

# Meet the most demanding HPC and AI needs with the help of Microsoft Azure

Luka Debeljak

Sales Manager, Azure Applications & Infrastructure  
Microsoft APAC

Fumiki Negishi

HPC/AI Sales Director  
Intel Asia-Pacific Japan

# Agenda

- Requirements when moving Big Compute to the Cloud
- Azure Infrastructure for Big Compute and HPC workloads
- Azure Batch for SaaS and New Cloud-native Apps
- Big Compute and Artificial Intelligence

# Big Compute and HPC workloads

## Industries



Aerospace



Architecture



Automotive



Financial Services



Science



Manufacturing



Energy



Media & Entertainment

## Workloads

Car crash simulations

Rendering

Audio & video transcoding

Test execution

Financial risk simulations

Genomics

Data ETL

Video pre- & post-processing

R at scale

Oil reservoir simulation

OCR

Compiled MATLAB

Deep learning & AI training

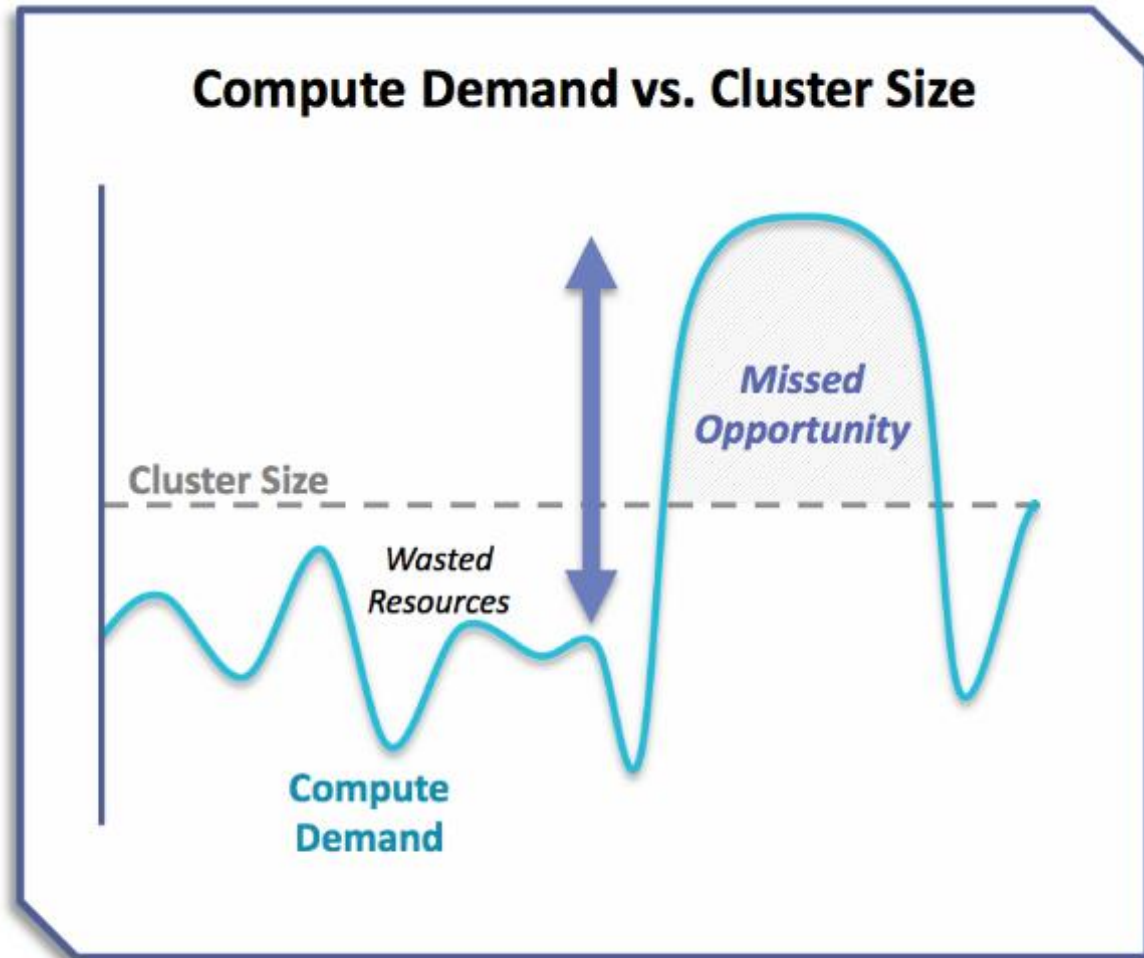
Engineering stress analysis

Image analysis

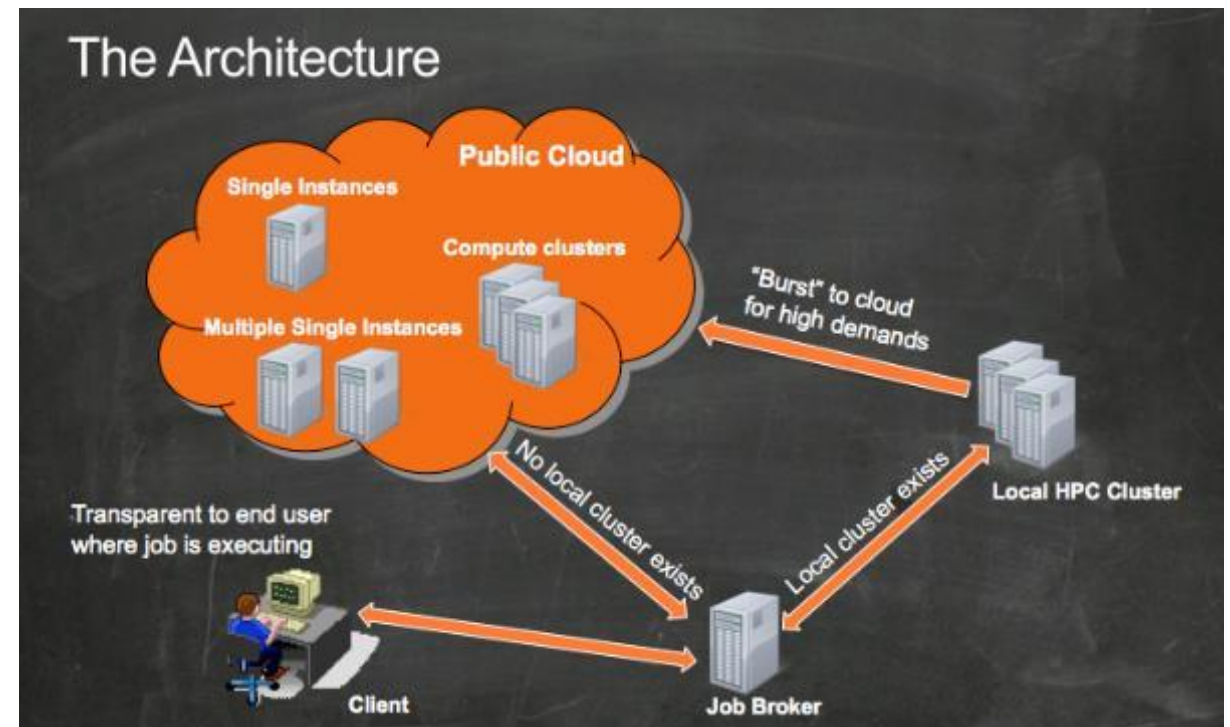
Computational fluid dynamics

Motorsports simulations

# #1 – Zero waiting in line for compute



Johnson & Johnson



# #2 – Ask questions of any scale

Ask the right question,  
regardless of scale

Customers use 100s to 1,000s  
Of cores to answer business-critical  
Questions they couldn't have done before.



# #3 – Users with unique requirements are OK

Trivial to support different use cases

Different RAM ratios, GPU, FPGA, Application/OS needs

Move workloads that don't fit internally to Cloud



#4 – Cloud gets faster & cheaper over time...

