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

AI is on course to double the electricity data centers use by 2030. The gear that delivers that power, and the grid connections behind it, now take months to years to arrive, so buyers secure factory output before they place orders. This report measures the queue and how it should clear.

950 TWh

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

30–60 weeks

Typical order to delivery for a US data hall's building-level electrical and cooling gear, NRG estimate

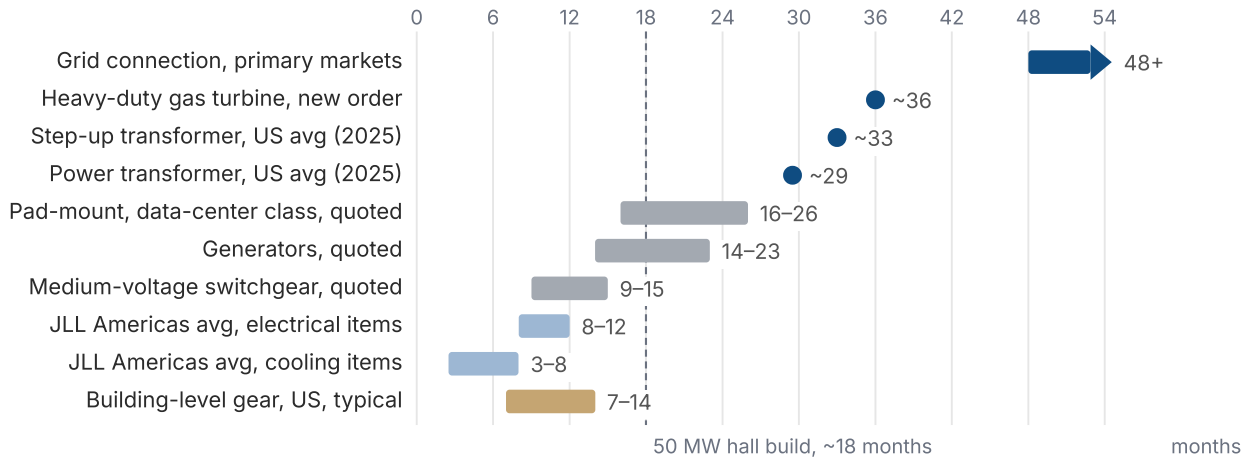
4+ years

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.

Order to delivery, months. Gold: NRG estimate. Grid row: wait for a connection



Source: JLL (2026); Wood Mackenzie (Q2 2025, US); Cushman & Wakefield (2026); company filings and announcements; NRG estimate; NRG calculation (weeks to months)

Data centers used about 485 TWh of electricity in 2025, 17% more than a year earlier, and the IEA expects about 950 TWh by 2030. It has not raised that path despite record spending, because turbines, transformers and grid connections are in short supply.

On an NRG estimate, the building-level electrical and cooling gear for a US data hall typically arrives 30–60 weeks after order, in line with JLL's 2026 Americas averages. Quoted ranges for the slowest items run longer, and a utility-class power transformer took about 128 weeks in 2025. A 50 MW hall takes about 18 months to build.

So the slot is booked first and the price settled later. In one published reservation, a power producer agreed in 2026 to pay a non-refundable deposit for a turbine slot with delivery by 2030. At one large turbine maker, reservations made up roughly half of gas commitments through 2026.

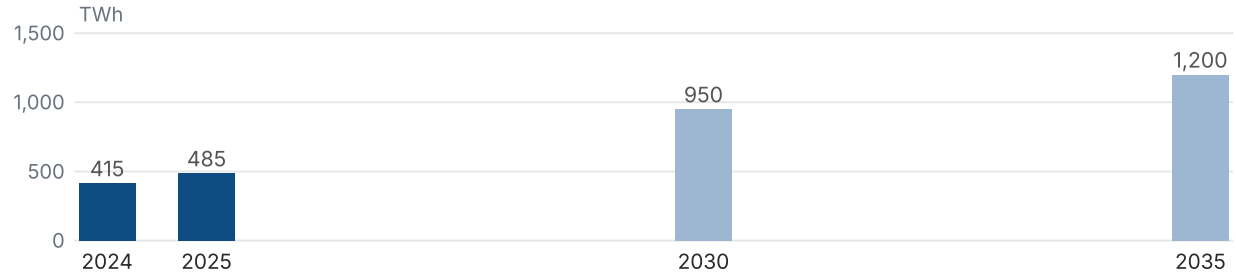
Texas has about 474 GW of large-load requests, mostly from data centers, and in Korea's capital region 3% of the capacity data centers asked for has cleared grid review. NRG estimates that grid access stays the binding constraint until about 2029–2032.

An order placed now for typical building-level gear lands in 2027, a utility-class transformer in 2029 and a new grid connection after 2030 (NRG calculation). 800-volt DC, due with the 2027 rack generation, then opens a new round of bookings.

