March 4, 2021. UK-Durham Univ.

Terry Leatherland, Cognitive Systems Sales Leader
US Public Market. (Education, Life Sciences, Healthcare, Govt.)
Dr. Ruzhu Chen, IBM Customer Center

Cognitive Solutions Update

DNA-Genomics

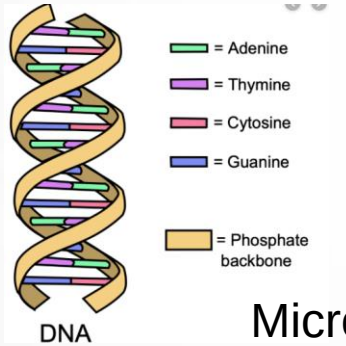

Microscopy

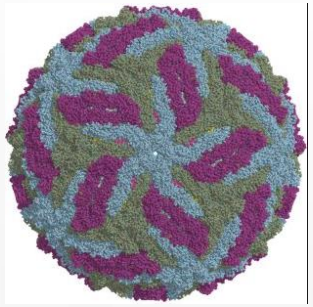

Agenda Topics

End to end solution design/portfolio for HPC/AI for Life Sciences from IBM Systems.

- Using Tuned and Optimized Server/Storage/Network/Software pipeline to achieve 2,5,10x faster “Time to Answer”.
- Genomics Solutions- 10x performance over standard Intel operation.
- High Resolution Microscopy – 2X faster Time to Answer
- Structural Biology Simulation- 5,10,20x faster simulation
- NA Supercomputer Systems in Academia and Government
- Server Solutions for Computational Science
- Storage solutions for High density or High-speed low density
- Parallel File System optimized for HPC
- Bayesian Optimization for HPC
- Hybrid Cloud- Orchestration
- Power10 Memory Clustering

We build fast computational “cars”. End-to-End

“Time to Answer”
Is a competitive and scientific advantage.

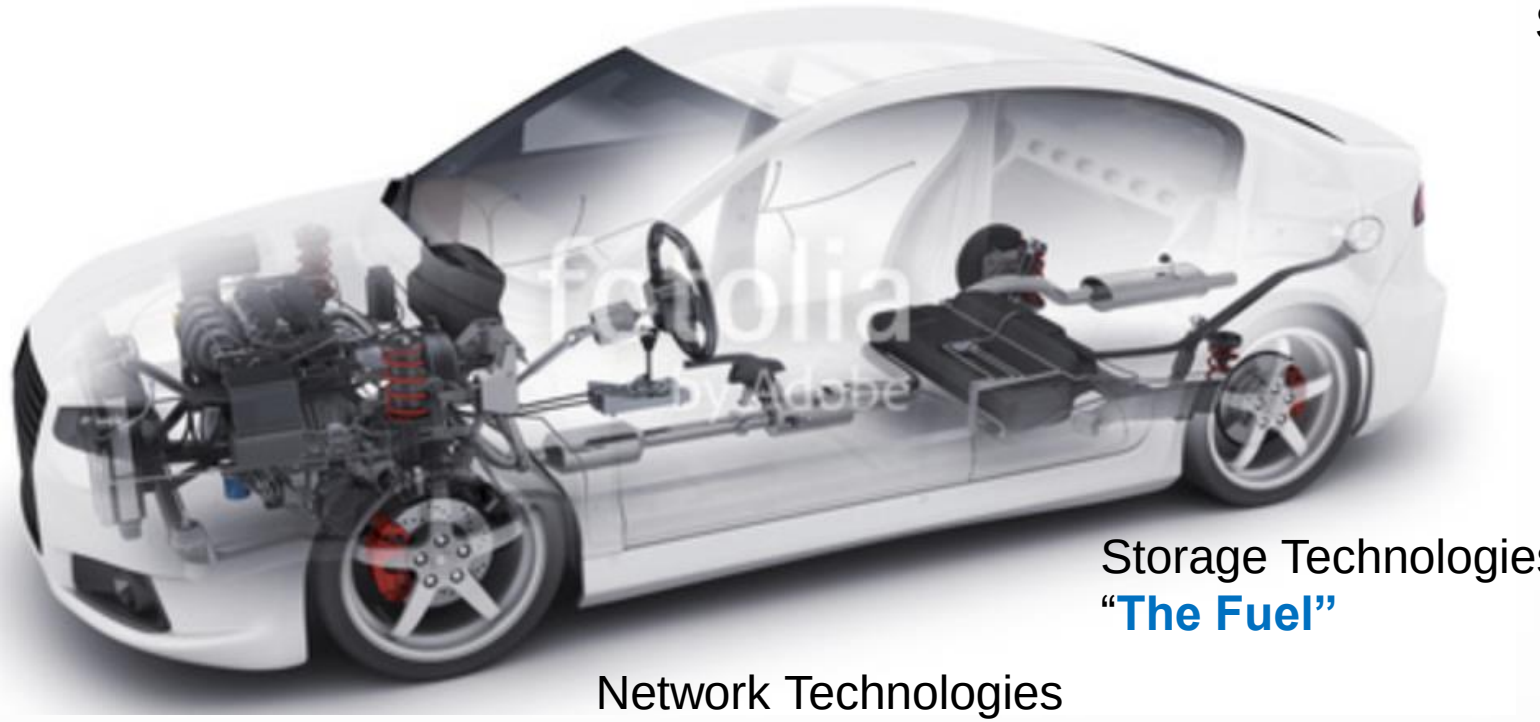

Single Vendor Support

Storage Technologies
“The Fuel”

Network Technologies
“The Fuel Lines”

Server Technologies
“The Engine”

Software Defined Infrastructure/Scheduler/Orchestrator
“The Drive Train”

Systems tuned front to back to produce results instead of economical parts assembly.

Improve time to results. Impact more lives.

We are living through an unprecedented moment for humanity that touches every aspect of business, technology and culture. The global pandemic caused by COVID-19 demands answers now.

IBM delivers faster insights with greater efficiency to impact more lives.

High-performance Data & AI deploys against this problem at massive scale and reduces time spent delivering life-saving insights through unique load balancing and model optimization technologies delivered by cutting-edge IBM lab research.

Smart loves problems, and there has never been a bigger problem facing our world.

Genomics

Molecular Simulation

Biomolecular Structure

Medical Diagnostics AI

Data Fusion and AI

Bio-Informatics

Artificial intelligence and high-performance computing have already begun to attack the virus, assisting in molecular drug discovery, genomics and medical image processing.

Five key challenges to progress remain despite advances

- **Data Overload**
- Oceans of data arise from rapid digitization and instrumentation of healthcare.

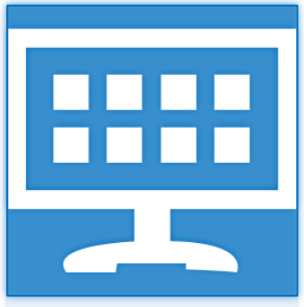

- **App Chaos**
- Thousands of applications, workflows and models are not all following the same rules.

- **Adoption**
- Vertically integrated toolsets with heavy customization and vendor lock-in create work silos.

- **Performance**
- When scaling up or out, most institutions cannot diagnose or analyze the performance problems they face.

Cost

Demanding workloads require well-orchestrated infrastructure to manage, monitor and control costs.

Power Systems

Optimizing Medical Imaging

Enhance image identification with deep learning to assist physicians and benefit patients

20x faster

Washington University in St. Louis School of Medicine

VANDERBILT SCHOOL OF MEDICINE

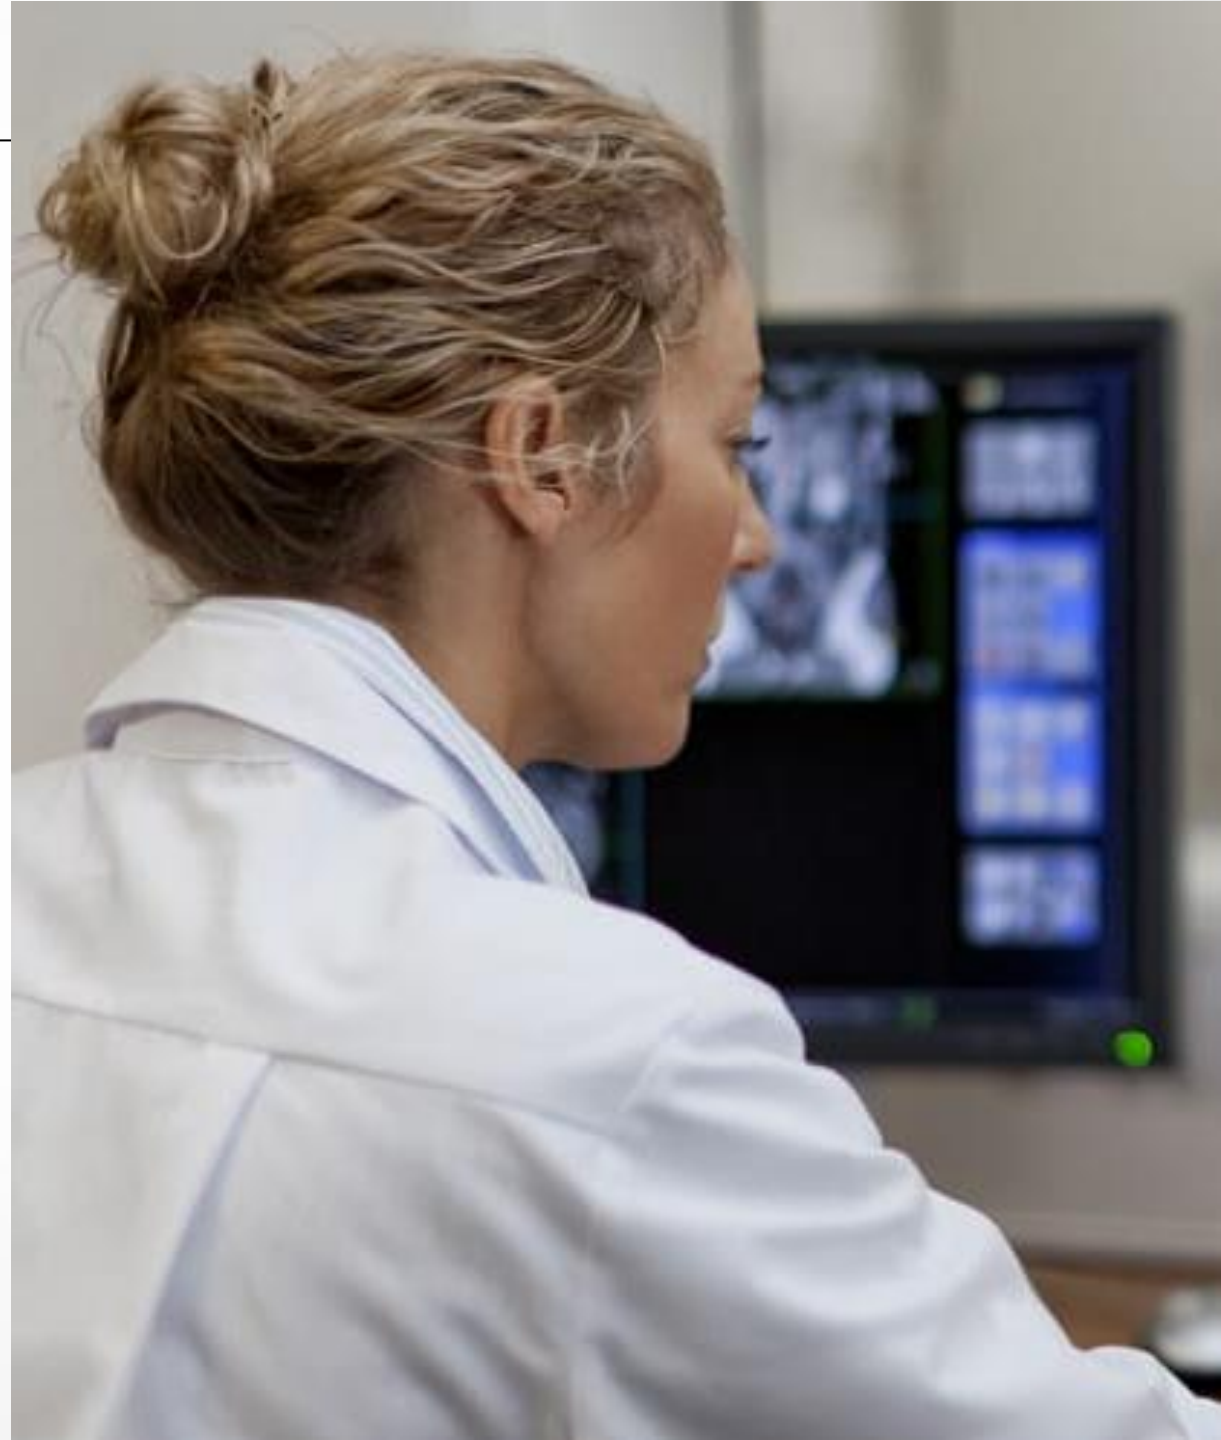

Power Systems

Optimizing Precision Genomics

Reduced time-to-completion for long-running jobs while increasing resource

100,000+

سدرة للطب
Sidra Medicine

IBM Accelerates Medical Research Tasks

Genomics

Biomarkers detection, biodata modeling and statistics data visualization

Diagnostics

Image classification using AI with flexible, targeted models in open frameworks

Molecular Simulation

Drug discovery via modeling of macromolecule receptors and small-molecule ligands

Data Fusion

Synthesize and model diverse data using data fusion, natural language processing, and machine learning

Biomolecular Structure

Cryo-EM image restoration and refinement analysis for drug design and discovery

BioInformatics

Cross Discipline Informatics Relatioanships AI

From Data to Insight with IBM HPDA Reference Architecture

A framework for designing, deploying, growing and optimizing infrastructure for HPC, AI and Cloud, created in collaboration with world's leading healthcare and life sciences institutions, and using Red Hat OpenShift, IBM Power Systems, IBM Storage and open API endpoints.

Medical Tasks

Genomics, molecular simulation, structural analysis, diagnostics, data fusion, manufacturing quality inspection.

Apps & Models

Large-scale and high-throughput workloads such as HPC, AI and Cloud computing

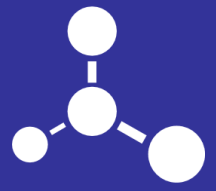

Orchestration

High Performance Computing & AI Platform Capable of Orchestrating Thousands of Servers and GPUs

DataHub

High Performance Data Fabric & Catalog Capable of Handling Exabytes of Data and Trillions of Objects

The Data: Biological Data Analytics

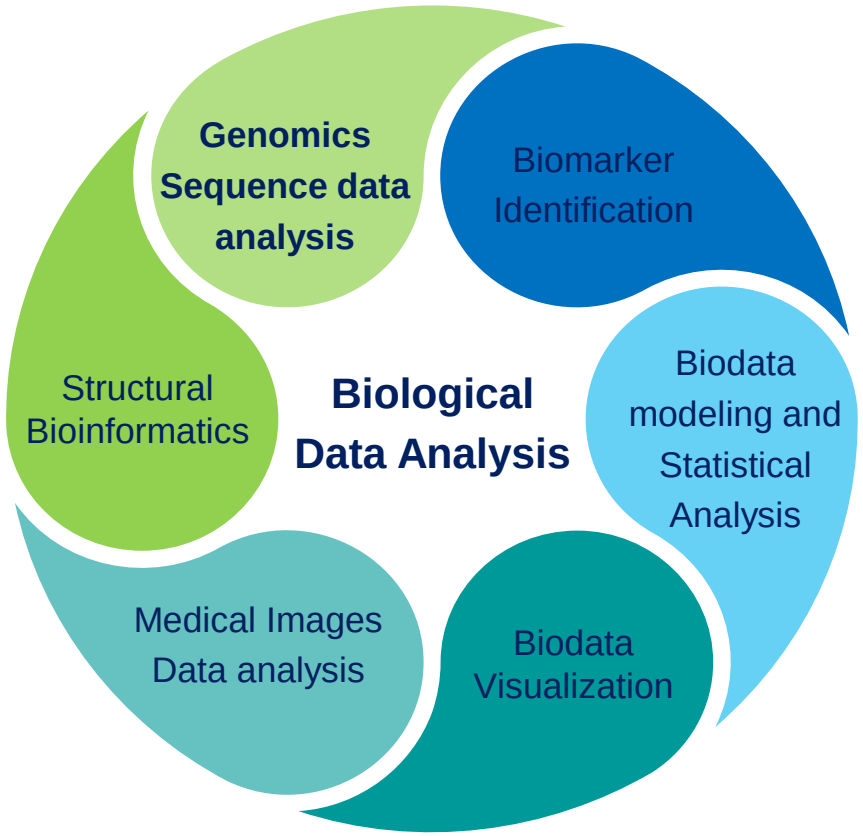

Biological Data Analytics
❑ Genomic Sequence Data: an explosive growth of biodata ❑ Sequence alignment ❑ Variant discovery and characterization ❑ Genomic profiling and pattern discovery
❑ Biomarker Identification: gene expression profile, RNA-seq, ChIP-seq, microarray identification and validation, etc.
❑ Structural Bioinformatics: identify and predict 3D biomolecule structures, such Cryo-EM data refinement, molecular dynamic simulation, NMR, x-Ray crystallographic data, etc.
❑ Biodata Modeling & Statistical Analysis: biological pathways analysis, Gene, clinical data cohorts study, data extraction, etc.
❑ Medical Image Processing: image segmentation, registration, statistic modeling.
❑ Biodata Visualization: 3D molecule structures, genomics sequences visualization, etc.

The Challenges: Analyzing Explosive biological data

Data Explosion 	Data Storage 	Computation 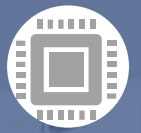	Solutions 
Large volume and variety of data around genomic sequences, gene expression, images, structural biomolecules, clinical and healthcare information, personized medicine data	High performance and high throughput storage hierarchy required for data loading, extraction and computation. Tertiary storage required for archive and store. Storage tools for data indexing, discovery and governance.	High performance and efficiency of software tools and applications for genomic variants and biomarkers analysis, drug discovery, medical image processing and molecule structure modeling, data visualization.	High throughput and optimized workload pipelines to accelerate biodata analysis with highly optimal and parallel I/O, memory, CPU and GPU computations.

What are the components of the IBM Solution

- Same components that are used at the Oak Ridge and Lawrence Livermore U.S. National Laboratories.
- IBM Power9 Servers w GPUs
- IBM Elastic Storage Systems
- Spectrum Scale Parallel File System
- IBM Spectrum Discover MetaData Catalog
- Advanced AI Tools

IBM Systems @ Oak Ridge

OAK RIDGE National Laboratory

#1 TOP500 SUPERCOMPUTER SITES

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
- **Spectrum Scale V5**
- **77 Elastic Storage Servers**

2.5 TB/sec
Throughput to storage architecture

250 PB file system
HDD storage capacity

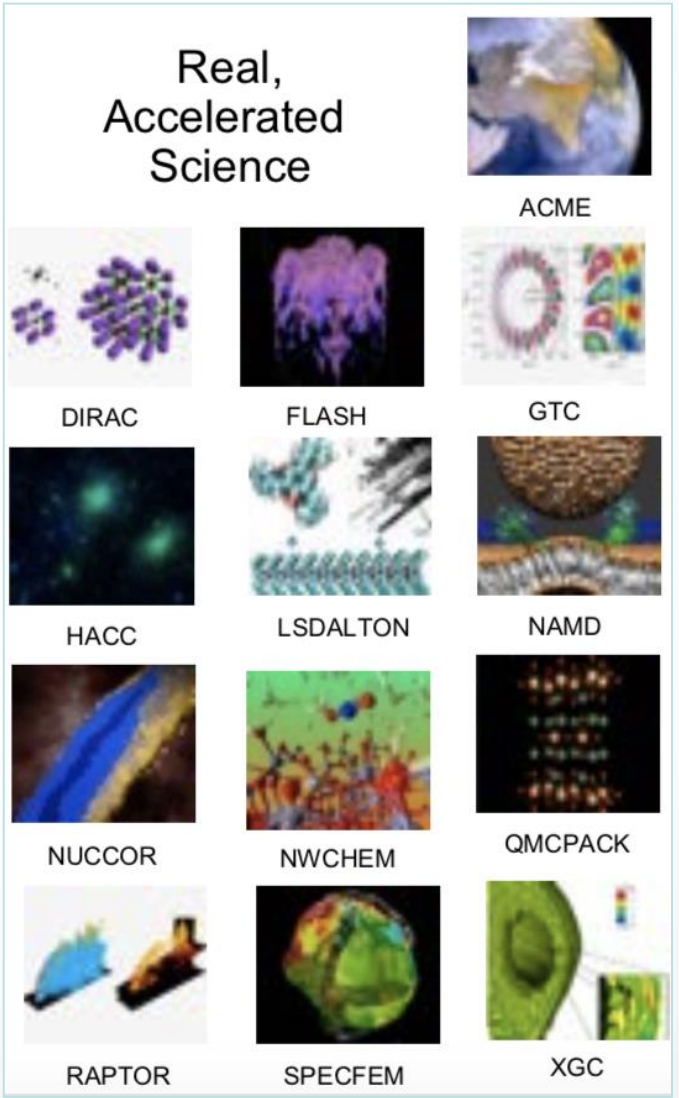

Challenges for the world's smartest supercomputer for open science

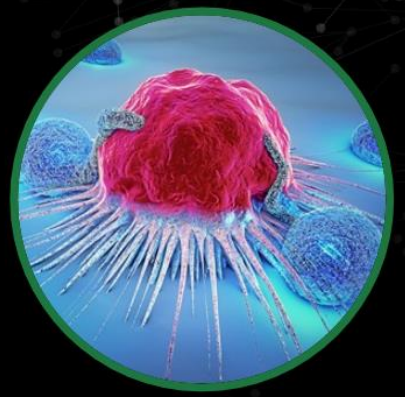

Combating Cancer

Through the development of scalable deep neural networks, scientists at the US Department of Energy and the National Cancer Institute are making strides in improving cancer diagnosis and treatment.

Predicting Fusion Energy

Predictive AI software is already helping scientists anticipate disruptions to the volatile plasmas inside experimental reactors. Summit's arrival allows researchers to take this work to the next level and further integrate AI with fusion technology.

Deciphering High-energy Physics Data

With AI supercomputing, physicists can lean on machines to identify important pieces of information—data that's too massive for any single human to handle and that could change our understanding of the universe.

Identifying Next-generation Materials

By training AI algorithms to predict material properties from experimental data, longstanding questions about material behavior at atomic scales could be answered for better batteries, more resilient building materials, and more efficient semiconductors.

IBM Power University HPC Adopters in the US.

Power Systems

Welcome to the waitless world

IBM

TACC

Massachusetts Institute of Technology

Lawrence Livermore National Laboratory

OpenPOWER™

OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY

MIAMI

ARC·TS ADVANCED RESEARCH COMPUTING TECHNOLOGY SERVICES UNIVERSITY OF MICHIGAN

IBM Power System AC922

Provides the data and compute-intensive infrastructure needed to deliver faster time to insights, creating an incredibly powerful training platform.

IBM Spectrum Scale

VIRGINIA TECH

Oregon State University

Princeton Research Computing

IBM Elastic Storage Server

Performance and growth for clusters, cognitive workloads and analytics

LSU CENTER FOR COMPUTATION & TECHNOLOGY

Interdisciplinary | Innovative | Inventive

NYU Langone Health

THE UNIVERSITY OF TENNESSEE KNOXVILLE

ILLINOIS

THE CENTER FOR GENOME ARCHITECTURE

Baylor College of Medicine

Rice University

Baylor College of Medicine

CWU

CENTRAL WASHINGTON UNIVERSITY

Rensselaer

why not change the world?®

AUBURN UNIVERSITY

THE UNIVERSITY OF TENNESSEE CHATTANOOGA

CLEMSON COMPUTING & INFORMATION TECHNOLOGY

** Hyperlinks to customer websites within the logos.

IBM

Cryo-EM Application Enablement on Power9

Current on Power9

RELION

CTF estimation (ctffind4, ctftilt)

CISTEM

Near term

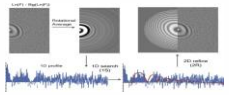

Gctf

crYOLO

MotionCor2

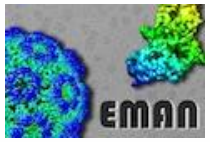

SPIDER

System for **P**rocessing **I**mage **D**ata from **E**lectron microscopy and **R**elated fields

Future

Frealign
Imagic
XMIPP
IMOD
ProTomo
ACE
findEM
DogPicker
TiltPicker
RMeasure
EM-BFACTOR
Chimera
CryoSparc. (** Not opensource)

The Solution: Accelerated Computing Platform (ACP)

Hardware Building Blocks

IB and Ethernet switches & UFM (Mellanox) shared

AC922
8335-GTG
2-4 GPUs

4-18 compute nodes
0-5 login nodes
Per rack

L922
9008-22L

1-5 Mgmt. nodes
Per rack

IBM Spectrum Storage

0-1 ESS per cluster
P8 any model
(optional)

IBM TOR Switch
Enet TOR Switch
Optional UFM Server

Compute Nodes

XCAT/Manager
Login Node

ESS Mgmt. Node

ESS* Storage

* Any p8 model

Solutions Stacks

- **High performance & scalability**
  - RHEL7.6, CUDA, ESSL, SMT,GPFS, NVMe …
  - Optimized workload pipelines, e.g., GATK4, RELION …
  - PowerAI, TensorFlow, Caffe …
- **Application interoperability & manageability**
  - Docker container, k8s, LSF ..
  - Anaconda, python, Jupyter …
  - MPI, XLC/C++, GCC, AT …
- **Data management & governance**
  - Spectrum Discover

ACP Solutions for Cryo-EM

Power Systems

Welcome to the waitless world

IBM

Small	Medium	Large
3-5 samples per day	10-15 samples per day	15+ samples per day
1 40-core IC922 node, GPUs + 512 GB Mem	2 40-core IC922 nodes, 512GB Mem + 4 GPUs each	4-8 40-core IC922 nodes, 512 GB Mem + 4 GPUs each
GL2 (GPFS) with IB adapter	GL4 (GPFS) with IB adapter	GL4 (GPFS) with IB adapter
One SSD (NVMe)	2 SSDs per node (NVMe)	2 SSDs per nodes (NVMe)
Ethernet	IB	IB
Optimized RELION software stack running on RHEL 7.6 with 4 SMTs per core.		

