

IBM Storage Scale & Scale System

Overview



David Bechtold
IBM Principal Storage Solutions Architect
Financial Services Sector
dkbechto@us.ibm.com
<https://www.linkedin.com/in/david-bechtold-222578/>



IBM Storage Scale

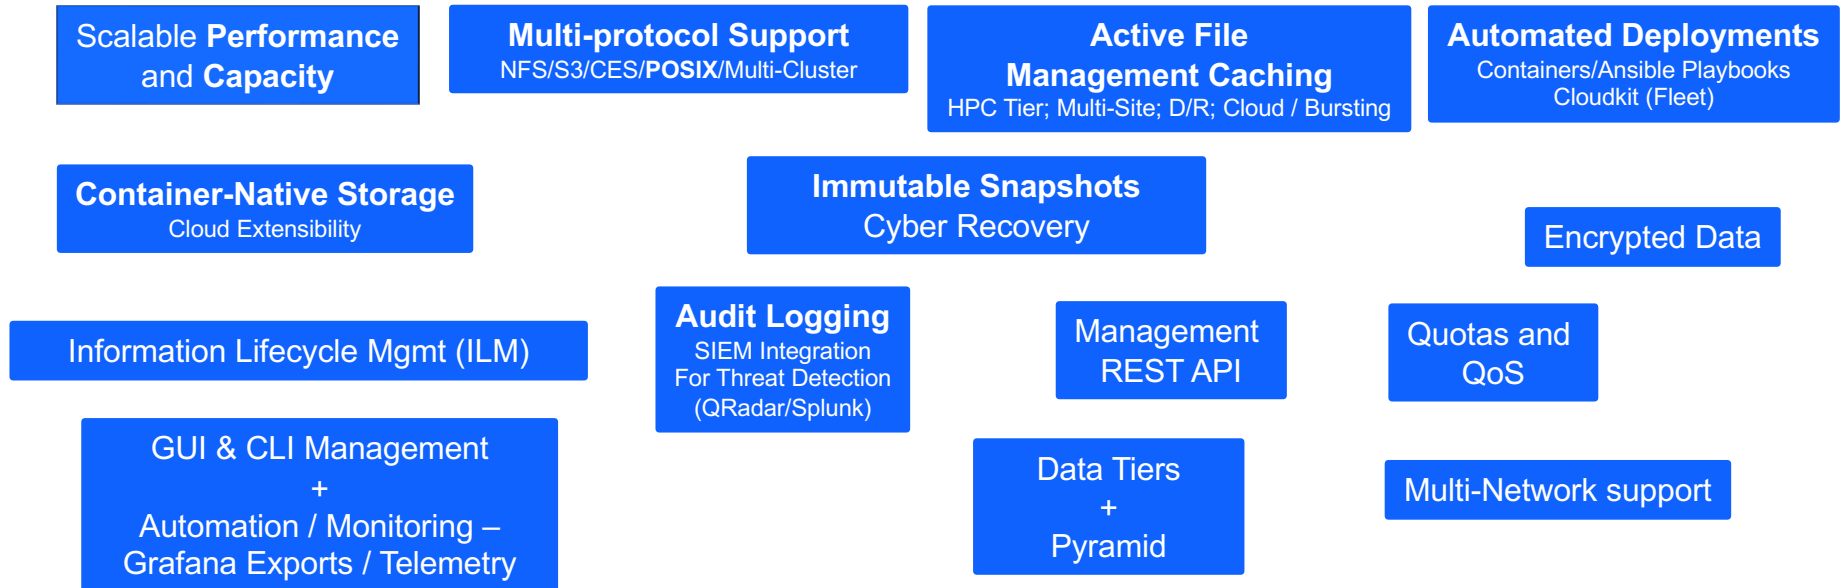
Introduction



What is IBM Storage Scale?

