1FINITY
a Fujitsu company

AI is Ready. Network is Not!

17th August 2026
1Finity Inc.
Solyman Ashrafi (*Solyman.ashrafi@fujitsu.com*)
Junichi Sugiyama
Katsumi Fukumitsu

Ecosystem Problem

1. Hyperscalers and model owners have compute capacity problems
2. Beyond the top NeoCloud providers, most have stranded compute problems
3. Nvidia & Hyperscalers solved Scale UP and Scale Out. Scale Across at physical layer needs more attention
4. How to connect data centers of the same owner (NeoCloud or hyperscaler) as a gigantic supercomputer
5. Optical Scale Across problem can be addressed by 1Finity products including 1Finity SmartNIC

Compute Racing Faster than Scale Across Interconnect

- AI is ready to leave the lab, but the network is not ready to carry it
- Bottleneck is shifting from compute capacity to network capability
- Compute (GPUs, accelerators, memory) and architecture have scaled faster than interconnect by 10^3
- Network not a passive background utility anymore
- AI performance now depends on BW, latency, distance, jitter, path control, observability (1F S-NIC !)
- NeoClouds have stranded compute
- DCI is too static, too rigid, and too “one-to-one.”
- AI needs something much more dynamic, distributed with compute continuum

HW FLOPS, DRAM BW, Interconnect BW

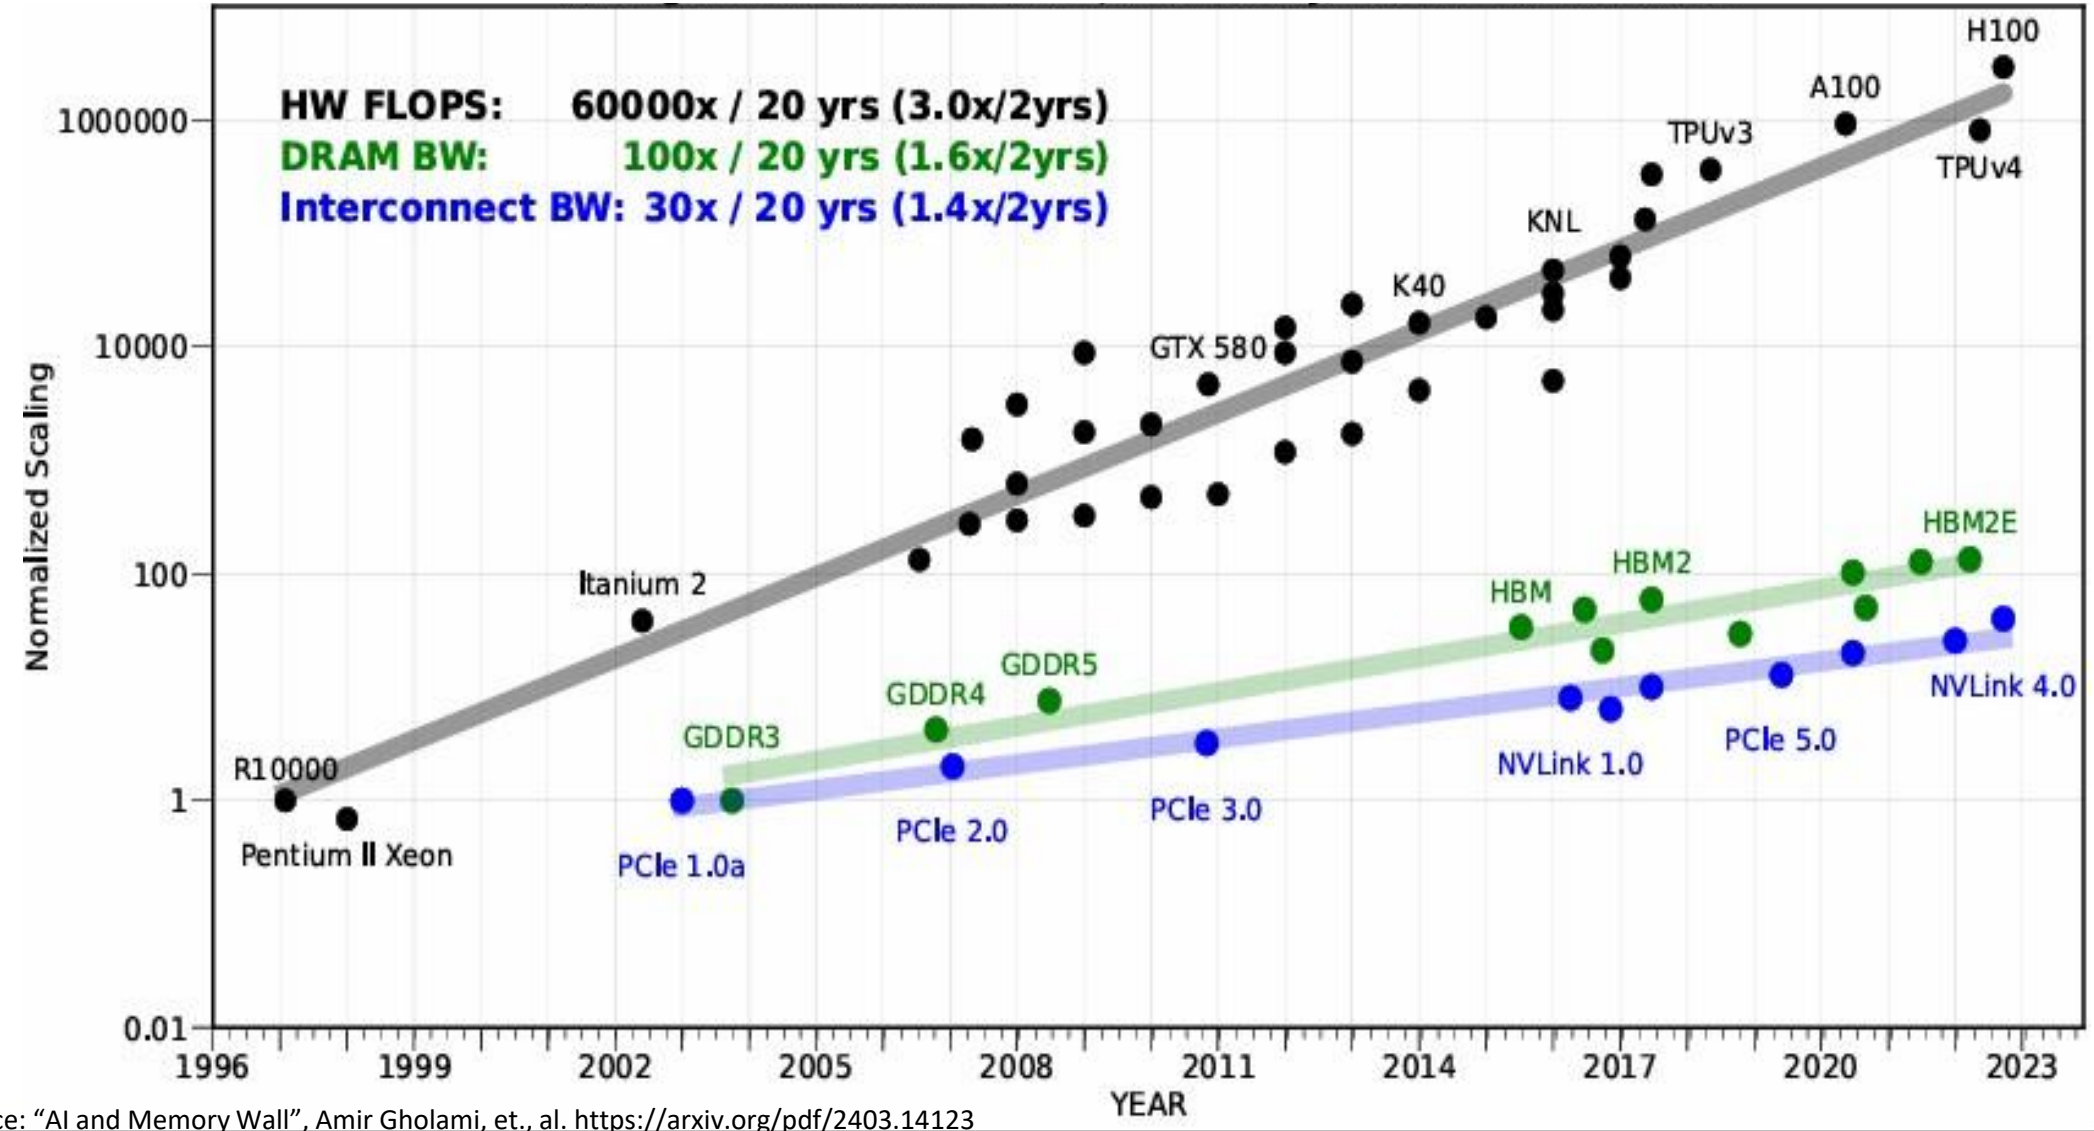

$> 10^3$

Reference: "AI and Memory Wall", Amir Gholami, et., al. https://arxiv.org/pdf/2403.14123

Scale Cross ≠ DCI

- Stitching Compute Continuum with Scale Across
- Scale Across extends backend beyond traditional DCI
- DCI focused on disaster recovery, storage replication, and workload migration
- AI Scale Across supports multiple backend networks for training
- Training depends on utilization, collective communication, and selected model-parallel approaches
- Inference depends on latency-aware routing, KV-cache and low latency predictable time-to-first-token
- Fluid Interconnect creates a programmable, xAware fabric for distributed AI infrastructure
- xAware = traffic-aware, distance-aware, latency-aware, jitter-aware, network-aware, performance-aware, and tier-aware

Scale Across Interconnect **is no longer plumbing**—it is a first-order system design variable that determines viability at scale

Scale Across = Scale Out + Long Distance

Bits/Sec/Hz/m/Latency/Prompt/Query/Token/Watts/$

1F SmartNIC to Maintain
Long Distance Throughput

Can the Ecosystem Improve this?

What is required to 2X this every 2 years?