RELION Class3D Standard Benchmark Runtime

Dataset	Power8 4 GPUs	Power9 4 GPUs	Power9 16 GPUs
Pf. ribosome	71 min.	60 min.	44 min.

Perf

RELION Class3D

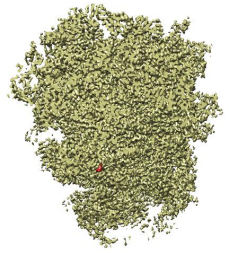

Steps	P8 + 4 P100	P9 + 4 V100	P9 + 16 V100
Estimated GPUs	35	28	9
Estimated FFTW	12	11	11
Others	24	21	24
Total Runtime (minutes)	71	60	44

* Running on GPFS with SSD as scratch directory

(https://www.emdataresource.org/EMD-2660)

RELION: cryo-EM Restructure Refinement

- RELION (for Regularized LIkelihood OptimizatioN)
  - Widely-used processing software package for macromolecular structure determination by single-particle analysis of cryo-microscopy (cryo-EM) images.
  - Use an empirical Bayesian approach to refinement of 3D reconstructions or 2D class averages from cryo-EM data
- RELION pipeline
  - Use pipelined approach for the entire single-particle workflow
  - GPU and CPU acceleration to reduce its computational costs
  - Its functionality was expanded with the incorporation of sub-tomogram averaging and helical reconstruction.
- RELION IBM Power Support
  - Latest version (3.1) is supported on Power9 with GPU + CPU
  - Optimized performance on CPU + GPU cluster environment
  - Use HPDA solution including IBM ESS, Power and IBM Spectrum LSF

IBM Spectrum Scale

POWER9

IBM CLOUD PRIVATE

ReliOn

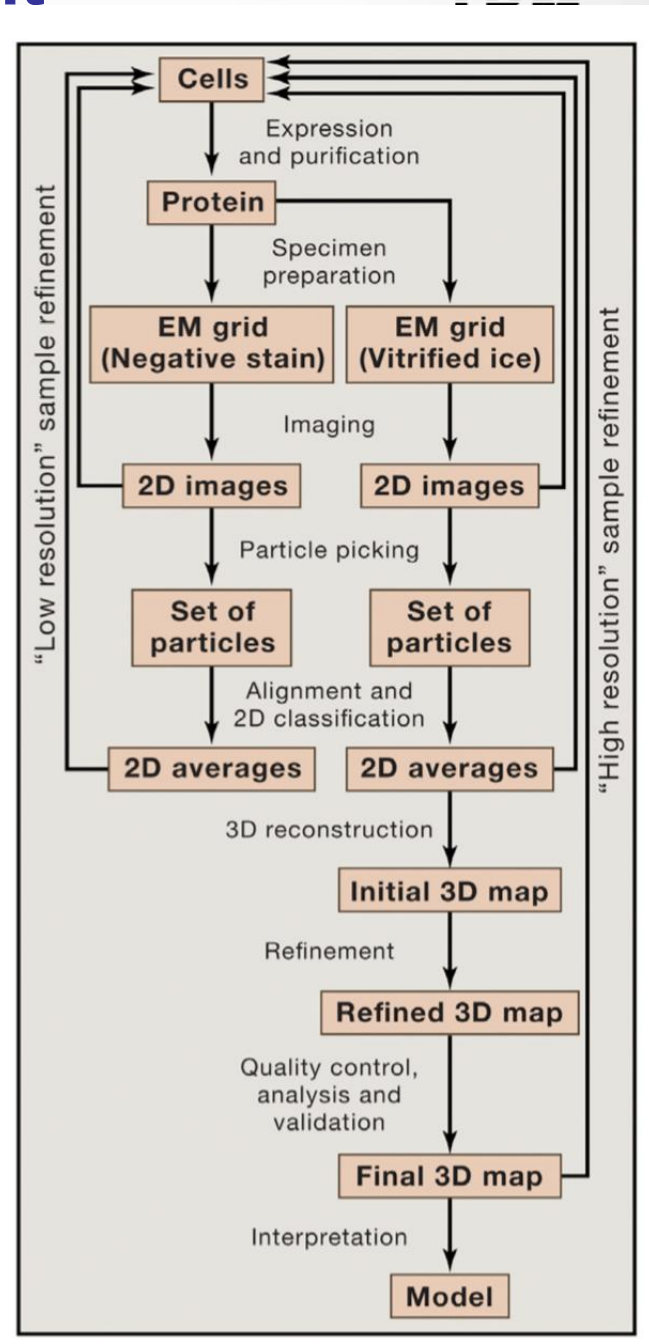

RELION: High Level Flowchart

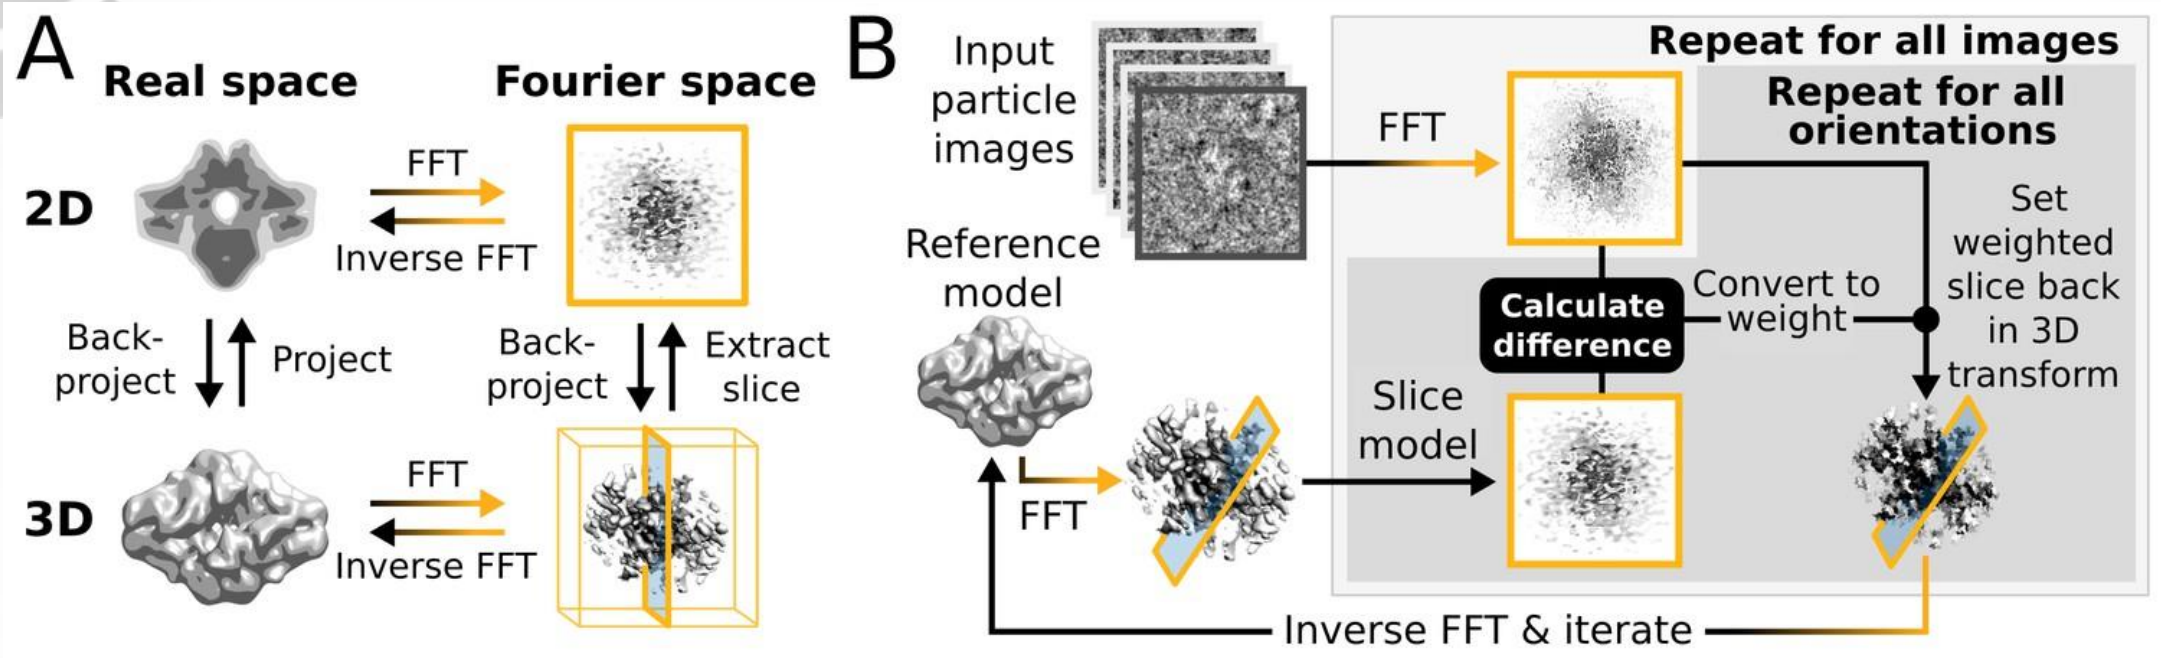

A) Operations and the real vs. Fourier spaces used during B) image reconstruction in RELION.

Source: https://elifesciences.org/articles/18722

Benchmark Data: *Plasmodium falciparum* 80S ribosome

105,247 particles

360 x 360 pixels

Reference Map

Map File Size	182 MB
Voxel Dimensions (Å)	1.34 x 1.34 x 1.34
Map Dimensions (voxels)	360 x 360 x 360
Map Data Type	Image stored as Reals
Density Statistics	**Min** -0.545 **Max** 0.962 **Avg** 0 **Std** 0.044
Recommended Contour Level	0.18

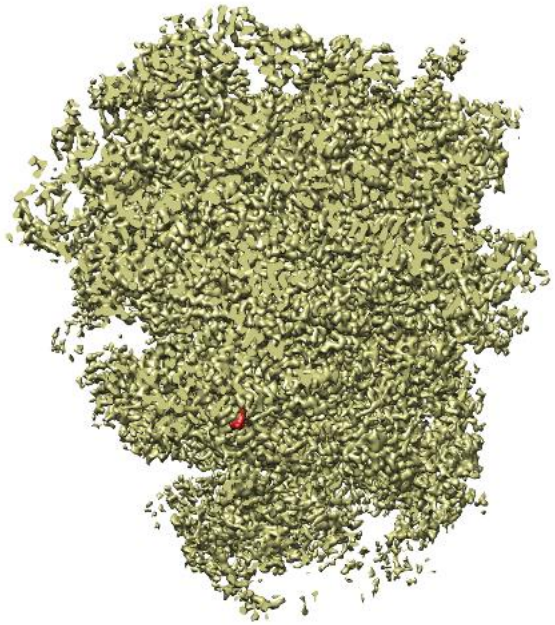

Resolution 3.2 Å

(https://www.emdataresource.org/EMD-2660)

RELION performance: Power9 T4 vs. x86 V100

- Power9-T4: 40 cores, 1TB DDR4 memory, 4 Tesla T4 GPUs
- X86: Xeon Gold 6140 (Skylake), 40 cores, 0.5TB DDR4 memory, 4 Volta 100 GPUs

Table C. RELION Single Node Performance

# No. GPUs	P9-T4	X86-V100	P9 vs. x86
2	61.69	119.13	1.93
4	45.67	83.00	1.82

- RELION benchmark dataset is downloaded from official website[1]: 105,247 particles; 360 x 360 pixels
- X86 benchmark results were obtained from the Ward Lab of The Scripps Research Institute [2]
- GPUs connected with PCIe ports. The P9 results are optimized with IBM ESSL and Spectrum MPI and runtime options for SMTs. X86 results are optimized with MKL.
- P9 with Tesla T4 are more than 80% better than x86 with Volta 100. **P9 + T4 delivers superior price performance.**

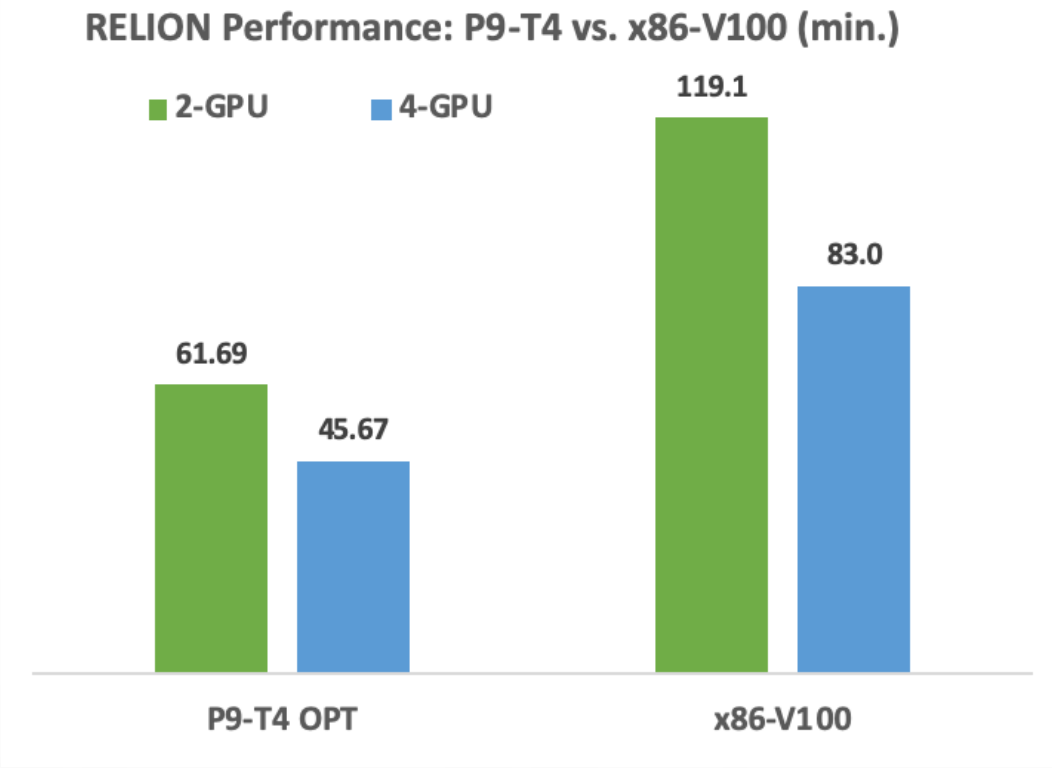

(lower is better)

RELION performance: Power vs. x86 vs. DGX

- Power9: AC922, **1 node**, 4 V100 GPUs.
- Power9-T4: 1 node, 4 Tesla T4 GPUs
- Power8: Minsky, 1 node, 4 P100 GPUs
- X86: Xeon Gold 6140, **1 node, 4** V100 GPUs.
- DGX: DGX-1, **one node** with 8 GPUs

Table C Benchmark Performance Comparisor

P8-P100	P9-T4	P9-V100	X86-V100	DGX-1 V100
71.02	45.67	42.33	83.00	77.00

- RELION benchmark dataset is downloaded from official website[1]: 105,247 particles; 360 x 360 pixels
- X86 and DGX-1 benchmark results were obtained from the Ward Lab of The Scripps Research Institute [2]. DGX-1 not using "—preread_images", which requires additional memory.
- Power9 outperforms x86 and DGX-1 with the same number of GPUs.

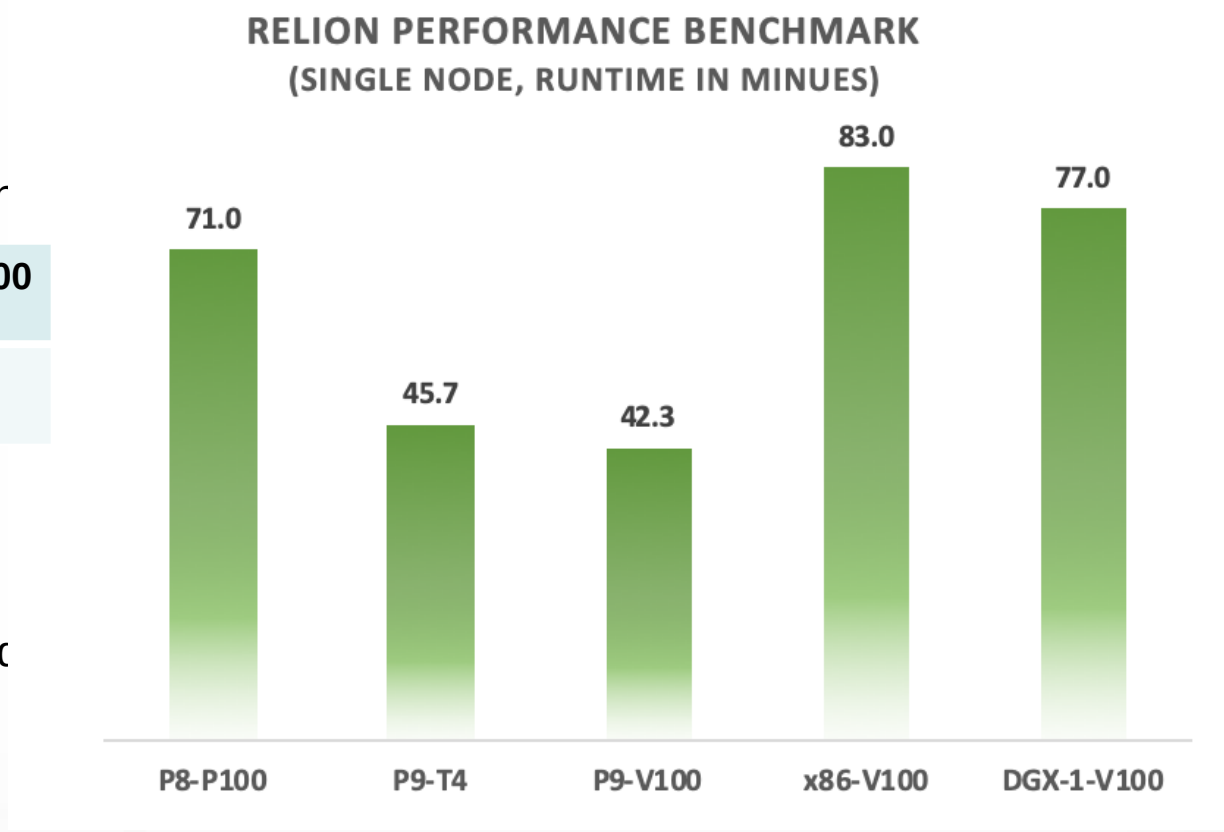

(lower is better)

Anaconda Environment for Applications

- Use anaconda enterprise network (AEN) to manage cryo-EM software repository on server.
- Easy to use and update software

Anaconda Architecture for Cryo-EM Analysis

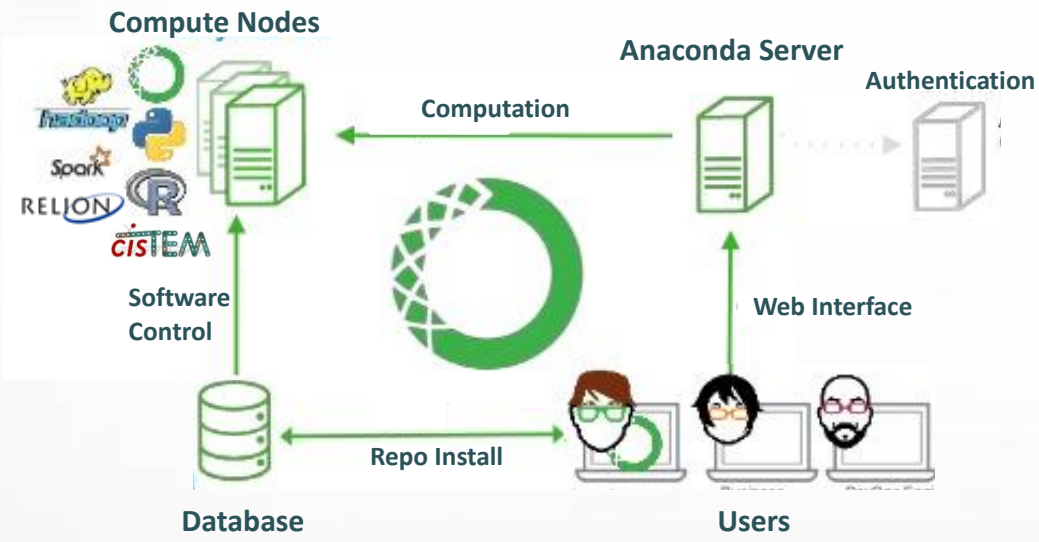

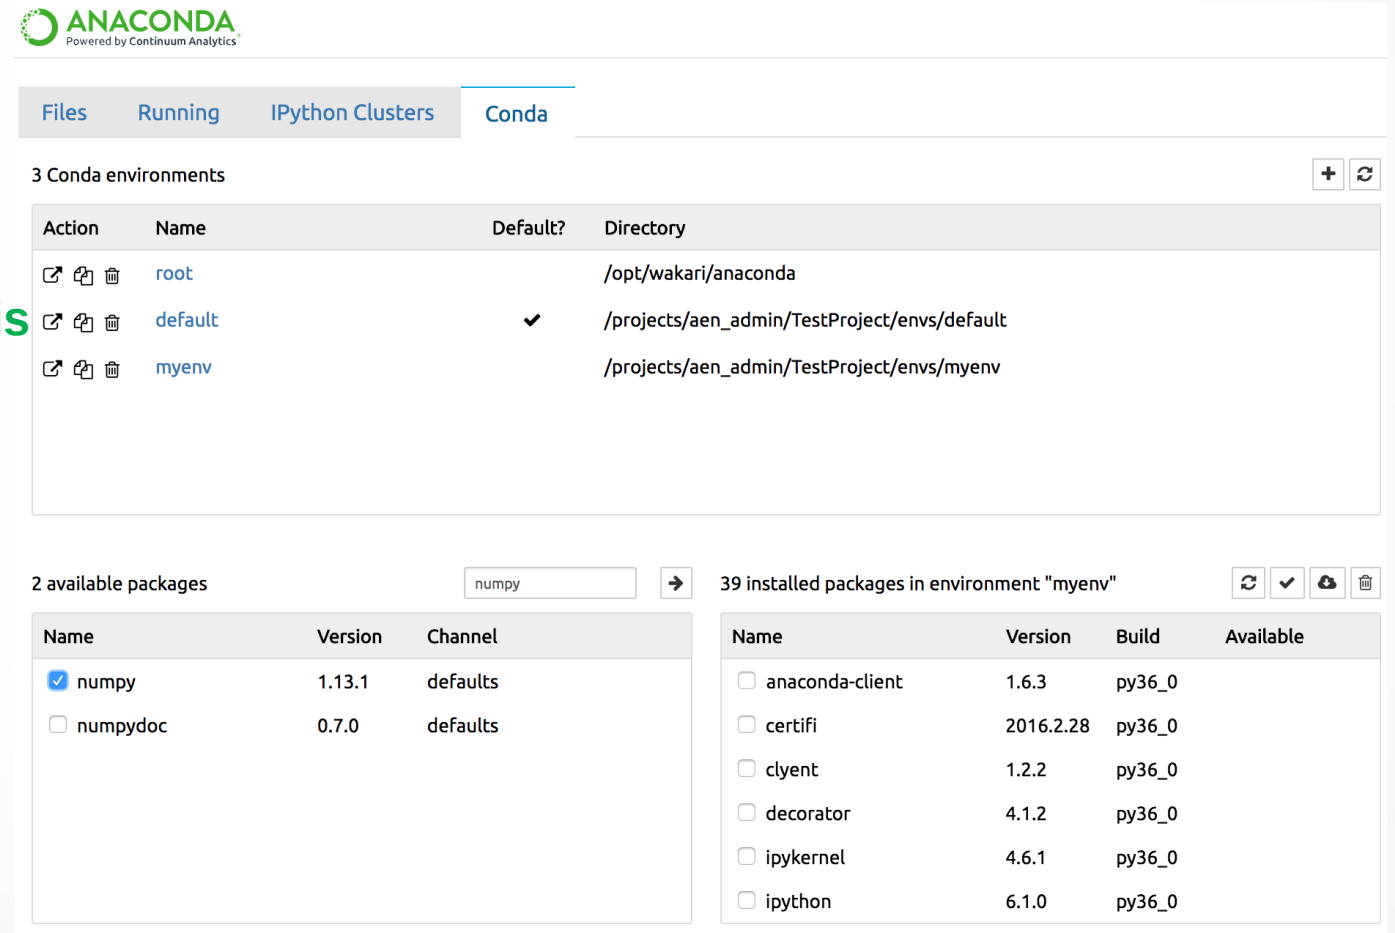

ACP Solutions for Genomics

Power Systems

Welcome to the waitless world

IBM

Small	Medium	Large
1-2 samples per day	4-6 samples per day	10+ samples per day
1 40-core AC922 node, 512 GB Mem	2 40-core AC922 nodes, each 512GB Mem	4-8 40-core AC922 nodes, each 512 GB Mem
GL2 (GPFS) with IB adapter	GL4 (GPFS) with IB adapter	GL4 (GPFS) with IB adapter
Ethernet	IB or Ethernet	IB or Ethernet
Optimized genomic software stack running on RHEL 7.6 with 4 SMTs per core.		

