In AI power, capacity is booked before it is bought.

AI is on course to nearly double the electricity data centers use from 2025 to 2030. Equipment and grid connections run on different clocks, and some buyers pay to reserve factory slots before a final order. This report separates those commitments from delivery and usable power.

950 TWh

Global data-center electricity use in 2030, IEA base case, up from 485 TWh in 2025

30–60 weeks

Illustrative US planning band for selected building electrical items and chillers, NRG estimate; not a package-wide delivery promise

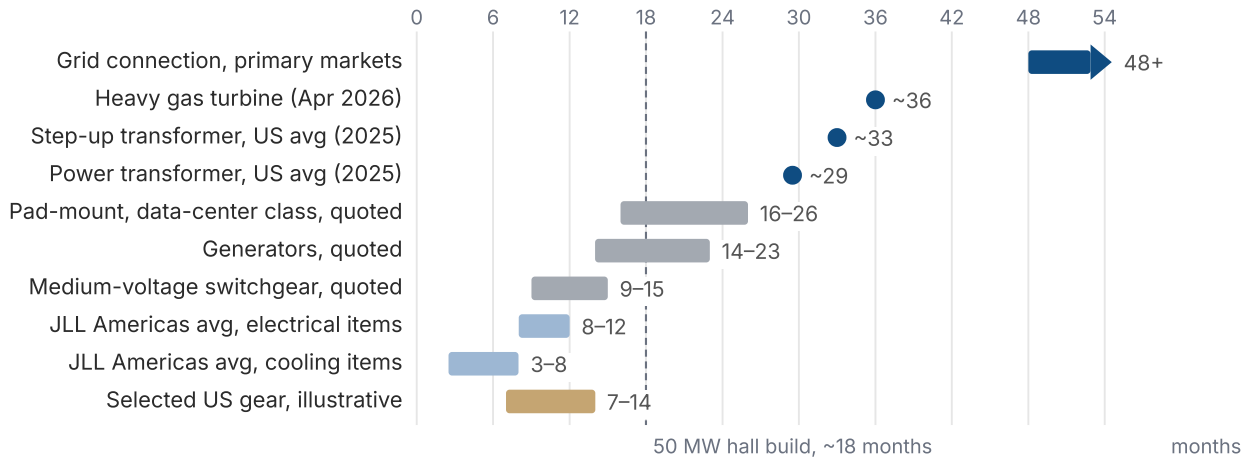
4+ years

Average wait for a grid connection in primary data-center markets, JLL, 2026

SUMMARY

The slowest pieces now take longer to arrive than an AI hall takes to build.

Different clocks in months, not sequential tasks. Gold: illustrative NRG estimate; grid row: connection wait



Source: JLL (2026); Wood Mackenzie (Q2 2025); C&W via press (2026); issuer April 2026 call; NRG estimate and calculation

Data centers used about 485 TWh of electricity in 2025, 17% more than a year earlier. The IEA's April 2026 update expects about 950 TWh by 2030. These are global estimates and forecasts, not installed generating capacity.

NRG estimate: an illustrative 30–60 weeks for selected US building electrical items and chillers. This is not a measured market distribution. JLL's 2026 Americas item averages and longer quoted generator and pad-mount ranges describe different products and methods.

A February 2026 public turbine reservation targets delivery by 2030, subject to a later purchase agreement and price negotiation. The deposit commits cash; it does not guarantee an operating date. At one large turbine maker, reservations were nearly half of gas commitments at September-end.

Korea's March 2026 capital-region snapshot shows 3% of cumulative requested capacity approved to date, not a final success probability. Texas's September operations report separately records 9,456 MW approved to energize and 4,132 MW of non-simultaneous peak load.

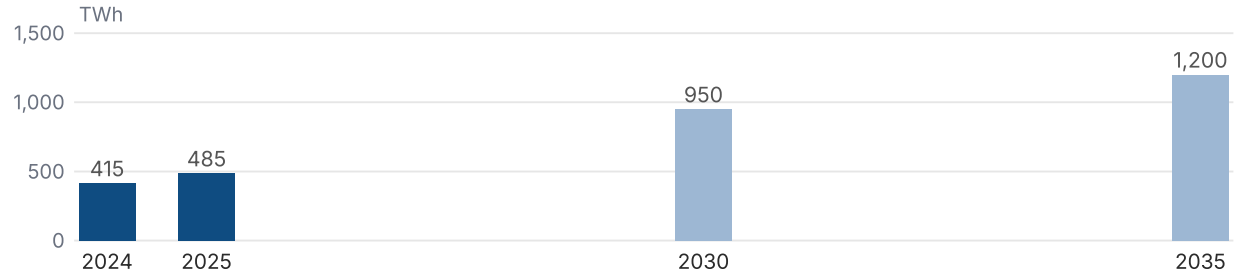
Equipment arrival, site readiness, fuel or utility supply, and integrated operation each need evidence. Grid waits can outlast equipment queues, but no market-wide easing date follows. Hybrid 800-volt DC designs can retain existing building AC infrastructure; a wholesale order reset is conditional.

01 · THE LOAD

Data-center electricity use nearly doubles by 2030.

The IEA puts global use at 415 TWh in 2024, 485 TWh in 2025 and about 950 TWh in 2030. In the United States, Berkeley Lab's 2030 range is wide, and most of the spread comes from how much hardware gets installed.

Global data-center electricity use, TWh, on a true time axis; 2030 and 2035 (lighter) are the IEA base case



Source: IEA, Energy and AI (2025); IEA, Key Questions on Energy and AI (2026)

Use rose 17% in 2025, against roughly 3% for electricity as a whole, and use by AI-focused data centers grew 50%. The IEA expects AI-focused use to triple by 2030. The 2026 path implies about 14% a year from 2025 (NRG calculation). The older 2025 edition's longer-horizon base case reaches about 1,200 TWh by 2035.

In the IEA's 2025 report, data centers account for nearly half of US electricity demand growth to the end of the decade. Japan and Korea together hold about 5% of global use, a share the IEA expects them to keep.