#5 – Time and Cost are the sole metrics that matter






# #6 – Accelerating answers, accelerates people

## SCALABLE COMPUTING (in hours)



2880 hours /  
120 Days to  
Decision


## ANTICIPATED BENEFIT (in hours)



1456 hours /  
60.6 Days to Decision

# #6 – Accelerating answers, accelerates people

## SCALABLE COMPUTING (in hours)



2880 hours /  
120 Days to Decision


## POST ADOPTION: AGILE DESIGN PROCESS

### Computing & Analysis



Higher Quality Output,  
Iterative Analysis,  
Less Context Switching

# #7 – Every smart person gets their own sandbox



## Old: Shared internal cluster

- Competition for resources
- Waiting in line for compute
- Shared downtime

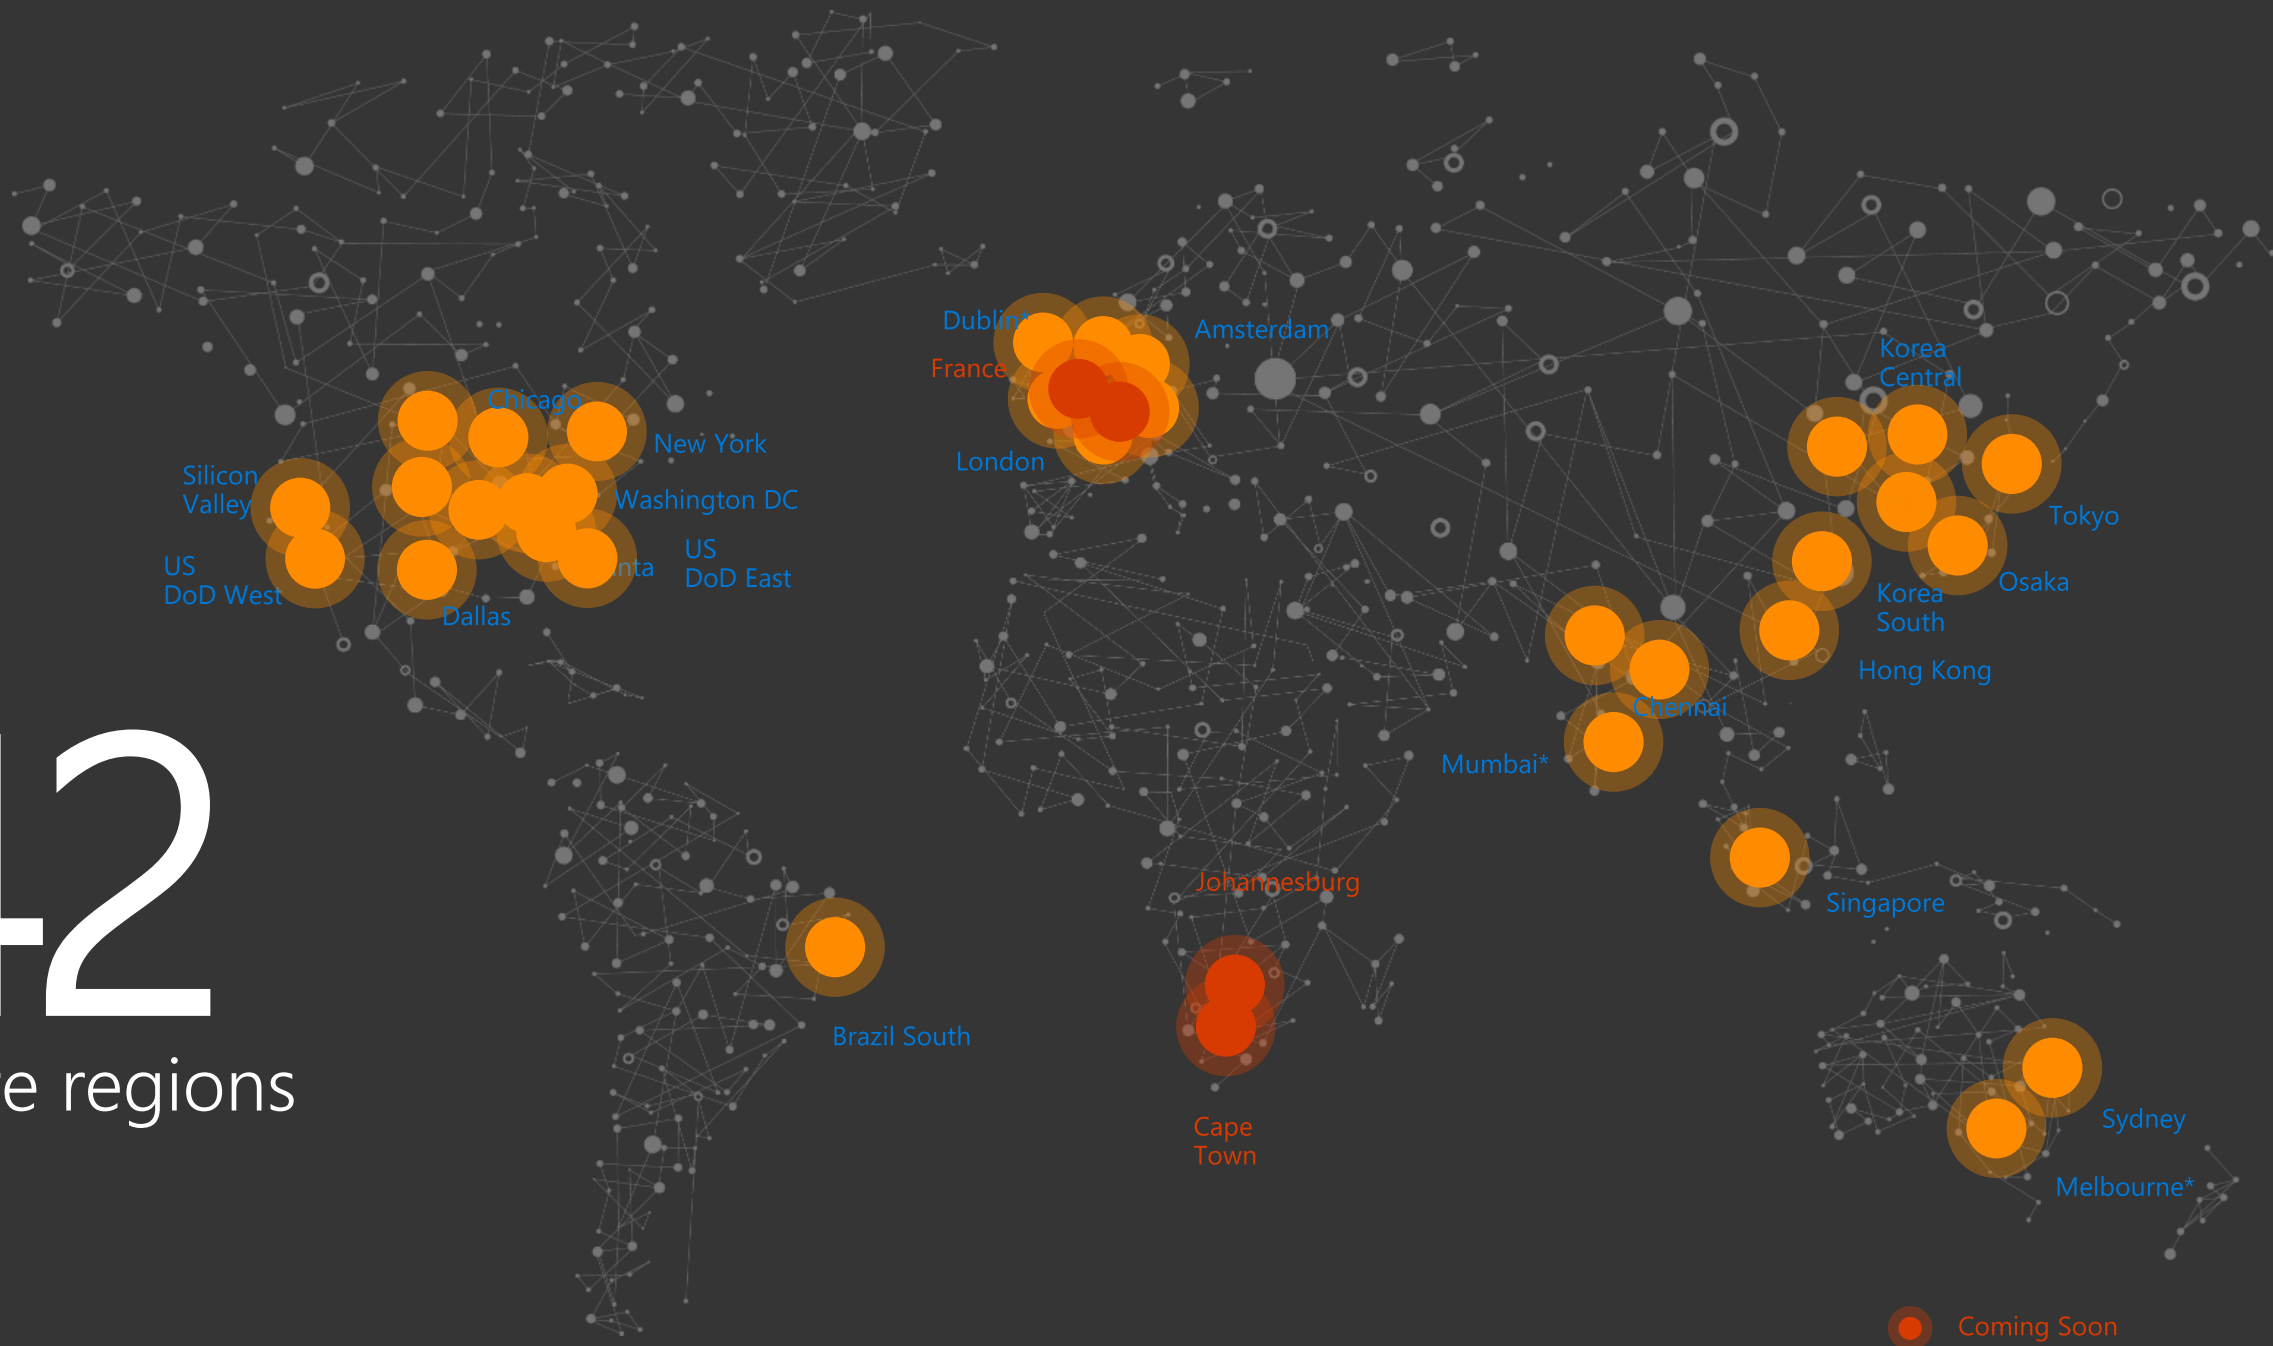
## New: Cluster Per Researcher

- Remove bottlenecks
- Cost controls to manage \$
- No waiting = 2x faster users

# Microsoft Azure and Big Compute

# 42

Azure regions





## Tools

Developer tools

DevOps

Portal + scripting



## Advanced workloads – Platform-as-a-Service (PaaS)

Web + Mobile + Media

Identity

Internet of Things

**Data + Analytics**

Microservices

**Artificial intelligence**

Containers

Cognitive services

Serverless

**High performance computing**



## Core infrastructure – Infrastructure-as-a-Service (IaaS)

Security

Management

Compute

Storage

Networking





Azure Stack + Hybrid

# Application types



## Embarrassingly parallel:

- Applications do not communicate
- May share common store & data
- May have dependencies
- E.g. Monte Carlo simulations, transcoding, rendering



