



# ConnectX-8 SuperNIC


Idan Burstein | Hot Chips 2025

# AI Training and Inference Require Fungible End-to-End Network Policies



## Inference

- Disaggregated partitioned workload
- Latency sensitivity
- On demand job scheduling, efficiency requires scale
- Large interface with the outside world (KV cache, agentic)



## Training

- Synchronized long lasting workload
- Large scale single job, campus or intercampus scale
- Tail latency impacts efficiency
- Minimal interface with the outside world

# ConnectX-8 Isn't Just Another NIC – It's a SuperNIC!



RDMA technology deployed over millions of GPUs

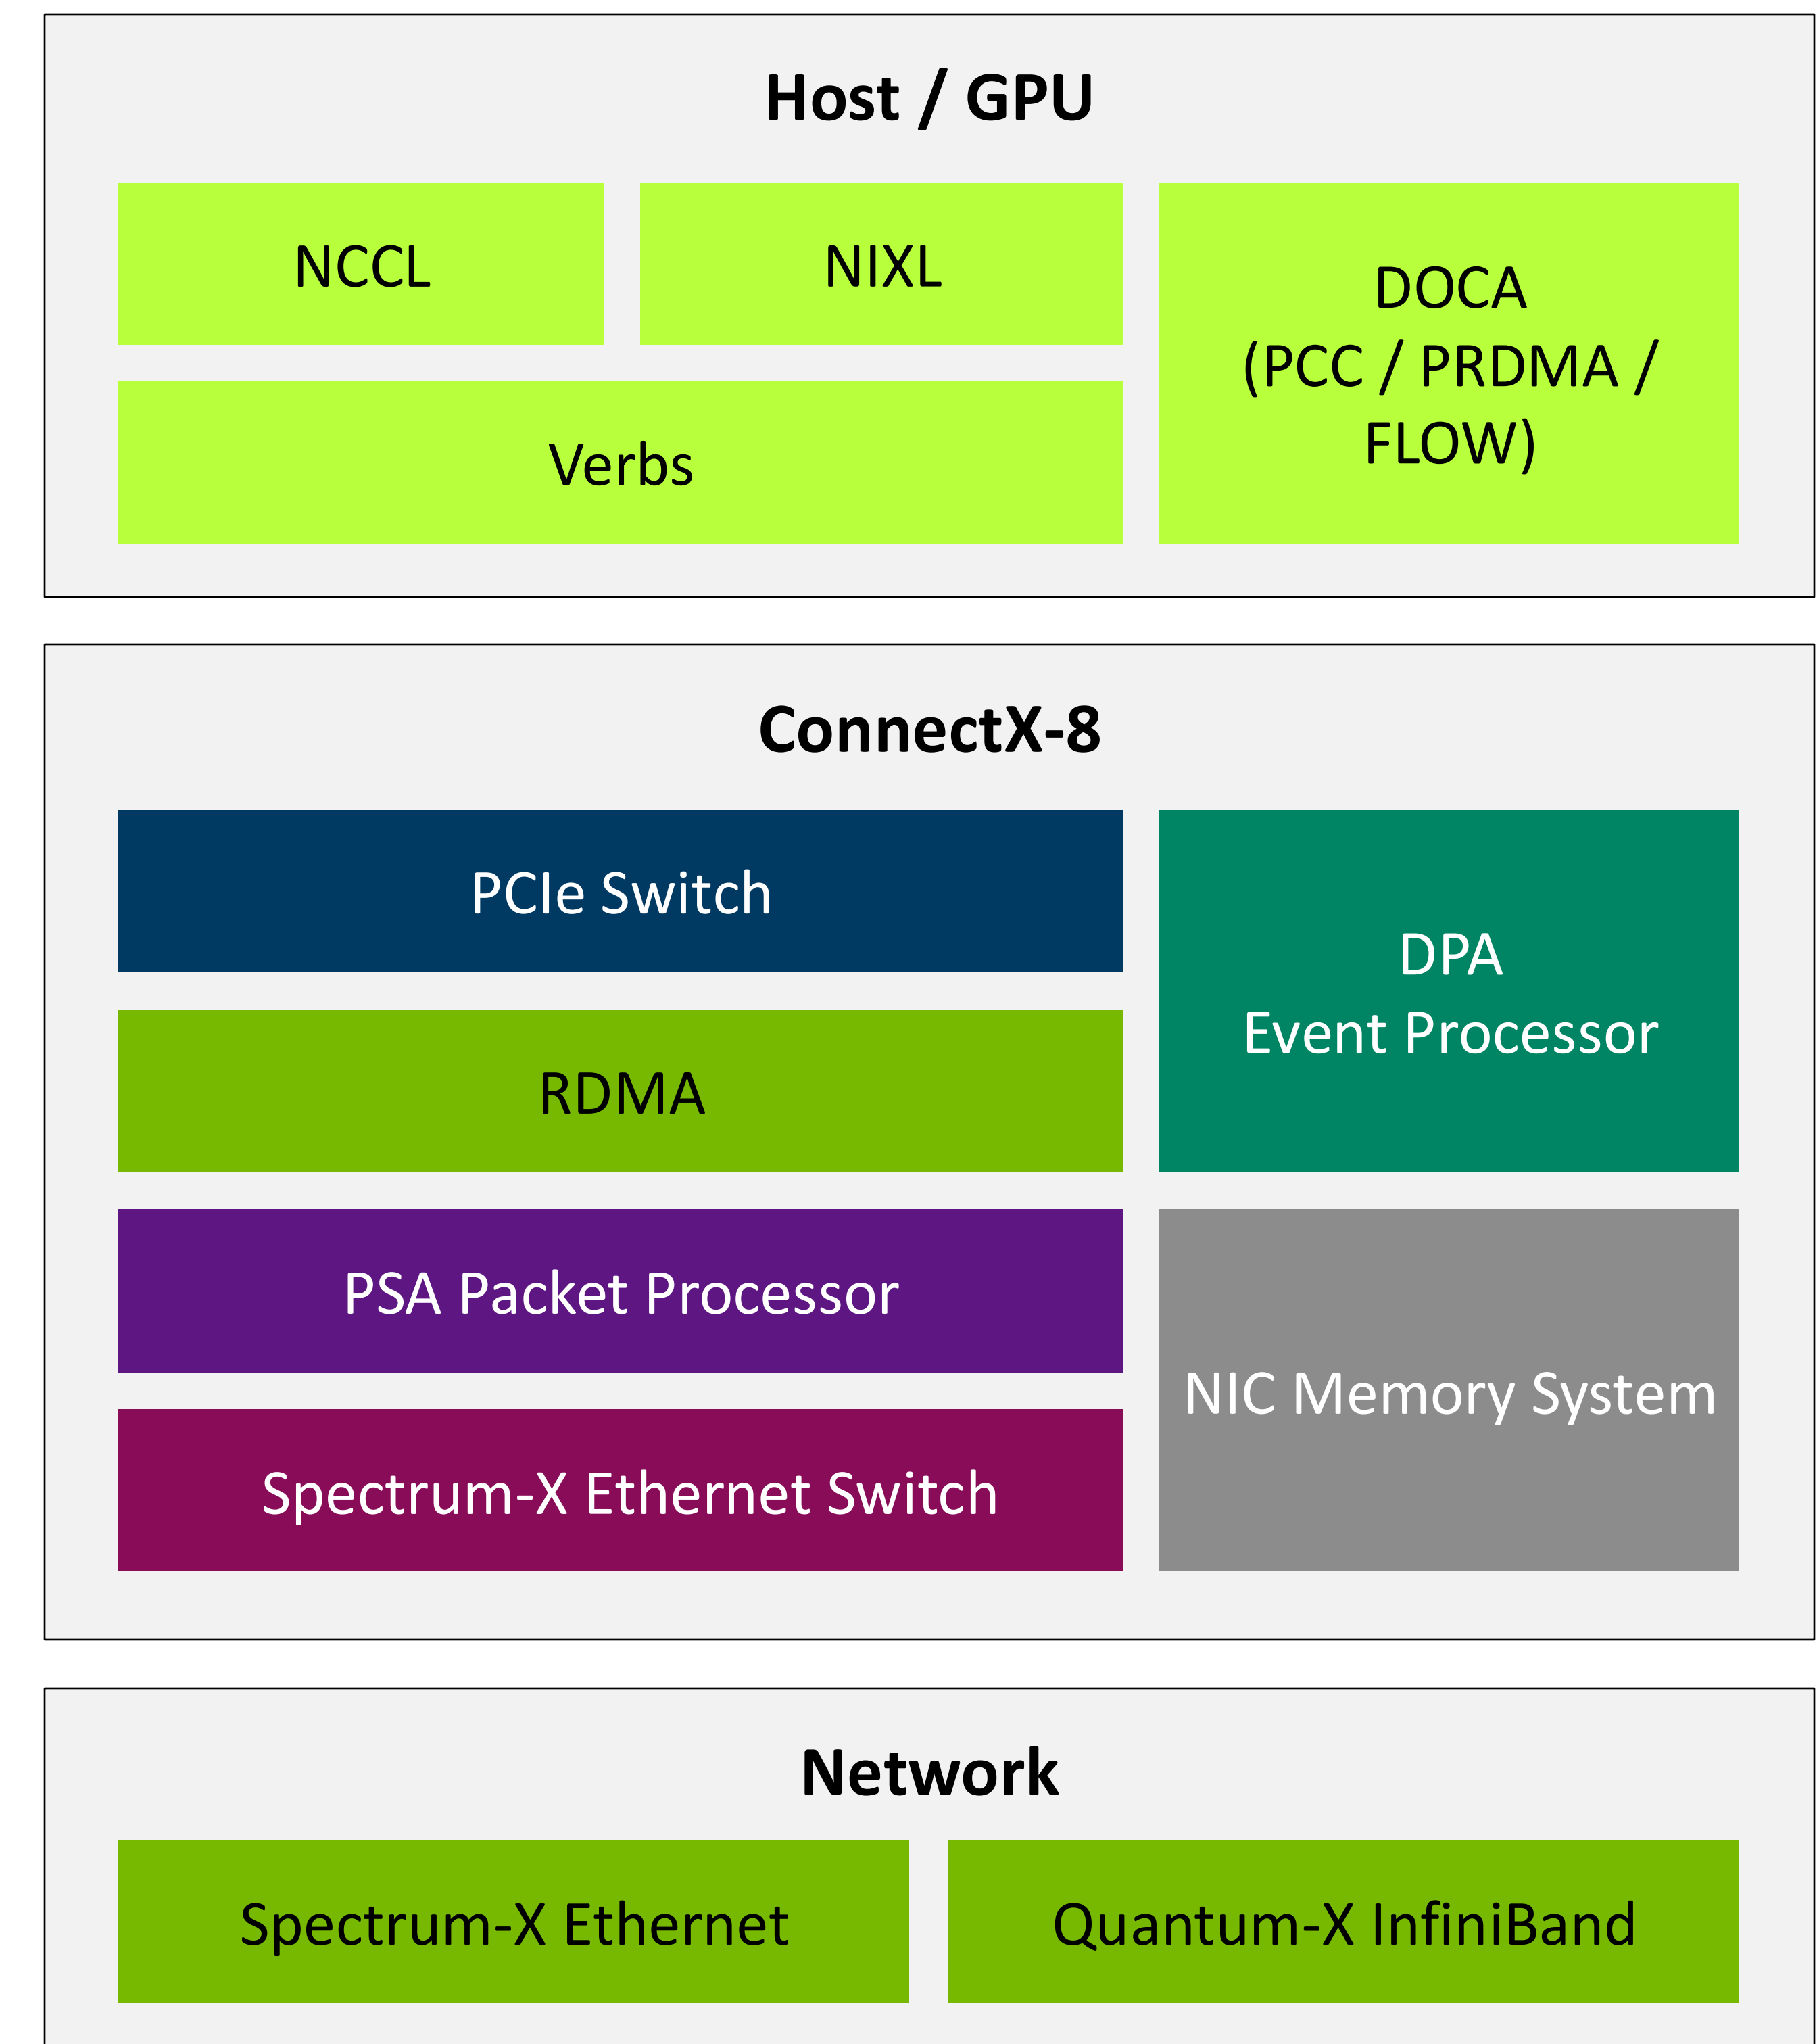
800G RDMA hardware pipeline, designed for AI bandwidth, latency and scale