Optimized GATK4 Best Practice Pipeline Runtime

Dataset	30x WES	50x WGS	Reference
NA12878	12 min.	8.7 hrs.	GRCH37 (HG37)

Perf — GATK4 Best Practice Pipeline

STEPS	WES	WGS
Bwa + samtools	1.92	215.4
markDuplicates	1.74	63.0
gatkBaseRecalibrator	0.58	19.8
gatkApplyBQSR	0.40	24.6
gatkHaplotypeCaller	0.99	150.1
gatkCombineGVCFs	1.73	-
gatkGenotypeGVCFs	1.40	3.0
gatkVariantRecalibratorSNP	2.23	31.8
gatkApplyVQSRSNP	0.15	1.8
gatkVariantRecalibratorIndel	0.25	9.5
gatkApplyVQSRIndel	0.14	1.8
Total Runtime (minutes)	112.27	520.8

Ruzhu Chen @ 2019

GATK4 WES Pipeline Performance: Power9 vs. X86

Using 30x coverage WES dataset (gcat_set_025) to run benchmark GATK4 Germline pipeline with reference genome GRch37 on both Power9 (PowerNV 8335-GTC) and x86 (Gold 6148 @ 2.40GHz). The performance results shown below.

GATK4 Germline Runtime

Dataset	Power CPU	X86 CPU	Ref Genome
30x WES	12 min	133 min	GRCH37 (HG37)

Power9 run **11x** faster than x86 Skylake with optimized scripts, while the output VCF files are **95 - 99%** concordance between the two architectures.

STEPS	Power9	X86 std
Bwa + samtools	1.92	5.07
markDuplicates	1.74	5.29
gatkBaseRecalibrator	0.58	9.89
gatkApplyBQSR	0.40	2.88
gatkHaplotypeCaller	0.99	101.92
gatkCombineGVCFs	1.73	-
gatkGenotypeGVCFs	1.40	2.59
gatkVariantRecalibratorSNP	2.23	3.57
gatkApplyVQSRSNP	0.15	0.14
gatkVariantRecalibratorIndel	0.25	0.98
gatkApplyVQSRIndel	0.14	0.08
Total(minutes)	12.27	133.32

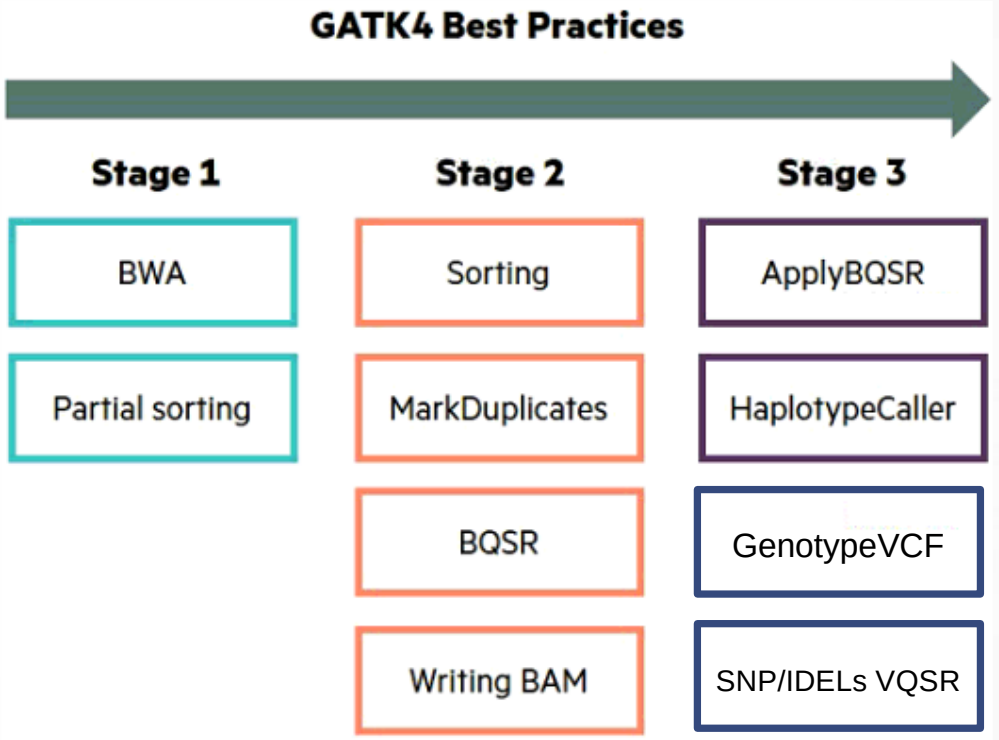

Source : https://software.broadinstitute.org/gatk/best-practices/workflow?id=11145

GATK4 WGS Pipeline Performance: Power9 vs. Skylake (x86)

- GATK: Genomic Analysis Toolkit, a JAVA-based tool used for genomic sequence variant analysis
- Runs poorly "out of the box" on POWER9 relative to Skylake, 8% slower, core-to-core comparison
- Improves performance by software and runtime optimization nearly **10x times**

Dataset	Power9 CPU	X86 CPU	Ref Genome
50x NA12878	85h	79h	GRCH38 (HG38)

>>Most tools run with single threads
- Poor single core/thread performance on Power9

>>SIMD capacity
- 128-bit P9 VSX vs. 512-bit x86 AVX

>> Thread scheduling
- Active threads are fewer on P9
- Thread waiting

GATK Pipeline STEPS	Power9	X86
BWA + SortSAM	4.66	12.61
markDuplicates	8.65	8.15
gatkBaseRecalibrator	9.49	8.39
gatkApplyBQSR	9.83	6.11
gatkHaplotypeCaller	51.73	42.88
gatkGenotypeGVCFs	0.48	0.29
gatkVariantRecalibratorSNP	0.40	0.38
gatkApplyVQSRSNP	0.03	0.02
gatkVariantRecalibratorIndel	0.15	0.10
gatkApplyVQSRIndel	0.02	0.02
Total(hours)	85.44	78.95

After Optimization

Dataset	Power9 CPU	X86 CPU	Ref Genome
50x NA12878	8.7h	23h	GRCH38 (HG38)

>> Benefit from SMTs
- Split data to use all SMT threads concurrently.

>>Use optimal code
- Modify marking duplicates tool

>> Workloads
- Rebalance workloads

GATK Pipeline STEPS	Power9	X86
BWA + SortSAM	3.59	10.61
markDuplicates	1.05	8.21
gatkBaseRecalibrator	0.33	0.30
gatkApplyBQSR	0.41	0.29
gatkHaplotypeCaller	2.50	2.71
gatkGenotypeGVCFs	0.05	0.26
gatkVariantRecalibratorSNP	0.53	0.36
gatkApplyVQSRSNP	0.03	0.02
gatkVariantRecalibratorIndel	0.16	0.10
gatkApplyVQSRIndel	0.03	0.02
Total(hours)	8.68	22.88

Molecular Dynamics

Using computational approach to simulate atomic motions, to model macromolecule receptors and their small-molecule ligands in Covid-19 drug discovery, design and prediction.

Molecular dynamics simulation

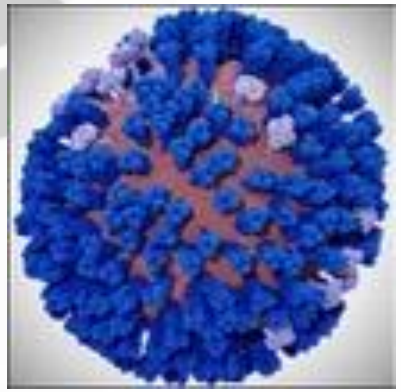

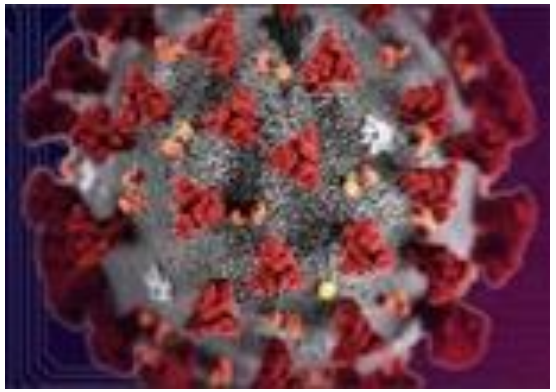

A) Using NAMD to simulate influenza virus (left)and Covid-19 (right)

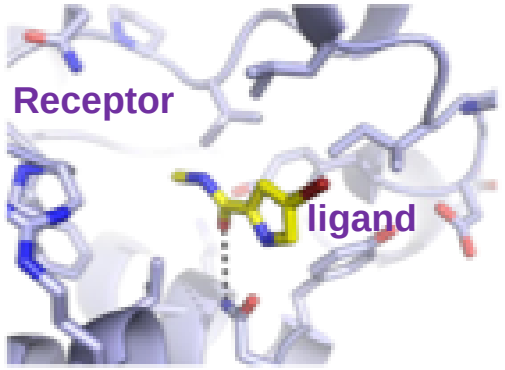

B) Drug discovery: protein receptor

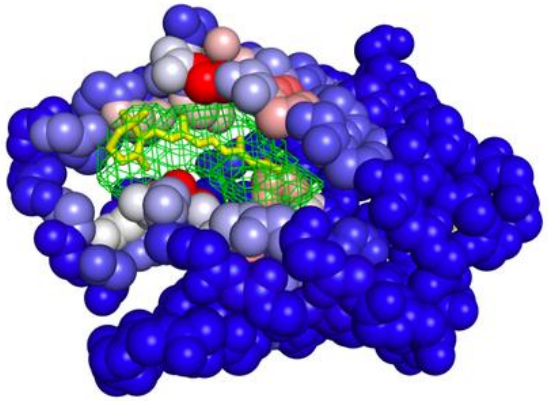

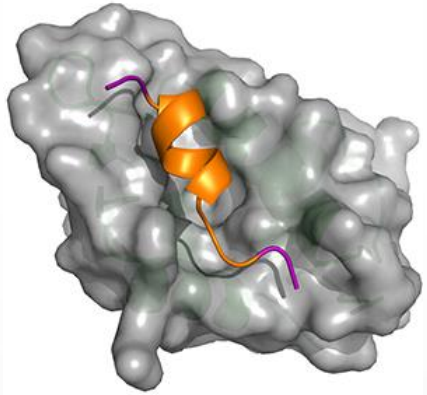

C) *In silico p*rediction of protein cryptic binding site

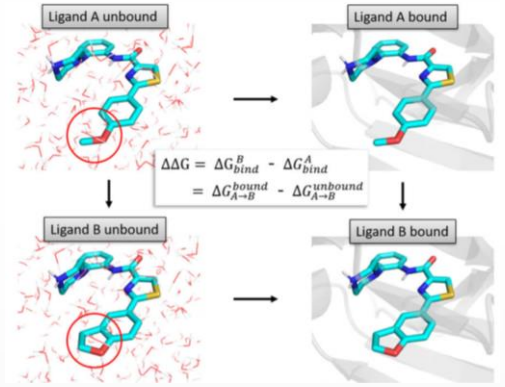

D) Predicting protein receptor ligand binding energy

NAMD STMV benchmark

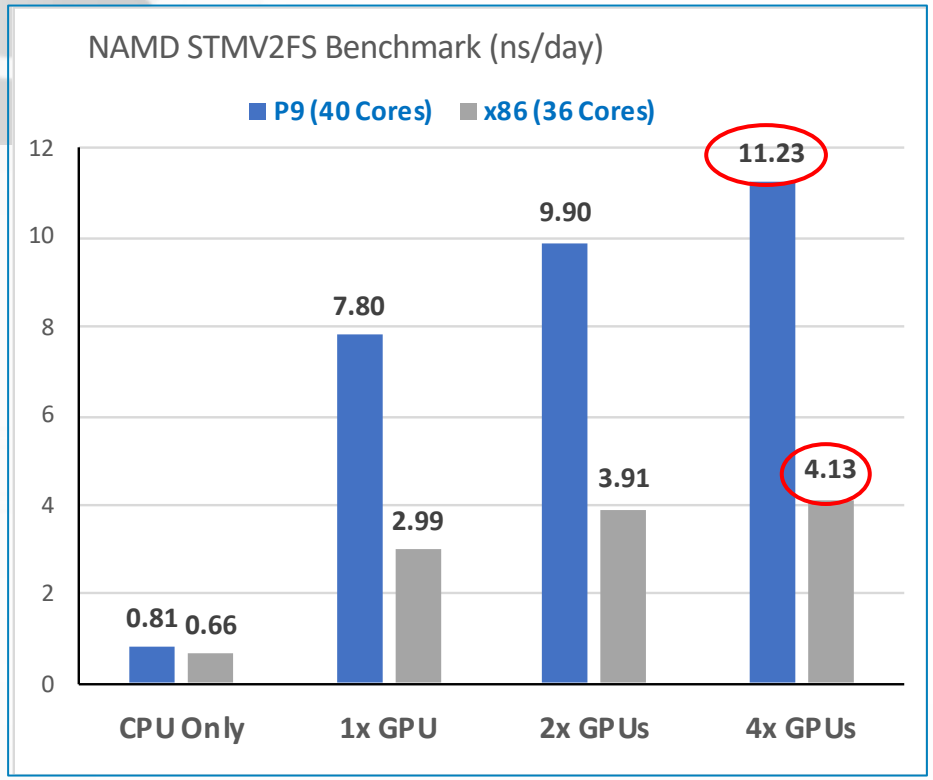

- ~1 million atoms are simulating on Power9 and x86 with the same configurations.
- Higher is better.

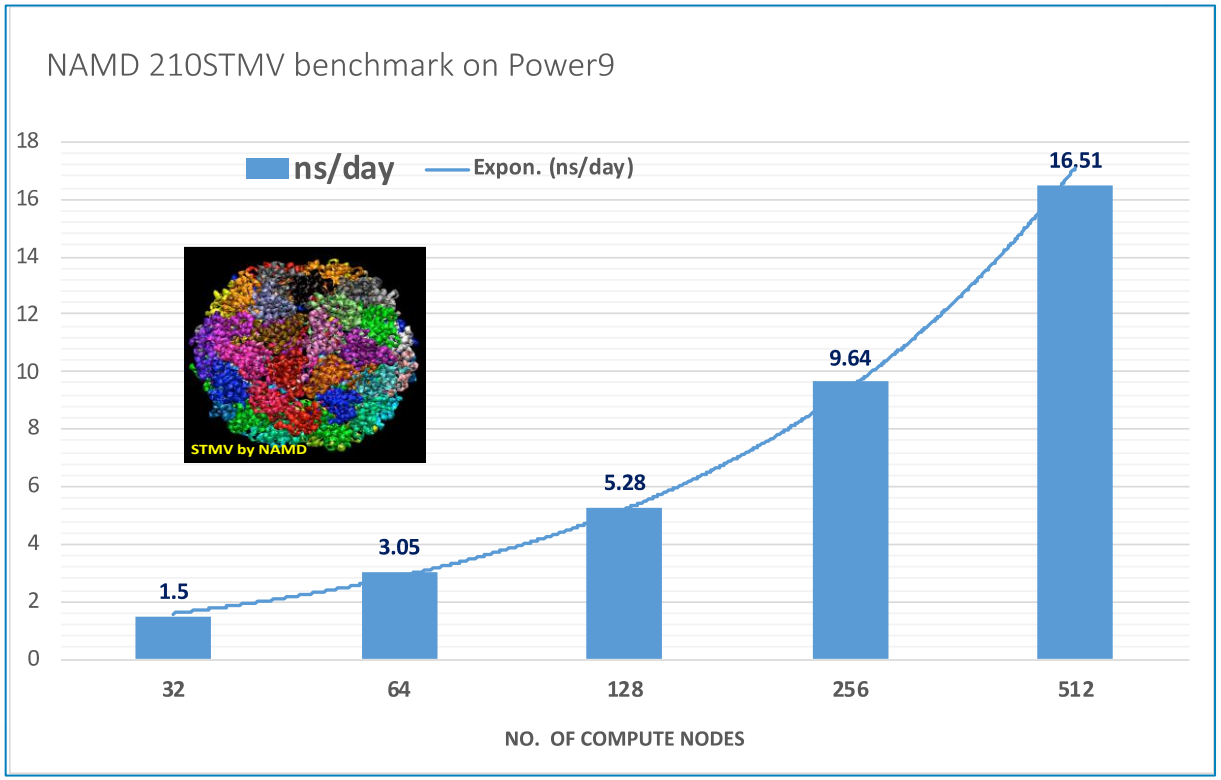

- 221 million atoms are simulating on Summit cluster with memory optimization.
- The simulation uses 4 GPUs per node.

AMBER 18 performance comparison

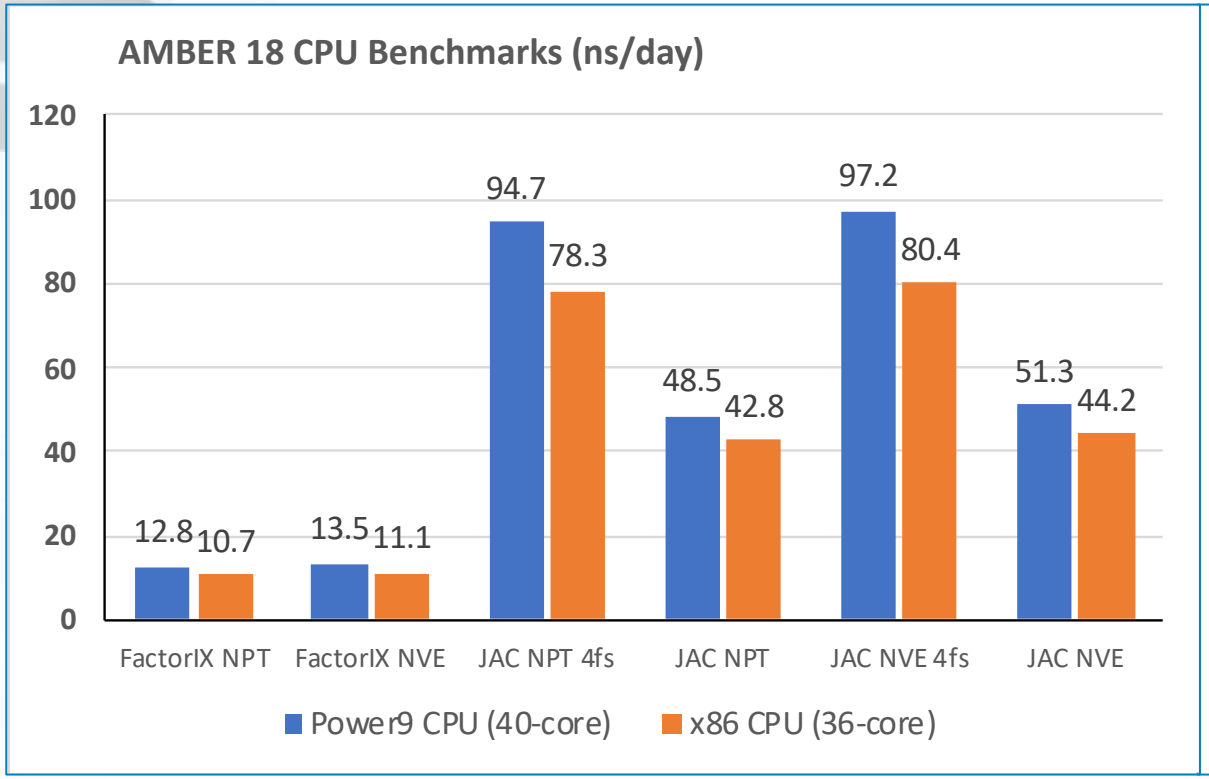

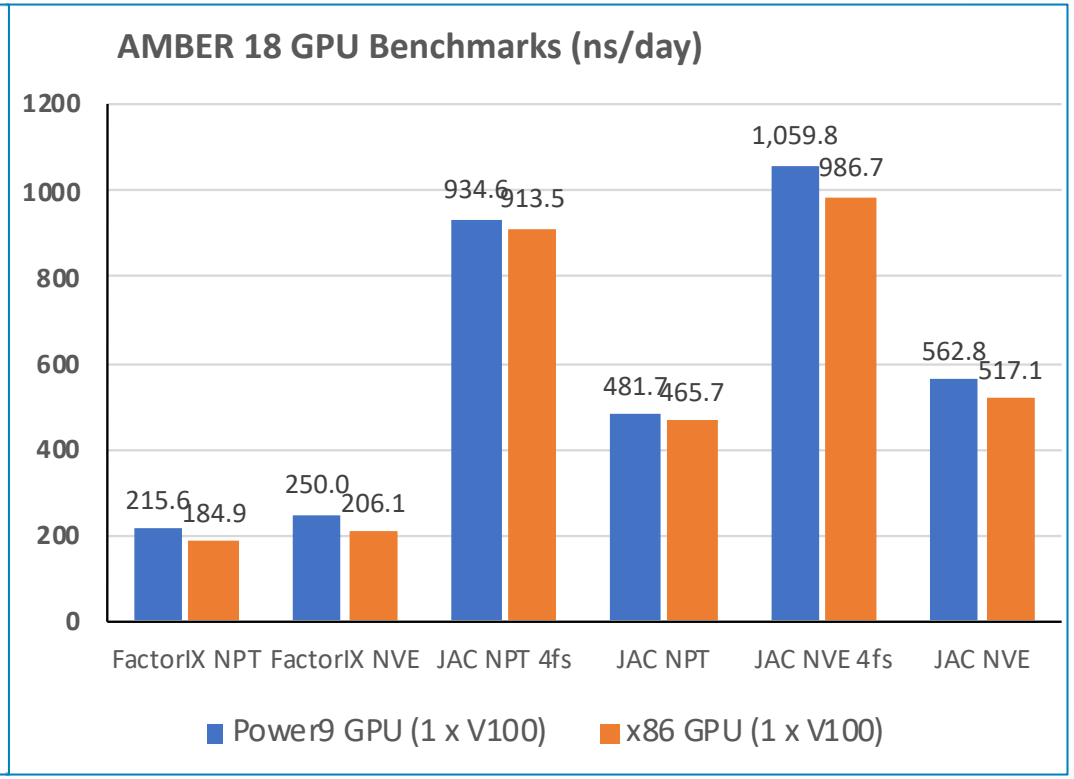

- Power9 AC922 outperforms x86 on both CPU and GPU in all benchmark executions.

Brief Intro of IBM Bayesian Optimization Accelerator.

- BOA is an AI Software package.
- Subscription Service license per year.
- Licensed by the # of BOA Server(s) (only runs on IBM AC922).
  - Sizing is dependent on the number of simultaneous Jobs BOA is processing against.
- BOA software generally completes the AI analysis in under 1 minute after the input from the “interface” is received.
- Jobs on external HPC systems can be paused to wait for BOA input for the next run.
- BOA is very specific to each scientific discipline and use case.
  - The Interface is “ re-usable” custom Python code per analysis type.