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 that AI-focused use to triple by 2030. Its base case implies growth of about 14% a year from 2025 (NRG calculation) and 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 agrees: the dates that set the pace belong to factories and grid operators.

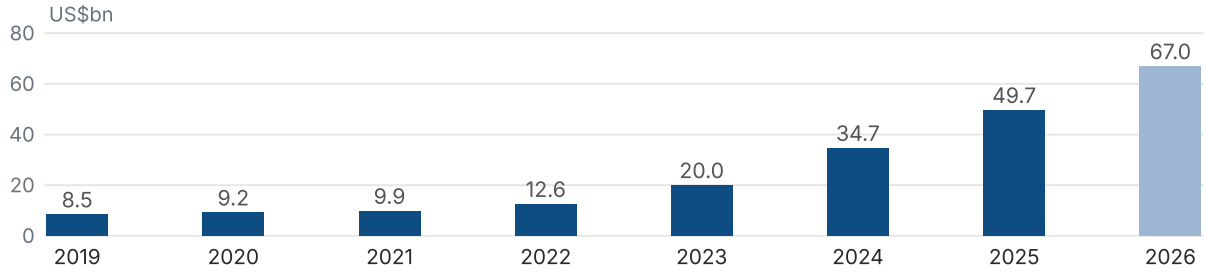
US data-center electricity use in 2030, Berkeley Lab cases

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

US data-center construction spending has quintupled since 2021.

US spending on data-center buildings ran at an annual rate of about $85bn in August 2026, the latest month available. The Census counts the buildings and their plant, not the servers inside.

US private data-center construction, US$bn, yearly average of seasonally adjusted annual rates; 2026 (lighter) is January to August



Source: US Census Bureau; NRG calculation

On yearly averages of the Census series, spending was $9.9bn in 2021 and $49.7bn in 2025, about five times as much (NRG calculation). Most of the rise came after 2022, and the August 2026 rate was 73% above a year earlier. Worldwide, JLL counts 103 GW of data-center capacity in 2025 and expects 97 GW more by 2030.

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 2026 guide finds US and Canadian costs up 21% per MW since late 2024 and blames supply limits, skilled-labor wages and component prices. Its all-in greenfield figures run from $8.9m to $23.3m per MW before chips. Power is the largest line; core and shell with sitework take 17%, contingencies 16% and site acquisition 7%.

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.

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

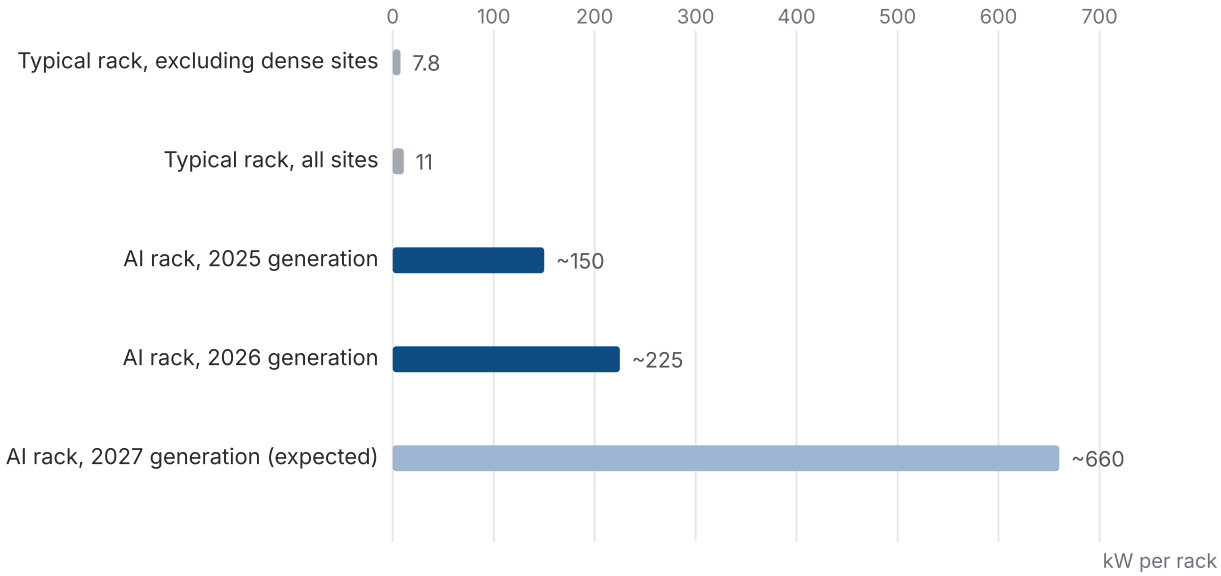
Sources: US Census Bureau; JLL (2026); Cushman & Wakefield (2026); NRG calculation

03 · THE HALL

A current AI rack draws about 29 times the power of a typical rack.

