

Technology, Media & Telecommunications Practice

The \$7 trillion data center build-out: How industrials can capture their share

Capital is pouring into data center development, but there are real constraints on growth. Incumbents can't meet demand for power and thermal equipment, creating room for new codesign entrants.

This article is a collaborative effort by Maria Goodpaster, Mark Patel, Pankaj Sachdeva, and Shih-Yung Huang, with Haley Chang and Wendy Yu, representing views from McKinsey's Industrials and Technology, Media & Telecommunications Practices.



The race to scale AI has triggered one of the largest infrastructure build-outs in modern history. By our estimates, [global spending on data centers could reach \\$7 trillion](#) by 2030. Whether a build-out is successful depends on many nuances, including the availability of capital and energy resources. But one issue is emerging as a critical growth constraint: the traditional operating models of industrial equipment suppliers that manufacture many of the long-lead critical components in the data center value chain. Many incumbents developed their processes serving utilities and other regulated industries with long planning cycles and predictable demand—an approach now misaligned with the speed and scale required in today’s data center market.

In this article, we use “equipment” to mean the full set of electrical, thermal, and mechanical infrastructure required to build and operate a data center—across both white space and gray space¹—excluding semiconductors. Suppliers of this equipment, often called original equipment manufacturers, or OEMs, are central to the build-out. They provide the transformers, switchgear, cooling systems, and other critical equipment that keep data centers operational; many are experiencing record demand and historically high valuations as a result. But even as revenues accelerate, many OEMs are struggling to keep pace—not just with delivery volumes, but with innovation timelines. Few can integrate fast-changing chip designs into their new products fast enough.

This lag in traditional industrial output is creating an opening for adjacent players in sectors such as automotive and aerospace: Drawing on their manufacturing prowess and expertise in electronic and thermal components, they could deliver the equipment that data centers need to scale—and step into a multitrillion-dollar market that is expanding faster than the traditional supply base can serve. But traditional industrial equipment makers that act decisively can still defend—and grow—their share of this expanding market before new entrants gain a foothold.

What’s breaking legacy data center supplier models?

We see three main forces shaping a shake-up in the industrial sector.

Power and critical components are now pacing items

Industrial equipment suppliers were not built to respond to today’s volatile data center demand (see sidebar “The high bar for data center equipment”). These companies’ planning cycles, procurement processes, and commissioning models evolved for steadier growth and more predictable utilization. Data centers are a massive opportunity for traditional suppliers, but the shape of demand—faster ramps, bigger step changes, and more simultaneous builds—challenges their long-held business models. Lead times in North America for critical equipment such as medium-voltage switchgear and transformers can now reach 80 weeks and 50 weeks, respectively—showing just how overtaxed industrial suppliers are (Exhibit 1).


For hyperscalers that purchase equipment to build data centers, delivery delays are not just an inconvenience but also a strategic risk. When power access units, transformers, switchgear, cooling systems, or the skilled labor needed to install, test, and commission this equipment become roadblocks, the entire build-out slips—potentially stranding billions in deployed capital.

¹White space is the secure, raised-floor area where IT equipment—servers, storage, and networking hardware—is actively housed. Gray space is the back-end area for facility infrastructure, such as generators, switchgear, chillers, and UPS systems, that support the white-space operations.

Exhibit 1

Lead times for critical equipment are growing, reaching more than 50 weeks in some cases.

Reported lead times for selected data center power equipment in North America,¹ weeks



¹Lead-time definitions by product: 3-phase transformers, UPS, and busway reflect catalog lead times, while MV switchgear and LV switchgear reflect engineered-to-order lead times. This distinction was made based on expert interviews indicating which definition is more commonly used for specific products. Source: McKinsey expert interviews with buyers at hyperscaler and colocation providers

McKinsey & Company

Equipment delays slow build-outs in the short term and can also derail hyperscalers' long-term business objectives. Thus, hyperscalers need to know if their equipment suppliers can reliably deliver on time and at scale.

The high bar for data center equipment

AI is reshaping data centers from collections of servers into tightly integrated power-and-thermal systems built to support high-density, highly variable compute. Several shifts define this transition—and raise the bar for equipment suppliers:

- *Liquid cooling becomes standard.* As rack densities increase, air cooling alone is no longer sufficient for many AI workloads. Liquid cooling has moved from a niche solution to a baseline requirement in new builds, introducing stricter qualification expectations across pumps, cold plates, connectors, sensors, and monitoring systems. For suppliers, this means proving long-term reliability, serviceability, and leak tolerance at scale—not just peak thermal performance.
- *More integration and testing happen before deployment.* Prefabricated modular data center designs increasingly integrate and test power and cooling systems in controlled factory environments or standardized modules before on-site installation. This shift enables faster and more predictable deployment, but it also pushes more system-level responsibility upstream to equipment providers. Suppliers are now expected to deliver components that fit cleanly together, work reliably as part of a larger system, and require minimal rework during installation.
- *Power systems are designed for dynamics, not averages.* AI workloads create sharp, rapid swings in power demand rather than steady-state loads. As a result, data center power architectures must absorb frequent spikes and drops without compromising uptime or equipment life. This has increased demand for tighter controls, faster response times, and, in many cases, on-site buffering and redundancy-raising expectations for both hardware performance and system integration.
- *Performance is managed across the life cycle.* Operators now manage large portfolios of data centers across regions and across different workloads, such as training and inference. Instead of optimizing infrastructure solely for initial efficiency, they increasingly tune performance over time to match power, cooling, and redundancy levels as workloads evolve. This shift emphasizes reliability, flexibility, and ease of modification over the asset's full life, rather than maximizing utilization at launch.

Silicon road maps dictate downstream innovations

Hyperscalers are increasingly selecting equipment suppliers based on whether they can prove a road map aligned with the future architecture, not just deliver products that meet today's specs.

Technical road maps for graphics processing units (GPUs) and custom silicon now set the rules for the entire supply chain. Because chip architecture defines power density, thermals, and interconnect requirements years in advance, equipment procurement no longer starts with an

RFP; it starts with data center builders soliciting suppliers that can engage at the semiconductor-engineering level. Hyperscalers and GPU providers use codified reference architectures from leading semiconductor companies to define the technical and operational requirements (that is, power density, cooling approach, interfaces) that suppliers must meet years before a shovel hits the ground. This landscape is further complicated by the need for suppliers to navigate multiple reference architecture ecosystems—each with distinct technical requirements and qualification pathways.

In other words, industrial manufacturers can no longer rely on delivering fully built equipment that slots into data center build-outs. They must instead design their products directly alongside semiconductor companies—iterating quickly and continuously to incorporate the latest chip designs.

The pace of data center innovation is accelerating

Suppliers today face another challenge: Data center architectures are evolving at a breakneck pace. Designs that once evolved over multiyear cycles are now refreshing annually—or even sooner—compressing the window for equipment suppliers to redesign, requalify, and scale their products based on these architectures.

The shift is visible in rising rack power densities, which are tightening the coupling between electrical and thermal systems. Transitions such as higher-voltage distribution—for example, 800 volts of direct current at the rack level—can drive redesigns across both IT equipment and facility infrastructure. Such changes can reshape how companies need to design, test, qualify, manufacture, and service products.

This acceleration resets how suppliers should approach new-product development and introduction (NPDI). Products can no longer be updated every few years; they must evolve fast enough to remain qualified across successive reference architecture generations—without sacrificing reliability or manufacturability. In practice, that compression is colliding with the already extended lead times described above. In response, hyperscalers are reserving production slots earlier and paying premiums for schedule certainty—while increasingly deprioritizing suppliers that cannot demonstrate aligned NPDI cadences and credible next-generation road maps (Exhibit 2).


What these three forces mean for suppliers: An operating-model reset

Many suppliers now understand that hyperscalers value reliability, specification compliance, and on-time delivery—but even those stringent demands are becoming table stakes. Suppliers must also demonstrate execution readiness and the ability to co-engineer products. To navigate this shift, industrial supplier incumbents can rebuild their operating models around three capabilities: an engineering-forward organization model, systems-level product development, and strategic investment to improve readiness.

Exhibit 2

Hyperscalers prioritize suppliers that are reliable, spec compliant, and fast.

The 12 most important key buying factors, as ranked by hyperscalers overall,¹ weighted net score



¹Question: Score each key buying factor on a scale of 1 to 10 when selecting a supplier (n = 15) (1 = not important, 10 = very important).
Source: McKinsey expert interviews with buyers at hyperscaler and colocation providers

McKinsey & Company

Organizational model: From sales-led motions to engineering-led engagement

Hyperscalers increasingly expect suppliers to participate early in the design phase of data center architectures, learn quickly through pilots, and take clear ownership of delivering complex systems on demanding timelines. Suppliers that want to earn that opportunity will need to establish an engineering-forward organizational model. They will need a cross-functional team to own the data center customer end to end—from early technical engagement and solution design through delivery, installation, testing, commissioning, and ongoing service.

According to our analysis, some suppliers overlay a lean data center market organization on top of existing product business units. This new organization leads front-end sales, engineering

engagement, and solution definition. Others move to a matrixed model, with a larger data center organization that unifies customer engagement, solutions, and services while products remain in product lines. A smaller set of suppliers has gone further, establishing a fully market-focused data center division with end-to-end P&L ownership (see sidebar “How equipment suppliers can adopt engineering-first organization models”).

Product development: From components to system integration

Data center customers increasingly favor equipment that fits cleanly within defined power, cooling, mechanical, and control environments.