Berkeley Lab's June 2026 update puts US data centers at 649 TWh in 2030 in its reference case, or 11.8% of national use, inside a range of 521–843 TWh. Its 2024 edition had put the 2028 share at 6.7% to 12.0%. The sensitivity cases in the table show which assumptions move the number most, and installations matter more than chips.

The IEA left its 2030 path almost unchanged in 2026 even as spending climbed, because supply bottlenecks hold back the near term. NRG's research points to the same operating constraint: demand can stay strong while equipment and connection dates limit delivery.

US data-center electricity use in 2030, Berkeley Lab cases

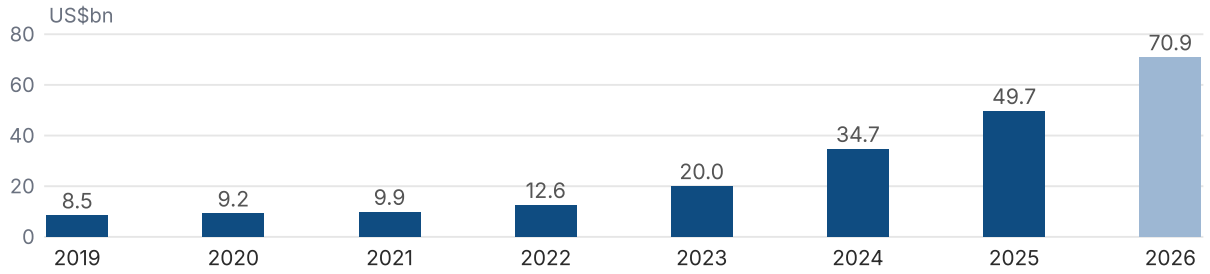
Case	TWh in 2030
Reference case	649
Fewer IT installations	578
More graphics chips shipped	664
Higher AI-server idle power	782

Sources: IEA (2025, 2026); Lawrence Berkeley National Laboratory (2026); NRG calculation

US data-center construction spending has quintupled since 2021.

US private data-center construction ran at a seasonally adjusted annual rate of $96.5bn in September 2026, the latest observation checked November 5. This measures construction put in place, not server purchases or actual year-to-date spending.

US private data-center construction, US$bn, mean seasonally adjusted annual rate; 2026 is January–September, not cumulative spend



Source: US Census Bureau; NRG calculation

The Census series averaged $9.9bn at an annual rate in 2021 and $49.7bn in 2025, about five times as much (NRG calculation). Its January–September 2026 mean was $70.9bn SAAR. September's preliminary $96.5bn was 90.3% above September 2025. The September rate and the partial-year mean answer different questions.

JLL puts the global average cost of shell and core at $7.7m per MW in 2020 and $10.7m in 2025, and expects $11.3m in 2026. Tenants add the fit-out on top, which JLL says can reach $25m per MW for AI.

Cushman & Wakefield's September 2026 public release reports US and Canadian development costs up 21% per MW against its late-2024 edition. Power takes 21% of the budget, the largest single line. Detailed cost ranges reproduced by the press have not been independently re-extracted from the interactive guide; they are not needed to establish that construction is becoming more expensive.

JLL found that 57% of projects ran three months or more late in 2025, against a typical build of 18 months for a 50 MW hall. A hall that waits on one late item also waits to earn rent, while the equipment already delivered sits idle. Worldwide, JLL counts 103 GW in 2025 and expects 97 GW more by 2030; that is a forecast of capacity, not construction spending.

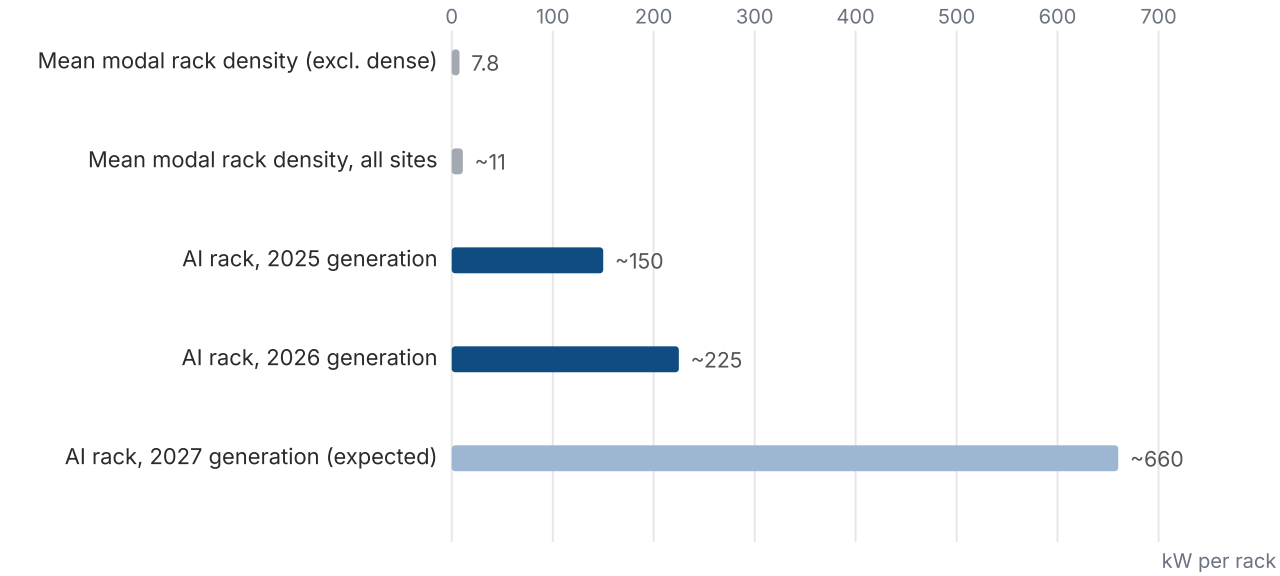
> **21%**
>
> Share of a US and Canadian greenfield data-center budget that goes to power infrastructure, the largest single line (Cushman & Wakefield, 2026).

