Microsoft

Microsoft Azure

Microsoft

Evolving inter-connect fabric and SDN in the cloud to meet the needs of AI workloads

Deepak Bansal
CVP and GM
Azure Networking

AI is a game changer for more and more applications

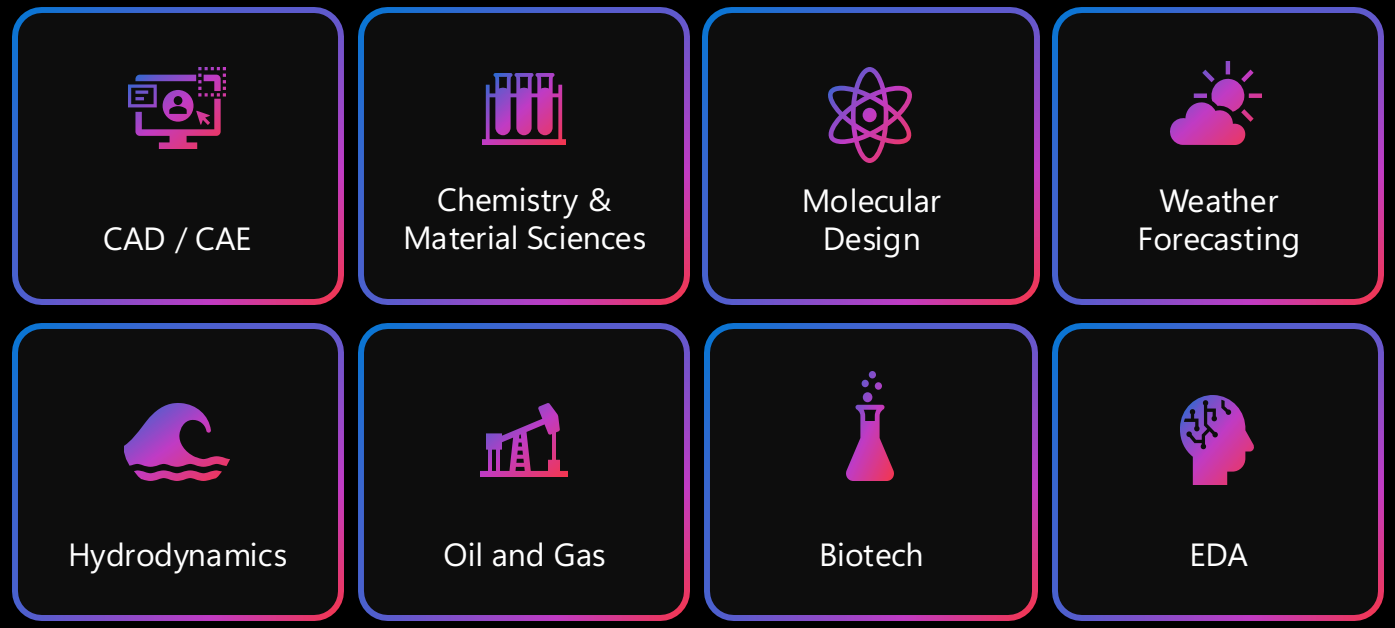

AI is the Heart of a Larger Cloud System

"Existing Apps" are driving the exponential growth of AI

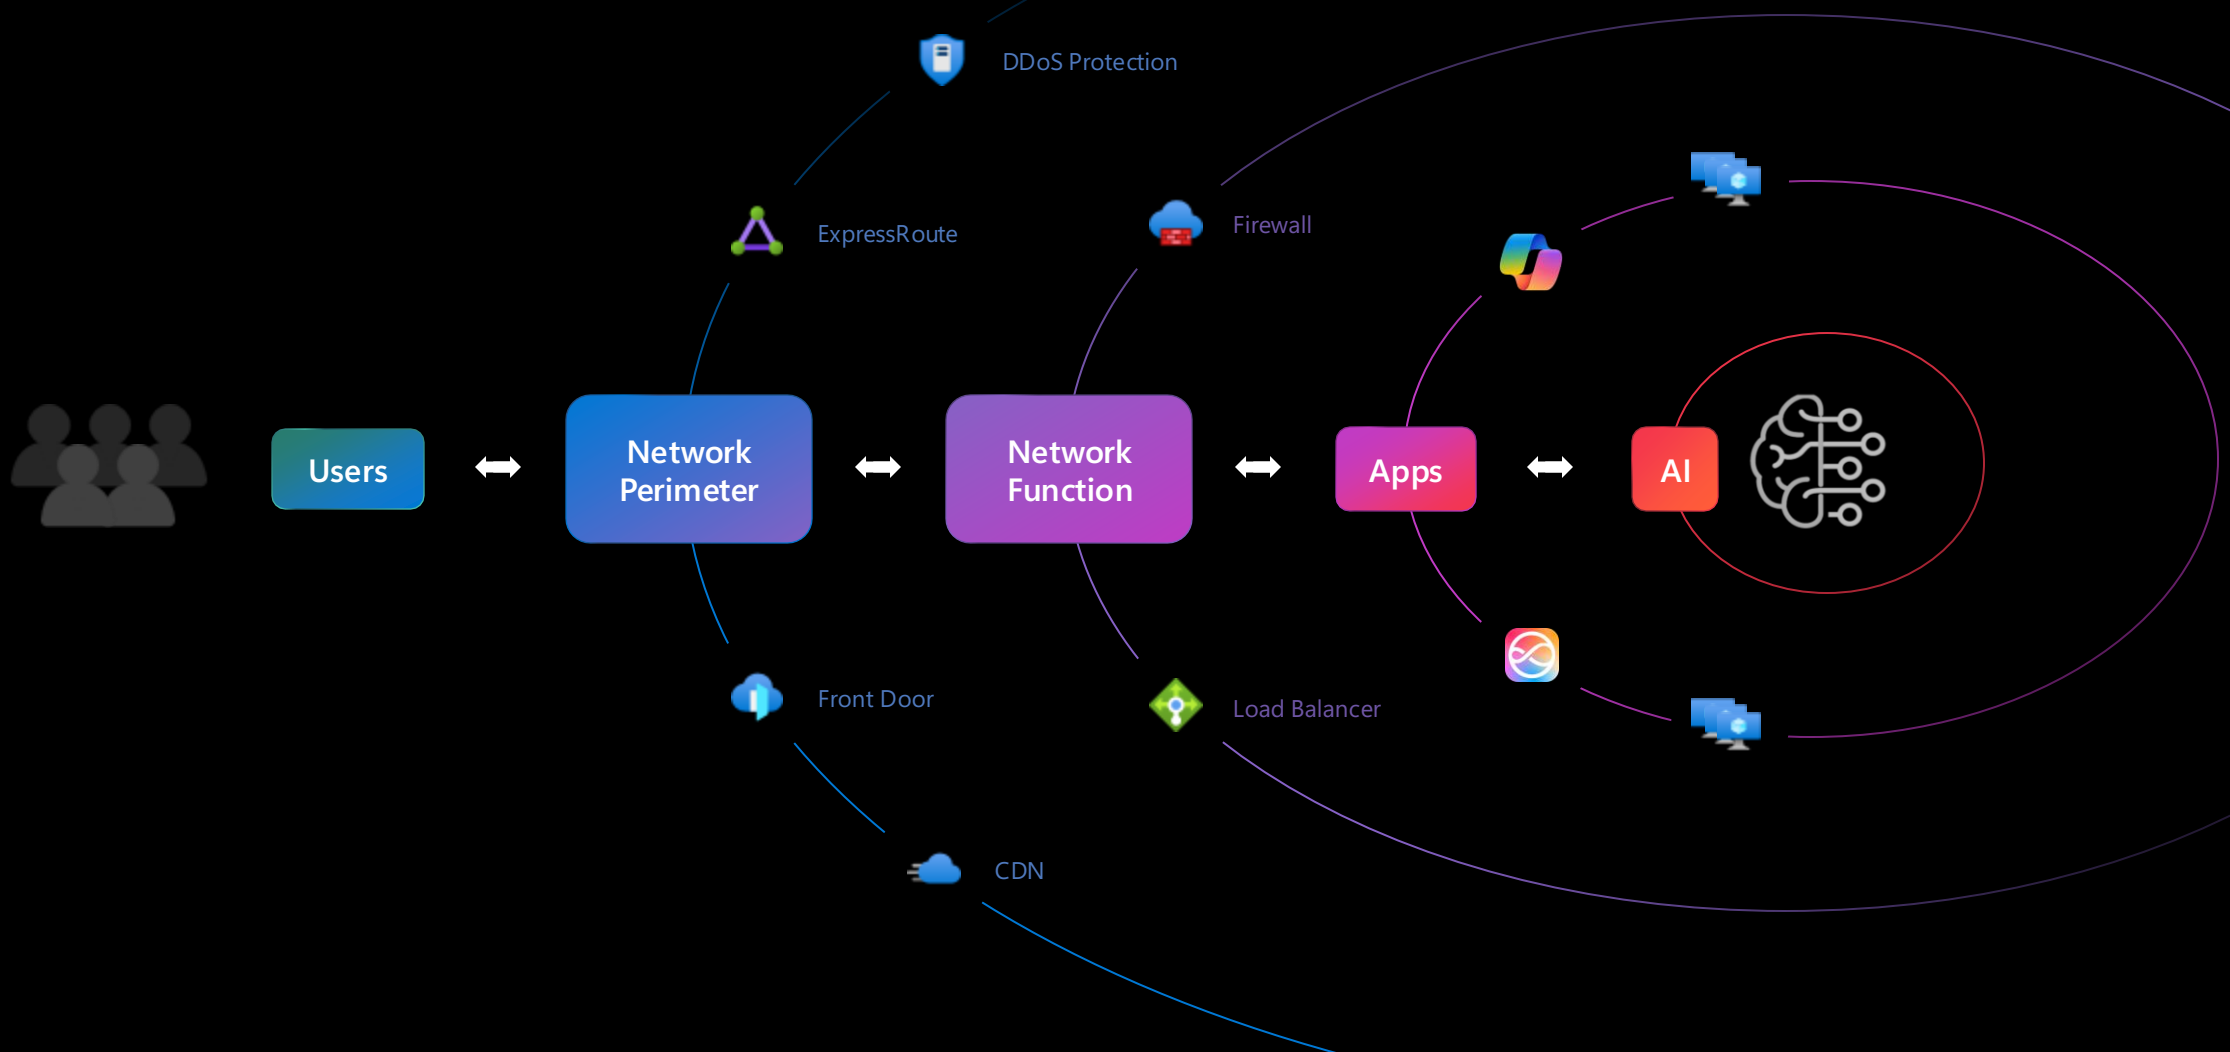

AI cloud providers must provide

Private, secure and reliable infrastructure to connect users to AI enabled apps

Flexibility to match the right AI technologies to already established or newly created apps

