Public

AI OPEN
DATA CENTER



HPE and AMD: BUILDING AI RACK SYSTEM WITH OPEN STANDARDS

Mike Wong *Senior Director, Product Management HPE*

Shangxin Du *Director Product Management AMD*

OCP PLATINUM® Member

Public

AMD Helios AI Rack by HPE

– Open Rack scale AI Infrastructure

72 GPUs Per rack	**31 TB** HBM4 Memory	**260 TB/s** Scale Up Bandwith
2.9 Exaflops AI Compute	**1.7 PB/s** HBM4 Bandwidth	**43 TB/s** Scale Out Bandwidth



OPEN Compute Project®

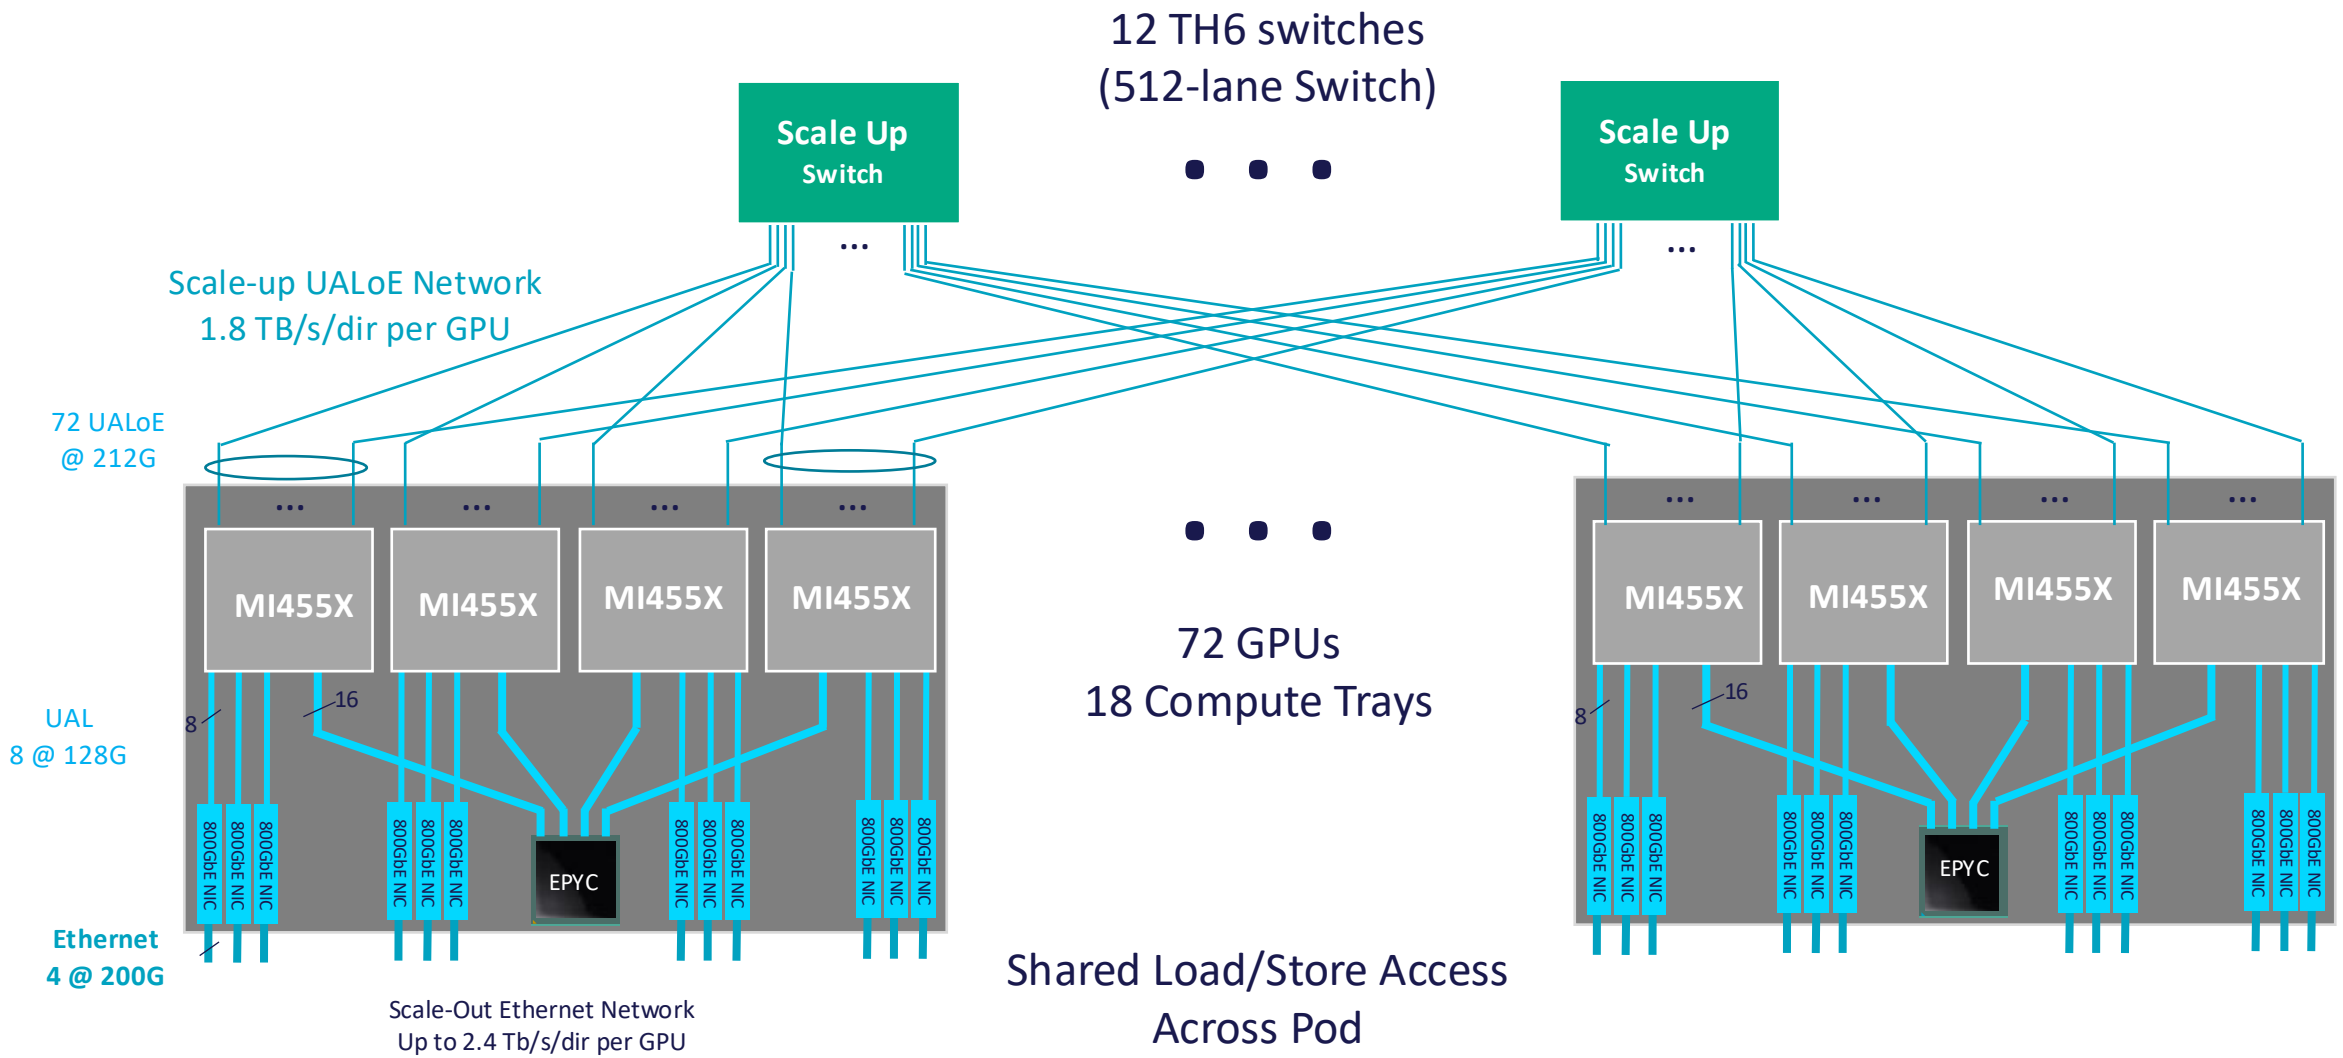
ULTRA ACCELERATOR LINK™

Ultra Ethernet Consortium

SONiC

Public

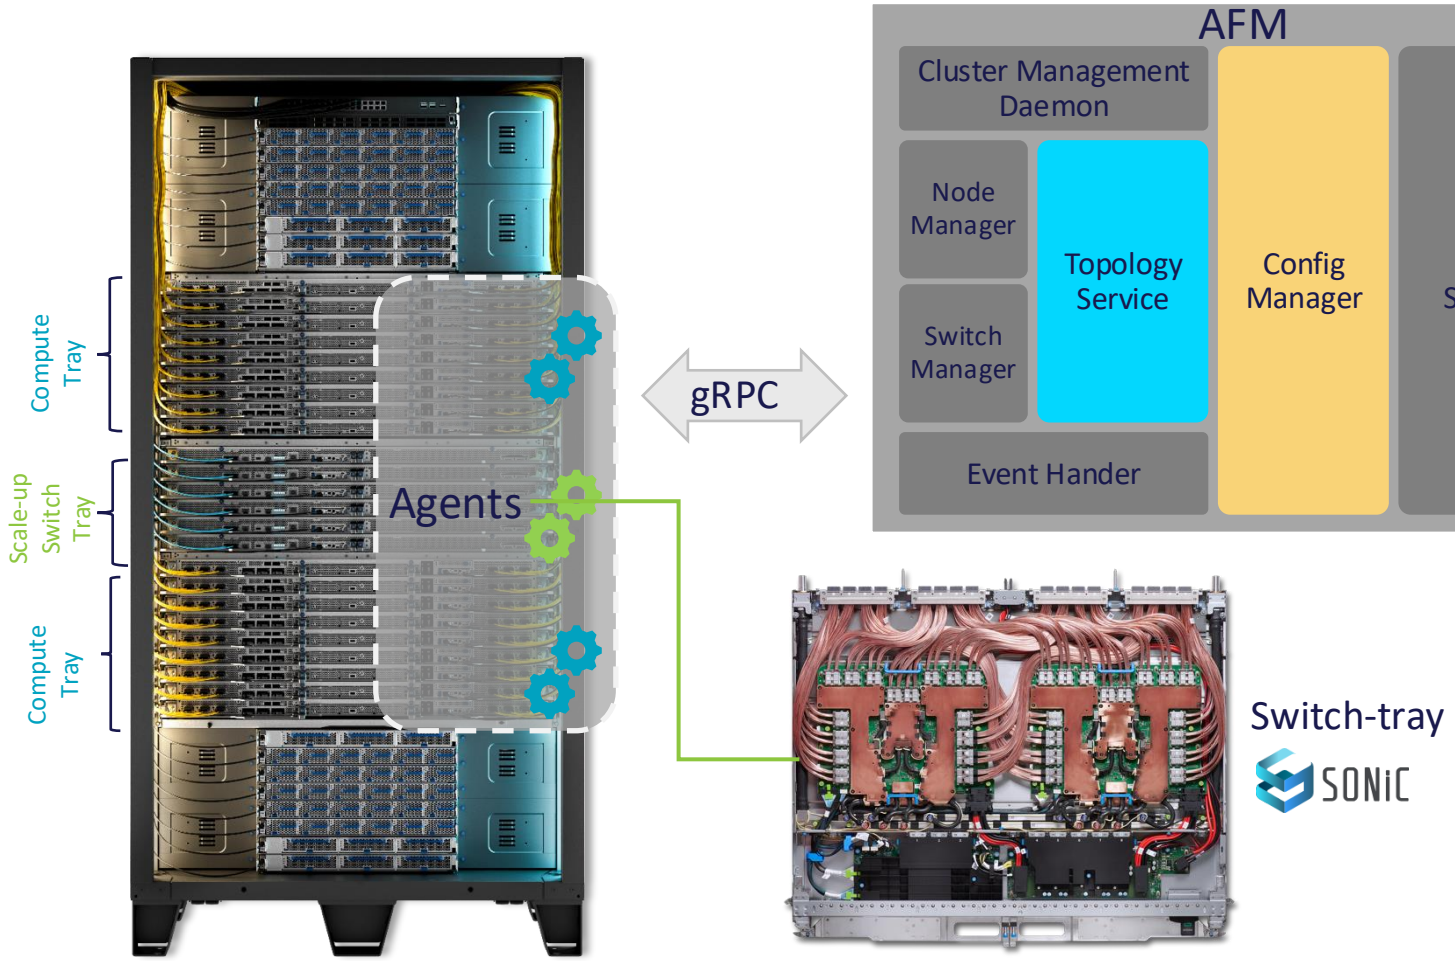
Helios Scale-up Pod Architecture



Public

Helios Scale-up Fabric

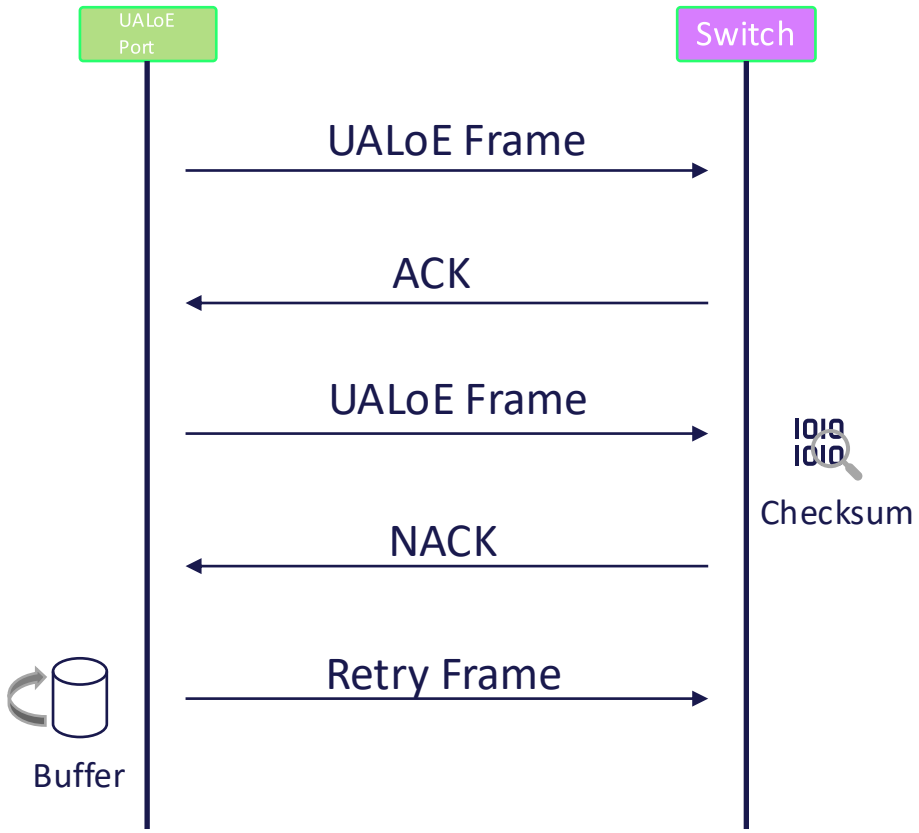
AMD Fabric Manager, Rack-wide Intelligence



- Streamlined workflow for scale-up network provisioning
- Fabric-level perspective for effective control, monitoring, and troubleshooting
- Simplified provisioning with zero-touch scale-up network deployment
  - Faster implementation
  - Reduced need for specialized skills
  - Consistent adherence to best practices
- Support for both advanced and turnkey deployment scenarios

