AI Data Centers

Michal Styszynski & Mahesh Subramaniam
Juniper Networks

ID 4950

February 13th, 2024

NANOG 90 – Charlotte, NC

AGENDA

- DC architectures for Existing & Modern workloads
- Lifecycle of an AI DC Network
- AI DC technologies
- Key takeaways

DC Arch. : Existing & Modern Workloads

DC Architectures

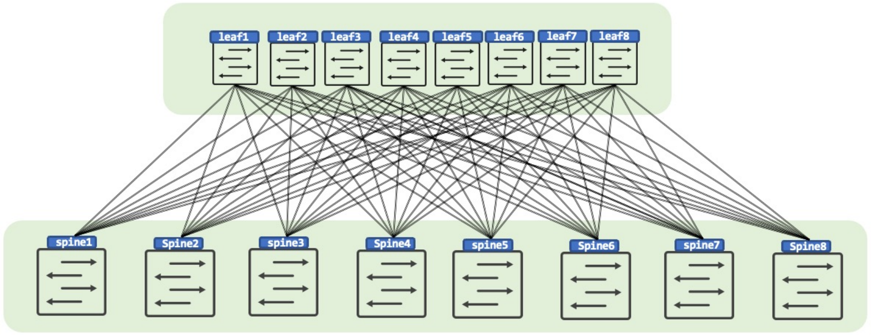

Focus

DC Workloads

#1 Hyper scaler

#2 Storage

#3 Media

#4 Enterprise

#5 AI/ML DC

#6 5G & RPHY

#7 SP/IX

#8 Gaming

#9 T2 Cloud

#10 Financial

Why is AI DC now?

- Maturity of AI ML models development:
  - AIML models became more accurate, more fluent, and more creative
  - Availability of opensource AI models increased recently
- The increasing availability of data:
  - As the amount of data available to AI models grows, so does the ability of those models to learn and improve
  - The more data an AIML model must learn from, the better it will be at generating natural language responses

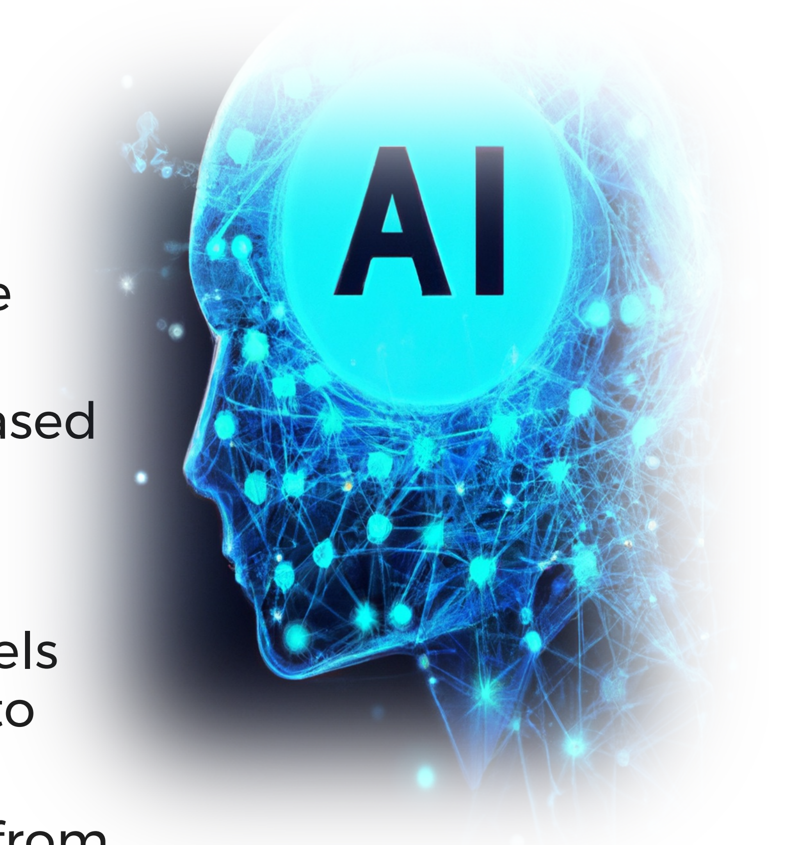

Why is AI DC now?

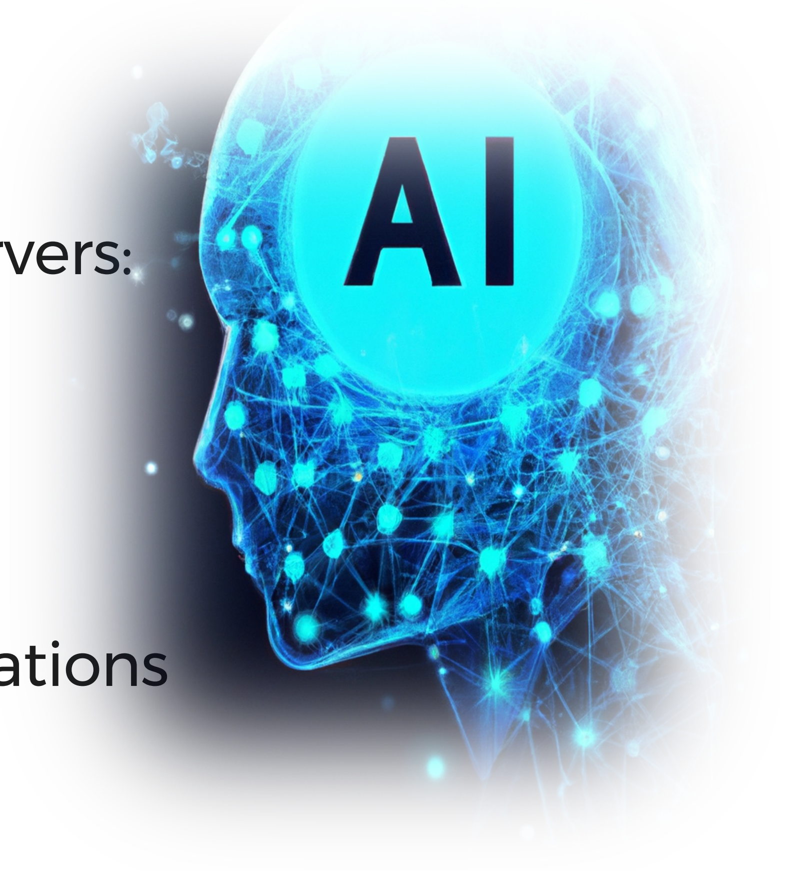

- Technological advancements at the servers:
  - parallel processing of the data requirement
  - use GPU instead of serialized CPU processing
- Quick adoption of Generative AI applications by the users