Integrated load balancing, congestion control and reliability

Deep data-path programmability for AI workload and data center versatility

Tight integration to system architecture

Enterprise-class security

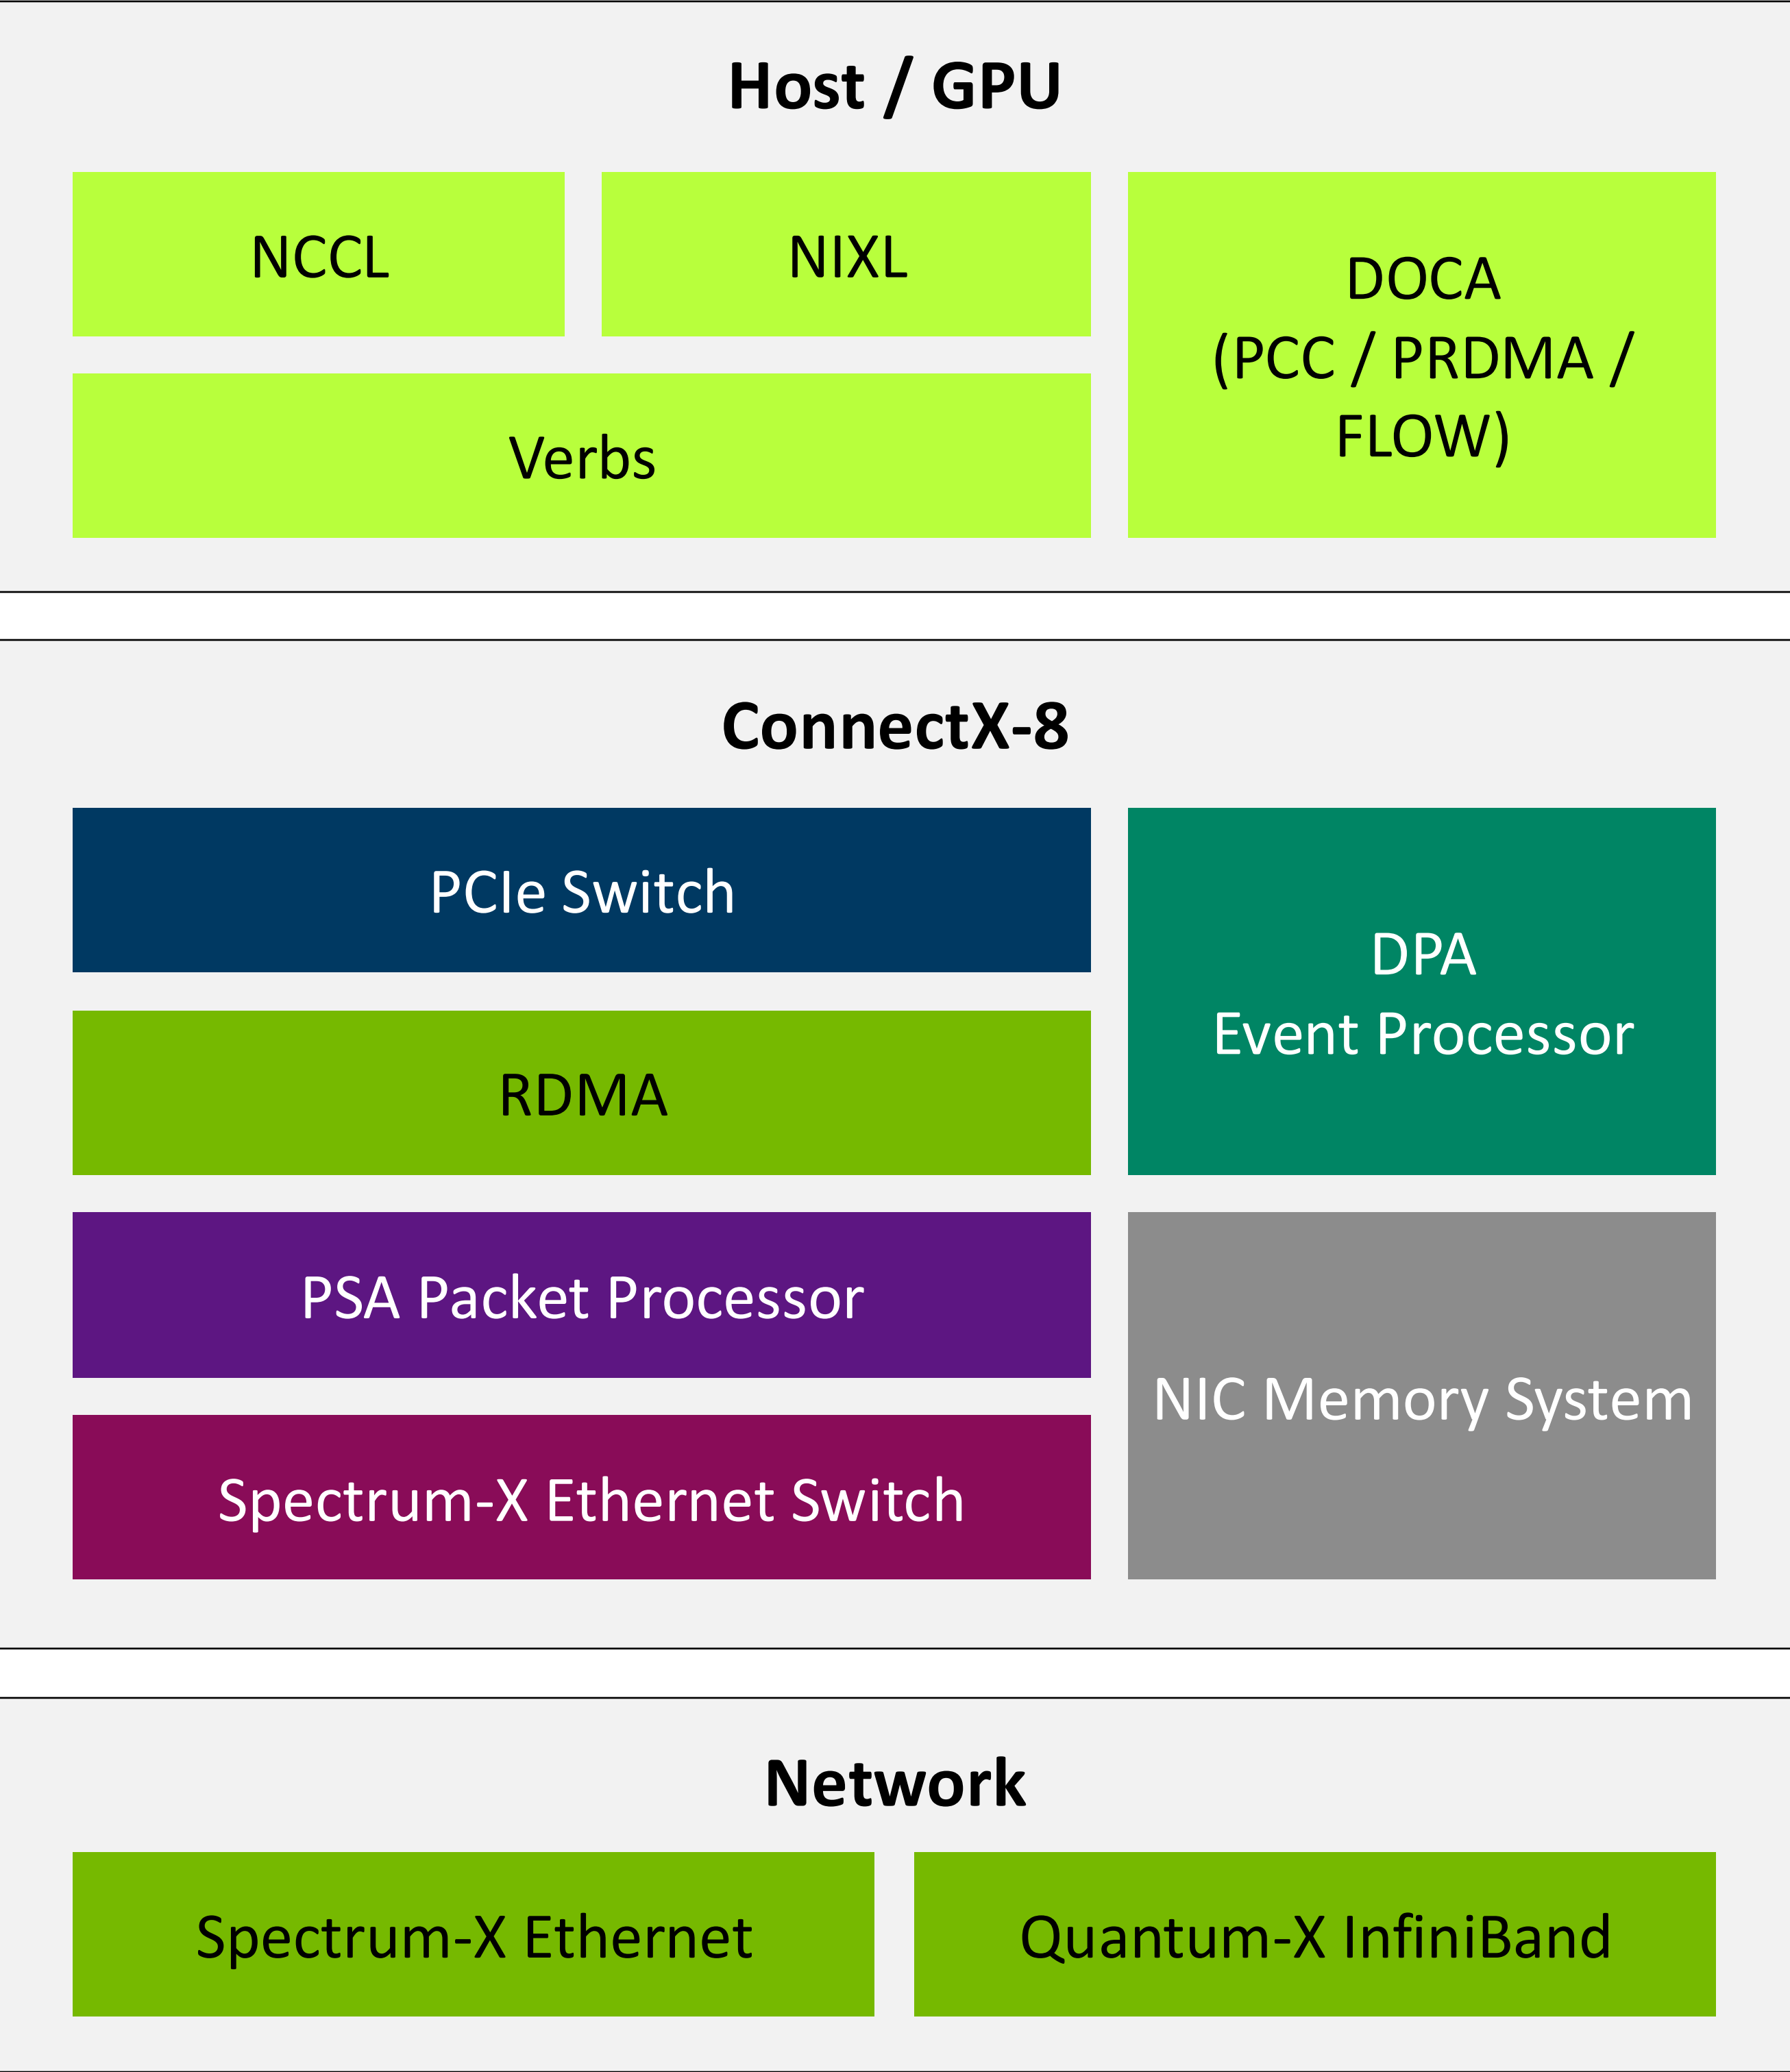


# ConnectX-8 is a SuperNIC!

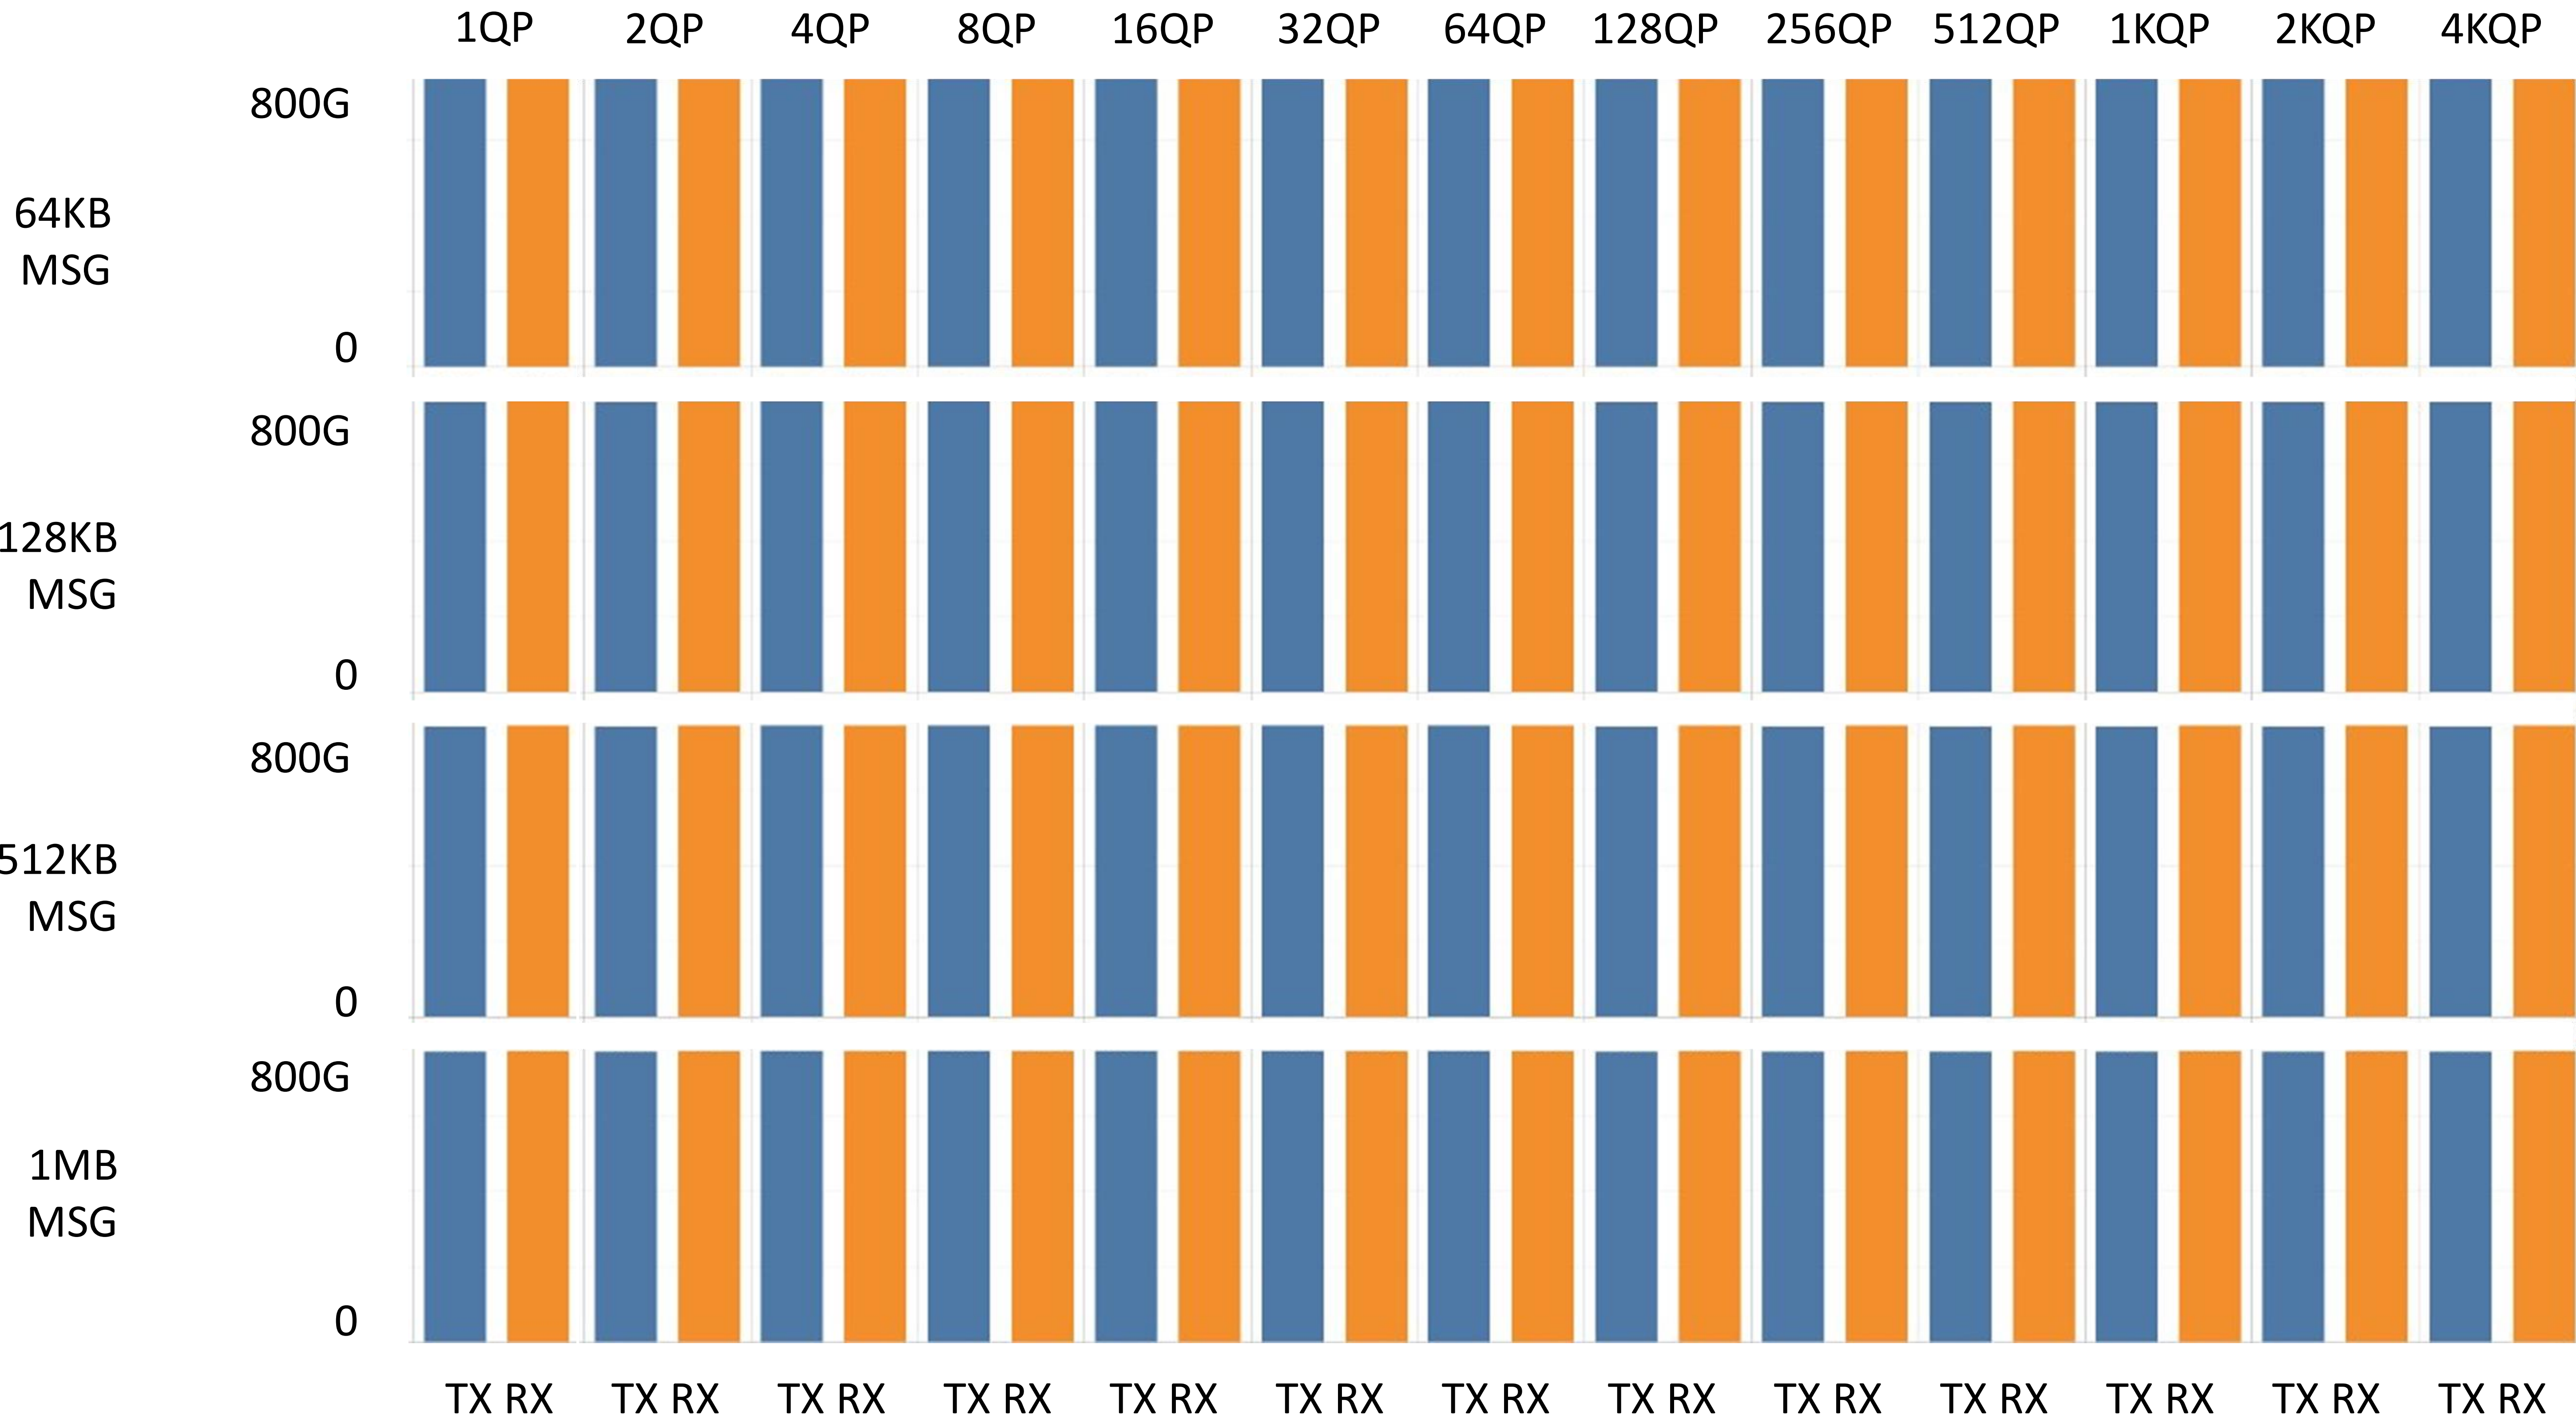


**ConnectX-8 800G SuperNIC**  
Redefining networking for AI factory

<b>Application API</b>	Verbs, NCCL, NIXL, DOCA
<b>Network interface</b>	800Gb/s InfiniBand XDR or 2x400G Ethernet Integrated Spectrum-X Ethernet up to 8 ports
<b>Host/IO interface</b>	48 lanes PCIe G6, integrated PCIe switch
<b>RDMA</b>	Line rate RDMA transport with out of order delivery Programable congestion control and retransmission
<b>Packet processing</b>	Line rate network encryption and multitenancy Programable routing
<b>Programmability</b>	DPA – RISC-V transport event processing PSA – Multithreaded packet processor
<b>Software API</b>	DOCA-PRDMA, DOCA-PCC, DOCA-FLOW

