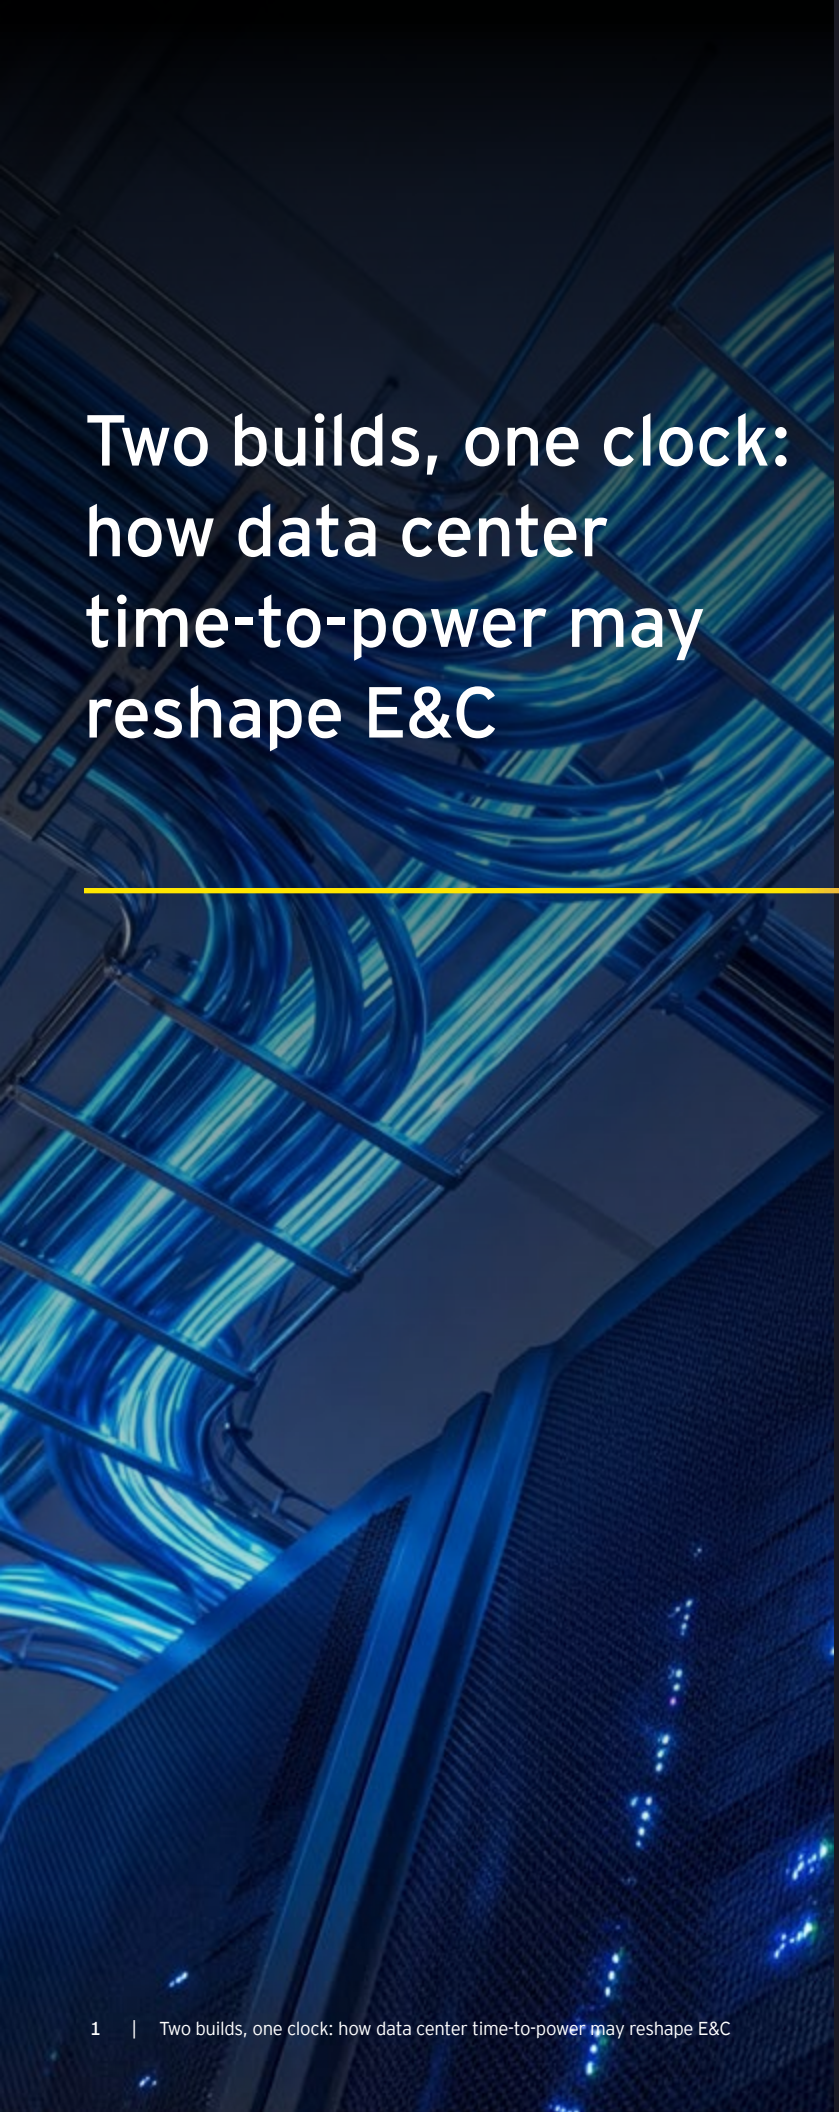
Two builds, one clock: how data center time-to-power may reshape E&C



The better the question. The better the answer.
The better the world works.

EY
Shape the future with confidence



Two builds, one clock: how data center time-to-power may reshape E&C

Data center construction in the US can finish years before power arrives. Boards should measure that exposure now.

In brief

- Completed private-side capital can sit idle for years. Energization, not construction completion, determines when it can begin generating a return.
- Federal reform of large-load connections is under way, but there is no national dataset for them. The generation queue, the nearest proxy, ran to a 2025 median of more than five years.
- A two-build integrator can own the interface between the campus and private-power scopes, but it cannot control utility timelines. Any risk transfer must therefore be bounded.
- Boards should quantify backlog tied to third-party power dates, assign interface ownership and stage investment in power-side capability as evidence emerges in requests for proposals (RFPs), pricing and win rates.

