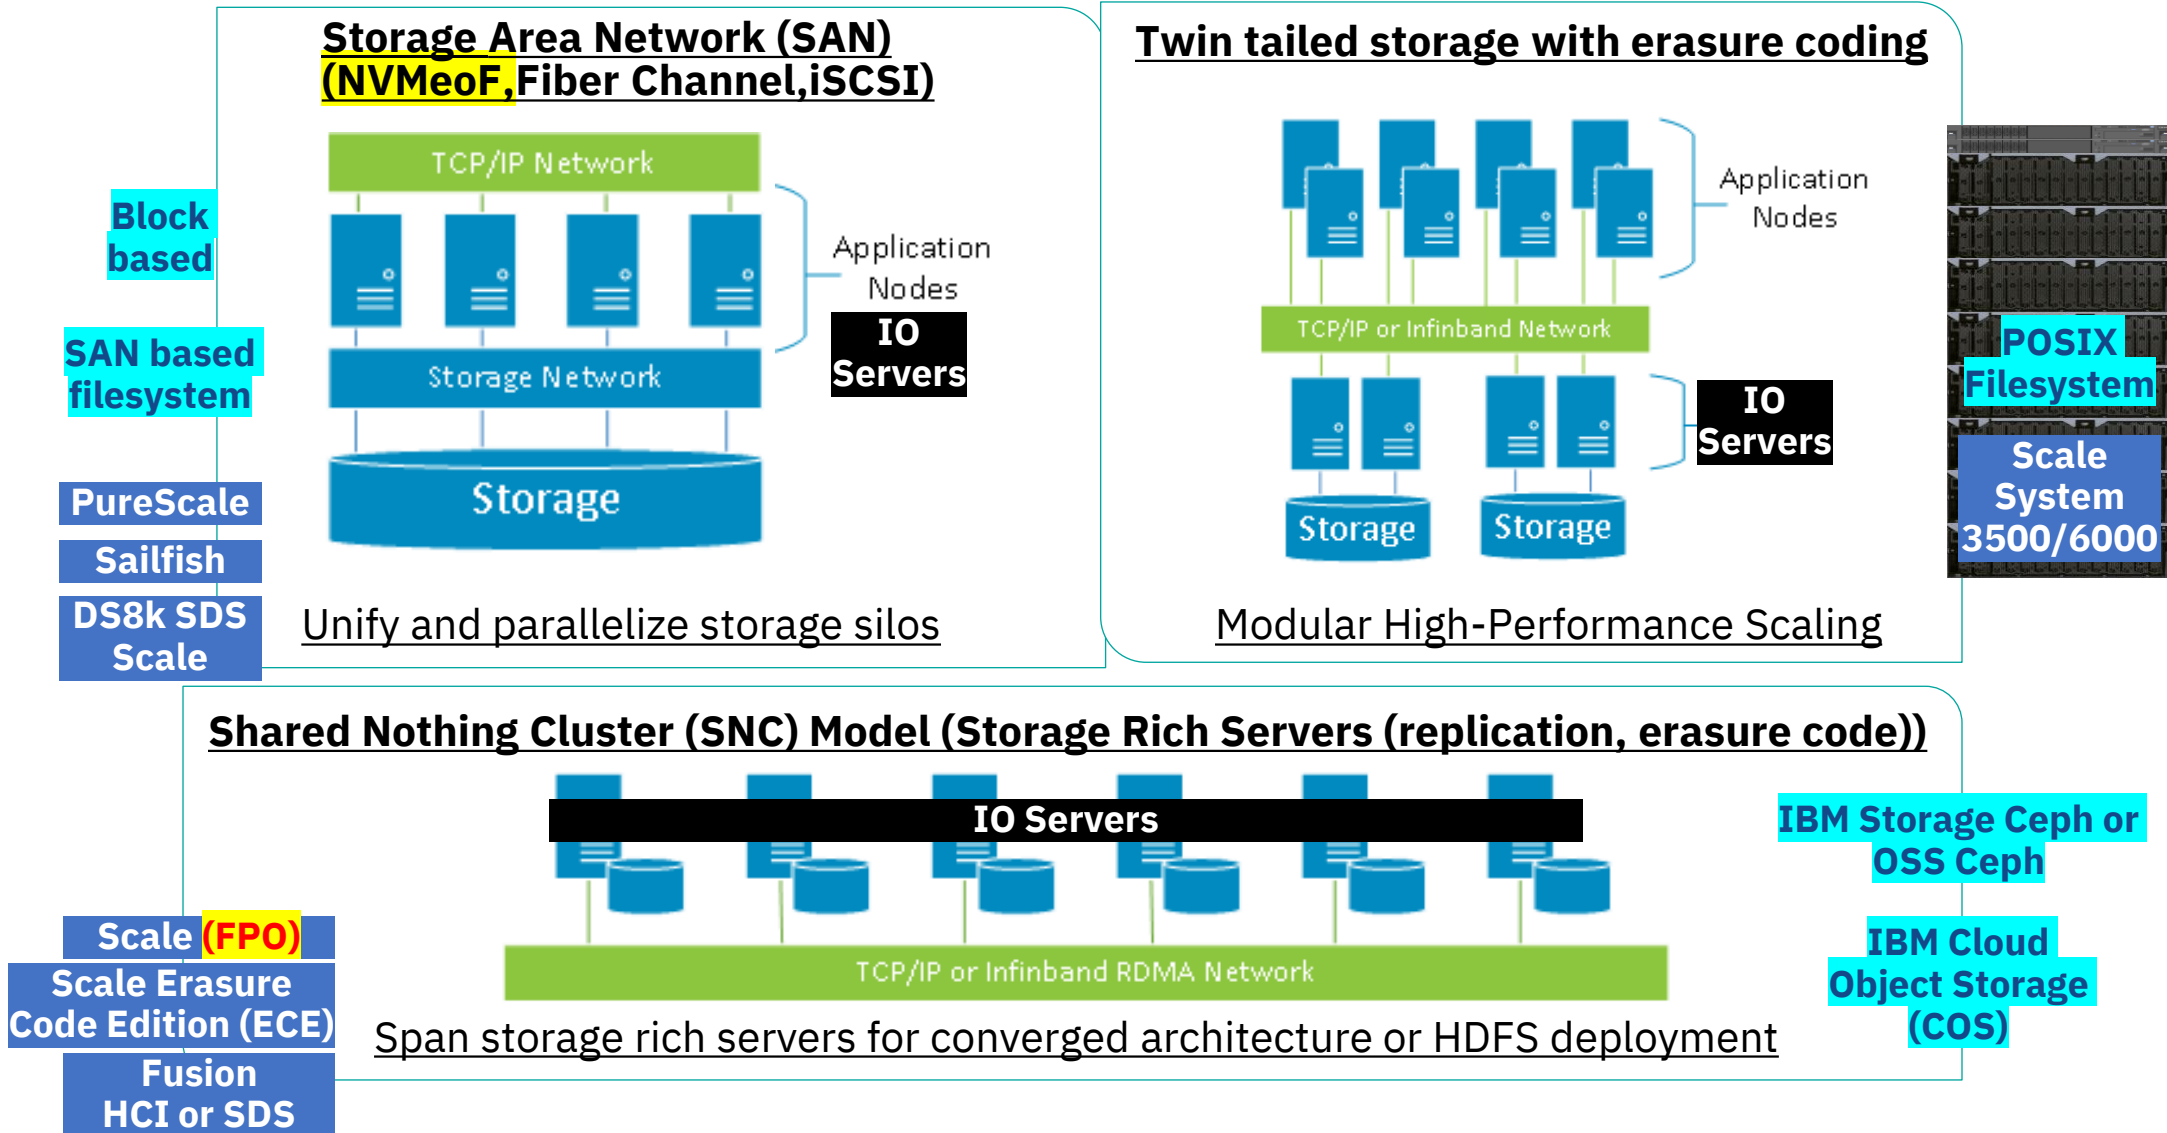


Disclaimer



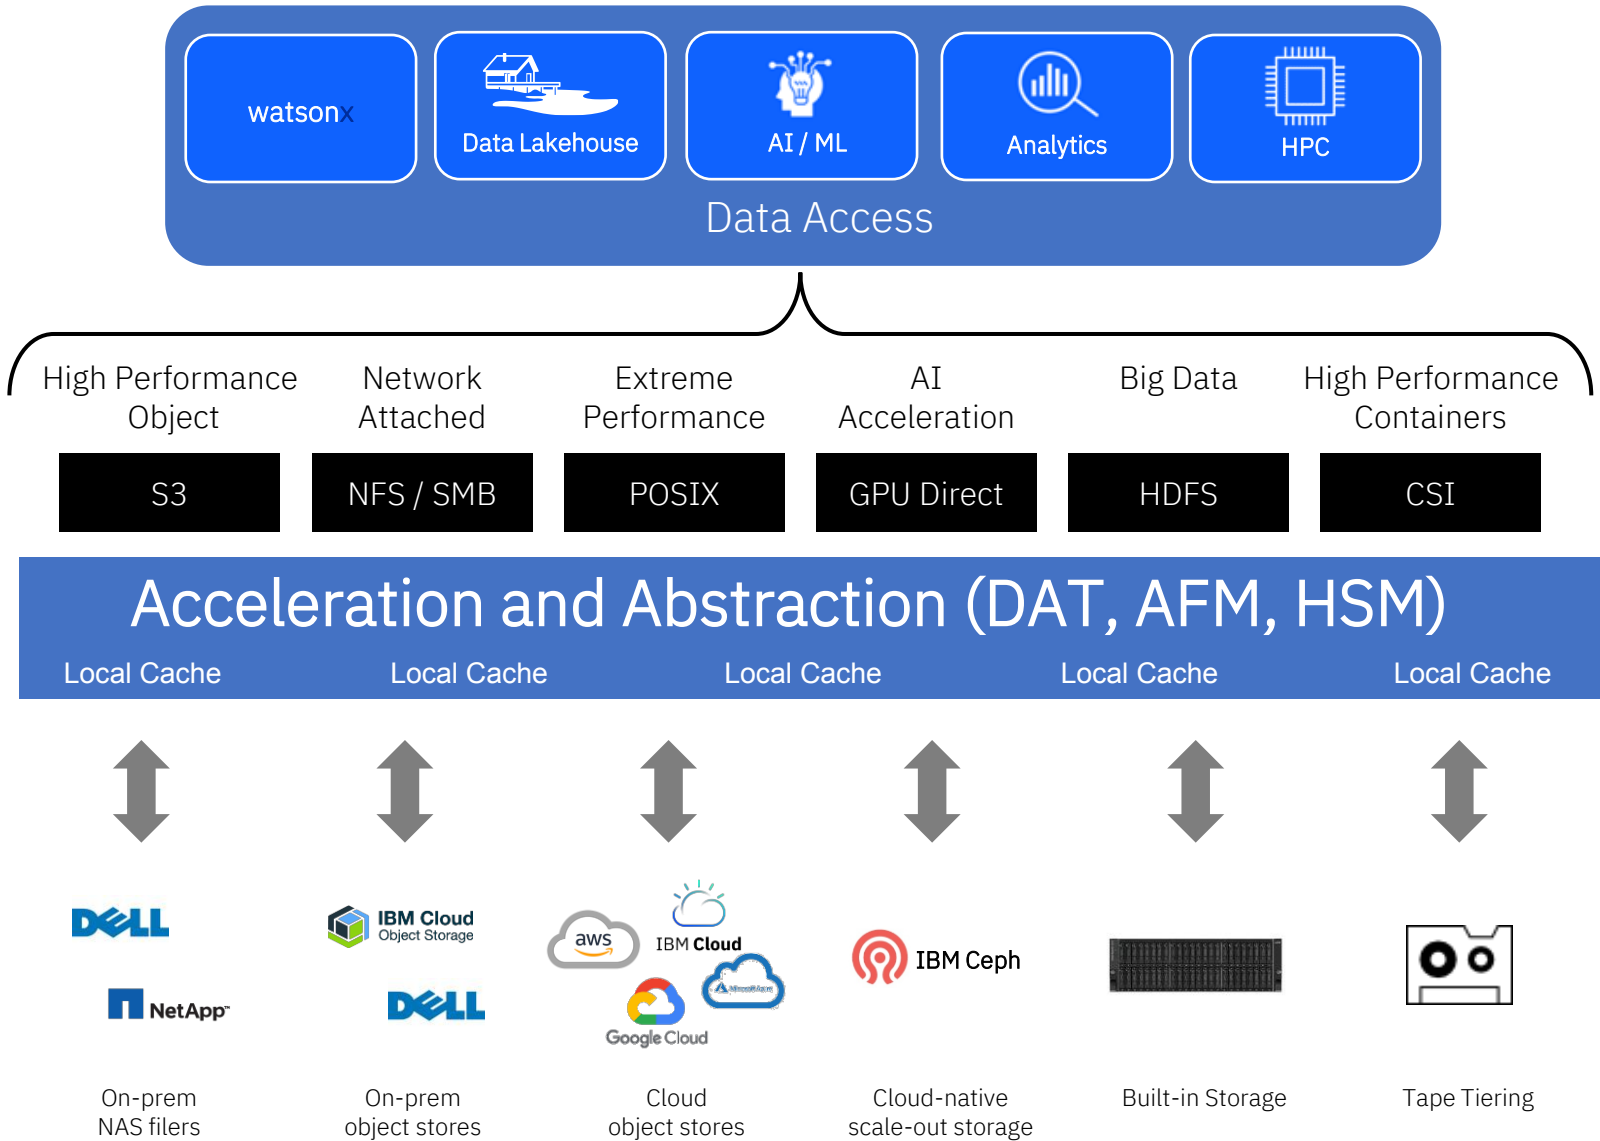

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.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

Scale Deployment model comparison



IBM Storage Scale a Global Data Platform

Global Data {Access, Acceleration, Abstraction and Assurance}



Multi-Protocol Data Access
Simultaneous multi-protocol access including GPU Direct 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 Assurance
Data security from source to destination with governance and

Outcome: Data accountability and integrity ensuring business continuity under any circumstance

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

JUPITER + IBM

IBM

A new class of
supercomputers
for AI-driven
scientific
breakthroughs

Extreme-scale computing for
AI powered by the NVIDIA
Grace Hopper™ and IBM
Storage Scale System

Hosted at the Forschungszentrum Jülich facility in Germany, JUPITER, the world's most powerful AI supercomputer, is being built in collaboration with NVIDIA, ParTec, Eviden and SiPearl to accelerate the creation of foundational AI models in climate and weather research, material science, drug discovery, industrial engineering and quantum computing.



6000 compute nodes to achieve the compute performance of 1 ExaFLOP/s, featuring NVIDIA Grace-Hopper processors. To keep the data flowing, a 21 Petabyte Flash Module (ExaFLASH) is provided based on the IBM *Storage Scale* software and a corresponding storage appliance based on IBM ESS 3500 building blocks.

More information @ <https://www.fz-juelich.de/en/ias/jsc/jupiter/tech>

IBM Storage Recent News



IBM Storage Scale System 6000 Now a Certified NVIDIA Cloud Partner



<https://community.ibm.com/community/user/storage/blogs/mike-kieran/2025/01/10/ibm-storage-scale-system-6000-now-a-certified-nvid>

