intel | Microsoft

SCAsia 2019

Supercomputing

Gathering the **Best of HPC** in Asia

March 11–14, 2019
Singapore

Azure HPC & AI: State of the Art

Powering everything from HPC to Quantum Simulation to Autonomous Driving

Rob Futrick
Principal Program Manager
Azure Compute

intel | Microsoft

~~Magic!~~

Demos!

Customers + Workloads

Autonomous Driving

SILICON Therapeutics

Genomics

Computational Chemistry

intel | Microsoft

SILICON Therapeutics

Novel drug discovery in the cloud

Unique quantum physics simulation technology to identify targets and design drugs to fight diseases that have been considered difficult for traditional approaches

- Molecular dynamics simulations on thousands of targets
- Five years of GPU compute-time run in 20 hours on 2048 NCv1 GPU VMs
- 80TB BeeGFS parallel filesystem
- 50+ TB output data

intel | Microsoft

→ Article

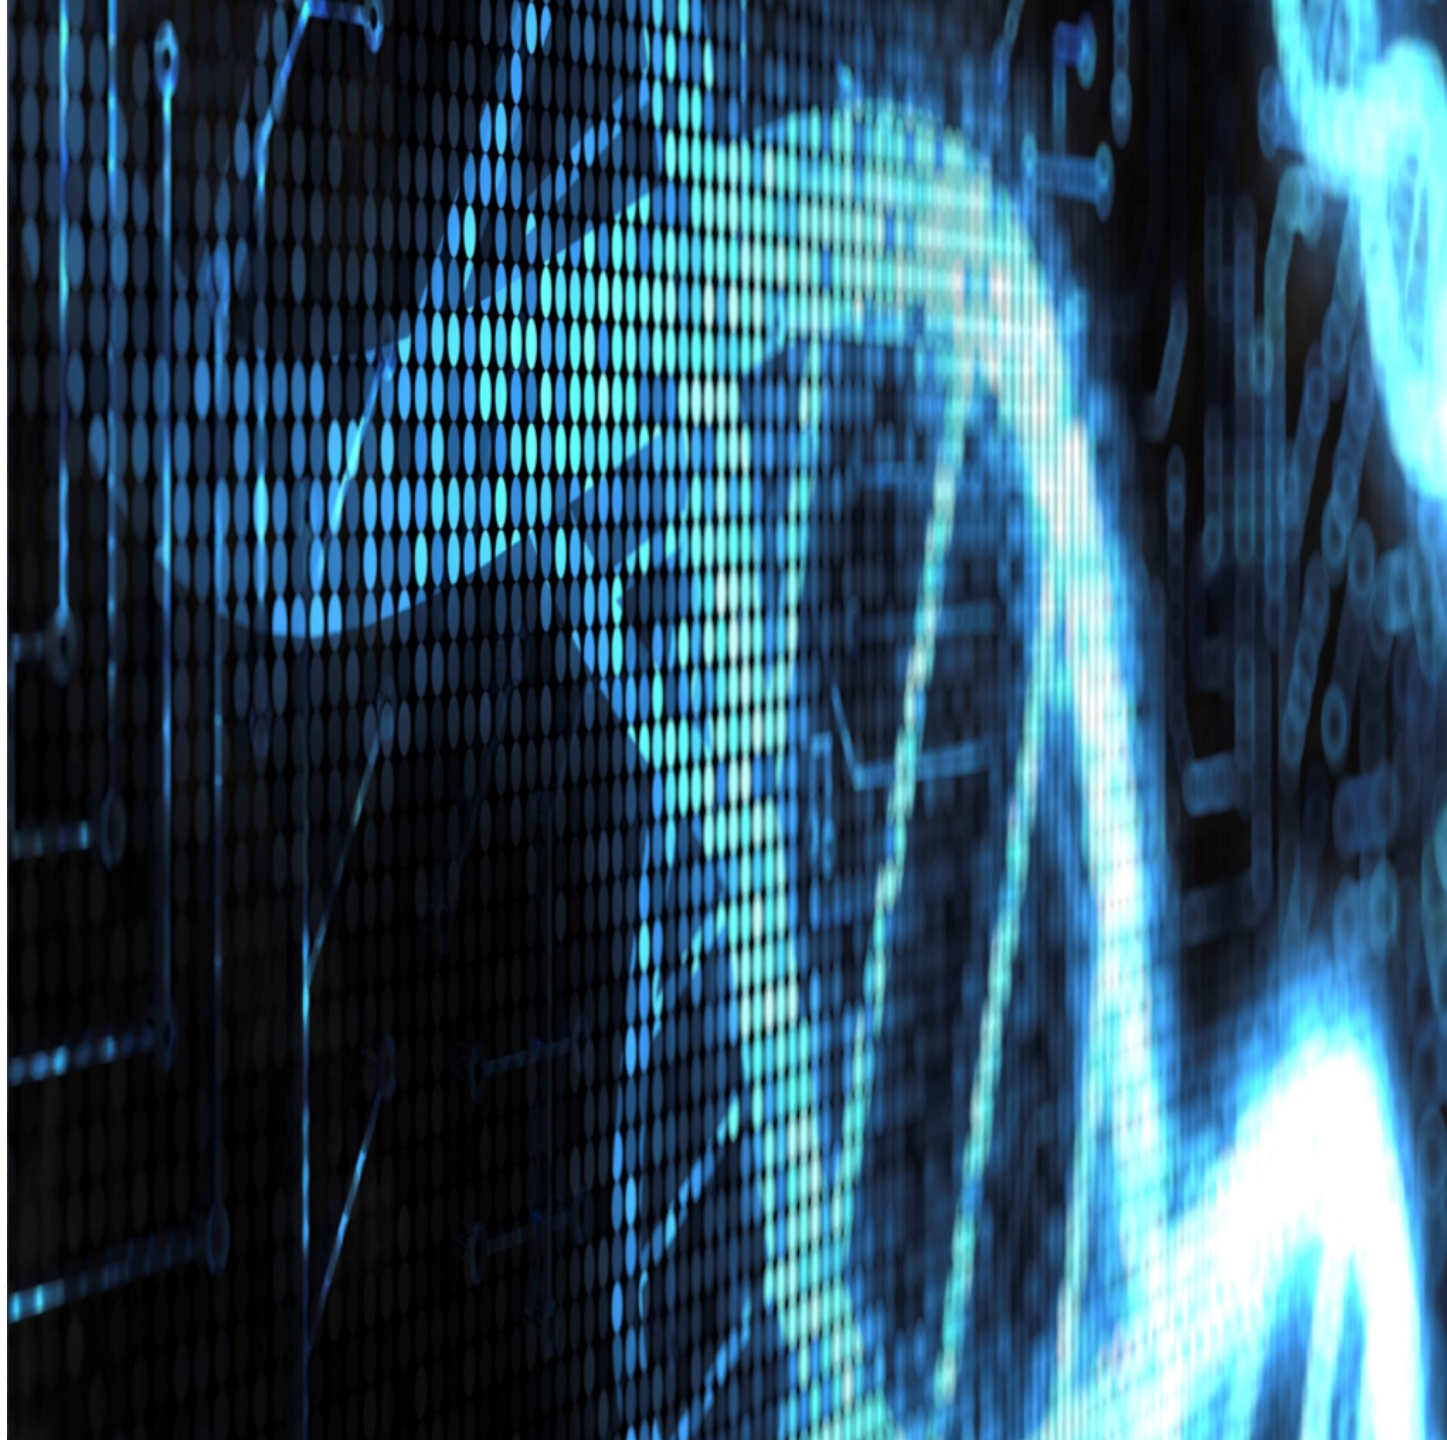

Roche

Neusomatic: First CNN approach to somatic mutation detection

Customer Roche Bioinformatics Researchers have developed a novel "Deep Learning" (Convolutional Neural Network) approach to detecting mutations for Cancer.

Challenge To perform (& publish) a study which tests Neusomatic across a range of real and synthetic data sets – in particular, against 261 whole-exome sequenced cancer samples (from TCGA) – and to demonstrate scalability & cost effectiveness on the public cloud (Microsoft Azure).

Solution Azure CycleCloud was used to automatically create & destroy Linux clusters using ~100 Azure pre-emptible compute virtual machines, upon which Neusomatic was used to analyze the TCGA cancer samples.

