NO<IA

Ethernet-based Al data center fabrics Open, scalable and future ready!

Alperen Akpinar



The life of a large language model (LLM)



AI data center fabrics



- Front-end networks is a typical CLOS-based data center fabrics
- Handles inference requests, cluster management, and data ingress/egress.

- **Back-end** networks are dedicated to GPU-to-GPU, optionally GPU-to-Storage communications
- A separate storage network may be considered for large clusters

VO<IY

Networking of AI clusters GPU-based workloads

The **backend** fabric connects one BIG compute cluster





Distributed GPU training with parallelism (simplified)

AI model parallelism techniques

Data parallelism



Collective Communication Library (CCL) Communicators



NOKIA

AI network traffic is different from traditional apps



What's important for AI training networks? Network latency story

Ethernet versus InfiniBand

- Transmission / Tail latency?
- High-capacity: 400G/800G per interface
- Lossless fabrics (PFC/ECN or credit-based)
- Non-blocking rail-optimized designs
- Al-aware automation

| Feature | Ethernet | InfiniBand | Key insight |
|---------------------------|--|---------------------------------|--|
| Latency | A few µs | < µs | InfiniBand leads at the µs level |
| Tail latency (Jitter) | Scheduled fabrics*, optimized with PFC/ECN | Minimal | Ethernet is competitive with scheduling and tuning |
| Scalability | Powers large AI clusters | Optimized for HPC | Both scale with AI workloads |
| Flexibility and ecosystem | Multi-vendor, low-cost, future-proof with UEC ** | Premium pricing, proprietary | Ethernet offers cost-efficiency and vendor flexibility |

* Broadcom's latest Ethernet switching ASICs

**UEC: "Ultra Ethernet Consortium", an industry initiative to enhance Ethernet for AI/HPC workloads

Job Completion Time (JCT): Total time of GPU computation and synchronization overhead for an AI training job.





VOVIA

Nokia reference design based on Ethernet Backend network: Rail stripe

Platforms: Nokia 7250 IXR-6e/10e/18e, Nokia 7220 IXR-H4, Nokia 7220 IXR-D4/D5 Automation: Event-driven Automation (EDA)



- Rail-optimized, lossless backend network design
- Multi-tenancy support
- Cost-optimized with low power profile
- Al-aware fabric automation
- Design approach that supports small to medium, fixed to modular configurations with tens to hundreds of GPUs, and server connections from 100GE up to 800GE speeds.

NO<IA

Nokia reference design based on Ethernet Back-end network: Rail-optimized large-scale

Spine: Nokia 7250 IXR-6e/10e/18e Leaf: Nokia 7250 IXR-X1b/X3b, Nokia 7220 IXR-H4 Automation: Event-driven Automation (EDA)



- Scalable multi-tier designs with multiple rail-stripes
- Medium to very large configurations that support thousands to tens of thousands of GPUs, and connections at 100GE up to 800GE speeds

NO<IA







Copyright and confidentiality

The contents of this document are proprietary and confidential property of Nokia. This document is provided subject to confidentiality obligations of the applicable agreement(s).

This document is intended for use by Nokia's customers and collaborators only for the purpose for which this document is submitted by Nokia. No part of this document may be reproduced or made available to the public or to any third party in any form or means without the prior written permission of Nokia. This document is to be used by properly trained professional personnel. Any use of the contents in this document is limited strictly to the use(s) specifically created in the applicable agreement(s) under which the document is submitted. The user of this document may voluntarily provide suggestions, comments or other feedback to Nokia in respect of the contents of this document ("Feedback").

Such Feedback may be used in Nokia products and related specifications or other documentation. Accordingly, if the user of this document gives Nokia Feedback on the contents of this document, Nokia may freely use, disclose, reproduce, license, distribute and otherwise commercialize the feedback in any Nokia product, technology, service, specification or other documentation.

Nokia operates a policy of ongoing development. Nokia reserves the right to make changes and improvements to any of the products and/or services described in this document or withdraw this document at any time without prior notice.

The contents of this document are provided "as is". Except as required by applicable law, no warranties of any kind, either express or implied, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose, are made in relation to the accuracy, reliability or contents of this document. NOKIA SHALL NOT BE RESPONSIBLE IN ANY EVENT FOR ERRORS IN THIS DOCUMENT or for any loss of data or income or any special, incidental, consequential, indirect or direct damages howsoever caused, that might arise from the use of this document or any contents of this document.

This document and the product(s) it describes are protected by copyright according to the applicable laws.

Nokia is a registered trademark of Nokia Corporation. Other product and company names mentioned herein may be trademarks or trade names of their respective owners.