IBM Storage Scale System 6000 is now a certified NVIDIA Cloud Partner (NCP) for HGX H100/H200/B200 systems. As a certified high performance storage partner for NCP, IBM Storage Scale System 6000 has demonstrated that it can deliver scalable high-performance IO to the most demanding AI training and inferencing workloads deployed on NVIDIA HGX GPUs in the cloud.



“The supercomputer will leverage **IBM Storage Scale System 6000** technology to deliver high-performance storage for AI, data analytics, and other demanding workloads.

As part of this agreement, CoreWeave customers can access the IBM Storage platform within CoreWeave’s dedicated environments and AI cloud platform.”

CoreWeave Partners with IBM to Deliver New AI Supercomputer for IBM Granite Models



NEWS PROVIDED BY
CoreWeave →
Jan 15, 2025, 08:00 ET

<https://www.prnewswire.com/news-releases/coreweave-partners-with-ibm-to-deliver-new-ai-supercomputer-for-ibm-granite-models-302351465.html>

- One of the first deployments of NVIDIA GB200 NVL72 at supercomputing scale
- Supercomputer will leverage IBM Storage Scale System to power AI research and development



<https://www.ibm.com/think/news/deepseek-r1-ai>

NVIDIA GTC presentation for Content Aware Storage (CAST)!

In-Person

Talks & Panels

Enable Intelligent Storage to Process Data for AI Applications [S71937]



Vincent Hsu, VP, IBM Fellow, CTO for IBM Storage, IBM

Rob Davis, VP Storage Technology, NVIDIA

The common implementation of AI pipelines today is to bring data to AI. This works well when your dataset is relatively small and co-located. When we look at the next step of AI journey, we know one thing for sure: there will be a lot more data in a lot more locations. The effective way to address this challenge is to push AI processing closer to where the data is. This concept is "AI Content-Aware Storage (AI CAST)." The vision of content-aware storage is to enable intelligent storage to process data for AI applications. We'll demonstrate the architecture of AI CAST by leveraging NVIDIA Blueprints and NIMs to accelerate the retrieval-augmented generation (RAG) pipeline by incorporating storage and storage metadata in the Continuous Data Ingest and vector DB management.

