

# Dell Networking

**Fabrics Ethernet para  
Inteligencia Artificial  
Generativa con SONIC**

**Gonzalo de Antonio**

*Dell Technologies Networking Sales Executive  
Western Europe South*





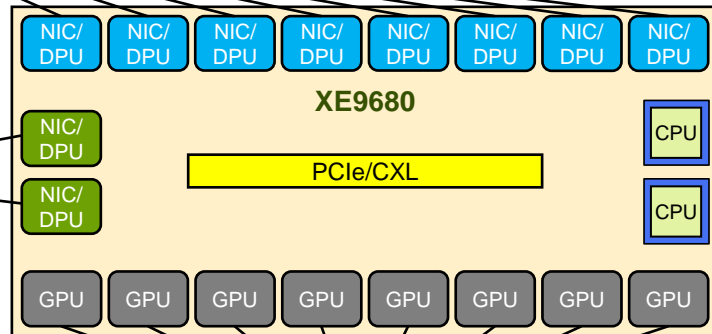



# AI workload types and design considerations

	Building the AI model		Using the AI model
Workload	<p><b><u>Training</u></b></p> <ul style="list-style-type: none"><li>• Massive consumption of data to create a foundation of general knowledge</li><li>• Processing can take days or weeks of elapsed time</li><li>• Often the realm of “AI factories” or service providers given the size and scale of the infrastructure requirements</li></ul>	<p><b><u>Fine Tuning</u></b></p> <ul style="list-style-type: none"><li>• Refines base models with specialized knowledge for specific applications</li><li>• Often run on-prem to ensure data security and protection of internal IP</li><li>• Smaller clusters are in reach of enterprise customers</li></ul>	<p><b><u>Inferencing</u></b></p> <ul style="list-style-type: none"><li>• Leveraging the models for business applications</li><li>• Often runs on-prem given the criticality of the information and results to a business</li><li>• Must deliver results in a conversational flow, requiring excellent system responsiveness</li></ul>
Requirements	<p>Highest Scale, thousands of GPU's</p> <ul style="list-style-type: none"><li>• Bursty Elephant Flows running iteratively over many thousands of cycles</li><li>• Any congestion or delay will drive extended tail latencies – dramatically extending job completion time</li><li>• <b>Fabric needs highly optimized topologies and load balancing strategies</b></li></ul>	<p>Lower Scale, dozens to hundreds of GPU's</p> <ul style="list-style-type: none"><li>• Derived from a general base model as a starting point</li><li>• Similar workload pattern to training, but smaller data sets and fewer iterations required</li><li>• <b>Less sensitive to BW / Latency in relation to overall job completion time</b></li></ul>	<p>Lowest Scale, can be CPU or GPU</p> <ul style="list-style-type: none"><li>• Least demanding load on the fabric</li><li>• Ability to serve requests and provide rapid response</li><li>• <b>Most inferencing environments can execute in a single host</b></li></ul>

# AI Fabrics Model


Other names used in the industry

- Host Fabric: { Front-End Fabric
- GPU Fabric: { Back-End Fabric



# How Networking in Gen IA is different ?


## DC "Classical" Network



iDRAC

- ❑ 2 Network Fabric @25G
- ❑ Design with oversubscription

## Gen IA Network



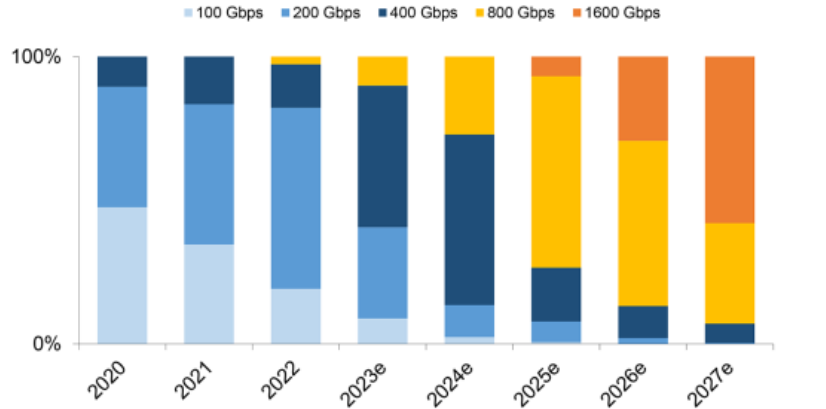
- ❑ 3 Network Fabric @400G & 100G
- ❑ Design without oversubscription