1:1 Interconnects

Hyperscalers

Neoscaler

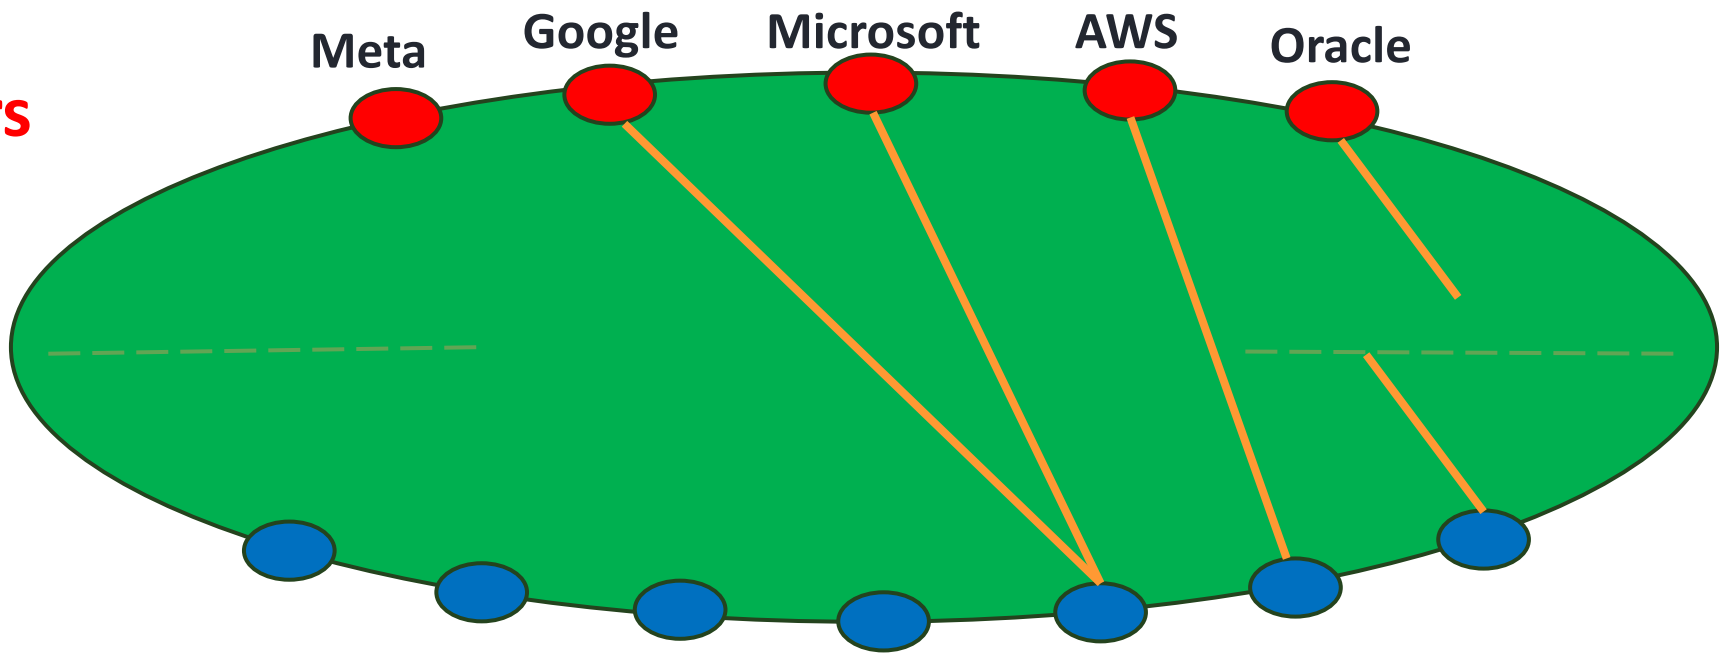

CoreWeave

ORACLE NEBIUS Azure

Crusoe FluidStack

together.ai Lambda Google Cloud aws

Scaleway Cirrascale VULTR VOLTAGE PARK GCORE firmus GMO TENSORWAVE

n:m Fluid Interconnect

*interconnect as a **programmable system layer***

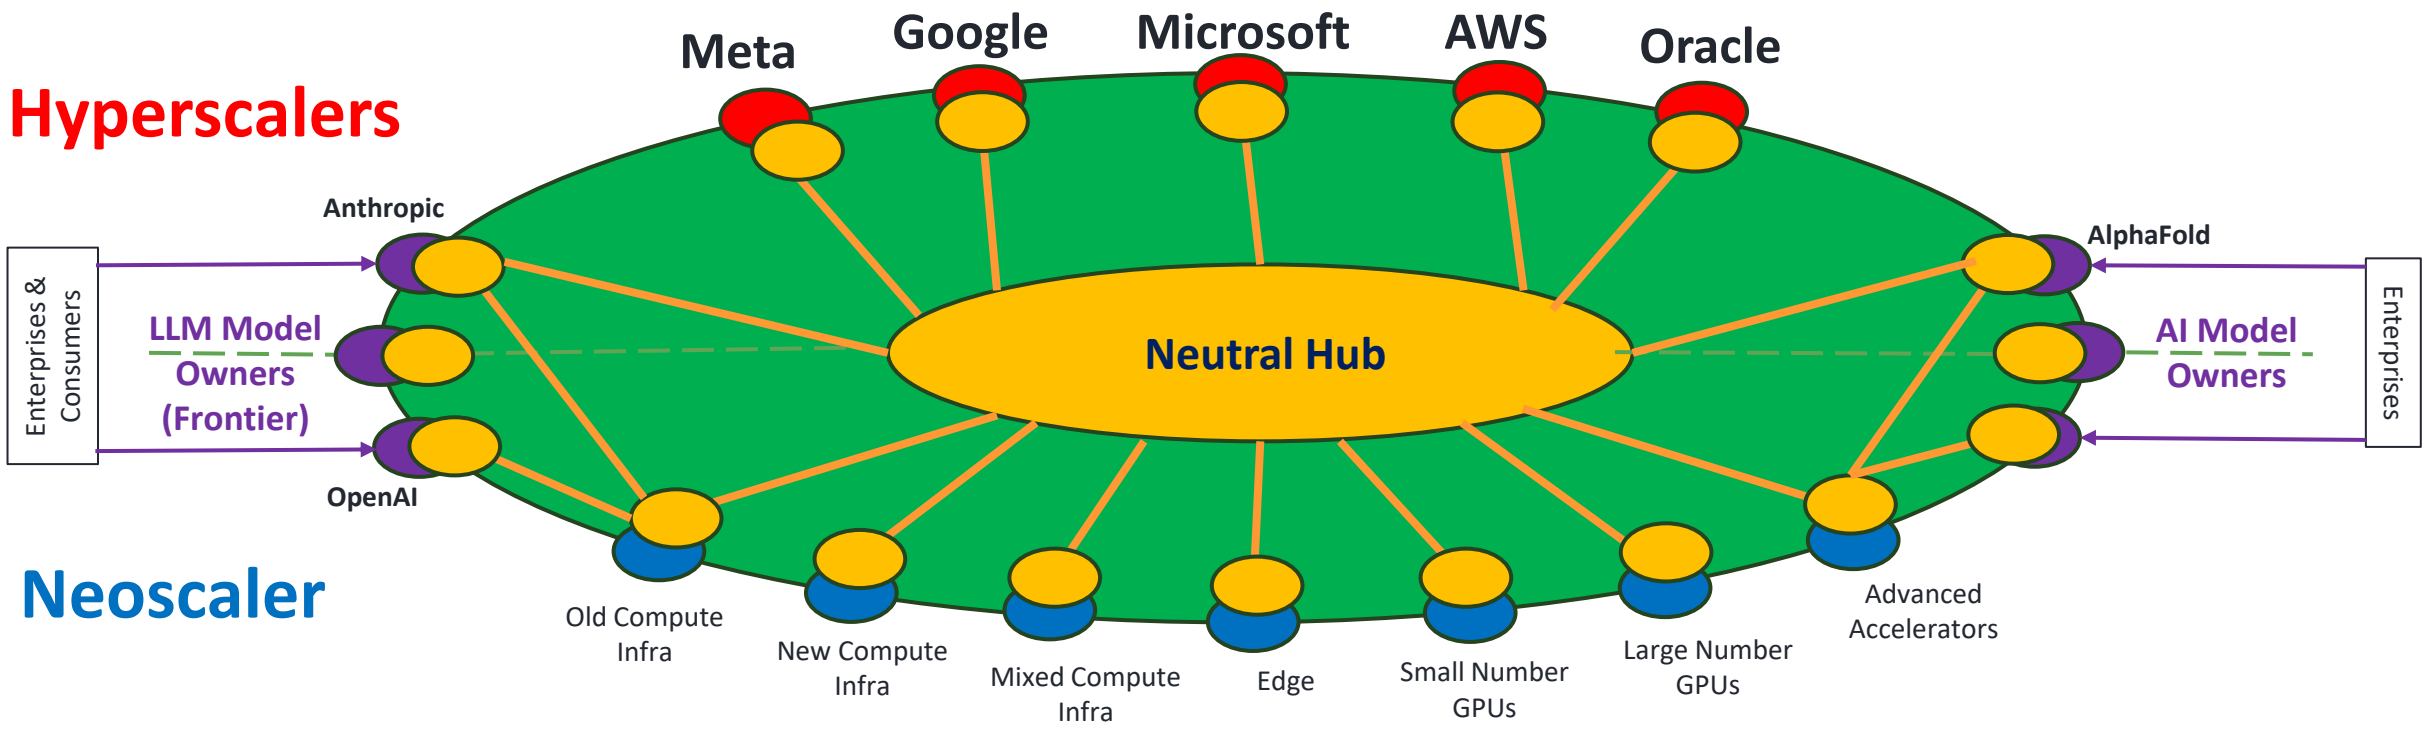

1. Fabric is Programmable, Traffic Aware, Distance Aware, Latency Aware, Jitter Aware, Network Aware, Performance Aware & Tier Aware
2. Optimizes for Lowest Cost, Highest Performance, Operationality Levels & Efficiency
3. Support Legal Authority for Sovereign AI
4. Support Geographic Residency for Data and Geo-Fencing

Multiple Hubs by Telcos, COLOs, CDNs

*interconnect as a **programmable system layer***

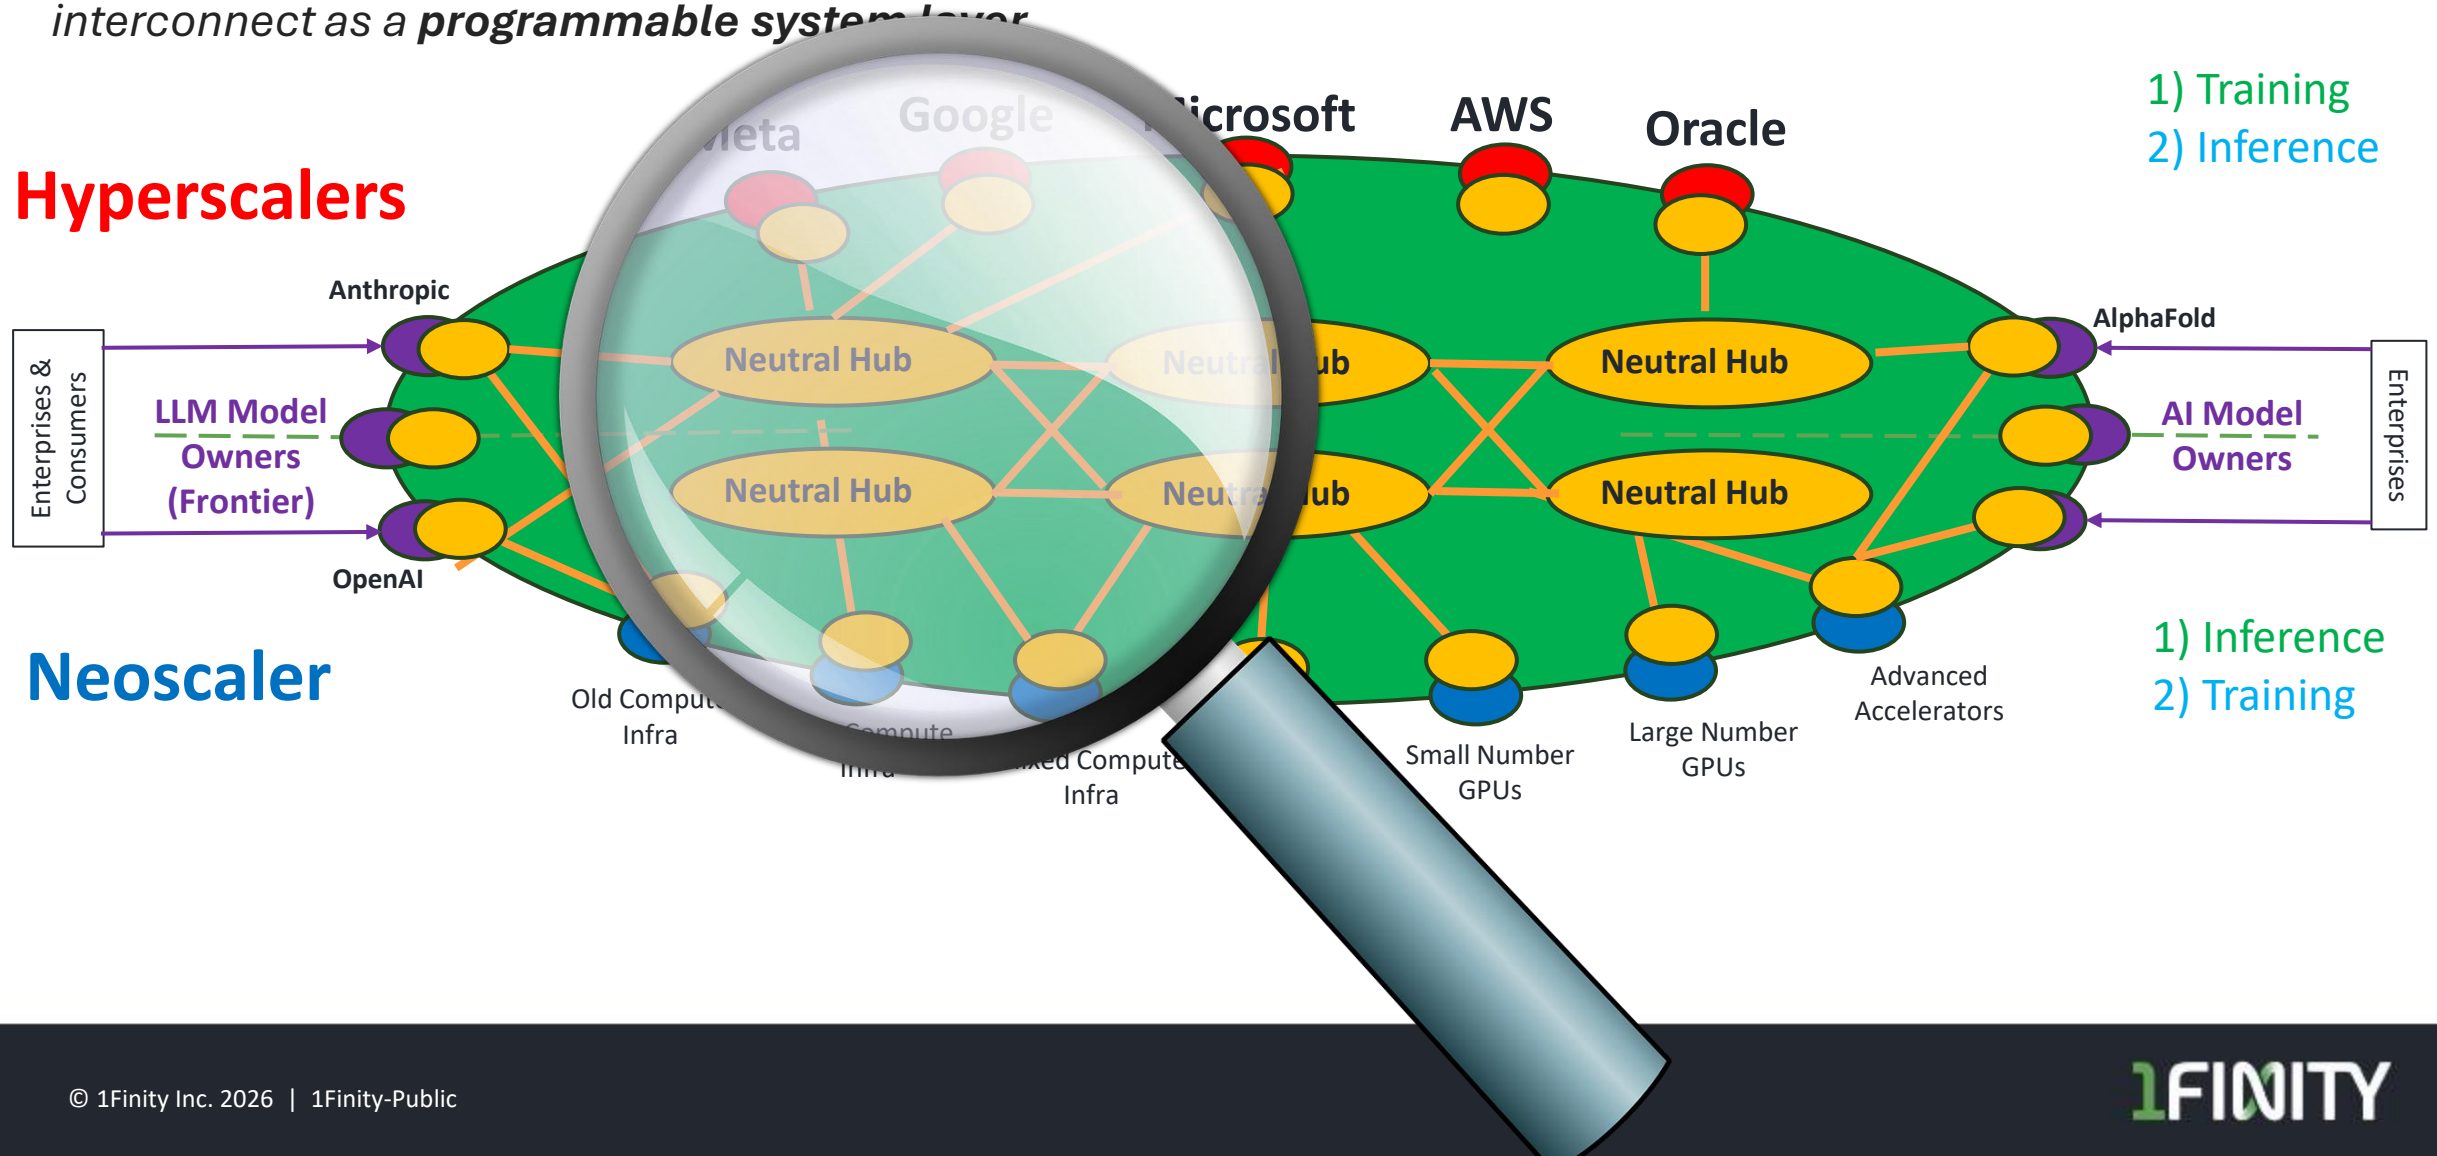

With 1Finity SmartNIC

(Hyperscaler Data Center Interconnect - A giant supercomputer infrastructure)

