AI Infrastructure over Ethernet: Secure, Scalable, High-Performance Design

Exploring how Ethernet-based fabrics are transforming AI workloads across training and inference environments

Speaker : Nazim Khan

Agenda

- AI Landscape & Overview
- AI Clusters
- Networking for AI Clusters
- Ultra Ethernet Consortium (UEC)
- Security for AI

The Elements of AI

CPU : The Brain of the Computer

Acts as the brain of the computer, managing general-purpose tasks and coordinating system operations for software execution

GPU : Powering Graphics & Computations

Specialized for parallel processing, excelling at rendering graphics and accelerating computations like AI and simulations

DPU : Optimizing Data & Networking

Dedicated to data processing and networking tasks, offloading these from CPU/GPU to optimize performance in data centers and cloud environments

AI Clusters are Interconnected networks of high-performance GPUs

What is GPU Collective Communication ?

- A method where a group of GPUs exchange data simultaneously to function as a single unit.
- Synchronizes model gradients across thousands of nodes during AI training.
- Utilizes topology-aware algorithms (Ring, Tree) to maximize bandwidth and minimize latency.

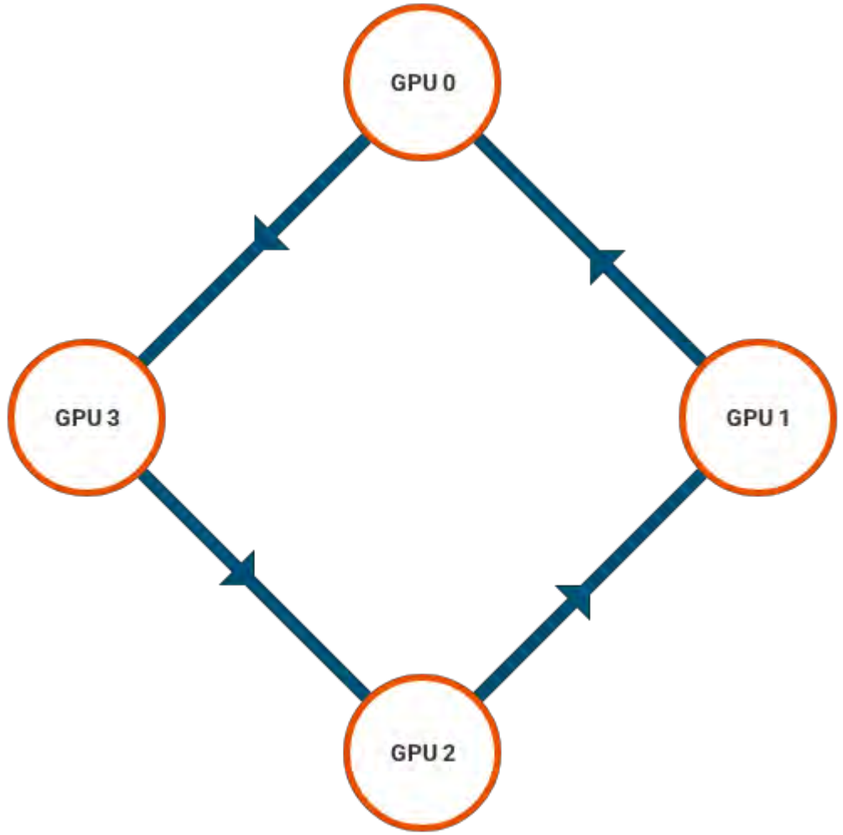

AI Cluster types

Training

Inference

AI Cluster types & requirements

	Training	Inference
Node to Node Bandwidth	High	Low
Key Metric	Training time	High Availability & Latency
Operational Mode	Offline	Online
Infrastructure	Large network GPU/CPU hosts	Smaller network CPU/GPU hosts