Uptime Institute's 2026 survey puts the typical rack at 7.8 kW outside the few very dense sites. Current AI racks draw about 225 kW, and the 2027 generation is expected to reach 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 ratio in the headline is an NRG calculation from the two surveys. Even the average is moving: Uptime's typical rack rose from 9 kW in 2025 to 11 kW in 2026 as a handful of very dense sites came online, and JLL expects average density to reach 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.

The 2027 generation, at about 660 kW, would draw about 85 times a typical rack (NRG calculation), which is why the power architecture inside the rack has to change. Rack-scale systems ship with published power and cooling requirements, so a hall's plant is sized to that platform and must change when the platform does.

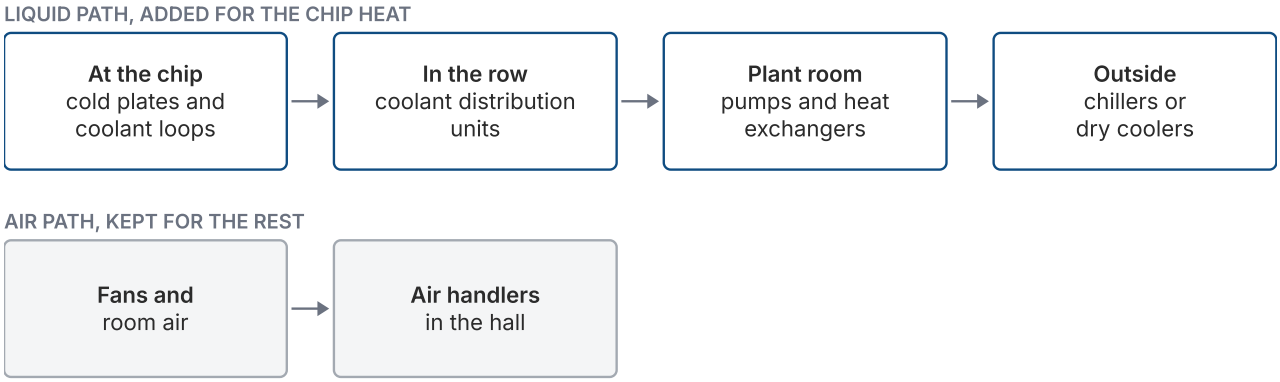
Liquid-cooled halls cost about 10% more to build than air-cooled ones, by JLL's estimate, and the premium buys plant rather than floor space. Cooling moves from mostly air to liquid at the chip, and power inside the rack from 54-volt DC to 800-volt DC from 2027–28.

Sources: Uptime Institute (2026); TrendForce (2026); JLL (2026); company filings and announcements; NRG calculation

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

Market research on data-center infrastructure puts liquid-cooling revenue at about $3bn in 2025 and about $7bn in 2029, roughly 24% a year. That outpaces the capacity it serves, because each AI hall needs more of it.

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 sizes the whole physical-infrastructure market, power and cooling together, at $120bn by 2030, growing 22% a year, against 14% a year for global data-center capacity in JLL's base case. Thermal management is the fastest segment and liquid cooling the fastest technology, and water-cooled chillers should outgrow air-cooled ones, which moves more of the cooling plant into the same factory queues as the electrical gear.

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.

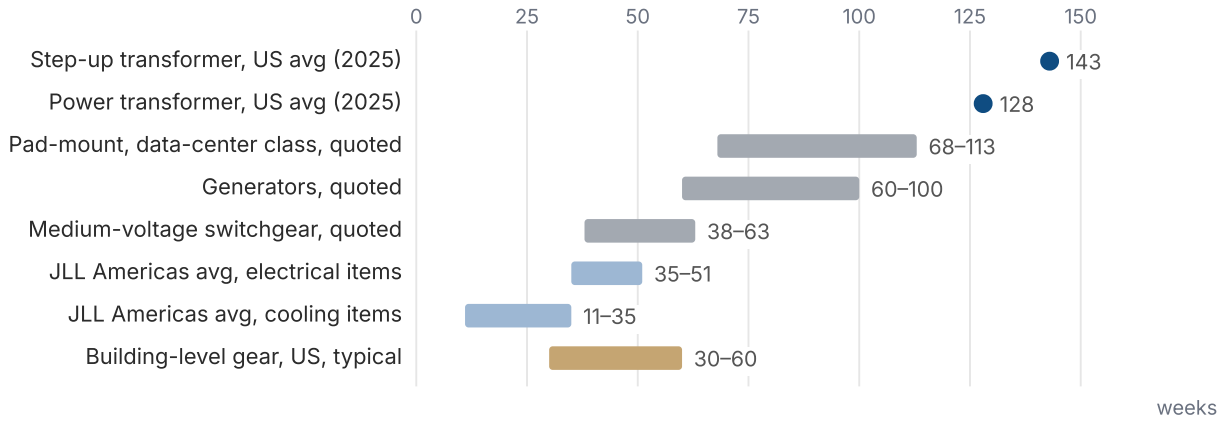
Cooling plant is building-level gear: it is sized to the rack generation and made to order. Its price has risen too, and Cushman & Wakefield puts cooling equipment up 32% between December 2021 and early 2026, a little more than transformers. The rest of the cooling plant is quicker: JLL's Americas figures for cooling items run 11–35 weeks, against about 35–51 weeks for the main electrical items.