AI Model – Lifecycle

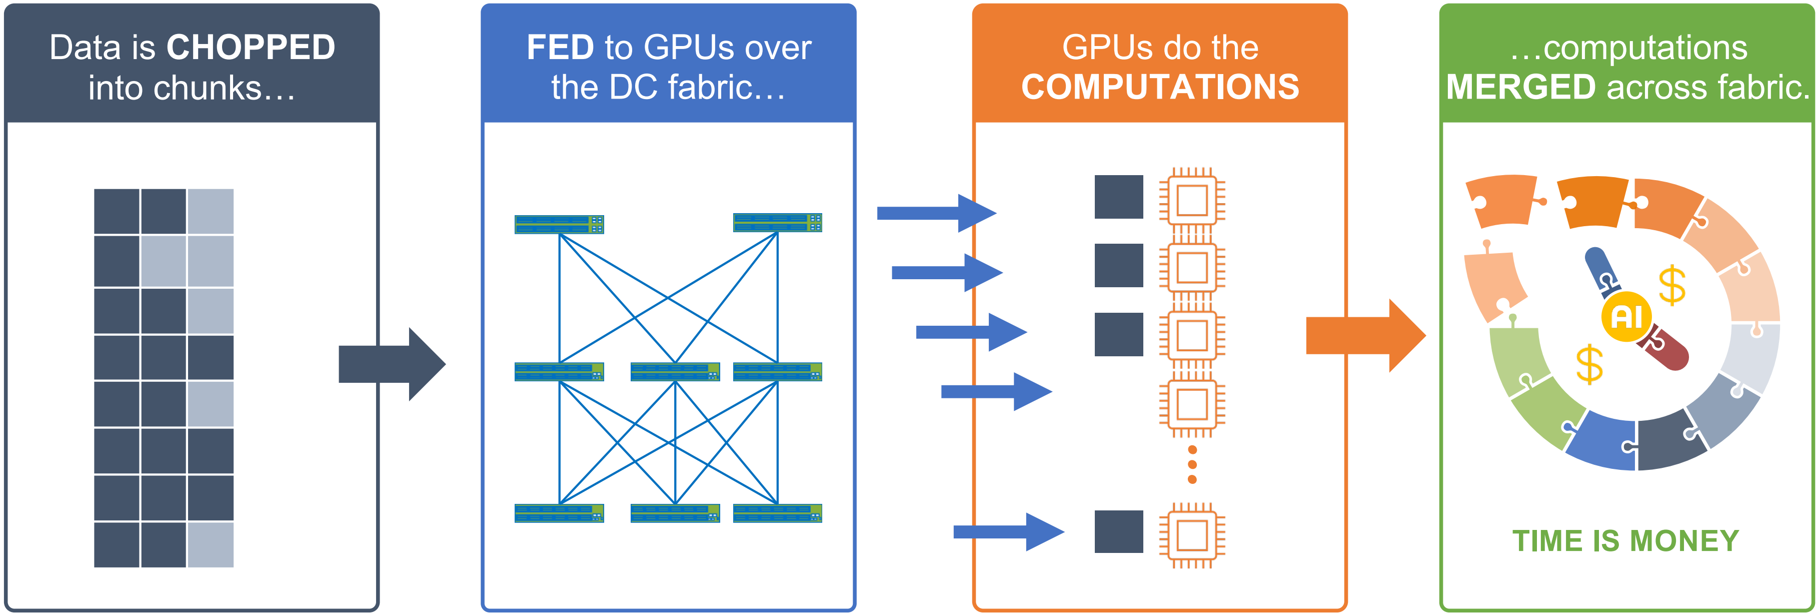

Data is **CHOPPED** into chunks…

FED to GPUs over the DC fabric…

GPUs do the **COMPUTATIONS**

…computations **MERGED** across fabric.

TIME IS MONEY

AI Model – Lifecycle

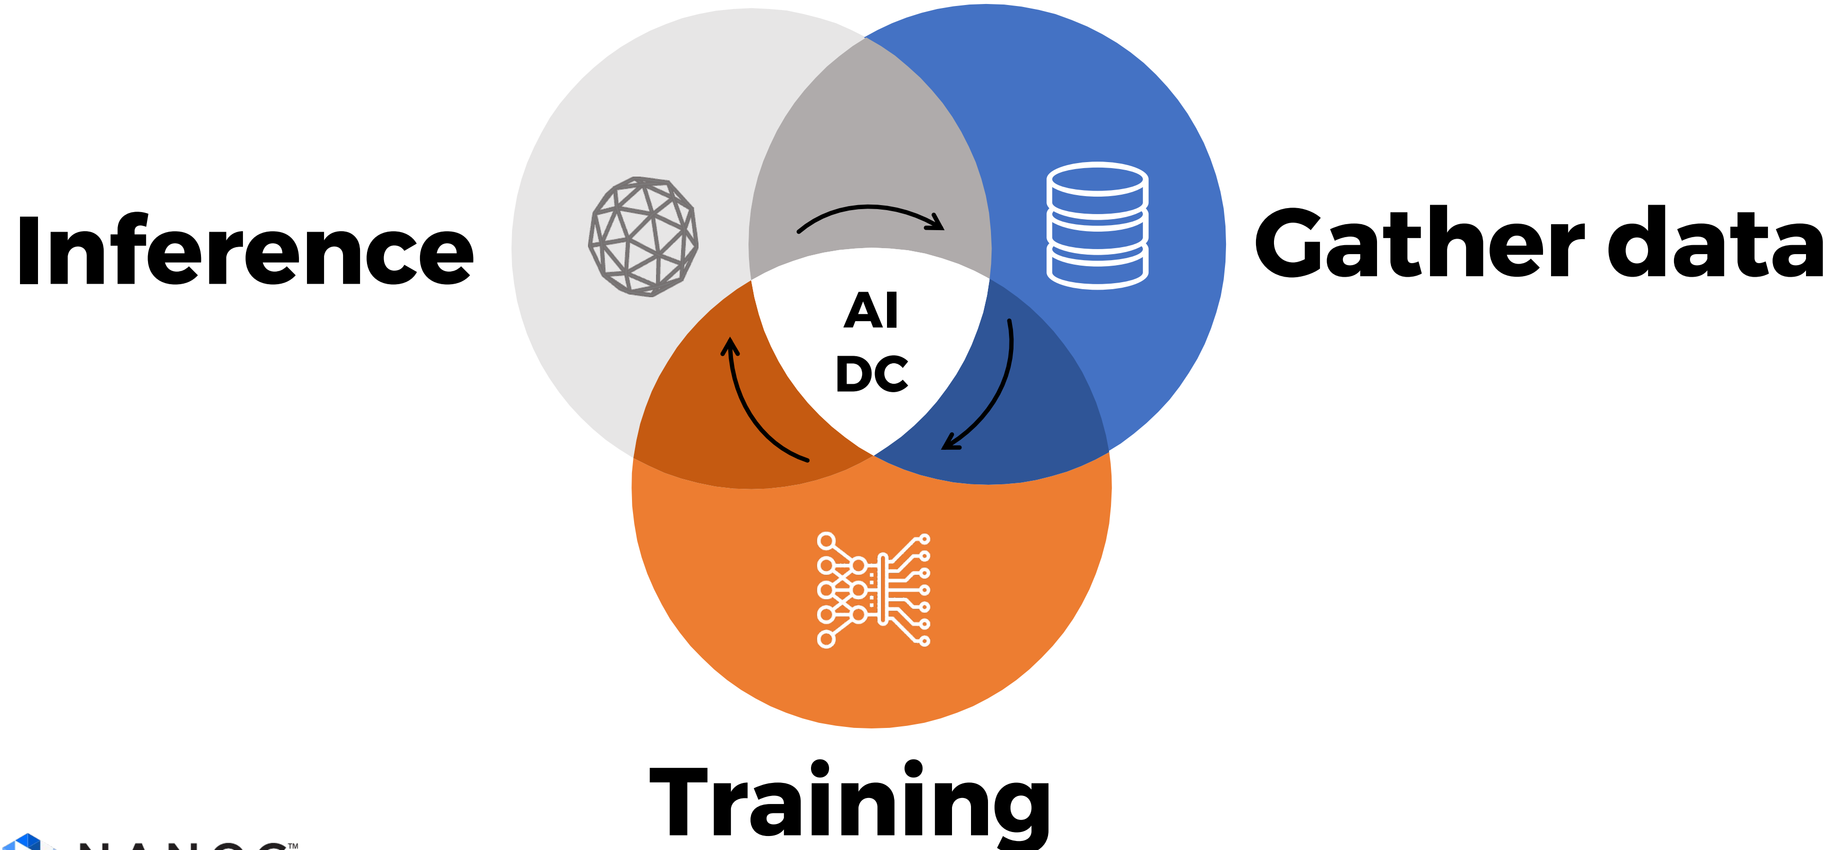

NANOG™

Anatomy of an AI DC Network

AI Cluster Networks

“Frontend”

- Inference clusters
- Shared storage pools
- Management

“Backend”

