ARISTA

Cloud and HPC/AI

Networking at scale (but not only for Cloud …)
Teratec 2024

Christophe Compain – Systems Engineering
May 2024

Arista At-a-Glance

- Launched in 2008 (IPO in 2014)
- 2023 facts
  - 4100+ employees
  - ~$5.86B
  - Fastest growing networking company
  - 9000+ Customers
  - 28+% Data Center Market Share*
    - #1 : 40% Port Share in 100G
    - #1 : 50% of 400G Share
- 92% Sat and high NPS Score
- Lowest CVEs in this industry
- French Subsidiary (since 2013)
  - 25+ employees

Market Sector Trends*

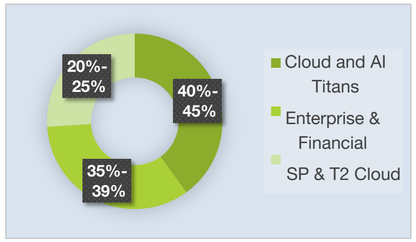

Product Trends*

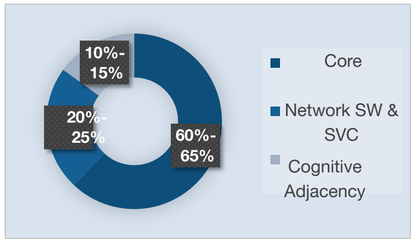

Core – Data Center & Cloud Networks
Cognitive Adjacencies – Campus & Routing
Cognitive Network - SW & Services