Ability to scale with insatiable demand both regionally and globally

Requirements

High bandwidth and **reliable networks**

Low latency to connect users to AI enabled apps

Security to safeguard data and identity

Hollow core fiber (HFC)

47% improvement in speed

Innovative
hollow core fiber (HCF)

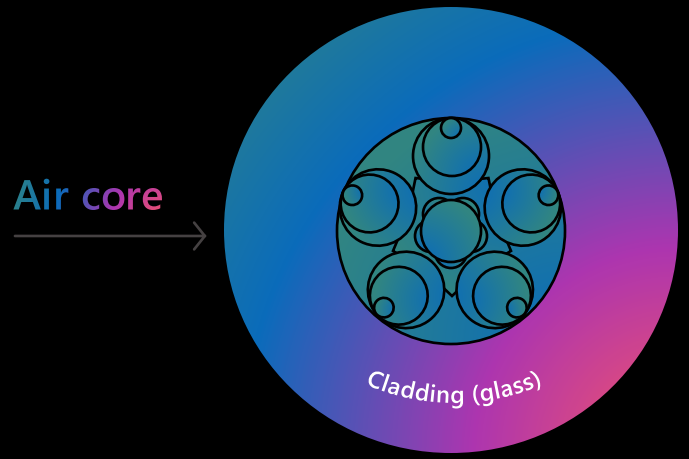

1.5X Faster

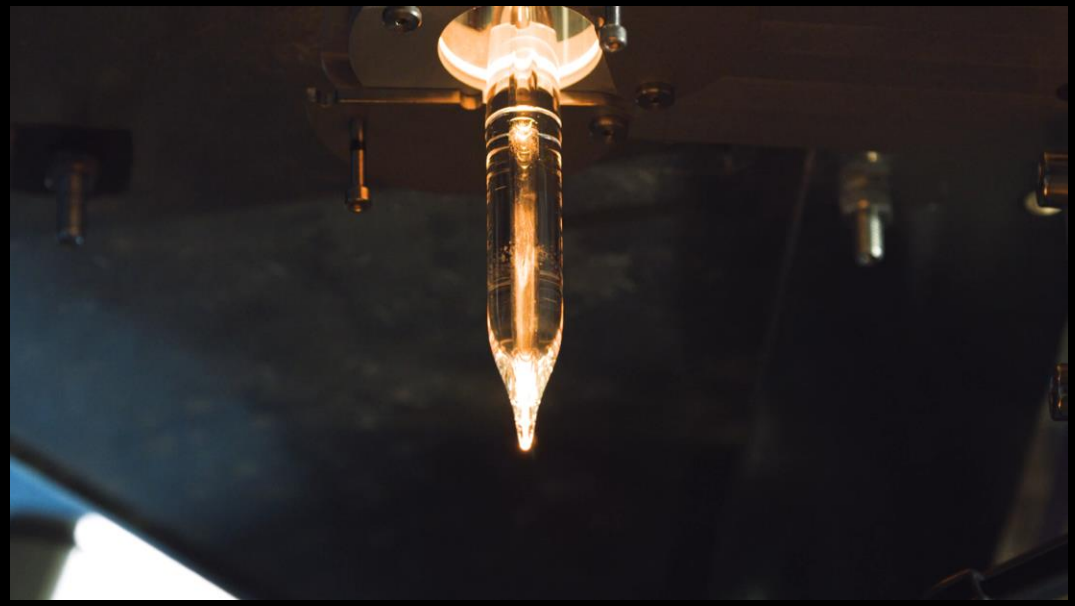

HCF in fab facility

Massive bandwidth required to interconnect storage, apps and AI

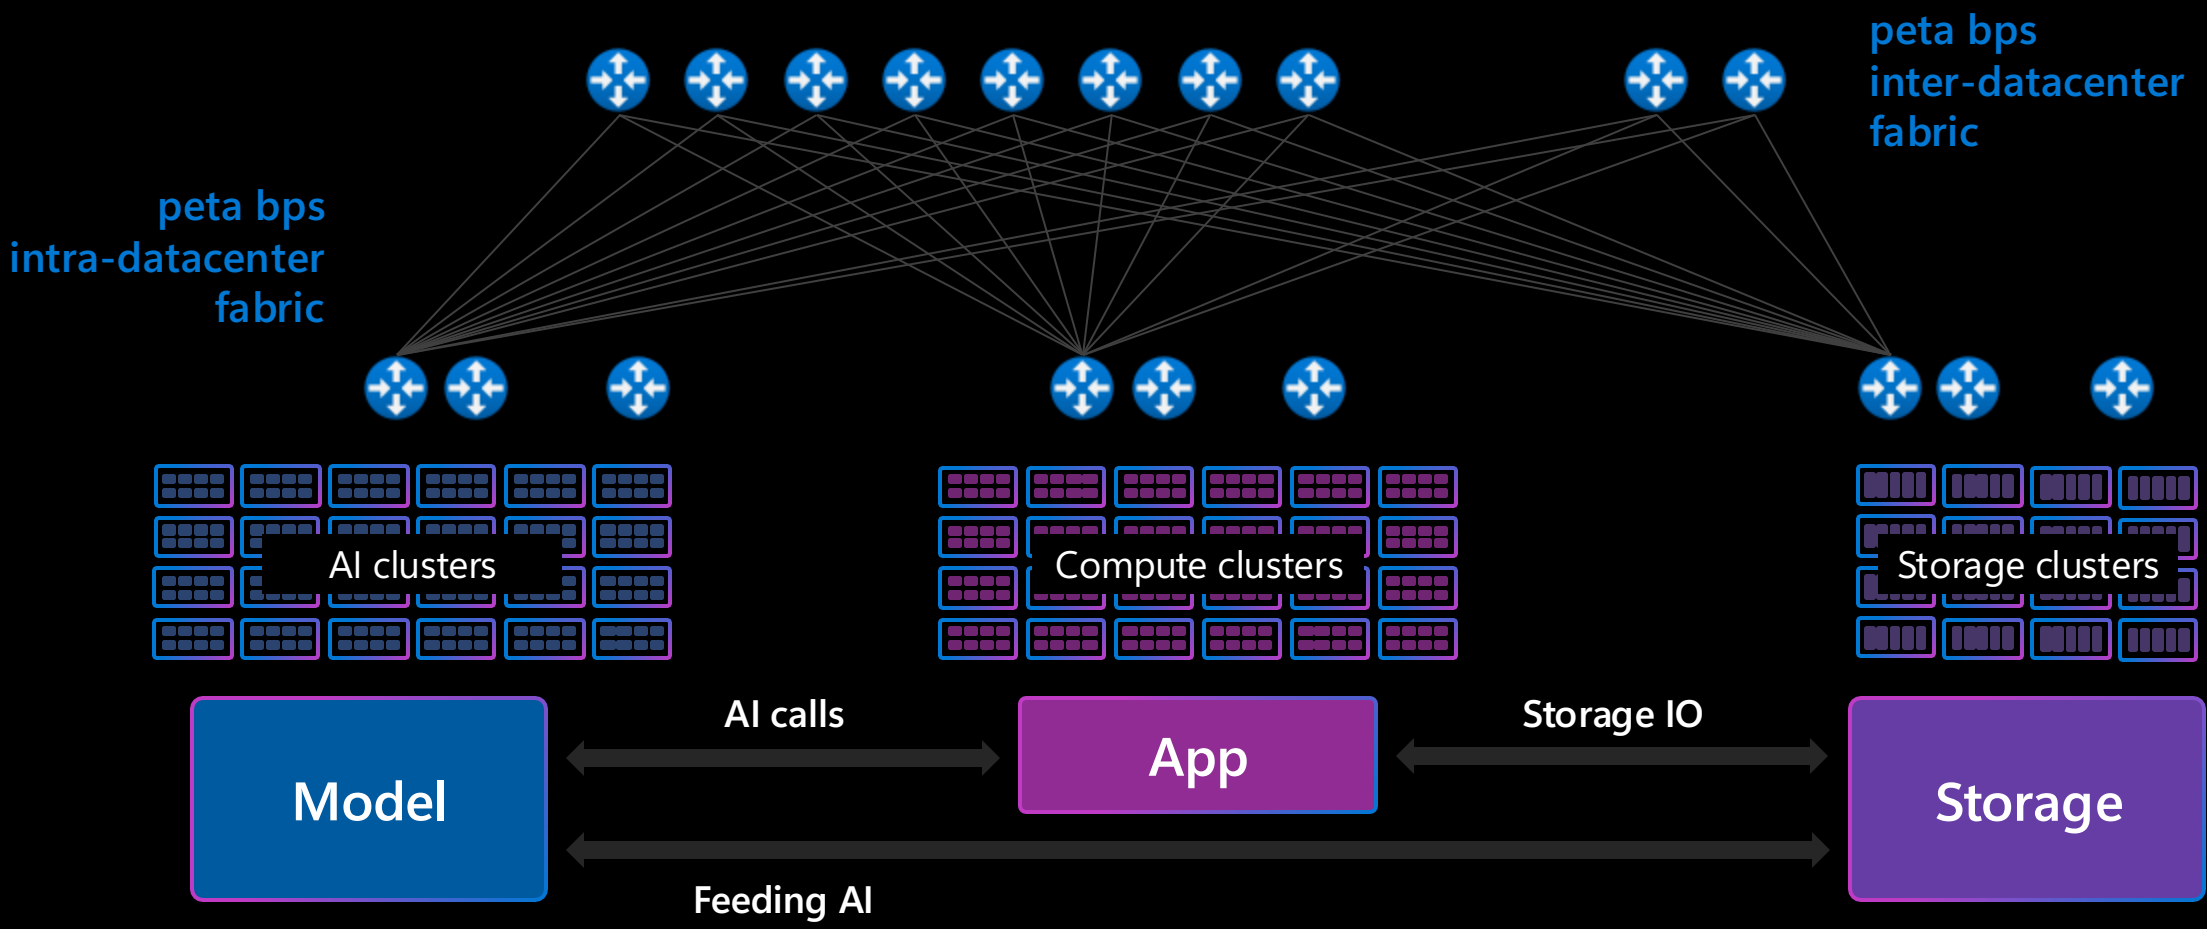

AI Backend Network in the Cloud

Raising the Bar for Hyperscale Datacenter Networks

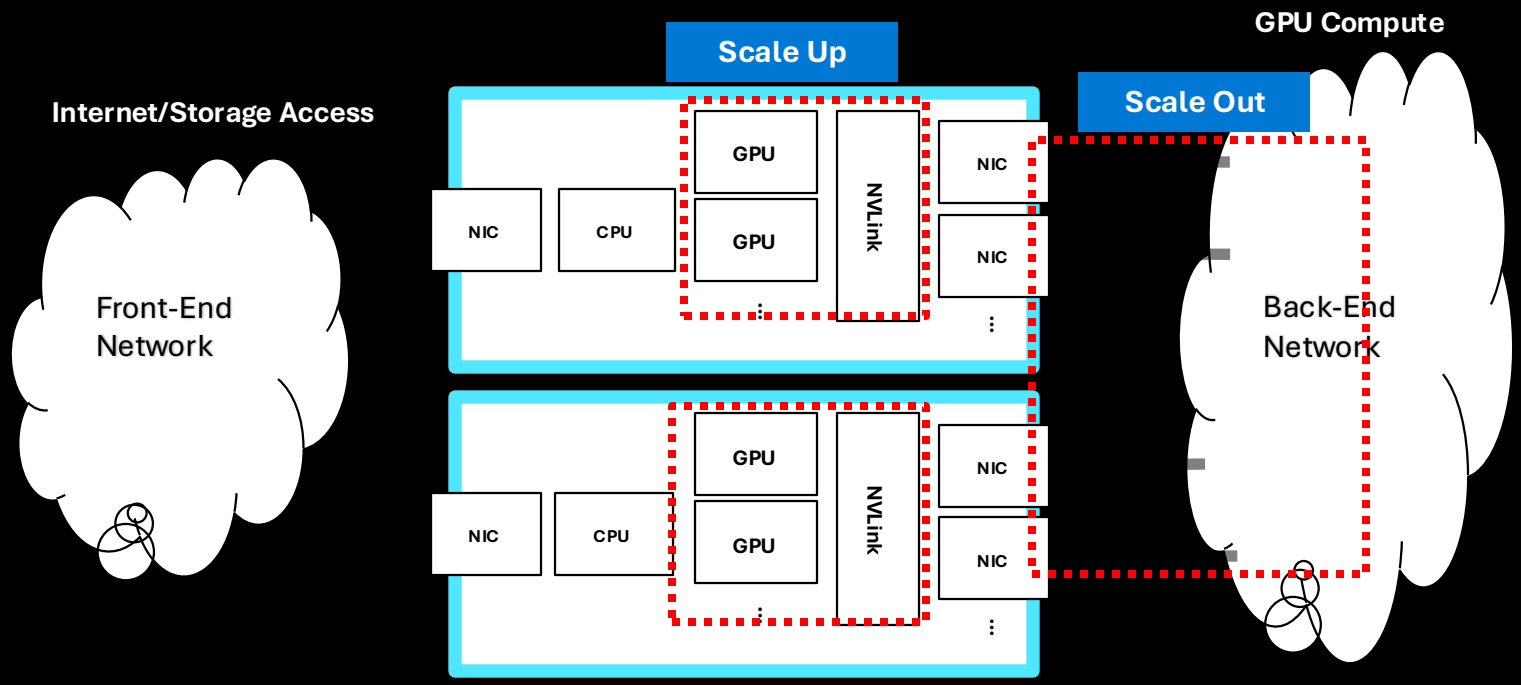

- Long lasting and large flows of training data
- Large bursts of data sent synchronously
- Long training time demands reliable networks
- Retries of failed jobs increased costs → Efficient traffic management, monitoring and visibility
- AI applications need fast processing and response → Lossless traffic with low latency
- Traffic using RoCEv2 has low entropy for ECMP → Traffic engineering technology for AI backend network