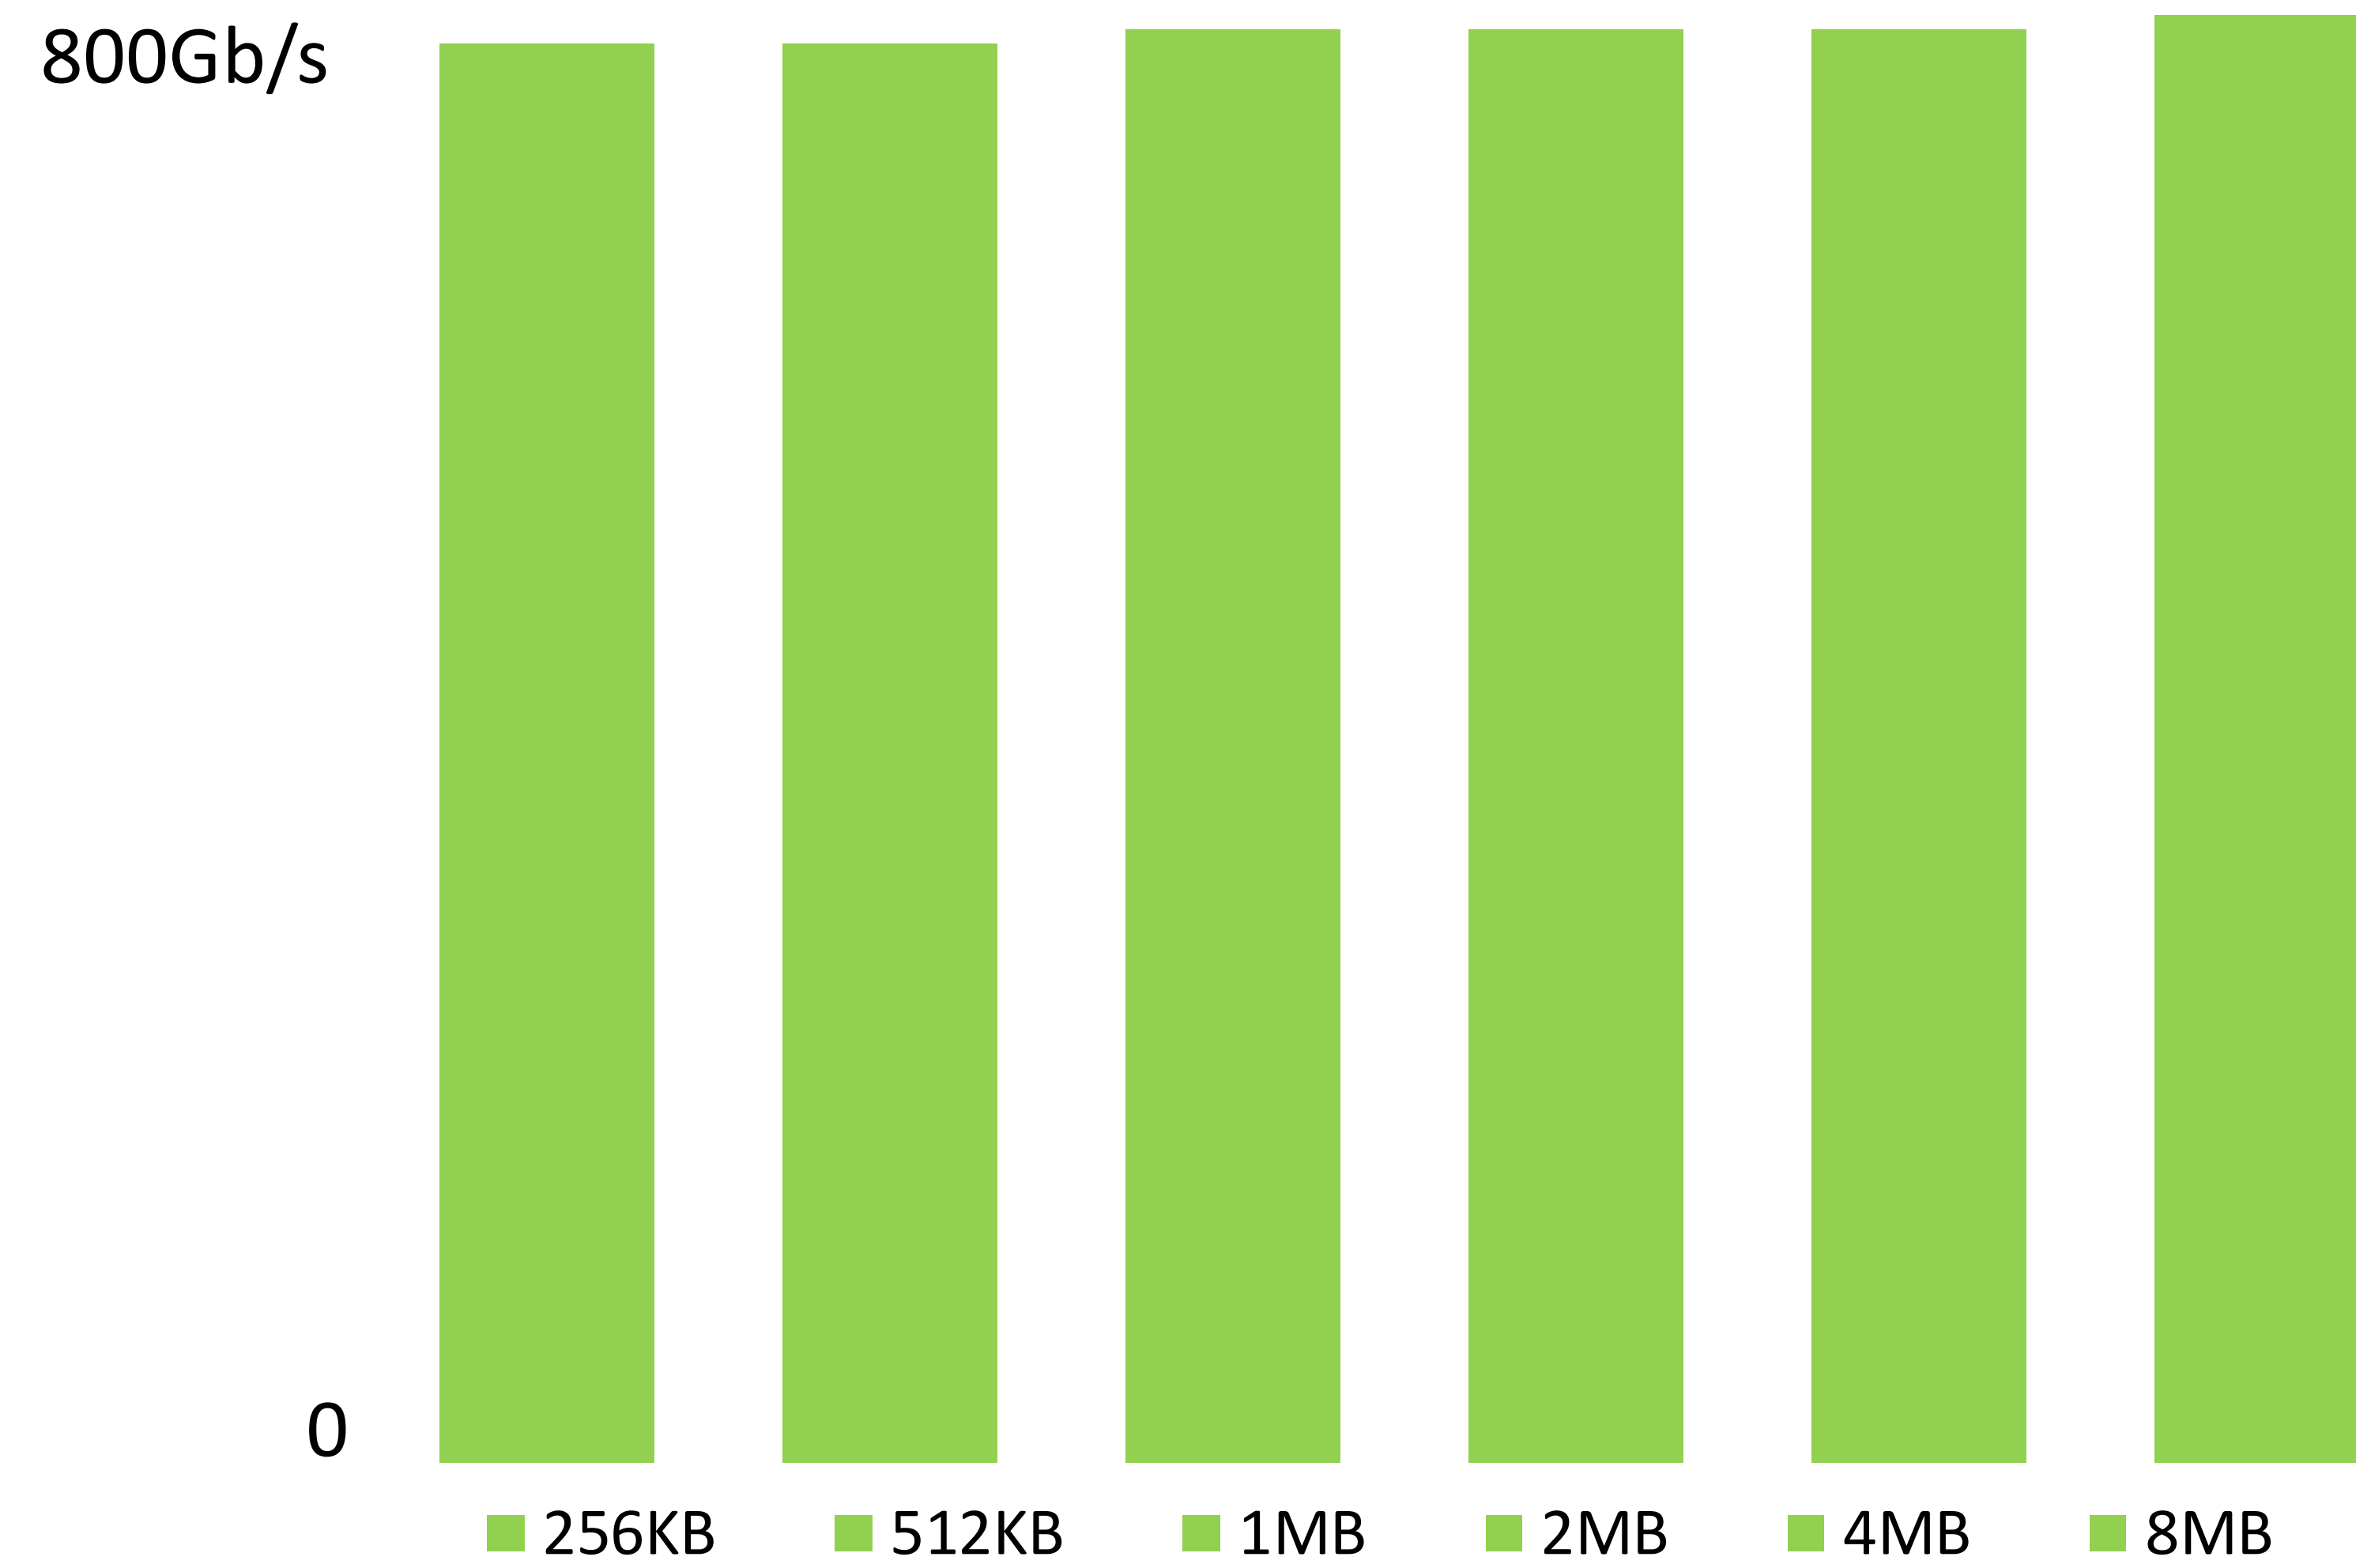


# ConnectX-8 RDMA Limitless Scaling at 800G




# ConnectX-8 NCCL Performance

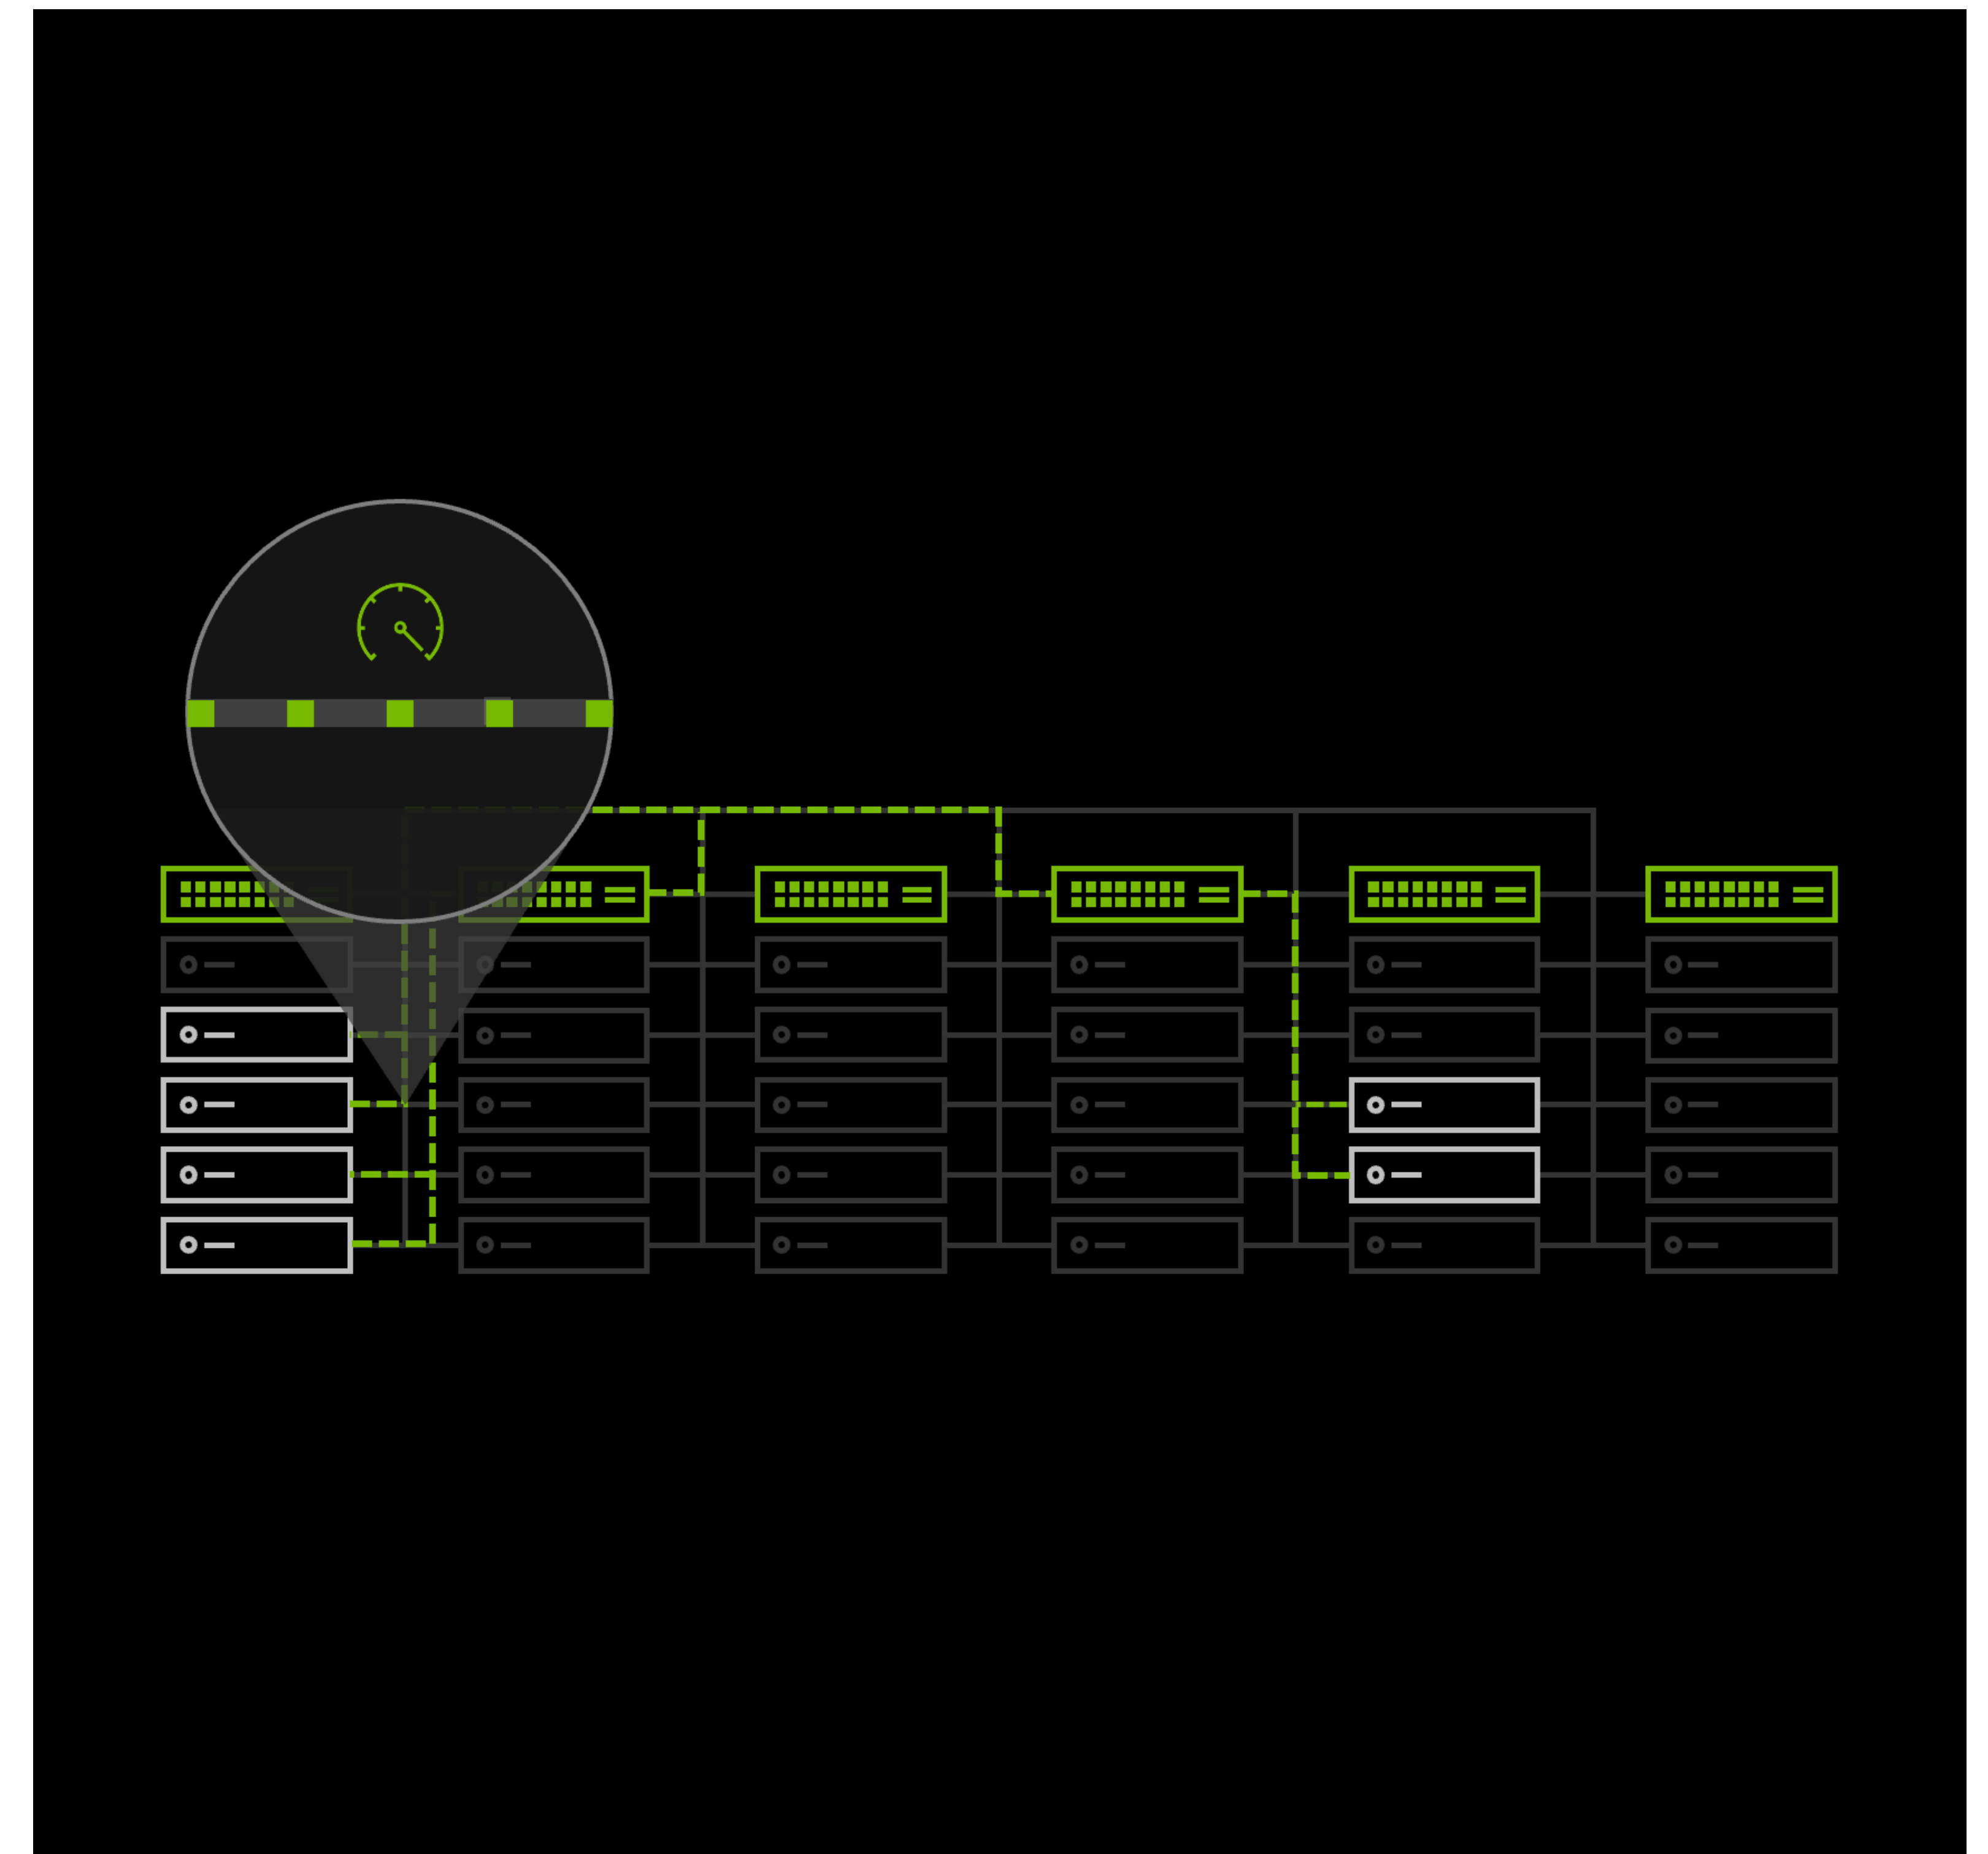
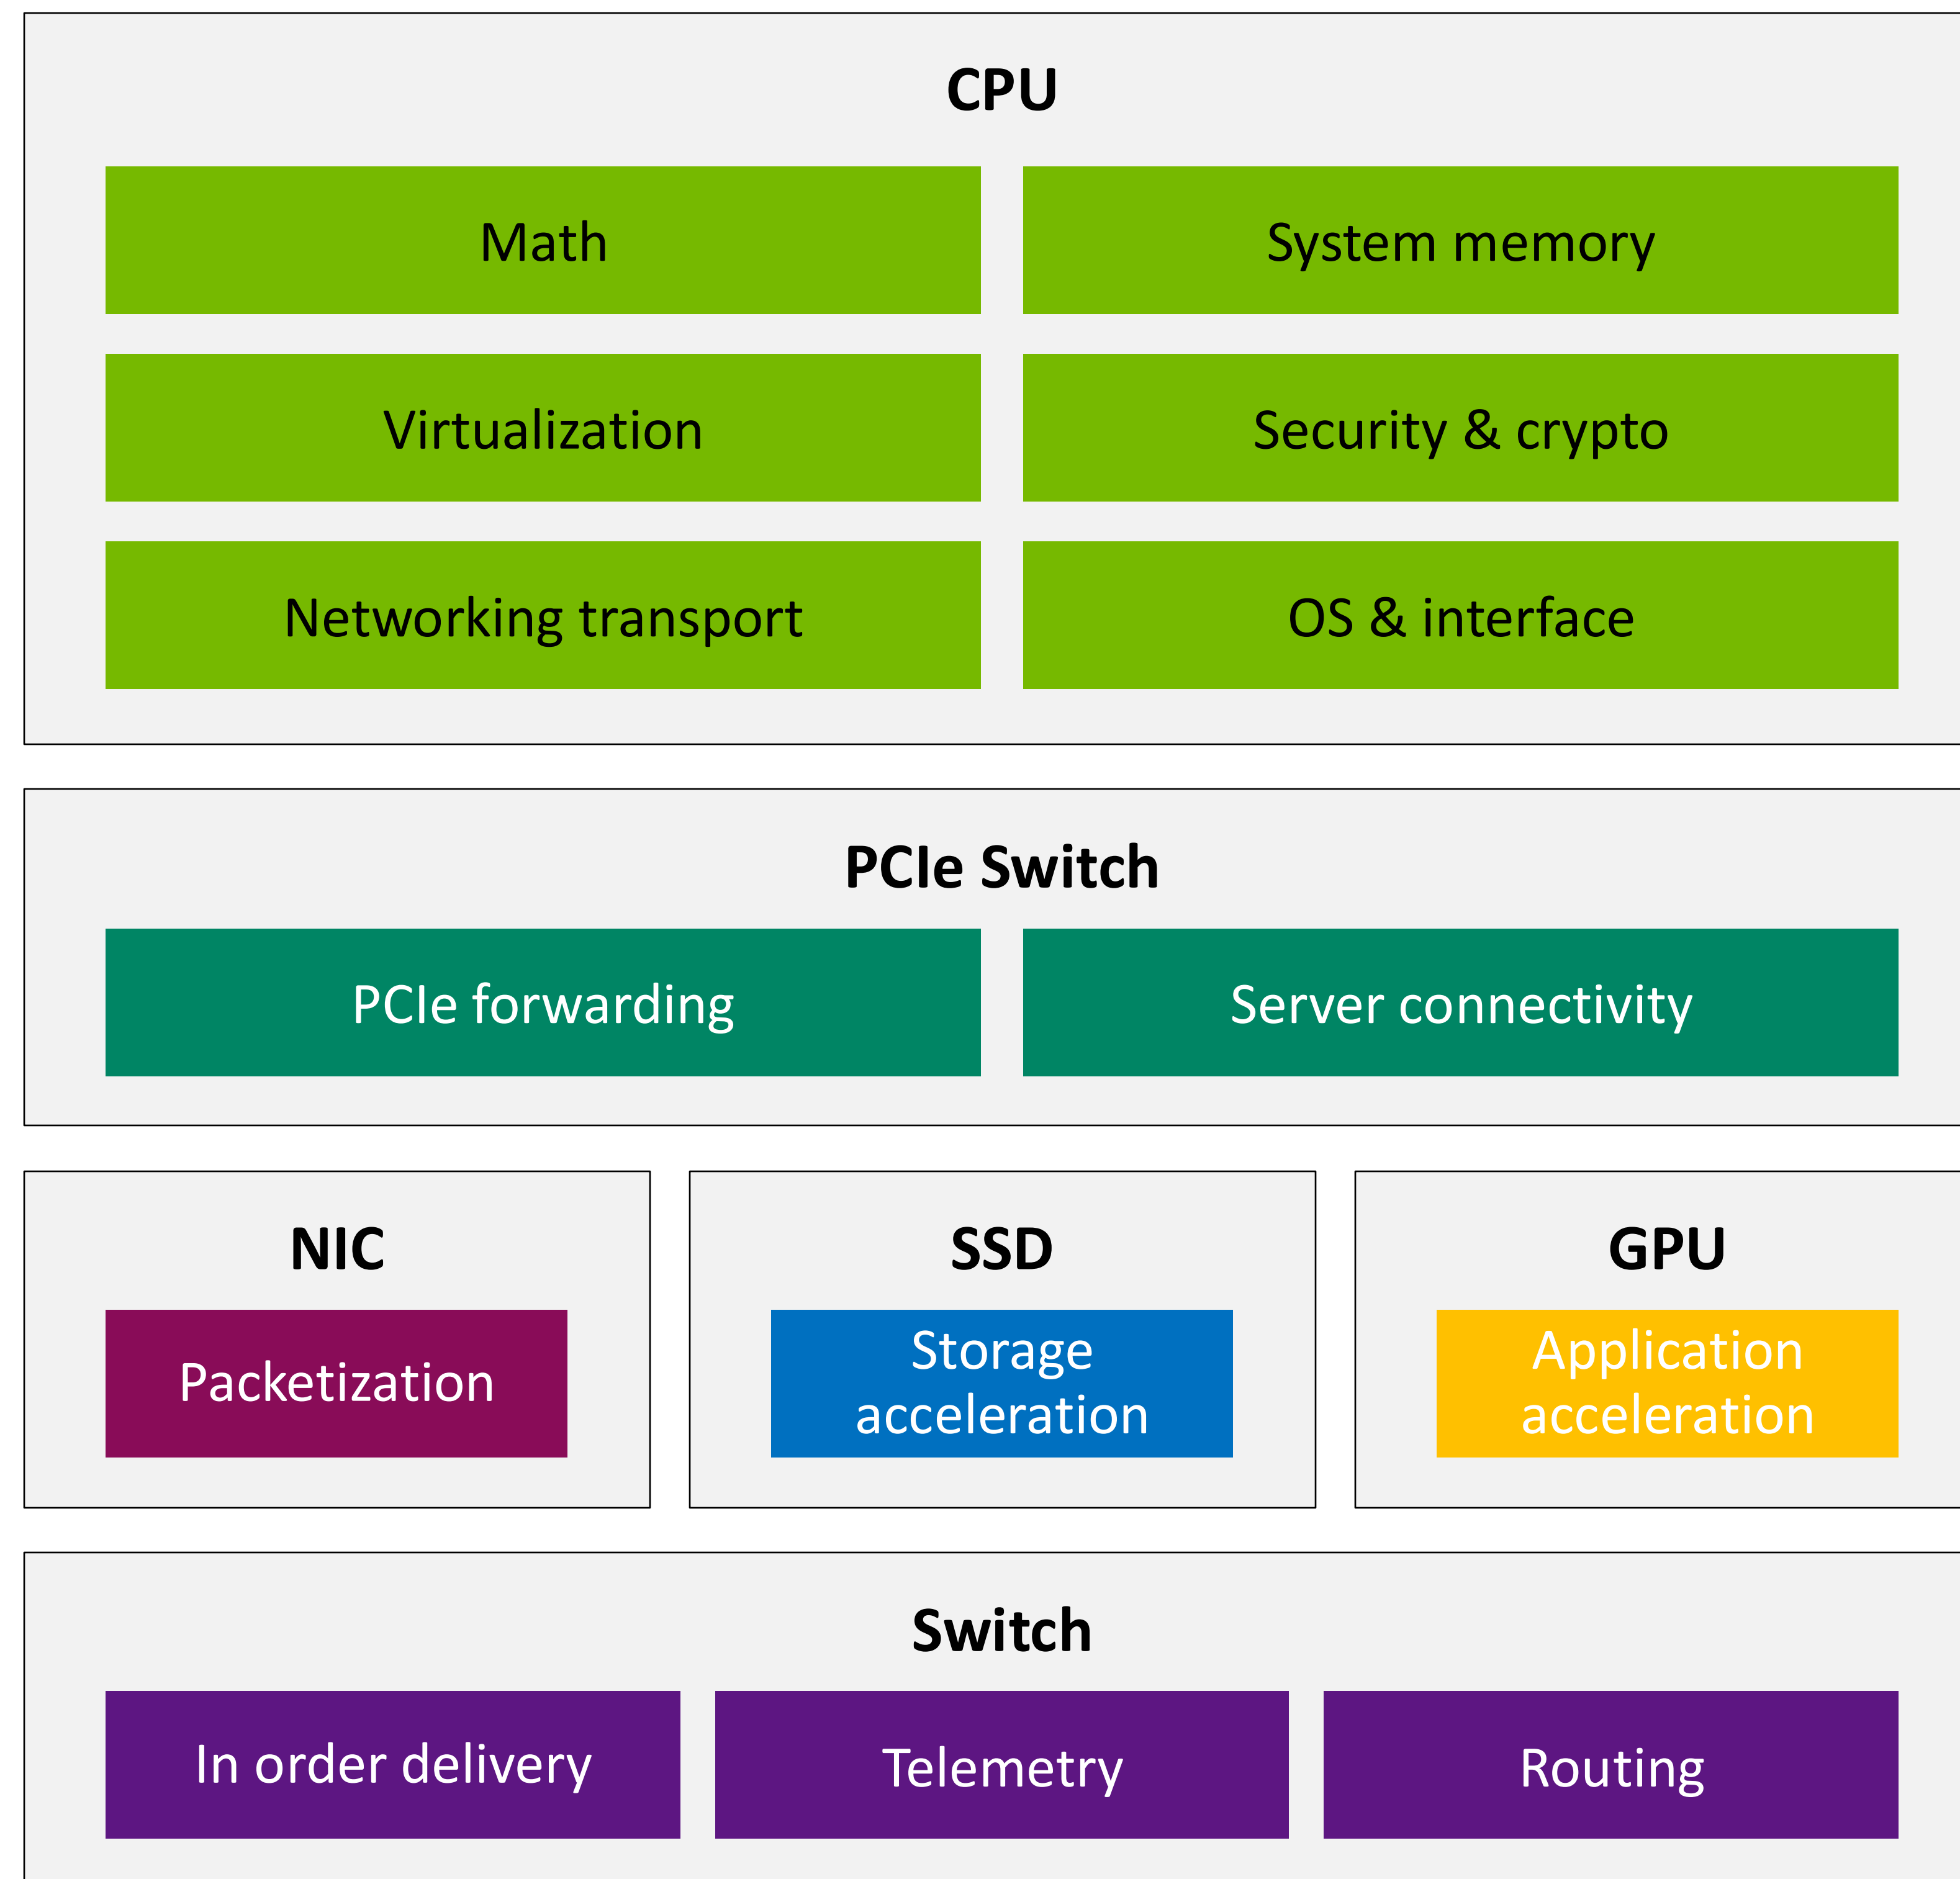
## NCCL AllReduce Performance



