

The background of the slide is a dark blue-tinted photograph of a city street with tall buildings and a road with lane markings.

# IBM Spectrum Scale IBM Elastic Storage Server

Andrew Beattie

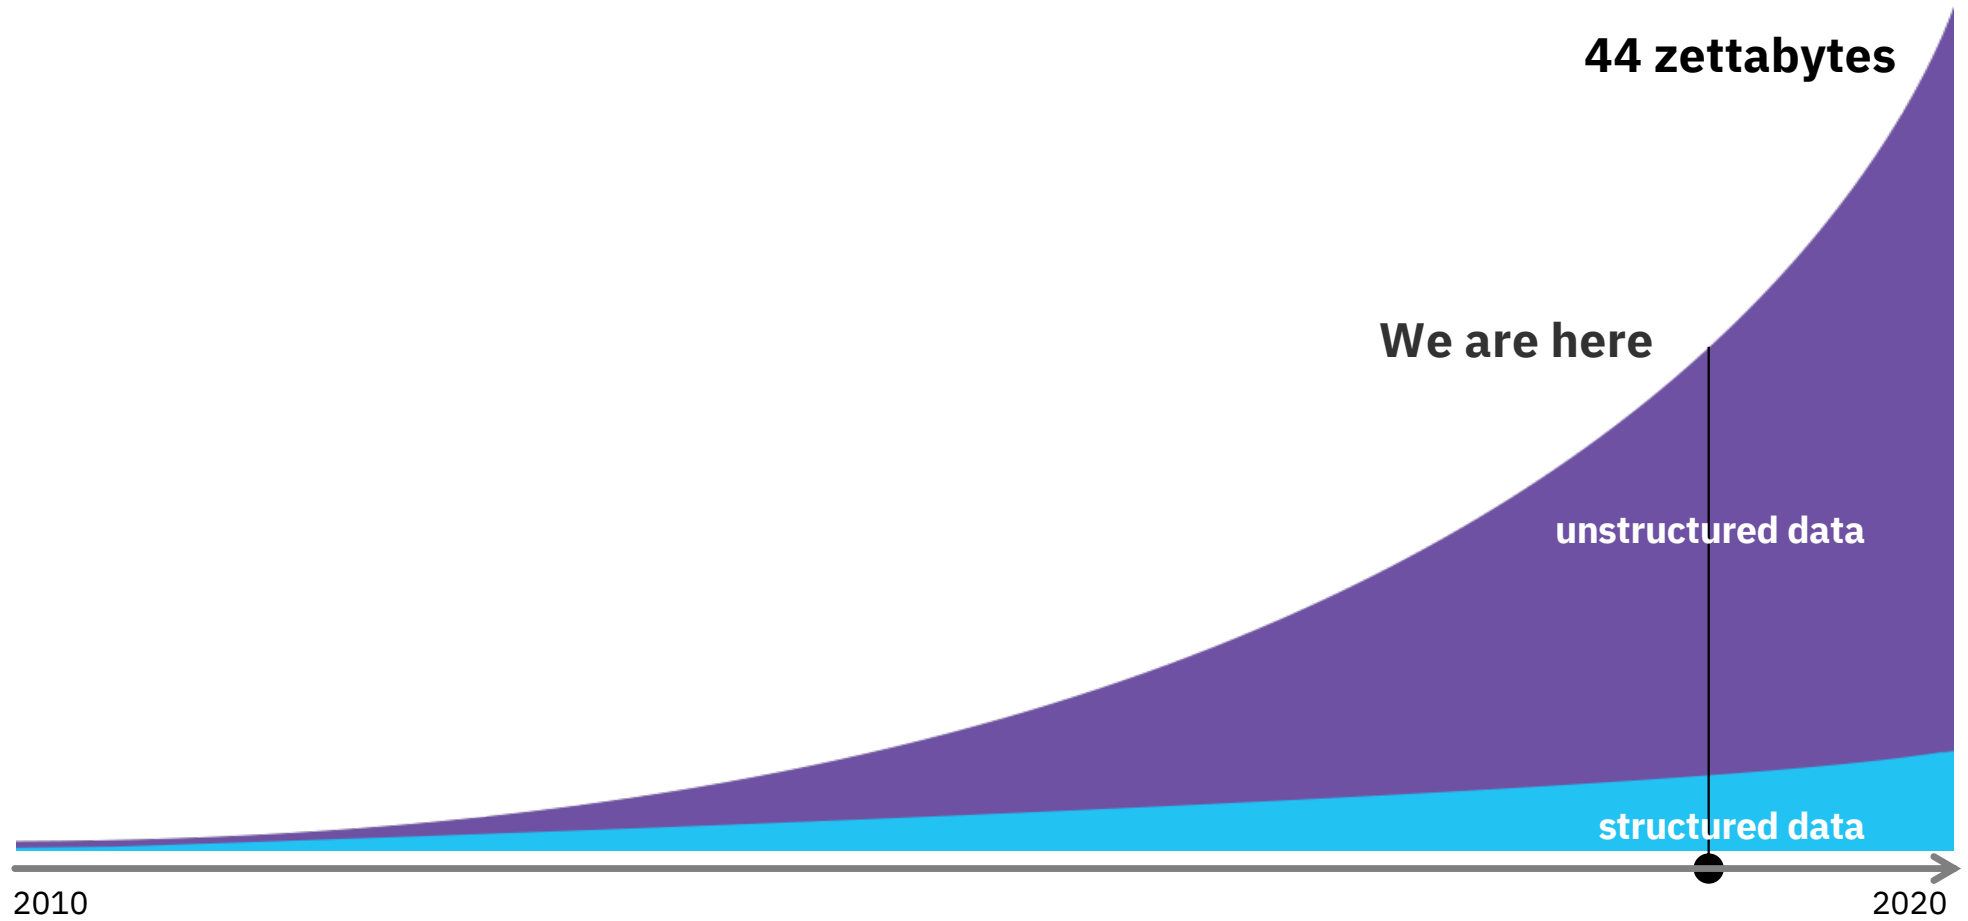
File and Object Storage – Technical Lead