BOA Topology

- BOA servers are dedicated to running BOA only
- Physically co-located with the HPC environment
- Large BOA systems include multiple GPU enabled nodes for throughput
- Multi-user, multi simultaneous experiments being simulated in the HPC

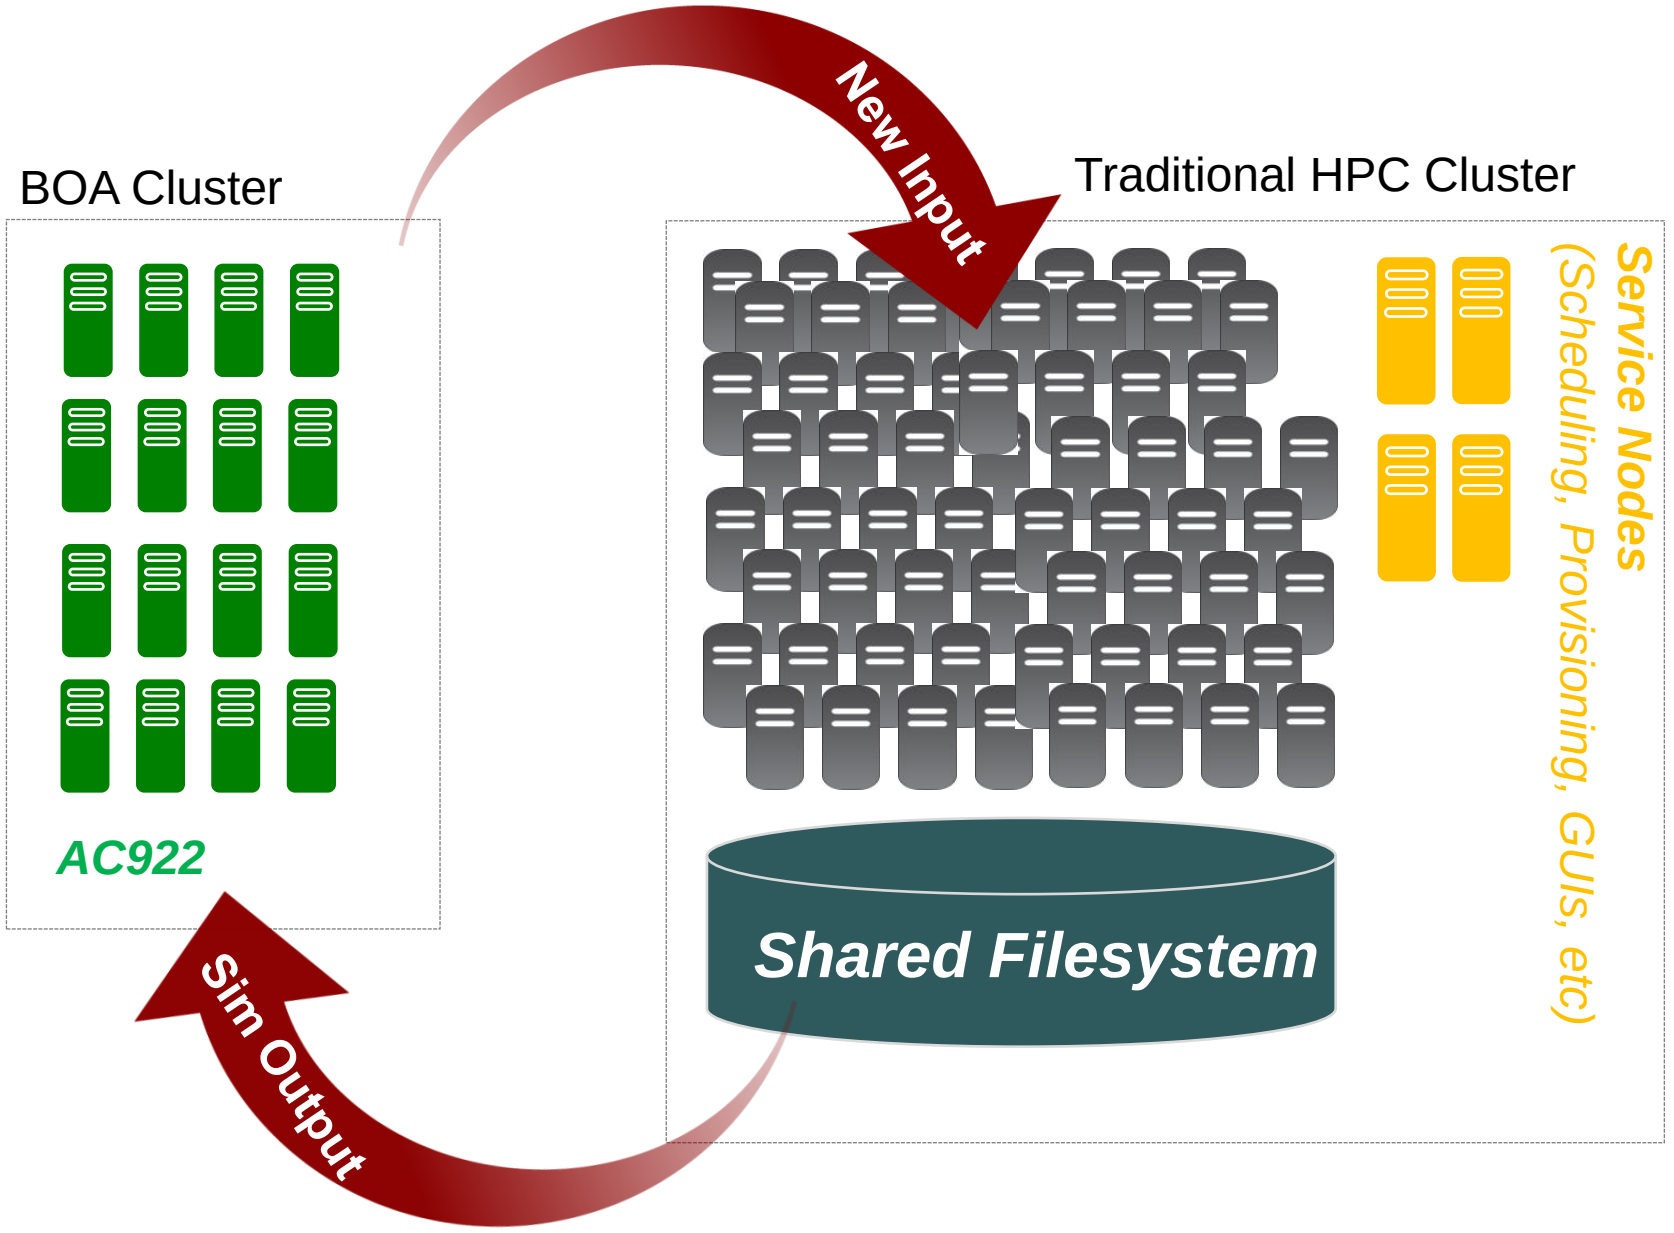

Real Examples of HUGE Value with IBM BOA

Computational Chemistry

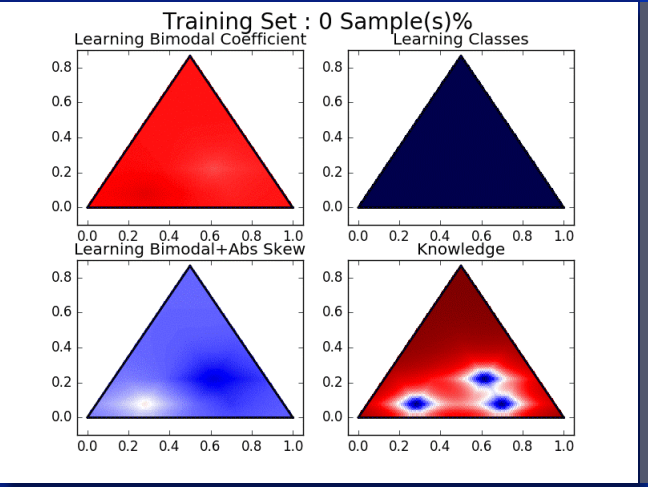

BOA accelerated workflow uses **1/3 of the calculations to achieve 4 orders of magnitude resolution increase**

Computational Fluid Dynamics

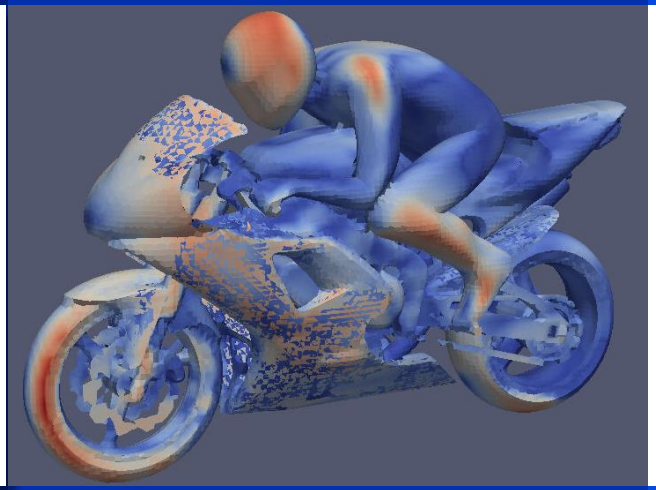

The BOA enabled design process accomplishes in hours what a designer does in weeks

Chip Design

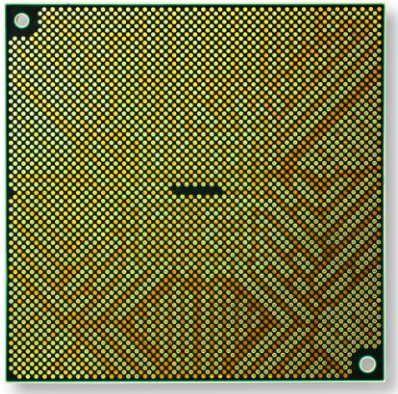

Orders of magnitude fewer simulations required to minimize jitter for signal integrity on the Power10 server

Oil & Gas

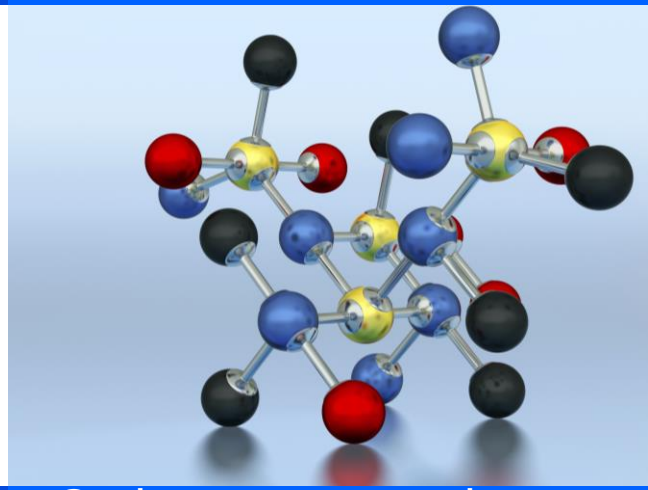

Carbon sequestration use case where BOA identified an answer 1.5% better in 1/3 the time

