

ARISTA

Investor Presentation

May 2026

Safe Harbor


Forward-Looking Statements - This presentation contains “forward-looking statements” regarding our future performance, including but not limited to statements regarding Arista’s growth drivers, total addressable market, total addressable data center ethernet switch revenue, and campus ethernet switch revenue. Forward-looking statements are subject to known and unknown risks, uncertainties, assumptions and other factors that could cause actual results, performance or achievements to differ materially from those anticipated in or implied by the forward-looking statements including but not limited to risks associated with: escalated or escalating U.S. tariffs and countermeasures and retaliatory actions taken by other countries; some of the key components in our products come from sole or limited sources of supply, which increases the risk of supply shortages, extended lead times or supply changes; key component supply constraints and inventory management; our reliance upon a predominant merchant silicon vendor; enhanced import/export restrictions, as well as countermeasures taken by affected countries; large purchases by a limited number of customers who represent a substantial portion of our revenue; adverse economic conditions, continuing uncertain economic conditions or reduced information technology and network infrastructure spending; volatility in our revenue and revenue growth rates; variability in our gross margins; variations in our results of operations; the rapid evolution of the networking market; failure to successfully carry out new products and service offerings and expand into adjacent markets; intense competition and industry consolidation; expansion of our international sales and operations; investments in or acquisitions of other businesses, products or technologies; industry cyclicity; fluctuations in currency exchange rates; failure to raise additional capital on terms satisfactory to us; our inability to attract new large customers or sell additional products and services to our existing customers; inability to grow sales of switching and routing platforms which generate most of our product revenue; inability to increase market awareness or acceptance of our new products and services; decreases in the sales prices of our products and services; long and unpredictable sales cycles; inability to offer high quality support and services offerings; declines in maintenance renewals and support contracts by customers; product quality problems, defects, errors or vulnerabilities in our products; failure to anticipate technological shifts; our dependence on third-party manufacturers to build our products; assertions by third parties of intellectual property rights infringement, misappropriation or other violations; failure or inability to protect or assert our intellectual property rights; cybersecurity incidents and breaches of our cybersecurity systems, or other security or privacy breaches or incidents; failure to detect cybersecurity incidents; failure to comply with government law and regulations; issues in the development and use of artificial intelligence, combined with an uncertain regulatory environment; future decisions to reduce or discontinue repurchasing our common stock pursuant to our stock repurchase programs; and other future events. Additional risks and uncertainties that could affect us can be found in our most recent filings with the Securities and Exchange Commission, including, but not limited to, our annual report on Form 10-K and quarterly reports on Form 10-Q. You can locate these reports through our website at <https://investors.arista.com/> and on the SEC’s website at <https://www.sec.gov/>. All forward-looking statements in this presentation are based on information available to the company as of the date hereof, and we disclaim any obligation to publicly update or revise any forward-looking statement to reflect events that occur or circumstances that exist after the date on which they were made.

ARISTA

Corporate Overview

The Making of Arista's Decade

10+ Years | 10K+ Customers | 150 Million+ Ports



Founders and Senior Leadership



Jayshree Ullal
CEO &
Chairperson



Andy Bechtolsheim
Founder, Chief Architect



Chantelle Breithaupt
Chief Financial Officer &
SVP



Ken Duda
President, Founder &
Chief Technology Officer



Hugh Holbrook
Chief Development
Officer



Todd Nightingale
President & Chief
Operating Officer



Mark Foss
SVP, Global Operations
& Marketing



John McCool
SVP & Special
Advisor



Ashwin Kohli
Chief Customer
Officer



Chris Schmidt
Chief Sales Officer



Chris Bellmare
SVP, Americas



Tyson Lamoreaux
SVP, Cloud/AI

Leader in Data-Driven Networking



Addressing a broad TAM across diverse customer business priorities: AI and data center, campus and branch networking, security and observability



With a differentiated architecture that combines reliable, high-performance hardware with high-quality software



Enabling a modern operating model that lowers the cost of operations and delivers real-time visibility into an autonomous and software-driven network



With satisfied customers adopting the broader Arista platform and driving increased market share across multiple use cases

The Arista Software Differentiators



NetDL

AI/ML-driven automation and analytics based on a single, centralized, and high-fidelity repository that aggregates diverse Arista and third-party datasets



EOS

Modern publish-subscribe-based state-sharing operating system that delivers industry-leading performance, reliability, and availability with capabilities such as in-service software upgrades performed with <1 sec traffic impact



Platform Dependent
Infrastructure

Hardware abstraction interface for diverse silicon chipsets delivering superior route scale, fastest network convergence times, lossless load balancing, and extensive hardware telemetry that combine to drive positive business outcomes, such as speeding up job completion times