Backend Network Key Technology Trend

- Key Technologies Supported by SONiC for AI Networking
  - **Segment Routing over IPv6 (SRv6)**: enabling AI application control of network path
  - **Packet Spray**: enable packet level enhanced **load balancing** cross large-scale paths and fault tolerance
  - **Packet Trimming**: for **fast retransmission** and **congestion control**
  - **High Frequency Telemetry**: microsecond level telemetry for **fast detection** & diagnostics

SRv6

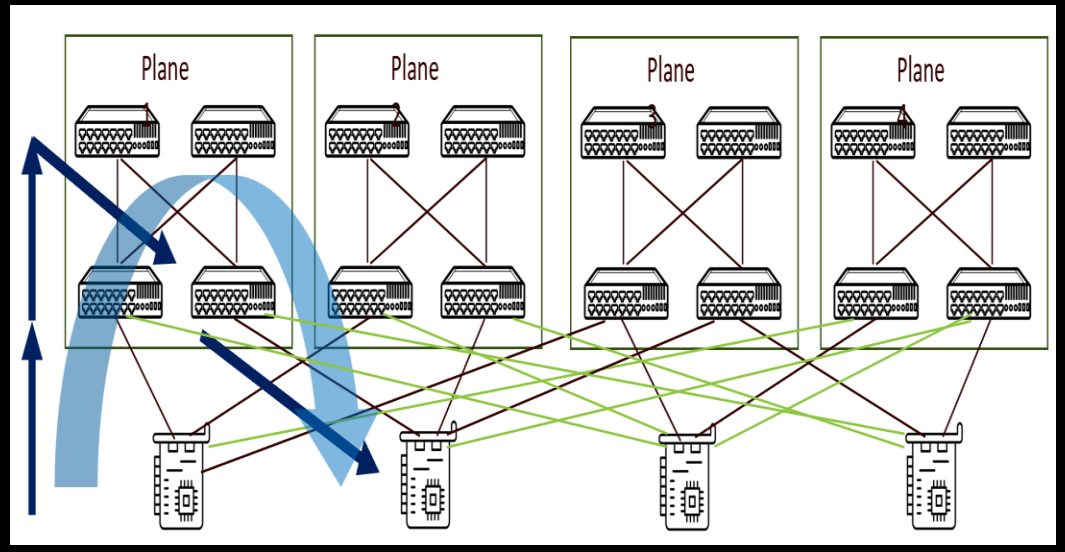

Packet Trimming

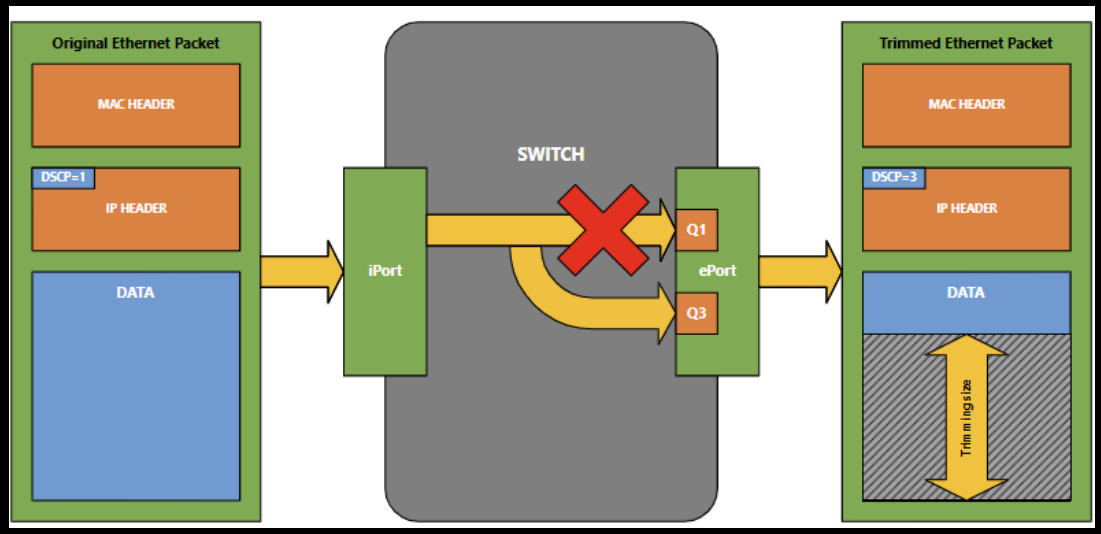

Some clusters now contain more than 100k GPUs

Reliable infrastructure is key to avoid re-runs

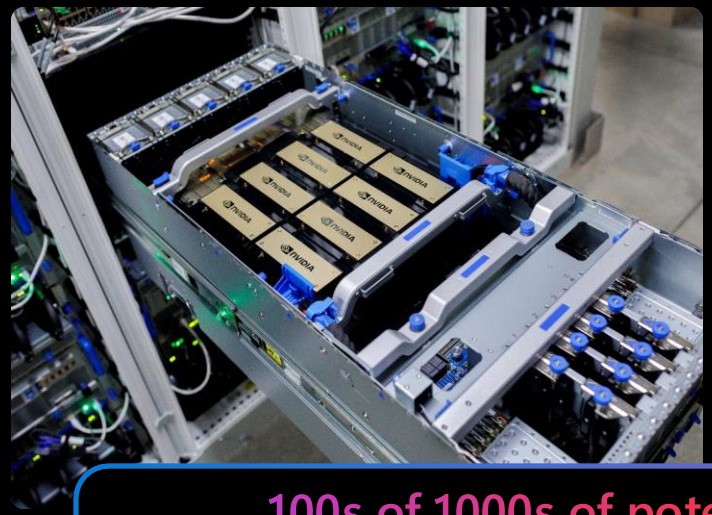

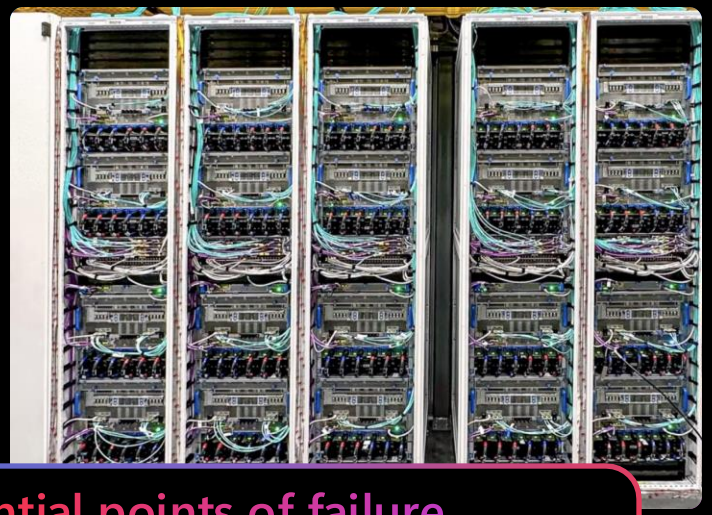

100s of 1000s of potential points of failure

Instability from inadequate cooling addressed with OCP Chassis and Rack design

Defect free builds achieved with high levels of automation before commissioning

Ethernet/InfiniBand interconnect addressed with defect detection and streaming telemetry

