



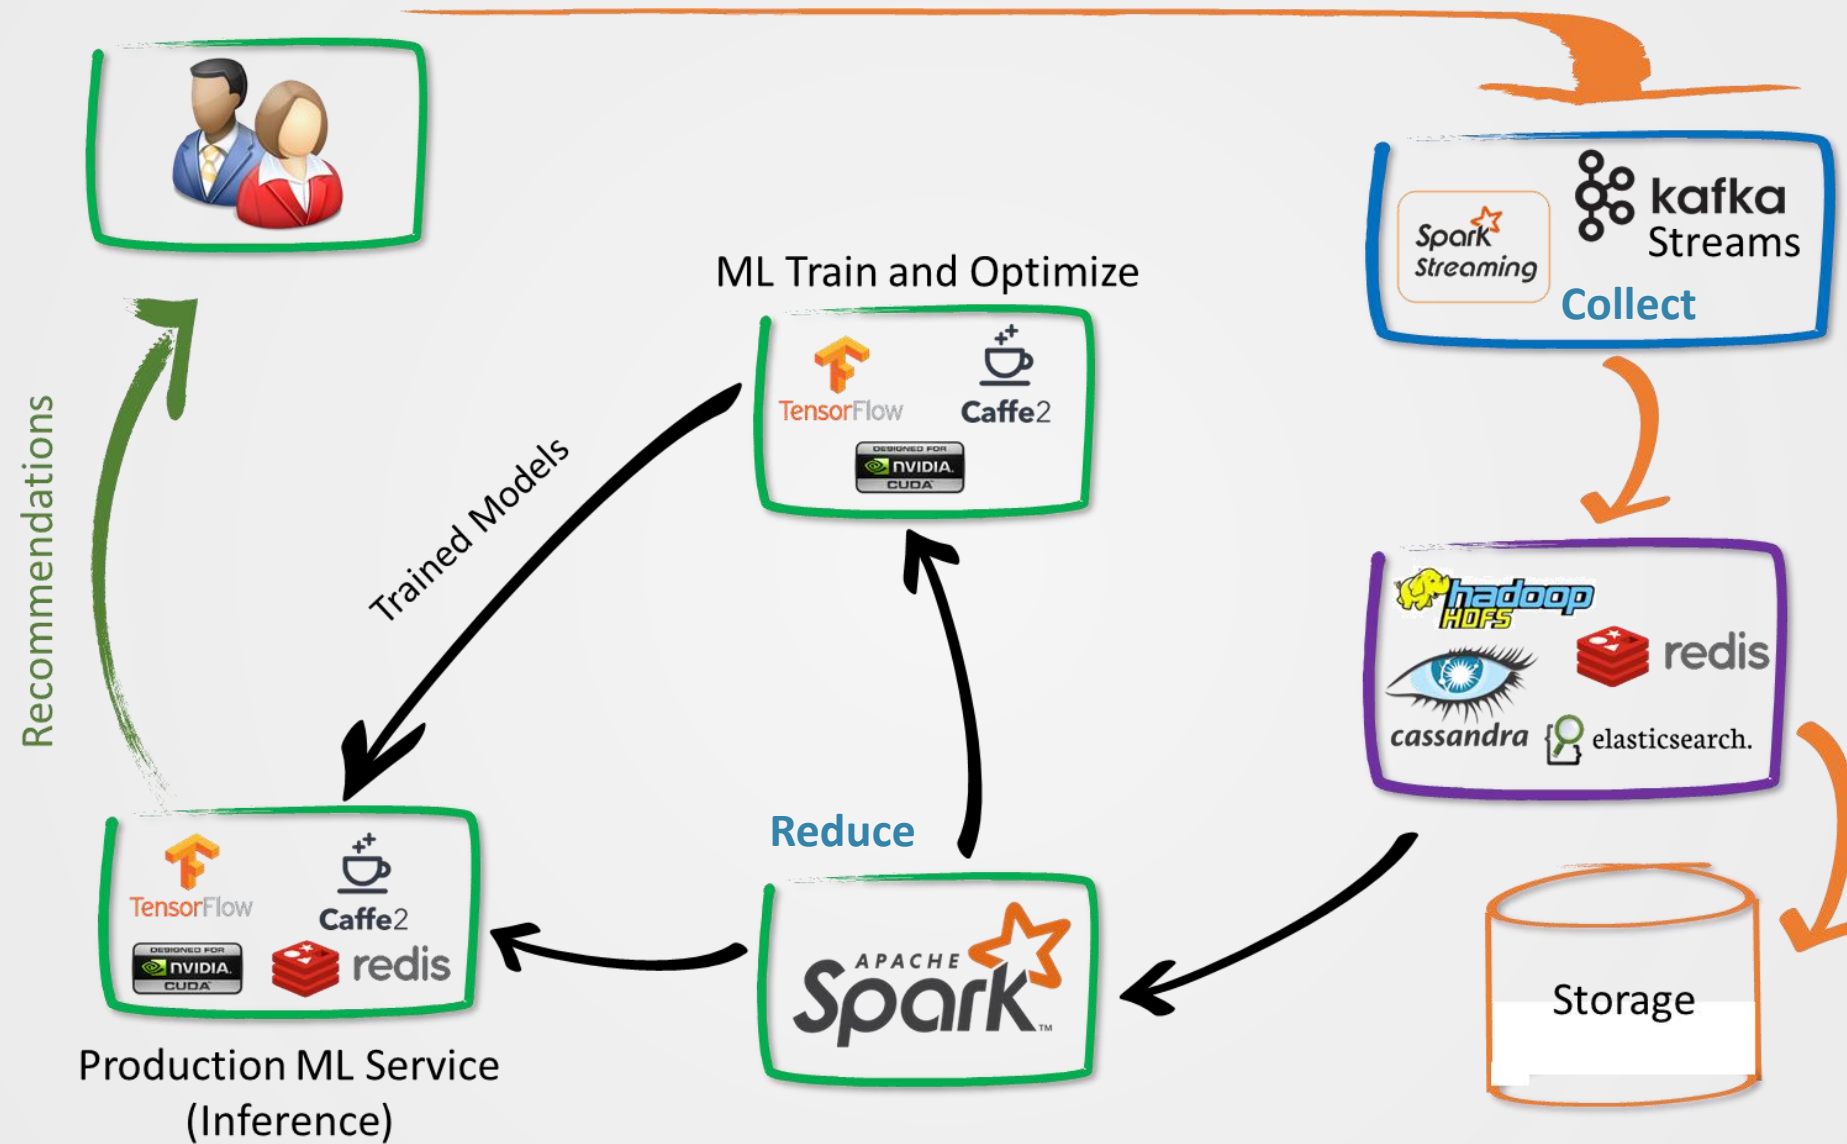
Network Accelerated AI

Elad Wind

Mar 28, 2018

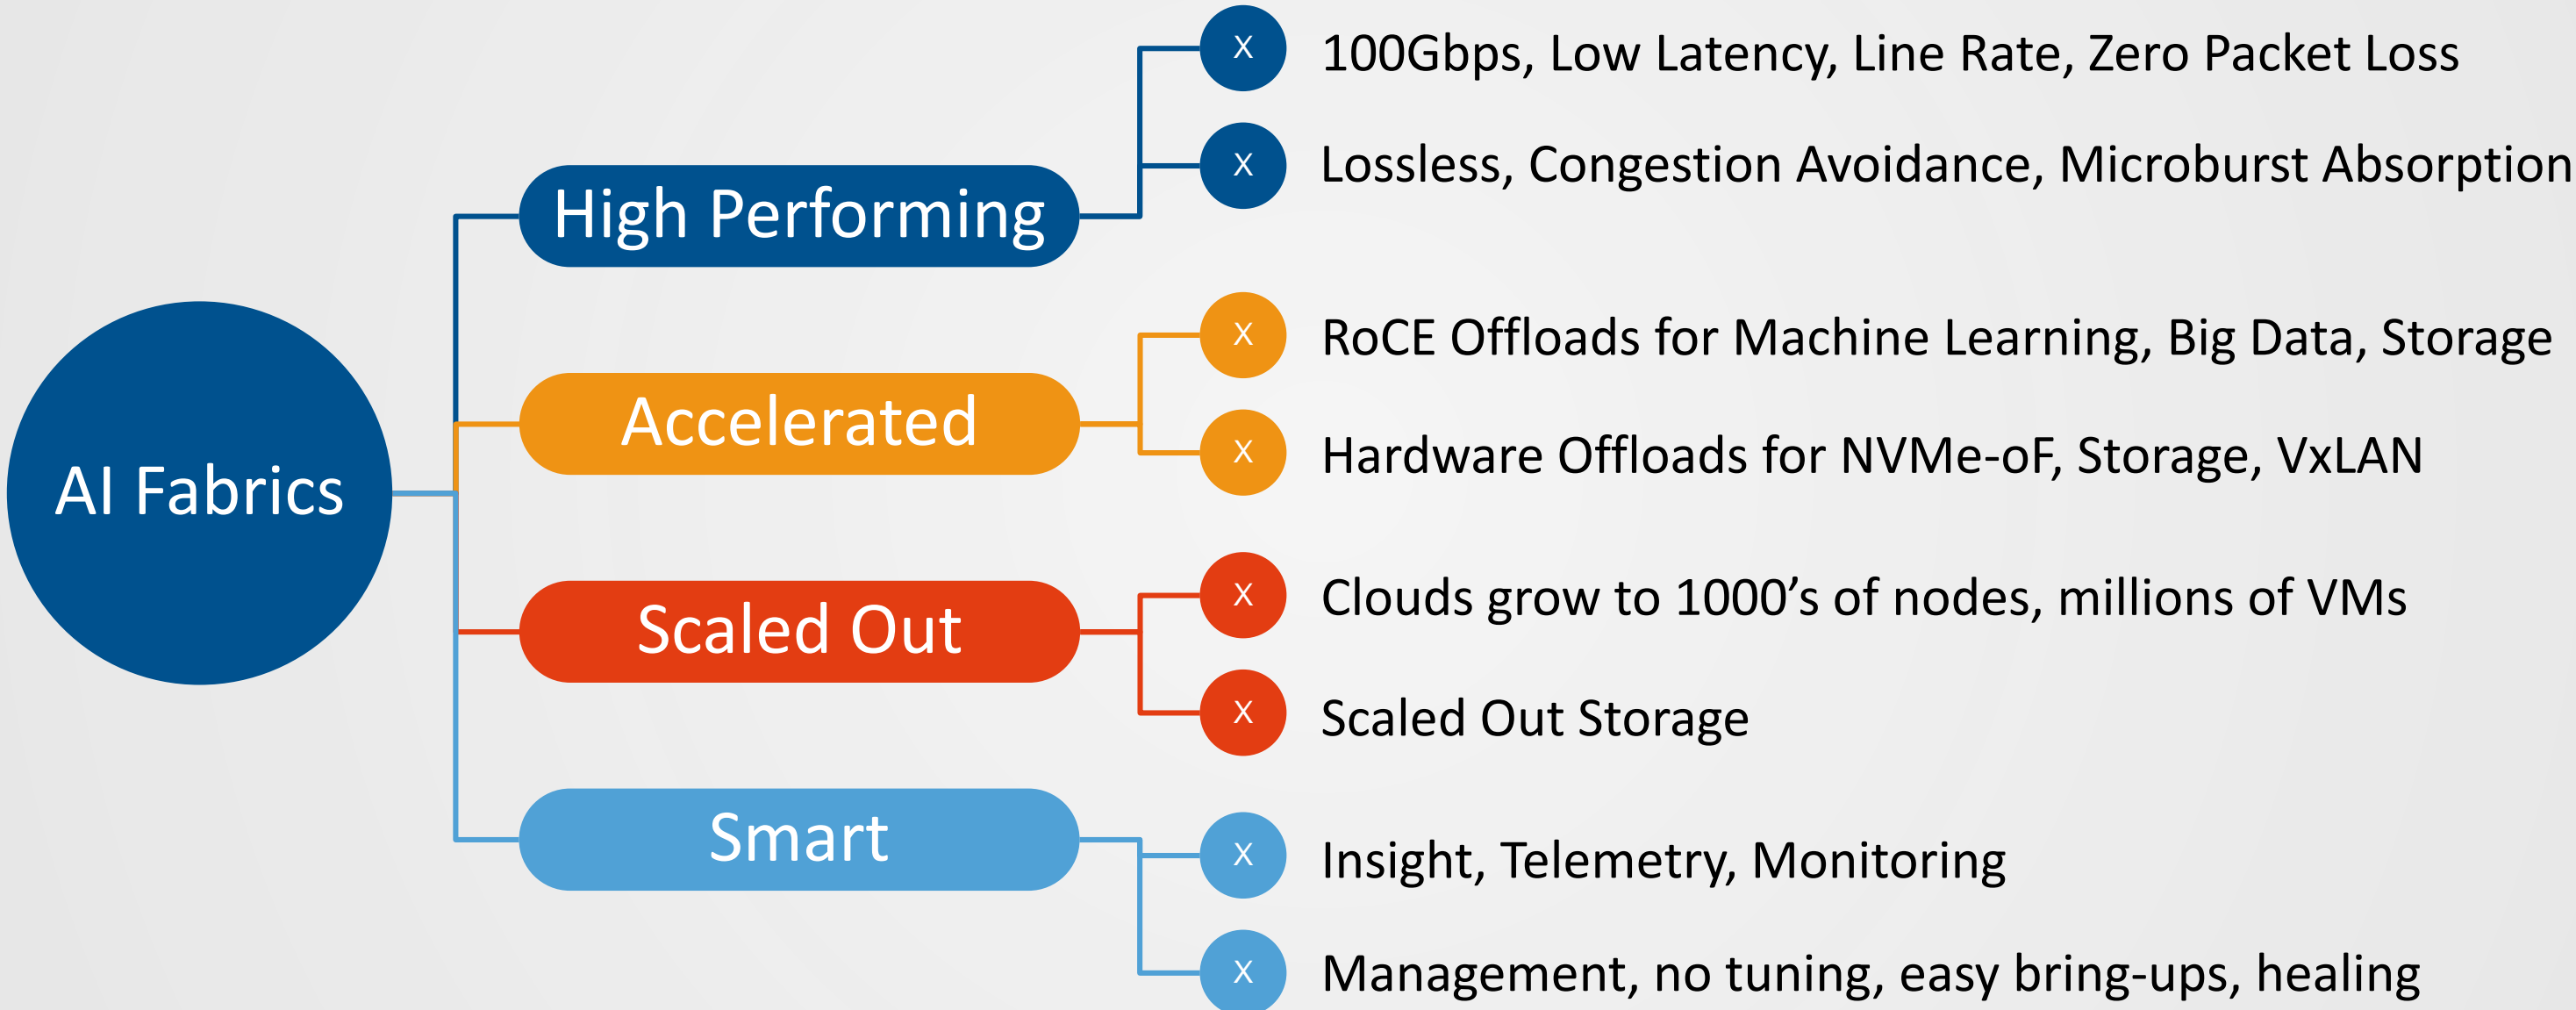


AI / Deep Learning Data Pipeline




Smooth Data Flow For Pipeline Productivity