- GPU Compute Fabric
- Dedicated Storage Fabric

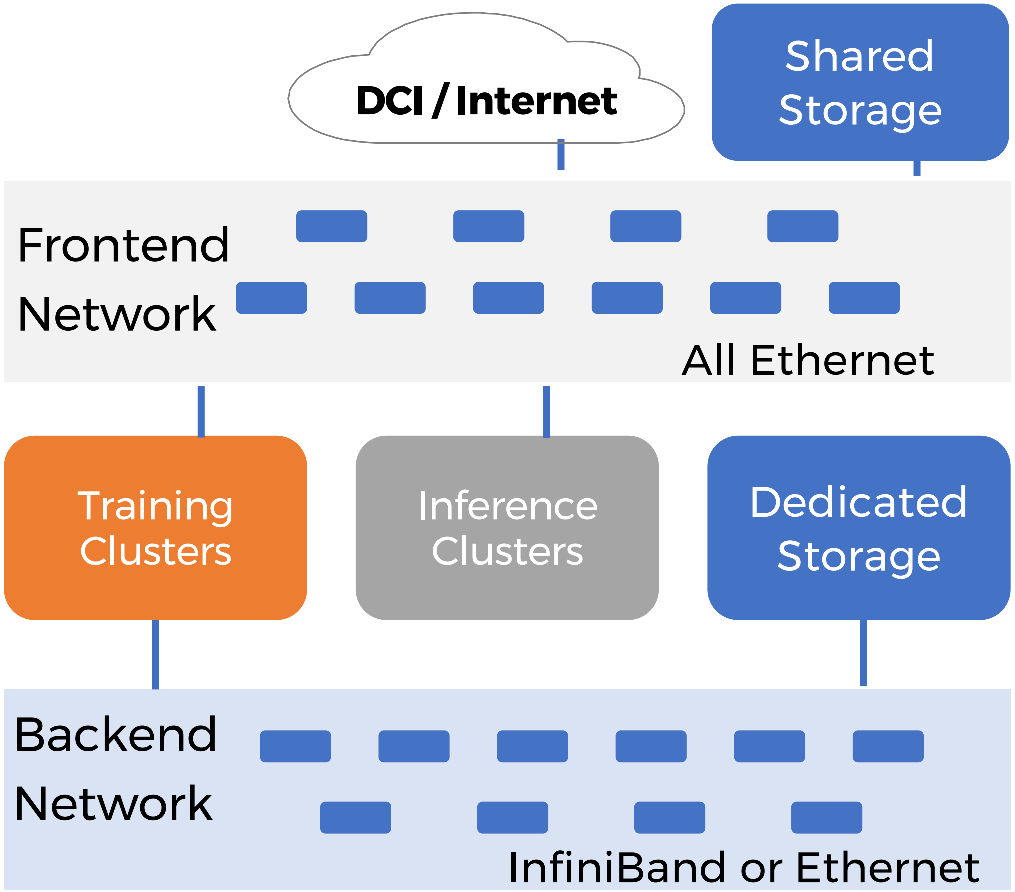

NANOG™

AI DC - Architectures

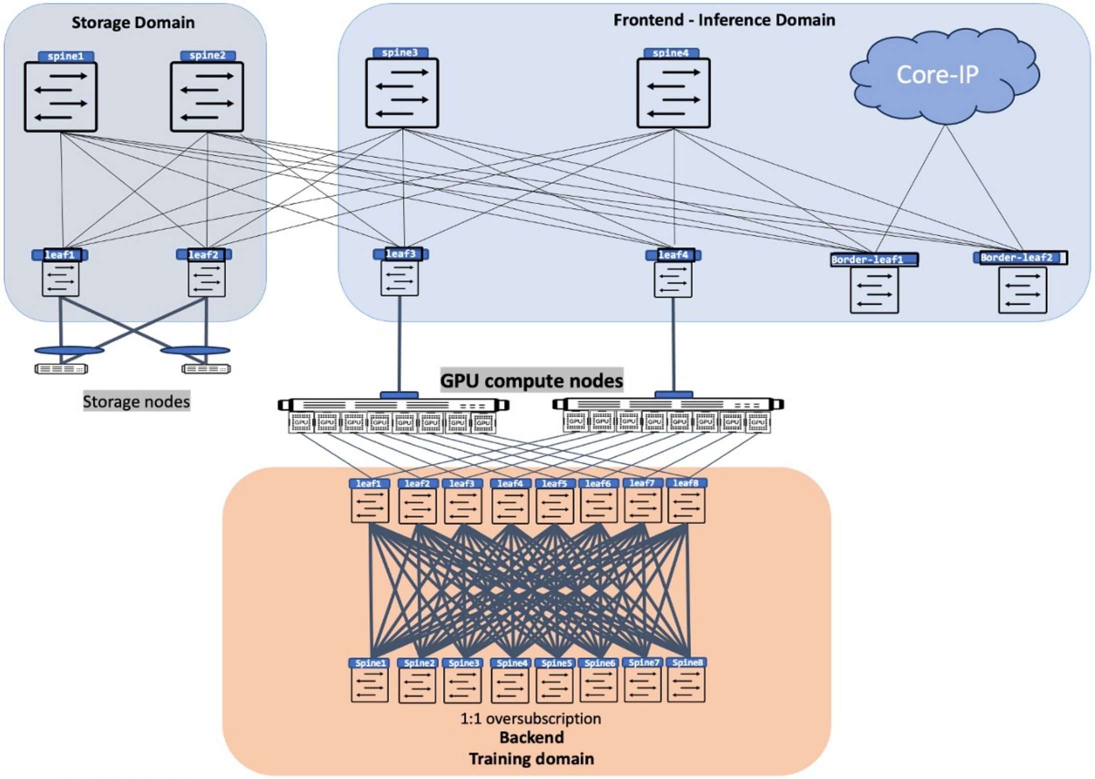

AI DC: Key capabilities:

- Efficient Load Balancing
- ROCEv2 Transport
- Congestion Mgmt
- Adaptive IP Routing
- Monitoring