Benefits

- Comparison of "standalone" and "ensemble" modes
- Fast compute time per sample: 2.42 hrs and 0.72 hrs
- Low cost per sample: 0.77 USD and 0.23 USD
- High accuracy overall: 98.9% and 97.2%

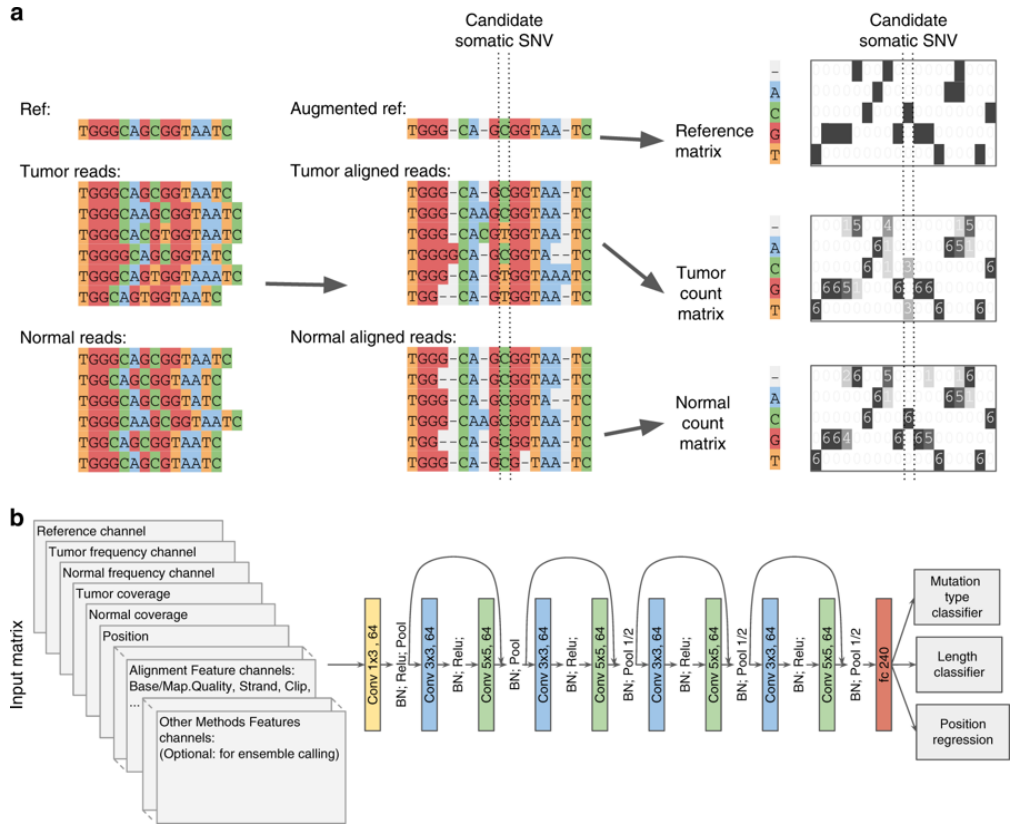

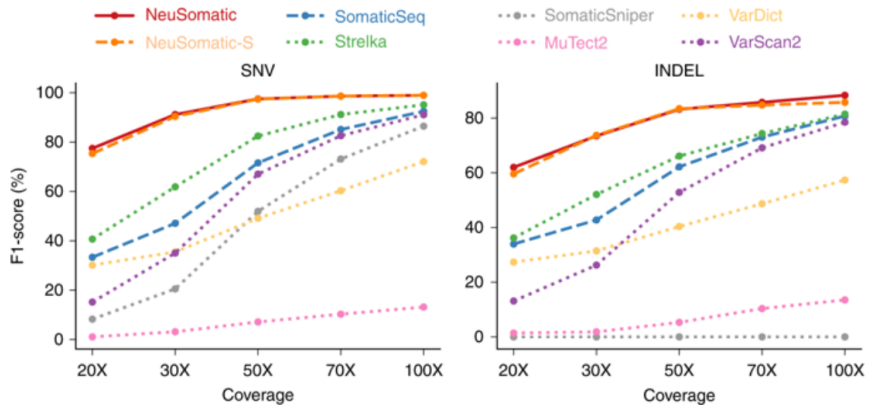

Performance analysis of the sequence coverage impact on the whole-exome sample mixture dataset. In this example, tumor has 50% purity and normal has 95% purity. Y-axis illustrates the highest F1-score achieved by each algorithm for sample alignments coverages ranging from 20× to 100×