How to build a lossless Network fabric

Remote Direct Memory Access - RDMA

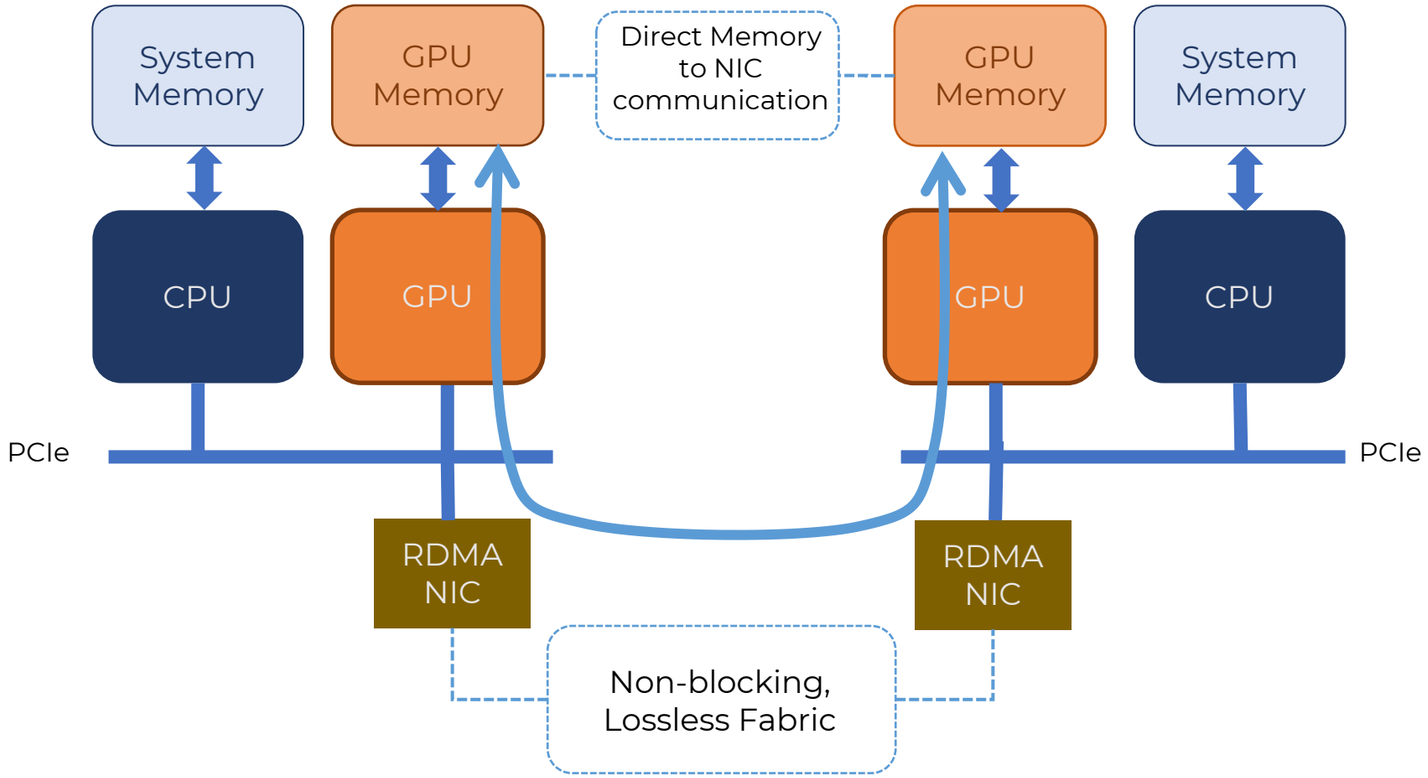

- RDMA allows AI/ML nodes to directly exchange data over a network by directly accessing system and GPU memory remotely
- Latency is very low, as CPU and kernel can be bypassed