Sources: US Census Bureau; JLL (2026); Cushman & Wakefield (2026); NRG calculation

One current AI platform is rated at about 29 times the mean modal rack density.

Uptime Institute's 2026 survey puts the average of respondents' most common rack densities at 7.8 kW after excluding the few very dense sites. TrendForce's June roadmap rates one current AI platform at about 225 kW and a future generation at about 660 kW.

Power per rack, kW; AI racks are one chip maker's rack-scale systems by generation



Source: Uptime Institute (2026); TrendForce (2026)

The headline is an NRG calculation comparing a platform power rating with a survey statistic, not a comparison of measured fleet electricity use. Including dense sites, the average of respondents' most common rack densities rose from 9 kW in 2025 to above 11 kW in 2026. JLL forecasts average density of 45 kW by 2030.

Racks of more than 50 kW are typical at only 3% of the sites Uptime polled, and the most common design density is still 4–5 kW. A hall built at AI densities buys more electrical and thermal plant per square meter, sized to the rack generation it will house rather than to the floor area.

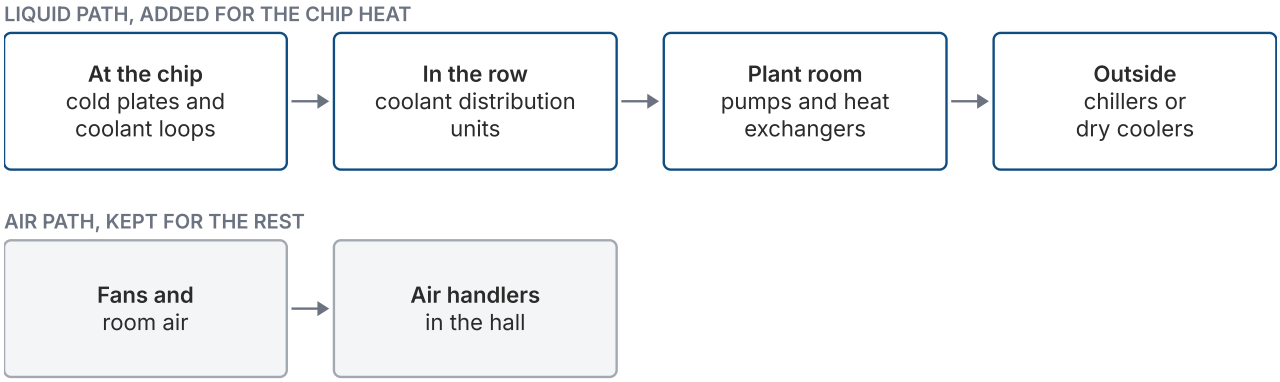
TrendForce's future 660 kW platform would be about 85 times the 7.8 kW survey figure (NRG calculation). Its expected second-half 2027 timing is a roadmap, not proof of deployment. Published platform requirements help size a hall's electrical and thermal interfaces; compatibility does not allocate factory output.

Liquid-cooled halls cost about 10% more to build than air-cooled ones, by JLL's estimate. Liquid reaches the chip while air still handles other heat. Higher-voltage rack designs may alter conversion and distribution equipment; hybrid designs can keep existing building AC infrastructure rather than replacing the entire power chain.

Liquid cooling is growing faster than any other part of data-center plant.

An infrastructure-research forecast puts global liquid-cooling revenue at about $3bn in 2025 and about $7bn in 2029, roughly 24% annual growth (NRG calculation). Revenue reflects equipment intensity, volume, mix and prices; it is not a measure of installed cooling capacity.

What a liquid-cooled hall adds, from the chip to the outside air



Source: data-center infrastructure market research (2026); NRG analysis

Most liquid capacity now arriving uses single-phase loops that carry coolant to a cold plate on the chip; immersion remains a small share. Fans and room air still remove the rest of the heat, and all of it must leave the building, so a liquid loop adds plant, as the diagram shows. Most new halls run both systems side by side.

The same research forecasts physical-infrastructure revenue, power and cooling together, at $120bn by 2030, growing 22% a year. JLL's 14% capacity-growth base case measures GW instead of revenue. More thermal plant per dense hall is one mechanism behind demand, but comparing those growth rates cannot isolate its contribution. The infrastructure forecast identifies thermal management as its fastest segment and liquid cooling as its fastest technology.

Leading-edge chips are expected to dissipate more than 4,000 W each by 2029, and JLL expects 80% of new facilities to be liquid-cooled by 2030. Chillers are the slowest item on the cooling side; JLL's 2026 Americas figure for them is about 35 weeks, close to the electrical gear.

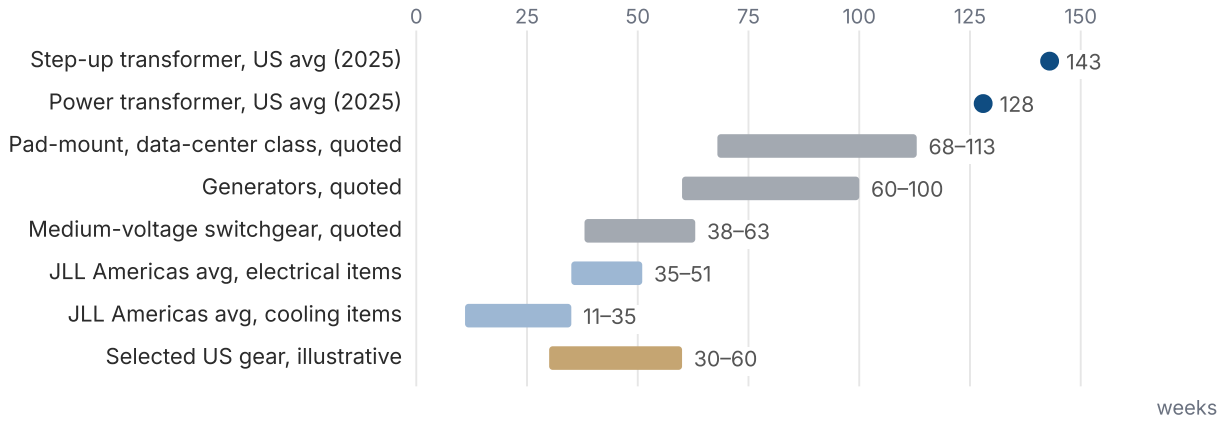
Cooling equipment and electrical equipment share project deadlines without sharing one lead time. JLL's 2026 Americas chart places cooling-item averages around 11–35 weeks, against roughly 35–51 weeks for the main electrical items. Those chart readings are approximate regional averages. A late chiller can still delay a hall even when other cooling components arrive quickly; an average for the whole equipment list cannot identify that risk.