IBM Systems – Storage


IBM A/NZ

# Data Growth is Exponential

*Demands New Technology and Strategy*



# Biggest Unstructured Data Challenges




Number of companies with 1,000 TB+ unstructured data stores **tripled**

Source: Forrester Survey, 2016 and 2017



**39%** of firms see sourcing, gathering, managing and governing data as their biggest challenge

Source: Forrester Survey, 2017



Data Scientists spend about **two-thirds** of their time finding data

Source: Gartner Market Guide, 2016

# IBM Software-Defined Storage portfolio



IBM  
Spectrum  
Storage

IBM's comprehensive set of award-winning storage software delivered across appliance, converged and cloud.



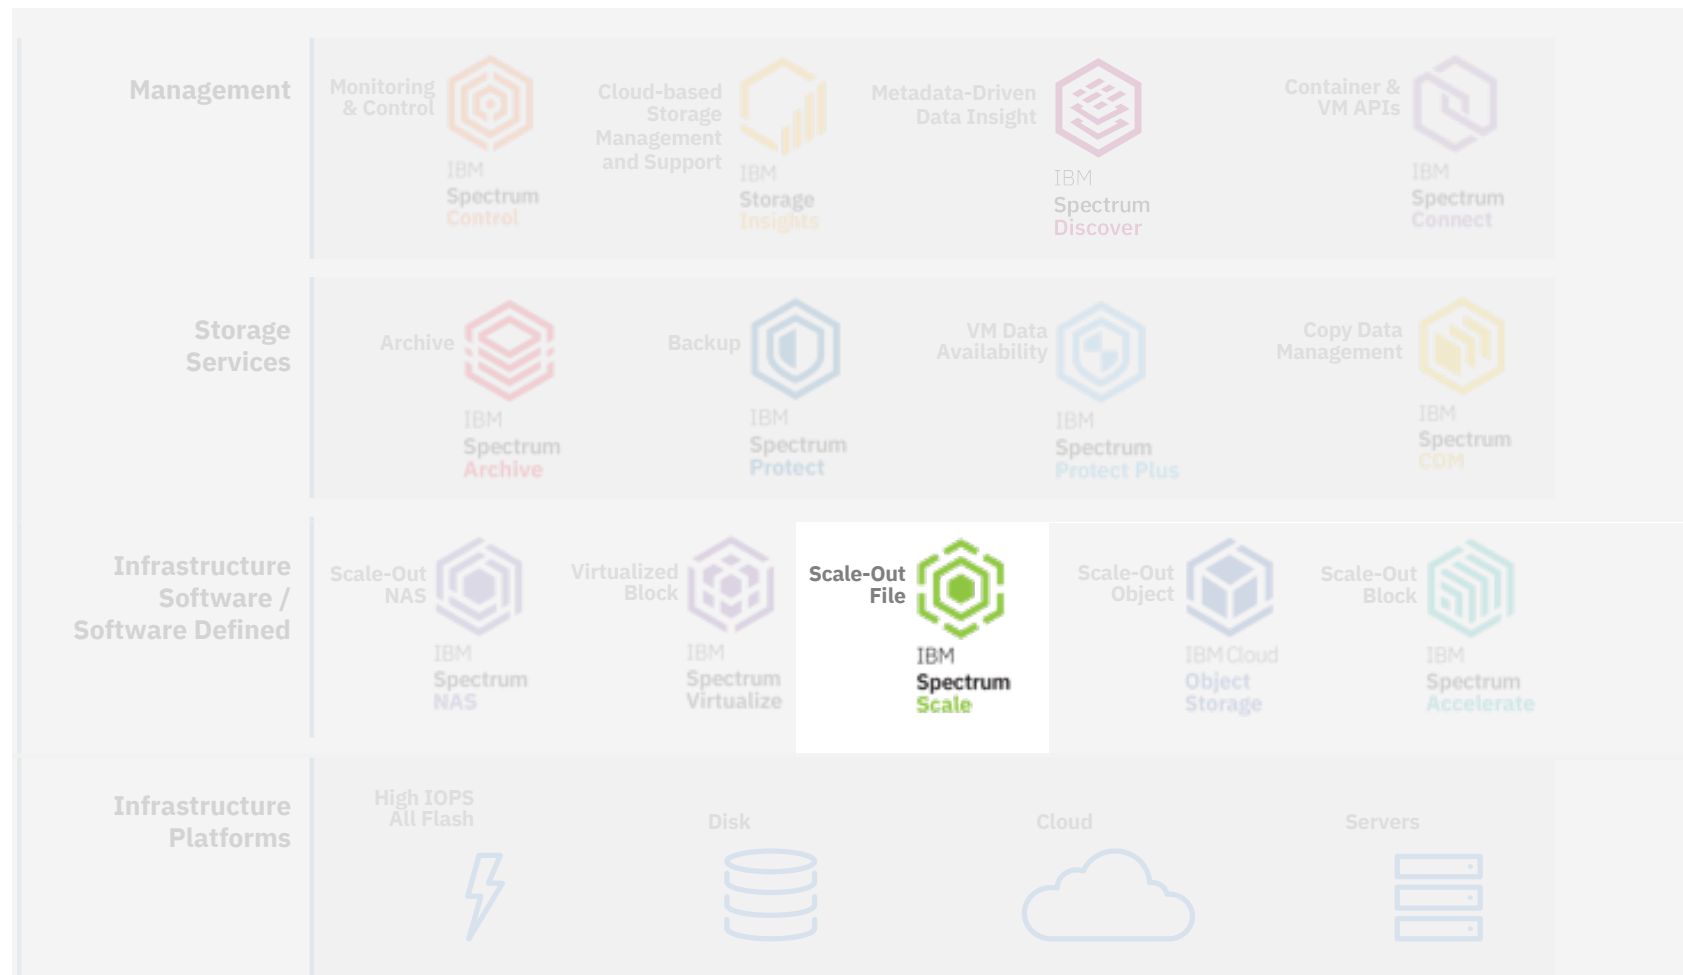
# IBM Spectrum Scale

## Store everywhere. Run anywhere.



IBM  
Spectrum  
Storage

IBM's comprehensive set of award-winning storage software delivered across appliance, converged and cloud.



# IBM Spectrum Storage Client Adoption

## 87 of the Fortune Global 100 use IBM Spectrum Storage

- 49 of the top 50 banks
- 14 of the top 15 telecommunications companies
- 18 of the top 20 energy companies
- 9 of the top 10 global healthcare companies
- 8 of the top 10 automobile manufacturers

**80 clients pick IBM Spectrum Storage every week**



# IBM Spectrum Scale



# Spectrum Scale simplifies data management at scale



## Those that want to get value from their data choose IBM Spectrum Scale

Grow and share the storage infrastructure while *automatically moving file and object data to the optimal storage tier as quickly as possible.*