Network Architecture for AI Clusters

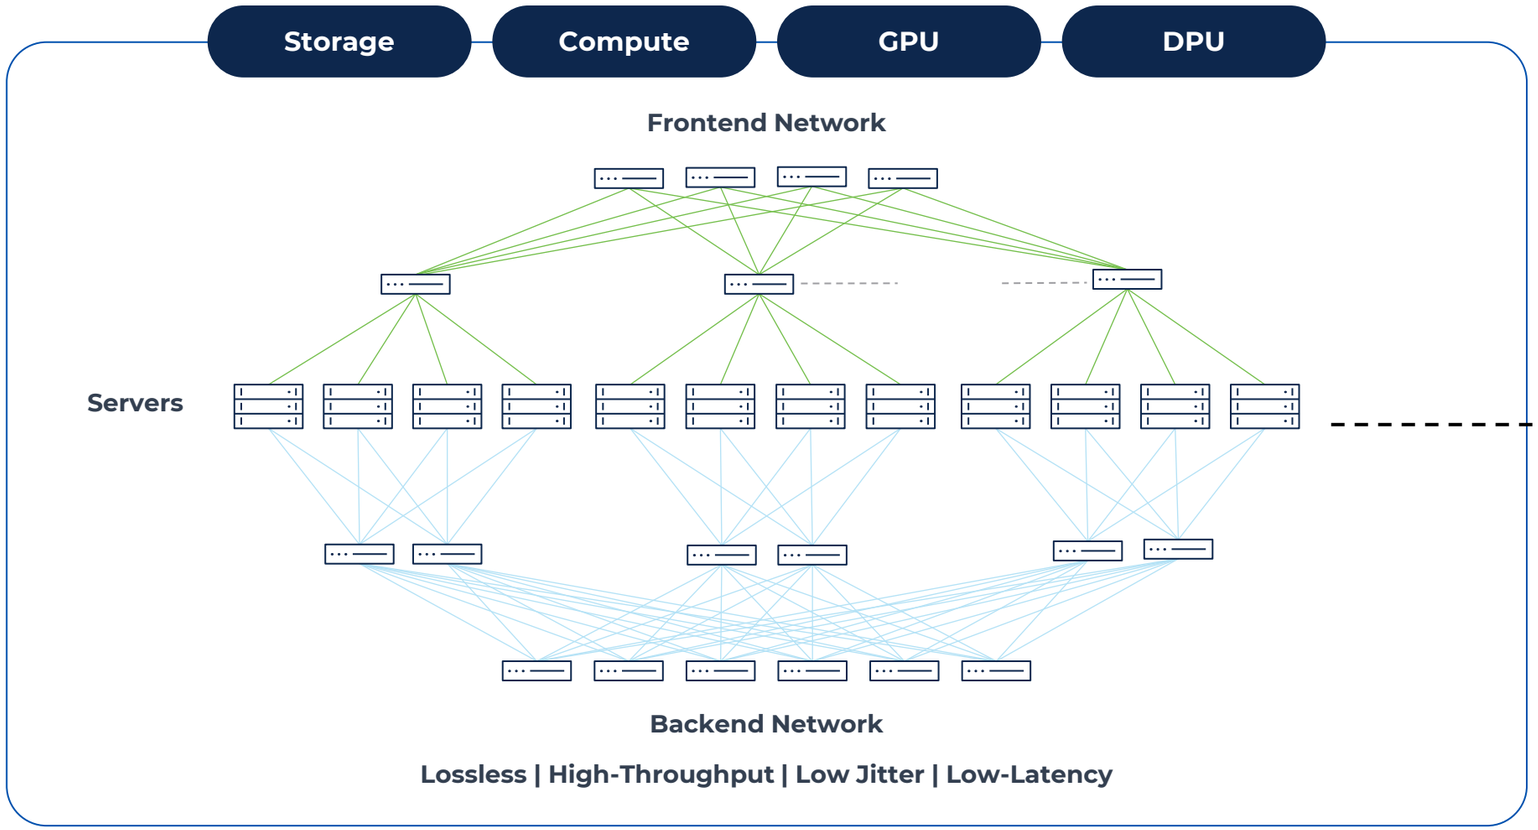

The **frontend network** is used for all other communications and can be used for IP storage access as well

The **backend network** is used only for GPU-to-GPU connectivity

AI Cluster Backend Network

Traditional InfiniBand

- High-performance clusters
- Proprietary ecosystem
- Established HPC standard
- Higher deployment costs

Modern Ethernet Alternative

- Cost-effective solution
- Open standards-based
- Enhanced with RoCEv2
- Broad ecosystem support

