

# Azure HPC

**Diversifying Your HPC Technology with Azure:  
A New Era of Scientific Computing**

*Mike Kiernan, Principle TPM Azure HPC Public Sector*

*"Together, AMD and Azure enable world class capabilities across high-performance computing, visualization, and enterprise workloads."*



# Supercomputing at Microsoft – a dual role

Significant deployment and internal usage of supercomputing infrastructure (HPC, AI, and other use cases)

Microsoft is a user/operator of HPC (like a HPC center)

Microsoft makes some of our HPC & AI capabilities available to customers as Azure HPC & AI (“cloud”)

Microsoft is also a provider of HPC to others (like most HPC centers)




# Microsoft Azure has 6x supercomputers in the Top50 of Top500 [so far...]

First permanent  
genuine supercomputers  
to be listed by a  
hyperscale company

- 11 **Explorer-WUS3** - ND96\_amsr\_MI200\_v4, AMD EPYC 7V12 48C 2.45GHz, AMD Instinct MI250X, Infiniband HDR, Microsoft Azure West US3 United States
- 16 **Voyager-EUS2** - ND96amsr\_A100\_v4, AMD EPYC 7V12 48C 2.45GHz, NVIDIA A100 80GB, Mellanox HDR Infiniband, Microsoft Azure Azure East US 2 United States
- 42 **Pioneer-EUS** - NDv4 cluster, AMD EPYC 7V12 48C 2.45GHz, NVIDIA A100, Infiniband HDR, Microsoft Azure Azure East US United States
- 43 **Pioneer-SCUS** - NDv4 cluster, AMD EPYC 7V12 48C 2.45GHz, NVIDIA A100, Infiniband HDR, Microsoft Azure Azure South Central US United States
- 44 **Pioneer-WEU** - NDv4 cluster, AMD EPYC 7V12 48C 2.45GHz, NVIDIA A100, Infiniband HDR, Microsoft Azure Azure West Europe Netherlands
- 45 **Pioneer-WUS2** - NDv4 cluster, AMD EPYC 7V12 48C 2.45GHz, NVIDIA A100, Infiniband HDR, Microsoft Azure Azure West US 2 United States


# The HPC Cloud Opportunity

## On Premises



 Compute Needs  
 Available Capacity

## Cloud



 Compute Needs  
 Available Capacity

Cloud elasticity provides a compute-on-demand model that matches IT consumption directly with organizational needs, thereby optimizing overall IT spend.

# Don't look up!


## Compelling Events (Examples)

- Need...GPU's!!
- Innovation agenda
- Science-to-operations acceleration
- Data/Archive exploding out of control
- Network bandwidth / Collaboration
- Long user wait times
- Out of (clean) power
- DC infrastructure ageing out
- Procurement fatigue
- Terrible experience with supplier
- Great cloud framework deal
- Staff retention
- Cyber threats
- Risk transfer
- Disaster Recovery
- <...>



Don't just run HPC in the cloud.

## Use the cloud built for HPC



### Purpose-built HPC

A full range of CPU and GPU capabilities that help applications scale to 80K+ cores