Store **Everywhere.** Run **Anywhere.**



# Introducing IBM Spectrum Scale

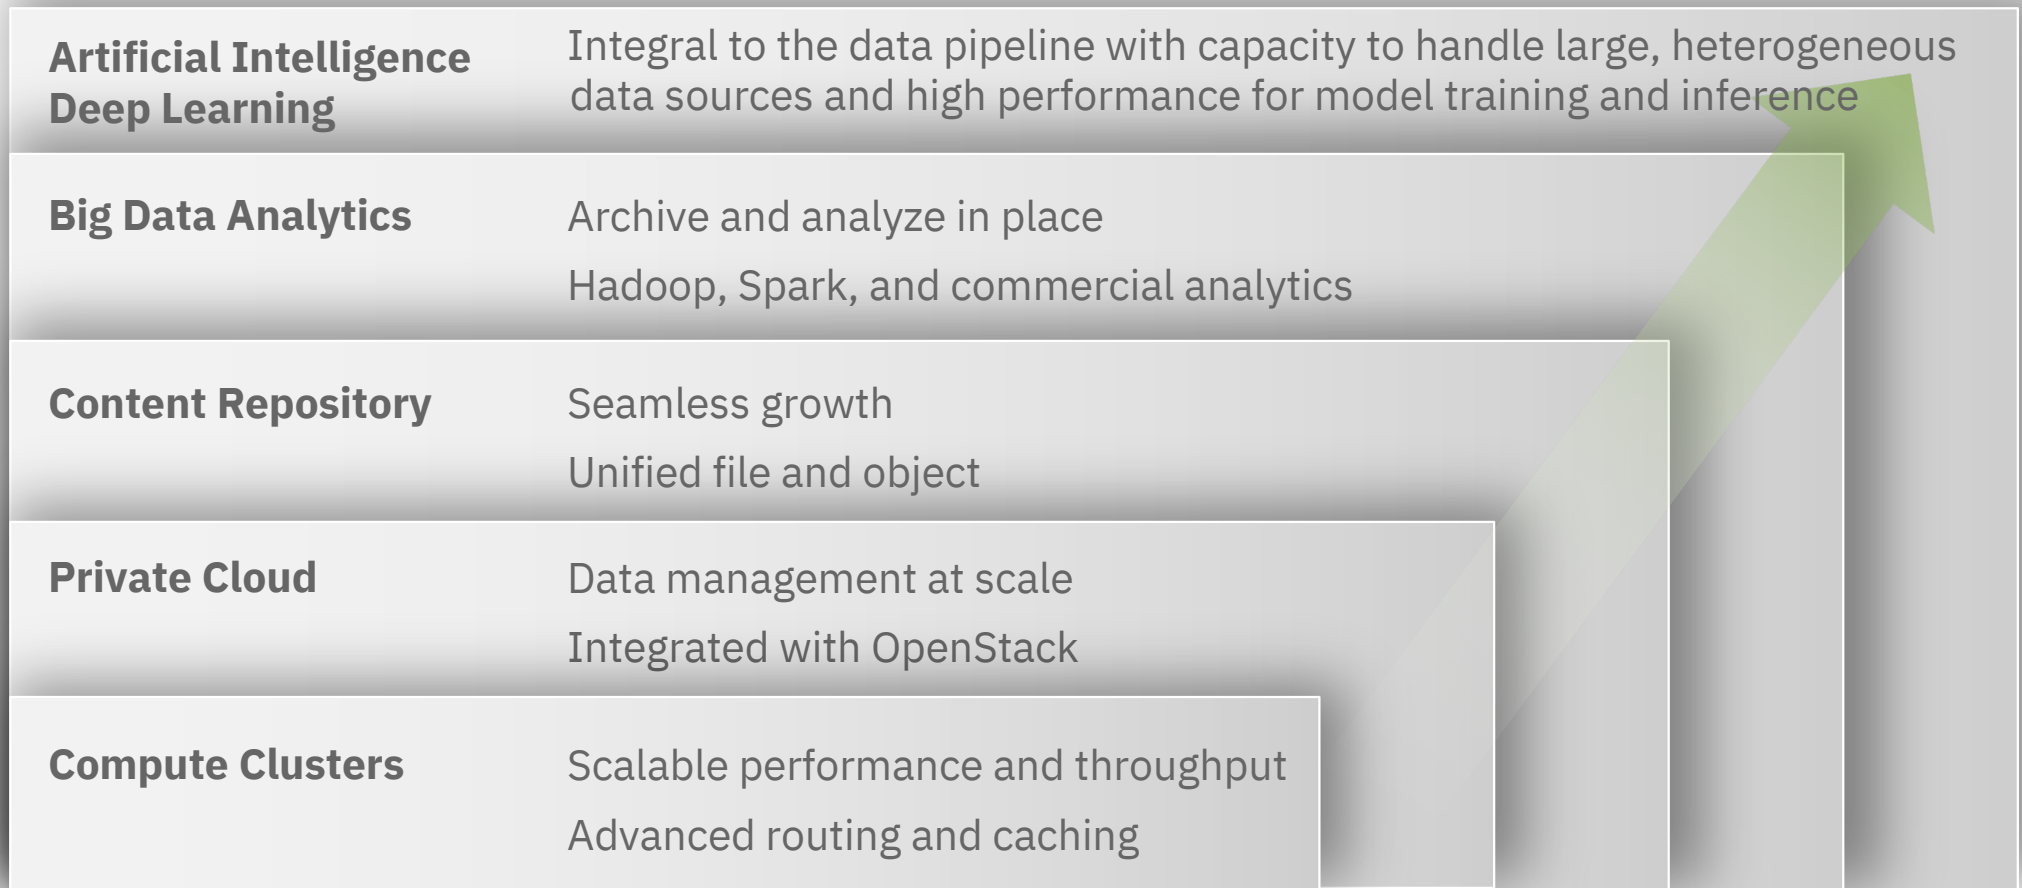


**Highly scalable high-performance unified storage**  
for files and objects with integrated analytics

- **Remove data-related bottlenecks**  
2.5TB/s demonstrated throughput
- **Enable global collaboration**  
HDFS, files and object across sites
- **Optimize cost and performance**  
Automated data placement
- **Ensure data availability, integrity and security**  
End-to-end checksum, Spectrum Scale RAID, NIST/FIPS certification



