

Market Guide for Data Center Switching

16 March 2026 - ID G00840744 - 16 min read

By: Andrew Lerner, Jorge Aragon

Initiatives: Strategy, Risks and Opportunities; Reinvent I&O as an Enabler of AI Value;

Accelerate Enterprise AI Value Realization at Scale; Financial

The data center switching market has splintered, which creates new challenges for the majority of enterprise buyers. As vendors shift their focus to AI, most enterprises will experience higher costs, longer lead times and lower quality. Heads of I&O must take action to sustain reliable and automated data center networks.

Overview

Key Findings

- The market is diverging into two buying patterns: general-purpose data center networks and AI network fabrics. This is because AI workloads demand different network designs, features and performance compared to general-purpose data center buildouts.
- Vendors are aggressively pivoting resources toward AI network fabrics and the largest buyers. While this benefits organizations building AI network fabrics, it introduces substantial challenges for organizations buying general-purpose data center networks (including reduced software quality, slowed feature velocity, higher prices, and longer lead times).
- There is a wide degree of differentiation for AI network fabrics, including support for GPU-to-GPU connectivity within a single system, performance, scale and automation.
- For general-purpose data center buildouts, most vendors have “good-enough” hardware platforms. The differences between vendors reside in software, including their ability to deliver automation and agentic NetOps, and natively embedded security capabilities.

Recommendations

- For general-purpose data center workloads, deploy rightsized physical infrastructures using a leaf-spine design with fixed-form-factor switches and 25G/100G-capable interfaces.
- Prepare for elongated lead times, reduced feature velocity and degraded support, and higher prices, as vendors pivot their focus on AI buildouts and/or larger buyers.
- Build out separate, dedicated physical switching fabrics for AI/GPU-based workloads (i.e., AI network fabrics) to ensure performance and prevent GPU starvation.

Strategic Planning Assumptions

- Prior to 2027, more than half of data center switching spend will support AI workloads.
- Through mid-2027, hardware costs for enterprises deploying general-purpose data center switches will increase 15% to 40% compared to their 2025 levels.
- Through mid-2027, enterprises deploying general-purpose data center switches will experience lead times of three to nine months, up from one to two months in mid-2025.

Market Definition

This market covers data center network switches and the requisite management and automation platforms for them. Data center switches are Ethernet switches installed in a data center environment intended to provide connectivity for endpoints, including servers, firewalls, and Layer 4 through Layer 7 appliances and mainframes.

Data center switches provide foundational connectivity mostly for compute resources in the data center. This is required to enable applications in support of business requirements. Emerging use cases that drive investments on data center networks include both AI and edge compute workloads.

Mandatory Features

- Network operating system (NOS) software – operating system software meant for installation on a physical Ethernet switch that controls both hardware (i.e., LED, physical ports, fan power) and logical components (i.e., routing and switching).

- Fabric management — software that enables multiple Ethernet switches to be managed as a single construct (a single logical entity) with a common API. This includes automation, Day 1 provisioning and Day 2 operational tasks. Fabrics can be standards-based or proprietary in nature.
- Centralized life cycle management — software that includes role-based access control (RBAC) to enable life cycle management of switches and NOS. This includes an administrative UI that enables provisioning, configuration management, troubleshooting, analytics and reporting. (Note: Centralized management can be embedded with fabric management.) This also includes turnkey integrations — those with commonly deployed data center or automation technologies.

Common Features

- Physical fixed-form factor network switches suitable for deployment at the data center network access (i.e., leaf), and core (i.e., spine) layers that provide IEEE 802.3 Ethernet network connectivity and services to endpoints, including servers, storage, user devices and peripherals. Physical port speeds are typically 10Gbps or higher.
- Advanced physical capabilities such as (a) physical modular network switches suitable for deployment at data center network access (i.e., leaf) and core (i.e., spine) layers that provide IEEE 802.3 Ethernet network connectivity and services to endpoints, including servers, storage, user devices and peripherals; or (b) physical ultrahigh performance network switches, including IEEE 802.3 Ethernet switches designed for the most stringent performance requirements and workloads such as high-performance computing (HPC), high-frequency trading or AI.
- Multivendor fabric management — fabric management supporting NOS and/or physical hardware from multiple vendors.
- Agentless microsegmentation.
- Commercial support for the SONiC network operating system.
- Natural language, conversational UIs, and/or chatbots.
- Advanced UIs, including AI assistants and AI agents.
- Virtual switches and/or container switches that operate in conjunction with a broader virtualization/container platform and handle the switching responsibilities for virtual machines or containers.
- Physical interface transceivers.

