Building out Modern and Scalable DC Fabric

Alperen Akpinar

May ’26 - CHNUG

Data Center Network Evolution

Speed, Tech, Ops

10/100/1000M	10/25G	100G	400G	800G+	1.6T+
STP	VLANs	VPLS	VXLAN	BGP/EVPN	UET/uSIDs/DCQCN

Legacy Ops

SNMP / CLI

Manual configs

Box-by-box management

Simple Automation

Config templates

Scripts

Basic monitoring

Intent-Based Automation

Programmable NOS

gNMI, YANG APIs

Event-driven workflows

AI-Driven Operations

Telemetry at scale

Closed-loop automation

AIOps integration

Data Center Network Evolution

Speed, Tech, Ops

You are here

10/100/1000M	10/25G	100G	400G	800G+	1.6T+
STP	VLANs	VPLS	VXLAN	BGP/EVPN	UET/uSIDs/DCQCN

Legacy Ops

SNMP / CLI
Manual configs
Box-by-box management

Simple Automation

Config templates
Scripts
Basic monitoring

Intent-Based Automation

Programmable NOS
gNMI, YANG APIs
Event-driven workflows

AI-Driven Operations

Telemetry at scale
Closed-loop automation
AIOps integration

2010s 2020s 2030s

Data Center Network Evolution

Speed, Tech, Ops

You are here

Coming soon

10/100/1000M	10/25G	100G	400G	800G+	1.6T+
STP	VLANs	VPLS	VXLAN	BGP/EVPN	UET/uSIDs/DCQCN

Legacy Ops

SNMP / CLI

Manual configs

Box-by-box management

Simple Automation

Config templates

Scripts

Basic monitoring

Intent-Based Automation

Programmable NOS

gNMI, YANG APIs

Event-driven workflows

AI-Driven Operations

Telemetry at scale

Closed-loop automation

AIOps integration

2010s 2020s 2030s

AI & Cloud Data Centers Are Getting Bigger!

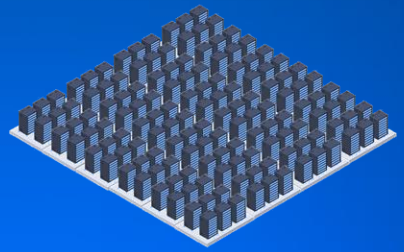

AI is resetting the expectation of what a "large Data Center" is

2013

Then…largest DC in Europe

Meta's Lulea Data Center

30 MW*

84,000 m^2

Source: Baxtel

2025

Largest DC in the World

China Telecom IMIP

150 MW

1,000,000m^2

AI & Cloud Data Centers Are Getting Bigger!

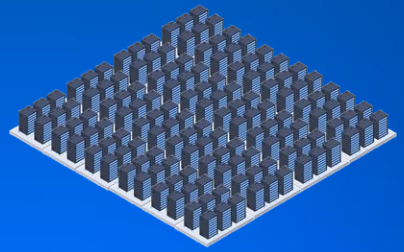

AI is resetting the expectation of what a "large Data Center" is

2013

Then…largest DC in Europe

Meta's Lulea Data Center

30 MW*

84,000 m^2

Source: Baxtel

2025

Largest DC in the World

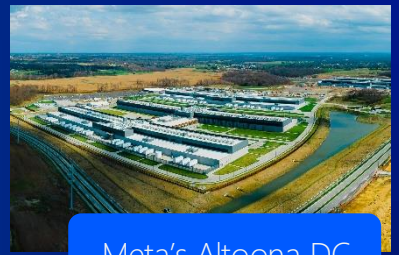

Meta's Altoona DC

500 MW

465,000 m2

AI & Cloud Data Centers Are Getting Bigger!