Suggested Audience Level: Technical, All

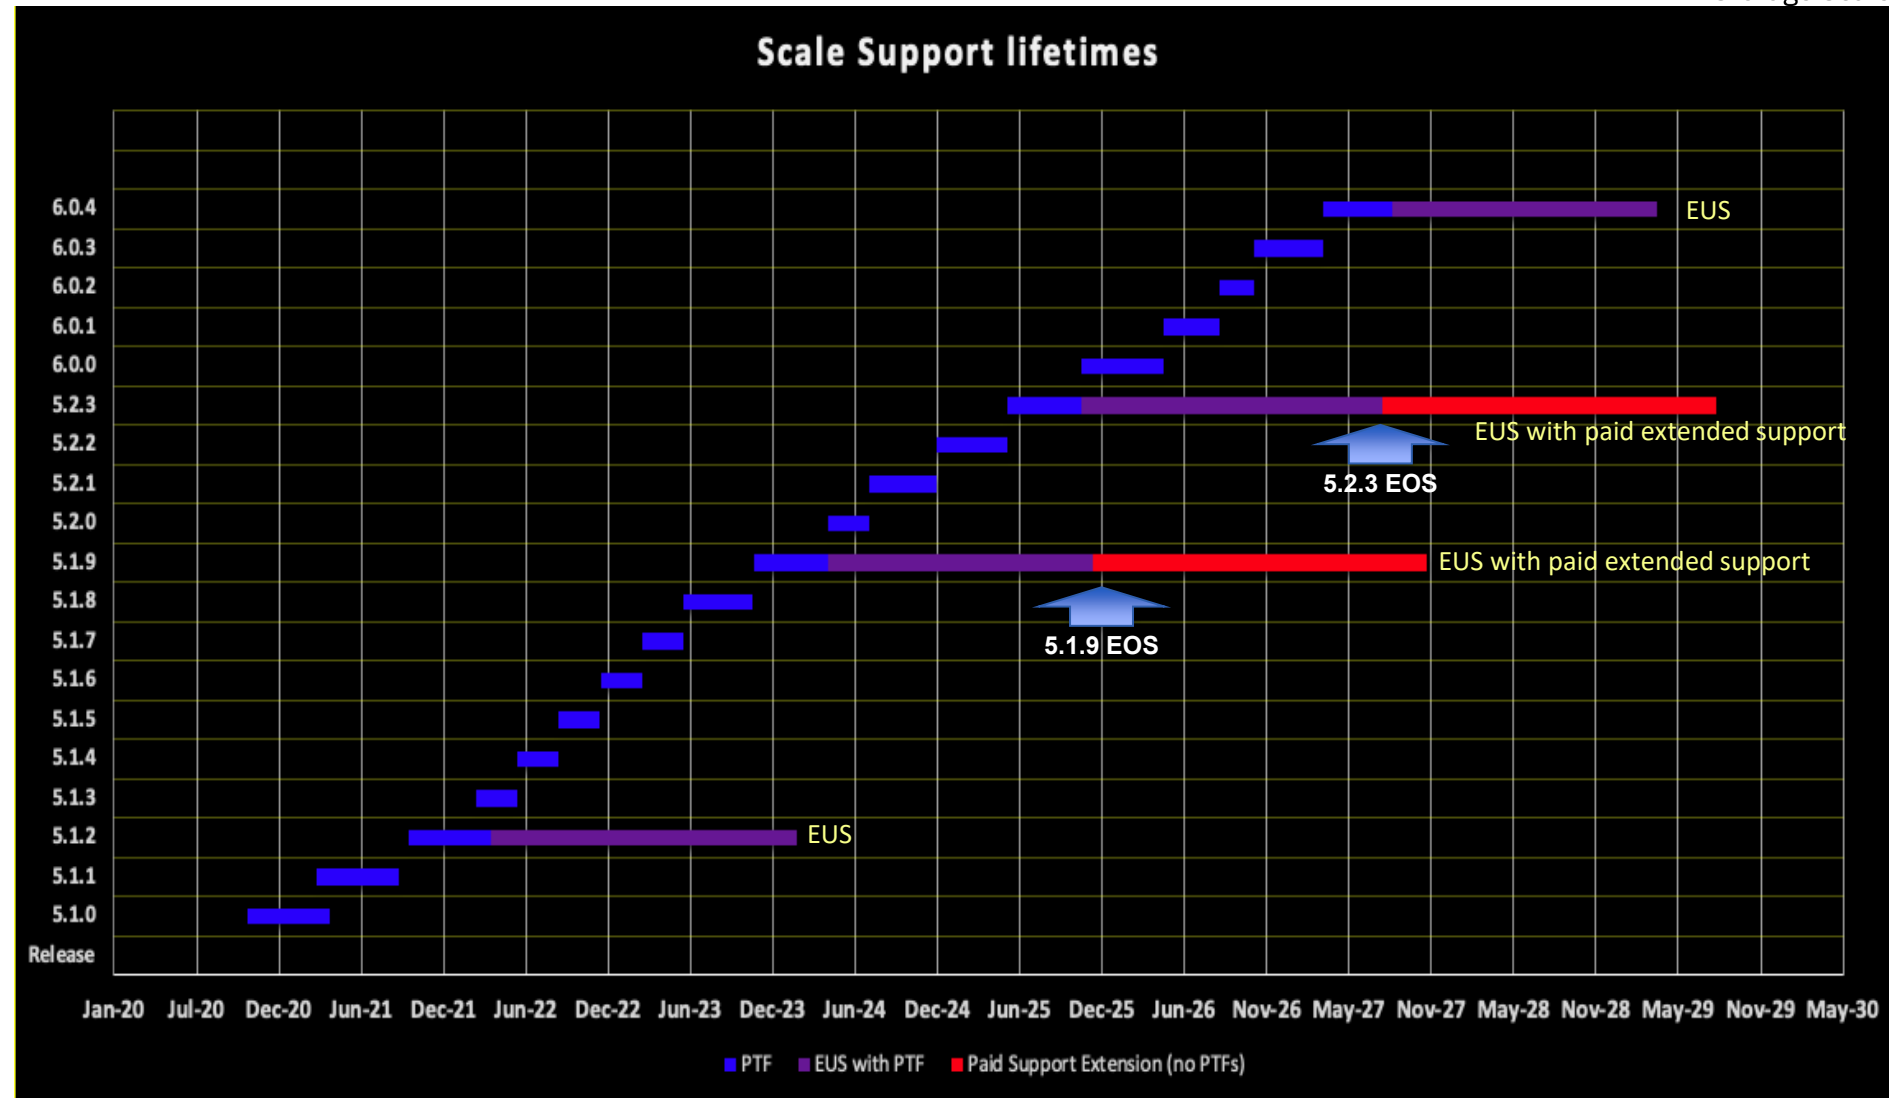
Add to Schedule ⌚

Monday, Mar 17 | 1:00 PM - 1:40 PM PDT

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.

Access Services –NFS, SMB, HDFS

Support and Currency:

- Samba 4.20 release
 - Improved Share ACL management performance
- NFS-Ganesha support for 5.7 code base
 - Improved health check
 - **Improved stability**
 - **GA support of NFS v4.2**


Improved performance:

- NFS “meta data cache” component was revised resulting in significant performance improvements

<https://community.ibm.com/community/user/storage/blogs/mara-miranda-bautista/2024/05/20/ibm-storage-scale-ces-nfs-520-performance-eval>

- HDFS transparency metadata redesign
 - Full parallelism for RPC calls (GPFSNamesystem)
 - No more lock contention in NameNode
- **Continued partnership with Tuxera for high-performance SMB**

After



1. Request arrives at the NameNode
2. NameNode collects metadata from Scale
3. NameNode serves request directly from Scale

Access Services – High Performance Object



Support and Currency:

- Swift is being Discontinued
- You can use 5.1.8 Swift code in CES of 5.1.9
- [New CES S3 is here!](#)
- <https://www.ibm.com/support/pages/node/7145681>



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

ILM support including Tiering to Tape support via RPQ

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

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 (Tech preview) 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>

Scaling limits for S3:

- Up to 10TB single object size
- Up to 5000 S3 accounts
- Up to 5000 S3 buckets
- Up to 100M objects per bucket (tested limit)
- Up to 3K client connections per CES node

Higher
scaling limits
as compared
to HPO 1.0 !

Deployment Requirements:

Storage Scale Cluster:	Storage Scale 5.2.1
Operating System:	RHEL8.x or RHEL9.x
Architecture:	x86_64, Power(ppc64le), Z(s390x)
Storage Scale CES Cluster Size:	Up to 10-node CES cluster (tested limit)

*No support for upgrade from CES S3 Tech Preview to CES S3 MVP GA

NO
Openshift
cluster
required !

Access Services – Container Native Storage Access (CNSA)



Improvements introduced in CNSA 5.2.x

Wider support to use the latest CNSA functionality.

Support for Red Hat OpenShift 4.16, 4.17, 4.18

Vanilla Kubernetes support

Support for parallel core pod upgrade

Avoid node reboots during upgrade

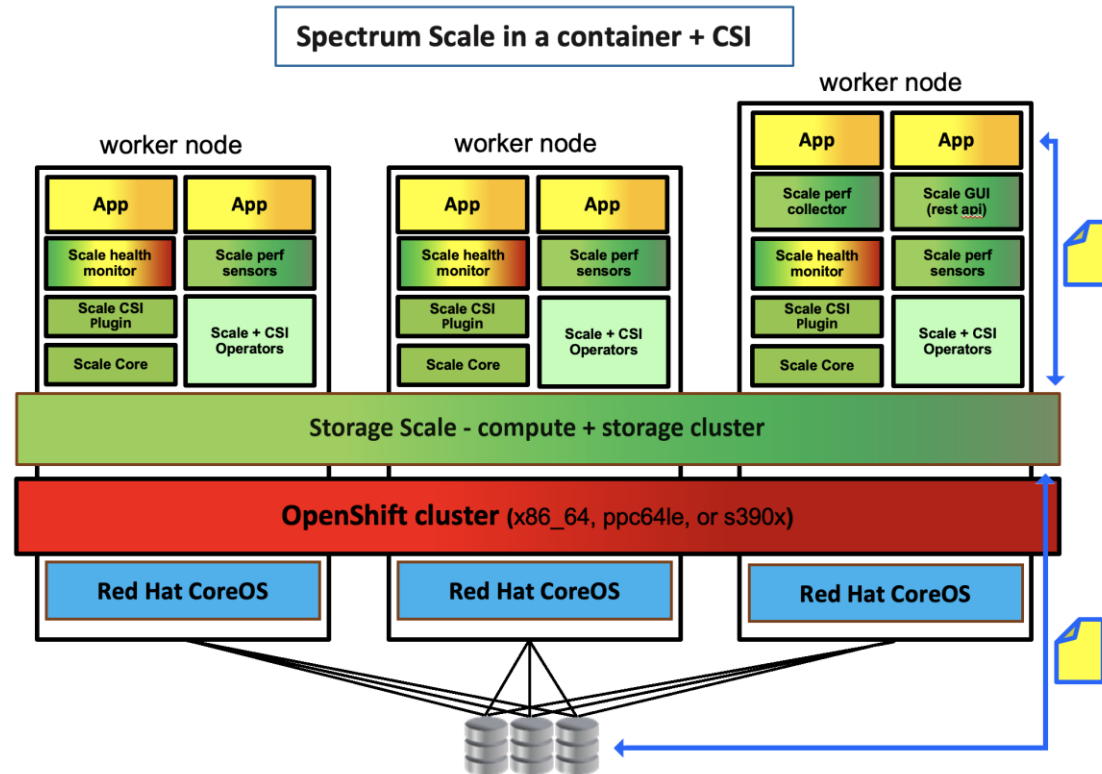
Multiple GUI hosts can be specified for CSI.
CNSA 5.2.1 will use multiple hosts in operator