Two builds, one clock

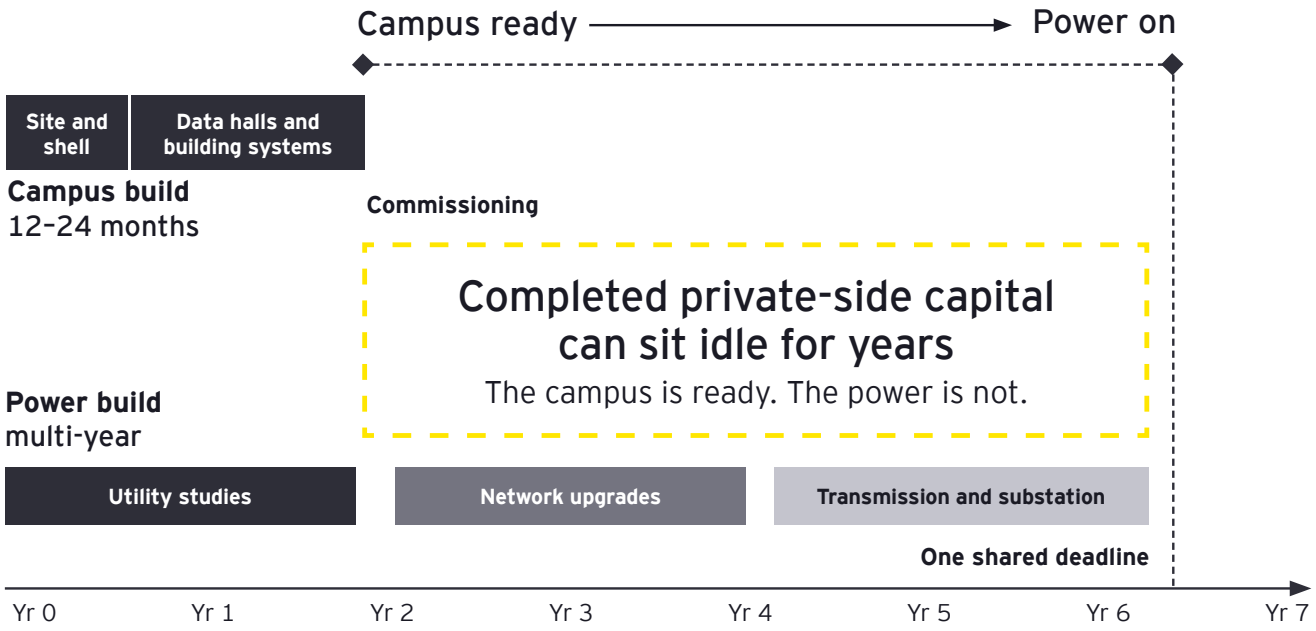
As of late 2025, a roughly 430,000-square-foot data center building in Santa Clara, designed for 48 MW of critical IT load, remained a vacant shell awaiting full energization. The local utility said a broader US$450 million system upgrade, serving that campus and nearby facilities as well as other customers, was scheduled for 2028, while the operator was seeking the remaining power sooner.[1]

A facility with power usage effectiveness above 1.0 requires more than 1 GW at the meter for every 1 GW of IT load, and that power must arrive on nearly the same schedule as the building itself. Figure 1 shows how the campus and power schedules diverge. A first data center facility or campus phase can be delivered in an illustrative 12 to 24 months, while the connection, substation and upstream power works needed to energize it can take several years longer. The campus build is the private side of the project; the power build is everything needed to energize the campus. The two are often contracted and procured separately, so the interface between them absorbs delay that no single party fully controls. This article takes a US market perspective and draws on UK evidence for comparison.

Load interconnection is the process of connecting a large electricity consumer, such as a data center campus, to the grid. Generation interconnection is the parallel process for connecting new supply. Load interconnection has historically been a state and local matter, so there is no long-running national dataset comparable to the generation queue. Generation interconnection data show how much the grid-development clock has lengthened. These queues measure generation and storage projects, not data center load connections, but they reflect the grid-development cycle on which large loads also depend. Where Lawrence Berkeley National Laboratory (LBNL) has duration data, the median project reaching commercial operation in 2025 waited more than five years from its interconnection request, nearly three times the 22-month median in 2008.[2] More than 2,000 GW of generation and storage sat in active transmission interconnection queues at year-end 2025, though most requested capacity is ultimately withdrawn. The total active queue shrank about 10% year over year as withdrawals outpaced new requests: volume fell while the wait lengthened. For boards, queue volume matters less than the lengthening grid-development cycle.[3]

Figure 1 — Two clocks, one deadline

The campus is built in months. The power that switches it on can take years.



Source: EY schematic (conceptual). Build durations are illustrative ranges.

In its February 2026 global outlook, a data center research firm estimated that 30% to 50% of capacity slated to come online in 2026 could be delayed, with power and construction constraints among the causes. Its tracker covers 777 facilities larger than 50 MW, totaling 190 GW, that have been announced since 2024.[4]