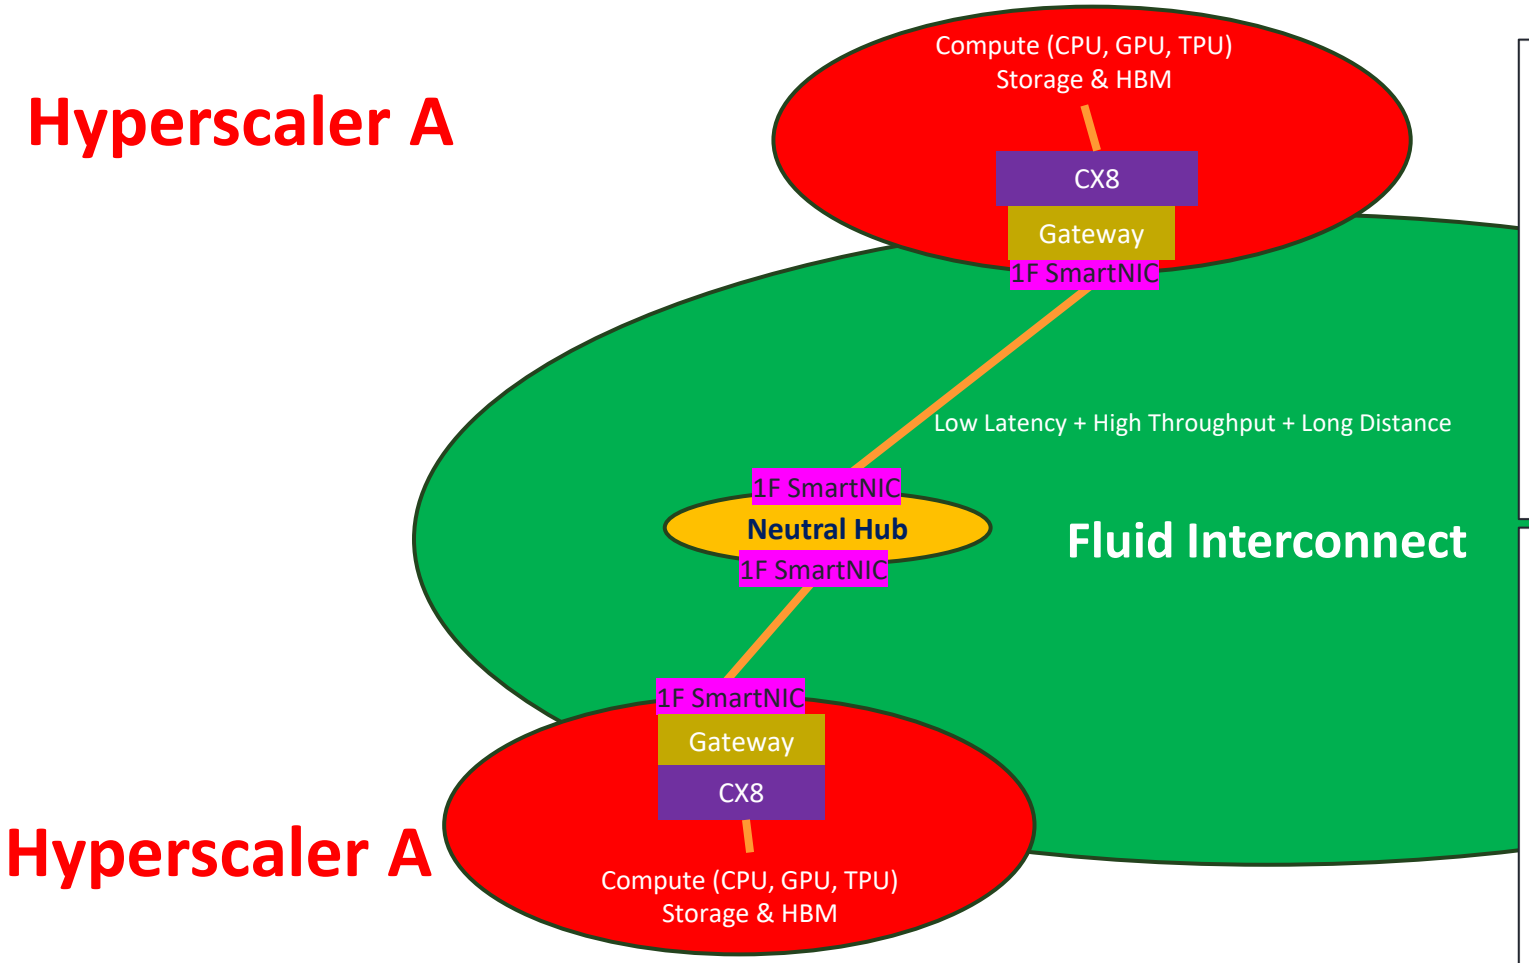

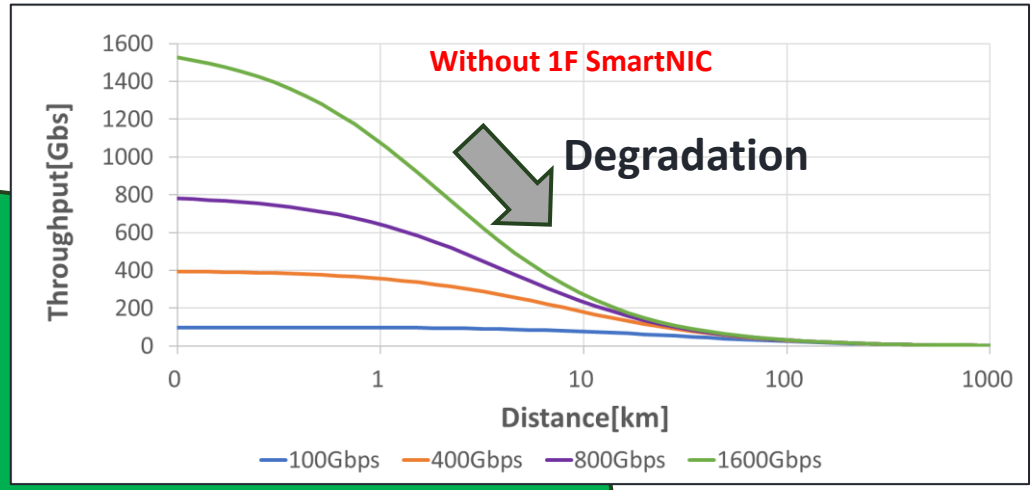

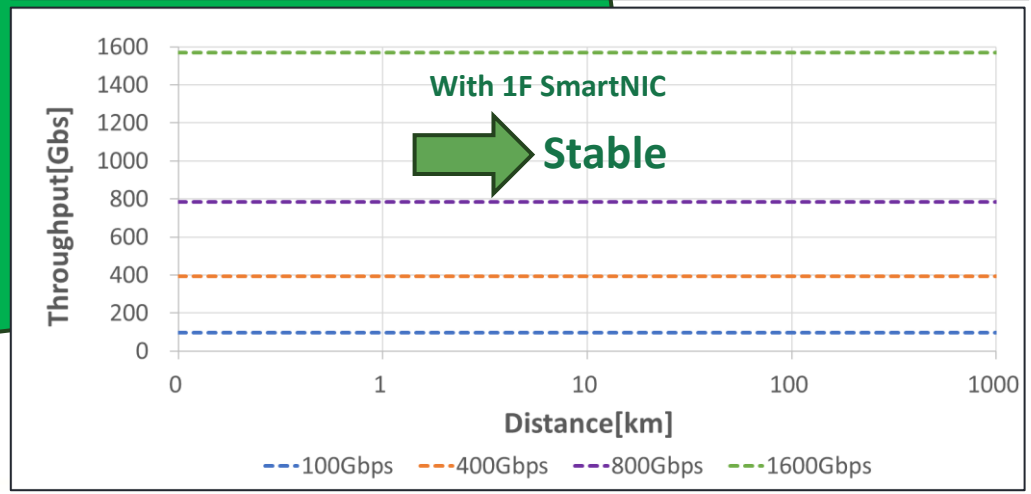

With 1Finity SmartNIC

(NeoCloud Data Center Interconnect - A giant supercomputer infrastructure)

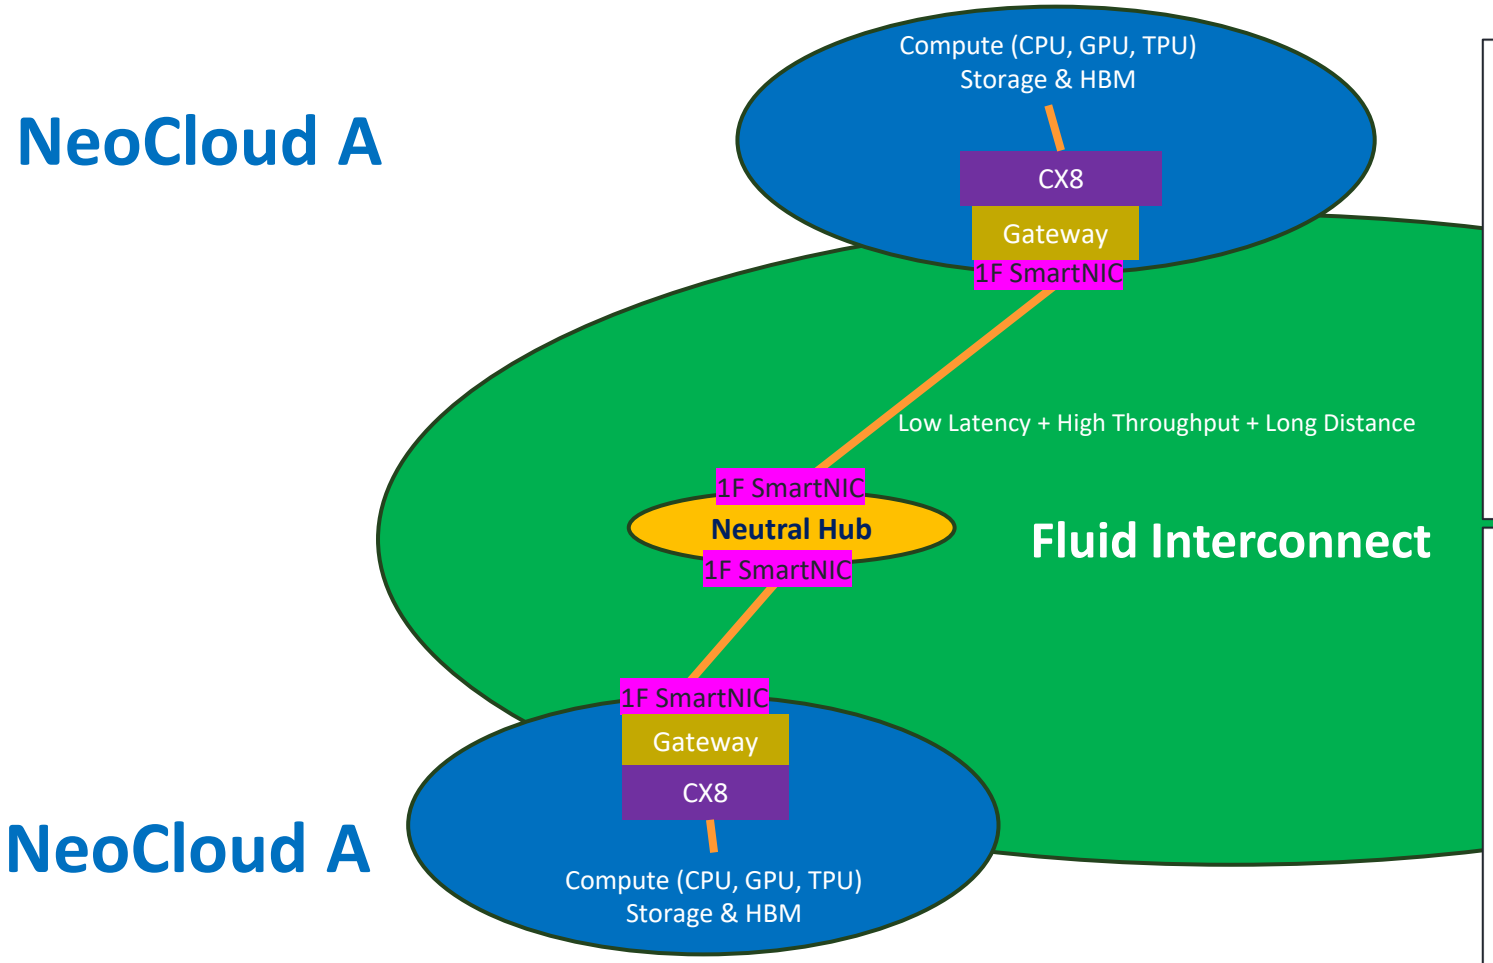

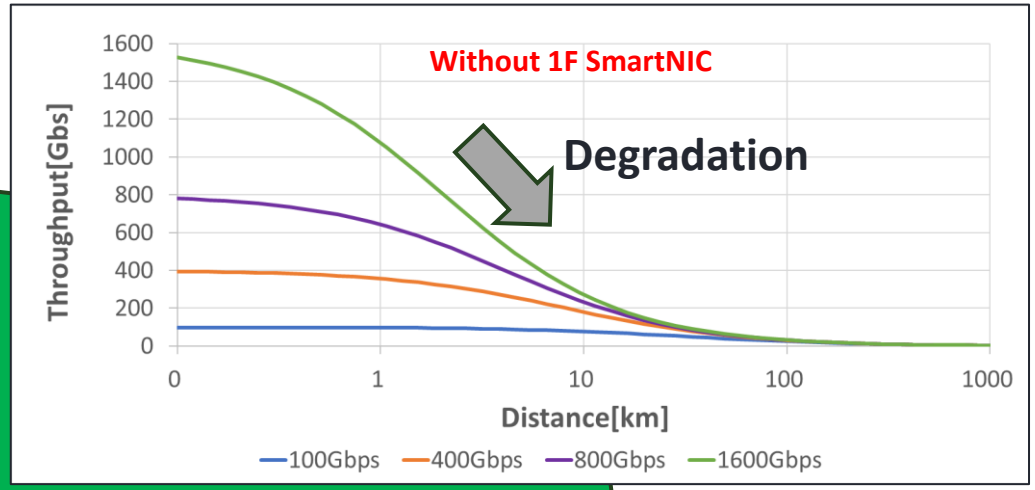

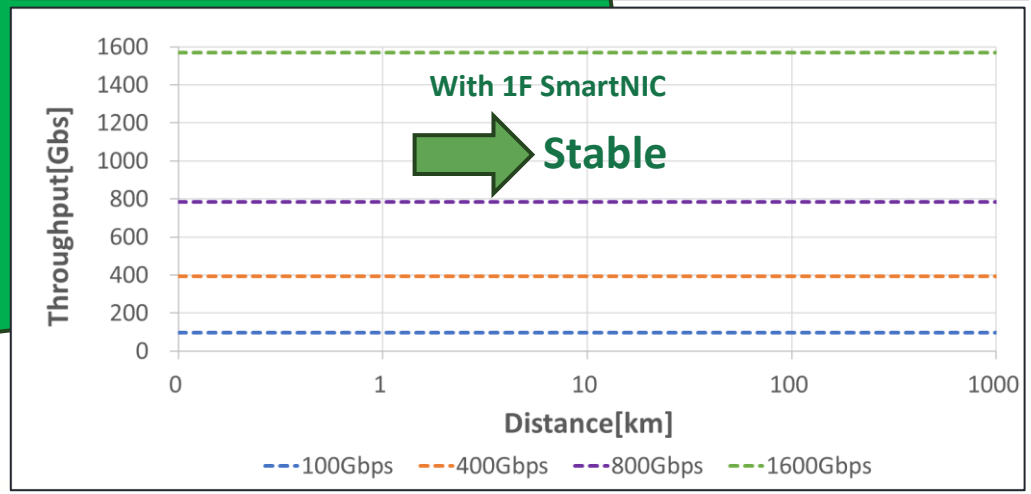

With 1Finity SmartNIC

(NeoCloud & Hyperscaler Interconnect)