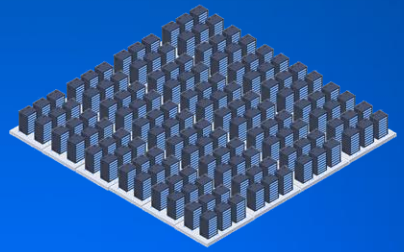

AI is resetting the expectation of what a "large Data Center" is

2013

Then…largest DC in Europe

Meta's Lulea Data Center

30 MW*

84,000 m^2

Source: Baxtel

2025

Largest DC in the World

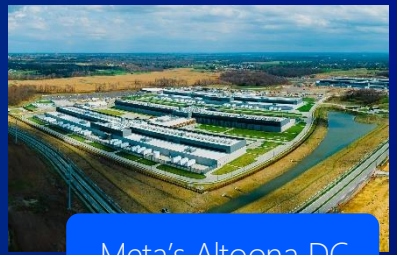

Meta's Altoona DC

500 MW

465,000 m2

2026

Largest DCs

GigaWatts

Networking for the AI era

Builds on and leverages **kubernetes**

Event-Driven-Automation

Intent-Based Network management, automation and orchestration

DDoS Protection

DDoS Detection, Mitigation and scrubbing

AI for network operations

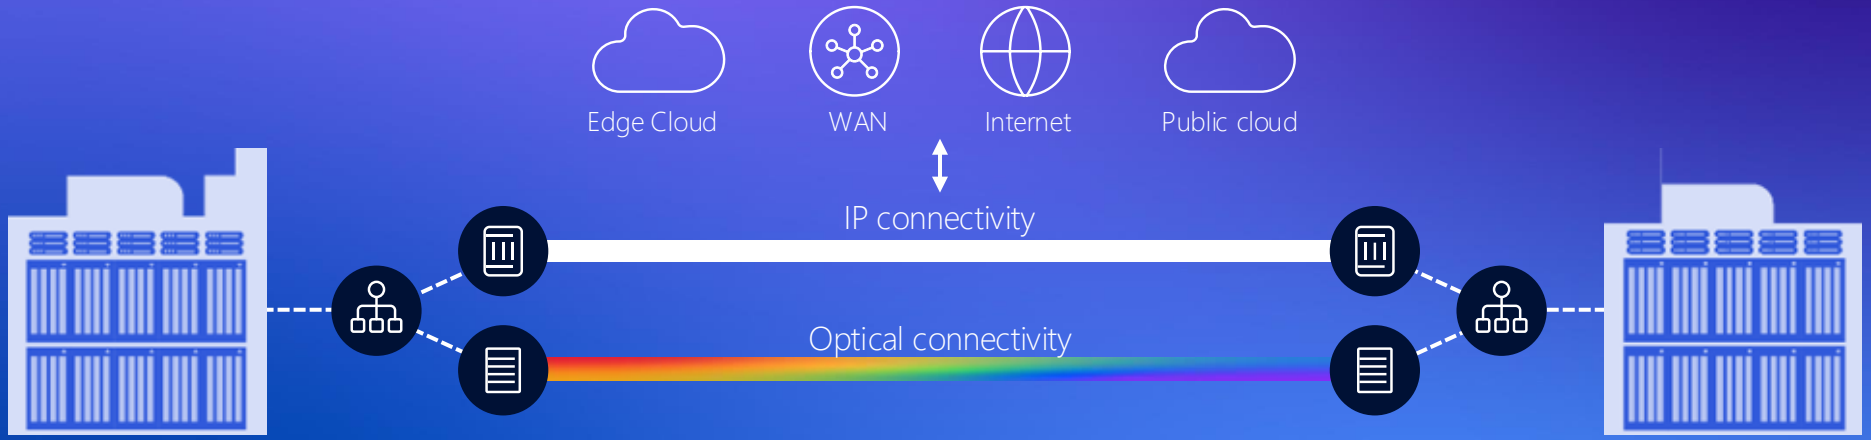

Data center fabric

Reliable data center switching with predictable and simplified operations for your data center environments

Data center gateway

High performance DC to DC, DC to internet, DC to WAN and DC to clouds interconnect and DDoS mitigation

Optical DCI

Dedicated optical networking enables high bandwidth, very low latency and highly secure data transmission for business-critical applications

Scaling the number of GPUs

Scale-Up

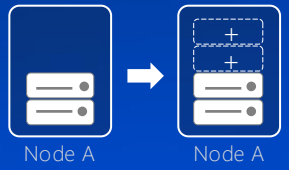

AMD UAlink
Broadcom SUE
Nvidia NVSwitch
PCIe Gen5 / Gen6
OCP ESUN

Scale-Out

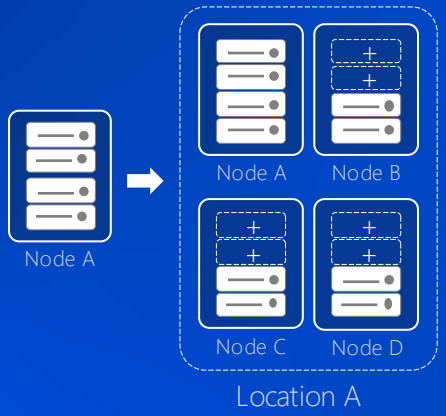

InfiniBand
ROCEv2
UEC

Scale-Across

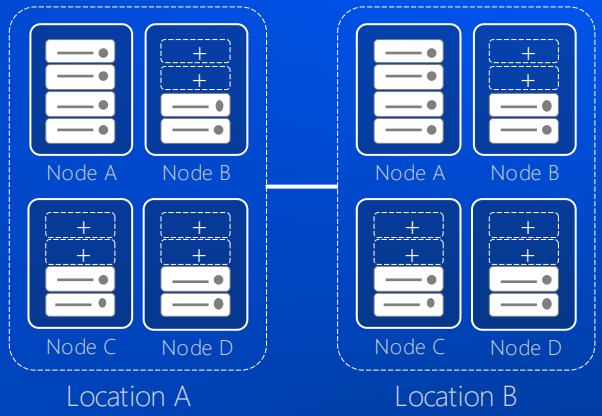

DCI (Data Center Interconnect) with optical or IP solution
RDMA over WAN
Encryption

GB200 NVL72 Intra-Rack Connectivity

Nvidia's scale-up network:

- **9 NVLink switch (18 chip)** each with **144 ports at 100GB/s** in a NVL72 rack
- Every GPU is connected to NVLink switch with 18x 100GB/s = **1.8TB/s** via NVLink Spine
- **NVLink Spine:** a copper fabric in the rack that connects compute trays (GPUs) to NVLink switches