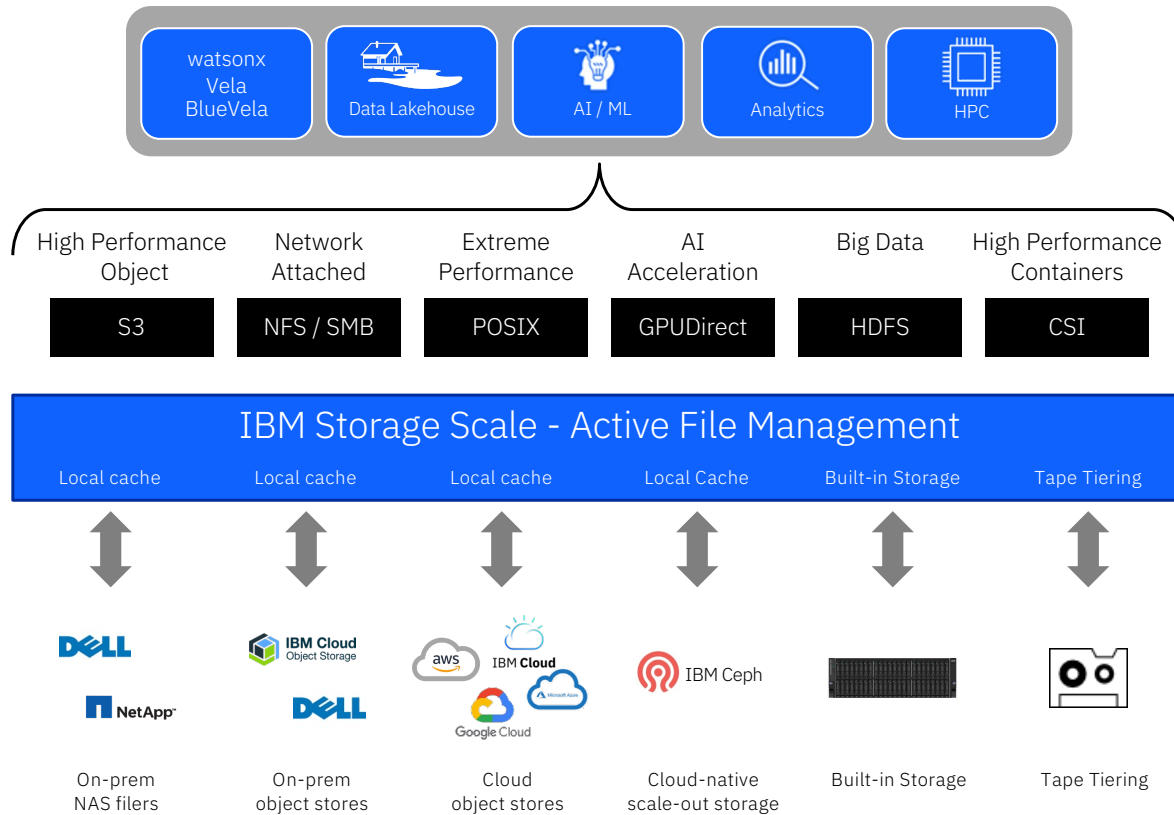
IBM Storage Scale is an enterprise-grade parallel file system that provides superior resiliency, scalability, and control. IBM Storage Scale delivers scalable capacity and performance to handle demanding data analytics, AI/ML, content repositories (including S3 data storage leveraging Nooba endpoints within Scale) and technical computing workloads at scale.

What Are Some of the Key Features of Storage Scale



IBM Storage Scale – Capability Highlights

Storage Access, Abstraction and Acceleration to maximize CapEx investment



Multi-Protocol Support Access

Simultaneous multi-protocol access including GPUDirect 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

IBM Storage Scale

Active Protection for Cyber Resiliency (NIST)

IDENTIFY

- Cyber Resiliency Assessment Tool, Probes 100s of different controls and best practices


Governance

- Data Catalog allowing for data orchestration and data migration control and accountability
- Watson Knowledge catalog

RECOVER

Recover Operations and Data Quickly

- Instant Restore with Storage Scale AFM
- Storage Scale and Storage Protect – recover multi-petabyte filesystems in hours
- QRadar Incident Forensics



PROTECT Active Protection against cyber attacks

- Multifactor Auth, RBAC, Privileged Access Monitoring (IBM Security Verify)
- Safeguarded Copies via immutable snapshots, logical air gap
- Scan snapshots for signs of ransomware
- Log all Admin & user actions

DETECT

Detect Suspicious Behavior

- QRadar and Splunk **SIEM integration**
- **File Audit Logging**, Watch Folders
- Analyze backup data for signs of ransomware (Spectrum Protect, Defender)
- Reporting: QRadar User behavior analytics
- IBM Flash Core Modules entropy detection

RESPOND

Alert and take action

- Automated action upon threat detection (QRadar)
 - Snapshot, Block Session , Etc..
- Alerts automatically prioritized based severity of the threat and criticality of the assets involved

IBM Storage Scale

Deployment Methods



IBM Storage Scale previously known as *Spectrum Scale* or *GPFS* is a software storage solution which can be deployed on-prem and in the cloud on existing SAN storage infrastructure, using IBM Storage Scale system, using Storage-Rich servers, or all the above!