Sources: infrastructure market research; JLL; Cushman & Wakefield (all 2026); NRG calculation

A hall has several equipment clocks.

JLL's 2026 US equipment average is 42 weeks. Component means, quoted ranges and historical utility-transformer benchmarks differ. One late required item can determine the project's date.

Order to delivery, weeks: 2026 regional means / quoted ranges; utility-transformer means Q2 2025. Gold: illustrative NRG estimate



Source: Wood Mackenzie (Q2 2025, US); Cushman & Wakefield (2026, US and Canada); JLL (2026); NRG estimate

Wood Mackenzie's Q2 2025 US means were 128 weeks for power transformers and 143 for generator step-up units. Demand was 116% and 274% above 2019 respectively; power-transformer demand exceeded supply by 39%. These are dated benchmarks. The IEA found transformer and cable waits had doubled in three years.

JLL's 2026 US equipment average is 83% longer than in 2019. Its Americas electrical-item means run about 35–51 weeks, and cooling means 11–35 weeks. These are approximate chart readings, with regional and product definitions distinct from the US utility-class benchmarks.

Press quotations from Cushman & Wakefield's 2026 guide give 38–63 weeks for medium-voltage switchgear, 60–100 for generators and 68–113 for data-center pad-mounts. The detailed interactive table was not re-extracted. Quoted ranges are not proven tail observations or a distribution comparable to JLL's means.

Wood Mackenzie's August 2026 estimate of a 15% power-transformer shortage uses a different definition from demand-over-supply. US bulk-power restrictions depend on implementation and specified findings. They do not prohibit every Chinese-made component. Procurement eligibility and delivery need item-level confirmation. Current supplier quotations remain essential.

> **30–60 weeks**
>
> NRG estimate: illustrative 30–60 weeks for selected US building electrical items and chillers. Limited research, not a distribution or package promise; excludes utility-class units and longer generator or pad-mount quotations.

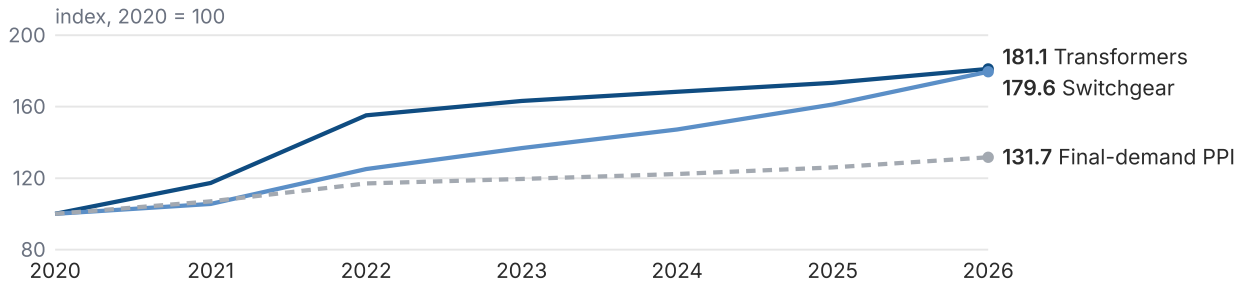
Sources: Wood Mackenzie (2025, 2026); Cushman & Wakefield (2026); JLL (2026); IEA (2025); NRG estimate

06 · THE PRICE

Transformer and switchgear indexes have outpaced final-demand prices since 2020.

BLS industry indexes for transformers and switchgear stood 86% and 88% above their 2020 means in September 2026. The final-demand PPI was 35% higher. Different index coverage makes this a price comparison, not an identical market basket.

Producer price indexes rebased to 2020, yearly averages; 2026 is January to September (NRG calculation)



Source: US Bureau of Labor Statistics via FRED; NRG calculation

Transformers moved first. Their index rose by more than half between 2020 and 2022 and has climbed about 4% a year since, in line with the IEA's finding that the power-transformer price index rose 1.5 times from 2020.

Switchgear took longer. Its index climbed steadily and then sped up: the 2026 average so far is 11% above 2025 (NRG calculation), the fastest rise since 2022.

In September 2026 the transformer industry index stood at 475, against a mean of 255 in 2020. It covers a broader industry than utility-class units alone. Producer prices record factory selling prices; they omit a buyer's lost revenue from project delay and do not reveal every contract quotation.

A flatter industry index does not establish that every utility-transformer quote has flattened. Copper costs affect windings and busbars, but pass-through varies with contracts, product mix and timing. Shorter delivery quotes can coexist with rising prices when old orders, materials costs or a changing equipment mix feed into current sales.

Lead time and price therefore need separate readings. Uptime's 2026 survey found 72% of operators at least somewhat concerned about supply-chain disruption. A buyer needs current item quotes, delivery terms and cancellation conditions alongside published indexes; an easing average does not establish that a required unit is available.

September 2026 endpoint versus the 2020 mean, NRG calculation

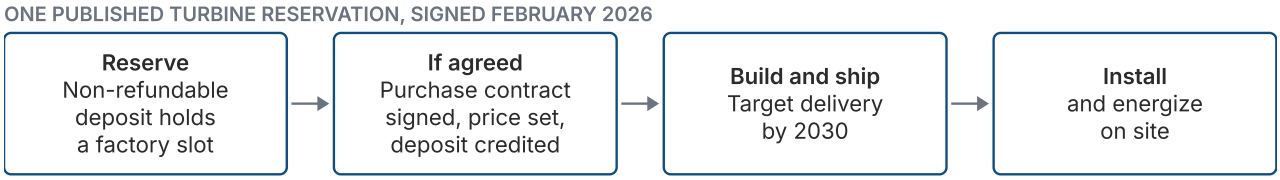
BLS series	Change
Transformer industry	+86%
Switchgear industry	+88%
Final demand	+35%

Sources: US Bureau of Labor Statistics via FRED; Cushman & Wakefield (2026); IEA (2025); NRG calculation

A reservation holds a target slot; the final order and operating date remain separate.