Robust Optical
Connectivity

Support for connectivity via a diverse range of optics that leverage a common driver model, ensuring seamless plug-and-play as well as extensive telemetry that transforms raw optical data into actionable insights



NOS Independent
HW Diagnostics

Network operating system (NOS)-agnostic, robust hardware diagnostics and validation suite for every Arista Platform that ensures a consistent and superior experience

Leader Across Networking Segments



Gartner Magic Quadrant for Data Center Switching 2025

Gartner Magic Quadrant for Enterprise Wired and Wireless LAN Infrastructure 2025

Gartner Magic Quadrant for SD-WAN 2024


Source : Gartner®, Magic Quadrant™ for Data Center Switching, Andrew Lerner et al., 31 March 2025

Gartner®, Magic Quadrant™ for Enterprise Wired and Wireless LAN Infrastructure, Mike Leibovitz et al., 25 June 2025


Gartner®, Magic Quadrant™ for SD-WAN, Jonathan Forest et al., 30 September 2024

GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally, and MAGIC QUADRANT is a registered trademark of Gartner, Inc. and/or its affiliates and are used herein with permission. All rights reserved. *This graphic was published by Gartner, Inc. as part of a larger research document and should be evaluated in the context of the entire document. The Gartner document is available upon request from Arista Networks.* Gartner does not endorse any vendor, product or service depicted in its research publications, and does not advise technology users to select only those vendors with the highest ratings or other designation. Gartner research publications consist of the opinions of Gartner's research organization and should not be construed as statements of fact. Gartner disclaims all warranties, expressed or implied, with respect to this research, including any warranties of merchantability or fitness for a particular purpose.


Arista Enables a Modern Network Operating Model



Why 10,000+ Customers Choose Arista



Strengthened by the foundation of robust quality assurance and diagnostics



Through an extensible platform that uses open standards to avoid vendor lock-in



Driven by real-time telemetry and AI Ops that lower manual effort for the network operator

Arista Support NPS




 **World-Class • Top 1% of companies**


Arista NPS Survey 2026

Compared to Other Vendors

"Customers who rated us equal to or better than other vendors"



Customer Satisfaction Rating




What are our Customers saying?




Arista Customer Feedback

*"Arista TAC's biggest differentiator is its **knowledge, speed, and capability**. Supporting every customer and use case is not easy, and delivering a consistently world-class TAC experience is a real challenge. Yet in every interaction I've had, Arista TAC has gotten it right. I **couldn't be happier with the support**."*




Arista Customer Feedback

*"They have been **available 24/7** for me. Every single time. I love the **close integration** they have with the **account and engineering teams**. It makes it convenient for me to navigate high-stakes issues."*



Arista Customer Feedback

*"What I value most about Arista TAC Support is their ability to **resolve issues quickly**. The engineers **consistently demonstrate strong technical expertise** and are able to guide us to solutions even when facing highly complex problems. Their willingness to work flexibly with our schedule and support us until the issue is fully resolved **has earned our deep trust**."*



Customer Savings with Arista's Modern Operating Model: A Case Study

Network Provisioning with Arista AVD	14 Days to 4 Hours	\$108K
Zero Touch Provisioning	40 Hrs/Device to 6 Hrs	\$520K
Unique Code Configuration	112500 lines to 1500	\$755K
Reduced Troubleshooting Hours	1350 Hrs to 540 Hrs	\$165K
Time Spent on Change Mgmt	7650 Hrs to 4200 Hrs	\$702K
Change Mgmt Team Hours	30 Mins/Week to 4 Mins	\$312K
Network Incident Downtime (CODE)	6/Year to 2/Year	\$1.1M
Network Incident Downtime (Human Error)	75% Reduction	\$1.8M
Automated ANTA Network Testing	770,000 Test Created	\$105K

\$5.65M

Reduced Cost of Network Operations over Three Years

Arista's Commitment to Corporate Responsibility



Environmental

Science-Based Targets Initiative-approved targets:

42% Reduction in absolute Scope 1, 2, and 3 emissions by 2030

Net Zero by 2050

Global Net Zero Supplier and Customer Engagement Program

7th Year

Reforestation across the globe through the Arista Tree Planting Initiative



Social

#15 America's Best Companies to Work For and Invest In by Forbes

7th Consecutive year of record number of Arista Foundation grants to impactful global non-profits

Ongoing partnership with Arizona State University to create new pathways into the IT workforce

Supporting food insecurity, education, disaster relief, essential medical assistance, and sustainability worldwide



Governance

Board oversees strategy, business plans and risk management

Lead Independent Director for independent oversight of management

77% independent directors
100% independent committee members

One of the few Fortune 1000 companies with **Female CEO & Female CFO**

Active year-round stockholder engagement



View all our [awards](#)



ARISTA

Quarterly Highlights

First Quarter 2026

Company Highlights

Arista announced XPO high-density liquid-cooled pluggable optics

Designed for next-generation AI data centers, XPO reduces networking racks by up to 75% and saves up to 44% of floor space compared to traditional pluggable optics, enabling shorter cable runs and lower-power AI scale-up interconnect technologies such as copper and RF.

