Washington Systems Center - Storage

Accelerate with IBM Storage: AI Scalability for a Decade

Mastering Massive Hybrid-Multi-Cloud Data Management

R. Kent Koeninger | IBM Washington System Center | Big Data and AI

@rkkoenin rkkoenin@us.ibm.com +1 603 589 0846

Dave McDonnell | Global Sales Exec for AI, BD and HPC, IBM Storage

@daveymcd davemcd@us.ibm.com +1 281 900 8176

Accelerate with IBM Storage Webinars

The Free IBM Storage Technical Webinar Series Continues in 2020...

Washington Systems Center – Storage 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 IBM Storage webinars, Clients, Business Partners and IBMers are welcome to send an email request to accelerate-join@hursley.ibm.com.

Located on the Accelerate with IBM Storage Site: https://www.ibm.com/support/pages/node/1125513

Also, check out the WSC YouTube Channel here:
https://www.youtube.com/channel/UCNuks0go01_ZrVVF1jgOD6Q

2020 Upcoming Webinars:

May 14 - Building a Data Protection Solution for Cyber Resiliency
Register Here: https://ibm.webex.com/ibm/onstage/g.php?MTID=e59cfcff05215c95622cb022b00bd00c7

May 19 – Spectrum Scale – Stretched Cluster Design
Register Here: https://ibm.webex.com/ibm/onstage/g.php?MTID=e21d6394be1f4e99e2f28866e7d2b5c88

May 21 - Storage Insights, Storage Insights Pro or Spectrum Control, which one is right for me?
Register Here: https://ibm.webex.com/ibm/onstage/g.php?MTID=eeb831096e67ff598348a5bb301d3038d

June 2 – Spectrum Scale ESS 3000
Register Here: https://ibm.webex.com/ibm/onstage/g.php?MTID=e6920e411fed595003800af92ecffe68e

June 4 - TS7700 Systems and zOS - Two Partners Better Together!
Register Here: https://ibm.webex.com/ibm/onstage/g.php?MTID=efdf15a2fcf8a4582d87a6e73d3ac9544

WSC Accelerate Survey

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

You can access it via

Menti.com 22 37 47

Big Data for AI Business Value

Scalable Machine Learning and Deep Learning

Improve insights

Increase sales
Enter new markets

Improve time to inference
with higher fidelity
and higher throughput
ML-DL models designed per month

Improve
product quality,
features, and innovation

Expand Rapidly

Grow as fast as needed
To as large as needed
TB → PB → EB
GB/s → TB/s

Reduce Costs

Cloud Storage
+
Online Storage
+
Deep Archives
=
30x to 100x
cost reduction
over the next decade

Convergences of AI, Analytics and HPC Scalability for a Decade

Rapidly Expanding AI, Machine Learning, and Deep Learning

Scalable Infrastructures for Explosive Big Data Growth

GPU Computing

Cognitive Applications require a modern data infrastructure

Attributes:

Performance
Scalability
Simple Management
Ease of Growth
Low Cost
Hybrid Multi-Cloud