Public

Operational Remediation – Link Level Retry

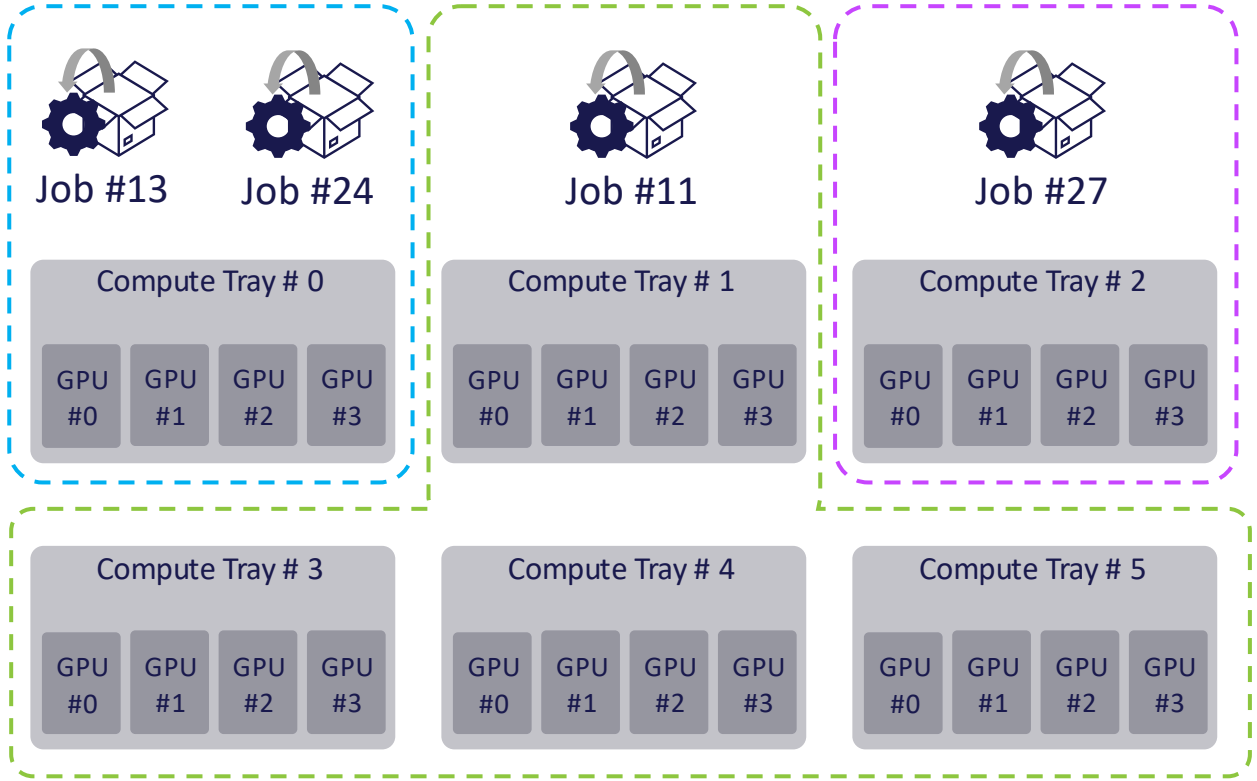


- Frame Transmission with Sequence Numbers
  - Each frame exchanged between nodes is assigned a local sequence number, which remains hidden from the upper protocol layers.
- Frame Buffering
  - The sender stores transmitted frames in a local retry queue, maintaining them in a buffer until confirmation of receipt is received from the receiver.
- Acknowledgment
  - Upon successful reception of each frame, the receiver sends an acknowledgment (ACK) back to the sender.
- Loss Detection
  - If a frame is lost or corrupted, the receiver detects a discontinuity in the sequence numbers and responds with a Negative Acknowledgment (NACK).
  - Additionally, if the sender does not receive an ACK or NACK within a specific time interval, it triggers a timeout mechanism to retry transmission.

Public

Operational Remediation – vPod

- A vPod can be created by combining any arbitrary set of GPUs.
  - Configuring a vPod necessitates DF reconfiguration.
  - At FCS, allocation is restricted to the granularity of the compute tray.
- Communication is confined within the vPod.
  - Traffic is segregated in scale-up switches through VLANs.
  - A link failure within one vPod does not affect GPUs in other vPods.
- Job assignment to a vPod:
  - Each vPod serves as an isolated failure domain.
  - Failure recovery within a vPod affects all jobs running on it.
  - Assign a vPod to a job when it requires equal or more than 4 GPUs.