Arista received a 2026 NPS Score of 89

The updated net promoter score indicates that 94% of customers are strongly positive about the company. These industry-leading metrics are a testament to the company's focus on driving customer success and satisfaction.

Powering AI Centers with AI Spines

The universal AI spine powered by the Arista 7800 delivers massive scale, predictable performance, and high-speed interface support. Powerful features such as Virtual Output Queuing (VOQ) eliminate head-of-line blocking, and large buffers absorb AI microbursts and prevent PFC storms.

Arista's FY 2025 Achievements

EOS

- Cluster Load Balancing (CLB) in Arista EOS® maximizes AI workload performance with consistent, low-latency network flows that maximize AI cluster performance and efficiency.
- By combining the power of open standards like MCP with the robust data and programmability of Arista's EOS and NetDL, Arista's AI agents streamline network operations.

Platforms

- Arista's expanded set of AI-Driven campus and branch networking offerings deliver an expanded set of switching, Wi-Fi 7 access point, and SD-WAN capabilities.
- The R4 series of Arista platforms for AI, data center, and routed backbone deployment deliver high performance, low AI job completion times, low power consumption, and integrated security.

Innovation

- Arista AVA® offers agentic AI capabilities to help customers eliminate human effort and reduce the cost of network operations.
- Arista Networks is collaborating with industry leaders to deliver Ethernet for Scale-Up Networks (ESUN), an open OCP workstream committed to the goal of open standards-based solutions for scale-up, based on Ethernet, and open to all.

Investor Conference Participation

Arista will attend the following conferences
with the financial community in May and
June 2026

21st Annual Needham Technology, Media, & Consumer Conference

Chantelle Breithaupt, Chief Financial Officer

Rudolph Araujo, Area Vice President of Product and Investor Advocacy

Thursday, May 14, 2026, 11:45 AM - 12:25 PM ET

Webcast: <https://investors.arista.com>

J.P. Morgan 2026 Global Technology, Media and Communications Conference

Ashwin Kohli, Chief Customer Officer

Tyson Lamoreaux, Senior Vice President, Cloud and AI Networking

Tuesday, May 19, 2026, 2:15 PM - 2:50 PM ET

Webcast: <https://investors.arista.com>

William Blair 46th Annual Growth Stock Conference

Chantelle Breithaupt, Chief Financial Officer

Rudolph Araujo, Area Vice President of Product and Investor Advocacy

Thursday, June 2, 2026, Time: 1:00 PM - 1:40 PM ET

Webcast: <https://investors.arista.com>

Bank of America Global Technology Conference 2026

Todd Nightingale, President and Chief Operating Officer

Brendan Gibbs, Area Vice President of Product Line Management


Wednesday, June 3, 2026, 10:45 AM – 11:15 AM ET

Webcast: <https://investors.arista.com>

ARISTA

Growth Drivers and Market Opportunity

DATA AS INTELLIGENCE: THE NETWORK'S NEW ROLE




Arista's Growth Drivers Addressing \$105B TAM



- Core - DC & AI/Cloud Networks → Fueled by large AI investment increases
- Strategic Expansion – Campus & Routing → Broad set of Campus/Branch offerings including SD-WAN
- Cognitive Networks – Software & Services → Increasing Observability market opportunity

Continued Growth Ahead In Switching Market Opportunity


Data Center Ethernet Switch Revenue



Source: 650 Group Ethernet Switch Data Center Forecast Q4 2025

US Top 5 Cloud Providers: Amazon, Apple, Facebook, Google, Microsoft
Chinese Top 3 Providers: Alibaba, Baidu, Tencent

Campus Ethernet Switch Revenue




Source: 650 Group Campus Ethernet Market Forecast Q4 2025

Large Business - Fortune 2000
Medium Business - Rest of Enterprise Market
Small Business - Less than 50 Employees

Continued Growth Ahead In Switching Market Opportunity


Data Center Ethernet Switch Analysis and Forecast




Source: Dell'Oro January 2026 – Long-Term Ethernet Data Center Switch Forecast

Delivering Consistent Market Share Gains vs Cisco

High Speed Data Center Switching Market Share in Dollars



Share in Ports




Source: Crehan Research Data Center Switch Market Share Report 4Q'25
Note: 10GbE and Higher - Excludes blade switches