## NCCL AlltAll Performance

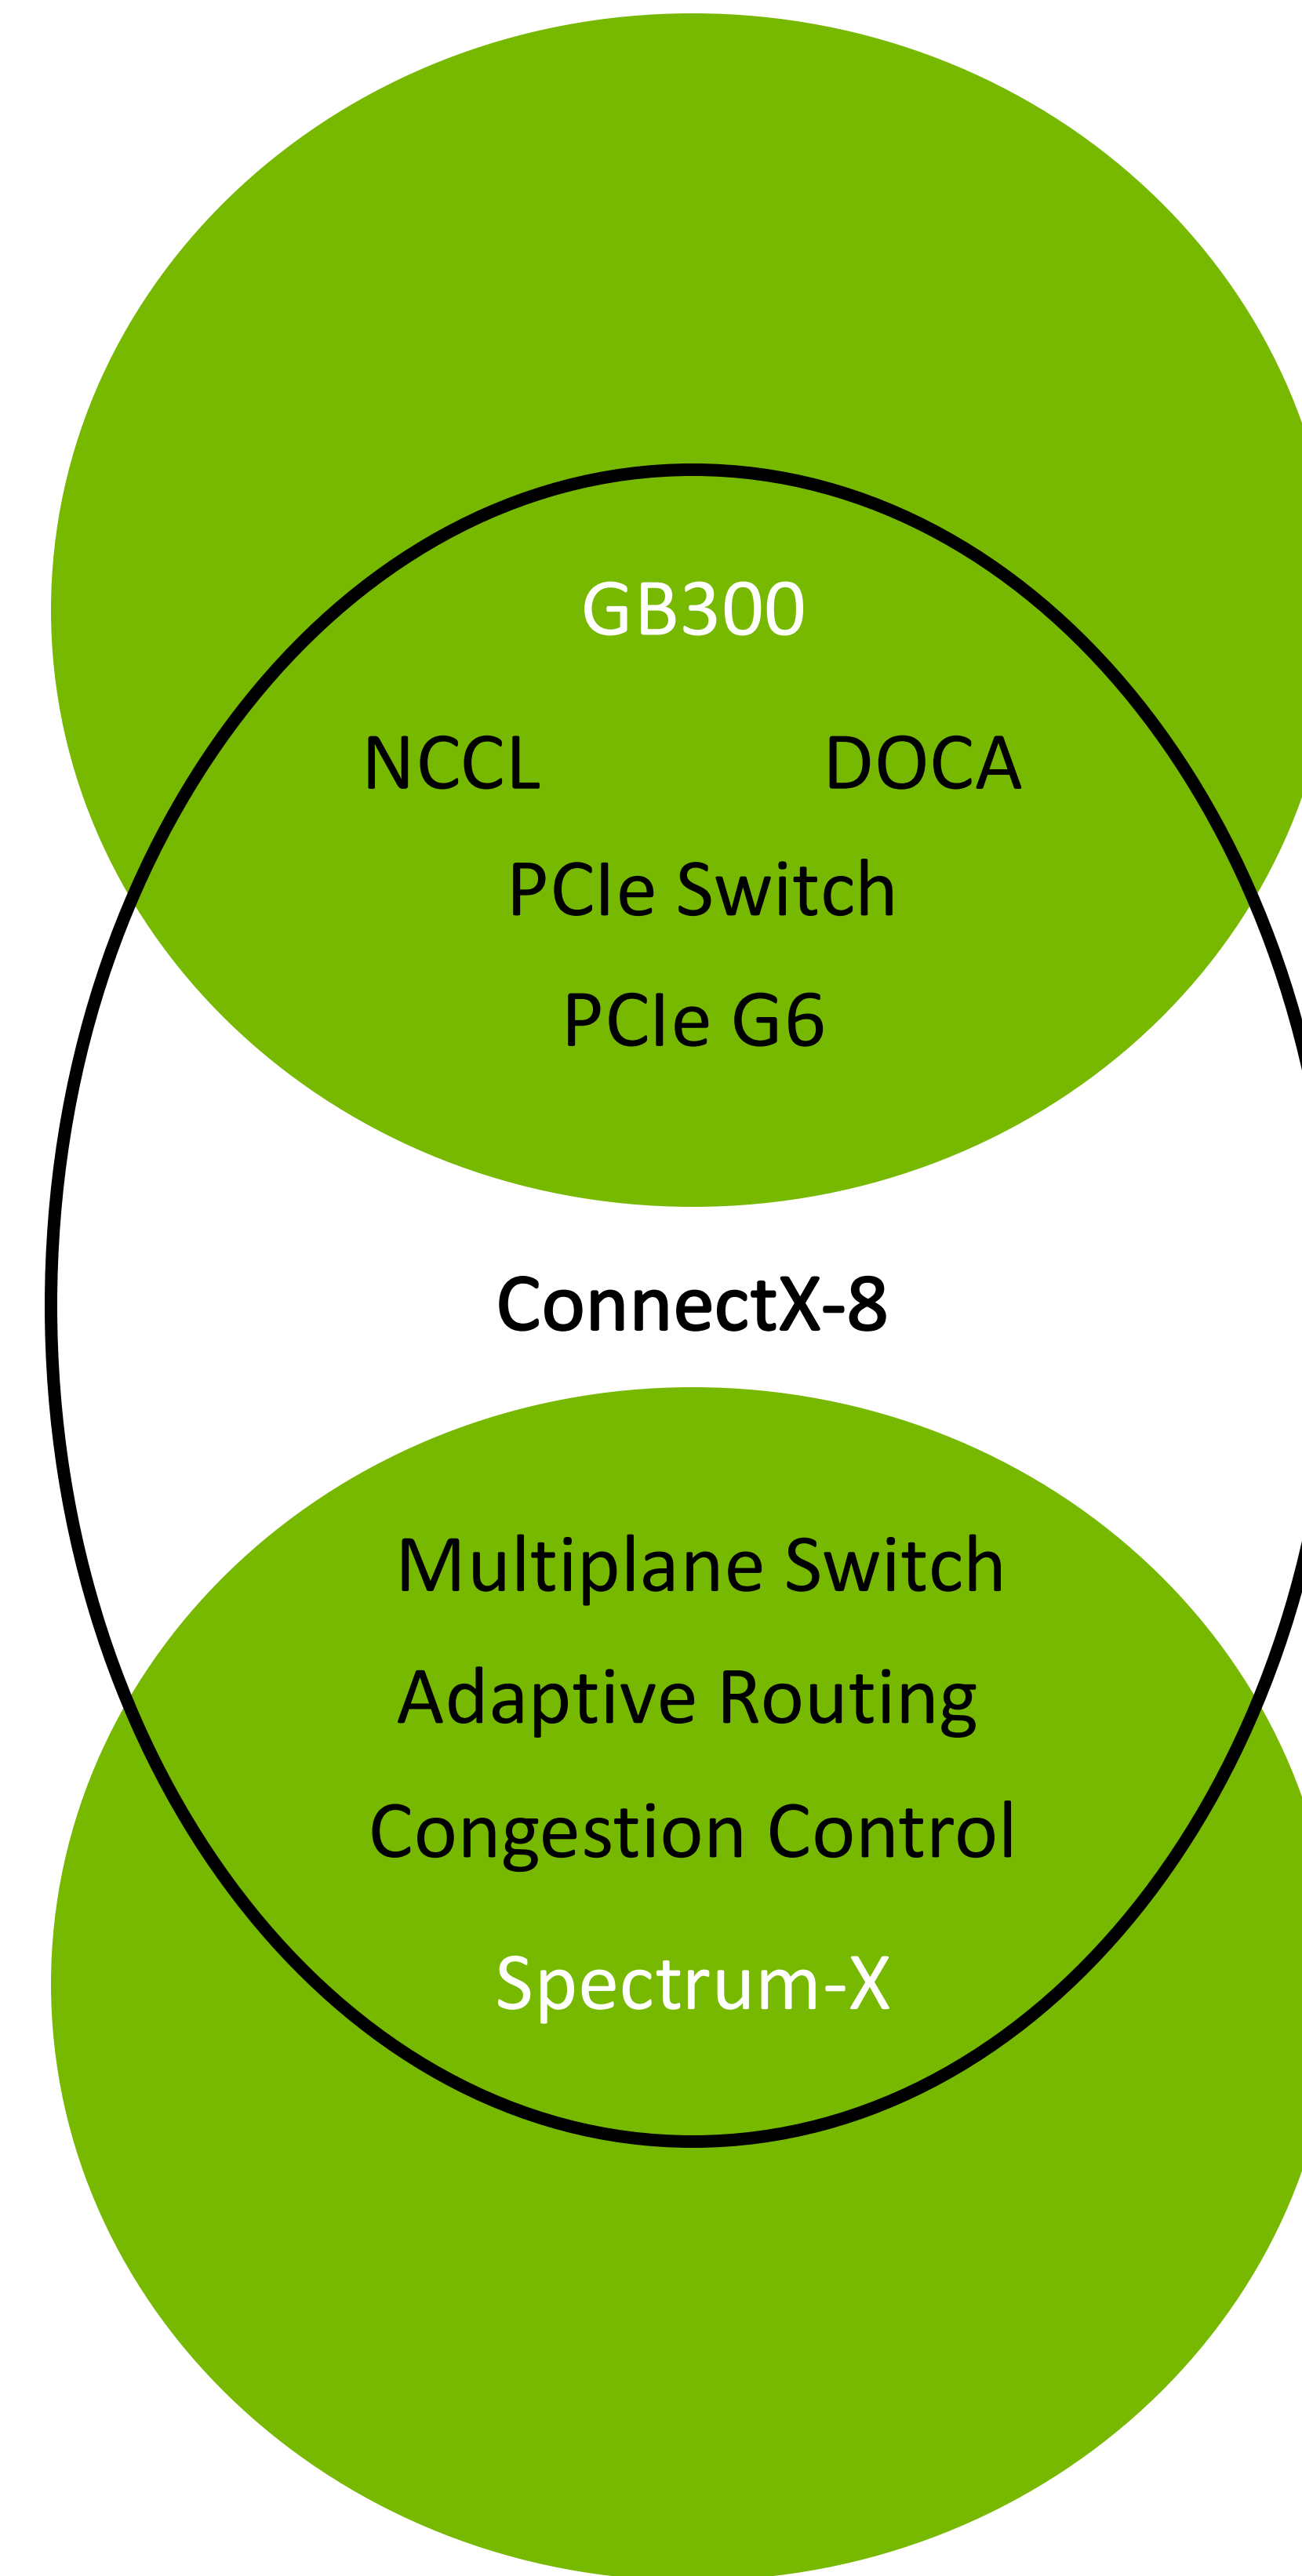


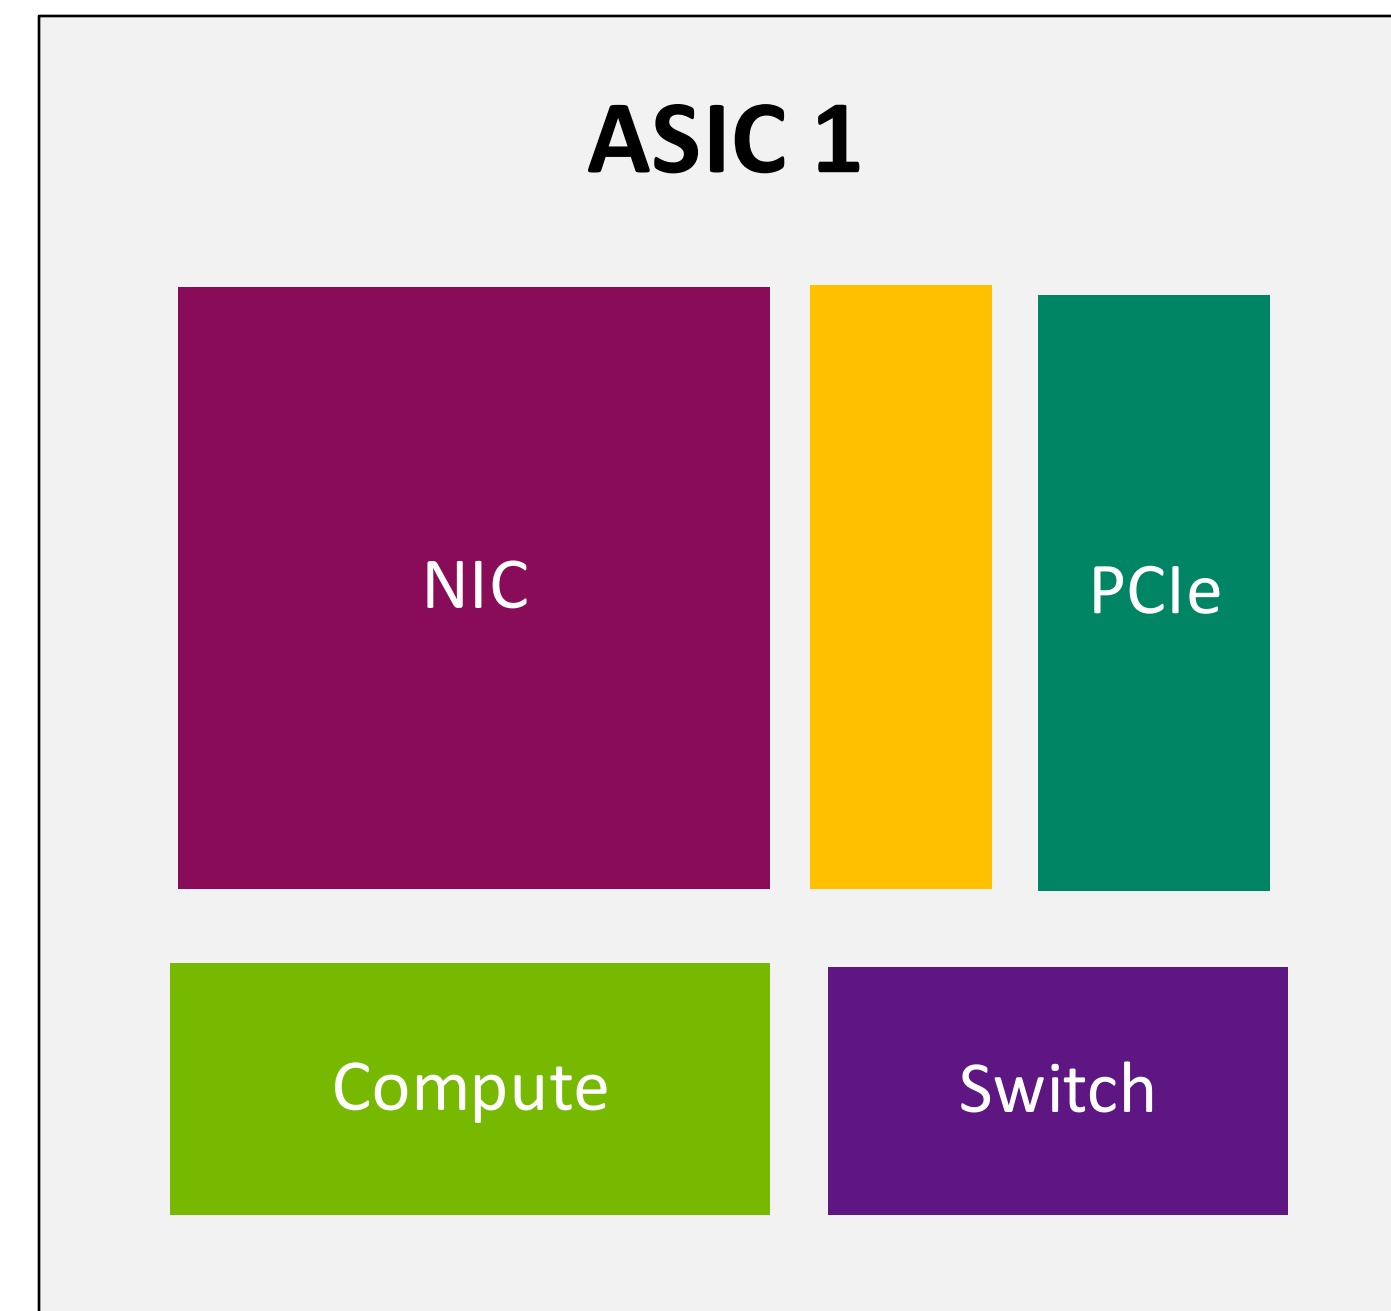


# Traditional Data Center – Collection of Pre-Defined Components



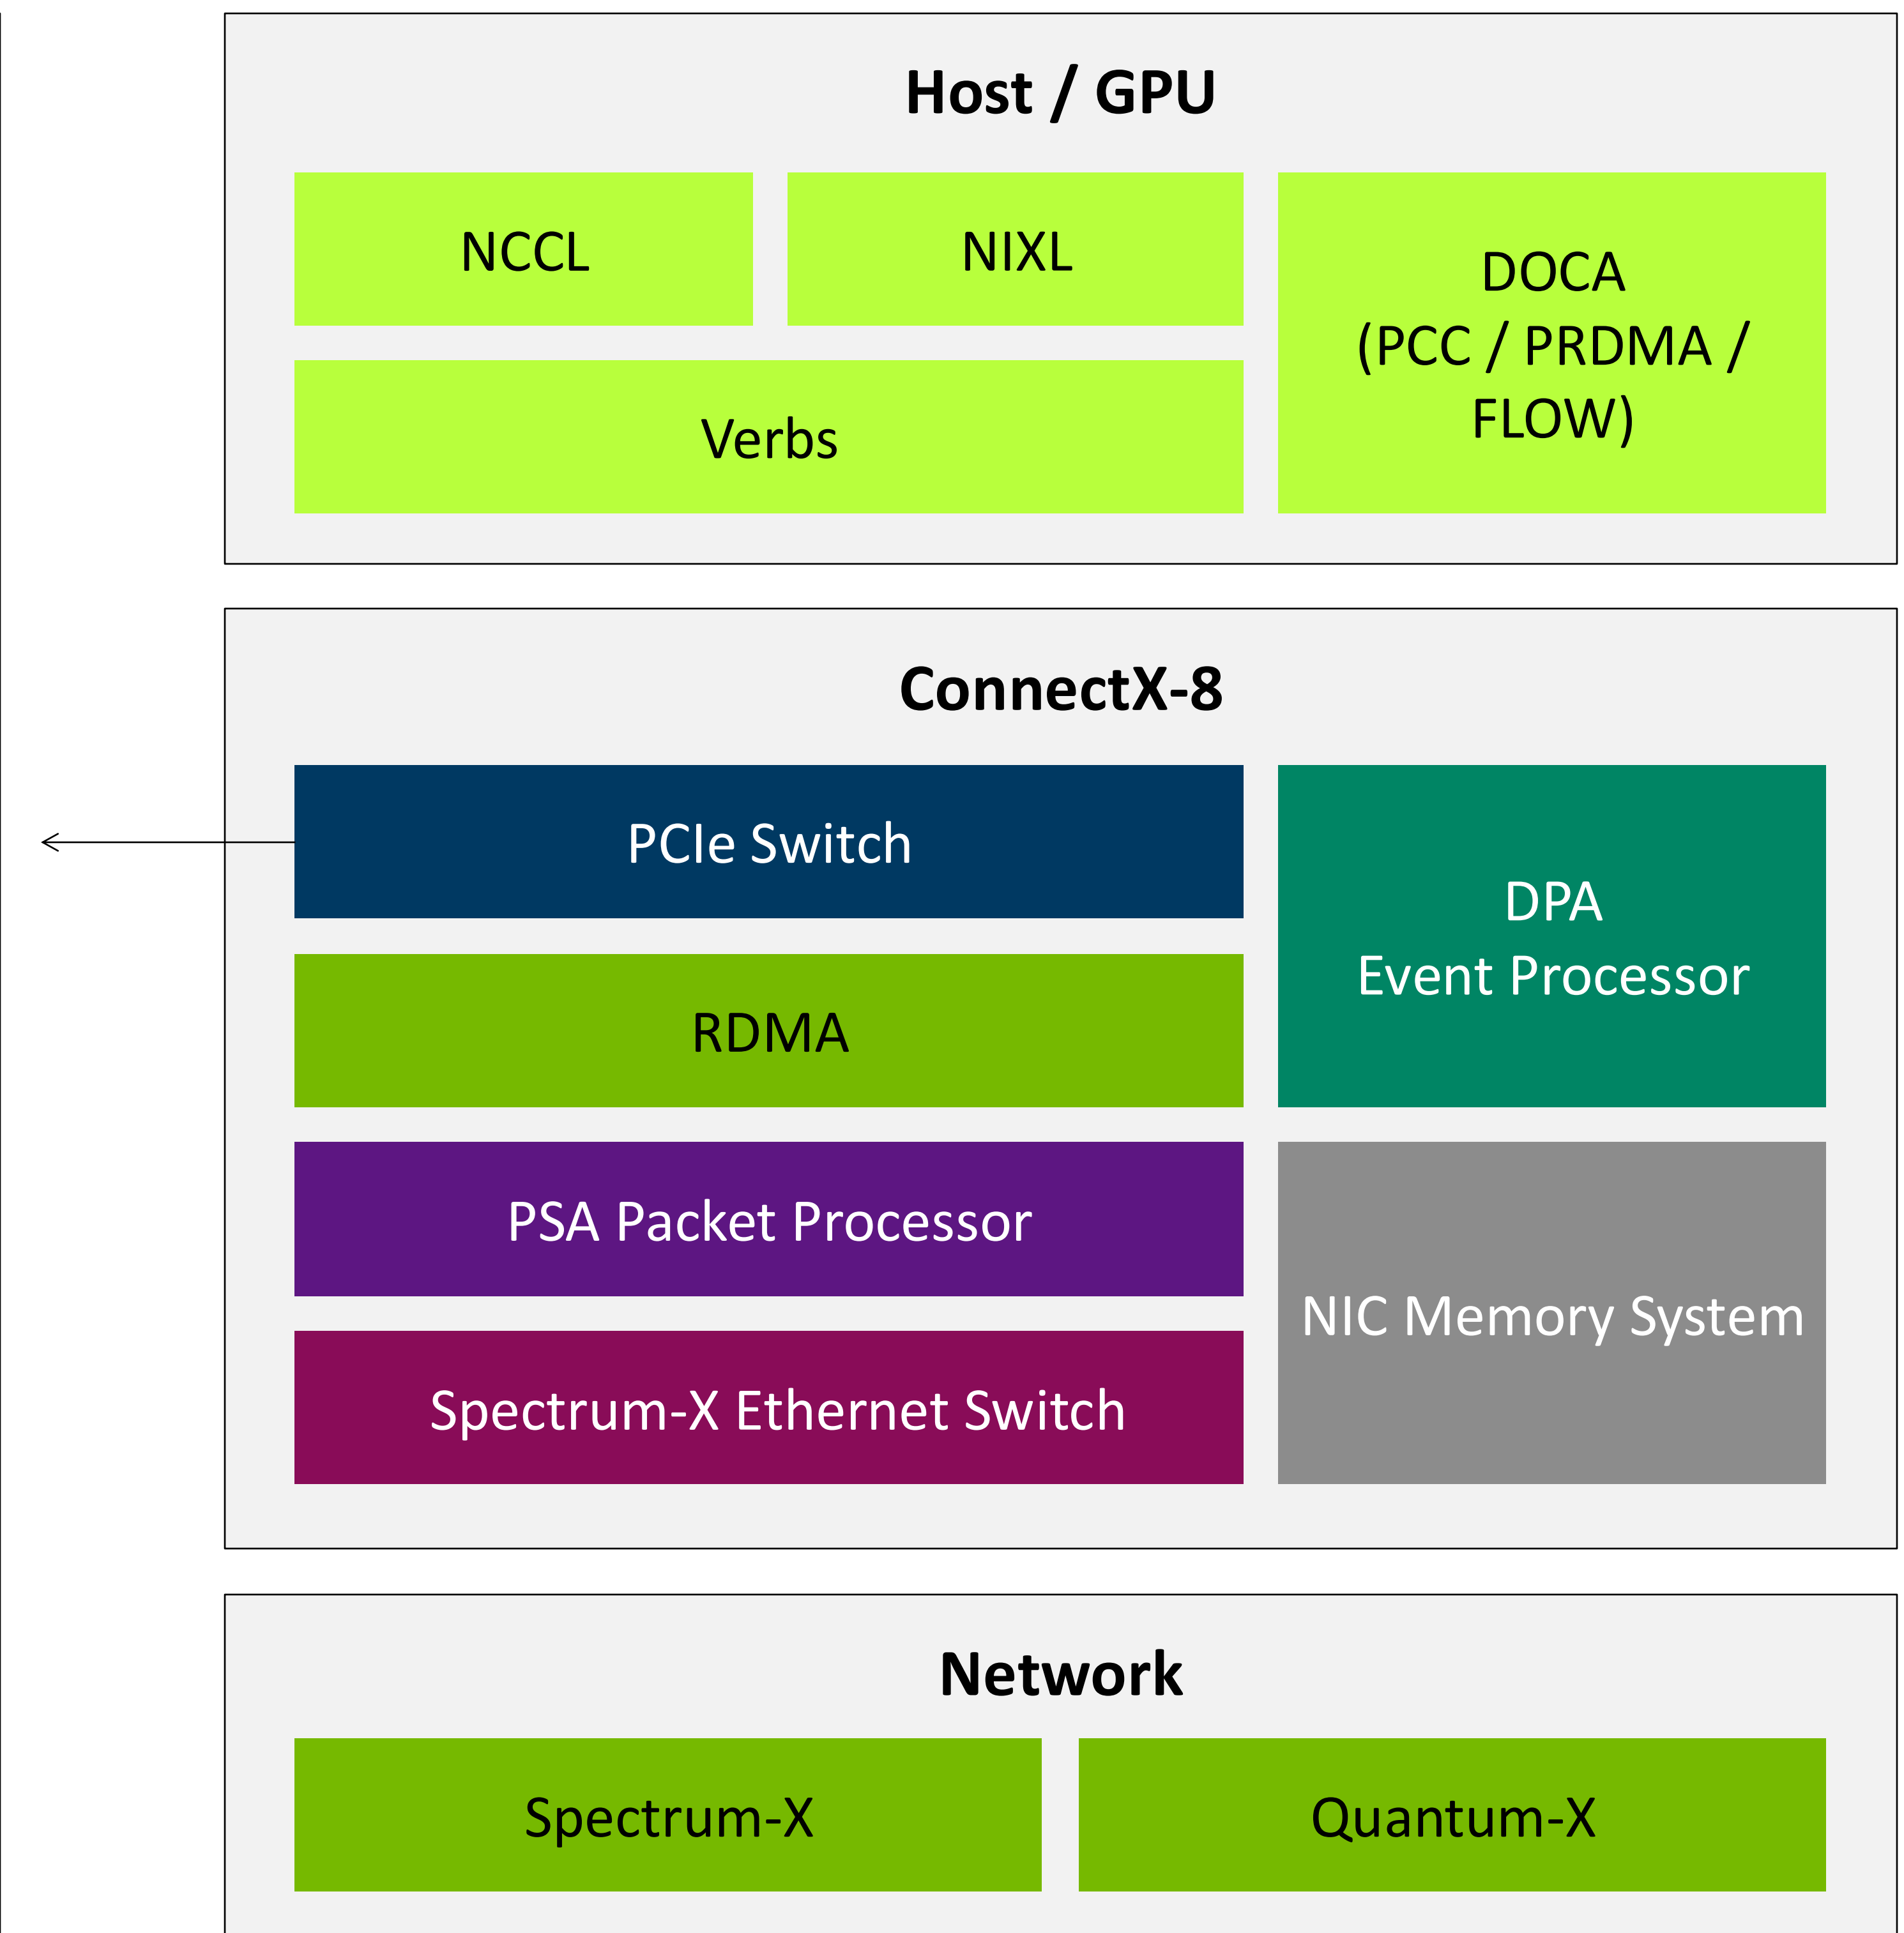
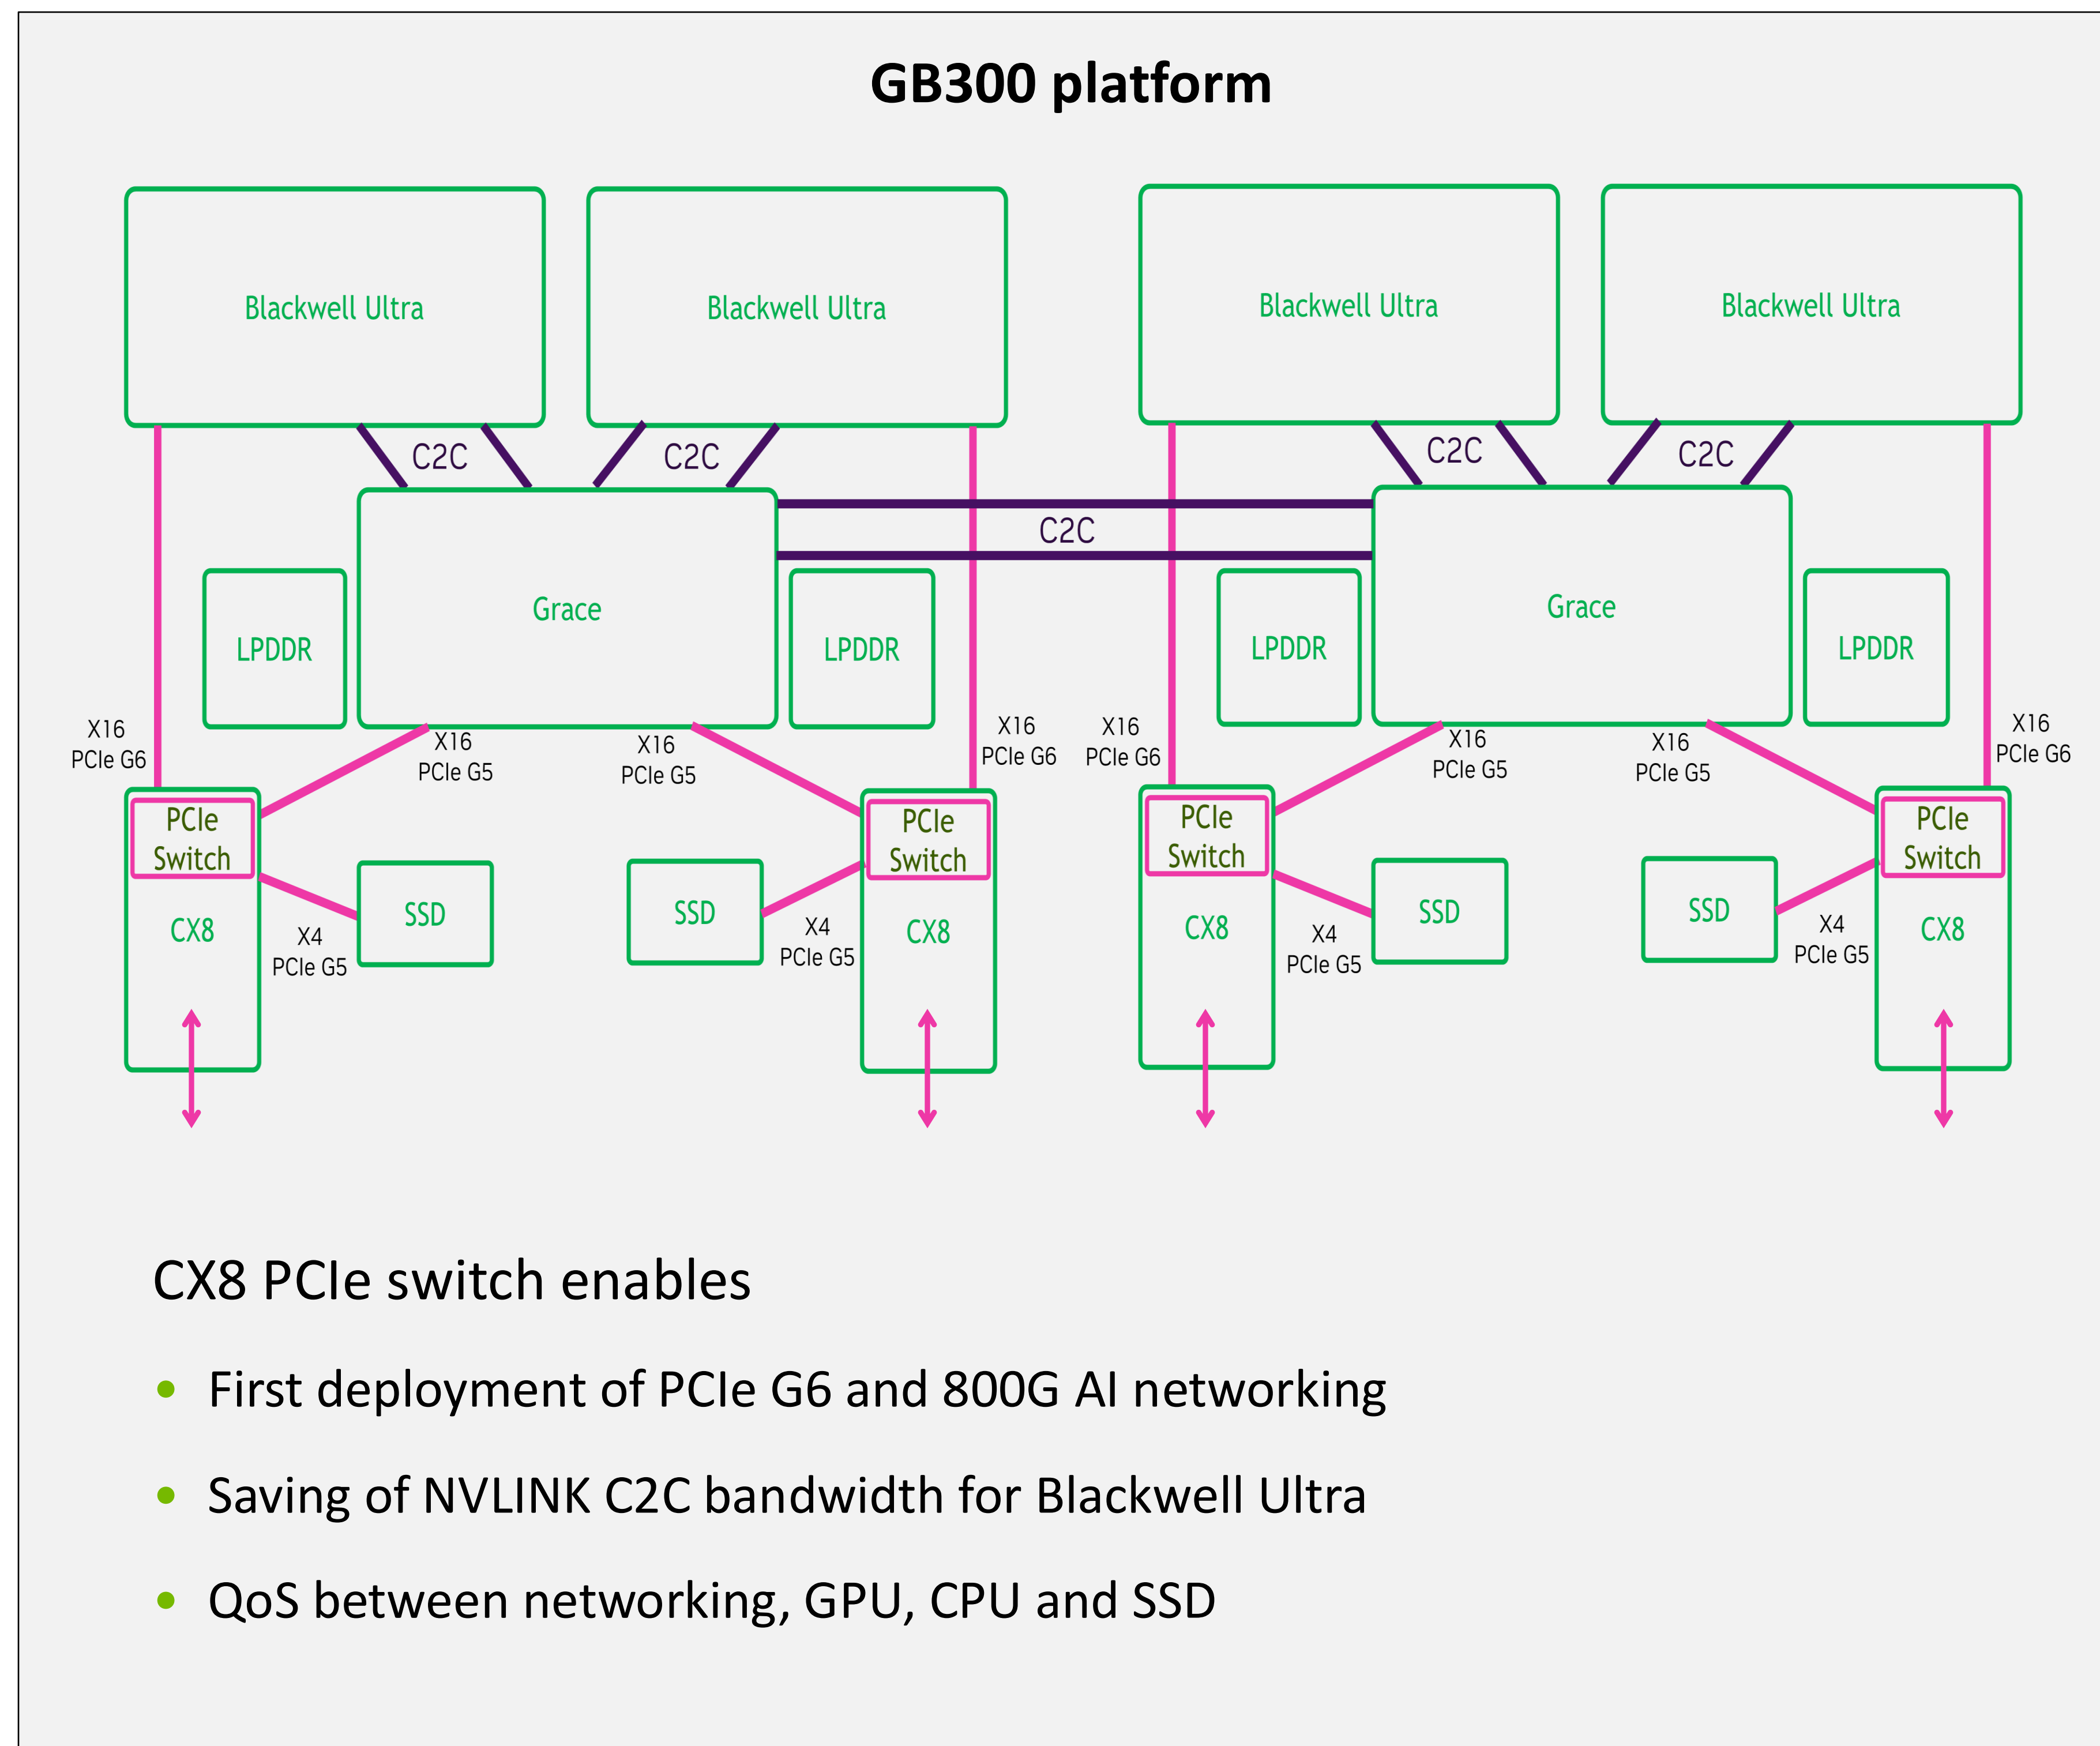
# Data Center Is the Unit of Computing – ConnectX is not just a NIC anymore