### High Performing Storage

A range of storage capabilities to support simple-to-complex storage needs



### Fast, Secure Networking

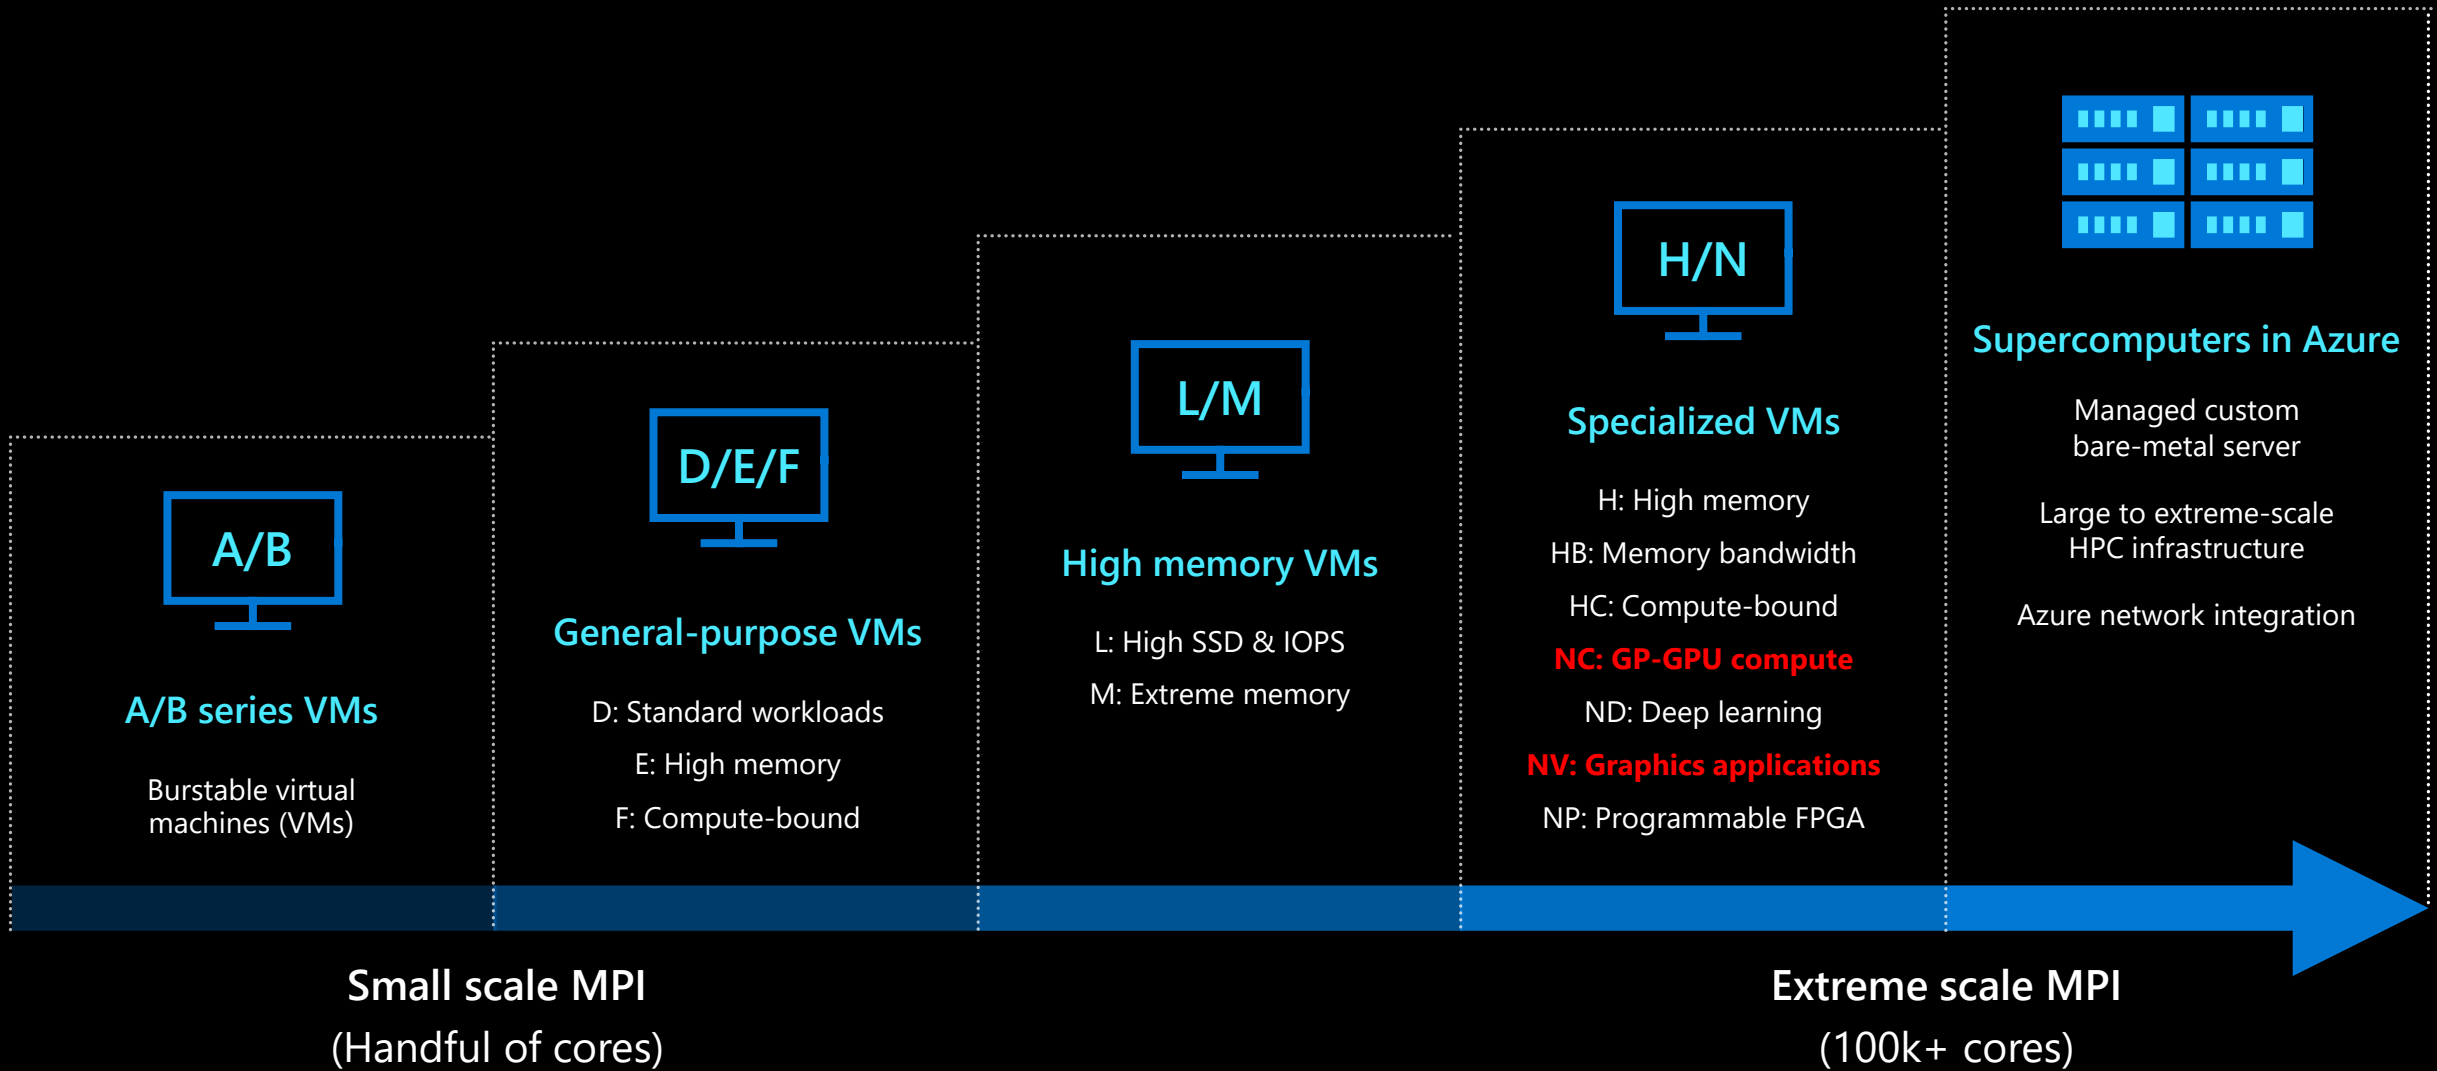
Fast InfiniBand inter-connects as well as edge-to-cloud connectivity