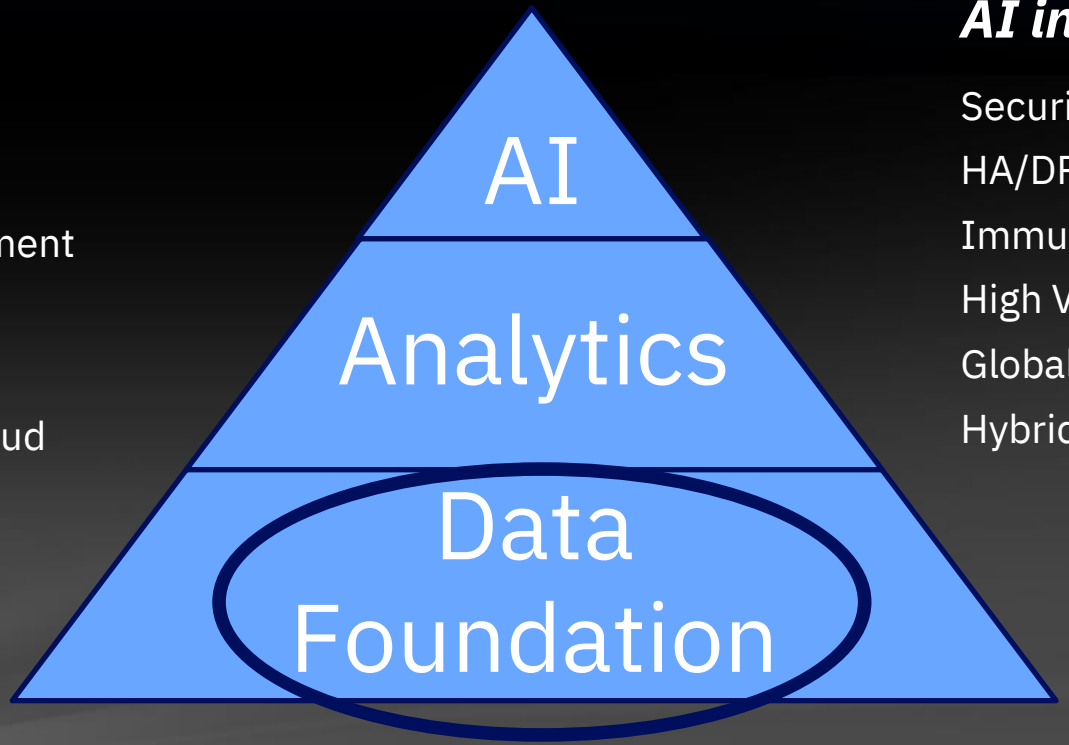

AI in Production:

Security
HA/DR
Immutability (GDPR+)
High Volume Global Ingest
Global Sharing of Data
Hybrid Multi-Cloud

Hybrid and multicloud are the new normal

94%

Share of enterprises using a mix of cloud models

67%

Share of enterprises using more than one public cloud provider

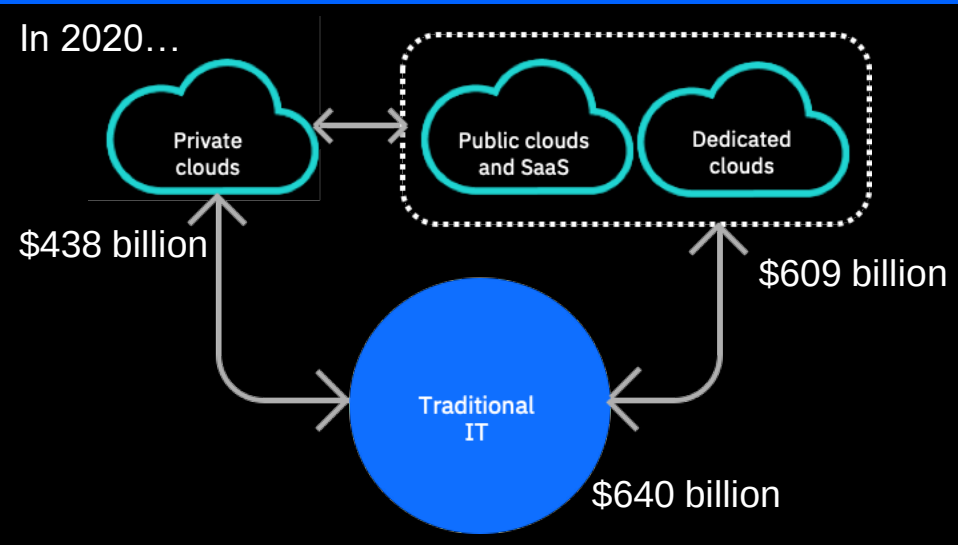

Learn more about the implications for storage

Read the blog series

ibm.biz/storage**hybridmulticloud**

How Does Hybrid-Cloud Apply to BD-AI?