Paper: https://www.nature.com/articles/s41467-019-09027-x

Code: https://github.com/bioinform/neusomatic

Autonomous Driving: End-to-End Workflow

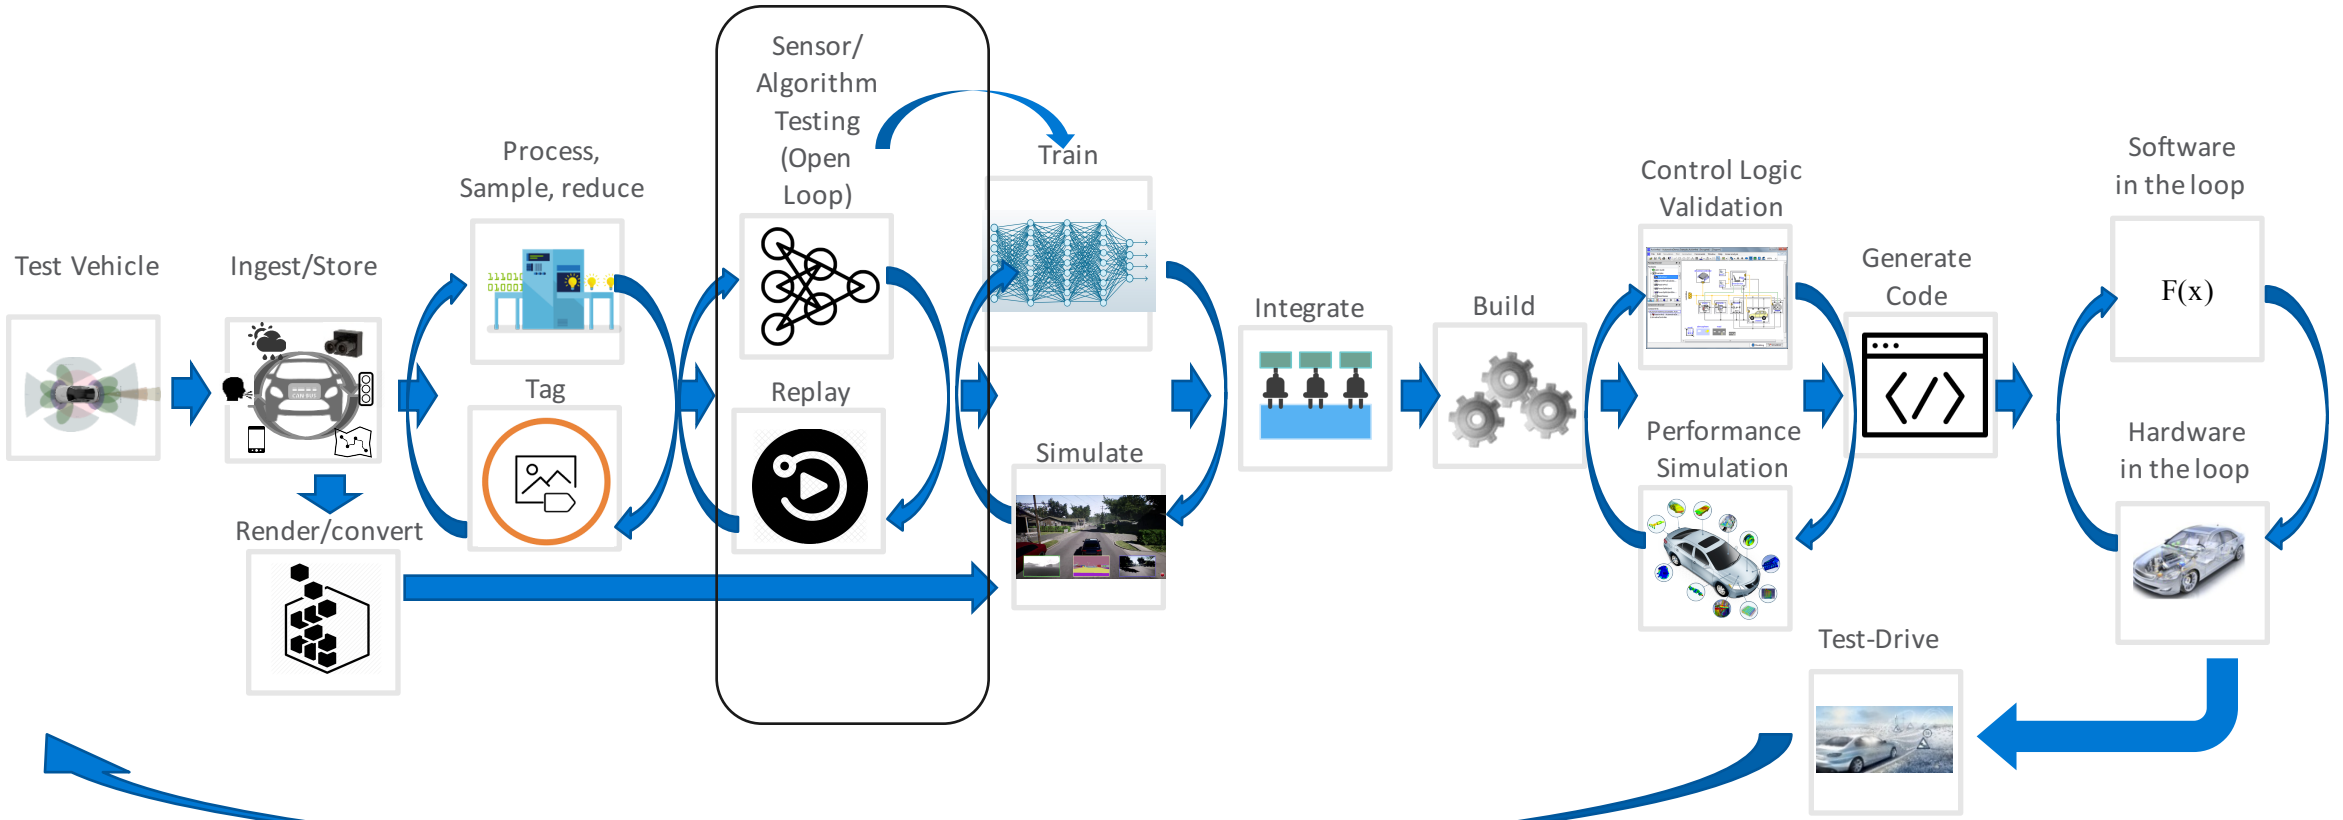

intel | Microsoft

Autonomous Driving: Open Loop or "Re-sim" Explained

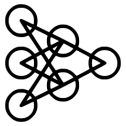

- Industry standard practice for autonomous dev/test
- Uses data recorded from test fleet vehicle sensors to test algorithm performance (mostly perception).
- Can be SiL or HiL – increasingly SiL as HiL rigs are expensive and can't scale to meet the explosion in data ingest volumes
- SiL job VM family selection dependent on the on-board SoC.
  - CPUs vs mix of GPU & CPU vs GPGPUs due to CUDA dependencies
- Must complete in 2~3 days.

Re-simulation Environment	CPU vs GPU	Estimated Volume
#1	Both	40K CPU cores 1800 GPUs
#2	CPU	130K CPU cores
#3	Both	60K CPU cores 2600 GPUs

Resim / Open Loop Examples:

- Vehicle detection
- Pedestrian detection
- Sign detection
- Lane detection
- Light detection

intel | Microsoft

The Wide Reach of HPC

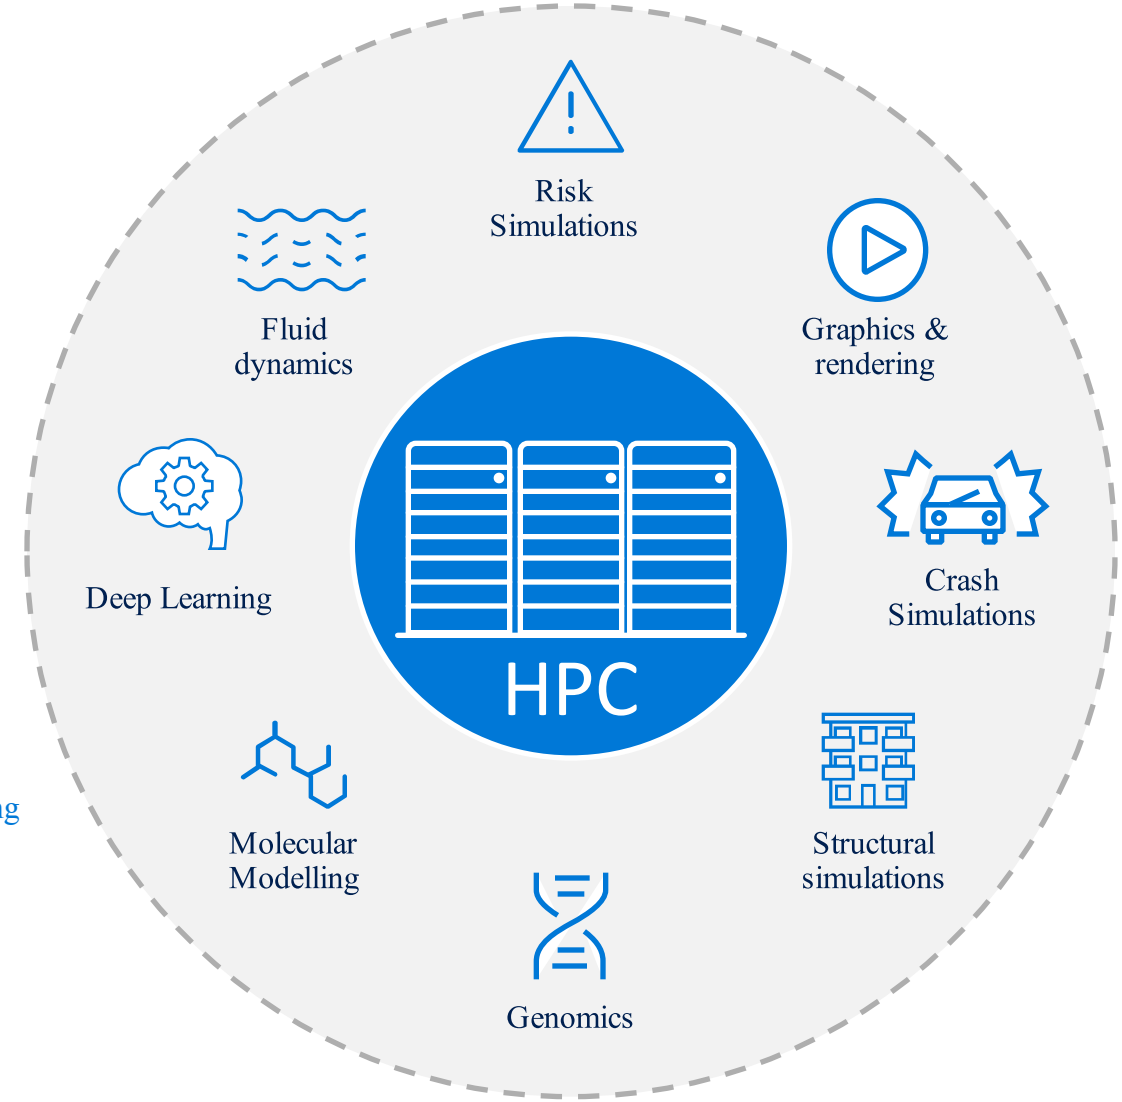

Food & beverages

Automotive

Weather forecasting

Academia

Banking

Oil & Gas

Defense & Aerospace

Oceaneering

Insurance

Ship engineering

Pharmaceutical

Healthcare

Chem Eng

Energy

Engineering & construction

Bio Science

intel | Microsoft

Why do HPC customers look to Cloud?