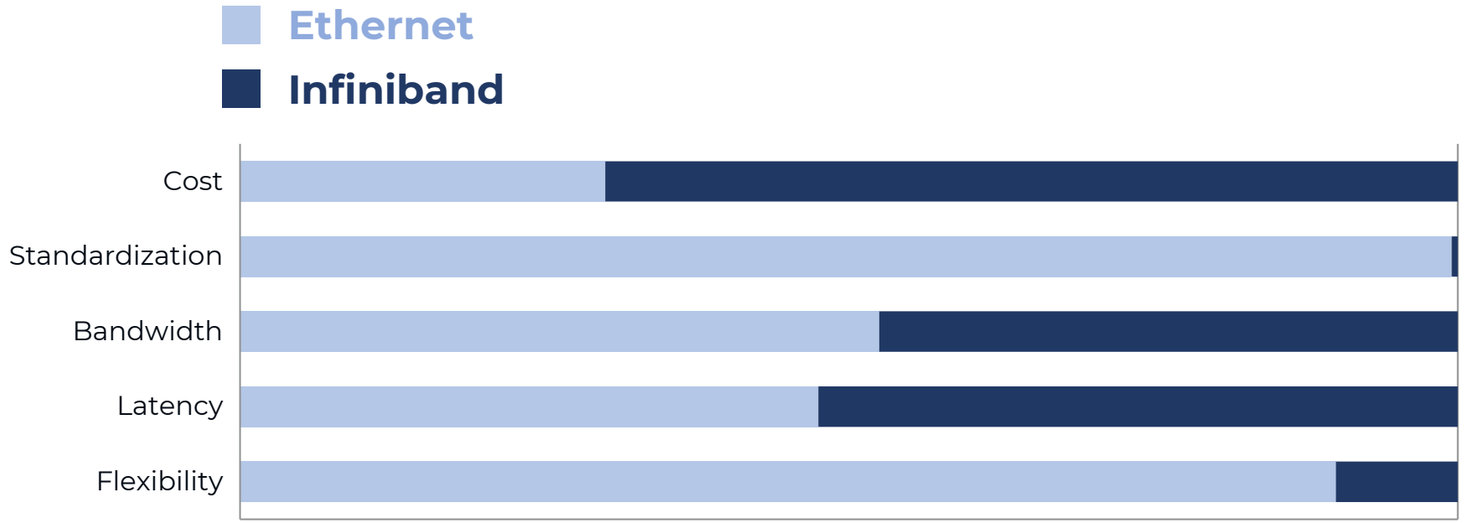

RoCE – RDMA over Converged Ethernet

- Extension of RoCE protocol that involves a simple modification of the RoCE packet format.
- Carry IP header and UDP header that serves as a stateless encapsulation layer for RDMA transport over IP.

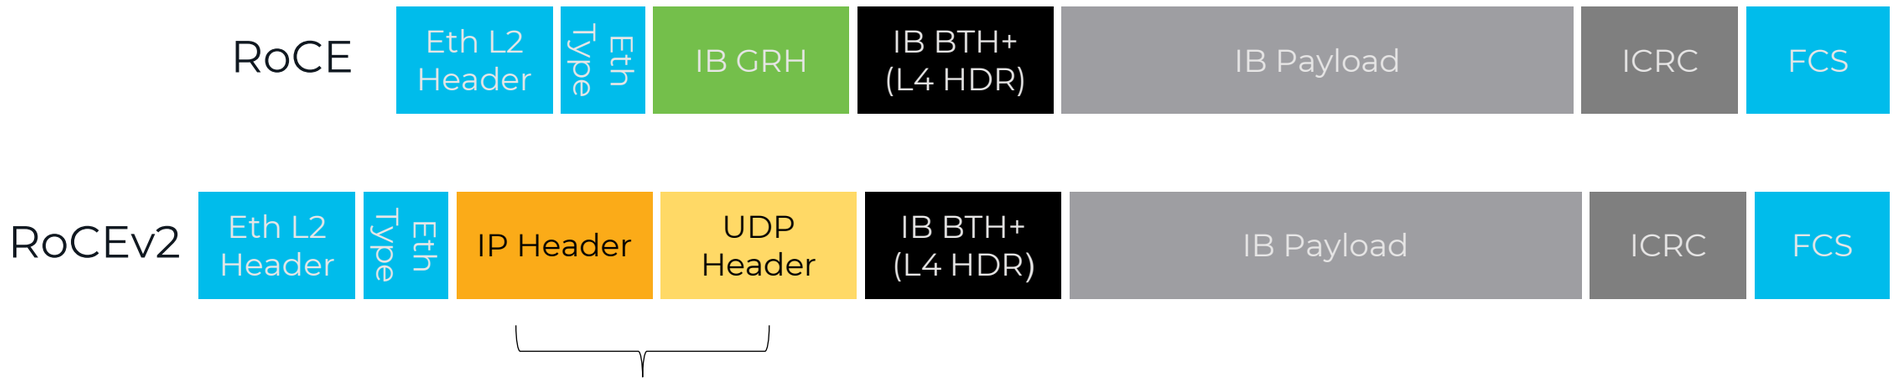

Infiniband Specification - Annex A16 and A17
https://www.infinibandta.org/ibta-specification/ 5