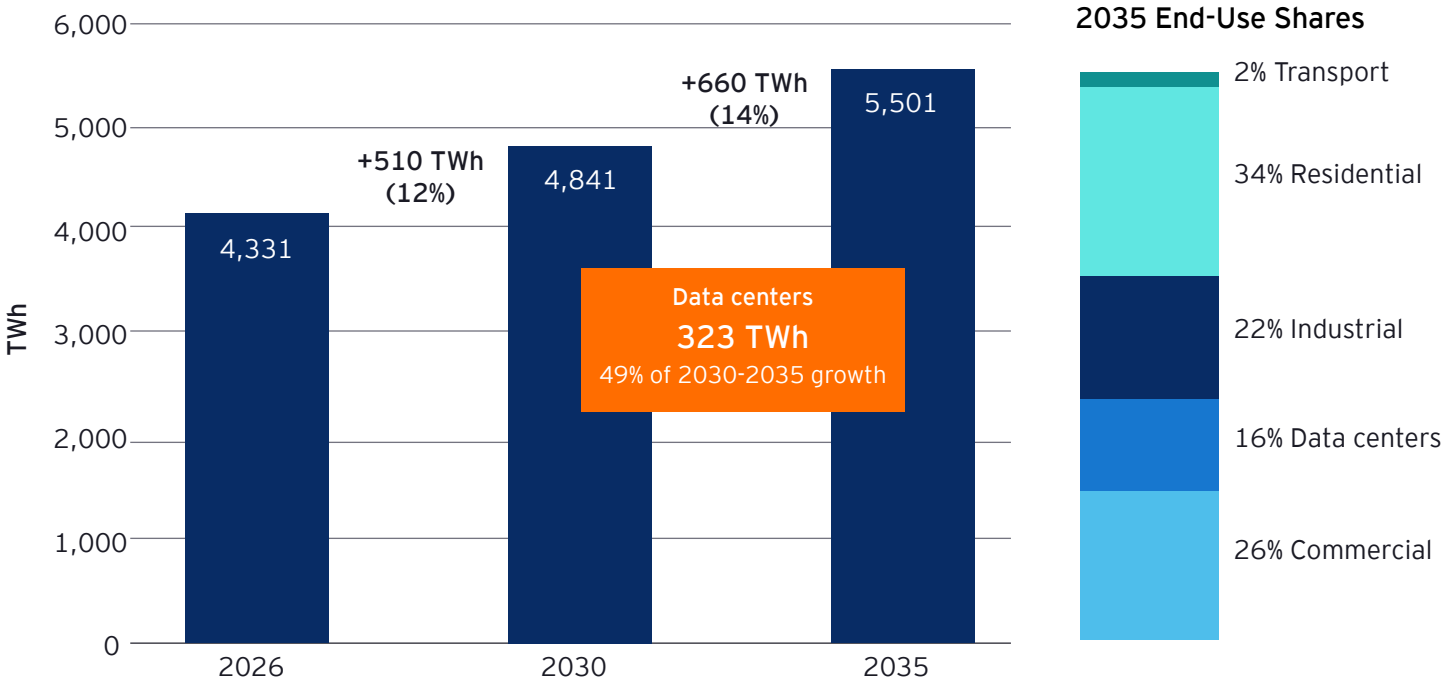
This is no longer only a utility problem. The divergence between the two builds is increasingly relevant to the engineering and construction (E&C) backlog. In the EY current-trajectory model, US electricity demand rises about 12% from 2026 to 2030 and a further 14% from 2030 to 2035. Data centers account for about 49% of the increase in that later period, making the sector the single largest contributor to demand growth. Figure 2 shows where that growth comes from.[5]

In one stylized 1 GW model, servers and networking are two-thirds of a data center's up-front capital, with the remaining third composed of the facility, land and utility works. However, this remaining third gates the return of the project.[6] The interface between the campus and its power is the element that a delivery firm can act on. Project owners carry most of this schedule exposure today. Contractors absorb a growing share of it, in backlog they may have to resequence or hold open until someone else's power date arrives. The question is whether owners will retain that exposure or transfer a defined share and, if so, to whom.

This is no longer only a utility problem. The divergence between the two builds is increasingly relevant to the engineering and construction backlog.

Figure 2 — US electricity demand and sources of growth, 2026-35 (TWh)

Data centers are 49% of the 2030–35 growth but only 16% of 2035 consumption.



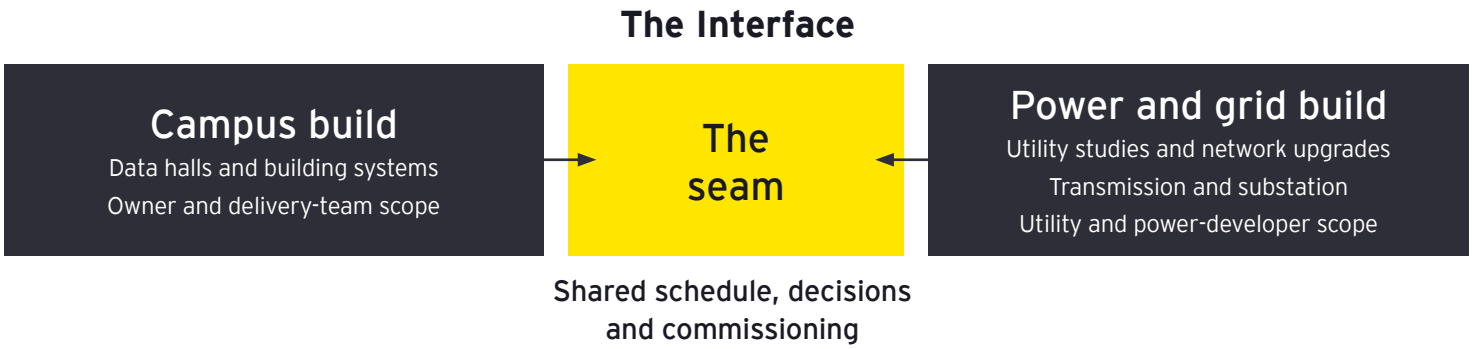
Source: EY analysis using the EY ERTA model, current-trajectory scenario. August 2026. Percentages may not sum due to rounding.

Where the interface breaks

Campus and utility schedules are usually managed separately, and neither side alone can manage the seam between them. The exposure appears at four recurring failure points, summarized in Figure 3.

Figure 3 — Where it breaks: the interface anatomy

The schedule risk lives in the seam between the two builds, not wholly within either one.



Four failure points — all in the seam

1	2	3	4
Misaligned dates	**Split procurement**	**Orphaned handoff**	**Divided authority**
Each contract follows a different clock.	Long-lead equipment falls out of sequence.	No single owner spans commissioning and energization.	Everyone sees the schedule; no one can enforce it.
The campus contractor works to the owner's energization date; the utility works to its own capital plan, queue and regulatory calendar. Neither side carries full responsibility for the other's slippage.	Separate procurement by campus developers, generation providers and utilities leaves transformers, switchgear and generation packages out of sequence with the delivery schedule.	Commissioning authority may sit with the campus team while final energization sits with the utility. Interconnection studies, transformer availability and transmission upgrades constrain that last step.	An owner may keep a combined master schedule, but its milestones carry no decision rights across the contract boundary. A schedule no one can enforce is simply a reporting tool; there is no control.

Source: EY schematic (conceptual). See note 7.

