Spectrum Scale AI ecosystem and how it supports GPU based workloads including Power AI

Tomer Perry
Scalable I/O development

Includes content provided by Piyush Chaudhary, Ted Hoover, Andreas Koeninger, Simon Lorenz

IBM

Disclaimer

The information in this document is **IBM CONFIDENTIAL**.

This information is provided on an "AS IS" basis without warranty of any kind, express or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Some jurisdictions do not allow disclaimers of express or implied warranties in certain transactions; therefore, this statement may not apply to you.

This information is provided for information purposes only as a high level overview of possible future products. PRODUCT SPECIFICATIONS, ANNOUNCE DATES, AND OTHER INOFORMATION CONTAINED HEREIN ARE SUBJECT TO CHANGE AND WITHDRAWAL WITHOUT NOTICE.

USE OF THIS DOCUMENT IS LIMITED TO SELECT IBM PERSONNEL AND TO BUSINESS PARTNERS WHO HAVE A CURRENT SIGNED NONDISCLUSURE AGREEMENT ON FILE WITH IBM. THIS INFORMATION CAN ALSO BE SHARED WITH CUSTOMERS WHO HAVE A CURRENT SIGNED NONDISCLOSURE AGREEMENT ON FILE WITH IBM, BUT THIS DOCUMENT SHOULD NOT BE GIVEN TO A CUSTOMER EITHER IN HARDCOPY OR ELECTRONIC FORMAT.

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.

Outline

- AI Pipeline
- HDP goes Mainstream
- Containers
- Storage for AI
- Getting the data closer to the GPU

Outline

- **AI Pipeline**
- HDP goes Mainstream
- Containers
- Storage for AI
- Getting the data closer to the GPU