Huge deployments require automated health monitoring, ticketing and mitigation

Backend Networking

Performance and Defect Streaming

Cloud

AI Based Monitoring and Defect Analysis

Automated Mitigation, Ticketing and Resolution

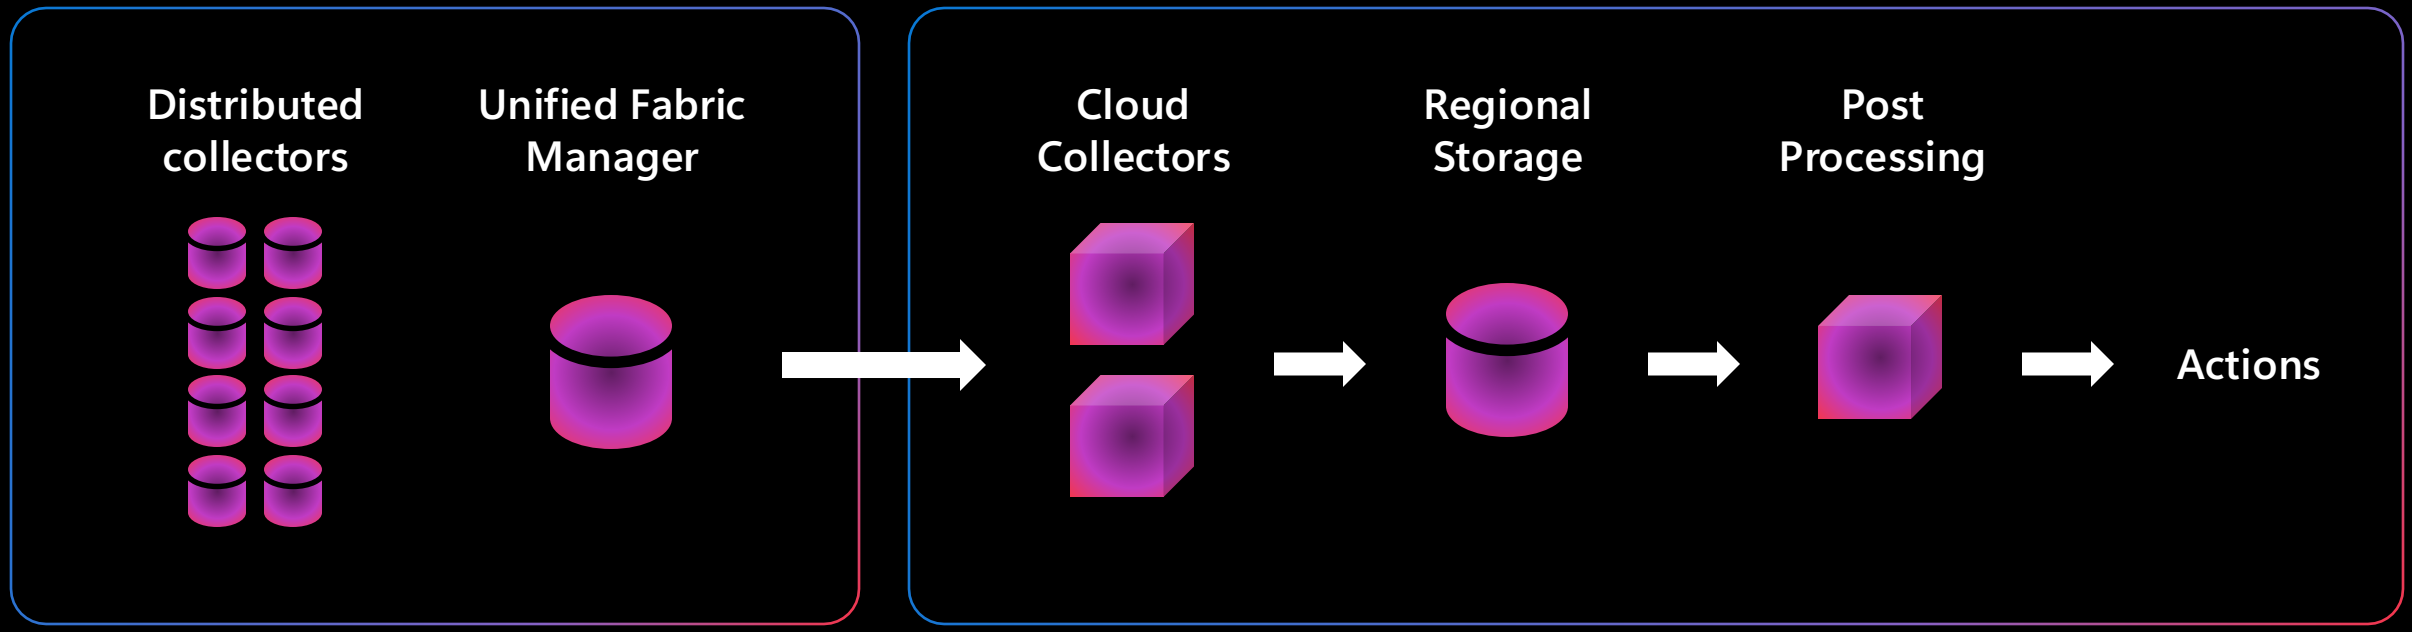

Training and inferencing have difference performance requirements

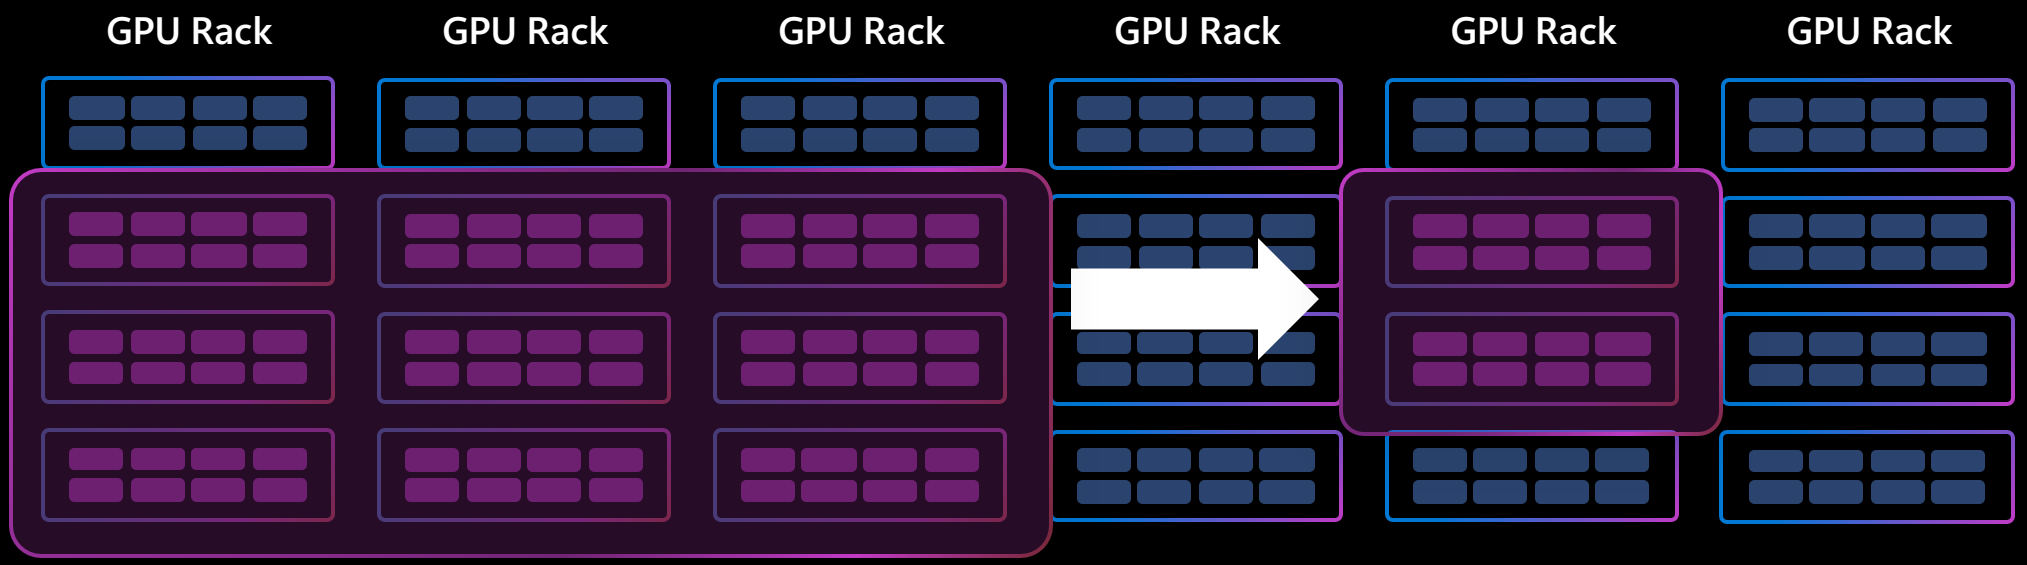

Scale-out arbitrary amounts of AI GPU VMs to train your model

Batch, share, and overlay jobs to optimize cost vs performance

Scale-in AI tech for the lighter needs of inferencing

Regional or global presence is key for some customers

60+ AI regions | **275k+** Miles of Fiber | **2Pbps+** WAN Capacity | **200T** Peering Capacity | **40k+** Peering Connections