A view from the UK

The same pattern shows up in the UK, where operators and a hyperscaler described the gap to Parliament in 2026. One UK operator was quoted a 2037 connection date for its next tranche of grid power despite holding land and planning permission. In evidence to the UK Parliament, a hyperscaler described two timing gaps. It said it was £2.5 billion to £3 billion into a project due online at the end of 2026, but a delayed substation meant full supply would arrive 18 months later. It also said it could build planned capacity in northern England within 18 months but could not energize it until 2033.

The same company told Parliament that during the 18-month gap it is obliged to turn off the data center "in its entirety, with 20 minutes of notice," calling that "a pretty unsustainable situation." The UK operator added that a medium-sized 36 MW facility can require £500 million of infrastructure investment before GPUs and computing platforms. Together, the testimony shows that energization is not always binary: a connected facility can still face interruptibility, capital exposure and operating constraints.[8]



Who can own it

Better management of the interface is where value can be created, but that work is rarely contracted as an accountable, risk-bearing scope. So who owns it? Call that role a two-build integrator: a party or consortium that sits between the campus build and the power build and takes contractual ownership of that interface. It can own everything on the private side, up to and including private generation. It cannot own the utility's timeline. That is the line between the delays an integrator can price and the ones that must remain relief events.

The integrator can also influence utility coordination and joint sequencing, and it treats utility and regulatory events outside its control as grounds for defined contractual relief. The value is earlier visibility into schedule divergence and coordinated long-lead ordering, so fewer surprises land at the interface.

A creditworthy owner's early commitment can improve the case for customer-specific investment and make the power date more credible. It does not confirm an earlier date. One agreement between a large generator and a hyperscaler, under which supply ramps to 1,920 MW no later than 2032, with the contract running through 2042, shows a large buyer aligning long-term power supply with data center capacity.[9]

Integrator control boundary

Can contractually control	Can coordinate or influence	Should remain outside scope or qualify for relief
■ Campus delivery	■ Utility studies	■ Utility queue position
■ Private generation	■ Approval milestones	■ Regulatory decisions
■ Procurement interfaces	■ Outage windows	■ External transmission completion
■ Commissioning integration	■ Readiness evidence	■ Unbounded utility delay



Figure 4 — Three models are emerging around the campus-power interface

E&C has the strongest position only under a defined set of conditions.

E&C-led integrator

Owns the campus and private-power interface and carries a bounded share of delivery risk

Energy-led player

Combines generation or power development with the campus where energy is the scarce asset

Large-buyer-led model

Captures the interface through dedicated supply, including co-located generation and long-term offtake

E&C's strongest position, when all three are true

- A material private-power scope exists
- The owner values accountable risk transfer
- Utility and regulatory delays are bounded as relief events

Alternative models strengthen when:

Energy-led: Generation is the scarce asset and is secured with the campus

Large-buyer-led: The buyer captures the interface through dedicated or co-located supply

All three contend for the same interface. E&C's opportunity is segment-specific.

Source: EY analysis (conceptual)

E&C-led integrator

When one contractor closed its approximately US$1.65 billion acquisition of a specialty utility business, its CEO described the logic as serving data center and power infrastructure "both outside and inside the fence." The fence is exactly the seam this article is about.[10] Others are building adjacent capabilities. Another contractor has already delivered a high-performance computing campus that combined the computing environment with a new 160 MVA primary utility substation, 21 MW of generator and UPS backup capacity, and cooling in a single design-build project.[11] And a third has established a power and energy group providing engineering, procurement and construction (EPC) services for solar and storage, another sign that construction firms are moving deeper into the power scope.[12]

Energy-led player

Where generation is the scarce asset, the developer may hold the interface. An energy services firm and a power developer plan to align campus development and generation from the outset.[13]

Large-buyer-led model

This is the strongest alternative model, and it need not mean self-build. This category also includes data center and co-location operators, some owned by real estate investment trusts, and telecommunications companies. Under one announced project, a major energy company's power venture has agreed to develop roughly 2.67 GW of dedicated generation for a hyperscaler-operated data center under a 20-year power agreement, with first power expected in 2028 and a final investment decision expected by the end of 2026. At launch, the co-located structure is designed to supply the campus behind the meter, so the buyer captures the interface commercially through long-term offtake without owning the build. The operator says it anticipates connecting both the power facility and the data center to the broader grid over time. In that model, the private-power scope expands while grid interconnection becomes less central to the initial energization date. That project sharpened our view of the competitive window: if that model becomes the norm, the independent integrator role is locked up before it forms.[14]

The market is still testing these models, and we did not find public evidence that any model is proven at scale. The evidence supports three different levels of confidence:

Established	Emerging	Unproven
Campus and power schedules can diverge materially, and that gap already shapes delivery.	Contractors, energy developers and hyperscalers are assembling adjacent power capabilities.	Integrated delivery improves energization dates or earns a distinct risk-transfer premium.

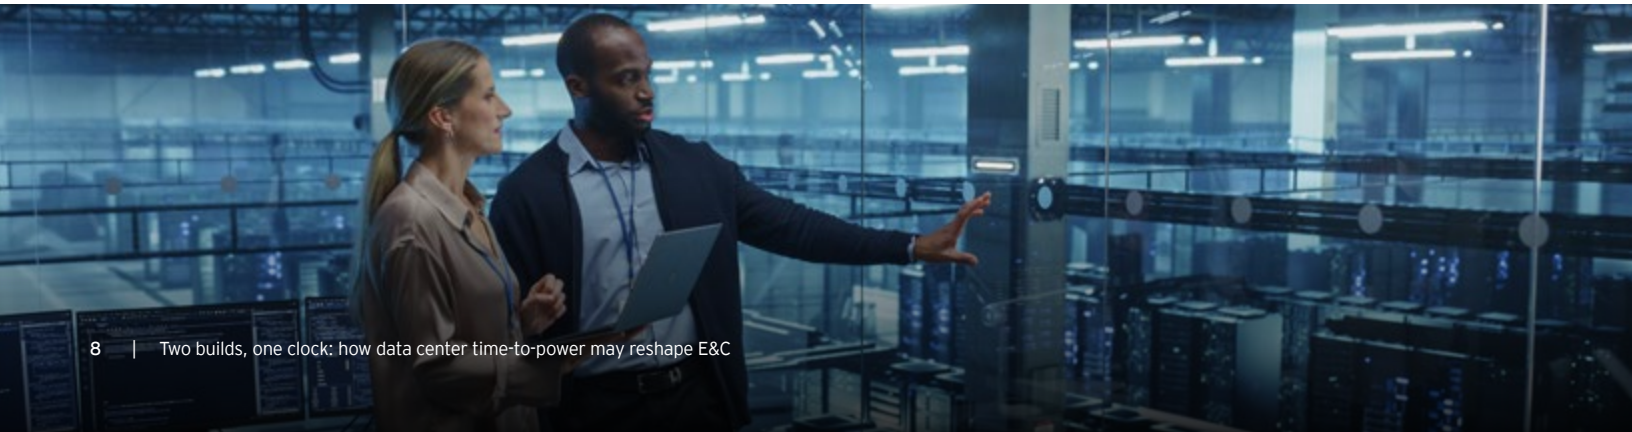
Our view is that measuring this exposure is already justified. Investing in a risk-bearing capability is a bet the evidence does not yet support. Those decisions belong on different timelines.