### Workload Orchestration

End-to-end workflow agility using known, familiar tools & processes

# Solve any HPC, AI workload—at any scale



# Azure H-series for HPC



	<b>HX</b>	<b>HBv4</b>	<b>HBv3</b>	<b>HBv2</b>	<b>HC</b>
<b>Processor</b>	176 cores AMD Genoa-X	176 cores AMD Genoa-X	120 cores AMD Milan-X	120 cores AMD Rome	44 cores Intel Skylake
<b>Memory</b>	1.4 TB	700 GB DDR5	448 GB DDR4	448 GB DDR4	352 GB DDR4
<b>DRAM Bandwidth</b>	780 GB/s	780 GB/s	350 GB/s	350 GB/s	190 GB/s
<b>InfiniBand</b>	400 Gb/s	400 Gb/s	200 Gb/s	200 Gb/s	100 Gb/s
<b>SSD</b>	3.6 TB NVMe	3.6 TB NVMe	1.8 TB NVMe	900 GB NVMe	700 GB SSD
<b>Availability</b>	Preview	Preview	Available	Available	Available

(+Azure Boost)



# Azure HBv4 Series Specifications (+Azure Boost)



Perfect for applications with...



- ✓ High memory bandwidth
- ✓ High CPU clock speed
- ✓ High per-core licensing costs
- ✓ High single-core performance



Powered by AMD EPYC™ 9V33X with 3D V-Cache



400 Gb/s NVIDIA Quantum-2 InfiniBand




Virtual NUMA mapped to physical NUMA topology



Physical cores only (no simultaneous multi-threading)

VM Name	HB176rs_v4	HB176-144rs_v4	HB176-96rs_v4	HB176-48rs_v4	HB176-24rs_v4
Size	176 CPU cores	144 CPU cores	96 CPU cores	48 CPU cores	24 CPU cores
Peak CPU Frequency*	3.7 GHz				
RAM per VM	700 GB				
RAM per core	4.3 GB	5.3 GB	8 GB	16 GB	32 GB
Memory B/W per VM	DRAM: 780 GB/s   3D V-Cache (L3): 5.7 TB/s   Effective Blended Average: 1.2 TB/s				
Memory B/W per core	6.8 GB/s	8.3 GB/s	12.5 GB/s	25 GB/s	50 GB/s
L3 Cache per VM	2304 MB				
L3 Cache per core	13 MB	16 MB	24 MB	48 MB	96 MB
SSD Perf per VM	2 * 1.8 TB NVMe – 12 GB/s (Read) / 7 GB/s (Write)				



Highest Perf per VM



Highest Perf per Core

# Azure HX Series Specifications



## Perfect for applications with...

- ✓ High CPU clock speed
- ✓ High memory-to-CPU ratio
- ✓ High per-core licensing costs
- ✓ High single-core performance



Powered by AMD EPYC™ 9V33X with 3D V-Cache



400 Gb/s NVIDIA Quantum-2 InfiniBand



Virtual NUMA mapped to physical NUMA topology



Physical cores only (no simultaneous multi-threading)

VM Name	HX176rs	HX176-144rs	HX176-96rs	HX176-48rs	HX176-24rs
Size	176 CPU cores	144 CPU cores	96 CPU cores	48 CPU cores	24 CPU cores
Peak CPU Frequency*	3.7 GHz				
RAM <b>per VM</b>	1.4 TB				
RAM per core	8 GB	10 GB	15 GB	29 GB	59 GB
Memory B/W <b>per VM</b>	DRAM: 780 GB/s   3D V-Cache (L3): 5.7 TB/s		Effective Blended Average: 1.2 TB/s		
Memory B/W per core	6.8 GB/s	8.3 GB/s	12.5 GB/s	25 GB/s	50 GB/s
L3 Cache <b>per VM</b>	2304 MB				
L3 Cache per core	13 MB	16 MB	24 MB	48 MB	96 MB
SSD Perf <b>per VM</b>	2 * 1.8 TB NVMe – 12 GB/s (Read) / 7 GB/s (Write)				