- Predictive analytics, anomaly detection and automated root-cause analysis for the fabric implementation.

Market Description


Product Description

The core aspects of a vendor offering in this market include fabric management software, network operating system (NOS) and hardware platforms. Fabric management software allows multiple switches to be managed as a single logical construct (fabric) and is programmatically accessed via API. Fabric managers are also used to integrate with other data center systems such as virtualization, container and automation tools. The NOS provides the operating system for the switch, while the hardware platform provides physical interfaces to connect hosts and other infrastructure.

Figure 1 illustrates the core aspects of a vendor offering in this market.

Figure 1: Components of Data Center Switching Offerings

Components of Data Center Switching Offerings



Source: Gartner
840744

Gartner

Customer Adoption

Gartner estimates there are 100,000 customers with data center switches, distributed across 400,000 physical sites. The two main drivers for investment in this market are: (a) refresh of general-purpose data center networks to support applications and cloud environments, and (b) AI network fabrics to support AI infrastructure buildouts.

General-Purpose Workloads

This includes providing connectivity for general-purpose data centers that primarily support compute, storage, and virtualization. Client interest in general-purpose data center networks is high, while volume is flat year over year. ¹ The most common physical design is a two-tier leaf-spine topology. The primary buyers in this segment are enterprises, including both midmarket organizations and larger enterprises. Investments are predominantly driven by the need for infrastructure refreshes and new data center buildouts.

AI Workloads

AI network fabrics provide connectivity for AI systems, including GPU clusters. The primary buyers in this segment include hyperscalers, neocloud providers, AI-focused companies, governments, and a limited number of extra-large enterprises. Large-scale model training is typically carried out by hyperscalers and major technology companies with dedicated AI teams and massive GPU clusters that require ultra-high-performance networking. However, most organizations are not building nor fine-tuning large foundation models; they focus on running inference on pretrained models to extract business value from their data. As a result, enterprise AI infrastructure strategies are primarily driven by inference requirements rather than large-scale training workloads. See Table 1 below for a summary of data center switching buying patterns for general-purpose versus AI infrastructure buildouts.

Table 1: Key Data Center Switching Buying Patterns

(Enlarged table in Appendix)

↓	General-purpose buildouts ↓	AI infrastructure buildouts ↓
Buyer profile	Enterprises constitute over 90% of the installed base and over 70% of the revenue. Over 100,000 buyers.	A small number of large buyers including hyperscalers, neoclouds and AI companies constitute over 90% of the revenue. We estimate that 150 companies account for 95% of spend.
Client Interest ¹	Slight decline Y/Y	Massive growth Y/Y
Key buyer criteria	■ Hardware/software reliability ■ Vendor support quality ■ Vendor incumbency ■ Agility/automation ■ Security/microsegmentation ■ Cost	■ Integration with xPU vendors ■ Performance/bandwidth ■ Cost ■ Lead times ■ Scale
Network traffic characteristics	■ Large number of flows; flows have varying sizes ■ Low levels of packet loss have minimal impact on application performance. ■ Typical interface requirements are 10Gbps, 25Gbps and 100Gbps. ■ Mix of transmission control protocol (TCP) and user datagram protocol (UDP) traffic	■ Smaller number of large flows (elephant flows) ■ Higher sensitivity to packet loss ■ Typical interface requirements are 400Gbps, 800Gbps and higher ■ Primarily UDP, including remote direct memory access (RDMA)/RoCEv2/NVME-oF

Source: Gartner (March 2026)

Market Direction

The major trends driving the market include: AI network fabrics, agentic NetOps software (AI agents being infused into networking products), SONiC, and sovereignty.

