Elastic Storage System (ESS) Update
Spectrum Scale User Group - Germany

Mar 3, 2020

Doug Petteway

ESS in a Nut Schell

24TB to 13PB/rack

- IBM Integrated Solution for Spectrum Scale
  - Spectrum Scale for ESS
    - The latest Scale Release (Target within 1 Month)
    - ESS Unique Software RAID
      - Erasure Code, Checksum, Disk Hospital, . . .
  - Integrated Red Hat Enterprise Linux
  - Proven IBM Hardware: Servers, Storage (NVMe, SSD, HDD, NVDIMM)
    - Optional: Rackless systems, Switches, etc)
  - Comprehensive IBM Test & Integration
  - Award-Winning IBM Support

IBM Spectrum Scale

redhat

The Fastest Performance and Implementation of Spectrum Scale
#1 & #2 SuperComputer

Differentiator: Spectrum Scale RAID for ESS

Erasure Code software provides enterprise storage performance using standard and inexpensive disk drives

Faster rebuild times

- More disks are involved during rebuild
- Approximately 3.5x faster vs. RAID-5

Minimal impact of rebuild on system performance

- Rebuild is done by many disks
- Rebuilds can be deferred with enough protection

Better fault tolerance

- End to end checksum reduces / eliminates file system checks
- Stable, consistent high performance at all times
- Disk Hospital
- Expanded Call Home