# Gen AI Fabrics - Market Insights – TAM

## Data Center Ethernet Switch Projections

- Addressable TAM is more than **\$100 B** in Cumulative Ethernet Data Center Switch Sales **Over Next Five Years**<sup>1</sup>
- **AI Fabric** (backend switching for GPU-to-GPU connectivity) forecasted to grow from **\$1.2B (2022) to \$15.2B (2027)**, five-year CAGR of 65%<sup>2</sup>
- 20 Percent of Ethernet Data Center Switch Ports Will Connect to AI Servers By 2027

## Migration to High-Speeds in AI Clusters (AI Back-End Networks)



\*Includes both Ethernet and InfiniBand

\* Source: Dell'Oro Group AI Networks Report December 2023



## Back-end AI Network Projections

- InfiniBand and Ethernet will coexist
- Expect Ethernet to gain share driven by
  - Improvement on Ethernet
  - Customers desire to multi-source
- Mix of Ethernet vs InfiniBand will depend on
  - Customers distinct choices
  - Size of clusters
  - Nature of workloads



# Ultra Ethernet Consortium

- Formed under the Linux Foundation July 2023
  - based on a private consortium created previously
- <https://ultraethernet.org/>



Michael Dell  
@MichaelDell

I am excited about collaborating with the industry to drive innovation in high-speed, high-capacity Ethernet fabrics for #AI at scale. 🚀🔗



5:22 AM · Nov 15, 2023 · 14.3K Views

DELLTechnologies

## Working Groups



### Physical Layer

Develop specifications that enhance the performance, latency and management of the Ethernet physical layer, the underlying medium and physical layer clients (link layer).



### Link Layer

Develop specifications that enhance the performance, latency, and management of the Ethernet link layer.



### Transport Layer

Develop specifications for an AI/HPC transport with enhanced throughput, latency, scalability, and management for Ethernet networks.



### Software Layer

Develop specifications and / or software API and / or open-source code for various AI/HPC use cases or applications. Scope includes but is not limited to: remote memory access optimizations, In Network Computations (INC) as well as security, management and storage.

## UEC Public Statements Summary

*“The UEC specifications offer significant improvements by addressing the following critical networking requirements for the next generation of applications”:*

- Average message completion time & “Tail latency” should both be minimized.
- Maximize concurrency
- fully utilizing the AI network
- Avoid causing congestion or head-of-line blocking or final hop Incast
- Minimize use of centralized traffic engineering or algorithm specific tuning

*By using...*

- Multi-pathing and packet spraying
- Flexible delivery order
- Modern congestion control mechanisms
- Efficient and fair rate control algorithms
- End-to-end telemetry
- Larger scale, stability, and reliability (is  $10^6$  endpoints too many? ☺)
- Physical layer of 800G, 1.6T, and faster Ethernet speeds
- End to end security

# Ethernet evolving to be the preferred choice for backend AI fabrics

- Market inflection points for Ethernets powered by AI fabrics
  - **Availability** of High Radix switching with next-Gen silicon technologies – 64x400G (25.6T), 64x800G(51.2T), 102.4T...
  - Improved **congestion monitoring, flow control, and Transport (RoCEv2)** protocol availability in NOS
  - Community effort to drive Ethernet Standards – **Ultra Ethernet Consortium**
  - Desire for **no-vendor lock-in** infrastructures
  - **Silicon and supplier diversity**
  - Lower **Total Cost of Ownership (~3x lower)**
  - **Latency improvements** with next Gen Silicon **from 800ns to 200ns** and **Cut-Through** switching

# AI fabric

Bringing it ALL together

## Back-End (GPU Fabric)

has most demanding requirements for raw performance, lossless attributes and lowest latency