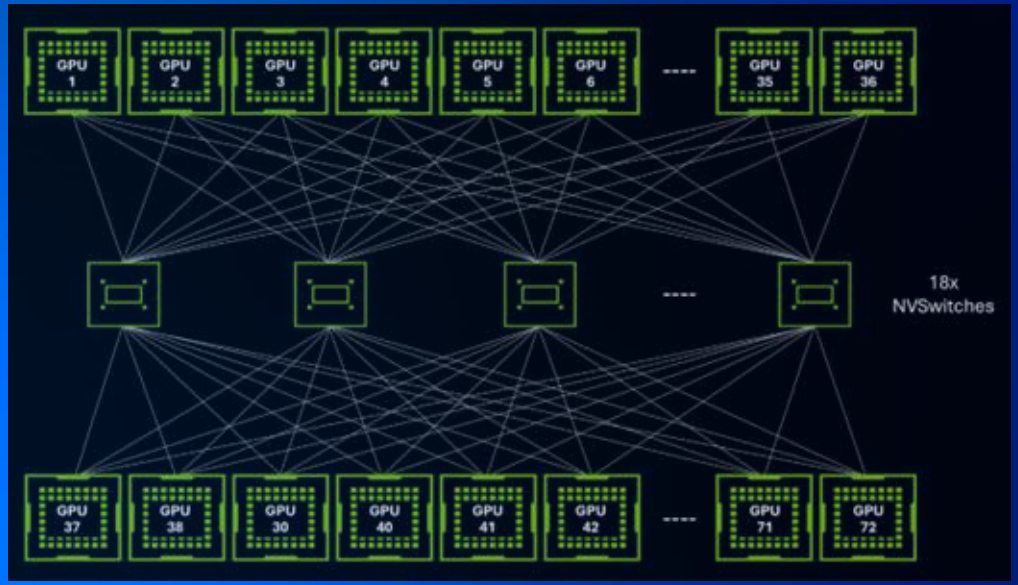

NVLink Interconnect

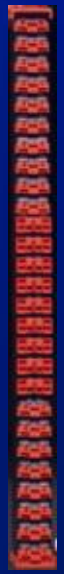

NVLink Spine

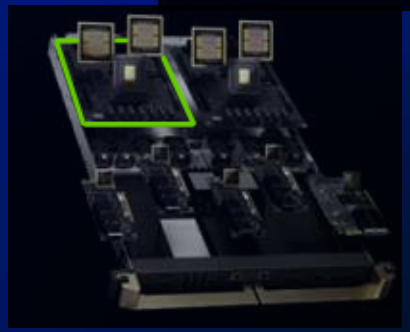

Compute Tray (2x GB200)

Scaling the number of GPUs

Scale-Up

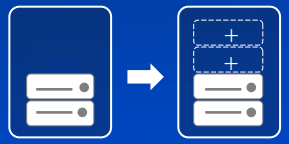

AMD UAlink
Broadcom SUE
Nvidia NVSwitch
PCIe Gen5 / Gen6
OCP ESUN

Scale-Out

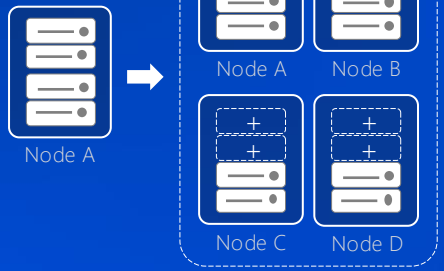

InfiniBand
ROCEv2
UEC

Scale-Across

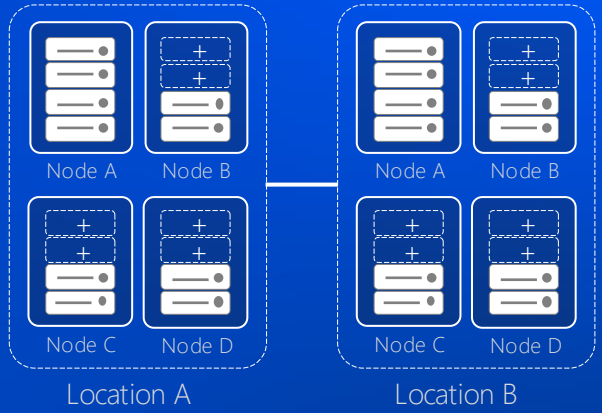

DCI (Data Center Interconnect) with optical or IP solution
RDMA over WAN
Encryption

Scale-out: AI traffic is different from traditional apps

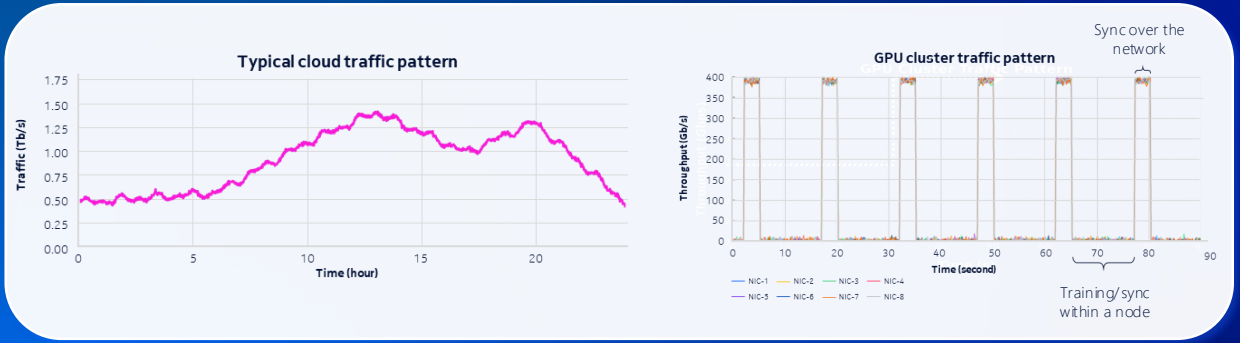

Job Completion Time (JCT): total time taken for data to be transferred across that network during the training process