1 Software Defined Storage

- Bring your own servers and storage
- Deploy On-Prem or in the cloud
- Supports variety of system types x86, Power, s390x
- Highly Scalable


2 Storage Scale System

- Building block storage solution based on GNR
- NVMe, HDD, or Hybrid building blocks
- Highly Scalable – add more building blocks for more Capacity and Performance
- Six 9's of Availability

3 Erasure Code Edition

- Bring your own storage rich servers
- NVMe and HDD support
- Utilize GNR RAID
- Add more servers for more capacity and performance

IBM Storage Scale
System 6000



IBM Storage Scale System

Integrated scale-out solution for **file** and **object** storage

Optimal scalable building block for **files** and **objects**:


- Pre-integrated, tuned appliance with full s/w stack installation & upgrade
- Scale-up and scale-out options
- Can manage billions of files/objects- scalable high perf Metadata engine
- Simplified administration: Scale GUI and Storage Control integration

One solution for all your **file** and **object** storage needs:

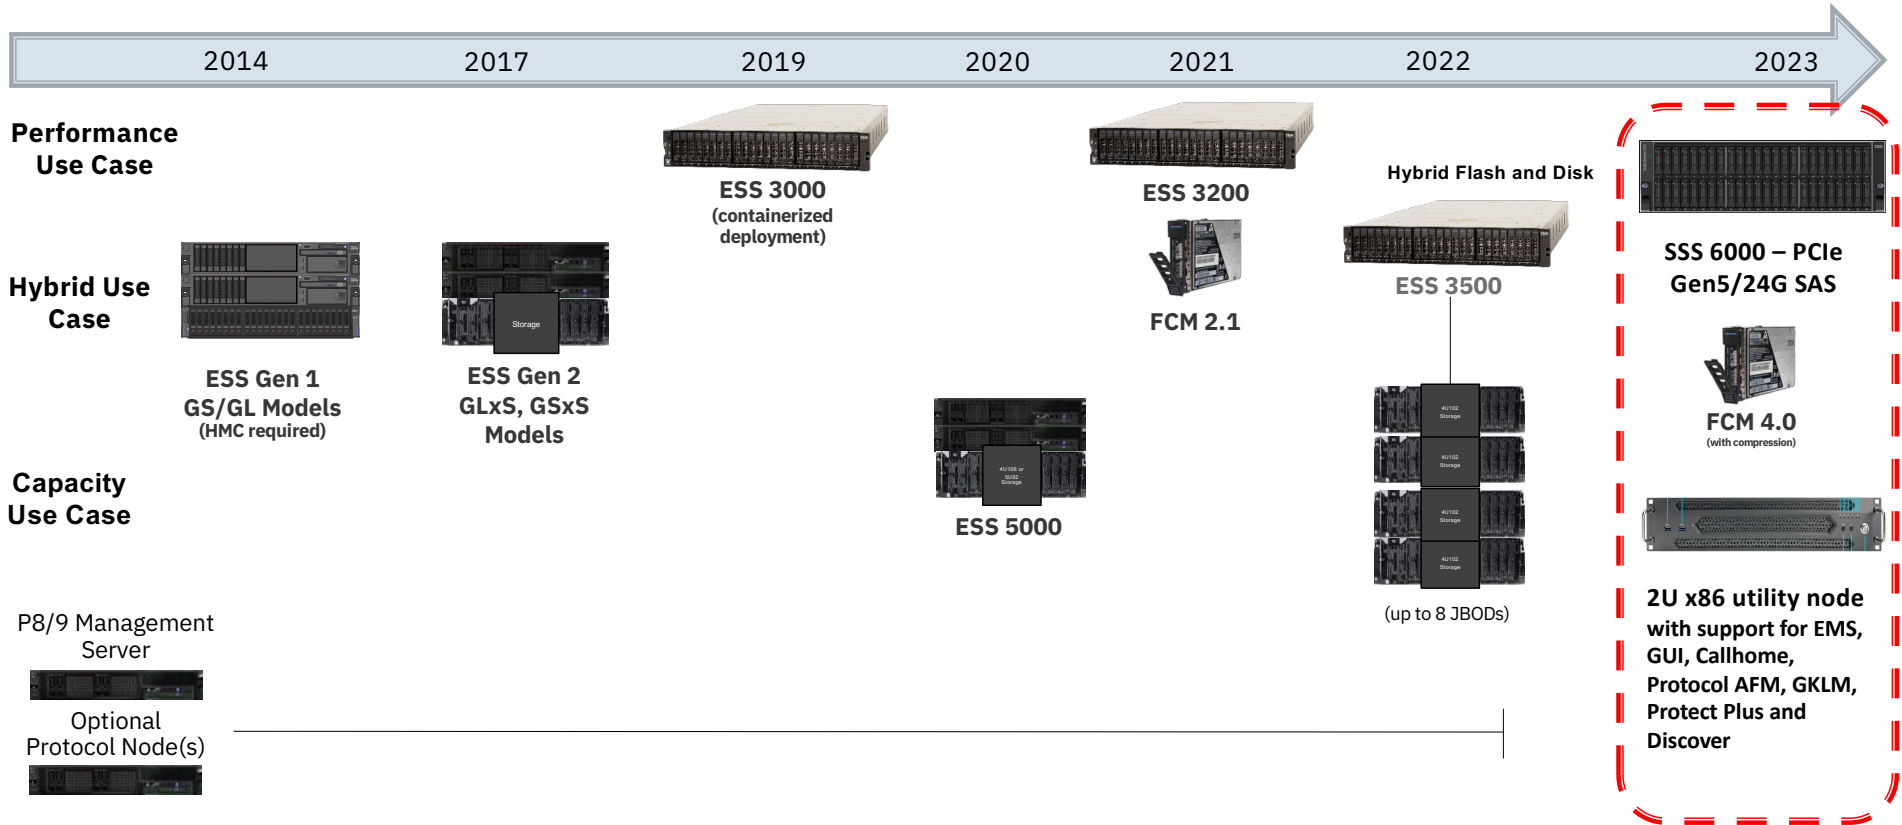
- Single shared copy of data- file and object
- Combine flash and HDD for right cost & performance, data placement
- Multi-protocol support plus container access

