



# ADVANCED TECHNOLOGY GROUP (ATG)

---



## **Accelerate with ATG Webinar: FlashSystem The Next Generation/IBM Storage Virtualize 9.1.2 Technical Update**

Matt Key

Principal Storage Technical Specialist

IBM Advanced Technology Group – FlashSystem

[mkey@us.ibm.com](mailto:mkey@us.ibm.com)

Byron Grossnickle

IBM Storage Virtualize/FlashSystem SME

[byrongro@us.ibm.com](mailto:byrongro@us.ibm.com)



## Accelerate with ATG Technical Webinar Series

---

*Advanced Technology Group* experts cover a variety of technical topics.

**Audience:** Clients who have or are considering acquiring IBM Storage solutions. Business Partners and IBMers are also welcome.

To automatically receive announcements of upcoming Accelerate with ATG webinars - Clients, Business Partners and IBMers are welcome to send an email request to [accelerate-join@hursley.ibm.com](mailto:accelerate-join@hursley.ibm.com).

### 2026 Upcoming Webinars – Register Here!

[IBM Z Solutions Aligned to the NIST Cyber Resiliency Framework - Thursday, March 26th, 2026](#)

[Detecting Cyber Attacks with IBM Storage Defender - Tuesday, March 31st, 2026](#)

[IBM Ceph 9.0: Policy-Based S3 Cloud Tiering — Live Demo & Deep Dive - Thursday, April 2nd, 2026](#)

[Understanding the Cyber Vault for IBM I blueprint - Tuesday, April 7th, 2026](#)



### *Important Links to Bookmark:*

[Accelerate with ATG](#) ← Click here to access the Accelerate with ATG webinar schedule for 2026, view presentation materials, and watch past replays dating back two years. <https://ibm.biz/BdSUFN>

[ATG MediaCenter Channel](#) ← This channel offers a wealth of additional videos covering a wide range of storage topics, including IBM Flash, DS8, Tape, Ceph, Fusion, Scale, Cyber Resiliency, Cloud Object Storage, and more. <https://ibm.biz/BdfEg0>

## Offerings

---

### Client Technical Workshops

- Cyber Resilience with IBM Storage Defender
- IBM Fusion & Ceph
- IBM Storage Scale & Storage Scale Advanced Functions
- IBM DS8000 G10 Advanced Functions
- IBM FlashSystem Deep Dive & Advanced Functions

### TechZone Test Drive / Demo's

- IBM Storage Scale and Storage Scale System GUI (and hands-on clusters)
- IBM Storage Virtualize Test Drive
- IBM DS8900F Storage Management Test Drive
- Managing Copy Services on the DS8000 Using IBM Copy Services Manager Test Drive
- IBM DS8900F Safeguarded Copy (SGC) Test Drive
- IBM Cloud Object Storage Test Drive - (Appliance based)
- IBM Cloud Object Storage Test Drive - (VMware based)
- IBM Storage Protect Live Test Drive
- IBM Storage Ceph Test Drive - (VMware based)

Please reach out to your IBM Representative or Business Partner for more information.

**\*IMPORTANT\*** The ATG team serves clients and Business Partners in the Americas, concentrating on North America.

## Meet the Speakers – Matt Key

---



Matt Key is an 20-year veteran of solid-state storage and came into IBM through the 2012 acquisition of Texas Memory Systems, the Houston-based engineering group behind the FlashCore Modules in IBM's distributed storage portfolio. Currently, he is the technical lead for Flash in the Advanced Technology Group (ATG), the client-facing group of subject matter experts across the portfolios of storage, servers, and software. Matt graduated Texas A&M (WHOOOP!) in 2006 with an engineering degree in telecommunications.

## Meet the Speakers – Byron Grossnickle

---



Byron Grossnickle is an IBM Storage Technical Specialist concentrating on Storage Virtualize software. This includes FlashSystem, SVC, and Storage Virtualize for Public Cloud. Byron has been with IBM 19 years exclusively in storage. Prior to working for IBM, Byron spent 6 years engineering storage in the Telcom Industry. Prior to that he worked 8 years in healthcare IT. Byron lives in the Kansas City area and is available to travel to customer engagements.

## Please take our survey at the end!

---

Please take a moment to share your feedback with our team!

You can access this 6-question survey via this direct link: <https://ibm.biz/Bdpcj3>

Or you can use this QR Code:



## IBM Storage Virtualize 9.1.2 Agenda

---



- Release Schedule
- Messaging
- Hardware Updates
  - 5<sup>th</sup> Generation FlashCore Modules
  - FlashSystem 9600
  - FlashSystem 7600
  - FlashSystem 5600
- Performance
- FlashSystem.ai
- Software Updates
- Ansible Updates

# Release Schedule

- Marketing/RFA Announce – 9.1.2 – Feb 10, 2026
- eGA – 9.1.2 – Mar 06, 2026
- pGA – March 06, 2026
- eGA – FlashSystem.ai – Mar 27, 2026

9.1.2 is a Non-LTS Release. Non-LTS releases are tested identically to LTS releases. Non-LTS releases generally do not get any patches and will not be maintained long term. Those patches will be available in the next Non-LTS or LTS release.

As 9.1.2 ships with new hardware, machine-hours and exposure to runtime will establish a near-LTS visibility.



# IBM FlashSystem

Engineered to Adapt, Built to Perform, Designed to Protect

IBM FlashSystem can help you **transform your storage** from a passive and reactive repository to AI-Driven Dynamic Storage. It achieves this through **pervasive intelligence** enabled by IBM **FlashSystem.ai** and **breakthrough capacity** with our **fifth generation FlashCore Module**, helping you to turn your storage operations into a **strategic advantage**.



## Adapt

Overcome skills shortages and challenges

Accelerate time to business value



## Perform

Lower storage costs

Meet increasing performance demands



## Protect

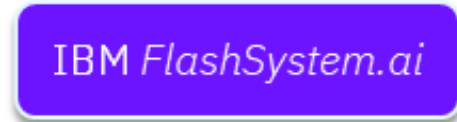
Reduce business disruption risks

Reduce compliance risks

# Summary

## Key product deliverables

### IBM FlashSystem 1Q'26 Launch



FlashSystem.ai  
Self-managing storage  
powered by AI.

- For 1Q'26 customers will be able to provision volumes by a request on natural language.
- Runs on-premises and leverages IBM watsonx for AI-driven orchestration on cloud.
- Radical simplicity: Ask, approve, and go.



Next-Gen Hardware  
Industry-leading performance,  
density and efficiency

- Controllers delivering up to 7 times more capacity than competitors
- Systems that pack 32 drives in just 2U.
- Ultra density for market lead TCO. Footprint, Energy consumption, Performance.



FlashCore Modules  
Intelligent density on a  
new form factor

- Announcing the 5<sup>th</sup> generation
- New form factor with enhanced data reduction by deduplication
- Increased capacity, more density on each drive.

# Adapt: Amplify expertise and accelerate decisions and action to deliver more agility, powered by FlashSystem.ai

***The Challenge:** Enterprise IT teams struggle to keep pace with rapidly changing business demands and emerging opportunities. They must adapt infrastructure, scale resources, and align data strategies quickly—without sacrificing reliability, compliance, or cost control—creating constant tension between agility, governance, and operational complexity.*

FlashSystem.ai transforms your storage into an intelligent, autonomous system that delivers more business value.

## **Overcomes skills shortages and challenges**

Amplify and scale human expertise with a natural language interface that can interpret and respond to user inquiries, AI-enforced data contracts, and autonomous data services.

## **Accelerate decision making & response**

Reduce time and effort needed to identify issues and take corrective actions with autonomous, intent-aware data services with managed oversight

**Amplify your storage team's expertise** with AI that continuously learns from feedback and suggests opportunities for improvement.

**Deliver seamless, self-service operations** that performs a wide range of tasks on your behalf with minimal manual effort.

**Align your storage with business intent** through Data Contracts, which helps to consistently deploy and manage storage aligned to business intent.

# Protect: Secure and compliant storage that can improve your operational resilience, powered by FlashSystem.ai

***The Challenge:** Enterprise IT teams face growing pressure to protect data from ransomware, malware, and other threats while meeting regulatory and organizational compliance requirements. Balancing security, governance, and operational efficiency is difficult as threats evolve, regulations tighten, and data volumes surge—creating risk & complexity.*

FlashSystem.ai transforms your storage into an intelligent, autonomous system that delivers more business value.

## Context-Aware Security

Mitigate threats with FlashSystem's AI-driven ransomware detection and autonomous healing actions for faster, smarter protection and recovery that deliver improved operational resiliency.

## Dynamic Compliance

Enable consistent adherence through AI-driven automation that identifies geo-specific and industry requirements, streamlines compliance processes, and reducing manual intervention.

**Protect critical data** with immutable safeguarded copies and reduce disruption risks with ransomware threat detection in **less than a minute.**

**Detect cyber attacks in real-time** with continuous monitoring of statistics from every single I/O with **zero performance penalty.**

**Recover quickly after an attack** from known good, safeguarded copies within **60 seconds or less. Guaranteed\***

# Perform: Breakthrough speed, improved efficiency, and sustainable storage built to last, powered by FlashSystem.ai

***The Challenge:** Enterprise IT teams face constant pressure to deliver high-speed performance for critical applications while maintaining operational efficiency. Meeting these demands requires optimizing resources, controlling costs, and minimizing complexity—without compromising reliability or scalability, a persistent challenge as workloads grow and expectations rise.*

FlashSystem.ai transforms your storage into an intelligent, autonomous system that delivers more business value.

## **Address performance demands**

Improve storage response times, lower latency, and accelerate throughput with the new higher capacity fifth generation FlashCore Module.

## **Deliver more efficiency**

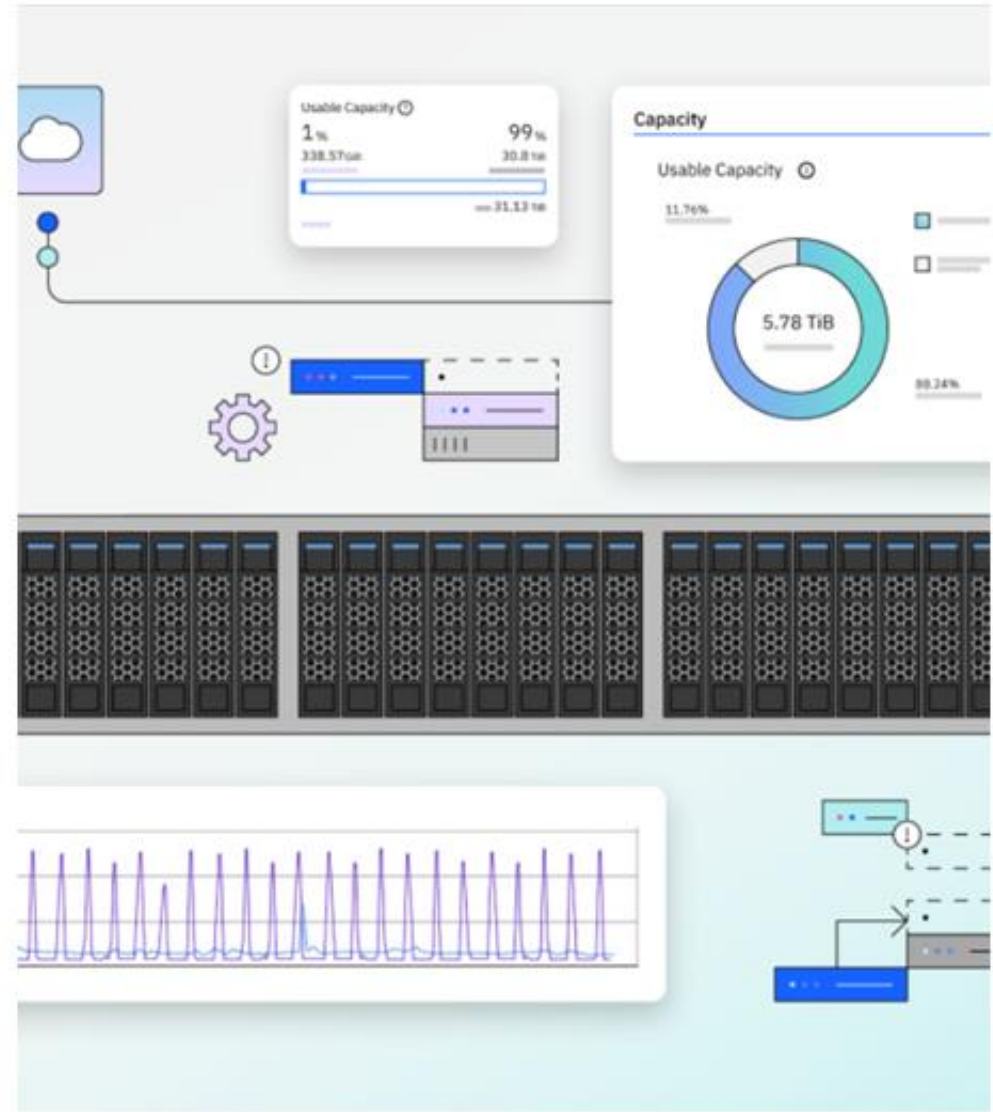
Proactively optimize storage with autonomous tuning, intelligent data placement, machine & human actionable optimizations, while consuming less energy, more drive durability, and hardware-level deduplication.

**Autonomously optimizes performance** based on real-time workload analysis, ensuring consistent and efficient data access.

**Utilizes AI to automatically move data** between different storage tiers based on access patterns, optimizing performance and cost.

**Dynamically scales resources up or down** based on workload demands, delivering cost efficiency and flexibility.

## IBM FlashSystem Hardware



# IBM FlashSystem 2026



## FlashSystem C200

Capacity optimised, fixed config (2.3PBe)  
for lower performance requirements  
<200K IOPS, 1-2ms



## FlashSystem 5600

1U entry level NVMe for general  
purpose workloads  
>380K IOPS, <0.25ms, >30GB/s,  
2.4PBe in 1U



## FlashSystem 7600

2U balanced performance and value  
workhorse system  
>700K IOPS, <0.25ms, >55GB/s  
7.2PBe in 2U



## FlashSystem 9600

Extreme performance and scaling for  
mission critical workloads  
>1.7M IOPS, <0.25ms, >85GB/s  
**11.8PBe in 2U**



IBM FCM Gen4



IBM FlashCore Gen5 EDSFF



High Density, Low Footprint, Lowest Power in the industry - sustainable, less hardware to achieve performance  
Best and consistent performance with FCM compression & deduplication, and 276-layer NAND with FlashCore endurance  
Real-time Ransomware Detection, Cyber Vault, and Safeguarded Copy - unique in the industry

Powered by IBM Storage Virtualize



Managed, monitored and proactive support by IBM Storage Insights



Storage Virtualize  
for Public Cloud

Hybrid Cloud integration  
TCT scheduled offload to S3

## SAN Volume Controller (SVC)



SA2 Storage Engine



SV3 Storage Engine

## FlashSystem 5015 & 5045



# Introducing FCM 5<sup>th</sup> Generation

## Performance

*Better power delivery  
and signal integrity*

## Cooling

*Optimised shapes allow  
for more efficient airflow*

## Scalability

*Supports future needs for  
higher capacities, faster speeds*

PSID: 49Q9VN1NAE9PCTA6954RN  
WWN: 000255012F002140

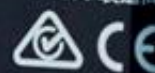


(11S)PN SN:03NK794YS124D5CH00



IBM

制造商/製造商



New Orchard  
http://



*Evolving the form factor from  
U.2 to E3.L to meet the next  
generation of computational  
storage requirements with a  
form factor designed for NAND  
rather than spinning disk...*

15—16

17—18

19—20

## FCM 5th Generation

- New Enterprise Datacenter Form Factor (EDSFF, E3.L1T)
- Latest 2Tb QLC NAND Flash Technology
- Increased Raw Capacity
- Increased Logical Capacity
- Increased Performance
- Increased Computational Offload
- Increased Security
- Increased Data Reduction
- Same FlashCore Reliability



*\*FCM2XL is capable of 6:1 addressability, however FlashSystem DRAID addressability is restricted to 4.88:1  
 \*\*Requires DRAID code changes to support max DRAID geometry*

Capacity Points in TB	Small	Medium	Large	XL	2XL
RAW Capacity	6.6	13.2	26.4	52.8	105.6
Physical Capacity (1:1)	5.6	11.2	22.4	44.8	89.6
Logical Capacity Ratio	6:1	6:1	6:1	6:1	6:1*
Logical Capacity	33.6	67.2	134.4	268.8	438.0*
Platform Support	x600	x600	x600	x600	9600

## FCM Security

### *IBM Quantum Safe Technology*

FCM 5<sup>th</sup> Generation uses a hybrid implementation of Post-Quantum Cryptography algorithms for Key Transport and Image Authentication:

- FIPs 203 and 204 Compliant Algorithms
- Hybrid use of ML-KEM (“CRYSTALS-Kyber”) for secure key transport of unlock PIN transmitted by FlashSystems storage controller to FCM
- Hybrid use of ML-DSA (“CRYSTALS-Dilithium”) signatures for verification of firmware authenticity

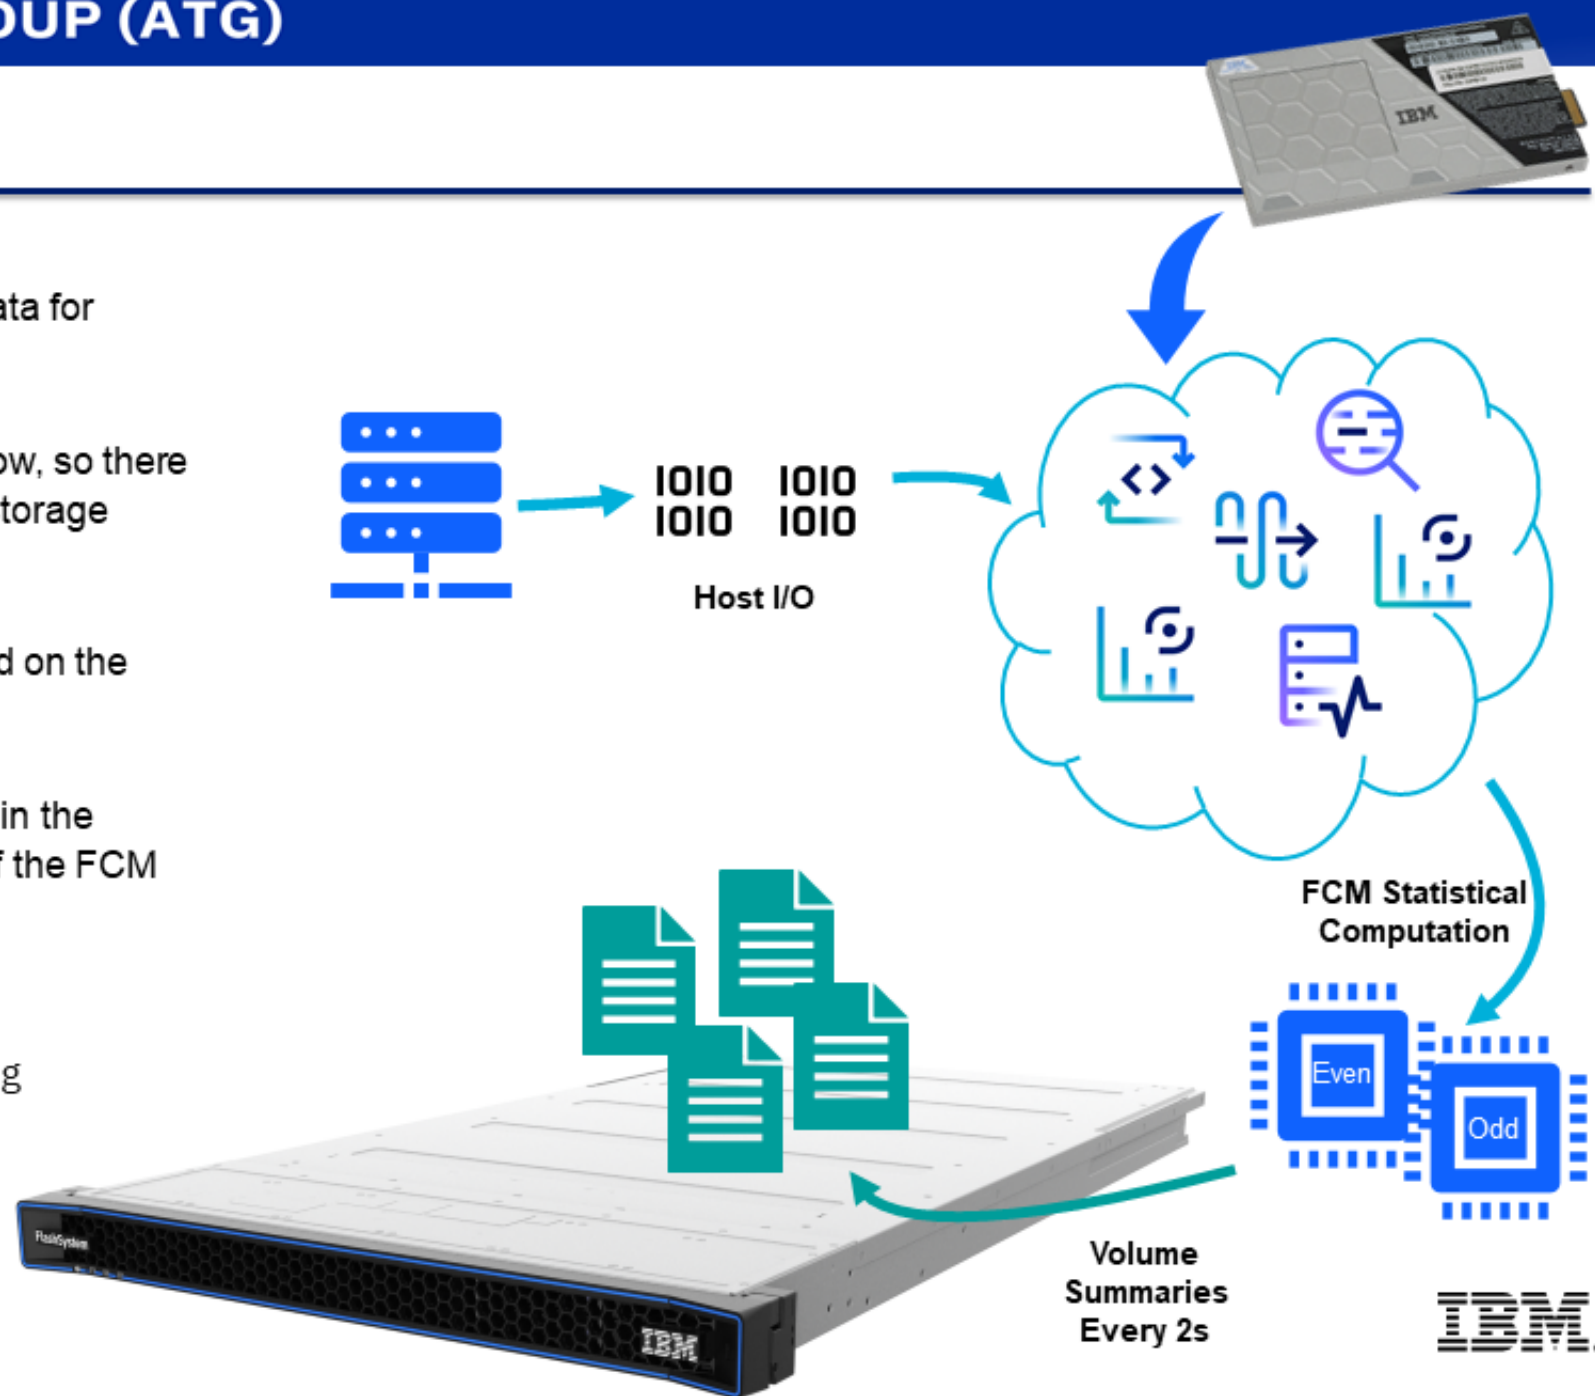
Customer data is stored in flash memory after being encrypted by a hardware XTS-AES-256 encryption engine.\*\*



## FCM Protection

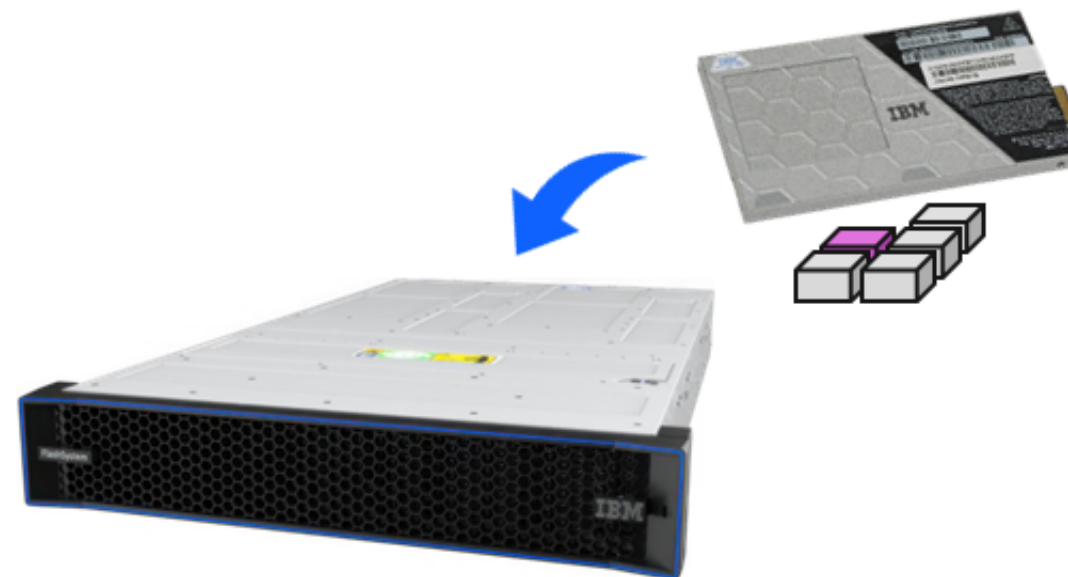
### Ransomware Threat Detection

- FCM is the ideal place to track and analyze data for intrusion detection
- Data is analyzed alongside the normal data flow, so there is no performance penalty, and no metadata storage needed.
- Statistics are collected and summarized based on the volume for every I/O
- Data analysis is fault resilient, stalls or issues in the analysis will not affect the storage functions of the FCM
- **New Statistics added to FCM5 NEW**
- Current models *still* do a great job of detecting encryption attacks
- Future models aim to target Exfiltration and Wiperware



## FCM Deduplication

- Provide True Data Reduction value to Customers with FCM Hardware Compression and Deduplication
- Storage Virtualize and FCM Co-design
- No additional appliance cost increase for computation
- No additional memory cost in controller for metadata
- Scalable at any capacity point
- No performance impact



## Deduplication

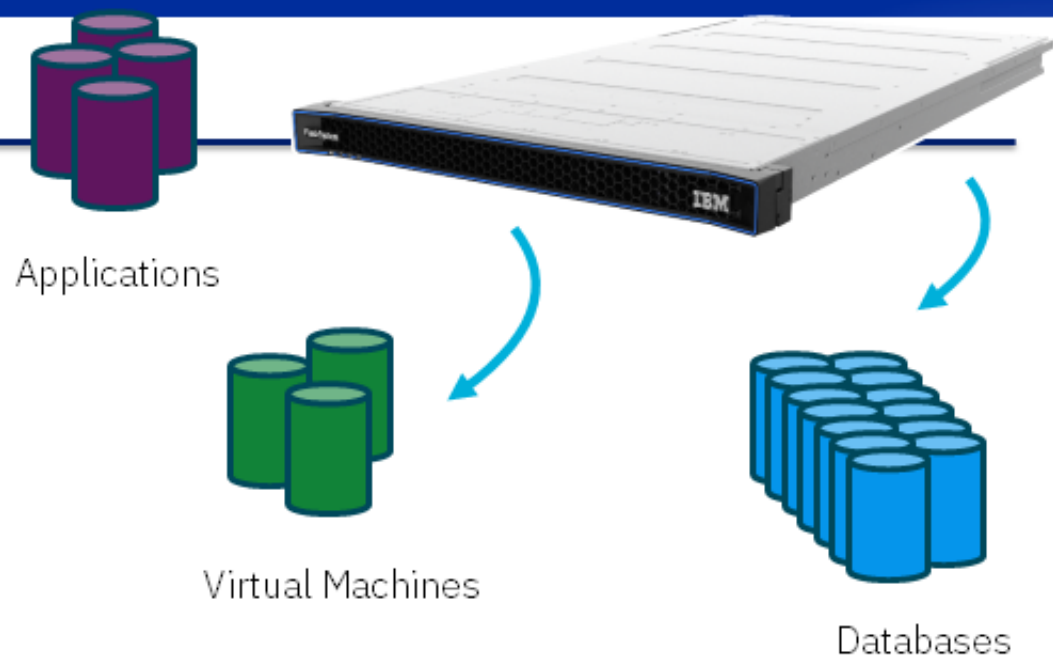
### *The Concept and the Challenge*

#### Concept

- Deduplication conceptually is straightforward
  - Hash chunks of data
  - compare hashes to find duplicates
  - Save fewer copies of the data
  - Keep a list of duplicates that reference the saved data
- VM / VDI / VSI typically have the best deduplication ratio

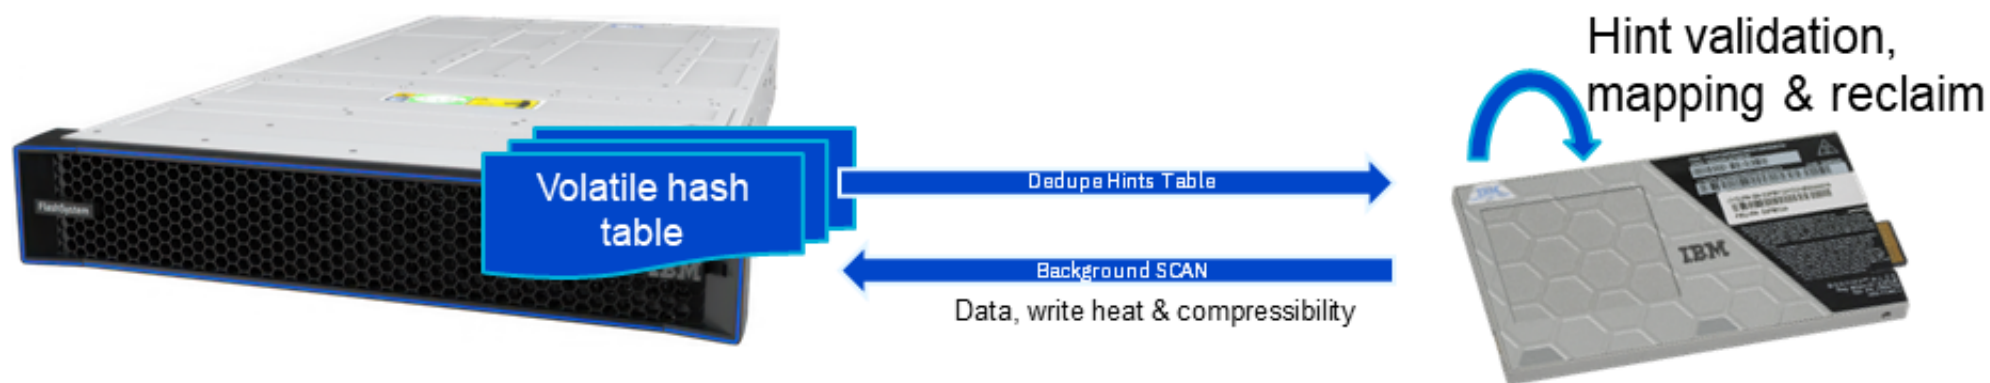
#### Challenge

- As System and SSD capacity grows, there is more and more data to analyze
- “Brute Force” dedupe requires a LOT of processing and a LOT of memory for metadata
  - Ex: 1PiB dataset, 32B SHA-2 Hash, 4KiB Data Chunk per Hash
    - Requires ~8TiB of memory to store the hashes
    - New host writes need to search through a LOT of hashes to dedupe inline
- Increased system cost and performance degradation quickly erode the value of deduplication
  - Vendors look to techniques of paging, block merging, and hash scaling for efficiency and simply dropping dedupe to maintain performance



## Deduplication in FCM 5<sup>th</sup> Generation

- Deduplication scanning, analysis, and mapping all done in the background to preserve host performance
  - Incoming data will still get the benefit of inline hardware compression, **no added Write Latency**
- Appliance controller uses **minimal processing** to perform dedupe scan and analysis
  - Controller only uses ~2.5GiB DRAM
- All Dedupe Mapping information is contained in existing FCM Logical to Physical Tables
- Reads to pages that have been deduped are done in-line by the FCM, **no added Read Latency**
- **No workload restrictions**
- **No capacity restrictions**
- **Unique to IBM: Even as capacity grows, no performance implications**



## FCM Assisted Deduplication Detailed Notes

---



### *Continual processing:*

- System will scan only 1 FCM per node at a time, sequentially.
- If something is degraded or non-redundant - processing will pause or abort :

### *Pause :*

- When one node goes away (any non committed scanning work the missing node was doing is discarded)

### *Resume :*

- One minute after the node re-joins the system

### *Abort :*

- If the DRAID array is unhealthy (i.e. not fully redundant)
- If the DRAID is expanding
- If software upgrade has started - re-assessed two minutes after upgrade is committed

No scanning interim data (within the nodes) is held in NV memory - any non committed work, or scanning results will be discarded and the process will begin again from scratch as appropriate.