Simplified Storage Portfolio – More ESS Focus

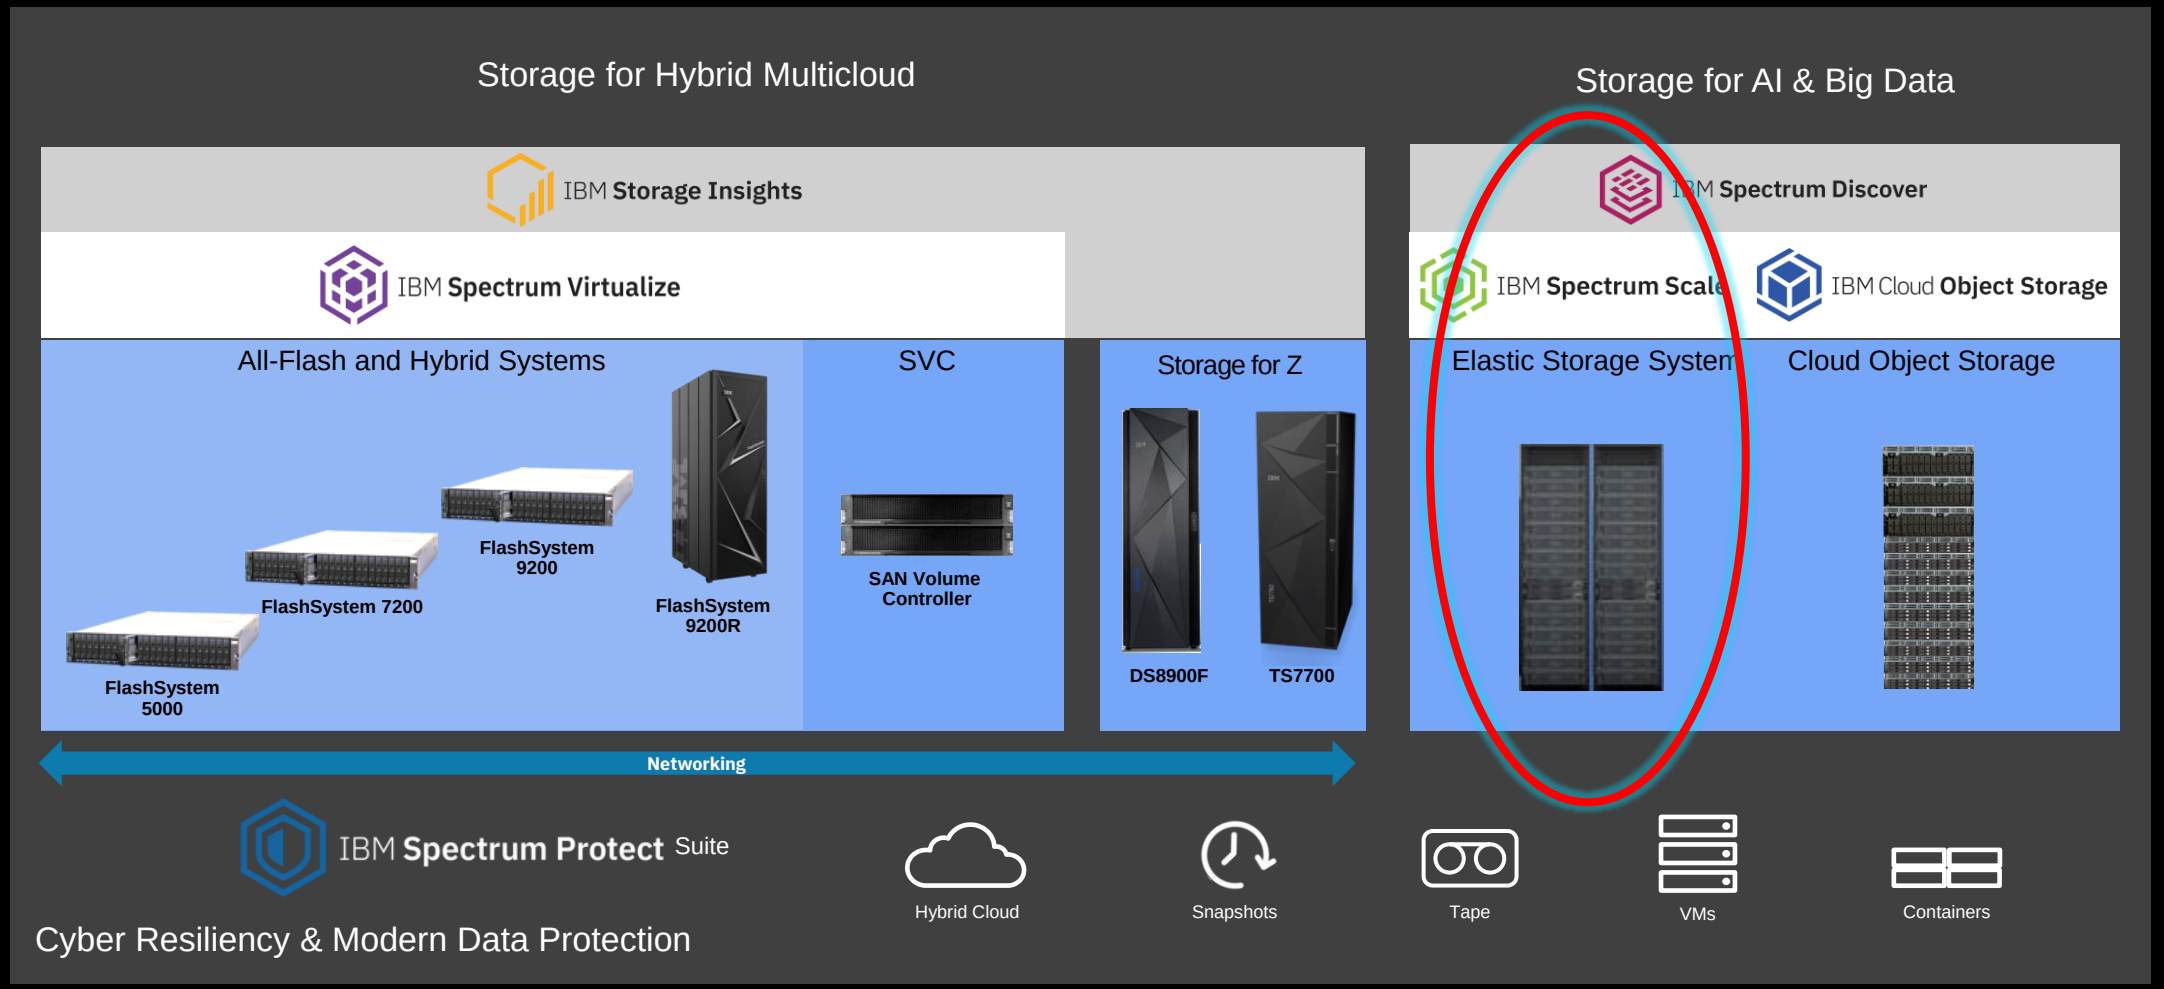

Spectrum Discover: Metadata management for an AI

Unified metadata and data insights for file and object storage on-premises and in the cloud

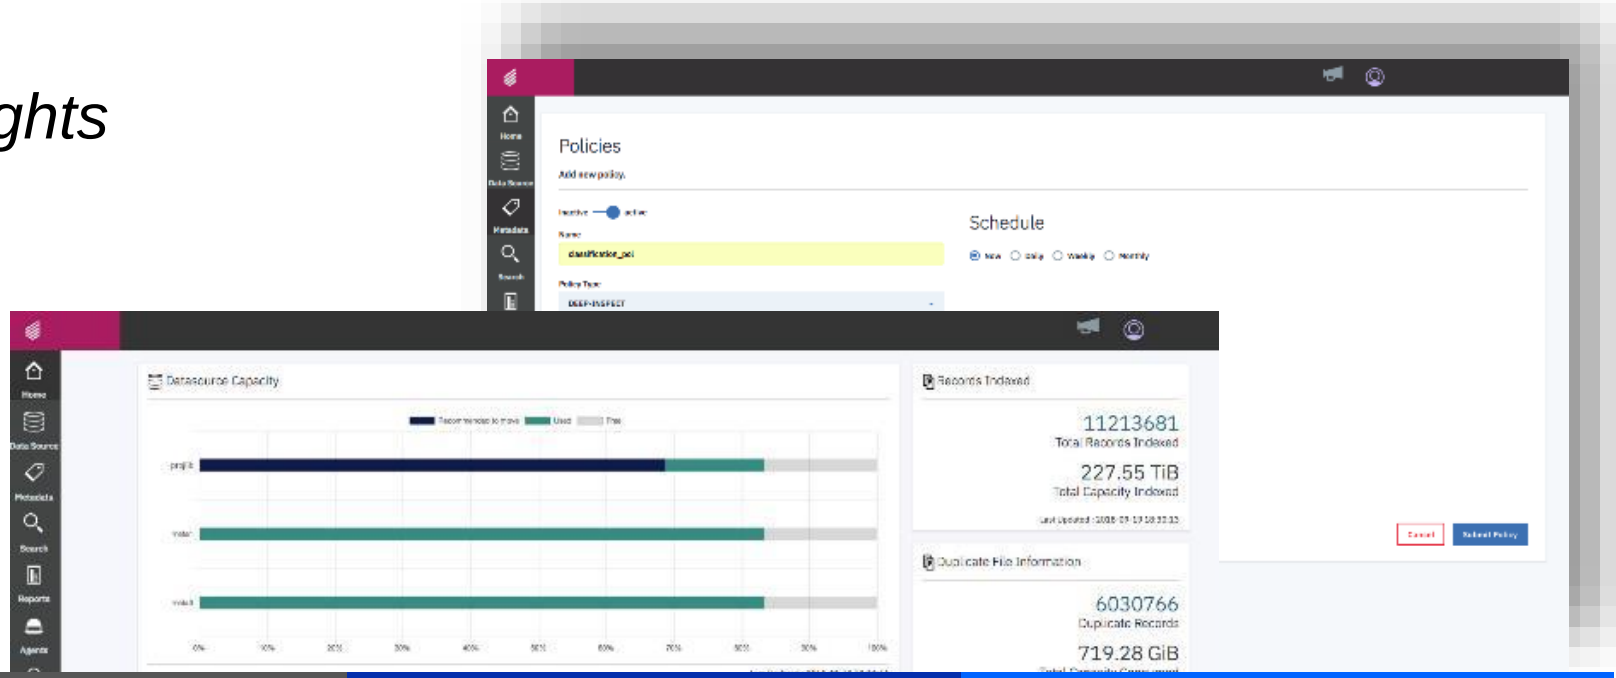

Discover

Automatically ingest and index system metadata from heterogeneous file and object storage systems on-prem and in the cloud

Classify

Automatically identify and classify data, including sensitive and personally identifiable information

Label

Enrich data with system and custom metadata tags that increase the value of that data

Find

Find data quickly and easily by searching catalogs of system and custom metadata

The Smarter and *Faster* Storage approach

The IBM ESS Family

Field proven, highly efficient and super reliable

Extreme
Performance & Scalability

- Over *1350* ESS purchased
- Over *300* ESS customers
- Over *4,000* Spectrum Scale clients
- Over *150%* Growth in 4Q19 YtY
- The leading Storage Solution for AI/ML and Analytics
- #1 and #2 Fastest Super Computer in the world
- Used by major Financial, Life Sciences, ADAS and Research & Government institutes across the world

#1 IBM is the world leader in Software Defined Storage Environments.