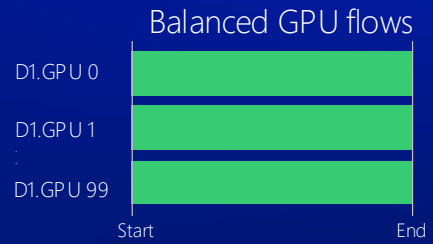

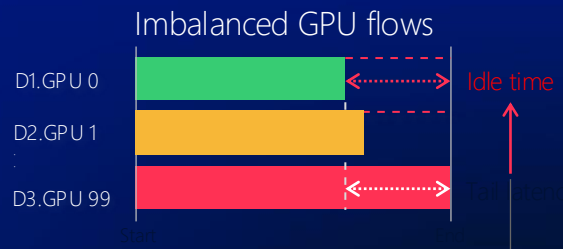

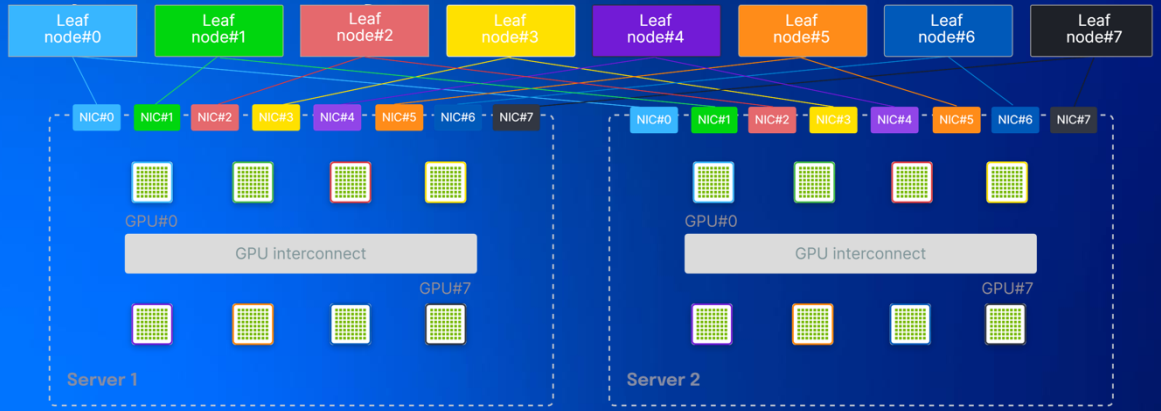

Directly affects training step times; slow GPUs delay all others!