It is the ASIC that connects the GPU to the outside world

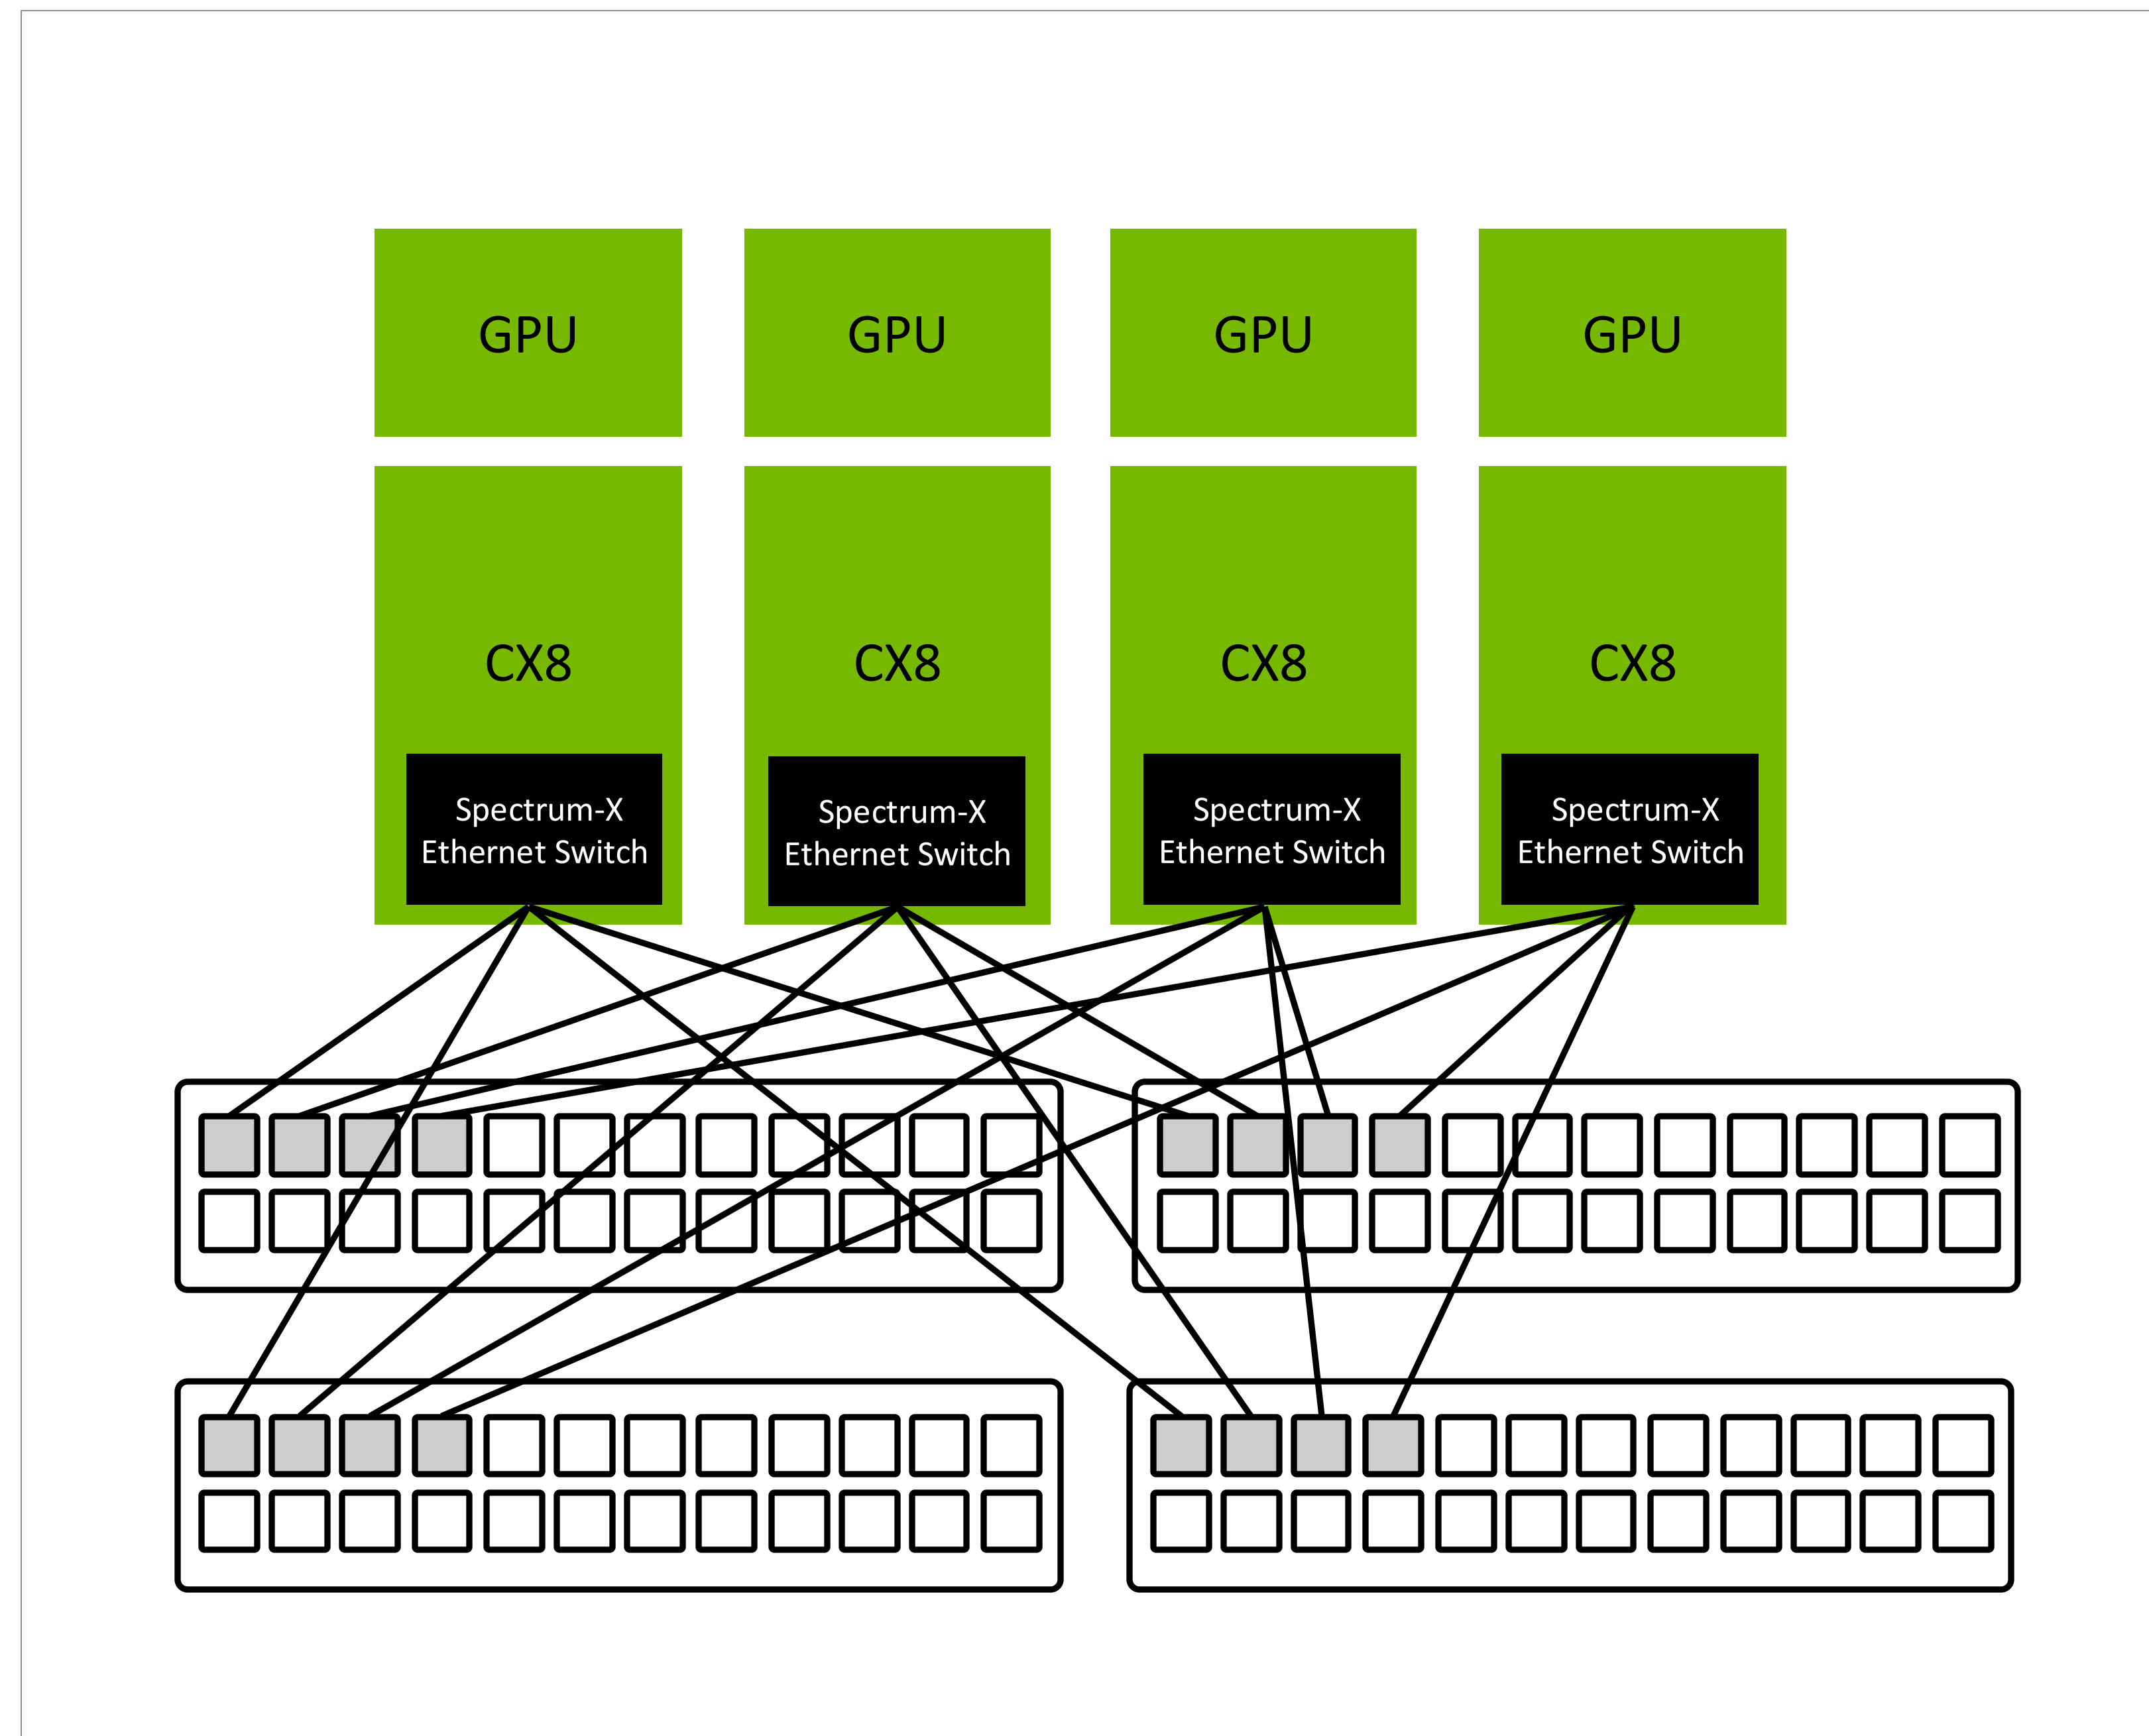


Grace Blackwell NVLink72

# ConnectX-8 PCIe Switch – First Deployment of PCIe G6 800G SuperNIC

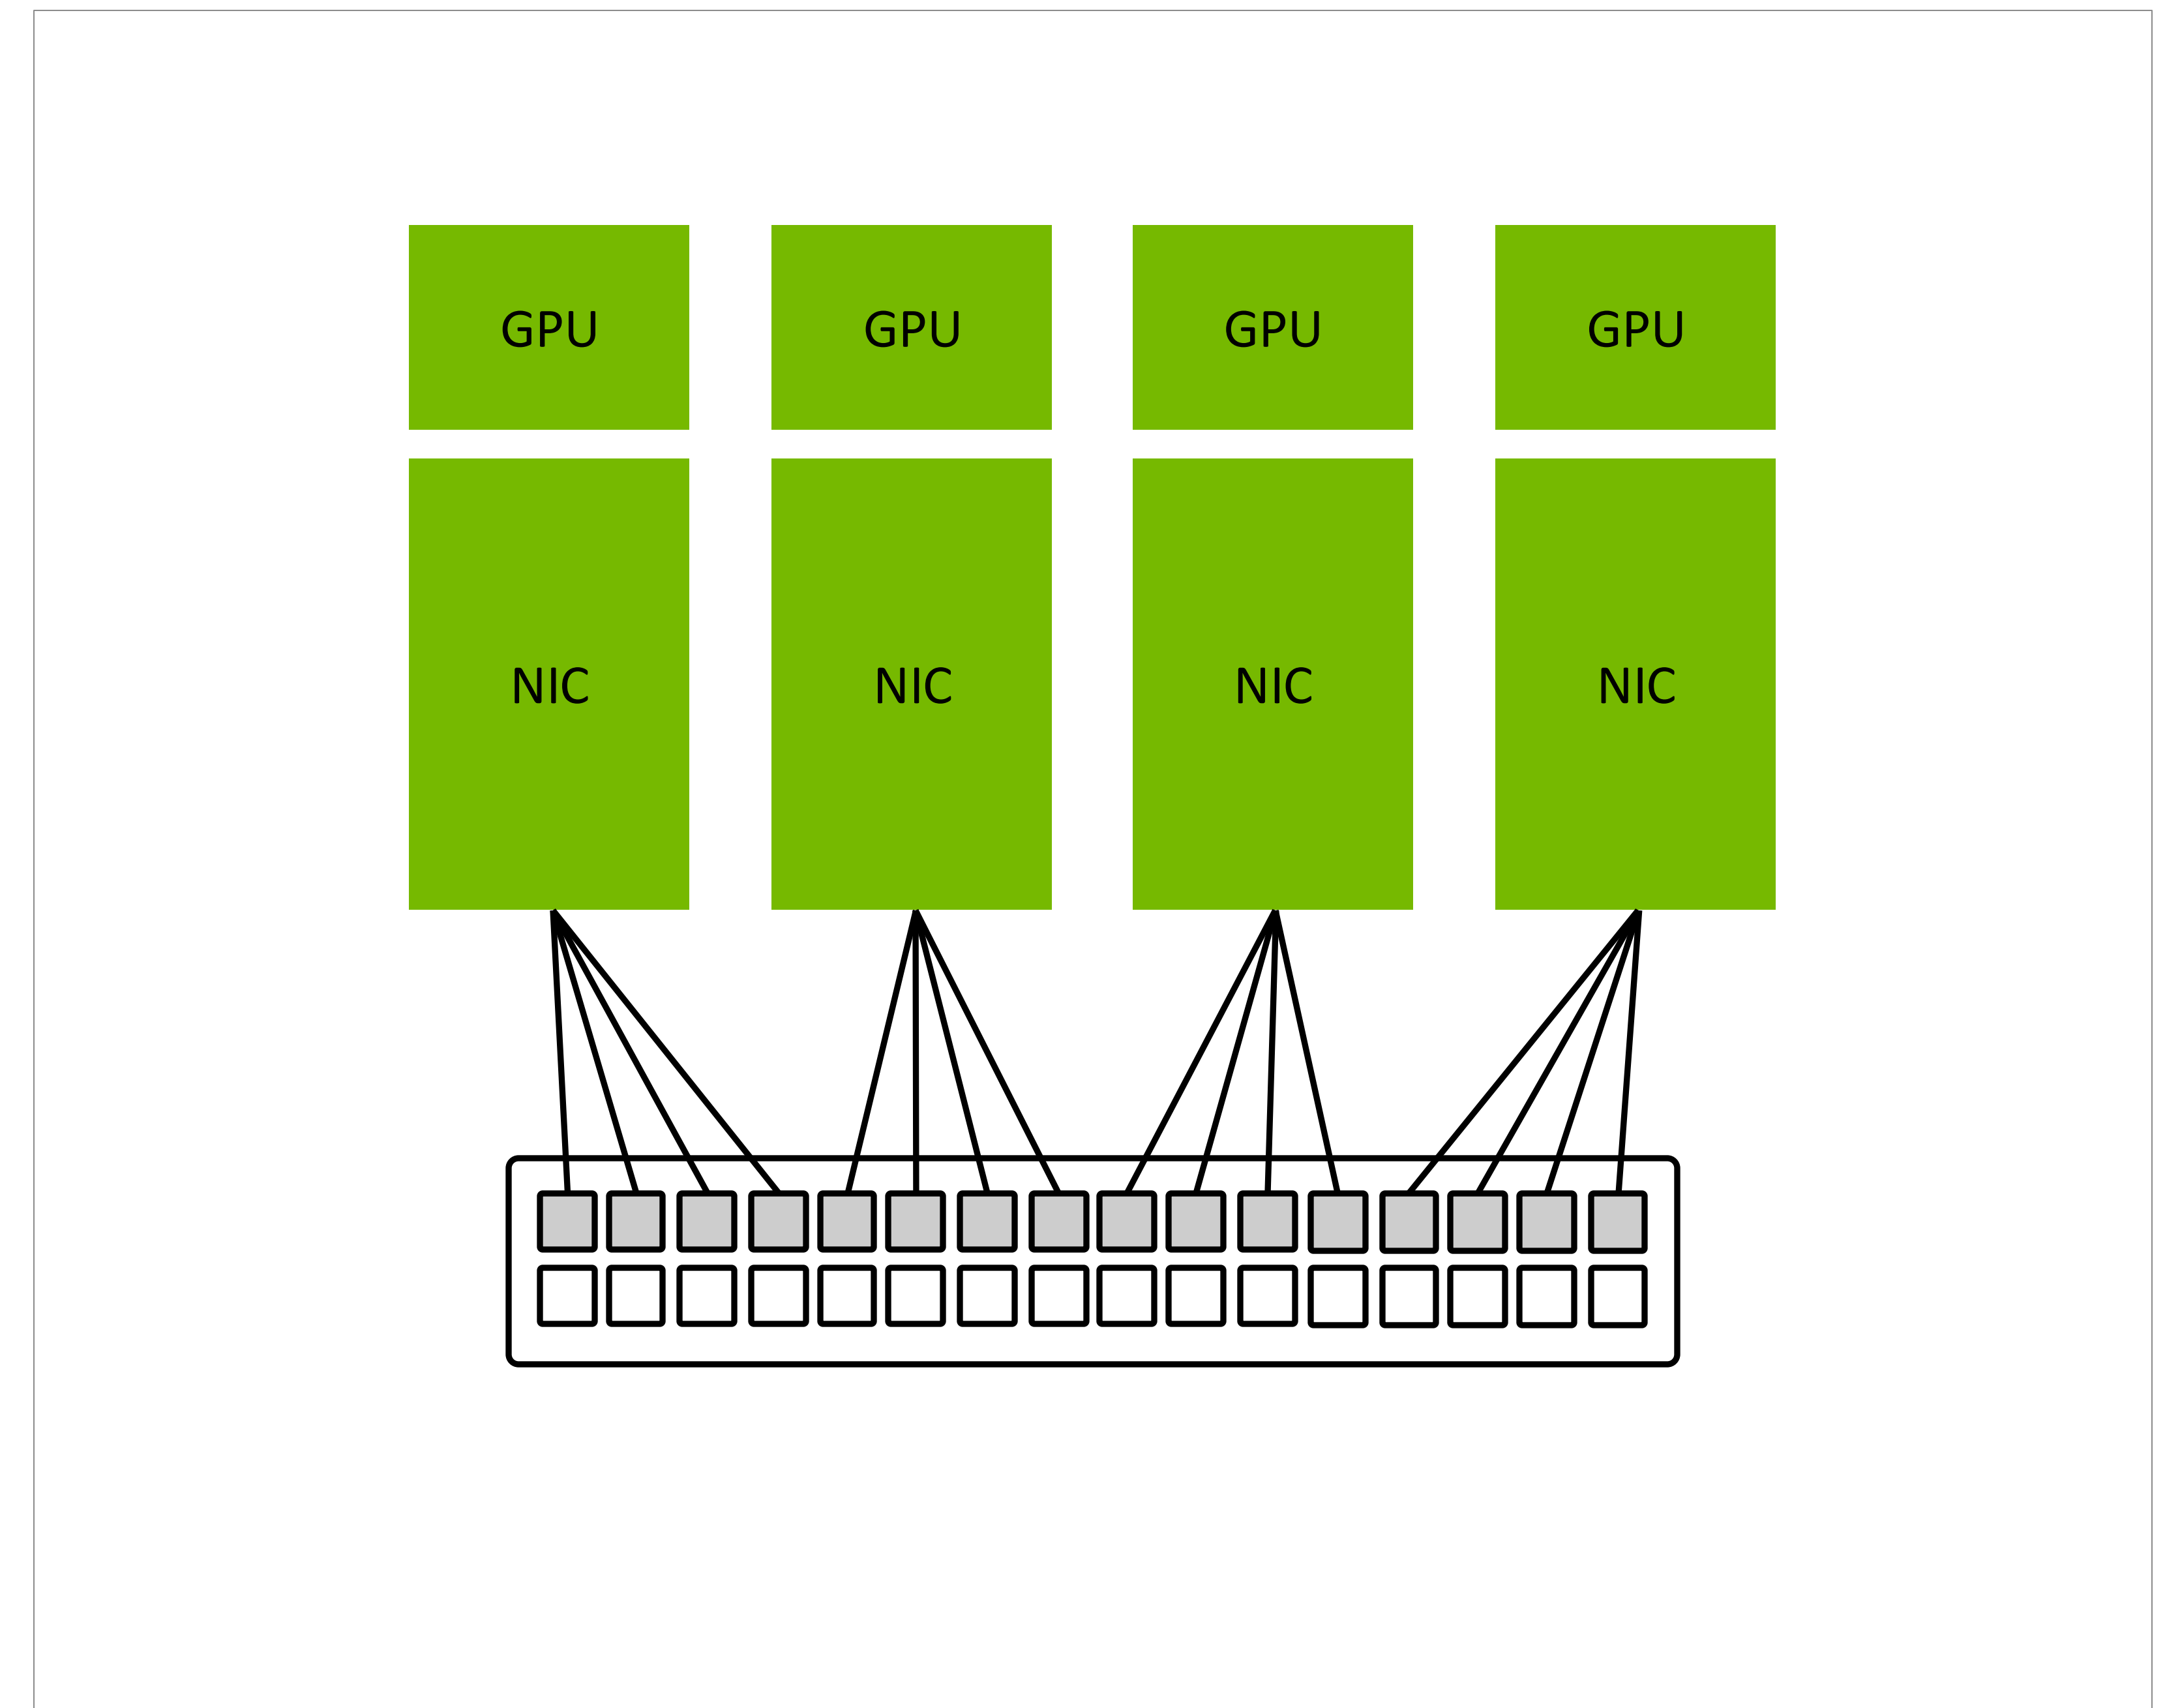


# ConnectX-8 Integrated Spectrum-X Ethernet – 64X GPU Scale



Multiplane (4 plane illustration)

Single GPU bandwidth: 800G, 8x100G  
Switch radix: 512  
GPU scale in two-level non-blocking fat-tree:  $256 \times 512 = 128K$

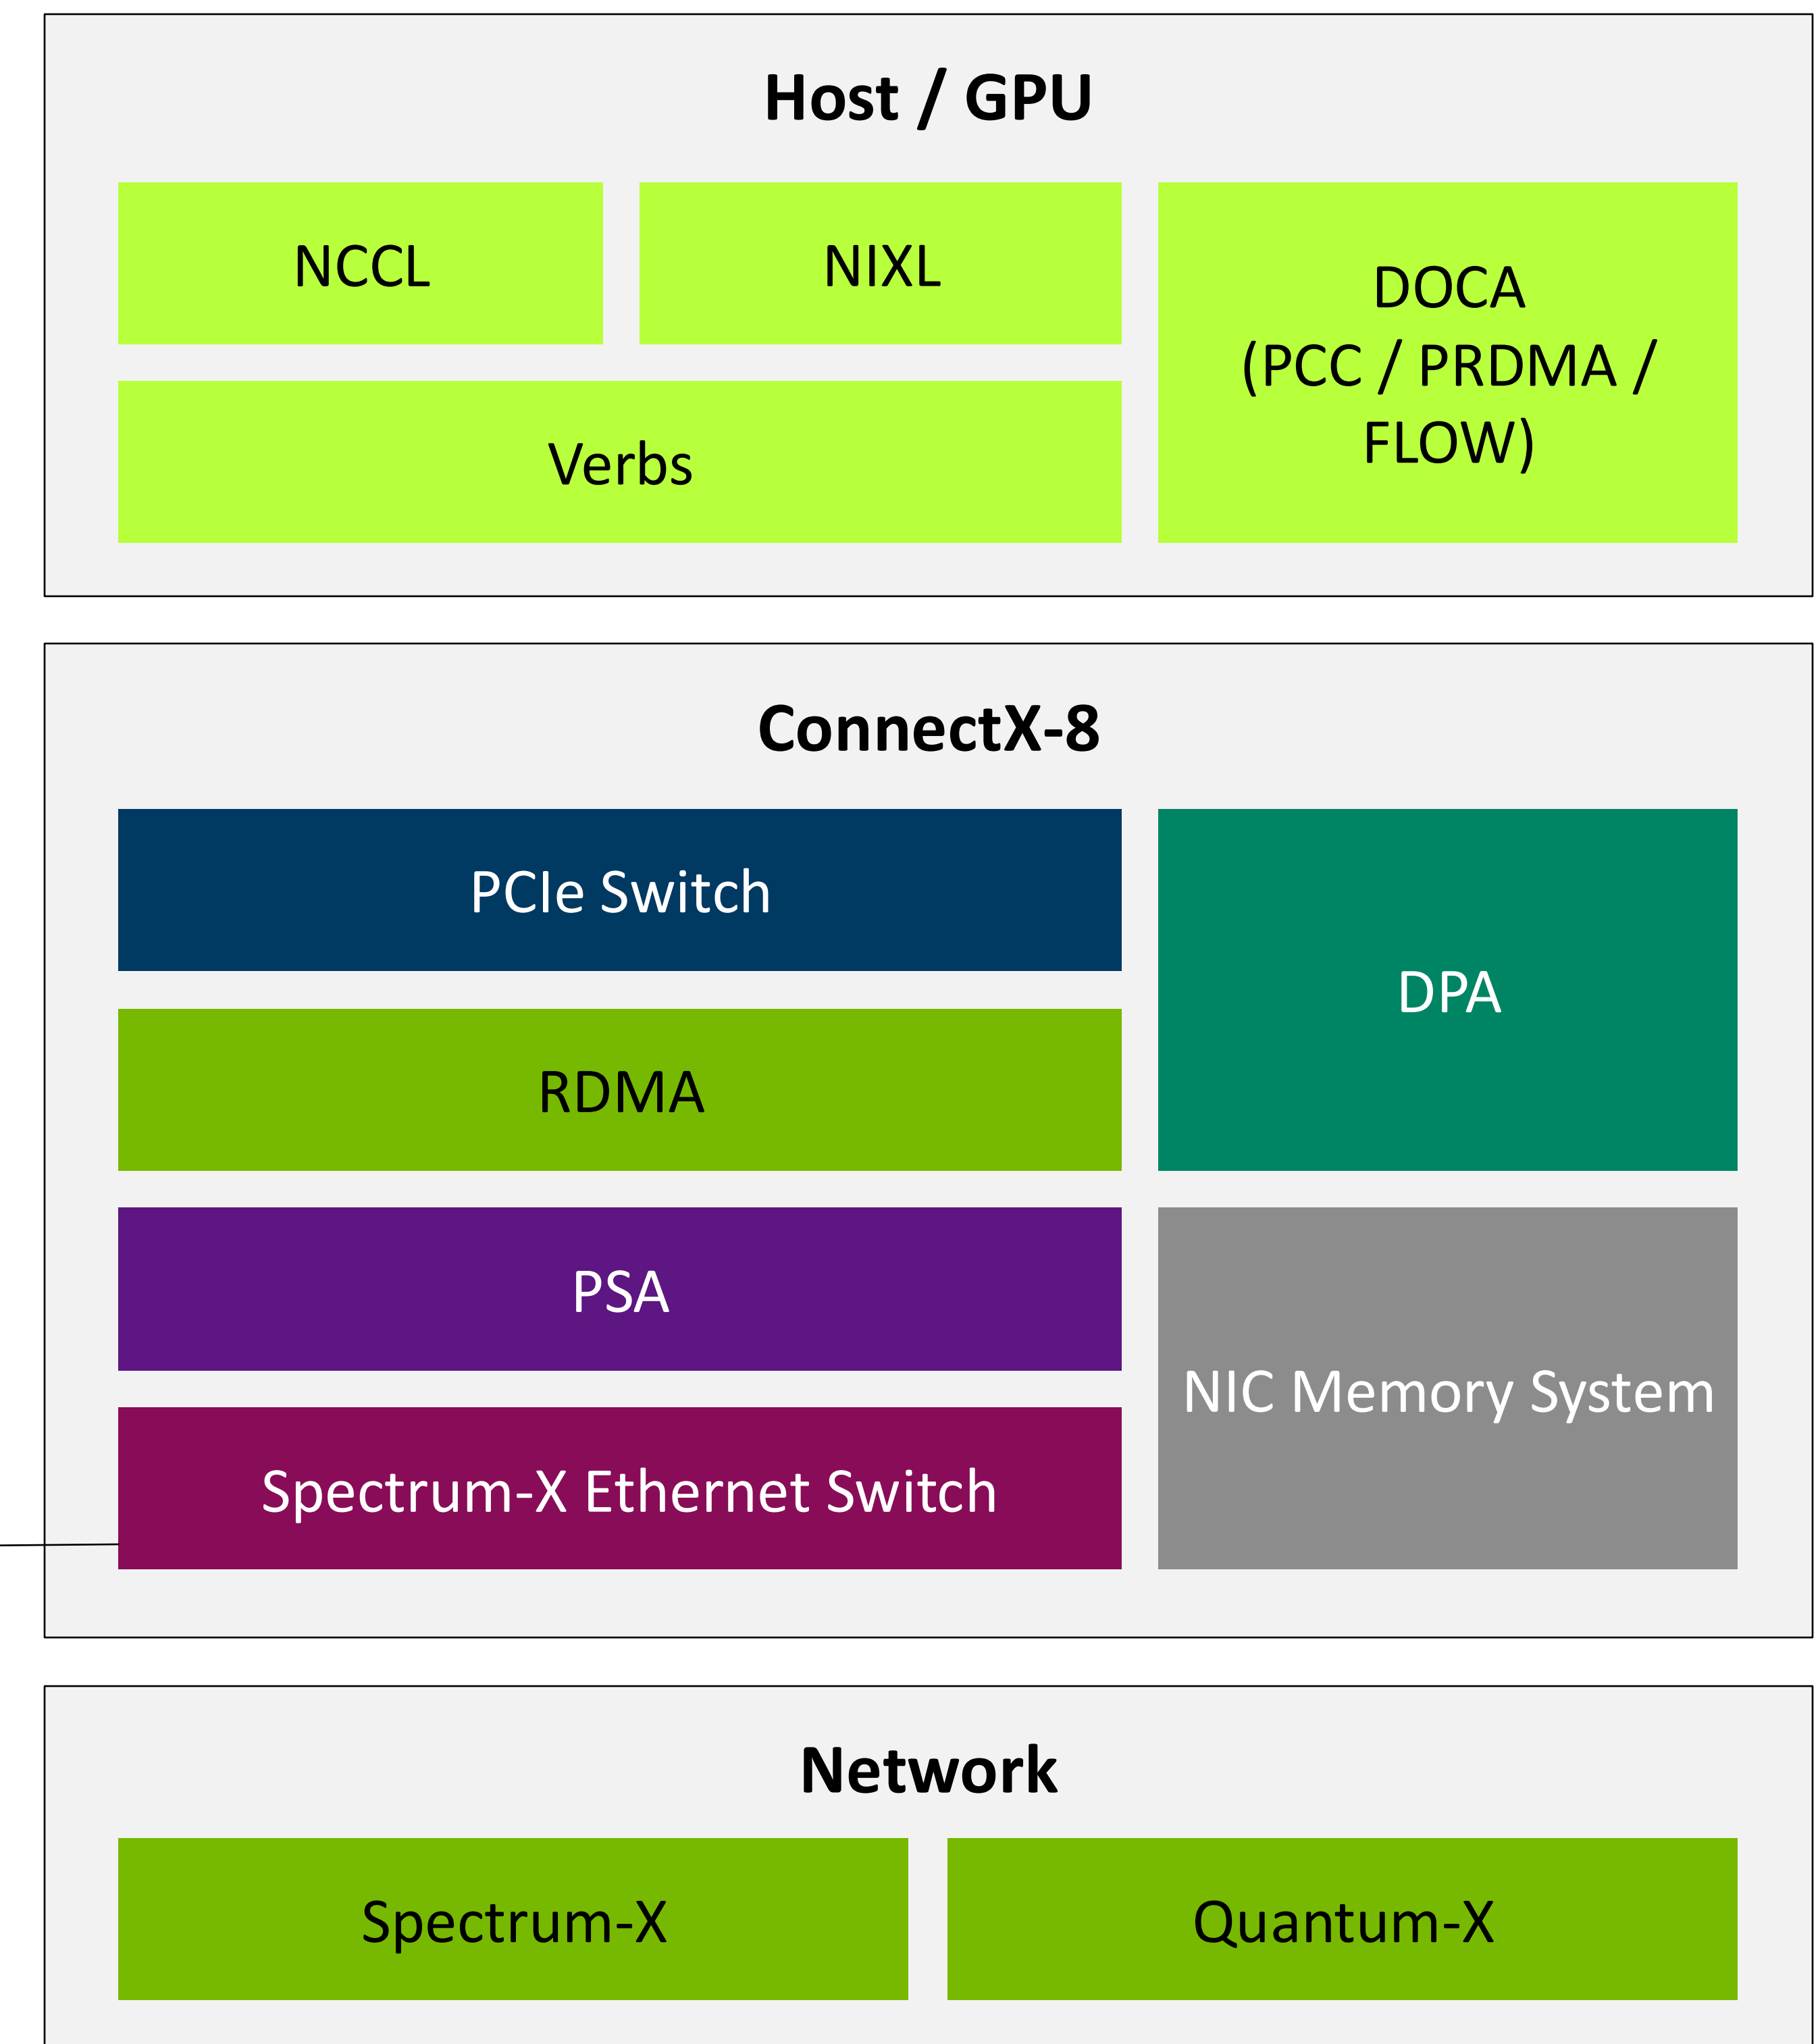
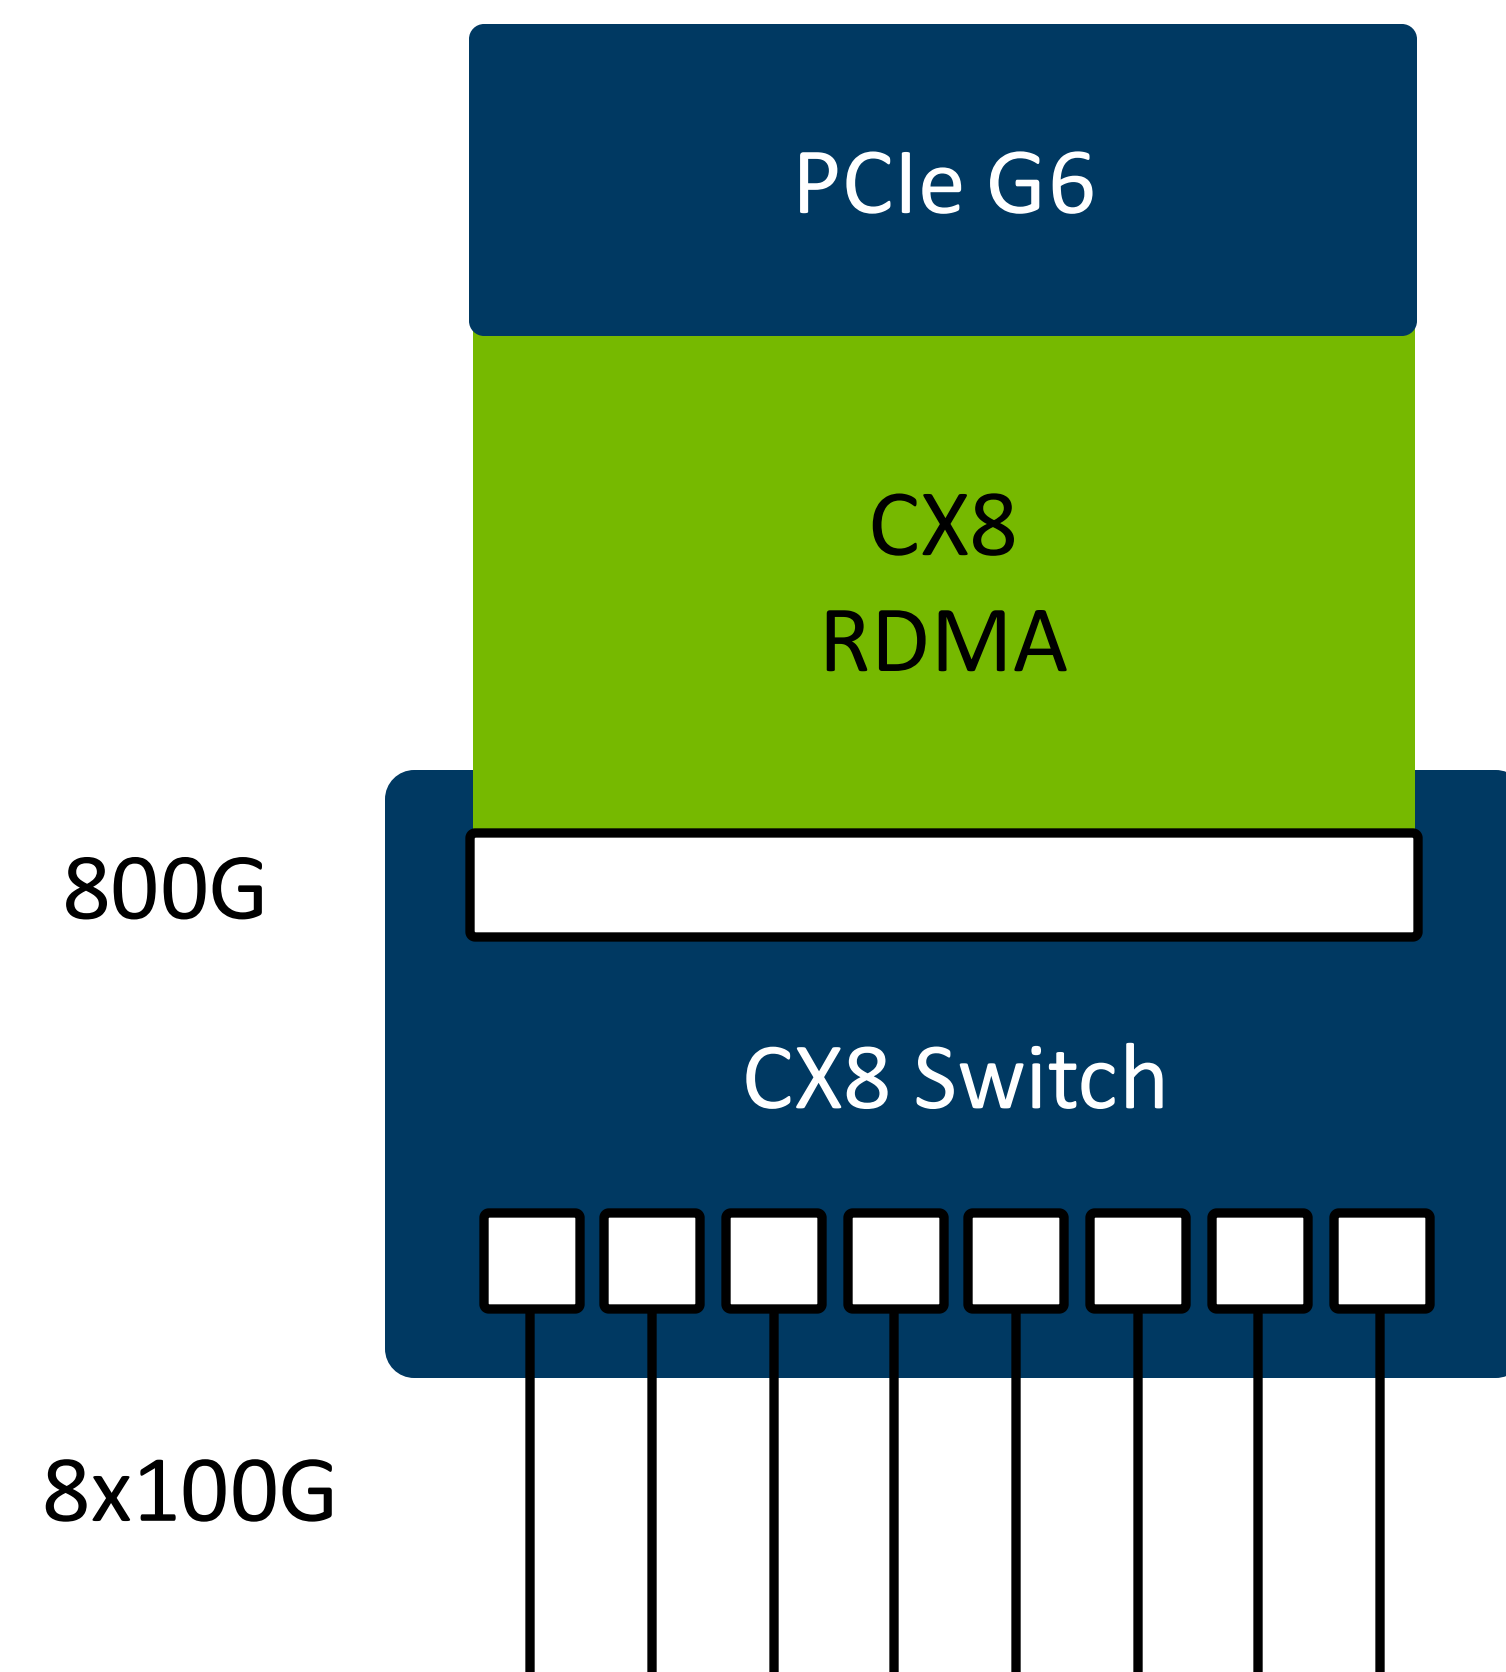


