The 1.6TbE Transition Will Reshape Data Center Architecture from the Ground Up

Moving to 1.6 Terabit Ethernet Forces Infrastructure Teams to Rethink Physical Topology, Power Delivery, and Cooling Systems Simultaneously – and Most Facilities Are Not Ready

Analyst(s): Tom Hollingsworth
Publication Date: August 19, 2026
Document #: AIOTH202608

Key Points

- The shift from 800GbE to 1.6TbE networking is not merely a bandwidth upgrade but an architectural inflection point. It demands redesigning physical rack layouts, switch placement, and cable plant topology to accommodate latency and signal integrity requirements at these speeds.

- Rack power densities now exceeding 80kW per cabinet make air cooling physically impossible for AI training clusters. This forces operators to deploy liquid cooling infrastructure that fundamentally changes facility plumbing, safety systems, and maintenance workflows.

- Traffic dynamics in 1.6TbE fabrics shift toward east-west dominance with tighter latency budgets, requiring switches to be positioned physically closer to GPU accelerators – a constraint that upends traditional top-of-rack and end-of-row network design philosophies.

Recommendations

1. Vendors delivering 1.6TbE switching platforms must provide clear co-packaged optics (CPO) roadmaps with explicit compatibility guidance for existing pluggable infrastructure, enabling data center operators to plan phased transitions without stranding cable plant investments.

2. Vendors offering liquid cooling solutions should integrate thermal management planning tools directly with network topology design software, allowing infrastructure teams to co-optimize switch placement, coolant distribution, and cable reach in a single workflow.

3. Vendors building GPU-dense server platforms must publish validated reference architectures that account for 1.6TbE cable reach limitations, power distribution at

80kW+ rack densities, and cooling zone requirements to reduce deployment risk for enterprise operators. Customers need to be ready for architectural changes in their data center.

What You Need to Know

- The IEEE 802.3dj standard ratifies 1.6TbE specifications using 200Gbps-per-lane signaling across eight lanes, doubling per-port throughput compared to 800GbE deployments now entering production.
- Current-generation AI training racks built around NVIDIA GB200 NVL72 and comparable platforms are pushing rack power consumption above 80kW, well beyond the thermal dissipation capability of any air-cooled facility design.
- CPO technology is expected to become the dominant interconnect approach for 1.6TbE and beyond, replacing traditional pluggable transceivers to reduce power consumption and heat generation within switch ASICs.
- Major switch ASIC vendors, including Broadcom and Marvell, have announced 1.6TbE silicon with initial sampling in late 2026 and volume production targeted for 2027.
- Hyperscalers and AI-focused colocation providers are already redesigning floor plans around shorter cable runs and distributed switching topologies to prepare for 1.6TbE deployment constraints.

Analysis

The networking industry is approaching a transition that will prove far more disruptive than any previous speed upgrade. Moving from 800GbE to 1.6TbE is not a simple doubling of bandwidth. It is an architectural reset that forces data center designers to reconsider physical topology, power delivery, thermal management, and traffic engineering simultaneously. The compounding effects of these requirements will strain organizations that treat this as a straightforward equipment refresh.

The Physics Problem Changes Everything

At 200Gbps per lane, signal integrity becomes exceptionally unforgiving. Copper reach distances shrink, optical loss budgets tighten, and the electrical interfaces between switch ASICs and transceivers become a primary source of power dissipation and heat. This is precisely why CPO has emerged as a critical enabling technology. As Futurum's earlier analysis on CPO noted, AI workloads are driving exponential bandwidth demand that outpaces the capabilities of current pluggable optics, and future scaling requirements simply cannot be met with existing technology[1]. The implication for 1.6TbE is direct: switches must be redesigned

[1] Co-Packaged Optics: The Key to Unleashing AI Networking's Full Potential

around optical engines integrated into the package substrate, fundamentally altering the form factor, thermal profile, and serviceability model that network operators have relied upon for two decades. This physical constraint cascades into topology design. When cable reach is limited and optical power budgets are tight, switches cannot sit in centralized aggregation rows hundreds of meters from compute. They must move closer to GPUs – in some emerging designs, directly into the compute rack or an immediately adjacent networking pod. This shatters the traditional three-tier network hierarchy and demands a far more distributed switching architecture where hundreds of smaller switch elements replace a handful of chassis-based aggregation platforms.

Power and Cooling Reach a Breaking Point

The power challenge at 1.6TbE speeds compounds an already severe problem. According to Futurum's 1H2026 Data Center Semiconductor Decision Maker Survey, power and cooling availability ranks as the third-largest constraint in scaling data center compute, cited by nearly 15% of respondents, while networking lead times, which will only worsen during a generational transition, represent the second-largest bottleneck at nearly 17%[2]. These constraints are not independent; they amplify each other. A rack drawing 80kW or more generates thermal loads that no air-cooled design can dissipate, mandating liquid-cooling infrastructure that, in turn, requires significant facility investment in plumbing, pumps, coolant distribution units, and leak detection systems. Futurum's research on data center energy constraints found that the massive capital expenditure on AI infrastructure faces a structural power generation gap because new grid-connected power generation cannot come online quickly enough to keep pace with data center construction[3]. When 1.6TbE switches add their own power and thermal contributions on top of already power-hungry GPU racks, facilities designed around 15–20kW per rack face wholesale reconstruction. Cooling system limits are already identified by 7.5% of survey respondents as the primary factor limiting AI cluster expansion, and grid interconnection and utility capacity constraints affect another 16.6%. The 1.6TbE transition will intensify both of these pressures.