High-speed GPU interconnect: *NVLink, InfinityFabric etc.*

AI-ready networks

Between data centers

Inside data centers

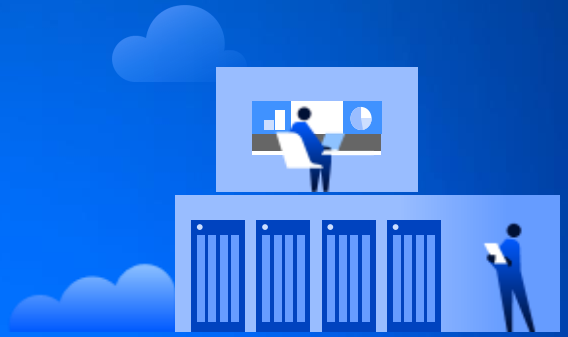

Data center

AI-ready networks

Between data centers

Inside data centers

- Scalable
- Lossless
- Automated

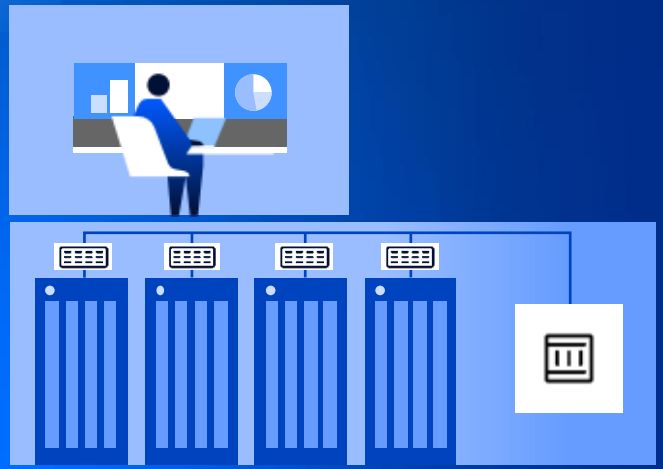

800G / 1.6T AI backend fabrics

- AI features on 7220 IXR-H5/H6 and 7250 IXR Gen3+

AI-ready networks

Between data centers

Inside data centers

Scalable

Lossless

Automated

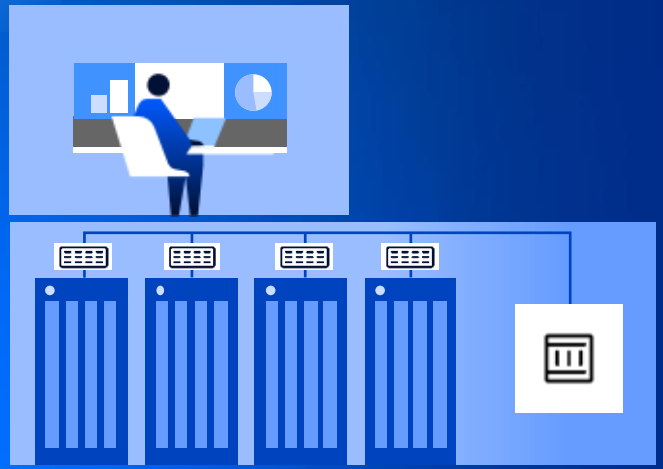

800G / 1.6T AI backend fabrics

- AI features on 7220 IXR-H5/H6 and 7250 IXR Gen3+

Load balancing and QoS

- DLB and RDMA aware (Q-pair) hashing
- DCQCN (PFC/ECN), MRC…

AI-ready networks

Between data centers

Inside data centers

Scalable

Lossless

Automated

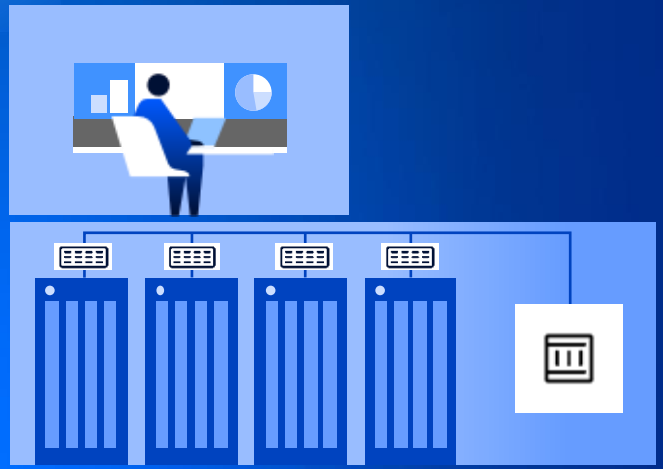

800G / 1.6T AI backend fabrics

- AI features on 7220 IXR-H5/H6 and 7250 IXR Gen3+

Load balancing and QoS

- DLB and RDMA aware (Q-pair) hashing
- DCQCN (PFC/ECN), MRC…

Advanced AI telemetry

- Congestion indicators
- Real time insight

AI-ready networks

Between data centers

Inside data centers

Scalable