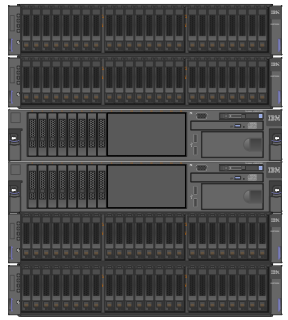

IBM Storage

Enabling enterprise data management for the future

Linux is the foundation for open. Red Hat is the leader.

54%

of public cloud apps run on Linux

#1

most used platform for development today

79%

of container hosts are Linux

71%

paid market share is Red Hat Enterprise Linux

#2

contributor to Kubernetes open source

The New ESS Architecture

Leverage the new integration with Red Hat

- Embedded Red Hat Enterprise Linux
  - Integrated and Streamlined
  - Everything we need and nothing we don't
- Integrate OpenShift technologies
  - Redefine ESS Deployment and Upgrade leveraging Containers (RHEL, Spectrum Scale, Scale RAID, IBM integration code)
  - Leverage Ansible playbook bundles (APBs) to create a standard mechanism for automating complex deployments.

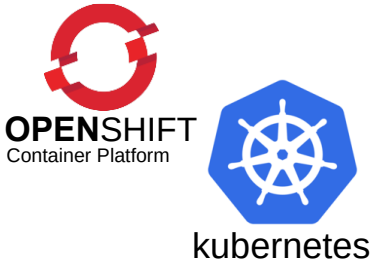

Simplified Deployment & Upgrade

High Performance for AI Training & Inference

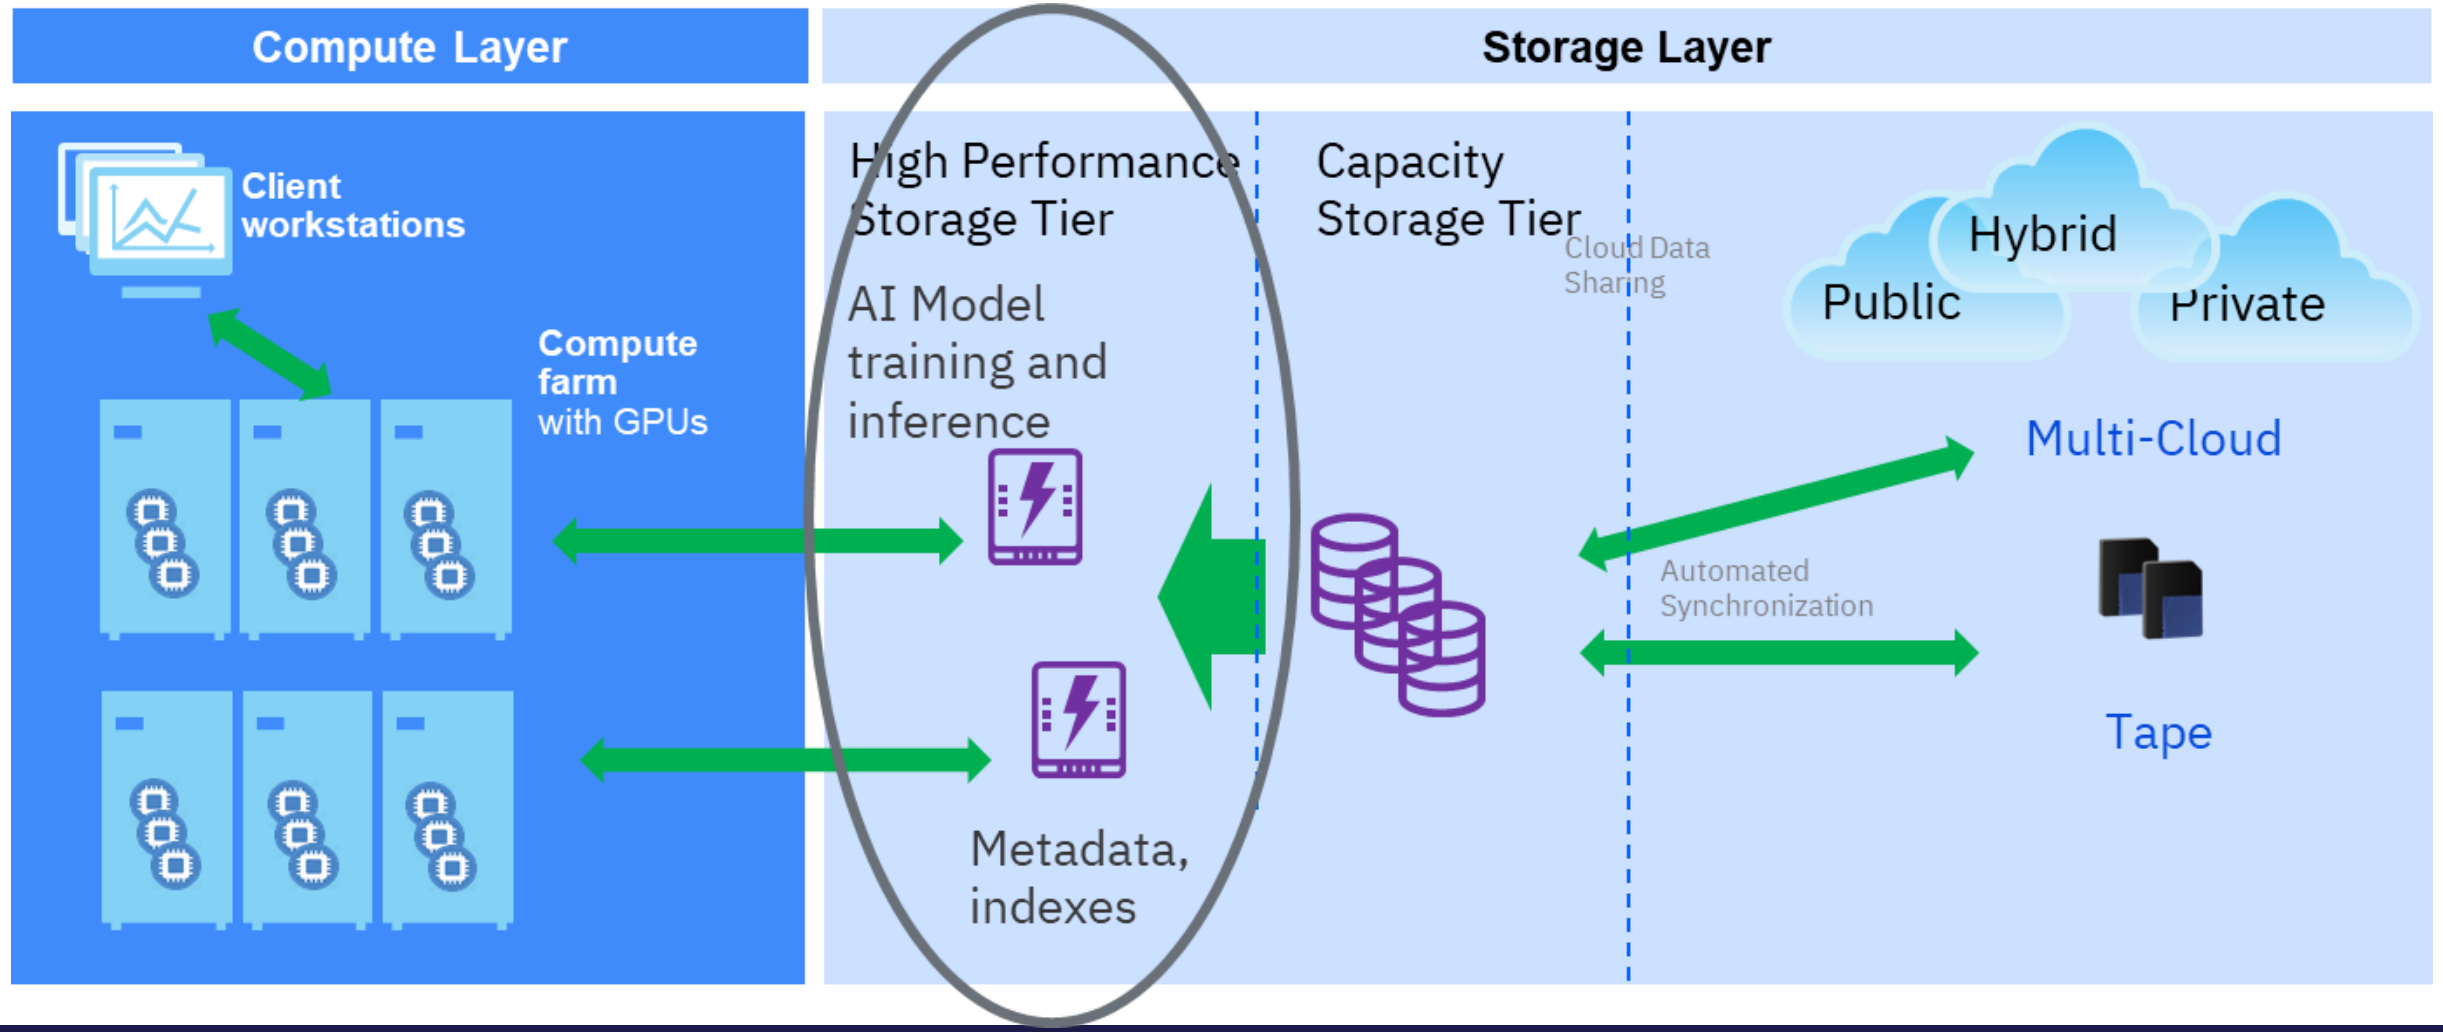

ESS 3000 Overview - *Now Shipping*

Core Design

- Self-Contained Storage Building Block
  - Dual active-active server running Spectrum Scale
  - 12 or 24 2.5-inch NVMe flash drives:
    (1.92 TB, 3.84 TB, 7.68 TB or 15.36 TB)*
  - Est: 40 GB/s sequential BW in a 2U enclosure
  - 280TB User capacity (8+2P)
  - 4 to 12 ports: IB EDR, 100GbE or mix
  - 384GB or 768GB memory per internal server

- Simple Install and Upgrade
  - Containerized Install/Upgrade