Agility/
Speed

Scale

Datacenter
Shutdown

Azure enables these benefits,

Through product focus on these needs

Why don't they always migrate HPC? (among others)

Historical Cloud
Cost-performance

Simple, Performant
Hybrid File Access

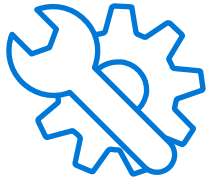

Hard to Migrate, Manage
& Control Cloud/Hybrid

intel | Microsoft

Most performant HPC Infrastructure

Azure Specialized infrastructure for HPC and AI

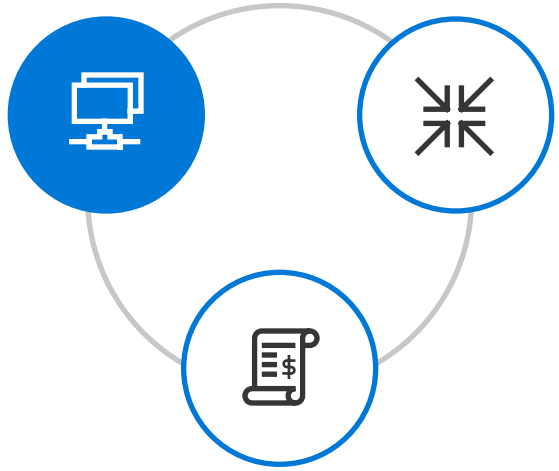

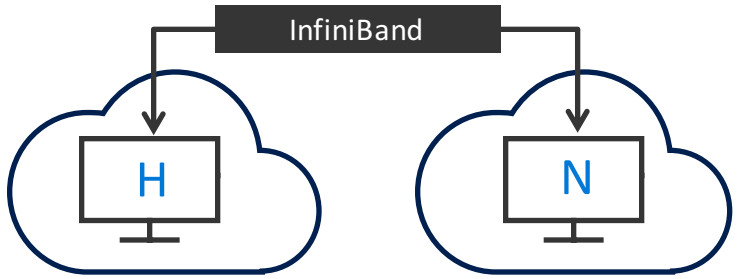

High-performance VMs

Tightly coupled parallel jobs

GPU-enabled VMs

NV—Graphic-based applications

NC—Advanced simulation

ND—Artificial Intelligence (Deep Learning)