Lossless

Automated

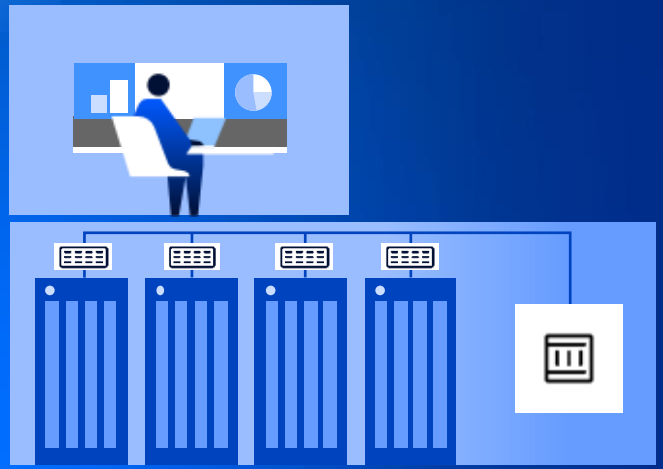

800G / 1.6T AI backend fabrics

- AI features on 7220 IXR-H5 and 7250 IXR Gen3

Load balancing and QoS

- DLB and RDMA aware (Q-pair) hashing
- DCQCN (PFC/ECN), MRC…

Advanced AI telemetry

- Congestion indicators
- Real time insight

Backend multi-tenancy

- Lightweight and secure segmentation

AI-ready networks

Between data centers

Inside data centers

Scalable

Lossless

Automated

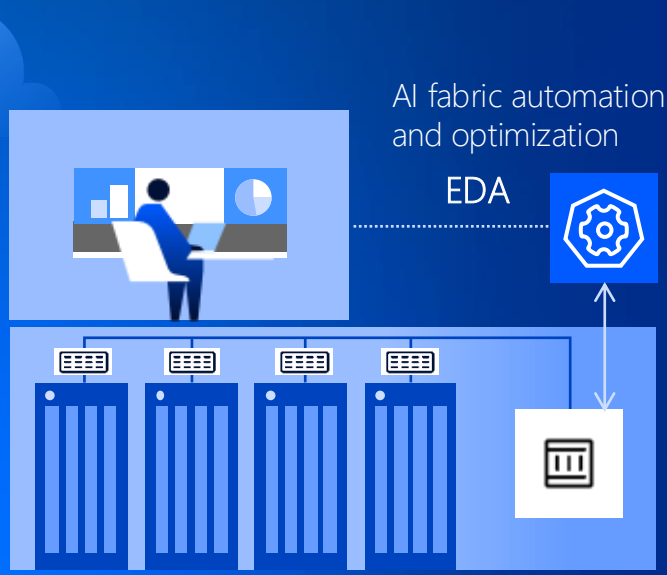

800G / 1.6T AI backend fabrics

- AI features on 7220 IXR-H5 and 7250 IXR Gen3

Load balancing and QoS

- DLB and RDMA aware (Q-pair) hashing
- DCQCN (PFC/ECN), MRC…

Advanced AI telemetry

- Congestion indicators
- Real time insight

Backend multi-tenancy

- Lightweight and secure segmentation

Scaling the number of GPUs

Scale-Up

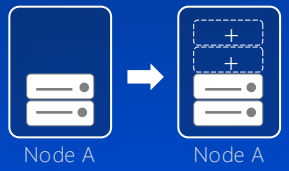

AMD UAlink
Broadcom SUE
Nvidia NVSwitch
PCIe Gen5 / Gen6
OCP ESUN

Scale-Out

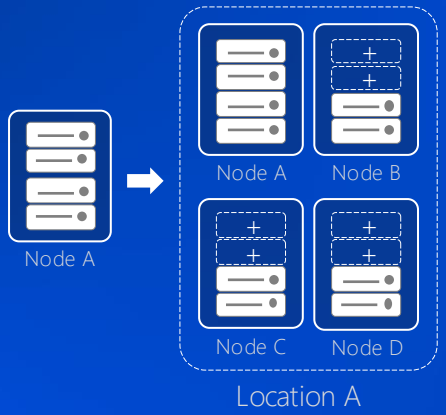

InfiniBand
ROCEv2
UEC

Scale-Across

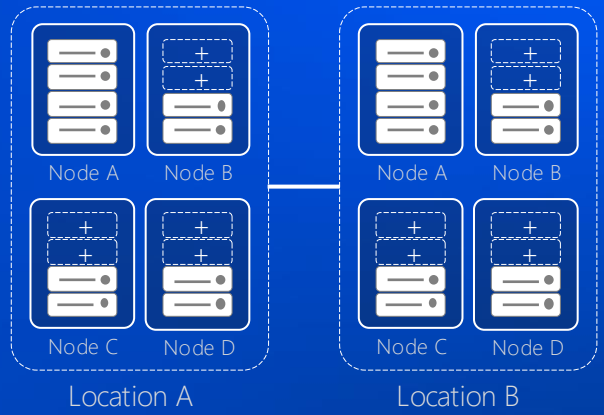

DCI (Data Center Interconnect) with optical or IP solution
RDMA over WAN
Encryption