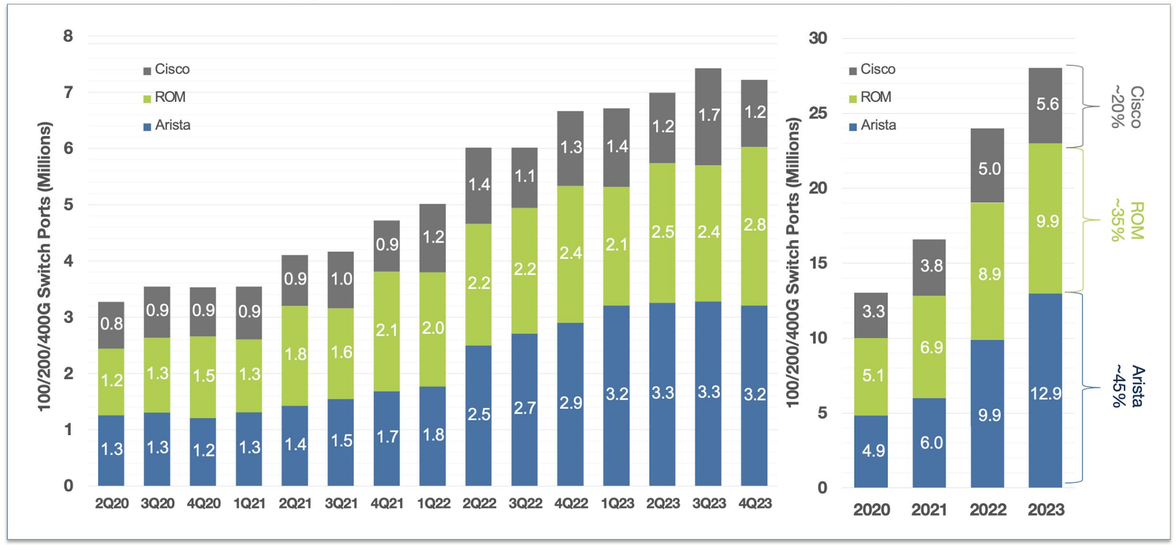

Source: Crehan Ethernet Switch Data Center Total Vendor Tables – 4Q'23

*Crehan Research: Q4 2023

Leading Disruption Across The Network
EOS Network Role Flexibility

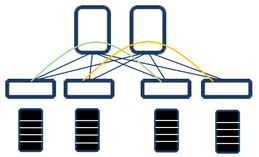

Resilient DC Leaf-Spine

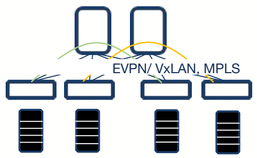

EVPN Services
VXLAN & MPLS

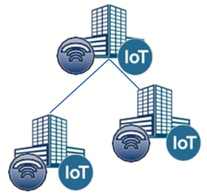

Campus Access & Spline

Data Center and Campus Evolution

DCI

DC Core Routing

CloudEOS

Data Center Core, Edge and Cloud Connectivity

IP Broadcast

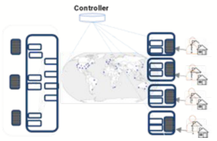

High Bandwidth
Content Peering

AI/HPC/IP Storage

Visibility and
Monitoring

Cloud Networking for Content, Peering, Big Data,
HPC, AI and Monitoring

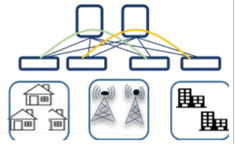

NFV Telco Cloud

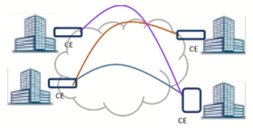

Carrier Edge
(L2VPN, L3VPN, MetroE,
Mobile)

Carrier Core
LDP, PWE, RSVP-TE

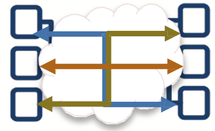

Carrier Core Evolution
Segment Routing

Carrier Core, Edge and Cloud Evolution

Arista Differentiation: Solve Real World Challenges

ARISTA

Hyperscalers infrastructure for HPC/AI

POSTED ON MARCH 12, 2024 TO AI RESEARCH, DATA CENTER ENGINEERING, ML APPLICATIONS

Building Meta's GenAI Infrastructure

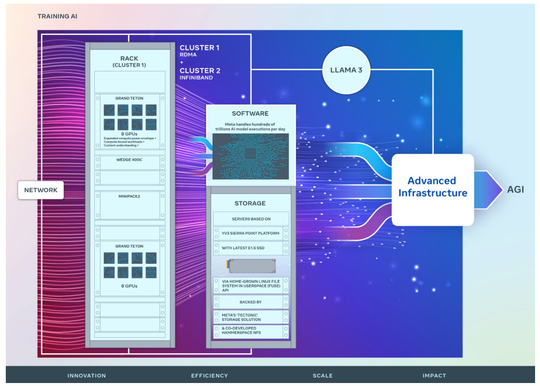

By the end of 2024, we're aiming to continue to grow our infrastructure build-out that will include 350,000 NVIDIA H100 GPUs as part of a portfolio that will feature compute power equivalent to nearly 600,000 H100s.

https://engineering.fb.com/2024/03/12/data-center-engineering/building-metas-genai-infrastructure/

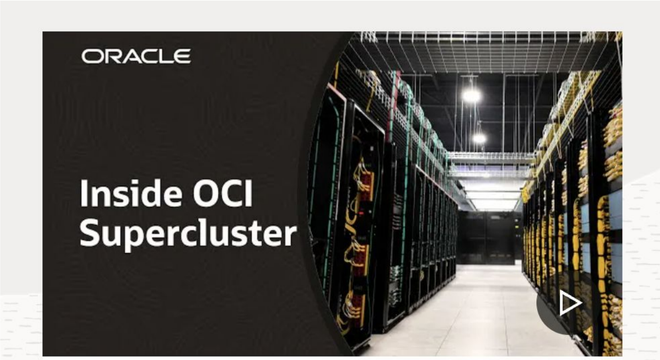

Run the most demanding AI workloads faster, including generative AI, computer vision, and predictive analytics, anywhere in our distributed cloud. Get the latest GPU compute, scaling up to the 32,768 GPU Oracle Cloud Infrastructure (OCI) Supercluster.