AI DC – Stripe Optimized Design – SOD

GPU compute nodes rail optimized – Stripe 1
Each server with 8 x GPU

GPU compute nodes rail optimized – Stripe 2
Each server with 8 x GPU

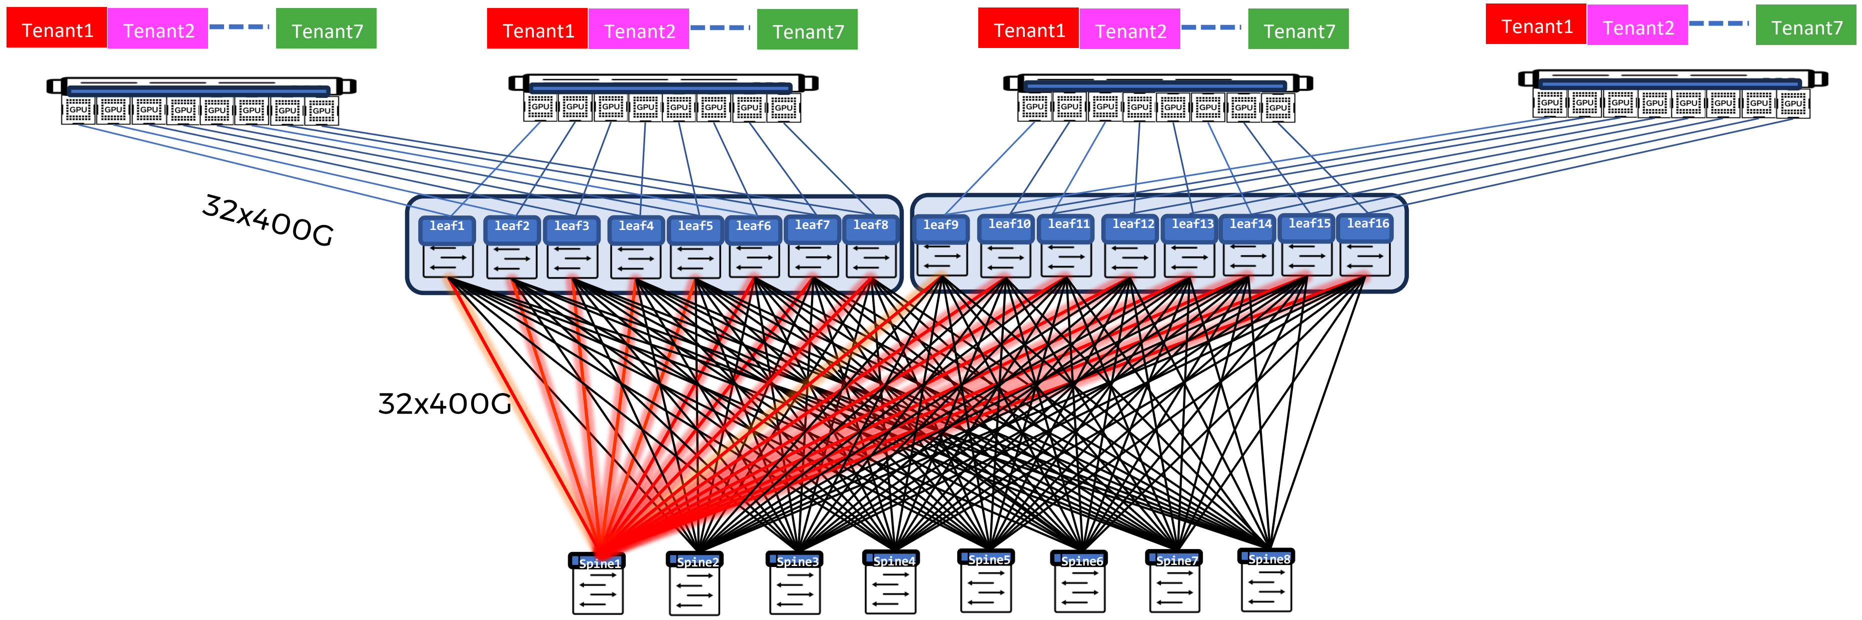

Stripe Optimized Design - SOD

+ Multi-tenancy

AI DC – Stripe Unified Design - SUD

Each server with 8 x GPU

Each server with 8 x GPU

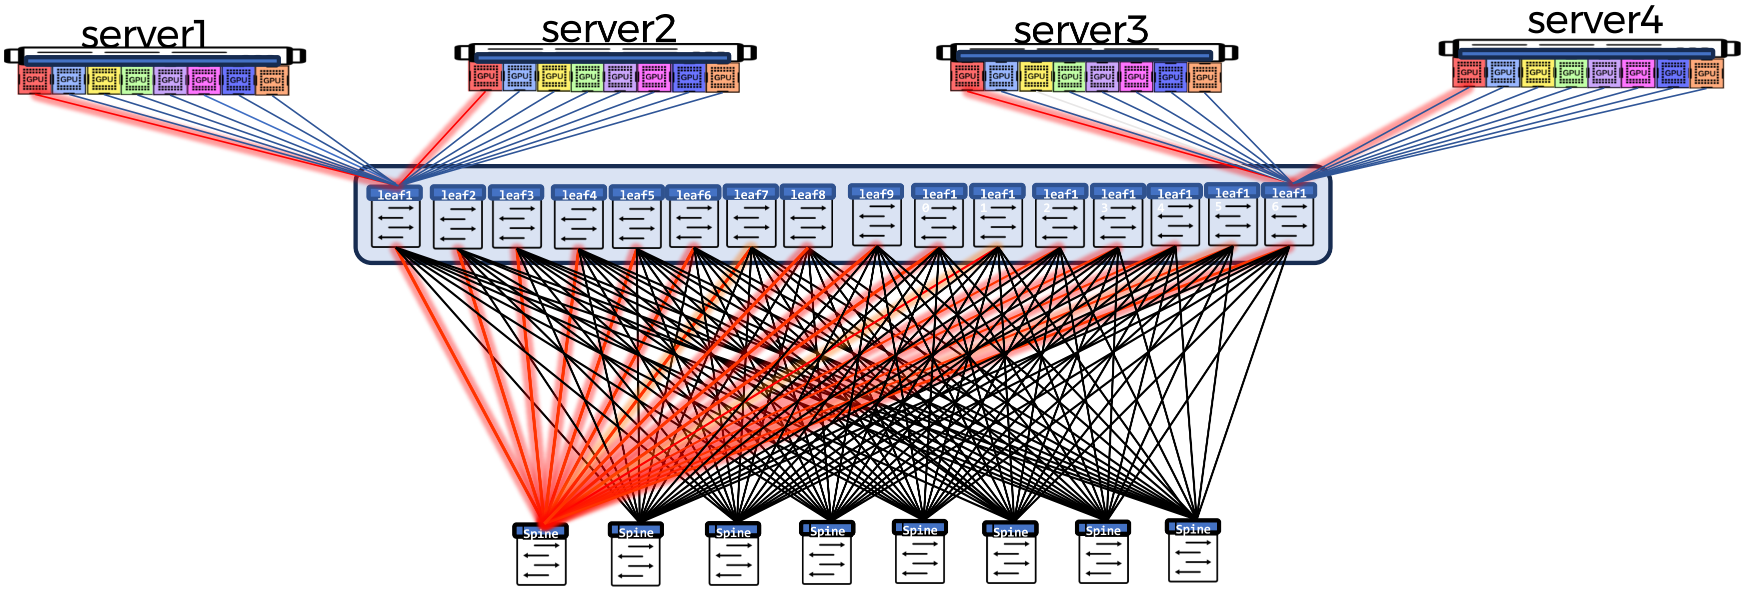

Stripe Unified Design - SUD

AI DC: Requirements

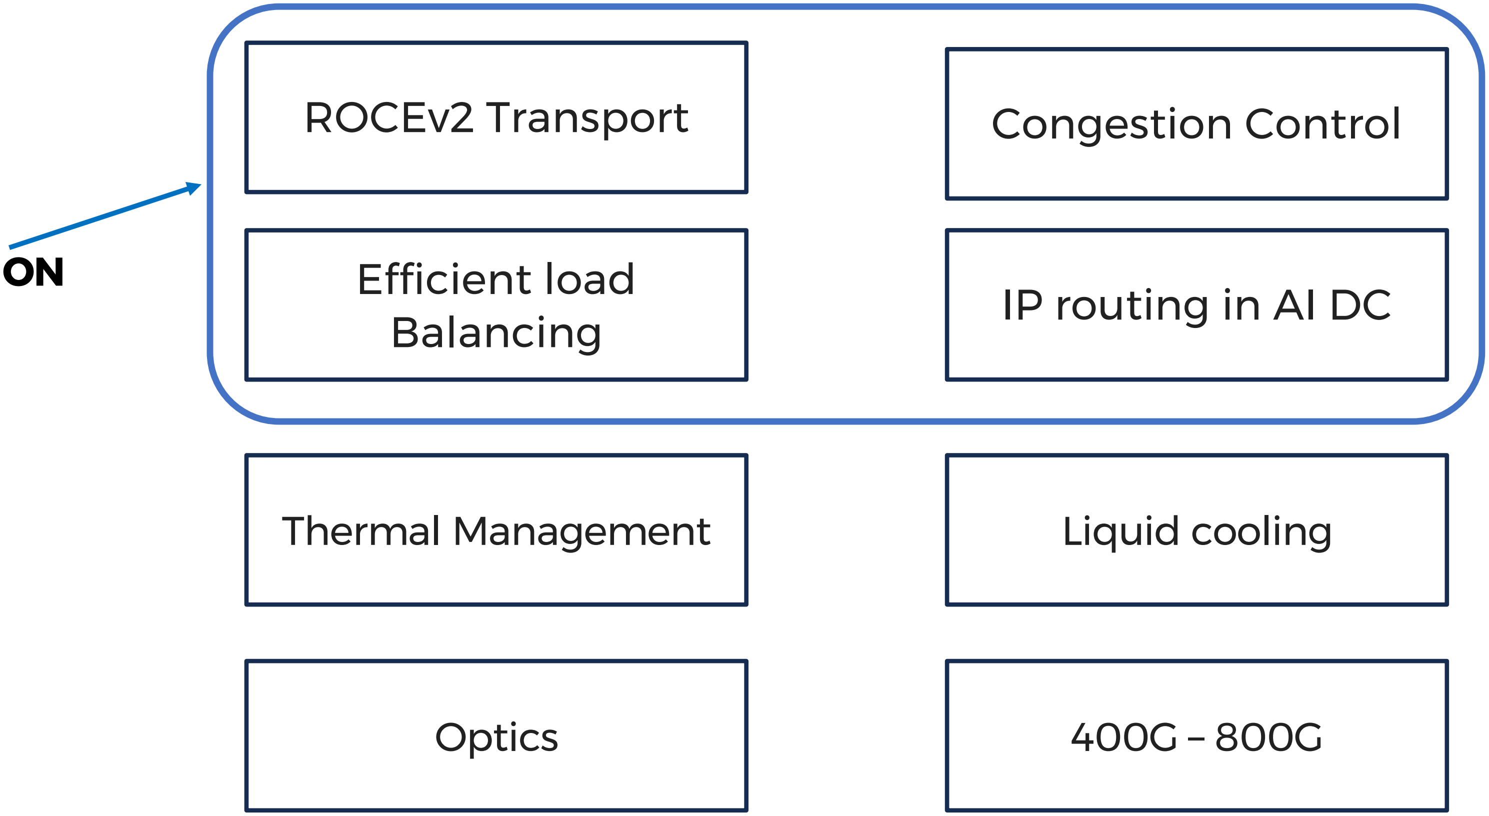