A Canadian producer's February 2026 release describes a non-refundable deposit for a turbine manufacturing slot. Target delivery is by 2030, with a definitive purchase agreement and price still to be negotiated. These are one transaction's disclosed terms.

Anatomy of a slot reservation, from one published example



Source: company filings and announcements

The deposit is credited toward the purchase price if the purchase proceeds. That avoids counting it twice as equipment expense; it does not remove the financing or opportunity cost of cash paid early. A non-refundable commitment also leaves downside if the project or contract does not proceed as expected. No final purchase agreement was established by this review.

JLL reports that large developers preorder selected materials and hold critical inventory. This supports early procurement as an industry mechanism, not the claim that all equipment uses the same deposit, price or delivery terms. Each contract determines what capacity is reserved, what can change and what happens after a delay.

A target slot can reduce exposure to factory availability without guaranteeing arrival or operation. Permits, site readiness, fuel or utility supply, and integration can remain open after the equipment is reserved. The economic benefit depends on whether those other conditions are ready when the equipment reaches the site.

Suppliers receive an earlier demand signal and cash commitment, but reservations still differ from firm orders. At one turbine maker, reservations briefly exceeded firm backlog in 2026. They move commitment earlier in the process; they do not by themselves shorten manufacturing or make every proposed project executable.

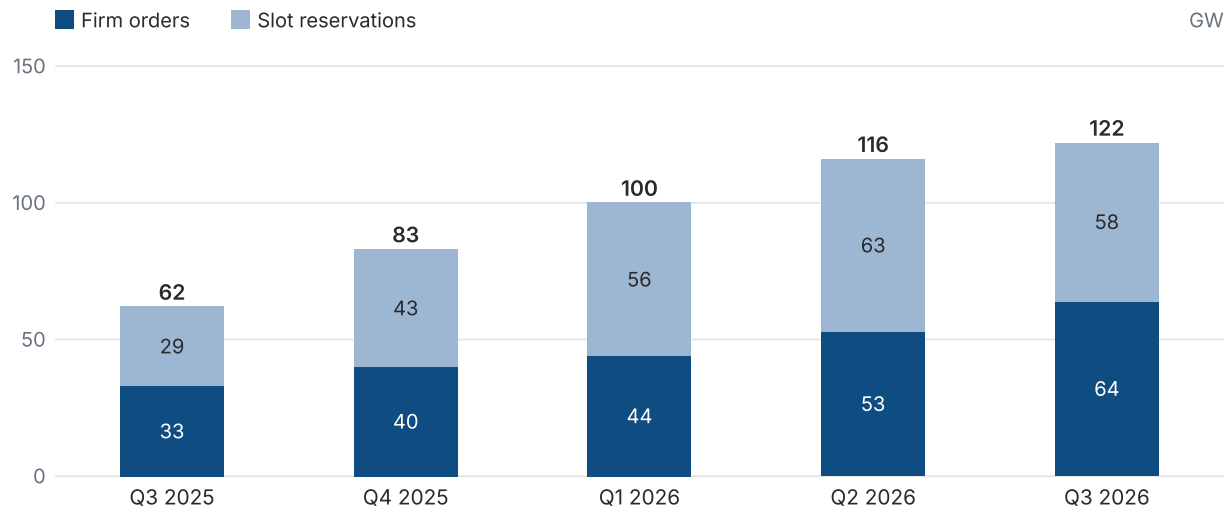
> **24 months**
>
> How far ahead large developers preorder select materials, by JLL's 2026 count. They also keep 6–12 months of critical components in stock.

Sources: company filings and announcements; JLL (2026)

One turbine maker's commitments include a large reservation layer.

At one large US maker, gas-turbine commitments almost doubled in a year to 122 GW, and about half were reservations rather than firm orders.

Gas-turbine backlog and slot reservations at one large US maker, GW, end of quarter



Source: company filings and announcements

Reservations exceeded half the maker's commitments early in 2026 and have since begun to convert. Between July and September it converted 13 GW into orders and added 8 GW of reservations. September-end commitments were 64 GW firm backlog plus 58 GW reservations. Its at-least-130 GW year-end target remains guidance, not installed power.

About a fifth of its contracted gas volume was tied directly to data-center load in April. The rest goes to utilities, independent producers and industry. Its grid-equipment backlog rose 60% in a year to $42.2bn, partly through an acquisition, so the wires side of the business is filling too.

In April 2026 the issuer described roughly three-year heavy-turbine lead times and then-available 2029–30 capacity. This is a dated disclosure, not a November delivery offer or a current industry-wide horizon. It also said construction, permits and fuel supply can hold a project back after the turbine date. NRG's research supports that distinction between equipment delivery and usable power.

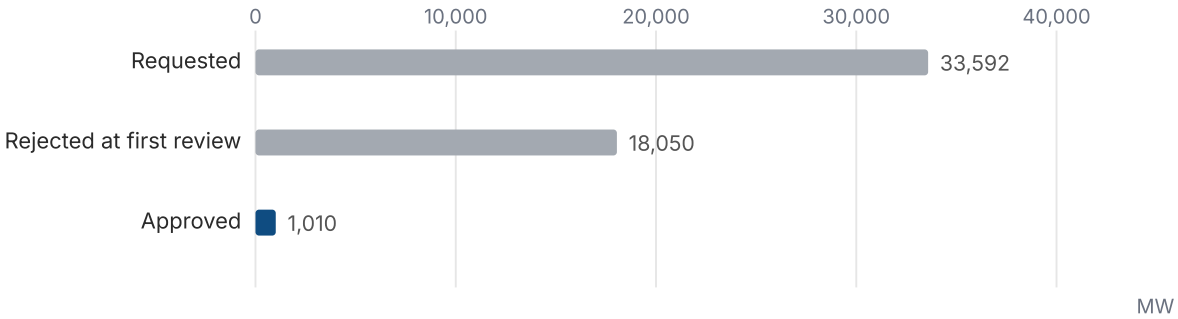
The IEA counts 47 GW of on-site gas projects for US data centers under development or proposed, and expects 15–27 GW of on-site gas to serve data centers worldwide by 2030. Orders for gas turbines to serve data centers rose 70% in 2025, by the agency's count. Those plants need transformers and switchgear of their own, so they join the same electrical queue. Our companion article on fuel cells for AI data centers looks at the other on-site options.