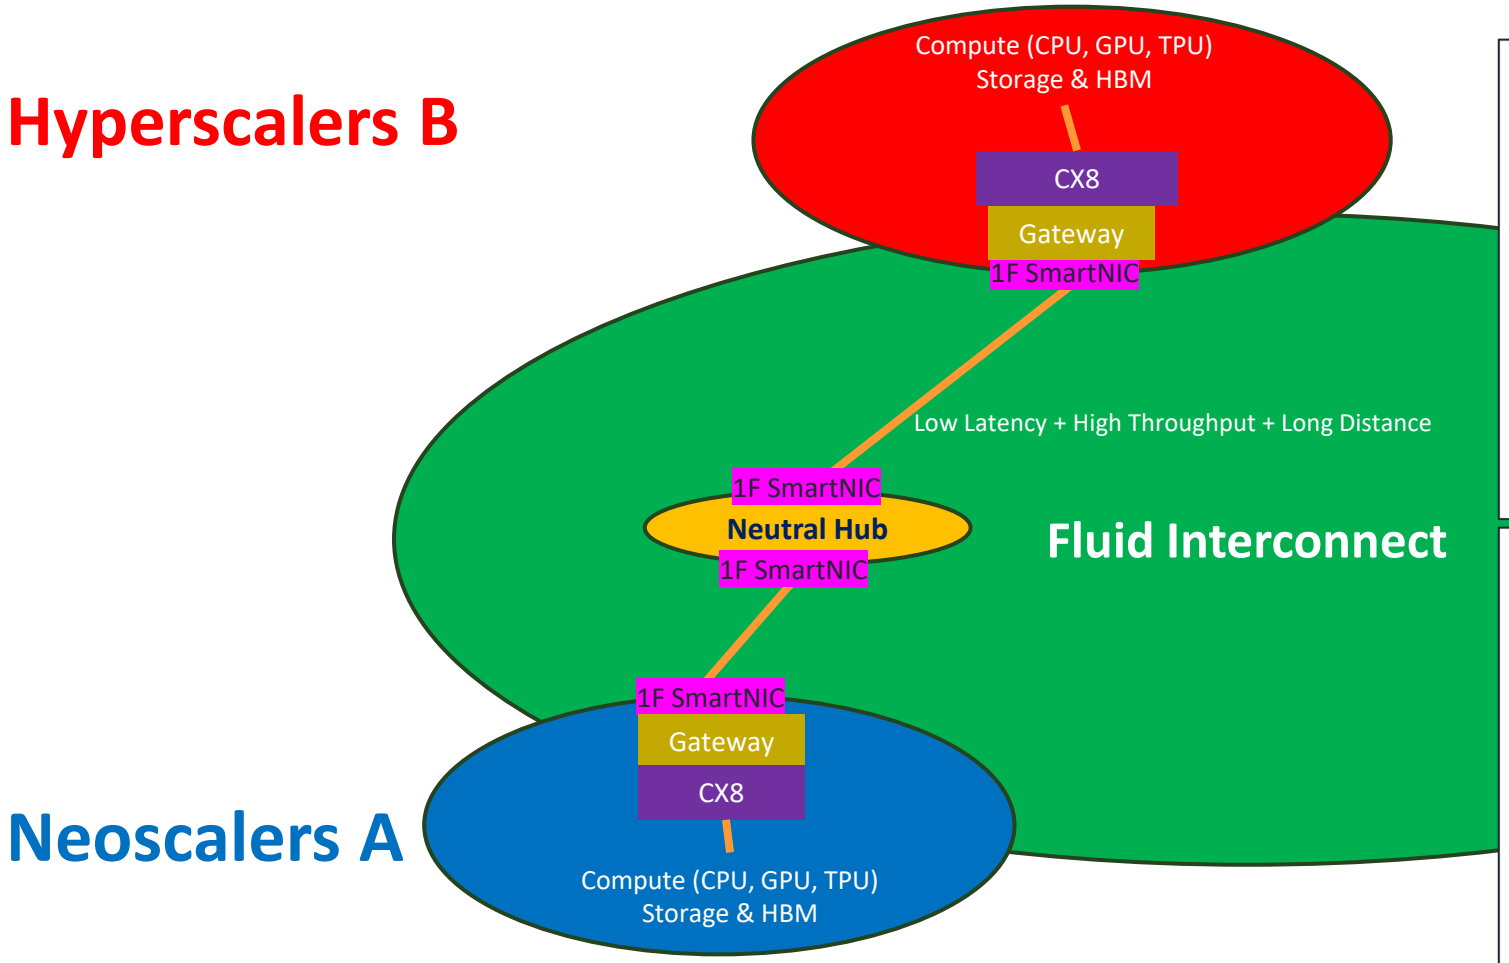

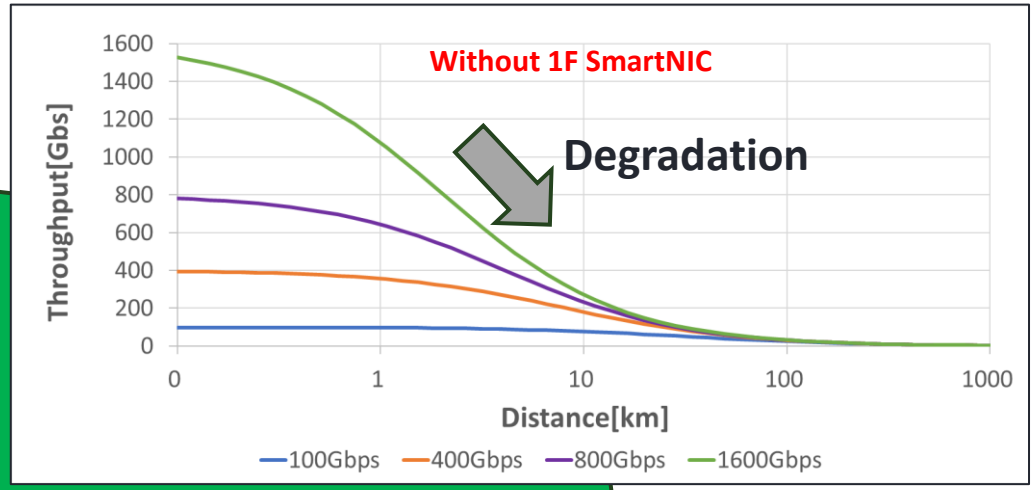

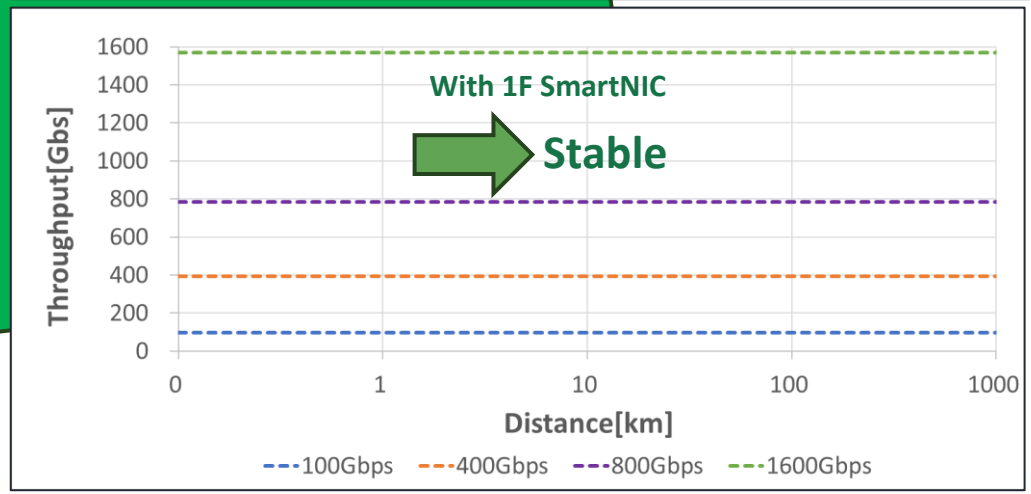

Emerging Solution Architecture

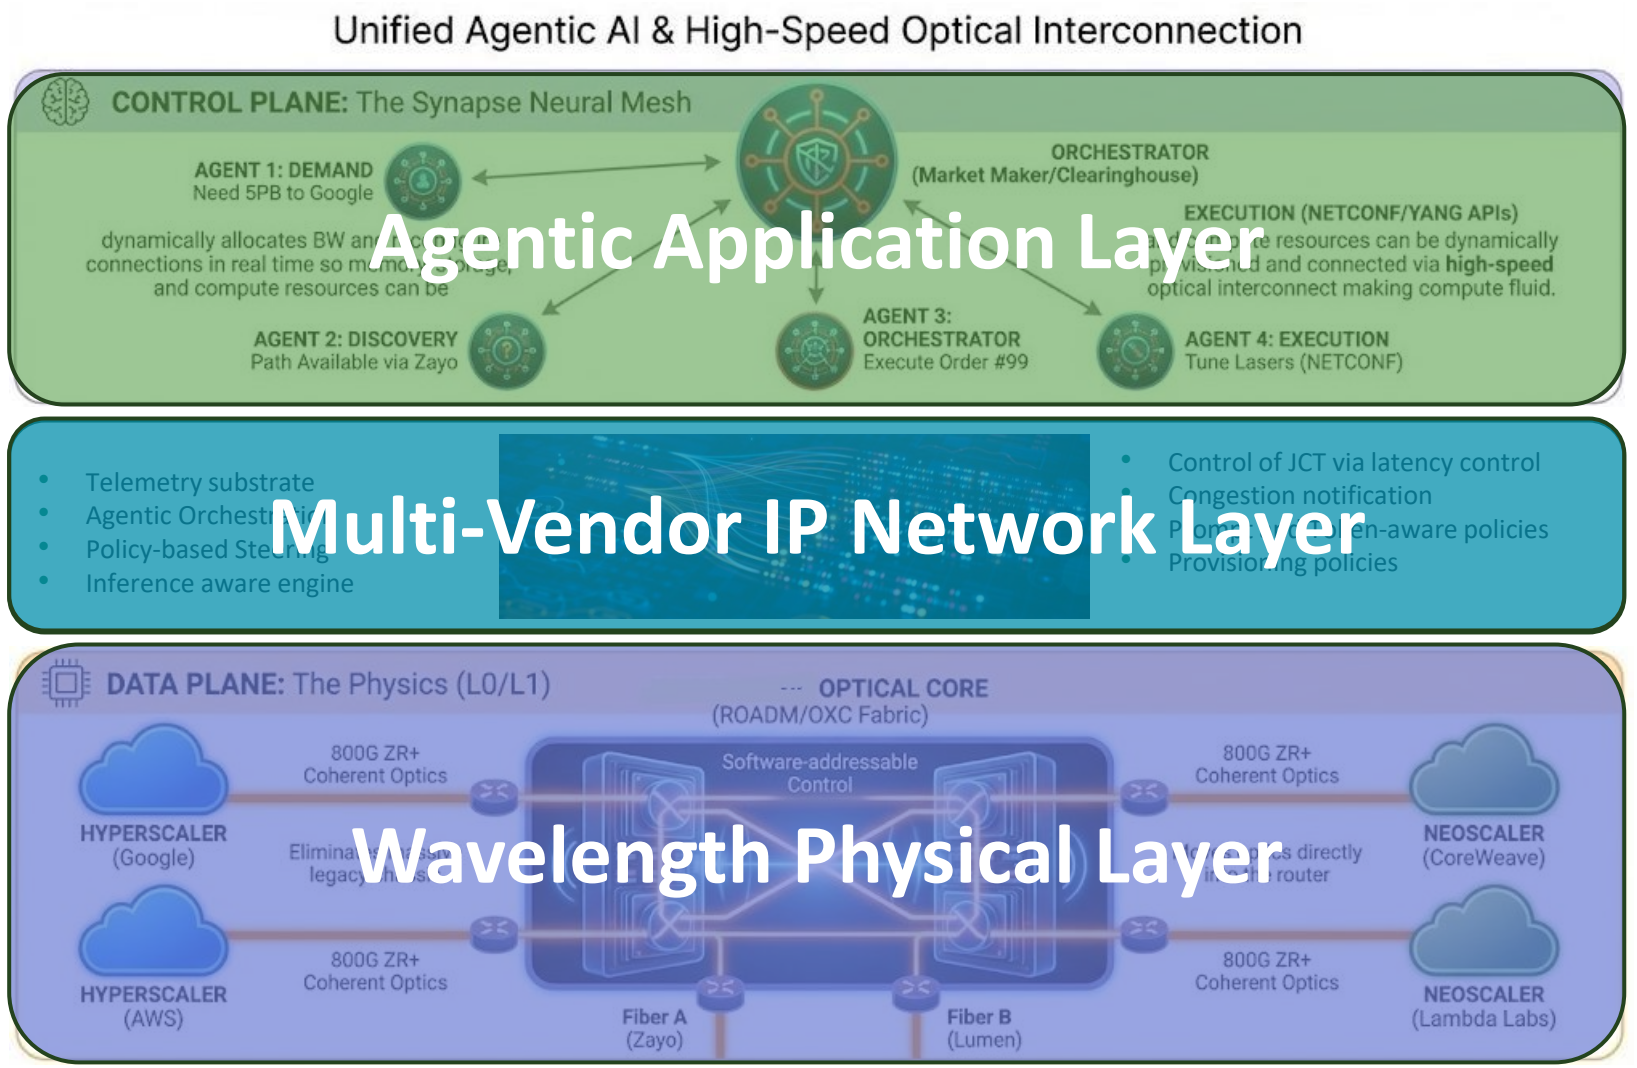

- Client SW
- Hub SW
- SI Services

- Peering Edge Router
- Provider Router
- Orchestrator
- AI Fabric

Access, Metro, Regional
- Open Line Systems
- Transponders
- Pluggables
- SmartNICs
- AI Automation
- Network Controller

Compute & Communication Convergence

Bridging the gap: The imperative of the Optical NIC

Critical Challenge

RDMA &
NVMe-oF

CPU

Datacenter (Intra-DC)

100km+ DCI link
(Inter-DC Connection)

CPU

CPU Overhead

Latency Spike

Throughput Degradation

Convergence Solution

Role of Optical NIC (Network Interface Card)

Computing/Network Convergence on NIC

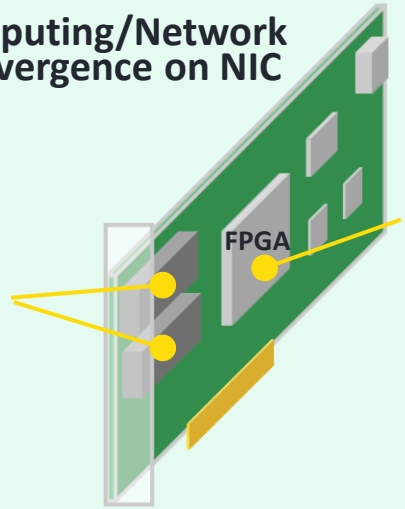

Optical NIC

Coherent Optical Pluggable Support

- ZR/ZR+ module

Function Integration (programmable)

- Long distance connectivity
- Continuation throughput and ultra low latency

WAN

Transforms WAN into local, high-speed backplane
WAN links become internal buses

Strategic Value

Enables "Virtual-DC"

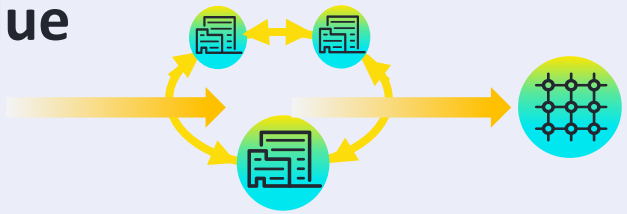

Positioning optical networking as the core element of AI computing

Flexibly Adaptable Networking Technologies

2023	2024	2025	2026	2027	2028	2029	2030
Ultra Ethernet Consortium	ULTRA ACCELERATOR LINK™ UALink™	ESUN (Ethernet for Scale-Up Networking)	HMO (Hierarchical Multi-domain Orchestration)			CAN (Cognitive AI-native Networking)	
Scale-out (Server-to-server / Rack-to-rack)	**Scale-up** (Accelerator-to-accelerator)	**Scale-up** (Accelerator-to-accelerator)	**Scale-Across, and beyond** Flexibly adaptable networking technologies are expected to emerge to meet the rapidly evolving requirements				

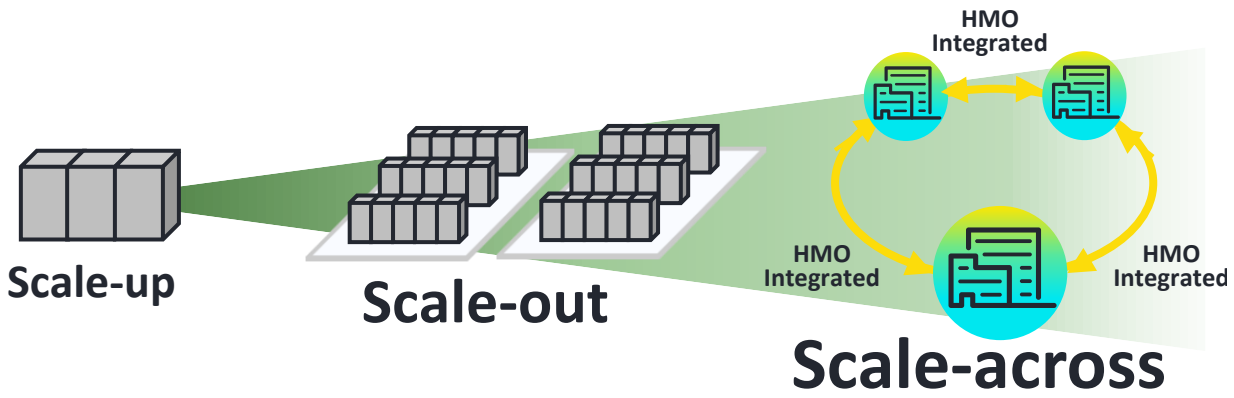

Next Generation: Cognitive / AI-native

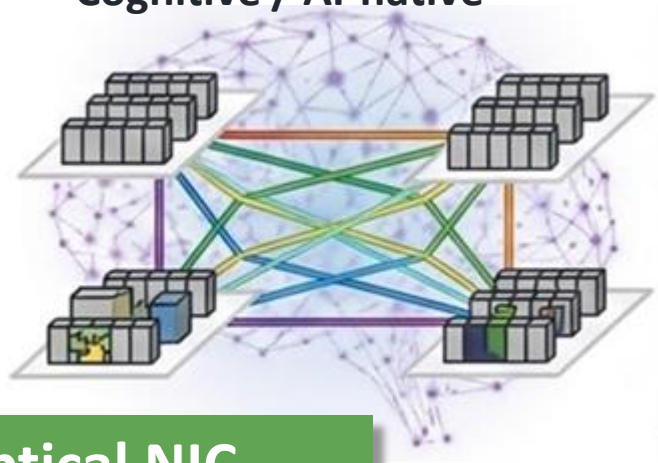

Core Requirements

Next-Gen Capabilities

Key Technology: Optical NIC

Challenge

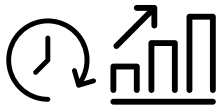

- **Latency accumulates** across multiple protocol layers.

- Queueing and buffering introduce **jitter**.

- **Maintaining throughput over long distance** is difficult at scale.

- Performance tuning **complicates deployment** and operations.

- Adapting infrastructure **to new workloads often requires redesigning** the entire network stack.

Need: Programmable Photonic Smart NIC

Lower latency

- **Lower latency**
  - Use a simpler protocol path with optical network basis.

Lower jitter

- **Lower jitter**
  - Reduce delay variation caused by buffering in electrical switches.

Long-Reach Throughput

- **Long-Reach Throughput**
  - Maintain high throughput even over long distances.

Simpler deployment

- **Simpler deployment**
  - Reduce the complexity of system deployment and operation.

Use case Flexibility

- **Use case Flexibility**
  - Enable faster adaptation to a wide range of use cases.