Networking for AI (AI Network Fabrics)

AI network fabrics are structurally reshaping the vendor ecosystem, pricing, innovation cycles, and operational models across all data center switching. Further, by 2027, more than half of data center switching spend will support AI workloads, surpassing general-purpose infrastructure. This is because large buyers are building AI network fabrics, which include:

- GPUs within a single system, commonly referred to as “scale-up”
- GPU systems within a data center, commonly referred to as scale-out
- GPUs across data centers, sometimes referred to as scale-across and/or data center interconnect (DCI)

AI network fabrics have unique requirements compared to general-purpose data centers, which are driving rapid innovation and adoption of several technologies including:

- **DPU:** Function accelerator cards (FACs), also known as DPUs or SmartNICs, are specialized devices that offload network, security, and storage tasks from servers to improve performance and efficiency. They matter to buyers because they boost server speed, reduce costs, and enhance security by handling demanding functions independently.
- **Speed:** 800 Gigabit Ethernet has become the volume standard for AI spines, with 1.6 Terabit Ethernet adoption expected to align with next-gen GPU releases.
- **Optics and cooling:** As switch ASICs hit thermal limits, the market is redesigning physical form factors. We expect increased adoption of co-packaged optics (CPO) and linear pluggable optics (LPO) to reduce power consumption by up to 50% per port. Furthermore, liquid cooling integration is transitioning from a niche requirement to a necessity for high-performance switches or densely populated racks/cabinets where air cooling is not efficient.

AI for Networking (Agentic NetOps Software) Moving From “Cute” to Critical

Network operations are evolving from human-assisted dashboards and automation to include AI agents. Vendors are introducing autonomous agents that support self-healing, predictive optimization, and goal-driven configuration without human intervention. This represents a shift from human-assisted and IaC to a degree of self-healing and optimization. The state of AI for networking operations in the data center has evolved through distinct phases:

- In 2023, the “cute” era introduced AI chatbots with limited incremental value, primarily aiding documentation and basic troubleshooting.
- From now through 2027, the “core” phase will see AI assistants providing task-level autonomy, predictive recommendations, granular hardware and software insights, and one-click automations.
- By 2029, the “critical” phase will emerge, featuring AI agents that are capable of autonomous operation and goal-driven operations, multiagent collaboration, and proactive automation — even when humans are not present (see [The Future of NetOps is Agentic](#)).

SONiC

Gartner continues to observe consistent interest in SONiC from large organizations. SONiC is a modular, open-source network operating system (NOS) originally developed and open-sourced by Microsoft and now governed by the Linux Foundation. We estimate that over 2,500 organizations are running SONiC in production, and roughly one-third are enterprises.^{1,2,3} Networking vendors are increasingly offering enterprise-ready SONiC distributions, including their own hardened versions complete with commercial support, validated hardware compatibility, and integration into their fabric management software.

SONiC is slowly and incrementally evolving from a niche white-box solution to a viable option on high-performance switches with full life cycle support. Consequently, the disaggregation of hardware and software is no longer limited to cloud deployments; it is becoming a viable option for enterprises seeking flexibility, cost control, and reduced vendor lock-in, particularly within modern leaf-spine architectures and scale-out AI networks.

Infrastructure Sovereignty

Sovereignty has evolved from a background compliance concern in 2024 to a strategic priority, fueled by new EU regulations and growing distrust of U.S. vendors. This shift is driving organizations in China, Europe, Canada, Australia, Taiwan, and Singapore to build sovereign infrastructure. The desire for sovereignty is directly impacting the data center market, as organizations are increasingly looking at data center switches from non-U.S. and non-China suppliers, and/or opting for localized data center options over the U.S.-based hyperscalers for hosting workloads.

Vendor Landscape Shifts

Sovereignty and the race to deliver AI network fabrics is dramatically altering the vendor landscape. This includes new vendors entering the market, shifting vendor relationships and new industry consortiums. NVIDIA has effectively become the dominant networking player for delivering AI network fabrics. However, the market is so lucrative that many vendors are angling to get their piece of the pie. Vendors are attempting to accelerate alternatives to NVIDIA's proprietary NVLink dominance in "scale-up" connectivity. Specifically, UALink and Broadcom's Scale-Up Ethernet (SUE) are challenging by creating alternative standards to foster multivendor interoperability and reduce lock-in. Simultaneously, the Ultra Ethernet Consortium is driving Ethernet as the primary choice for back-end "scale-out" clustering, enabling it to displace InfiniBand by addressing critical packet loss and performance limitations.