Global schedulers offer AI everywhere you need it

Transparent preempt, scale up, scale down, run

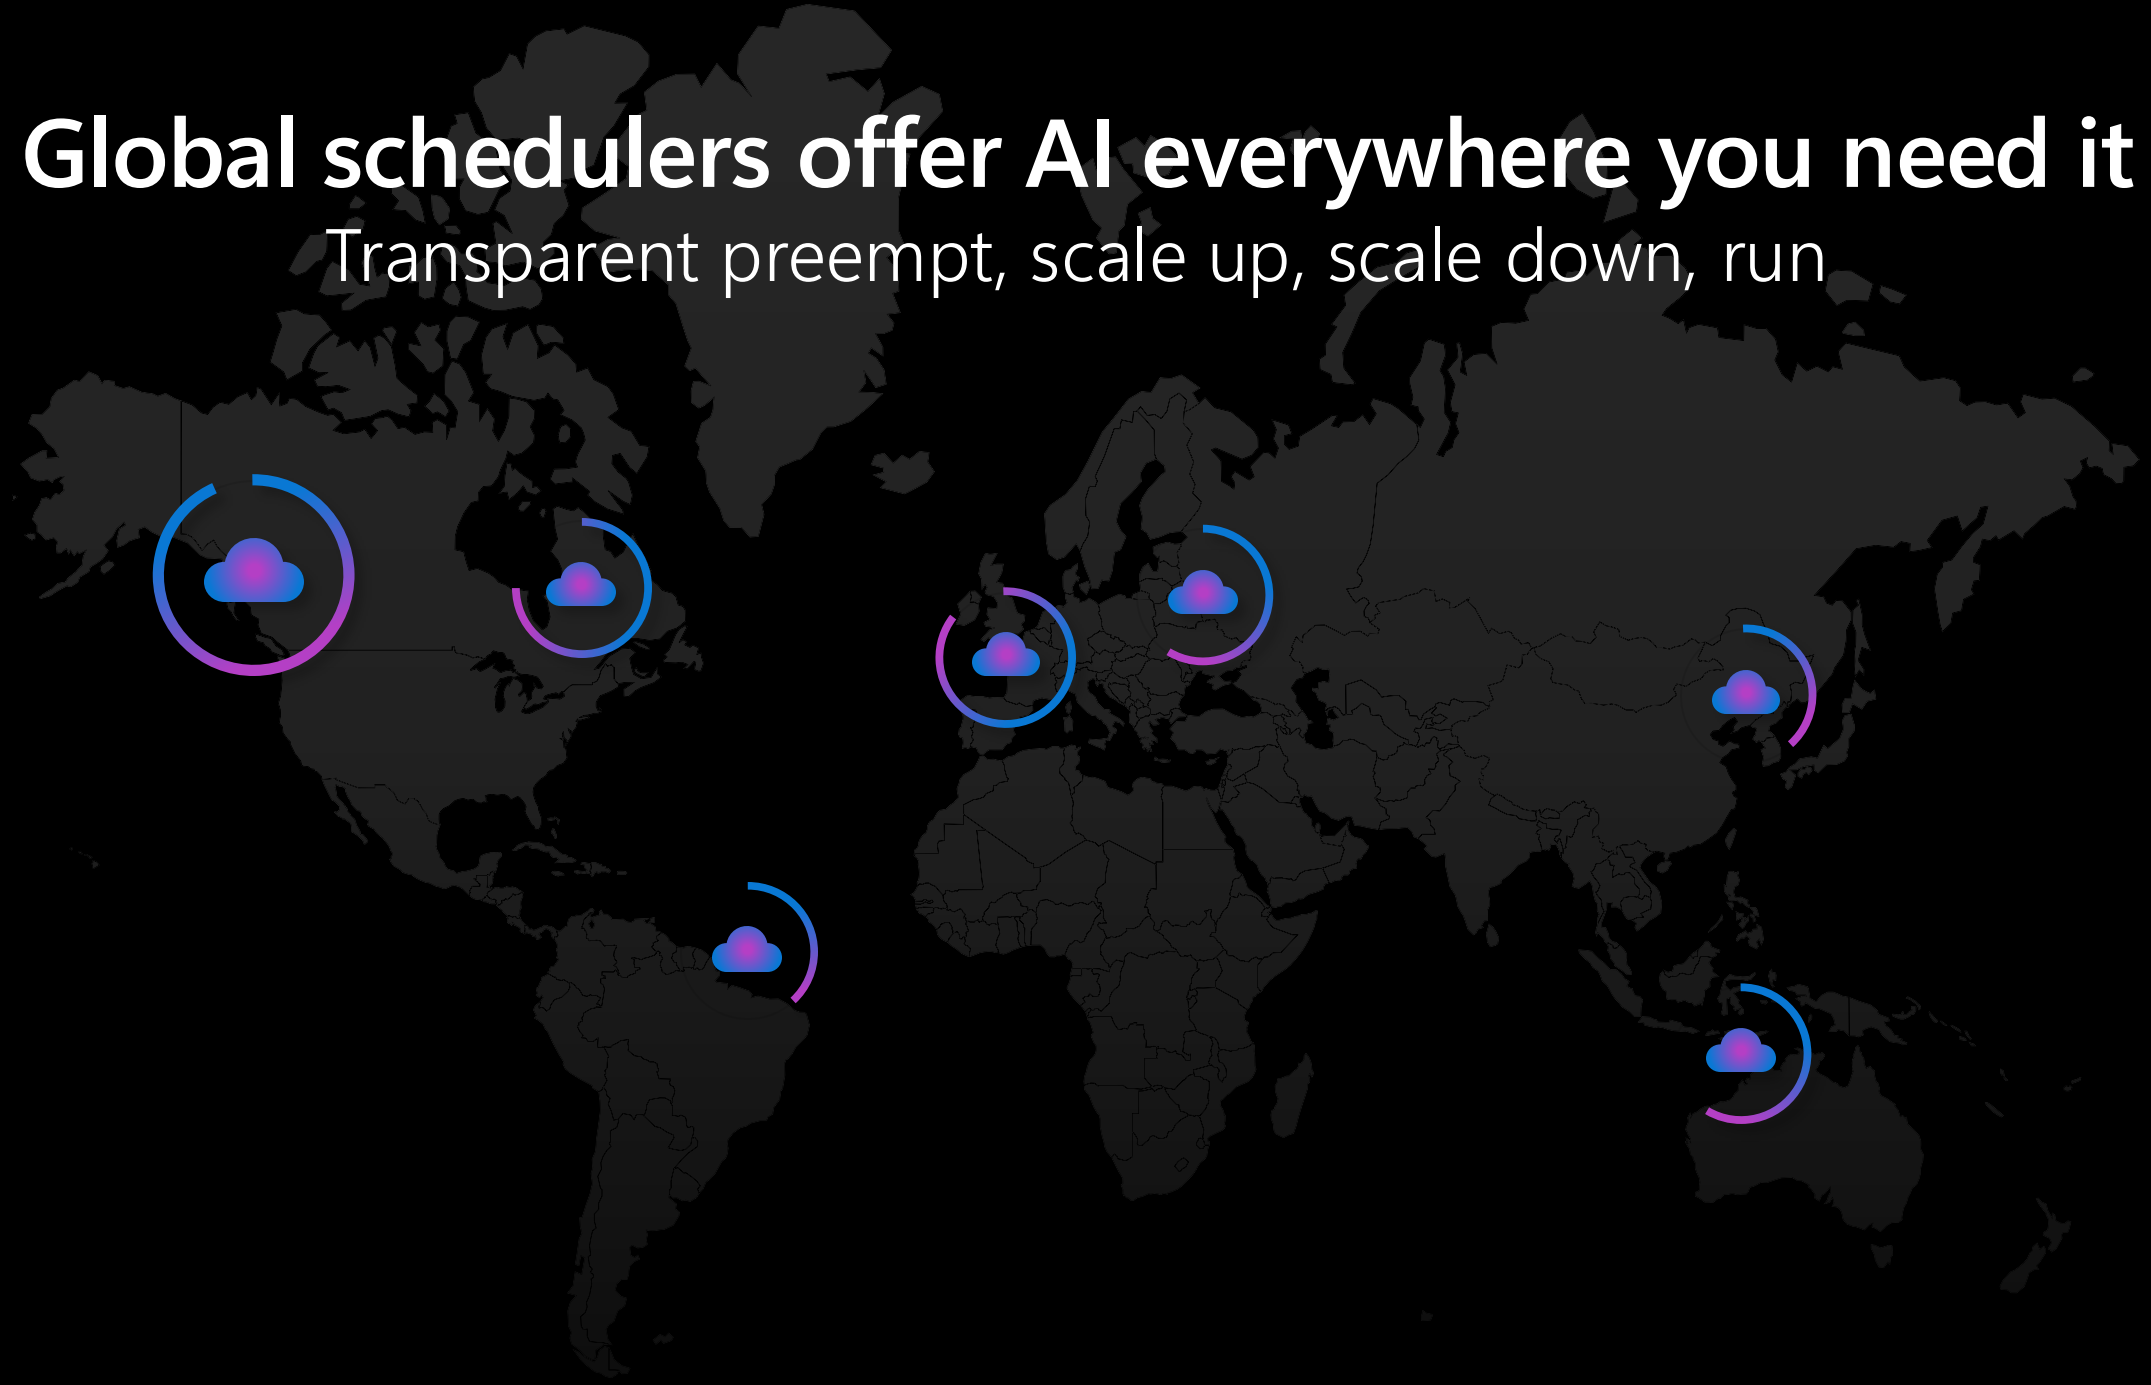

AI Appliance: Integrating distributed GPUs to Azure

- **Secure Access to Azure Storage**
  FastPath Private Link to Azure Storage over WAN
- **Extreme Network Scale & Perf**
  12.8Tbps+
- **Horizontally Scale DPUs**
  12.8 Tbps+ via ECMP
- **Lowers Training Time**
  Model load times significantly lowered
  Training checkpoint time reductions
- **Increased Efficencies**
  Run Azure Services Efficiently