Sources: company filings and announcements; Power Engineering; Turbomachinery Magazine; IEA (2026)

Korea's March snapshot records approvals to date, not final success odds.

Capital-region grid-impact-review applications totaled 33,592 MW through March 2026, with 1,010 MW approved to date. The cumulative 3% share includes cases at different stages. It does not measure the eventual approval rate of completed applications.

Capital-region grid-impact reviews, cumulative to March 2026, MW; status bars are not a complete partition



Source: Electronic Times (July 2026)

By count, 10 of 522 applications had approval, or 1.9% (NRG calculation). The chart shows all requested capacity, first-stage rejections and approvals as separate status measures, not an exhaustive partition. Other cases may still be pending or at later review stages. A first-stage rejection describes that application and assessment; it does not establish that supply is permanently impossible.

National utility applications are a different series. Press reporting records 88 in January–August 2026 versus 46 in full-year 2025, not a same-period growth rate. Its reported capacity mix also differs from the capital-region review population. The provisional national demand plan remained in consultation in October; draft projections are not final supply commitments.

Texas press reports put large-load requests near 474 GW in mid-2026. A separate official September operations report records 9,456 MW approved to energize and 4,132 MW of non-simultaneous observed peak in September. These are different populations and stages; dividing one by the earlier request total would not establish a project conversion rate.

JLL's 2026 average connection wait in primary data-center markets exceeds four years. The IEA's 2025 assessment says constraints could delay about 20% of planned capacity to 2030. Such evidence supports persistent site-specific grid risk, not a year when the global queue clears. Equipment delivery alone cannot prove a hall can run.

Separate gates

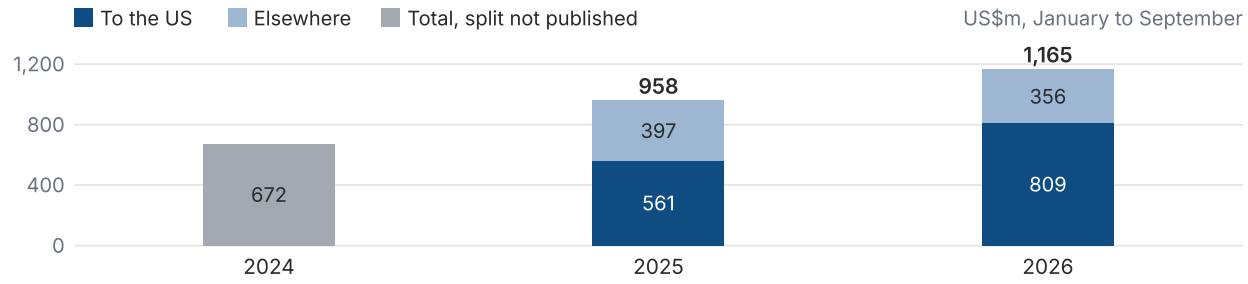
Request, technical review, approval, permission to energize and observed operation require separate evidence. No universal grid-easing date is inferred.

Sources: CBRE; Electronic Times; Segye Ilbo; Utility Dive (2026); ERCOT September operations; JLL; IEA; NRG calculation

Korea's large-transformer exports have shifted to the United States, which now takes over two-thirds.

Press reporting of customs data puts Korea's exports of transformers above 10,000 kVA at $1,165m in January–September 2026, up 21.6% year on year. The underlying HS extraction was not independently reproduced, so this remains a dated press-carried series.

Korea's exports of transformers above 10,000 kVA, January to September, US$m; the 2024 split is not published



Source: Korea Customs Service via The Korea Economic Daily (October 2026); NRG calculation (elsewhere = total less US)

The full year 2025 was Korea's best for large-transformer exports since records began in 1977, at $1,299m. US-bound shipments reached $809m in the first nine months of 2026, against $561m a year earlier, while total exports were 73.4% above the $672m of the same months of 2024.

The US share rose from about 57% in full-year 2025 to 69.4% in January–September 2026 (NRG calculation), across different-length periods. Export value reflects both prices and units. These large transformers serve many grid and industrial uses; the customs series does not isolate AI demand or installed equipment. The 2024 destination split is unavailable in the cited report.

Wood Mackenzie counted nearly $1.8bn of North American transformer-plant investment announced since 2023 and expects a supply response. Its 2025 demand-over-supply figure and 2026 shortage measure have different definitions. They cannot establish the size or direction of a comparable year-on-year deficit without a common denominator.

Building-level equipment could ease before utility-class transformers if new lines are commissioned, qualified and able to keep pace with demand. That is an NRG scenario, not a demonstrated clearing order. Announced investment is not operating output. Factory completion, accepted units, cancellations and new load all matter; materials and contract lags can keep prices high after delivery quotes shorten.