Proven and ready for **critical data** and **workloads**:

- HA, DR & data caching/sharing with sync or async replication
- Reliability, consistent performance, fast rebuild times
- Data durability from intelligent erasure coding rebuilds and drive hospital
- Five 9s of availability (99.999%)



IBM Storage Scale System Evolution



IBM Storage Scale System 6000



Gen5 Dual Canister 4U(A/A)-48 NVMe

AMD Genoa EPYC, dual-socket 48 cores per canister

Cache: 24 x 64GB (1536GB) per Canister; 3.072TB per System

NVIDIA CX7 adapters (8):

- 400Gb single port (IB only) x16 Gen5
- 200Gb VPI dual port (IB/ETH) x16 Gen5

48x U.2 NVMe SSDs (1.5PB)

48x U.2 IBM FlashCore Modules (1.8PB or up to 5.5PB @ 3:1)

x86 utility node (management & cluster export services)

NVMe Drives Options -

- 3.84 TB
- 7.68 TB
- 15.36 TB
- 30.74 TB
- 19.2 / 38 TB FCM 4.0

HDD SED Drive Options

- 12 TB SAS HDD
- 16 TB SAS HDD
- 20 TB SAS HDD
- 22 TB SAS HDD

Performance and Sustainability

- **2x throughput improvement**
- **NVMeoF support**
- **Hybrid performance and capacity support**
- **Containerized protocol support on IO nodes**

Why IBM Storage for NVIDIA GPU deployments?