BD-AI can place high demands on storage infrastructures

Varies by use cases
Varies by BD-AI throughput
Varies by purpose-build efficiency

Speed matters.
– Storage infrastructure can impede or improve time to inference by 10x to 100x.
– Slow storage infrastructure can impede or improve fidelity.
– *Analysts personnel can dominate costs*

Public-cloud bandwidth costs
– Sometimes limit BD-AI throughput

Spectrum Scale on bare-metal clouds and purpose-built-managed clouds
– Lower-cost bandwidth = higher throughput
– Faster time to inference with higher fidelity
– More frequent and higher-quality models per analyst

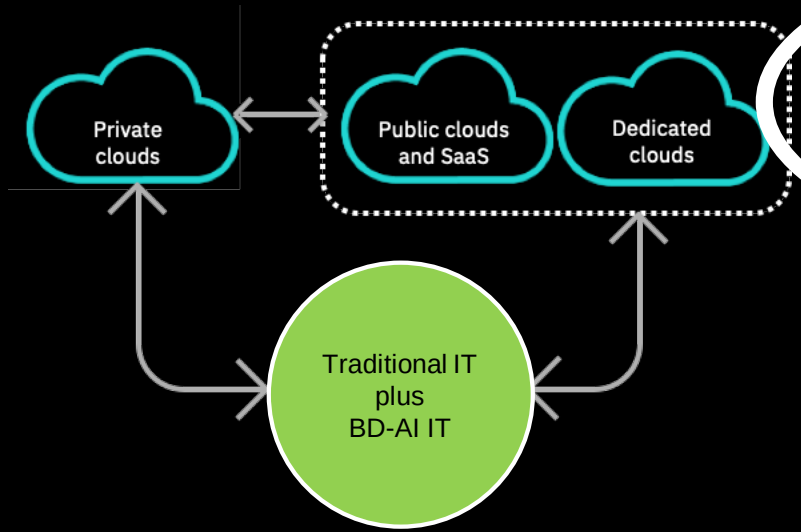

Hybrid-Multi-Cloud BD-AI

Balancing high throughput and low cost

Public Cloud Economics

Spectrum Scale Speed in the Cloud

Bare-Metal BD-AI Clouds

ORACLE

IBM Spectrum Scale in the Cloud

Re-Store | A HPC autonomous

Spectrum Scale Speed w/o Datacenters

BD-AI SSaaS

Spectrum Scale as a Service

IBM Spectrum Scale off prem

IBM Spectrum Scale

IBM Spectrum Archive

Speed

Datacenter BD-AI

IBM Spectrum Scale Elastic Storage Server

Storage: PB → EB in This Decade

Datacenter BD-AI Archives

IBM Spectrum Scale Online Archive SSOA

High-performance HDD archive | At the price of private clouds | Multi-prem HA and DR

Archives Files and Objects (S3)

IBM Spectrum Archive Simple Scalable Tape IBM Enterprise Tape

30x to 100x less expensive than HDD over the next decade | **Scalable** capacity, throughput, and access | ***Simpler*** than private-cloud archives

Datacenter BD-AI

Predominance of Scalable BD-AI

Fast
Customizable, Scalable, Efficient

Datacenter BD-AI

IBM Spectrum Scale
Elastic Storage Server

Purpose-built for speed

Optimized BD-AI pipelines
Customized for your particular BD-AI needs

E.g., Varying designs for Automated driving (ADAS),
sentiment analysis, and financial information systems (FIS)

Choices matter:

Choice of scalable compute servers:
– Processors, memory, # of GPUs…
Choices among scalability, performance, and cost optimized
– Flash ←→ Disk ←→ Cloud ←→ Tape
Choices among high-speed networks:
– 1 GB/s → 40 GB/s per server to feed the beastly GPUs

E.g.,
– BD-AI datacenter rapid ingest:
– *10s to 100s of times faster than in public clouds*