Strong, Sustained Share Gains Continued Through 2025

Arista's Market Leadership in 100G/200G/400G/800G

AI and Data Center High Speed Ethernet Market Analysis




Source: 650 Group Ethernet Switch Data Center Branded Vendor Shares Q4 2025

ARISTA

Centers of Data

Powered by Arista Networks

Arista Platforms for the Centers of Data



7060X
AI Leaf



7280R, 7800R
AI Spine



7700R
Distributed Etherlink Switch

AI Center (Etherlink)



CloudVision



EOS



AVA

Access Points
Wi-Fi 7, 6E, 6



710P, 720X
Fixed Access
PoE Switches



750
Modular Aggregation
PoE Switches



Velo Edge
Unified Branch

Campus Center



**7050X, 7060X
7130**
Shallow Buffer &
Low Latency Leaf



7020R, 7280R
Fixed Deep Buffer
Leaf or Spine



7300X
Modular Shallow
Buffer Spine



7800R
Modular Deep
Buffer Spine

Data Center



AWE 7200
Route Reflector,
Cloud Routing



Velo Edge
Secure SD-WAN



7020R, 7280R
Access & Edge
Routing, Peering



7800R
Core Routing

WAN Center

Built on the Foundation of Netdi

Network Diagnostic Infrastructure



Open
Standards Commitment

Meticulous
Design (L0/L1)

Rigorous
Validation

Integrated
Production

Resilient &
Secure Platform

Support &
Expertise



Tools

HW Diagnostics

Signal Integrity

Quality Control

Control Trackers

L0 Passive & L1
Active Events

Every Arista platform has the same meticulous validation and robust diagnostics, whether an operator chooses to run Arista’s flagship EOS or an open-source NOS

Data-Driven Network-as-a-Service Platform



Arista CloudVision®



Network Automation Simplicity

Complete Network Telemetry

Technology Partner Centric

One Network Management System

Arista Extensible Operating System (EOS®)



Trusted by 10K+ Customers

Highest Quality Network Operating System

Complete Switch & Router Capabilities

Software Programmability

One Network Operating System (EOS) and One Network Data Lake (NetDL) powered by Netdi for Diagnostics Infrastructure

Edge Network



WAN Network



Campus Network



Data Center



AI Cloud



Private Cloud



Public Cloud



Our Platform Mission

Deliver best multidomain data-driven networking for client to cloud /AI deployments

ARISTA

Arista EOS, NetDL, AVA, and CloudVision


Top Five Reasons EOS is the Best



Superior Quality




Single OS Image



Real-time Visibility




APIs & Automation




Rich Features &
Platforms

The EOS Architectural Advantage



Arista EOS Stack Evolution



AI-Enabled Network Operations 2020s

Centralize all network state data and enrich with 3rd party integrations to build a foundation for AI NetOps



NetDL - adds enriched data sources to packet, flow, and network data to enable powerful AI/ML applications

SDN and Intent Systems 2010s

Centralized models for network automation and data



NetDB - centralizes the network state of multiple devices, enabling automation and provisioning orchestration

Federated Routing Protocol 1990s/2000s

Development of BGP/OSPF and multiple distributed topology construction methods - achieved unprecedented scale



SysDB – aggregates network state at the device-level to improve reliability and efficiency via a publish/subscribe model

EOS: A Network Operating System Optimized for AI



A robust and resilient network OS

- In-service software upgrades
- Low defects and CVEs

A lossless network

- Cell-based fabric
- Optimal load balancing
- Adjustable buffer allocation schemes
- Advanced congestion management



Scalability and multi-tenancy

- Mature feature set with VXLAN + EVPN
- Native support for encryption on the wire

Network observability

- Extensive counters
- Real-time state streaming
- End-to-end visibility across AI-workload/server, NIC, and switch




Optimizing Software for TTFJ in an AI world

- TTFJ: **Time to First Job** is a critical metric
 - Every day of bring-up is a day the GPUs are idle
- EOS is architected for TTFJ:
 - Availability is critical to bring up time
 - Deep telemetry
 - Deep platform sharing across silicon families
 - Optics and Phy management
 - Platform commands and telemetry
 - Test infrastructure and tests
- This can not be bolted on after the fact