Most current evidence is company or analyst commentary, and some reflects integrated campus construction rather than integrated campus-plus-power delivery. An advisory or technology-led entrant could assemble the same capability through acquisition or partnership, so E&C's first-mover advantage is thinner than it looks.

Jurisdiction matters. Weak utility readiness raises the need for an integrator. It also reduces the risk an integrator can prudently accept.

Policy may shift cost without fixing coordination. The Federal Energy Regulatory Commission (FERC) Order No. 2023 reforms generation queues; large-load connections remain a separate policy problem. That gap is now being addressed. FERC opened a large-load interconnection docket in October 2025 at the direction of the Secretary of Energy, and on June 18, 2026 it issued show cause orders under section 206 of the Federal Power Act directing all six regional grid operators it regulates, and their transmission owners, to justify or reform their large-load rules across five categories. Those categories include transparency into transmission upgrade costs, and a cost recovery agreement requiring large load customers to make a minimum secured financial contribution toward the costs incurred to serve them. The Ratepayer Protection Pledge began with voluntary commitments by participating technology companies to pay for required generation and power-delivery infrastructure; the July 23, 2026 expansion added 23 governors and roughly 187 companies, including 55 utilities and 27 data center developers. The White House said the expanded pledge covers organizations responsible for roughly 80% of the power delivered to US homes and businesses. The operators in the Santa Clara example were among the developer signatories. That shift could enlarge the privately contractible power scope without resolving interface governance. A related legislative proposal, the Ratepayer Protection Act (H.R. 9340), passed the House 417 to 3 on September 16, 2026 and is now before the Senate. It would require state regulators and nonregulated utilities to consider a standard under which a data center drawing at least 100 MW at a single site or campus pays the full incremental cost of the generation, transmission and distribution upgrades needed to serve it, backed by financial assurances or contributions. States would only have to consider that standard, not adopt it. The bill had not been enacted as of September 24, 2026.[15]

E&C firms can compete for this role, though only under specific conditions. E&C's case is strongest where project owners want real risk transfer, the private-power scope is large and utility delays can be bounded as relief events. Where generation is scarce or a hyperscaler captures the interface through dedicated supply or long-term offtake, alternative models strengthen. In competitive bids, owners may resist a full exclusion for utility delay; a capped liability paired with defined relief is often the more realistic landing zone.



Whether owners will pay for it

If owners begin contracting for energization accountability, campus specialists without power-side capability could lose scope and leverage. That risk belongs on the board's agenda today.

Three sources of value are easy to blur: scarce capacity, integration and transferred risk.

- **Scarcity.** Supplies of long-lead grid equipment remain constrained, and energy-sector employers reported hiring difficulty in the 2025 US Energy and Employment Report.[16]
- **Integration.** A single accountable party may reduce coordination costs and surface divergence earlier.
- **Risk transfer.** The next buying cycle will test whether owners will pay an integrator to carry a bounded energization obligation.

The public data confirm the constraint but say little about price. One electric power infrastructure contractor reported that operating margin in its electric segment increased 140 basis points year over year, from 10.1% to 11.5%, and attributed the gain to increased project scope, self-performed solutions and execution rather than scarcity pricing. A mechanical and electrical contractor reported that margin in its electrical construction business increased 210 basis points, from 11.8% to 13.9%, largely on favorable project mix and field execution, alongside record remaining performance obligations of US$17.14 billion, up 43.9% year over year, a reminder that record backlog deserves an execution-risk lens as well as a growth lens.[17] The power contractor has separately said it limits participation in combined-cycle and simple-cycle EPC work unless contract terms bound the risk, and has described its power-generation strategy as combining an integrated delivery model, disciplined execution and a cost-effective contracting approach designed to mitigate execution risk.[18] We found no like-for-like scarcity or integration premium in the public pricing reviewed for this analysis, and any risk-transfer premium would have to emerge in future bids. Treat the economic proposition as a testable hypothesis. Credible proof would appear in recurring RFP language, observable pricing and measurable win-rate or margin evidence, rather than in announcements alone.

If an owner asks an integrator to commit to energization, the contract will have to separate delays within the integrator's control from utility or regulatory events outside the integrator's control. The negotiation should define the commercial operation date, notice requirements, time and cost relief, and termination rights. It should also address concurrent delays, where integrator-controlled and utility-caused events overlap, and identify who owns schedule float. Notice deserves particular emphasis: excusable-delay relief can be lost when contemporaneous records and contractual notices are incomplete. The contract should also identify the delay-analysis method that will govern any later dispute.

A broader energization promise may change how an entity identifies performance obligations, accounts for contract modifications, estimates and constrains variable consideration, evaluates transfer of control, and measures progress under ASC 606. The conclusion is contract-specific and requires accounting judgment. We deliberately stop at internal board reporting here; external concentration or risk-factor disclosure requires a separate assessment.

The hardest commercial issue is the energization date, because it still depends on the utility. Before standing behind it, a board should identify which retained risks are genuinely insurable or bondable, price the residual exposure, and determine whether owners will pay enough to cover that risk. If owners will not pay enough, integration should remain a coordination service rather than a paid, risk-bearing scope.