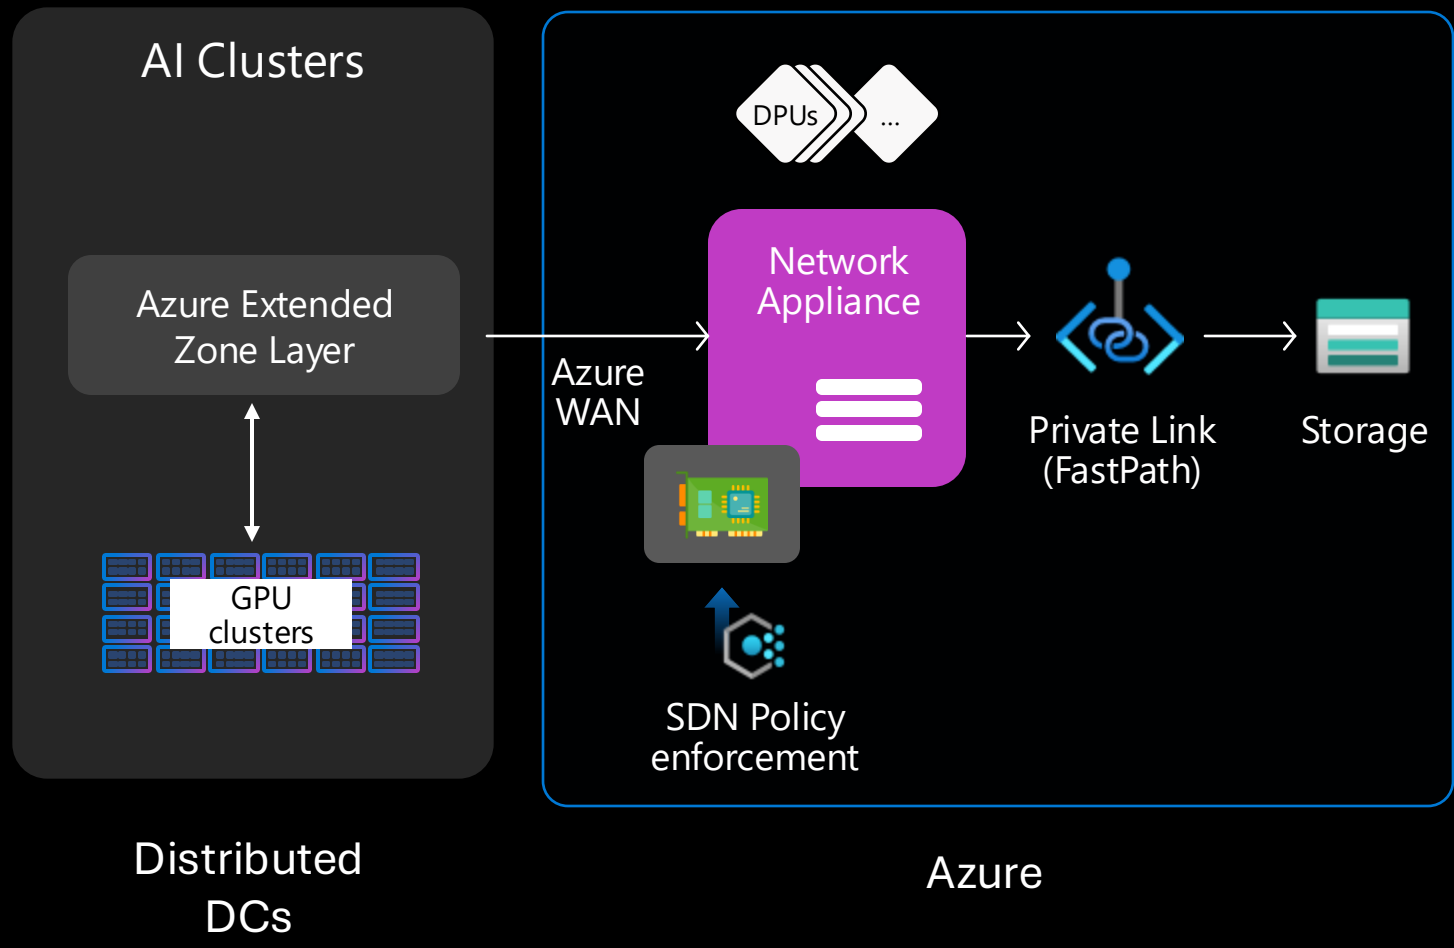

Partners are essential in keep up with demand

DASH based SmartSwitch allows for Bare Metal Virtualization

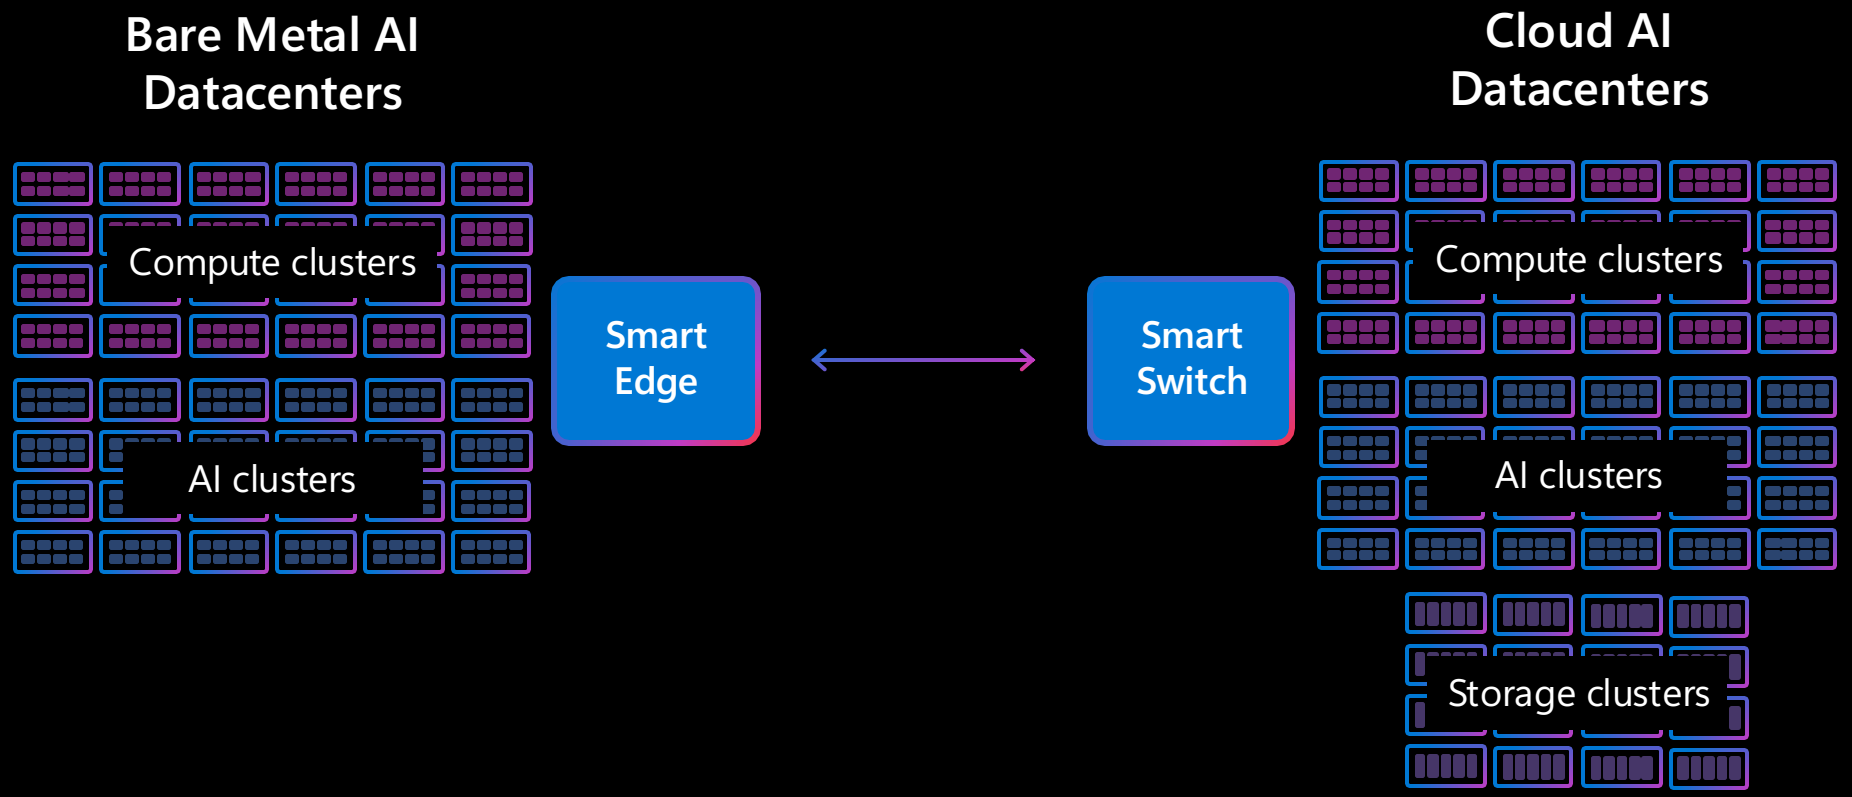

Application Networking for AI

AI workloads

- Training
- Inference
- AI Apps & Edge AI
- Data Pipelines

Networking demands

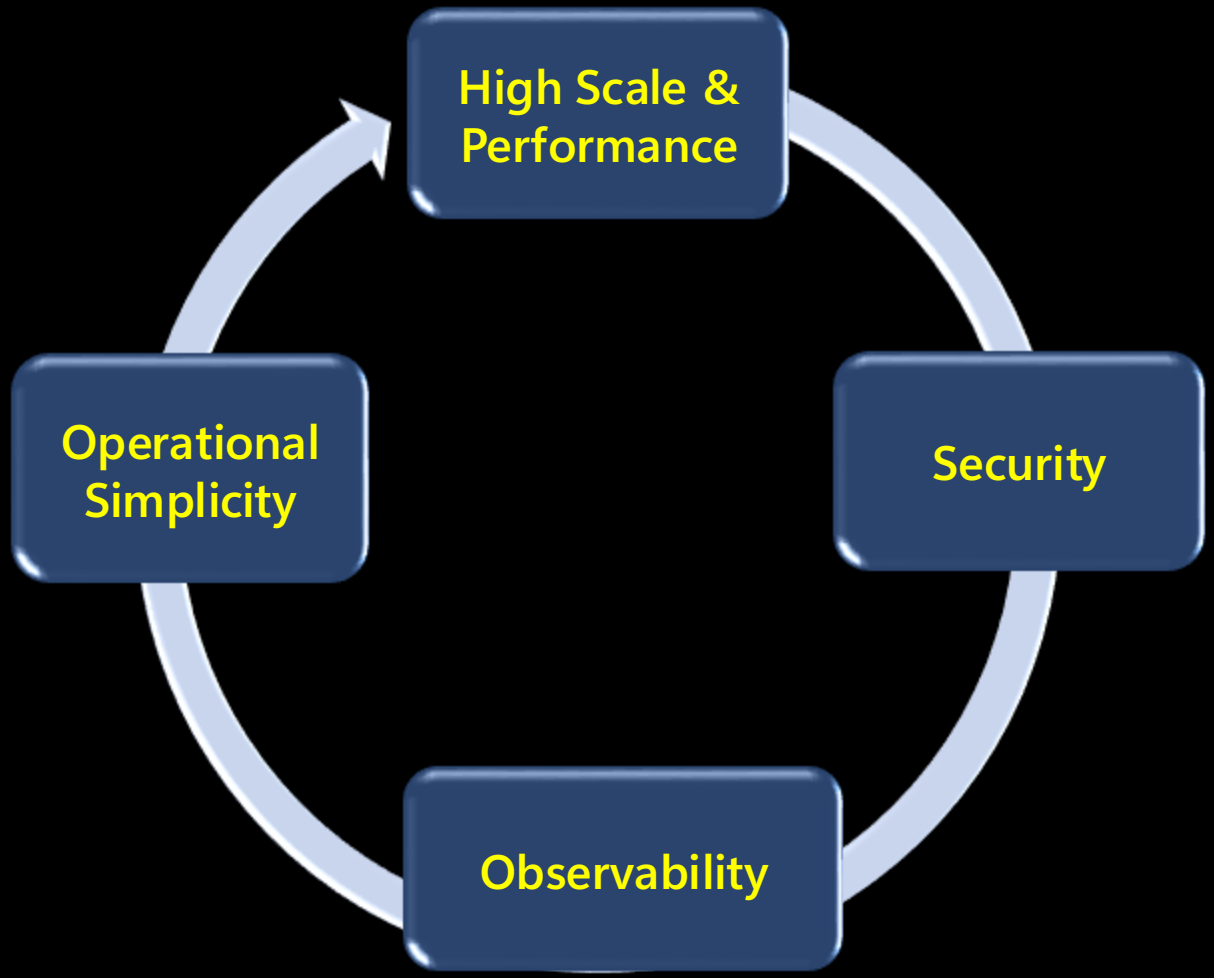

Azure container networking solutions

- **High Performance, scale:** Azure Managed CNI (Container Networking Interface)
- **Security:** Advanced L3-L7 Policies and encryption
- **Observability:** Advanced Network analysis at workload level