## Deduplication Status for 9.1.2

---

- Only scanning drives and fingerprint database processing
- No dedupe work will be done
- No dedupe savings – size on compression only
- Will still have potential and achieved field populated in lsadinternals





Deduplication reclaim will be enabled in a future software release

## New FlashSystems in 2026

---



## Spec-sheet Summary - FlashSystem 2026

	FlashSystem grid			
	 FlashSystem C200	 FlashSystem 5600	 FlashSystem 7600	 FlashSystem 9600
Type of media	NVMe	NVMe	NVMe	NVMe
Rack Units	2U	1U	2U	2U
Drive slots	24 drive slots	12 drive slots	32 drive slots	32 drive slots
Physical / Effective	1.1 PB / 2.3 PBe	403 TB / 2.42 PBe	1.22 PB / 7.29 PBe	2.43 PBU / 11.88 PBe
Achieved with	FCM4 – 46 TB (XL)	FCM5 – 52.8 TB (XL)	FCM5 – 52.8 TB (XL)	FCM 5 – 105.6 TB (XXL)
CPU info per system	4x10 core	2x12 core Ice Lake D	2x16 core AMD Epyc	2x48 core AMD Epyc
CPU cores per system	40 cores	24 cores	32 cores	96 cores
High-speed bus	PCIe Gen 4	PCIe Gen 4	PCIe Gen 5	PCIe Gen 5
Memory per enclosure	256 GB	256 GB to 512 GB	768 GB to 1.5 TB	1.5 TB to 3.0 TB
Max Cache Hit IOPs	1.5M	2.6M	4.33M	6.37M
Max 70/30/50 IOPs 16KB	-	388K	700K	1.75M
Max Read Miss Bandwidth	23 GB/s	30 GB/s	60 GB/s	86 GB/s
Max HBA slot per system	4	4	8	8
Max FC ports/system	16 – 32Gb FC ports	16 - 32Gb FC ports 8 - 64Gb FC ports	32 - 32Gb FC ports 24 - 64Gb FC ports	32 - 32Gb FC ports 24 - 64Gb FC ports
Max Ethernet ports/system	8 - 10GbE on-board 8 – 25GbE ports	4 - 25GbE on-board 16 – 25GbE ports 8 – 40/100GbE ports	32 – 25GbE ports 16 – 40/100GbE ports	32 – 25GbE ports 16 – 40/100GbE ports

## FlashSystem 9600



Two 2T-E3 slots for potential future use-cases

32x 1T-E3 slots for FCM5s or for small capacity (1.9TB 3.84TB, 7.68TB) industry-standard commodity NVMe E3 SSDs

16x 2T-E3 compatibility from day1!

Two 2T-E3 slots for potential future use-cases

## 9600 has stunning density and sharable features



Unique

Capacity Point	Small	Medium	Large	XL	XXL
FCM5 Raw Capacity	6.6 TB	13.2 TB	26.4 TB	52.8 TB	105.8 TB
FCM5 Physical Capacity	5.6TB	11.2 TB	22.4 TB	44.8 TB	89.6 TB
FCM5 Logical Capacity	33.6 TB	67.2 TB	134.4 TB	268.8 TB	438.0 TB
x32 FCM5 Physical w/ RAID	151.9 TB	303.8 TB	607.6 TB	1215.2 TB	2430.4 TB
x32 FCM5 Logical w/ RAID	911.4 TB	1822.8 TB	3645.6 TB	7219.2 TB	11880.8 TB

## 4U FlashSystem 9500 vs 2U FlashSystem 9600

### 9500 Facts :

Rack Space : 4RU

Max capacity :  
Physical : **1.5PB**  
Logical : **4.7PB**

Per RU :  
**394.8 TB** physical  
**1.18 PB** logical



### 9600 Facts :

Rack Space : 2RU

Max capacity :  
Physical : **2.43PB**  
Logical : **11.8PB**

Per RU :  
**1,215 TB** physical  
**5.9 PB** logical



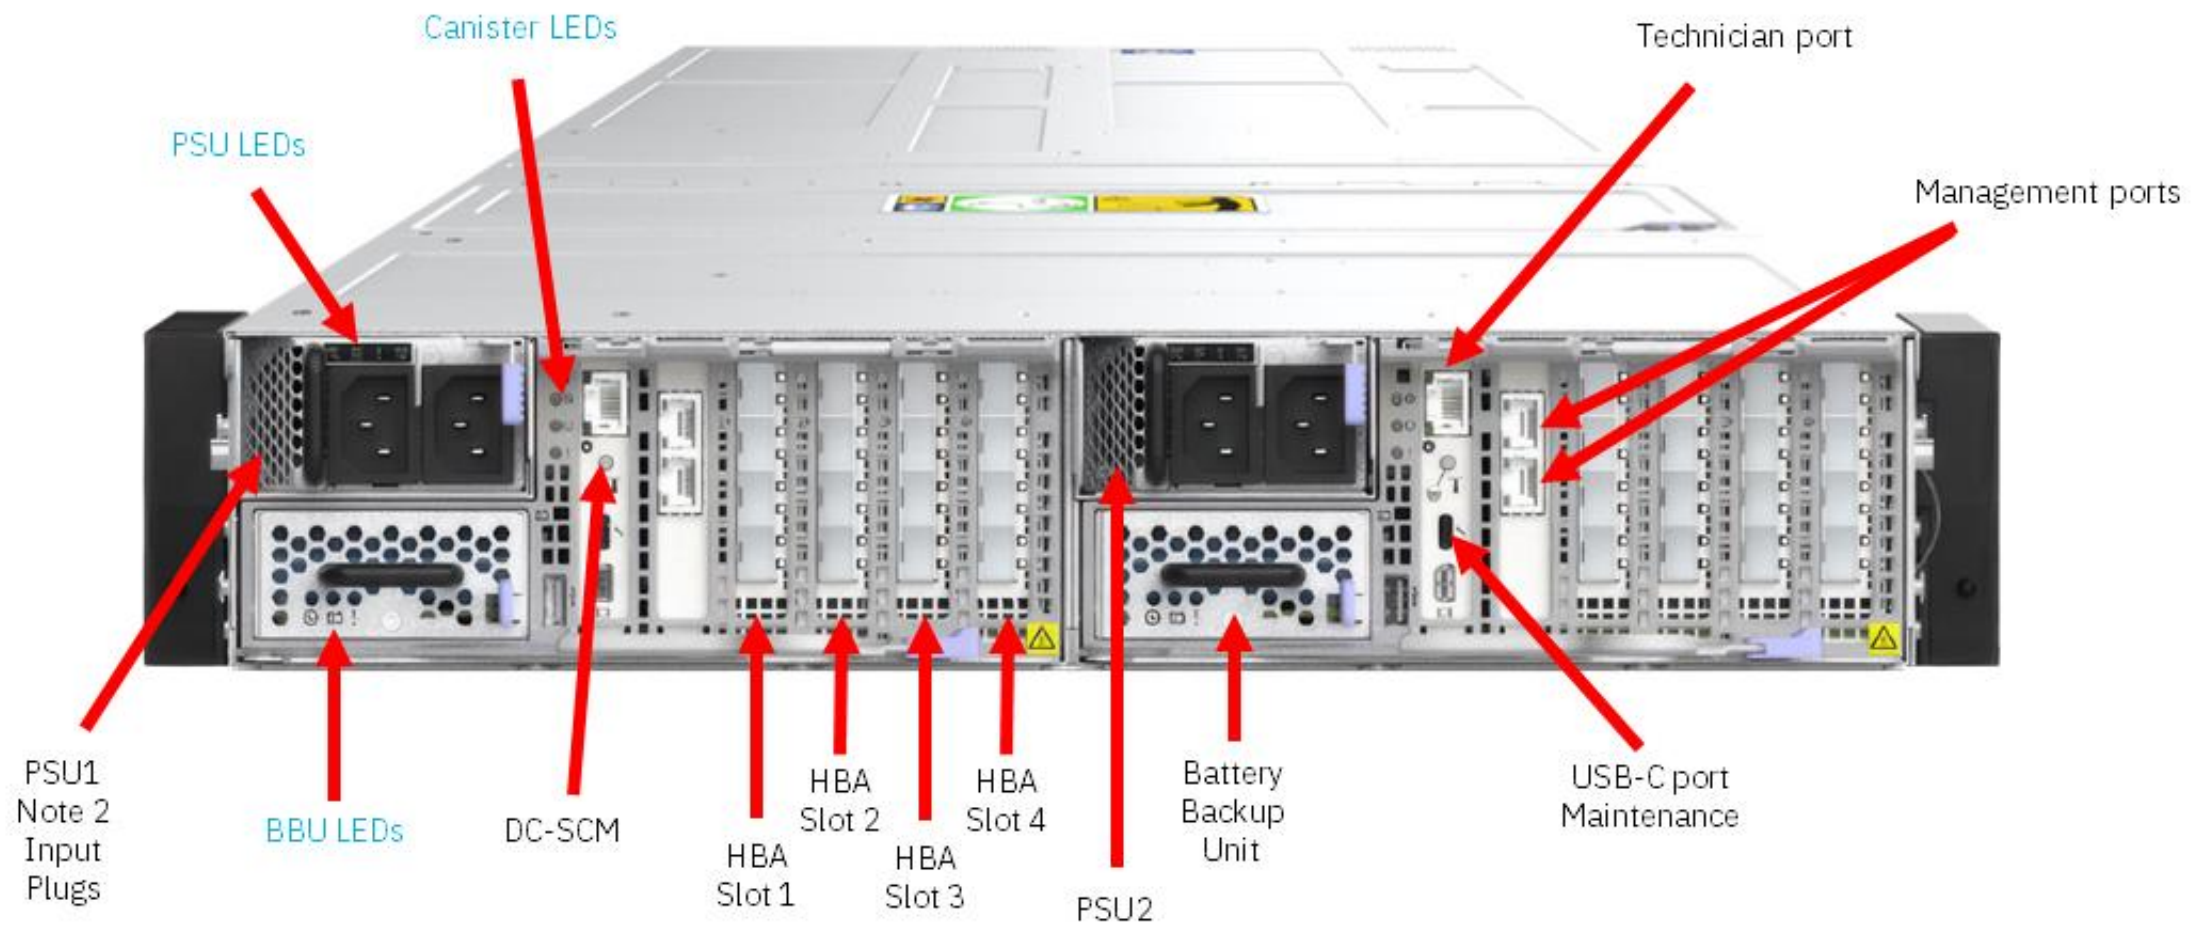
### Comparison :

Rack Space : **50% reduction**

Max capacity :  
Physical : **54% increase**  
Logical : **151% increase**

Per RU :  
Physical : **207% increase**  
Logical : **400% increase**

## FlashSystem 9600 - Rear View





## FlashSystem 9600 adapter configs

Slot	Adapter	Speed	Max Ports <sup>(SEP)</sup> per system	Connector <sup>(SEP)</sup> Type	Supported <sup>(SEP)</sup> protocol	Function
On-board (T)	DC-SCM	1Gb	2	RJ45	<ul style="list-style-type: none"> <li>Technician Port</li> </ul>	<ul style="list-style-type: none"> <li>Initial Configuration</li> <li>Recovery for locked accounts/missing password</li> <li>DA only – No switch attach</li> </ul>
On-board (M)	Mgmt. Ports	10Gb	4	RJ45	<ul style="list-style-type: none"> <li>Management</li> </ul>	Management Functions
1, 2, 3, 4	Fibre Channel	32Gb	32	SFP+	<ul style="list-style-type: none"> <li>SCSI</li> <li>NVMe/FC</li> </ul>	<ul style="list-style-type: none"> <li>Host I/O</li> <li>Replication and HA</li> <li>Migration</li> <li>Virtualized Storage Attach</li> </ul>
1, 2, 3	Fibre Channel	64Gb	24	SFP+	<ul style="list-style-type: none"> <li>SCSI</li> <li>NVMe/FC</li> </ul>	<ul style="list-style-type: none"> <li>Host I/O</li> <li>Replication and HA</li> <li>Migration</li> <li>Virtualized Storage Attach</li> </ul>
1, 2, 3, 4	Ethernet	10Gb 25Gb		SFP+ SFP28	<ul style="list-style-type: none"> <li>iSCSI</li> <li>NVMe/TCP</li> </ul>	<ul style="list-style-type: none"> <li>Host I/O</li> <li>Replication with TCP/IP</li> </ul>
1, 2, 3, 4	Ethernet	40 100Gb*	16	QSFP+ QSFP28	<ul style="list-style-type: none"> <li>iSCSI</li> <li>NVMe/TCP</li> </ul>	<ul style="list-style-type: none"> <li>Host I/O</li> <li>Replication with TCP/IP</li> <li>Short distance Replication with RDMA</li> <li>HA with RDMA (iWARP)</li> </ul>

\*Denotes oversubscription

## FlashSystem 7600



Two 2T-E3  
for future use

32x 1T-E3 slots  
*(Only supports one drive type/capacity per enclosure)*

FCM5 - all capacities up to **XL**  
OEM NVMe E3 SSD- small capacity (1.9TB 3.84TB, 7.68TB)