THIS SESSION FOCUS

ROCEv2 Transport

Congestion Control

Efficient load Balancing

IP routing in AI DC

Thermal Management

Liquid cooling

Optics

400G – 800G

RDMA Workload : AI DC

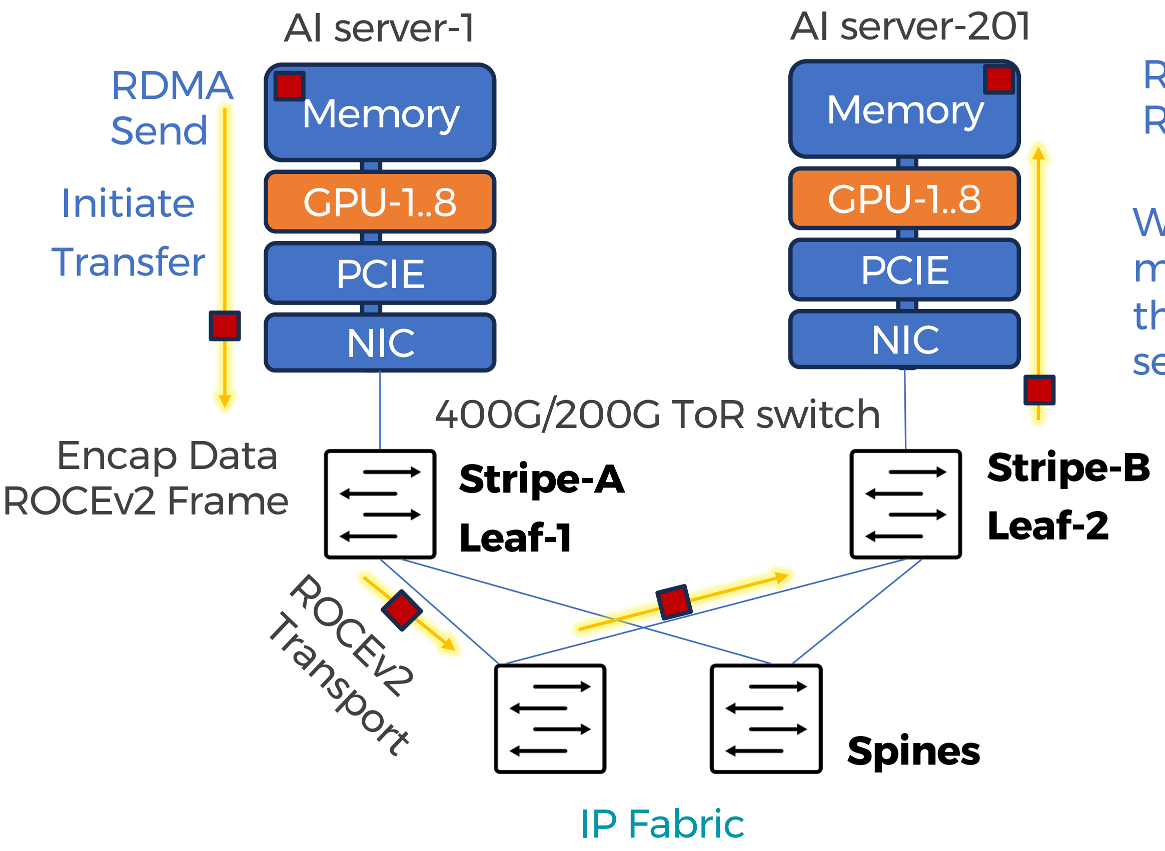

Write directly to the GPU memory without going through the CPU of the server

NANOG™

ROCEv2 - Transport for AI DC

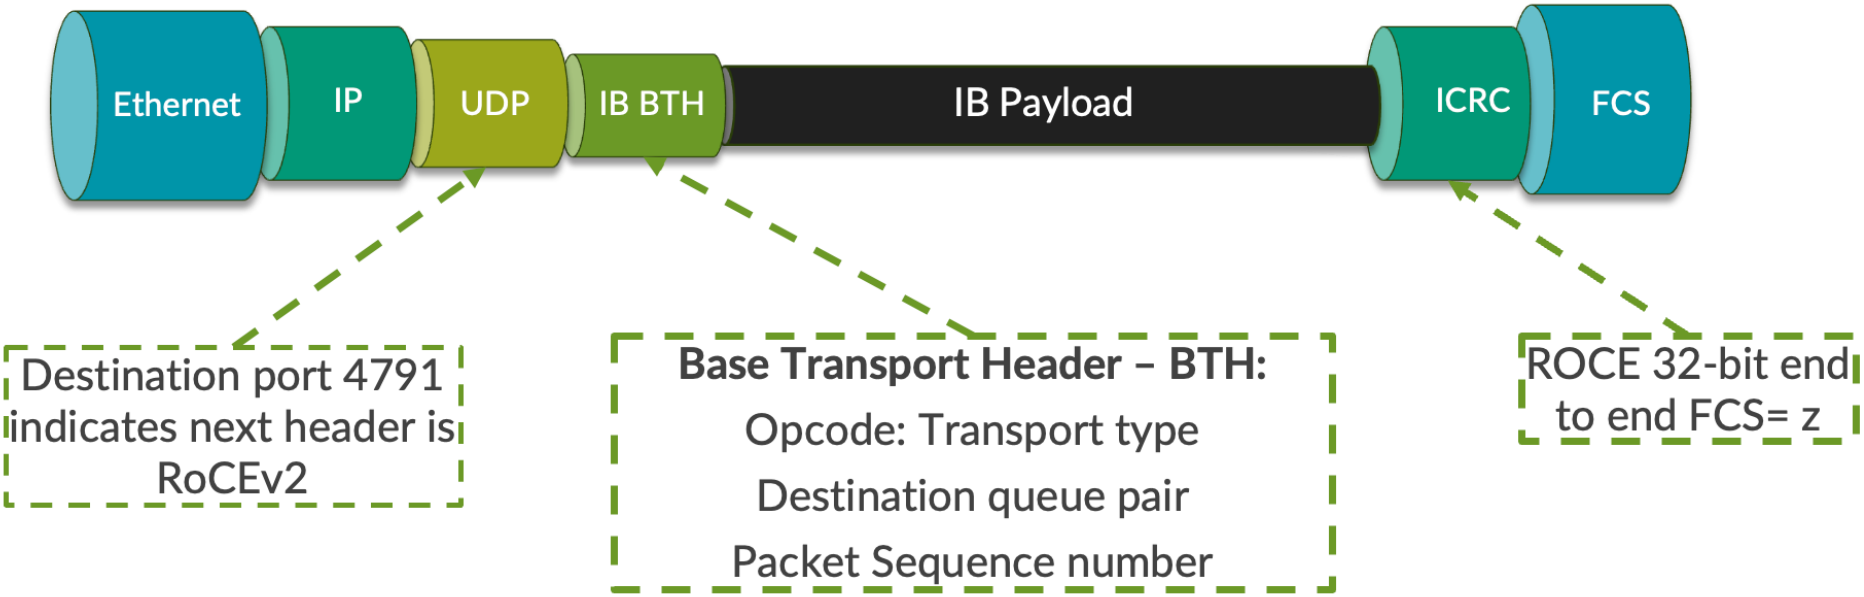

Destination port 4791 indicates next header is RoCEv2

Base Transport Header – BTH:

Opcode: Transport type

Destination queue pair

Packet Sequence number

ROCE 32-bit end to end FCS= z

ROCEv2 – session establishment

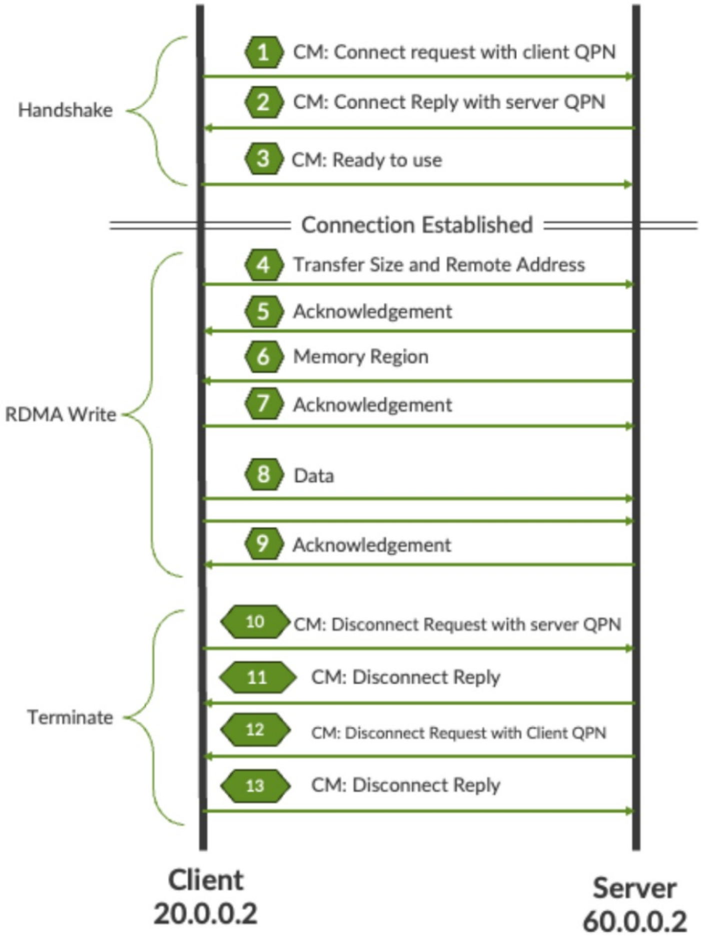

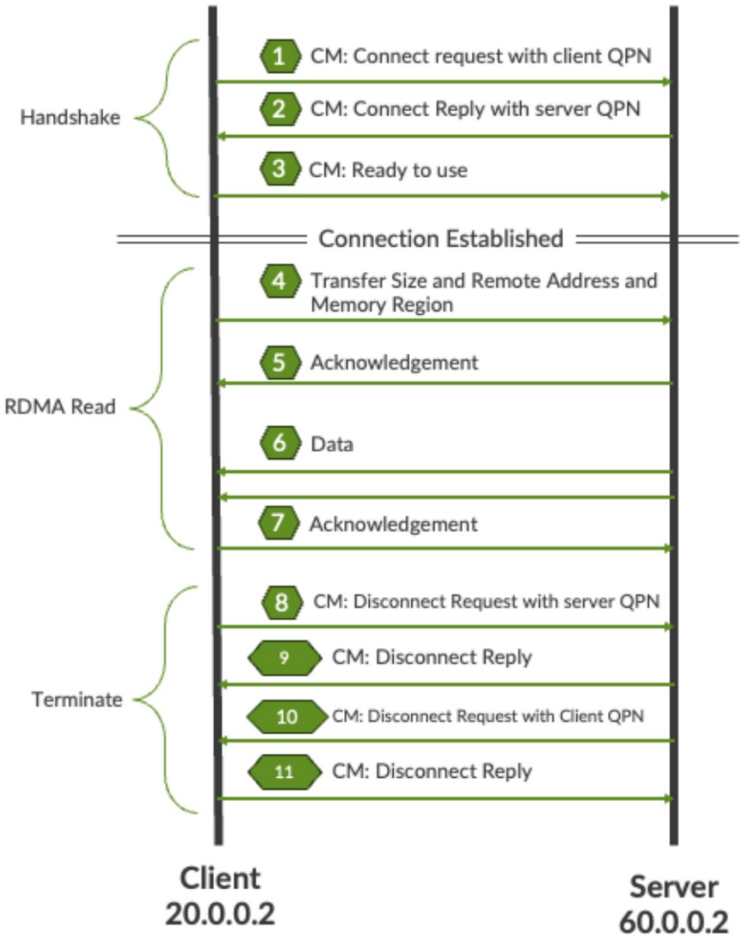
Handshake
1 CM: Connect request with client QPN
2 CM: Connect Reply with server QPN
3 CM: Ready to use
Connection Established
RDMA Read
4 Transfer Size and Remote Address and Memory Region
5 Acknowledgement
6 Data
7 Acknowledgement
Terminate
8 CM: Disconnect Request with server QPN
9 CM: Disconnect Reply
10 CM: Disconnect Request with Client QPN
11 CM: Disconnect Reply
Client
20.0.0.2
Server
60.0.0.2

DCQCN – PFC-DSCP vs ECN

PFC-DSCP
Pause-level-1

↓

PFC-DSCP
Pause-level-2

↓

PFC-DSCP
Pause-level-3

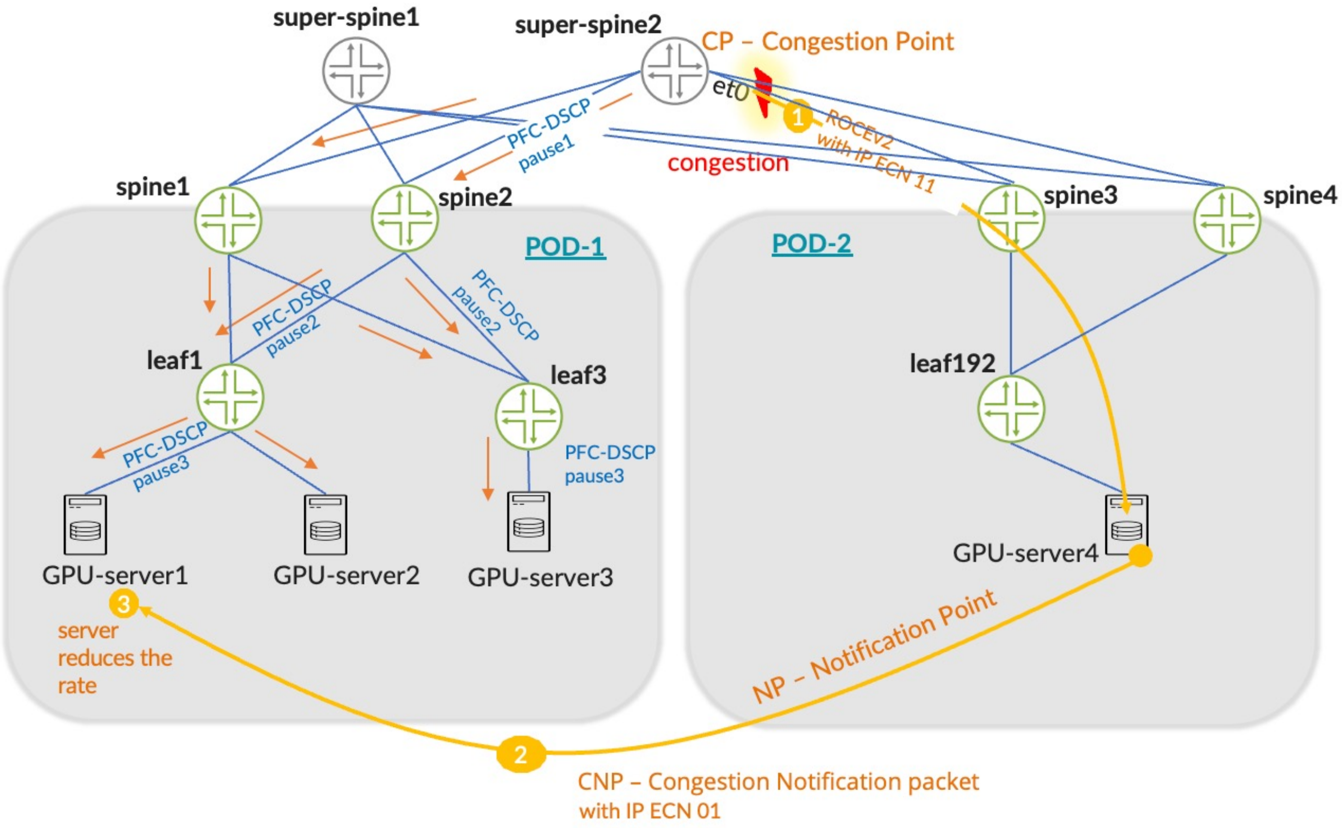

AI DC - Dynamic Load Balancing

DLB (Dynamic Load Balancing) - per-packet optimal spraying

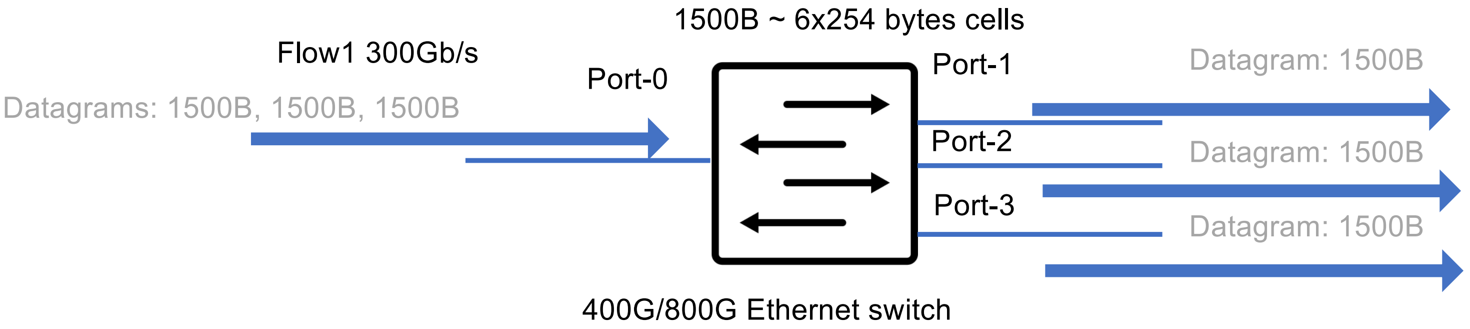

Packet **re-ordering may happen** at the destination NIC Card connected to the GPU

AI DC - Dynamic Load Balancing

DLB (Dynamic Load Balancing) – “flowlet” mode

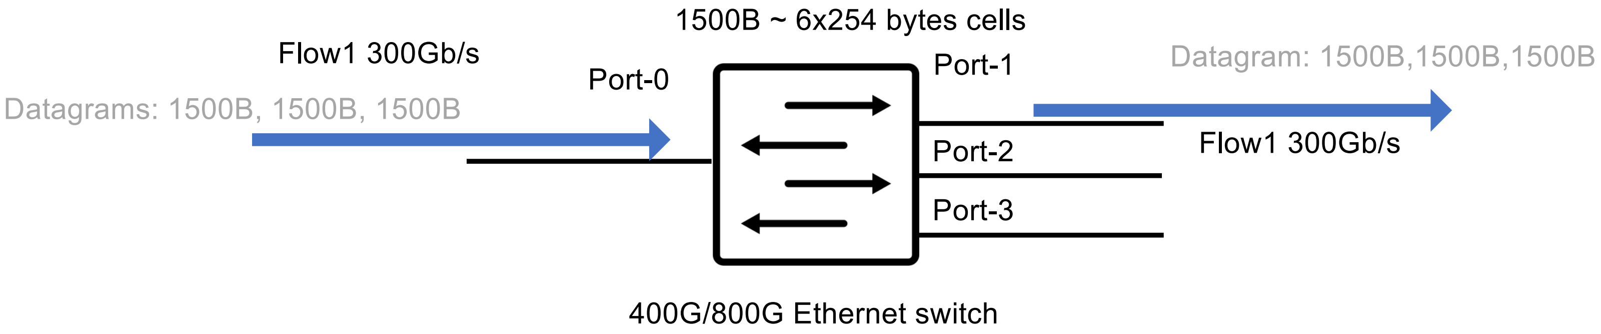
1500B ~ 6x254 bytes cells
Flow1 300Gb/s
Port-0
Port-1
Datagram: 1500B,1500B,1500B
Datagrams: 1500B, 1500B, 1500B
Port-2
Flow1 300Gb/s
Port-3
400G/800G Ethernet switch