Configure Resource limits of core pods

Vela only: AFM caching via StorageClass addition of volumeType: “cache” as well as cacheMode options

Local disk attachment utilizing a direct disk attachment configuration, replacing prior technology preview of a shared nothing local disk configuration.

Infiniband RDMA support (RoCE coming soon)



Access Services – Container Storage Interface

Improvements introduced in CSI 2.14

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

Support for CSI specification 1.9

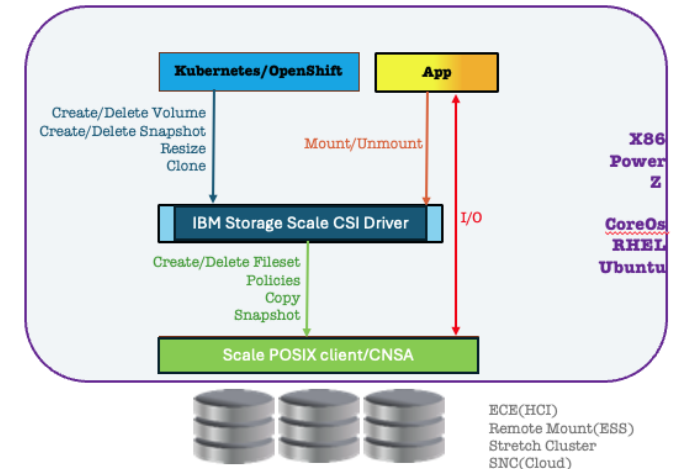

Support to customize the volume name prefix

Support for Kubernetes 1.31 and Red Hat® OpenShift® 4.17 / 4.18

Improvements in the script for debug data collection

storage-scale-driver-snap.sh [-l | -n | -o | -p | -s | -v | -h]

Important: Starting with IBM Storage Scale Container Storage Interface 2.13.x, the support for OpenShift with RHEL worker nodes is discontinued.



Dynamic Provisioning - Create/Delete Volume

Static Provisioning

Volume Snapshot

Volume Expansion

Shallow Copy

Volume Cloning

Compression

Tiering

ConsistencyGroup


Remote Mount

Multiple Filesystem

FsGroup

GUI HA

Lightweight Volume/Fileset Based Volume



Access Services – ARM

GA! -The official Architecture name is aarch64

Wider support to use ARM functionality Data Processing Units (DPU)

Current goal: ARM client

compute nodes (Grace Hopper)

Make it a platform for Scale like any other

DPU (Blue Field-3) for exploitation
research spike

- BF-3 can be used as NIC with Scale as any other supported NIC given OS and MOFED supports it
- Work in Progress for further exploitation
- IOR results from Grace Hopper
 - Max Write: 46437.87 MiB/sec (48693.63 MB/sec)
 - Max Read: 47281.41 MiB/sec (49578.15 MB/sec)



• **Included**

- SE package / install toolkit / rpm based install
- NSD client
- Scale base functionality (IO, policies, remote mounts, snapshots, quotas, etc.)
- Manager roles: file system manager / token manager / cluster manager
- RDMA (IB or RoCE) including GDS
- Health Monitoring
- Target OS: RHEL 9.3 and Ubuntu 22.04 (ask to open RFE for customers askign for RHEL 8)
- File audit logging, watch folders folders
- Call home
- GUI (can display ARM node, but cannot run on ARM)


- The NSD server functionality is now supported on arm64 platforms.

• **Excluded**

- SNC
- Protocols
- BDA / HDFS
- CNSA
- TCT (discontinued)
- HSM




QuantaGrid S74G-2U



- We need to learn whether there are ARM designs that need code changes
 - so far the only one has been Raspberry Pie ;-)
 - ... and that has been fixed but is still not supported

Performance update from 6000

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







ECE - New EC code 16+2P/3P – 1

more efficient use of capacity and some improvement with full track writes

Number of nodes	3WayReplication	4WayReplication	4+2P	4+3P	8+2P	8+3P	16+2P	16+3P
3	1 Node + 1 Device *	1 Node + 1 Device *	Not recommended 1 Node	Not recommended 1 Node	Not recommended 2 Devices	Not recommended 3 Devices	Not recommended 2 Devices	Not recommended 3 Devices
4	1 Node + 1 Device *	1 Node + 1 Device *	Not recommended 1 Node	1 Node + 1 Device #	Not recommended 2 Devices	Not recommended 1 Node	Not recommended 2 Devices	Not recommended 2 Devices
5	2 Nodes	2 Nodes *	Not recommended 1 Node	1 Node + 1 Device	Not recommended 1 Node	Not recommended 1 Node	Not recommended 2 Devices	Not recommended 2 Devices
6	2 Nodes	2 Nodes *	2 Nodes #	2 Nodes	Not Recommended 1 Node	1 Node + 1 Device #	Not recommended 2 Devices	Not recommended 2 Devices
7	2 Nodes	2 Nodes *	2 Nodes	2 Nodes*	Not Recommended 1 Node	1 Node + 1 Device	Not recommended 2 Devices	Not recommended 2 Devices
8	2 Nodes	2 Nodes *	2 Nodes	2 Nodes*	Not Recommended 1 Node	1 Node + 1 Device	Not recommended 2 Devices	Not recommended 2 Devices
9	2 Nodes	3 Nodes	2 Nodes	3 Nodes	Not Recommended 1 Node	1 Node + 1 Device	Not Recommended 1 Node	Not Recommended 1 Node
10	2 Nodes	3 Nodes	2 Nodes	3 Nodes	2 Nodes #	2 Nodes	Not Recommended 1 Node	1 Node + 1 Device #
11+	2 Nodes	3 Nodes	2 Nodes	3 Nodes	2 Nodes	3 Nodes	Not Recommended 1 Node	1 Node + 1 Device
18	2 Nodes	3 Nodes	2 Nodes	3 Nodes	2 Nodes	3 Nodes	2 Nodes #	3 Nodes
19	2 Nodes	3 Nodes	2 Nodes	3 Nodes	2 Nodes	3 Nodes	2 Nodes	3 Nodes

- To protect data from disk failure, all failure tolerances that are marked with # in the table need to be paid attention to for the spare disk space other than the erasure code. You can change the number of spare disk space to the same or bigger number than the node number before creating vdisks. For example, for 6 nodes with 4+2p erasure code, you can change all DA's spare disk space to 6 before creating vdisks.
- All failure tolerances that are marked with * are limited by recovery group descriptors rather than by the RAID code.

Performance improvement for cached objects cleanup for snapshot delete



Linear batched deletions
of mmdelsnapshot

Online Filesystem Check Updates!

mmfsckx improvement

Mitigate inode update slowness during online fsckx scanning on a large inode 0 files

Improvement Summary:

Reduce impact to inode update performance during mmfsckx on a filesystem with large inode 0 files

Performance evaluation:

Evaluate file create/delete rate during mmfsckx
Monitor mmfsckx time

Configuration:


- x10 x86 client nodes
- ESS3500 performance model with 24 NVMe Drives
- 200Gib IB with RDMA
- 1B inode allocated
fileset with 1M 1024KB files in 10 subdirectories, 4000 10GB files

Performance Results

	mmfsckx (min)	File create KIOPs	File delete KIOPs
R5.2.0	13		
		65	78
	13	IO aborted after being paused > 3m	
R5.2.1	14		
		66	78
	14	51	72

- mmfsckx time remains equivalent
- File create/delete was impacted slightly with mmfsckx


Abstraction and Acceleration Services – Dynamic Page Pool



Dynamic workload management!

Scale detects a shortage of the pagepool memory, then attempts to increase the pagepool size.

When the Linux kernel detects the memory pressure, it requests Scale to shrink the size of the pagepool.



Configuration:

`mmchconfig dynamicPagepoolEnabled=yes -N node1`

`mmchconfig pagepool=default -N node1`

`mmshutdown -N node1`

`mmstartup -N node1`

`mmdiag -pagepool`

`GPFSBufMgr monitor pagepool size via zimon`

Config parameter	Allowed values	Default	Description
dynamicPagepoolEnabled	yes/no	no	Enable dynamic pagepool vs. static pagepool
pagepoolMinPhysMemPct	1-50	5	Minimum size of dynamic pagepool as percentage of physical memory.
PagepoolMaxPhysMemPct	10-90	75	Maximum size of dynamic pagepool as percentage of physical memory.
pagepoolChangeGracePeriod	1-86400	10	The grace period for growing the dynamic pagepool, in seconds. The dynamic pagepool grows only once every grace period.