*Future 16x 2T-E3 ready!*

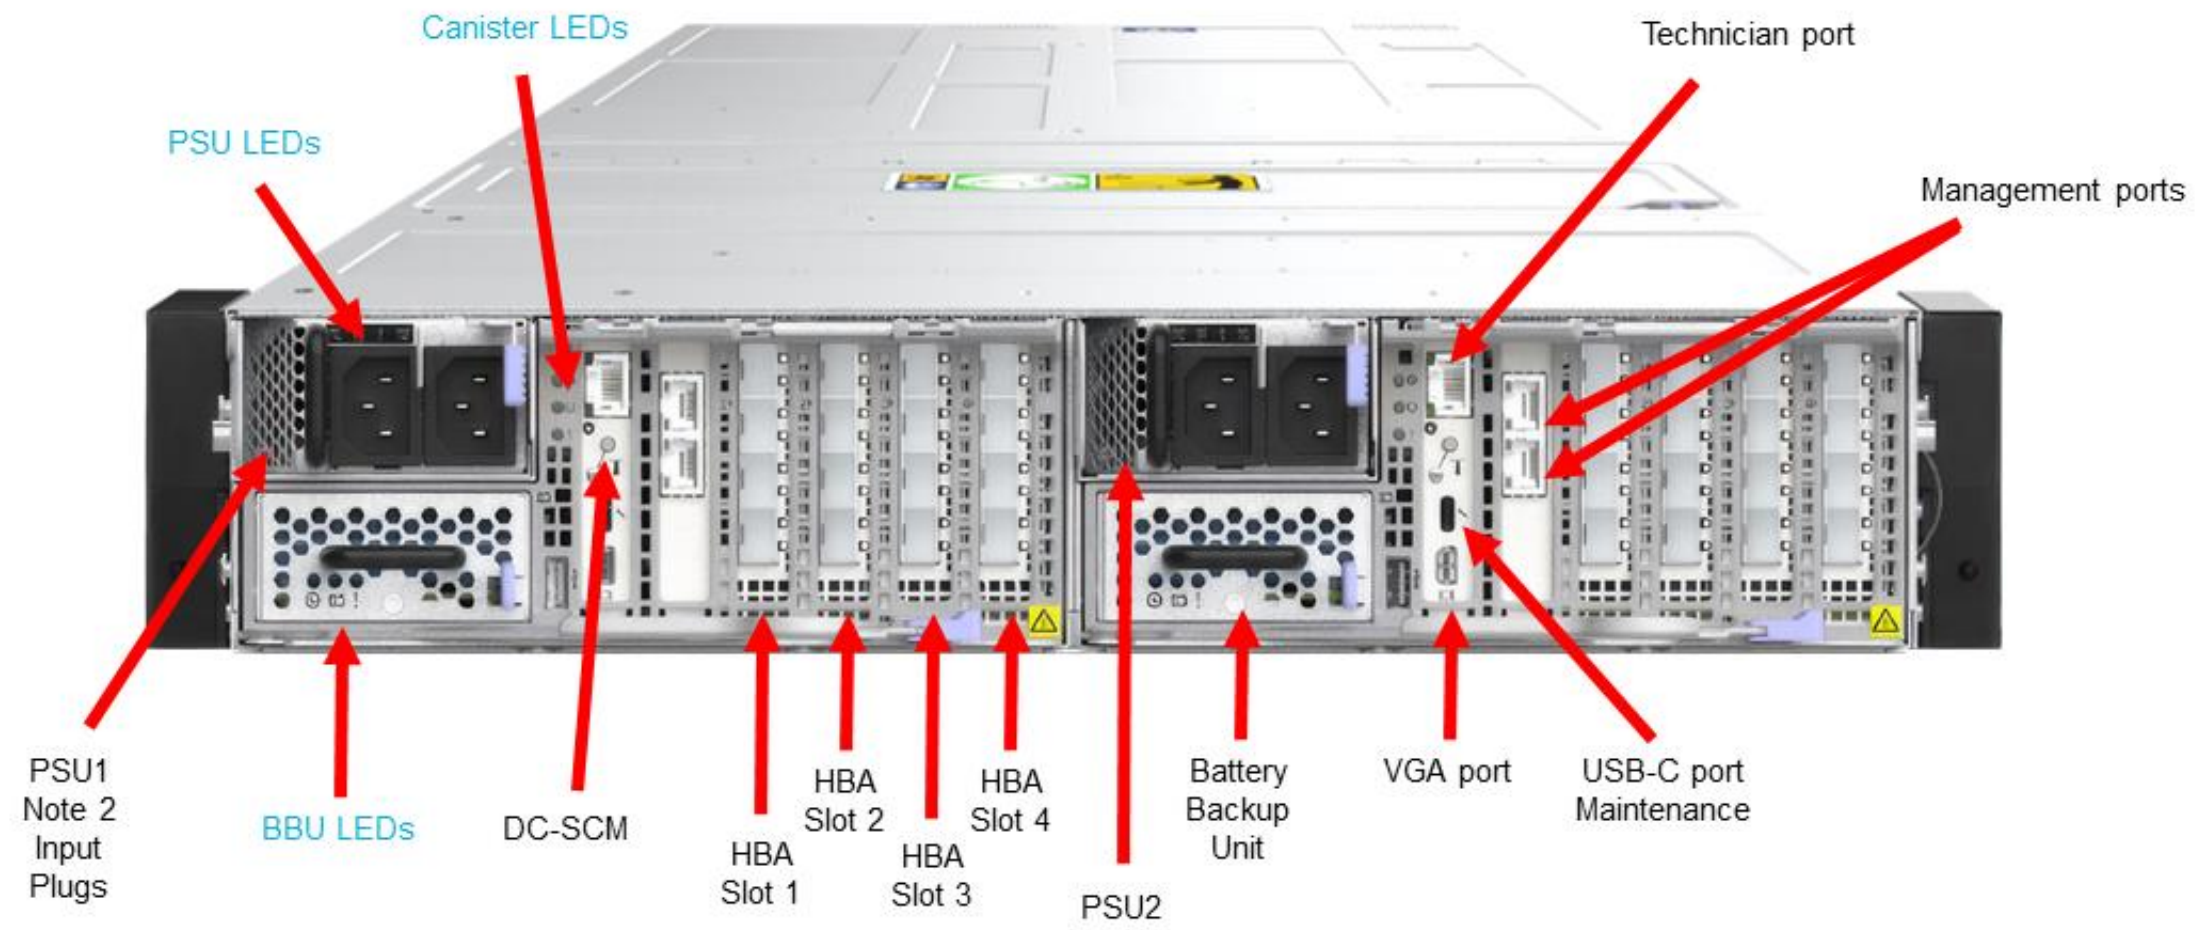
Two 2T-E3  
for future use

## FlashSystem 7600 - FCM5 Capacity Examples



Capacity Points in TB	Small	Medium	Large	XL
x1 FCM5 RAW	6.6	13.2	26.4	52.8
x1 FCM5 physical	5.6	11.2	22.4	44.8
x1 FCM5 effective	33.6	67.2	134.4	268.8
x32 FCM5 physical w/RAID	151.9	303.8	607.6	1215.2
x32 FCM5 effective w/ RAID	911.4	1822.8	3645.6	7291.2

## FlashSystem 7600 - Rear View





## FlashSystem 7600 adapter configs

Slot	Adapter	Speed	Max Ports <sup>(SEP)</sup> per system	Connector <sup>(SEP)</sup> Type	Supported <sup>(SEP)</sup> protocol	Function
On-board (T)	DC-SCM	1Gb	2	RJ45	<ul style="list-style-type: none"> <li>Technician Port</li> </ul>	<ul style="list-style-type: none"> <li>Initial Configuration</li> <li>Recovery for locked accounts/missing password</li> <li>DA only – No switch attach</li> </ul>
On-board (M)	Mgmt. Ports	10Gb	4	RJ45	<ul style="list-style-type: none"> <li>Management</li> </ul>	Management Functions
1, 2, 3, 4	Fibre Channel	32Gb	32	SFP+	<ul style="list-style-type: none"> <li>SCSI</li> <li>NVMe/FC</li> </ul>	<ul style="list-style-type: none"> <li>Host I/O</li> <li>Replication and HA</li> <li>Migration</li> <li>Virtualized Storage Attach</li> </ul>
1, 2, 3	Fibre Channel	64Gb	24	SFP+	<ul style="list-style-type: none"> <li>SCSI</li> <li>NVMe/FC</li> </ul>	<ul style="list-style-type: none"> <li>Host I/O</li> <li>Replication and HA</li> <li>Migration</li> <li>Virtualized Storage Attach</li> </ul>
1, 2, 3, 4	Ethernet	10Gb 25Gb		SFP+ SFP28	<ul style="list-style-type: none"> <li>iSCSI</li> <li>NVMe/TCP</li> </ul>	<ul style="list-style-type: none"> <li>Host I/O</li> <li>Replication with TCP/IP</li> </ul>
1, 2, 3, 4	Ethernet	40 100Gb*	16	QSFP+ QSFP28	<ul style="list-style-type: none"> <li>iSCSI</li> <li>NVMe/TCP</li> </ul>	<ul style="list-style-type: none"> <li>Host I/O</li> <li>Replication with TCP/IP</li> <li>Short distance Replication with RDMA</li> <li>HA with RDMA (iWARP)</li> </ul>

\*Denotes oversubscription

## FlashSystem x600 Ethernet HBAs

BYOC rules as previous - details in docs website

Quad-port  
25Gbps or  
10Gbps



Dual-port  
40Gbps or  
100Gbps



Platform	Adapter	Slot	iSCSI host attach	NVMeTCP host attach	Long distance Ethernet partnership over TCP	Short distance Ethernet partnership over RDMA
FlashSystem 5600	4x25G Ethernet	1/2	Y	Y	Y	N
	2x100G Ethernet	1/2	Y	Y	Y	Y
FlashSystem 7600	4x25G Ethernet	1/2/3/4	Y	Y	Y	N
FlashSystem 9600	2x100G Ethernet	1/2/3/4	Y	Y	Y	Y

### Configuration Restrictions

- NVMe/RDMA will not be supported on new platforms
- Long distance partnership over TCP will not support compression on new platforms  
*(Intel -> AMD means hardware accelerator for compression is not be available)*
- Asymmetric HighSpeed Ethernet partnership (100G-25G) is transitional only\*

### \* Asymmetric HSE Partnership Consideration Risks

- Asymmetric partnership (100G-25G) is transitional only
- Bandwidth over-subscription
- Likelihood of congestion under sustained or bursty workloads
- Effective max I/O bandwidth contained by minimum port-to-port bandwidth.
- Degraded I/O performance and HA loss
- Host side throttling during transition
- Flow control

## FlashSystem 5600

New EDSFF enclosure, the same node canisters as existing 5300



15 drive slots with  
3 currently reserved  
*designated by [X]*

10/25Gb I/O port  
Optionally: secondary system IP  
Internally: port 3

10/25Gb I/O port  
Internally: port 2

USB port  
Debug port

Canister  
LEDs

Technician port

Dedicated 1GbE Mgmt. port  
Internally: port 1  
Used for system IP and service  
IP

12x 1T-E3 slots  
*(Only supports one drive type/capacity per enclosure)*

FCM5 - all capacities up to **XL**  
OEM NVMe E3 SSD- small capacity (1.9TB  
3.84TB, 7.68TB)

## FlashSystem 5600 - Capacity Examples

*The existing node canister has only 12 drive connections, hence the three unused slots...*



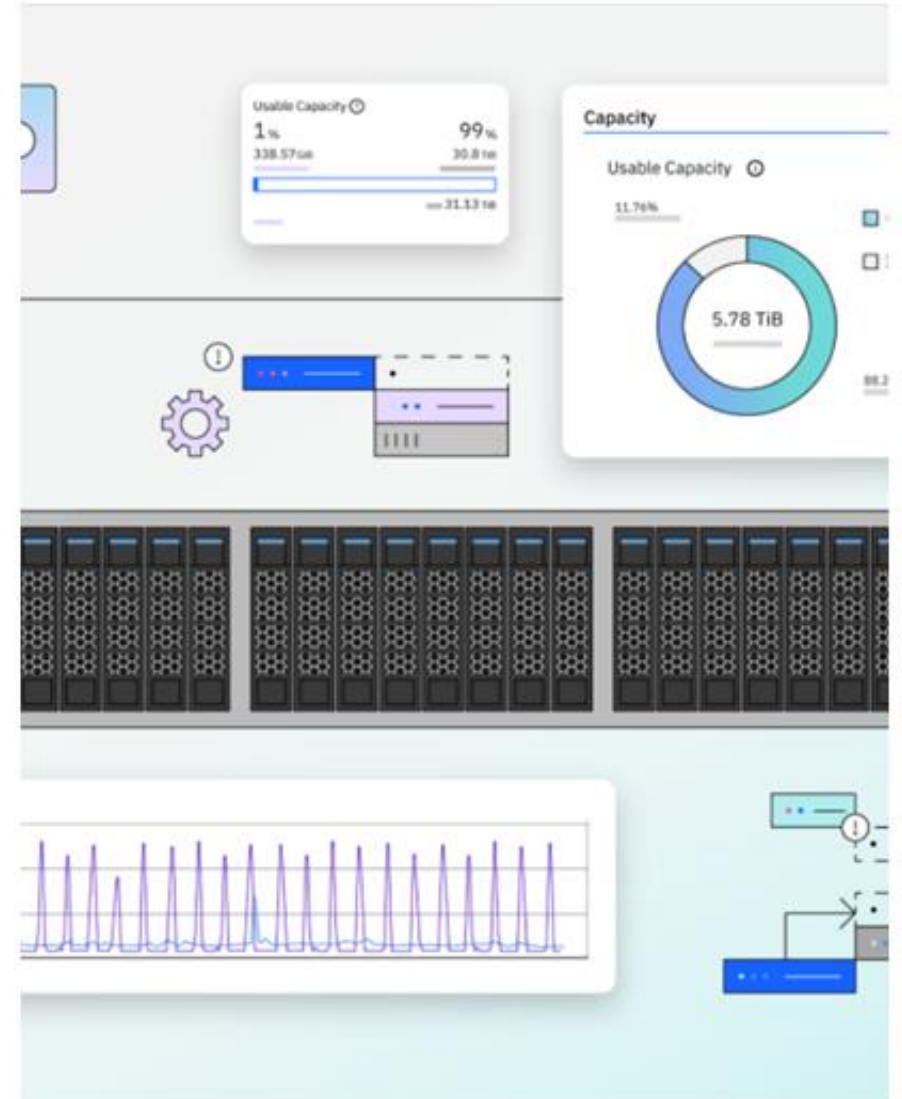
Capacity Points in TB	Small	Medium	Large	XL
x1 FCM5 RAW	6.6	13.2	26.4	52.8
x1 FCM5 physical	5.6	11.2	22.4	44.8
x1 FCM5 effective	33.6	67.2	134.4	268.8
x12 FCM5 physical w/RAID	50.4	100.8	201.6	403.2
x12 FCM5 effective w/ RAID	302.4	604.8	1209.6	2419.2

## Performance Highlights - 2026 FlashSystem

All measurements taken at a maximum of 1ms latency

Metric	Parameters	FS5300	FS5600	Gain	FS7300	FS7600	Gain	FS9500	FS9600	Gain	Gain per RU
Interface		16x 32Gb FC	16x 32Gb FC		12x 32Gb FC	12x 32Gb FC		12x 64Gb FC	12x 64Gb FC		
Drive Count	Type	12x FCM4-S	12x FCM5-S		24x FCM3-S	32x FCM5-S		40x FCM4-S	32x FCM5-L	-8	8
Max Hit IOPS	4KB Read Cache Hit	2.5M	2.6M	4%	3.3M	4.3M	30%	4.2M	6.3M	50%	214%
Max Miss IOPS	4KB Rdm Read Miss <i>from Flash</i>	650K	700K	8%	950K	1.2M	26%	2.4M	3.0M	25%	150%
OLTP IOPs	Measured at 1ms	388K	388K	-	530K	700K	32%	1.5M	1.75M	17%	133%
Interface		16x 32Gb FC	16x 32Gb FC		12x 32Gb FC	12x 32Gb FC		24x 32Gb FC*	12x 64Gb FC		
Drive Count	Type	12x FCM4-S	12x FCM5-S		24x FCM3-S	32x FCM5-S		48x FCM4-L	32x FCM5-L	-16	8
Max Bandwidth	256KB Rdm Read Miss <i>from Flash</i>	30 GB/s	30 GB/s	-	43GB/s	55GB/s	28%	97GB/s	86GB/s	-11%	77%
	Per rack U	30 GB/s	30 GB/s	-	21.5GB/s	27.5GB/s	28%	24.25GB/s	43GB/s	77%	

## FlashSystem.ai



---

# What if you could have a conversation with your Storage?

*...and your Storage could  
take action for you?*

FlashSystem grid

Dashboard

Monitoring

Pools

Volumes

Hosts

Policies

Access

Settings

Storage partitions

## Performance

Granularity

seconds

View historical data with **IBM Storage Insights**

Node Comparison

System

Latency

138  $\mu$ s

read

write

41  $\mu$ s 138  $\mu$ s

Bandwidth

1 MBps

read

write

0 MBps 1 MBps

IOPS

65 IOPS

read

write

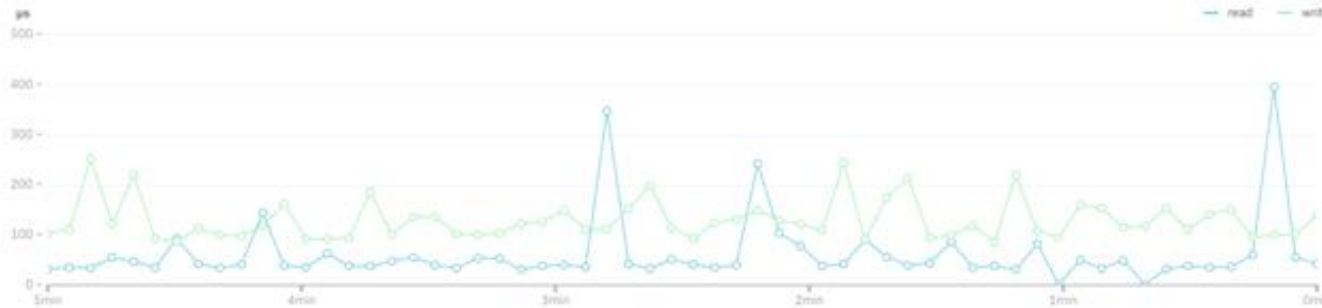
0 IOPS 65 IOPS

CPU Utilization

9 %

Power Consumption

442 Watts



## Capacity

Plan Capacity with **IBM Storage Insights**

123 %

### Physical capacity

21.39 TiB Total physical capacity



Physical used

1.94 TiB

Physical available

19.45 TiB

### Logical capacity

99.01 TiB Total logical capacity



Logical used

13.45 TiB

Logical available

85.56 TiB

### Capacity savings

4.2:1 Data reduction

6.22 TiB

1.6:1 Thin-provisioning

4.92 TiB

40.2:1 Total savings

11.14 TiB

1/3

## System Health

Ransomware Threat Detection: ✔

Version: 9.1.2.0 (build 191.31.2601021203000)

Cluster ID: 00000204A00102AC

### Hardware Components

All Online

Expand

### Logical Components

All Online

Expand

### Connectivity Components

All Online

Expand



The screenshot displays the IBM FlashSystem management interface. On the left is a navigation sidebar with options: FlashSystem grid, Dashboard, Monitoring, Pools, Volumes, Hosts, Policies, Access, Settings, and Storage partitions. The main content area is divided into several sections:

- Performance:** Shows a summary table and a line graph. The summary table includes:

Metric	read	write
Latency	138 $\mu$ s	41 $\mu$ s / 138 $\mu$ s
Bandwidth	1 MBps	0 MBps / 1 MBps
IOPS	65 IOPS	0 IOPS / 65 IOPS
CPU Utilization	9 %	
Power Consumption	442 Watts	

The graph shows read and write performance over a 5-minute period.
- Capacity:** Displays physical and logical capacity usage. Physical capacity: 21.39 TiB Total, 1.94 TiB used, 19.45 TiB available. Logical capacity: 99.01 TiB Total, 13.45 TiB used, 85.56 TiB available. Capacity savings: 4.2:1 (6.22 TiB), 40.2:1 (11.14 TiB).
- System Health:** Shows Hardware Components and Logical Components, both with "All Online" status and "Expand" buttons. Ransomware Threat Detection is active, and the version is 9.1.2.0.