AWS Teases 65 Exaflop 'Ultra-Cluster' with Nvidia, Launches New Chips

Empowering Azure Storage with RDMA : Today, around 70% of traffic in Azure is RDMA and intra-region RDMA is supported in all Azure public regions.
https://www.microsoft.com/en-us/research/uploads/prod/2023/03/RDMA_Experience_Paper_TR-1.pdf
Nvidia launches AI foundry on Microsoft Azure
https://www.sdxcentral.com/articles/news/nvidia-launches-ai-foundry-on-microsoft-azure/2023/11/

Google Cloud claims 'most powerful' publicly available machine learning cluster

Compute

Announcing Trillium, the sixth generation of Google Cloud TPU

May 15, 2024

Behind Omniva: The secretive GPU cloud startup that tried to build the world's largest crypto data center

Tesla Unveils Top AV Training Supercomputer Powered by NVIDIA
The cluster uses 720 nodes of 8x NVIDIA A100 Tensor Core GPUs (5,760 GPUs total) to achieve an industry-leading 1.8 exaflops of performance.

How to drive AI/HPC hyperscaling from the networking point of view?

Applications

Semi conductors

Protocols

Architecture

How to drive AI/HPC hyperscaling from the networking point of view?

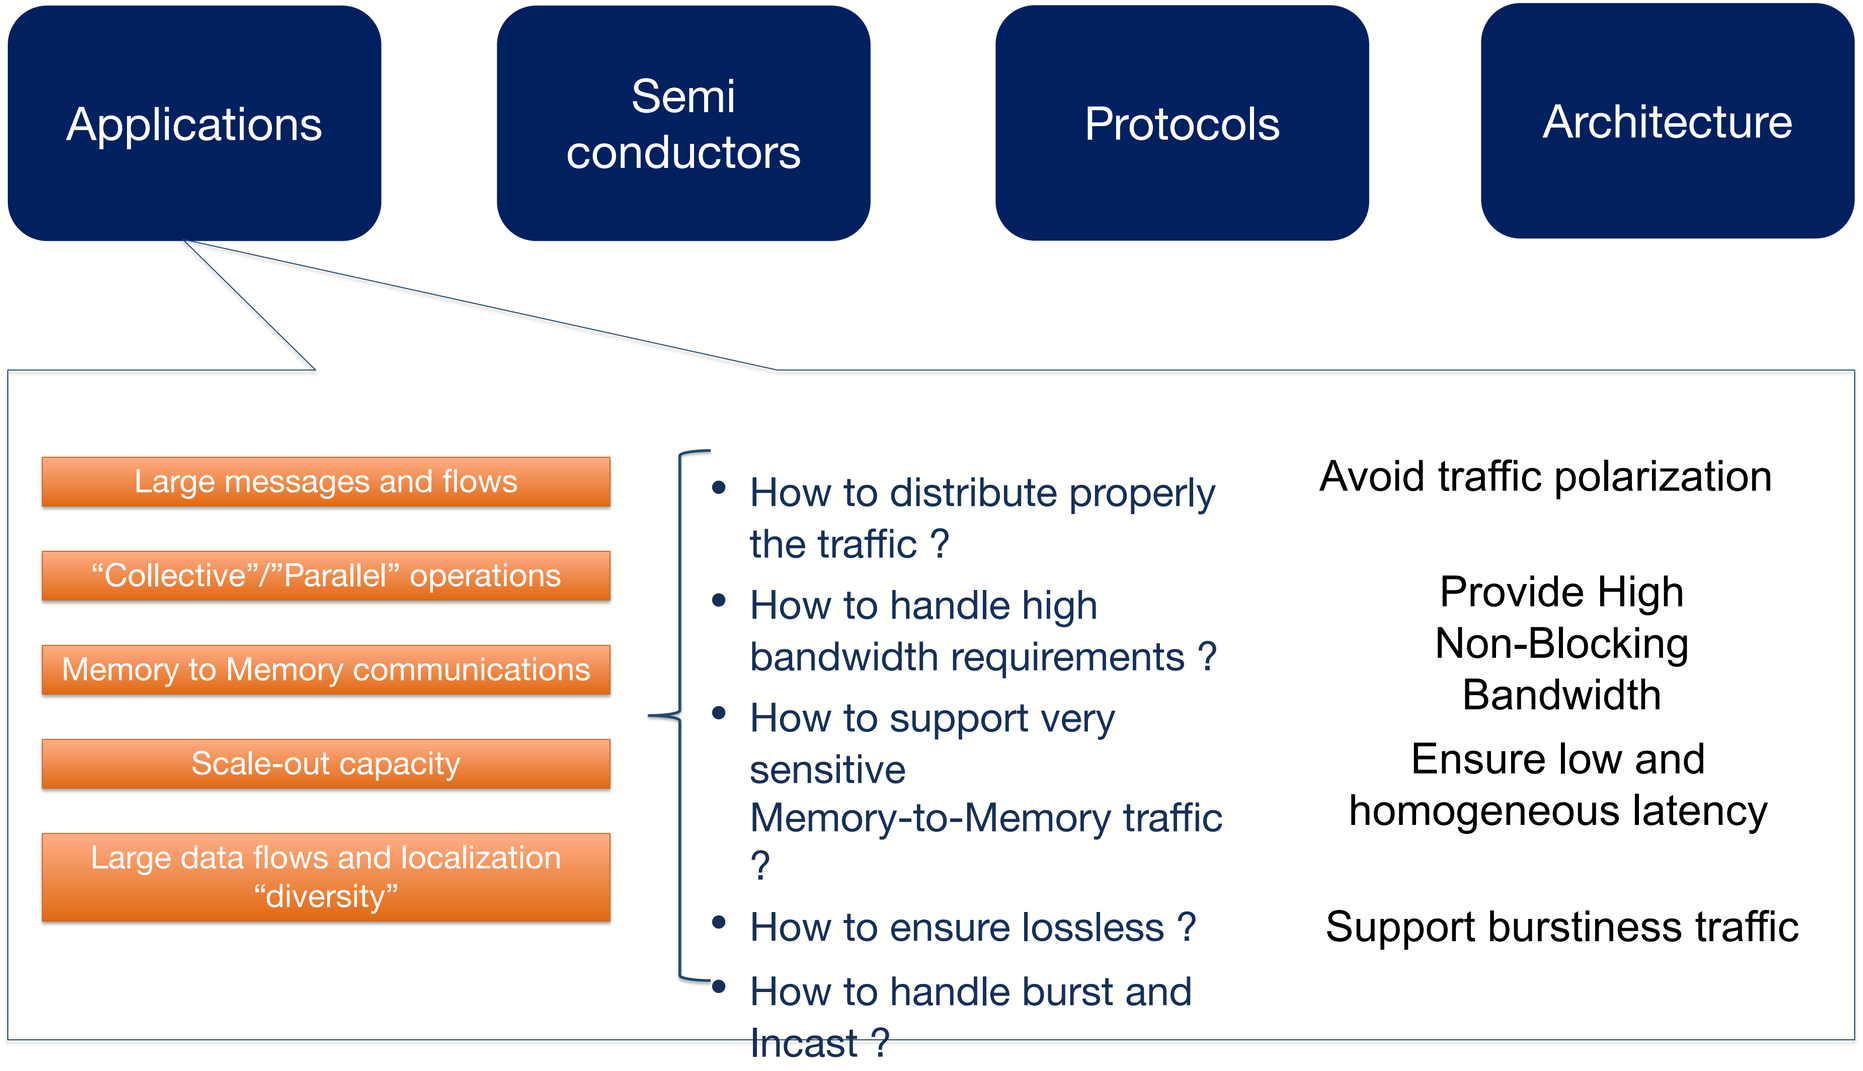

How to drive AI/HPC hyperscaling from the networking point of view?