# IBM Spectrum Scale for common workloads



# IBM Software Defined Infrastructure delivers value across industries



## Compute Clusters




### Nuance Communications

Runs live customer-interaction environments for its clients and refines the speech models that power them

**20% boost in storage performance** run 20 million jobs per month accessing over 600TB of data per day

## Global Storage Cloud




### Cypress Semiconductor

Eliminated data access bottlenecks and has **increased performance 10x** using the same hardware

Continuous data availability across hardware outages

## Content repository



### Caris Life Sciences

Correlates molecular data for 65,000 patients and supporting 7,000 oncologists worldwide

**Manages nearly a terabyte of data** per patient enabling precision cancer treatment

## Big Data Analytics



### Citi

**100X performance improvement** combined with on-demand access to compute power drastically speeds time to results

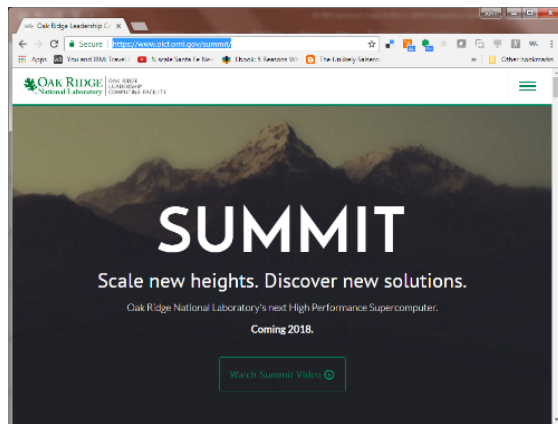
Computing resources used more effectively with **hardware utilization increasing from 20% to 80%**

# Performance engineering matters



## Imagine you need to meet these goals:

- 2.5 TB/sec single stream IOR as requested from ORNL
- 1 TB/sec 1MB sequential read/write as stated in CORAL RFP
- Single Node 16 GB/sec sequential read/write as requested from ORNL
- 50K creates/sec per shared directory as stated in CORAL RFP
- 2.6 Million 32K file creates/sec as requested from ORNL



## IBM Spectrum Scale innovations have have delivered these requirements

<https://www.olcf.ornl.gov/summit/>

# IBM Spectrum Scale V5



**New version of highly scalable high-performance unified storage** for files and objects with integrated analytics

*Rebuilt from the file system up to handle new workloads on flash from the densest storage to the largest scale*

- **A new level of storage performance and efficiency**  
Dramatic improvements in I/O performance
- **Enhanced security and compliance**  
Integrated file audit logging capability
- **Simpler, more powerful system administration**  
Faster and simpler out-of-the-box experience  
Enhanced GUI features for many capabilities



# Spectrum Scale v5.0



## A new level of storage performance and efficiency

### Dramatic improvements in I/O performance

#### Support for newest low-latency, high bandwidth hardware such as NVMe

- Significantly improved communication speed between nodes

#### Improved performance, space efficiency for mixed workloads

- Small and large block size workloads running simultaneously in same file system
- Optimize large block performance via new 4MB default block
- Simultaneously optimize small-file space efficiency with variable sub-block size

#### Improved IOP/sec and metadata performance

**IOP/s can improve 3x to 5x over previous releases\***

\* Performance numbers are estimates based on IBM internal lab tests and are subject to verification



Silverton Consulting White Paper:  
[IBM Spectrum Scale 5.0.0. IO performance](#)

# Spectrum Scale v5.0

## Improved security and compliance



### **New File Audit Logging capability with Data Management Edition**

- Track user accesses to filesystem and events
- Supported across all nodes and all protocols
- Parseable data stored in secure retention-protected fileset
- Events that can be captured are:
  - Open, Close, Destroy (Delete), Rename, Unlink, Remove Directory, Extended Attributed Change, Access Control List (ACL) change

### **File-level immutability**

- Independent KPMG certification

### **Data security following removal of physical media**

- Data Management Edition data protected by on-disk encryption

### **Protocols include encryption of data in motion**



# Spectrum Scale v5.0

## Simpler, more powerful system administration



### **Faster and simpler out-of-the-box experience**

- Easier setup and improved performance with less manual configuration.
- Setup for optimal performance aided with additional 20+ parameters now handled automatically
- Manual override of the automated settings is readily available for users who prefer to select their own settings