Spectrum Scale Online Archive (SSOA)

*Cost optimized, high-performance, d*atacenter BD-AI Archives

IBM Spectrum Scale Online Archive SSOA | High-performance HDD archive | At the price of private clouds

Online archives

At the cost of private clouds

High-Performance, Low-Cost Archive

Spectrum Scale Online Archive
SSOA

IBM Spectrum Scale
Elastic Storage Server

All-in-one Archive

– Configure capacity that would have been used in Private Cloud

– Archive data in IBM Spectrum Scale itself

Instant Access

– Online, compressed disk

BD-AI performance

– Simultaneous throughput and archive

Automated and transparent tiering

– Flash ←→ Disk ←→ Cloud ←→ Tape

Exascale, Automated BD-AI Archive

Performance and cost optimized

BD-AI for decades
30x to 100x less expensive than disk
Simple

IBM **Spectrum Archive**
Simple Scalable Tape
IBM Enterprise Tape

30x less expensive than HDD over the next decade

Scalable capacity, throughput, and access

Simpler than private-cloud archives

BP → EB
30x Less Expensive

Simple Scalable Tape

IBM Spectrum Archive

IBM Enterprise Tape

IBM Simple Scalable Tape (SST)
– IBM Spectrum Archive
– IBM Enterprise Tape

Simpler than private-cloud archives
– Simpler installation, operation, and expansion

IBM Spectrum Archive
– Automated operations
– Transparent access

Low Cost
– CAPEX, OPEX, PSC, labor…

Automated and transparent tiering
– Flash ←→ Disk ←→ Cloud ←→ Tape

21st-Century **Simple Scalable Tape**

Simple

SST

Massive Rows of Tapes
Always Stored
In Automated Libraries

SST

Labor Intensive
Quickly Insert Cartridges
Crazy Glue the Door Shut

SST

"Iron Mountain" Offsite Storage
Online

IBM Simple Scalable Tape

Set it and forget it

Scalable capacity, throughput, and access
Full automated
Less effort than with Private Clouds

Easy PB → EB

Scalable

Scalable capacity
– Add tape cartridges and frames, as needed
Scalable throughput
– Add tape drives, as needed
Scalable access (cartridge mounts per minute)
– Add libraries as needed (additional robotics)

Easy installation and expansion

Install 10s to 100s of PB in a day or three
– Install the library.
– Install the drives.
– Load the cartridges.

Install cartridges to the required capacity
– Crazy Glue the door shut

Easy to increase capacity
– Do nothing the first year: Automatic 50% boost
– (Open the door with acetate) and insert new cartridges

IBM Simple Scalable Tape

Automated Product Lifecycle Management

Simple installations
Simple expansions
Automated migrations

Automatically Migrate Data Before Obsolete

$12,500/PB decreases to $785/PB tape media over a decade

Scalable Capacity
– Add tape cartridges and frames, as needed
Scalable throughput
– Add tape drives, as needed
Scalable access (cartridge mounts per minute)
– Add libraries as needed (additional robotics)

Hands-free migrations
– Tape drives for production
+ tape drives for migration
– Spectrum Archive automatically migrates data among tape pools before End of Service

Quick and Easy Upgrades
A few hours to upgrade tape capacity
– Swap tape cartridges
Vs. days to upgrade disk capacity
– Remove, replace, and add disk drives

Price and Density Improvements Over the Next Decade

From 100 Petabytes (PBs) to 1 Exabyte (EB)