An AI assistant overlay is positioned on the right side of the dashboard. It is titled "FlashSystem.ai" and shows a chat window with the following content:

You 11:51 AM  
Hello! what are your capabilities?

FlashSystem.ai Agent  
**IBMFlashSystemAssistant – What I Can Do**  
I'm a specialized virtual administrator for IBM Storage FlashSystem environments. My focus is on managing every aspect of the FlashSystem stack, using a set of dedicated worker agents and system tools to ensure actions are accurate, up-to-date, and safe.

**Core Capabilities**

Area	What I Can Do
Volumes	List, create, r

Type something...

IBM FlashSystem 7200 | chicagobear-c | Dashboard

FlashSystem.ai

You 15:50

What can you do?

FlashSystem.ai agent 15:50

I'm a virtual assistant that can manage IBM FlashSystem storage for you. Using the built-in workers, I can:

- **Volumes** – list, create, rename, resize, move between volume groups, and show storage-pool details.
- **Volume groups** – list, create, rename, and adjust replication or snapshot policies.
- **Snapshots** – create snapshots, list volume- or group-level snapshots, and work with snapshot policies.
- **Replication** – list replication policies and answer questions about them.
- **Partitions** – list partitions and provide information about them.
- **Hosts** – list hosts, create FC/iSCSI/NVMe hosts, modify host settings, manage port-sets, and map volumes to hosts.
- **Host clusters** – list, create, modify clusters and map volumes to clusters.
- **System health** – retrieve health status and answer health-related queries.

In short, I can perform most day-to-day administrative tasks on a FlashSystem—provisioning storage, configuring protection, and checking the system's health—while keeping the interaction simple and conversational. Let me know what you'd like to do!

Type something...

## Talk to your storage where you have an admin that never sleeps

---

Three fundamentals:

Human in the loop

- Human must initiate and authorize the action via natural language

Actions are on-prem

- One way communication to pull relevant storage information from the cloud

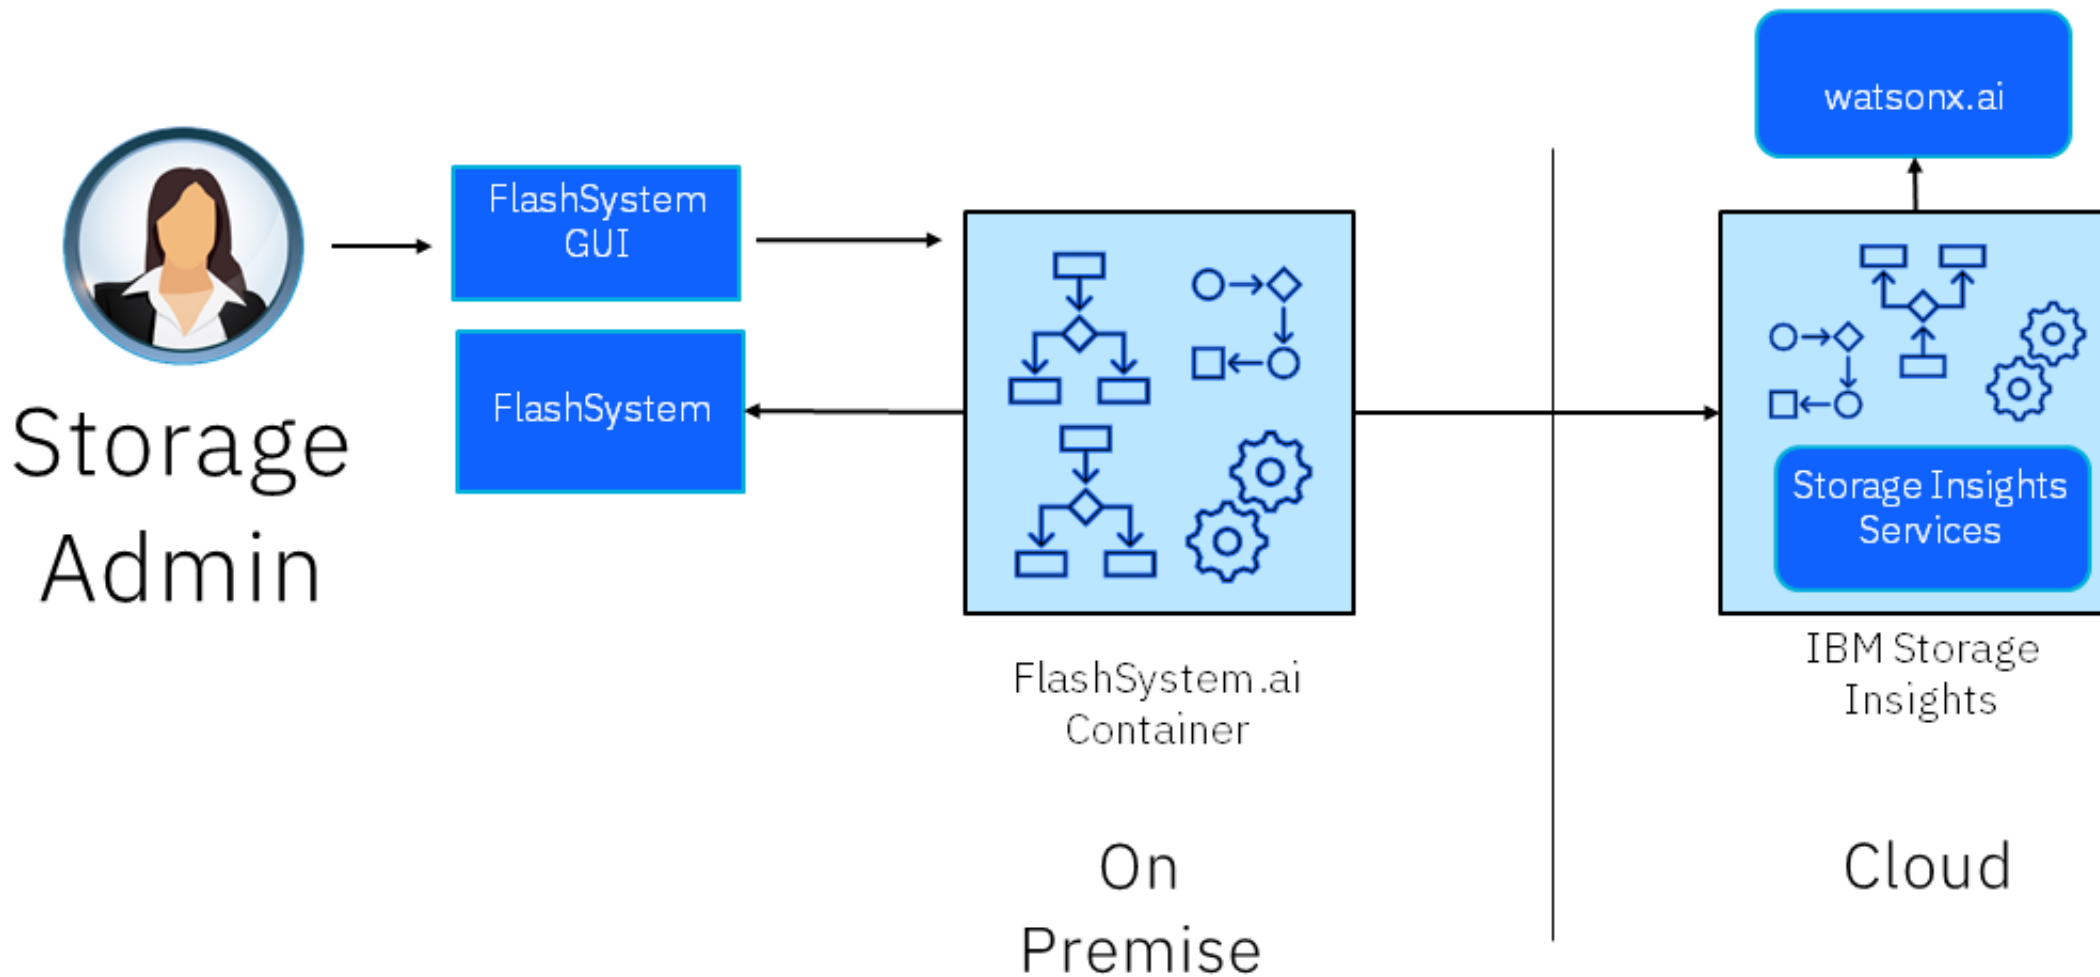
FlashSystem.ai never sleeps

- ... so you can!

*FlashSystem.ai is just part of the FlashSystem and it's powered by watsonx.ai*

## FlashSystem.ai Architecture

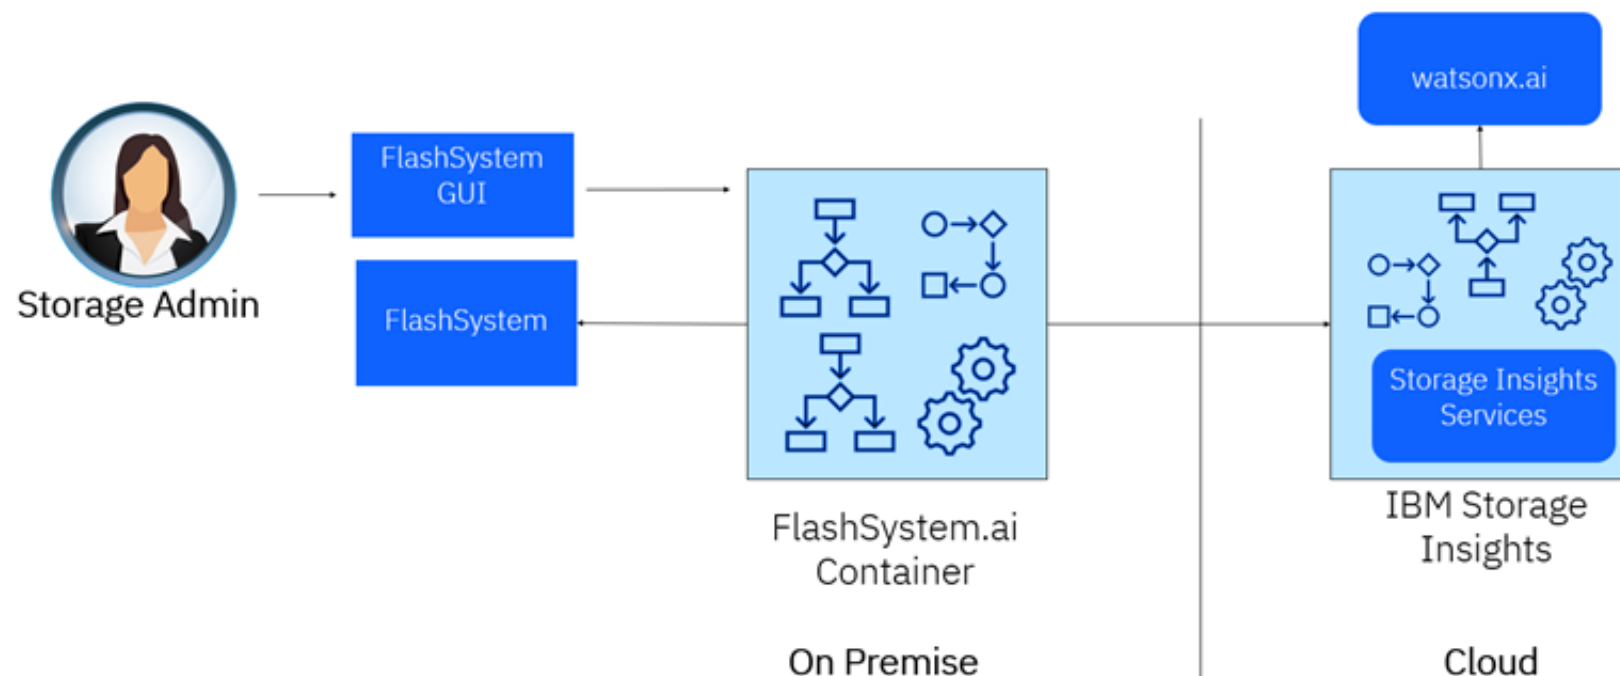
powered by **watsonx.ai**



*The FlashSystem.ai container provides a set of services that can only be used to interact with the FlashSystems and your IBM Storage Insights instance*

## FlashSystem.ai Self-hosted AI container

- A container hosted in the client's compute provides the AI capabilities for a FlashSystem or grid
- No on-prem GPUs needed or watsonx.ai subscription needed
- FlashSystem must be able to reach the FlashSystem.ai container and the AI container must be able to access the REST API of the FlashSystem / SVC
- The container must be able to contact Storage Insights
- Administrator user account for the AI needed on each system, to make any changes, but can have lesser rights.
- Container needs registering with each FlashSystem / SVC
- All actions taken on the FlashSystems are on prem ONLY



## Initial FlashSystem.ai Agentic capabilities

It is like benefiting from a storage admin who never sleeps.

**Outline;** inform, curate your information

- Show detailed system overview
- List volumes
- List volume groups
- List snapshot policies
- List replication policies
- List hosts
- List snapshots
- List storage pools
- List partitions and portsets
- List host-volume mappings

Zero Click Provisioning. You can ask...

**Create;** act for you – respond to request with ‘human in the loop’

- Create / Add volume
- Create / Add host
- Create / Add volume and attach to host
- Create / Add volume group with snapshot policy
- Create / Add volume group with snapshot policy safeguarded
- Create a snapshot

**Update;** act for you – respond to request with ‘human in the loop’

- Rename/resize volumes, change volume groups (Can only increase in size can't reduce)
- Modify volume group policies
- Host cluster management (create, list, map volumes)

## IBM FlashSystem.ai

powered by **watsonx.ai**

### What you can't do with FlashSystem.ai

- Delete any object on the system
- Resize any object on the system to be smaller than it is.
- Carry out any configuration change without human in the loop approval
- Does not work with any FlashSystem that is below 9.1.0
- Does not work with the 5015/5045 family

Why – remember this is based on ‘cause no harm’ first principal

- FlashSystem.ai cannot perform any actions which can affect the availability of data, such as deleting system objects or performing any security-related tasks.