Standard NIC connectivity (1 ports 4 lanes illustration)

Single GPU bandwidth: 800G, 8x100G  
Switch radix: 512  
GPU scale in two-level non-blocking fat-tree:  $64 \times 32 = 2K$

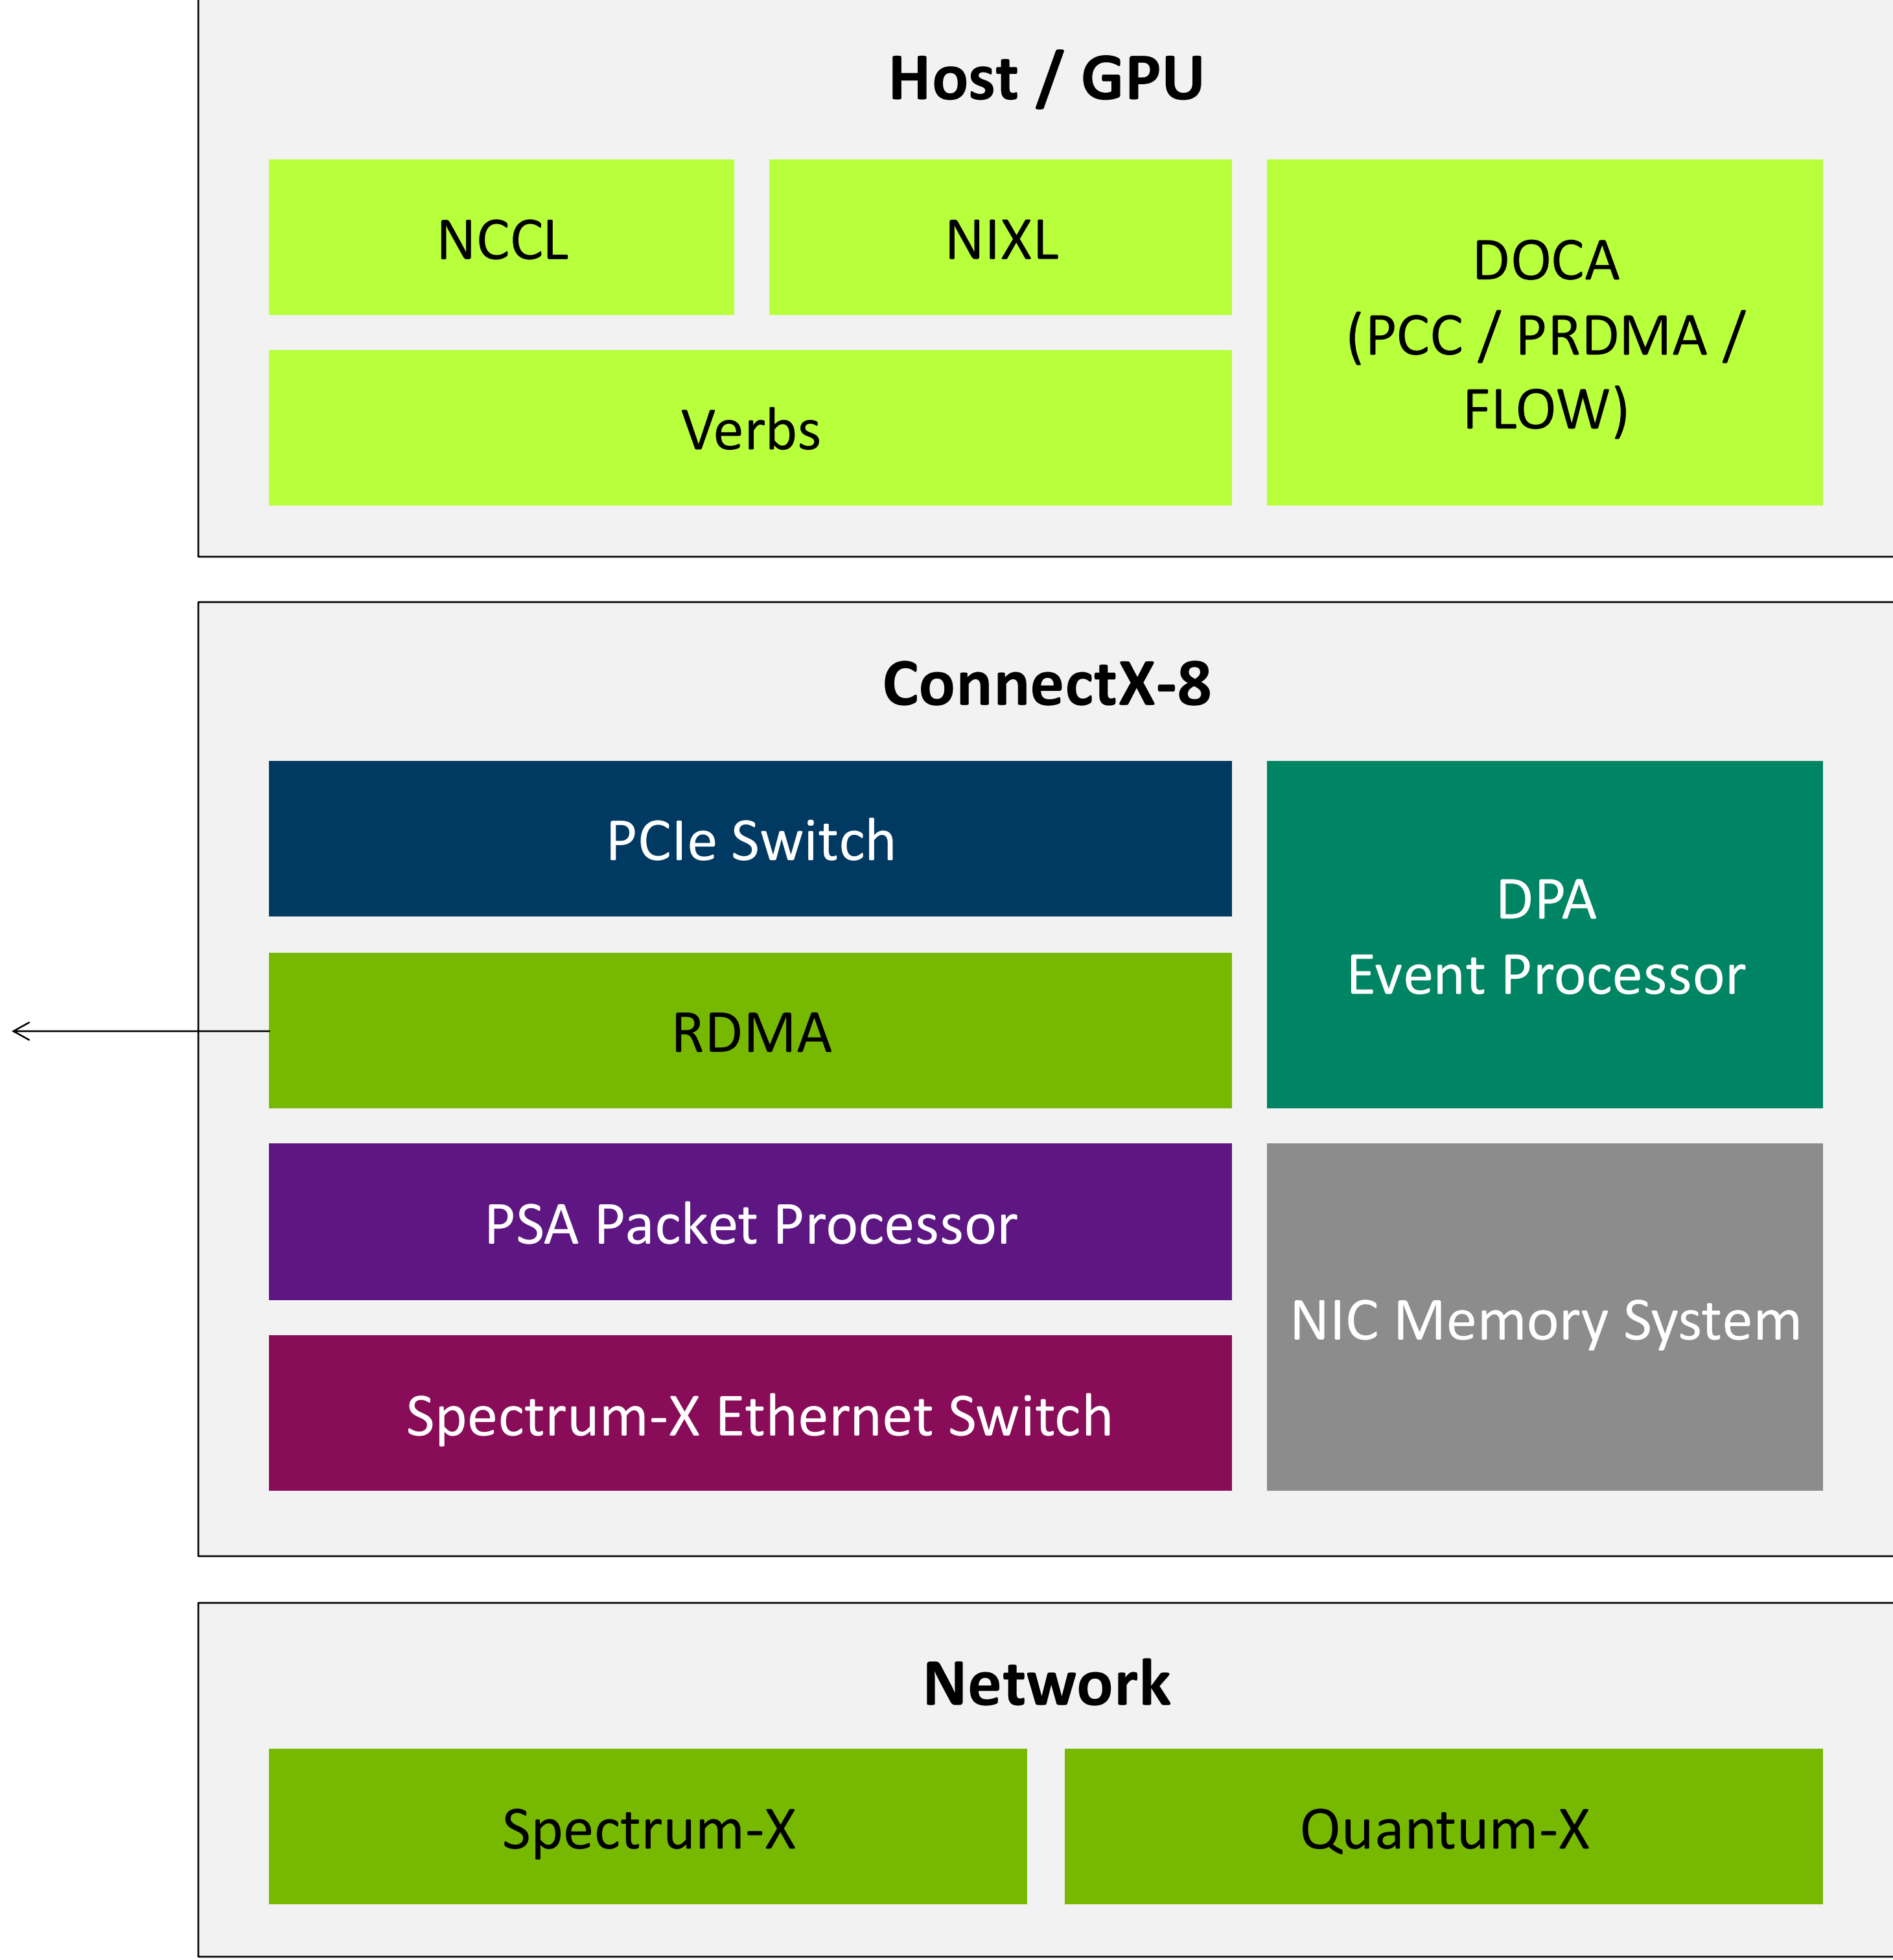
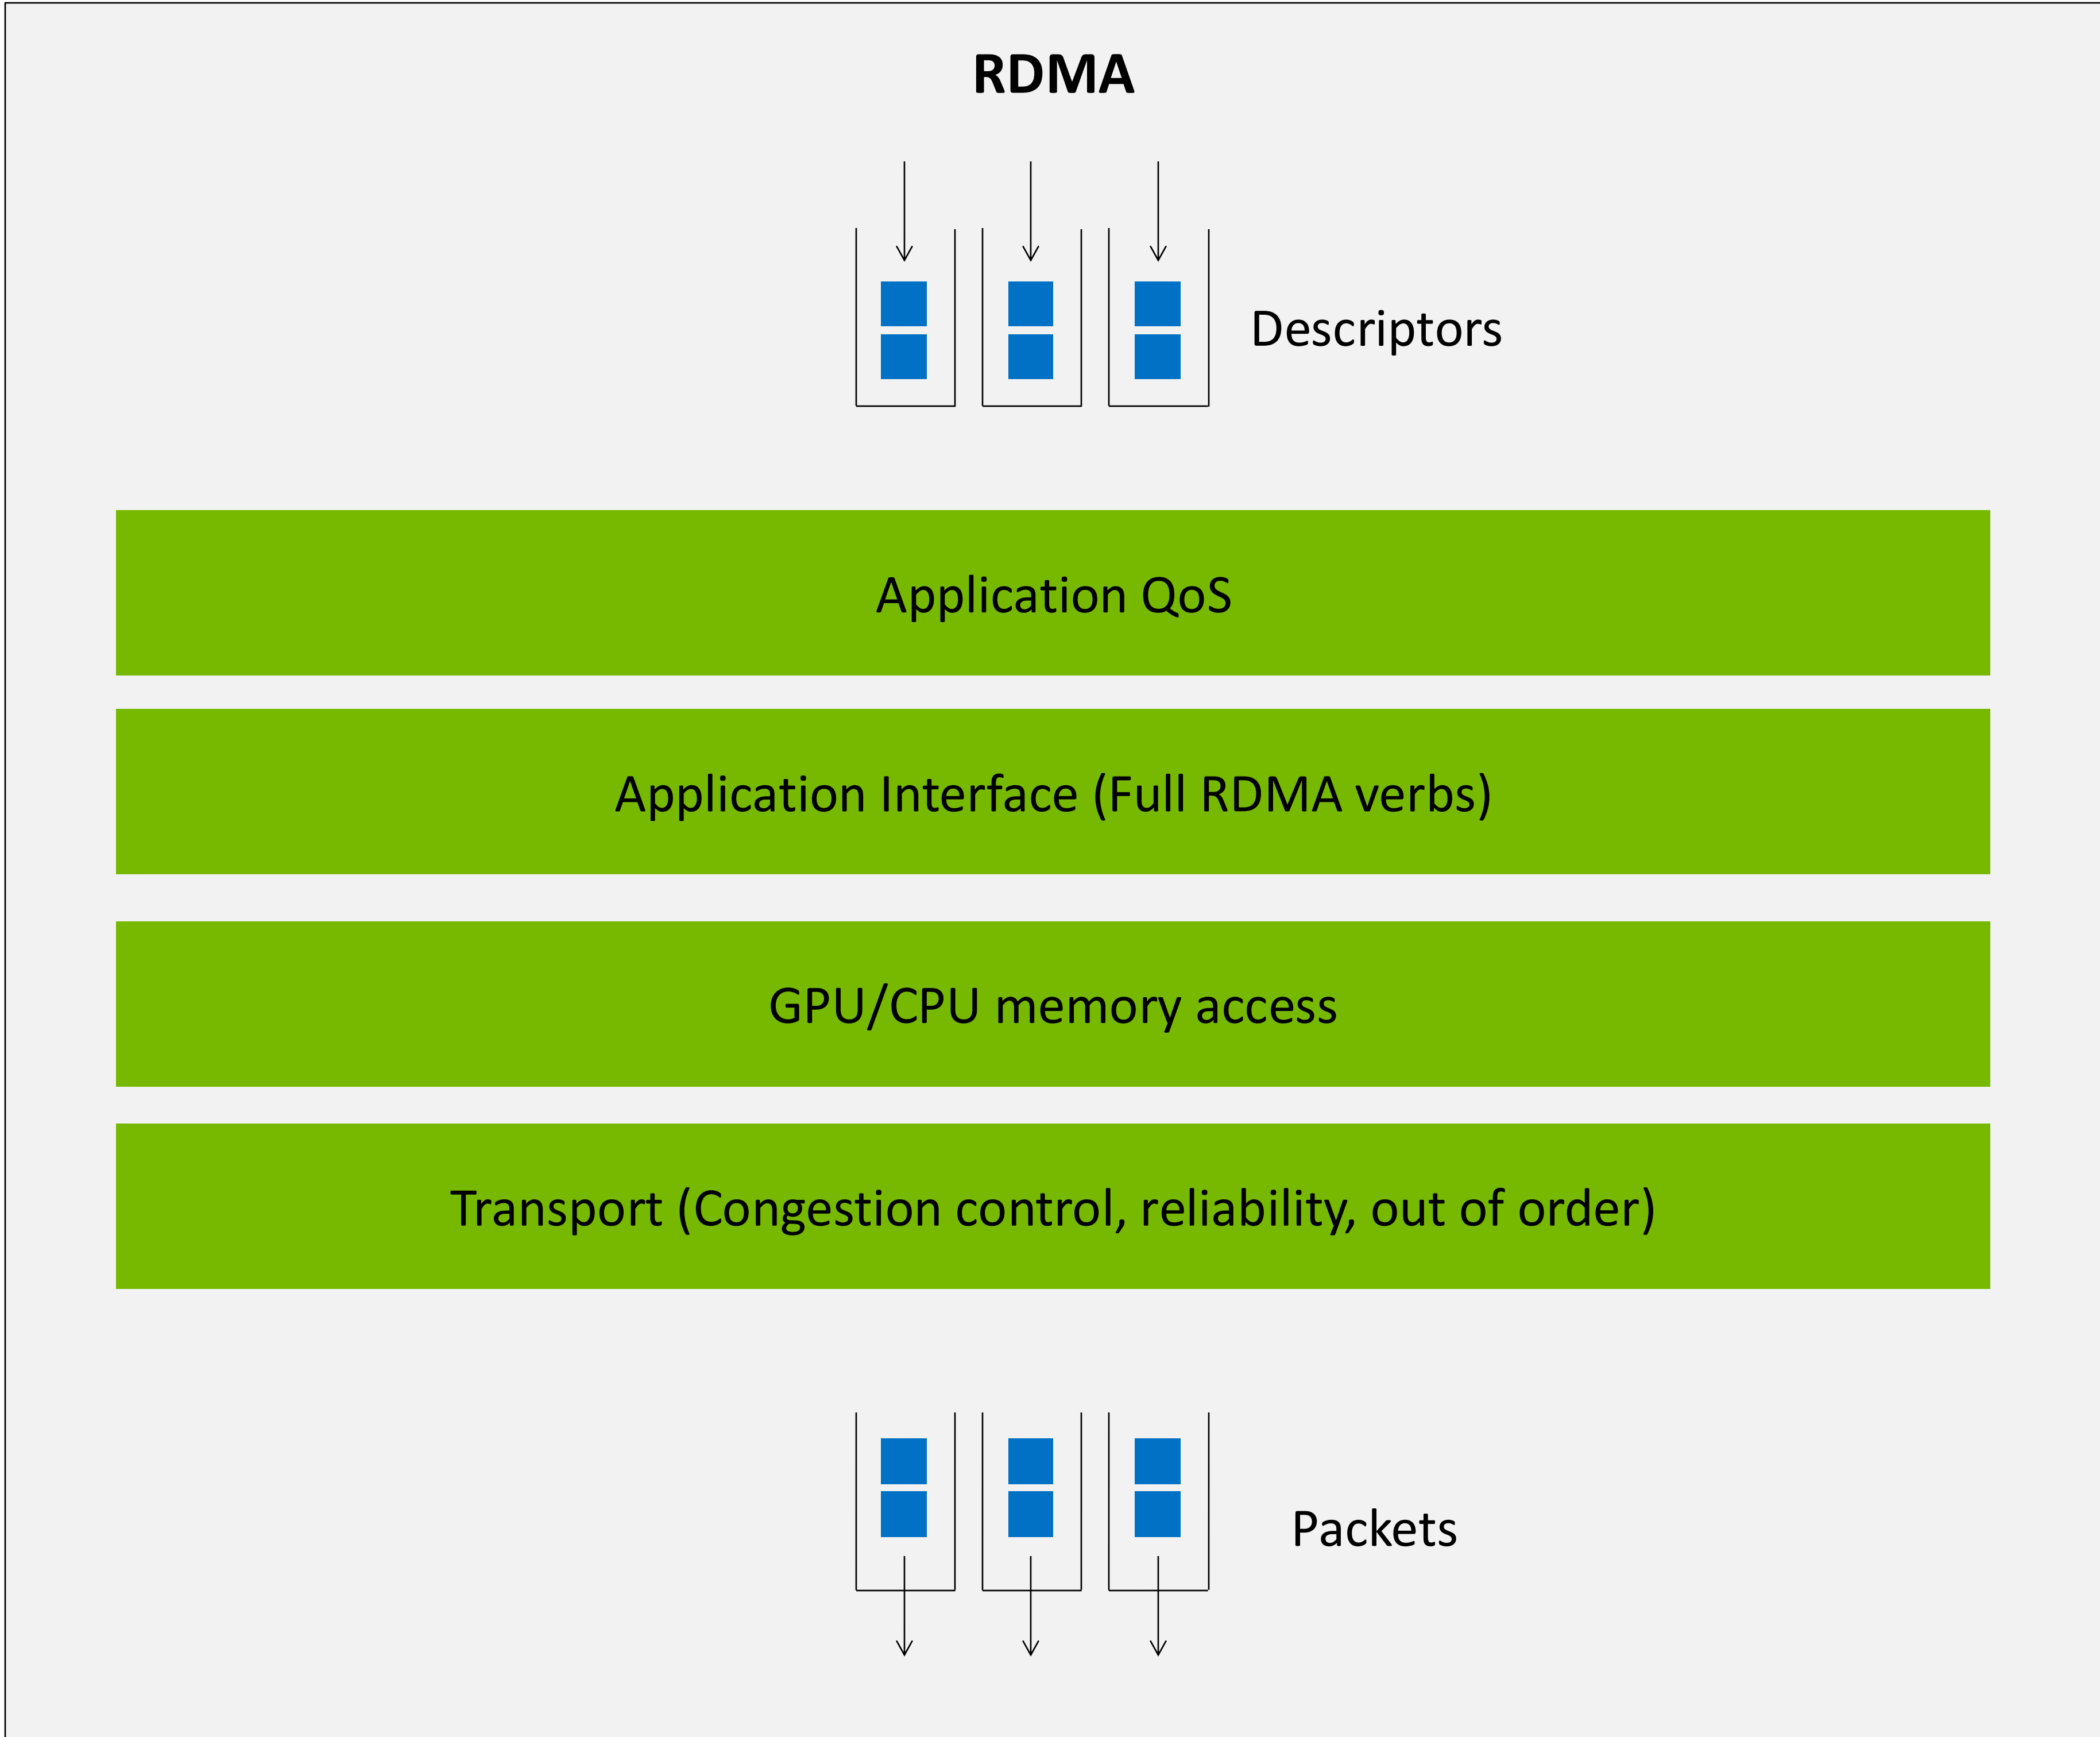
# ConnectX-8 Integrated Spectrum-X Ethernet Switch for AI Scalability