Interface Functions

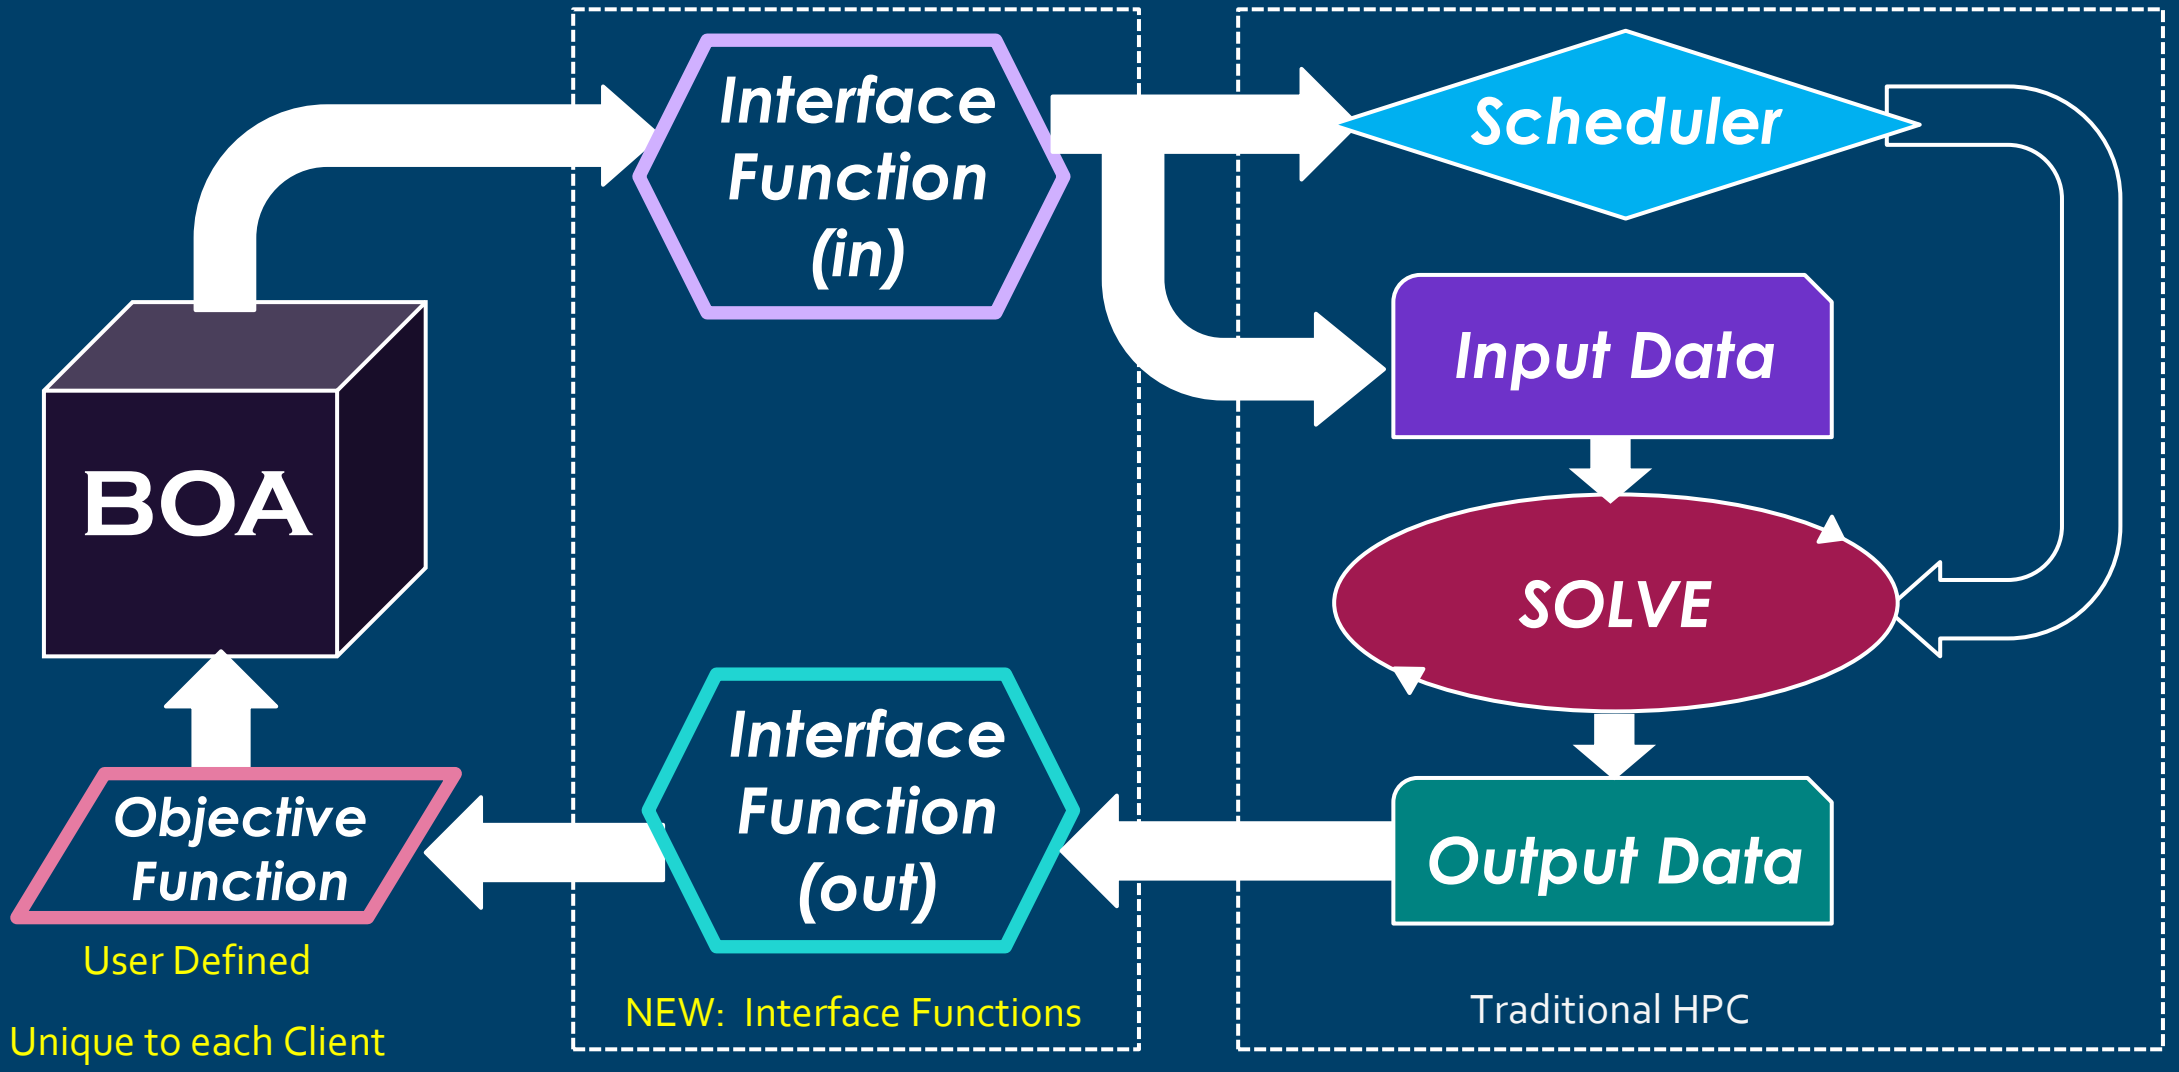

Value Statements

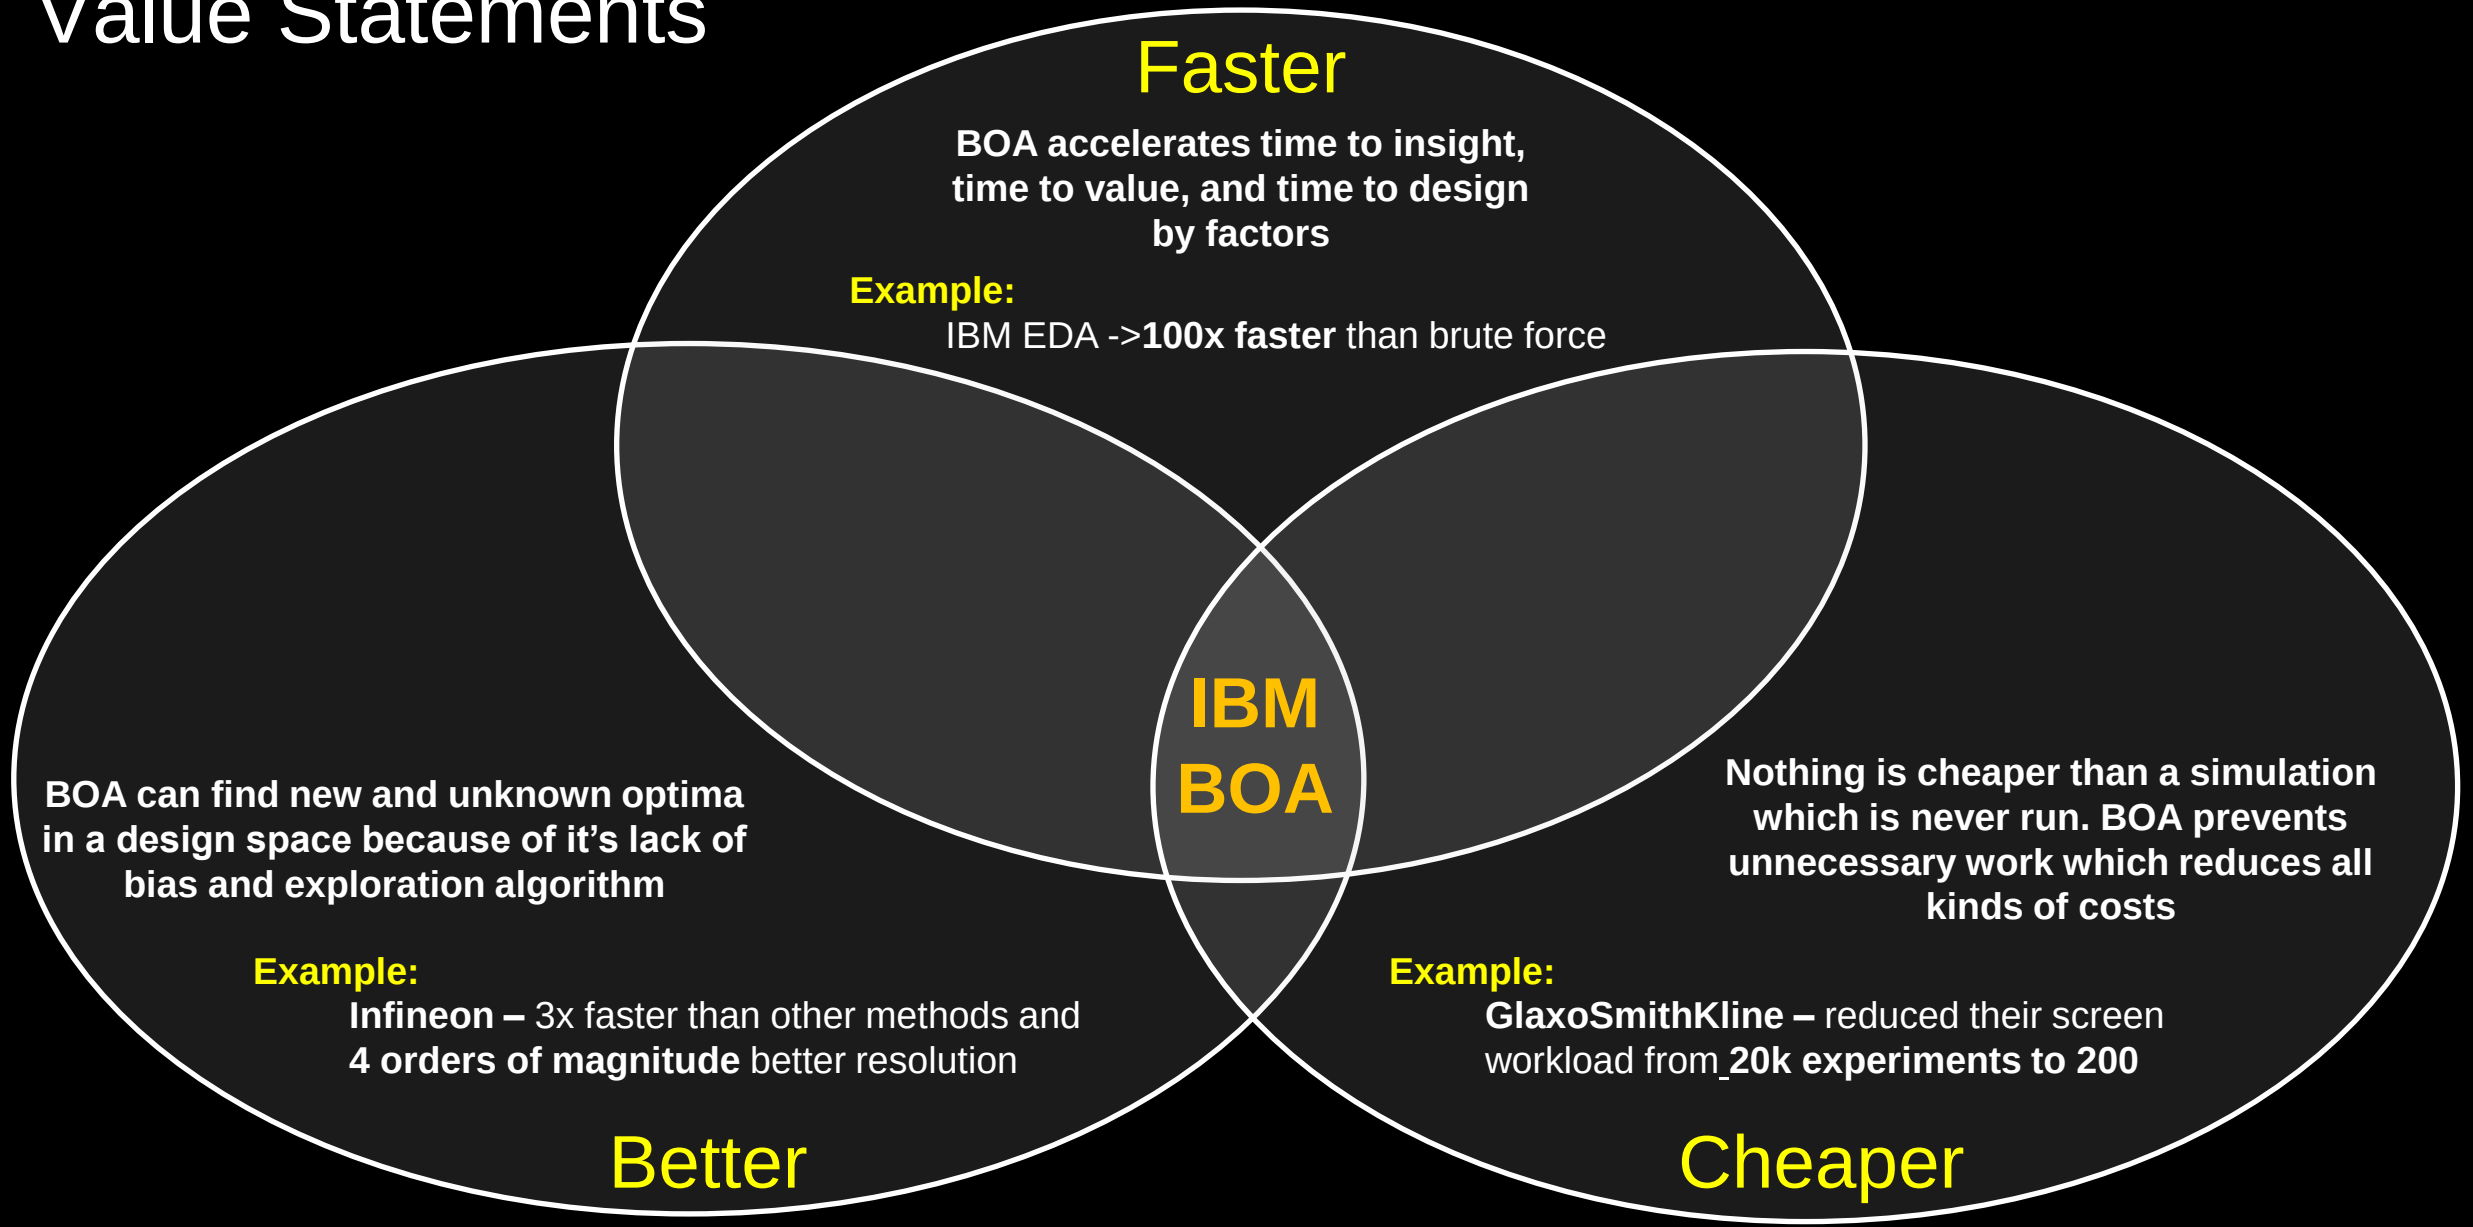

Formula 1 Racing

simplified nose design

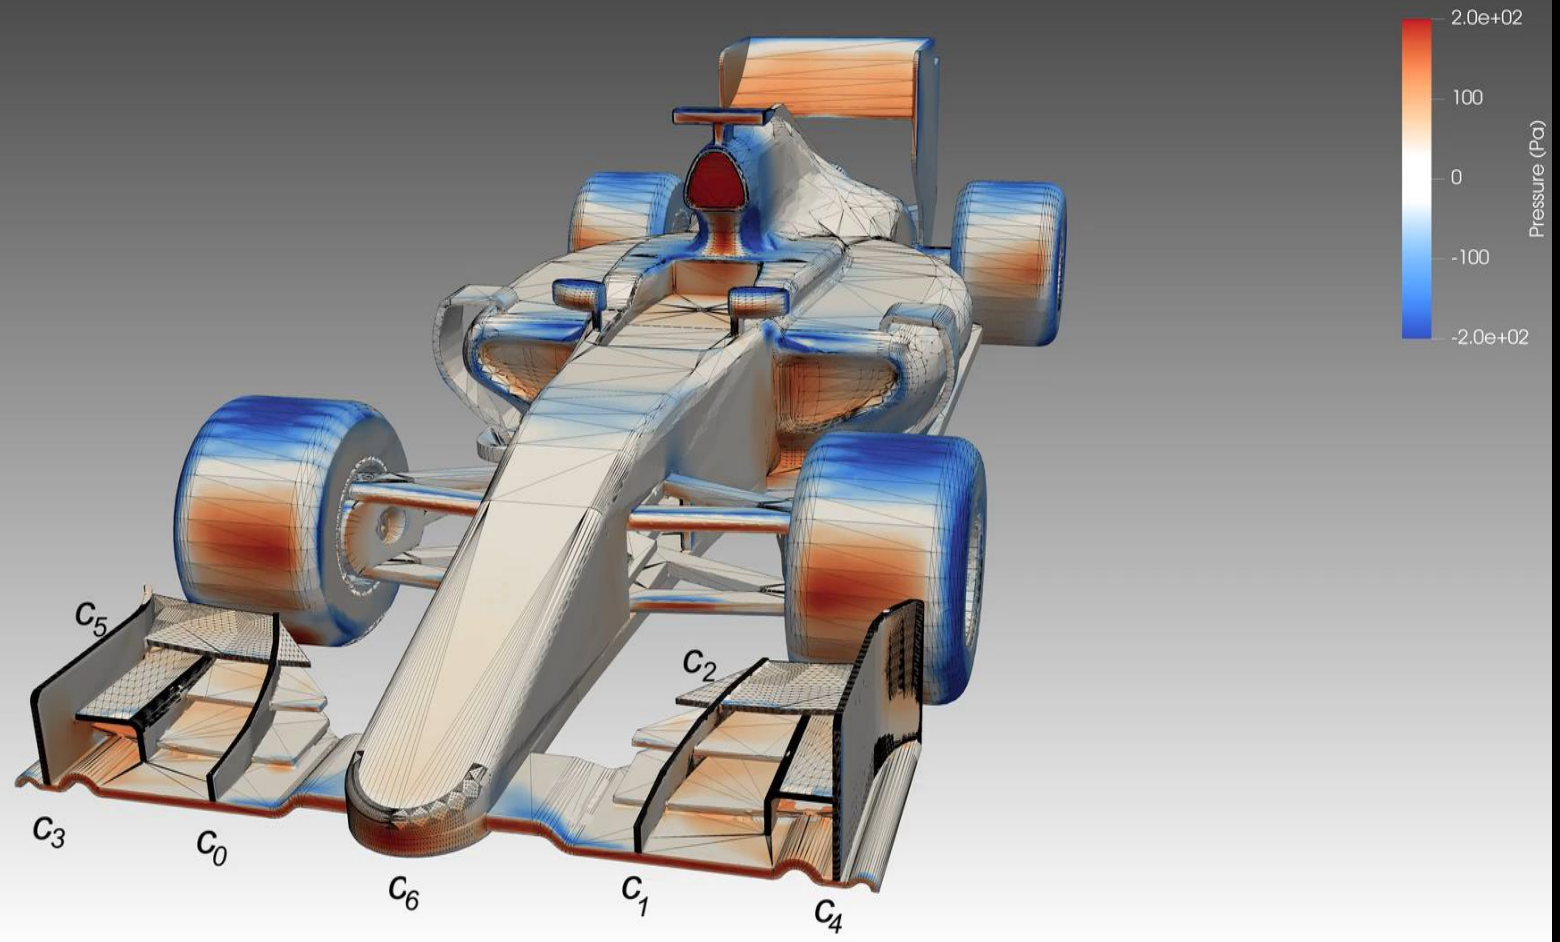

- 7 design variables corresponding to geometry of vanes in the wing
- openFOAM CFD simulator used to calculate DRAG @ each iteration
- Normally this kind of design is done by hand by aerodynamics experts

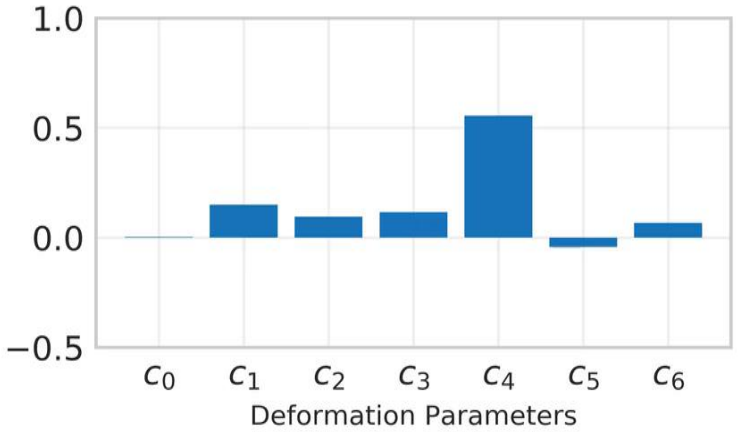

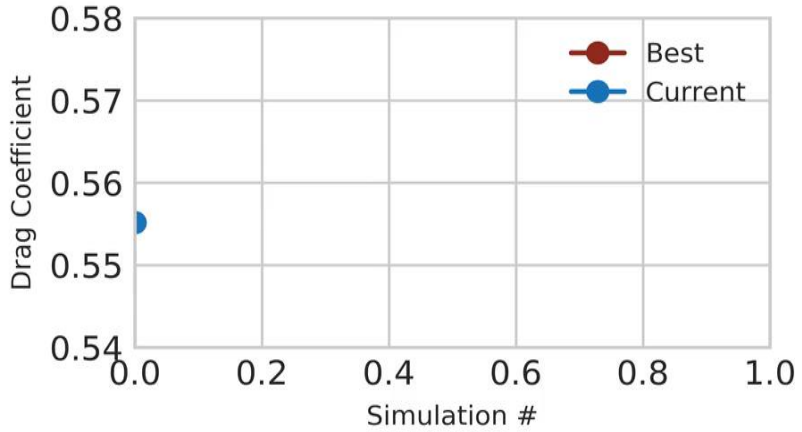

IBM Power Electronics Design

Real world use case

Signal Integrity Design for Power10

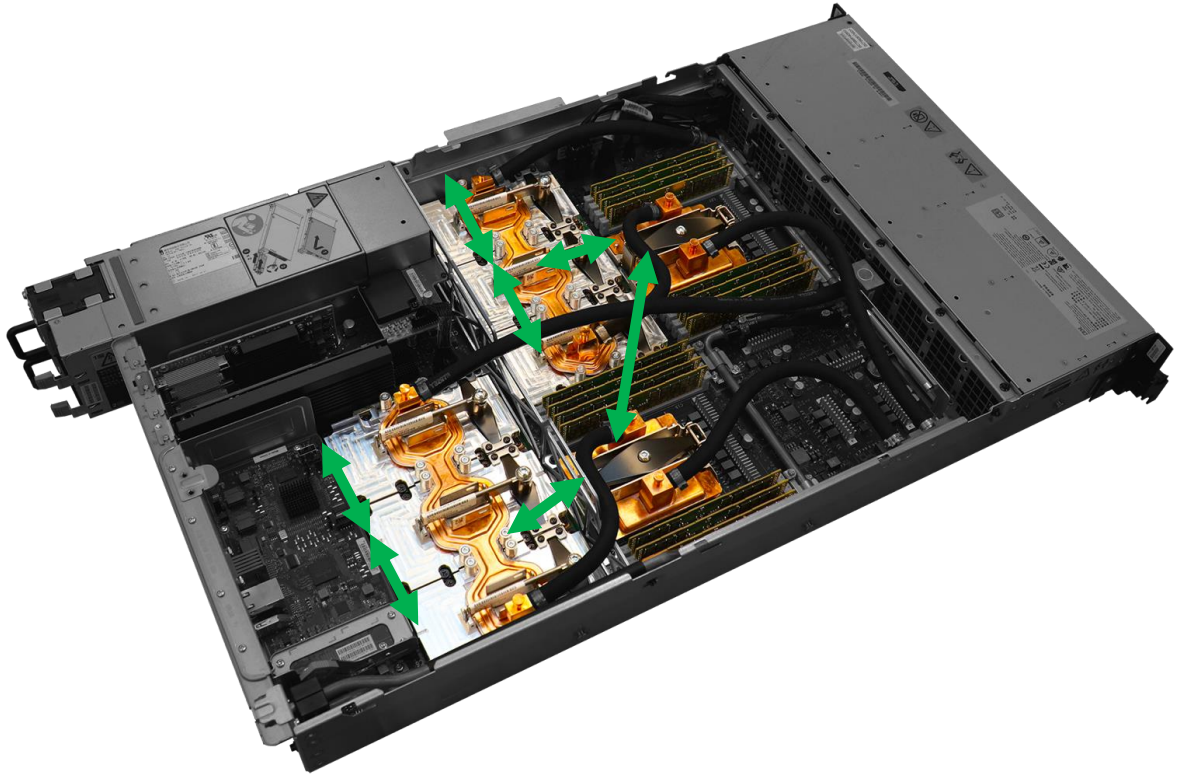

Design Challenge: Optimize the communication speed of the links shown in green using minimum time & compute resources

Brute force method for design space exploration requires ~5600 simulations, each taking ~20 minutes (IBM Simulator: HSSCDR)

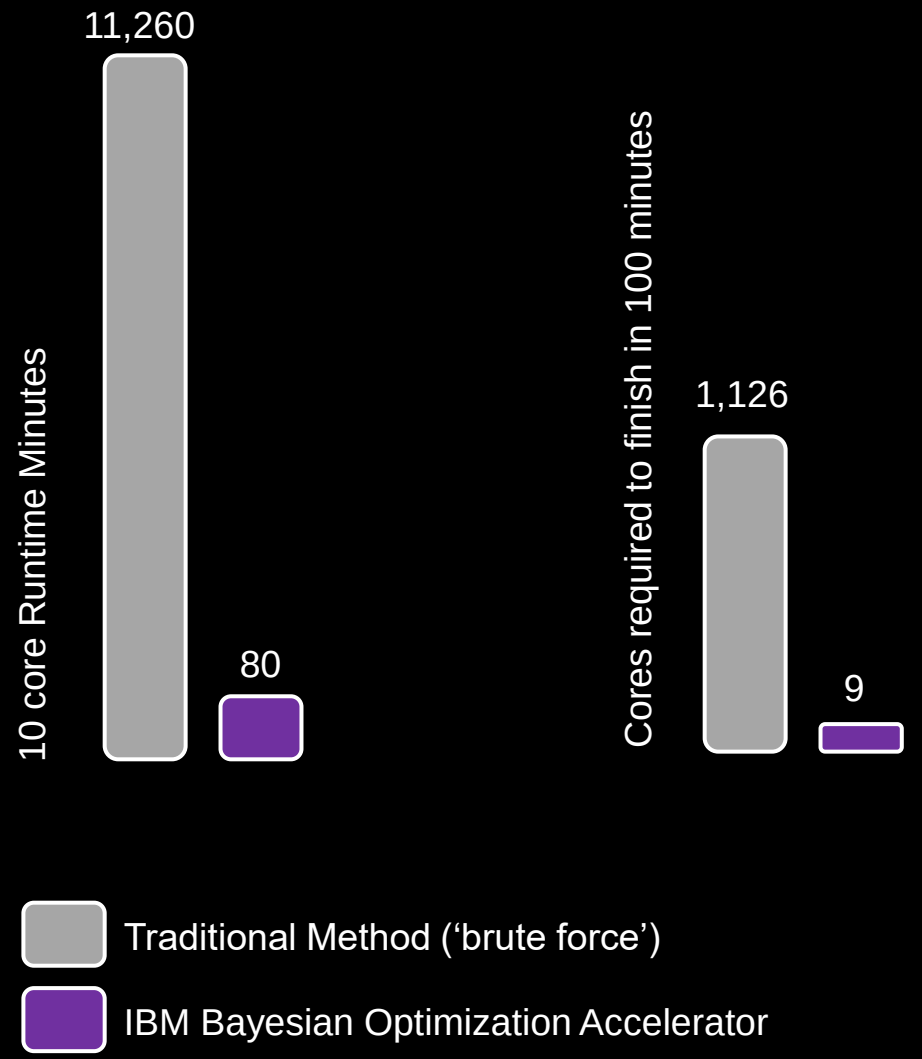

Server Technologies
“The Engine”

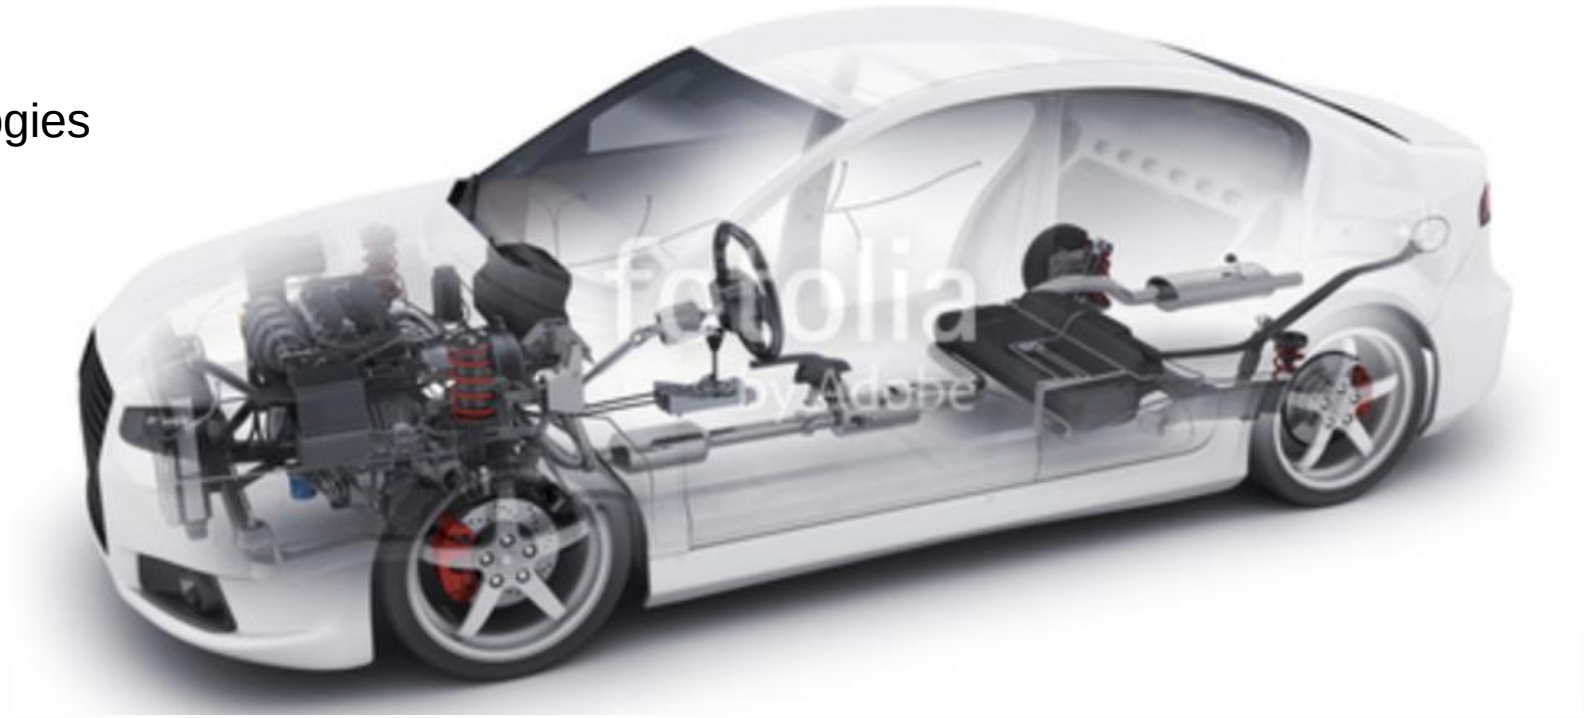

A Portfolio for the AI Era

PowerVM and high RAS

L922

- Industry leading reliability and computing capability
- PowerVM ecosystem focus for outstanding utilization
- Focus on memory capacity with up to 4TB of RAM

Accelerated Compute

AC922

- Industry first and only in advanced IO with 2nd Generation CPU - GPU NVLink delivering ~5.6x higher data throughput
- Up to 4 integrated NVIDIA “Volta” GPUs air cooled (GTH) and up to 6 GPUs with water cooled (GTX) version
- OpenCAPI support
- Memory coherence

Data, Inferencing, and Cloud

IC922

NEW!

- Storage dense, high bandwidth server – up to 24 NVMe or SAS/SATA in 2U[1]
- Advanced IO with PCIe Gen4
- Optimized inferencing server with up to 6 Nvidia T4 GPUs at GA and additional accelerators in roadmap[1]
- OpenCAPI support[1]
- Price/performance server

Big Data

LC922 / LC921

- Big Data server with up to 120TB storage capacity, large form factor support, KVM support, leveraging P9 compute for a composable design
- 1U and 2U form factors
- Advanced IO with PCIe 4.0/CAPI 2.0
- Up to 44 cores (2U) or 40 cores (1U) at lower frequency

IBM AC922 GPU Server

Realize unprecedented performance and application gains with POWER9 and NVLink 2.0

- Up to 4 GPUs per server for Deep Learning Training
- Contains all the advanced IO – NVLink 2.0, PCIe 4.0 and coherence
- Water cooled option improving data center and computing efficiency

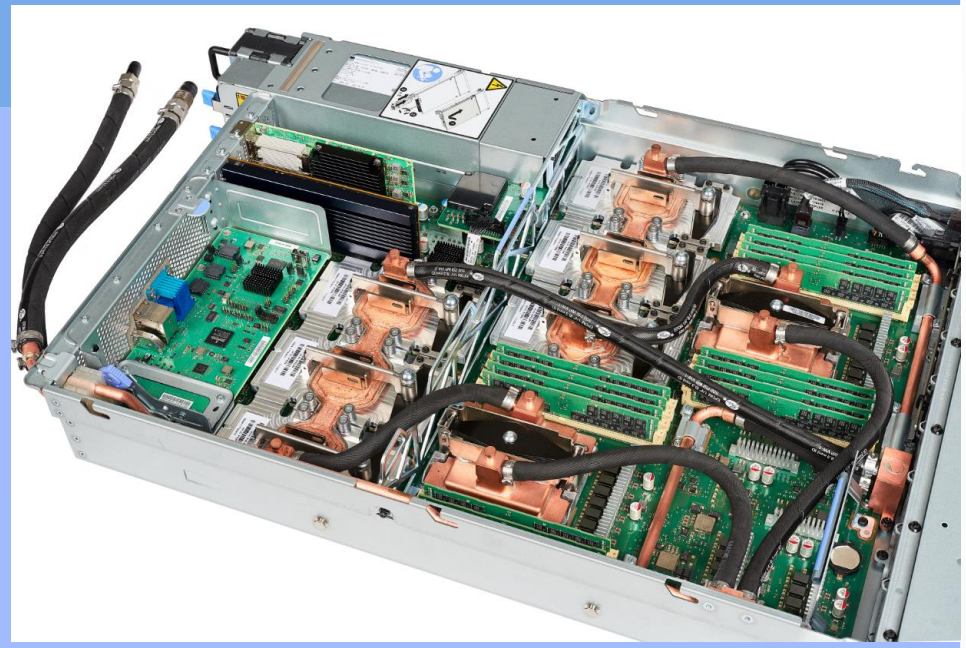

- 2-socket, 2U Packaging
- Processors - up to 40 POWER9 cores
- Up to 4 NVIDIA Volta GPUs
- 2 TB Memory (16 DIMMs)
- 4 PCIe Gen4 Slots
- 2x SFF (HDD/SSD), SATA, 7.7 TB storage
- Supports 1.6TB and 3.2TB NVMe Adapters
- Redundant Hot Swap Power Supplies and Fans
- Air cooled Processor and GPU's for maximum performance
- Default 3 year 9x5 warranty, 100% CRU

GPU Accelerator comparison

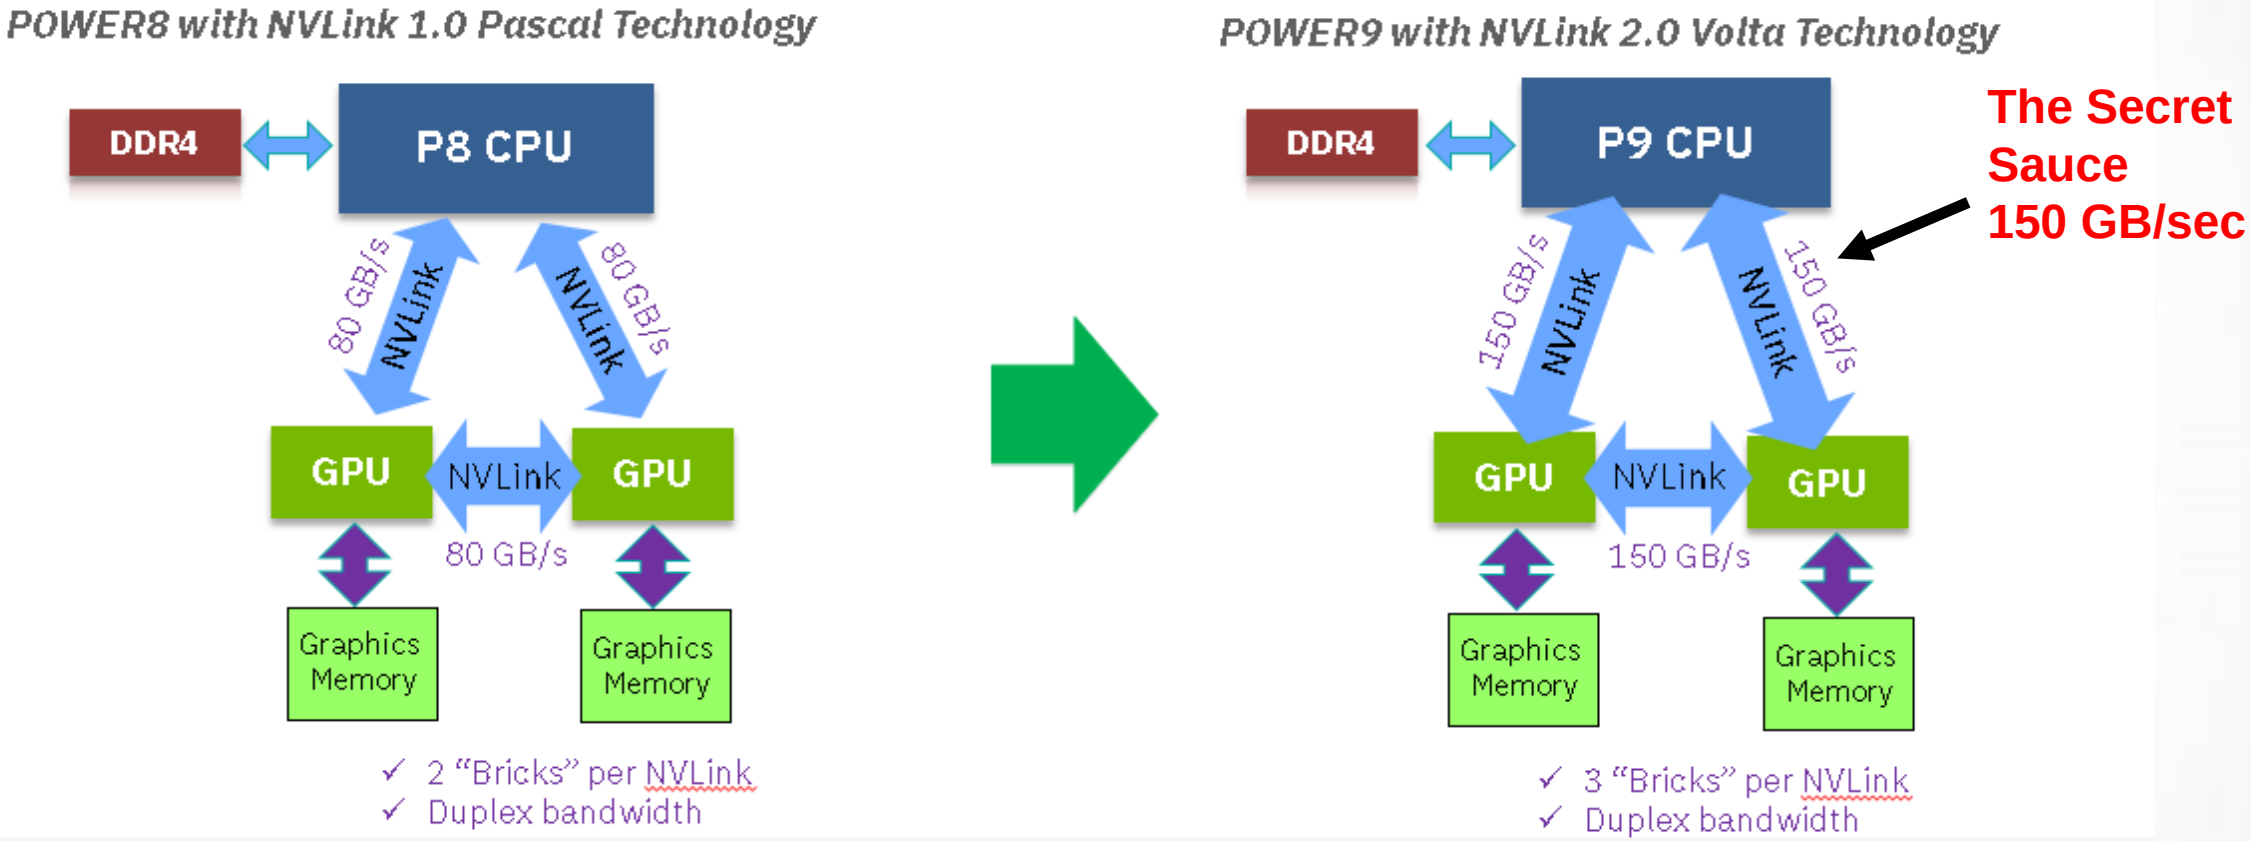

This slide highlights the outstanding advantage in performance the P9 and Volta GPU combination provides over it's predecessor.
The performance bump is achieved in 2 ways.