### **Enhanced GUI features for many capabilities**

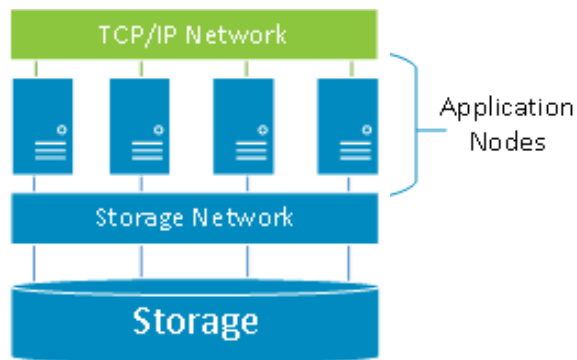
- Performance, capacity, network monitoring
- AFM (multi-cluster management)
- Transparent Cloud Tiering
- Enhanced maintenance and support, including interaction with IBM remote support
- Network troubleshooting






# Spectrum Scale deployment models

## Enterprise Integrated Model




Unify and parallelize storage silos

## Network Shared Disk (NSD) Model



Modular High-Performance Scaling

## Shared Nothing Cluster (SNC) Model – Enhanced with Scale Native Raid



Span storage rich servers for converged architecture or HDFS deployment



# Use Cases for IBM Spectrum Scale



## New Generation Workloads

AI, Machine Learning and Deep Learning  
Big Data and Analytics  
Selected Solutions

- ISV Solutions – SAS Grid, SAP HANA
- Industry Solutions
  - *Healthcare, Genomics*
  - *Video, streaming media, media and entertainment*

## Modernize and Transform

“Data Ocean” Unified Storage  
Data-intensive Technical Computing

## Lifecycle Management

Archive of data - across disk and tape  
Information Life Cycle Management

## Resiliency

High Performance Back-up / Restore



IBM  
**Spectrum**  
**Scale**

SAS Grid on Spectrum Scale/ESS:

<https://ibm.box.com/s/d7z8poi3eoymge60unjopi912xhforpr>

# Storage for the world's most powerful supercomputers



World's most powerful supercomputer



## Summit System

- **4608 nodes**, each with:
  - 2 IBM Power9 processors
  - 6 Nvidia Tesla V100 GPUs
  - 608 GB of fast memory
  - 1.6 TB of NVMe memory
- **200 petaflops** peak performance for modeling and simulation
- **3.3 ExaOps** peak performance for data analytics and AI

## IBM Spectrum Scale IBM Elastic Storage Server

**2.5 TB/sec** throughput to storage architecture  
**250 PB** HDD storage capacity



World #2 supercomputer

## Sierra System

- **4320 nodes**, each with
  - 2 IBM Power9 processors
  - 4 Nvidia V100 GPUs
  - 320 GB of node memory
  - 1.6 TB of NVMe memory
- **IBM Spectrum Scale**
- **IBM Elastic Storage Server**

**125 petaflops** peak performance  
**154 PB** HDD storage capacity

# IBM Spectrum Scale Value



## Storage management at scale

- Powerful GUI and health monitoring
- Unified File, Object and HDFS
- Distributed metadata and high-speed scanning
- QoS management
- 1 Billion Files and yottabytes of data
- Multi-cluster management with Spectrum Control

## Store everywhere. Run anywhere.

- Advanced routing with latency awareness
- Read or Write Caching
- Active File Management for WAN deployments
- File Placement Optimization
- End-to-end data integrity
- Snapshots
- Synchronous or Asynchronous DR

## Improve data economics

- Tier seamlessly
- Incorporate and share flash
- Policy driven compression
- Data protection with erasure code and replication
- Native Encryption and Secure Erase compliance
- Target object store and cloud
- Leading performance for Backup and Archive

## Software Defined Open Platform

- Heterogeneous commodity storage: flash, disk and tape
- Software, appliance or Cloud
- Data driven migration to practically any target
- File/Object In/Out with OpenStack SWIFT and S3
- Transparent native HDFS
- Integration with cloud

# Get It Your Way



Software



Integrated Solution



Cloud service

# IBM Spectrum Scale on Amazon Web Services (AWS)

## Bring Your Own License



### **Easily deploy a high-availability scalable cluster filesystem on AWS infrastructure**

- For Scale customers: exploit AWS elasticity for your workload
- For AWS customers: an HPC-capable file system

### **Purchase and deployment**

- Deploy through AWS Marketplace (same experience as any other AWS offering)
- IBM-provided Cloud Formation script
  - Deploys a cluster of EC2 instances (Amazon virtual servers)
  - Deploys Spectrum Scale across the cluster
  - Mounts the Scale filesystem across the instances
- IBM Spectrum Scale Support from IBM through normal channels
  - AWS infrastructure support from Amazon. If in doubt, call IBM first.