Public

AINIC

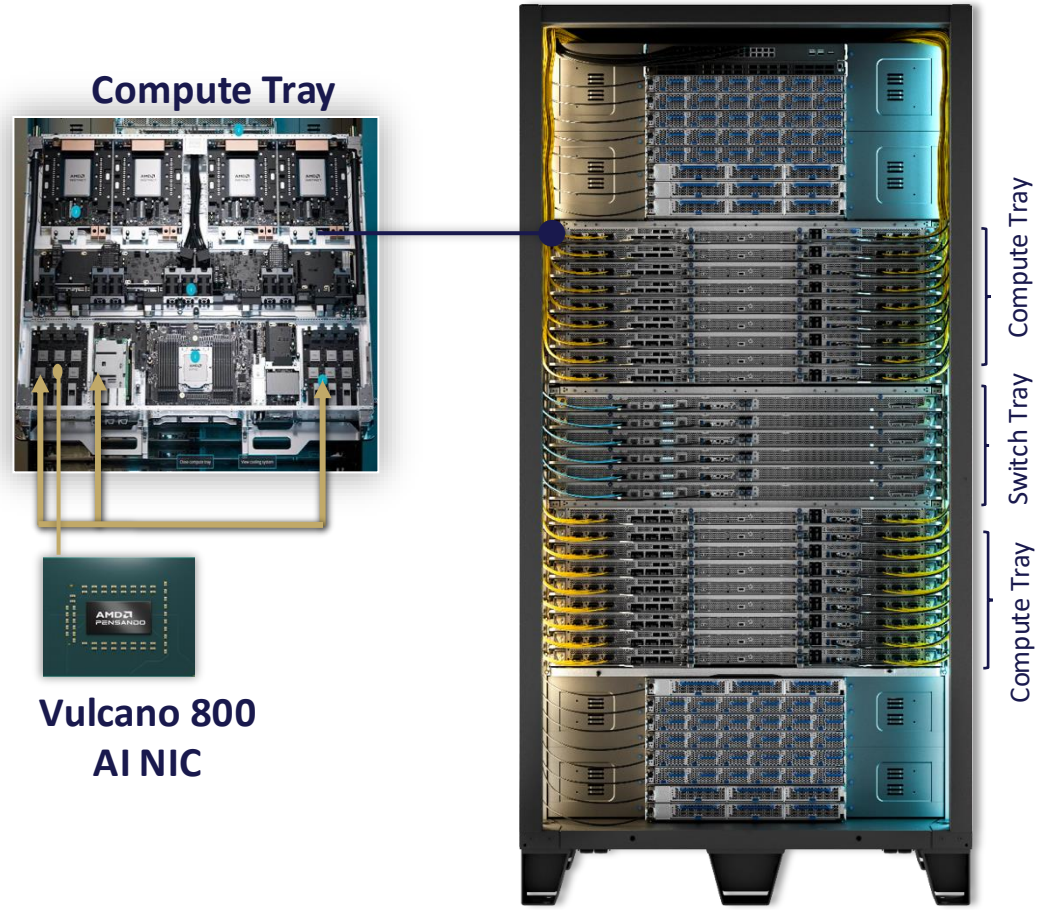
Programmable backend and frontend

Helios Rack AI Solution

- AI NIC is designed to run programmable RDMA over Ethernet and network offloads to support low