Networking Requirements




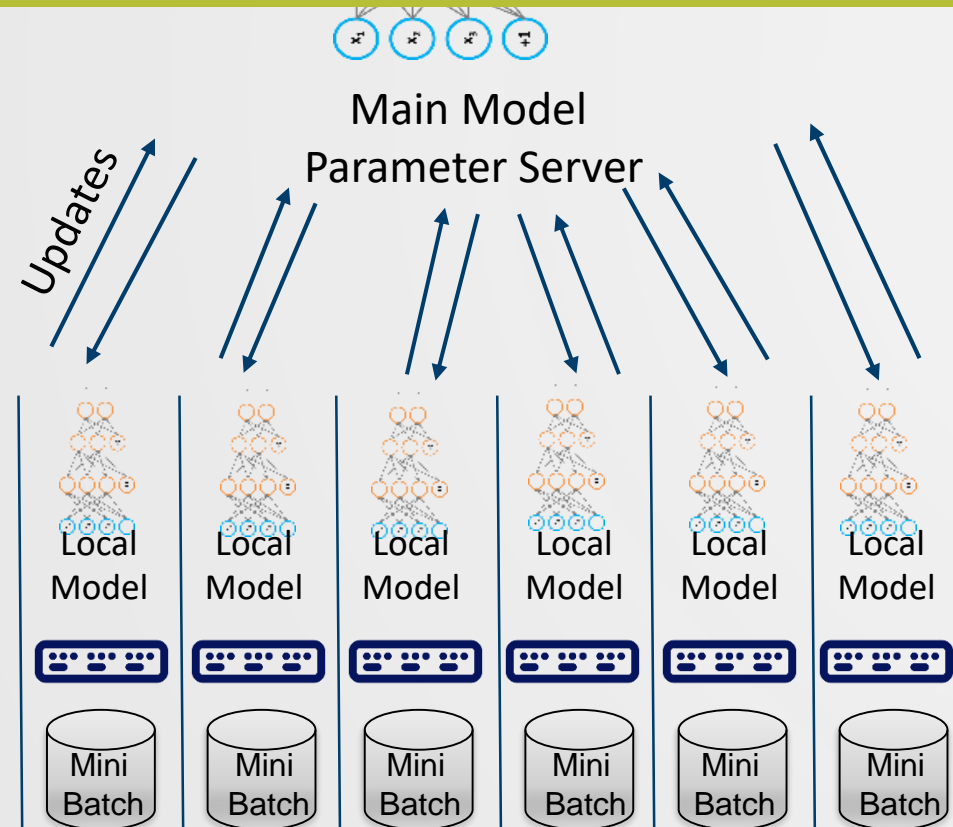
High Performing



Training Phase Is The Biggest Bottleneck

Scale Out To Reduce Training Time

- ↑workers (nodes)
- ↑CPUs, ↑GPUs
- Distribute the load
- Data parallelism



Network Requirements

- CLOS
- East-West vs. North-South
- Many-to-Many, Many-to-One
- ↓Blocking ↑Bandwidth ↓Latency
- Congestion Avoidance
- Storage is everywhere