Workflow and Data Flow is complex

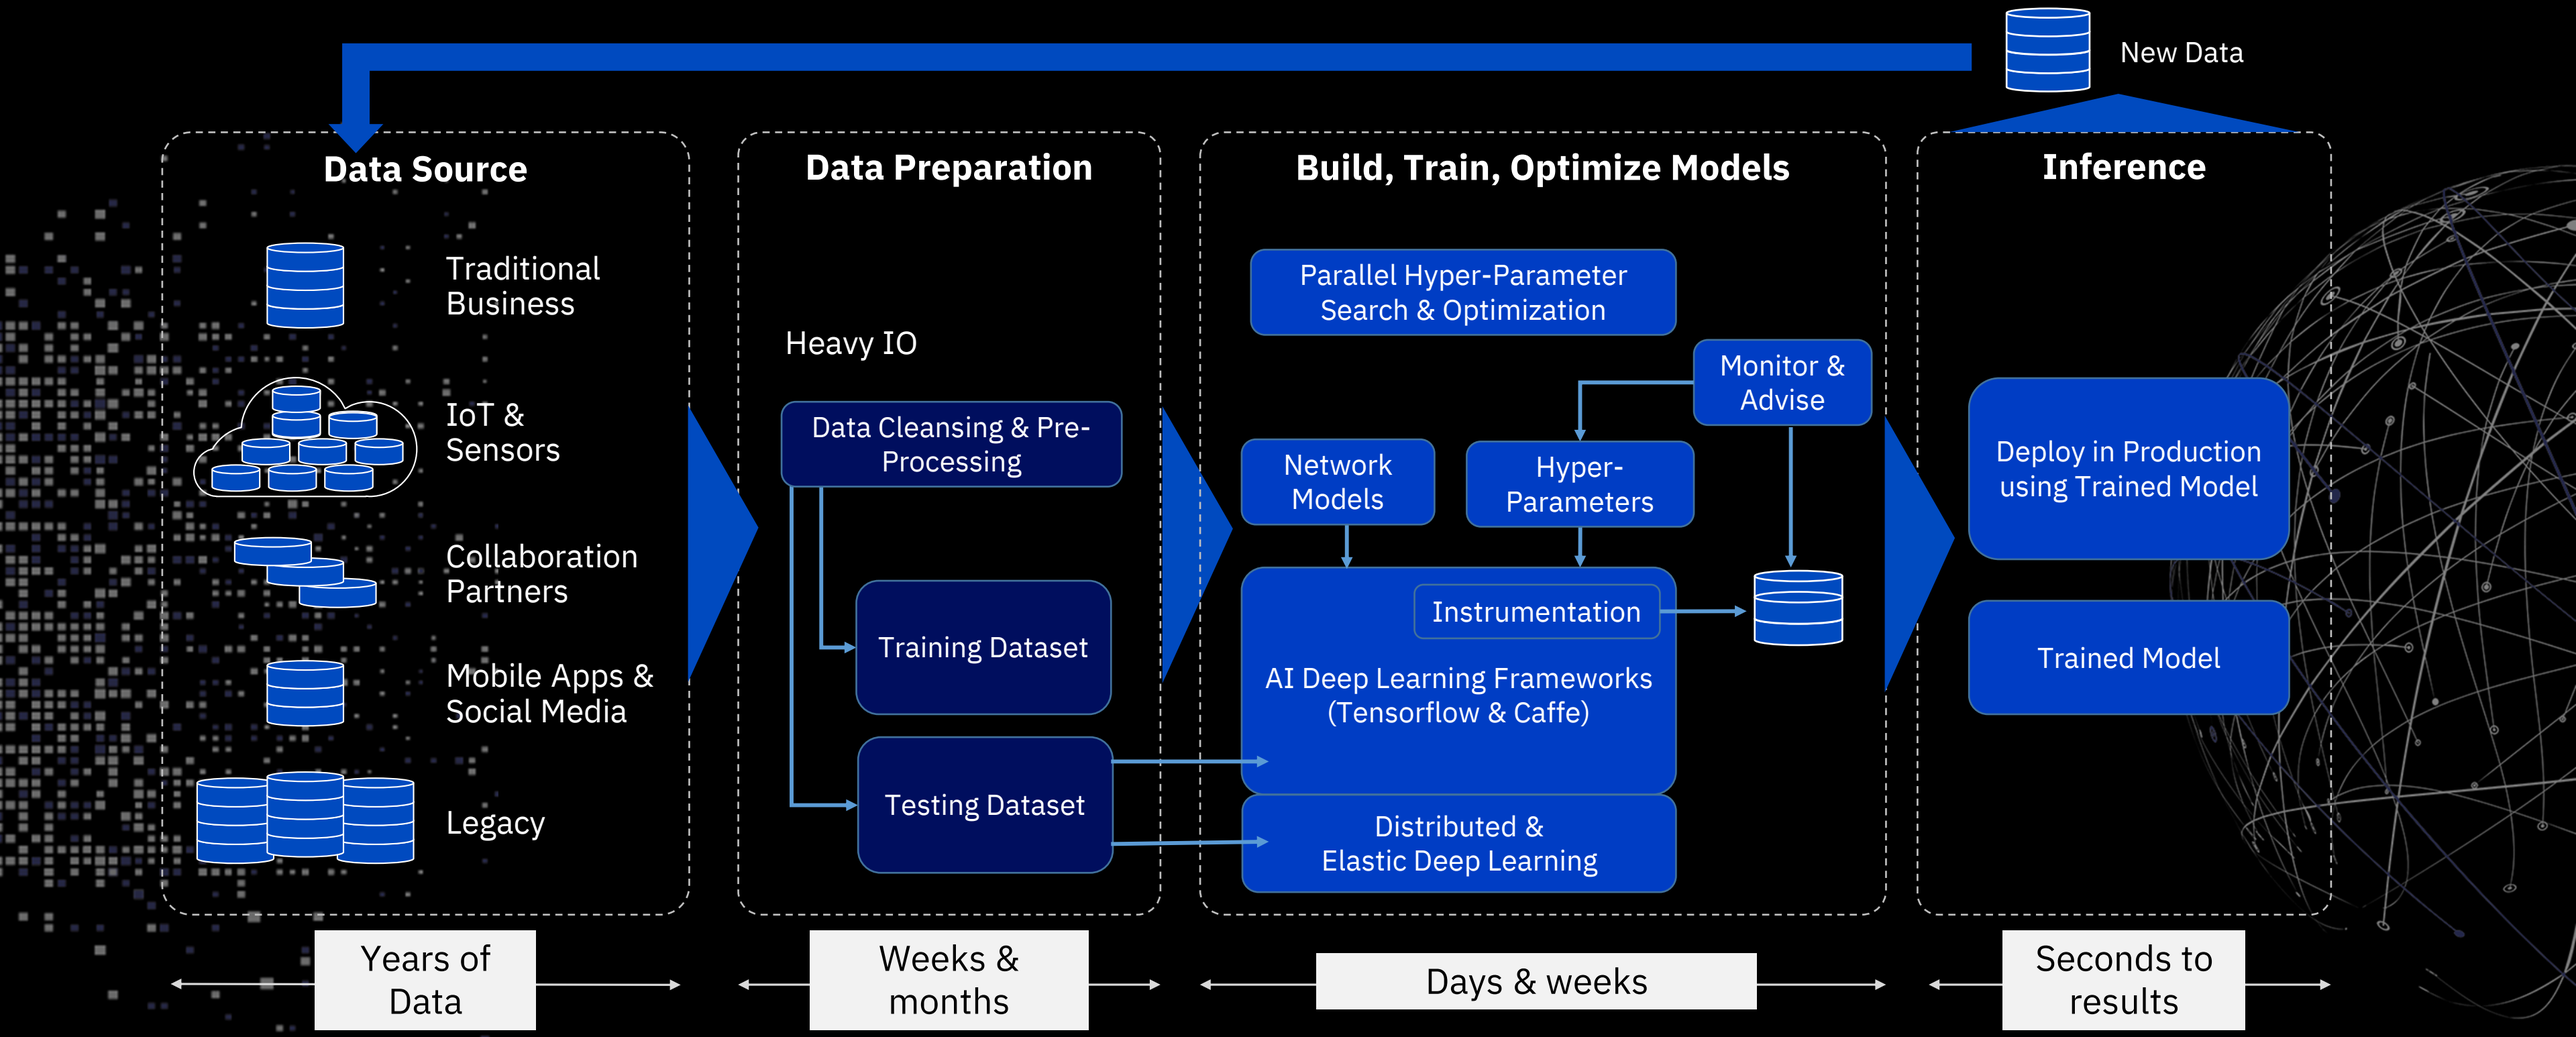

Enterprise Data Pipeline with IBM Spectrum Storage