latency, high bandwidth storage and collective communications

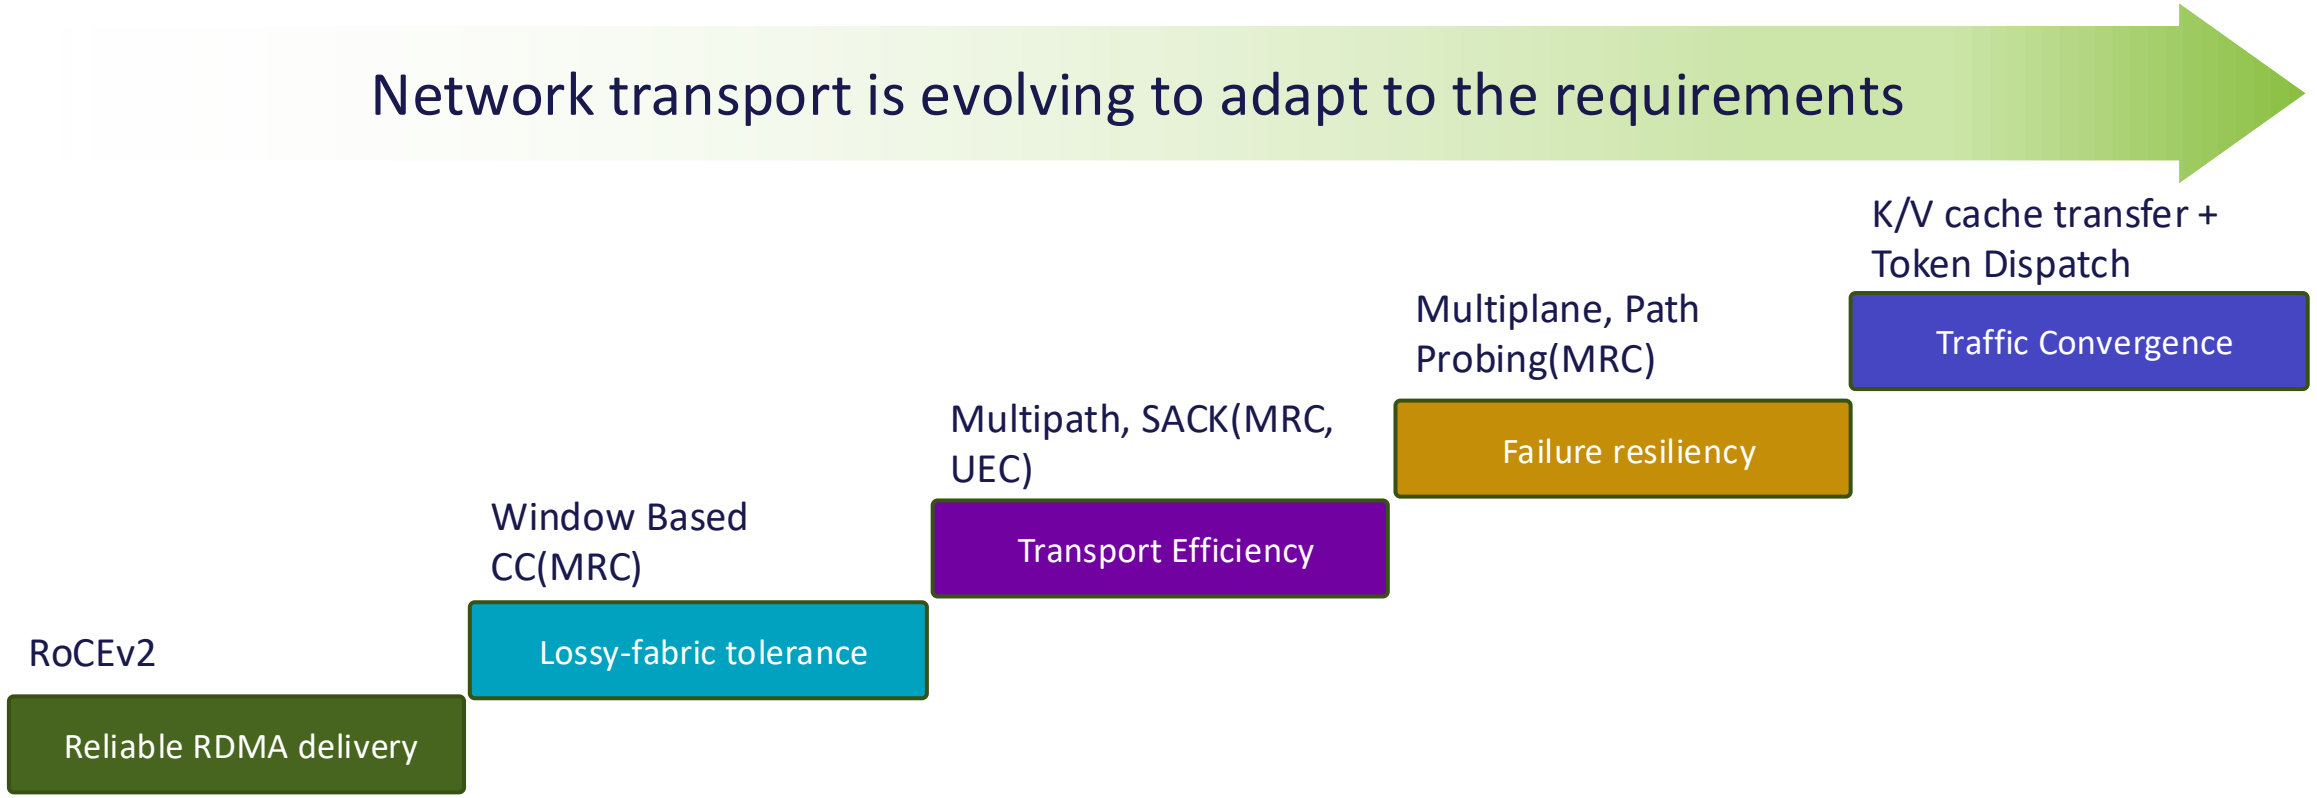
- Examples:
  - Software-Defined Networks (SDN)
  - RDMA (RoCEv2, MRC, MetaRoCE, UEC)
  - Security
  - Storage (NVMeoF)