Networks Optimized for AI





Congestion Management

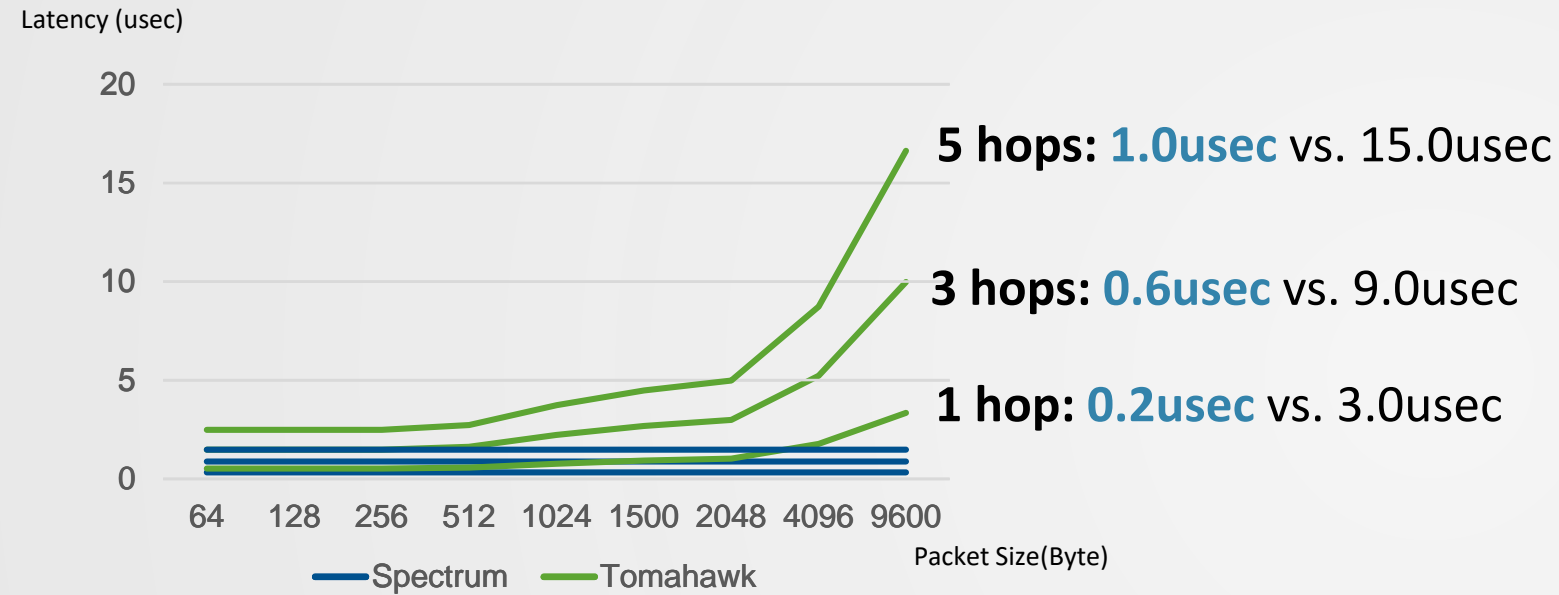
Fairness

Avoidable Packet Loss


Guaranteed Low Latency



Distributed Systems: Network Latency Builds Up



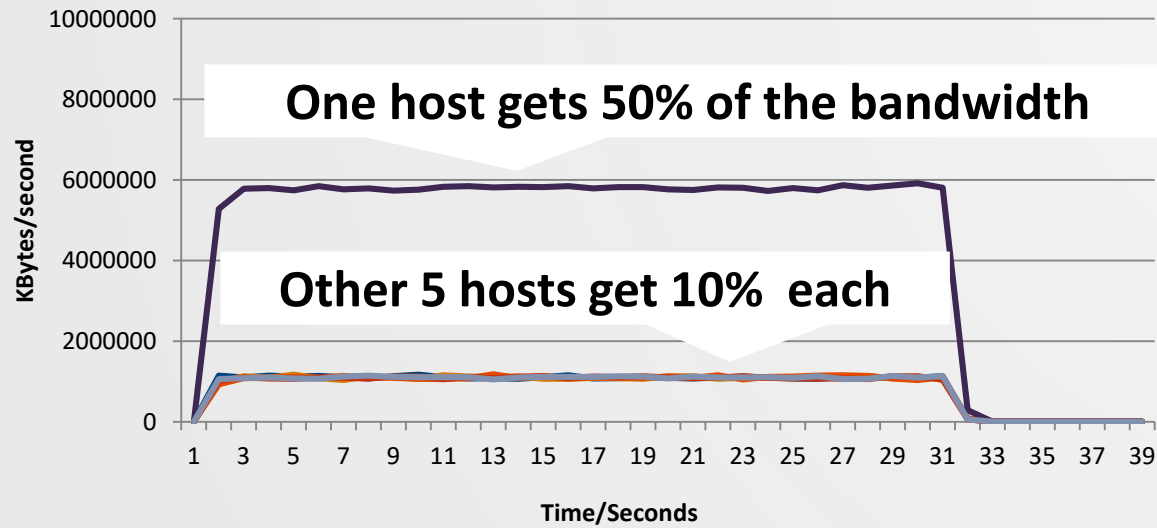

Mellanox: Consistent Low Latency



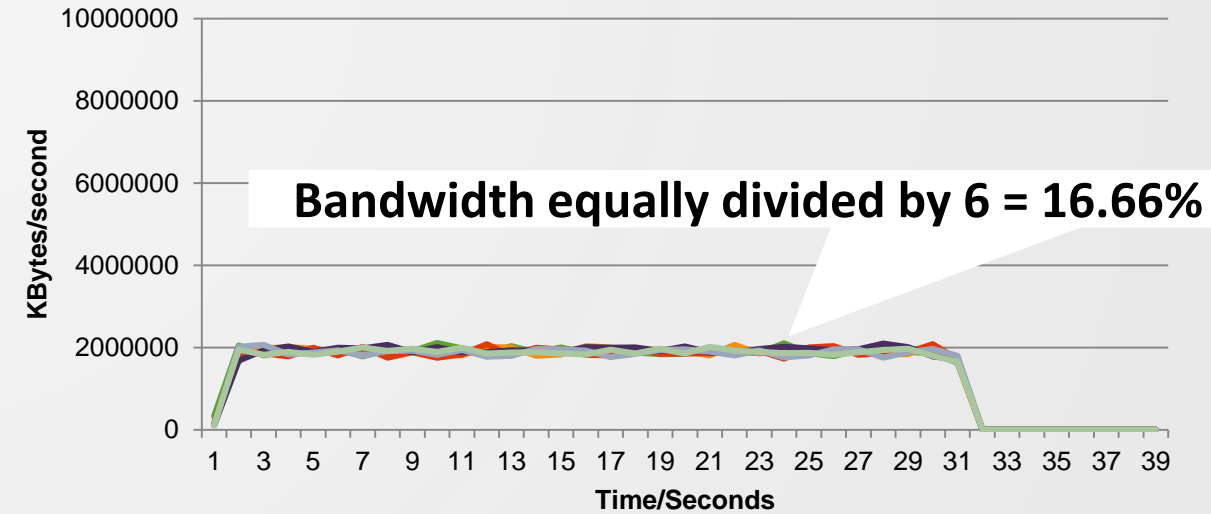
Spectrum Predictable QoS with RoCE

**Fair Bandwidth Distribution:
Distributed Learning is I/O intensive**


6 to 1 communication, ECN and PFC enabled



Competition

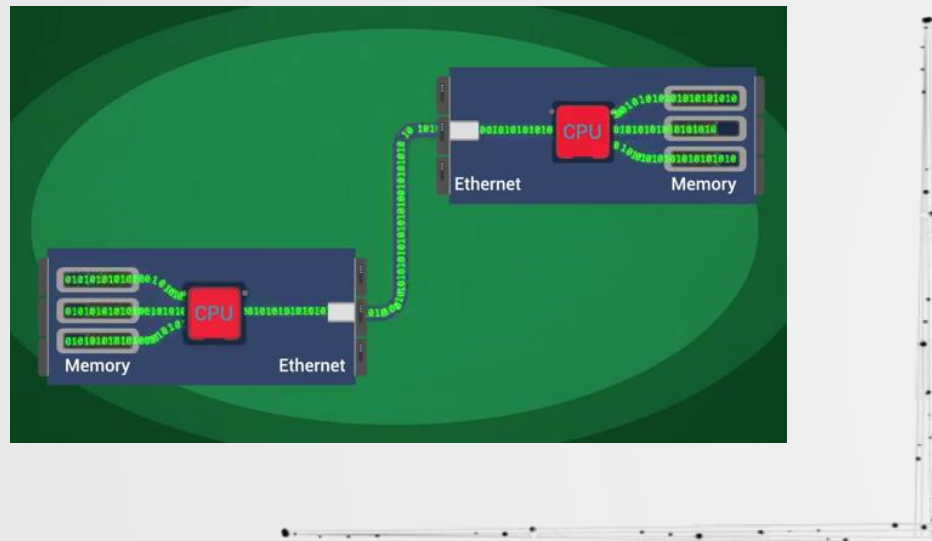


Accelerate!