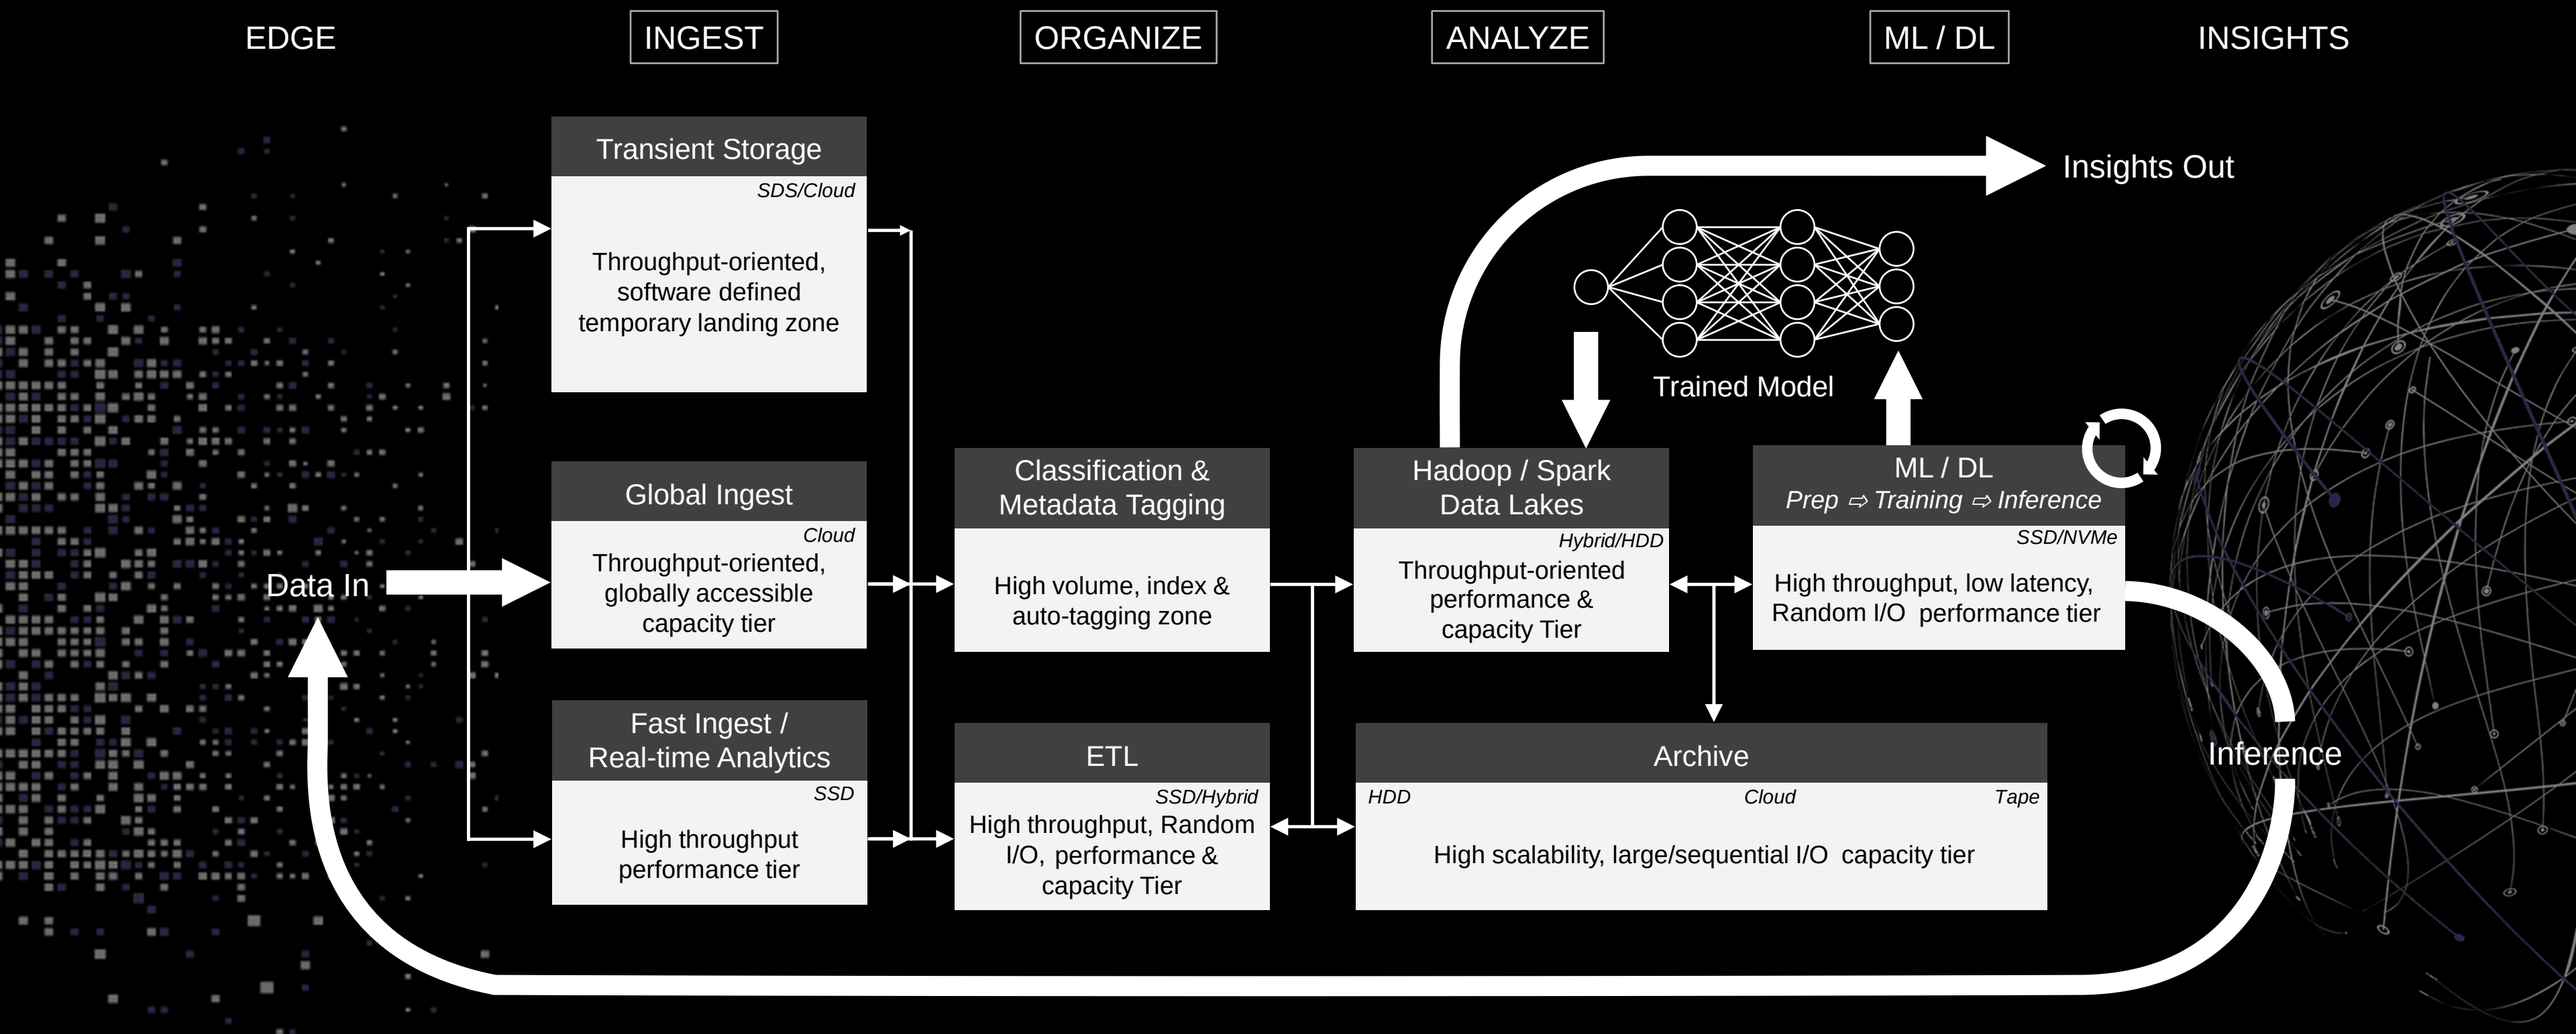

Enterprise Data Pipeline with IBM Spectrum Storage