Default configuration changes with 5.2.N


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

Abstraction - Management and Orchestration



Monitoring and Health

Call Home:

- TIPS for new mmchconfig defaults
- Unified Call Home for all ESS (no need for ESA agent)
Hardware-related health events will automatically create a call-home ticket
- Upload to pmr with: `gpfs.snap --pmr XXXXXX`
- Automatic group expansion for new ESS nodes
- mmcallhome test connection: better analysis of the root cause of connectivity issues

mmhealth:

- Added cluster certificate expiration monitoring
- Extended SCALEMGMT health monitoring
- Added swap memory usage monitoring

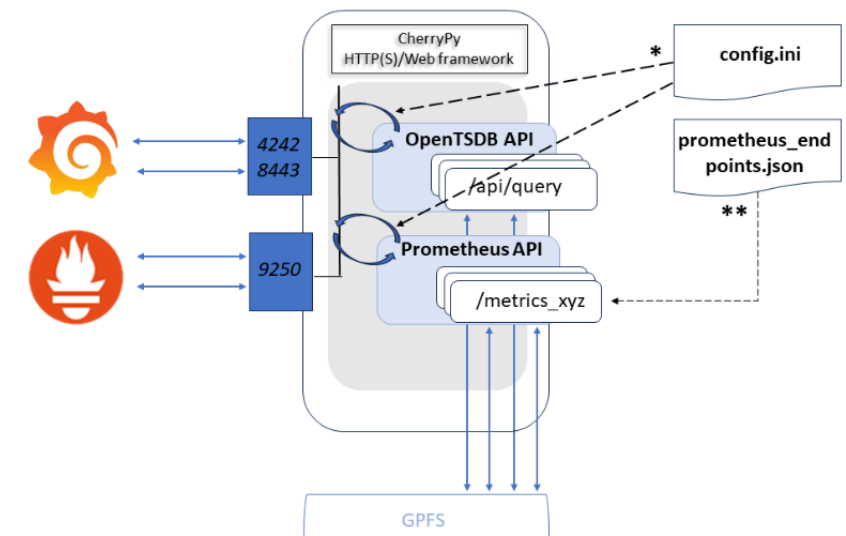
Perfmon

- Prometheus Exporter for Zimon Data
- New Network Sensor

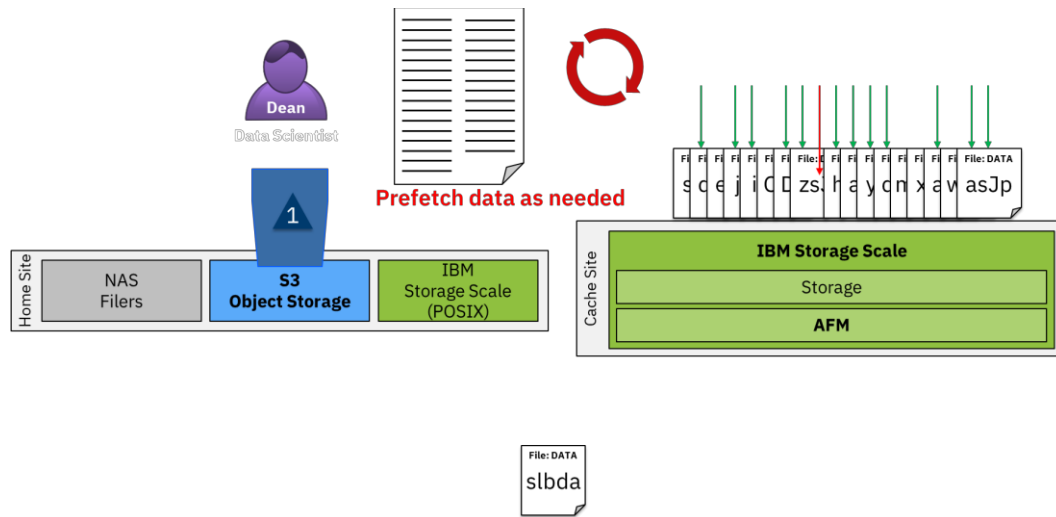
FTDC:

- Improved data collection for S3

Callback	Resulting Events
postRGTakeover	gnr_rg_takeover gnr_rg_takeover_warn
postRGRelinquish	gnr_rg_relinquish gnr_rg_relinquish_warn
rgPanic	gnr_rg_server_panic
daRebuildFailed	gnr_da_rebuild_failed



Abstraction and Acceleration Services – Active File Management (AFM)



- Provide option to delete objects using Non-MDS gateway
- Inode eviction from AFM cache
- Simplification of migration commands improves user experience. One command will ensure all critical steps are completed internally.
#mmafmctl gpfs11 startCutover -j ro1
- Support of CES S3 for AFM to cloud object storage fileset.
 - AFM to cloud object Storage already supports IBM COS, Amazon AWS S3, Seagate Lyve Cloud, and Microsoft Azure Blob storage as Cloud Object Storage backend.
 - AFM cloud object storage fileset can use CES S3 endpoints (Nooba backend) seamlessly.
- Provide option to enable Automatic eviction after upload/reconcile
mmafmcosctl fs1 fileset1 /gpfs/fs1/fileset1 upload --all --evict-after-upload
- User should be able to perform Range Eviction on individual fileset.
- Added fileset level parameter to enable NFS V4 support for specific fileset(s) instead of all **# mmchfileset fs1 testfset -p afmNFSv4=yes|no**
- Provide custom storage location support for recovery/reconcile events at fileset level.
- Support of multi site replication in AFM to cloud object storage MU mode
- New callback added for AFM DR RPO sync: **afmDRRPOSync**
- New cluster parameter : '**afmObjMUCheckFName**' & '**afmFastLookup** & **afmLookupMapSize**'
- Allow numeric uid/gid for mmafmcosctl and mmafmcosconfig
- Support for downloading object and directory with single list with COS backend

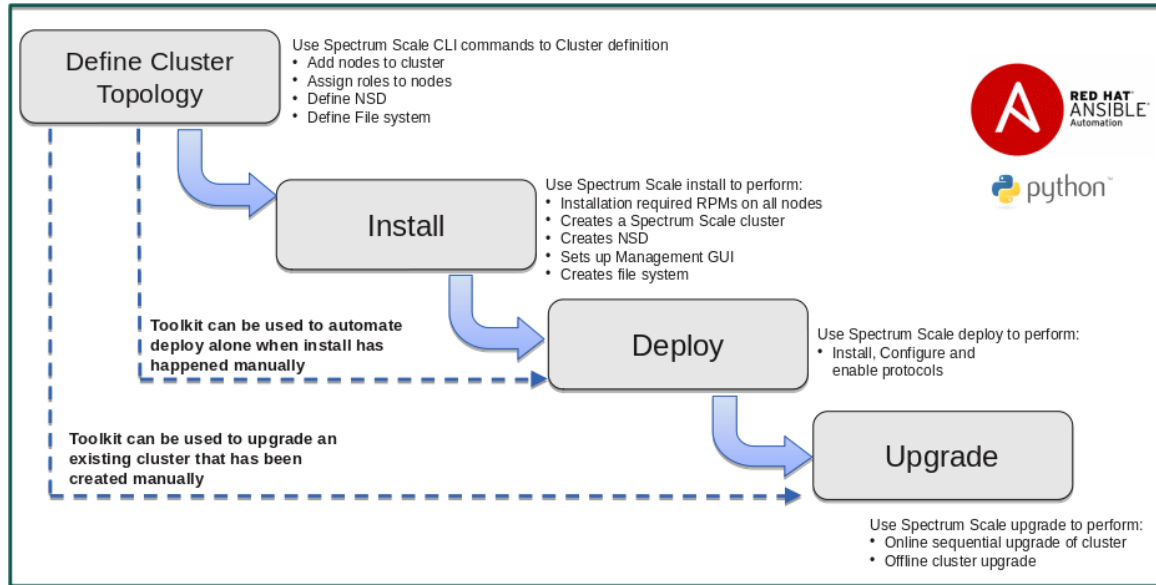
Validation Tool using REST API and
Certification Tool using:
mmafmtransfer

Migrate TCT enabled fileset to AFM-
S3 MU (Tiering only)

mmexpelnode Persistence

- The previous behavior was that nodes expelled with **mmexpelnode** were allowed to rejoin the cluster (“unexpelled”) once the cluster manager node changes for any reason.
- Only available on (**and is default on**) clusters with `minreleaselevel >=5.2.1` (both newly created and after `mmchconfig release=LATEST`)
- Can restore old behavior with `mmchnode disablepersistExpelList=yes`
 - Don’t forget “-i” if not restarting the cluster
- Use **mmexpelnode -l** to list the already-expelled nodes. Expelled nodes that are still running `mmfsd` will show up as “arbitrating” in **mmgetstate** output.
- This does NOT affect what GPFS does internally/automatically as far as expelling nodes that are unresponsive to RPCs.

Storage Scale Deployment Toolkit



[CES S3] CES S3 based Object protocol toolkit support for X86.