1) The bus speed increases from 20Gb/s to 25Gb/s.
2) There is a wider bus, increasing from 2 Bricks to 3 Bricks. (1 Brick contains 8 lanes)

The combination of bus speed and increased bus width allows for the 87.5% performance jump.

Designed for data to deliver performance at scale

4X
Threads per core*

Parallel Processing

4X
Mem. Bandwidth*

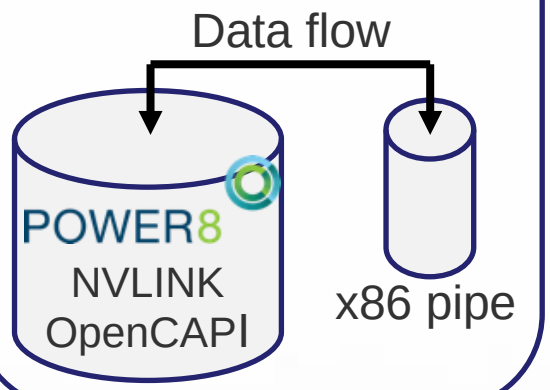

4X
More cache* @ lower latency

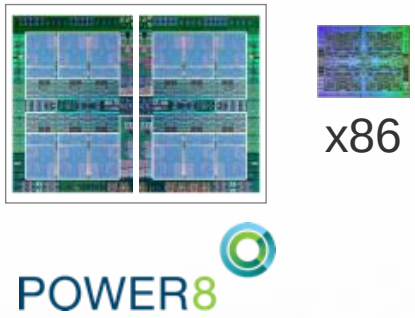

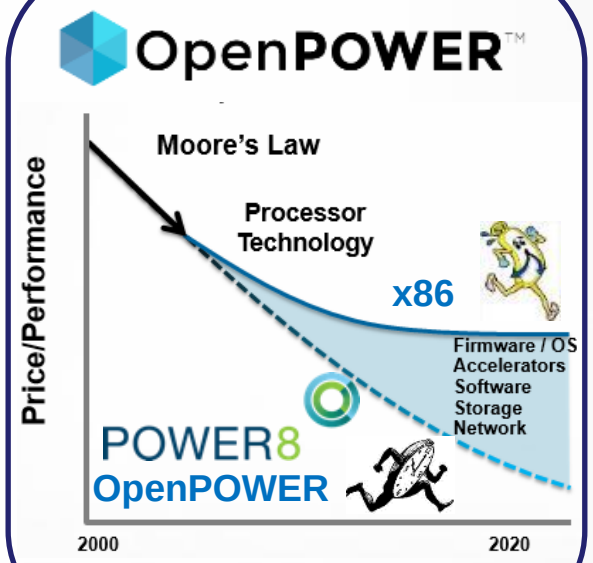

These design decisions result in best performance for data centric workloads like:

NoSQL and Relational Databases, Big Data Analytics, OLTP, ML/DL

SMT=Simultaneous Multi-Threading

Power System IC922 for DATA

NVMe and PCI Gen4 capability designed to be the fastest compute and data server available

Genomics and CryoEM Recommended Server

- **Balanced storage, network, and memory design for optimized storage rich solutions**
  - 33% more bandwidth (340 GB/s DDR BW on IC922 vs. 255 GB/s BW on x86)
  - Better memory capacity capability with 32 DDR4 RDIMM slots (competition needs bigger-sized, higher cost DIMMs)
- **Rich storage capacity – up to 24 SAS/SATA drives**
- **Total 10 PCIe slots – PCIe Gen4 slots available to support high speed network connectivity**
  - 2x throughput capability for high performance tiers

Better Nvidia T4 GPU Hosting Capability

2U IBM IC922 [a]

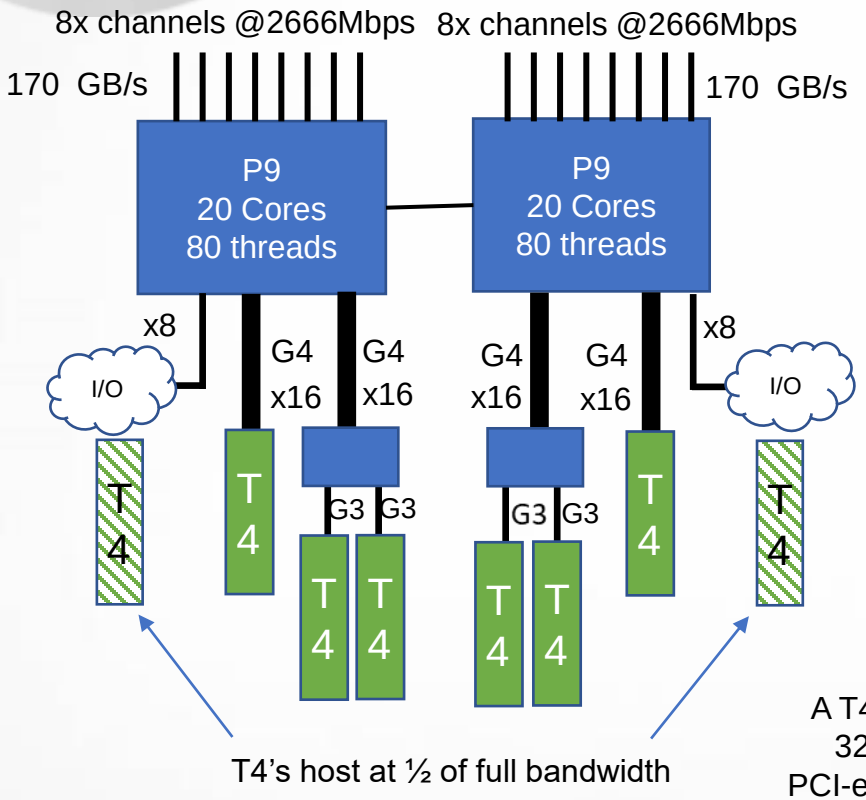

224 GB/s of T4 Bandwidth &
340 GB/s DDR Bandwidth

A T4 requires 32GB/s of PCI-e bandwidth

Reference 2U Intel Xeon Gold 6XXX [b]

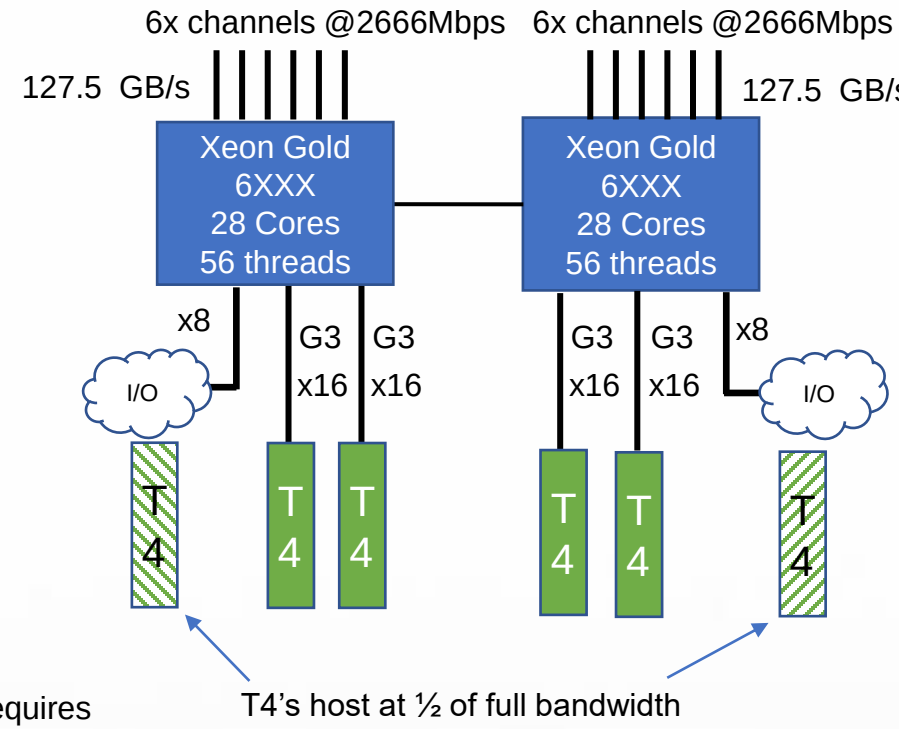

160 GB/s of T4 Bandwidth &
255 GB/s DDR Bandwidth

IC922 scales up to 8 Nvidia T4 GPUs[1] with 6 hosted at full PCI-e bandwidth, for a total of 40% greater PCIe bandwidth and >30% greater memory bandwidth than a reference Intel Xeon Gold 6xxx System

[a] - IC922 optional PCI-e Gen4 riser
[b] - Cisco c240 M5 with Intel Xeon Gold 6xxx CPU

Brief Intro to Power10 Memory Clustering Technology

Slides taken from the 2020 Hot Chips conference proceedings.

Socket Composability: SCM & DCM

POWER10

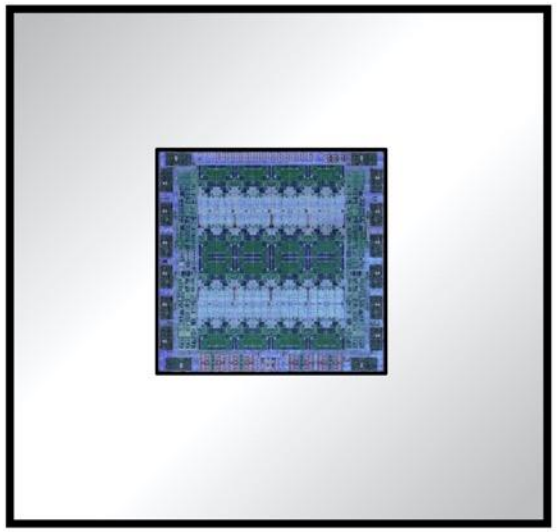

Single-Chip Module Focus:

- 602mm² 7nm (18B devices)
- Core/thread Strength
  - Up to 15 SMT8 Cores (4+ GHz)
- Capacity & Bandwidth / Compute
  - Memory: x128 @ 32 GT/s
  - SMP/Cluster/Accel: x128 @ 32 GT/s
  - I/O: x32 PCIe G5
- System Scale (Broad Range)
  - 1 to 16 sockets

Up to 16 SCM Sockets

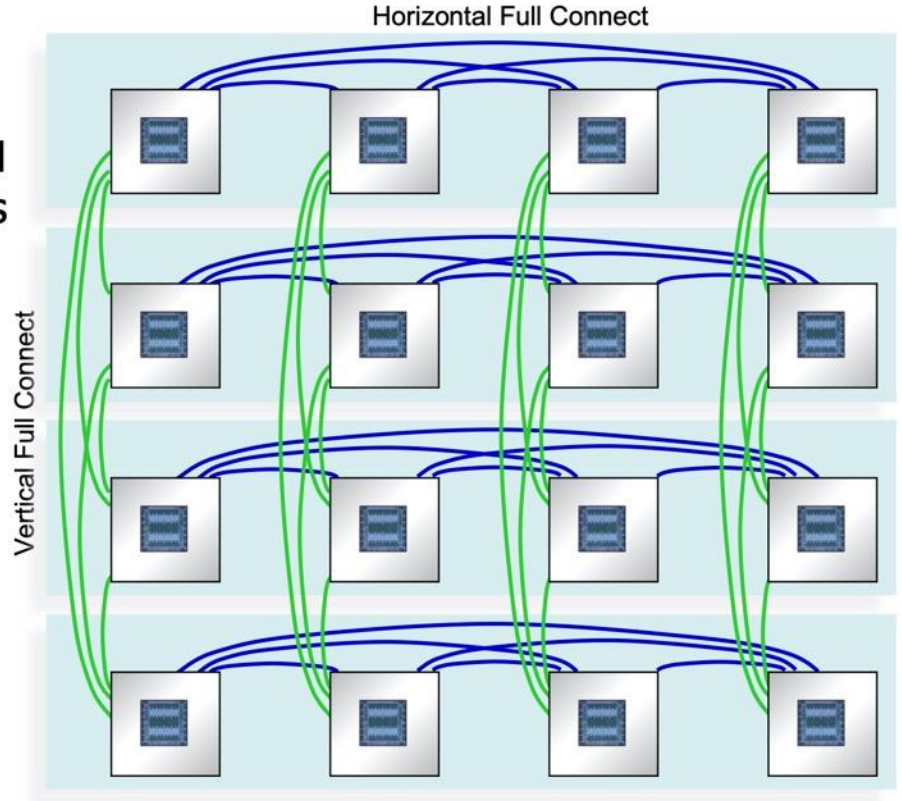

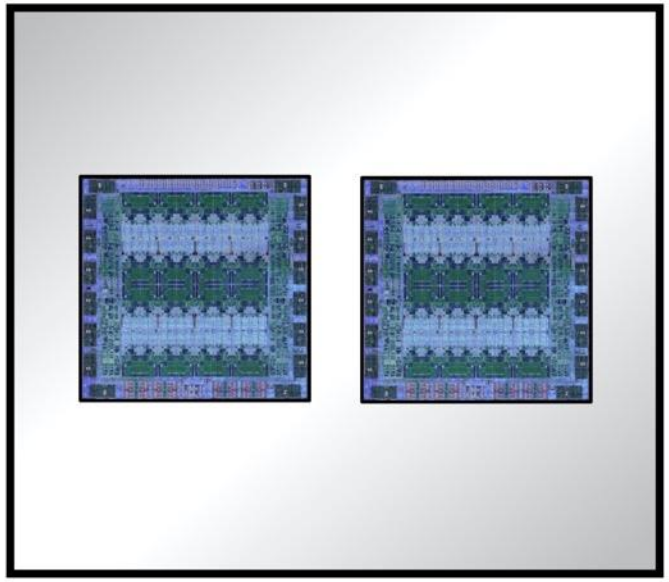

Dual-Chip Module Focus:

- 1204mm² 7nm (36B devices)
- Throughput / Socket
  - Up to 30 SMT8 Cores (3.5+ GHz)
- Compute & I/O Density
  - Memory: x128 @ 32 GT/s
  - SMP/Cluster/Accel: x192 @ 32 GT/s
  - I/O: x64 PCIe G5
  - 1 to 4 sockets

Up to 4 DCM Sockets

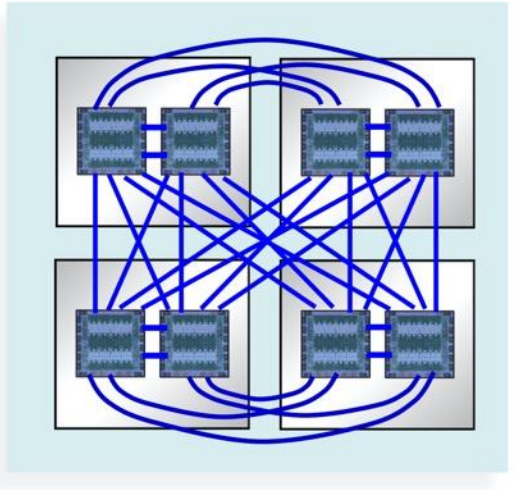

(Multi-socket configurations show processor capability only, and do not imply system product offerings)

Memory Clustering: Distributed Memory Disaggregation and Sharing

POWER10

Use case: Share load/store memory amongst directly connected neighbors within Pod
Unlike other schemes, memory can be used:
- As low latency local memory
- As NUMA latency remote memory

Example: Pod = 8 systems each with 8TB
Workload A Rqmt: 4 TB low latency
Workload B Rqmt: 24 TB relaxed latency
Workload C Rqmt: 8 TB low latency plus 16TB relaxed latency

All Rqmts met by configuration shown

POWER10 2 Petabyte memory size enables much larger configurations

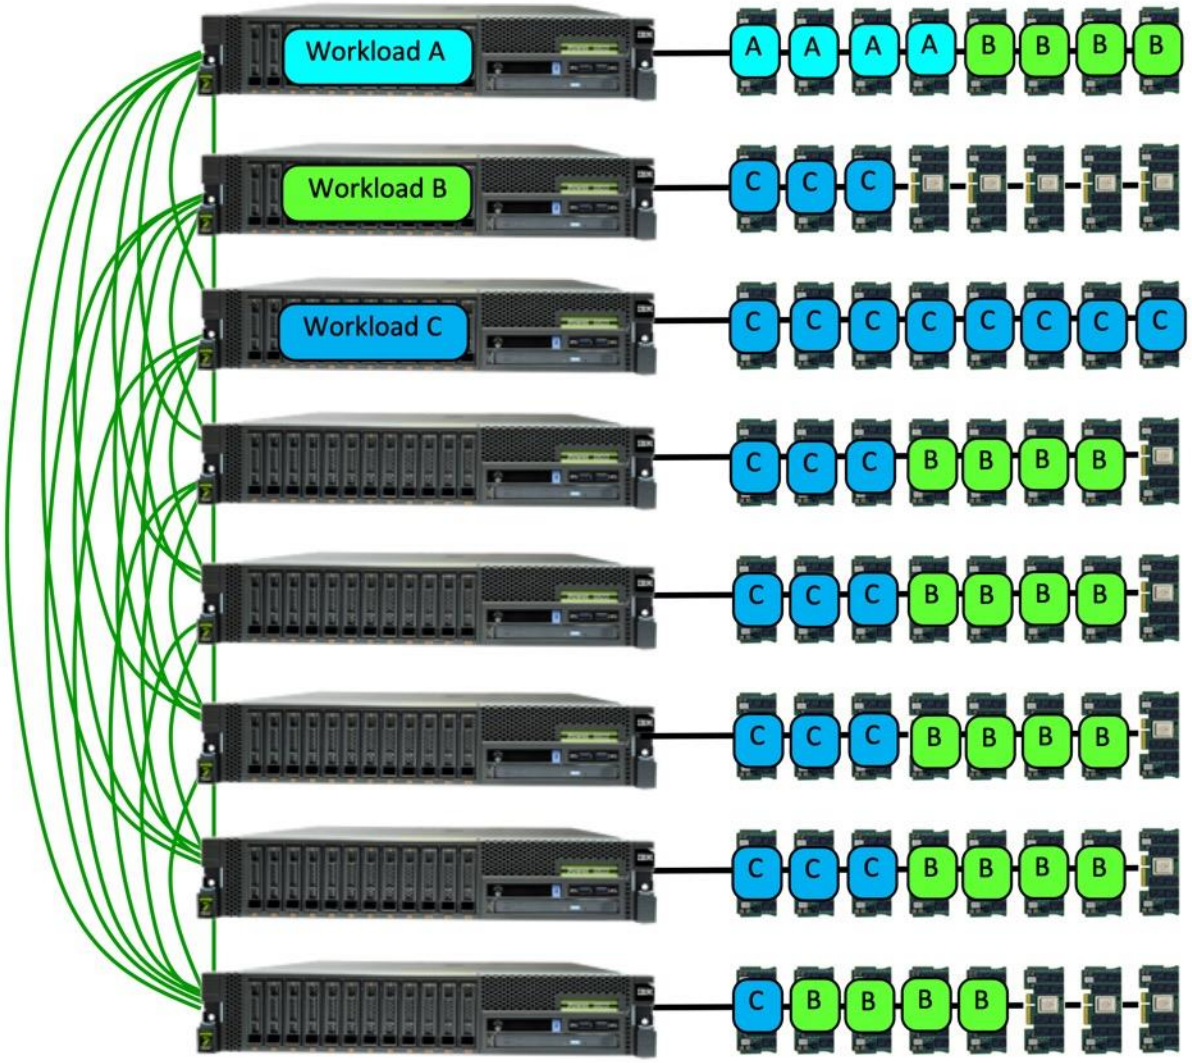

(Memory cluster configurations show processor capability only, and do not imply system product offerings)

Memory Clustering: Enterprise-Scale Memory Sharing

POWER10

Pod of Large Enterprise Systems
Distributed Sharing at Petabyte Scale

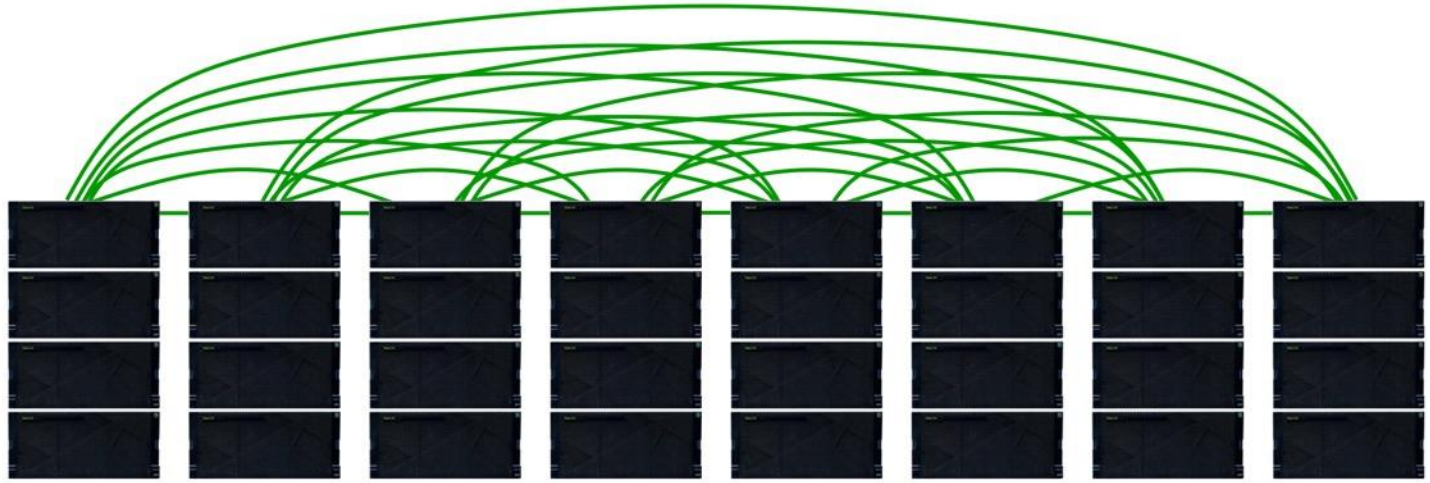

Or Hub-and-spoke with memory server and memory-less compute nodes

6400 THREADS 800 CORES, SHARED MEMORY

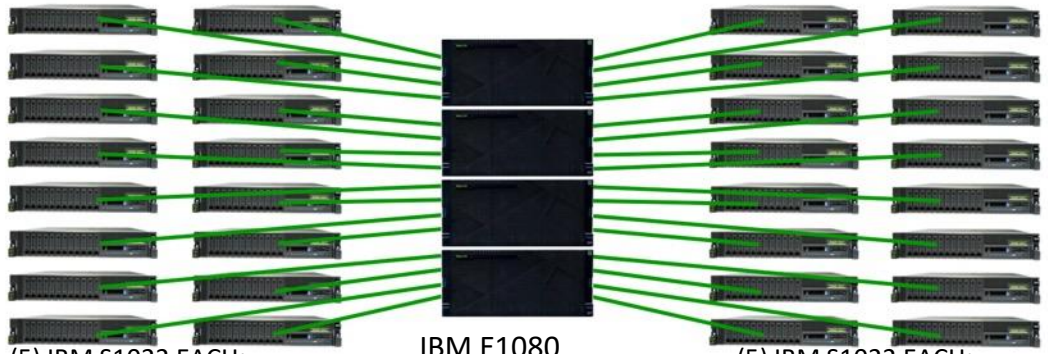

(5) IBM S1022 EACH:
2 Proc.,56 CORE, 4TB MEM
448 THREADS each

IBM E1080
32TB 16P/240 CORE,
1920 THREADS

(5) IBM S1022 EACH:
2 Proc.,56 CORE, 4TB MEM
448 THREADS each

(Memory cluster configurations show processor capability only, and do not imply system product offerings)

Memory Clustering: Pod-level Clustering

POWER10

Use case: Low latency, high bandwidth messaging scaling to 1000's of nodes

Leverage 2 Petabyte addressability to create memory window into each destination for messaging mailboxes