Two decisions for E&C boards

Decision 1. Measure and govern power-date exposure. In a delivery-constrained market, a record backlog is as much a risk measure as a growth measure. Boards should assess backlog through an execution-risk lens and, for internal board reporting, separately identify how much depends on power dates outside the company's control. Figures 5A and 5B present a conceptual, rule-based diagnostic for internal board discussion. The example programs, values and ratings are illustrative only; the framework has not been validated as a predictor of schedule outcomes. Our review did not identify a public sector-wide measure of the share of backlog dependent on third-party power dates. We think the industry should begin measuring it internally. For portfolio prioritization, pair the control rating with financial materiality; a small red program should not automatically be treated as more consequential than a much larger amber program.

Three questions belong on the agenda:

- How much of our backlog depends on a power date we do not control?
- Who owns the interface on our largest programs?
- Where in our contracts is energization accountability priced?

What each control rating means for action: **Green** means proceed and monitor. **Amber** means assign an interface owner and close the dependency or procurement gap. **Red** means stress-test and escalate the exposed assumption. Figure 5A then converts the worst control rating into the program verdict: On track, Watch or At risk.

Figure 5A — Delivery-certainty dashboard

Control ratings and financial materiality (see definitions on page 14). One verdict per program.

Program	Power-date confidence	Interface authority	Supply readiness	Contract exposure	Verdict	Financial materiality
Program A Coastal AI campus, 320 MW	G	G	A	G	Watch	US$0.9b exposed backlog; US$0.09b gross profit at risk; US$0.3b owner capital before firm power; delay 2/4 months (base/stress)
Program B Inland hyperscale, 500 MW	A	R	A	A	At risk	US$2.4b exposed backlog; US$0.43b gross profit at risk; US$0.8b owner capital before firm power; delay 6/10 months (base/stress)
Program C Brownfield expansion, 180 MW	G	G	G	G	On track	US$0.6b exposed backlog; US$0.03b gross profit at risk; US$0.15b owner capital before firm power; delay 0/1 months (base/stress)

Base/stress: Delay months under the approved schedule assumption (base) and a downside power-date assumption (stress). Illustrative sensitivities, not forecasts.

One rule for the verdict

- On track: all four controls green
- Watch: one or more amber, no red
- At risk: any control red

Portfolio backlog by delivery-certainty status



A rating shows control weakness; exposure shows consequence. Size it, do not just rate it.

Illustrative example, not actual company or EY portfolio data

Source: EY framework (conceptual). All programs, values, ratings and delay assumptions are illustrative; the dashboard is not a forecast or a validated predictor of schedule outcomes.

Board posture	Signal to act on
Monitor	Integrated-scope and energization-accountability language starts appearing in priority RFPs.
Prepare	Multiple live bids require power-side scope or bounded energization obligations.
Act	Repeated pricing, win-rate and margin evidence supports investment in risk-bearing capability.

Decision 2. Stage the capability decision. Stay a strong campus specialist, or build toward the two-build integrator role, coordinating both builds? Start with actions that are cheap to reverse. Name an interface owner on priority programs. Then decide which power-side capabilities to own, which to obtain through partners and which contractual boundaries to standardize. The build-or-buy decision should follow only when evidence from bids, pricing and win rates supports it. The postures below pair each board response with the signal that should trigger it.

A board does not need to pick the winning delivery model. What it must know is who owns the interface on each major program and how much backlog hangs on a third-party power date. A completed facility without firm power is idle capital. Every additional month increases its carrying cost and puts more of the expected return at risk. Buyers will increasingly reward firms that can measure that exposure, price it, staff the integrator role and carry a bounded share of the schedule risk. The rest may still build the campus. On current evidence, they will not own the clock.

Figure 5B — How to rate delivery certainty

Four controls and what Green, Amber and Red mean for each.

Control	Green	Amber	Red
Power-date confidence Firmness of the utility energization date	Committed date; milestones and dependencies documented	Conditional date; material dependencies remain open	No credible committed date, or date already slipping
Interface authority Who governs the campus-power seam	Named interface owner with defined escalation and decision rights within the contracting structure	Owner named, limited authority	No single custodian
Supply readiness Long-lead equipment and connection works	Critical equipment secured to schedule	Vendor capacity identified but not secured	Critical equipment or manufacturing slots unresolved
Contract exposure How utility delay is treated in contract	Relief, caps and exclusions defined and within approved risk appetite	Material terms or caps remain open	Exposure unbounded, or outside insurable or bondable coverage

Use with Figure 5A. Verdict = worst control: all green = On track ◆ any amber, no red = Watch ◆ any red = At risk

Source: EY framework (conceptual)



Notes

1. Santa Clara: Bloomberg coverage of Digital Realty's SJC37 and nearby STACK facilities awaiting additional Silicon Valley Power capacity; Silicon Valley Power said a broader US$450 million system upgrade serving the Digital Realty and nearby STACK facilities, as well as other customers, was scheduled for 2028. John Gittelsohn and Michelle Ma, "Data Centers in Nvidia's Hometown Stand Empty Awaiting Power," Bloomberg, November 10, 2025, cited here via the Energy Connects republication, https://www.energyconnects.com/news/utilities/2025/november/data-centers-in-nvidia-s-hometown-stand-empty-awaiting-power/. See also The Real Deal, "Santa Clara data centers await power supply" (November 2025), https://therealdeal.com/san-francisco/2025/11/10/santa-clara-data-centers-await-power-supply/.
2. Lawrence Berkeley National Laboratory, "Queued Up: 2026 Edition," published June 2026, https://emp.lbl.gov/publications/queued-2026-edition-characteristics (data through year-end 2025): in the regions with available duration data, the median project reaching commercial operation in 2025 waited more than five years from its interconnection request, nearly three times the 22-month median in 2008. Describes the 2025 operational cohort in those regions, not a current queue-wide wait. Interconnection queue data compiled by Lawrence Berkeley National Laboratory in partnership with GridTracker; the Queued Up data file is licensed CC BY 4.0.
3. Lawrence Berkeley National Laboratory, "Queued Up: 2026 Edition," published June 2026, https://emp.lbl.gov/publications/queued-2026-edition-characteristics (data through year-end 2025): 2,061 GW of generation and storage (1,312 GW generation and approximately 749 GW storage) actively seeking interconnection. The data set covers generation requests and excludes load interconnection, distribution-connected and behind-the-meter projects. Most requested capacity is ultimately withdrawn; total active capacity fell roughly 10% year over year as withdrawals outpaced new requests. See note 2 on attribution and licensing.
4. Currence's February 2026 global outlook tracks 190 GW across 777 data center and AI facilities larger than 50 MW announced since 2024 and estimates that 30%–50% of capacity slated to come online in 2026 could be delayed, with power and construction constraints among the causes. Currence outlook, https://www.currence.ai/blog/data-center-outlook.
5. EY analysis using the EY Energy and Resources Transition Acceleration (ERTA) model, current-trajectory scenario, extracted August 2026.
6. Data center capital split: in Epoch AI's stylized 1 GW model (which assumes grid power, not an owned generation system), the facility, land and utility works are roughly 31% of up-front capital, with servers and networking the rest. That share is an EY calculation from the up-front CapEx breakdown published with that analysis. This is a share of up-front capital expenditure, not of annualized total cost of ownership (in which servers are the larger recurring line); its share is larger on an up-front capex basis than on an annualized cost basis because facility capital is spread over a longer assumed life than the IT equipment it houses. Amelia Michael and Ben Cottier, "Servers account for 60% of the total cost of ownership of a one-gigawatt AI data center," Epoch AI, May 14, 2026, licensed CC BY 4.0, https://epoch.ai/data-insights/ai-datacenter-cost-breakdown.