  - Embedded Red Hat OS

- Leverages proven Flash System technology
  - NVMe standard flash drives & IBM Flash Core FCM-2 (future)
  - Processing, NVMe Storage, RAS capability & dense packaging

*Raw capacity

Delivers:

Fast NVMe Storage
High Performance Storage Tier (HPT)

Fully Utilize Compute GPUs

Simple Installable Solution
IBM SSR delivery
Optional Lab Services

Fast Time-to-Value
hours vs. days/weeks

ESS 3000: Ideal Customers

- **AI & big data analytics Workloads** in Commercial or Federal accounts

- Clients who want to **Accelerate Scale Out business applications** with NVMe and IBM Flash technologies
- AI workloads requiring a High Performance Tier of storage to ensuring **full GPU resource utilization**
- Extend an **existing Spectrum Scale environment** with a rich set of capabilities and performance in existing clusters

Stronger Together: Enterprise AI

- #1 in GPUs, #1 in AI Compute Servers
- AI Solution orientation
- Mindshare with ML/DL developers and data scientists

- #1 in SDS, #1 in Tape, #1 in Analytics, #2 in Storage
- SDS and Fastest DGX I/O performance
- Mindshare with IT decision makers

NVIDIA and IBM Together

IBM SDS (SW Defined Storage) Solution Complements NVIDIA DGX
Global partner network in IBM Storage to match to Deep Learning partners
Acceleration of AI Adoption in the market
Broadest portfolio of data management solutions to complement DGX
Deep Learning partner and NVIDIA support