Applications

Semi conductors

Protocols

Architecture

Ethernet and IP

5nm: -x2 density

HW Accelerator

>50% power reduction with LPO

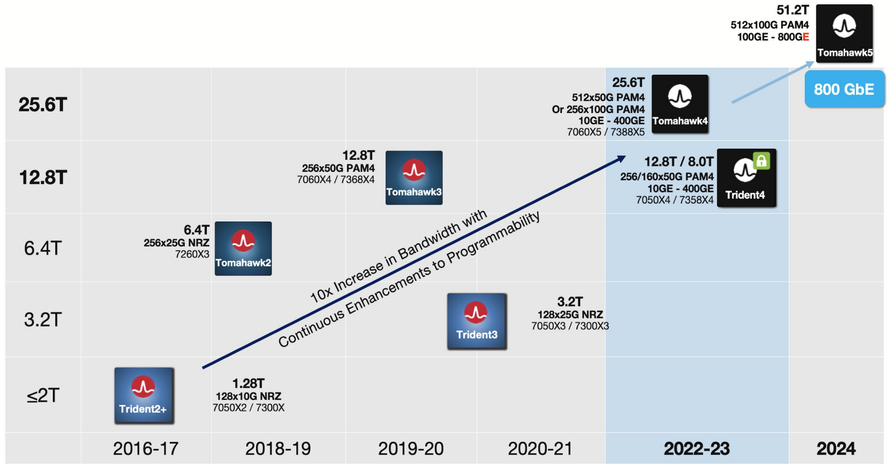

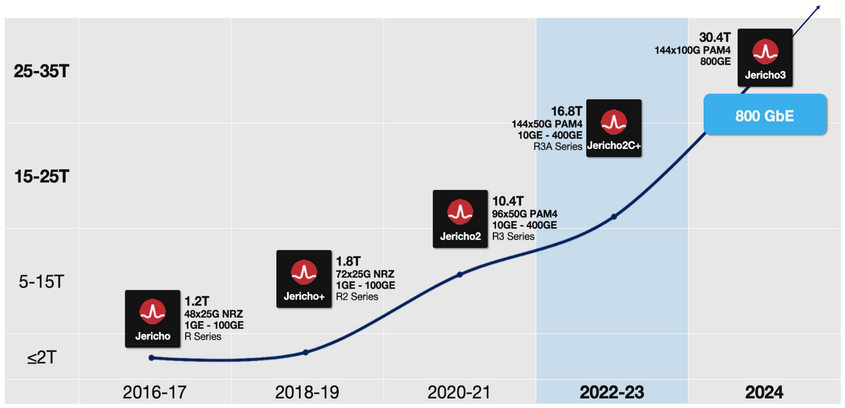

ARISTA

How to drive AI/HPC hyperscaling from the networking point of view?

Applications

Semi conductors

Protocols

Architecture

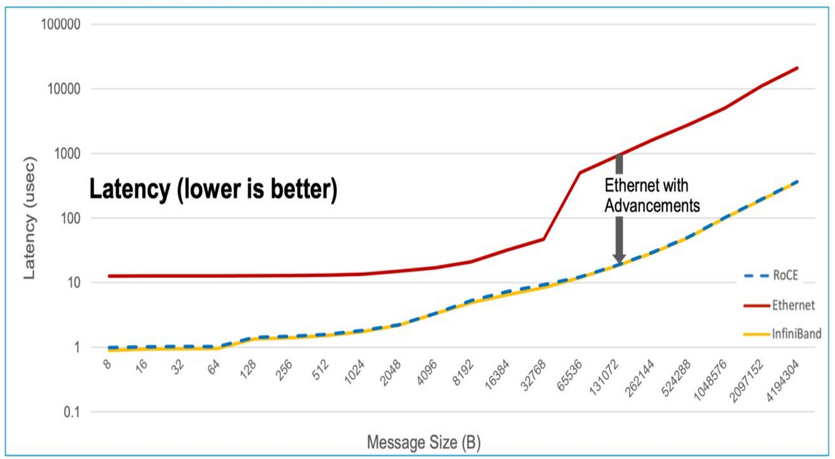

High performance and standardized transport

- IP Transport (v4 or v6)
- RDMA over Converged Ethernet
- Smart load balancing
- Back-pressure mechanisms
- Fast convergence
- Hitless upgrade
- Real-time telemetry