The world's fastest systems need the world's best storage.
IBM has the best storage for NVIDIA GPUs

IBM Storage Scale System 6000

A single 4U node with active-active controllers and redundant hardware to maximize always on data

Highest Performance Platform

- Fastest performance for reads, writes, and density
- Linearly scalability for future growth

A Robust Enterprise Platform

- Six 9's for all apps: AI, Analytics, HPC, Back-up, Archive, Cloud
- Cyber-resilient, encryption, WORM, and immutability

Collapse Layers & Simplify

- Eliminate extra copies and share data globally with all protocols
- Data cataloging and tiering for economics and data flexibility



Ultimate Performance

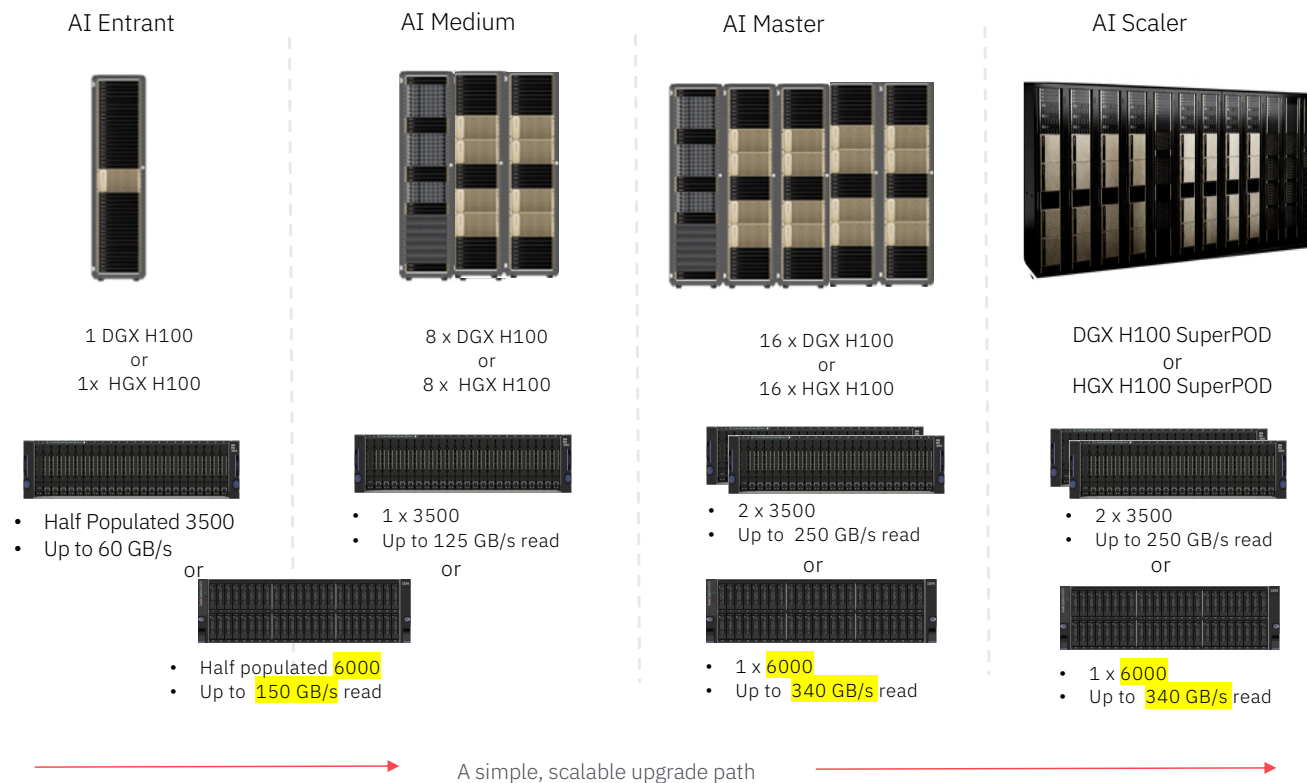
up to **340 GB/s** read performance per node
up to **175 GB/s** write performance per node

Software-defined flexibility

IBM Storage for Data and AI & NVIDIA GPU Solutions

A full spectrum of scalable AI solutions

Start small and scale predictably in response to business demand with the same IBM Storage Software