Speed up TTFJ and drive better XPU utilization


Rich EOS Suite for AI Centers

- Robust EOS foundations
 - State-rich, high quality and in-service update availability
- Enhanced load balancing for AI
 - RDMA-aware load balancing and QoS
 - Dynamic cluster load balancing
- Dynamic congestion control
 - ECN, PFC, buffer mgmt
- AI instrumentation
 - Including agents for hosts/DPUs/NICs and job-oriented observability
- Front End EOS-scale for AI Centers
 - STANZ traffic analyzer
 - Multitenancy for AI Centers interconnect
 - Large Radix 600 way
 - Multiprotocol scale



Powerful EOS Suite across Back End and Front End Etherlink networks

Extensible Network Data Lake with a Broad Ecosystem of Industry Partners



Optimized Networking: From Design to Operational Excellence



Design

Prescriptive architectures enabled as software with broad support for a myriad of design patterns



Arista Validated Designs (AVD)
Arista Cloud Test
Source of Truth integrations
Security / Identity Policy



Operations

Cloud-like service interface and catalog to provision topologies, overlays, and policies and controls



Enterprise-wide Automated Provisioning
Network-wide Upgrades
Quality of Experience
Real-time Monitoring



Support

Enabling a shift from reactive TAC to proactive software-led troubleshooting and multi-channel support



Better Data → Better Insights
Automated Bug/PSIRT
Autonomous Virtual Assist (AVA)
Direct TAC Access
Proactive RMA




Full Stack Lifecycle

HW and SW lifecycle management.
From design, to HW selection, to procurement, to software versioning



Automated Lifecycle Mgmt
Compliance / Risk Mgmt
Broad Operator Types
Software Recommendation
3rd Party Integrations



Network as-a-Service built on the NetDL Foundation

Arista AIOps for Networking Stack

Agents

Ask AVA

AVA Insights

AVA Agents

Services



AVA Runtime

LLMs | Context Engine | Tools & Actions | Policy & Safety | MCP

Network Intelligence

Event Engines

Topology Telemetry

Provisioning & Intent

Health & Compliance

Applications & Security

Best in Class Data



EOS Network Data Lake (NetDL)

Unified, Multi-domain Network Telemetry, 3rd-party Data Sources, Application/Compute

Powered by EOS



OT



Campus



Data Center



IOT



SDWAN



Remote Users



IaaS /
PaaS



SaaS

ARISTA

CloudVision: Multi-Function, Multi-Domain NetOps



ARISTA

Arista Data Center, Cloud and AI Solutions

AI Traffic Differs from Cloud Networking

Datacenter Type	CPU Datacenter	AI Datacenter
Connectivity per XPU	100G	400G/800G
Network Design	Oversubscribed	Non-Blocking
Bandwidth per 128K Chips	1 Petabit/Sec	100 Petabit/Sec
Traffic Pattern	Random	All-to-All, All-Reduce
Flow Duration	Short-lived	Long-lived
Congestion Impact	Minor	Very Significant

AI Networks Need a New Architecture to Optimize Performance


AI Deployment Landscape

	Dedicated Training	Multi-tenant / Hosted	Inference
Purpose	Building a model based on very large data set	Training / Experimentation / Tuning / Inference	Application of the trained model
Where it happens	In the Cloud	Specialist Providers / Large Enterprise	In the Cloud / Distributed / On Prem
Scale	10,000s to 100,000s XPU	100s to 1,000s XPU	10s to 100s of XPU
Workload Mix	Minimal, Model Specific	Generic, Highly Variable	Back-End: Model Specific Front-End: Transactional
Predictability	Predictable	Dynamic	Back-End: Predictable Front-End: Dynamic
Networking Profile	Very large scale Very high bandwidth	Large scale, high bandwidth	Back-End: High Performance Front-End: Load Dependent
Network Optimization	Highly Optimized	Fully Connected	Back-End: Optimized Front-End: Data Center+

Networking for the AI Center: Training & Inference

Front End Leaf Spine

Connecting CPUs and Storage to the Data Center, Campus and Users




Thousands of XPU hosts, high speed storage and support applications

Back End AI Fabrics

Scale Out


Back-End Connecting Racks of XPUs Together to Form a Cluster



Inter-rack scale out to ~100k XPUs

Scale Across


Connecting Distributed Clusters to Increase Scale



Inter-site scale across to >1M XPUs

Scale Up

Connecting XPUs within an Enclosure



Intra-rack scale up for ~1k XPUs

Domain-Specific Products for Data Center Networks

Consistent High Performance and Extensible EOS

Highest Bandwidth



7060X Series

Cloud and AI application stacks
Highest Switching performance
Lowest Latency
Scale Out & High Radix
High density 400/800G Fixed Spines
Multiple Choices - 25G to 800G

