

Arista Networks - Where to Place the Next Bet?

Internal ID: S&F-STR-067

Topic: Platform Strategy, Competitive Strategy, Innovation

Industry: Communications Equipment

Source: Strategy & Finance

Publication date: 5 October 2026

License: Not specified

Decision date and information cutoff: 5 August 2026

Case convention: This is an author-created classroom decision exercise based on public information. The assignment and choices are hypothetical; company disclosures and competitor announcements are historical. No private deliberations are represented.

The decision

On 4 August 2026, Arista Networks reported its first quarterly revenue above \$3 billion and announced new 1.6 Tbps AI fabric platforms. The following day, its Form 10-Q described demand for networking as strong while warning that the long-term trajectory of AI networking remained unknown. It also reported customer concentration, large-account discounts, long testing and acceptance periods, and gross-margin pressure from silicon and supply conditions. [S1, MD&A, pp. 15-18, 21-24; S2]

For this exercise, you are a member of a hypothetical Arista product-strategy leadership team asked to recommend the emphasis for the next product cycle. The company has a fast-growing AI networking opportunity and a broader plan to connect AI centers, data centers, campuses, and wide-area networks through a common software and telemetry stack. It cannot maximize every product, engineering, and customer-facing effort at once.

Assignment: Recommend whether Arista should put greater emphasis on deepening its AI Ethernet fabric leadership or accelerating its broader client-to-cloud “Centers of Data” platform. Explain what it should prioritize, which risks it must accept, and what evidence would cause it to adjust course. This is a classroom choice constructed from public evidence, not a claim that Arista’s board faced this exact allocation decision. Use only information available by 5 August 2026.

A network company broadens its ambition

Arista began as a supplier of high-performance Ethernet switching and routing systems for large networks. Its proposition combines hardware with EOS, its state-oriented operating system, and a growing suite of network applications. The company describes Network Data Lake (NetDL) as a source of telemetry and analytics that can support operations across different network locations. Arista’s “Centers of Data” strategy links four domains: AI centers, data centers, campus centers, and WAN centers. Management presents open, standards-based technology, reliability, and automation as distinguishing features. [S1, Part I, Item 2, MD&A, “Our Strategy,” p. 15]

The scope of the platform has grown beyond the data center. In June 2025, Arista completed its \$300 million cash acquisition of VeloCloud from Broadcom. Arista described VeloCloud’s secure,

AI-optimized cloud WAN portfolio as complementary to its data-center and campus wired and wireless offerings. [S3, Part I, Item 1, “Business,” VeloCloud acquisition]

Arista reports as one operating segment. Its public filings do not disclose revenue or gross margin by AI networking, cloud, campus, or WAN product. The company’s platform language describes a strategic direction; it does not provide a measured cross-domain adoption rate or prove that a campus customer will buy AI networking, or vice versa. [S1, Note 10, “Segment Information”; MD&A, pp. 15-18]

The AI networking opening

AI computing clusters move data rapidly among processors, memory, and storage. Network performance can affect the effective use of expensive compute equipment. Arista’s August 2026 release introduced 1.6 Tbps AI fabric platforms, including liquid-cooled configurations, for scale-up, scale-out, and scale-across networks. It characterized second-quarter demand as strong and reported that product shipments had increased across its customer base. [S2; S1, MD&A, “Revenue, Cost of Revenue and Gross Margin,” p. 17]

The opportunity is contested. NVIDIA describes Spectrum-X as an AI Ethernet platform integrating switches and networking adapters. In May 2026, it announced that its Multipath Reliable Connection protocol, first optimized on Spectrum-X hardware, was being released as an open specification through the Open Compute Project. NVIDIA named OpenAI, Microsoft, and Oracle among deployments. These are NVIDIA’s own product and deployment claims, not independent comparative tests. [S4]

The choice is not simply proprietary versus open. Both vendors make open-Ethernet claims, while combining proprietary software, hardware, and system-level designs. Arista can emphasize its network operating system and multi-vendor interoperability; NVIDIA can offer close integration with its compute stack. The public sources available here do not provide comparable customer adoption, AI-specific revenue, product-level margins, or independently audited performance across identical configurations.

A broader enterprise contest

Arista’s broader platform competes with established enterprise-networking vendors. Cisco’s May 2026 release reported year-over-year growth of more than 50% in networking product orders, more than 25% in campus networking orders, and more than 40% in data-center switching orders for its fiscal third quarter, ended 25 April. Those are Cisco-reported order-growth rates, not revenue growth or market share. Cisco also described a multi-year campus refresh cycle. [S5]

In November 2025, Cisco announced cloud-managed fabrics and Unified Branch capabilities spanning campus, branch, and industrial networks, alongside workflows across Meraki, Catalyst Center, Catalyst SD-WAN Manager, ISE, and Nexus. This was a product announcement, not evidence of realized customer adoption or superior performance. It nevertheless illustrates the breadth of the portfolio and operating tools a large incumbent can assemble. [S6]

Arista’s own filing notes that large competitors may bundle products, develop their own silicon, and combine network hardware and software. It also names white-box networks and open-source operating systems as alternatives. Suppliers, partners, and customers can become competitors. [S1, Part II, Item 1A, “We face intense competition,” pp. 32-33]

Momentum with exposure

Exhibit 1. Arista's reported quarterly results

USD millions, except percentages; fiscal quarters ended 30 June. Company-wide, not AI-specific.