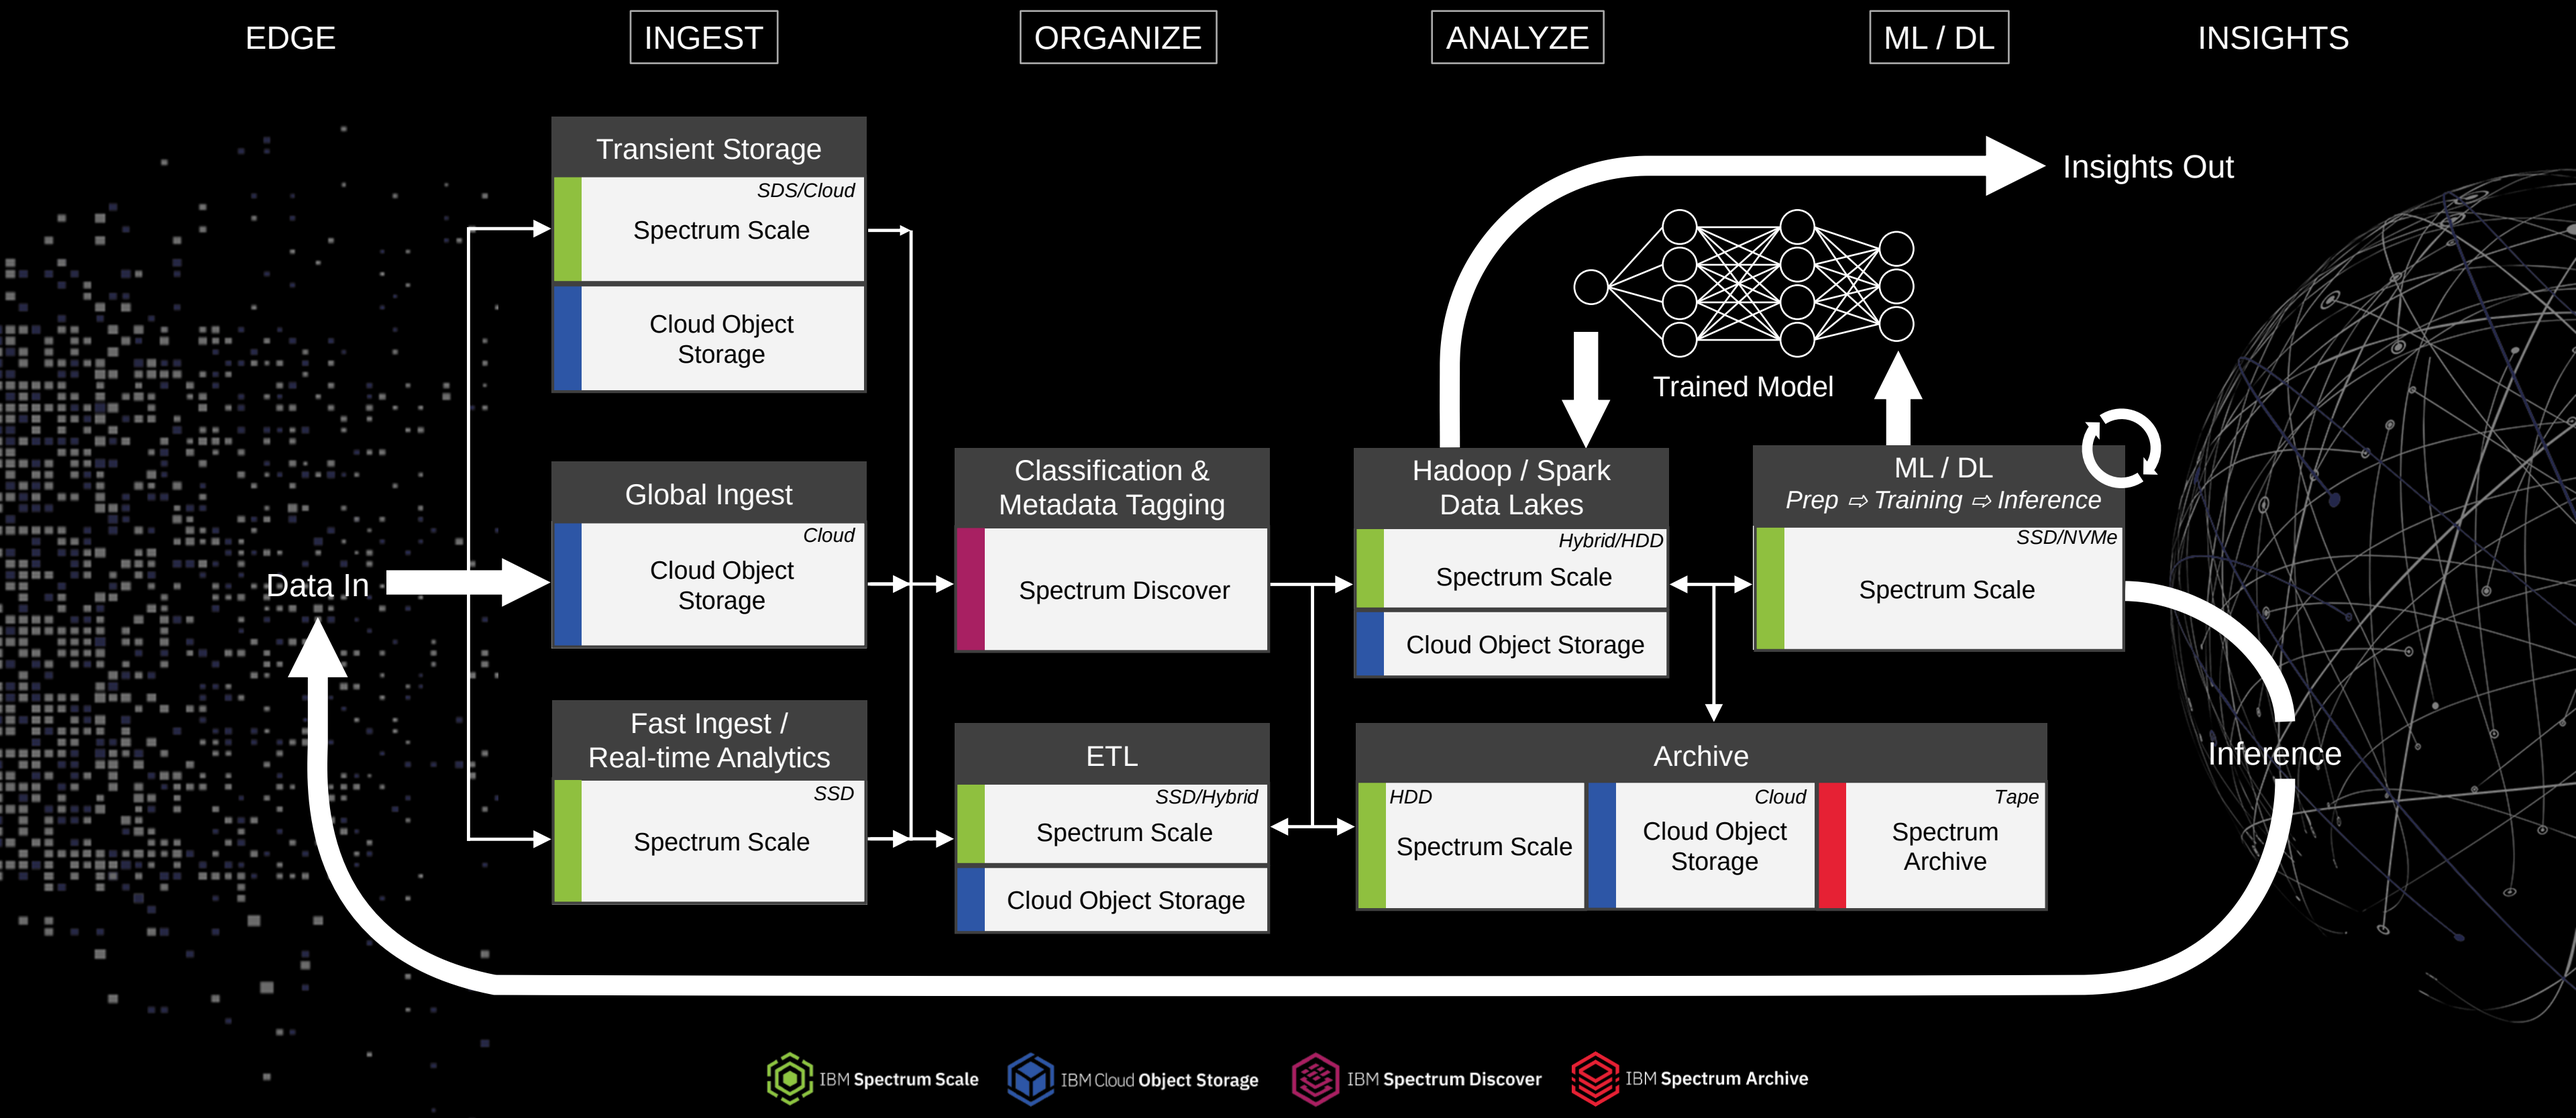

IBM Spectrum Scale IBM Cloud Object Storage IBM Spectrum Discover IBM Spectrum Archive