RoCEv2

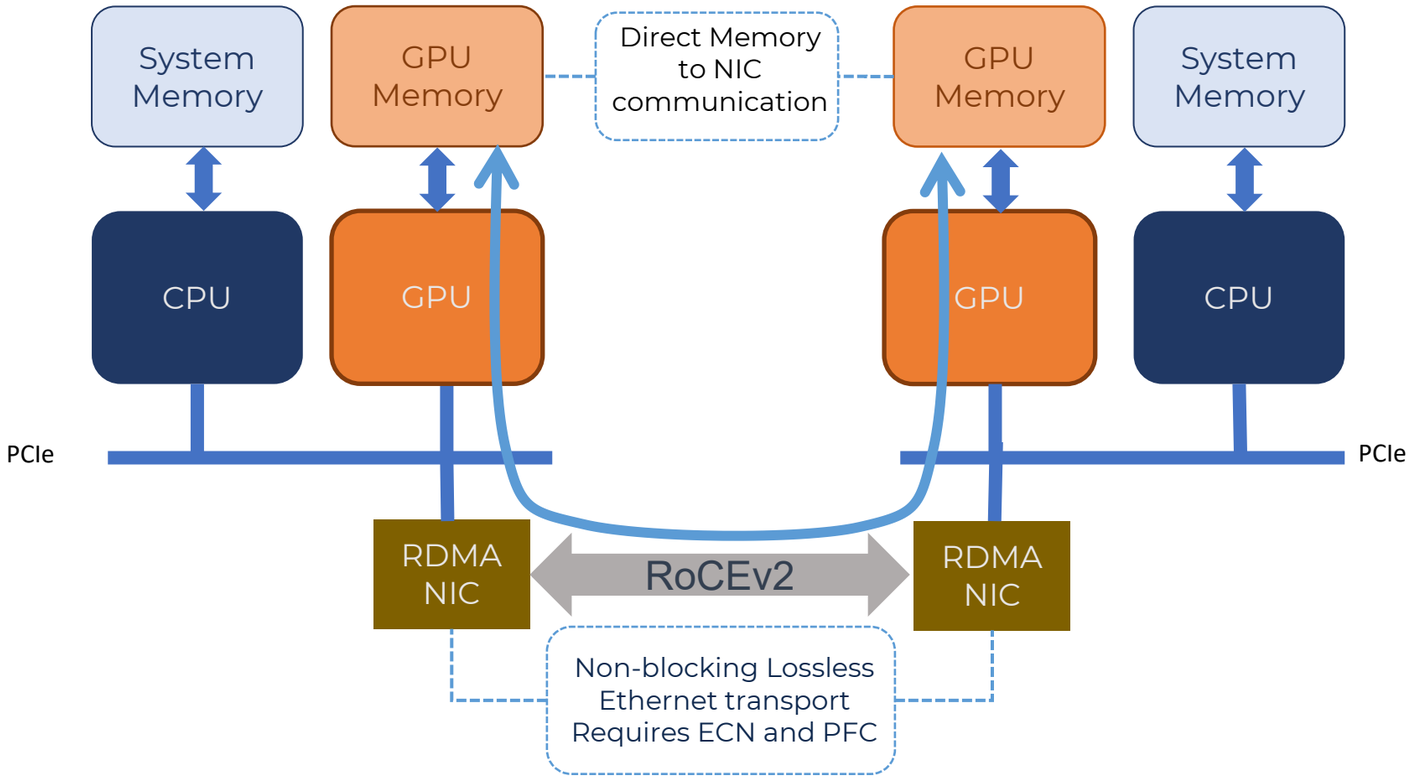

- Uses well-known UDP Destination Port: 4791
- Supports both IPv4 and IPv6
- Makes use of ECN field in IPv4/6 header for signaling of congestion
- Requires PFC to be enabled for RoCEv2 transport and lossless network