Further, this is forcing traditional incumbents to adapt in unfamiliar ways. This is rewriting the rules of competition as vendors that used to directly compete head-to-head are now partnering and simultaneously competing. For example, NVIDIA, Dell Technologies and Cisco all compete for AI network fabrics with long-standing, well-established portfolios. However, Cisco now supports embedding NVIDIA Spectrum-X silicon in Nexus switches, while Dell both resells NVIDIA Spectrum-X and its own PowerSwitches. In addition, new vendors have emerged to address high-performance requirements, including Aria Networks, Upscale, Unifabrix, and Nexthop.

Market Analysis

Overview

Cisco drove enterprise data center networking in the late 1990s and early 2000s, while Arista Networks drove the market during the cloud era. Now, NVIDIA is poised to take over in the AI era. Specifically, AI network fabrics are emerging as a distinct market segment, and potentially its own market, as evidenced by a different set of buyers, differing requirements and a specialized set of suppliers.

AI Network Fabrics


AI network fabric is defined as a collection of data center purpose-built hardware and software to address AI workload communication requirements. For example, packet loss sensitivity, high-speed bursting and in-order packet delivery aren't typically met by traditional Ethernet, within or across GPU clusters. The key product areas associated with an AI network fabric include:

- **Scale-Up AI Fabrics (SAIF):** Used for connecting GPUs within a single system/rack to share memory. While NVIDIA's proprietary NVLink dominates, open standards like Ultra Accelerator Link (UALink) and Broadcom's Scale-Up Ethernet (SUE) are emerging to offer vendor-neutral alternatives.
- **Scale-Out Networks:** Used for connecting multiple GPU systems. The battle between InfiniBand and Ethernet is settling, with Ethernet projected to handle over 65% of GenAI workloads by 2029. The **Ultra Ethernet Consortium (UEC)** is driving specifications to make Ethernet lossless and performant enough to replace InfiniBand for most clusters.
- **Scale-Across (DCI):** A growing requirement for connecting data centers over the WAN to train massive models across distributed sites, driving demand for high-performance data center interconnects.

Figure 2 illustrates SAIF and scale-out network to interconnect GPUs for AI workloads.

Figure 2: Scale-Up AI Fabric (SAIF) and Scale-Out Network

Scale-Up AI Fabric (SAIF) and Scale-Out Network



Source: Gartner
840744

Gartner


General-Purpose Data Center Networking

The most common physical design we observe for new general-purpose deployments is a two-tier leaf-spine topology (sometimes referred to as a Clos topology). However, in extra-large environments (hundreds of switches in the same location), an additional switching tier above the spine is often deployed, sometimes referred to as a super spine. In smaller environments (500 virtual machines [VMs] or fewer), a two-switch data center (with a shared campus core) is often optimal, although we see few vendors lead with this approach.

Figure 3 illustrates a leaf-spine topology; including underlay and overlay networks, which has become the standard on data centers.

Figure 3: Clos Topology and Building Blocks of Spine-Leaf Fabrics

Clos Topology and Building Blocks of Spine-Leaf Fabrics



Source: Gartner

IDS/IPS: intrusion detection system/intrusion prevention system; HCI = hyper-converged infrastructure

818649_C

Gartner

Prepare for the AI “Squeeze”: Nearly all data center networking vendors are dramatically shifting their focus toward AI infrastructure buildouts. Vendors are in a “race” to grab share as hyperscalers, neoclouds, AI companies, and a small number of enterprises are deploying AI in their data centers. Consequently, enterprises buying switches for non-AI workloads face a “squeeze,” whereby we expect:

- Software quality and feature velocity will decline.
- Support will decline
- Hardware lead times will increase by three to nine months
- Hardware prices will increase 15% to 40% beyond typical levels

As a result, organizations must plan for extended lead times when purchasing general-purpose data center switches. This means either getting “to the front of the line” via placing orders over six months earlier than normal and/or “extending the life of assets” longer than usual, via avoiding purchases through the end of 2027. This already requires adjusting budgets to account for price increases.

Representative Vendors

The vendors listed in this Market Guide do not imply an exhaustive list. This section is intended to provide more understanding of the market and its offerings.

The vendors listed in this Market Guide do not imply an exhaustive list. This section is intended to provide more understanding of the market and its offerings.