## Tightly coupled:


- Applications communicate; mainly use MPI
- Requires low latency, high bandwidth networking for scale
- E.g. car crash simulation, fluid dynamics, AI training



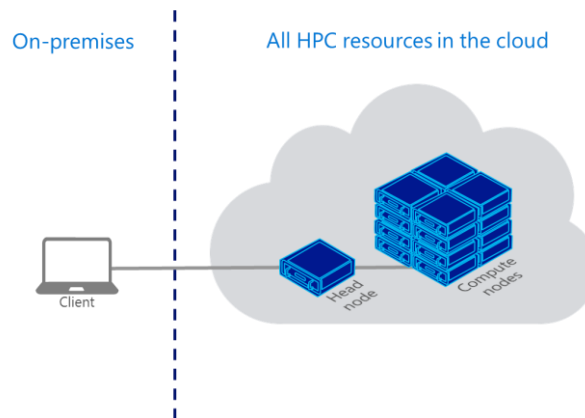
# Azure HPC: Two main types

## Self-managed

### Cloud burst

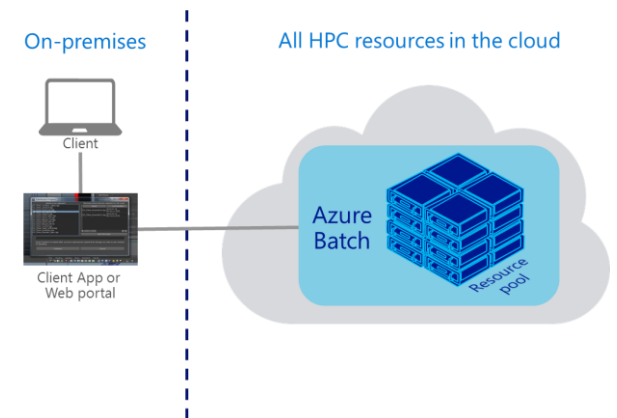


### Cluster on the cloud



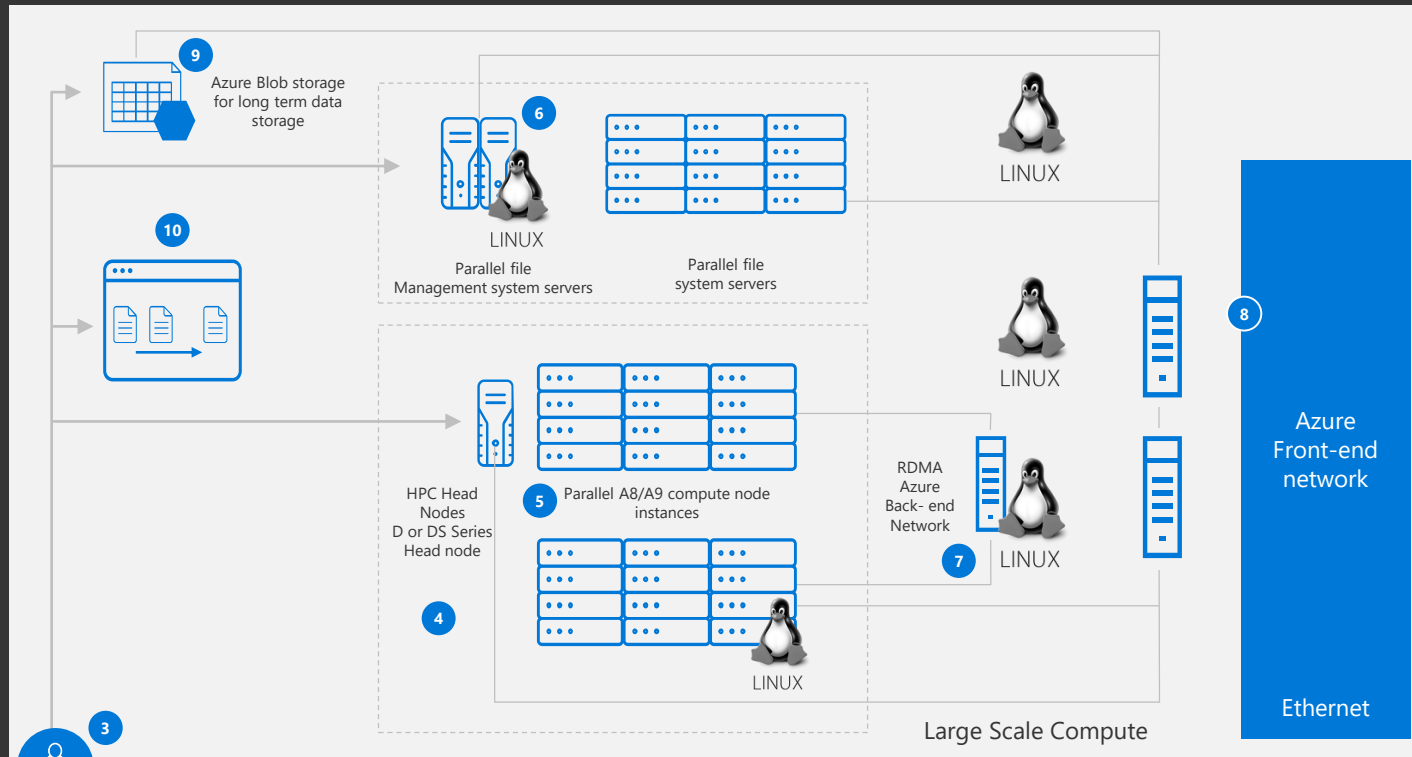
## Fully-managed

### HPC as a service


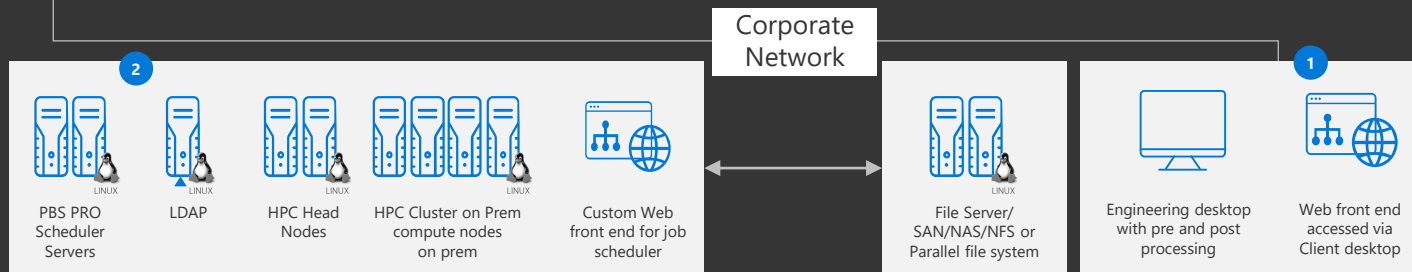




# Big Compute Conceptual Architecture



- 1 End User Infrastructure
- 2 On Prem HPC
- 3 Connectivity to Azure
- 4 HPC Head Node
- 5 HPC Compute Nodes
- 6 Lustre Parallel File System
- 7 RDMA High Speed Networking
- 8 Azure Front End Network
- 9 Blob storage
- 10 Job Submission Web Interface



# No-compromise HPC and AI VMs




- Up to 16 cores, 3.2 GHz E5-2667 V3 Haswell processor
- Up to 224 GiB DDR4 memory
- **FDR InfiniBand** (56 Gbps, 2.6 microsecond latency)
- 2 TB of local SSD



- Up to 72 cores, 3.7 GHz Intel Xeon Scalable (Skylake)
- Up to 144 GiB DDR4 memory
- **Accelerated Networking (30 Gbps VM-to-VM)**
- 500 GB of local SSD



- Up to 4 NVIDIA Tesla K80 GPUs
- Up to 24 cores
- Up to 224 GiB memory
- Up to 1440 GiB of local SSD
- **FDR InfiniBand**



- Up to 4 NVIDIA Pascal P100 GPUs
- Up to 24 cores
- Up to 448 GiB memory
- Up to 3 TB of local SSD
- **FDR InfiniBand**



- Up to 4 NVIDIA Pascal P40 GPUs
- Up to 24 cores
- Up to 448 GiB memory
- Up to 3 TB of local SSD
- **FDR InfiniBand**