IBM Spectrum Storage for AI with NVIDIA DGX

- **Converged solution**
  - From 1/3 rack to largest supercomputers
  - IBM ESS 3000 + NVIDIA DGX + Mellanox

IBM leads in AI Infrastructure

- #1 in performance
- Multiple extensible SDS solutions to accelerate your data pipeline

ESS 3000 for AI with NVIDIA DGX

Near Linear Scaling by adding 40GB/s per 2U appliance

No need for downtime or reconfiguration

Best in class throughput potential

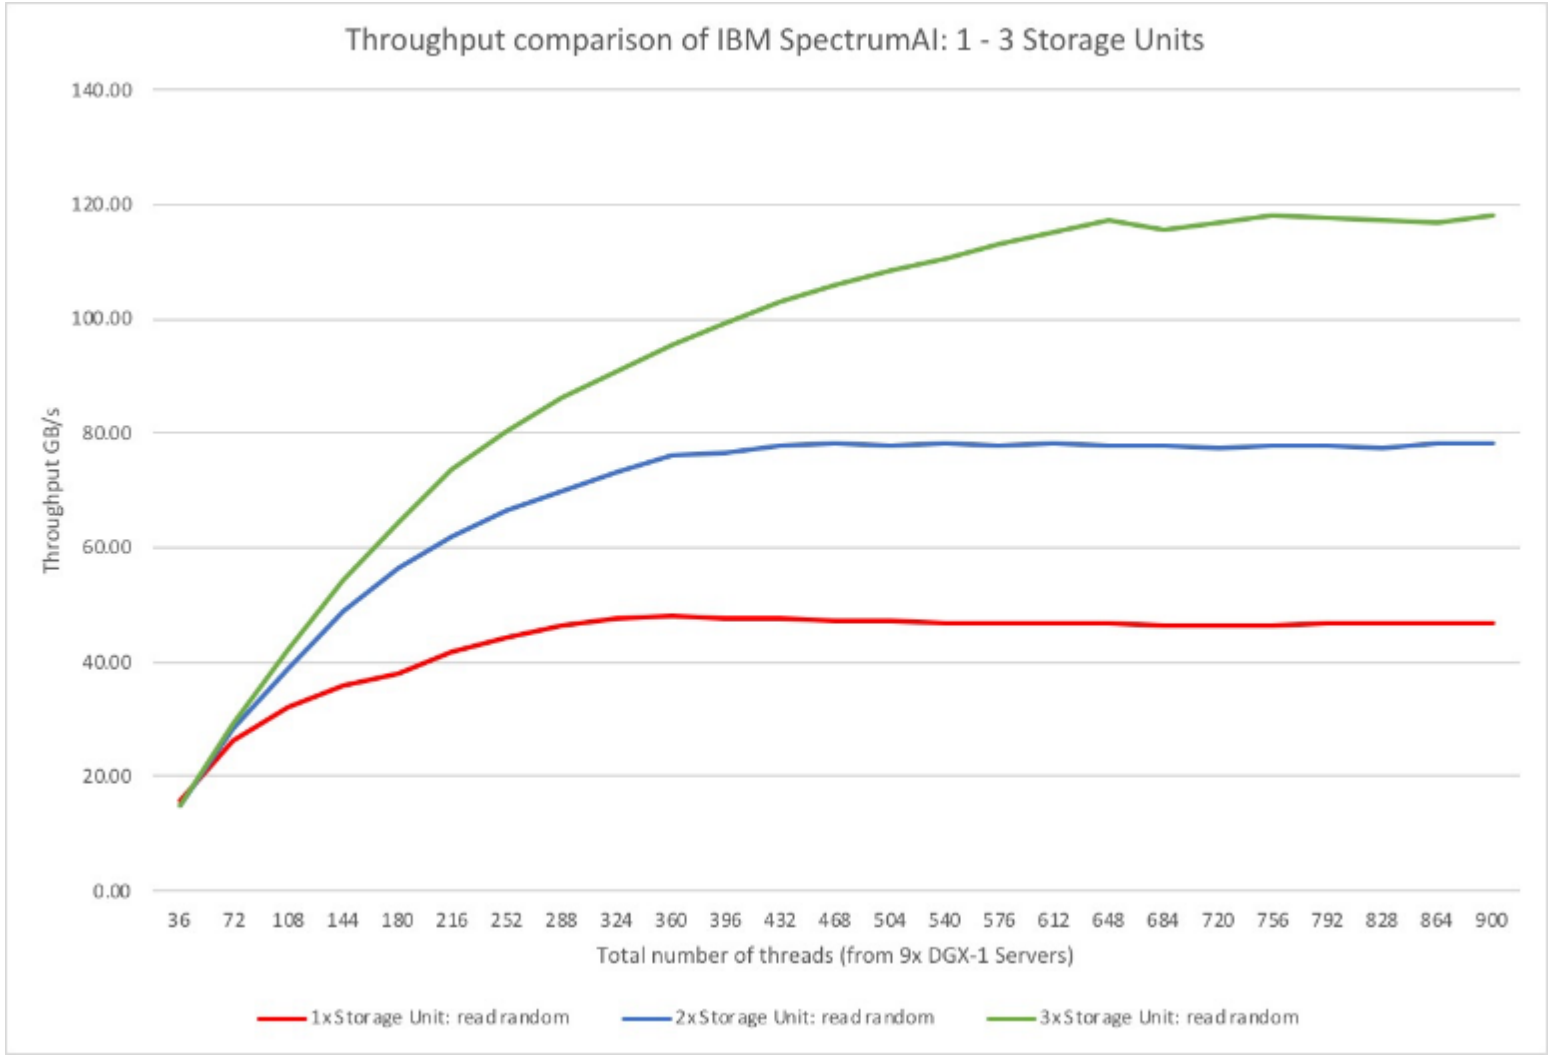

ESS 3000 for AI with DGX

Keeping GPUs busy regardless of IO access patterns

Headroom for data centric workload growth

Maximize GPU Utilization

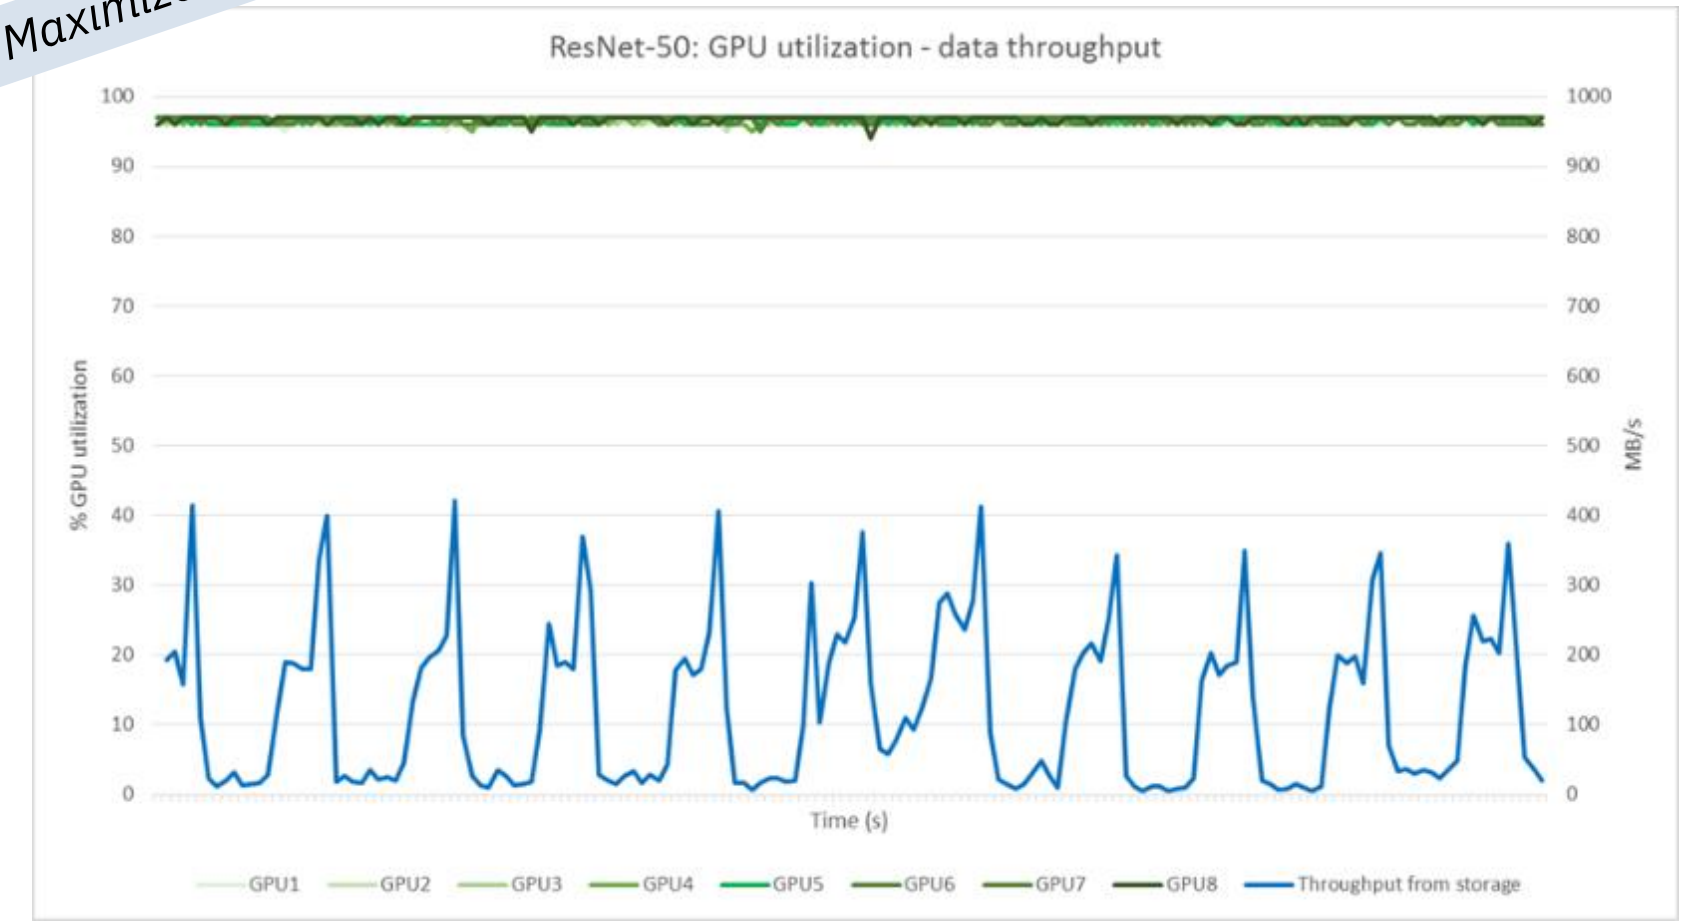

ESS Storage System Tiers

- Start Small and Grow as Needed
- One Management Server per Scale Cluster & Optional Protocol Nodes