Core Technology 2: Transport Acceleration

Long-Reach Throughput

Conventional

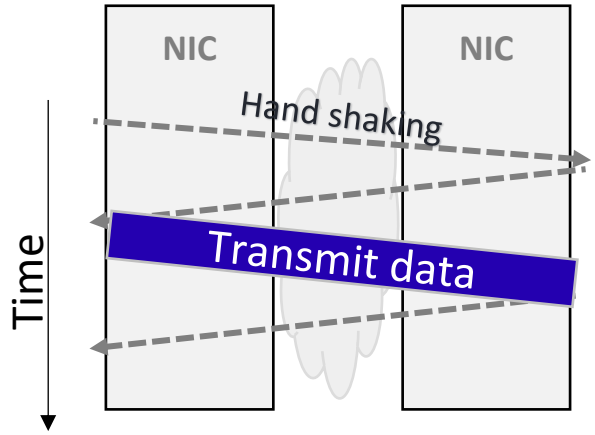

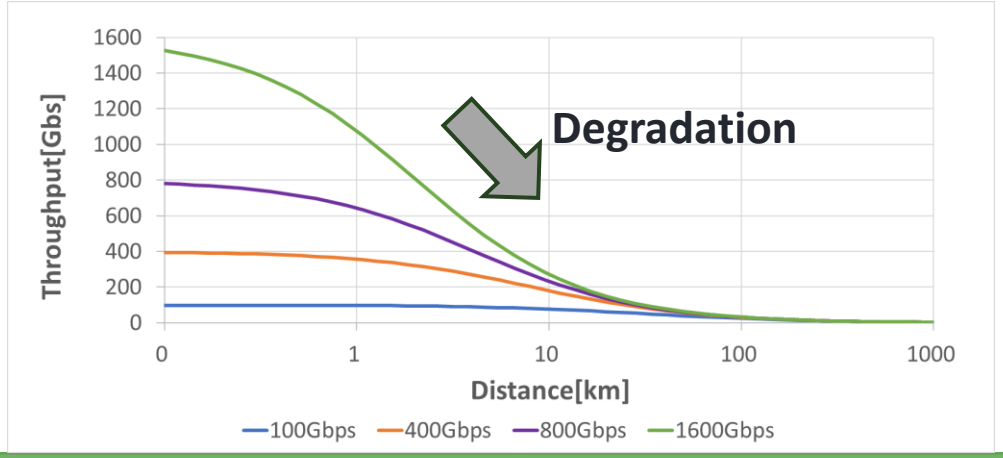

- **Higher line rate** and **longer distances reduce throughput.**
- **Even intra-DC is significantly affected at the 1.6T NIC generation.**

1Finity Solution

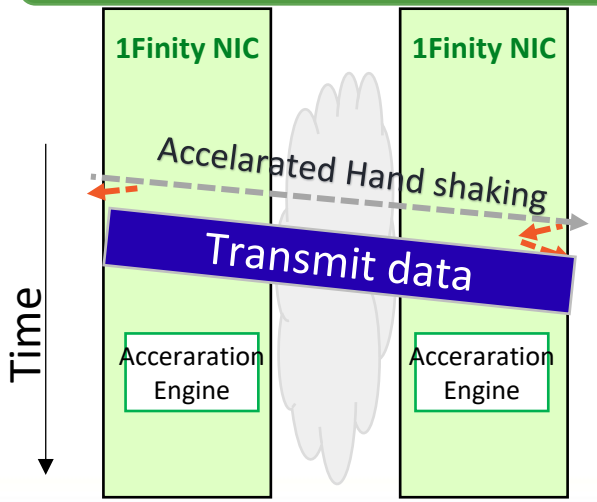

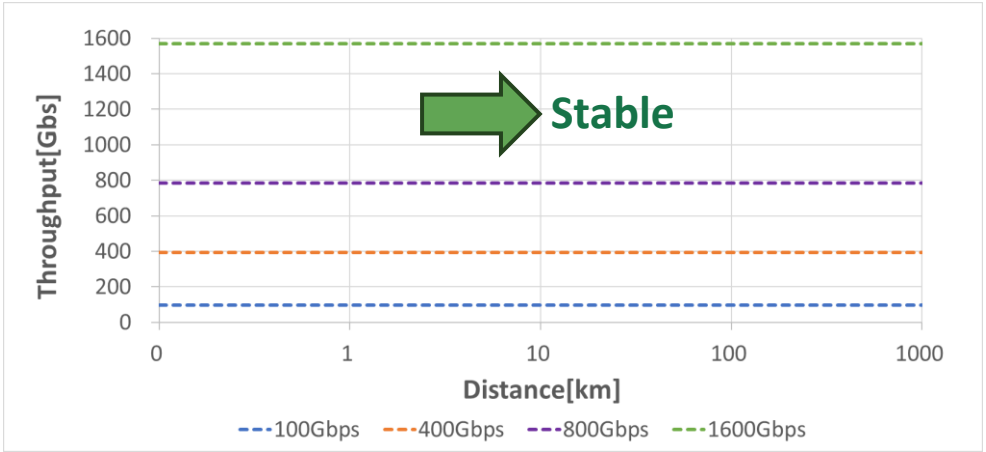

- **No degradation**
  - Handshake-free, ultra-low-latency transfer
  - Removes two-RTT setup overhead.

Core Technology: FPGA The Programmable Advantage

Usecase Flexibility

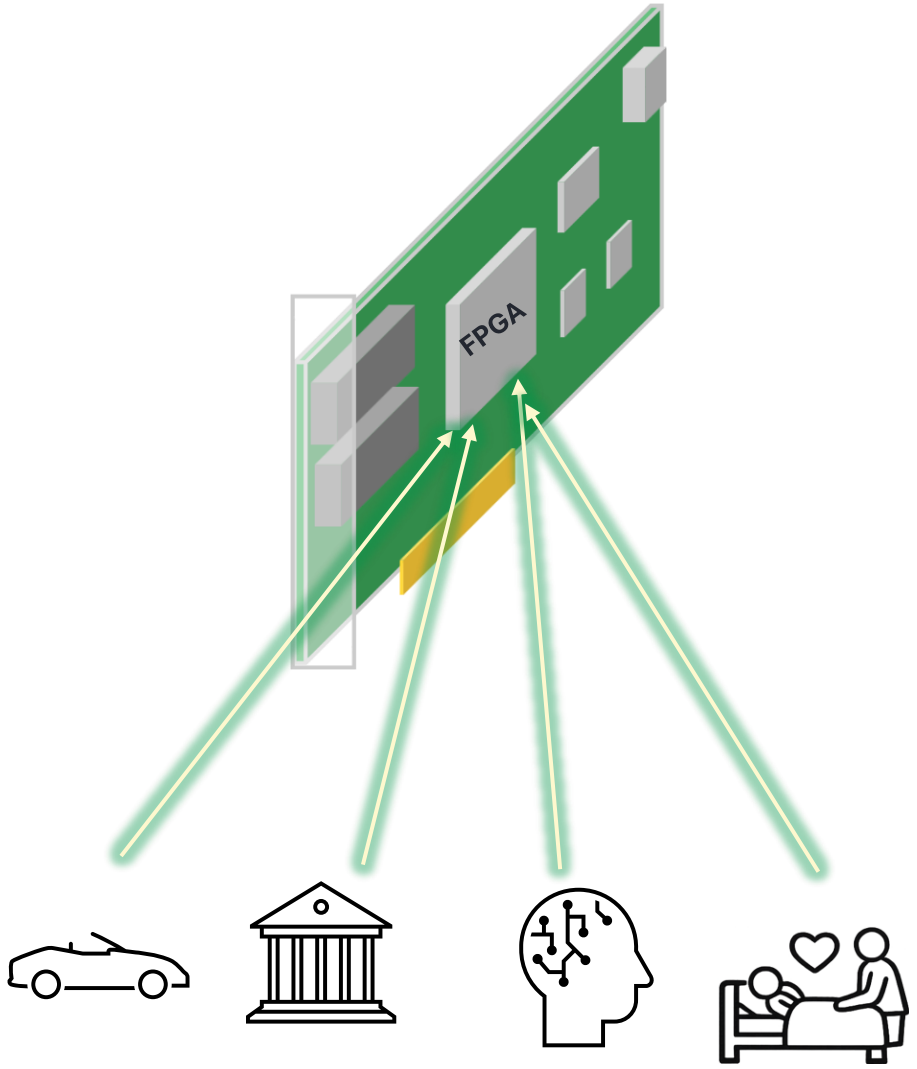

- **Agile Protocol Implementation**
  - Immediate hardware deployment for pre-standard specs and congestion control.
- **800G Ultra-High-Speed Processing**
  - Flexible, zero-latency processing even at 800Gbps bandwidth.
- **Dynamic AI & Topology Adaptation**
  - Seamlessly tracks AI workloads and network changes via data/control plane separation.
- **Smart CPU Offloading**
  - Accelerates the entire system by delegating network processing and freeing up CPU.

A Seamless Distributed Database Across Multiple Data Centers Enabled by a NIC Co-Designed with DB Middleware

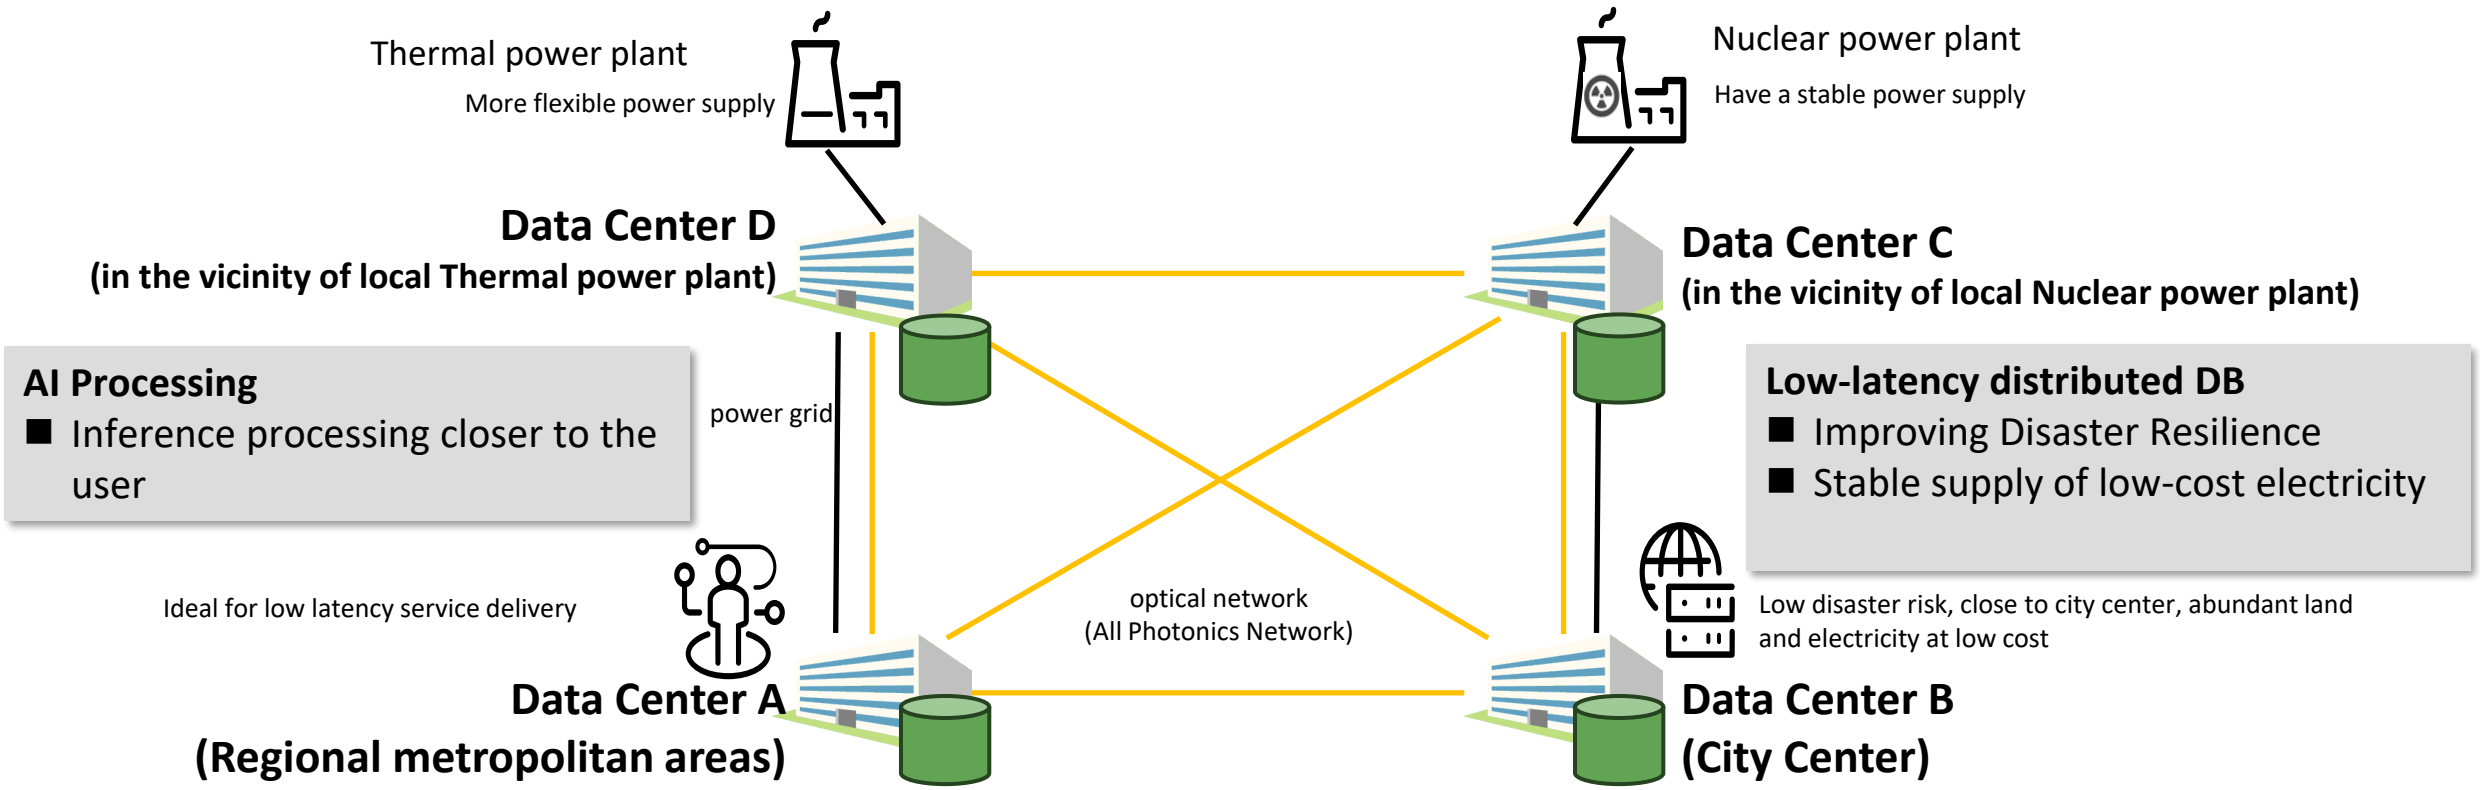

Gain the benefits of distributed data centers without complex user tuning.
Enabled by a NIC tightly coupled with distributed DB middleware.