- Up to 4 NVIDIA Tesla M60 GPUs
- Up to 24 cores
- Up to 224 GiB memory
- Up to 1440 GiB of local SSD



- Up to 4 NVIDIA Tesla V100 GPUs
- Up to 24 cores
- Up to 448 GiB memory
- Up to 1344 GiB of local SSD
- **FDR InfiniBand**


# Azure is an Intel-powered Platform



**Makes clouds faster**  
Intel® Xeon® processors for  
Azure compute and storage



**Makes cloud smarter**  
Intel® Field-Programmable  
Gate Arrays (FPGA)



**Makes clouds safer**  
Intel® SGX enhances security with  
encryption data during computation



**Accelerates networking for  
more efficiency:**  
Intel® Silicon Photonics 100G  
PSM4



**Maximizes performance  
across operating systems:**  
Clear Linux\* OS for Intel®  
Architecture



**Enables the future of AI:**  
Intel® Open Source machine learning  
frameworks and libraries

# More compute for your Powerful Applications

Be more productive running your business applications using the right workloads



## High-performance compute

High-performance compute workloads; modeling; simulations; genomic research

Intel® Xeon® processor  
E5-2667 v3 with  
DDR 4 memory  
  
Intel® Xeon® processor  
E5-2670

Azure H and A8-11 Series



## Compute intensive

High CPU-to-memory ratio; massive large-scale computation; deep learning

Intel® Xeon® Platinum 8168  
processor

Fv2 VM family



## SAP workloads

SAP applications across Dev/Test and production scenarios. SAP NetWeaver; SAP S4/HANA; SAP BI

Intel® Xeon® E7-8890 V4  
processors

SAP HANA VM family



## Memory optimized

Large database workloads; ERP; SAP; data warehousing solutions

Intel® Xeon®  
E5-2673 v4 processors

Azure GS, G, DSv3, Ev3  
and DS Series



# High Performance Computing in Azure

Running H-series and Fv2-series VMs on Intel architecture delivers scalable, unparalleled performance for your complex engineering and scientific workloads.



Break free from the limitations of on-premises infrastructure



Analyze large-scale data



Run simulations and financial models



Reduce time to market

# Intel® Xeon® Scalable processors in Azure

Intel® Xeon® Platinum 8168 is Intel's fastest processor in the public cloud. Combine it with the new Fv2-series VMs, and you get the fastest VMs in Azure.



Intel® Xeon® Scalable processor

Intel® AVX-512






Intel® QAT

Intel® Arria® 10 FPGAs

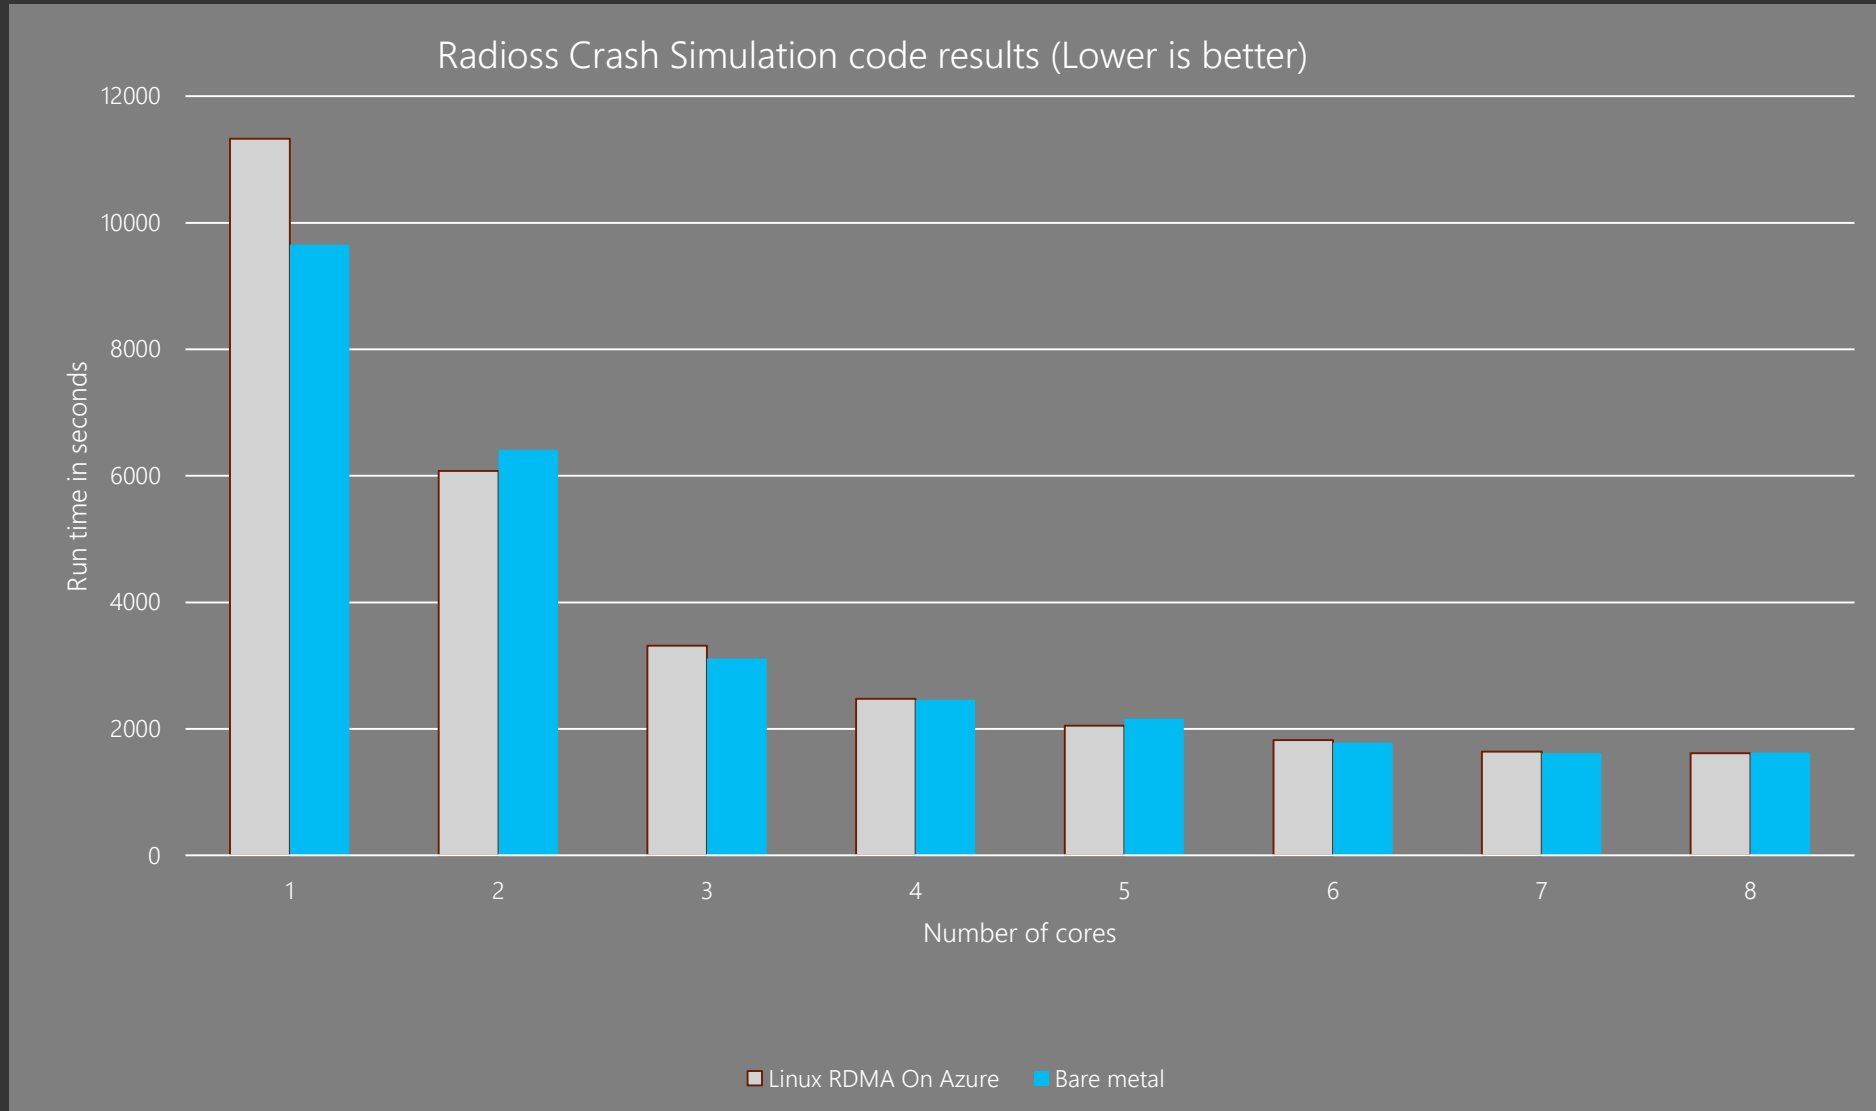


for the most high-demand apps  
for workload-optimized performance  
to speed up data compression and cryptography  
for ultra low latencies