### Using FlashSystem.ai

- Administrator account to take action
- Will support IBM FlashSystem 5x00 / 7x00 / 9x00 / SVC systems at the appropriate firmware level(s)
- At least 1 FlashSystem/SVC must be at 9.1.2
  - That FlashSystem hosts the FlashSystem.ai interface
  - Additional FlashSystem/SVCs can be at 9.1.0 and above (they can have Administrator or lower rights such as Monitor if no action is expected)
  - Participating FlashSystems can be stand-alone or part of an FlashSystem grid.
- There is a simplified experience for clients will be using FlashSystem.ai with a FlashSystem grid, as they are already a set of established trusted systems
- watsonx.ai subscription is not required
- *Note: Storage Insights Pro is not a pre-requisite for FlashSystem.ai, but Storage Insights connection is a requirement*



Embedded FlashSystem.ai discussion window

## IBM FlashSystem.ai working with Storage Insights

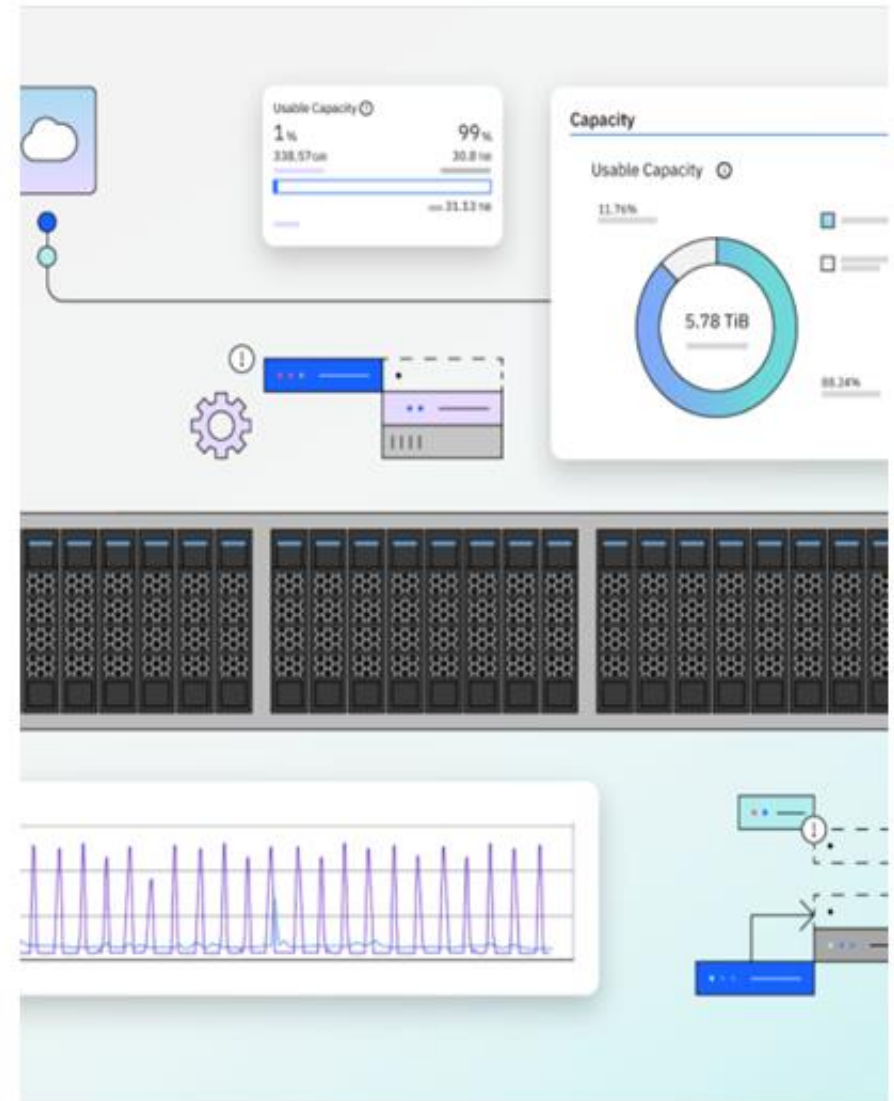
### Storage Insights integration

- Storage Insights connection is a requirement for IBM FlashSystem.ai
- Storage Insights will have a services infrastructure to work with FlashSystem.ai
- To integrate Storage Insights with FlashSystem.ai – a Storage Insights Tenant ID and REST API Key are required
  - Cloud Call home is a requirement for the participating FlashSystems
- Note: Storage Insights Pro is not a prerequisite for FlashSystem.ai, however will provide additional capabilities*

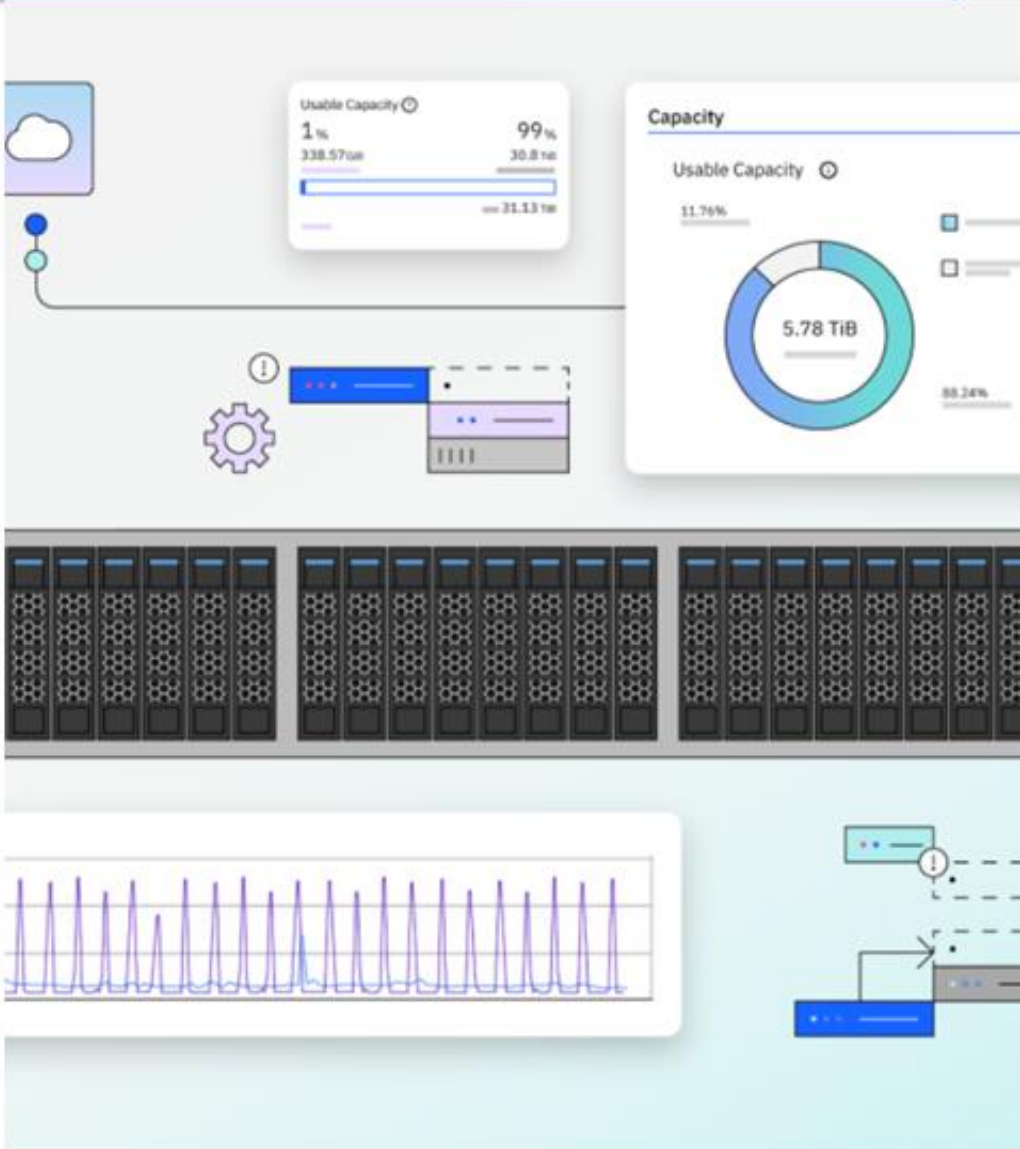


FlashSystem.ai	IBM Storage Insights (Foundation / Free)	IBM Storage Insights Pro
Natural language management	Yes	Yes
Requests per month	Limited	Unlimited
Autonomous Operations	No	Yes
Proactive actionable insights	No	Yes

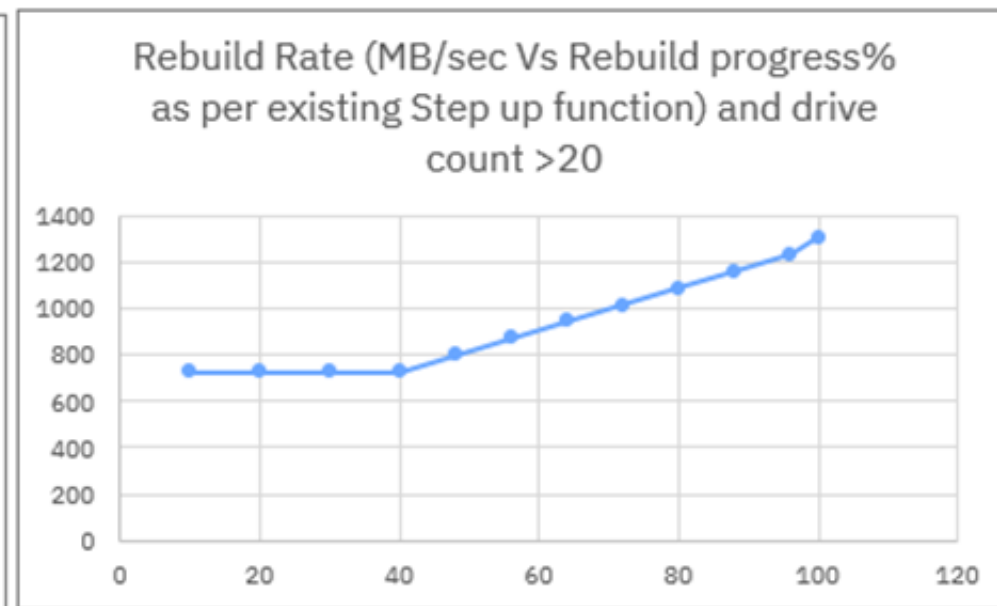
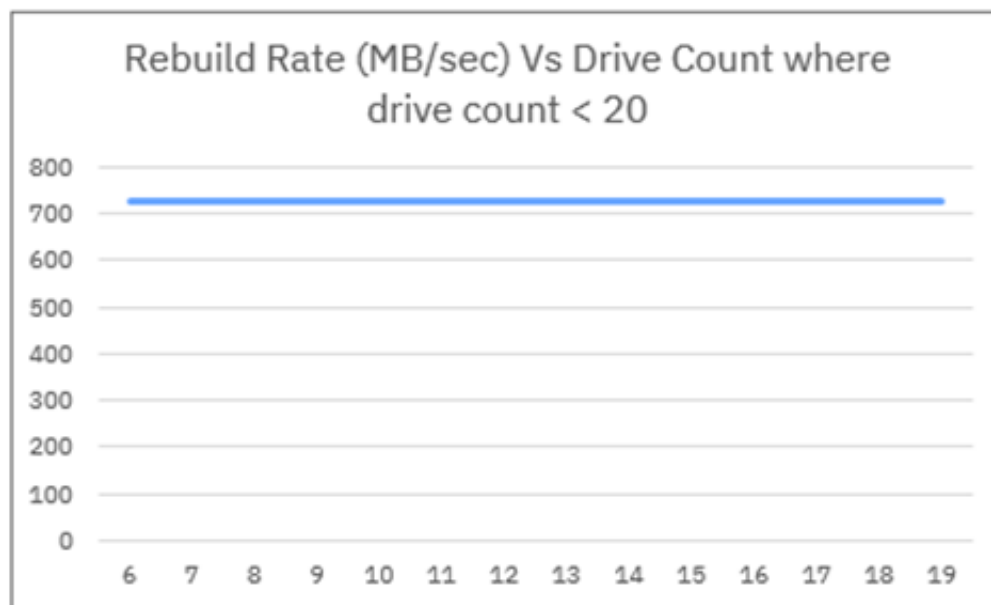
## Software Updates



# DRAID Support on New Platforms



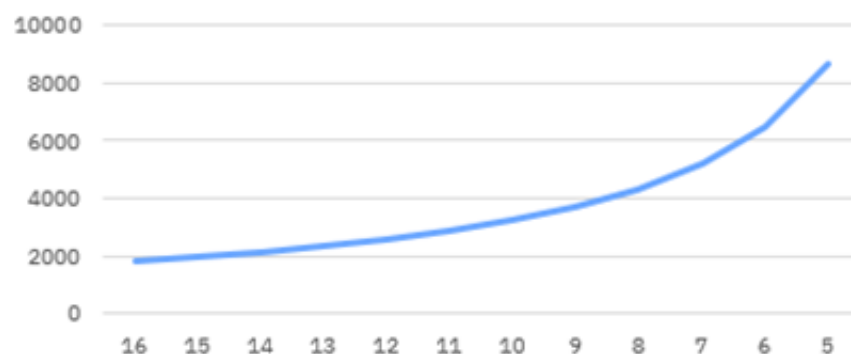
## Previous Rebuild Rate for L/XL FCM 3/4 Drives



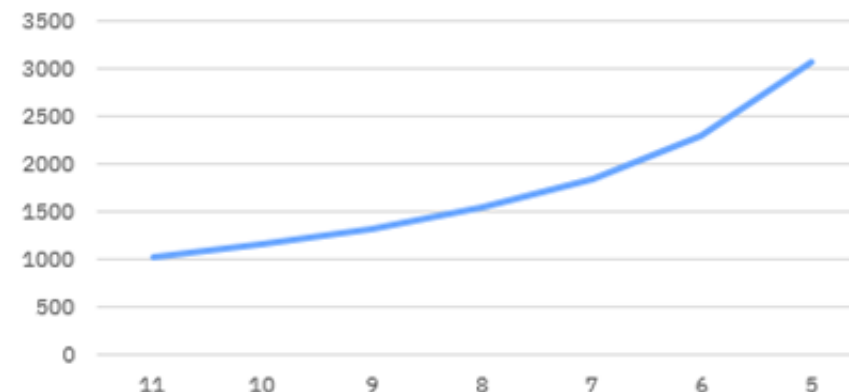
- Step up function works on a base static rebuild rate of 728 MB/sec
- Step up function does not consider effect of stripe width and drive count on drive utilization%
- As such drive can be underutilized or over utilized for certain type of DRAID geometry

## Geometry-based Rebuild Rate for FCM 3/4/5 L/XL/XXL Drives (Code Level 9.1.2+)

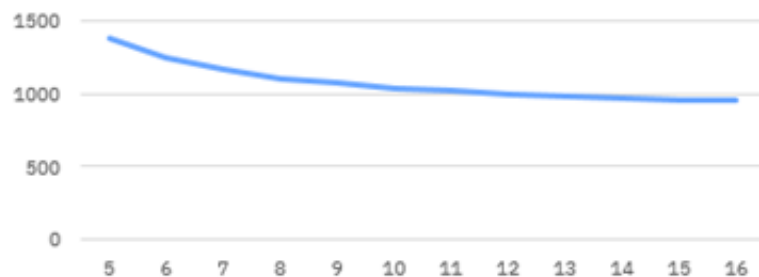
Rebuild Rate Vs Stripe Width for Drive Count  
32(FS9600/FS7600) with 25% drive  
utilization



Rebuild Rate Vs Stripe Width for Drive Count  
12(FS5600) with 25% drive utilization



Rebuild Rate Vs Stripe Width Where Drive  
Count = Stripe Width + 1 with 25% drive  
utilization



rebuildRate =

$$\frac{\text{driveRebuildPercentage} \times (\text{totalDrives} - \text{failedDriveCount}) \times (\text{driveBandwidth})}{100 \times (\text{stripeWidth} - 2)}$$

driveRebuildPercentage = 25%

Benefit: optimized rebuild performance covering various configurations

## Other Rebuild Changes for L, XL, XXL

---

- Rebuild Rate

- Rebuild rate will be increased depending upon the host workload
  - For higher host workload increase will be lesser
  - For lower host workload increase will be higher



- Scrub Rate

- Default scrub period for FCM5e XL will be 4 weeks and for FCM5e XXL will 7 weeks.
- Above changes are done so that scrub consume drive read bandwidth around 3% only.
- Other drives will have scrub period equal to or below 2 weeks.

## FlashSystem DRAID array configurations and enhancements

FlashSystem 5600
<b>DRAID-1</b>
<ul style="list-style-type: none"> <li>Max drive size FCM 5-L, limit between 2 and 12 drives, SAS is 2-16 drives</li> </ul>
<b>DRAID-5</b>
<ul style="list-style-type: none"> <li>Max drive size FCM 5-L, limit between 4 and 8 drives</li> </ul>
<b>DRAID-6</b>
<ul style="list-style-type: none"> <li>Limit between 6 and 12 FCM 5 drives</li> </ul>
<ul style="list-style-type: none"> <li>Max distributed arrays - 8</li> </ul>
<ul style="list-style-type: none"> <li>Homogeneous drive type/capacity per enclosure</li> </ul>
<ul style="list-style-type: none"> <li>Min default extent size = 2GB</li> </ul>

FlashSystem 7600
<b>DRAID-1</b>
<ul style="list-style-type: none"> <li>NA</li> </ul>
<b>DRAID-5</b>
<ul style="list-style-type: none"> <li>NA</li> </ul>
<b>DRAID-6</b>
<ul style="list-style-type: none"> <li>Limit of between 6 and 32 drives per enclosure</li> </ul>
<ul style="list-style-type: none"> <li>Minimum of 9 FCM XL drives</li> </ul>
<ul style="list-style-type: none"> <li>Max distributed arrays - 4</li> </ul>
<ul style="list-style-type: none"> <li>Homogeneous drive type/capacity per enclosure</li> </ul>
<ul style="list-style-type: none"> <li>Min default extent size = 8GB</li> </ul>

FlashSystem 9600
<b>DRAID-1</b>
<ul style="list-style-type: none"> <li>NA</li> </ul>
<b>DRAID-5</b>
<ul style="list-style-type: none"> <li>NA</li> </ul>
<b>DRAID-6</b>
<ul style="list-style-type: none"> <li>Limit of between 6 and 32 drives per enclosure</li> </ul>
<ul style="list-style-type: none"> <li>Minimum of 9 FCM XL drives</li> <li>Minimum of 13 FCM XXL drives</li> </ul>
<ul style="list-style-type: none"> <li>Max distributed arrays - 4</li> </ul>
<ul style="list-style-type: none"> <li>Homogeneous drive type/capacity per enclosure</li> </ul>
<ul style="list-style-type: none"> <li>Min default extent size = 8GB</li> </ul>

### Fast Drive Rebuilds

IBM's extent awareness technology gives the IO stack knowledge of where the client has written data and uses this knowledge to dramatically reduce rebuild times. This is important because FCM5 drives have the ability to roughly double effective capacity and the impact to rebuild time is proportional, but now the array will only rebuild where client data has been written. Therefore, even extents allocated to volumes can be skipped if unwritten or unmapped.