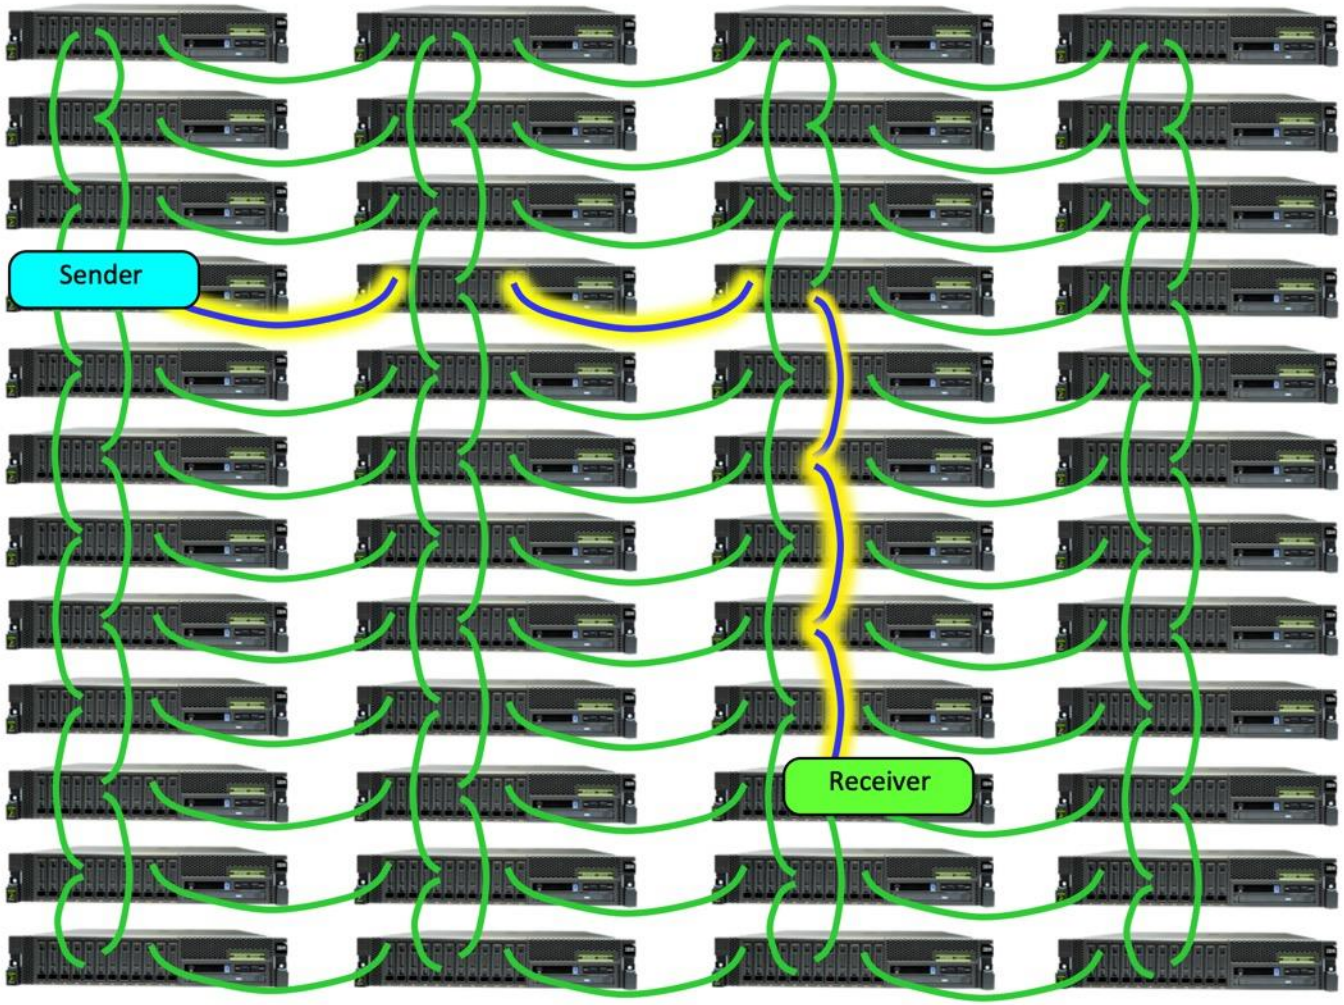

(Memory cluster configurations show processor capability only, and do not imply system product offerings)

Storage Solutions in the Life Sciences Platform

Storage Technologies
"The Fuel"

Storage Landscape for AI/Research Data

Power Systems

Welcome to the waitless world

IBM

Data Ocean (Ingest)

Instrument + Data

Spectrum Protect

Data Backup

ESS Backup Target

IBM Cloud Object Storage

Hybrid Cloud

Data Lake

Spectrum Scale File System For HPC/Ai Single File System Across Tiers

Elastic Storage System

S3 Object Vault

NAS Gateway

SAS Disc 36GB/Sec (GL6S)

Policy Based Tiering

Tier 1 File High Speed

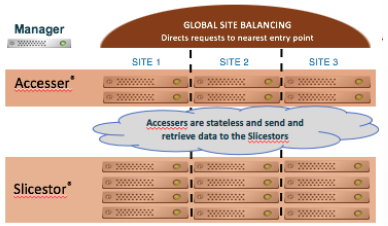

On-Premise 3 Site Protected Object Store

Policy Based Archive

Tier 3 Object (On-Premise/Cloud)

AI Processing

Nvidia DGX Systems

IBM Power9 Nvidia Volta Watson MI

Any Intel GPU OpenSource AI

Spectrum Discover MetaData Catalog

PreFetch data to be Processed

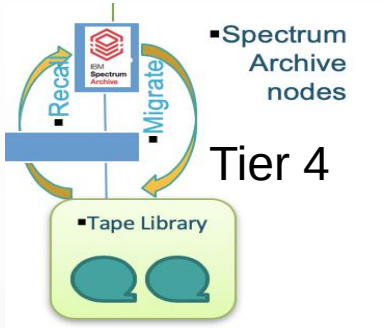

▪Spectrum Archive nodes

▪Tape Library

Tier 4

Tier0 ESS3000

Spectrum AI Flash Compute Cache NVME Flash. Tier0

40/80/120GB/sec

High Performance Tier

High Speed 100Gb Eth/IB

Where is the data?

IBM Spectrum Discover

- Metadata is the structured data about the unstructured object
  - Who, what, when, where, and why of account, container, object, stream, dir, file
  - Perfect for indexing and searching
- Metadata may be separate from the data, stored with the data, or derived from the data
  - Posix inode plus extended attributes
  - Standard document headers (doc, ppt, mp3, dicom, pdf, jpeg, GeoTIFF)
  - Custom metadata tags
  - AI derived metadata

System Metadata

Location
Size
Owner
Group
Permissions
Last-Modified
...

Biomedical

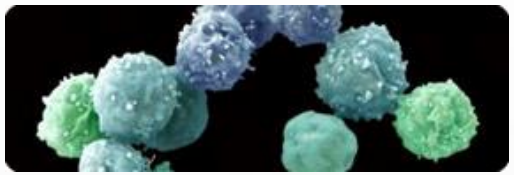

Age, Biomarkers, Developmental Stage, Cell Surface, Markers, Cell Type/Cell Line, Disease State, Extract Molecule, Genetic Characteristics, Immunoprecipitation, antibody, Organism,

Image

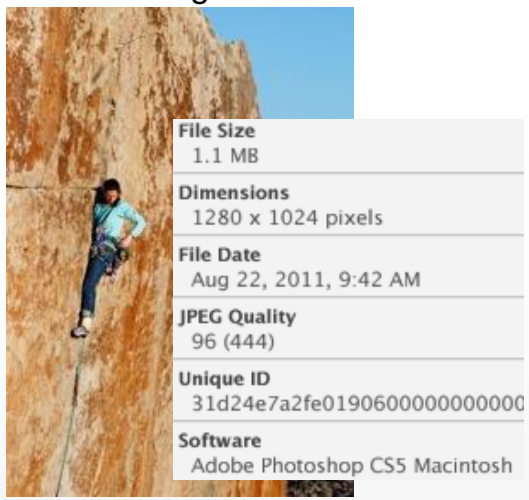

Natural Language Processing

PYTORCH

TensorFlow

Spectrum Discover - Metadata Management Software

IBM Spectrum Scale

IBM
Spectrum
Discover

Provides unified metadata management and insights for heterogeneous file and object storage, on-premises and in the cloud.

Discover	Classify	Label	Find
Automatically ingest & index system metadata from multiple file & object storage systems on-prem & in the cloud	Automatically identify and classify data, including sensitive and personally identifiable information	Enrich data with system & custom metadata tags that increase the value of that data	Find data quickly and easily by searching catalogs of system & custom metadata

IBM Spectrum Discover Capabilities

IBM Spectrum Scale

File & Object Storage

IBM Spectrum Scale

IBM Cloud Object Storage

ISILON

amazon web services S3

NetApp

ceph

Data Insight

IBM Spectrum Discover

Search | Reporting | Dashboard

- Simple to deploy (VMware virtual appliance)
- Metadata curation
- Custom metadata tagging
- Automatic indexing
- Policy-Engine
- Action Agent API

Scanning & Event Notifications

Activation & Optimization

Large-Scale Analytics
- Data discovery
- Dataset identification
- Data pipeline progression

Data Governance
- Data inspection
- Data classification
- Data clean-up

Data Optimization
- Archive / tiering
- Duplicate data removal
- Trivial data removal

Spectrum Discover Dashboard

Monitor storage utilization and data recommendations (Move/Archive)

Preview capacity use by data facet
- Classification
- Owner
- File Type
- Etc.

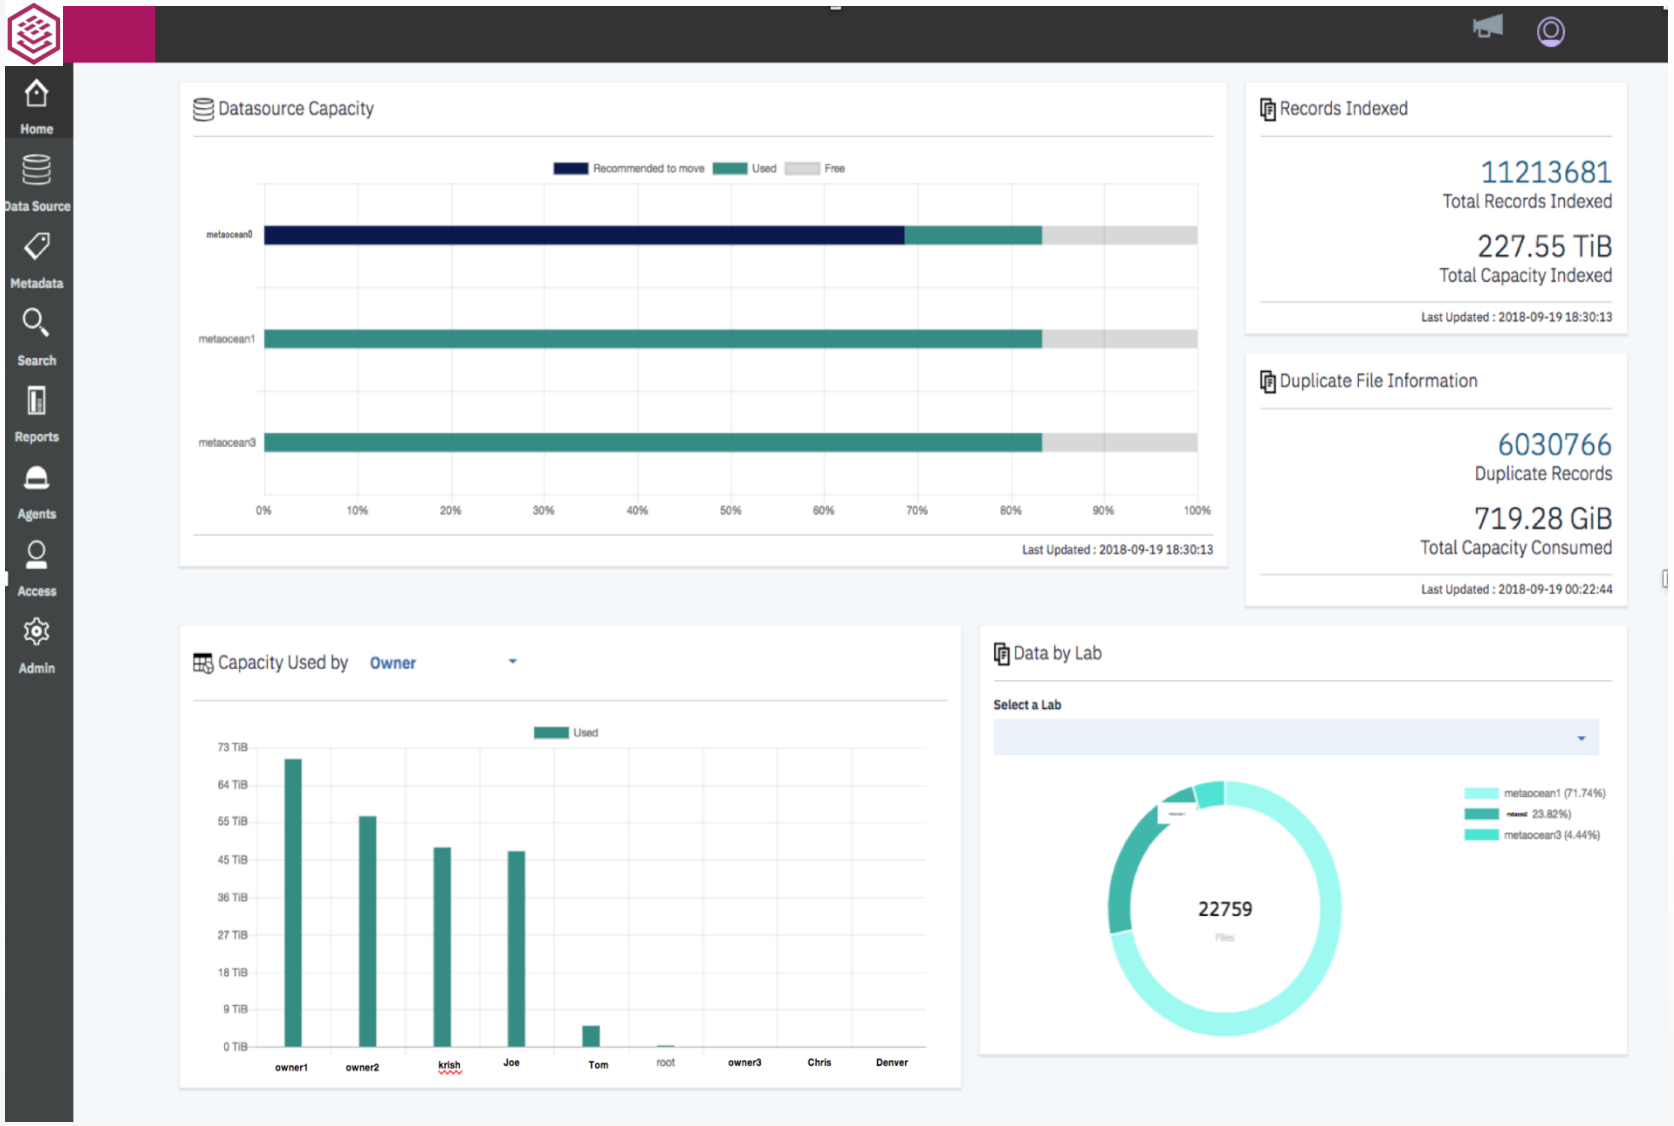

Total indexed data and capacity

Duplicate file or object candidates
- Number
- Capacity used

Data capacity by group/collection
- Customer defined
- Lab/Project/etc.

Data lake with IBM Spectrum Scale

Power Systems

Welcome to the waitless world

IBM

4000+ clients

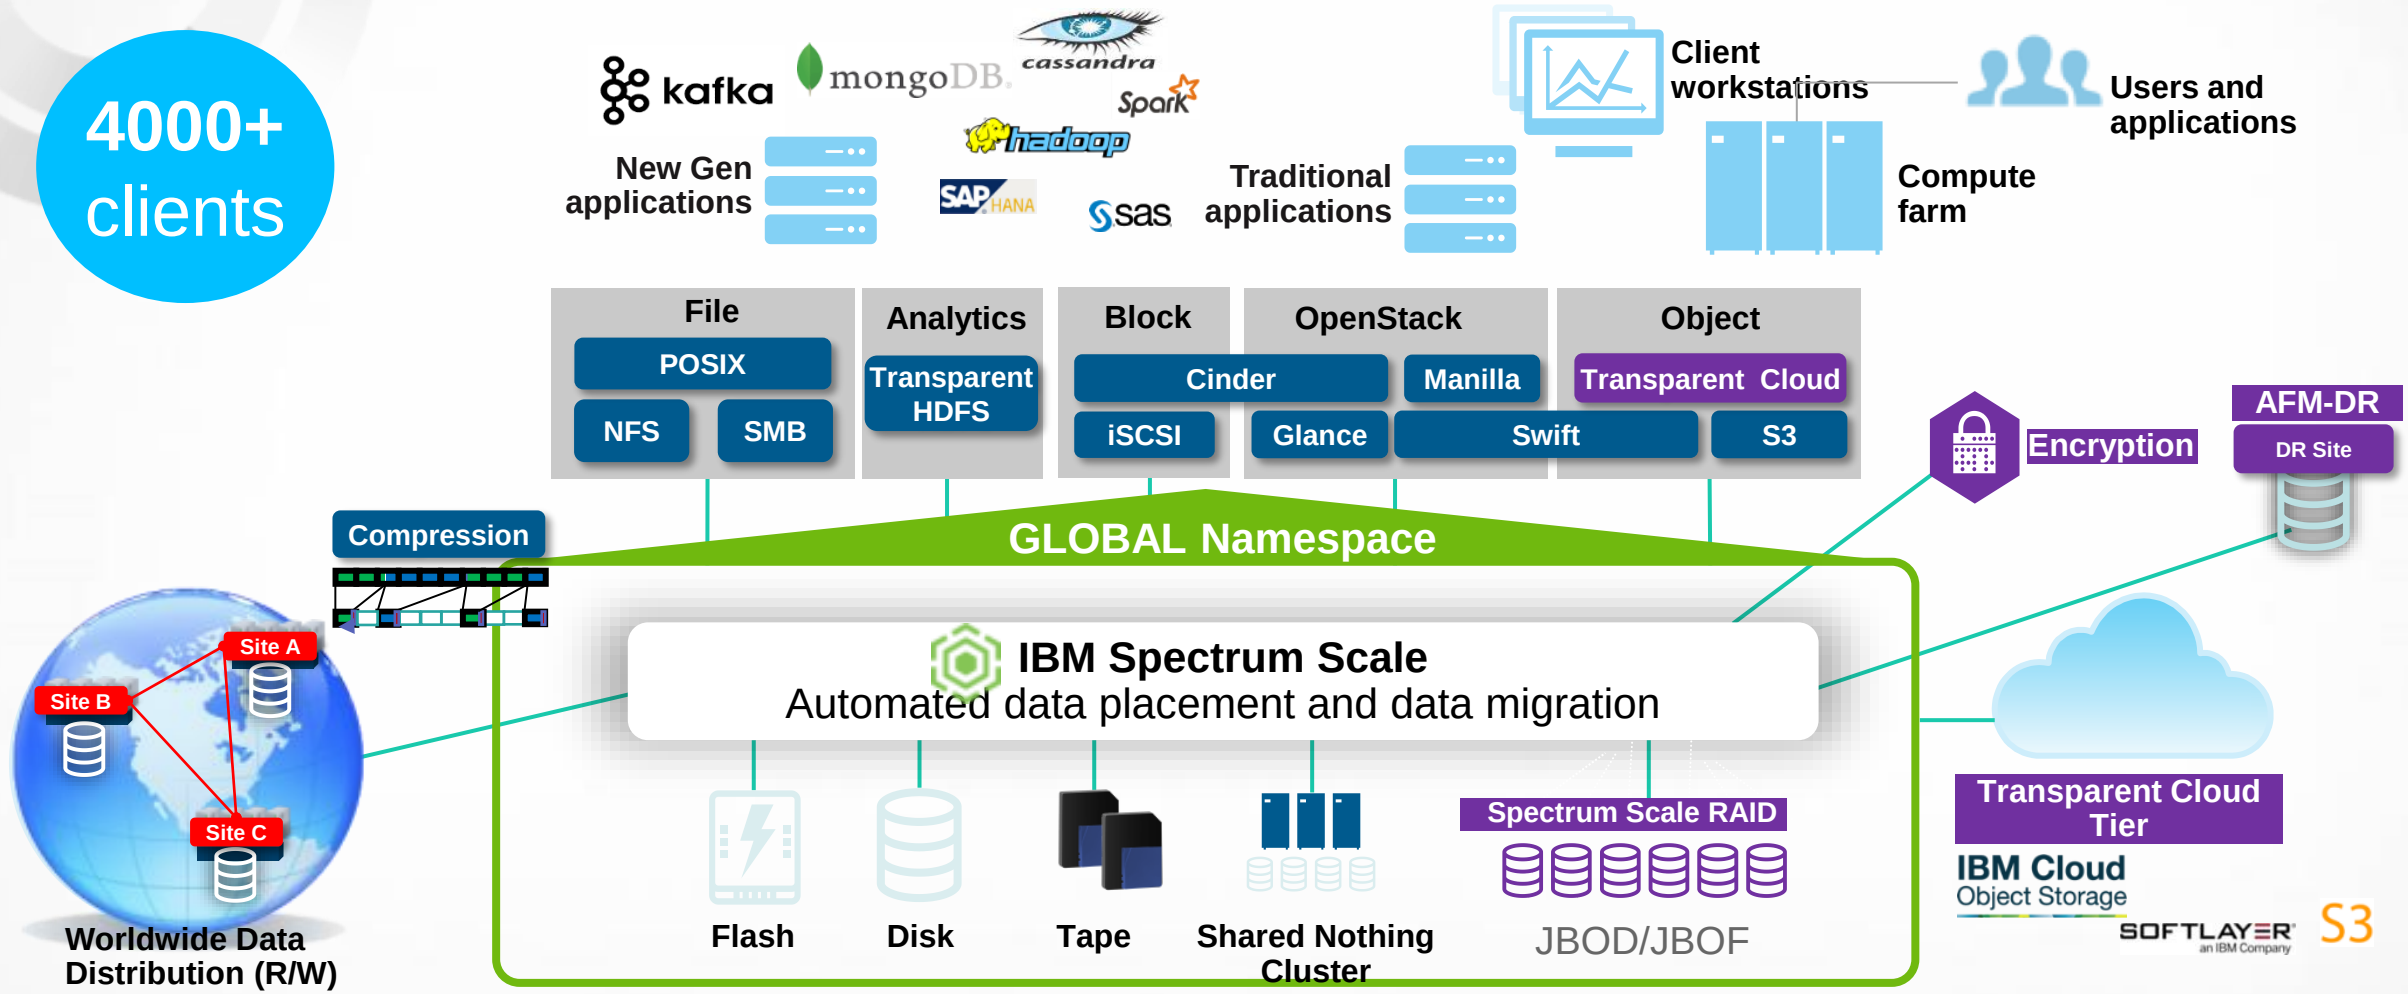

Consolidate all your unstructured data storage on Spectrum Scale with unlimited and painless scaling of capacity and performance

Unleash new storage economies on a global scale.