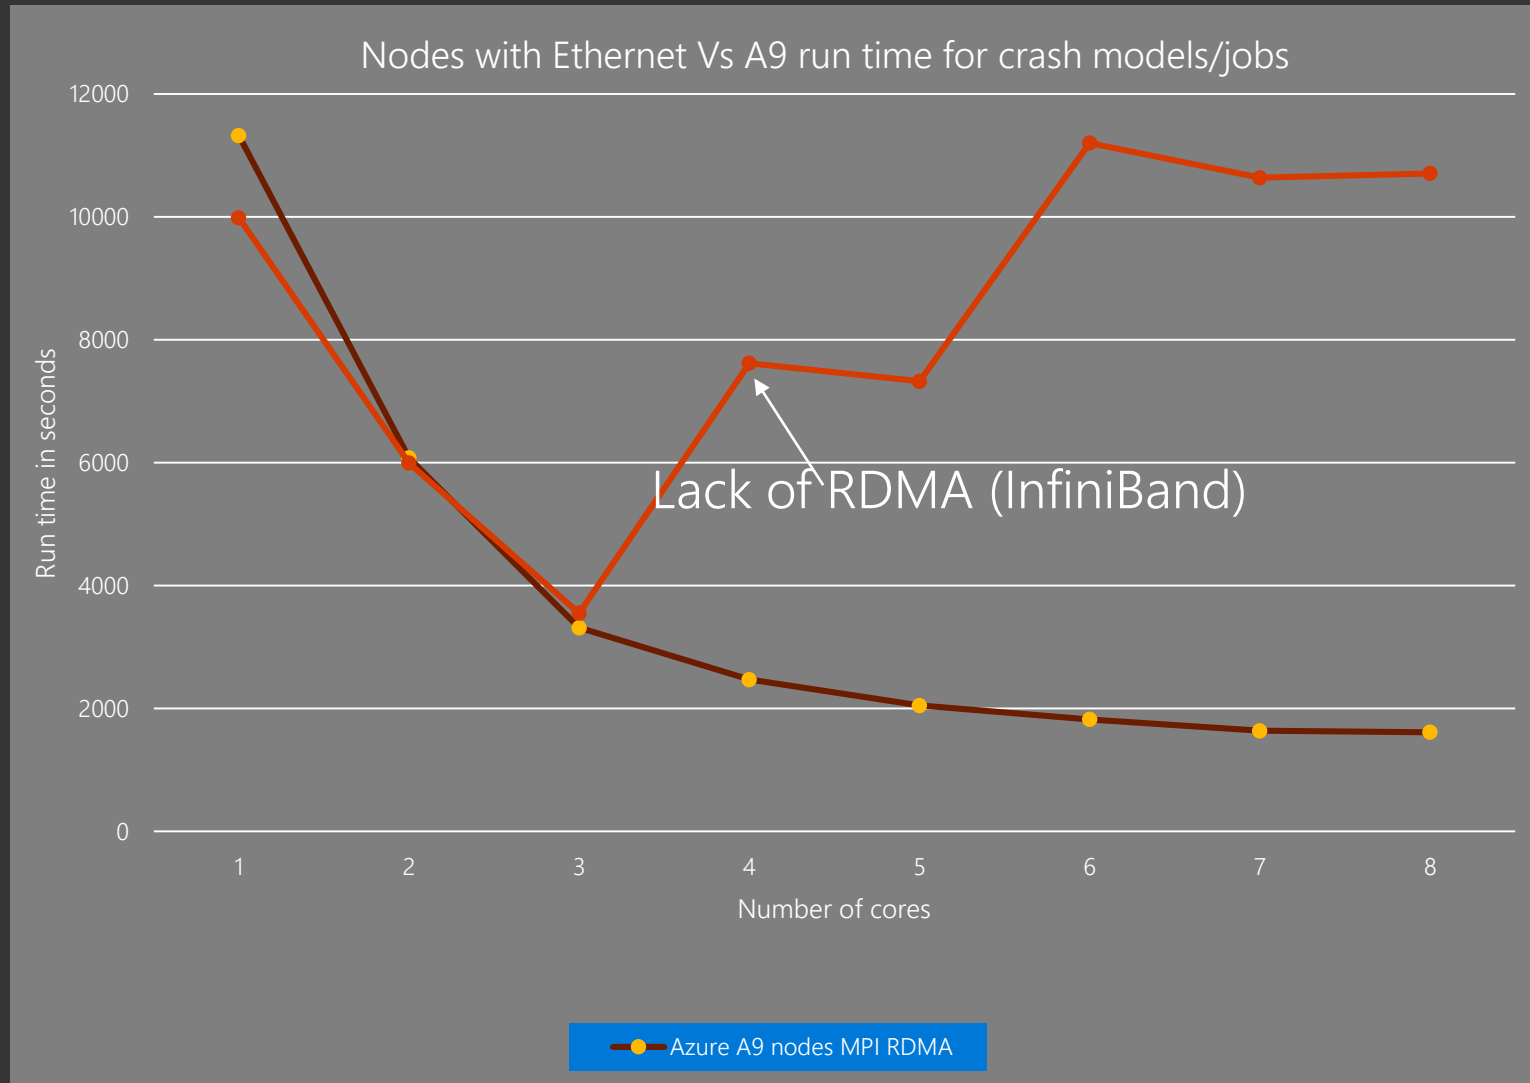
## Ideal for compute-intensive workloads

-  Financial workloads
-  Scientific analysis
-  Genomics
-  Geothermal visualization
-  Deep learning

# Why InfiniBand RDMA matters?



# Why InfiniBand RDMA matters?





# Supported applications, solvers, services, platforms and frameworks

## Cloud Workstation:



## Deep Learning and AI Training:



Caffe



## Cloud Rendering:



## Supported OS:



## HPC Simulation and Analysis:



# Cycle Cloud: HPC cluster as a service



Provisioning



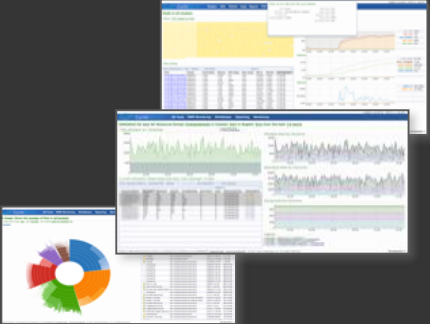
Cluster Configuration



Monitoring




Optimization



# Azure Batch

# Azure Batch

Enable applications and algorithms  
to easily and efficiently run in  
parallel at scale



Rendering

Media transcoding & pre-/post-  
processing

Test execution

Monte Carlo simulations

Genomics

Deep Learning

OCR

Data ingestion, processing, ETL

R at scale

Compiled MATLAB

Engineering simulations

Image analysis & processing

# How these services are built in Azure: Using Azure Batch

## User application or service

Get and manage VMs

Install task applications

Manage and authenticate users

Start the tasks

Move task input and output

Queue tasks

Task failure? Task frozen?