Rapid Density Growth
20 TB to 330 TB per cartridge
by 2030

Rapid Price Reductions
$12,500/PB to $785/PB
by 2030

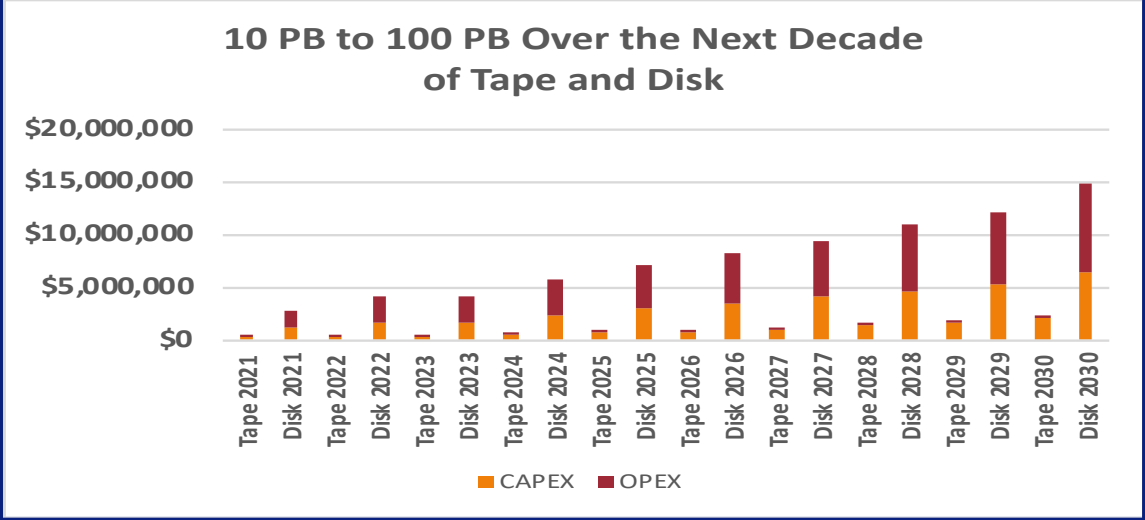

10x cost savings

$13,000,000 saved over this decade

$1.5 M for tape storage
vs.
$15 M for disk storage

Lower CAPEX and OPEX
Lower cost
Lower PSC
Less Labor

Price and Density Improvements Over the Next Decade

From 100 Petabytes (PBs) to 1 Exabyte (EB)

Rapid Density Growth
20 TB to 330 TB per cartridge
by 2030

Rapid Price Reductions
$12,500/PB to $785/PB
by 2030

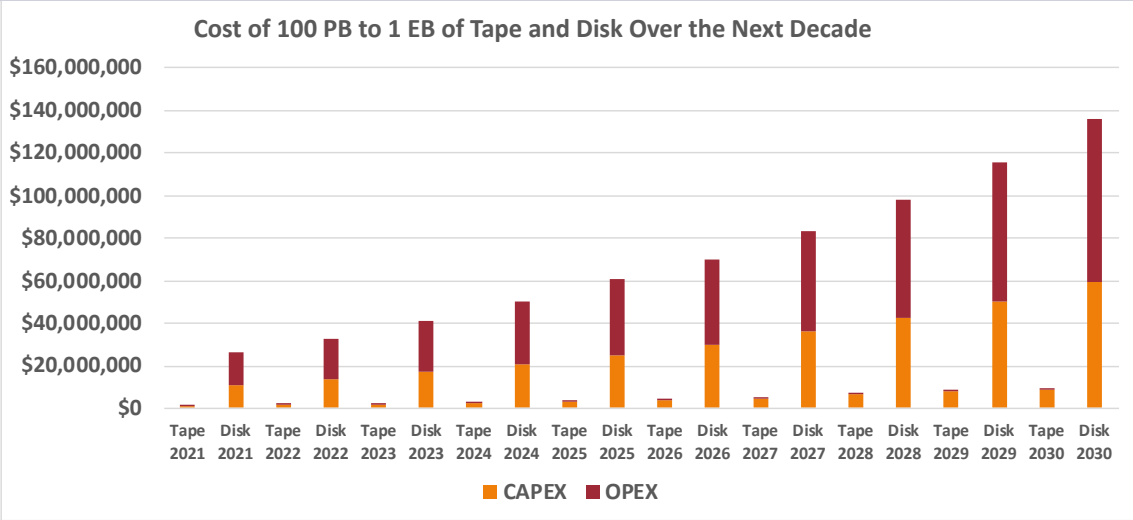

10x cost savings

$135,000,000 saved over this decade

$15 M for tape storage
vs.
$150 M for disk storage

Lower CAPEX and OPEX
Lower cost
Lower PSC
Less Labor

Simple Scalable Tape OPEX

Power, space, cooling, and support low percentage of CAPEX

Wondering about implications beyond capital cost?