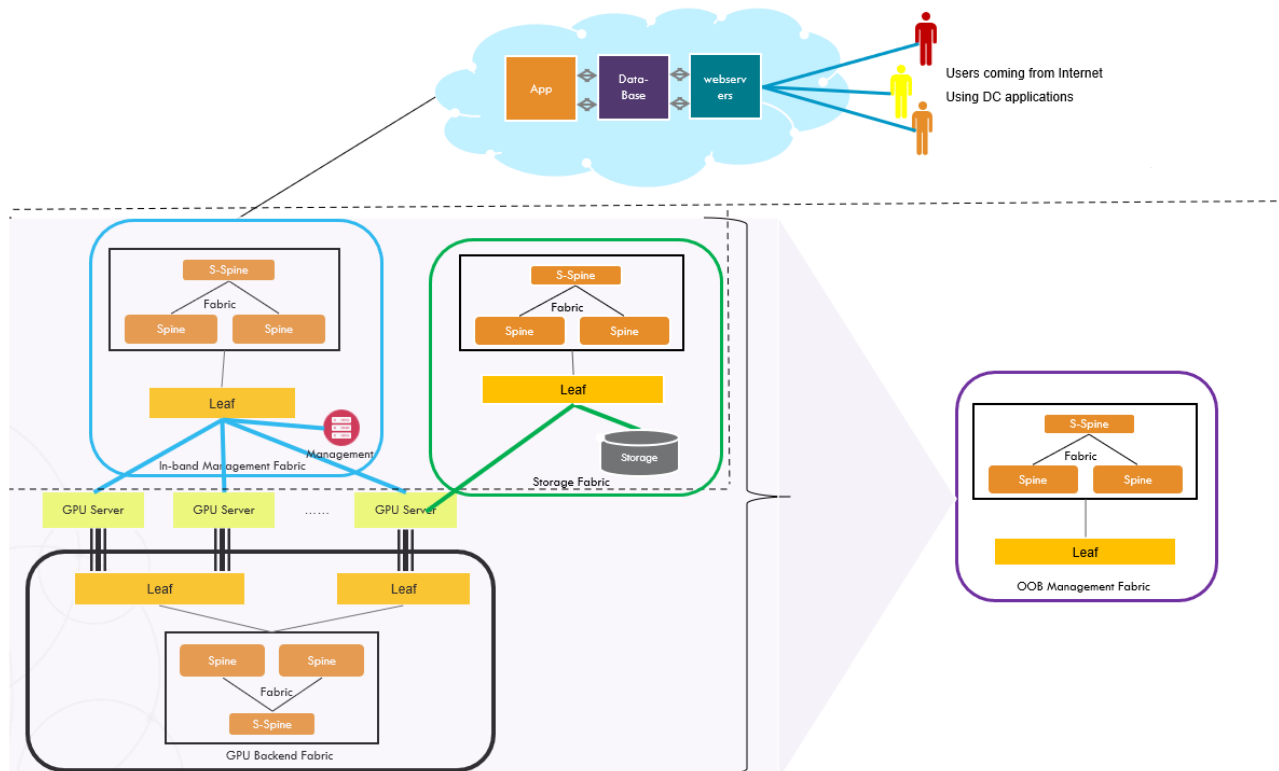
## Front-End fabrics

support application traffic, storage access and connection to the general network

**OOB Mgmt Network** for administration and fabric management


# Dell Networking

## Delivering Ethernet Solutions across all use cases within AI Fabrics





LEAF SPINE



# Dell Enterprise SONiC and GenAI use case



Connection to storage and application traffic

OOB Management

Created by: [/magdiel.ribeiro](#)

# Dell Technologies World 20-23 May

## Dell Networking for AI Solutions

High-bandwidth, low-latency advanced Ethernet fabrics



### Z9864F-ON

2U form factor for twice the performance of the current generation PowerSwitch



### Enterprise SONiC 4.4 with SmartFabric Manager

Simplified management, validated blueprints and better resiliency with advanced AI Ethernet features



### Broadcom Thor 2 NIC

2x End-to-End Bandwidth  
400 Gbps bi-di throughput  
Low Latency

High performance Ethernet for modern AI workloads

Adaptive Routing and RoCE congestion control

Detailed performance insights with advanced analytics and monitoring

## Next-gen Ethernet fabrics for modern workloads like GenAI

2X  
density of 400GbE  
switching ports<sup>1</sup>

64  
800GbE ports

8K  
supported GPU nodes<sup>2</sup>

<sup>1</sup> vs. the previous generation 400GbE switch, the Z9864F-ON  
<sup>2</sup> Two-socket