Scale up and down

PaaS  
Cloud Services

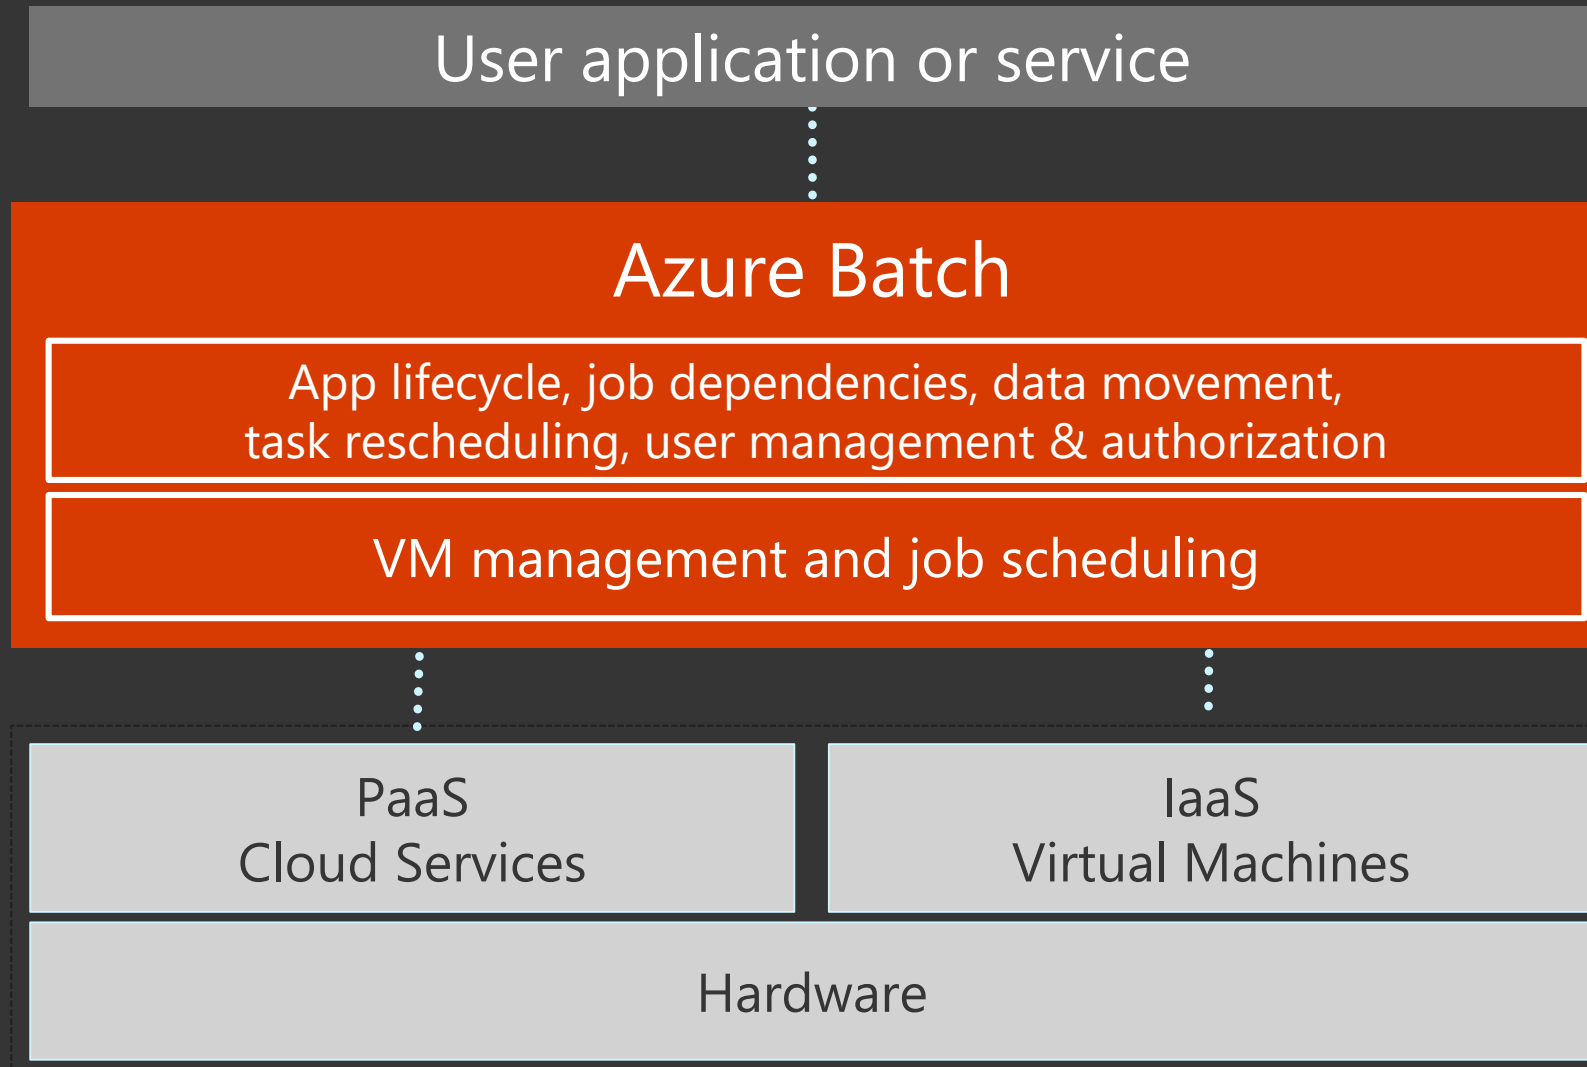
IaaS  
Virtual Machines

Hardware

Significant amount effort spent managing compute resources, security, data movement, job running, and application lifecycle, not related to your actual workload or business

Provided by the cloud platform


# Azure Batch: HPC as a Service



Provided by the cloud platform

- Don't worry about the "plumbing"
- Focus on the workload/app
- Access higher-level capabilities
- Minimize the required cloud or Azure experience

# Azure Batch focus areas




Elastic

Capacity on demand  
Jobs on demand



Efficient

Scale according to load  
Pay by the minute



Scale

1 to 10,000's VMs  
1 to millions of tasks



Cost  
effective

No charge for Batch;  
pay for used resources  
No head node  
Use low-priority VMs

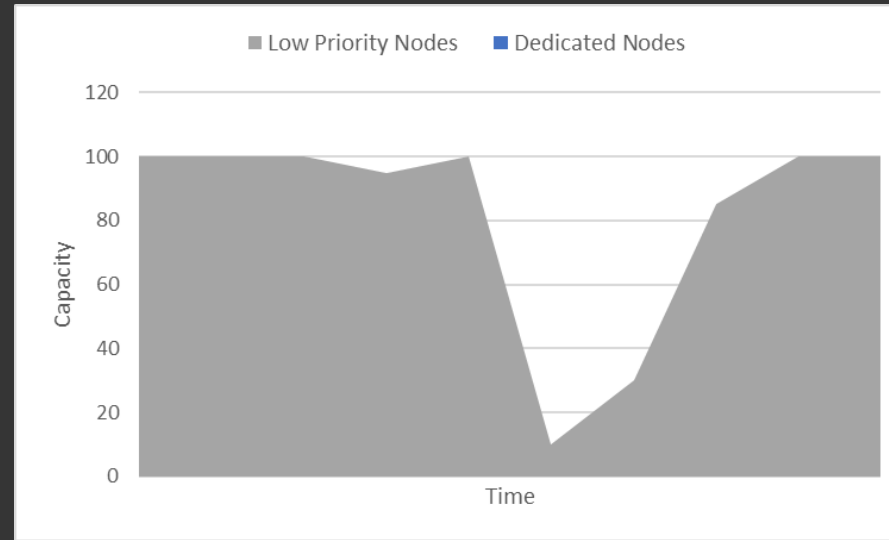
# Azure Batch capabilities

- Access via API's, CLI's, and UI's:
  - .NET, Java, Node.js, Python, REST
  - PowerShell, x-plat Azure CLI
  - Azure Portal, Batch Labs x-plat client UI
- Choice of VMs:
  - Windows or Linux
  - Standard or custom images
  - Windows pool can use AHUB
  - Use low-priority VMs
- Rich app management:
  - Get apps from blobs, Batch app packages, package managers, custom VM images
  - Docker container images
- Pool scaling:
  - Manual or automatic
- VM networking:
  - Pool VMs can be in a VNET
- Job scheduling:
  - Supports both embarrassingly parallel and tightly coupled MPI jobs
  - Run > 1 task in parallel per node
  - Detect and retry failed tasks
  - Can set max execution time for jobs and tasks
  - Task dependencies
  - Job prep and cleanup tasks
- Monitoring:
  - VM monitoring and auto-recover
  - Metrics and logs available via Portal and API




# Batch low-priority flexibility

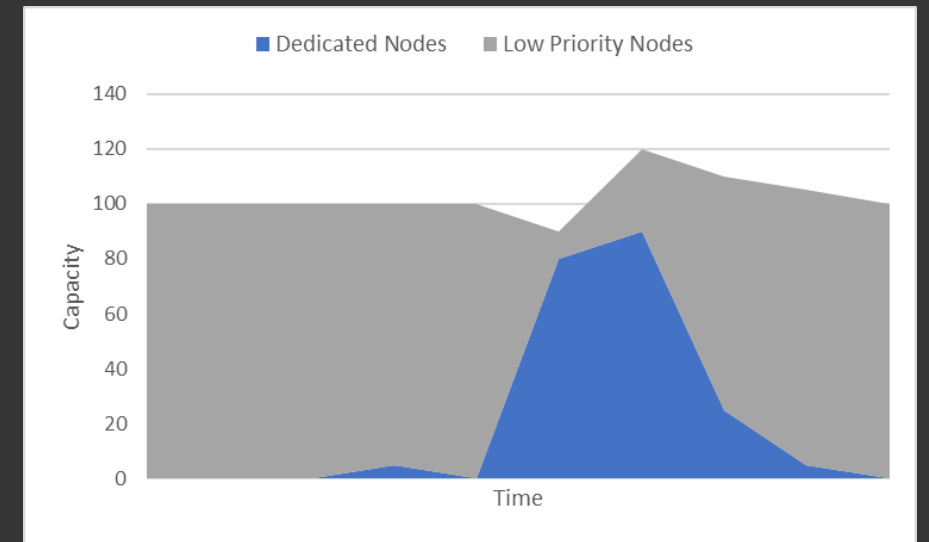
## Lowest cost



## Lower cost, with guaranteed baseline capacity



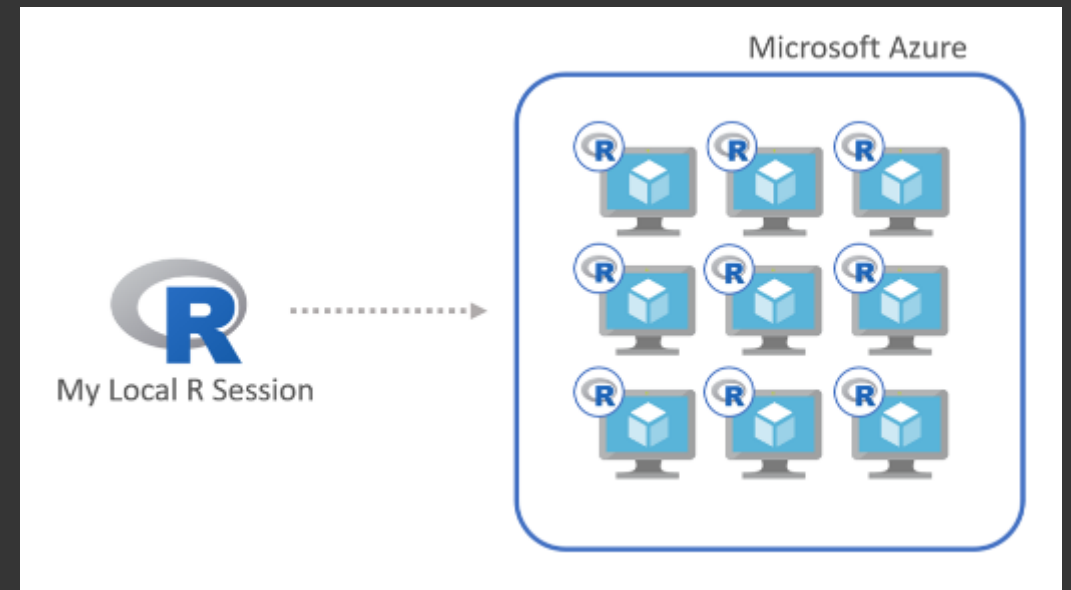
## Lowest cost, while maintaining capacity