The vendors identified in Table 2 include those who are either (a) continuing to invest in the market, and/or (b) relevant to Gartner clients as evidenced by mentions in inquiry.

Vendor Selection

Table 2: Representative Data Center Switching Vendors

(Enlarged table in Appendix)

Vendor	Headquarters	Product(s)
Alcatel-Lucent Enterprise	France	OmniSwitch Series, AOS, OmniVista
Aria Networks	U.S.	Aria Switch
Arista Networks	U.S.	7000 Series, EOS, CloudVision
Arrcus	U.S.	ArcOS, ACE-AI, ArcIQ
Asterfusion Data Technologies	China	AsterNOS, CX Series, X Series
Aviz Networks	U.S.	ONES, Aviz Certified Community SONiC, Broadcom Enterprise SONiC
Cisco	U.S.	Nexus Dashboard, Nexus 9000, Cisco 8000, Hyperfabric
Dell Technologies	U.S.	PowerSwitch Series, SmartFabric Manager, Enterprise SONiC
DriveNets	Israel	Network Cloud-AI
Edgecore Networks	Taiwan	AIS, APS, DCS switches, Open SONiC
Extreme Networks	U.S.	Extreme 8000/SLX Series, IP Fabric, ExtremeCloud Orchestrator
Fortinet	U.S.	FortiSwitch
FS	U.S.	N Series Switches, PicOS
H3C	China	S Series Switches, Comware, SeerEngine
Hedgehog	U.S.	Open Network Fabric
HPE (includes Juniper and Aruba)	U.S.	CX Series, AOS-CX, Fabric Composer, QFX/PTX Series, Junos, Apstra
Huawei	China	CloudEngine switches, iMaster NCE-Fabric
IP Infusion	U.S.	OcNOS, IP Maestro
Netris	U.S.	Netris Controller
Nexthop	U.S.	No specific product name.
Nokia	Finland	7000 Series, SR Linux, Event-Driven Automation (EDA)
NVIDIA	U.S.	Spectrum Switches, Cumulus Linux, NetQ
Ruijie Networks	China	RG-N Series Data Center Switches
Unifabrix	Israel	MAX
Upscale AI	U.S.	No specific product name.

Source: Gartner (March 2026)

Market Recommendations

- For general-purpose workloads, prefer fixed-form-factor switches using a leaf-spine design with fixed-form-factor switches and 25G/100G-capable interfaces for traditional workloads.
- For general-purpose workloads, prioritize the management and automation capabilities of the fabric management software rather than hardware characteristics when selecting a vendor.
- Adjust data center network strategies for general-purpose workloads to mitigate risk associated with extended lead times, higher costs and reduced software quality and support as vendors shift their focus toward large AI buildouts.
- For AI workloads, build out separate AI network fabrics; don't reuse existing switching capacity. Use SAIF for connectivity within a single AI "system" and Ethernet-based switches to "scale-out" when connecting multiple AI systems together.
- For AI network fabrics, prefer vendors that enable unified management and visibility of the full-stack including NICs.
- For AI workloads, prefer Ultra Ethernet Consortium-based products when they become generally available, which we expect in 2026.
- For SONiC-based deployments, build out a separate leaf-spine topology and then migrate the workloads to the new fabric since SONiC is not meant to be a like-for-like replacement of proprietary data center switches.
- For security-focused enterprises, prioritize vendor offerings that include security and segmentation capabilities directly at the fabric level, to rely less on third-party firewalls or agents on servers.

Evidence

¹ Gartner has engaged in over 500 interactions with clients and/or prospects on the topic of data center networking in 2025.

² Gartner analysts are regularly briefed by data center networking vendors.

³ Gartner analysts regularly review publicly available information on the market including blogs, vendor spec sheets, quarterly financial statements, etc.

Recommended by the Authors

Some documents may not be available as part of your current Gartner subscription.