US share of Korea's large-transformer exports, by period

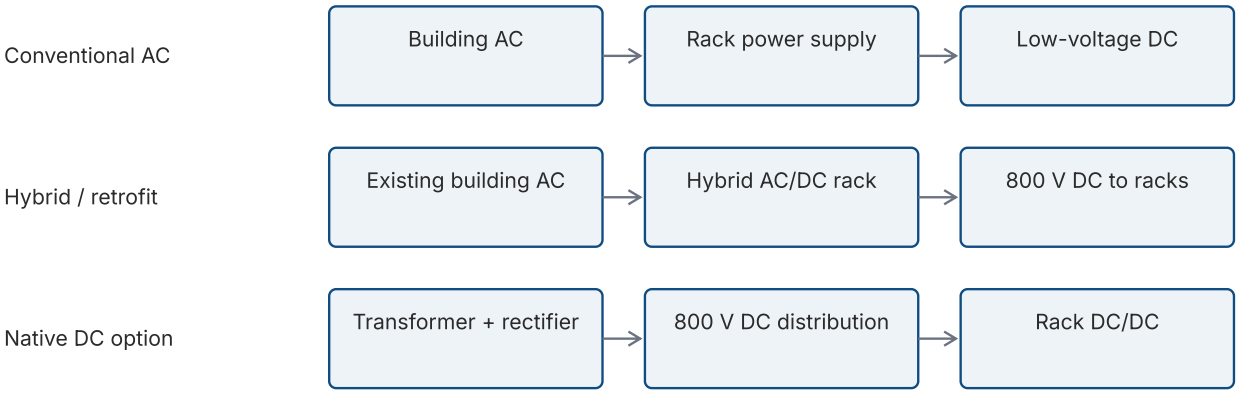
Period	US share of value
2025, full year	about 57% (NRG calculation)
2026, January to September	69.4%

Sources: Korea Customs Service and Korea International Trade Association via press; Wood Mackenzie; NRG calculation

800-volt DC has native and hybrid paths, with different equipment consequences.

Higher-voltage distribution reduces current at a given power. The August 2026 platform explanation includes a hybrid power rack retaining existing building AC infrastructure and a row power center expected in 2027. Architecture selection need not replace the whole building power chain.

Schematic alternatives, not complete electrical designs; SST is one possible upstream conversion route



Source: Platform technical notes (2025, August 2026); OCP (March / June 2026); NRG analysis

The earlier technical design note claims up to 5% better end-to-end efficiency and 45% less copper under its design comparison. It says a 1 MW rack at 54 volts could require up to 200 kg of busbar copper. These are vendor design claims, not measured gains across an installed fleet. Results depend on topology and the equipment compared.

The hybrid path converts existing AC supply near the racks and can preserve upstream building assets. A native DC path moves conversion upstream and changes distribution interfaces. Traditional transformers with rectifiers and solid-state approaches are alternatives; 800-volt DC does not universally require a solid-state transformer.

The August explanation expected hybrid power racks in the second half of 2026 and a row power center in 2027. TrendForce's June roadmap expected optional shipments from the third quarter of 2026 and wider adoption in 2028. Expected availability, shipment and a commissioned hall are separate milestones; no deployment is inferred here.

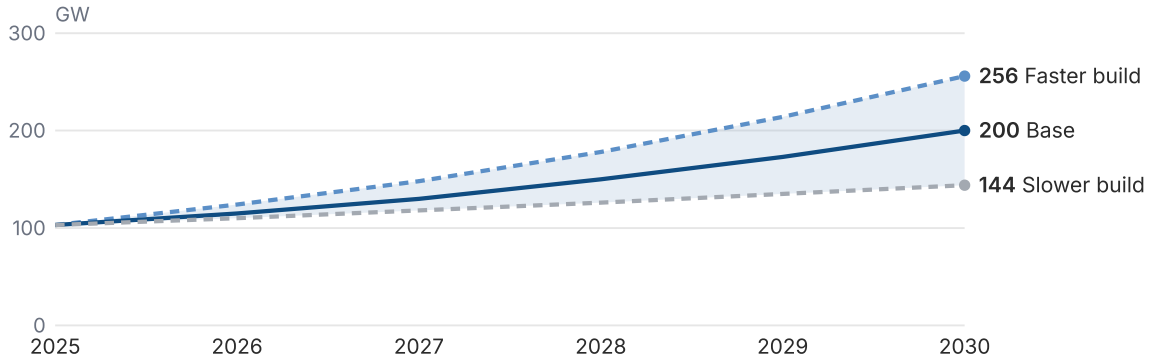
New converters and DC distribution can create fresh procurement needs where selected, while retained AC assets need no automatic rebooking. Buyers must establish ratings, protection, compatibility and certification at the interfaces. OCP's March white paper and June SST specification describe engineering work; neither establishes a universally mandatory topology or a completed safety-code transition.

Sources: Platform announcements (2025, August 2026); TrendForce (June 2026); OCP (March / June 2026)

How fast the queue clears depends on how fast the world builds.

JLL's published base case grows global capacity from 103 GW in 2025 to about 200 GW in 2030. Compounding its faster and slower rates for five intervals gives 256 and 144 GW (NRG calculation). These paths begin in 2025, not at this report's revision date.

Global data-center capacity, GW, 2025 to 2030; faster and slower paths are NRG calculations from JLL's growth rates



Source: JLL (2026); NRG calculation

FASTER BUILD	256 GW in 2030	At 20% annual capacity growth from 2025, demand could outrun commissioned and qualified equipment output. Lead times could lengthen again. On-site power would still require fuel, permits and integration.
BASE	200 GW in 2030	JLL's rounded endpoint is about 200 GW; exactly compounding 14% gives about 198 GW. Earlier easing in some equipment requires supply additions to beat demand. Grid risk remains site-specific.
SLOWER BUILD	144 GW in 2030	At 7% annual growth from 2025, financing limits or efficiency gains could relieve equipment demand. Whether prices follow depends on materials, product mix and contract lags; grid dates need independent evidence.