RoCE = RDMA Over Converged Ethernet

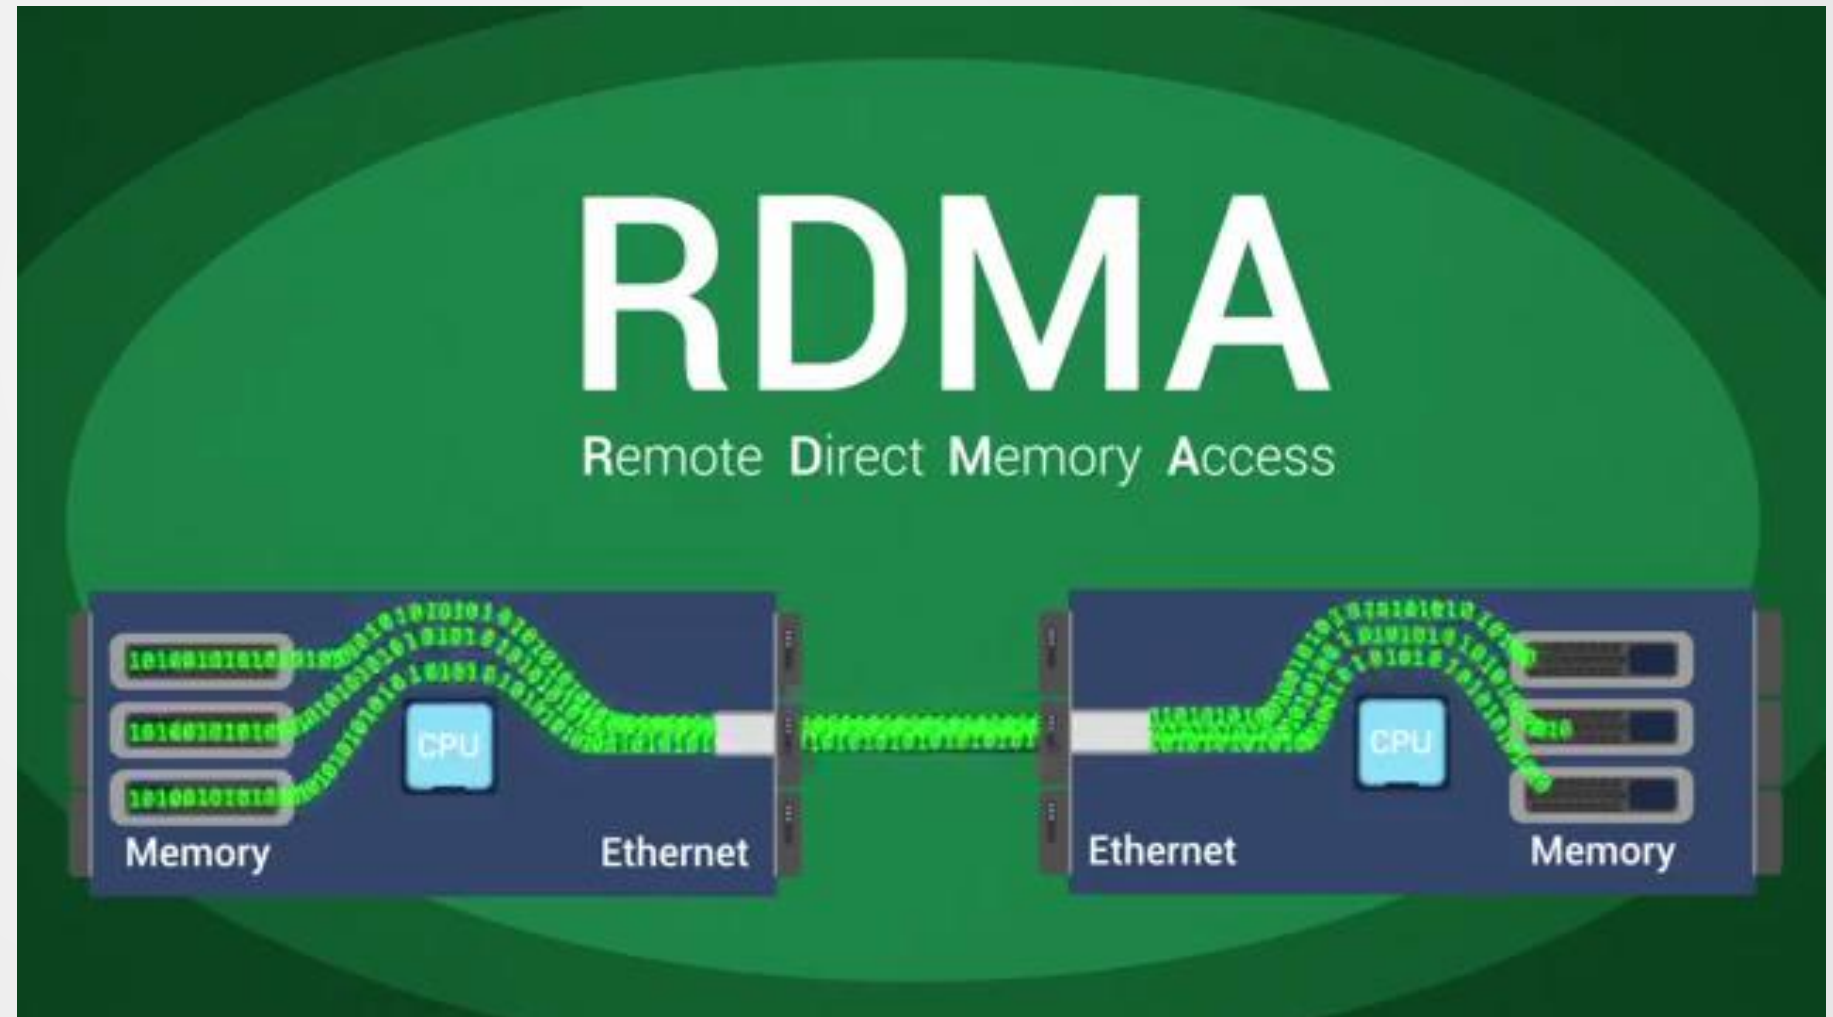
Bypass the CPU, transfer directly to memory