## DRAID Extent Awareness

---

### What is it?

- Extent awareness gives the IO stack (without DRP) knowledge of where the client has written data and uses this knowledge to dramatically reduce rebuild times

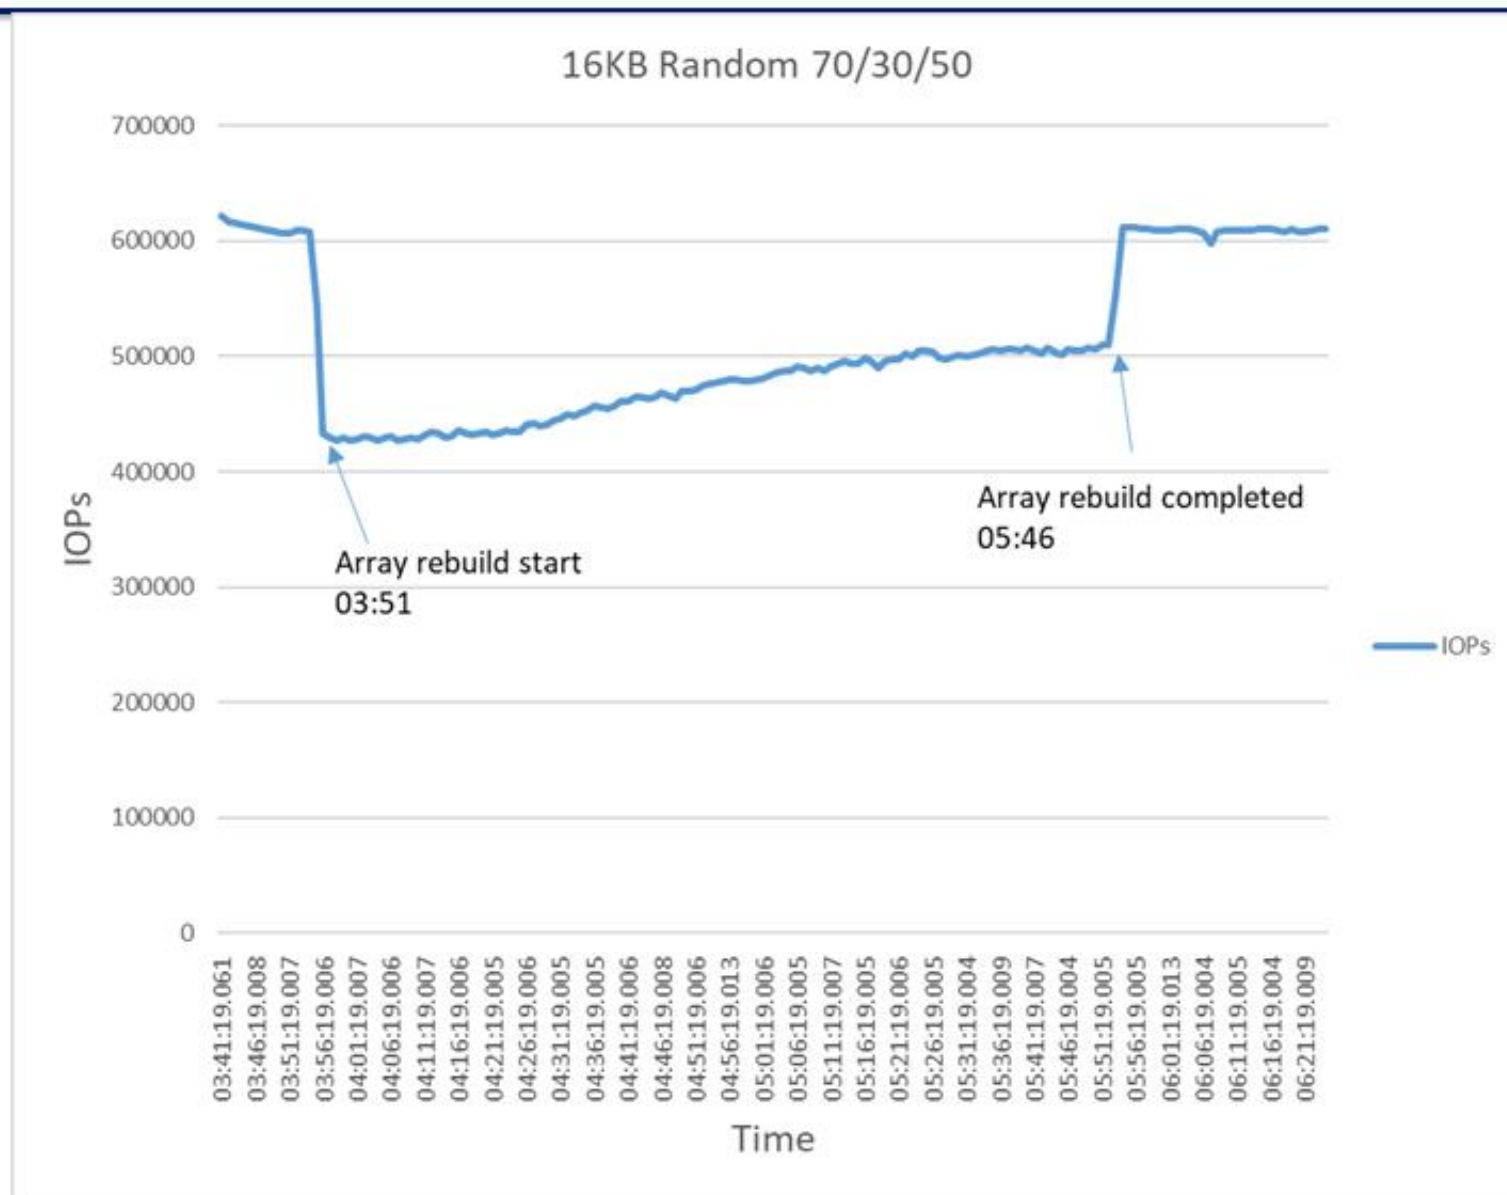
### Why it's needed?

- FCM5 drives roughly double effective capacity
  - Impact to rebuild times is proportional

### What you need to know

- FCM5 only initially
- Only extents with written data will be rebuilt
  - Rather than allocated extents today
    - This drastically reduces rebuild time for fully allocated volumes
- Uses DRAIDs hardened metadata to remember where extents are written
- May be used to optimize dedupe in the future

**200TiB+ Written  
2 hour rebuild  
vs 7 hours previously**



## Partition Split

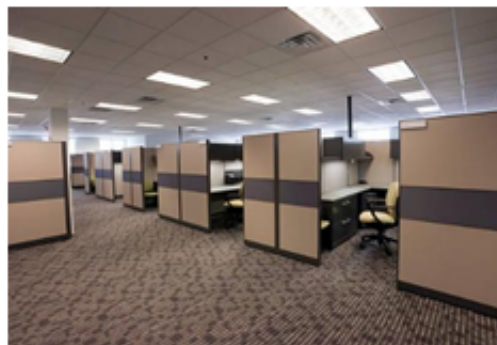
---

### What is it?

- Allows a user to split resources out of a partition into its own partition
  - Similar to a draft partition only it splits instead of bringing objects into a partition

### Why its needed

- To allow flexibility in PBHA and grid deployments to combine and split partitions are business needs evolve.



### What you need to know

- Split partition uses a draft partition that temporarily tracks resources until the draft partition is published
- HA is unaffected by a partition split
  - The resulting partition will be HA with no loss of sync
- DR policies must be dropped from volume groups to be split out and added back after the split is complete
  - Identical to merge partition today
- Some function are restricted while the draft partition exists
  - Partition migration
  - Merge partition



- Dashboard
- Monitoring
- Pools
- Volumes
- Hosts
- Policies
- Access
- Settings

**Storage partitions**

## All storage partitions

Total partitions: 1

### partition1

Hosts	Volumes
3	12
Volume groups	
3	

[View partition](#)



- Dashboard
- Monitoring
- Pools
- Volumes
- Hosts
- Policies
- Access
- Settings

**Storage partitions**

## All storage partitions

Total partitions: 1

### partition1

Hosts

3

Volume groups

3

[View partition](#)



Merge partition

**Split partition**

Delete partition

Rename partition

Manage IP address

Configure certificate

## Split partition ✕

The split partition action will let you select the existing resources from the source partition partition1 and create a new partition.

Enter the name of the new storage partition.

Storage partition name

Cancel

Create

## Select existing resources

Complete the steps to select existing resources for your new partition.

- Select volume groups
- Review volume groups
- Summary (vCenter partition)

### Select volume groups

Select from the following volume groups available in the source partition partition1 to add into the new partition splitpartition.

Search

Volume group	Volume Count	Host Count	Policies
<input checked="" type="checkbox"/> vol_rep_thick_0_group	5	2	
<input type="checkbox"/> vol_rep_thin_0_group	5	2	
<input type="checkbox"/> vg1	2	0	

Items per page: 10

1-3 of 3 items

1 of 1 page

## Select existing resources

Complete the steps to select existing resources for your new partition.

☑ Select volume groups

🔄 Review volume groups

🔄 Summary (vCenter partition)

### Review volume groups

Review volume groups selected for the new partition. Additional volume groups may be selected if required.

**i** **Additional volume groups were automatically added to the partition**  
 A partition must include all volumes mapped to any hosts included in the partition, therefore additional volume groups containing those volumes were added. Select the **Automatically included volume groups** tab to review.  
[More information](#)

**Selected volume groups (1)**

Automatically included volume groups (1)

🔍 Search

Select volume groups +

Volume group	Volume Count	Host Count	Policies
vol_rep_thick_0_group	5	2 mapped hosts	

Items per page: 25 ▾ 1–1 of 1 item 1 ▾ 1 of 1 page ◀ ▶

## Select existing resources

Complete the steps to select existing resources for your new partition.

- ✔ Select volume groups
- 🔍 Review volume groups
- 🔄 Summary (vCenter partition)

### Review volume groups

Review volume groups selected for the new partition. Additional volume groups may be selected if required.

Selected volume groups (1)

Automatically included volume groups (1)

🔍 Search

Volume group	Volume Count	Host Count	Policies
vol_rep_thin_0_group	5	2 mapped hosts	

Items per page: 25 ▾ 1-1 of 1 item 1 ▾ 1 of 1 page ◀ ▶

## Select existing resources

Complete the steps to select existing resources for your new partition.

✔ Select volume groups

🔄 Review volume groups

🔄 Summary (vCenter partition)

### Review volume groups

Review volume groups selected for the new partition. Additional volume groups may be selected if required.

Selected volume groups (1)

**Automatically included volume groups (1)**

🔍 Search

Volume group	Volume Count	Hosts
vol_rep_thin_0_group	5	2 m

Items per page: 25 ▾

1-1 of 1 item

vol\_rep\_thin\_0\_group

Close



#### Hosts

zooty-h0

zooty-h1

1-2 of 2 items



## Select existing resources

Complete the steps to select existing resources for your new partition.

- ✓ Select volume groups
- ✓ Review volume groups
- 📌 Summary (vCenter partition)

### Summary

#### ✓ Existing volume groups selected for splitpartition

Volume groups	Volume count
vol_rep_thick_0_group	5

#### ✓ Selected resources added to your partition!

Begin exploring your storage partition.

#### Review the remaining resources from the source partition partition1.

Volume groups: 1	Volume count: 2
------------------	-----------------

- Dashboard
- Monitoring
- Pools
- Volumes
- Hosts
- Policies
- Access
- Settings
- Storage partitions**

## All storage partitions

Total partitions: 2

### partition1

Hosts	Volumes
1	2

Volume groups

1

[View partition](#)

### splitpartition

Hosts	Volumes
2	10

Volume groups

2

[View partition](#)

## Automatic Drive Upgrades for all NVMe Drives

---

- All NVMe drives will be upgraded after drive formats and system upgrades
- Previously only FCM drives were upgraded

## Compression Over Native IP Replication

---

- Compression over native IP based replication will no longer be offered on the new platforms
  - The new platforms do not have the Intel QuickAssist chips/cards on them
  - Looking at adding compression in the future when it is re-designed
- No change in IPsec encryption feature (software supported)

## Port Level I/O Stats for iSCSI Hosts (API info for 3rd party collection)

---

### What is it

- Port level I/O stats will be generated for iSCSI hosts and displayed in Storage Insights
- These stats include:
  - hbt – bytes transmitted to hosts
  - hbr – bytes received from hosts
  - her – commands received from hosts

### What you need to know

- Collected in the Nn\* stats file
  - Existing file for node level stats
- Number of entries will depend on number of Ethernet ports available on the system
- type\_id field is the port number
  - Same as in the lspportethernet view
- lspportstats will show these port statistics in the CLI
- More stats could be added in the future
- These stats are available in Storage Insights

## Partition Migration Over High-Speed Ethernet (HSE)

---

### What is it

- Allows for partition migration over all Ethernet environments
- HSE (RDMA) support for short distance partnerships
- **Automated partnership creation**

### Why its needed

- Storage Virtualize needs to have a common set of migration features regardless of the transport used within the data center

### What you need to know

- Automated partition migration in:
  - All Ethernet based environments
  - All Fibre Channel based environments
  - Mixed FC and Ethernet based environments
- DR topology support using HSE (RDMA), long distance TCP/IP or FC
- UUID based partition migration
  - Compatible with older cluster-id based migrations
- Certificate exchange for secure communication
- iWARP only

## Changing FlashSystem grid Coordinator

---

### What is it

- Allows the FlashSystem grid coordinator to be changed without disrupting grid operations
- **Two types** of move:
  - **Coordinated** – The FlashSystem grid coordinator system is present
    - Can be done from the GUI
  - **Uncoordinated (Recovery)** – The FlashSystem grid coordinator is not present or down in some fashion
    - Must be done from the command line

### Why its needed

- Environments are dynamic and change
- Need the ability to change coordinator without destroying and re-creating the grid

### What you need to know

- Validations take place before to ensure smooth transition
- Handles unsupported/unavailable systems gracefully
- Self restoration/healing if new coordinator cannot make old coordinator a member
- If down, the old coordinator will make itself a member when it comes up
- CLI commands:
  - `managegrid –makecoordinator`
  - `managegrid –makecoordinator -removeinvalidmembers`

## Dedicated Quorum IPs

---

### What is it

- Adding the ability to have dedicated IPs for the quorum device
  - Currently, the service IPs are used

### Why its needed

- Customers have asked us for separate quorum IPs to reduce the security risk of multiple functions using the same ports for different purposes
- Adds flexibility to the system

### What you need to know

- New “Quorum” default portset added to the system
- mkip, chip, lsip, rmip commands are used to work with new quorum IPs
- To move to quorum IPs the client must:
  - Create a quorum IP for each node
  - Generate a new quorum app (it automatically uses the new IP addresses)
  - Deploy new quorum app
- Quorum app needs port 1260 open

## Fibre Channel Auto Zoning (Brocade)

---

### What is it

- Allows the storage array to automatically peer zone host to storage in a Brocade fabric

### Why its needed

- A large part of a storage administrators time is dedicated to zoning and zoning best practices
- This eliminates the need to spend the time to do these menial connectivity tasks

### What you need to know

- Implemented as auto-zoning portsets
- When defining a host, place it in the auto-zoning portset
  - Peer zoning is automatically implemented
- If automatic zoning is not desired, delete the auto-zoning portsets and use the default or normal host portsets
- Inline FlashSystem FC zoning can only zone other WWPNs to the storage
  - It cannot modify or delete existing zones
  - Feature can be turned off on the SAN ports
- Ansible playbook for Cisco zoning
  - Plans to tie into this in the future for auto-zoning on Cisco