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

Typical building-level gear arrives within about a year of order.

Equipment lead times stopped lengthening in 2026 but remain far above 2019 levels. The hall's own gear arrives in months, while the utility-class transformers behind it take two to three years.

Order to delivery, weeks; quoted ranges and averages are 2026, transformer averages Q2 2025. Gold: NRG estimate



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

Wood Mackenzie's survey put US power transformers at 128 weeks in the second quarter of 2025 and generator step-up units at 143. Demand for power transformers was 116% above its 2019 level, demand for step-up units 274% above, and annual demand for power transformers ran 39% above available supply. The IEA found waits for transformers and cables had doubled in three years.

Those are utility-class units at the substation. Inside the fence, lead times are shorter. JLL puts the 2026 US average for data-center equipment at 42 weeks, 83% longer than in 2019, and its Americas figures for the main electrical items run about 35–51 weeks.

Cushman & Wakefield's 2026 guide quotes wider ranges: 38–63 weeks for medium-voltage switchgear, 60–100 for generators and 68–113 for data-center pad-mount transformers. The averages describe a typical order, and the quoted ranges show how long the slowest orders can run.

A US order restricting Chinese bulk-power equipment could lengthen the line again. Wood Mackenzie puts the 2026 shortage of power transformers at 15% and expects the order, which touches an import market worth more than $22bn since 2025, to fall mostly on data centers.

> **30–60 weeks**
>
> NRG estimate: typical order to delivery across a US data hall's building-level electrical and cooling gear. It sits with JLL's 2026 Americas averages.

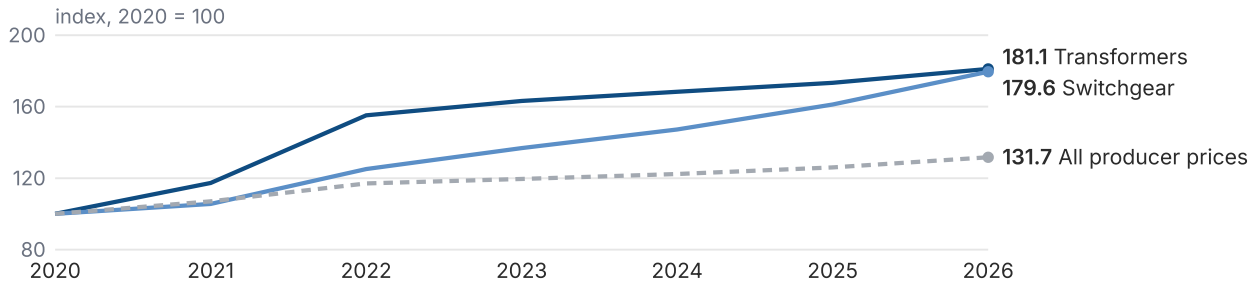
Sources: Wood Mackenzie (2025, 2026); Cushman & Wakefield (2026); JLL (2026); IEA (2025); NRG estimate

06 · THE PRICE

Since 2020, transformer and switchgear prices have risen far faster than producer prices.

Federal producer price indexes for transformers and switchgear stood 86% and 88% above their 2020 averages in September 2026, while producer prices as a whole were 35% higher. Switchgear, which lagged until 2024, has now caught up with transformers.

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 index stood at 475, against an average of 255 in 2020. Producer prices record what factories charge, not what a late project loses, so they understate what the queue costs a buyer.

Transformers are still scarce, but their prices have flattened, while the building-level electrical gear every hall needs is where pricing now moves fastest. Cushman & Wakefield's guide, in the table, ranks the items the same way, and copper, the metal in every winding and busbar, rose 46% over the same years, which feeds straight into both.

Prices and lead times are two readings of one shortage. In Uptime's 2026 survey, 72% of operators were at least somewhat concerned about supply-chain disruption. When both readings ease for an item, its line can be booked later.

Price change, December 2021 to early 2026

Item	Change
Switchgear	+60%
Generators and motors	+33%
Transformers	+29%

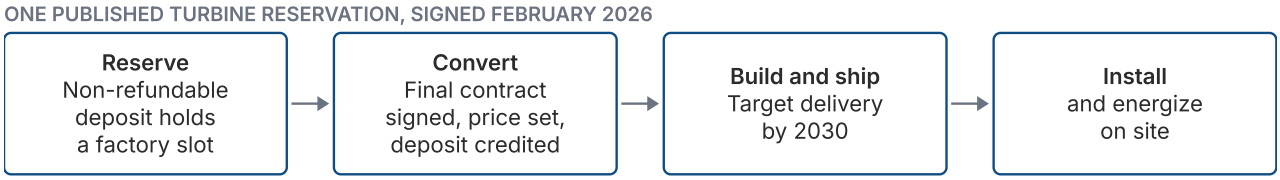
Sources: US Bureau of Labor Statistics via FRED; Cushman & Wakefield (2026); IEA (2025); NRG calculation

07 · THE RESERVATION

A reservation fixes the delivery date and leaves the price for later.