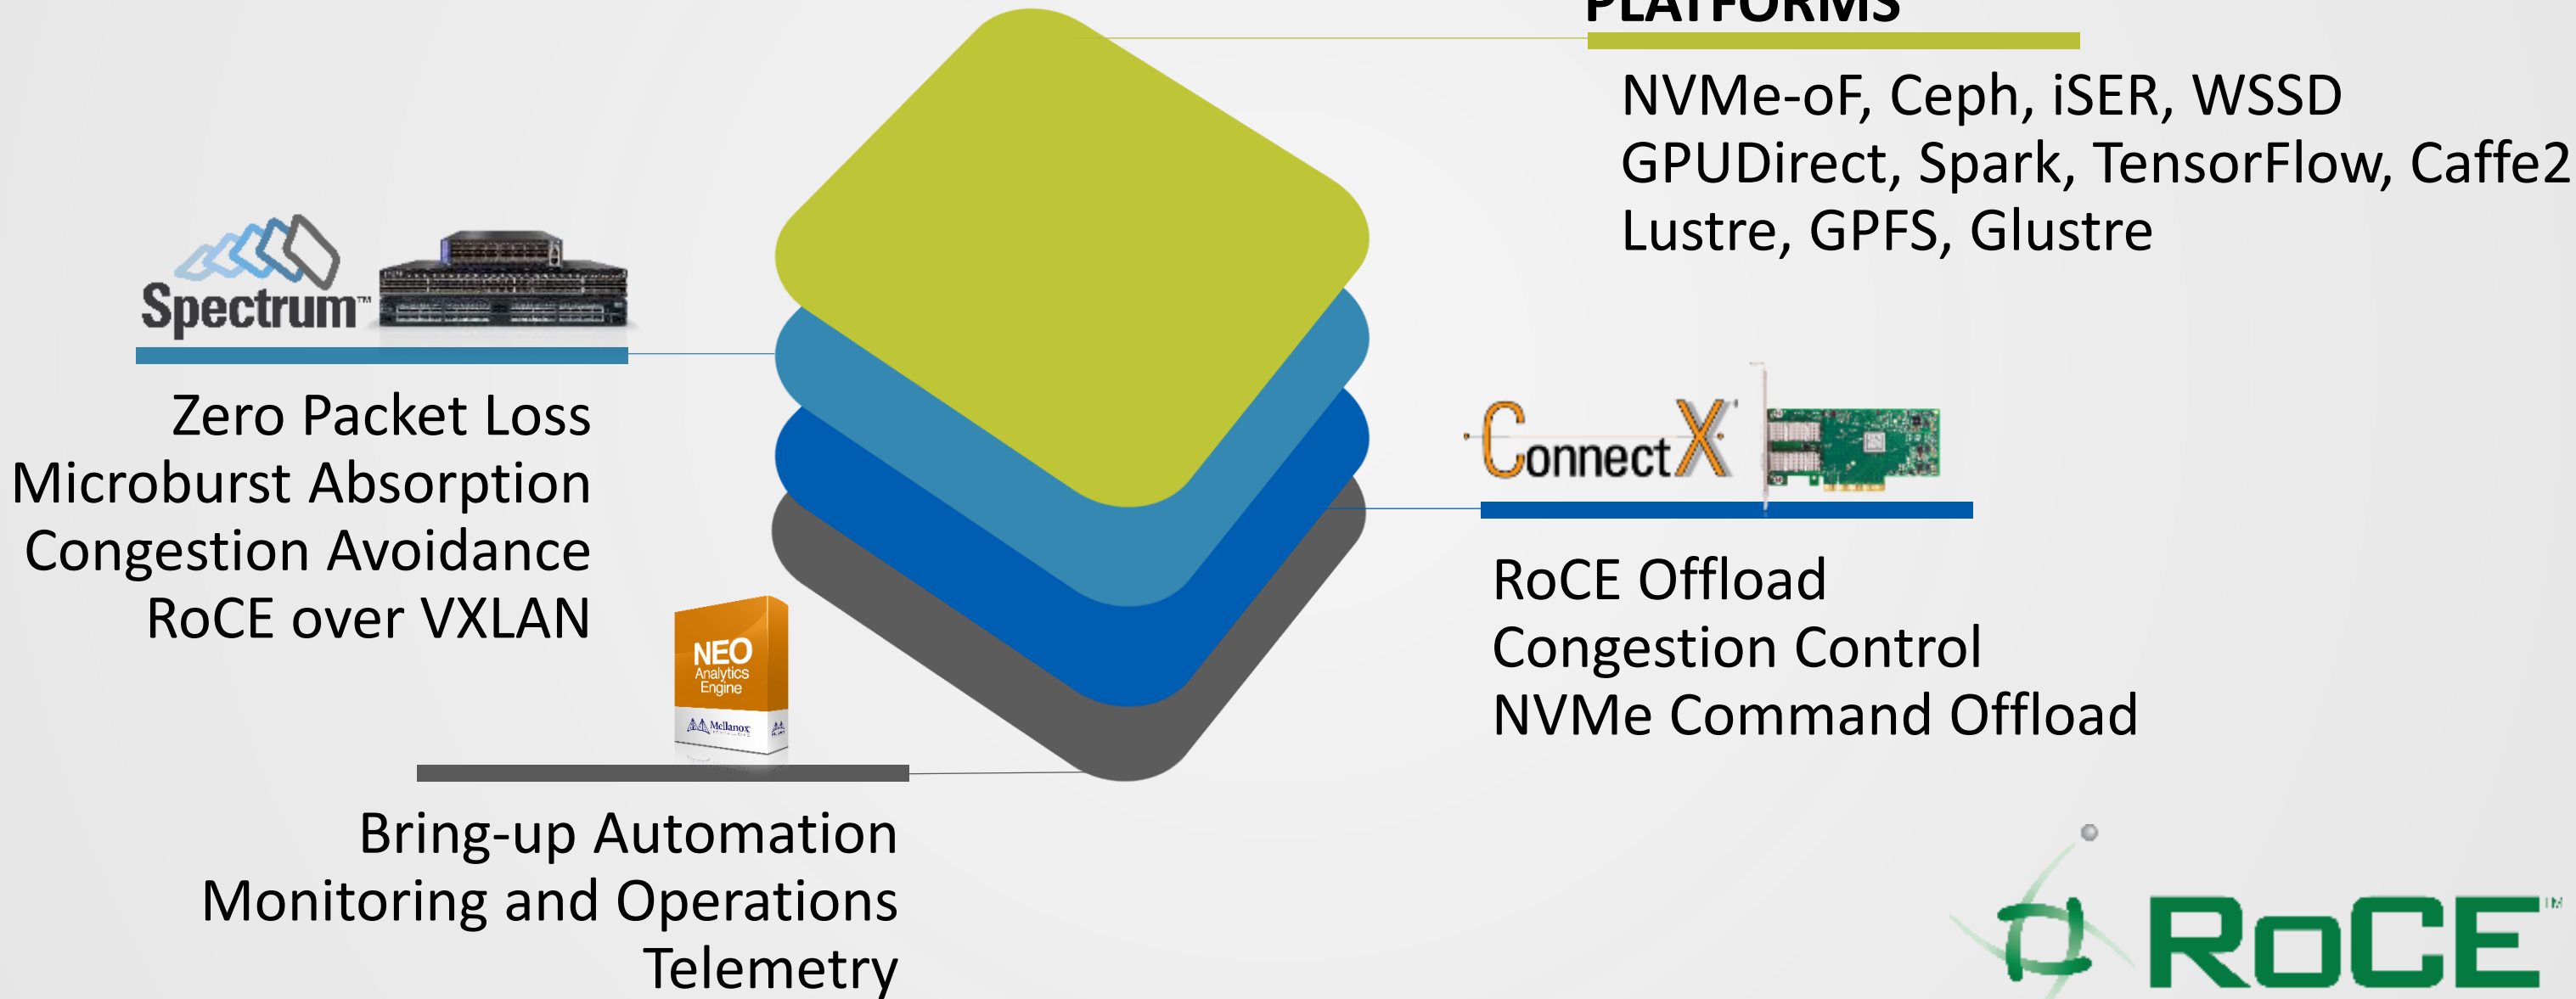
0 CPU Utilization

↑ Bandwidth


↓ Latency



The Novelty: End-to-End RoCE Acceleration



RDMA Supercharges Leading AI Frameworks





60%
Higher ROI

50%
Lower CapEx
& OpEx


RoCE By Default, No Extra Coding

Distributed Training Needs RDMA/RoCE



RDMA vs. TCP +50% Performance
Advanced RDMA vs. TCP +173% Performance

Scale Out!