Looking forward : Ultra Ethernet Consortium

- Formed in the summer of 2023, the UEC aims to develop a new standard for interconnection for AI and HPC datacenters needs
  - 55+ members
- Next generation HPC/AI transport
  - Multi-Pathing and Packet Spraying
  - Flexible Ordering
  - HPC/AI optimized congestion control
  - End-2-End telemetry
  - Security

Steering Members

AMD ARISTA BROADCOM CISCO EVIDEN an atos business

Hewlett Packard Enterprise intel Meta Microsoft ORACLE

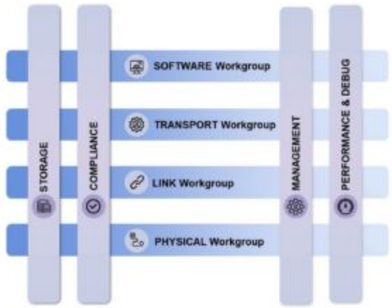

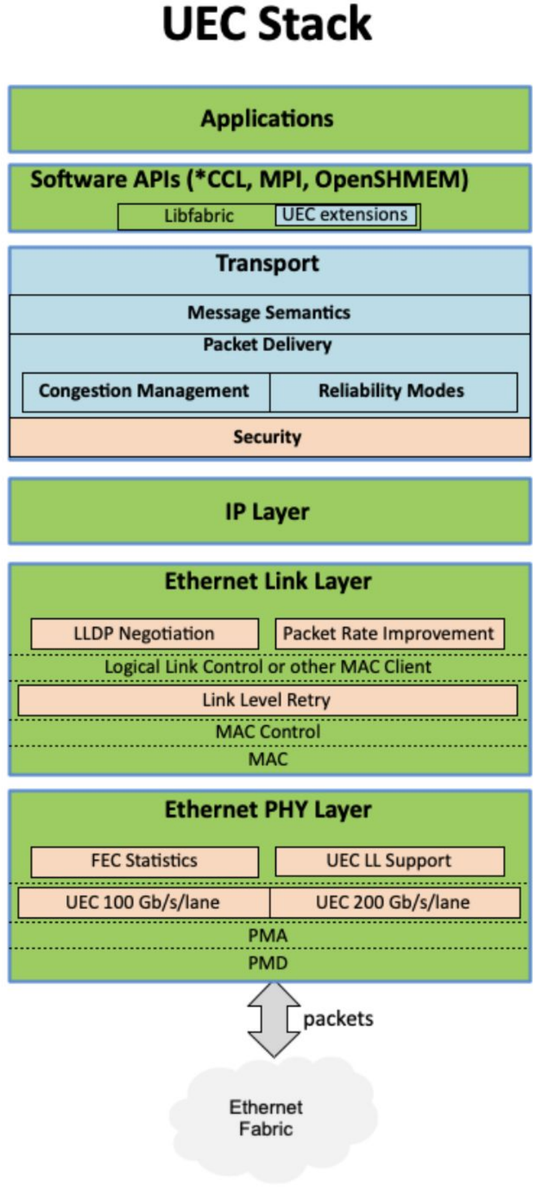

ARISTA

How to drive AI/HPC hyperscaling from the networking point of view?