Traffic Dynamics Demand New Thinking

The traffic profile within 1.6 TbE fabrics differs qualitatively from that of previous generations. AI training workloads generate massive east-west traffic flows between GPU nodes during gradient synchronization, with completion time directly tied to tail latency across the fabric. The 1H2026 Data Center Semiconductor Decision Maker Survey reveals that 55% of organizations now deploy hybrid InfiniBand-plus-Ethernet fabrics for AI training, while roughly 20% rely on Ethernet RoCE alone[4]. As 1.6TbE Ethernet matures, it will pull workloads away from InfiniBand, but only if the fabric can deliver comparable latency characteristics. This requires not just faster

[2] AI Chipsets DM: Bottlenecks (1H2026)
[3] AI Grid Constraints Will Push Over 33% of Data Centers Off-Grid by 2030
[4] AI Chipsets DM: Networking (1H2026)

ports but a fundamentally different topology with fewer hops, shorter physical paths, and deterministic congestion behavior. The deployment timeline data reinforces the urgency:

- 34% of organizations plan large-scale 800G or 1.6T networking deployments within the next six months
- 37% target the seven-to-twelve-month window
- Nearly 14% report current deployment.

This compressed timeline means infrastructure teams have months, not years, to resolve the architectural challenges that 1.6TbE imposes. Organizations deploying AI workloads in their own data centers, representing 36% of respondents according to survey data[5], face the most acute challenge because they cannot rely on hyperscaler-managed infrastructure to absorb complexity.

The Organizational Challenge

Perhaps the most underappreciated dimension of the 1.6TbE transition is organizational. Traditionally, network engineers, facilities teams, and server architects operated in relatively independent domains. At 1.6TbE, their decisions become deeply interdependent. Switch placement affects cable reach, which constrains rack layout, determines cooling zone design, and dictates power distribution topology. No single team can make decisions in isolation. Organizations that maintain siloed infrastructure planning will find themselves trapped in iterative redesign cycles as each team's choices invalidate another's assumptions. The vendors best positioned to capture value in this transition are those offering integrated design tools and validated reference architectures that span networking, compute, power, and cooling. Point solutions, such as a faster switch here or a better optic there, will not suffice when the challenge is systemic. Data center operators should demand that their networking vendors articulate not just port-speed roadmaps but complete architectural blueprints that account for the physical realities of deploying 1.6 TbE at scale.

What to Watch

- Will CPO achieve production readiness aligned with 1.6TbE switch ASIC volume timelines, or will CPO delays force operators into suboptimal pluggable designs with higher power and shorter reach?

- How quickly will liquid cooling infrastructure vendors scale manufacturing capacity to meet simultaneous demand from compute and networking equipment deployments?

- Will hyperscalers publish open reference architectures for 1.6 TbE fabric topologies, accelerating enterprise adoption, or will designs remain proprietary competitive advantages?

---

[5] AI Chipsets DM: Deployment Method (1H2026)

- Can Ethernet-based fabrics close the latency gap with InfiniBand at 1.6TbE speeds, or will AI training workloads remain split across dual fabrics for the foreseeable future?
- Which switch ASIC vendors will deliver production 1.6TbE silicon first, and will initial availability favor hyperscaler custom designs over merchant silicon available to enterprises?

Other Insights from Futurum

Co-Packaged Optics: The Key to Unleashing AI Networking's Full Potential — Tom Hollingsworth, February 2026 (Document #: AIOTH202601)

AI Grid Constraints Will Push Over 33% of Data Centers Off-Grid by 2030 — Brendan Burke, Nick Patience, Olivier Blanchard, March 2026 (Document #: AIOBBNPOB202603)

About Us

About the Authors

Tom Hollingsworth, Futurum Research Advisor, covers communications networks for The Futurum Group. He is an expert in data center networking technologies with over 20 years of experience. He is CCIE Emeritus #29213. He helps clients understand the technologies behind networking and how it can enable technology shifts toward new advances. He can be reached at thollingsworth@futurumgroup.com.

About The Futurum Group

Every day, The Futurum Group's analysts, researchers, and advisors help business leaders worldwide anticipate tectonic shifts in their industries and leverage disruptive innovation. Unlike traditional analysts, The Futurum Group works not only in analysis and research but also takes that insight and knowledge even further, engaging all the way through the go-to-market process.

Futurum Research provides in-depth research and insights on global technology markets using advisory services, custom research reports, strategic consulting engagements, digital events, go-to-market planning, and message testing. It also creates, distributes, and amplifies rich media content that all stakeholders read, watch, and listen to.

See more details on The Futurum Group at futurumgroup.com.

Copyright & Use License