Scaling Out Is A Balancing Act

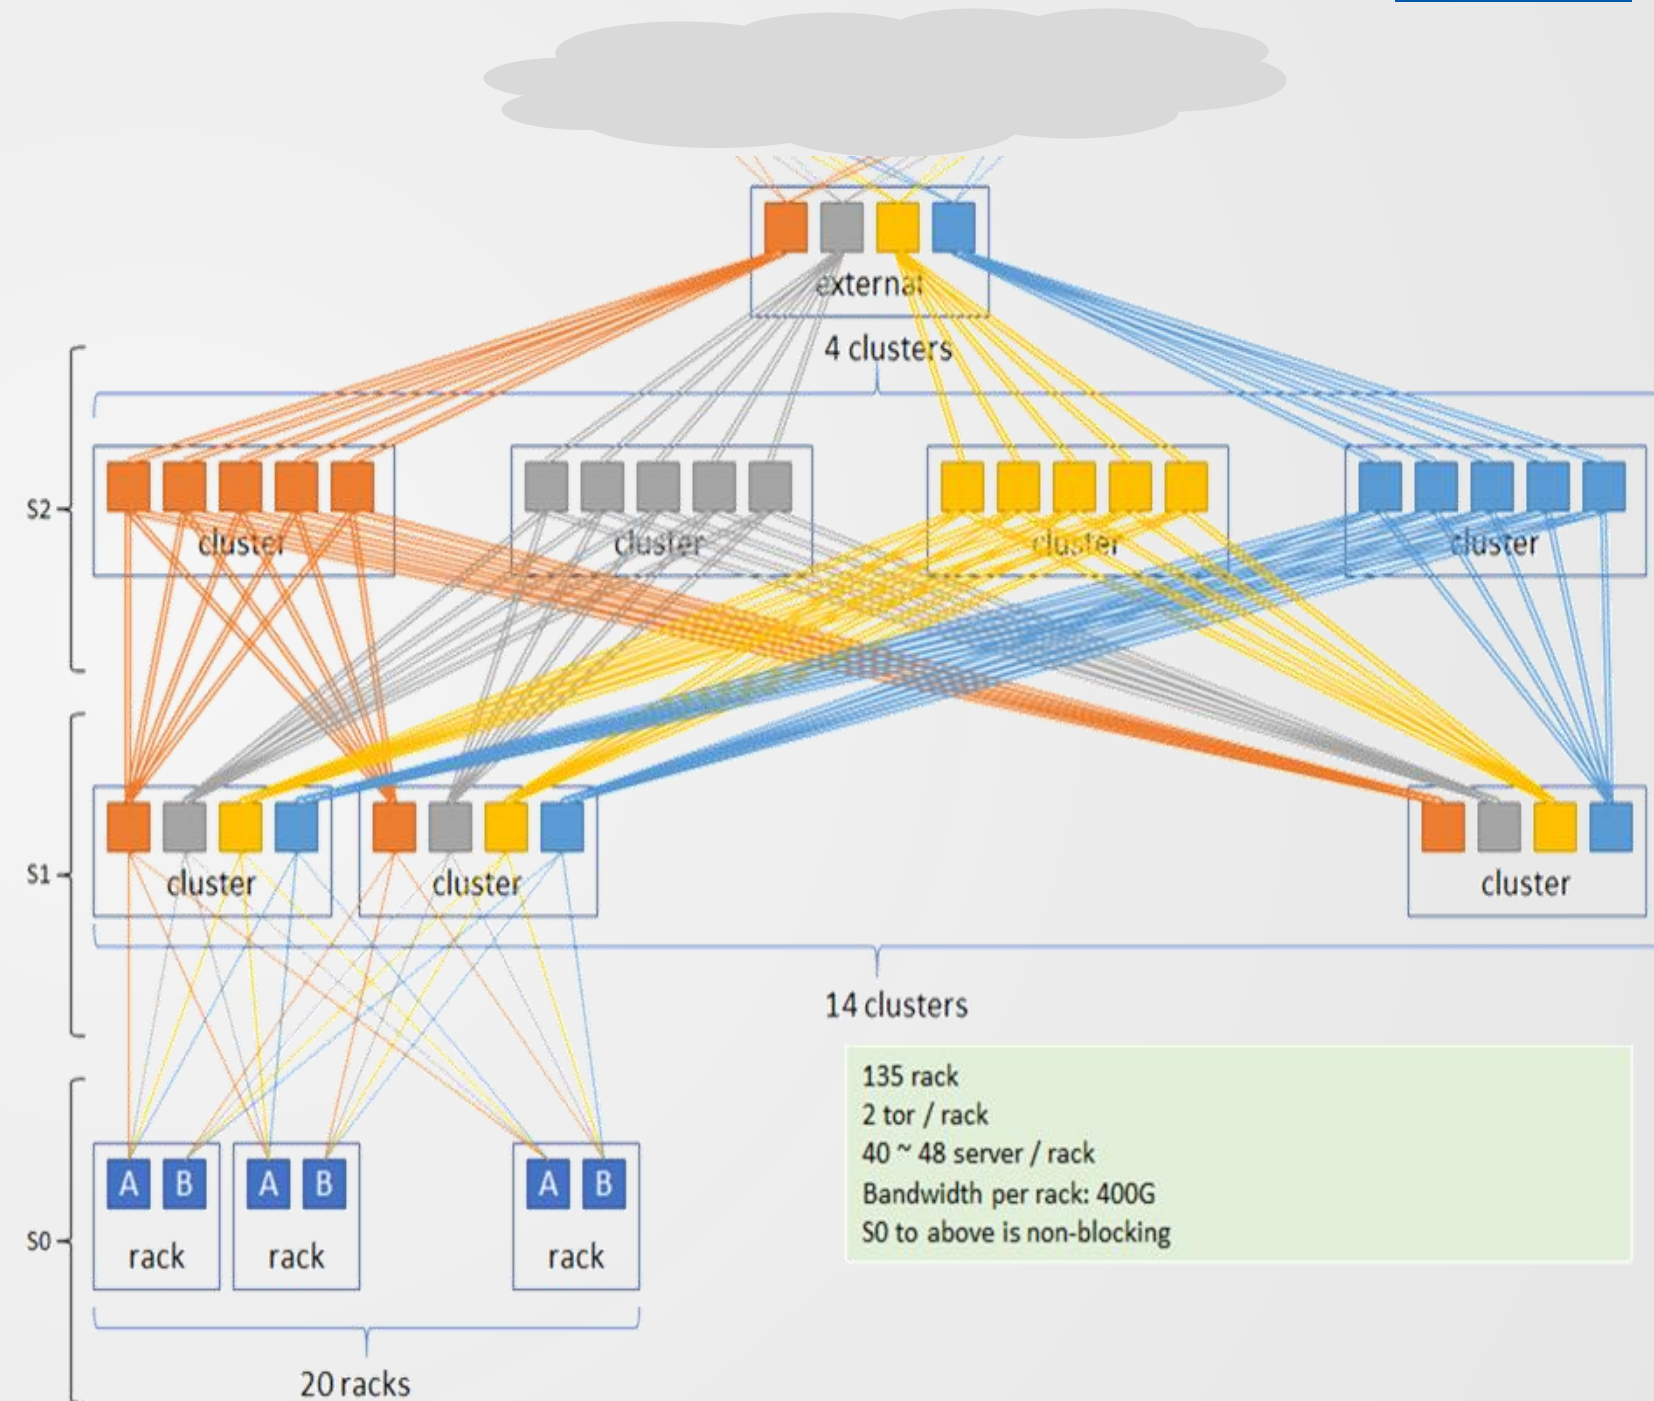
RoCE needs a Lossless Fabric
 Packet Drops = Lost Bandwidth
 Hardware Congestion Control (ECN)