AI-ready networks

Between data centers

Inside data centers

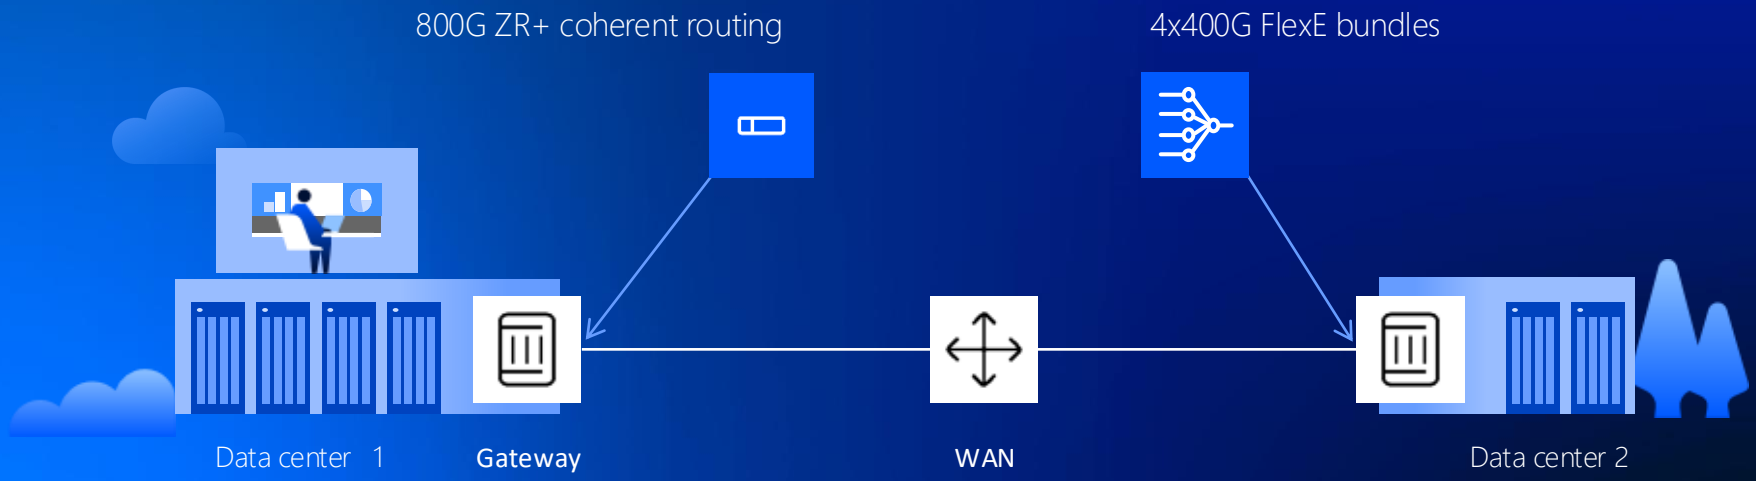

AI-ready networks

Between data centers

Inside data centers

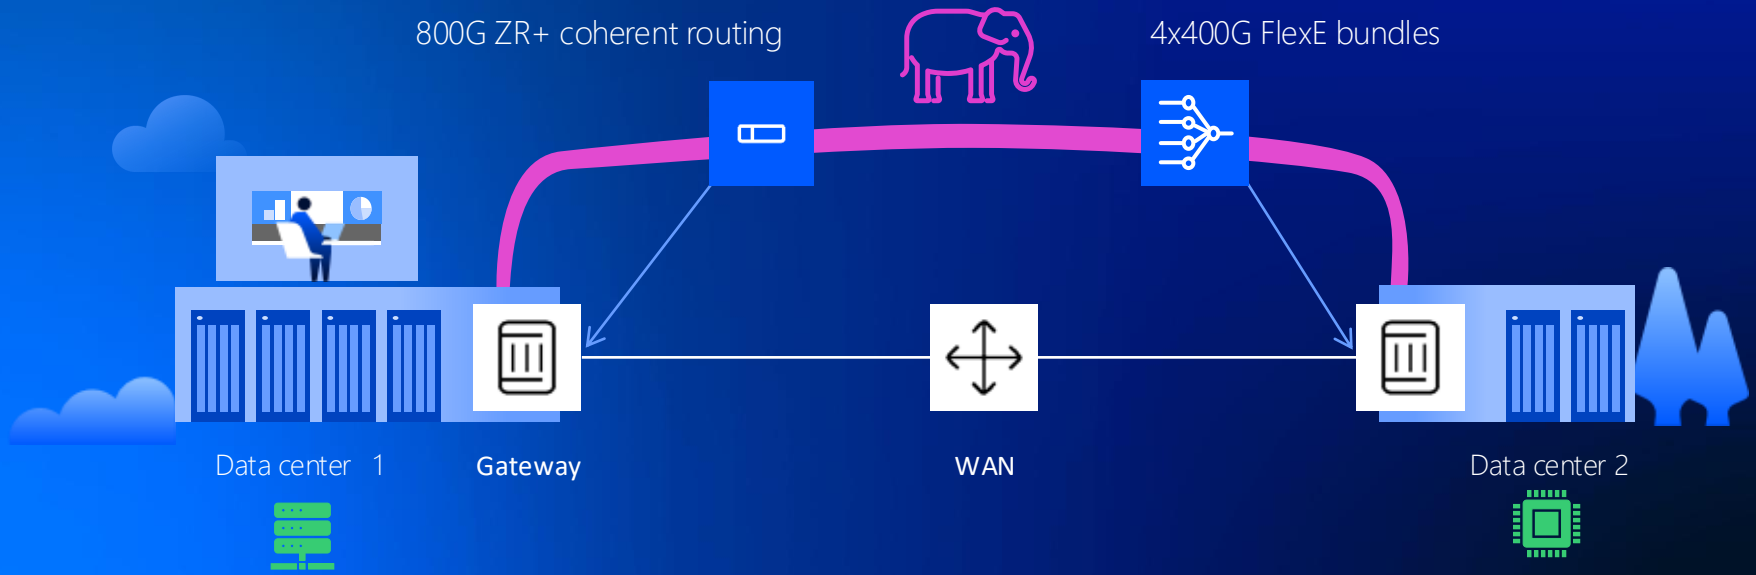

Scale across: Network requirements are similar to DCI

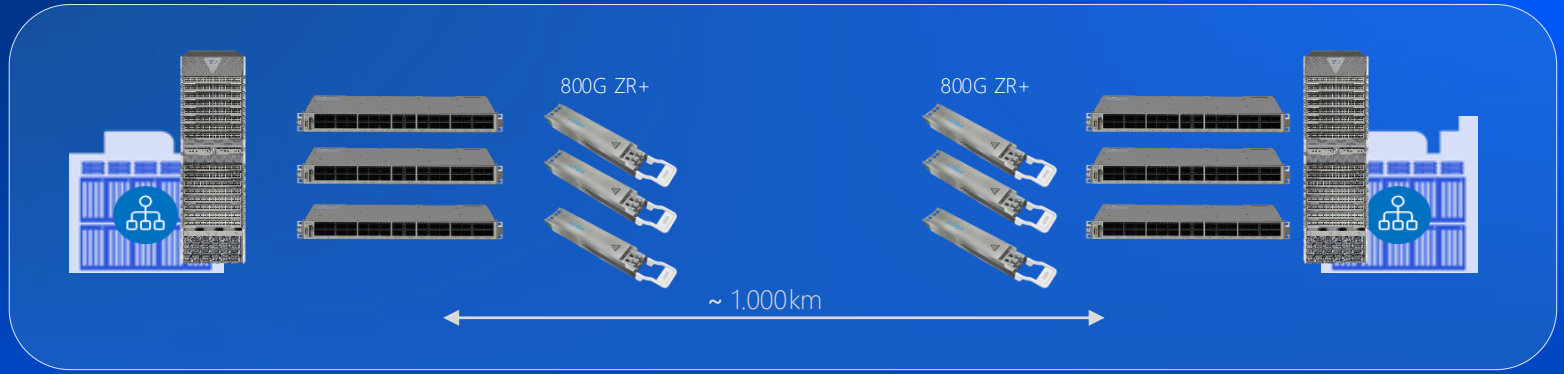

High-Density Scale-Across Router

Abstraction of internal DC network for scale

Managing WAN-realities with many paths

- Dynamic provisioning and steering of traffic
- High-frequency reporting of telemetry/congestion data
- Shared underlying

800G ZR/ZR+ (C band and L band)