Applications

Semi conductors

Protocols

Architecture

High performance and standardized transport

- Dedicated network (back-end)
- Fat tree topology
  - Small / Medium / High-Scale
- 400G and 800G and no oversubscription
- Consistent standards based Open Ethernet architecture

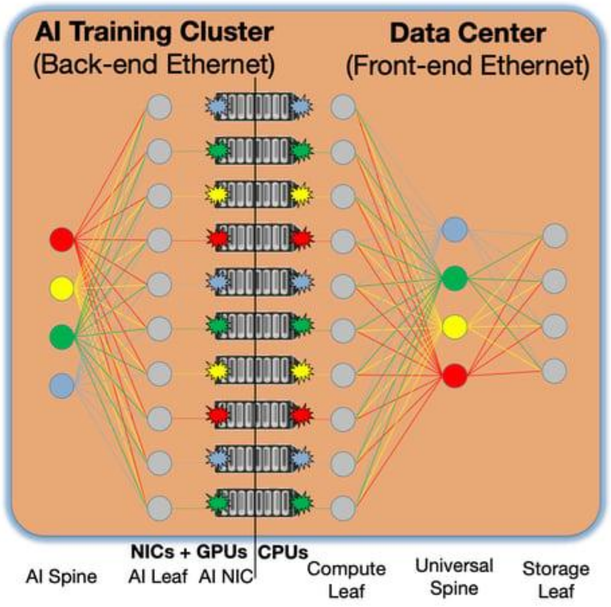

Data localization is key
How to move data securely and properly ?

- Data gravity/inertia challenges:
  - Move “temporarily” a *relevant* amount of data with a limited bandwidth
  - Move “increments” of data every days/hours/minutes or second
- Data transport challenges:
  - Congestion control
  - IPSec and FW

How to move data around securely and properly ?

1 Direct/Private Peering

- Enterprise/Large Enterprise class of solution
- Congestion control : it depends …
- Colocation model thru peering facilities
- Secure tunneling technology is optional
- Best expected quality of service : bandwidth/latency/jitter.

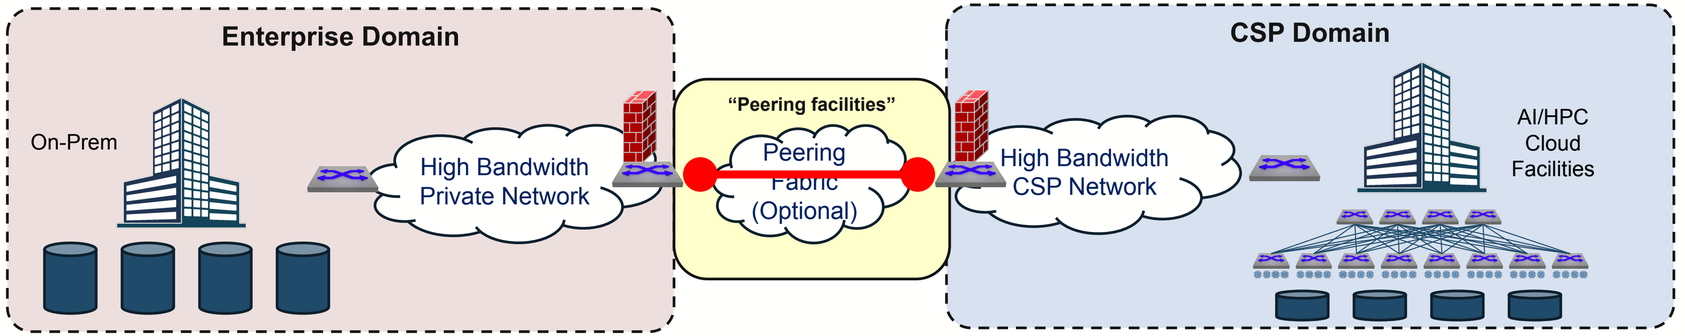

2 Internet As a Transport

- Secure tunneling technology is mandatory: IPSec is the standard
- Congestion control : TCP based
- Quality of service is relaying on internet : jitter could be very bad, fragmentation issue