Congestion Management: ECN

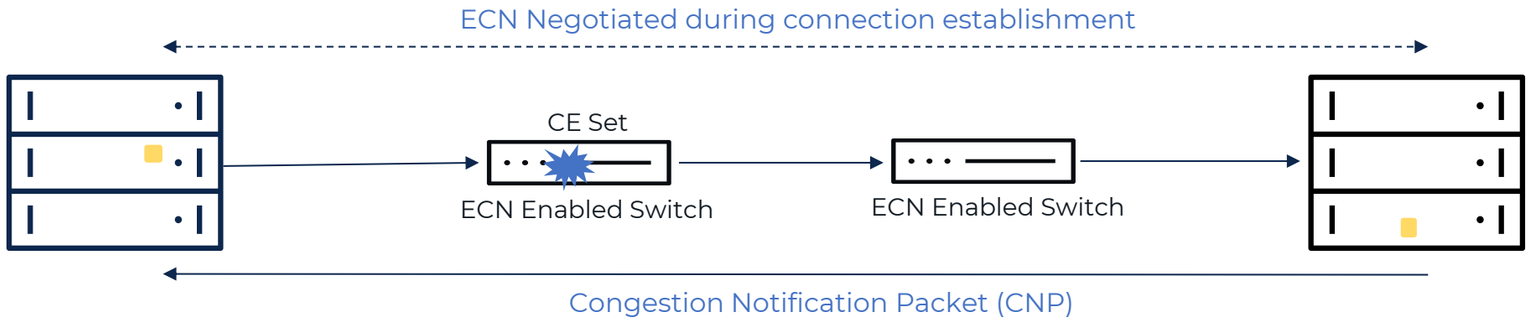

- Transmitter/Receiver slow down or speed up
- Multiple techniques for selecting which packet and how many packets get marked
- Per flow/connection end-to-end congestion management
- A flow can be very elastic as a result

ECN	Meaning
00	Non-ECN Capable
10	ECN Capable Transport (0)
01	ECN Capable Transport (1)
11	Congestion Experienced

Congestion Management: PFC

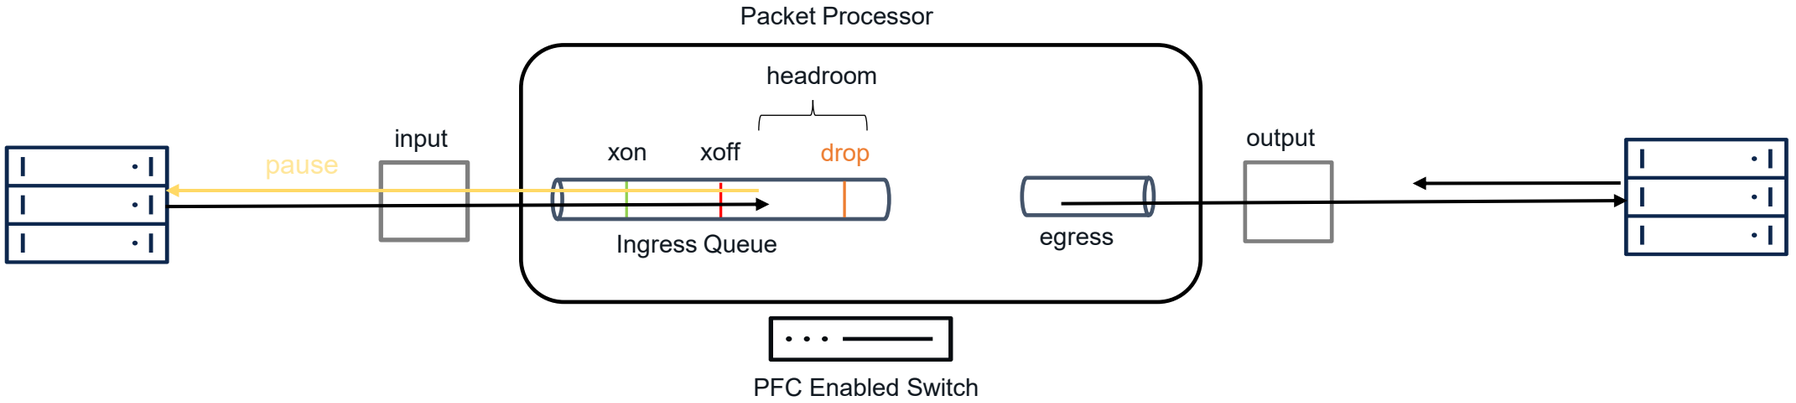

- A.k.a "Lossless Ethernet”
- PFC enables Flow Control on a Per-Priority basis
- PFC is also called Per-Priority-Pause
- There can be lossless and lossy priorities at the same time on the same wire

Ultra Ethernet Consortium (UEC) Initiative

Mission

Industry collaboration focused on evolving Ethernet specifically for AI and HPC applications

Key Focus Areas

- Advanced congestion control mechanisms
- Deterministic latency guarantees
- Enhanced telemetry capabilities

Industry Impact

Driving open standards that enable broader ecosystem adoption and innovation

Security for AI

Data Poisoning Attacks

Adversaries manipulate training data to inject malicious inputs and compromise AI models

Model Extraction Attacks

Unauthorized extraction of AI models to gain access to sensitive intellectual property

Adversarial Examples

Carefully crafted inputs designed to fool AI models and cause misclassification or incorrect decisions

Privacy Violations

Sensitive user data leaks or misuse due to vulnerabilities in AI systems

Insider Threats

Malicious insiders with access to AI infrastructure can compromise system security and data integrity