Spectrum Advantages

- Mix of TCP and RoCE
- Fast Notification
- Flexible Buffers
- Semi-Lossless

Spectrum Prevents

- Congestion Spreading
- Deadlocks



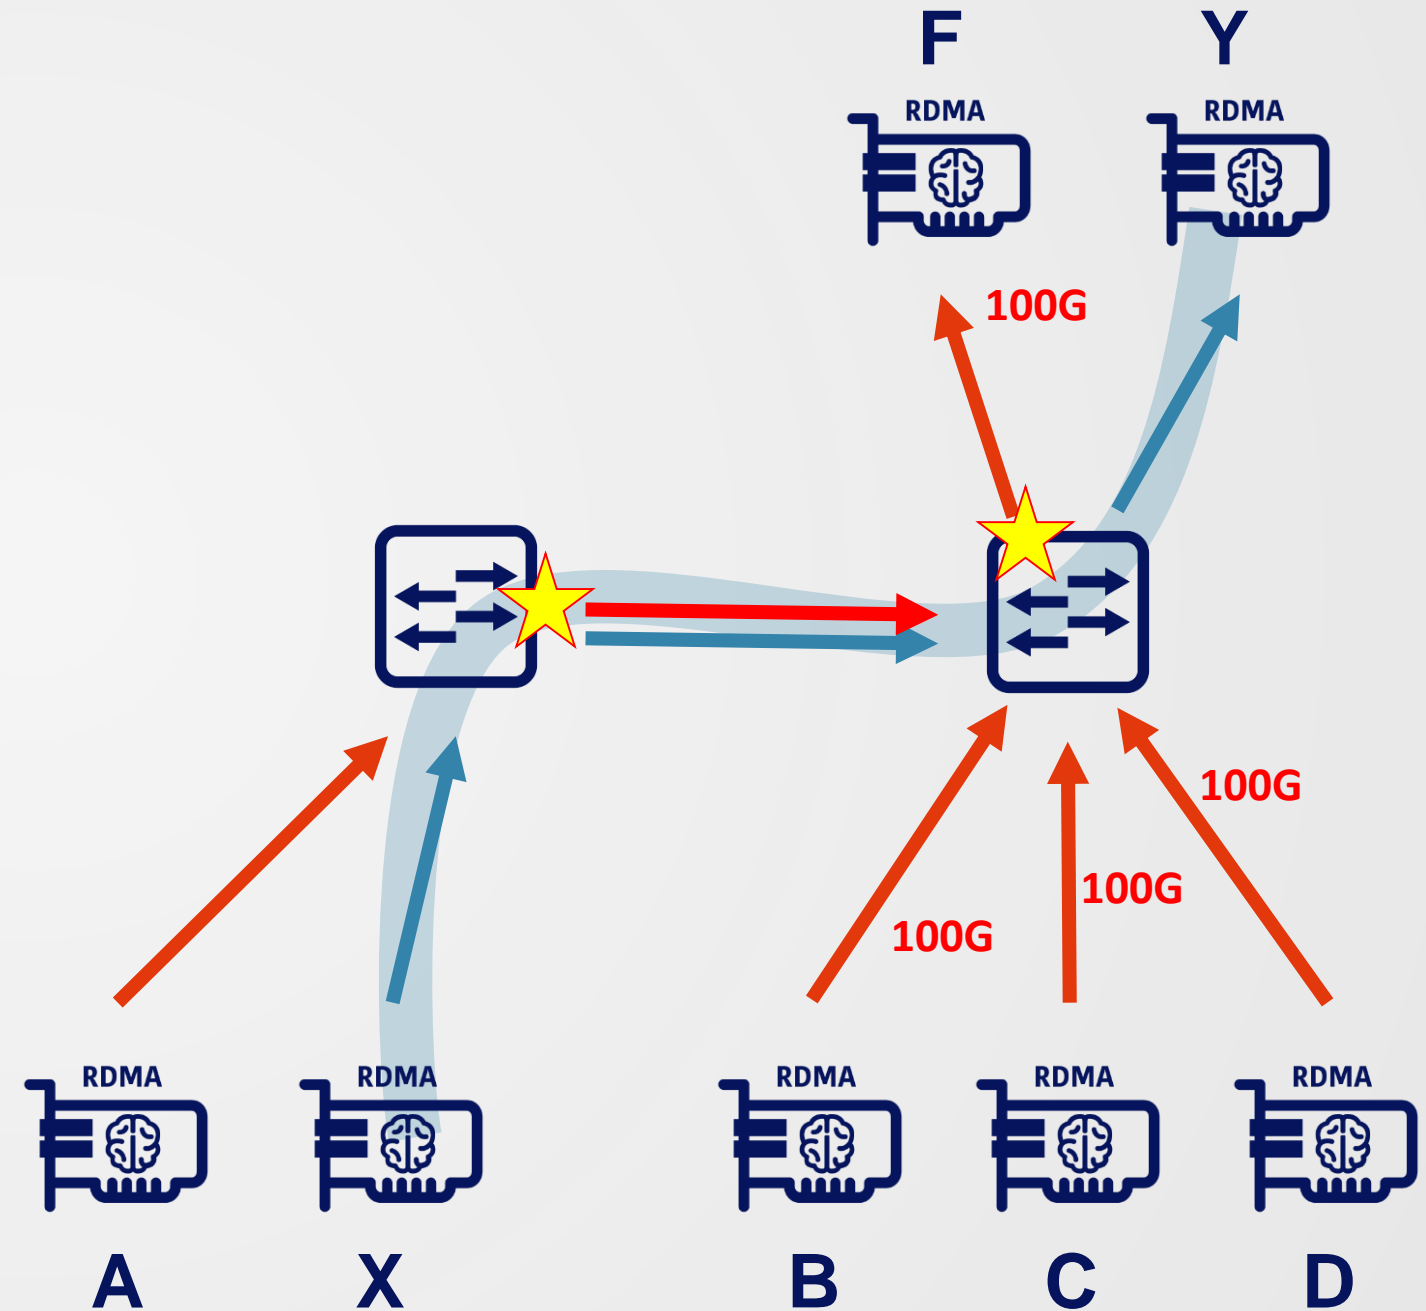
Scaling Lossless

Priority Flow Control Prevents Packet Drops Altogether

- Scaling PFC requires VLANs, QoS
- Cause Congestion Spreading

Explicit Congestion Control


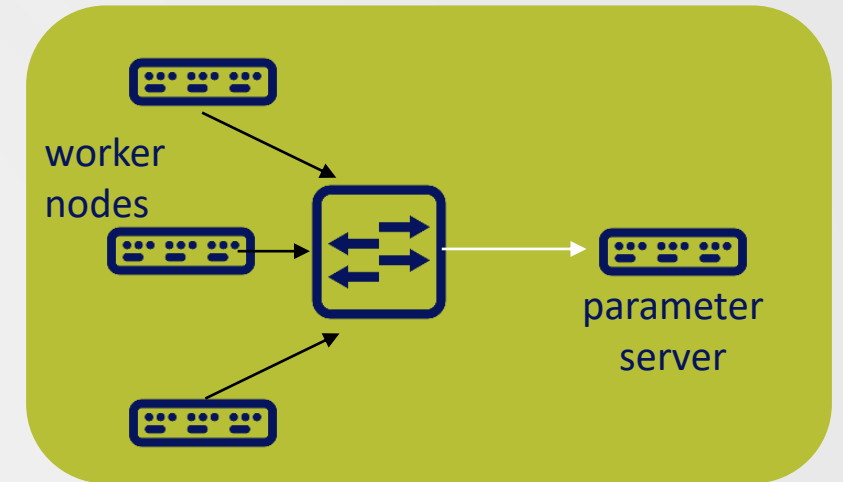
- ECN throttles A+B+C+D
- Victim traffic passes
- In heavy congestion: ECN+PFC as a backup




Spectrum Congestion Management For RoCE

**Optimal Congestion Management Avoids Pauses
Pauses Cause Bandwidth Degradation**

3 to 1 communication, ECN and PFC enabled




Competition



Multi-tenant Networking with EVPN

Primary Data Center




- Standards-based BGP EVPN
- Millions of VXLAN tunnels
- RoCE over VXLAN
- Secure between Tenants
- 100K Tunnels per switch

Remote Data Center


Massive Scale-out with RoCE

Smart!



NEO Management and Operations

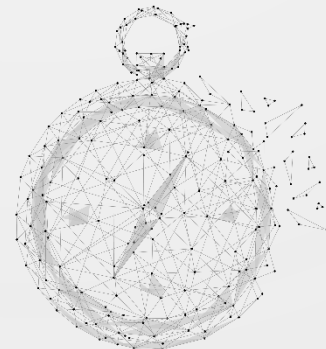
Day-1: Fabric Bring-up Made Easy



Day-2: End-to-end Network Operations



RoCE Diagnostics at Scale



All In One Place

Gain Visibility With Telemetry

In-ASIC Monitoring Tool


- No more black box guessing
- u-sec sampling generate histograms

Information per port/TC/LAG

- Packet Drop Counters
- Bandwidth (tx/rx) & queue depth
- Delays from flow control

Customize Actions

- Trigger events on thresholds
- Mirror and analyze packets discards
- Analyze start time and duration



Histograms – a Tool for Analysis

They are not just for shopping online



Summary

High Performing

Supports the concurrent requirements of AI

Accelerated

Hardware accelerates AI, Storage and Big Data

Scaled Out

Proven in scale, 1000's of nodes, up to millions VMs

Smart

Insight, telemetry, monitoring, auto-tuning, auto-provision

Mellanox Supercharges Leading AI Companies

facebook

Microsoft



nVIDIA

PayPal

Alibaba Group

Baidu

Tencent 腾讯

JD.COM
多·快·好·省



美团
meituan.com

SENSETIME
Focused on innovative computer vision
and deep learning technologies



Thank You