1Finity Optical NIC

Network rate : up to 800G

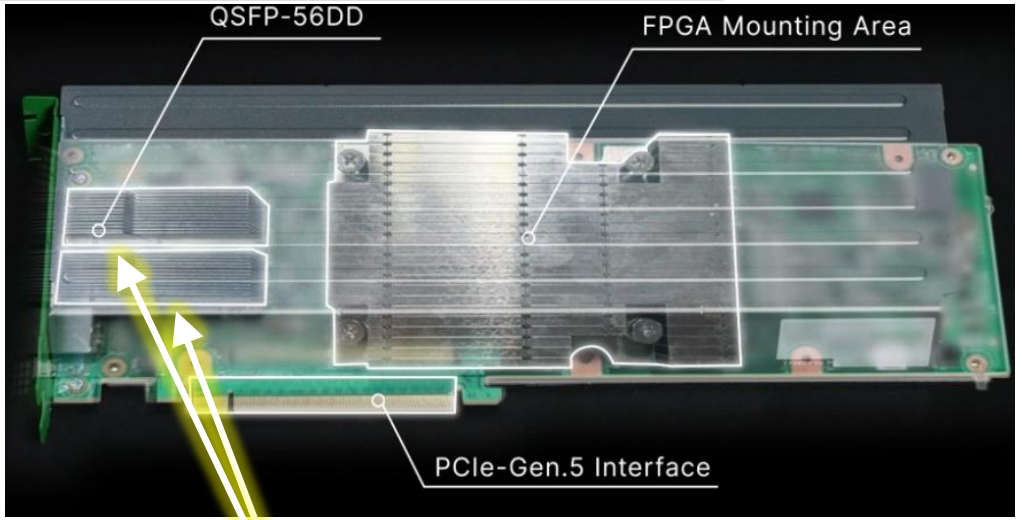

1Finity P300
800ZR+ Optics

- Signal Processer : **FPGA**
- We have started proof-of-concept trials with **multiple partners**.
- Leveraging the advantages of **FPGAs**, we optimize signal processing for a wide range of use cases.
- Coherent Optics can directly be used with this NIC

Performance Evaluation : High throughput

Conventional

Lab evaluation result by **Conventional NIC**

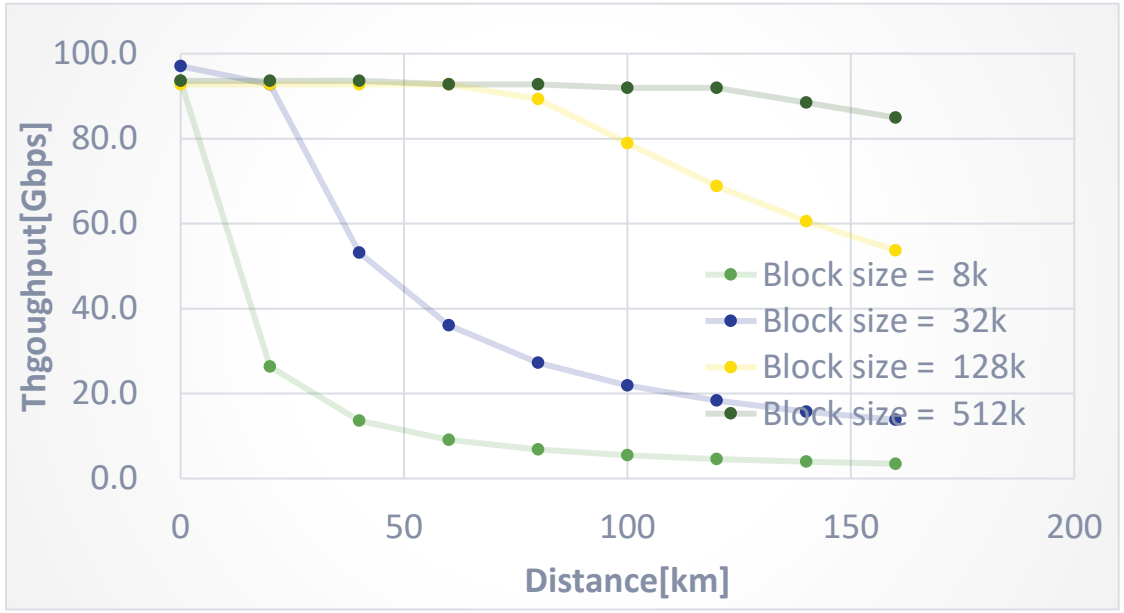

1Finity Solution

Lab evaluation result by **1Finity NIC**

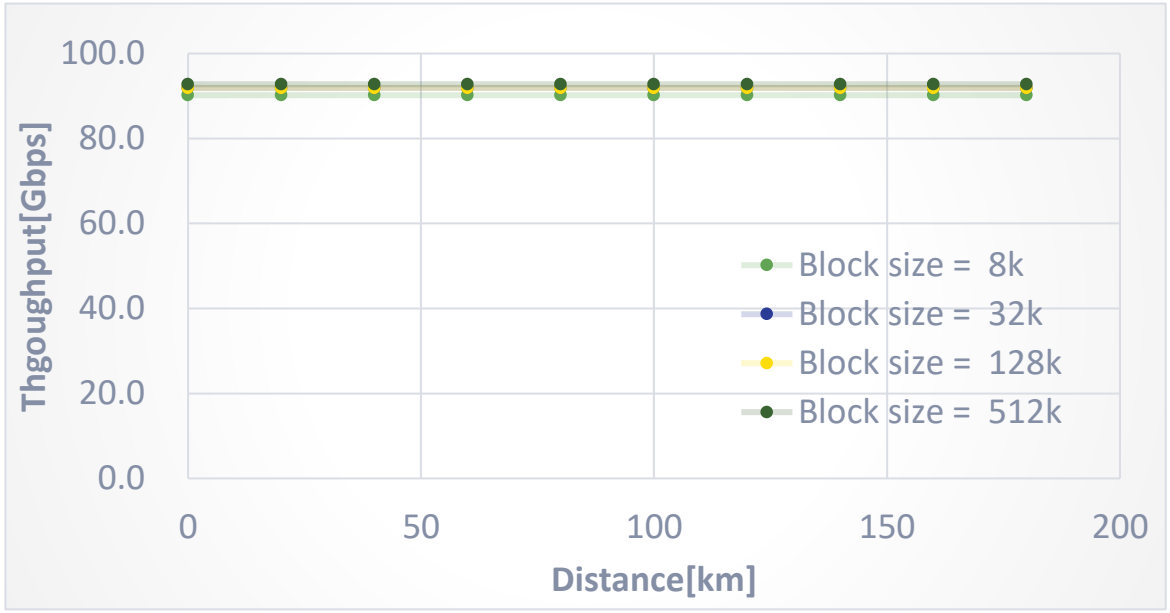

Experimental results demonstrate that
1Finity method sustains throughput
without measurable degradation.

Performance Evaluation : Low Latency-Delta

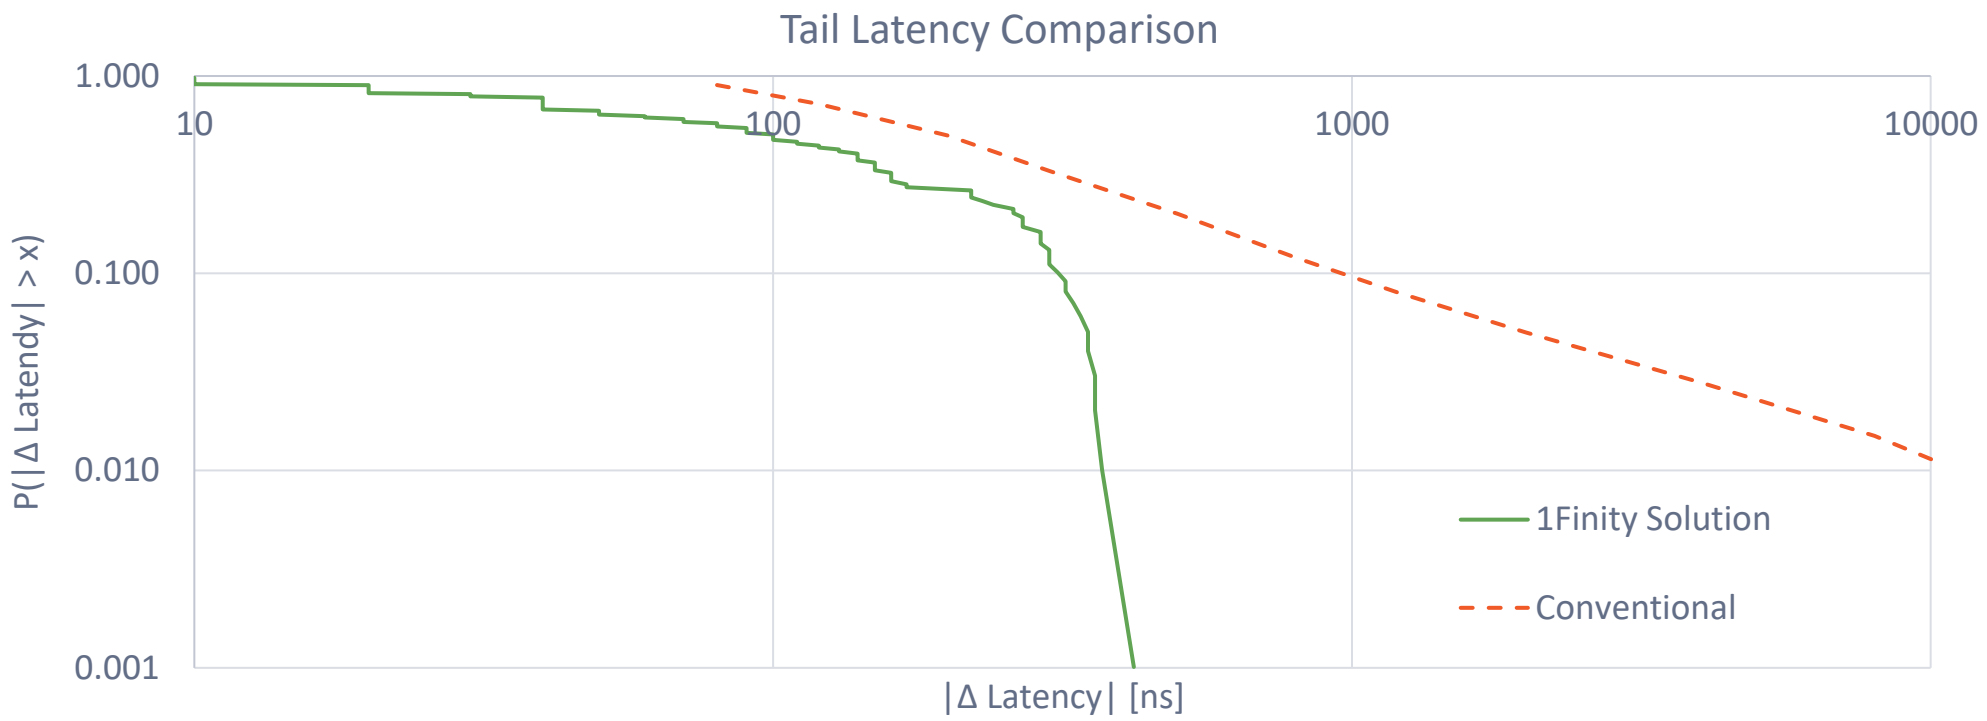

Published baseline shown as an approximate qualitative reference reconstructed from digitized data in:
A. Wolosewicz et al., INDIS 2025, Fig. 6(b). Original packet-level data are not publicly available.