### **Native AWS cloud experience**

- Cluster expansion and limited contraction
- Shutdown and restart of cluster without removing any Amazon storage volume resource
- Collection of AWS cluster management debug info

# IBM Elastic Storage Server

# IBM Elastic Storage Server (ESS)



## Integrated scale-out data management for file and object data

### Optimal building block for high-performance, scalable, reliable enterprise Spectrum Scale storage

- Faster data access with choice to scale-up or out
- Easy to deploy clusters with unified system GUI
- Simplified storage administration with IBM Spectrum Control integration

### One solution for all your Spectrum Scale data needs

- Single repository of data with unified file and object support
- Anywhere access with multi-protocol support: NFS 4.0, SMB, OpenStack Swift, Cinder, and Manila
- Ideal for Big Data Analytics with full Hadoop transparency

### Ready for business critical data

- Disaster recovery with synchronous or asynchronous replication
- Ensure reliability and fast rebuild times using Spectrum Scale RAID's dispersed data and erasure code
- Five 99999s of availability






# IBM Elastic Storage Server: Speed and Capacity




## Speed

**Model GS1S**  
24 SSD




14 GB/s

**Model GS2S**  
48 SSD




26 GB/s

**Model GS4S**  
96 SSD



40 GB/s

**Model GH12:**  
1 2U24 Enclosure SSD  
2 5U84 Enclosure HDD  
166 NL-SAS, 24 SSD



18 GB/s\*

**Model GH14:**  
1 2U24 Enclosure SSD  
4 5U84 Enclosure HDD  
334 NL-SAS, 24 SSD



38 GB/s\*

**Model GH24:**  
2 2U24 Enclosure SSD  
4 5U84 Enclosure HDD  
334 NL-SAS, 48 SSD



40 GB/s\*

\* Estimate of performance aggregated across SSD and HDD. All estimates assume EDR Infiniband connections

## Capacity

**Model GL1S:**  
1 Enclosures, 9U  
82 NL-SAS, 2 SSD



6 GB/s

**Model GL2S:**  
2 Enclosures, 12U  
166 NL-SAS, 2 SSD



12 GB/s

**Model GL4S:**  
4 Enclosures, 20U  
334 NL-SAS, 2 SSD



24 GB/s

**Model GL6S:**  
6 Enclosures, 28U  
502 NL-SAS, 2 SSD



36 GB/s


# IBM Elastic Storage Server: Hybrid Models

**Model GL1S:**  
1 Enclosures, 9U  
82 NL-SAS, 2 SSD

## ESS Hybrid Models: GH12, GH14 and GH24

- Provide combination of flash and HDD storage tiers in one ESS building block
- Combine all-flash performance with GLxS capacity
- Use cases:
  - Single system for both metadata services and high density
  - Use the Flash for Burst Buffer storage and have Spectrum Scale migrate data to high density spinning disk
  - Use Scale to automatically manage data location between flash and disk based upon heat maps and policy
  - Handle multiple kinds of workloads such as video and analytics in the same environment

**Model GH12:**  
1 2U24 Enclosure SSD  
2 5U84 Enclosure HDD  
166 NL-SAS, 24 SSD



**18 GB/s\***

**Model GH14:**  
1 2U24 Enclosure SSD  
4 5U84 Enclosure HDD  
334 NL-SAS, 24 SSD



**38 GB/s\***

**Model GH24:**  
2 2U24 Enclosure SSD  
4 5U84 Enclosure HDD  
334 NL-SAS, 48 SSD



**40 GB/s\***

\* Estimate of performance aggregated across SSD and HDD. Assumes EDR Infiniband connections

# IBM Elastic Storage Server: Ultra-Dense Storage Models

- ESS GLxC models
  - The latest Seagate and helium drives
  - The highest storage density
    - Approximately 24% per ESS
  - Faster drive access speeds\*
  - More drive enclosure connections
- Leveraging the CORAL supercomputer modular design



Model	4U106 drawers	Drives	Raw capacity	Software licenses
GL1C	1	104	1.04 PB	104
GL2C	2	210	2.1 PB	210
GL4C	4	432	4.22 PB	432
GL6C	6	634	6.34 PB	634

\* Final Benchmarks to be published

# IBM Elastic Storage Server 5.3.2

## Improved storage capacity and economy

- New ESS GLxC models use denser enclosures to deliver up to **26% more capacity in 17% less rack space\***

## Upgrade without disrupting operations

- Capacity upgrades can be performed without application disruption
- Software automatically rebalances data across all drives