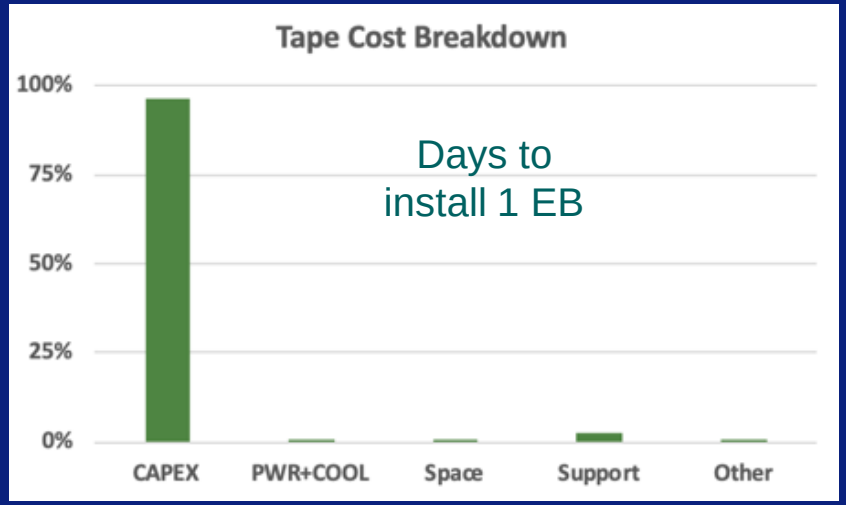

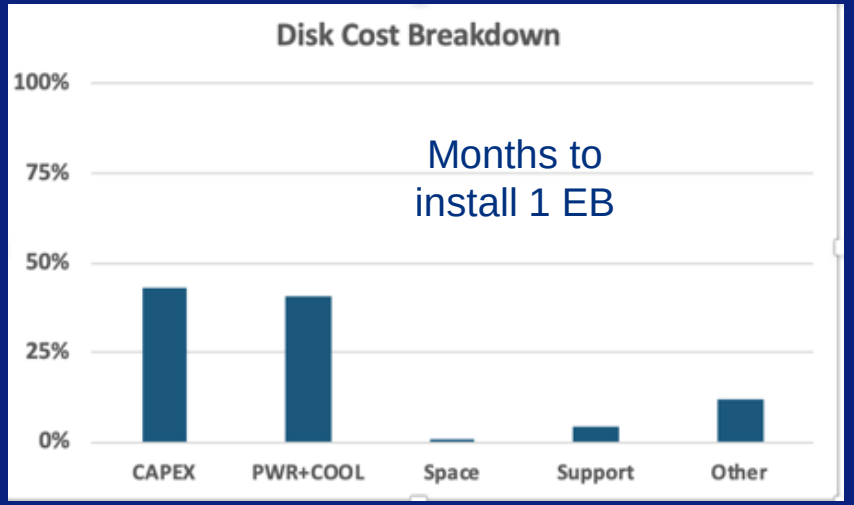

BD-AI in Virtualized Public Clouds

IBM Spectrum Scale in the Cloud

Public Cloud Economics

Easy to Install
Scalable capacity and performance
Standard Licenses

Existing BD-AI in the Cloud

S3 objects or files
Scalable GPUs
Scalable storage capacity
Scalable storage bandwidth?