A Canadian power producer's February 2026 release shows how a slot reservation works. The same logic, paying to hold factory output ahead of a firm order, now reaches the hall's own gear, where developers preorder months ahead.

Anatomy of a slot reservation, from one published example



Source: company filings and announcements

The buyer pays a non-refundable deposit to hold a manufacturing slot with a target delivery year. Both sides then negotiate a final contract, which sets the price, and the deposit counts toward it. The buyer has fixed when; what it pays stays open. Here the slot was booked in 2026 for a turbine due by 2030, roughly four years ahead. The deposit is credited to the final price, so it is a cost only if the project falls away.

Developers apply the logic to the hall's own gear, preordering materials and holding stock. A booking also gives the supplier an early read on demand, which helps it decide when to add a line and how to price it.

Buyers pay for the date because JLL names speed to power as the first test in choosing a site, and when the slowest item takes years, a late delivery costs far more than the deposit. The reservation turns an open-ended risk into a known payment, and the same terms suit any project whose start date hangs on one slow item.

For suppliers, reservations fill factory schedules before orders are firm; at one large turbine maker they briefly exceeded firm gas orders in 2026. Reservations move the queue earlier without shortening it.

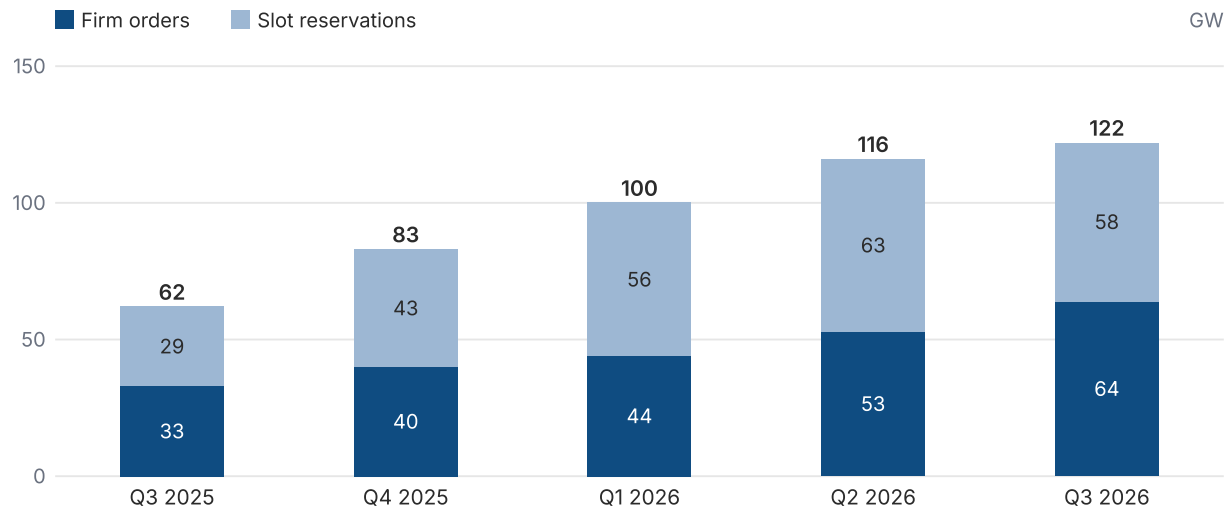
> **24 months**
> How far ahead large developers preorder select materials, by JLL's 2026 count. They also keep 6–12 months of critical components in stock.

Sources: company filings and announcements; JLL (2026)

Gas-turbine makers are booked into 2029–30, much of it by reservation.

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 rose to more than half of the maker's commitments early in 2026 and have since begun to convert. Between July and September it turned 13 GW of reservations into orders and signed 8 GW of new ones, and it expects at least 130 GW under contract by the end of the year. The largest makers' delivery slots now run into 2029–30.

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.

Its chief executive put lead times on new heavy-duty orders at about three years and said the turbine is often not what holds a project back; construction, permits and fuel supply are. NRG's research points to the grid date as what pushes a campus toward on-site generation.