Outline

- AI Pipeline
- **HDP goes Mainstream**
- Containers
- Storage for AI
- Getting the data closer to the GPU

Integration of HDFS in CES

- HDFS Transparency becomes an integral part of Spectrum Scale
- Easy setup of HDFS Transparency using existing CES mechanisms
  e.g. „mmces service enable“
- Only NameNodes will be managed by CES
- HDFS clients always talk to the same CES IP (for NameNode requests)
- CES monitors the NameNode and moves the CES IP to another available node if something goes wrong
- Multiple HDFS clusters supported through multiple CES groups

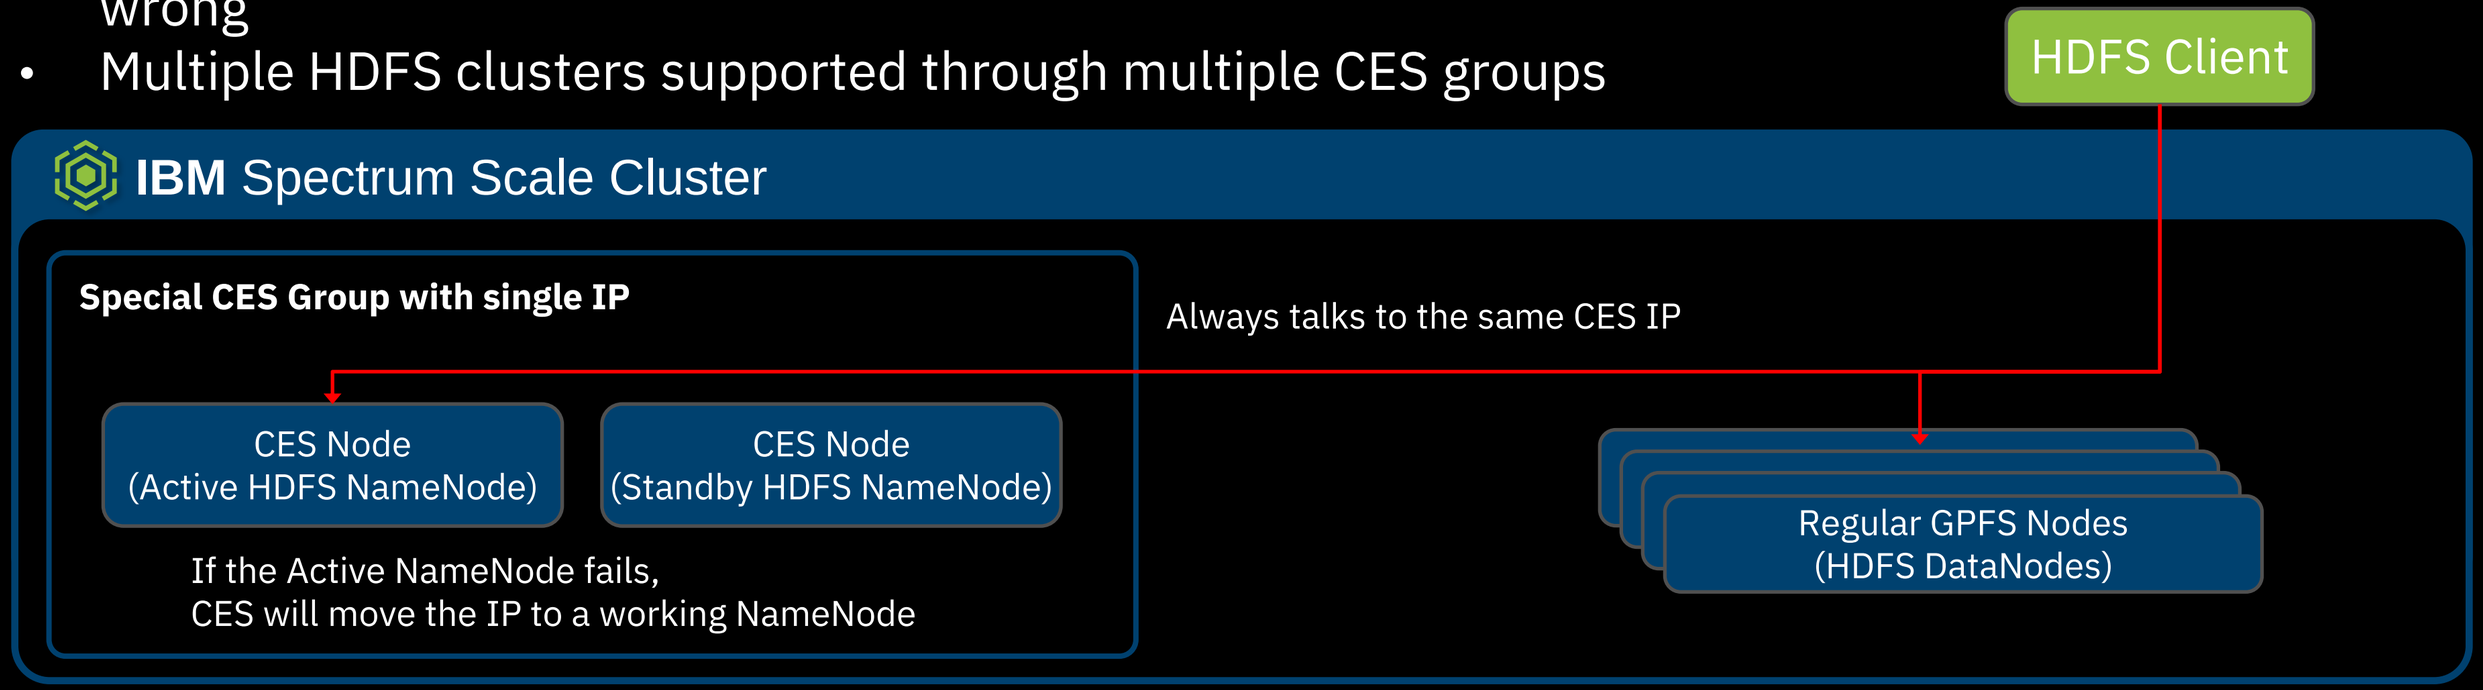

Outline

- AI Pipeline
- HDP goes Mainstream
- **Containers**
- Storage for AI
- Getting the data closer to the GPU

Journey to cloud requires an open, hybrid approach

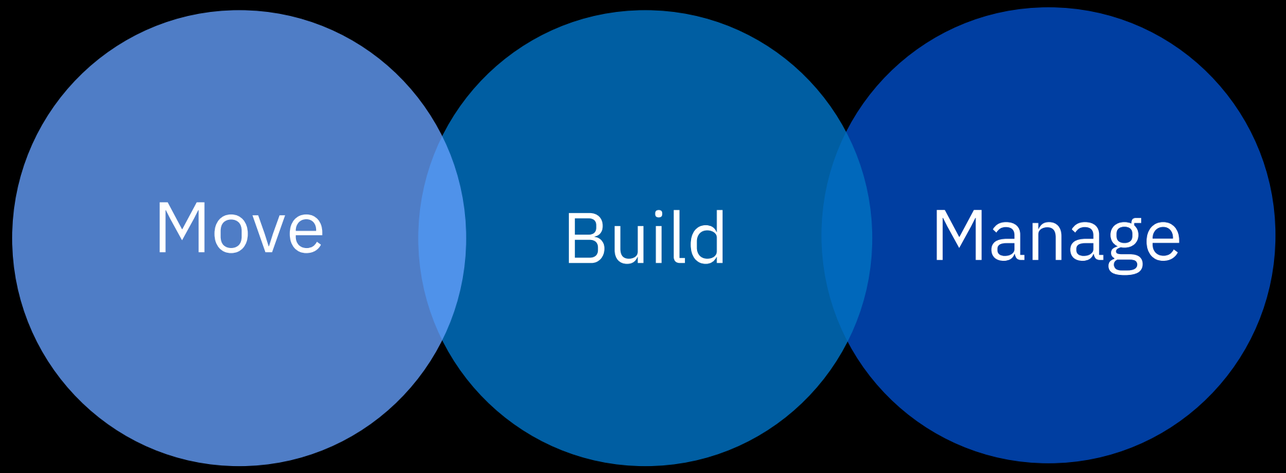

Portable

Future proof by building once, deploying anywhere for flexible data and workload placement

Container platform

Predictable

Open and integrated consistent management services that ensure operational integrity and reduce cost

Operational services

Productive

Integrated and secure containerized software for an agile, yet governed, enterprise

Containerized software secure by design

Goal: Deliver High Performance File Services to Containerized Application Workloads