7. On recurring constraints at the campus-to-grid handoff (interconnection studies, transformer procurement and transmission upgrades), see the public evidence cited elsewhere in these notes: LBNL on interconnection timelines (notes 2–3), https://emp.lbl.gov/publications/queued-2026-edition-characteristics. US Department of Energy on grid equipment lead times (note 16), https://www.energy.gov/oe/articles/strengthening-americas-grid-supply-chain.

8. Kao Data's Harlow campus was quoted a 2037 connection date for its next grid tranche. See UK Parliament oral evidence (April 14, 2026), https://committees.parliament.uk/oralevidence/17474/html/, including Hugh Milward's evidence that Microsoft is obliged to turn off a data center in its entirety with 20 minutes of notice during an 18-month power gap, and Clinton Hasell's evidence that a medium-sized 36 MW facility can require £500 million of infrastructure investment before GPUs and computing platforms.

9. Long-term power aligned with data center capacity under a power purchase agreement; a representative deal: Talen–AWS (supply ramping to 1,920 MW no later than 2032, under a contract running through 2042). Company release, https://ir.talenenergy.com/news-releases/news-release-details/talen-energy-expands-nuclear-energy-relationship-amazon.

10. MasTec's approximately US$1.65 billion acquisition of The Superior Group, framed as expanding its ability to serve data center and power infrastructure, both outside and inside the fence. Company release, https://investors.mastec.com/news-releases/news-release-details/mastec-completes-previously-announced-acquisition-superior-group.

11. Kiewit integrated high-performance computing campus: a delivered 160 MVA primary substation, 21 MW of generator and UPS backup capacity, and cooling combined with the computing environment in a single design-build project. Kiewit project description, https://www.kiewit.com/projects/high-performance-computer-center-hpcc-2/.

12. Clayco established a Power + Energy group providing EPC services for utility-scale solar and battery storage. Clayco company webpage, https://claycorp.com/power-energy-storage.

13. Liberty–PowerBridge, a planned approximately 2 GW campus aligning generation and campus development from the outset; further agreements and approvals are required. Company announcement: Liberty Energy and PowerBridge JV announcement, https://investors.libertyenergy.com/news-and-events/press-releases/2026/07-22-2026-230509078.

14. Large-buyer-led model: the Microsoft–Chevron development in West Texas, known as Project Kilby (approximately 2.67 GW of dedicated co-located power generation, anticipated first power in 2028, final investment decision expected by the end of 2026). Company releases of announced plans, subject to Chevron's final investment decision and other conditions. Microsoft's Pecos announcement and Chevron's generation agreement describe the same campus near Pecos, Texas. The power facility is designed for 100% behind-the-meter supply, avoiding utility delivery and ERCOT interconnection for that dedicated capacity. Microsoft Pecos announcement, https://blogs.microsoft.com/blog/2026/06/22/powering-the-next-wave-of-ai-expanding-capacity-with-our-new-datacenter-in-pecos/. Chevron newsroom, https://chevroncorp.gcs-web.com/news-releases/news-release-details/chevron-signs-20-year-power-agreement-microsoft-west-texas-data.

15. Federal Energy Regulatory Commission (FERC) Order No. 2023 reforms the generator interconnection queue (not the separate large-load connection process); it was upheld by the US Court of Appeals for the DC Circuit on July 31, 2026, which denied all petitions (Advanced Energy United, et al. v. FERC, No. 23-1282). On large-load connections, see LBNL "Speed to Power" (2026). The White House "Ratepayer Protection Pledge" commits signatory hyperscalers to build, bring or buy their own generation and to pay for the power-delivery infrastructure their data centers require: Proclamation 11014 of March 4, 2026 (Federal Register, 91 FR 11439, March 9, 2026), expanded July 23, 2026: White House. The proclamation sets out voluntary company commitments; it does not itself set retail rates, which remain within state regulatory processes. As amended by the committee, the Ratepayer Protection Act (H.R. 9340) applies to data center loads of at least 100 MW at a single site or campus, narrowed from the introduced bill's coverage of any nonresidential load at that threshold. The House Committee on Energy and Commerce voted 52–0 on July 21, 2026, to order the bill reported, as amended. The House passed the bill, as amended, by 417 to 3 on September 16, 2026, under suspension of the rules (Roll Call 312), and it was referred to the Senate. It had not been enacted as of September 24, 2026. Federal Energy Regulatory Commission, Interconnection of Large Loads to the Interstate Transmission System, Docket No. RM26-4-000, opened after the Secretary of Energy directed FERC on October 23, 2025 to consider an advance notice of proposed rulemaking, https://www.ferc.gov/rm26-4. The advance notice proposed a threshold of 20 MW or more, a different test from the 50 MW used in the show cause orders. On a separate track, FERC issued six show cause orders on large-load integration on June 18, 2026 under section 206 of the Federal Power Act, one to each of the six regional transmission organizations and independent system operators it regulates and their transmission owners, each in its own docket rather than in RM26-4-000; those orders define large load as peak load in excess of 50 MW interconnecting to transmission lines greater than 69 kV, https://www.ferc.gov/news-events/news/ferc-launches-aggressive-targeted-action-speed-large-load-integration. Sources: FERC Order No. 2023 court opinion, https://media.cadc.uscourts.gov/opinions/docs/2026/07/23-1282-2185998.pdf. Lawrence Berkeley National Laboratory, Speed to Power, https://emp.lbl.gov/publications/speed-power-solutions-accelerating. White House Ratepayer Protection Pledge proclamation, https://www.whitehouse.gov/presidential-actions/2026/03/ratepayer-protection-pledge-proclamation/. Federal Register publication, https://www.federalregister.gov/documents/2026/03/09/2026-04645/ratepayer-protection-pledge. White House July expansion announcement, https://www.whitehouse.gov/releases/2026/07/president-trumps-ratepayer-protection-pledge-secures-american-ai-dominance-protects-consumers/. White House pledge signatory list, https://www.whitehouse.gov/ratepayer-protection-pledge/. Associated Press coverage of 187 companies and 23 governors, https://apnews.com/article/490ea7e4c7227d5e550b00a0056c33c9. House Committee on Energy and Commerce, "Ratepayer Protection Act Advances from House Committee on Energy and Commerce," July 21, 2026, https://energycommerce.house.gov/posts/ratepayer-protection-act-advances-from-house-committee-on-energy-and-commerce. H.R. 9340 introduced text, https://www.govinfo.gov/app/details/BILLS-119hr9340ih. Amendment in the Nature of a Substitute to H.R. 9340, https://docs.house.gov/meetings/IF/IF00/20260720/119490/BILLS-119-HR9340-A000370-Amdt-4.pdf. Congress.gov, "H.R. 9340 — Ratepayer Protection Act," https://www.congress.gov/bill/119th-congress/house-bill/9340. Congressional Budget Office, "H.R. 9340,