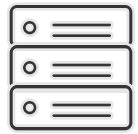

Cray in Azure

Managed, Custom Bare-metal

HPC or Supercomputing

On the Azure Network

>80,000 IOPs Premium Storage

Low latency, high throughput apps

FPGA

PGA

Microservices—AI/Edge

Compute-optimized VMs

Batch processing, Monte Carlo simulations

Large memory VMs

Large databases

intel | Microsoft

HPC Targets by Workload

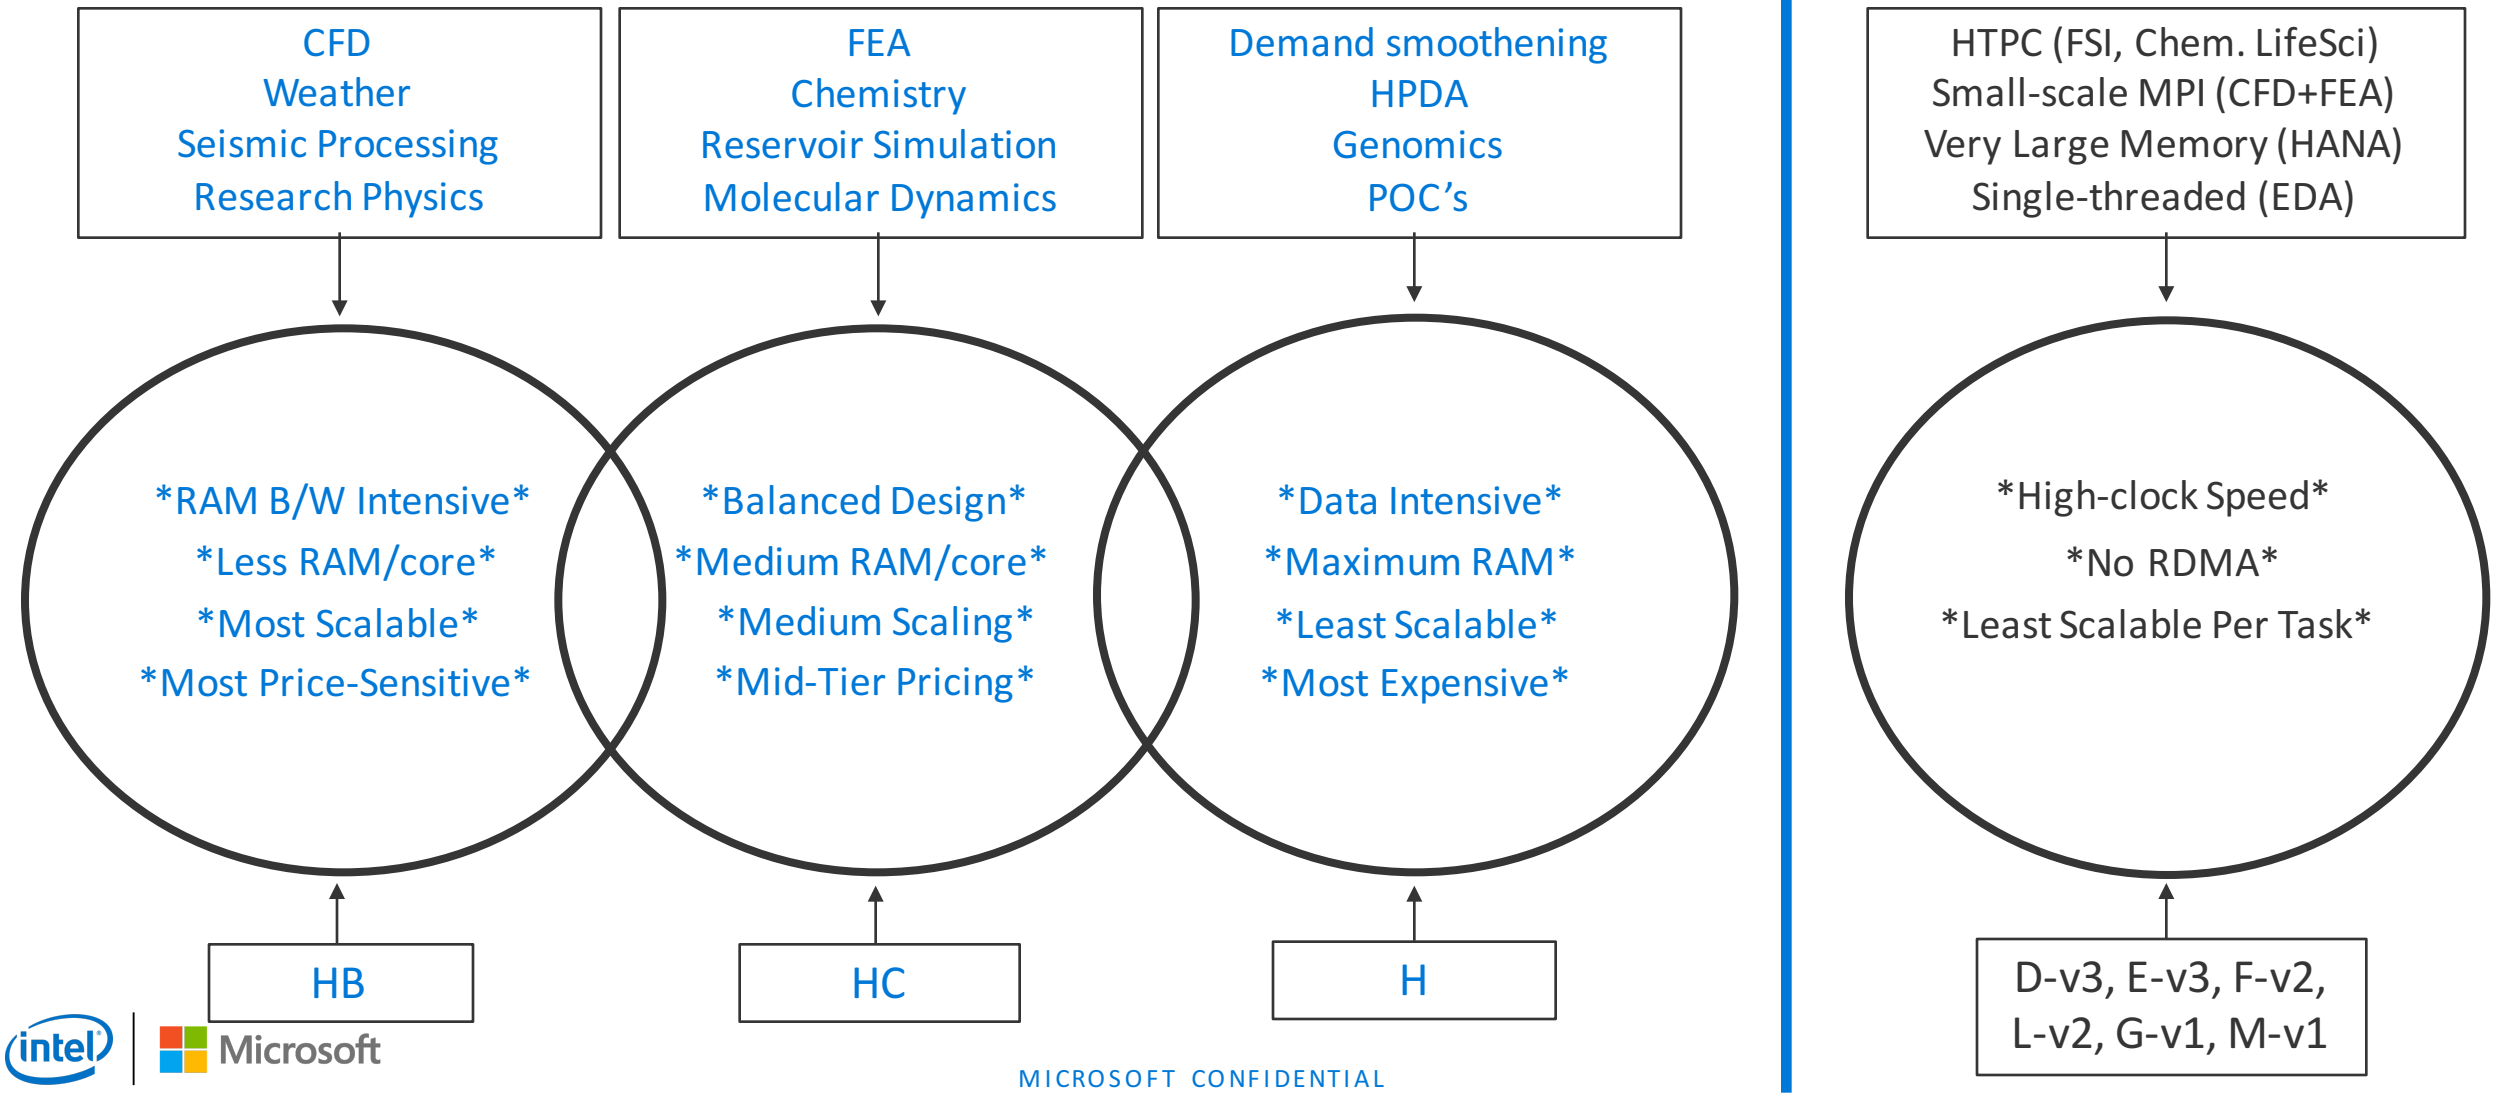

High-Performance Computing VMs (H)

	HB*	HC*	H
Targets	CFD, Seismic, Weather, Genomics	Genomics, FEA, MD, Chemistry	HPDA, Single-Threaded
Workload Driver	Memory Bandwidth	Raw Compute	Large RAM/core
Cores	60	44	16
CPU	AMD EPYC	Intel Xeon Platinum	Intel Xeon E5 v3
Memory Bandwidth	260 GB/sec	190 GB/sec	80 GB/sec
Memory	4 GB/core	8 GB/core	14 GB/core
Local Disk	700 GB NVMe	700 GB NVMe	2 TB SATA
Infiniband	100 Gb EDR	100 Gb EDR	56 Gb FDR
Network	40 Gb Ethernet	40 Gb Ethernet	12.5 Gb Ethernet

New, Q4 2018 Release

High-Performance Computing VMs (H)

Powered by AMD EPYC and Intel Xeon Platinum

	HB*	HC*	H
Workload Driver	Memory Bandwidth	Raw Compute	Large RAM/core
Max MPI Job Size	18,000 cores	13,200 cores	3,200 cores
MPI Support	All	All	Intel MPI 5.x
Burst Buffer Injection Bandwidth	> 1 TB/sec (whole cluster)		TBD
Azure Storage Support	Premium	Premium	Standard
OS Support for RDMA	CentOS/RHEL 7.6 SLES 12 SP4 WinServer 2016	CentOS/RHEL 7.6 SLES 12 SP4 WinServer 2016	CentOS/RHEL 6.x+ SLES 12+ WinServer 2012+

New, Q4 2018 Release

Cray in Azure: A Unique Offering

Dedicated & customized Cray, directly connected to Azure:

- Mission-critical Supercomputing directly connected to your data in Azure
- Custom-configured, bare-metal HPC cluster to supercomputer scale in Azure
- Solve today's simulation needs, while future-proofing for AI, IOT challenges
- Scale with Cray beyond anything else available in the cloud
- No Data movement! Your Cray HPC storage is on the same Azure network as elastic Azure RDMA, GPU, FPGA VMs
- Gain capital efficiencies, get more agility, capacity, science per pound
- Reduce risk with Cray administered & managed system leveraging Microsoft engineering & resources

Sentinel Pilot System for Private Invite-only Preview

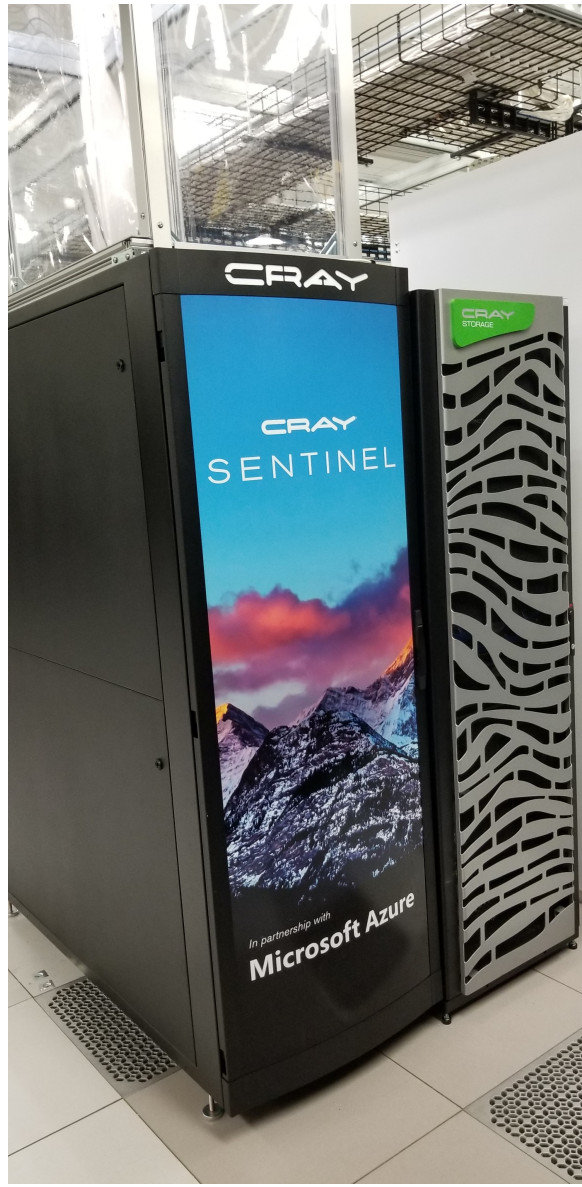

intel | Microsoft

Performant Hybrid Storage:

Avere, NetApp Files , Cray ClusterStor, VM Parallel FS Templates

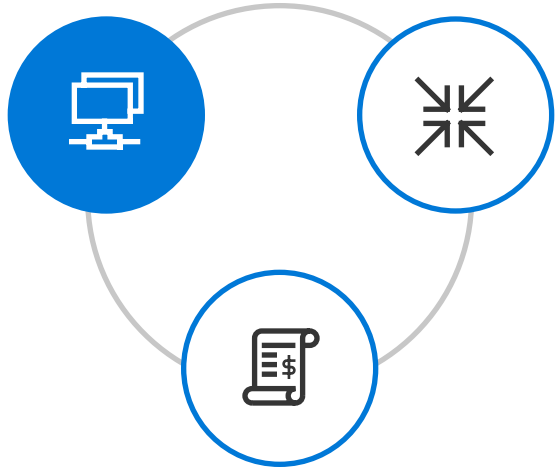

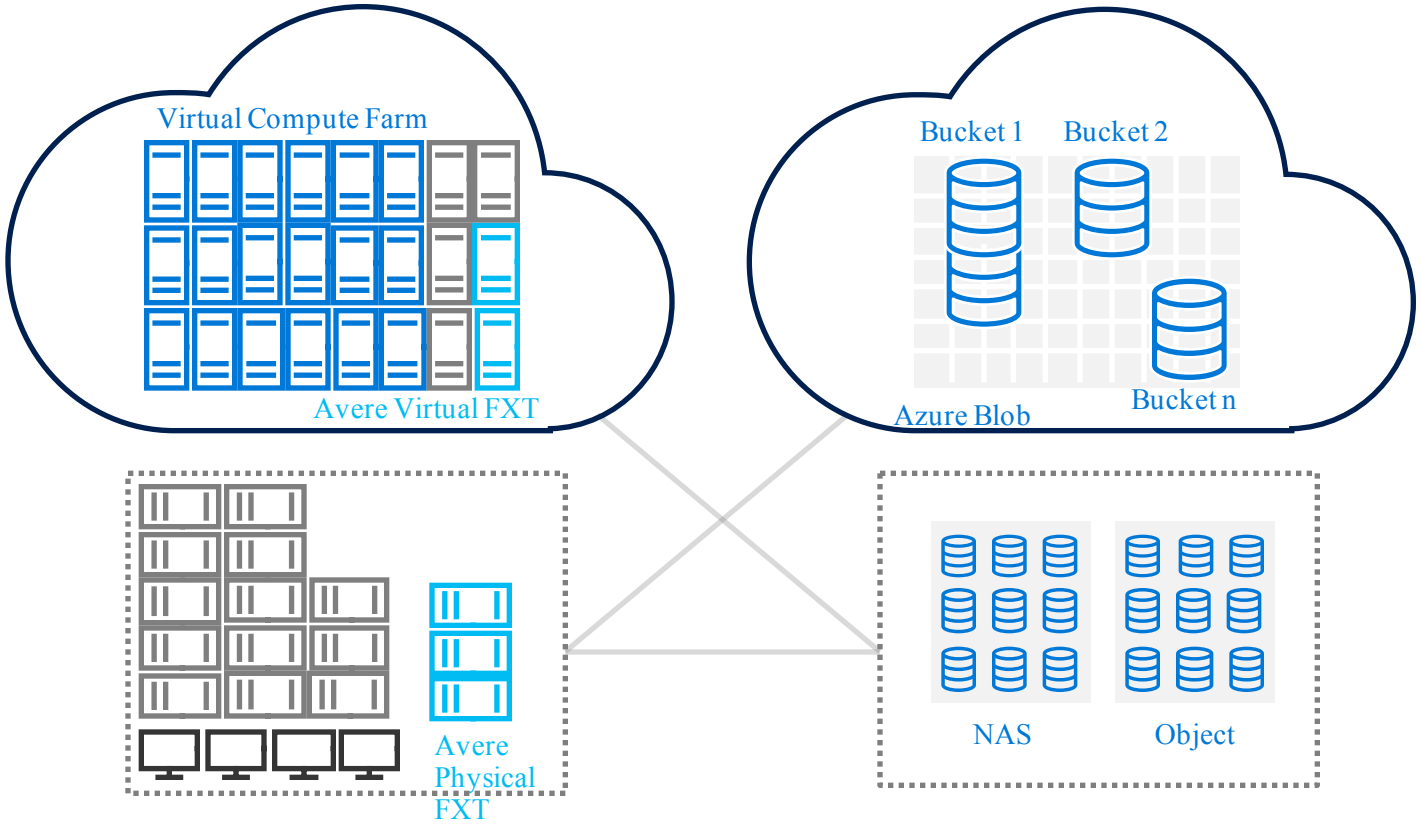

Customer needs	Azure delivers
Avere Low-latency file access	Edge-core architecture
Avere Manage as a single pool of storage	Global namespace (GNS)
Avere High security	AES-256 encryption (FIPS 140-2 compliant), KMIP
High Performance Bare Metal Storage	**Azure NetApp Files** and **Cray ClusterStor** solutions
Parallel Virtual File Systems for VM based write perf	Azure CycleCloud templates for **BeeGFS**, **GlusterFS** w/**Lv2**

intel | Microsoft

Services for Workload Mgmt

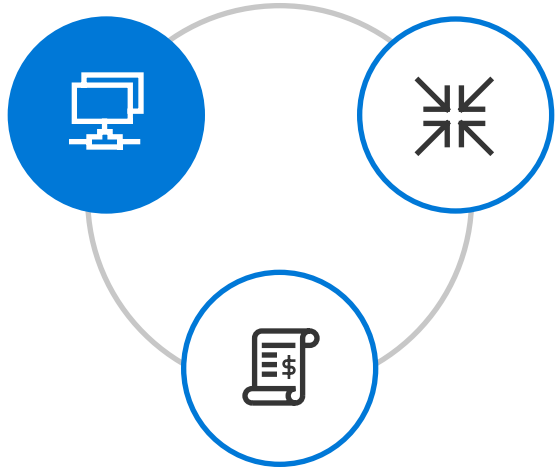

Azure Batch
running jobs

Azure CycleCloud
running clusters

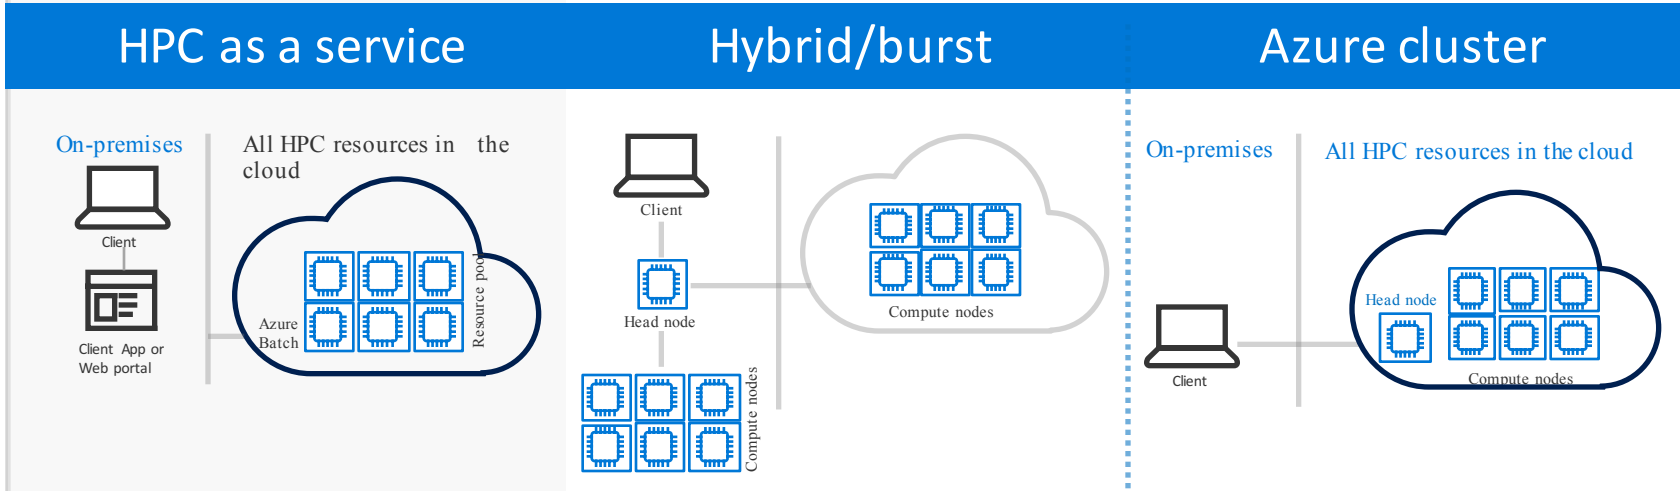

Azure Batch

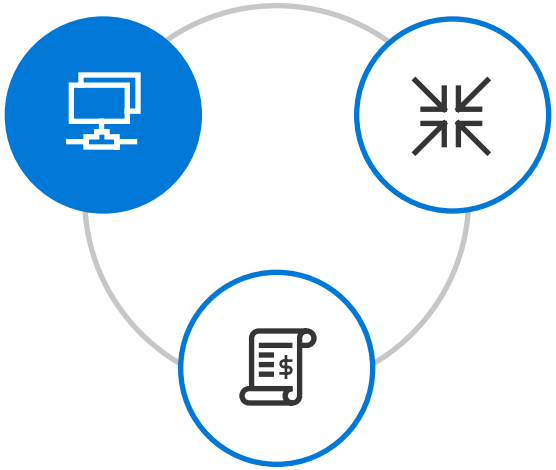

Batch pools

Configure and create VMs to cater for any scale: tens to thousands

Automatically scale the number of VMs to maximize utilization

Easy low-priority and VM sizing, suited to your application

Batch jobs and tasks

Task is a unit of execution; task = application command line (EXE, BAT, CMD, PS1, etc.)