## High performance connectivity

- Optional integrated 100Gb Mellanox Ethernet switch provides a lower cost high performance interconnect alternative to Infiniband

## Leverage the latest IBM Spectrum Scale releases

- Data Management Edition and Data Access Edition

In a single 42U Rack



>70GB/sec  
>8PB of storage  
>789TB per Rack Unit  
*Ultra-dense storage*

\* Model GL6C compared to GL6S; varies by model

# Non-disruptive upgrades


## Simple expansion of storage capacity

- Spectrum Scale will automatically rebalance data in the background
- System automatically puts the new capacity to use
- No need to Archive & Restore data
- No System disruption\*


Non-Disruptive Upgrades	
From	To
GS1S	GS2S
GS2S	GS4S
GL1S	GL2S
GL2S	GL4S
GL4S	GL6S

Example


**Model GS1S**  
With 24 SSDs



Install additional  
drawer with 24 SSDs



**Model GS2S**  
With 48 SSDs



\*Requires space available in the rack

# IBM Power Accelerated Computing Platform

Accelerated compute configurations for High Performance Computing (HPC) and Artificial Intelligence / Deep Learning (AI/DL)

- Provides ability to create your own installation based on the IBM CORAL installation – the world's most powerful and smartest supercomputer



# CORAL-based Configuration

- Simplified configuration of scale-out infrastructure
- Support for HPC and AI workloads
- Software preloaded and fully rack-integrated in IBM manufacturing

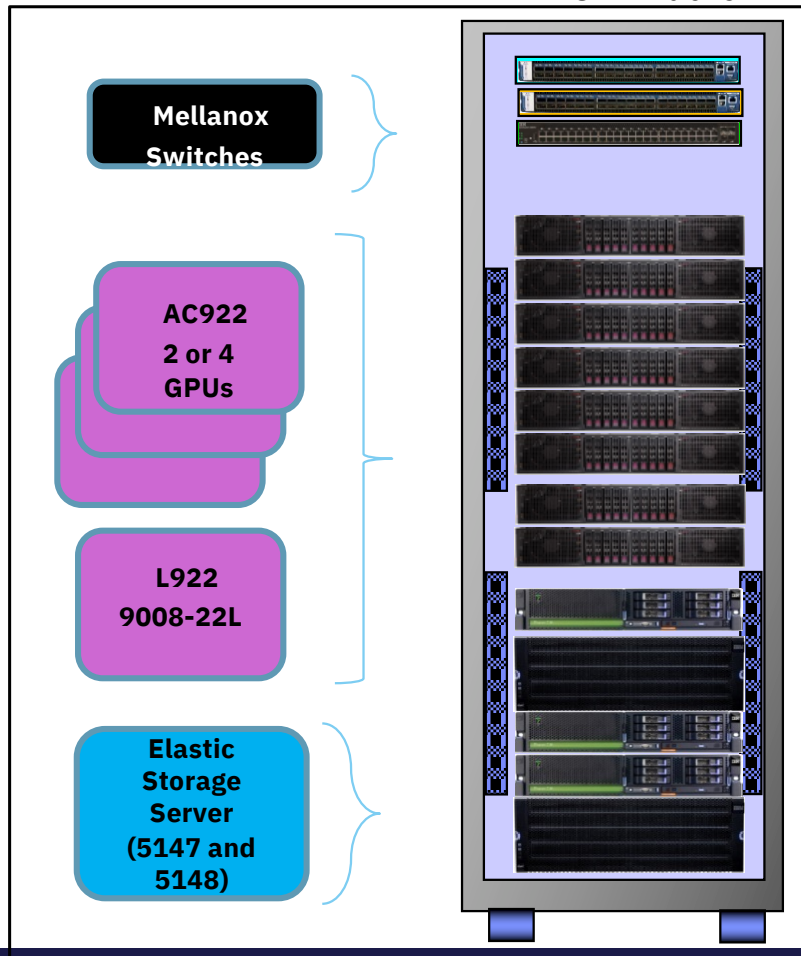
<b>Storage</b>	<b>Compute</b>	<b>Management</b>	<b>Switches</b>	<b>Rack</b>
Elastic Storage Server	AC922 (2 or 4 GPU) 8335-GTG	L922+ Management Server	Mellanox	One to four 42U Racks (S42)
Optional	Air Cooled Only Same Processors as in CORAL Servers		100Gb InfiniBand 40Gb Ethernet 10Gb Ethernet 1Gb Ethernet	If you really need more, let us know!

# Power Accelerated Computing Platform



## Sample Building Block View

1-4 S42 Racks



Configurable with Elastic Storage Server and Network Switches

Includes Login and Management Servers

Assembled in IBM factory for each client

Delivered ready to “plug in” to power