- Deployed in multiple modes
  - Front-End Ethernet NIC
  - Scale-out RDMA NIC
  - Storage Initiator/ Target NVMeoF



Public

Scale-out Transport is evolving



AMD Pensando P4DMA Engine

One programmable data path carries every layer

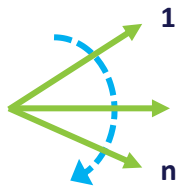
- configurable packet & memory parsing
- custom logic for new protocols
- custom collectives in the data path
- telemetry that adapts as transports change

Public

MRC delivering key performance enhancements

Multipathing & In-Order-Delivery

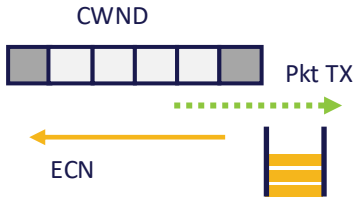
Packet Spray



Fast Job Completion Time

- Eliminating Hash Collisions
- Maximize fabric utilization with **1:1** subscription
- Microseconds failure recovery
- Reduce the tail latency

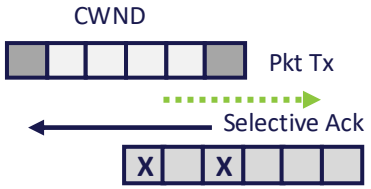
Path Aware NSCC



Open ecosystem

- Avoid the head-of-line blocking
- Path-aware throttling
- Fast ramping for bursty workload

