



Co-Designing AI Fabrics: Inside Google's Virgo Network and the Campus-as-a-Computer


Sunil Ahluwalia *Group Product Manager, Google*

Jeongkeun “JK” Lee *[Software Engineer, Google]*



OCP
PLATINUM
Member

25 years of Network Innovations




OCP
GLOBAL
SUMMIT

2026

Scaling Innovation for the AI Era.


AI Workloads are Network Intensive



Foundational models parameters growing exponentially

Training: Very compute- and network-intensive

Single datacenter → multi-data center → multi-campus



Exponential increase in scale-out network bandwidth

Latency sensitive workloads: strict delay bounds



OCP
GLOBAL
SUMMIT

2026

Scaling Innovation for the AI Era.

AI scale-out network design principles

Design for AI Hypercomputer

- Design for entire campus as a computer, Optimize for workload performance & reliability
- Deployment and operation lifecycle tightly coupled with the GPU/TPU

Design for scalability and reliability

- Training job should not crash for a single link or switch failure
- Telemetry is key to draw the best perf and availability

Design for speed

- Minimize the networking materials to deploy → hence, minimal layers and dense rack design




OCP
GLOBAL
SUMMIT

2026


Scaling Innovation for the AI Era.

Reimagining the data center network


Independent scale-out accelerator backend and front end networks



Virgo: Scale-out data center fabric



Virgo Architecture



Multi-plane topology

- SERDES-level “Skinny links”
- $N = \# \text{ planes} \sim \# \text{ of skinny links per accelerator}$
- One plane down $\rightarrow (N-1)/N$ BW available
- Multi-plane scaling w/o 3rd layer

Routing

- Packet level load-balancing
- Switch-driven adaptive routing with fast reaction to failures
- SRv6 source routing for bare-metal customers running their own host stack

Co-design Host and Fabric

End-to-End Technologies

Congestion Signalling (CSIG)

An in-band signaling protocol that gives end-hosts visibility into fine-grained switch congestion signals.

High-speed Switch Telemetry (HST)

Very fine-grained (sub-millisecond) visibility into network switch behavior, like port and queue status.

Firefly

u-sec level clock synchronization of endpoint NICs, critical for timestamping and debugging

Actionable Insights for AI Workloads

- ◆ Which communication-pairs in my workload are **starved for bandwidth**?
- ◆ Why is my network transfer slow (and **where is the bottleneck**) ?
- ◆ Does my workload have bandwidth **headroom** for more communication?
- ◆ Is my fabric **adequately provisioned** for the models I'm running?



OCP
GLOBAL
SUMMIT

2026

Scaling Innovation for the AI Era.

Google Falcon Transport on NVIDIA Networking

From 2023 OCP contribution to silicon-to-software co-design across any Ethernet network



Multipath Load Balancing

Sprays packets evenly



Programmable Congestion Control

Sub- μ s Swift & Congestion Signaling



Hardware Loss Recovery

Selective ACKs & RACK



Per-Connection Security

Line-rate PSP encryption

Extends from NVIDIA ConnectX to BlueField

Co-designed HW roadmap & DOCA SW

One Transport Across NVIDIA & Google NICs

Full wire-format interop + Standard IB Verbs

Scale-Out AI & Disaggregated ML Storage

Consistent low tail latency over shared, lossy Ethernet fabrics



Google Cloud



Call to Action

Virgo is the new network for AI era

- Optimized for performance @ scale, reliability and deployment speed
- Powered by High-speed switch telemetry and Congestion Signaling
- Enable new generations of GPU and TPU products in Google Cloud

Next steps

- Scale across – long-haul RDMA backend
- Extend CSIG for sub-RTT congestion control (UEC)

Multi-vendor fabric

- Call for protocol unification and interoperability

Learn more about
Google's AI networking at



Thank You!



OCP
GLOBAL
SUMMIT

2026