# R - doAzureParallel

<https://github.com/Azure/doAzureParallel>

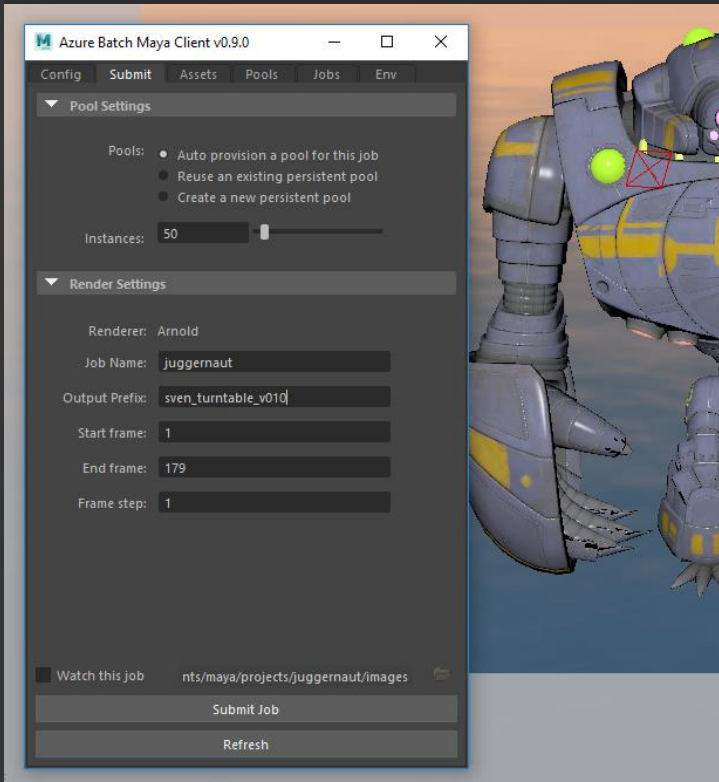
- Scale up R execution using Batch
- Parallel backend package for popular *foreach* package
- Each iteration of *foreach* loop runs as a Batch task



# Batch Rendering

Autodesk 3ds Max / Maya


Integrated Client Plugin



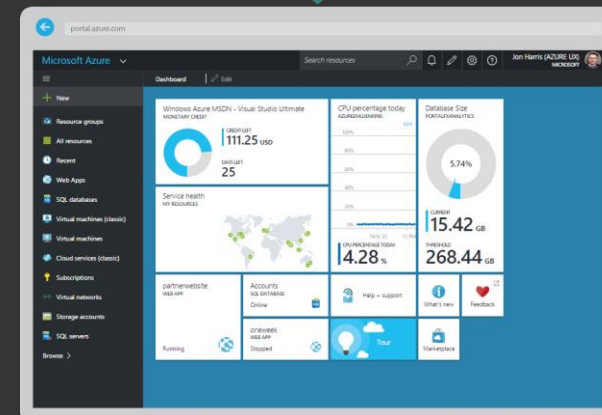
Upload assets

Submit job

Return outputs



- Monitoring
- Reporting
- Single bill



# Azure Batch AI Training

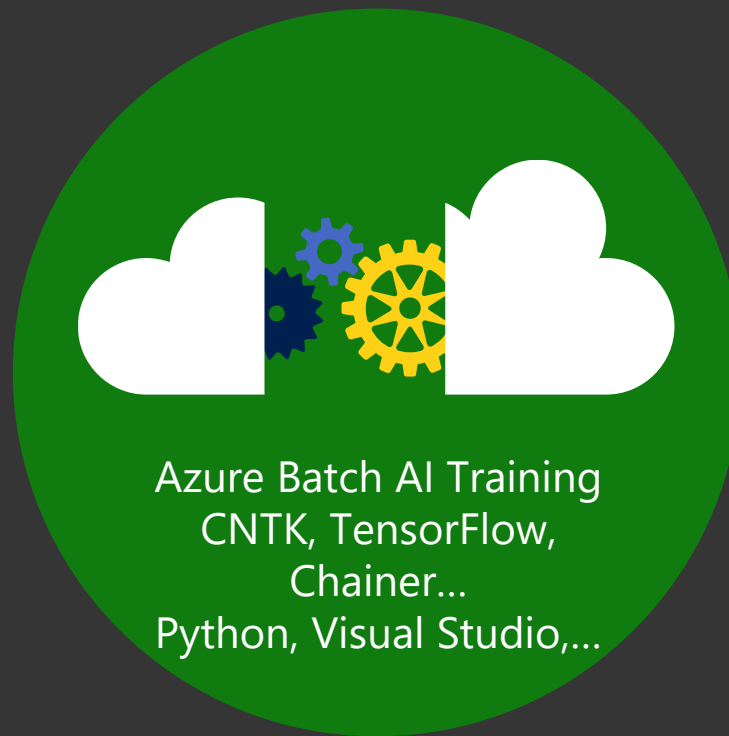
# Train and Deploy Custom AI End-to-End



Your Data (Images, Text, Logs, Time Series...)

Your Data

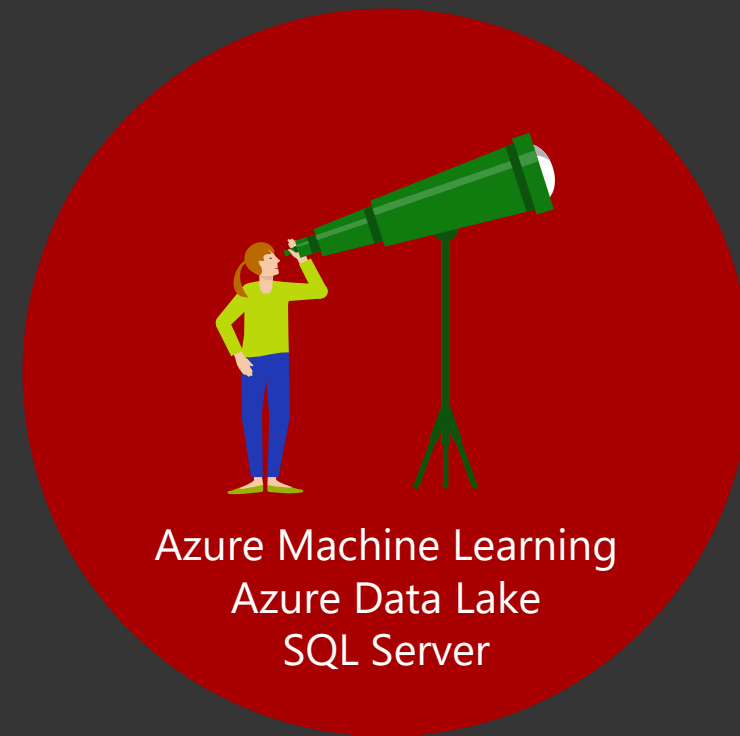
+



Azure Batch AI Training  
CNTK, TensorFlow,  
Chainer...  
Python, Visual Studio,...