Support Workloads that Require High Performance File Services

- Analytics & Cognitive
- High Performance Computing
- AI Data Pipeline

Support the Workload Ecosystem in the Cloud

- Containerized Applications, Storage
- Ephemeral and Persistent Storage Volumes

Flexible Deployment

- Dynamic Provisioning, Configuration, Upgrade

Support for Multiple Clouds

- Public, Private, Hybrid

Support Hybrid Use Cases

- Cloud Burst – Single Name Space
- Multi Cloud Data Sharing
- Archive
- High Performance Tiering

Solution Integration (Partners)

Spectrum Scale Containers Models

Storage for Containers

- Container Ready Storage

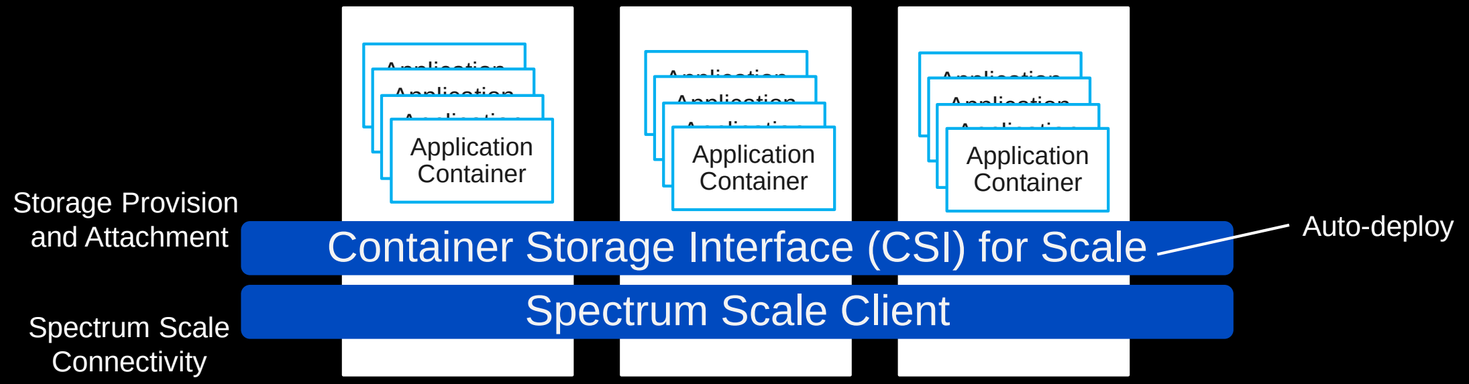

Storage in Containers

- Containerized Storage

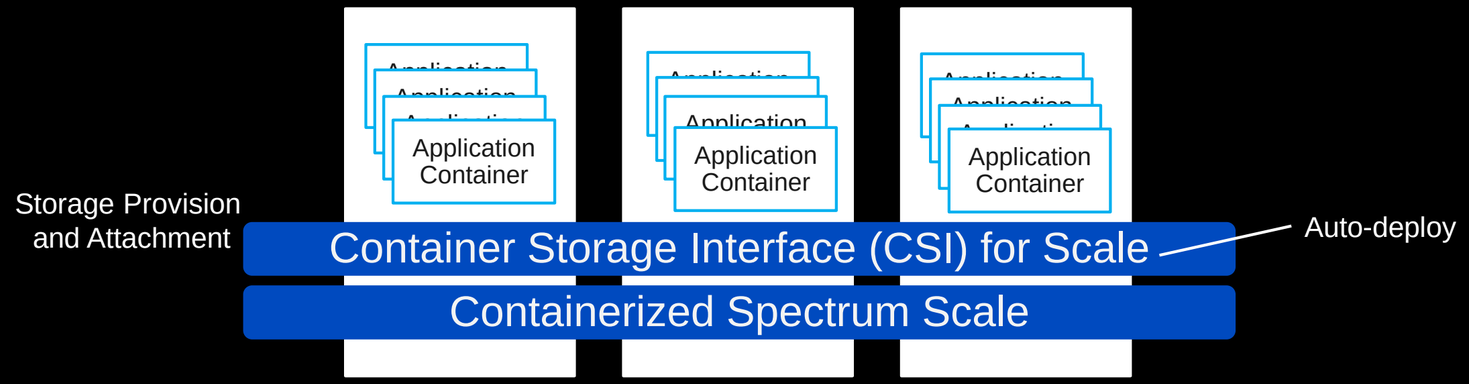

Evolution of IBM Spectrum Scale Containers

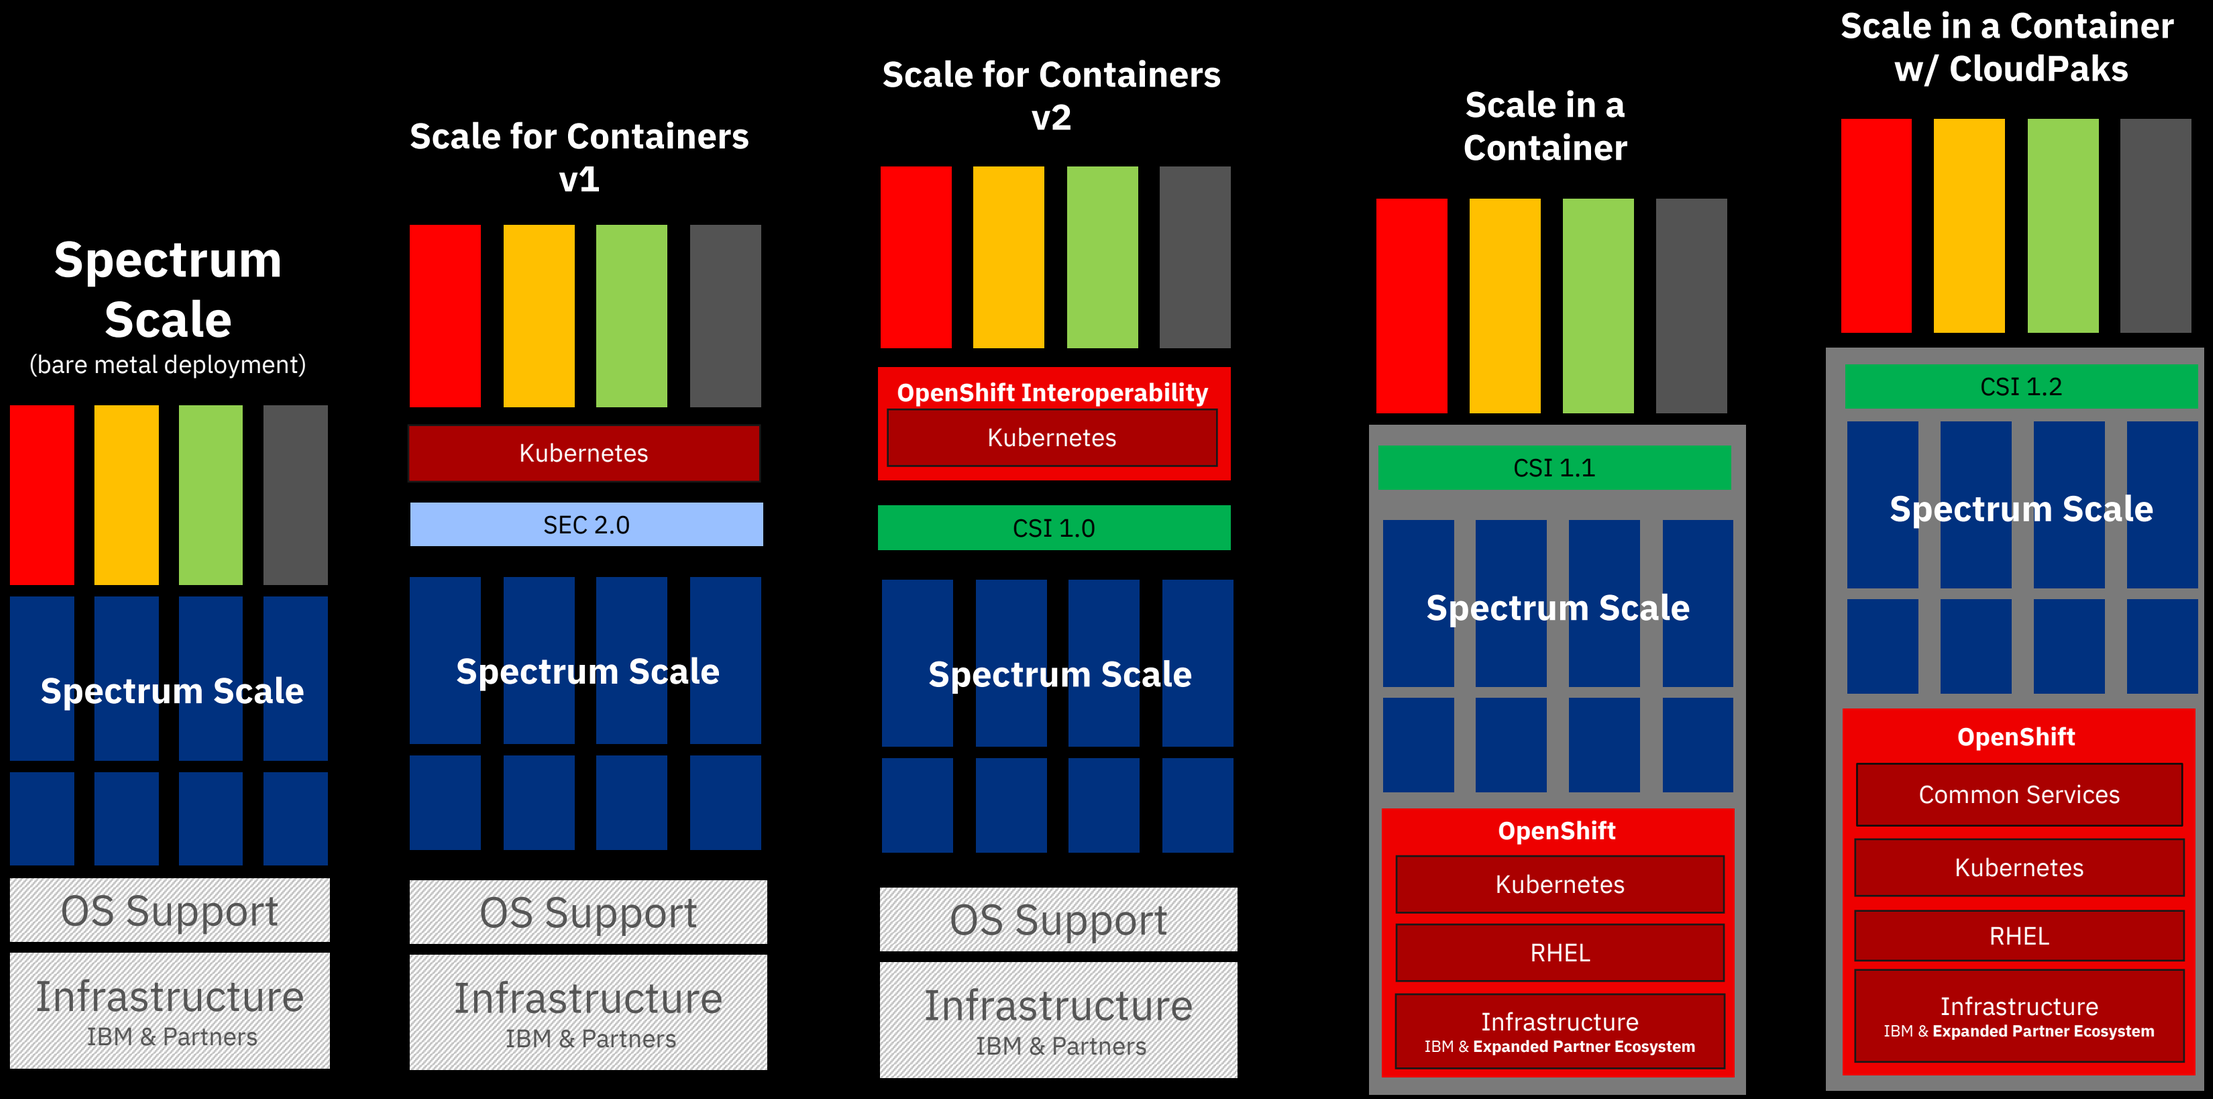

Evolution of IBM Spectrum Scale on Cloud

Current
Partner and Scale Offerings

Future
Partner and Scale Offerings

Spectrum Scale
(bare metal deployment)

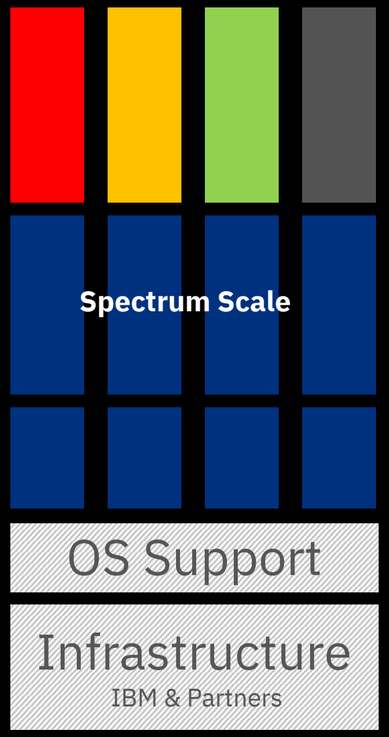

Spectrum Scale IBM Cloud

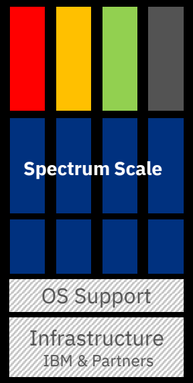

Spectrum Scale on AWS

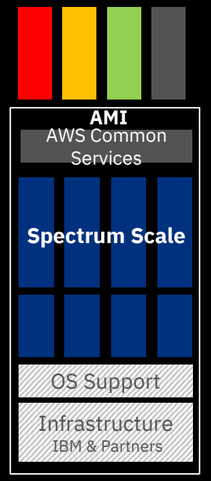

Spectrum Scale Partner Solutions

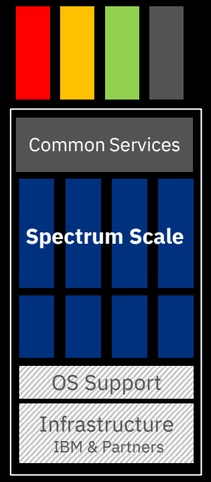

Scale in a Container On Cloud

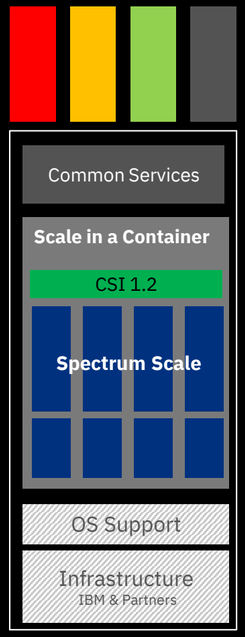

Scale in a Container w/ CloudPaks (Multi-Cloud)