[First Take: NVIDIA Dramatically Alters the Data Center Networking Landscape](#)

[What are "Scale-Up" AI Fabrics and Why Should I Care?](#)

[The AI Race Puts Your Network at Risk](#)

[Innovation Insight: Agentic NetOps Software Redefines Networking](#)

© 2026 Gartner, Inc. and/or its affiliates. All rights reserved. Gartner is a registered trademark of Gartner, Inc. and its affiliates. This publication may not be reproduced or distributed in any form without Gartner's prior written permission. It consists of the opinions of Gartner's Business and Technology Insights Organization, which should not be construed as statements of fact. While the information contained in this publication has been obtained from sources believed to be reliable, Gartner disclaims all warranties as to the accuracy, completeness or adequacy of such information. Although Gartner insights may address legal and financial issues, Gartner does not provide legal or investment advice and its insights should not be construed or used as such. Your access and use of this publication are governed by [Gartner's Usage Policy](#). Gartner prides itself on its reputation for independence and objectivity. Its insights is produced independently by its Business and Technology Insights Organization without input or influence from any third party. For further information, see "[Guiding Principles on Independence and Objectivity](#)." Gartner insights may not be used as input into or for the training or development of generative artificial intelligence, machine learning, algorithms, software, or related technologies.

Table 1: Key Data Center Switching Buying Patterns

↓	General-purpose buildouts ↓	AI infrastructure buildouts ↓
Buyer profile	Enterprises constitute over 90% of the installed base and over 70% of the revenue. Over 100,000 buyers.	A small number of large buyers including hyperscalers, neoclouds and AI companies constitute over 90% of the revenue. We estimate that 150 companies account for 95% of spend.
Client Interest ¹	Slight decline Y/Y	Massive growth Y/Y
Key buyer criteria	■ Hardware/software reliability ■ Vendor support quality ■ Vendor incumbency ■ Agility/automation ■ Security/microsegmentation ■ Cost	■ Integration with xPU vendors ■ Performance/bandwidth ■ Cost ■ Lead times ■ Scale
Network traffic characteristics	■ Large number of flows; flows have varying sizes ■ Low levels of packet loss have minimal impact on application performance. ■ Typical interface requirements are 10Gbps, 25Gbps and 100Gbps.	■ Smaller number of large flows (elephant flows) ■ Higher sensitivity to packet loss ■ Typical interface requirements are 400Gbps, 800Gbps and higher

↓	General-purpose buildouts ↓	AI infrastructure buildouts ↓
	■ Mix of transmission control protocol (TCP) and user datagram protocol (UDP) traffic	■ Primarily UDP, including remote direct memory access (RDMA)/RoCEv2/NVME-oF

Source: Gartner (March 2026)

Table 2: Representative Data Center Switching Vendors

Vendor	Headquarters	Product(s)
Alcatel-Lucent Enterprise	France	OmniSwitch Series, AOS, OmniVista
Aria Networks	U.S.	Aria Switch
Arista Networks	U.S.	7000 Series, EOS, CloudVision
Arrcus	U.S.	ArcOS, ACE-AI, ArcIQ
Asterfusion Data Technologies	China	AsterNOS, CX Series, X Series
Aviz Networks	U.S.	ONES, Aviz Certified Community SONiC, Broadcom Enterprise SONiC
Cisco	U.S.	Nexus Dashboard, Nexus 9000, Cisco 8000, Hyperfabric
Dell Technologies	U.S.	PowerSwitch Series, SmartFabric Manager, Enterprise SONiC
DriveNets	Israel	Network Cloud-AI
Edgecore Networks	Taiwan	AIS, APS, DCS switches, Open SONiC
Extreme Networks	U.S.	Extreme 8000/SLX Series, IP Fabric, ExtremeCloud Orchestrator
Fortinet	U.S.	FortiSwitch
FS	U.S.	N Series Switches, PicOS

H3C	China	S Series Switches, Comware, SeerEngine
Hedgehog	U.S.	Open Network Fabric
HPE (includes Juniper and Aruba)	U.S.	CX Series, AOS-CX, Fabric Composer, QFX/PTX Series, Junos, Apstra
Huawei	China	CloudEngine switches, iMaster NCE-Fabric
IP Infusion	U.S.	OcNOS, IP Maestro
Netris	U.S.	Netris Controller
Nexthop	U.S.	No specific product name.
Nokia	Finland	7000 Series, SR Linux, Event-Driven Automation (EDA)
NVIDIA	U.S.	Spectrum Switches, Cumulus Linux, NetQ
Ruijie Networks	China	RG-N Series Data Center Switches
Unifabrix	Israel	MAX
Upscale AI	U.S.	No specific product name.

Source: Gartner (March 2026)