BD-AI in Virtualized Public Clouds

Expensive bandwidth with or without Spectrum Scale

Public Cloud Economics

Attractive Cloud Capacity Prices
Unattractive Cloud Bandwidth Prices
Low performance density
– GB/s per TB
– High Cost per GB/s

– BD-AI bandwidth can matter.
– 100x cost of bandwidth can matter.
– Can impeded analytics and AI throughput

5-Year Comparison of AWS Elastic Block Storage (EBS) and IBM ESS GL6S Bandwidth Cost

IBM ESS GL6S	
GB/s	30
AWS	$ BW
MB/s/TB burst	10
$/GB/mo	$0.05
$/TB/5-years	$2,700
1 TB = 1 GB/s	
MB/s/5-years	$2,700
GB/s/5-years	$2,700,000
30 GB/s/5-years	$81,000,000

AWS bandwidth 100 to 200 times more expensive than on IBM ESS GL6Ses

Scalability Matters: Algo Stock Trading
Winning More Trades by Increasing the Degrees of Confidence

- **Slow Scale-Out NAS vs. Fast Big Data Pipeline**
- I/O task: Load 6 TB of data across 25 servers
- BD-AI pipeline: 32x faster loads
- BD-AI pipeline: 38 times more analyses per month on high-performance Big Data / AI pipeline

- **Millions of dollars increased profits**

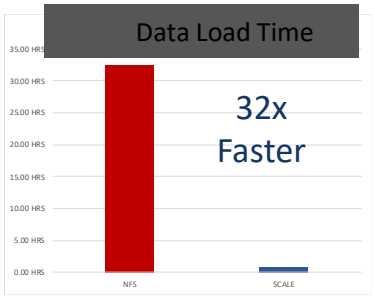

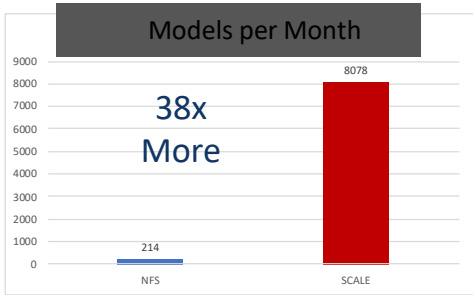

Faster data loads with Scalable BD-AI
= Larger sample sizes + more runs
= Higher degrees of confidence
= More trade wins

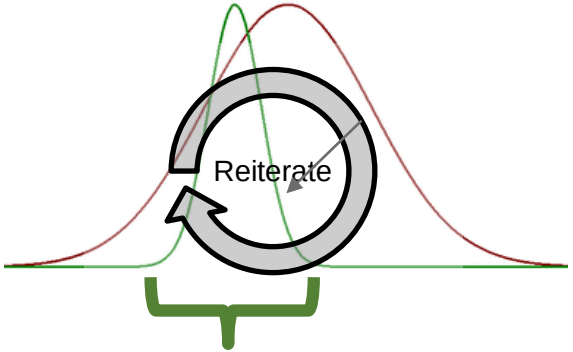

Narrower Range
= Higher degrees of confidence
Winning trades improved 7%
(Lucrative!)

BD-AI as a Service

IBM Spectrum Scale as a Service (SSaaS)

Datacenter BD-AI as a Managed Service

No datacenter needed

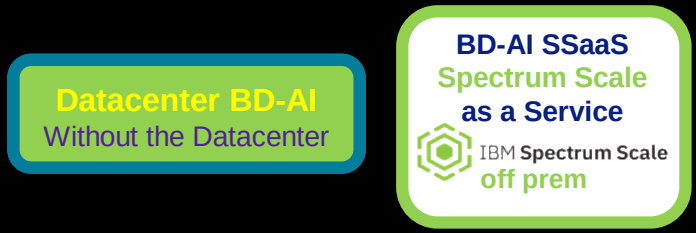

Scalability and efficiency of datacenter BD-AI
Without managing a datacenter
Without added personnel
Without the need for BD-AI storage expertise
Offered by experts on IBM Spectrum Scale