# SONiC

Open Source Networking  
Made Real



**THIS IS THE WAY**




## EVOLUTION OF NETWORKING

A move from closed & proprietary  
to flexible & open


# Open-Source Networking

Traditional  
Networking




Before  
2014

Open Networking



2014 to  
2020

Open Source  
Networking



2020



# Software for Open Networking in the Cloud



An open-source  
software project  
under Linux  
Foundation

Seed code  
provided  
by MSFT

Currently  
running on  
45,000+ Azure  
switches

An  
automated, intent-  
based, API-  
centric, and  
purpose-built  
containerized  
network

“By 2025, 40% of organizations that operate large data center networks (more than 200 switches) will run SONiC in production environments” ~ Gartner

# SONiC Provides the Foundation for Modern Network Fabrics



## Open & Extensible

with standards-based protocols and access to Linux tools & 3rd party apps.



## DevOps Ready

through Containerization, automation and API-centricity.



## Scalable

Created for cloud operations at scale.



## Innovative

Feature velocity, advanced telemetry and deep analytics.

# Why SONiC ?

## Avoid vendor lock-in

- **Open-source** solution to **improve innovation speed** and **code reliability**
- **Open-networking** to **choose** hardware and software and **optimize cost**
- **Easy integration** to various open tools (Automation, Monitoring, Backup, Analytics ...)



## Resiliency and modularity through a containerized system architecture

- Increased **security** and **stability** by **siloing** of the different functionalities within each **container**
- Possibility to **deploy home-made container**
- **Adding features** made easy
- **One solution** for all use cases depending microservices deployed (DC, Edge, SmartNic ...)

## Standardization that allows the sustainability of investment and skills:

- **Standard Protocols**
- **Standard Architecture**
- **Standard CLI & Automation**





# SONiC – Innovation of the industry > Linux of Networking

## SONiC, now a Linux Foundation Project

### Software for Open Networking in the Cloud (SONiC) Moves to the Linux Foundation

Leading open source network operating system enabling disaggregation for data centers now hosted by the Linux Foundation to enable neutral governance in a software ecosystem

**SAN FRANCISCO – April 14, 2021** – Today, the [Linux Foundation](#), the nonprofit organization enabling mass innovation through open source, announced the Software for Open Networking in the Cloud (SONiC, an open source networking operating system), is now part of the Linux Foundation. The Linux Foundation provides a venue for continued ecosystem, developer growth and diversity, as well as collaboration across the open source networking stack

- Leading Network Operating System for Disaggregated Hardware, deployed in large scale enterprise and cloud data centers globally
- Global Community of Cloud Providers, Enterprises, ODMs and Silicon Vendors
- Collaboration with Open Compute Project on SAI specifications, & new Hardware-Software co-design strategy

THE **LINUX** FOUNDATION



## SONiC Ecosystem

Merchant Silicon	Switch Platform	Adoption	System/Service

THE **LINUX** FOUNDATION

<p><b>Dennis Cai</b> SR DIRECTOR, HEAD OF NETWORK INFRASTRUCTURE, ALIBABA CLOUD</p>	<p><b>John DaCosta</b> VP NETWORKING SWITCH PRODUCT MARKETING, MARVELL</p>	<p><b>Ed Doe</b> VP &amp; GM, NETWORKING AND EDGE GROUP, INTEL</p>	<p><b>Mohammad Hanif</b> SENIOR DIRECTOR, ENGINEERING, BROADCOM</p>	<p><b>Saurabh Kapoor</b> DIRECTOR OF PRODUCT MANAGEMENT, NETWORKING SOLUTIONS AT DELL TECHNOLOGIES</p>
<p><b>Amit Katz</b> VP, ETHERNET SWITCH, NVIDIA</p>	<p><b>Dan Lenoski</b> VICE PRESIDENT, NETWORK INFRASTRUCTURE, GOOGLE</p>	<p><b>Xin Liu</b> PRINCIPAL, P4A MANAGER, MICROSOFT SOFTWARE COMMITTEE CHAIR</p>	<p><b>Dave A. Maltz</b> TECHNICAL FELLOW &amp; CVP, AZURE NETWORKING, MICROSOFT CHAIR OF THE GOVERNING BOARD</p>	<p><b>Vijay Tapaskar</b> VP ENGINEERING, CISCO</p>

**Lihua Yuan**



## SONiC Governing Board Leadership

## SONiC Community

2 releases/year	400 ~ 800 commits/month
1710 active code contributors	105 supported platform