SIGNPOSTS

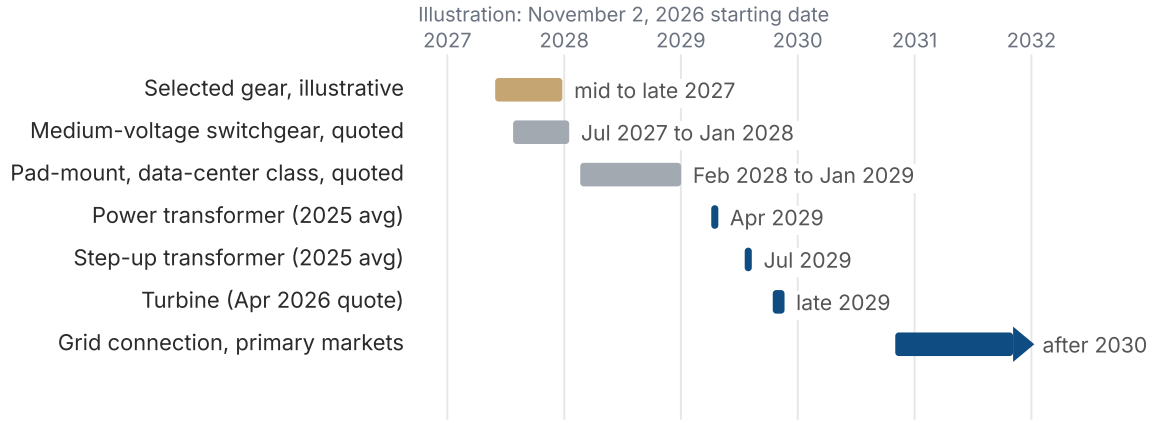
- US construction: September 2026 preliminary SAAR $96.5bn, up 90.3% year on year
- Commissioned and qualified equipment output versus new orders and cancellations
- Grid requests, approvals and observed operation measured separately
- Hybrid versus native 800-volt adoption and actual commissioning

Sources: JLL (2026); US Census Bureau; Utility Dive; IEA; NRG calculation

Published lead times illustrate delivery dates; usable power can arrive later.

Apply selected published lead times to a November 2, 2026 starting date (NRG calculation). Historical means and quoted ranges are not current offers. Delivery dates differ from an operating date, and the evidence gives no universal end date for grid constraints.

Illustrative delivery dates from dated public inputs; gold: NRG estimate. Grid row has no fixed upper bound



Source: Wood Mackenzie (2025); Cushman & Wakefield (2026); JLL (2026); company filings and announcements; NRG estimate; NRG calculation

SELECTED HALL GEAR	illustrative 2027	NRG estimate: 30–60 weeks for selected electrical items and chillers, not all equipment. The upper end is about 14 months. Generator and pad-mount quotations can run longer.
TRANSFORMERS	illustrative 2028–29	The dates reuse 2026 quoted pad-mount ranges and Q2 2025 utility-unit means. They are calculations, not current offers or a forecast of market-wide delivery normalization.
TURBINES AND GRID	separate constraints	An April 2026 turbine lead-time statement is applied illustratively. A connection wait above four years reaches beyond 2030 from this baseline; it provides no maximum or global clearance date.

SIGNPOSTS

- Wood Mackenzie's next transformer lead-time and shortage readings
- Turbine reservations set against conversions to firm orders in company filings
- Switchgear producer prices, up 11% so far in 2026
- Korea's grid-review approvals and its final demand plan

Sources: Wood Mackenzie; Cushman & Wakefield; JLL; BLS; company filings and announcements; NRG estimate and calculation

METHOD

How this report was built.

Originally published November 2, 2026; revised and checked November 5, 2026. Public inputs include international agencies, national laboratories and statistics agencies, market research, issuer disclosures and press-carried data. Observation periods and forecast horizons remain separate. Dated benchmarks are retained where a like-for-like current primary table could not be established.

An NRG estimate is a numerical assessment informed by Nathan Research Group's own research and cross-checked against public evidence. This report retains one NRG estimate: an illustrative 30–60 weeks for selected US building electrical items and chillers. It is based on limited research, not a representative distribution, and cannot cover an entire equipment package.

Public cross-checks include JLL's approximate 2026 Americas component means and longer quoted ranges carried by the press from Cushman & Wakefield's guide. They differ in geography, product and method. The latter detailed table, the customs HS extraction and the newest comparable grid-request totals were not independently reproduced. These series are dated and qualified rather than certified as the latest.

An NRG calculation uses public inputs alone: growth rates, rebased price indexes, ratios, unit conversions and illustrative dates. The 2025–2030 capacity paths have five compounding intervals; the report does not start a new five-year forecast in November. Annual-rate construction means are not cumulative expenditure. Different grid stages are not a conversion funnel.

General operating mechanisms are NRG analysis informed by research and public evidence, not estimates of prevalence. Equipment arrival does not prove usable power; reservations reduce selected risks without closing all project conditions. Companies are described rather than named. Public examples establish only the terms and periods disclosed.

This is business and market research, not investment advice.

WORKING WITH NATHAN RESEARCH GROUP

Partner With Nathan Research Group

Filings show how large the order books have grown and the IEA shows how large the demand is, but neither shows the queue from the inside: how long-lead equipment is planned, reserved and delivered. That knowledge sits with the people who design, make, buy and connect AI halls. Reaching them, compliantly, is what we do.

WHO WE PUT IN THE ROOM

Data-center operators and developers	Site selection, power procurement and phasing of AI campuses.
Electrical and cooling equipment makers	How electrical and cooling equipment is specified, built and delivered across the industry.
Grid and utility planners	Interconnection queues, substation build times and Korea's review of large new loads.
Power-distribution engineers	Electrical design for high-density data halls.
Supply-chain analysts	Bottlenecks in electrical steel, copper and power semiconductors.
Energy-market analysts	Data-center load forecasts and the demand assumptions behind national power plans.

If your team is evaluating long-lead power and cooling equipment, the AI data-center power and cooling chain, or the broader data-center infrastructure market, tell us the decision you're trying to make.

START THE CONVERSATION

Web	nathanrg.com
Email	team@nathanrg.com General enquiries nathan@nathanrg.com Nathan Choi · CEO
Phone	+82 10 3059 3630 Nathan Choi · CEO +82 2 939 3630 Office · Seoul

Based on public sources and NRG's own research. Figures marked NRG estimate are rounded ranges checked against public data. Nathan Research Group does not request or facilitate material non-public information, and runs every engagement through a documented compliance protocol.