1Finity method shows
very low latency-delta which can be very low tail latency

Reference: Standard Protocol Stack

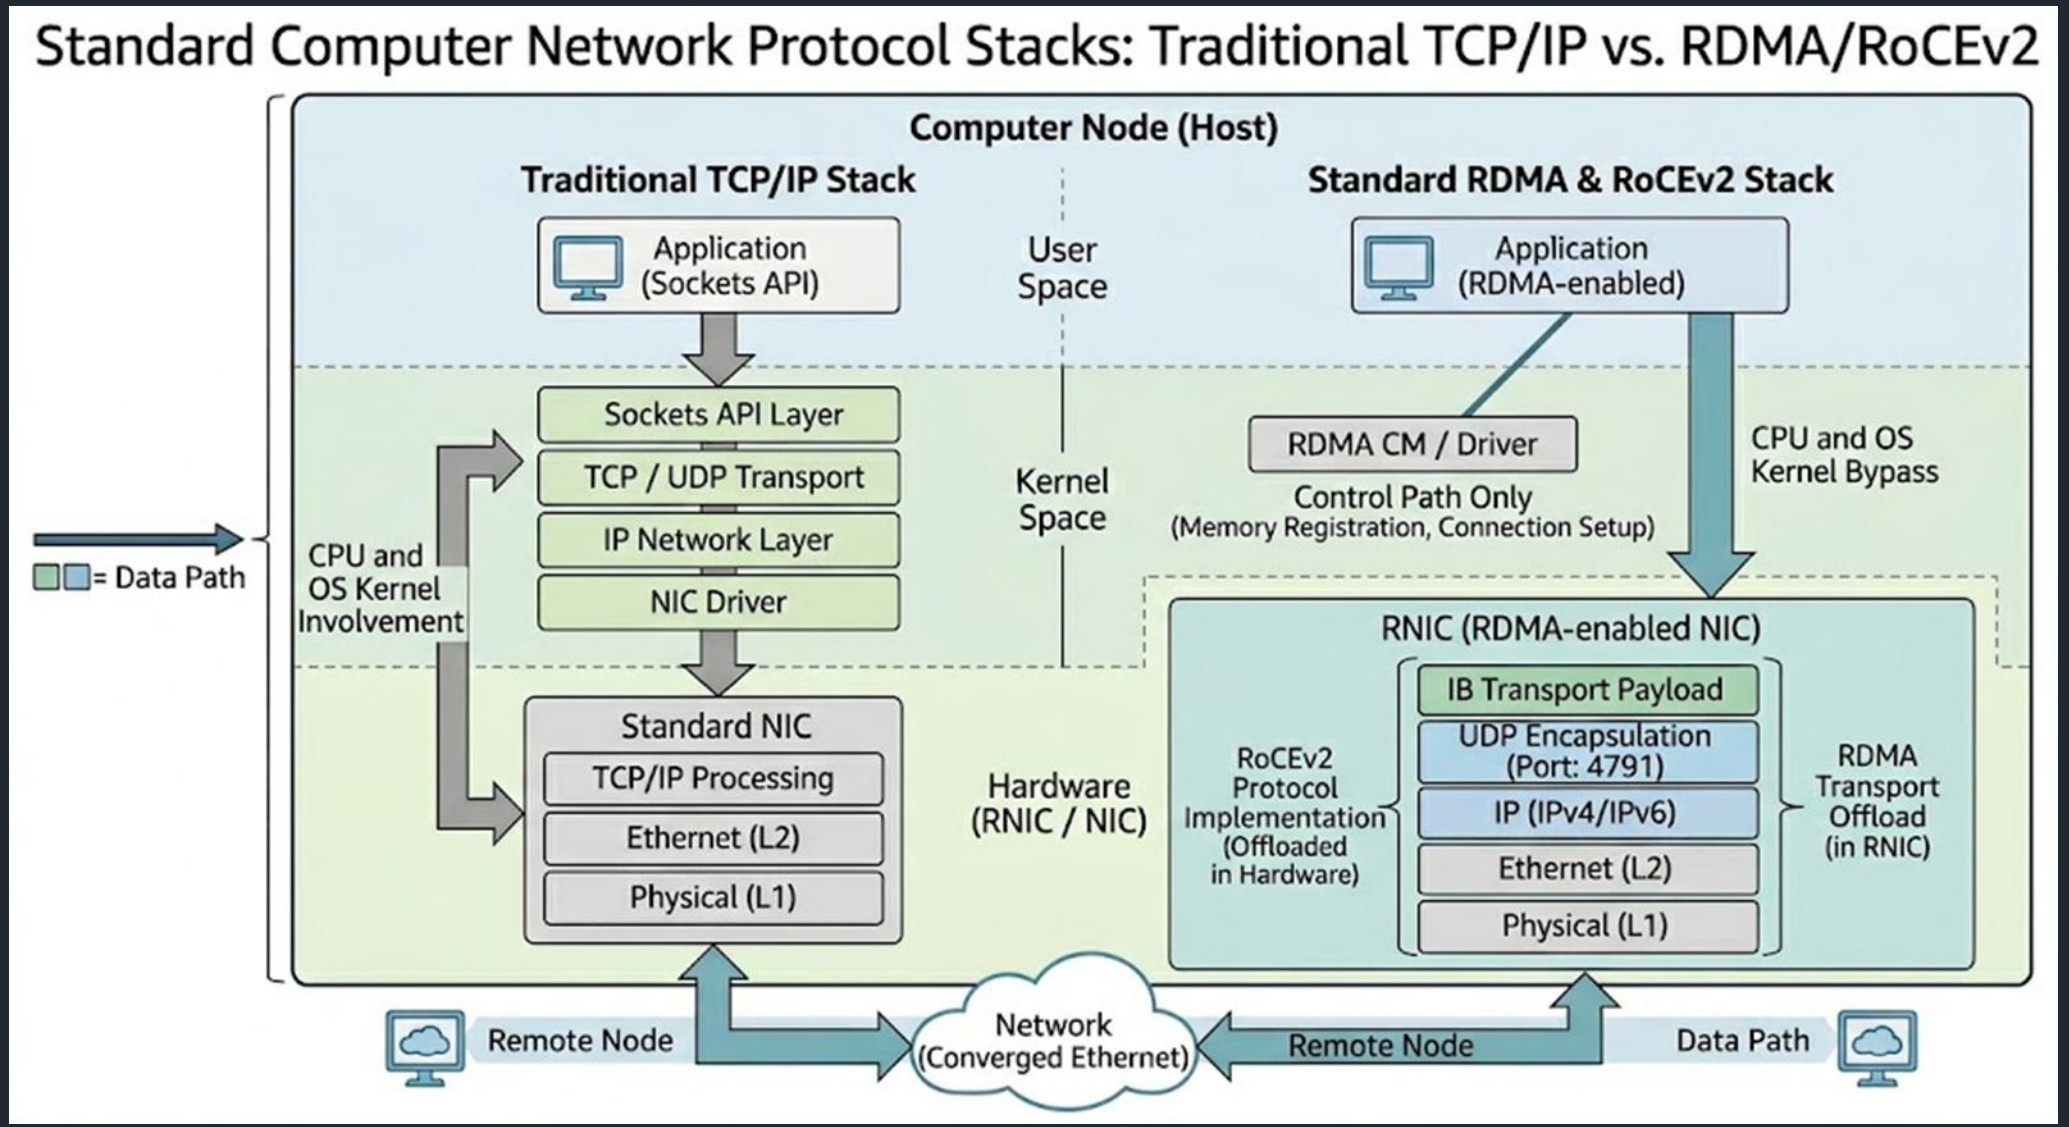

Reference: 1Finity NIC Protocol Stack

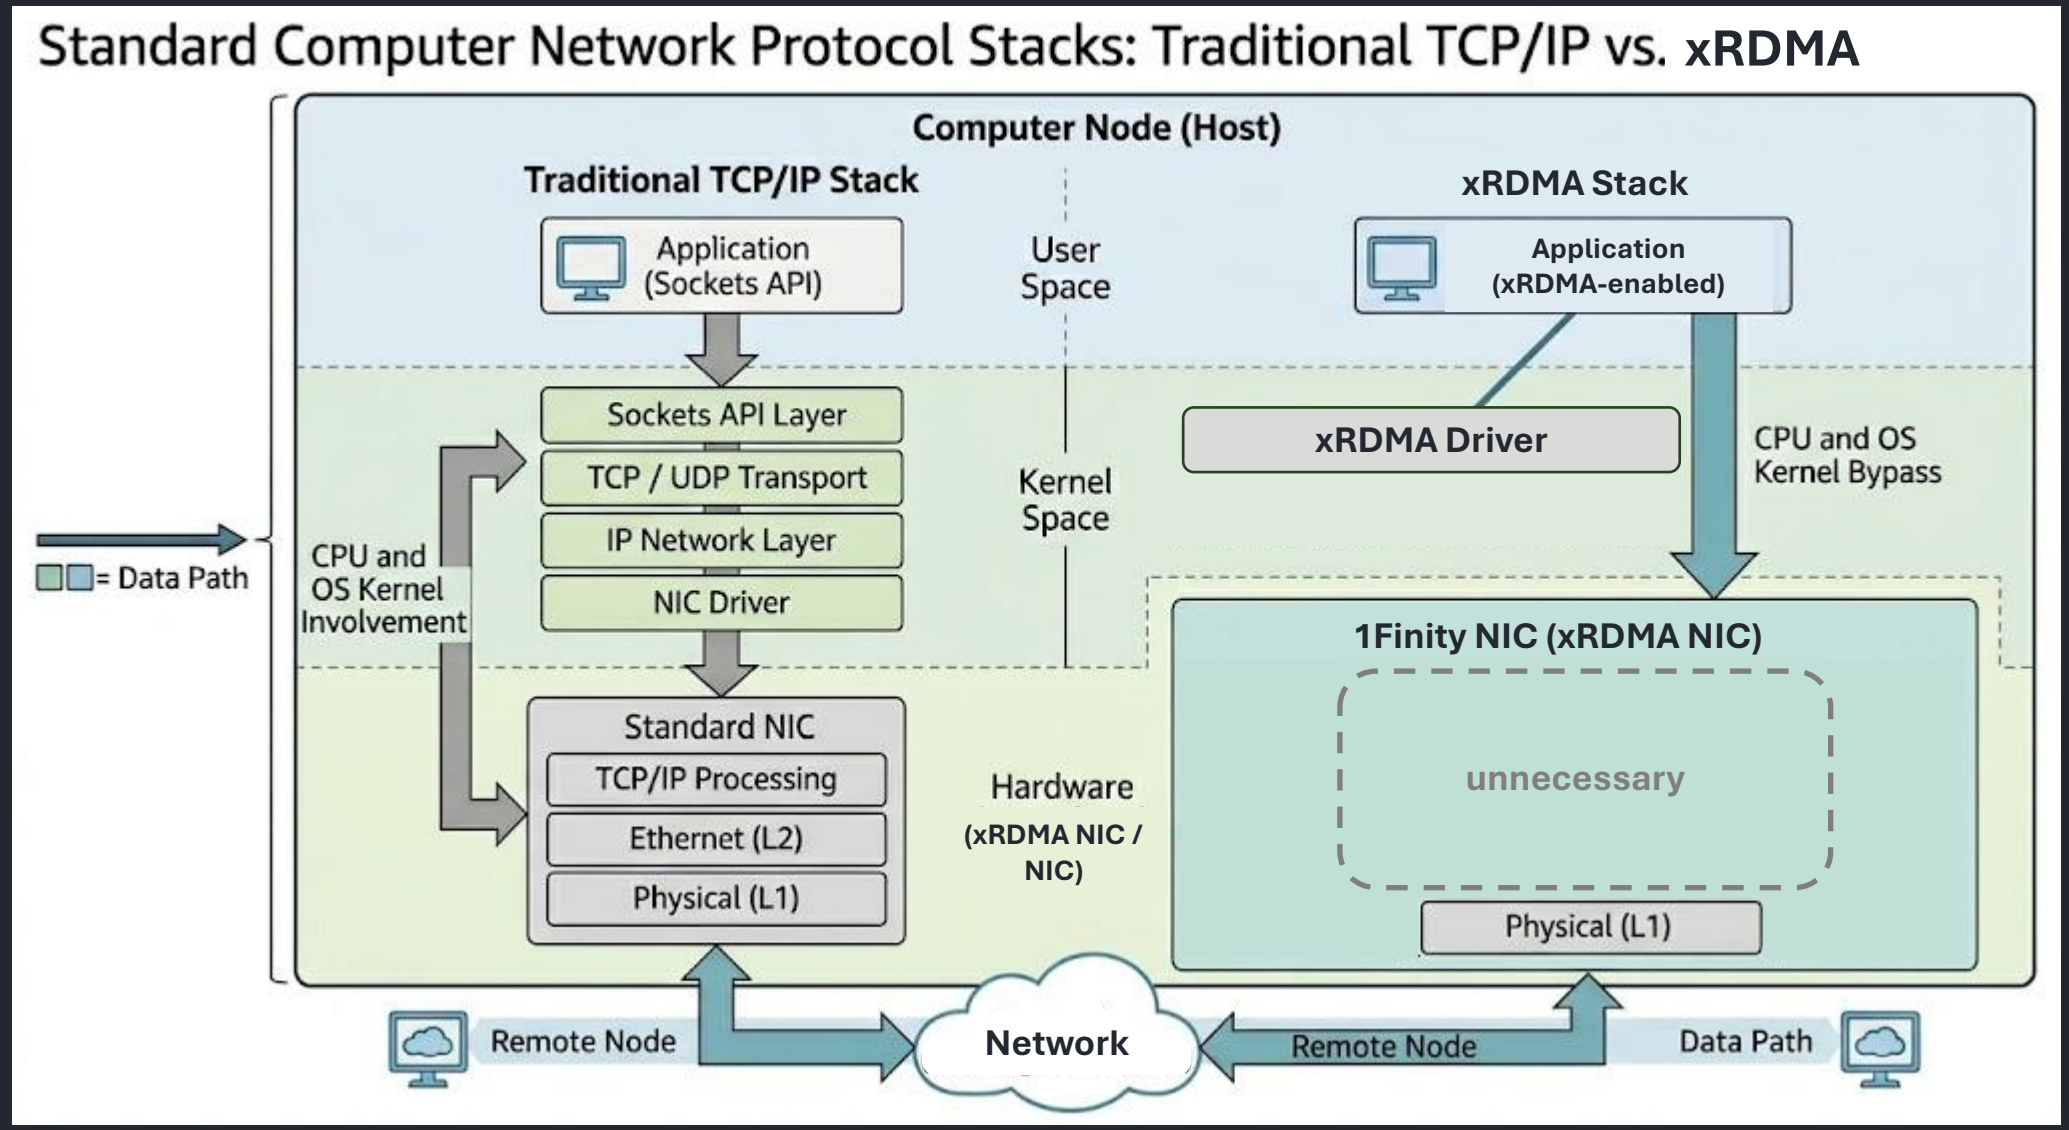

Core Technology 1: Simplified Transport Protocol

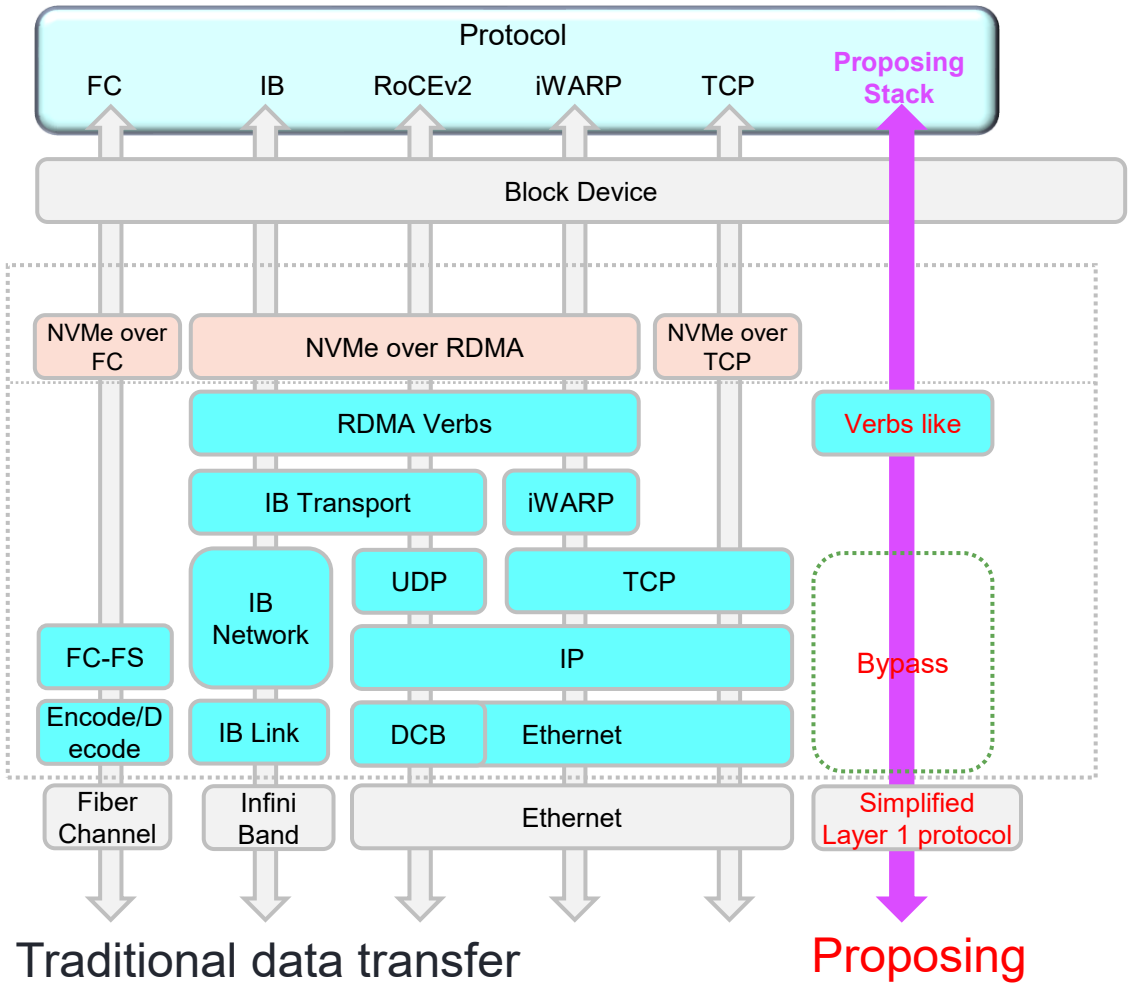

- **Protocol-stack bypass**
  - Removes TCP/UDP, IP, Ethernet, and IB transport overhead.
- **Simplified physical layer**
  - Lightweight Simplified Layer 1 protocol.
- **Verbs-like access**
  - RDMA-Verbs-style efficient access.
- **Vertically integrated direct path**
  - Straight-through bypass for ultra-low latency.

1Finity NIC & Broadcom TH6 + CPO

1Finity NIC (Optical Network Interface Card)

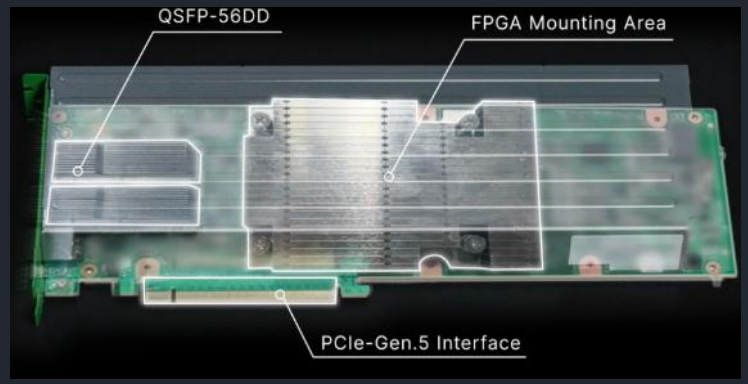

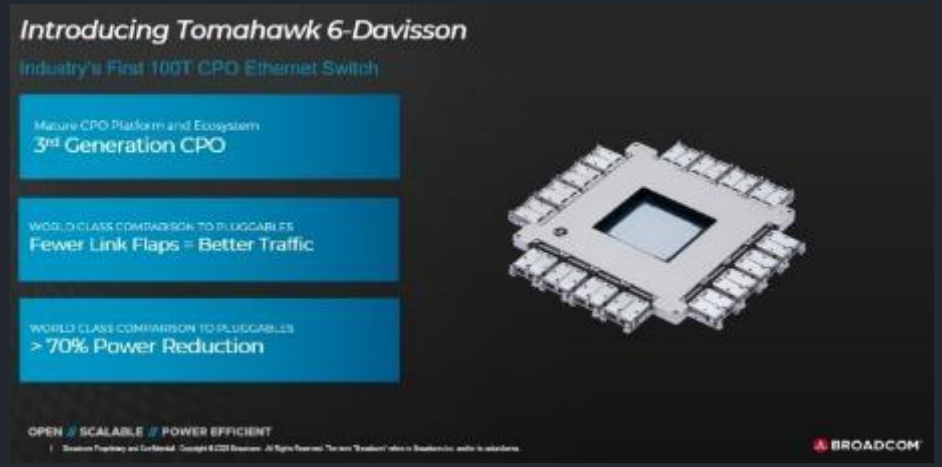

Architectural Positioning (Conceptual Diagram)

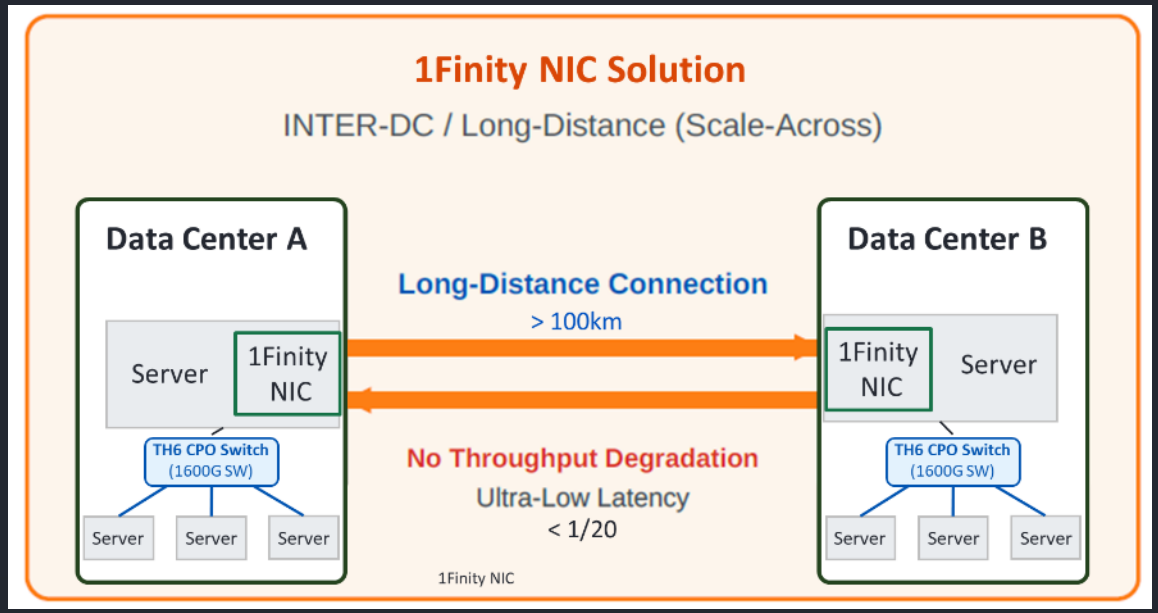

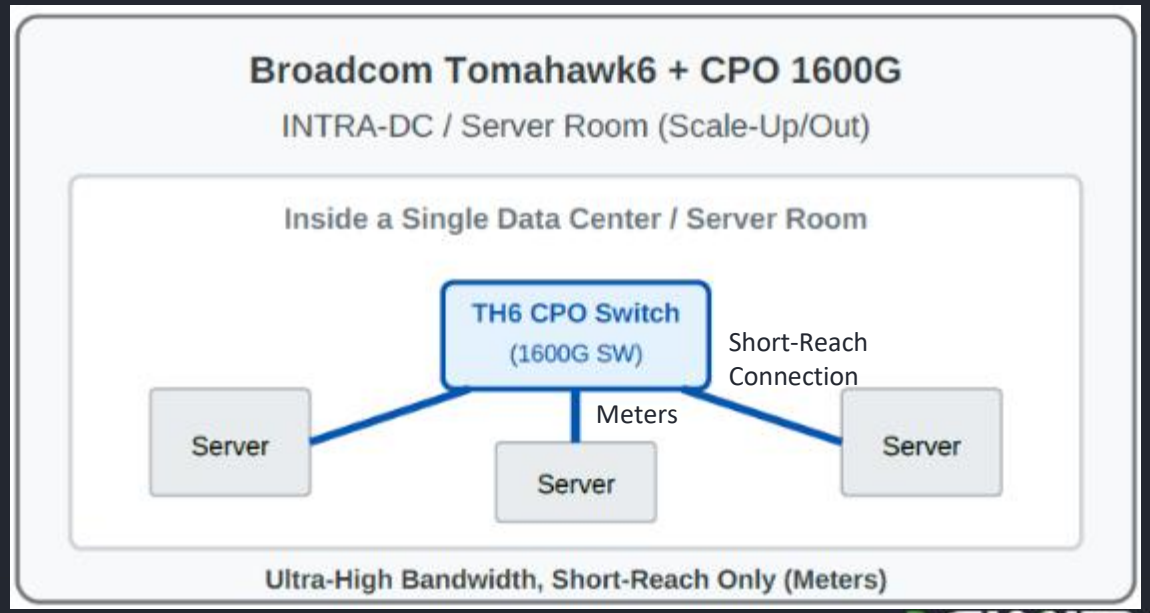

1Finity NIC & Nvidia BlueField/ConnectX

1Finity NIC (Optical Network Interface Card)