- **Operational Simplicity:** Multi cluster growth and Azure supported

Azure Container Networking for AI

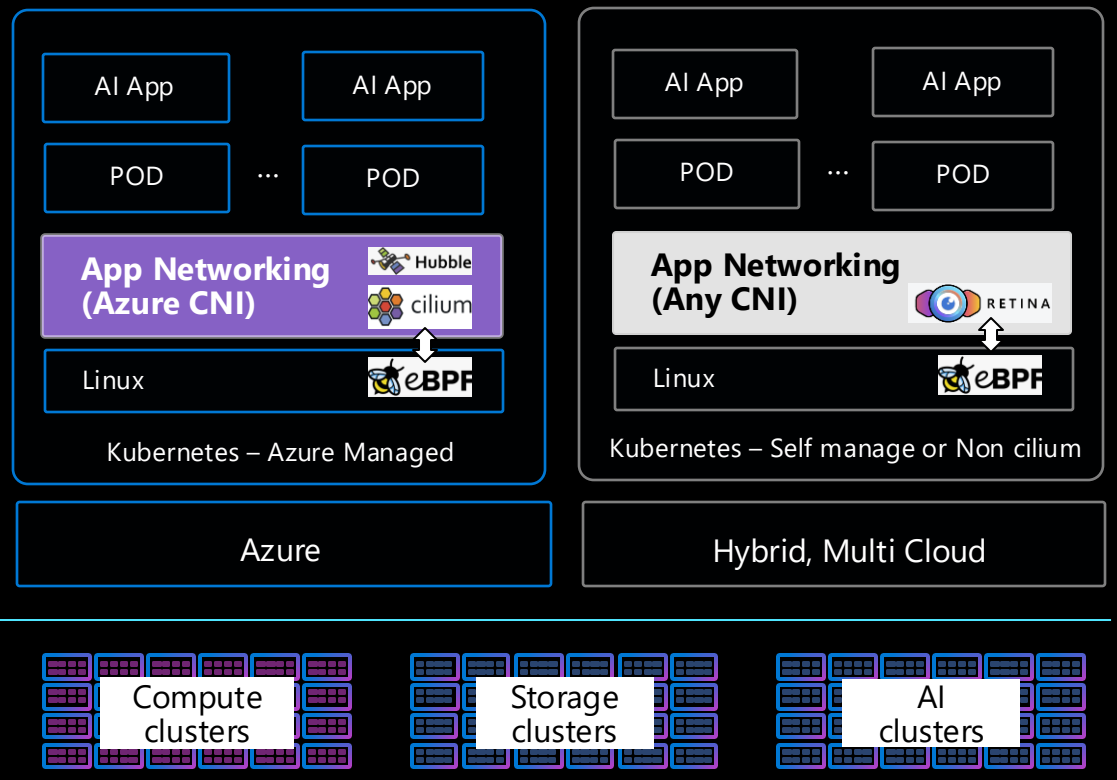

Azure CNI:

- Azure CNI powered by **Cilium** (built on eBPF Linux Kernal)
- High performance, scalable **data-plane** for AI workloads
- **Azure Integrations** for Flexible IPAM and VNET routing options scaling to 1M+ IPs
- Granular network **security** policies L4-L7, Filtering on FQDNs, Encryption
- **Observability and AI insights** with Pod and protocol metrics, flow logs

Any CNI:

- **Retina** (Open source) supports **observability** for many scenarios: customer self-managed, hybrid and multi cloud.