- New Subgroup on Special Topics**
- PINS subgroup
  - DASH subgroup
  - IPSec SAI subgroup
  - YANG model subgroup
  - Kubernetes subgroup
  - Chassis subgroup
  - MPLS subgroup
  - SONiC application subgroup

THE **LINUX** FOUNDATION

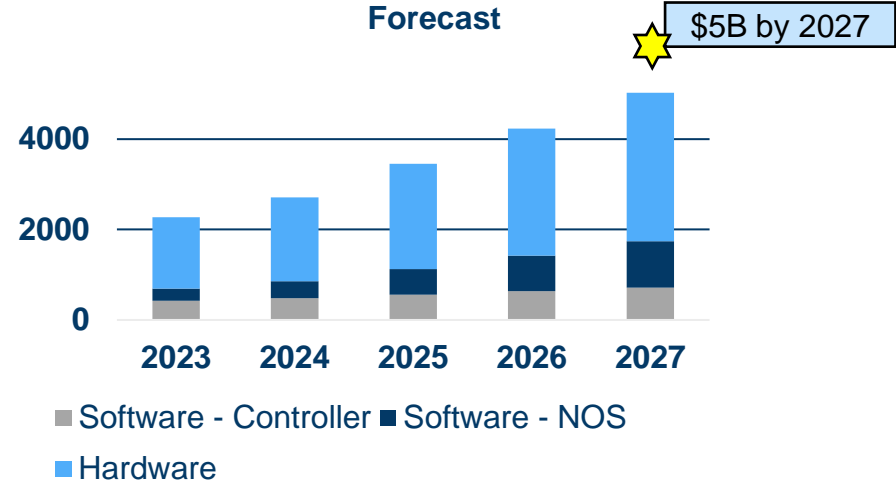


# Exponential Growth and Adoption of SONiC

Hype Cycle for Enterprise Networking, 2023



Dell'Oro Group - WW SONiC Market Forecast



Units in Billion US Dollars

## Gartner



### Disruption

SONiC is predicted to be a long-term market disruption driver




### Innovation

SONiC promises innovation like Linux transformation of server OS Market, by removing ties w/ hardware vendors & standardizing NOS



### Commercial Support Recommended

For mission-critical data centers, consider using commercially supported distributions of SONiC to ensure stability and reliability

A close-up, front-facing view of the Mandalorian helmet, which is highly reflective and detailed. The helmet is set against a dark, atmospheric background with some faint light sources. The helmet's visor is prominent, and the overall lighting is dramatic, highlighting the metallic textures and the iconic T-shaped design.

This is the way.



# Our Vision & Strategy for a Modernized Network Fabric



## Our Vision

Simplify Fabric Design & Management

Security, Reliability and Flexibility

Make SONiC the Linux of Networking

Interoperable Standards-based Technology



## Our Strategy

Commercial Distribution of SONiC

Unified NOS from Edge to Cloud

Complete Enterprise Feature Sets

E2E Support, Documentation, Training



## Result

**Delivering customers full control and choice of the technology stack** – HW, SW, and ecosystem tools, at optimal cost without vendor lock-in

# SONiC Differentiators

## Edge Use Case



**E3248PXE** : 48 x 10GbE PoE (90W) + 4x25GE + 2x100GE

**E3248P** : 48 x 1GE PoE (30W) + 4x10GE + 2x100GE

- Unique SONiC play at the Edge
- High-end access switches

## Run 3rd party container



- Huge databased of containerized applications
- Extend use cases beyond switching ...

## Ecosystem



Open source




Commercial

- No software tools lock-in
- No hardware lock-in
- Consistency from Edge to DC

# Summary

- We are entering the era of open source networking
- SONiC and its partners/community is leading the change: Question is not 'IF' but '**WHEN**'
- The risk is not in moving to SONiC. The risk is in **not moving** to SONiC.
- Dell Technologies is bridging the gap and bringing SONiC to the enterprise with open source innovation, enterprise and cloud data center and GenAI readiness and global support and services on Dell EMC PowerSwitch and Open HW platforms

A close-up, front-facing view of Mando'ao in his iconic silver helmet and black cape. The helmet has a visor with a black slit for eyes. The background is dark and slightly out of focus, suggesting an indoor setting with some ambient light.

This is the way.