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 scale-out
- Easy to deploy clusters with unified system GUI
- Simplified storage administration with IBM Spectrum Control integration

Entry
ESS 3000

ESS 3000
cluster

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 using protocol nodes - NFS 4.0, SMB, Object
- Ideal for Big Data analytics including 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 9s (99.999%) of availability

Elastic Storage
Server cluster

ESS Ecosystem

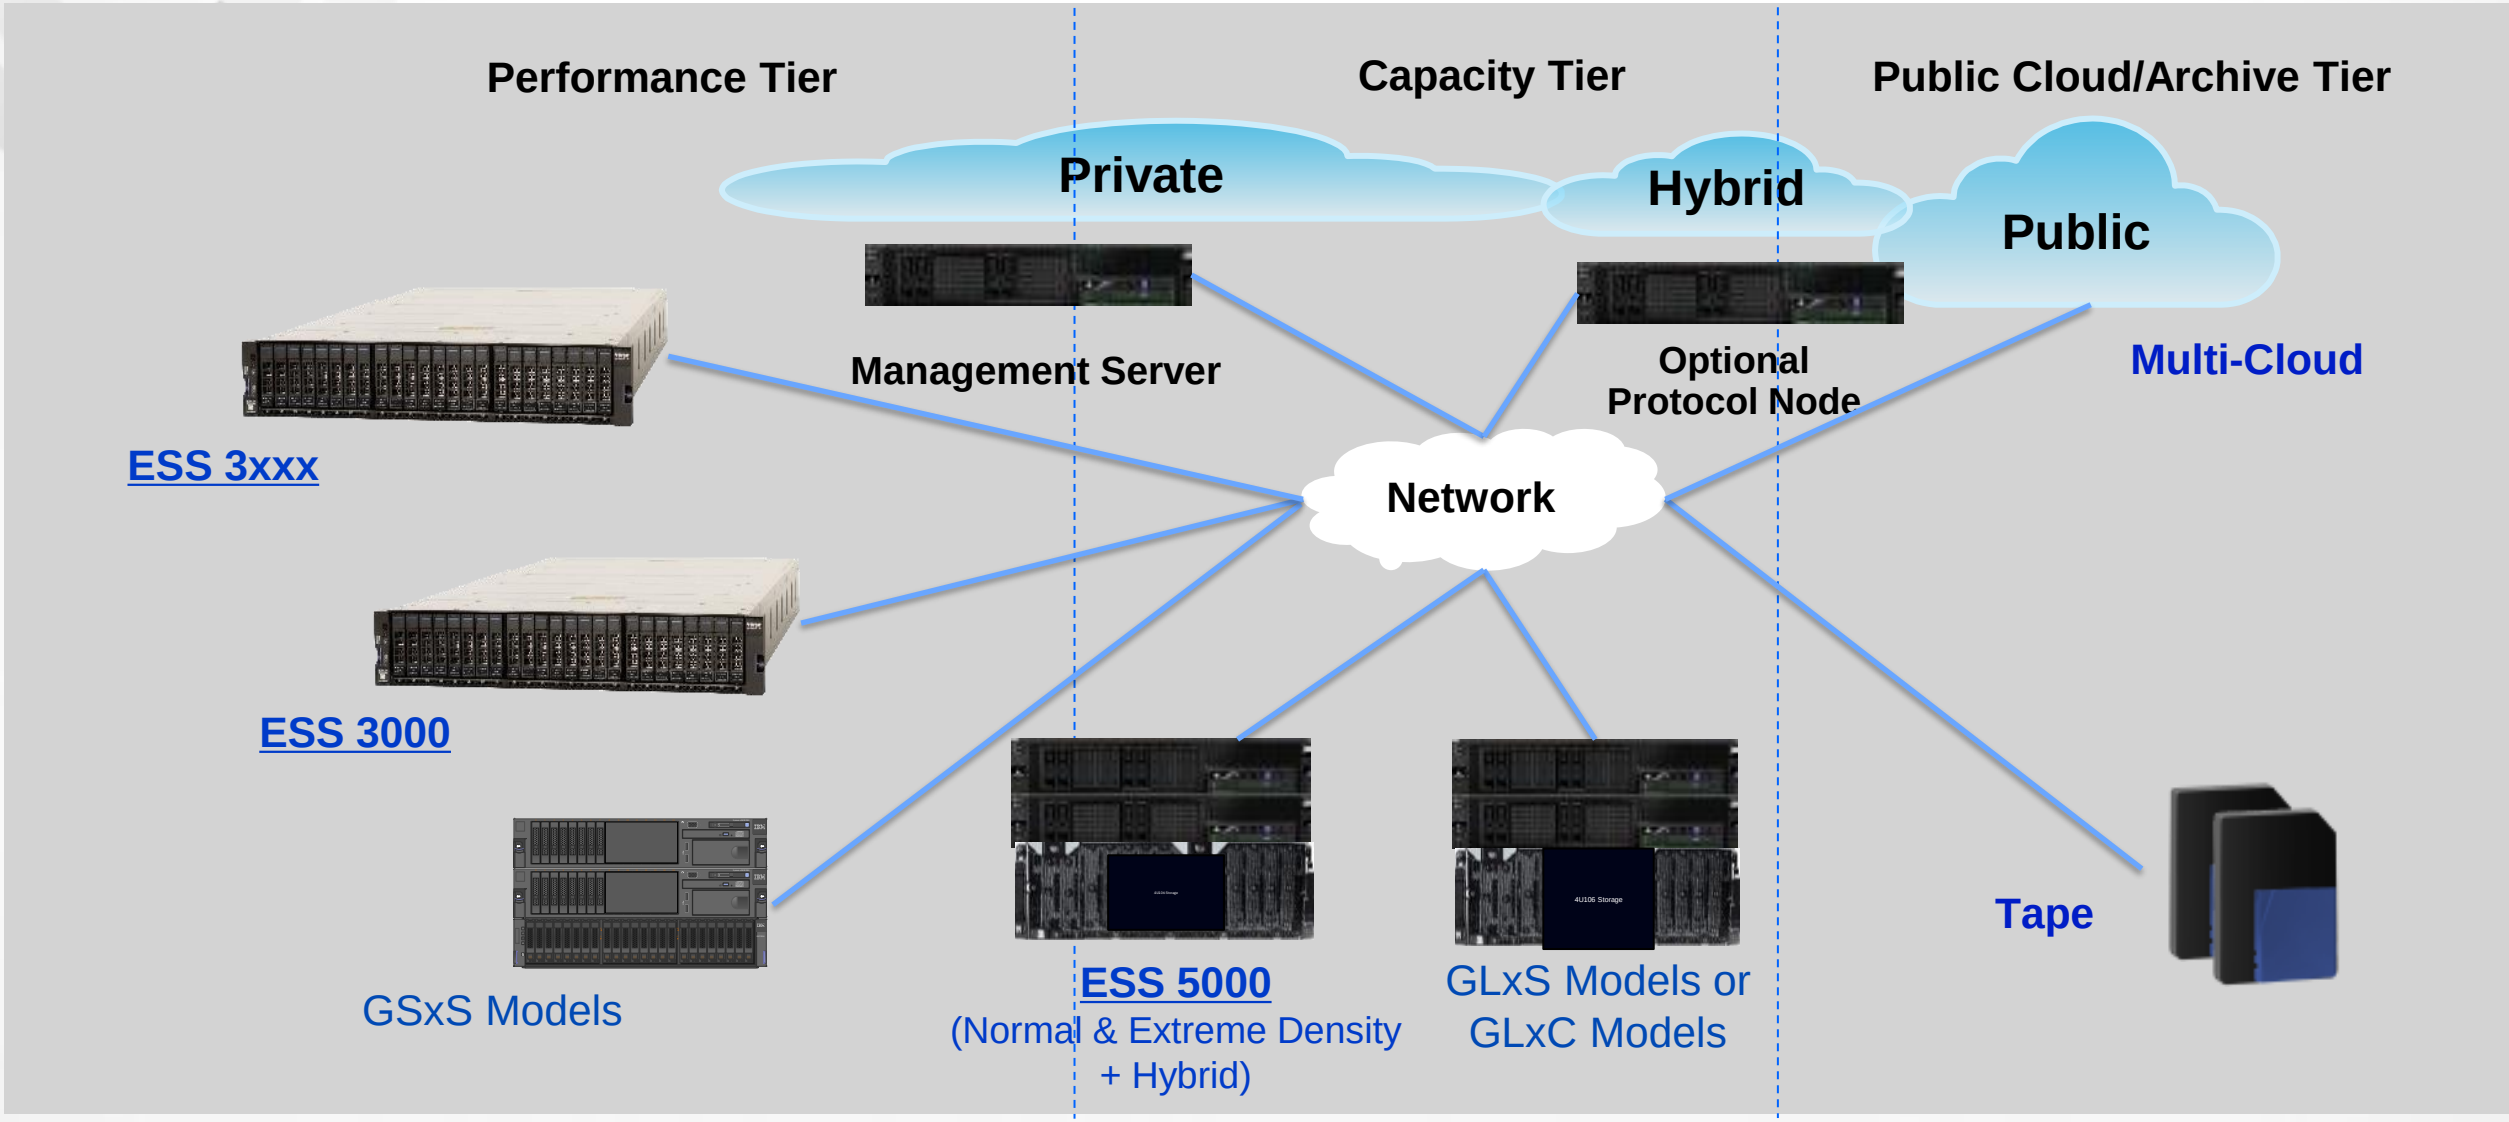

1Q21 ESS Release – New SL7 Model of ESS 5000

Announce: Feb 23, 2021
GA: Mar 12, 2021

ESS 5000
New Supported Models:

New SL7

- Up to 8.8 PB Raw Capacity in a single 42U rack*
- Up to 66 GB/s in a single 42U rack**

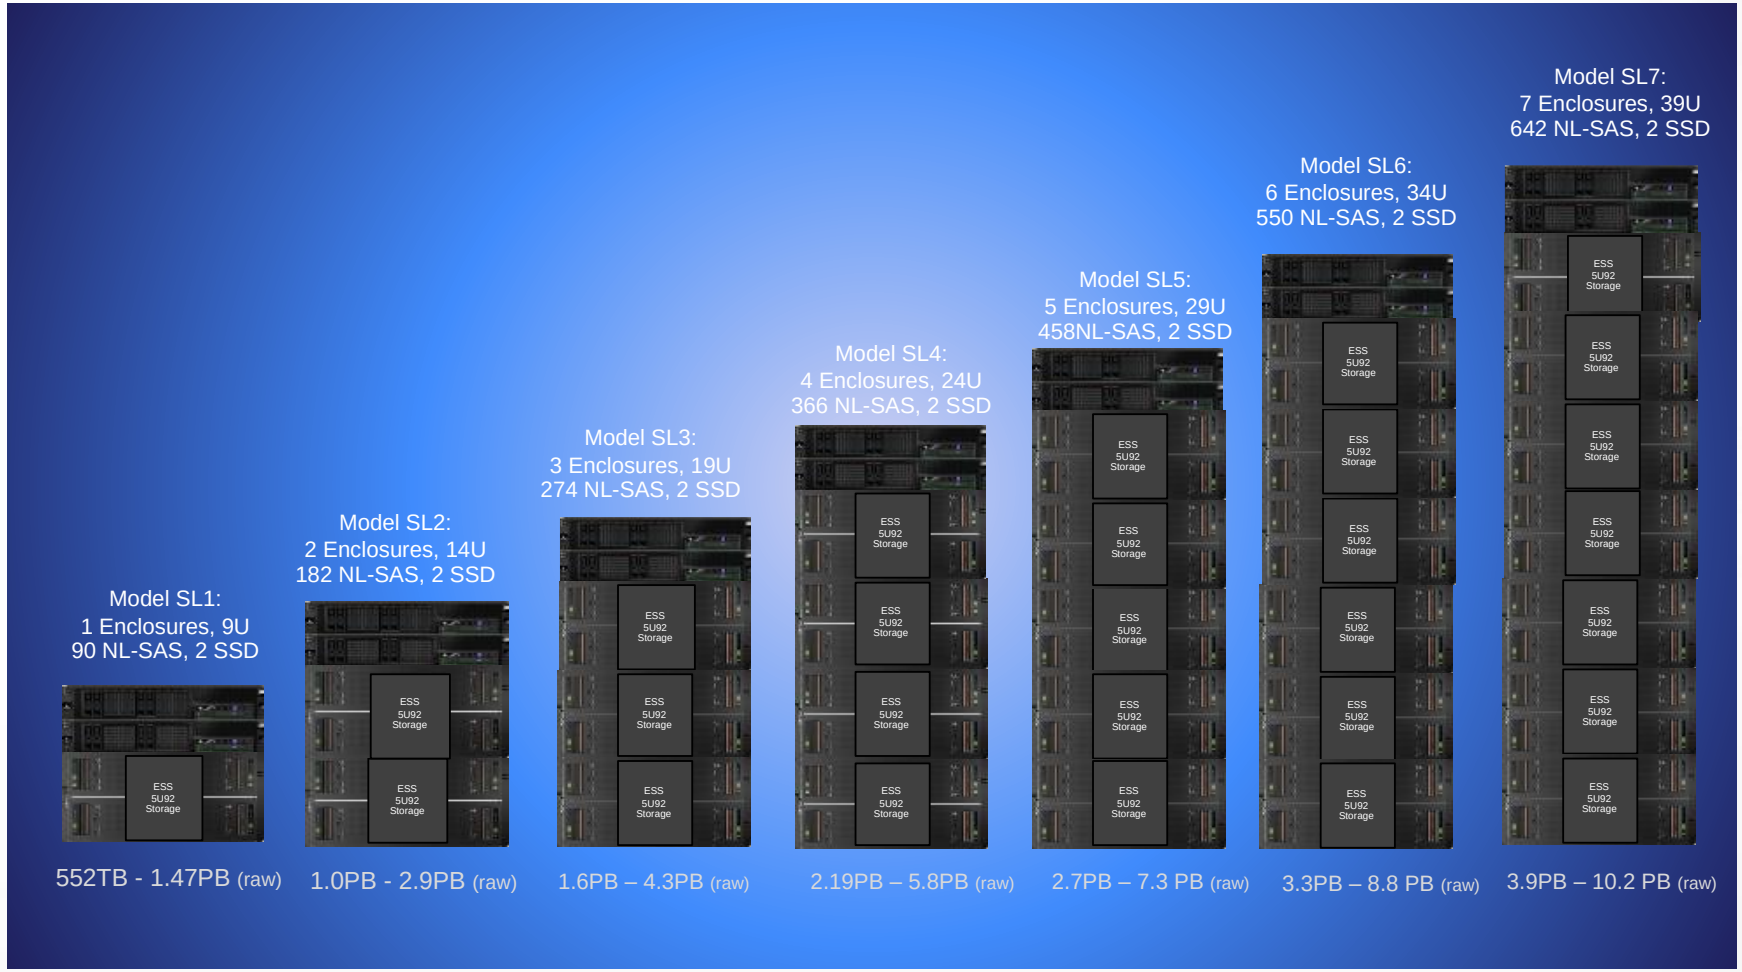

** SL7 **with a standard rack***

*** 16M file, 8+2p Read **with 2 - SL3 models***

1Q21 ESS Release – New SC9 Model of ESS 5000

Announce: Feb 23, 2021
GA: Mar 12, 2021

ESS 5000
New Supported Models:

New SC9

- Up to 15.2 PB Raw Capacity in a single 42U rack*
- Up to 100 GB/s in a single 42U rack**

** SL9 **with a rack extension***

*** 16M file, 8+2p Read **with 2 - SC4 models***

IBM Cloud Object Storage information dispersal

Redefining availability and economics of data storage

Traditional storage

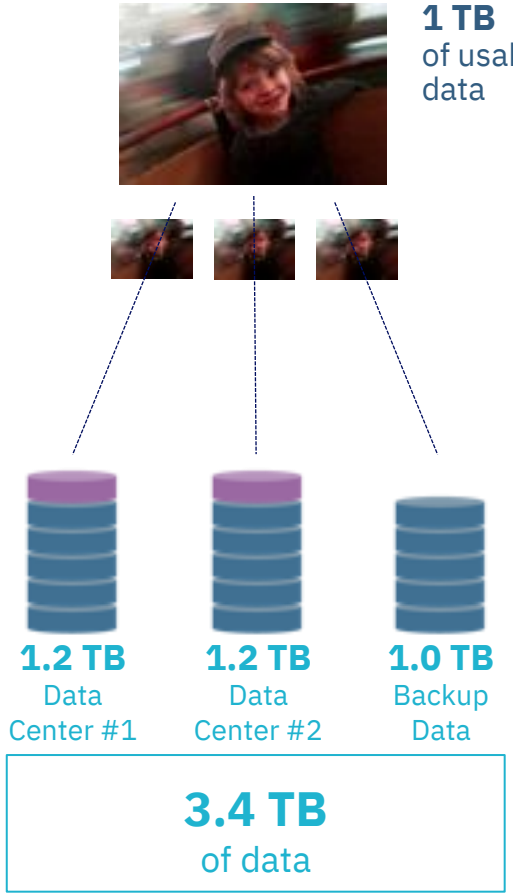

Traditional storage requires 3.4 TBs raw storage capacity for 1 TB of usable storage.

IBM Cloud Object Storage requires less than half the storage and 70% lower TCO*.

What does that mean to IT?

You can lose a disk, a server or even a whole site due to failure or disaster, and still quickly recover 100% of your data.

Slices are distributed geographically for durability and availability.

IBM Cloud Object Storage

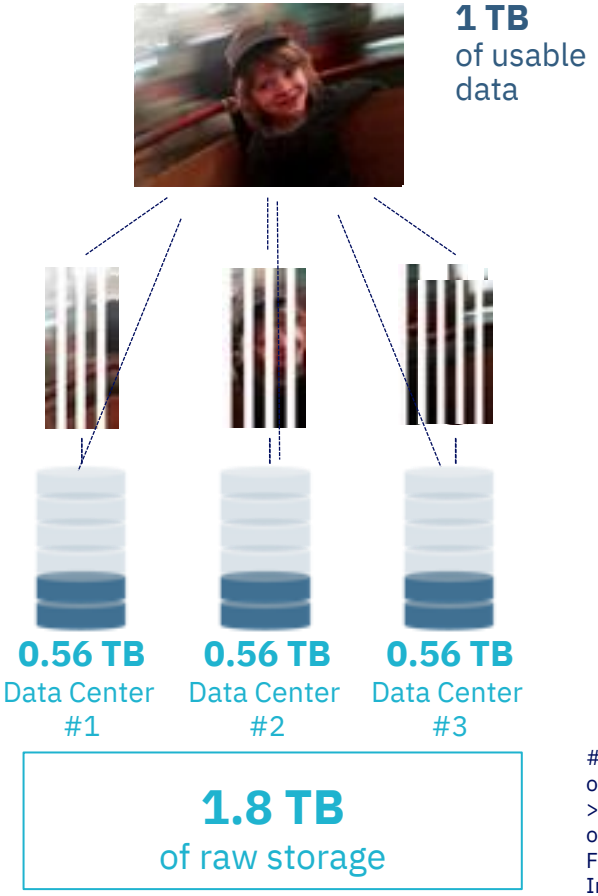

Our object storage requires only 1.8 TBs raw storage capacity for 1 TB of usable storage.

#1 Cloud Object Storage on-premises solution delivers >70% lower TCO than traditional on-premise storage. Source: Forrester Total Economic Impact Study

IBM Cloud Object Storage

New economics: Lower TCO and site-level redundancy

Traditional storage

- More storage
- More power
- More floor space
- More software
- More personnel

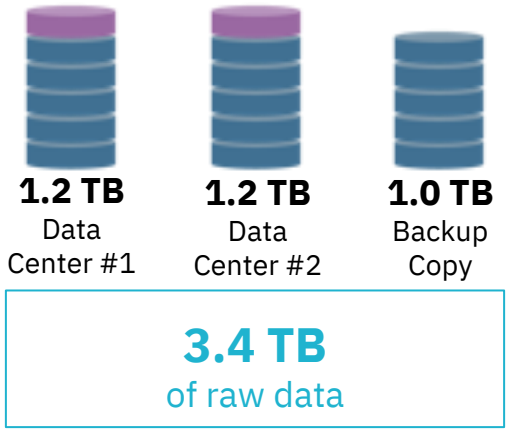

Traditional storage requires 3.4 TBs raw storage capacity for 1 TB of usable storage.

IBM Cloud Object Storage

- Less storage
- Less power
- Less floor space
- Less software
- No downtime
- Less personnel
- Less costly—up to 70% lower TCO

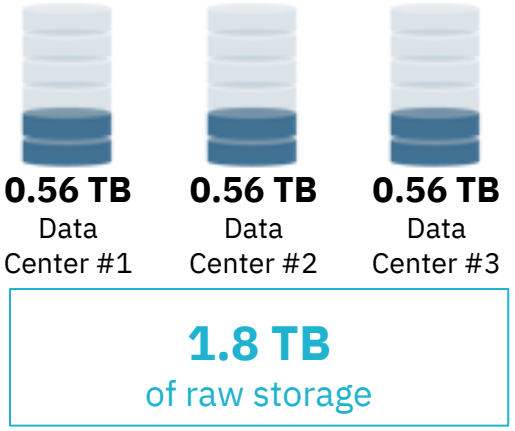

Our object storage requires only 1.8 TBs raw storage capacity for 1 TB of usable storage.

IBM Cloud Object Storage System

Components that are simple to manage and flexible to deploy

Manager

- Fault management
- Reporting
- Provisioning
- Performance monitoring
- Storage configuration
- Single pane of glass

IBM Accesser®

- S3 interface
- Encrypts data
- Slices data
- Disperses data
- Retrieves data
- Stateless

IBM Slicestor®

- Storage for slices
- Single site or multi-site
- Capacity-based pricing
- Data integrity management

Supported as an IBM appliance

Supported on IBM certified industry standard platforms – Software only

Software-defined storage with flexible hardware deployment options

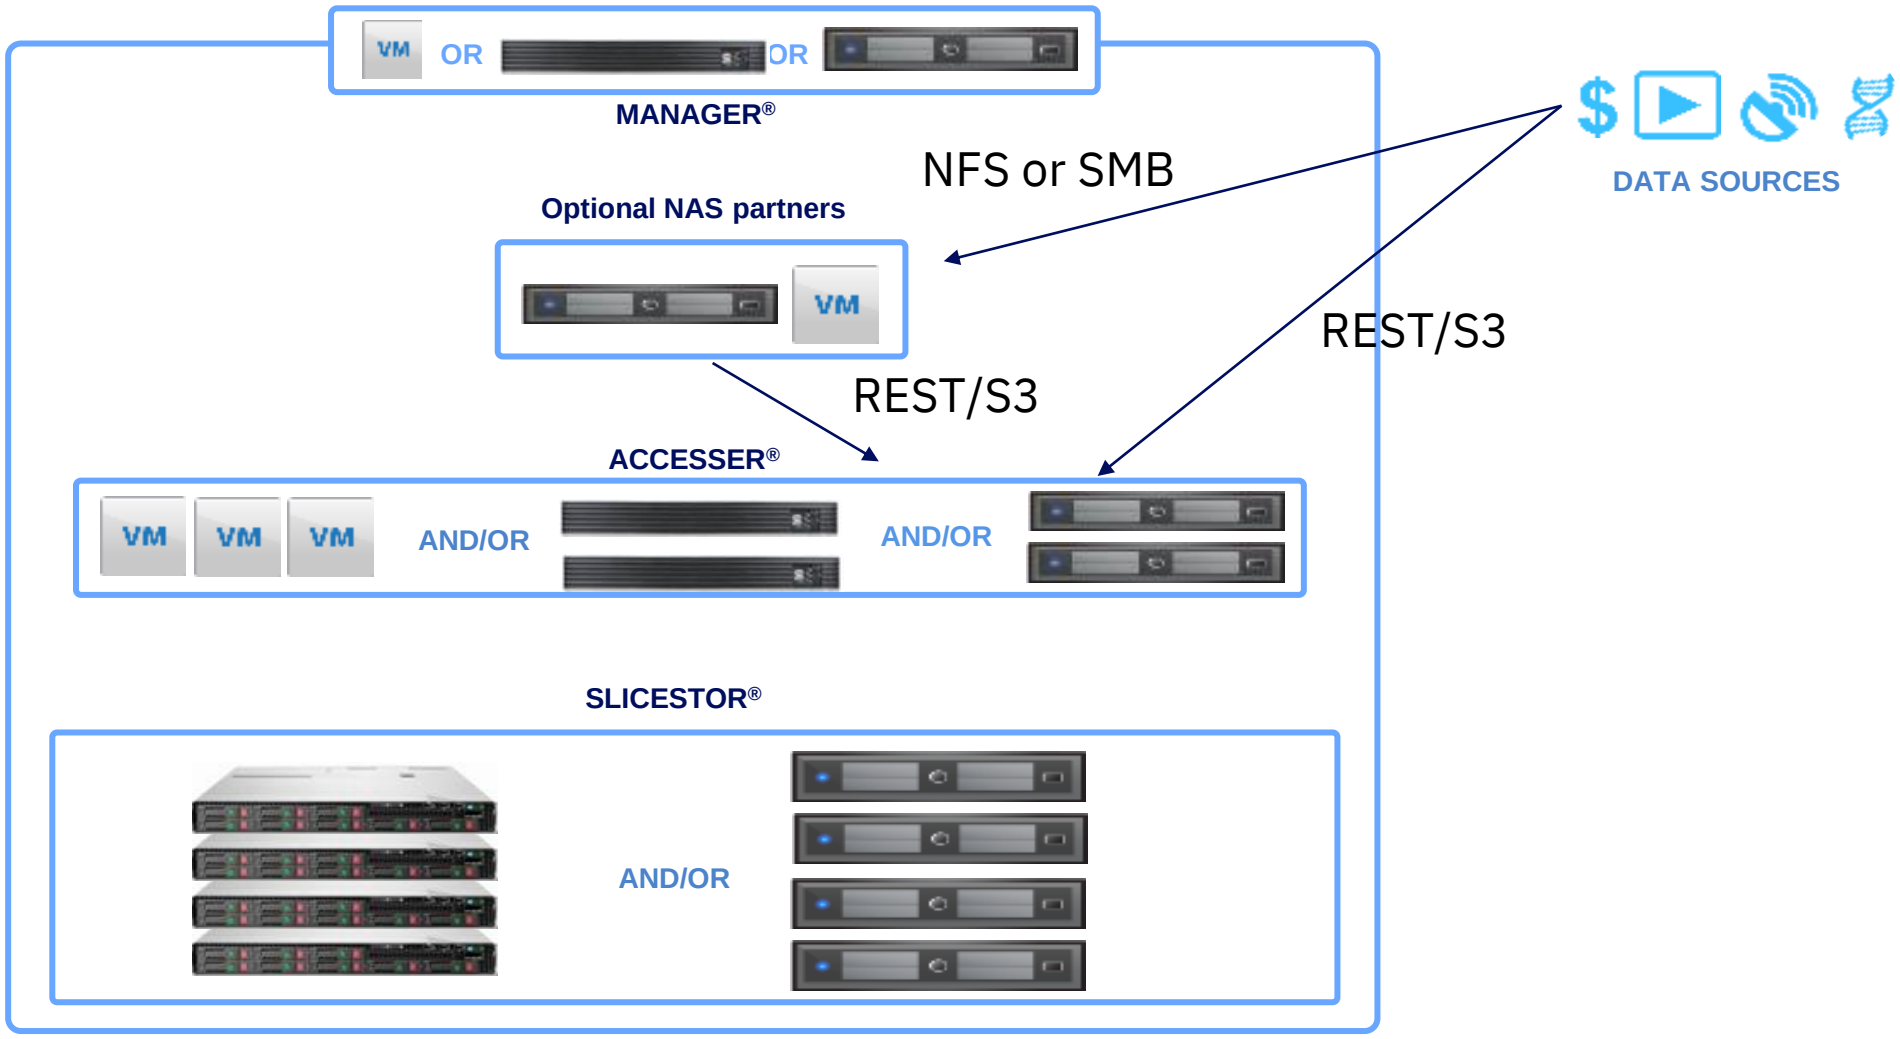

IBM appliance

Software only
using certified server

Software only
on VM or Docker container

Flexible deployment options across one, two or more sites

3,456TB of RAW Licensed Capacity

Single site most cost effective
2,304TB Usable Capacity

Site A

Two site traditional mirroring
1,152TB Usable Capacity

Site A *Site B*

Geo-dispersed multiple sites and most deployed by customers
1,920TB Usable Capacity

Site A *Site B* *Site C*

Any one or more sites can be located in IBM Cloud

The Bottom Line

- **IBM is in the Cognitive Systems business.**
- Cognitive systems demand the widest processing power available.
- The Power Systems architecture is the BEST for AI Workloads
  - Combined with 4,8 Multi-Threading and even Hadoop/Spark processes faster.
- The Elastic Storage System is unmatched in function, scale, and performance.
- IBM Hybrid Cloud software defined infrastructure allows you to grow on demand at the most economical rates and flexibility.

.

Give us your hardest workload!

How can IBM make you and the university successful?

Thank you….