Hybrid-multicloud
Augment Cloud economics
with BD-AI datacenter economics

Substantial savings for archived data
compared to AWS EBS, AWS Glacier,
Google Nearline, and Google Coldline

E.g.,
A large tier-2 cloud-based services company
operating without a datacenter
plans to archive their data
using IBM Simple Scalable Storage
with IBM Spectrum Archive and IBM Enterprise Tape

Bare-Metal IBM Spectrum Scale in the Cloud

Avoids complexities and inefficiencies of Cloud virtualization

No storage virtualization
High-bandwidth network
As if in a datacenter

High-performance cloud storage
Good storage-density performance
First high-performance

Substantial savings compute and storage
Compared to AWS EBS and AWS Glacier

E.g.,
An oil company is using Spectrum Scale
on bare-metal on Oracle Cloud Infrastructure
(OCI) for seismic simulations
(to identify potential oil repositories).
Re-Store designed, implemented,
and Is managing this service on OCI.

Top 20
Supercomputing
IO500

Spectrum Scale
Speed
in the Cloud

This is not an endorsement, just an example.

Improve Time to Inference and Fidelity w/o Building a Datacenter

Start in Public Cloud → Expand for faster time to inference

Public Cloud → S3 Among All Options
Plus Spectrum Scale Transparent Data Sharing
Add Options and Grow as Needed

Intuit Status Quo
Public Cloud Economics

Shared S3 and Files

Small Groups of Analysts

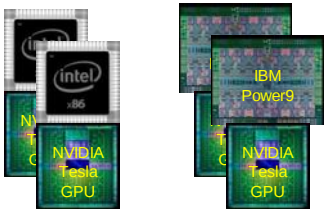

Spectrum Scale Starter Solution

BD-AI Quick Start
x86 + Power9
+ GPUs Computing
Spectrum Scale in a Box
w/o a datacenter

Shared

Shared

Shared

Spectrum Scale Performance in the Cloud

Bare-Metal BD-AI Clouds

ORACLE

IBM Spectrum Scale in the Cloud

Re-Store | AHPC

Spectrum Scale Performance w/o Datacenter

BD-AI SSaaS
Spectrum Scale as a Service
IBM Spectrum Scale off prem

Scale as needed to High-Performance analytics and AI

Summary: AI Scalability for a Decade

Hybrid-Multicloud Big Data - AI

IBM ESS
+ IBM Spectrum Archive
+ IBM Enterprise Tape

100 PB → 1 EB
30x less expensive than disk over the next decade

High throughput to support BD-AI GPU computing
Low cost performant and archive storage

Speed

Datacenter BD-AI

IBM Spectrum Scale
Elastic Storage Server

IBM Spectrum Scale
Elastic Storage Server

IBM Spectrum Archive
Simple Scalable Tape
IBM Enterprise Tape

Storage: PB → EB in This Decade

Datacenter BD-AI Archives

IBM Spectrum Scale
Online Archive
SSOA

High-performance HDD archive | At the price of private clouds | Multi-prem HA and DR

Archives Files and Objects (S3)

IBM Spectrum Archive
Simple Scalable Tape
SST

30x to 100x less expensive than HDD over the next decade | Scalable capacity, throughput, and access | *Simpler* than private-cloud archives

Start Small and Grow as Needed

Spectrum Scale Starter Solution

Spectrum Scale Speed in the Cloud

Bare-Metal BD-AI Clouds

ORACLE

IBM Spectrum Scale in the Cloud

Re-Store | AHPC

Spectrum Scale Speed w/o Datacenters

BD-AI SSaaS
Spectrum Scale as a Service

IBM Spectrum Scale off prem

IBM Spectrum Scale in the Cloud

Low-cost and performant cloud bare-metal BD-AI storage
Low-cost and performant Storage as a Service

Accelerate with IBM Storage Survey

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

You can access this 5 question survey via Menti.com with code 22 37 47 or

Direct link https://www.menti.com/mkg7a2x6q8

Or

QR Code


IBM