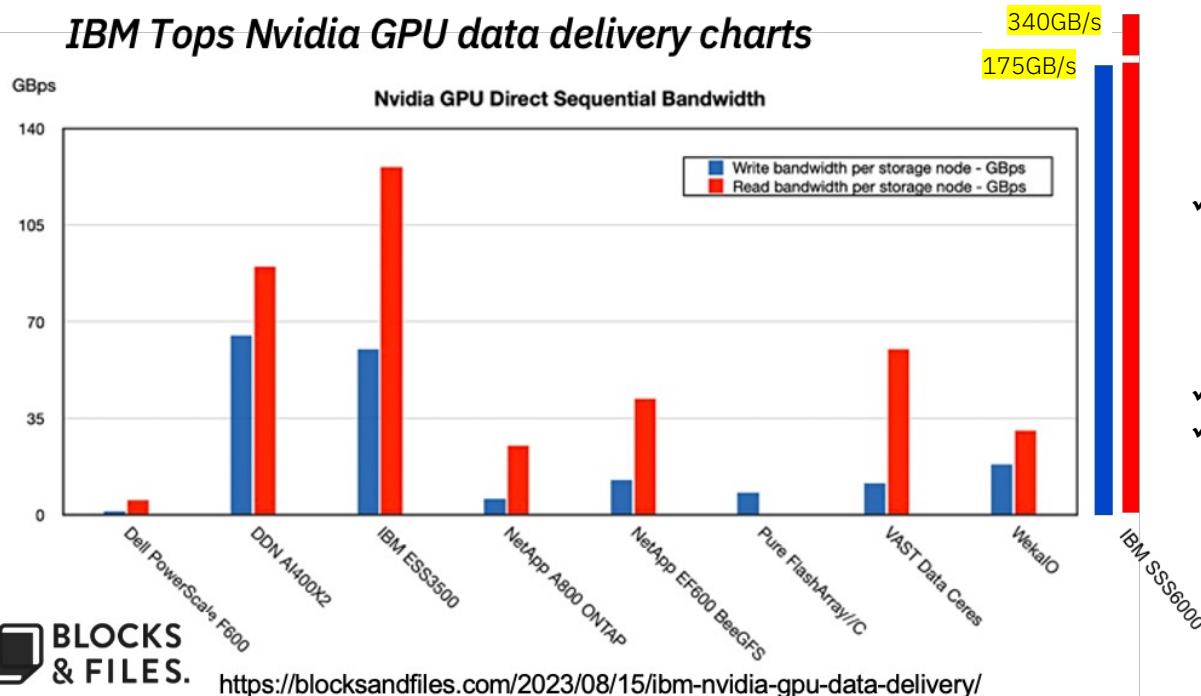


IBM Storage:

- Simple building blocks – scalable seamless storage upgrade path as needs grow from 1st DGX to AI CoE DGX SuperPOD
- Global Data Platform – Data fidelity capabilities to automate AI workflows.
- Data Economics – Eliminate copies and transparently tier
- Trusted, global enterprise level support and services.
- Successful deployments across the globe

IBM Storage sets new marks for performance

IBM Tops Nvidia GPU data delivery charts



- ✓ IBM SSS6000 is more than 2x more performant than the current ESS3500
 - ✓ 340 GB/s read
 - ✓ 175 GB/s write
- ✓ Latest in networking
- ✓ Ready to support GH200 platforms



<https://blocksandfiles.com/2023/08/15/ibm-nvidia-gpu-data-delivery/>




GPU Technology Conference 2024

<https://www.cbsnews.com/colorado/news/nvidia-ai-focus-under-jensen-huang-60-minutes/>



Accelerate AI with NVIDIA and IBM Storage

The IBM Storage Scale System 6000 is the fastest integration point for NVIDIA DGX to address the challenges of AI data optimization



-  **Changing technology**
GPUs used for AI are driving the need for larger data sets and faster data delivery
-  **Data silos**
Data is scattered and siloed throughout an organization making it difficult to gain access to relevant data for AI
-  **Unknown threats**
The veracity and accuracy of data are critical to AI and it must be protected from data breaches – accidental or otherwise
-  **Costs**
More data and faster delivery can mean new technologies and infrastructure that can strain budgets and sustainability strategies



IBM Storage Scale System 6000

-  **Accelerate AI and data delivery**
GPU Direct Storage with embedded AI accelerator delivering 310 GB/s and 15M IOPS
-  **Eliminate data silos**
Globally connect relevant data without data movement from across the organization (on-premises and off)
-  **Data and cyber resilient storage**
Six 9s of availability with globally dispersed erasure coding for always on and immutable data protection against accidents and threats
-  **Meet sustainability goals and lower costs**
Greater storage density on all-flash media with computational drives to offload CPU-intensive services across storage tiers

IBM