[CES S3] CES S3 based Object protocol toolkit support for PPC64LE.

[CES S3] CES S3 based Object protocol toolkit support for S390X.

[**Native Rest API Tech Preview**] Toolkit supported features for Native Rest API Tech Preview Deployment

[**Toolkit Arm**] Toolkit extended features support on ARM .

[**Python**] Smart Installer Revolutionises Python management, Automatically utilises Latest python installed version without user configuration.

[**Currency**] Extended OS currency

[**ESS Protocol Node**] ESS Protocol node certification with 5.2.1 Toolkit.

[**ECE**] ECE install toolkit enhancement to support in config populate with vdiskset in multi-DA and file system with multiple vdisk sets


[**Cloud**] NFS & SMB support for Cloud-Kit.

[**Cloud**] CES S3 Support for Cloud-Kit.

[**Open Source**] Open Source Ansible Role certification with 5.2.1.

[**Documentation**] Ansible tuning config for deployment consideration

Abstracting Cloud Service Deployment – Cloudkit!



What is Storage Scale Cloudkit?

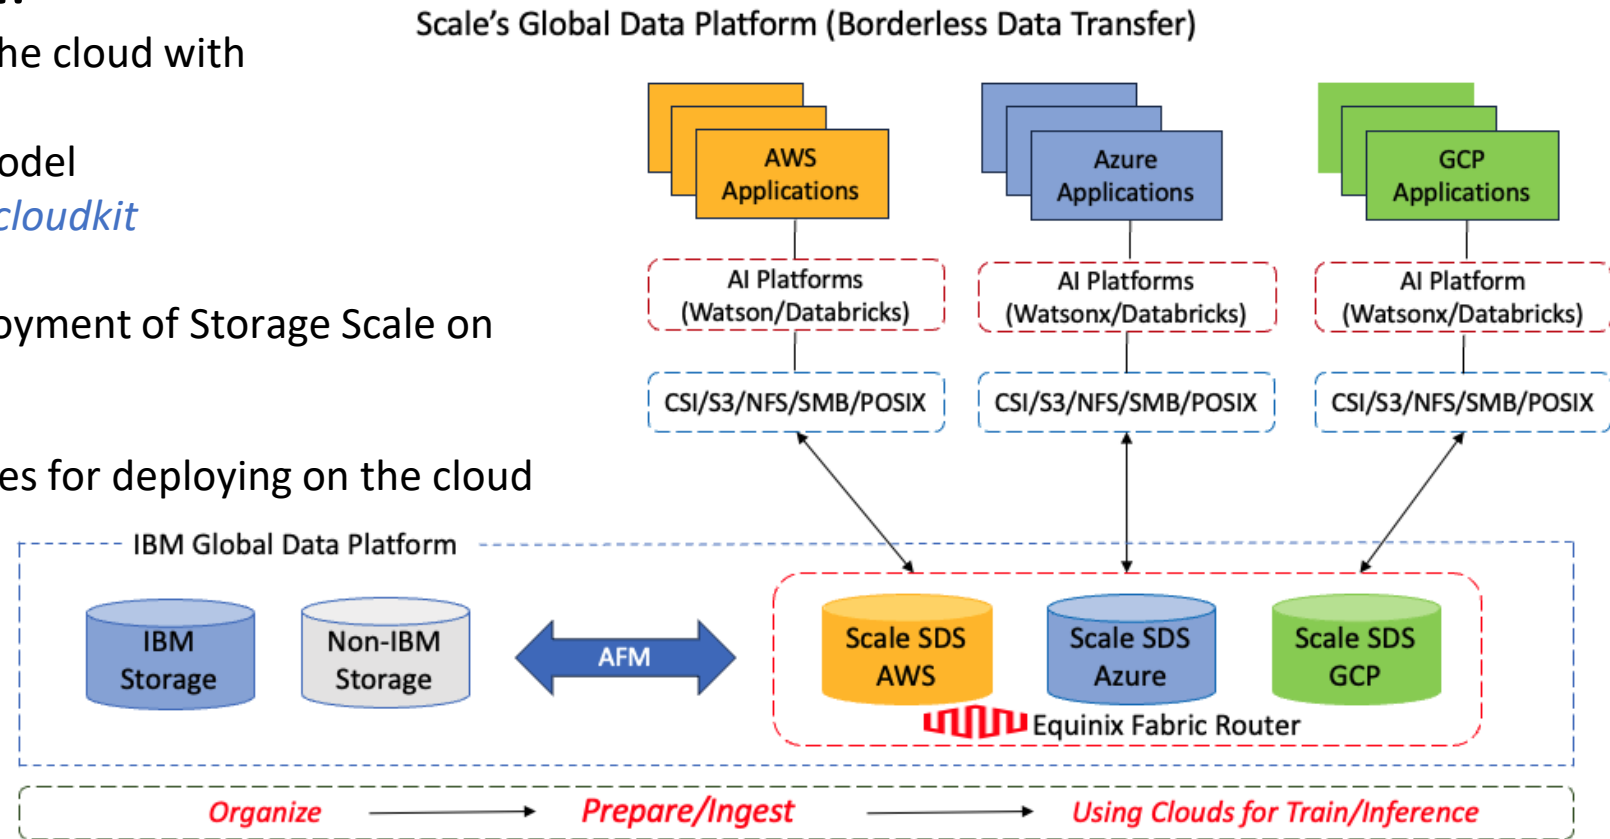
Create Storage Scale clusters on the cloud with

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




Advantages

Support for major public clouds Amazon (AWS) and Google (GCP)