Measure	Q2 2025	Q2 2026	Year-over-year change
Product revenue	\$1,877.0	\$2,605.2	+38.8%
Service revenue	\$327.8	\$430.5	+31.3%
Total revenue	\$2,204.8	\$3,035.7	+37.7%
GAAP gross margin	65.2%	62.9%	–2.3 percentage points

Source: [S1], Part I, Item 2, MD&A, “Results of Operations-Three and Six Months Ended June 30, 2026 Compared to Three and Six Months Ended June 30, 2025,” “Revenue, Cost of Revenue and Gross Margin,” p. 17. Figures are company-wide; gross-margin movement is not attributed to AI product mix. Arista said the margin decline was principally driven by a greater proportion of sales to large end customers, who generally receive higher discounts. [S1, MD&A, p. 18]

Arista disclosed that two end customers represented 16% and 26% of total revenue in fiscal 2025, compared with 15% and 20% in fiscal 2024. The filing does not name them in this discussion. Large orders may shift with customer capital budgets and project timing; customers can require lengthy product trials, testing, qualification, and acceptance. The company said AI demand had increased networking demand, but the long-term trajectory remained unknown and customer demand estimates were difficult to forecast. [S1, Part II, Item 1A, large-customer risk factor, p. 24; “The networking market is rapidly evolving,” pp. 30-31]

The company also identified merchant silicon and component supply as dependencies in its development and fulfillment process. It said that suppliers’ delays, shortages, or failure to meet technical requirements could force product redesign and delay launches. [S1, Part II, Item 1A, merchant-silicon and supplier risk factors, pp. 23-24]

Exhibit 2. Competing strategic demands

Emphasis	Potential strategic benefit	Publicly disclosed challenge
Deepen AI Ethernet leadership	Build on AI networking demand, new 1.6 Tbps fabrics, and customer qualification; develop differentiated performance and reliability for large clusters.	Large-customer concentration and discounts; uncertain AI cycle; qualification and acceptance risk; dependency on silicon, components, and manufacturing supply; competing integrated systems.
Broaden Centers of Data	Connect data-center, campus, and WAN products through EOS/NetDL; use VeloCloud capabilities to broaden enterprise coverage and potential customer relationships.	Less operating history in adjacent markets, longer enterprise/campus sales cycles, incumbent breadth and bundling, and no published evidence of cross-domain attachment or incremental platform economics.

Sources: [S1], “Our Strategy,” pp. 15-16, and Part II, Item 1A, pp. 30-35; [S3], VeloCloud acquisition discussion. Benefits are strategic possibilities, not established results.

The allocation question

One approach would direct the next product-cycle emphasis toward AI Ethernet: accelerate switch and software improvements, deepen customer engineering and qualification, and extend work with silicon and systems partners. This could capture an expanding opportunity and reinforce Arista's technical position. It would also increase exposure to a limited number of large customers, fast product transitions, and uncertain deployment schedules.

A second approach would place more emphasis on connecting Arista's data-center, campus, and WAN propositions. A broader customer base could reduce reliance on a few cloud and AI accounts and make EOS/NetDL more valuable across customers' networks. But winning enterprise deployments could require a different sales motion, more channel and support capability, and sustained investment before the company can demonstrate cross-domain adoption.

The disclosures do not quantify an available engineering budget, incremental return by product, or a demand forecast by market. A recommendation must therefore state what Arista should learn and which signposts should trigger adjustment, rather than treating current growth or management's platform claim as sufficient proof.

Discussion questions

1. What is Arista's distinctive capability in AI Ethernet, and which parts of the customer solution depend on suppliers, partners, or customer choices?
2. Does the company have evidence that its AI, data-center, campus, and WAN products reinforce one another? What evidence is missing?
3. How should customer concentration, discounts, qualification periods, and reported gross-margin movement affect the next product-cycle emphasis?
4. Which alternative would you recommend, and what near-term evidence would cause you to change course?

Sources

- **S1.** Arista Networks, Inc. Form 10-Q for quarter ended 30 June 2026, filed 5 August 2026, SEC accession 0001596532-26-000175. **SEC filing**. Cited locators identify Item, section, and printed page. Historical customer concentration percentages are fiscal-year figures; Q2 revenue and margin figures are quarter figures.
- **S2.** Arista Networks, “Arista Networks, Inc. Reports Second Quarter 2026 Financial Results,” 4 August 2026. **Company release**. Company statements and product announcements are attributed.
- **S3.** Arista Networks, Inc. Form 10-K for year ended 31 December 2025, filed 19 February 2026, SEC accession 0001596532-26-000013. **SEC filing**. Item 1, “Business,” VeloCloud acquisition discussion.
- **S4.** NVIDIA, “Spectrum-X - the Open, AI-Native Ethernet Fabric - Sets the Standard for Gigascale AI, Now With MRC,” 6 May 2026. **NVIDIA blog**. Vendor claims about product performance and deployments; no independent comparison supplied.
- **S5.** Cisco, “Cisco Reports Third Quarter Earnings,” 13 May 2026, fiscal quarter ended 25 April 2026. **Cisco investor-relations release**. Cisco company revenue and product-order disclosures are distinct; order growth is not revenue or share.
- **S6.** Cisco, “Cisco Supercharges its Secure Enterprise Network Architecture for the AI Era,” 3 November 2025. **Cisco Newsroom**. Product and architecture announcement, not a customer adoption or economic measure.