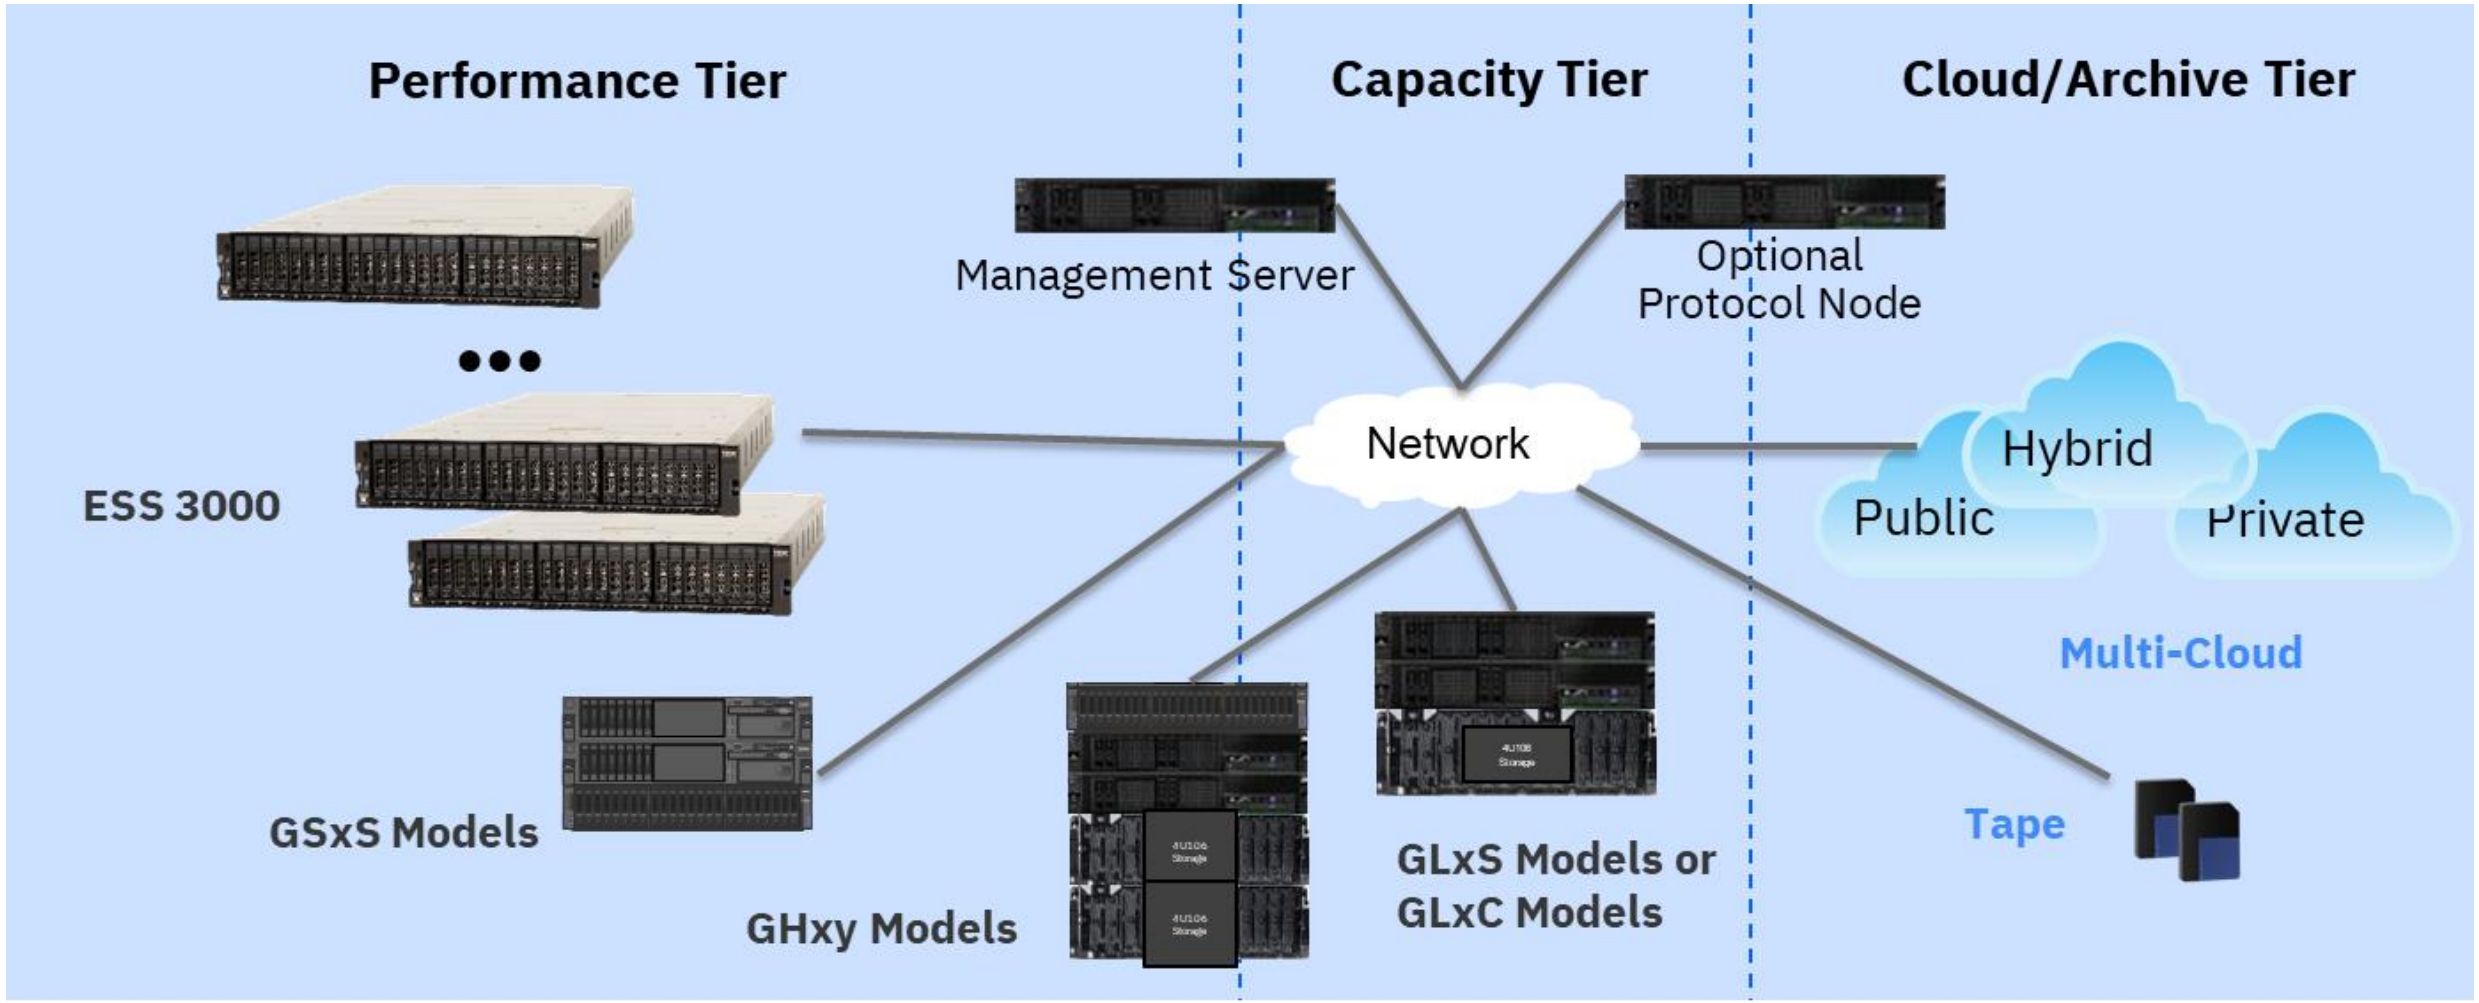

Back-up

Free! Introduction to IBM Elastic Storage Server and Spectrum Scale RAID

https://www.onlinedigitallearning.com/course/view.php?id=2173 (Log on with your IBM ID)