Feature Rich



7050X Series

Enterprise application stacks
Multicast, EVPN, VXLAN
Rich Telemetry for deep visibility
Compute TOR for 10/25/50/100/400G
Flexible Traffic Management
Integrates into Campus Architecture

Cloud Grade Scale



7280/7800R Series

Routing, Storage and AI Spine
EVPN, MPLS, SR
High capacity scaling
Metro & DCI with MACSec & ZR
Packet Spraying & VOQ
Fixed and Modular Choices

Arista Etherlink Evolution

- Evolution from current 100G to 200G SerDes platforms starting in 2026/27
- Single Chip TH5 → TH6 followed by VOQ Jericho/Ramon J3/R3→J4/R4
- Optics will evolve from the current 800G to 1600G and variants
- Etherlink TCO maximizes compute density - 65% performance advantage, 35% power and space savings
- Simplified network with 85% fewer switches to manage and 35% reduction in cables and optics



Arista Etherlink AI Portfolio

Single Tier



Fixed & Modular Systems


51.2T $\Rightarrow$ 460.8T

64-576 x 800G Nodes

128-1152 x 400G Nodes

Lowest Cost, Power &
Complexity

Leaf-Spine



2- or 3-Tier
Leaf-Spine or Plane-based
Multi-Petabit Bisectional B/W
DLB, PFC, ECN,
Scales to Hundreds of
Thousands of Accelerators

Distributed Etherlink Switch




Single Hop Interconnect

Fully Scheduled Lossless
100% Efficient, Cell Spraying
Multi-Petabit Bisectional B/W
Leading Edge Performance for
up to 30,000 Accelerators

Accelerator & NIC Agnostic, Open Standards, Smart AI Features

Scaling out the AI Spine

7800R4 AI Spine



Lossless Fully Scheduled VOQ
100% Fair And Efficient Traffic Spraying
Integrated Redundancy and Resilience
Optimized Pipeline for AI Workloads

50X Scale Out

7700R4 Distributed Etherlink Switch



Lossless Fully Scheduled VOQ
100% Fair And Efficient Traffic Spraying
Integrated Redundancy and Resilience
Optimized Pipeline for AI Workloads

Common Architecture – Re-packaged for AI Scale Out

The 7700R4 Distributed Etherlink System


- Single-tier distributed switching system
 - Scales to over 30k accelerators in single system
- Optimized for AI workloads
 - 100% efficient VOQ/Cell architecture
- Rich EOS and CloudVision
 - Managed as single network cluster
 - Independent device upgrades and replacement
- Standard Ethernet network connectivity options
 - Linear-drive Pluggable Optics $\geq$ 50% power savings



DES is the Purest Form of Scale-Out Networking for AI

Arista's Scale Optimized AI Networking Portfolio


Single Device



400G Hosts
7060X6: ≤ 128
780xR4: 129 – 1,152

800G Hosts
7060X6: ≤ 64
780xR4: 65 – 576

7060 Leaf and Spine




400G Hosts (400G L-S)
1,153 – 8,192

400G Hosts (800G L-S)
1,153 – 4,096

800G Hosts
577 – 2,048

7060 Leaf and 7800 Spine




400G Hosts (400G L-S)
8,193 – 73,728

400G Hosts (800G L-S)
4,096 – 36,864

800G Hosts
2,049 – 18,432

7060 3-Tier




400G Hosts (400G L-S)
73,729 – 262,144

400G Hosts (800G L-S)
36,685 – 131,072

800G Hosts
18,432 – 65,536

7800 2-Tier



400G Hosts (400G L-S)
262,145 – 663,552

400G Hosts (800G L-S)
131,073 – 331,776

800G Hosts
65,537 – 165,888

Distributed Etherlink Switch

Single Logical Hop



400G Hosts
Up to 31k


800G Hosts
Up to 27k

With Arista, Most AI Clusters Can Be Built With An Efficient 1 or 2-Tier Design

Scale without Complexity

Optimal power utilization • Easiest incremental growth • Minimal optics/cabling


Arista Bluebox: Offering Customer Options



Delivering Faster Job Completion Times

Accelerated High Scale RDMA for AI Powered by the Modular AI Spine Chassis

Feature Category	Fixed Form Factor	Modular Chassis Form Factor	Modular Chassis Advantage
Buffer Architecture	Shallow (Tomahawk)	Hierarchical Hybrid Buffers (Jericho) <i>Shallow on-chip + deep on-package buffering</i>	Very low latency during optimal conditions, while ensuring lossless RDMA delivery during transient congestion or long-distance transport
Congestion Management	Tiered/Hop-by-hop	Integrated Virtual Output Queuing (VOQ)	Fair and lossless transport with AI workload prioritization of RDMA traffic at high scale
Port Speeds	Up to 800 Gbps per port	Up to 3.2 Tbps per port with HyperPorts	HyperPort technology enables 4x faster port speeds for 44% faster JCT than 800GbE ports
Chassis Port Density	Up to 64x 800GbE per fixed switch <i>Inter-switch interconnect over optics</i>	Up to 576x 800GbE per modular chassis	Enables ultra scale in 2-tier AI cluster designs, for lower power, lower cost, and fewer optics
		Non-blocking, cell-based fabric <i>Low cost, low power interconnect over copper</i>	Perfectly efficient traffic distribution between all ports to maximize AI spine capacity utilization