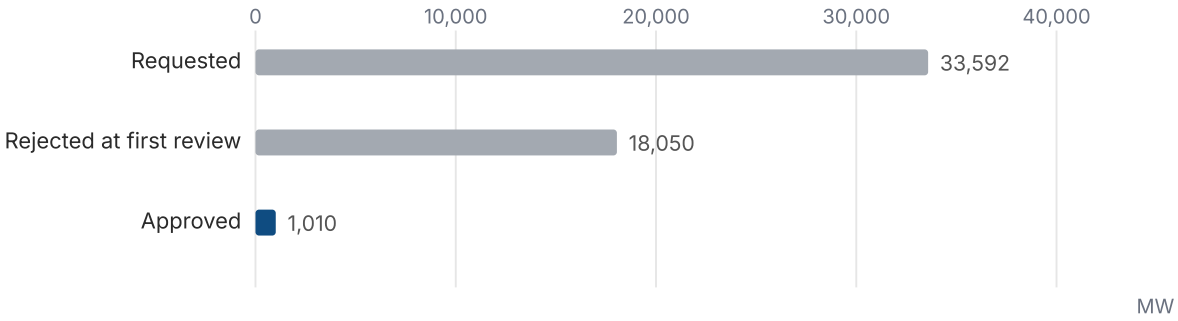
The IEA counts 47 GW of on-site gas projects for US data centers under development or proposed, and expects 15–27 GW of on-site gas to serve data centers worldwide by 2030. Orders for gas turbines to serve data centers rose 70% in 2025, by the agency's count. Those plants need transformers and switchgear 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)

Grid approvals are the tightest gate: Korea's capital region has cleared 3% of requested capacity.

In Korea's capital region, data centers asked the grid for 33,592 MW through March 2026 and 1,010 MW was approved. In Texas, requests to connect large loads have reached about 474 GW.

Korea, capital region: data-center grid-impact reviews, cumulative to March 2026, MW



Source: Electronic Times (July 2026)

The approved share by capacity, 3% of the requested, is an NRG calculation from that count. By number, 10 of 522 applications were approved, or 1.9%, and more than half the requested megawatts were turned down at the first review. A first-stage rejection means the grid cannot supply the load on current plans.

Data centers filed 88 applications for electricity supply with the national utility in the first eight months of 2026, against 46 in all of 2025, and two-thirds of the capacity they asked for now sits outside the capital region. In August the government raised its 2040 peak-demand forecast to 158–165 GW, from 132–138 GW in April.

Texas shows the same pressure at larger scale. Large-load requests total about 474 GW, roughly 90% from data centers and more than five times the grid's record peak, and the grid operator has delayed a transmission study while the state audits the data centers in the queue. Neither queue shows how many requests are firm projects.

The IEA estimated in 2025 that grid constraints could delay about 20% of the data-center capacity planned to 2030, and JLL puts the average wait for a connection in primary data-center markets above four years. For a developer the effect is the same in both places: the connection date, not the equipment, sets when a hall can run.

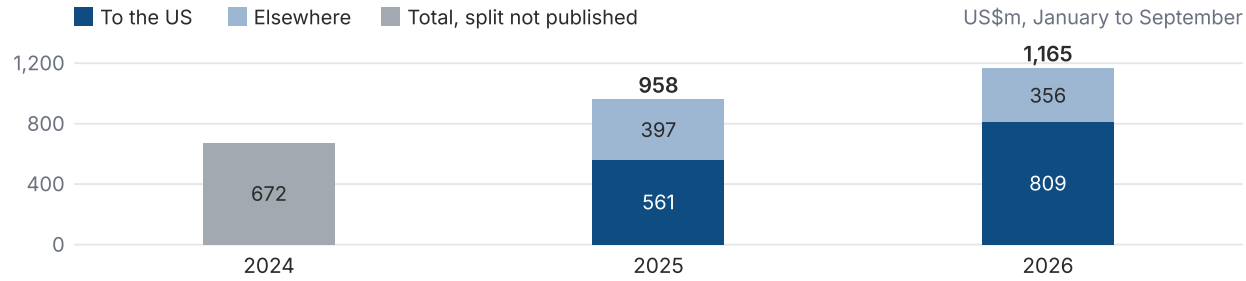
> ## 2029–2032
>
> NRG estimate: grid access remains the binding constraint for large AI campuses until about 2029–2032, judging by grid waits, turbine slots and transformer supply.

Sources: Electronic Times; Segye Ilbo; Utility Dive (2026); JLL (2026); IEA (2025); NRG calculation and estimate

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

Korea's exports of transformers above 10,000 kVA reached $1,165m in the first nine months of 2026, up 21.6% on a year earlier, and US-bound shipments grew faster still.

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 has more than doubled since 2023, as the table shows. Export value also rises with price, so part of the growth reflects higher transformer prices rather than more units. Large transformers go into AI campuses and into the lines that feed them, so both kinds of build pull on the same exports, and the customs series cannot tell them apart.

Wood Mackenzie counted nearly $1.8bn of transformer-plant investment aimed at North America since 2023 and expects the deficit to narrow, though its 2025 deficit and 2026 shortage figures are not strictly comparable. Read together, they suggest the gap is narrowing, but slowly.

NRG expects building-level gear to clear first, because its lines are smaller and quicker to add. Large power transformers need bigger plants and long test cycles, so high-voltage units come last. A new plant adds nothing until it is built and its first units pass testing, so investment announced in 2024 and 2025 reaches lead times only later in the decade.