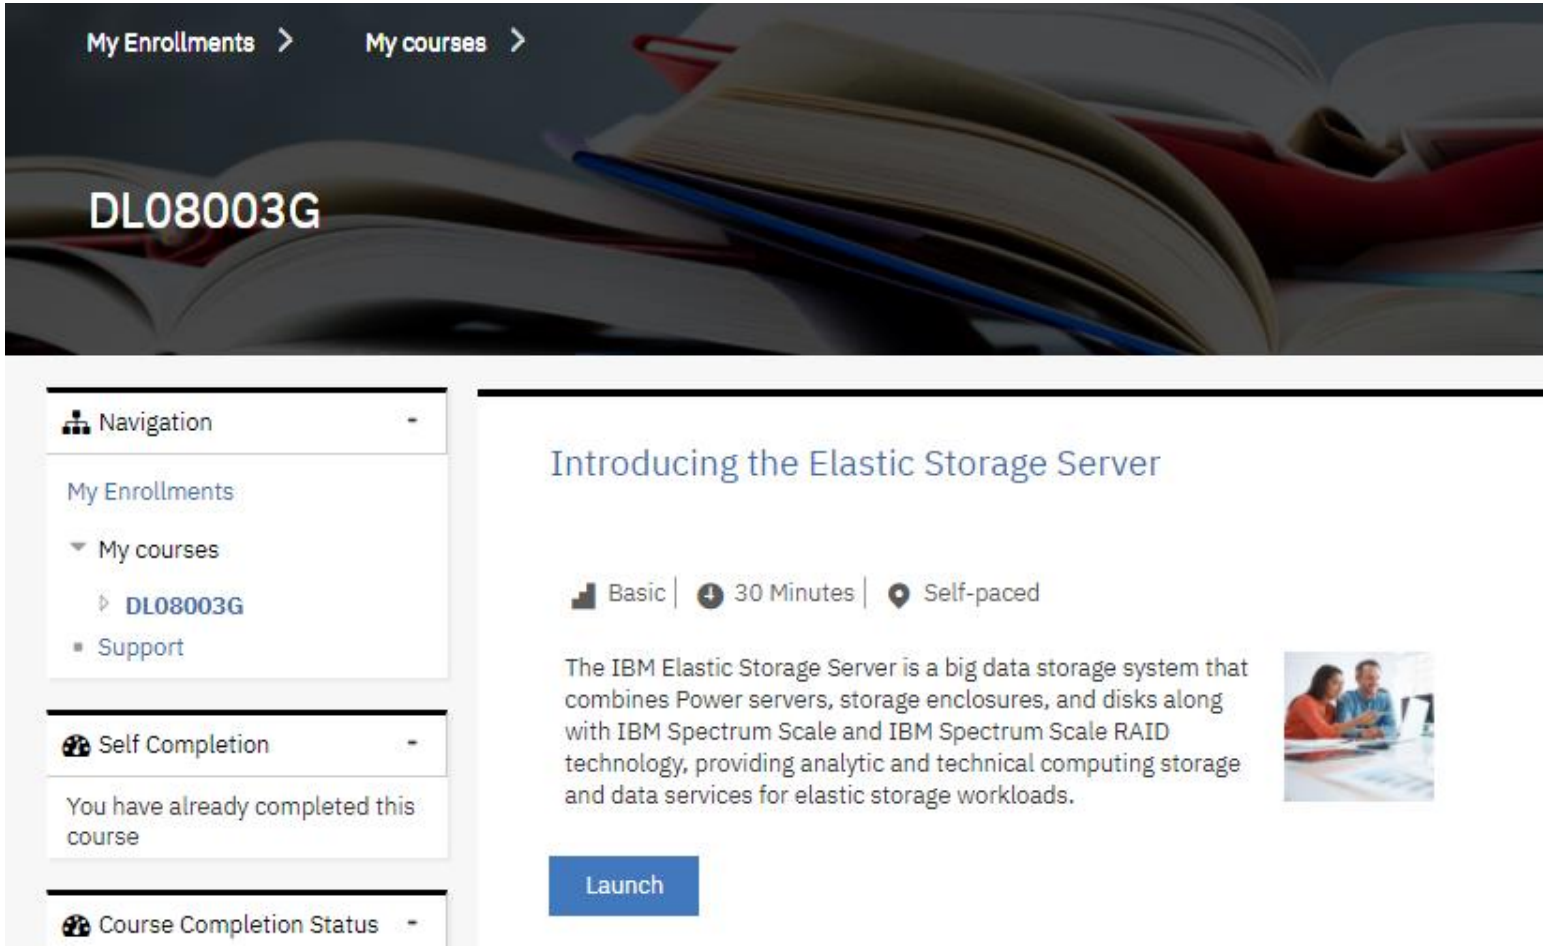

Free! - IBM Elastic Storage Server and Spectrum Scale RAID - gssutils

https://www.onlinedigitallearning.com/course/view.php?id=3570 (Log on with your IBM ID)

gssutils simulation for SSR Tools sub-menu

My Enrollments / My courses / DL08010G: SSR gssutils

About Simulation

This content is for installers or administrators for the IBM Elastic Storage Server.

You need 15 minutes or less to go through this content.

This simulation is for the **gssutils tool** that is pre-installed as part of the IBM Elastic Storage Server. The focus of this simulation is for the **SSR Tools s**
to use in order to verify the installation of an ESS.