Jobs are created and tasks are submitted to a pool. Next, tasks are queued and assigned to VMs

Any application, any execution time; run applications unchanged

Automatic detection and retry of frozen or failing tasks

intel | Microsoft

Introducing ANSYS Cloud

Available since Feb. 5, 2019

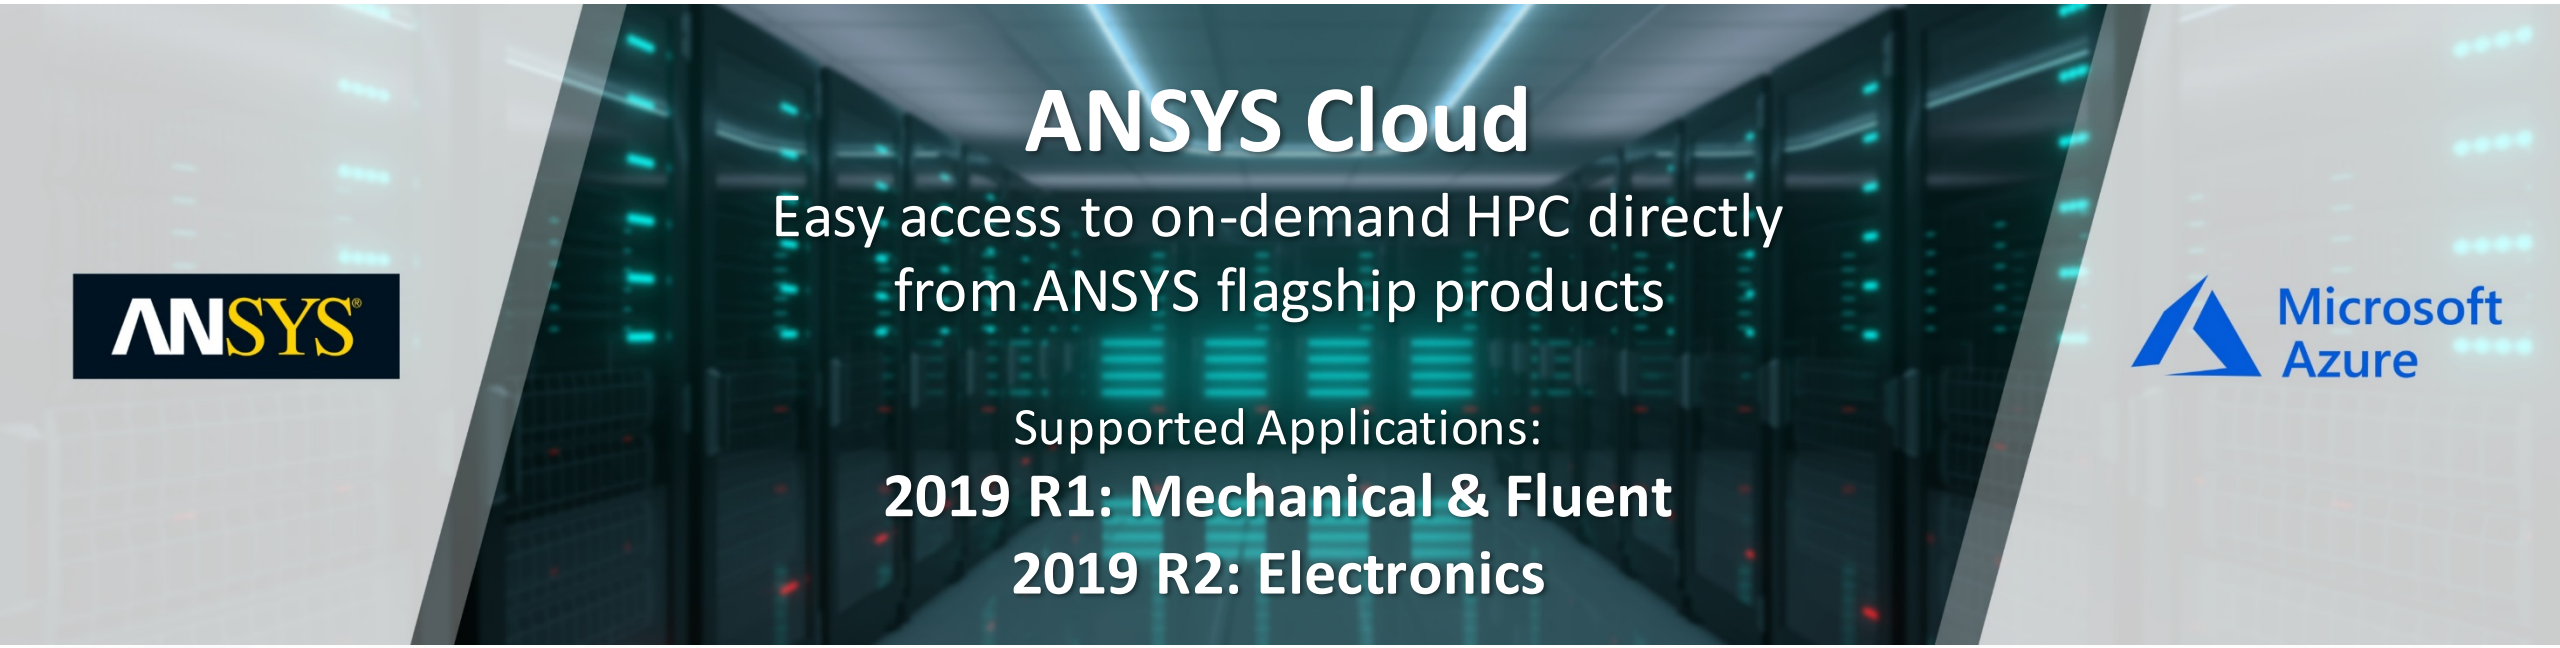

ANSYS Cloud

Easy access to on-demand HPC directly from ANSYS flagship products

Supported Applications:

2019 R1: Mechanical & Fluent

2019 R2: Electronics

ANSYS

Microsoft Azure

1-click burst-to-the cloud

Web-based 3-D Postprocessing

Highly optimized for ANSYS solvers

Single vendor solution for SW+HW

Learn more @ ansys.com/products/platform/ansys-cloud

Recording of ANSYS Cloud webinar

ANSYS Cloud security white paper

ANSYS Cloud overview white paper

Sign-up for an ANSYS Cloud seminar

Azure CycleCloud

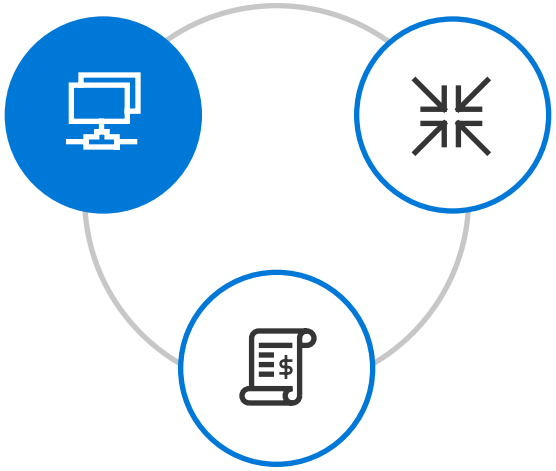

User empowerment

Able to cloud-enable existing workflows and schedulers

Enable instant access to resources

Provide auto-scaling, error handling

IT management

Link workflows for internal and external clouds

Use Active Directory for authentication and authorization

Provide secure and consistent access

Business management

Able to link usage to spend

Provide tools to manage and control costs

intel | Microsoft

Optimizing Infrastructure & Enabling Workloads

- Create a packaged end-to-end workload solutions with detailed product characteristics
- Focused reference architectures with compelling end-to-end experiences
- Partner with ISVs, MSPs, SIs
- Creating new infrastructure/HPC offerings that are workload optimized