AI Load balancing

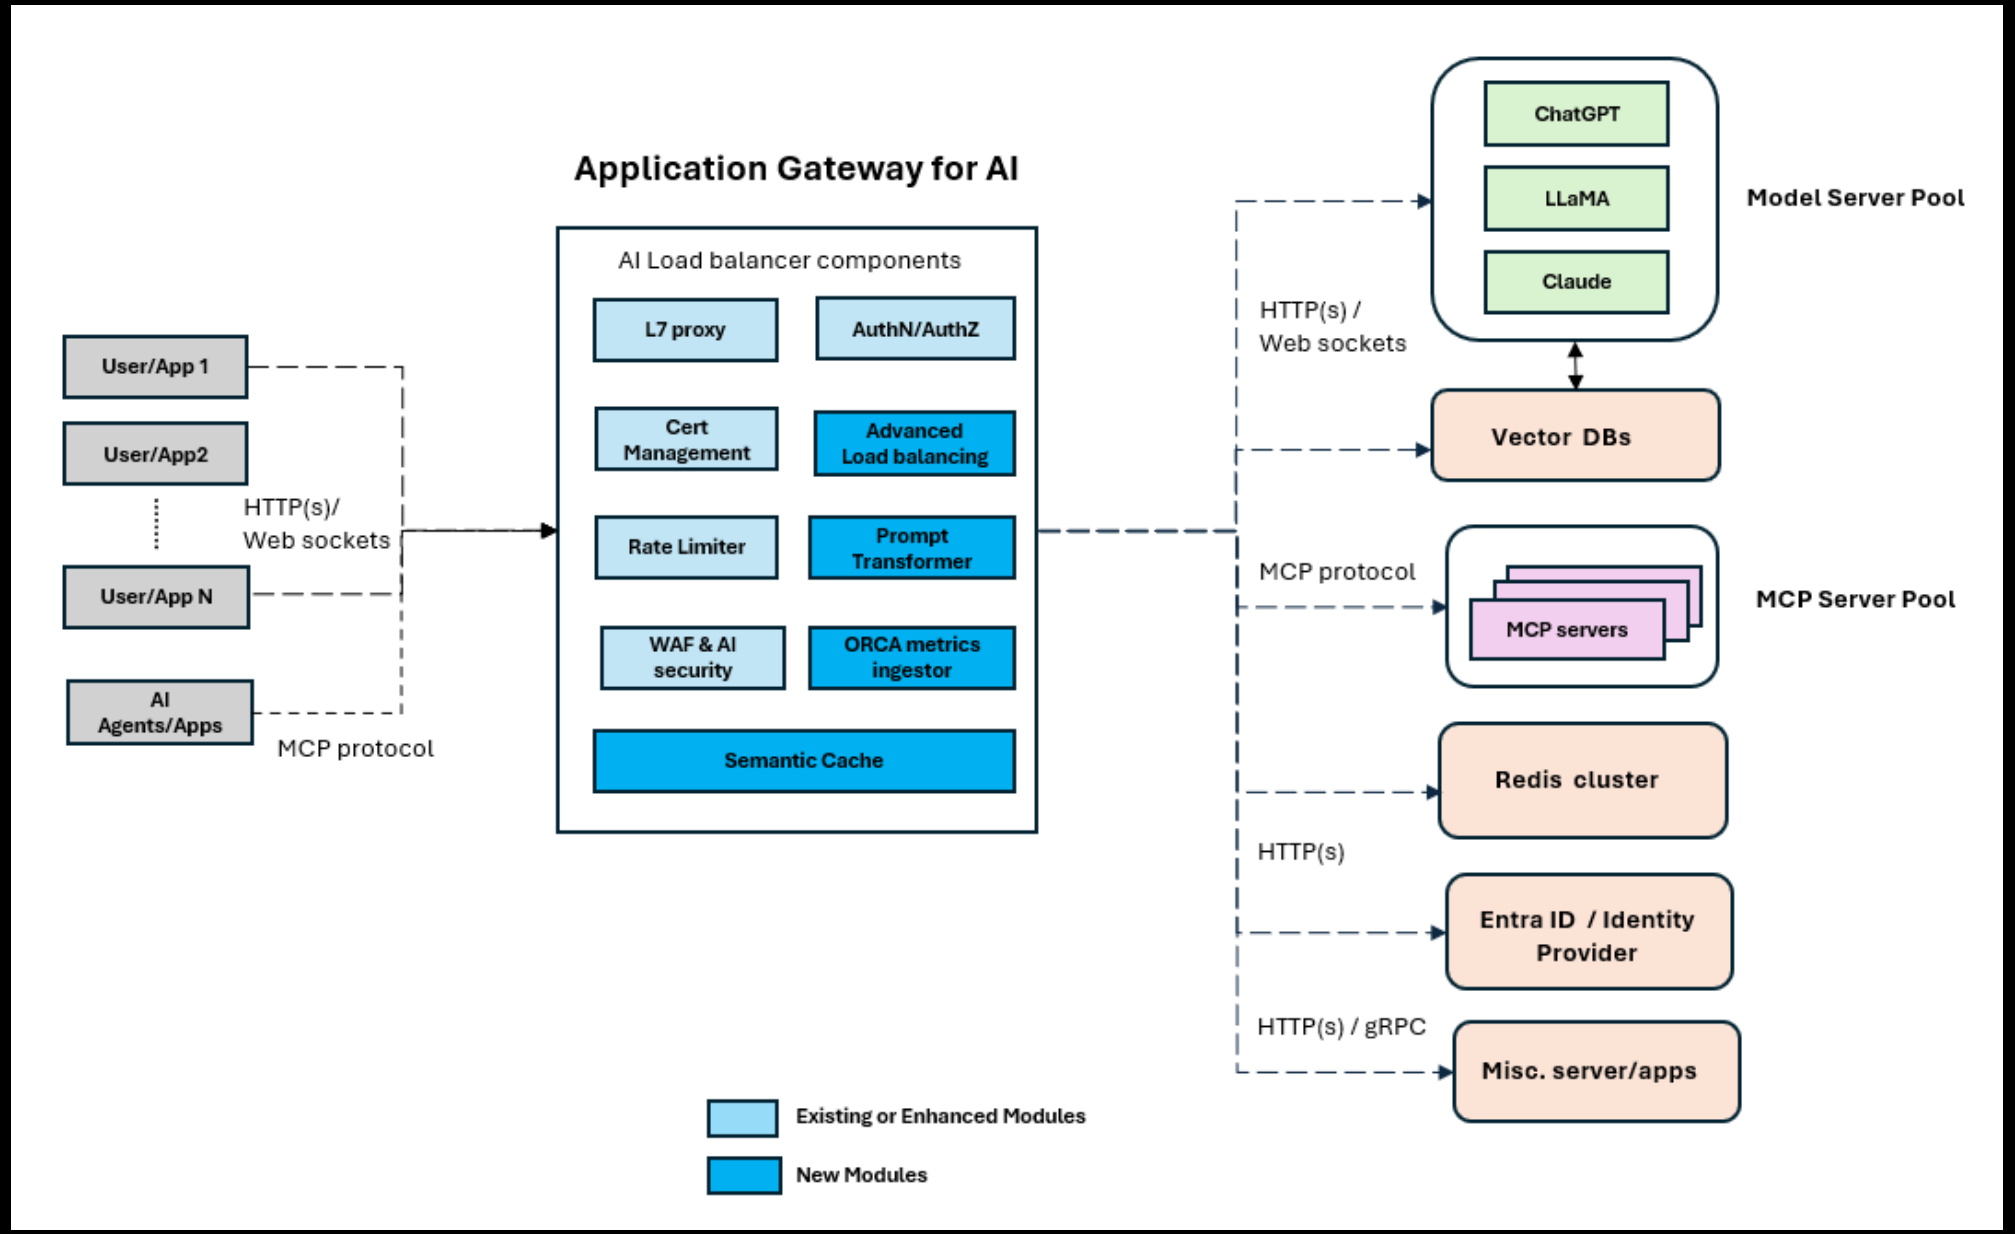

Security Perimeter

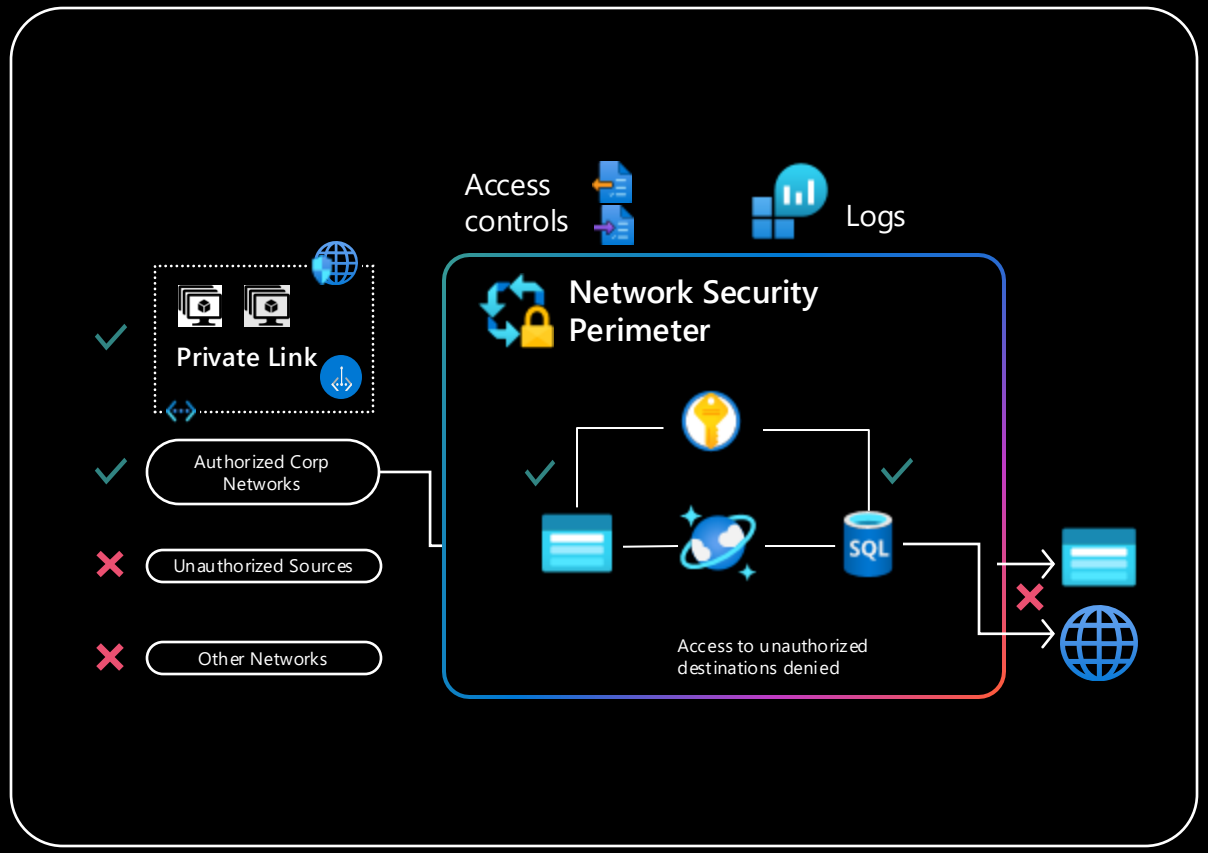

- Ensures data and models are isolated
- Unified firewall for multi-tenant resources
- Granular access control for PaaS to PaaS
- Egress controls for data exfil protection

Networking for AI

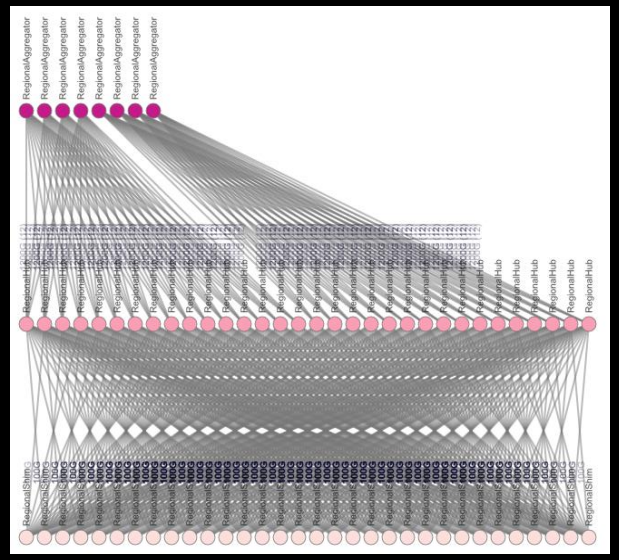

Extreme scaling
800G → 1.6T → 12.8T

Low latency
<2ms <1ms

High bandwidth
>12.8 Tbps

Leading open source
SONiC

- High **performance & security** network for
  (for all types of accelerators - Nvidia, AMD, Maia)
- Dedicated, **low latency** WAN for AI
- First large-scale co-operative transport
  (host + network) for **fast failover and load**
  balance via SRv6 on SONiC
- Millisecond-level, **high-frequency telemet**
  for troubleshooting
- BareMetal host with **SDN policies**
- Secure and high-bandwidth connectivity via
  Private Link to storage for training

100B+ near-term
AI infrastructure investments
enabling customers to

Build the future of their business on AI