Ratepayer Protection Act," September 9, 2026, https://www.cbo.gov/publication/62762. House Clerk roll call, https://clerk.house.gov/Votes/2026312; and the committee's release on passage at https://energycommerce.house.gov/posts/ratepayer-protection-act-passes-house-with-strong-bipartisan-support.

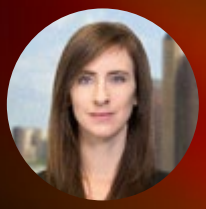
16. US Department of Energy, Office of Electricity, "Strengthening America's Grid Supply Chain," August 10, 2026, https://www.energy.gov/oe/articles/strengthening-americas-grid-supply-chain (reporting lead times of two or more years for critical grid equipment, including transformers, circuit breakers, substation components and power electronics); US Department of Energy, "2025 US Energy and Employment Report," 2025, https://www.energy.gov/sites/default/files/2025-08/National%20USEER_08282025.pdf (employer-reported hiring difficulty).
17. Quanta Services, Inc., "Quanta Services Reports Second Quarter 2026 Results," July 30, 2026, Supplemental Segment Data, https://investors.quantaservices.com/news-events/press-releases/detail/402/quanta-services-reports-second-quarter-2026-results (Electric segment operating margin of 11.5% in the second quarter of 2026, compared with 10.1% in the second quarter of 2025); EMCOR Group, Inc., "EMCOR Group, Inc. Reports Second Quarter 2026 Results," July 30, 2026, https://emcorgroup.com/application/files/3517/8535/6062/2Q26_Earnings_Release_Final.pdf (US Electrical Construction & Facilities Services operating margin of 13.9%, compared with 11.8%; remaining performance obligations of US$17.14 billion, up 43.9% year over year). See also the companies' Form 10-Q filings for the quarter.
18. Quanta Services, Inc., "Building Upon Decades of Power Generation Experience, Quanta Expands Its Total Solutions Platform and Announces Its Selection by NiSource to Provide Power Generation and Grid Infrastructure Solutions for a Large Load Customer," October 30, 2025, https://investors.quantaservices.com/sec-filings/all-sec-filings/content/0001193125-25-257378/d851878dex992.htm (describing an integrated power-generation approach, disciplined execution and a cost-effective contracting approach designed to mitigate execution risk); Quanta Services second quarter 2026 earnings call (describing selectivity on combined-cycle and simple-cycle generation EPC work unless contract terms are favorable).

Definitions

Terms used in Figure 5A (page 10).

- Exposed backlog: remaining contracted work whose timing or economics depend materially on a third-party power date.
- Gross profit at risk: forecast project-level gross profit in that backlog that could be deferred, reduced or lost if the power date slips. Based on an illustrative program-specific gross-margin assumption, not segment operating margin.
- Owner capital before firm power: owner or developer capital spent or committed to the campus before the power date meets the "committed" standard in Figure 5B. Excludes the E&C firm's working capital.
- Base delay: estimated slippage in months against the approved power-on date under the program's most likely assumptions.
- Stress delay: the same measure under a defined downside power-date assumption.

Contacts



Andi Pritchett
Partner, Engineering and Construction Assurance Leader
Ernst & Young LLP
andi.pritchett@ey.com



Brad Cunningham
Senior Manager, Assurance
Ernst & Young LLP
brad.cunningham@ey.com



Ajay Yeldandi
Global Industrial Products Analyst, EY Sector Insights
Ernst & Young LLP
ajay.yeldandi@ey.com

The views reflected in this article are the views of the author(s) and do not necessarily reflect the views of Ernst & Young LLP or other members of the global EY organization.

EY | Building a better working world

EY is building a better working world by creating new value for clients, people, society and the planet, while building trust in capital markets.

Enabled by data, AI and advanced technology, EY teams help clients shape the future with confidence and develop answers for the most pressing issues of today and tomorrow.

EY teams work across a full spectrum of services in assurance, consulting, tax, strategy and transactions. Fueled by sector insights, a globally connected, multi-disciplinary network and diverse ecosystem partners, EY teams can provide services in more than 150 countries and territories.

All in to shape the future with confidence.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. Information about how EY collects and uses personal data and a description of the rights individuals have under data protection legislation are available via ey.com/privacy. EY member firms do not practice law where prohibited by local laws. For more information about our organization, please visit ey.com.

Ernst & Young LLP is a client-serving member firm of Ernst & Young Global Limited operating in the US.

© 2026 Ernst & Young LLP.
All Rights Reserved.

US SCORE no. 32303-261US
2608-10757-CS

ED None

This material has been prepared for general informational purposes only and is not intended to be relied upon as accounting, tax, legal or other professional advice. Please refer to your advisors for specific advice.

ey.com