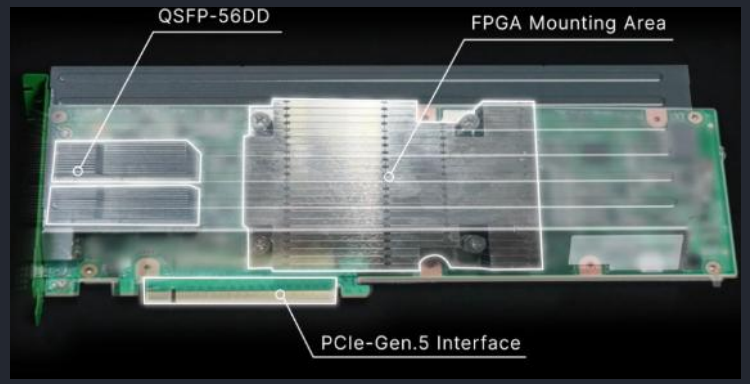

Nvidia ConnectX-7

Architectural Positioning (Conceptual Diagram)

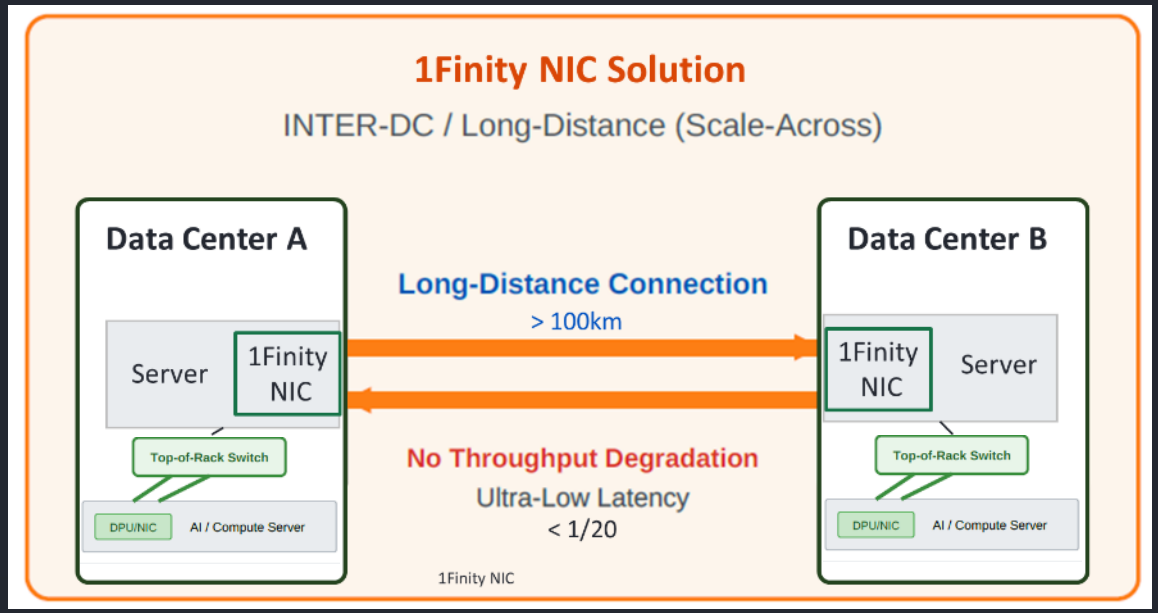

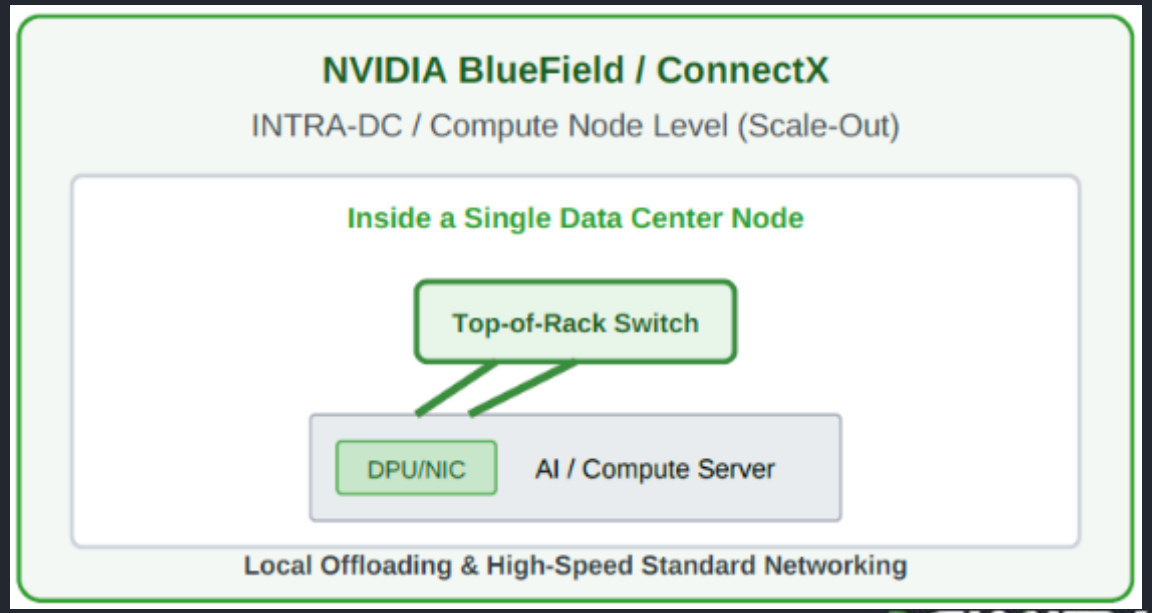

Ecosystem with Global Partners

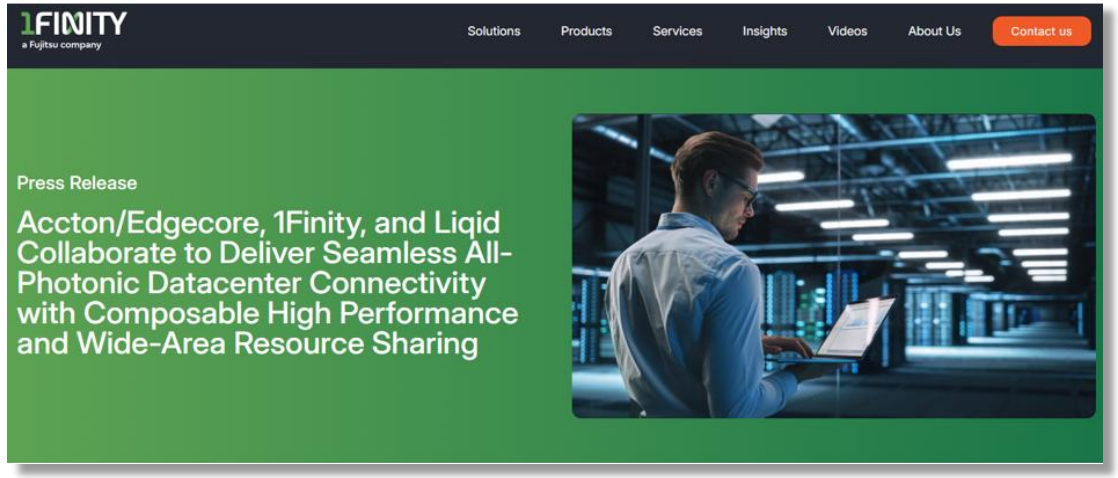

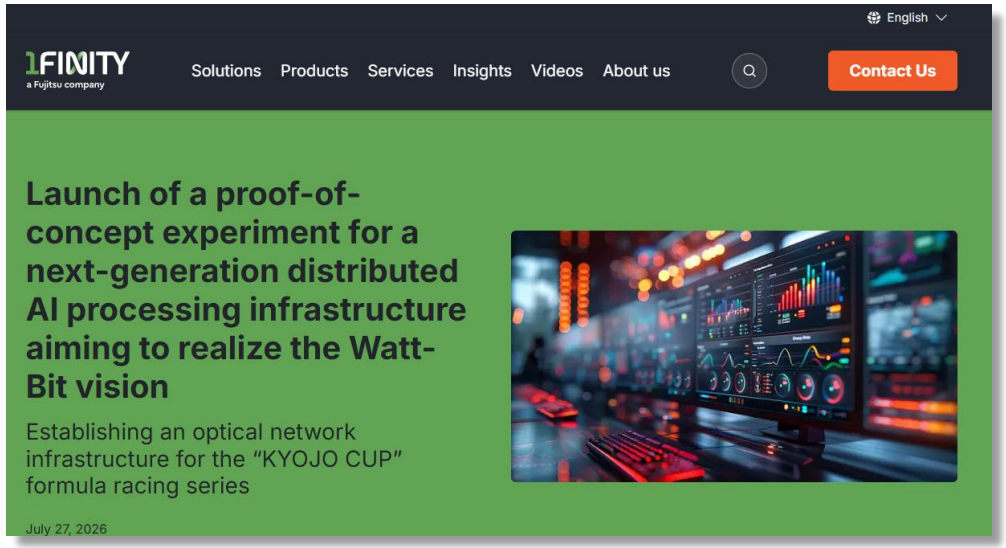

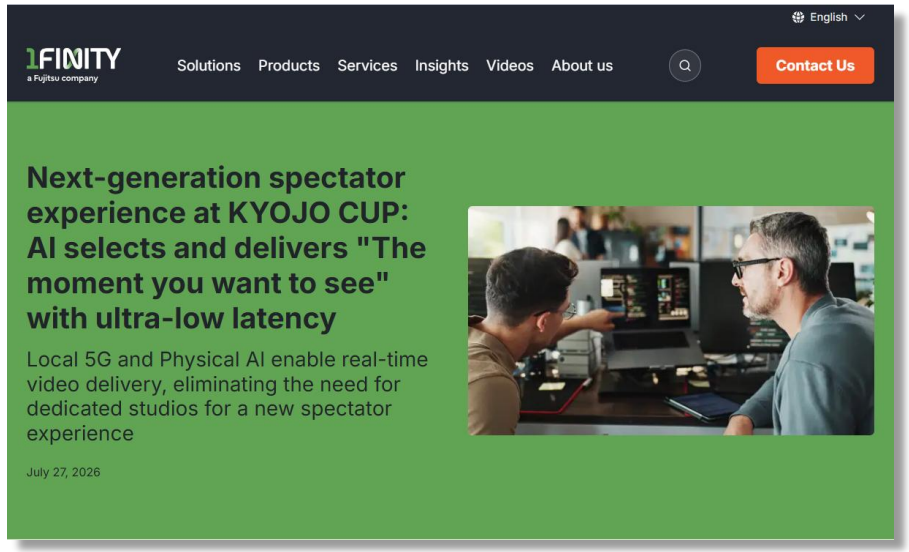

Building an ecosystem across diverse industries for real-world deployment

North American Expansion Through Partnerships with Accton, Edgecore, Liqid, and Others

- Establishing and advancing partnerships for North American expansion

Ecosystem

Accton Making Partnership Work

Edgecore Networks

LIQID

1FINITY

Integrated with Nexvec (Open Fabric)

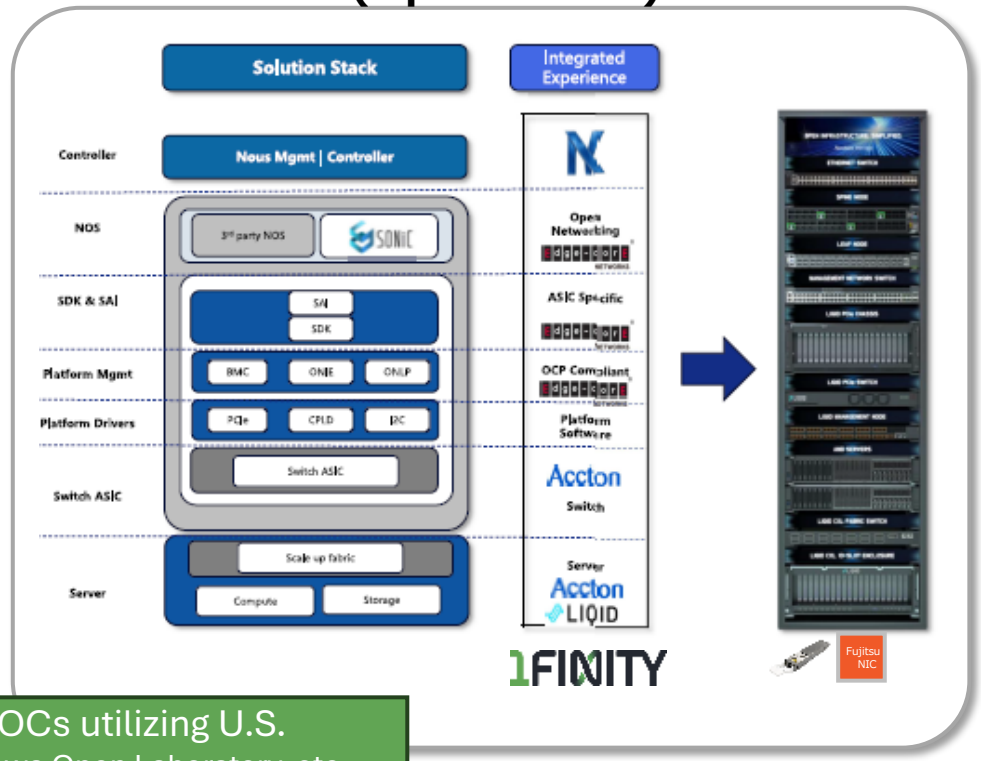

Considering joint POCs utilizing U.S.
ONIC@Sunnyvale, the Okinawa Open Laboratory, etc.
Joint POCs and customer demos utilizing these facilities are under consideration

In progress

Plan of Approach

Customer Proposal

Japanese Potential Customer

NTT

KDDI (Tomorrow, Together)

SoftBank

Rakuten Mobile

SAKURA internet

IIJ (Internet Initiative Japan)

QTnet

OPTAGE (What's next?)

@Tokyo

MUFG

SMBC

INTEC (TIS INTEC Group)

TOYOTA

SUPER FORMULA

NeoClouds +

US Target Customer

Google

Microsoft

amazon

Meta

EQUINIX

AT&T

verizon

Akamai

Looking Ahead

An ecosystem is forming around 1Finity, and together we are turning the vision of Distributed AI infrastructure into reality.

1FINITY
a Fujitsu company

Call For Action: Inviting ecosystem partners for a Proof of Value (POV)

Plesae Contact: *Solyman.ashrafi@fujitsu.com*