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

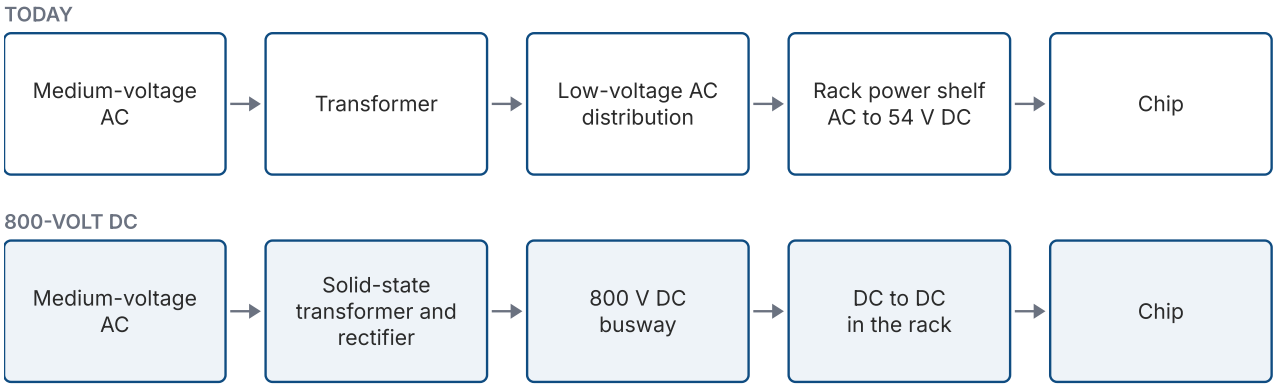
Period	US share of value
2023, full year	31.9%
2025, full year	about 57% (NRG calculation)
2026, January to September	69.4%

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

800-volt DC arrives with the 2027 rack generation and adds items to the equipment list.

Above about 200 kW, the 54-volt power used inside today's racks runs into physical limits. The leading AI-chip maker plans to move to 800-volt DC, which changes what a hall has to order.

Power path from the hall edge to the chip, today and with 800-volt DC



Source: company filings and announcements

The chip maker's design note puts the gain at up to 5% better end-to-end efficiency and 45% less copper, and says a 1 MW rack fed at 54 volts would need up to 200 kg of copper busbar. Current racks, at about 225 kW, still fit conventional power supplies. The 2027 generation changes the power architecture.

Converting grid power to 800-volt DC once, at the edge of the hall, removes most of the intermediate conversion steps, which is where the efficiency gain comes from. TrendForce said the first 800-volt racks were due to ship as an option from the third quarter of 2026. The next rack generation, at about 660 kW, follows in the second half of 2027, with wider deployment in 2028.

For the queue, a new architecture means new reference designs, and new designs bring a new round of reservations. Solid-state transformers, rectifiers and DC busways are items few halls have bought at scale so far. TrendForce estimates that one 800-volt power rack can serve one or two of the next-generation racks, depending on redundancy.

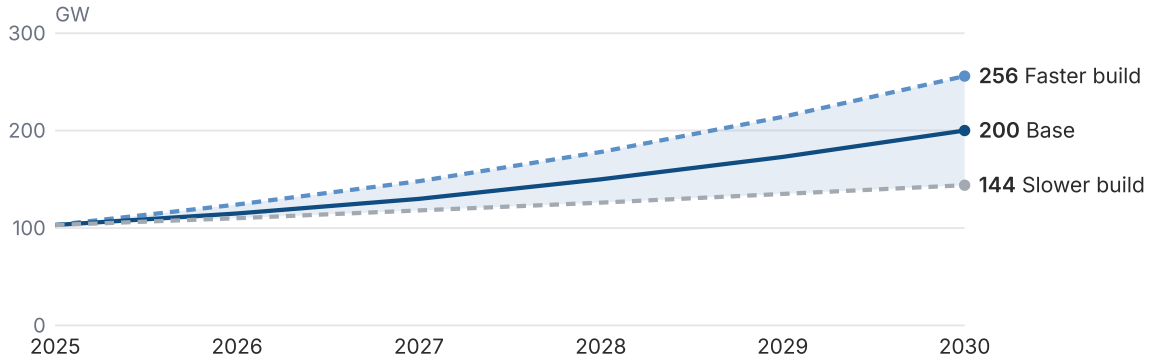
Buyers planning halls for 2028 and later must either book today's equipment now or wait for the new list and risk losing the slot. If the 2027 generation lands on time, part of the building-level queue starts again just as the old one eases. Until then, the gear booked for 2026 and 2027 halls is the AC equipment in the upper row of the diagram.

Sources: company filings and announcements; TrendForce (2026)

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

JLL's base case takes global data-center capacity from 103 GW in 2025 to about 200 GW in 2030; its faster and slower growth rates give 256 and 144 GW (NRG calculation). The cases are illustrative, not forecasts.