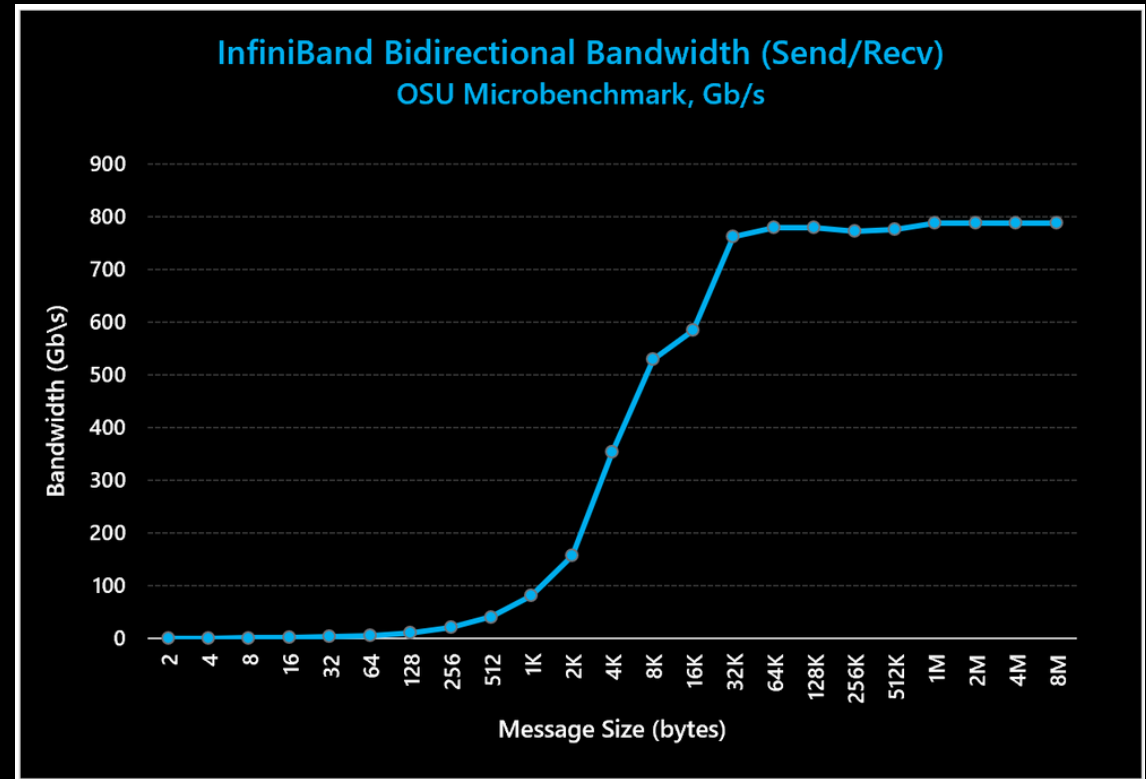



Highest Perf per VM



Highest Perf per Core

# Benchmarks



[Performance & Scalability of HBv4 and HX-Series VMs with Genoa-X CPUs \(microsoft.com\)](https://microsoft.com)