The Arista 7800 Chassis delivers **up to 44% faster job completion times*** in Scale Out and Scale Across Networks

Improved Route Convergence | Faster Recovery Time | Lower Power Consumption | Higher AI Cluster Scale

ARISTA

**Industry Innovation That Drives
Open Standards**

An Open Ecosystem for AI

Model Builders

ANTHROPIC
Meta

Google

Microsoft
OpenAI

NIC

AMD
BROADCOM
intel
NVIDIA
Ultra Ethernet Consortium


Storage

ddn NetApp
PURESTORAGE VAST
WEKA


AMD arm Google intel Microsoft
NVIDIA PENGUIN SOLUTIONS SambaNova SYSTEMS SUPERMICRO

Compute & Infrastructure

Arista AI Optimized Load Balancing Performance




NCCL All-to-All Performance



Driving Open Standards for AI Networking


Ultra Ethernet Consortium (UEC) and Ethernet for Scale Up Networking (ESUN)

UEC Steering Members



-  Native multipath aware end-to-end transport for very large-scale RDMA clusters
-  Rapid connection ramp up, order independence and fast recovery
-  Automated link level and end-to-end congestion control
-  High-definition telemetry and signaling mechanisms
-  Native hardware-based, high group-based encryption
-  Maximize network utilization and minimize latency

ESUN OCP Workstream



Focus Areas

- L2/L3 transport
- Efficient headers
- Error recovery
- Lossless network

XPO: Densifying the Optical Layer

1. Up to 75% Reduction

in Switch Racks

Large savings in structural costs (racks, bus bars, manifolds, plumbing, etc.)

2. Up to 44%


Footprint Reduction

Building half as many data centers saves a lot of time and money

4. Denser Switches & Routers

Up to 6.4 Pbps per rack

The density per rack is limited by the front-panel density of OSFP today



**8x
OSFP**



**1x
XPO**



4. Shorter Reaches

for Scale-Up Tech

Incredibly important to enable the lowest power options for the all-important Scale-Up use case

5. Optimize Revenue Potential

Reducing Structural "Tax"

No cloud/AI revenue is generated by rack packaging i.e. structural cost is a tax on revenue

XPO Addresses AI Requirements

1	Higher Density	204.8T per OU; 4X OSFP	✓
2	Liquid Cooling	Large AI clusters are liquid cooled	✓
3	Higher Reliability	Lower system-level failures in time	✓
4	Optics Universality	DR, FR, LR, SR, ZR/ZR+, next-generation coherent-lite, slow&wide, copper, and RF-Microwave	✓
5	Higher Power Efficiency	Linear-drive (LPO/LRO) support	✓
6	Rapid Time to Market	Leverages existing ecosystem	✓
7	Open Robust Ecosystem	Multi-vendor optionality	✓

Soldered-Down vs Open CPO

	Soldered-Down CPO	Open CPO	Open CPO Advantages
Connectorized	No	Yes	Allows repairability
Multivendor	No	Yes	Provides purchasing optionality via an open ecosystem
Laser Source	External	External/Internal	Offers the option of internal lasers that reduce power utilization
Fiber Attach	Fiber Array Unit	Pigtail	Lowers the cost through the use of pigtail connectors
Co-packaged Copper Support	No	Yes	Allows use of lower power copper connectivity

Open CPO = Lower Operational Risk • Lower Power Utilization • Lower Cost • Faster Time-to-Market

Arista: History of Innovations in Optics

OSFP




LPO



XPO



Open CPO (Future)



Key Customer Benefits:

- Enables high-speed optics: 400G, 800G, 1.6T +
- For any type of optic: DAC, DR, FR, LR, SR, ZR/ZR+
- At high power & thermal tolerance: up to 40W

Key Customer Benefits:

- 50% reduction in power usage per optic, with increased reliability, and reduced latency
- Yields ~40% reduction in power at the system level: 1kW less per 51.2T system

Key Customer Benefits:

- 4x density improvement
- 8x higher reliability; liquid cooling optics at 45C delivers higher longevity
- Supports 400W; 8x 1.6T ZR/ZR+ at 50W per channel


Key Customer Benefits:

- Serviceable, repairable via RTF
- Lower power consumption, lower cost, and simpler to manufacture than soldered-down CPO
- Open with multi-vendor options, supporting CPO and CPC


ARISTA

Arista Cognitive Campus


Campus Networking Ready for Modernization



Arista Strategy For The Cognitive Campus




Arista's Comprehensive Campus Portfolio



Arista Modern Operating Model

- **AVA AI** for Networking
- **AGNI** for Identity Management
- **NetDL** for real-time consistent state streaming
- **MSS** for consistent end-to-end segmentation
- **SWAG** for virtual campus switch aggregation
- **SD-WAN** for distributed Campus-to-Branch




Arista Architecture Blueprints

- **Common Spine** across the enterprise (DC, AI, Campus & Routing)
- **Single EOS Operating System**
- **Scalable L&S Architecture**
- **Eliminate** downtime (SSU, Hitless upgrades)

Arista Cognitive Campus Operations


CloudVision Campus for Wired + Wireless

Day-0: Build




Tenant Provisioning,
Device Pre-provisioning,
Design Templates,
Cloud Test

Day-1: Deploy



Zero Touch Provisioning,
Device Onboarding,
Image Download,
Campus Studios,
Connectivity Baselines,
Segmentation Policy,
PoE Management

Day-2: Operate




AIOps + Troubleshooting,
Unified Campus Dashboard,
Level 1 Operator Workflows,
Change Controls, Compliance,
Zero Touch Replacement
Application QoE
Radio Resource Management

Zero Touch Provisioning → Zero Touch Operations

ARISTA

Arista Routing Portfolio

Arista Routing Use Cases



Purpose Built for Intensive Cloud, Enterprise and Carrier Workloads


Spine and Routing R3 Series Portfolio

30

Fixed and Modular Products
Jericho2c+ and Qumran2c+




Universal TunnelSec™ Encryption



Dense 400G DWDM
Rich EVPN Services

50%

Reduction in Power
Consumption / 100Gbps



Cloud Grade Routing and an AI Spine from 1 to 460 Tbps

The Arista AI Spine (R4 Portfolio)

End-to-End Encryption: MACsec, IPsec, VXLANsec

Fixed 10G and 25G Leaf w/
100G Uplinks



Fixed 100G Leaf
w/ 800G Uplinks



Fixed 800G Spine




Modular 800G Spine
w/ 3.2Tbps HyperPort



7020R4



7280R4




7800R4

AI Spine: Consistent Rich Features, Deep Buffers, Scale with Quality

Arista and VeloCloud: Elevating the Branch

VeloCloud Delivers SD-WAN Connectivity for Branch Offices and HQ Alike



Arista


ARISTA

ARISTA


Arista Network Software and Services

Arista Zero Trust Networking: Enhanced!

Open Services with Interoperability across the Stack



Pervasive Observability with Zero Trust Security



Network Software and Services

Arista's Holistic Approach



CloudVision

Automation tool for configuration and management of wired and wireless networks; collects streaming telemetry from EOS state database for end-to-end network visibility and observability



DMF – DANZ Monitoring Fabric

Network monitoring, analytics, and recording, for capturing data inside the network and delivering it to analysis tools



AGNI - Arista Guardian for Network Identity

Next generation cloud-native solution that delivers identity-based network access control



ARISTA
NDR

Arista NDR

Proactive AI analysis of network data to identify an attacker based on their intent versus merely looking for signatures of an attack



EOS

CloudEOS

EOS that runs in AWS, Azure, and GCP via their marketplace and service catalogs



A-CARE

Arista A-Care – Software updates, features, and break/fix troubleshooting. Powered by Arista AI ops toolset designed to speed problem identification and resolution.

Our Competitive Strengths



Best-of-breed merchant silicon for high-performance, purpose-built platforms

Industry-leading capacity for diverse requirements including, low latency, high port density, and deep packet buffers

Netdi diagnostics and infrastructure service ensures high quality, signal integrity, and power efficiency



One operating system (EOS), one data lake (NetDL), and one management platform (CloudVision) compared to fragmented competitive approaches

Feature consistency across the entire portfolio from client to cloud

Modern operating model with AI Ops, real-time telemetry, and in-service upgrades to lower network operations costs



Arista-staffed Technical Assistance Center (TAC) that provides 24/7/365 support to drive customer success

Industry-leading Net Promoter Score that reflects deep customer trust in both hardware and software leadership

ARISTA

Thank You

www.arista.com