Training With Scale-Out  
GPU Clusters on Demand

=



Azure Machine Learning  
Azure Data Lake  
SQL Server

Intelligence In Your  
Apps and Data Services

# Azure Batch AI Training Service

- Managed Service
- Supports Role Based Access Control
- Hierarchical Quota Management
- Easily Provision VMs at scale
- Load based automatic scaling
- Run experiments in Parallel
- Run in Containers or directly on VM
- Run any toolkit (CNTK, Tensorflow, Caffe, Chainer...)
- Only compute cost. Service is free

# Azure Batch Shipyard

<https://github.com/Azure/batch-shipyard>

- Drive Batch using Python command line tool and JSON recipes (no development or API usage required)
- Supports Linux Docker container images & Singularity
- Data movement:
  - Azure Files, Azure Blobs, NFS, GlusterFS
- Create and manage NFS and GlusterFS file systems



# Genomics acceleration in Azure

A revolution in genomic analysis

① Application - HAMD  
② Concurrent / embarrassingly parallel HAMD (havit++) (back up on RDMA if needed)  
③ OS → Windows Server 2012 - Cent OS 7.1 (HAMD - MSMPZ) (Intel - MPI) (HPC - PACE)

## How

A Microsoft team worked with researchers at the Broad Institute to review the algorithms in the Burrows-Wheeler Aligner (BWA) and the Genome Analysis Toolkit (GATK)

## Results

Using Microsoft's expertise in software development, they discovered how to greatly increase efficiency and speed, without compromising accuracy

## Solution

A fully-managed service on Azure that enables clinicians and researchers to focus on getting the results they need, faster and reliably

## Benefits

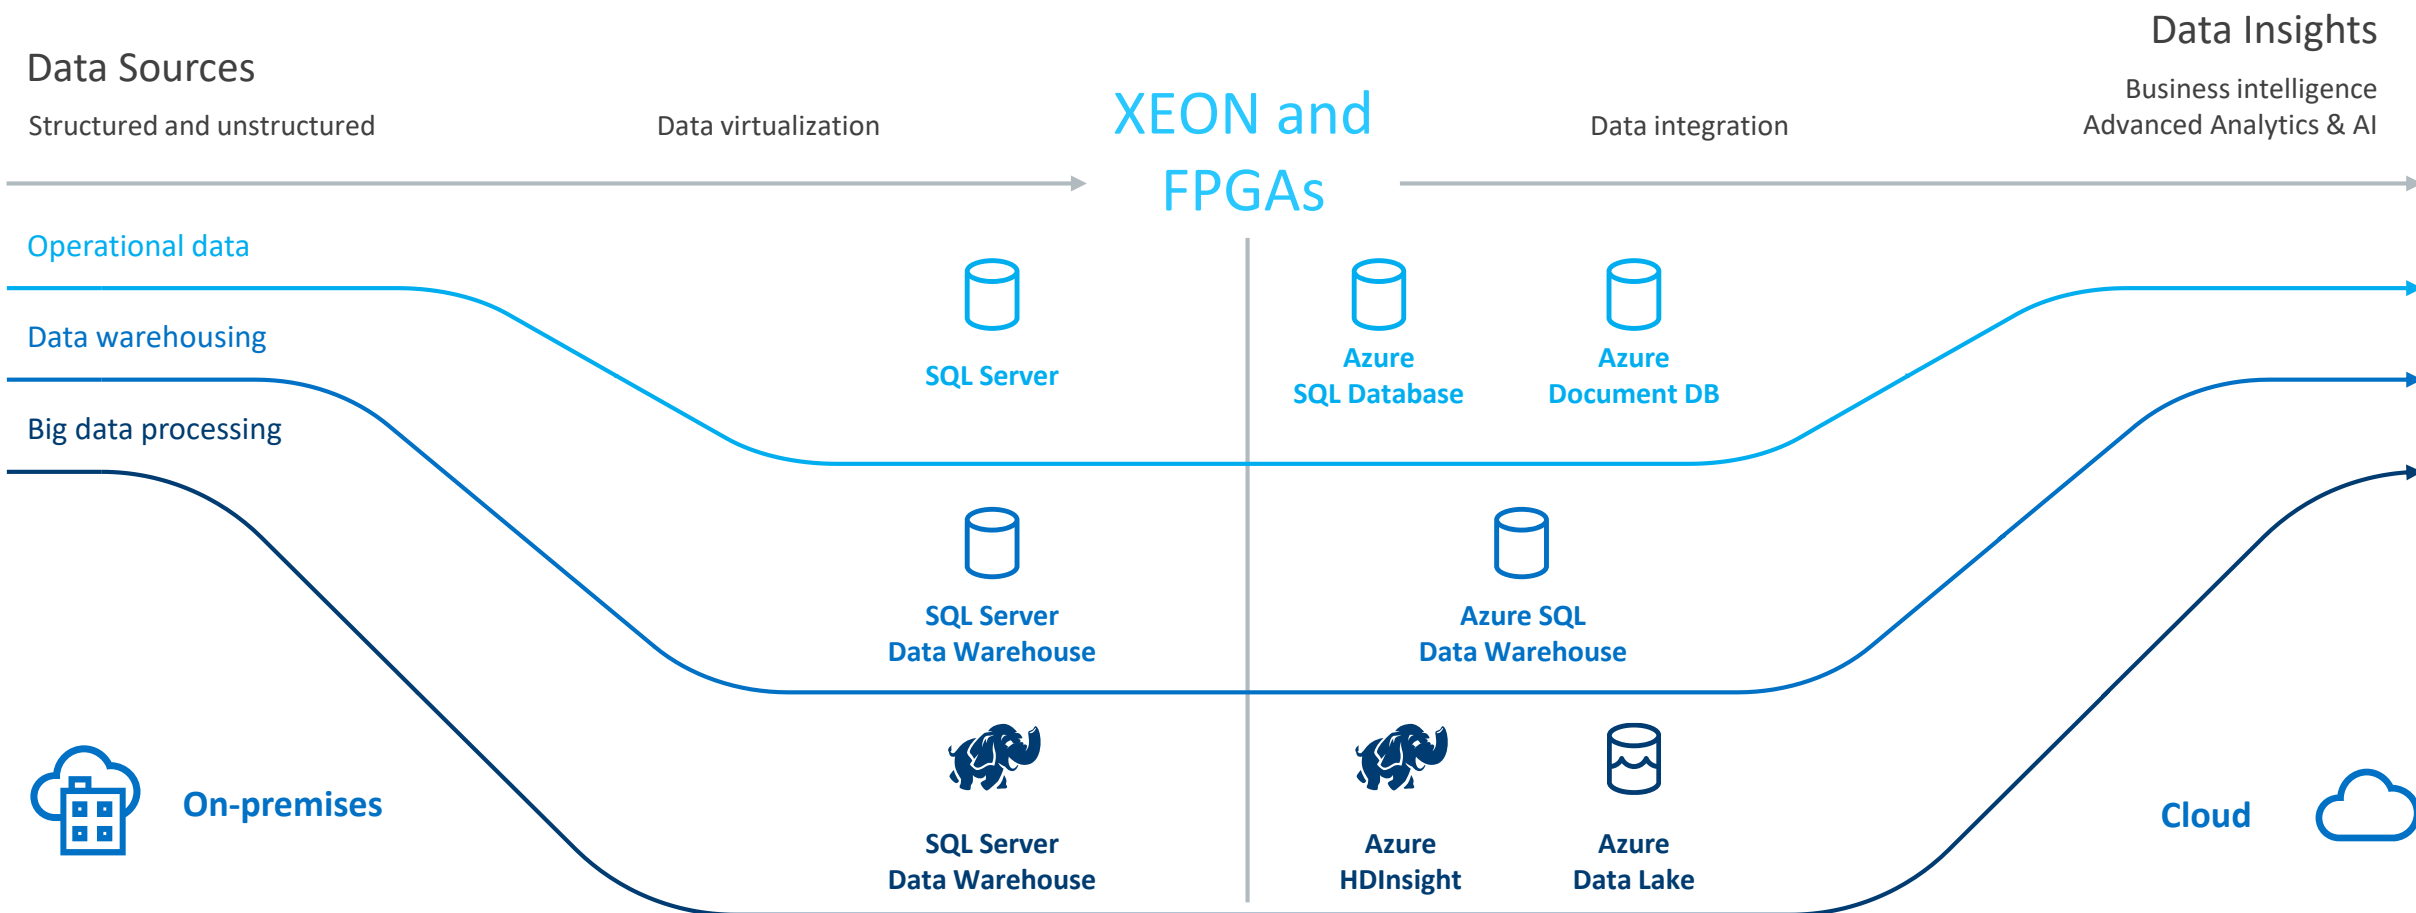
- Run BWA and GATK analysis up to seven times faster
- Run in parallel, at any scale, with a single line of code
- Leave behind the complexity of managing infrastructure

"As this type of information is used more often in the clinical setting, the emphasis on speed becomes much stronger." – Geraldine Van der Auwera, Broad Institute

⑥ Guest - Network - SSH ACCESS  
RDMA  
Phase 2 :: 12 POC  
⑦ Subscription - (Mandita)  
⑧



# Azure Data Platform optimized for Intel architecture and for Customers



# Project Brainwave

Project Brainwave is a powerful platform for an accelerated AI cloud

- **Deep-learning platform**  
Powered by Intel® 12NM Stratix 10 FPGAs
- **Record-setting performance**  
Over 130,000 compute operations per cycle



# Powerful Alliance for your Digital Transformation

## AZURE + INTEL



### Productive

Intel and Microsoft co-engineering to offer differentiated Azure services powered by the latest Intel Xeon processors



### Hybrid

Flexible and consistent hybrid cloud solutions with Intel Xeon Scalable processors, from Azure to Azure Stack



### Intelligent

Innovative AI, Data, and Analytics services optimized with Intel technologies



### Trusted

Unique Security Cloud Services enabled by Intel SGX technology

# Next Steps

---

Got some  
new ideas?

Microsoft Big Compute

<https://azure.microsoft.com/en-us/solutions/big-compute/>

Microsoft HPC

<https://azure.microsoft.com/en-us/solutions/high-performance-computing/>