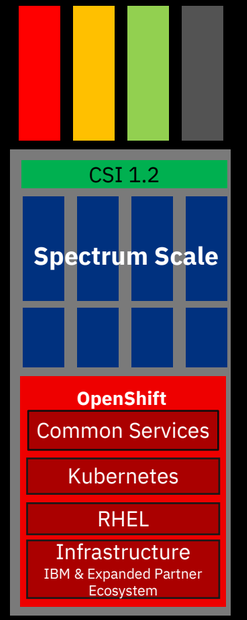

Evolution of Hybrid Cloud with IBM Spectrum Scale

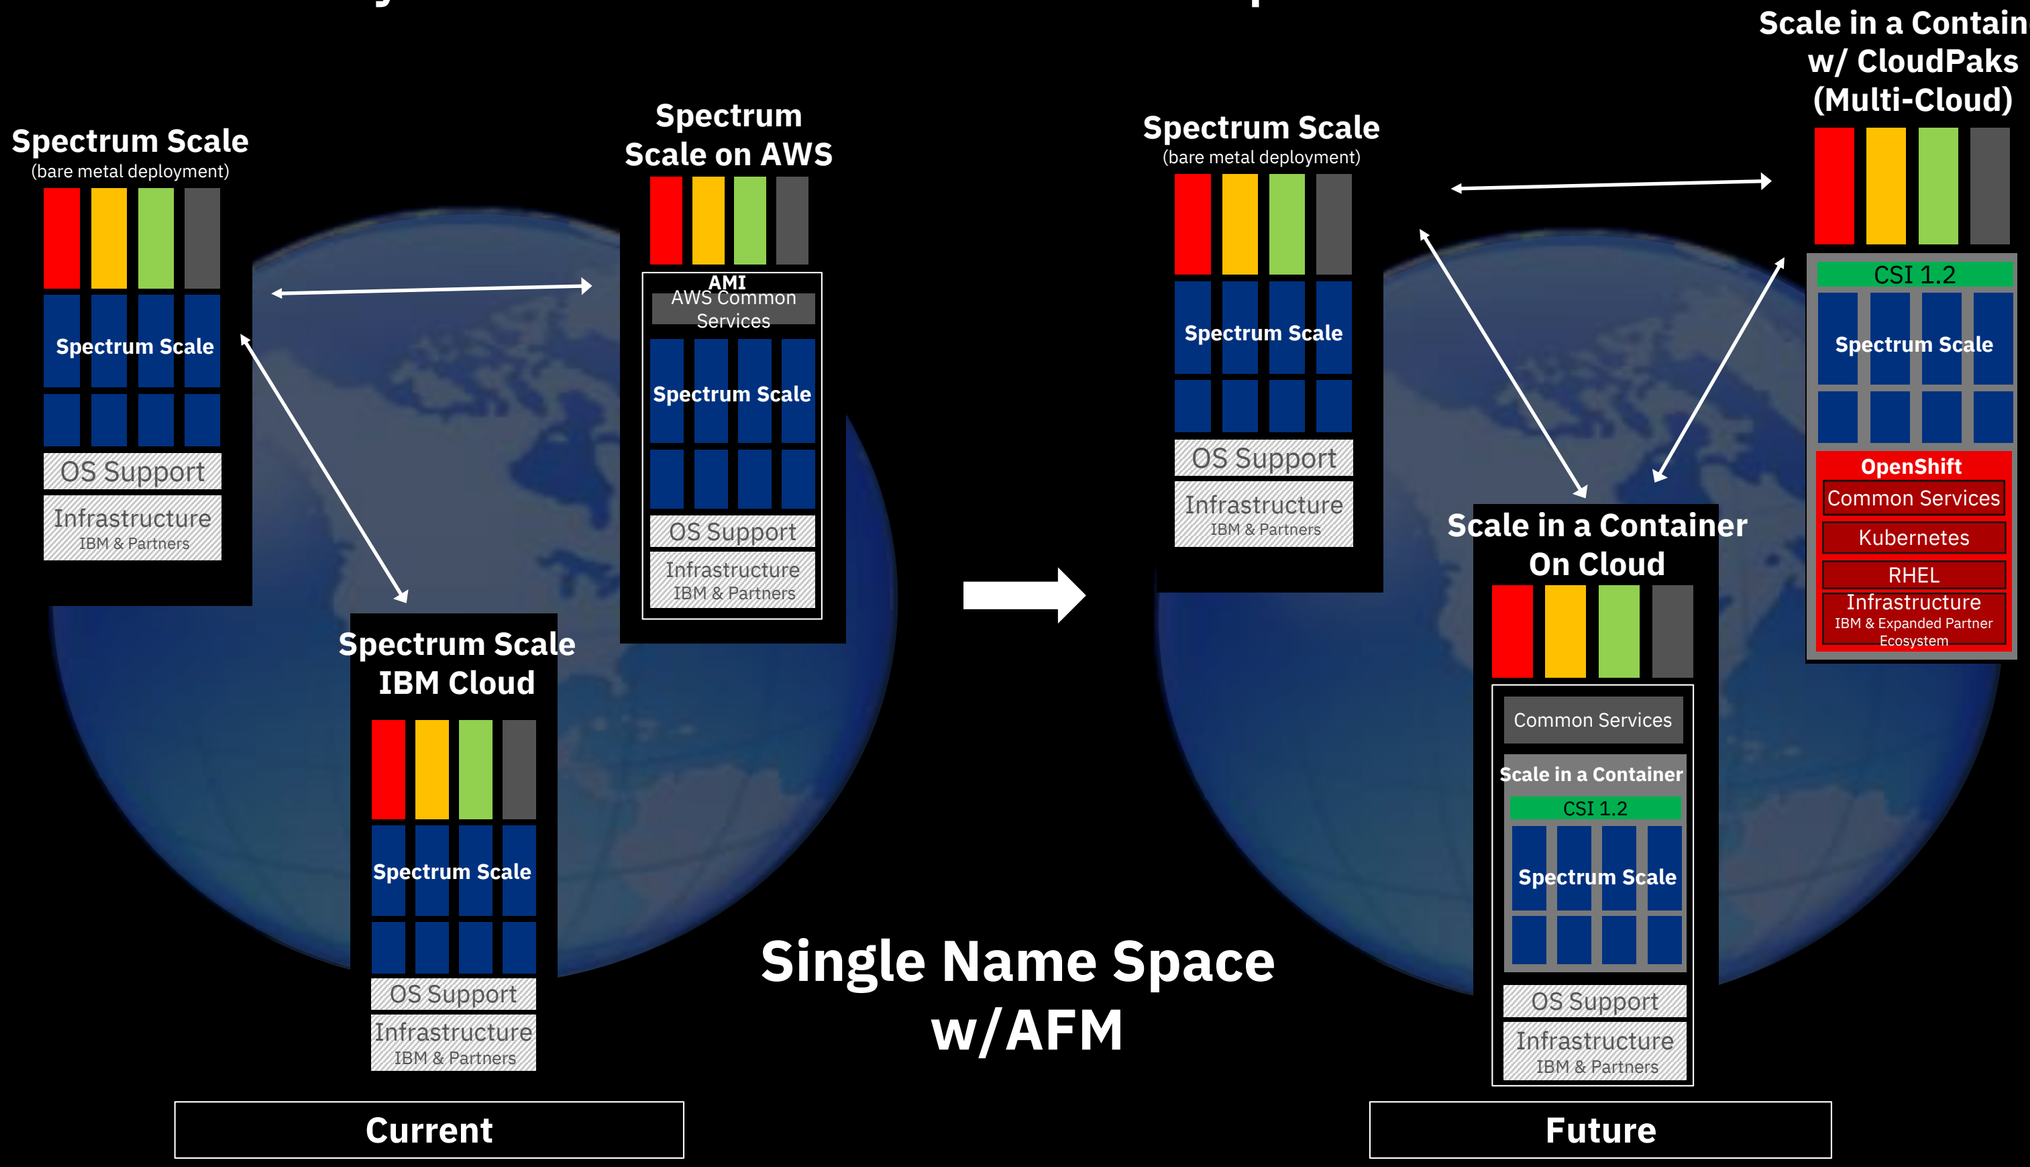

Outline

- AI Pipeline
- HDP goes Mainstream
- Containers
- **Storage for AI**
- Getting the data closer to the GPU

Start small and scale easily

from experiment to production at enterprise scale

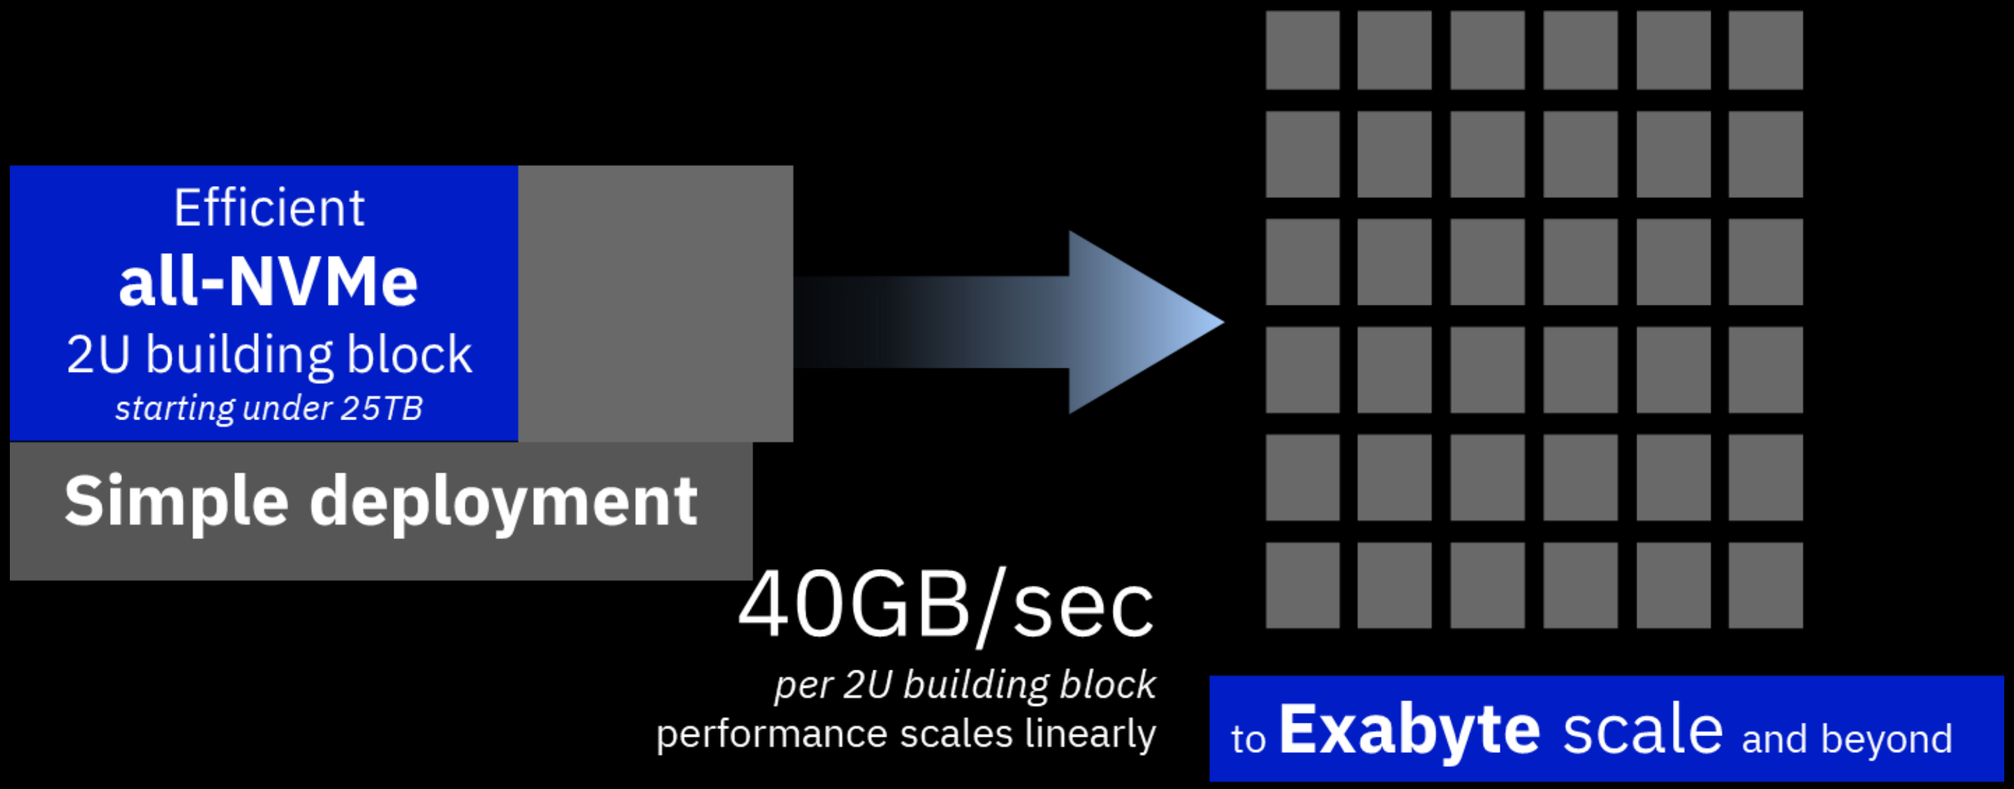

NVMe Flash for AI and Big Data Workloads

IBM Elastic Storage System 3000

All-new storage solution

- Leverages proven FS9100 technology
- Integrated scale-out advanced data management with end-to-end NVMe storage
- Containerized software for ease of install and update
- Fast initial configuration, update and scale-out expansion
- Performance, capacity, and ease of integration for AI and Big Data workflows