Maximum capacity
Single RU Router @ 25T
Rack systems @ ,576 PB

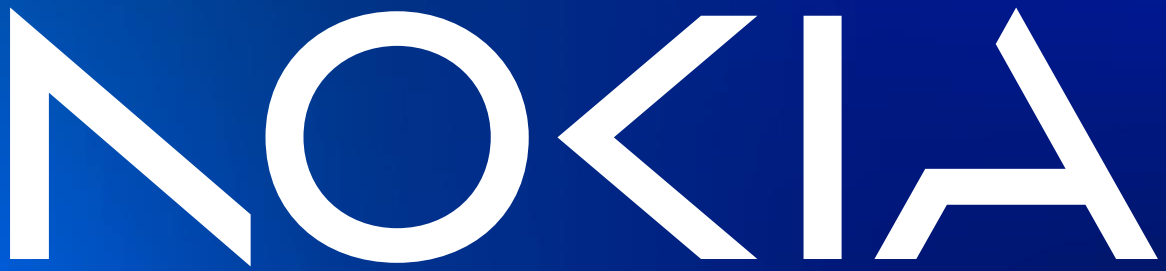
NOKIA

Nokia AI Deployment Examples

Based on Ethernet, open and future-safe networking

NSCALE

Press release

- "Smaller" Training Clusters @1K to 16K GPUs
  - Two-tier topologies

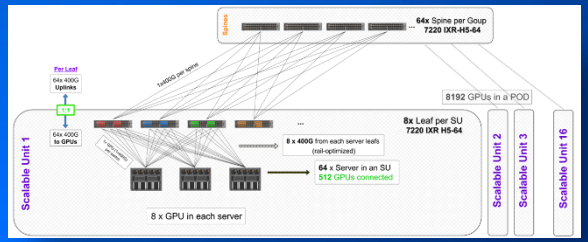

- "Giga" AI Factories @100K+ to 1M GPUs
  - Multi-plane topologies
  - Advanced routing, QoS, and load balancing

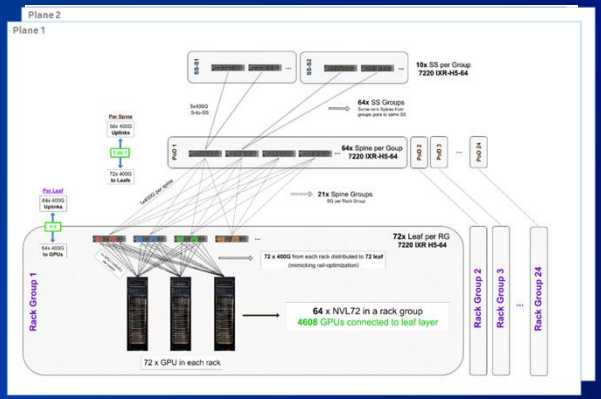

- Inference Edge Clusters
  - Optimized for size and consolidation
  - Integrated WAN interconnect

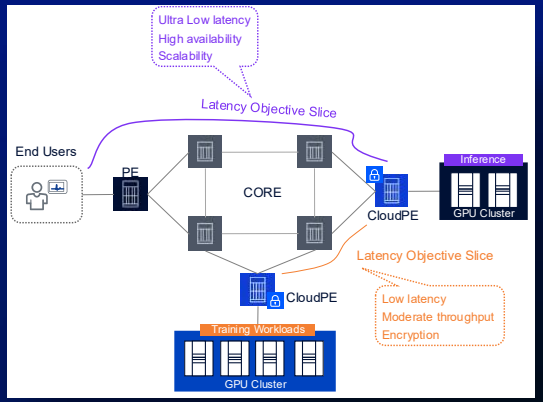

NOKIA

Nokia AI Deployment Examples

Based on Ethernet, open and future-safe networking

NSCALE

Press release

- "Smaller" Training Clusters @1K to 16K GPUs
  - Two-tier topologies

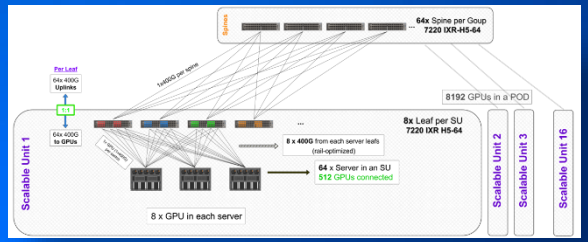

- "Giga" AI Factories @100K+ to 1M GPUs
  - Multi-plane topologies
  - Advanced routing, QoS, and load balancing

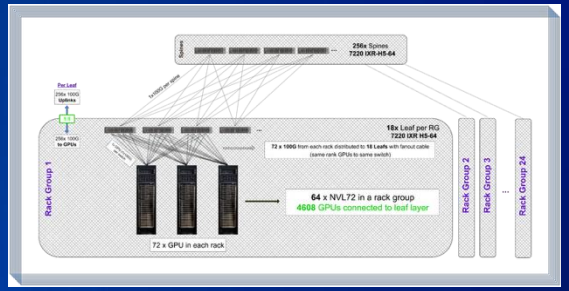

- Inference Edge Clusters
  - Optimized for size and consolidation
  - Integrated WAN interconnect

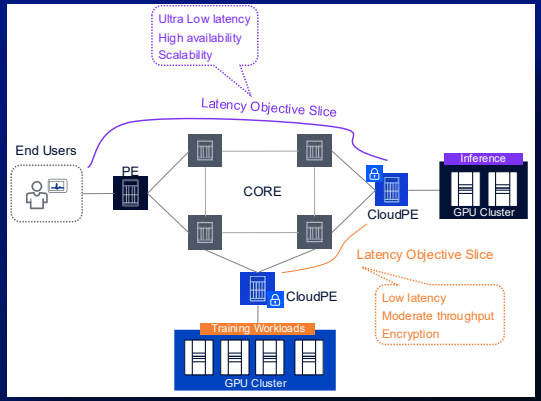

NOKIA

Building for open

Consistently raising the bar

Copyright and confidentiality