Packet **re-ordering won’t happen** at the destination NIC Card connected to the GPU

AI DC – Selective Load Balancing

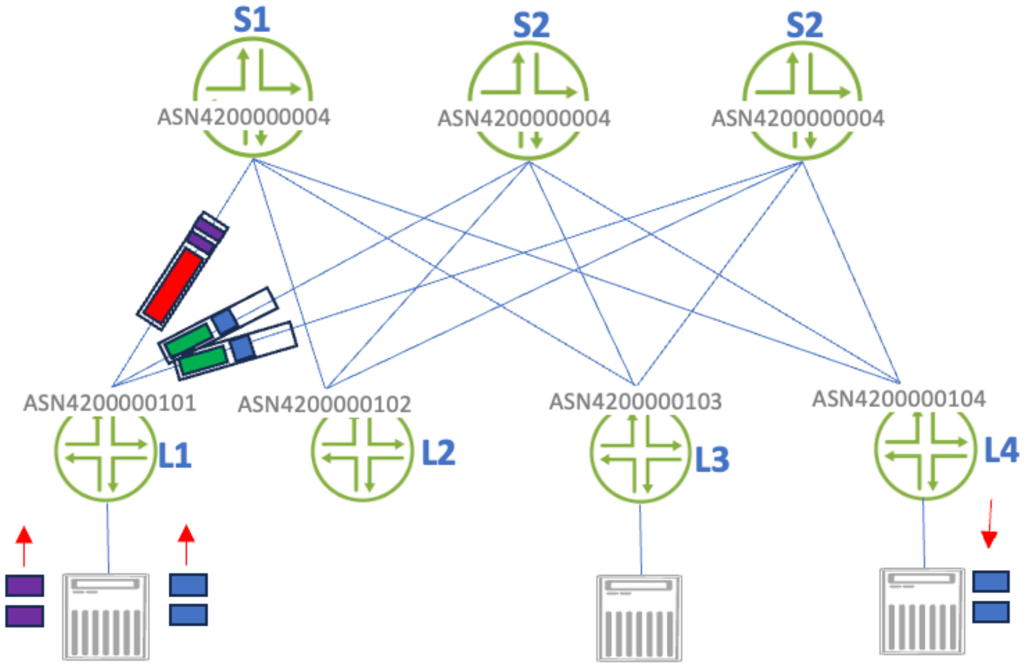

- Ability to selectively enable DLB via access lists for read/write operations
- It can handle out-of-order packets and enable DLB per packet mode for just that service

NANOG™

AI DC – Global Load Balancing

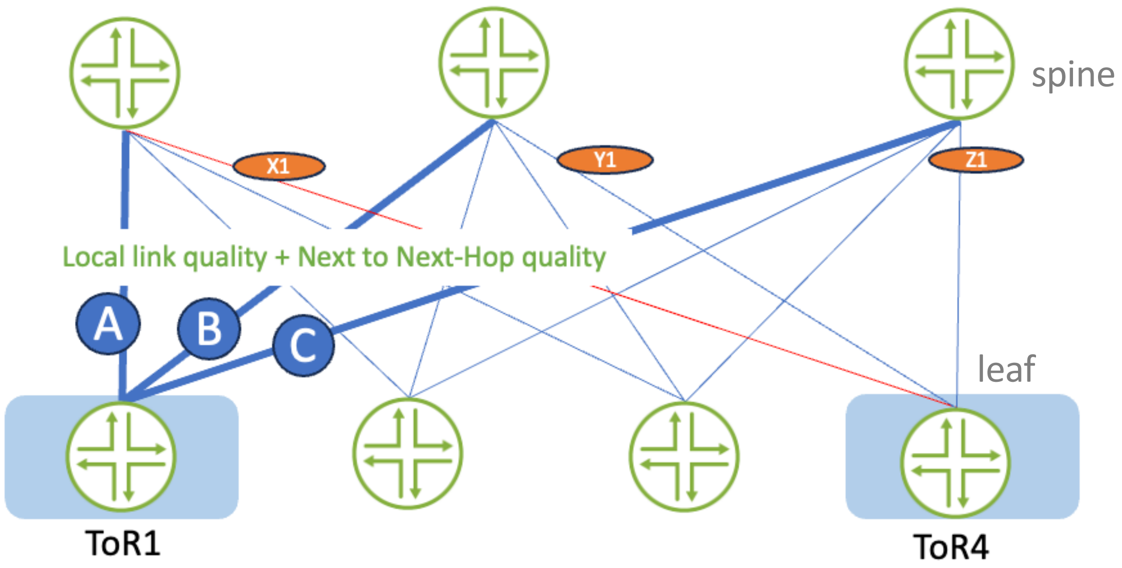

Route in HW at ToR1

NEXTHOP ID	Remote Quality Profile ID
ECMP_NH_ID1	SwitchID.ToR4

Local Port	Quality
A	Q(A)
B	Q(B)
C	Q(C)

Quality	Remote Port
Q(X1)	X1
Q(Y1)	Y1
Q(Z1)	Z1

- Global Load Balancing (GLB) uses **path quality**

- GLB selects a better end-to-end path.

AI DC: Efficient load Balancing summary

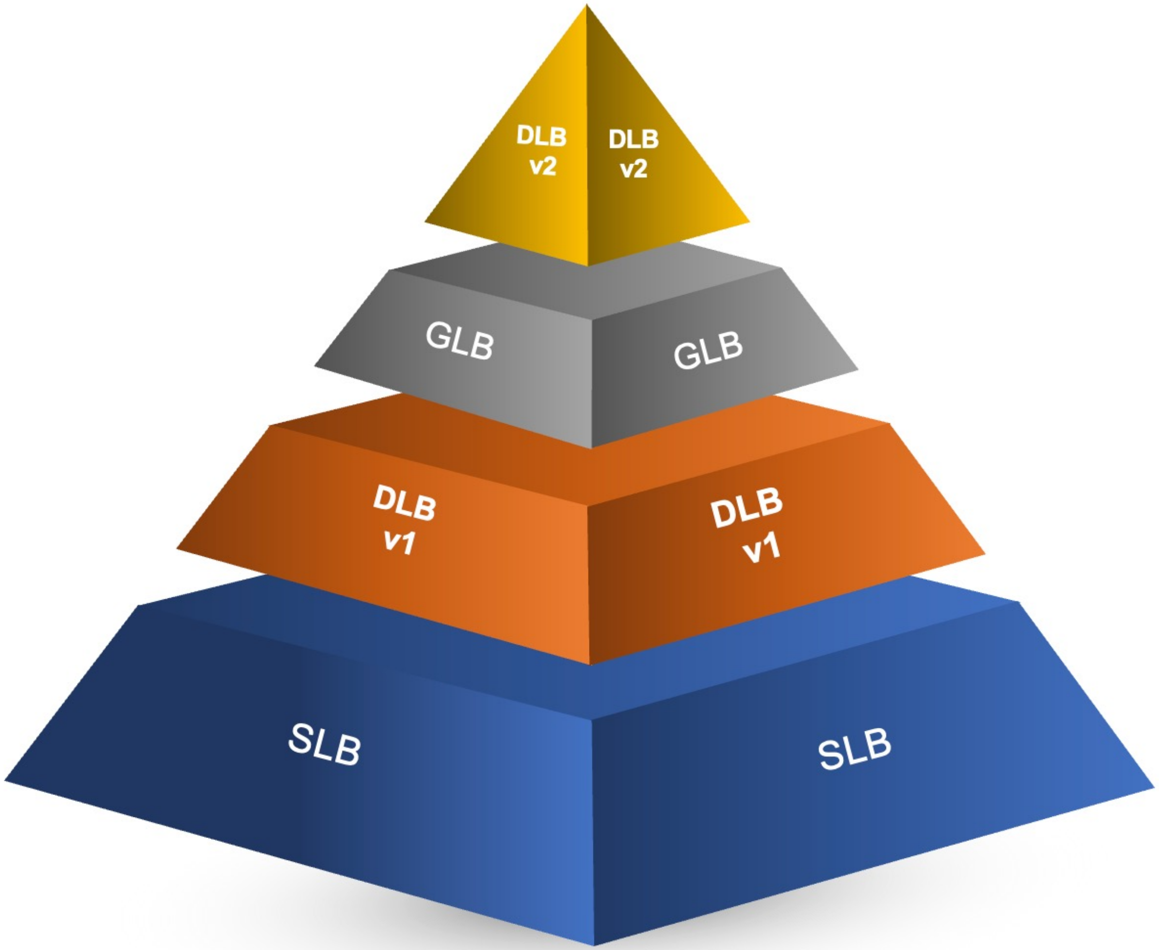

IP Routing for AI DC

Frontend	eBGP underlay/overlay: • underlay: eBGP unnumbered / RFC5549 • overlay: EVPN-VXLAN
Backend	• BGP unnumbered/RFC5549 • IGP protocols: RIFT or ISIS