The demand for plug-and-play equipment is driven in large part by the rise of prefabricated, modular data center design frameworks that require power, cooling, and control systems to be pre-integrated and tested in repeatable modules before deployment.

How equipment suppliers can adopt engineering-first organization models

Our analysis shows that incumbent equipment suppliers are redesigning their organizational models to better serve the pace and complexity of large-scale data center build-outs. This requires moving from a product- and function-led structure to an engineering-first account model.

In practice, suppliers are making this leap by creating data-center-specific team structures instead of organizing around product lines. They are hiring or reskilling dedicated talent for each data center account, mapping clearer ownership for product category sales, and reducing handoffs from design through commissioning. But this model only works if suppliers can ensure these teams clarify cross-functional accountability, improve issue resolution processes, and build commissioning expertise. When executed well, engineering-first organizational models allow equipment suppliers to design products that meet future-forward architecture requirements, align their own road maps more closely with customer priorities, and improve delivery predictability. Such an approach helps suppliers strengthen long-term partnerships with data centers and create competitive advantage.

“Data center build-outs are driving exciting growth for OEMs, but they’re also challenging the existing operating models of these companies. Customers demand higher-touch engagement, deep technical experts who can innovate alongside their teams, and innovation cycles at least two times faster than the current state. The pace and speed of innovation—and the agility required to meet it—are requiring suppliers to rapidly transform themselves to keep up.”

—Shih-Yung Huang, a McKinsey partner who advises industrial clients on strategic growth and commercial acceleration

As a result, equipment is now increasingly developed as integrated building blocks rather than as individual products. This means suppliers have to define technical interfaces between subsystems and collaborate early with semiconductor companies. In response, many incumbent industrial suppliers are attempting to bring synergistic talent under one roof through targeted M&A or partnerships—such as power equipment manufacturers partnering with cooling, controls, software, or automation providers to build system-level integration expertise.

Strategic investment: From efficiency to readiness

A supplier's manufacturing scale still matters to hyperscalers building data centers, but readiness to meet future architectures is becoming even more important. Suppliers that underinvest in validation and testing (such as high-voltage labs or AI-load-simulation risk) could miss qualification windows for data center build-outs. Readiness also depends on how suppliers expand capacity. Forward-thinking players are investing in digital tools and simplifying designs to accelerate production and increase output.

Talent is emerging as another readiness constraint. Experienced engineers and field service technicians are in limited supply and high demand. To improve talent readiness, industrial suppliers can partner with trade schools to build the pipeline, upskill existing teams, and, in some cases, acquire services organizations to bring field technicians in-house.

Finally, readiness now extends deep into the supply chain. Industrial suppliers are increasingly locking in access to critical tier-one and tier-two components well in advance, recognizing that shortages anywhere in the system—from high-voltage transformers to the electrical steel, breakers, and controls inside them—can delay delivery of entire solutions.

Action items for industrial CEOs

For industrial company CEOs who want to position their organizations as preferred partners for the next generation of data center build-outs, we recommend focusing on five priorities:

- *Establish joint product development and R&D with leading technology players.* Build direct co-engineering relationships with hyperscalers and GPU providers so your products are designed as reference architectures from the outset. The most strategic suppliers don't just respond to RFPs; they help shape them.
- *Build the ability to scale production quickly.* Ensure there is flexibility to move from pilot to volume production on hyperscaler timelines.
- *Organize around the customer, not the product line.* Stand up dedicated, cross-functional account teams that own the hyperscaler relationship end to end—from early technical engagement through delivery, commissioning, and service—rather than routing customers through siloed product organizations.

Find more content like this on the
McKinsey Insights App



Scan • Download • Personalize



- *Deploy capital fast but derisk it.* Stage gate investments to match demand signals, pursue innovative capital partnerships that reshape the risk profile, and diversify across adjacent industries to avoid overexposure to a single demand cycle.
- *Treat talent as a strategic investment, not an operational input.* Invest aggressively in workforce development—through trade school partnerships, upskilling programs, and targeted acquisitions of services organizations—to prevent field talent from becoming a binding constraint on growth.

The window for industrial suppliers to secure leading positions in the great data center build-out is narrowing. Those who move early—aligning with reference architectures, building qualification pathways, and delivering hyperscaler-level execution—can capture a disproportionate share of the trillions of investment at stake. Those that wait for the RFP may find the opportunity has already been allocated—without them.

[Maria Goodpaster](#) is a partner in McKinsey's Seattle office; [Mark Patel](#) is a senior partner in the Bay Area office, where [Haley Chang](#) and [Wendy Yu](#) are consultants; [Pankaj Sachdeva](#) is a senior partner in the Philadelphia office; and [Shih-Yung Huang](#) is a partner in the Southern California office.

This article was edited by Kristi Essick, an executive editor in the Bay Area office.

Copyright © 2026 McKinsey & Company. All rights reserved.