Selective Retransmission

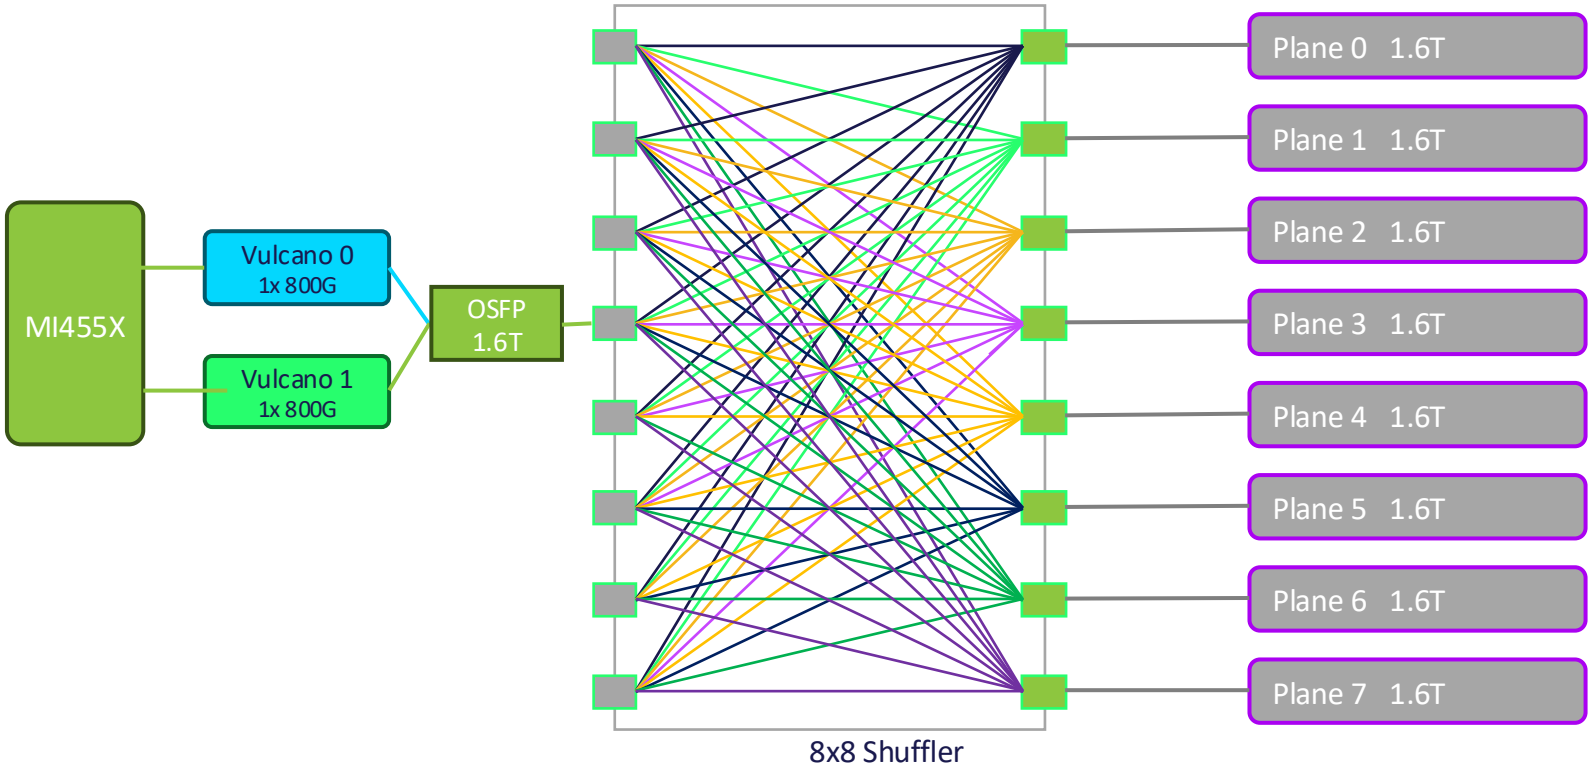


Operation Efficiency

- Massive bandwidth conservation
- Faster lost recovery

Public

Multi-Plane Improves Redundancy & GPU Utilization



- Port breakout to split 800G NIC into 4×200G, distributed over four independent fabrics/planes
- Packet spray & in-order delivery: achieve non-blocking GPU bandwidth for large clusters without a core layer
- Highly available with better fault isolation: Maximize GPU cluster scale at lower cost
- Fabric shuffler reduces the connectivity complexity

Public

Summary

- Massive growth in AI demand drives the need for an open rack-scale AI systems
- AMD Helios rack co-designed with emerging AI workloads
  - 72 GPUs operate as a unified compute engine
  - ORW rack provides integrated power, cooling, and management
- Scale-up network:
  - Shared memory ld/st access to over 30TB of HBM4 at 1.8 TB/s/dir per GPU
  - UALoE protocol provides flexible partitioning, security isolation, and multi-layer resiliency
- Scale-out network:
  - Pensando AI-NICs provide protocol flexibility, advanced load balancing & congestion control, and telemetry & error management
  - 64x 1.6T Ethernet switches support large radix, multi-plane, standard-based load balancing

Public

Thank You!