AFM-COS, Upgrades

Tech-preview support for fleet support on AWS and GCP cluster instance

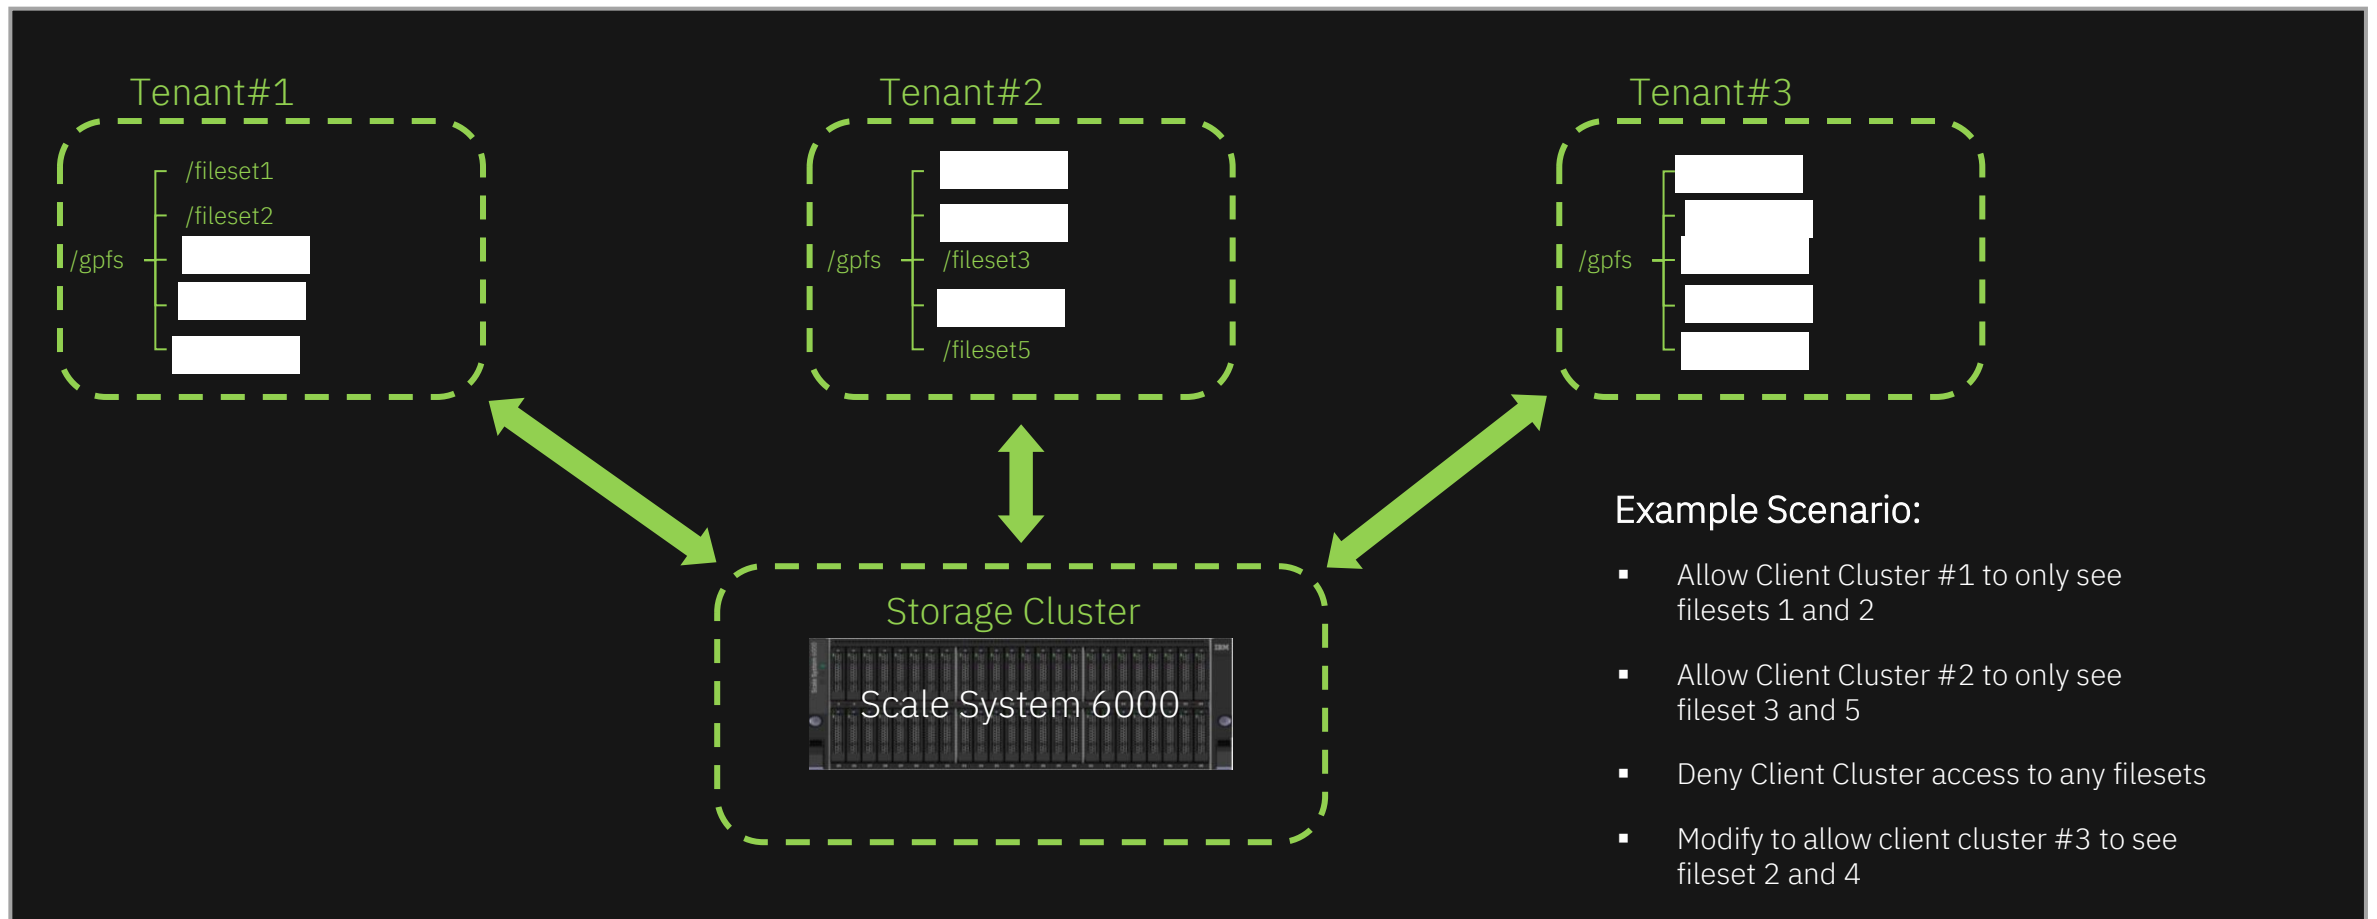


More news...



Remote Fileset Access Control

- Provides multi-tenancy capabilities for remote client clusters
- Define which remote clusters can see which filesets within a single filesystem namespace
- Dynamic ability to grant or deny fileset access to a remote cluster using *mmauth* allow or deny command
- Quotas and snapshots will only be visible for the authorized filesets, not all filesets within a filesystem



Scale System models are built for speed and capacity

Speed

Capacity

Hybrid

IBM Storage Scale
System 6000

4U48 Enclosure
24 or 48 NVMe drives

IBM Storage Scale
System 3500



2U24 Enclosure
12 or 24 NVMe drives

IBM Storage Scale System
3500 Hybrid (NVMe + HDD)

IBM Storage Scale System
6000 Hybrid (NVMe + HDD)

IBM Storage Scale System
3500 Capacity (HDD-Only)

IBM Storage Scale System
6000 Capacity (HDD-Only)



Performance
Machine type 5149-F48

Performance
Machine type 5141-FN2

Up to 8x 4U102 JBOD

Up to 9x 4U91 JBOD

Up to 8x 4U102 JBOD

Up to 9x 4U91 JBOD

3500 Hx
Machine type 5141-FN2

6000 Hx
Machine type 5149-F48

3500 Cx
Machine types
5141-FN2

6000 Cx
Machine types
5141-FN2

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

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

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

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

New Scale System Software Features



Scale 5220

RH9.4 (UN / 6000)

Gen5 Samsung drives

Red fish script support for ESS 6000/UN (read-only)

mmptop (live CPU/memory info)

Improved call home ticket lifecycle management

S3 support on utility node (non-protocol VM (3500/6000)

Firmware updates for 6K/4u91

MES for Falcon/HBA (1-9 enclosures) ESS 6000 4u91

MES 24 to 48 FCM drive support

Protocol VM ESS 6K (POC)

SED support with TPM (no key server) ESS 6K

ESS HW metrics in Zimon

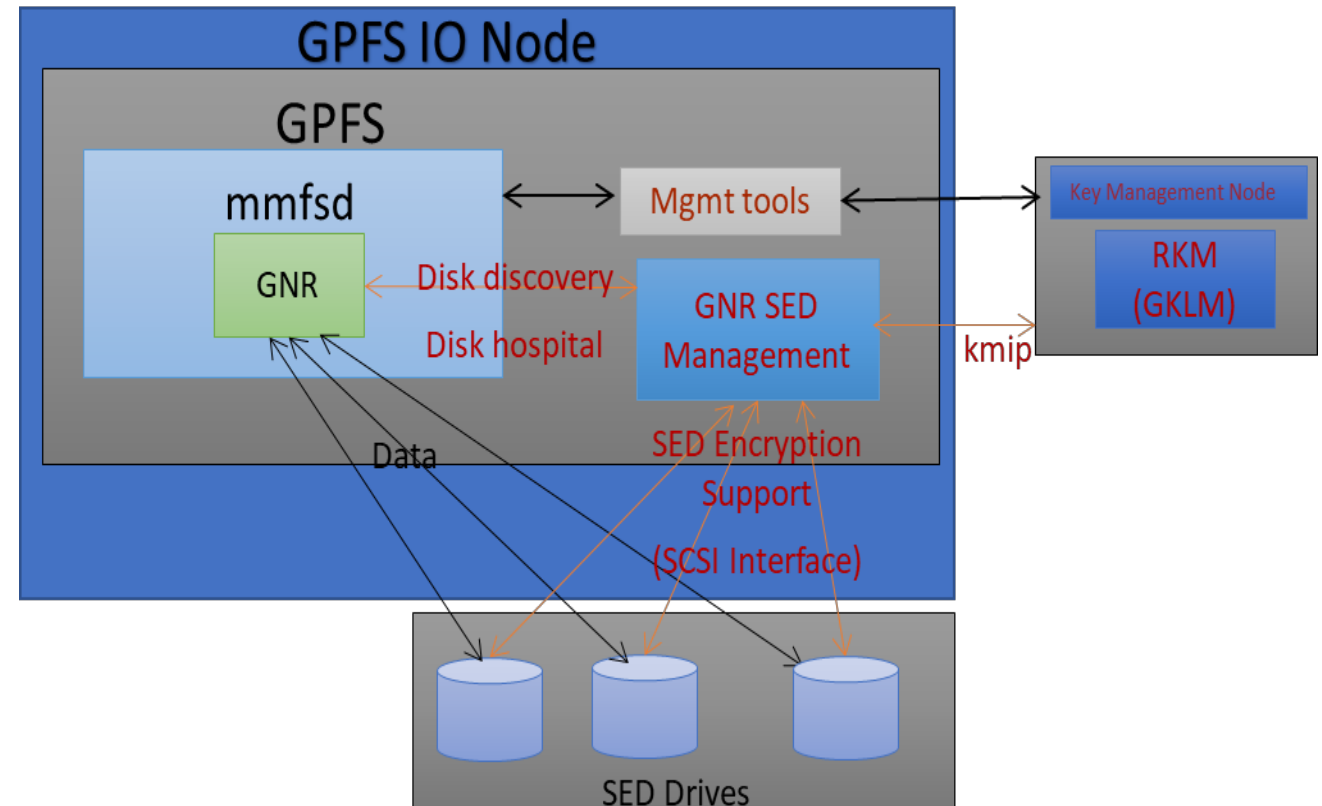
SED Support with GKLM : Overview

Background:

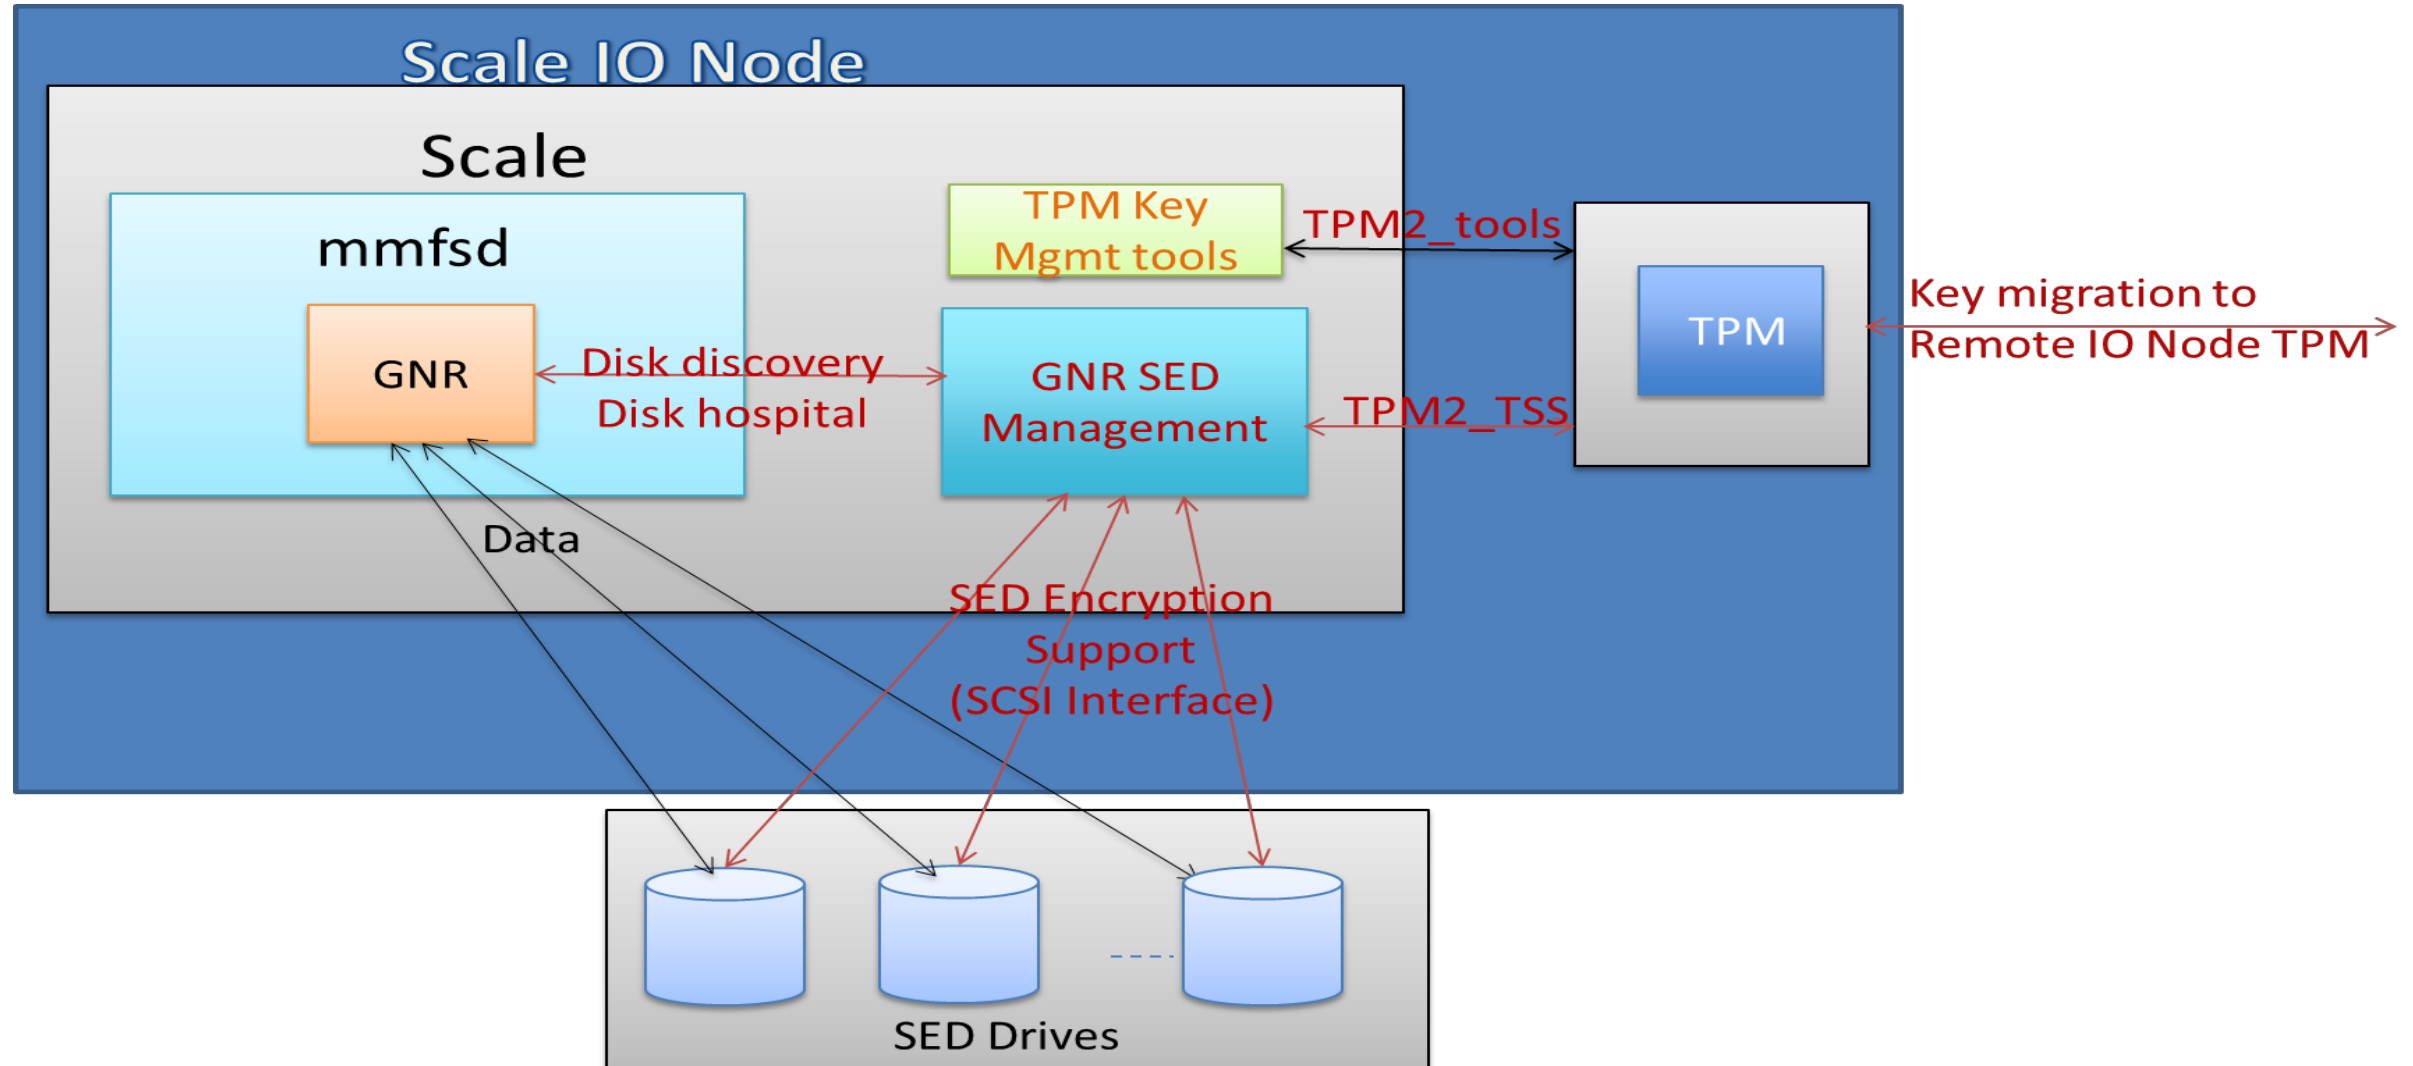
- ❑ SED enabled by enrolling with MEK
- ❑ Auto lock on power off
- ❑ Data Security at Rest
- ❑ Need to unlock at Power ON using MEK
- ❑ Crypto erase by changing DEK

Challenges:

- ❑ External Key Managers are expensive
- ❑ Different Key Managers



SED Support using TPM : Overview



Data Acceleration Tier (DAT): Ustore NVMeoF Monitoring

Mmhealth enhancement for Ustore / NVMeoF


- Nvme component (ESS side)
 - show non-gnr nvme devices in mmhealth nvme component (server side). Existing nvme checks are done on exported nvmes too (e.g. nvme_temperature_warn)
 - Additional smart check through HAL (nvmeof_raw_disk_smart_failed)
- NVMeoF component (ESS side)
 - Detect if node exports nvmes (mmvdisk nvmeof list -Y)
Noderole=NVMeoFTarget
 - Check packages, modules, multipath settings
- Disk Component (client side)
 - By default mmhealth shows NSDs which the node is NSD server only (not the case for ustore)
 - Evaluate nodeclass „nvmeofClients“ to show and monitor all NSDs incl. Ustore NSDs

```
# mmhealth node show nvmeof -v
Node name: c145f11san06b.gpfs.net
```


Component	Status	Status Change	Reasons & Notices
NVMeoF	HEALTHY	2024-03-21 17:23:30	-

Event	Parameter	Severity	Active Since	Event Message
nvmeof_modules_installed	NVMeoF	INFO	2024-03-21 17:23:30	NVMeoF modules are installed.
nvmeof_multipath_disabled	NVMeoF	INFO	2024-03-21 17:23:30	Native multipath is disabled for NVMeoF.
nvmeof_packages_installed	NVMeoF	INFO	2024-03-21 17:23:30	NVMeoF packages are installed.

Thank you for using



Storage Scale



Storage Scale
System