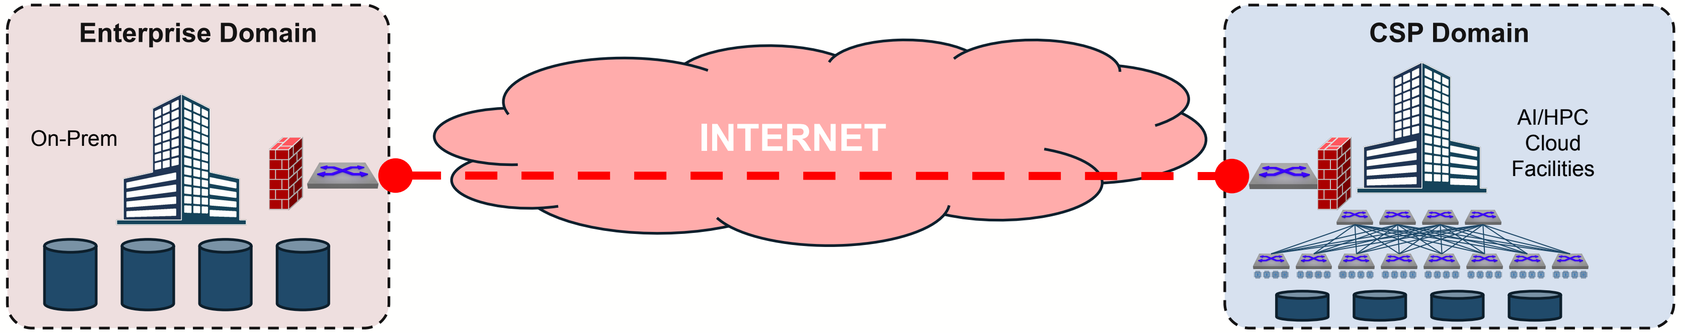

3 Physical Transport

- Where (1) or (2) aren't possible
- Need extraction and import mechanism

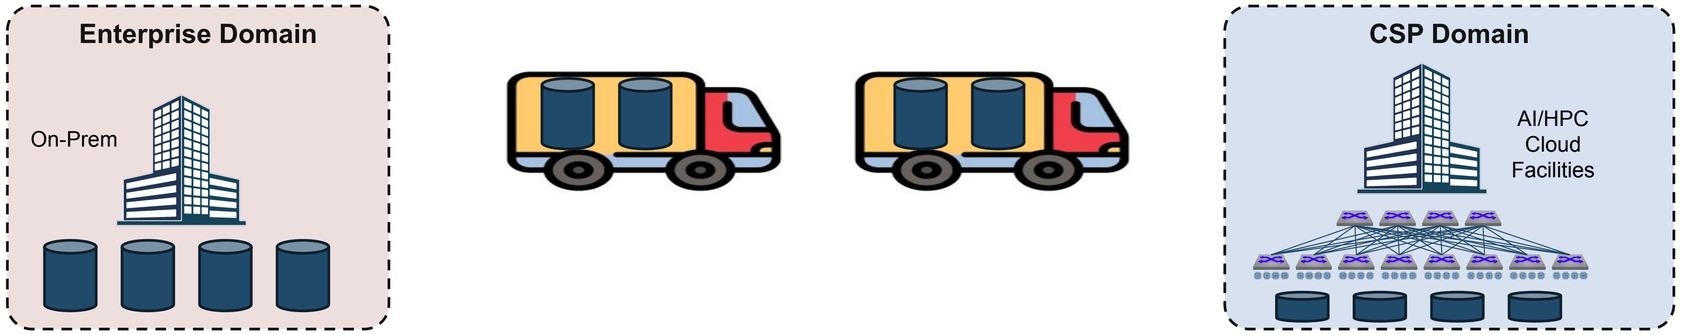

I’m not an hyperscaler … Can I run HPC/AI networks ?

OF COURSE YOU CAN

Simple design

IP Standards

Simplified Operations

https://solutions.arista.com/cea-dam-success-story

ARISTA

AI ready networks

https://www.arista.com/en/solutions/ai-networking

Thank You

arista.com