BGP unnumbered / RFC5549

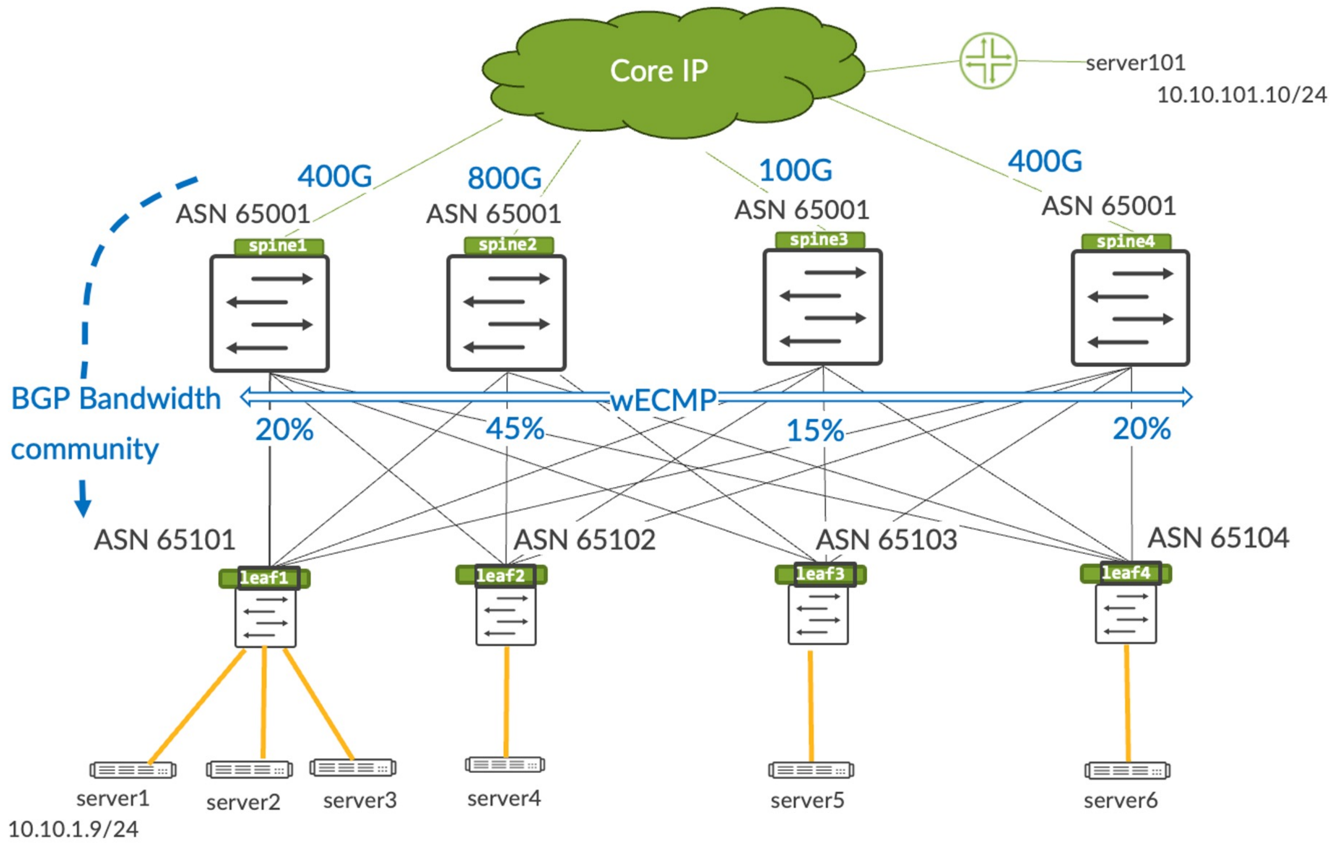

RIFT routing for backend network

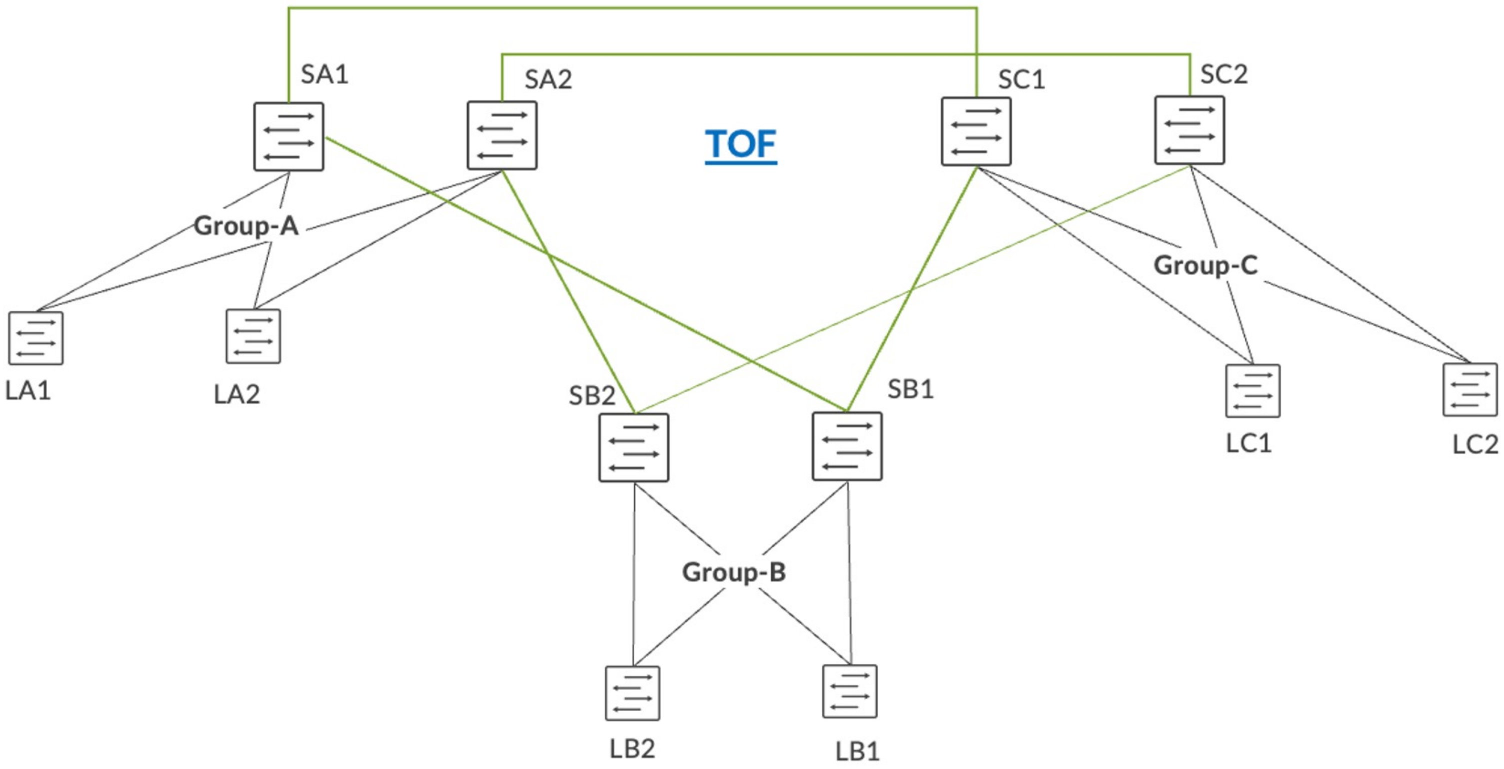

AI DC key takeaways

- The number of new AI applications is increasing over time.
- Dedicated AI DC infrastructures are built to accelerate parallel data processing.
- Ethernet 400G/800G adoption is increasing thanks to AI
- Congestion Management & Load Balancing efficiency are the key network components in AI DC

Thank you

Feb 12-14, 2024

NANOG™