# Azure HPC storage options

	<b>Roll your own (RYO)</b> Centos/Ubuntu/Debian	<b>Azure NetApp Files</b> Third-party hardware First-party interface	<b>Azure HPC Cache</b> Combination file cache and Blob-as-POSIX filer	<b>ClusterStor</b> Lustre-based, fully managed HPC environment
<b>Usage</b>	Fast, inexpensive network file system (NFS) server with no extra features	File sharing, databases, analytics, and enterprise applications	Edge caching (across WANs) Access to on-premises NAS data Portable operating system (POSIX) with massive data (Blob)	Lustre-based, single-tenant, bare-metal, and fully managed HPC environment in Azure
<b>Best for</b>	Small number of clients reading and writing	Read/write-many (unique) devices Write-once, read-once devices	Write-once, read-many devices Large number of clients (HPC) Diverse data locations (Global Namespace)	Environments running several applications with large, sequential input/output (I/O) workloads
<b>Fault tolerance</b>	RYO active/standby	Local redundancy	Distributed cache for high-availability (HA) access	HA server pair; 2U24 drive enclosure
<b>Performance</b>	A single VM	Up to 6.5 GB/s read/write (100TB volume)	2GB/s, 4GB/s, or 8GB/s throughput	Up to 460 GB/s throughput
<b>Access</b>	NFS and/or SAMBA	NFS or server message block (SMB) (mp mount Summer 2020)	NFS (SMB Summer 2020)	Lustre 2.7 file system and supported enhancements
<b>Total cost of ownership</b>	RYO	First-party service	First-party service	First-party service

# Azure Managed Lustre

Turbocharge your filesystem performance with a managed solution for HPC and AI workloads.

Azure Managed Lustre delivers a highly performant file system for HPC and AI workloads anytime and anywhere; stress-free in a pay-as-you-go model, simplifying operations and keeping costs low. Unlock the potential of the world's most popular parallel filesystem for your business.




## Accelerate HPC workloads

Provides a high-performance distributed parallel file system solution delivering hundreds of GB/s storage bandwidth and solid-state disk latency, with several performance options.



## Protect application investments

Enables HPC applications in the cloud without breaking application compatibility by providing familiar Lustre parallel file system functionality, behaviors and performance, securing long-term application investments.



## Azure Blob integration

Connect file systems with Azure storage containers for high - performance processing of blob data and archiving processed data into tiered blob storage for optimal data placement and cost management.



## Managed solution

Enables the benefits of a Lustre parallel file system for HPC workloads with the expected agility of cloud services, available in most Azure regions and without management headaches.

# Software ecosystem

## Orchestration

- **Azure CycleCloud**
- VM Scalesets
- Job schedulers



## Runtimes and libraries

- MPI implementations
- Development: Compilers, debugging
- Numerical and I/O libraries
- AI Frameworks



## Marketplace

- Operating system (OS) images
- Container images
- VM extensions
- Parallel file systems templates