## Simplified SAN management and host configuration

### 3 Roles, 1-hour Minimum

Server Admin



Gather OS details  
Gather HBA WWPNs

SAN Admin

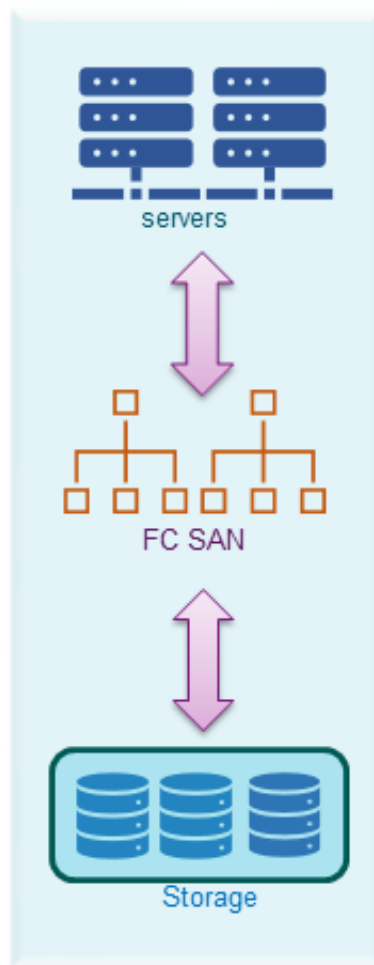
Create/manage SAN Zones  
Apply good pathing practices

Storage admin



Provision host configs  
Workload distribution across target ports  
Optimize host multi-pathing

Conventional SAN mgmt



### 1 Role, less than 2 min

FDMI & Fabric  
Discovery



Target Driven  
Zoning



Fabric monitoring &  
Zoning Engine



Single Click Host config & mgmt



Storage admin

- **Single Click host configuration:**  
Reduce host provisioning time from hours to a few minutes
- **Target driven auto-zone** creation with optimal multi-path count & distribution
- Use of **Peer zones** to avoid fabric noise
- Reduce maintenance window: Automatic zone updates on fabric maintenance & connectivity changes
- Supports **all operating systems**
- Seamless integration with Brocade fabrics
- Ansible playbook for Cisco fabrics

## UUIDs for All Partition Objects

### What it is

- All partition objects will have UUIDs that can be used to identify/reference those objects in CLI/RESTAPI commands
  - Previously, only partition volumes had UUIDs
  - Makes them unique across systems

### Why its needed

- From a server/automation perspective UUIDs are a common way to identify and work with objects on a system
- IBM is positioning partitions as the primary target for automated workflows
- Assists with migrations making sure objects have unique identifiers between systems

### What you need to know

- All objects compatible with partitions will have UUIDs generated
  - These UUIDs will also be generated on upgrade to 9.1.2
- RFC 4122 Format
- Partition objects that will have UUID include:
  - Volumes (done previously)
  - Hosts
  - Host Clusters
  - Volume Groups



## Re-raise Out of Space Alerts

---

### What is it

- Persistently logs low space alerts if no action is taken to correct the low space/out of space situation

### Why its necessary

- Previously a low space alert was raised, but if a customer acknowledged it no other alert was raised and the customer would generally run out of space because they had not dealt with the situation

### What you need to know

- The alert will persist until positive action has been taken
- Those actions can be
  - Free space
  - Change the warning threshold to be notified in the future



## Single Volume Type and Pool Type on New Systems

---

### What is it

- FCM5 systems will have a single pool type
  - No Data Reduction Pools (DRP) on these systems as the drives have deduplication
- FCM5 systems will default to a single volume type (Fully Allocated)
- Since the drives can be over provisioned to 6-1 there is little need for thin volumes at the logical layer

### Why its needed

- With FCM5 drives there is no need for DRP
- With FCM5 over provisioning the system is simplified by a single volume type
- Because FCM5 drives thin provision, compress and deduplicate, it makes it much simpler not having to manage this on a software level

### What you need to know

- By default, provisioning policies will not display and not be used as they are unnecessary
- Single volume type will automatically be deployed under the following conditions:
  - Only one drive class on the system
  - Drive type (FCM5) identifies as deduplicating
- Can be turned off in GUI/CLI at any given point in time
- No volume type/savings/transformation information will be presented on the screens
- Volumes still have a format process
  - In the future this may be modified
- Single pool type regardless of the drive type
  - Thin provisioning only for industry standard drives

## Patching Improvements

---

### What it is

- Patches have been redesigned to auto-install during software upgrades, whether to a new level or PTF release
- Patches can be viewed and subscribed to for auto-updating when new versions of the patch are released

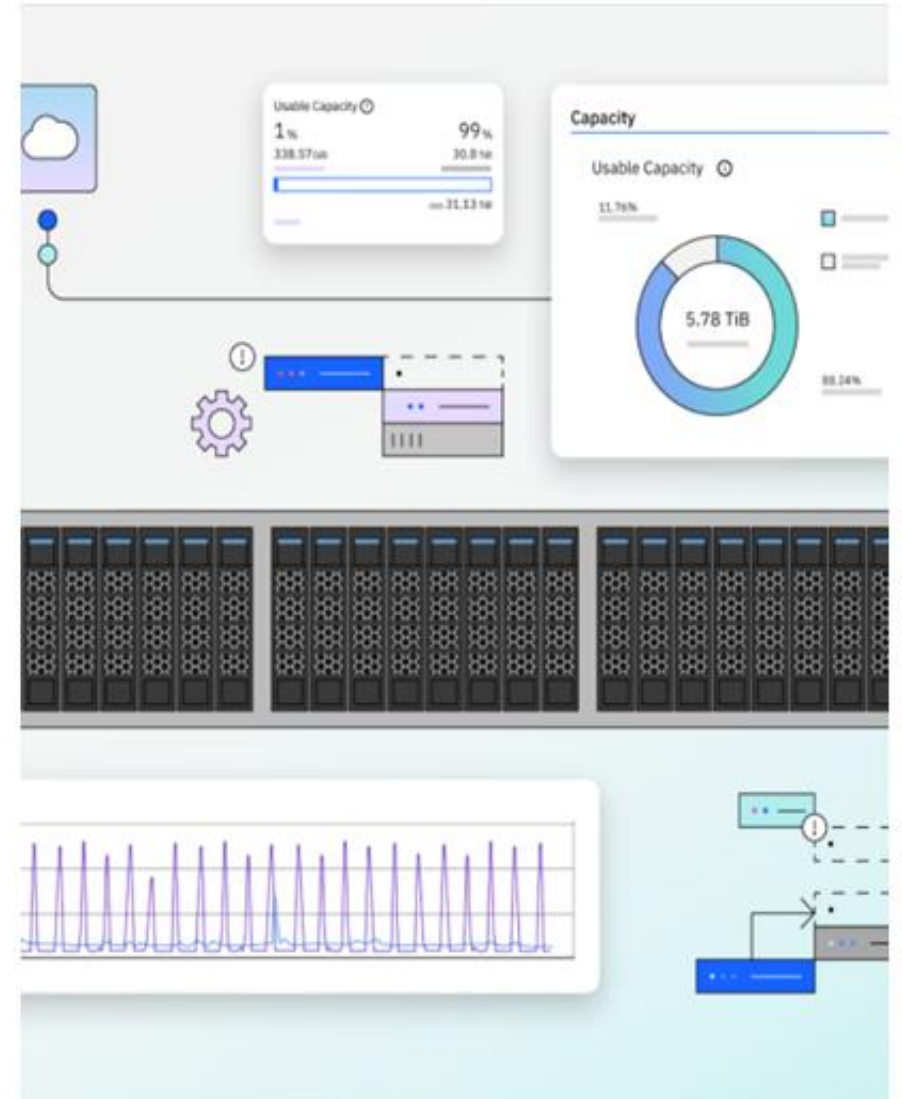
### Why its needed

- Patches had to initially be installed manually to begin the process creating an unnecessary burden on the storage administrator

### What you need to know

- A policy can be set system-wide or on an individual patch level. Policies are:
  - Policy = None (Ignore)
  - Policy = Autodownload – revised patches are downloaded but not installed
  - Policy = Autoupdate – revised patches are downloaded and automatically installed
  - Global policy can be set to none, autodownload, autoupdate
- Settings can be reset to factory default
- Checks for updates happen every 24 hours
- Cloud call home is a pre-requisite for this feature

## System/Easy Setup Redesign



## Setup Complete Summary With Detailed View

### System Setup

This setup wizard will guide you through configuring your system for first-time use.

 **Setup Complete!**



#### License agreements

Review and agree to terms and conditions



#### System details

Configure name, password, date and time, encryption recovery key and other system settings



[View summary](#) ▾



#### Register with IBM

Test the internet connectivity, connect to Call Home and IBM Storage Insights



[View summary](#) ▾



#### Support assistance

Opt for remote service, provide contact information and system location



[View summary](#) ▾

**i** This setup process creates auto zoning enabled fibre channel portsets by default during initial configuration for new systems.

If singlevolumentype then this text : This setup process creates a pool, array and auto zoning enabled fibre channel portsets by default during initial configuration for new systems.

[Cancel](#)

[Continue](#) 



#### License agreements

Review and agree to terms and conditions



#### System details

System Name : Cluster\_violet-cl.ssd.hursley.ibm.com

User name : superuser

Date and time : 01/22/2026 12:32 pm (GMT+5:30) Mumbai, New Dethi

System utility : Standalone system



[Hide summary](#) ^



#### Register with IBM

Call home : Connected, via internet

Storage Insights : Registered

IBMid : dipali.balasaheb.kamble@ibm.com

Tenant : 01f04c15-96cc-157d-993e-856dc6d1e56e



[Hide summary](#) ^



#### Support assistance

Company name : IBM

Country or region : IN

Street address : Pune

City : Pune

State : MH

Postal code : 110056

Asset or machine location : TECH PARK, 3RD FLOOR

Name : John Doe

Email : JohnDoe@abc.com

Phone (Primary) : 9999999999

Phone (Secondary) :



[Hide summary](#) ^

## Auto Port set Creation Message

- Dashboard
- Monitoring
- Pools
- Volumes
- Hosts
- Policies
- Access
- Settings
- Storage partitions

**Explore IBM FlashSystem grid**  
This system supports IBM FlashSystem grid configuration.

Add this system to an existing FlashSystem grid or create a new FlashSystem grid with this system as the first member. Within a FlashSystem grid, you can seamlessly migrate workloads among connected systems and optimize their infrastructure usage.

[Learn more](#)

**Performance**    Granularity: seconds    [View historical data with IBM Storage Insights](#)

Metric	read	write	Unit
Latency	0 µs	0 µs	µs
Bandwidth	0 MBps	0 MBps	MBps
IOPS	0 IOPS	0 IOPS	IOPS
CPU Utilization	5%		%
Power Consumption	565 Watts		Watts

**Capacity**

Physical capacity  ⓘ

Physical used ■

Physical available ■

Capacity Savings

0 bytes

**Suggested Actions** ✕

There are suggested actions for you to do below.

Actions (5)

**Auto portset creation unsuccessful**  
This can occur due to multiple connectivity reasons. You will need to manually review and configure ports and portsets. Not Now

**Actions**  ^

- View fibre channel ports
- View fibre channel portsets

**Create Host**  
Create a Host for providing access to Volumes. Not Now

**Create**

**Create Pool**  
Create a Pool in order to Create Volumes. Not Now

**Create**

**Multifactor authentication**  
Enhance security by adding multiple authentication factors for verifying user identity. Not Now

**Enable**

**Two person integrity**  
Require two authorized users to approve critical tasks for added security. Not Now

**Configure**

**Auto-portset creation successful**  
System has created optimal fibre channel portsets based on your system configuration. Not now

**Actions**  v

- View fibre channel ports
- View fibre channel portsets

## Name Change

---

### Becomes

- **IBM Storage Control Select Edition**

- 5608-PC2 – AAS (e-Config)
- 5725-G33 – PPA

- 1 license per 25 slots or 500TiB raw

- Includes:

- Storage Insights for Spectrum Control (Pro)
- Copy Services Manager

- **IBM Storage Intelligence Control**

- 5608-PC2 – AAS (e-Config)
- 5725-G33 – PPA

- 1 license per 25 slots or 500TiB raw

- Includes:

- Storage Insights for Spectrum Control (Pro)
- Copy Services Manager

**No licensing changes**

**No pricing changes**

**No entitlement changes**

# Ansible Updates

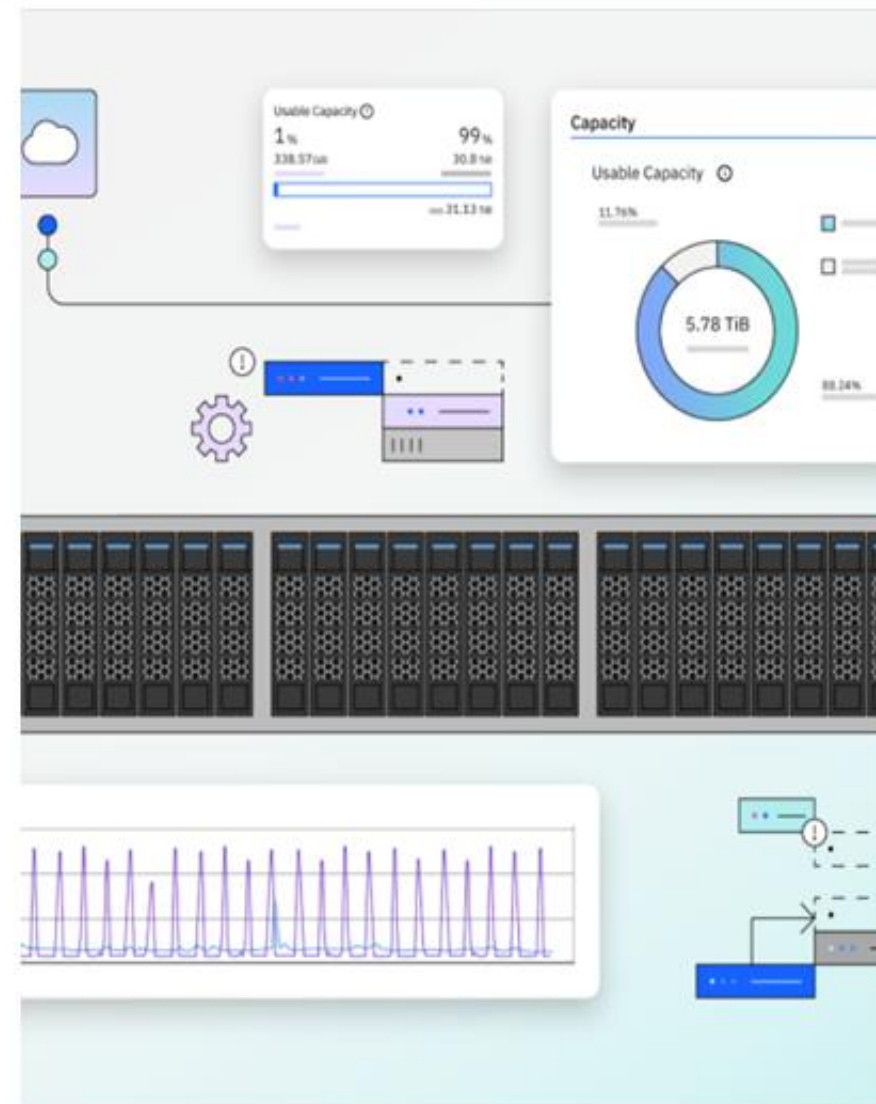




Figure: Ansible-Driven Automation for IBM FlashSystem

# Ansible Modules



Works with REST API and ssh



49 Modules and 11 set of playbooks



All basic Storage objects like volume, volumegroup, snapshots, hosts



Covers PBHA, PBR, PBRHA, FlashSystem Grid



Release: Every 3 months once, via github, Ansible automation hub (certified version), Ansible galaxy (community version)

# Playbooks

- GM/GMCSV migration to PBR
- HyperSwap migration to PBHA
- Setting up PBR/PBHA
- Setting up PBRHA (3-site)
- Set up host-rescan across FlashSystem grid (volume migration)
- FlashSystem Log collection
- FlashSystem Firmware upgrade
- Copy 1 FlashSystem config to another

## Ansible v3.1.0 and v3.2.0 Updates

---

### ibm.storage\_virtualize v3.1.0:

- Certificate management for PBR, PBHA and FlashSystem grid using new FlashSystem APIs
- Updating management IP parameters (VLAN, gateway, and IP address)
- Security enhancement by hiding authentication outputs

### ibm.storage\_virtualize v3.2.0:

- Clone, thin clone of volume and volumegroup in partition
- Support for logging in via partition IP for selected modules
- Enabled autoexpand, preferrednode and cache parameters for Volume
- Enhanced partition\_migration\_host\_action playbook for iSCSI hosts

## Accelerate with ATG Survey

---

Please take a moment to share your feedback with our team!

You can access this 6-question survey via this direct link: <https://ibm.biz/Bdpcj3>

Or you can use this QR Code:



# Thank you!