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	AI demand adds to the base at 20% a year and more campuses build their own power. Building-level lead times lengthen again, and reservations spread to the hall's own gear.
BASE	200 GW in 2030	Capacity grows 14% a year while grid waits stay above four years. Building-level gear eases first; utility-class transformers and grid connections stay tight to about 2030.
SLOWER BUILD	144 GW in 2030	Efficiency gains or tighter financing hold growth to 7% a year. Building-level lead times fall back toward pre-2020 levels, while transformers and grid access still bind.

SIGNPOSTS

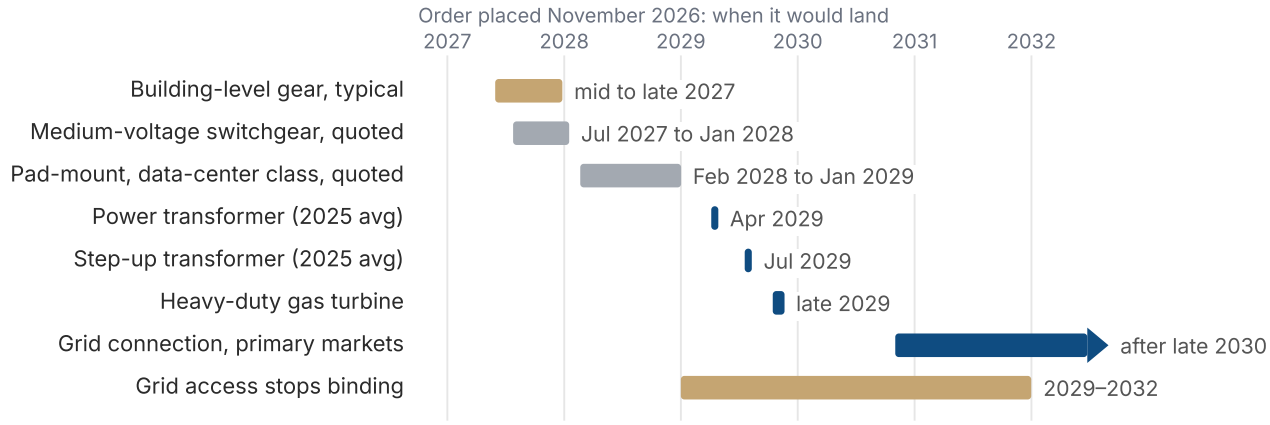
- US data-center construction spending, 73% above a year earlier in August 2026
- The Texas audit of about 474 GW of large-load requests
- The IEA's next demand update, against its 950 TWh path for 2030
- Adoption of 800-volt DC after the 2027 rack generation

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

An order placed now lands between 2027 and 2030, and grid access eases last.

The chart dates an order placed in early November 2026 against today's lead times (NRG calculation). Building-level gear lands within a year; transformers, turbines and grid connections take until the end of the decade.

When an order placed in November 2026 would land, by item; gold rows are NRG estimates, other dates NRG calculations from public lead times



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

HALL'S OWN GEAR	lands 2027–28	Typical orders land within a year (NRG estimate: 30–60 weeks). JLL found lead times flat in 2026, but switchgear prices rose another 11%, so price eases after lead time.
TRANSFORMERS	lands 2028–29	Data-center pad-mounts land in 2028 and utility-class units in 2029. The 2026 shortage is 15%, and new North American lines add output only later in the decade.
TURBINES AND GRID	eases 2029–32	Heavy turbines ordered now arrive about three years out. Grid connections take longer, and NRG estimates grid access binds until about 2029–2032.

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.

Figures with a cited source are public. They come from an international agency (the IEA), a national laboratory and statistics agencies (Lawrence Berkeley National Laboratory, the US Census Bureau, the US Bureau of Labor Statistics, Korea Customs Service and the Korea International Trade Association), market-research firms (Wood Mackenzie, JLL, Cushman & Wakefield, TrendForce, Uptime Institute and a firm that tracks data-center infrastructure), company filings and announcements, and press reports of grid and government data. Each figure is the newest release we could find on the data-as-of date; older series carry their year.

An NRG estimate is a figure from Nathan Research Group's own research, rounded into a deliberately wide range and published only where the newest public data point the same way. This report carries two.

The two NRG estimates and their public cross-checks: typical order to delivery for a US data hall's building-level electrical and cooling gear, 30–60 weeks, sits with JLL's 2026 Americas averages of about 35–51 weeks for the main electrical items; grid access binding until about 2029–2032 rests on JLL's four-year-plus connection waits, turbine slots booked into 2029–30 and a 15% power-transformer shortage in 2026.

An NRG calculation is a figure we derived from public numbers alone: growth rates, price indexes rebased to 2020, ratios, unit conversions, the dates on which an order placed now would land, and the faster and slower capacity paths built from JLL's growth rates. Its inputs are the sources named under each chart.

Where the text describes industry practice, it is NRG's own analysis of its research and the public evidence cited, written as general practice rather than about any one company. Where one company's public disclosure is used, the text says so. Companies are described rather than named.

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.