## Applications

- Installation recipes
- Performance tuning and benchmarking

# Managed supercomputing service for weather and climate

## Packages/Projects

A: Hosting

B: Supercomputers

C: Networking

D: Storage/MASS

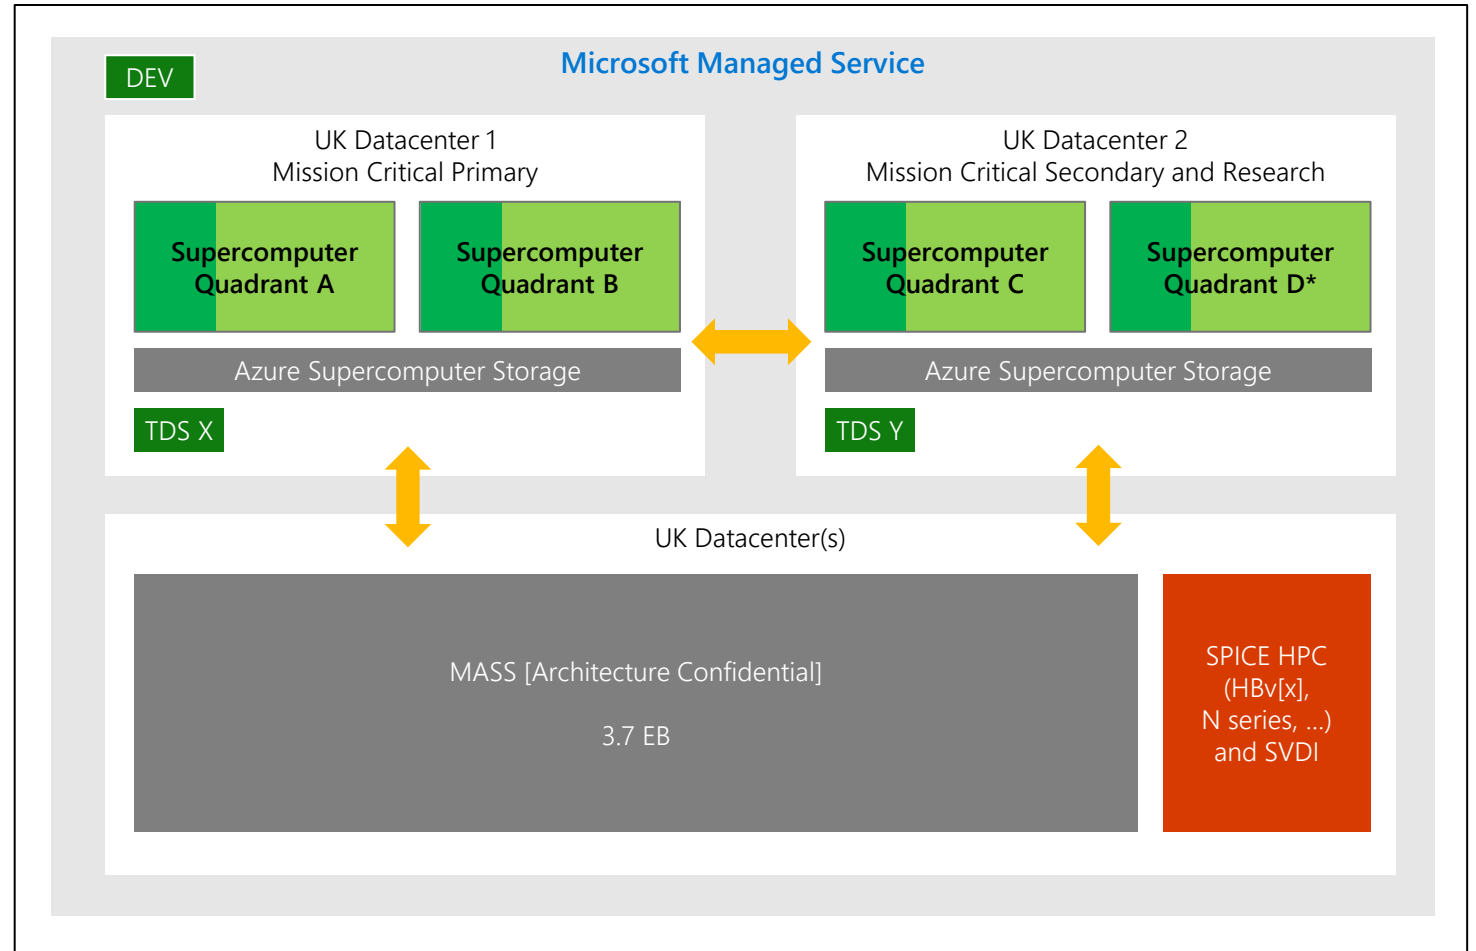
E: SPICE (SVDI & interactive HPC)

U: Umbrella (managed service)

### Quadrants

X nodes of Milan to give total of 1.1x current  
Add Y nodes of Genoa (4.9x) to give full 6x

**£1bn, 10 years, transformational**



~2660 nodes per quadrant, > 10,000 nodes total, ~1.8M AMD EPYC CPU cores, >60PF

# Take Aways



- Azure HPC can deliver to almost any scale and requirement
- Cloud HPC could complement what you do
- Look for compelling events, or try it anyway
- Focus on outcomes in your requirements
- Please do ask us to tender – chances are we will anyway 😊