- An extension of the Spectrum-X Ethernet switch infrastructure
- Integrated load balancing and congestion control
- Switch functions are abstracted from the application



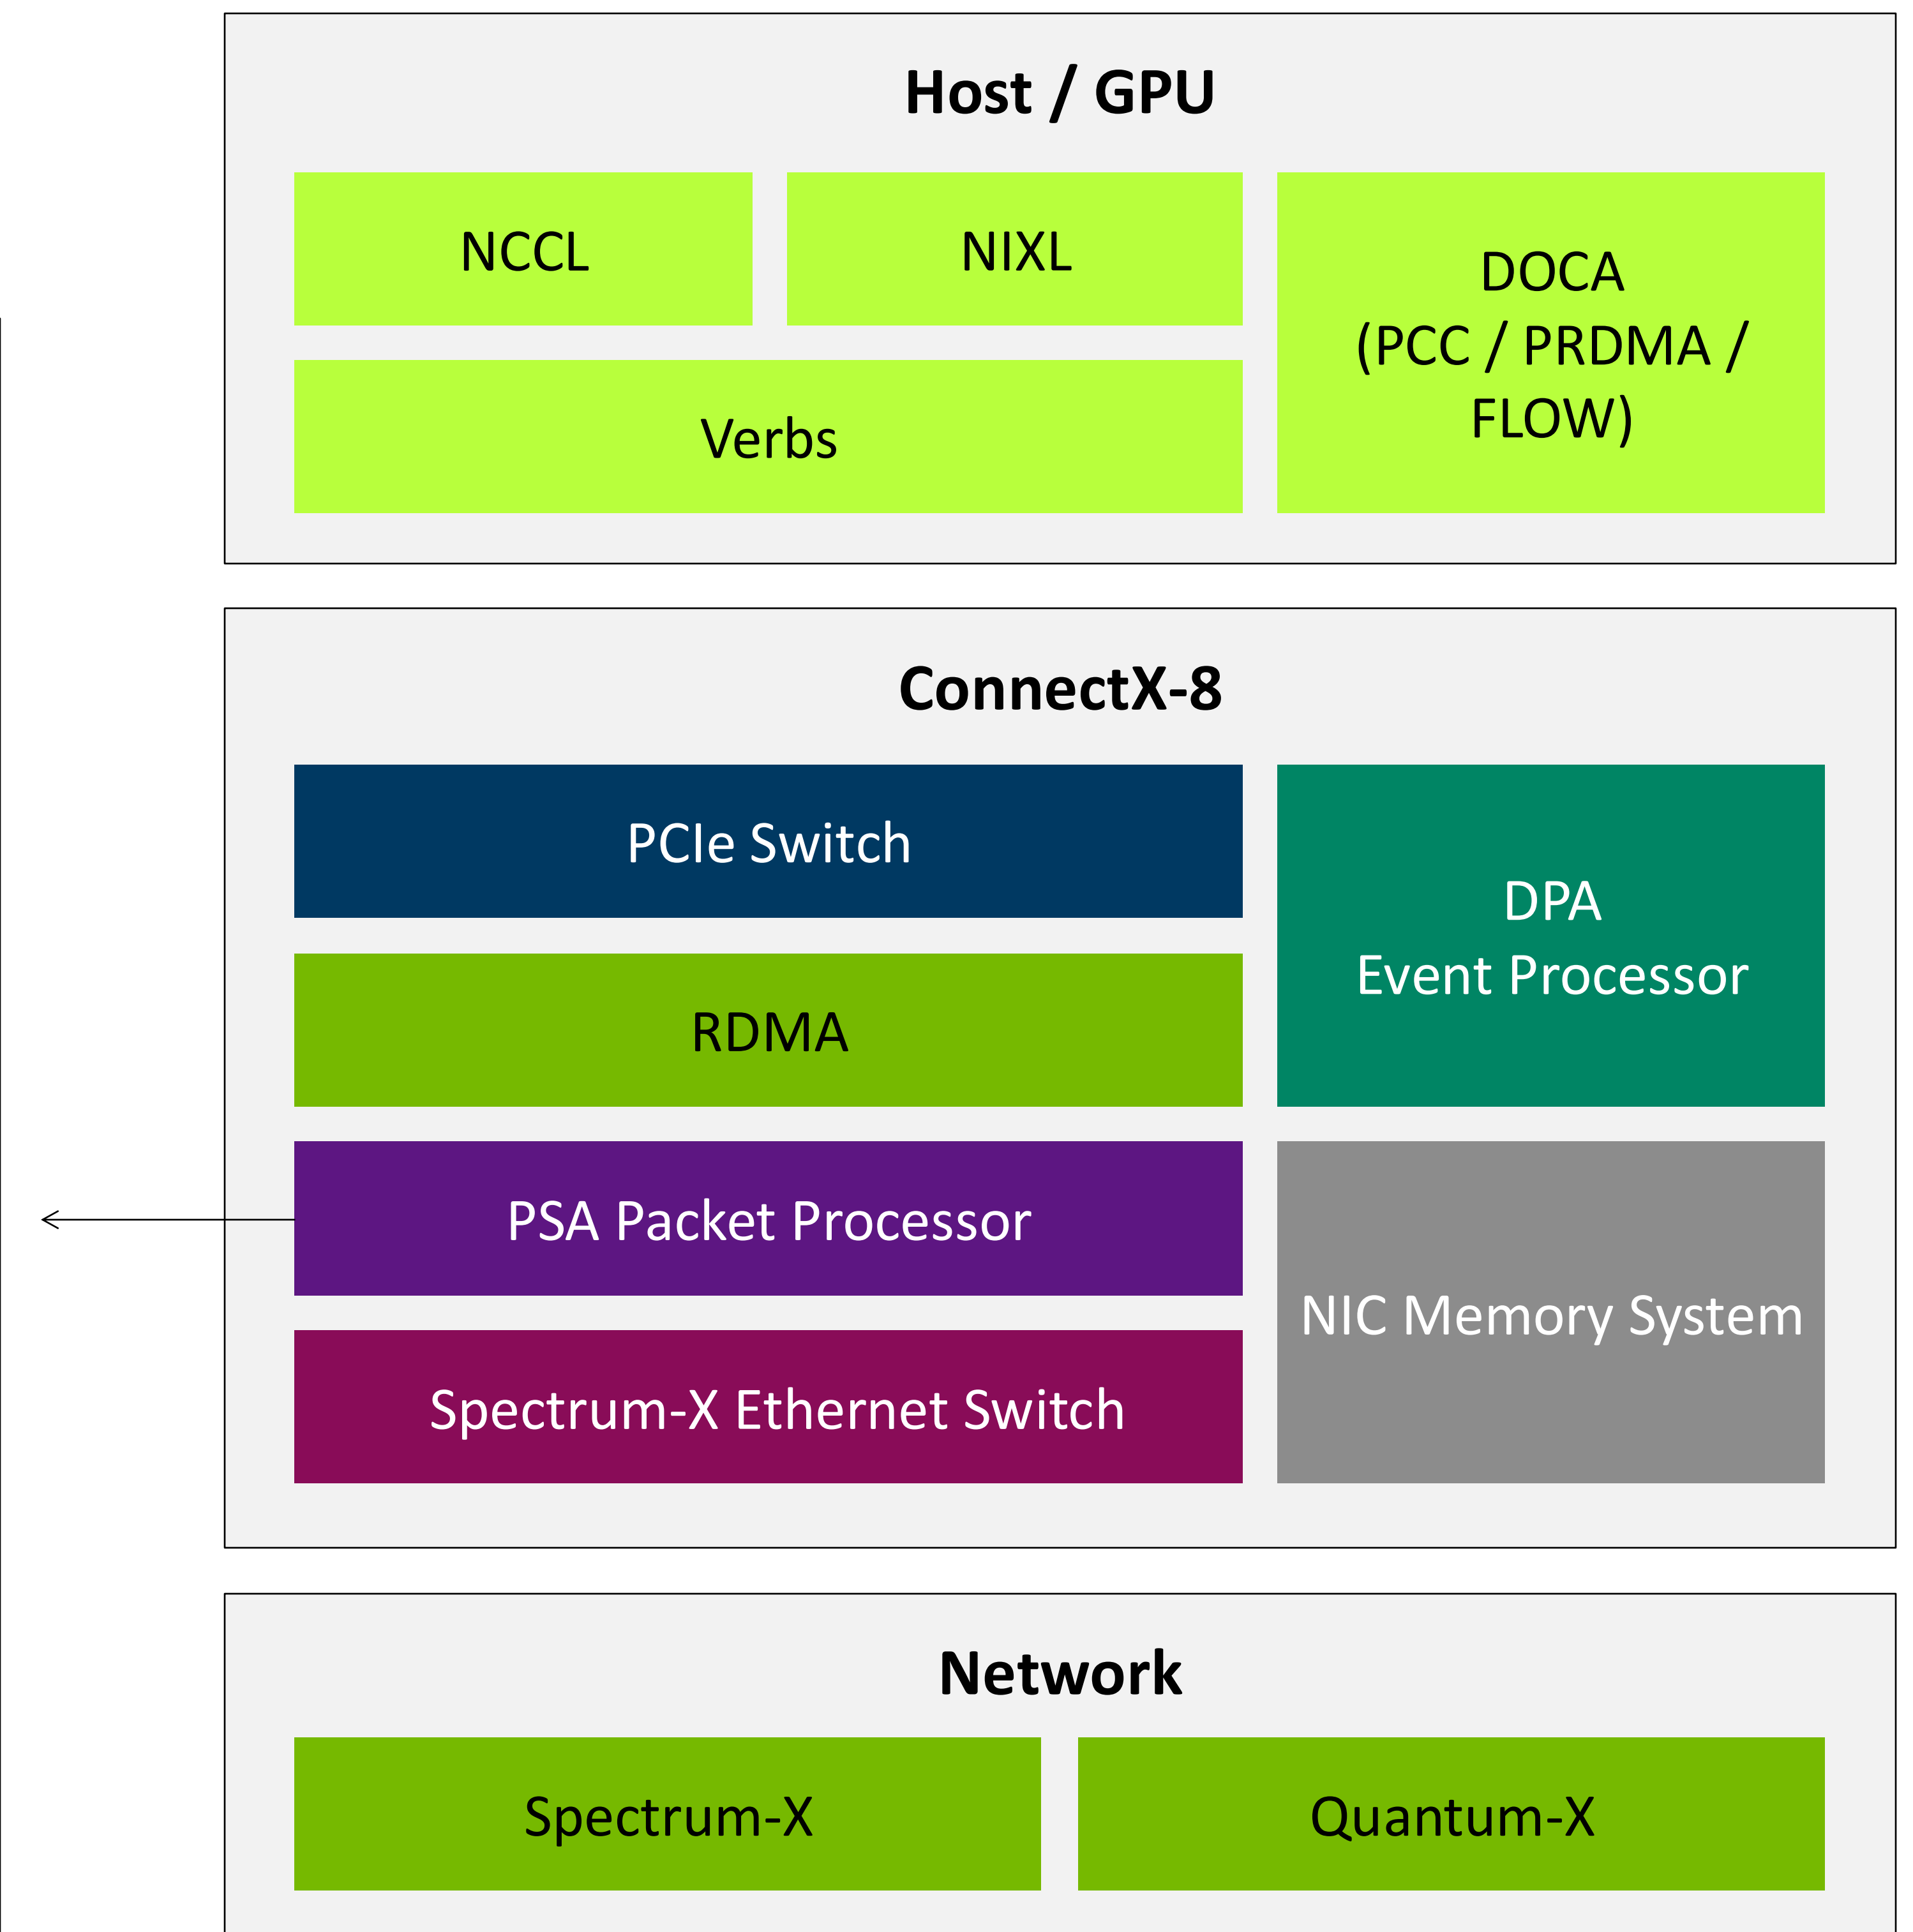
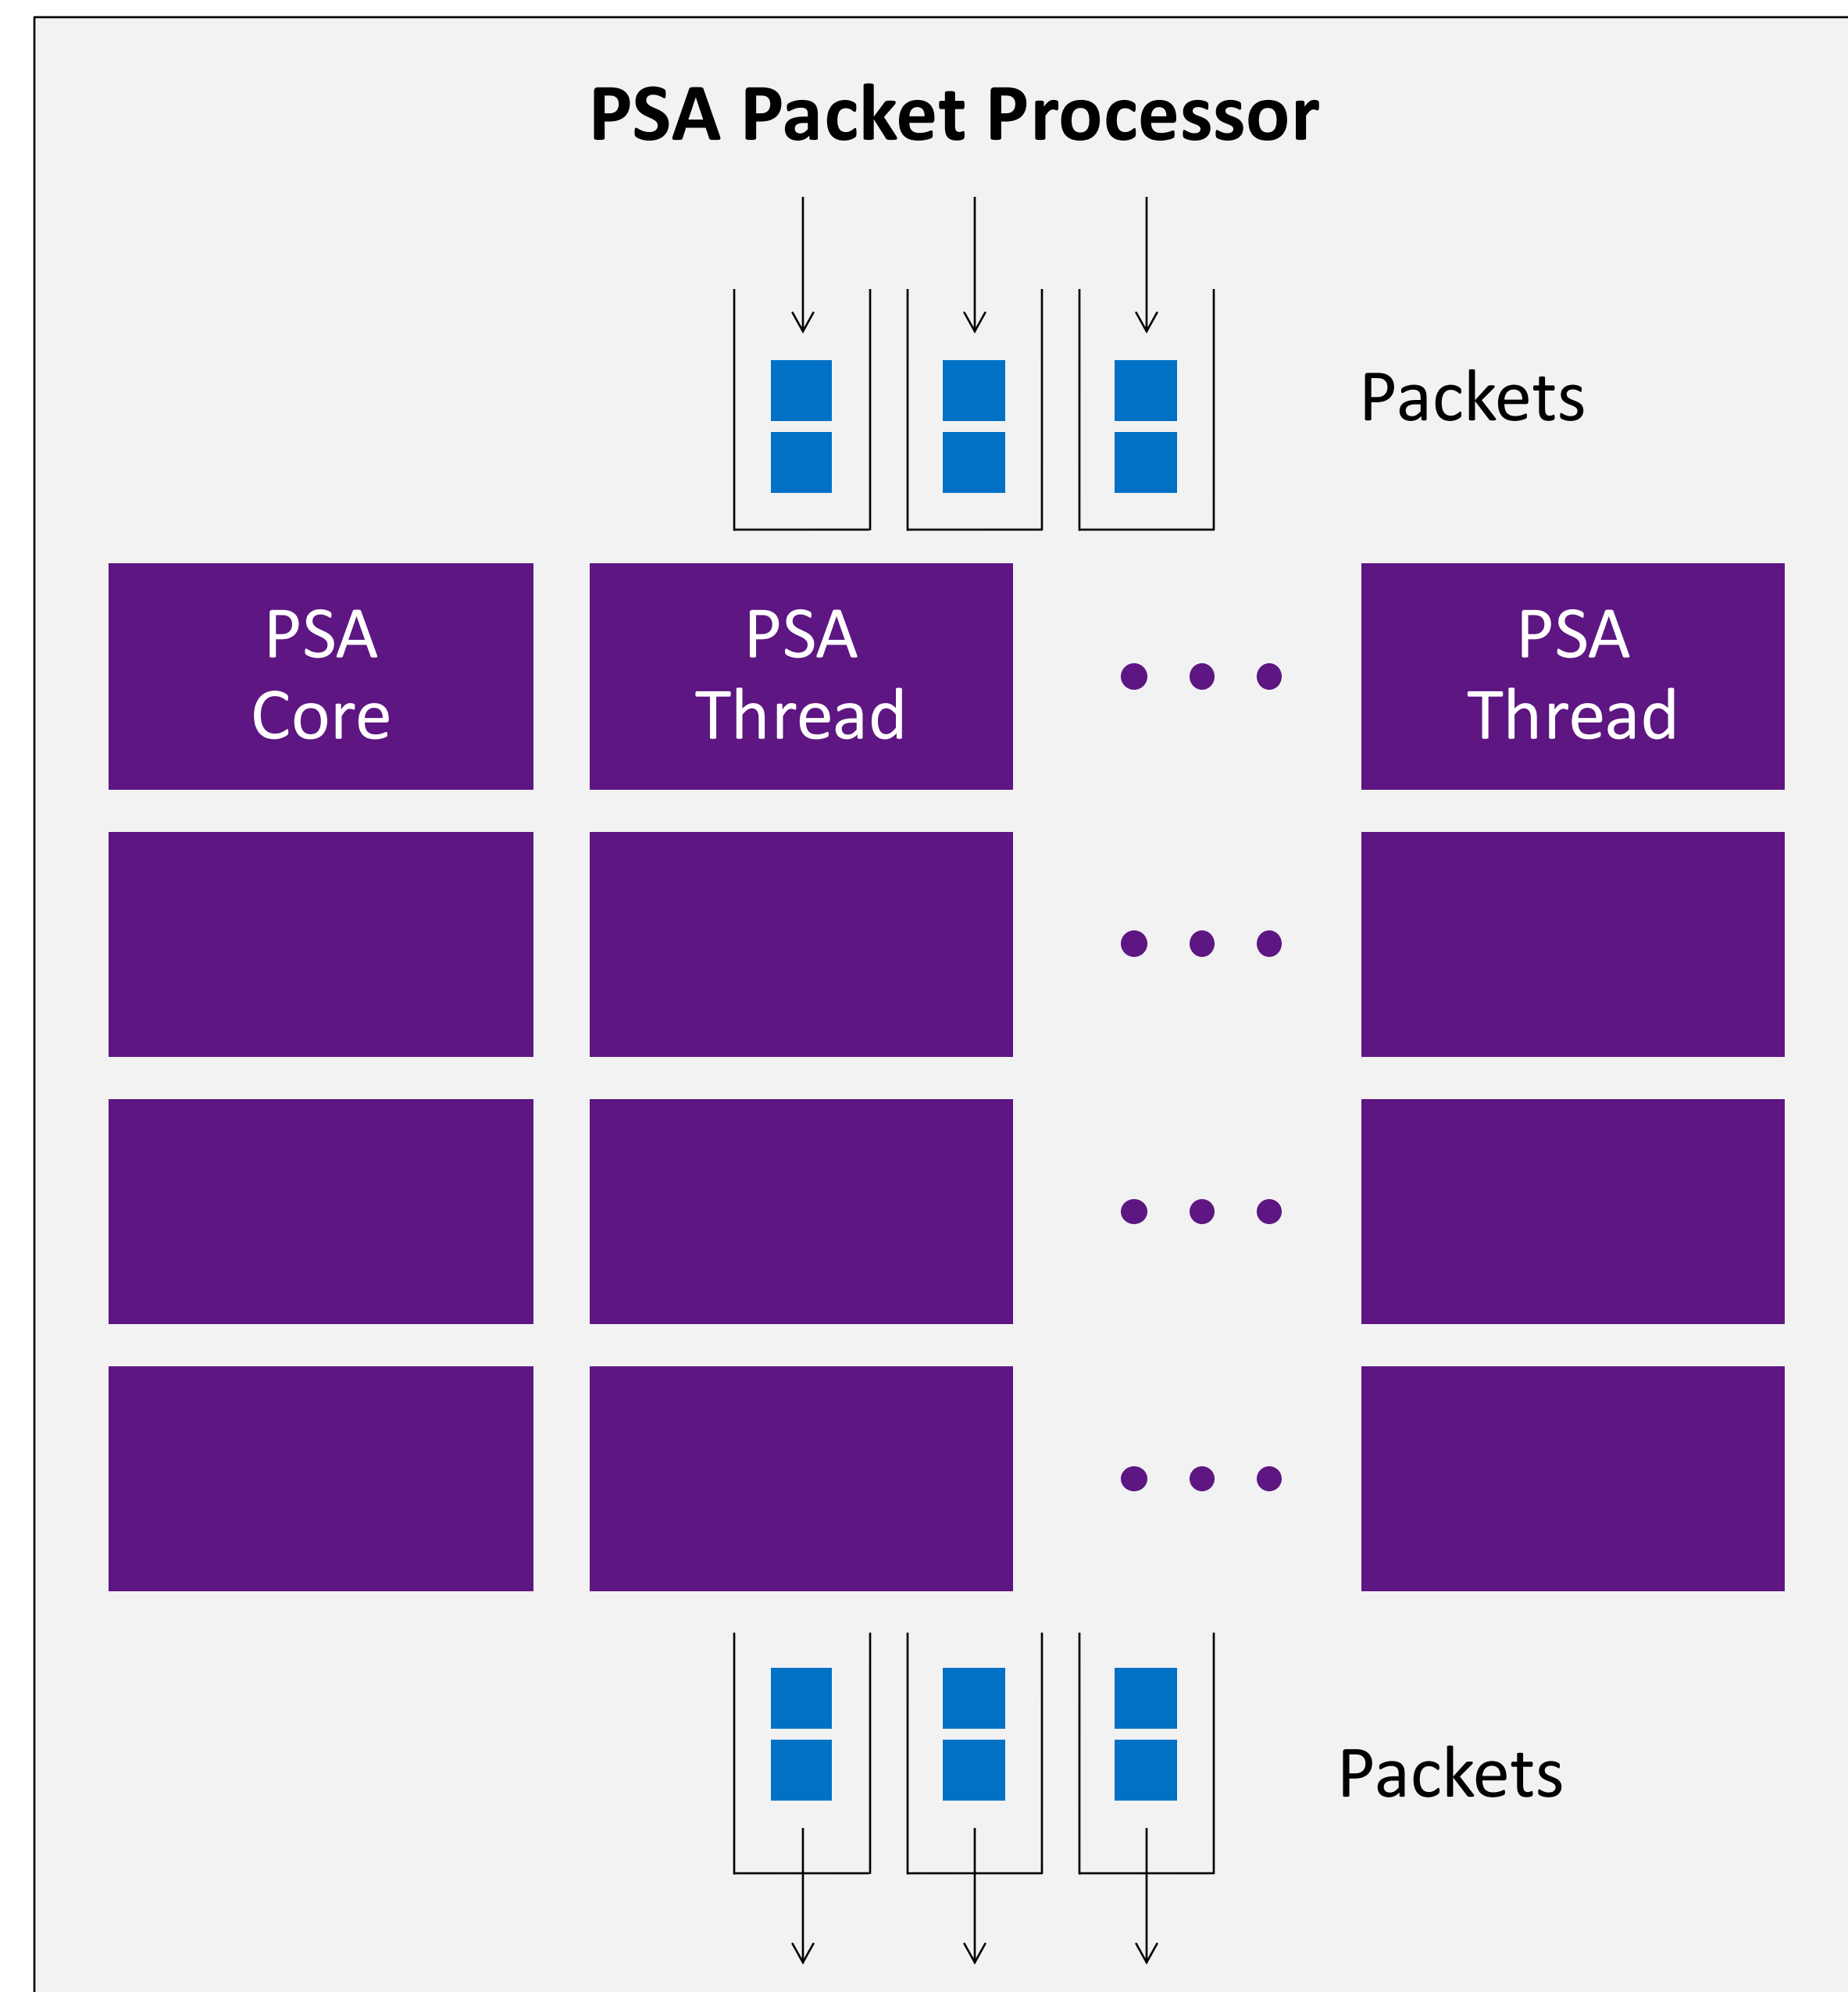
# ConnectX-8 RDMA – Multi-Layer Innovation for AI Networking