Security Principles for AI

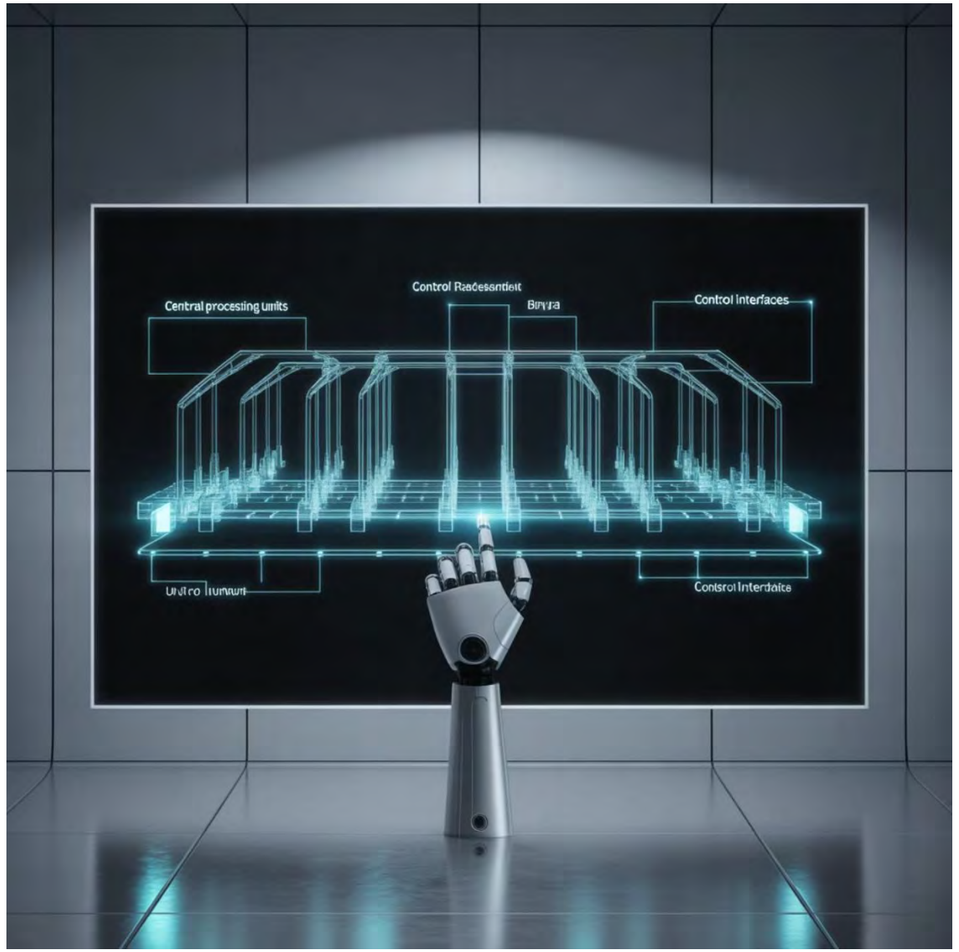

Data Encryption

Implement robust encryption techniques to protect AI data at rest and in transit

Secure Design Principles

Incorporate security best practices into the design and architecture of AI infrastructure

AI Model Hardening

Improve robustness of AI models against adversarial attacks, such as adversarial training and defensive distillation

Monitor & Detect Anomalies

Implement continuous monitoring and anomaly detection to identify & respond to security incidents

Robust Access Controls

Implement strong authentication and authorization mechanisms to manage access to AI resources

DPU-Enabled NICs and Accelerator Integration

Data Processing Units (DPUs)

- Offload network processing from main CPU
- Specialized packet processing capabilities
- Enhanced security and isolation features
- Programmable data plane acceleration

Model Partitioning Benefits

- Distributed model execution across nodes
- Reduced memory requirements per device
- Parallel inference processing
- Improved resource utilization