Autonomous Driving = GPU, CFD, FEA, AI

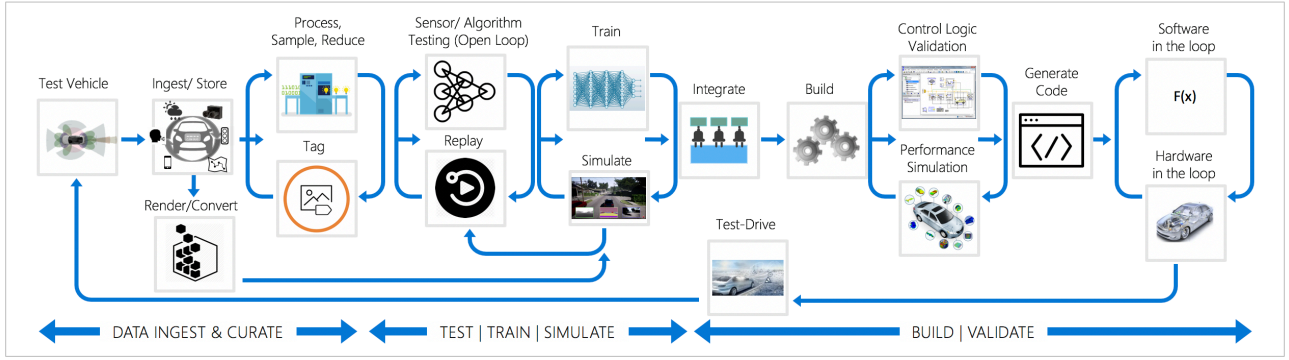

Silicon design – Understanding Workload Needs

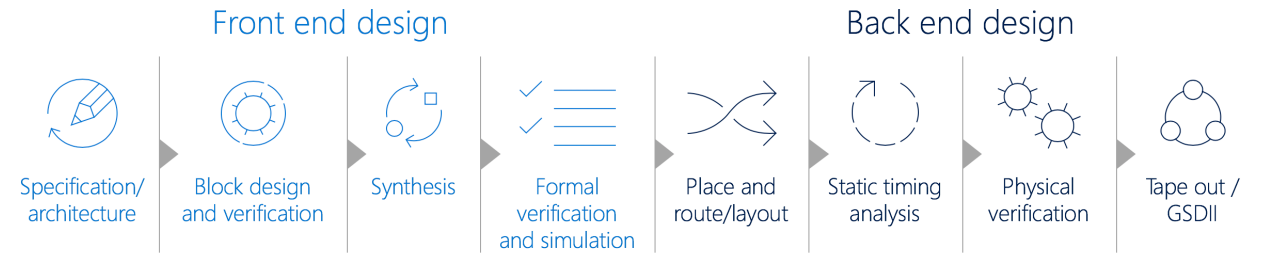

Azure compute mapping

HB-Series 4:1 Mem to Core, up to 100 Gbps high perf bandwidth

F-Series 2:1 Mem to Core, up to 12 Gbps bandwidth

HC-Series 8:1 Mem to Core, up to 100 Gbps high perf bandwidth

H-Series up to 16:1 Mem to Core, low core count, turbo

M-Series up to 28:1 Mem to Core, up to 4 TiB, up to 30 Gbps

Scheduling – Industry Scheduling Tools

High Performance Resource Management and Provisioning – Azure CycleCloud, VM Scale Sets

intel | Microsoft

Silicon design

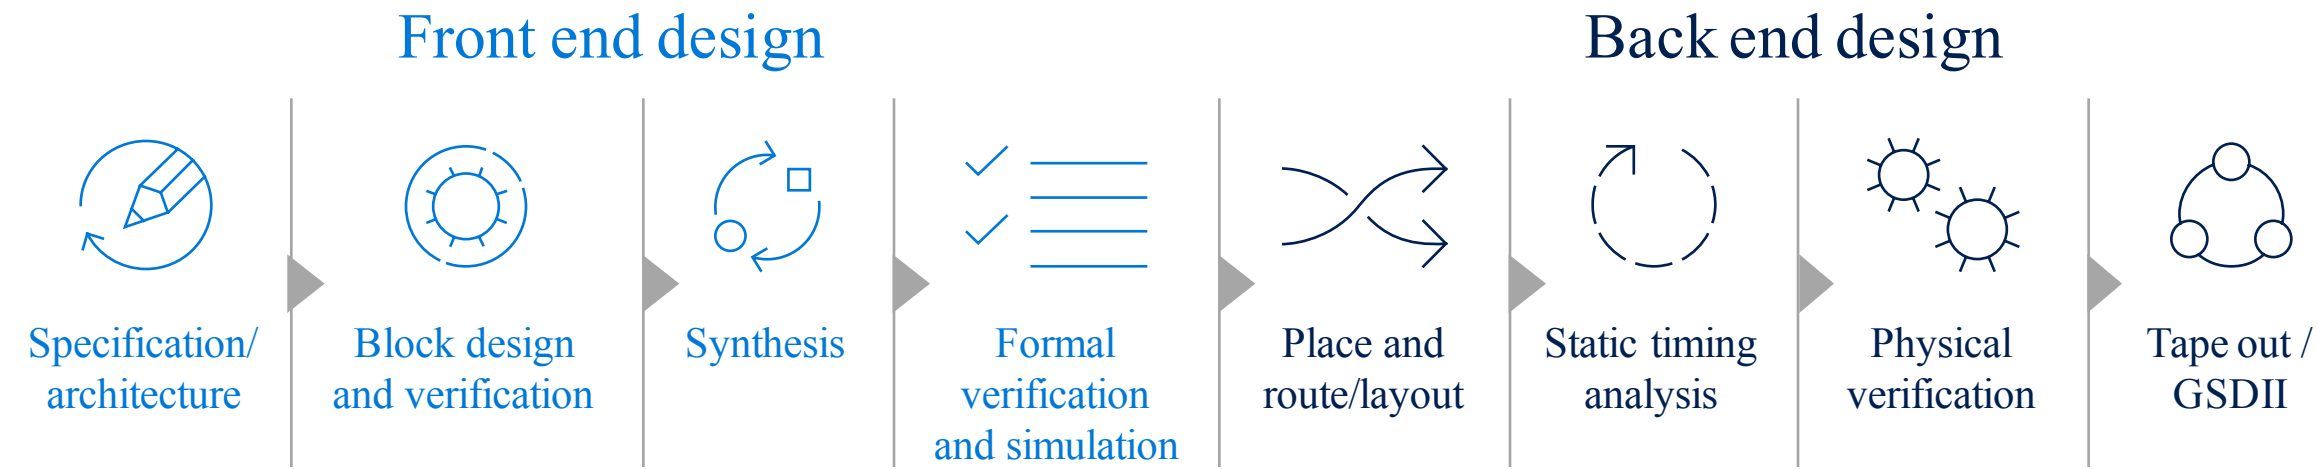

Azure compute mapping

D-Series 4:1 Mem to Core, up to 30 Gbps bandwidth

F-Series 2:1 Mem to Core, up to 12 Gbps bandwidth

HC-Series 8:1 Mem to Core, up to 40 Gbps ethernet bandwidth

H-Series up to 16:1 Mem to Core, low core count, turbo

M-Series up to 28:1 Mem to Core, up to 4 TiB, up to 30 Gbps

Scheduling – Industry Scheduling Tools

High Performance Resource Management and Provisioning – Azure CycleCloud, VM Scale Sets

Genomic Sequencing

Primary Analysis

Secondary Analysis

Tertiary Analysis

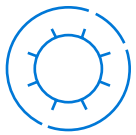

Base Calling

Sequence Reads & Quality Scores

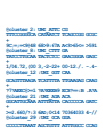

Filter Reads

Read alignment, sorting, duplicates

BAM File merging sorting

Variant Calling Read quantification

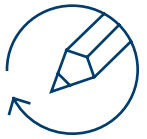

Annotation, Curation, Classification, Interpretation

Clinical Reporting

Azure compute mapping

HB-Series 4:1 Mem to Core, up to 100 Gbps bandwidth

HC-Series 8:1 Mem to Core, up to 100 Gbps high perf bandwidth

F-Series 2:1 Mem to Core, up to 12 Gbps bandwidth

H1-Series: 14 Mem to Core, 12.5Gbps high perf bandwidth. InfiniBand

FPGAs

M-Series up to 28:1 Mem to Core, up to 4 TiB, up to 30 Gbps

Cray CS & ClusterStor – Bare-metal Linux HPC & High Performance Storage

Scheduling – Industry Scheduling Tools (PBS, Grid Engine, LSF, etc)

High Performance Resource Management and Provisioning – Azure CycleCloud, VM Scale Sets, Azure Batch

Q&A