Scaling traditional operating system services for AI networking



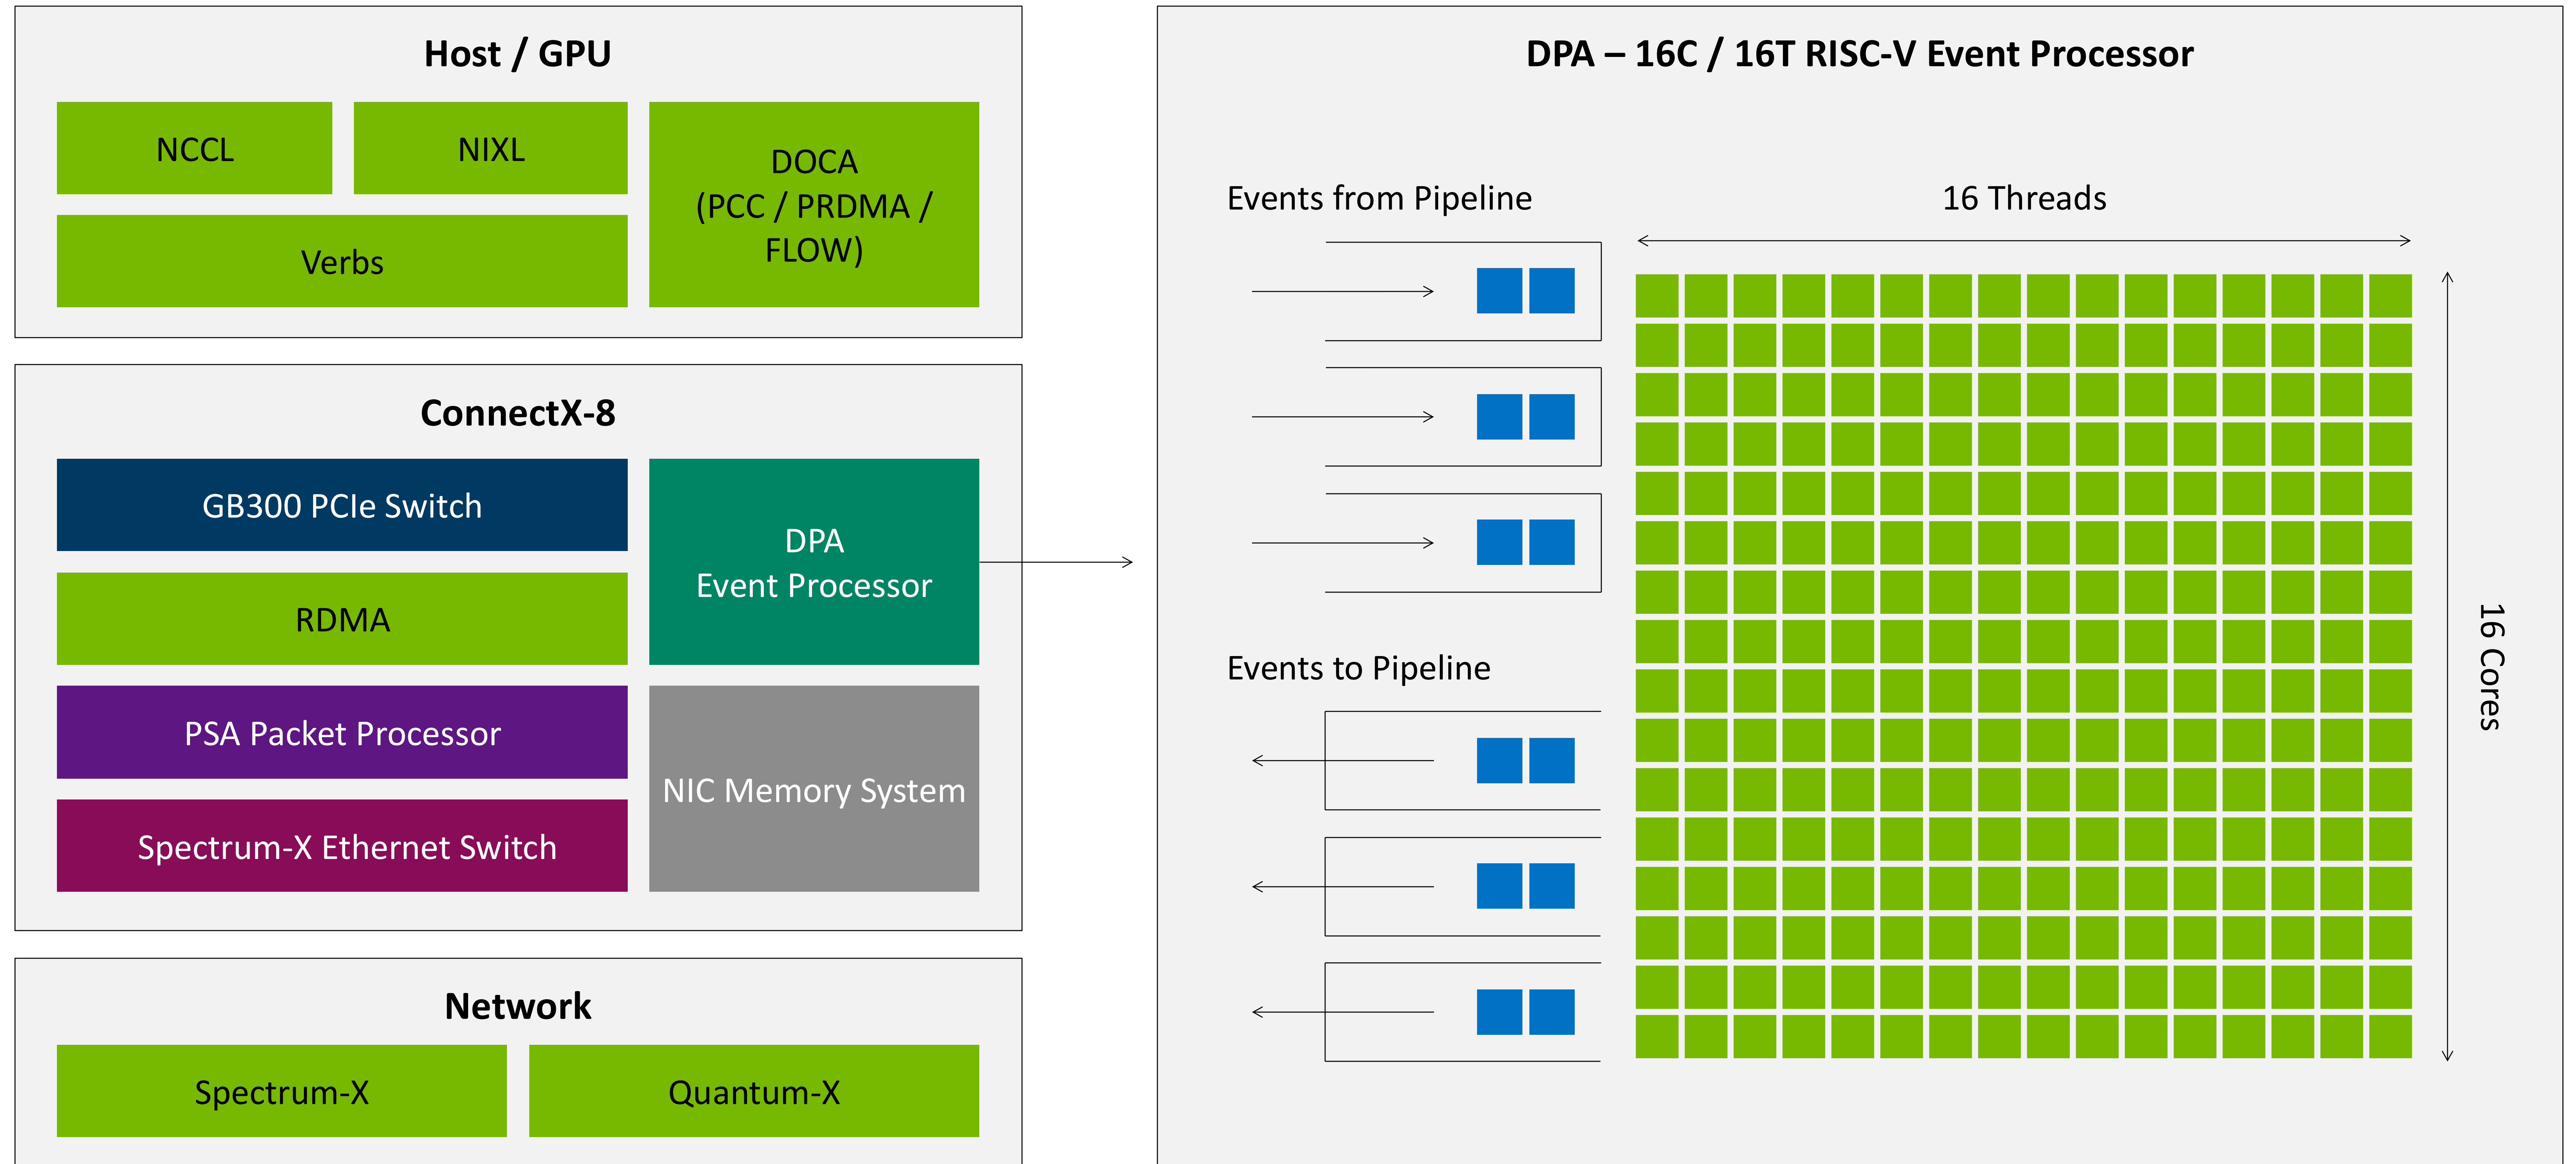
# ConnectX-8 Packet Processor – Securing and Routing AI Factories

- Spectrum-X Routing
- Spectrum-X Telemetry
- Network crypto
- Network multi-tenancy




# ConnectX-8 Data Path Accelerator

Designed for running the algorithms that keep the network efficient



# Spectrum-X Ethernet Congestion Control & Routing

Mountains of hardware and software co-design between the SuperNIC and the Switch



Failure detection & reaction  
A combination of both



Congestion control  
Finding the right rate to transmit to a given destination



Routing  
Finding the right path on a given network state

## Spectrum-X Congestion Control

## Spectrum-X Adaptive Routing

### AI workload

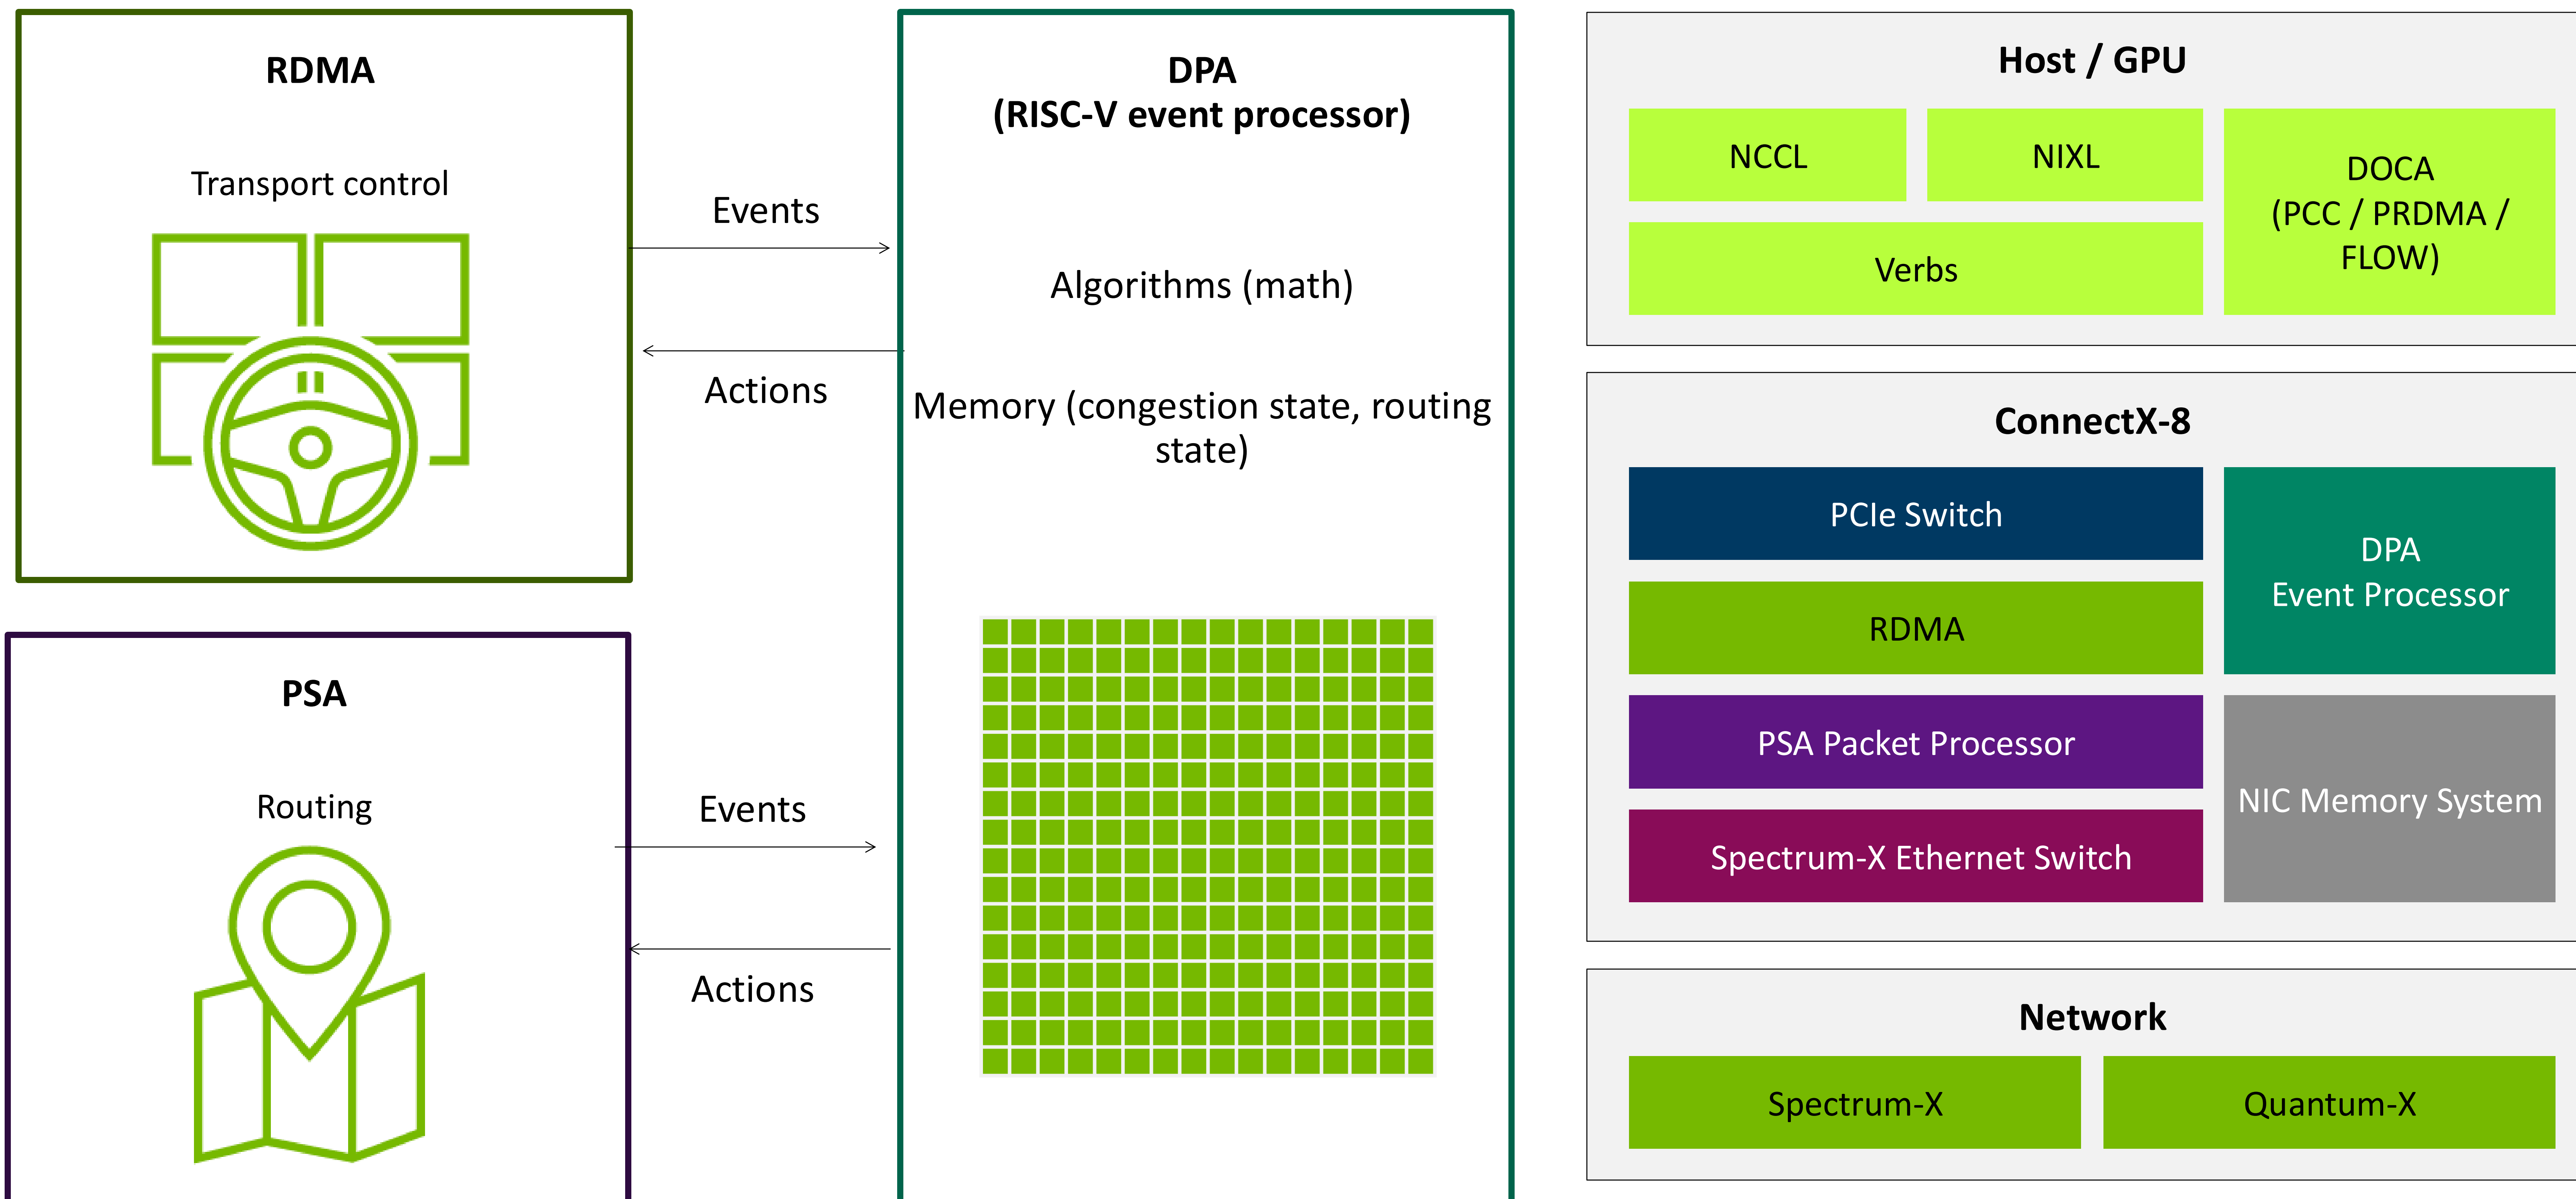
Single job over many GPUs—  
highly sensitive to tail latency

High bandwidth large scale collective requires  
high accuracy and fast convergence

Network round trip time - 5-10us  
Packet transmission time - 2ns  
Requires fast & frequent updates


# ConnectX-8 Congestion Control & Routing

Interaction Between Hardware Pipeline and Programmable Processors



# Demo


# Spectrum-X Ethernet Delivers 60% Shorter Training Step Time In Isolation



OTS RDMA NIC and Switch

LLM + Job1:  
Training Step Time: 17.2

LLM + Job1 + Job2:  
Training Step Time: 23.1

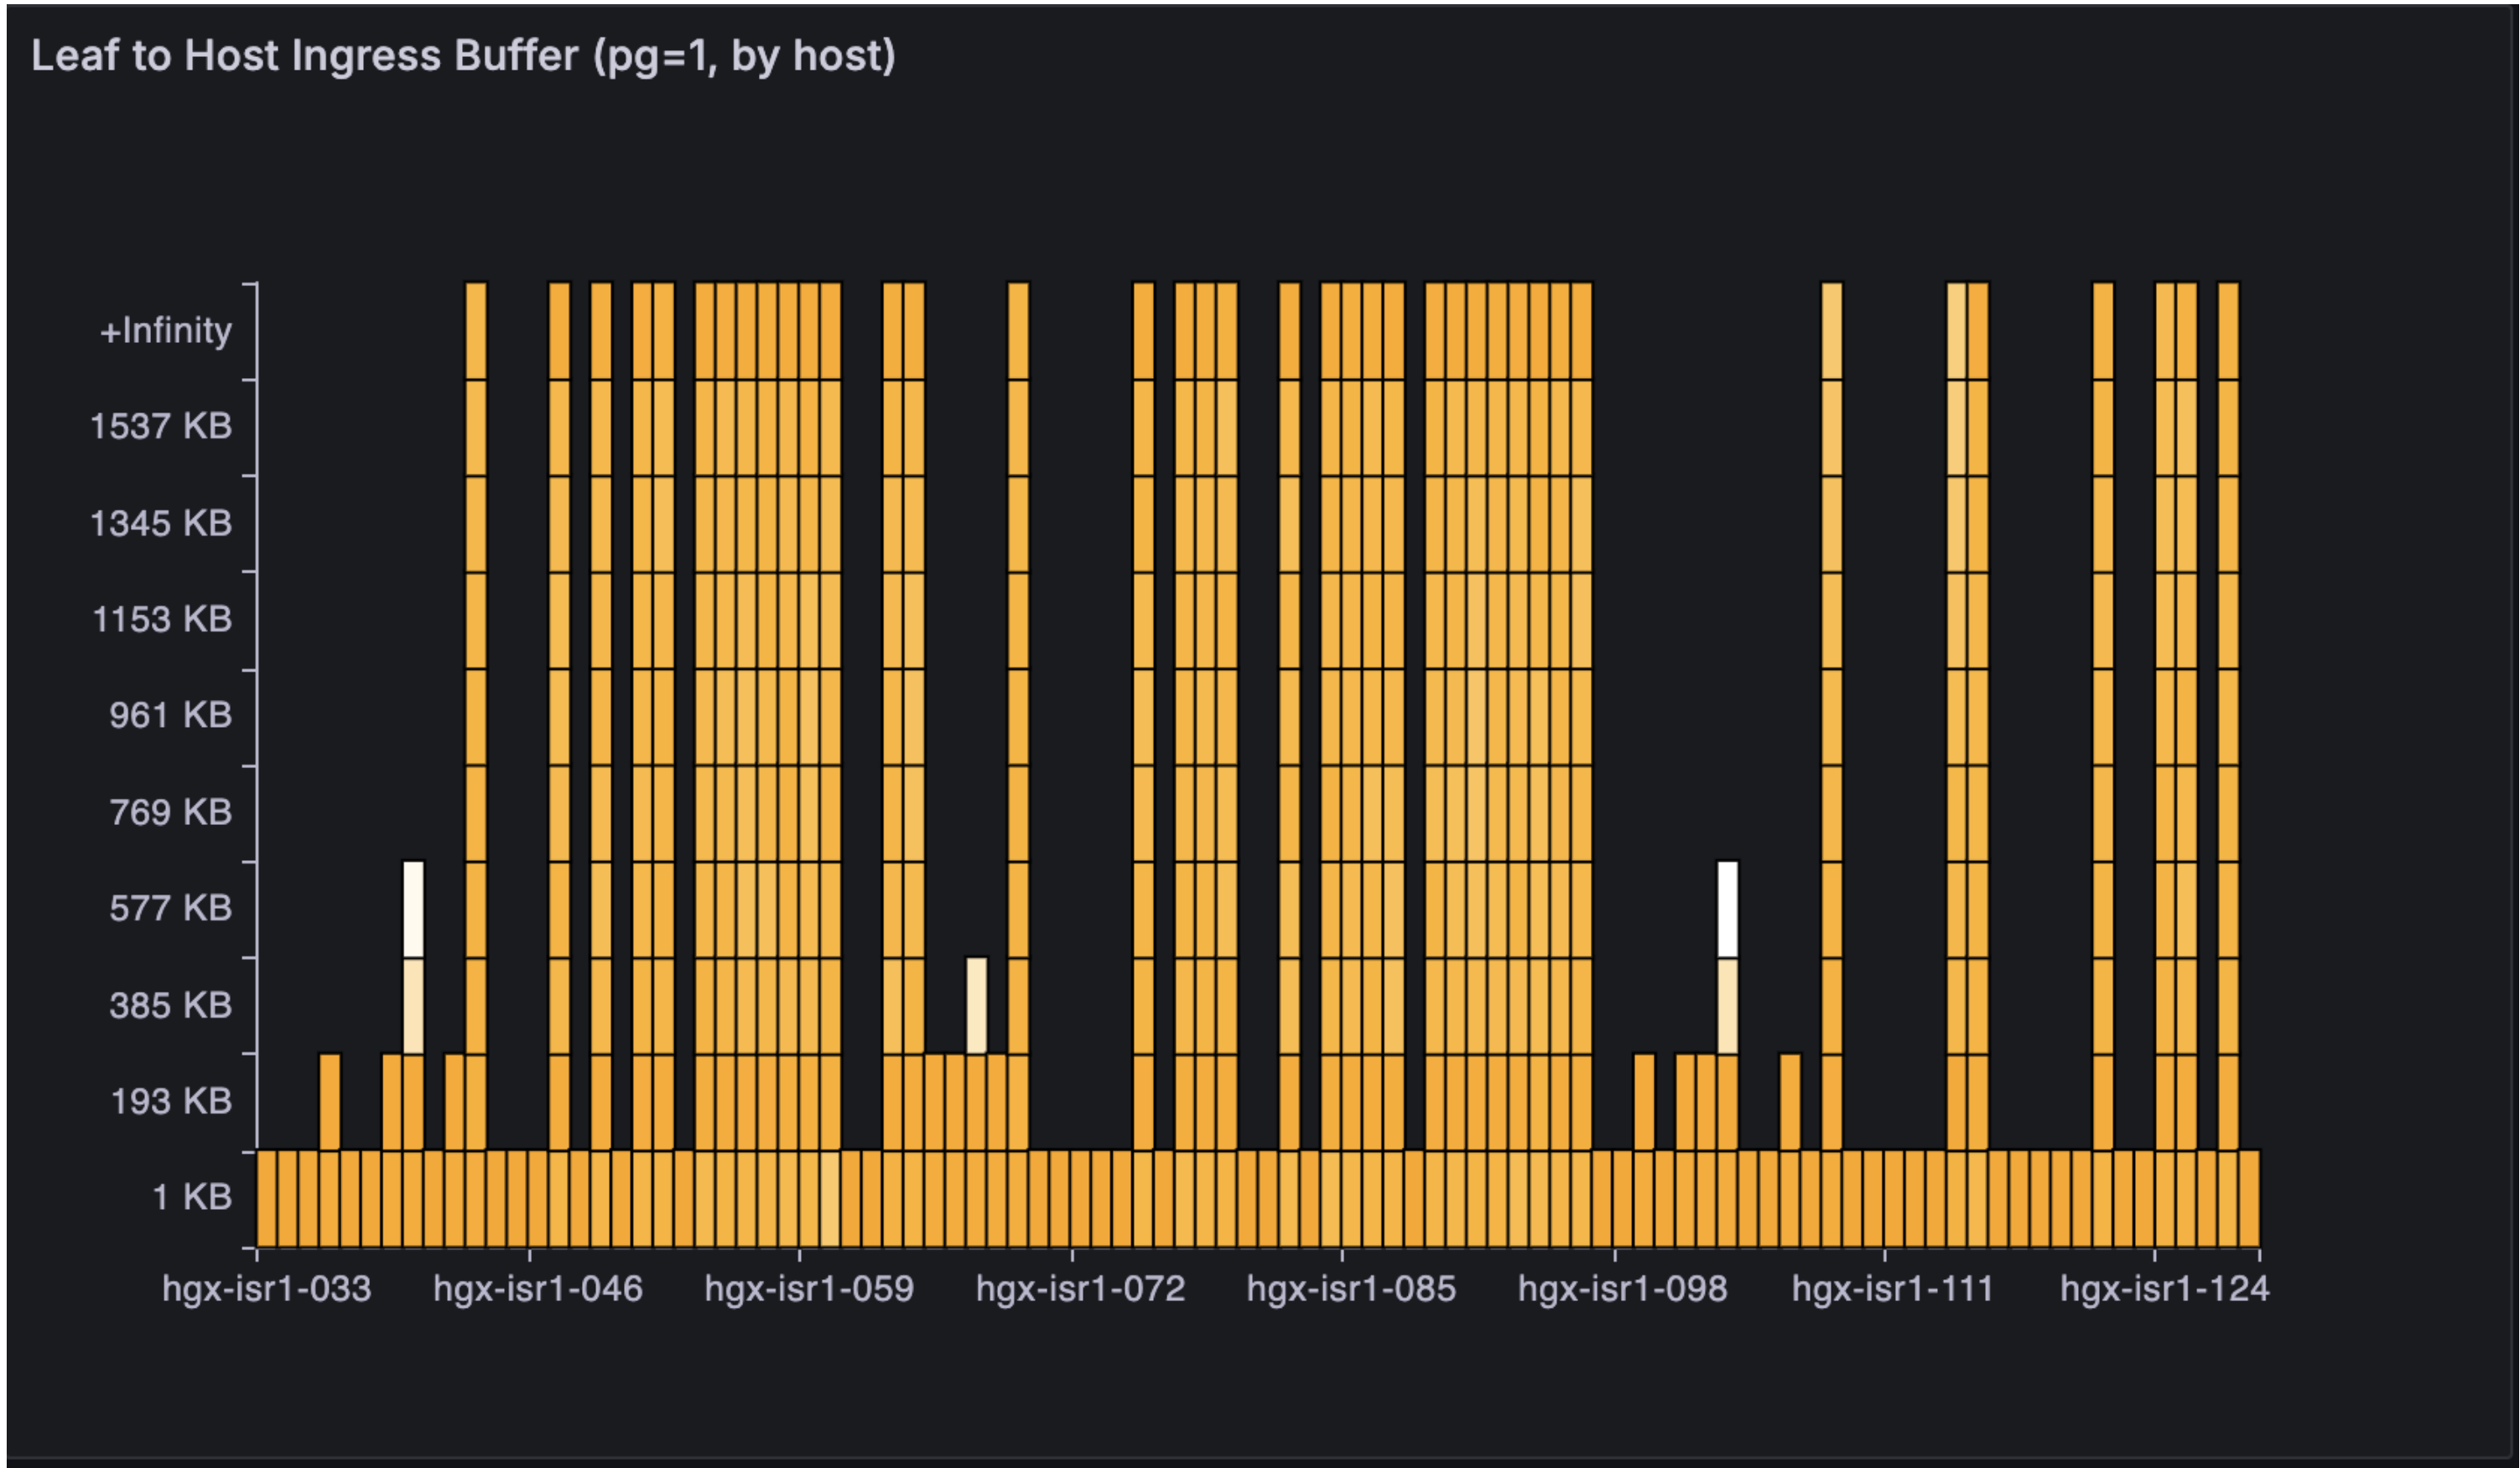


NVIDIA Spectrum-X Ethernet


LLM + Job1:  
Training Step Time: 13.9

LLM + Job1 + Job2:  
Training Step Time: 13.9

# Spectrum-X Ethernet Delivers Zero Tail Latency Under Load



OTS RDMA NIC and Switch



NVIDIA Spectrum-X Ethernet

# Spectrum-X Performance

Switch-to-SuperNIC, end-to-end network processing, bringing high performance to Ethernet