IBM Spectrum Scale

Outline

- AI Pipeline
- HDP goes Mainstream
- Containers
- Storage for AI
- **Getting the data closer to the GPU**

Data Accelerator for AI and Analytics Infrastructure

- **Performance Tier**
  - **Maximize performance** of storage: **$/IOP & $/GB/s are key**
    - Low latency random I/O & High bandwidth sequential
    - Relatively small compared to Capacity tier (say 5-25%)
    - Can be Lower Durability, Lower Availability, Lower Reliability,
      if Architected properly
    - No Geo-distribution

- **Capacity Tier** (aka “Data Lake”)
  - **Minimize** the **cost** of storage: **$/TB is key**
    - High Durability, Availability, Reliability, Geo-distribution

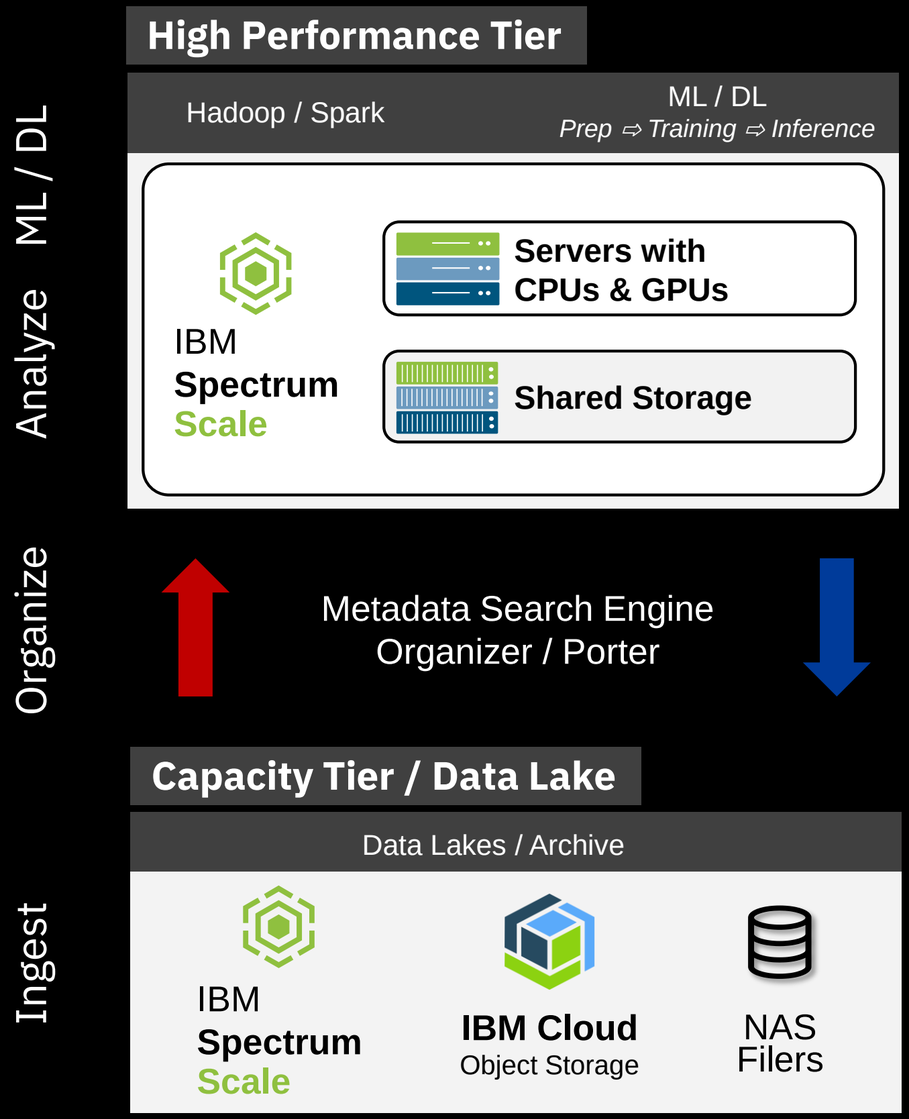

Accelerating data for NVIDIA GPUs

- NVIDIA Magnum IO is a collection of software APIs and libraries to optimize storage and network I/O performance in multi-GPU, multi-node processing environments. NVIDIA developed Magnum IO in close collaboration with storage industry leaders, including IBM.

- Collaboration with Nvidia continues to align Spectrum Scale's pagepool with GPU memory ("pagepool tiering")

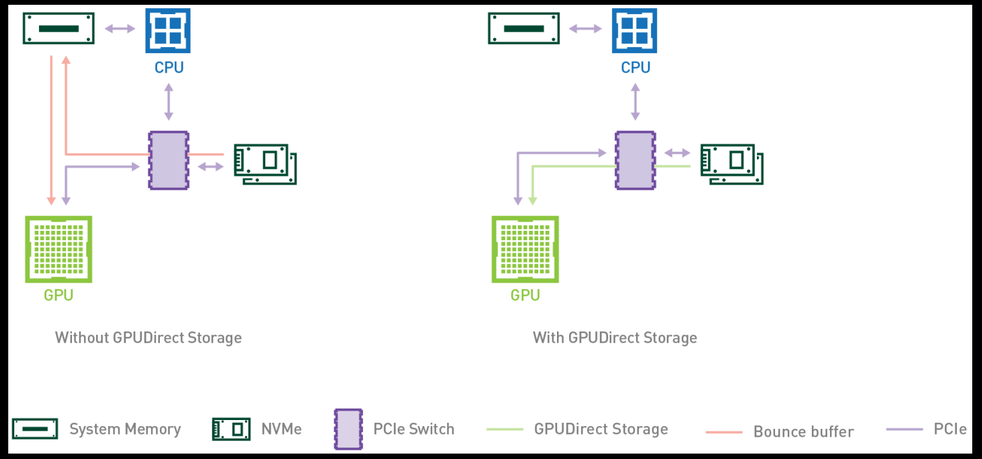

https://devblogs.nvidia.com/wp-content/uploads/2019/08/GPUDirect-Fig-1-New.png

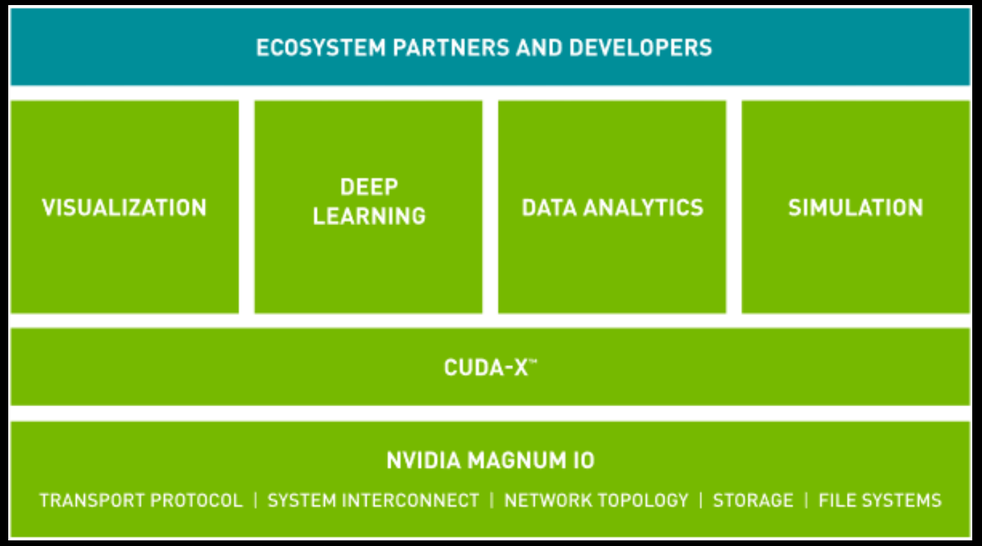

https://www.nvidia.com/en-us/data-center/magnum-io/

Questions?

Data Accelerator for AI and Analytics

Updated and Simplified IBM **Spectrum Scale** Picture