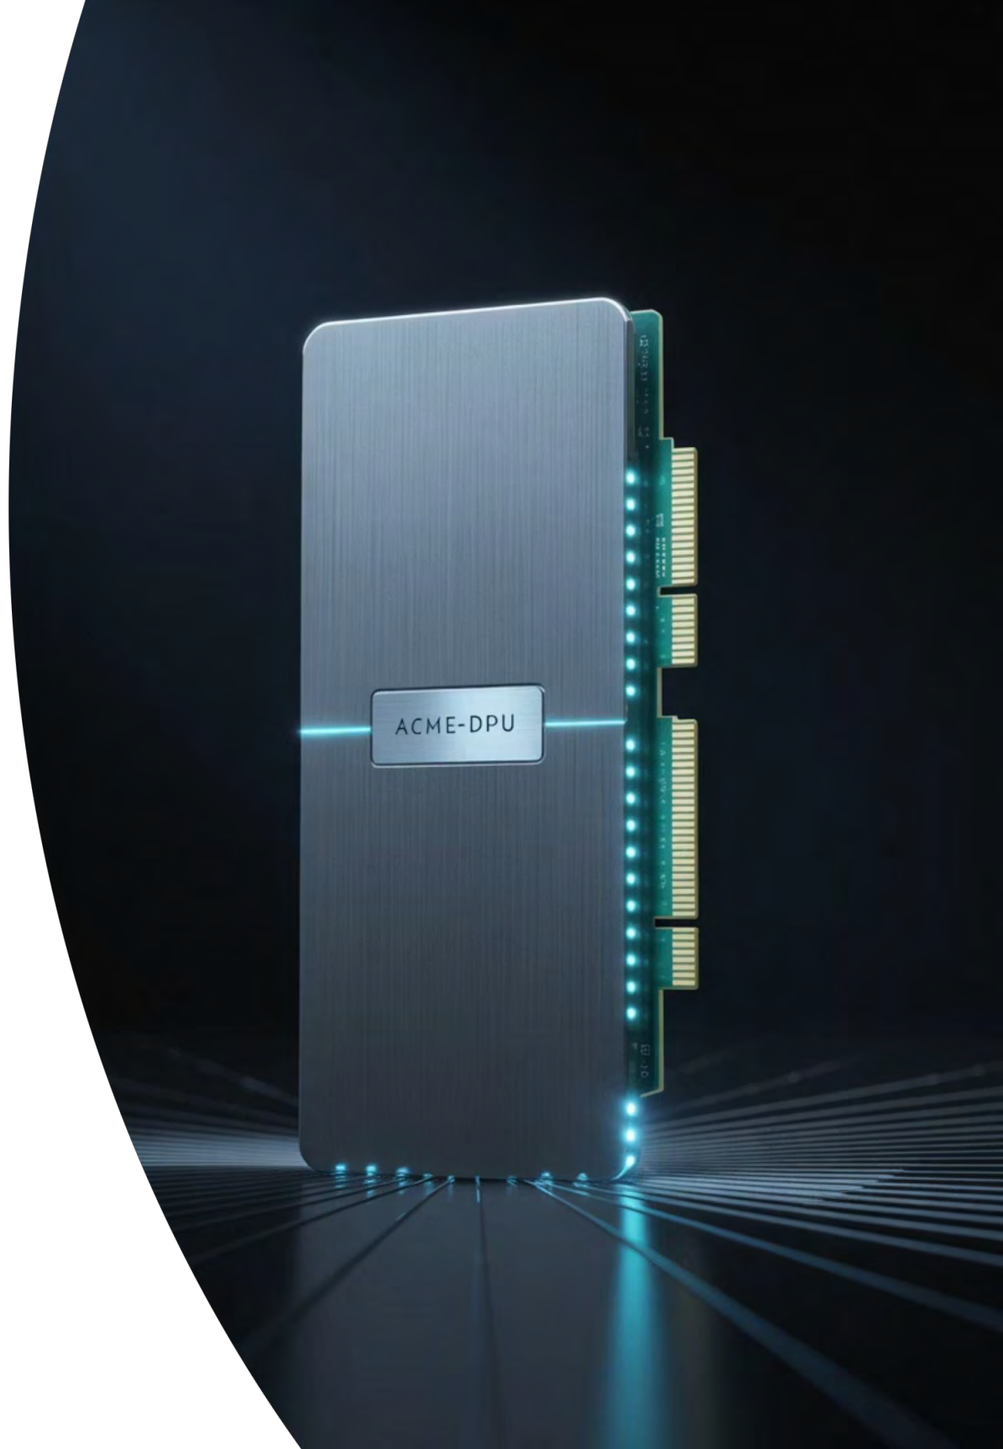

Building Ethernet-Based AI Systems: Key Takeaways

✓ **Ethernet viability for AI workloads**

Modern Ethernet with RoCEv2 provides cost-effective alternative to traditional InfiniBand for many AI applications

✓ **Architecture considerations matter**

Proper topology design, congestion management, and traffic optimization are critical for success

✓ **Industry collaboration drives innovation**

UEC and open standards initiatives are accelerating Ethernet evolution for AI and HPC applications

Thank you!