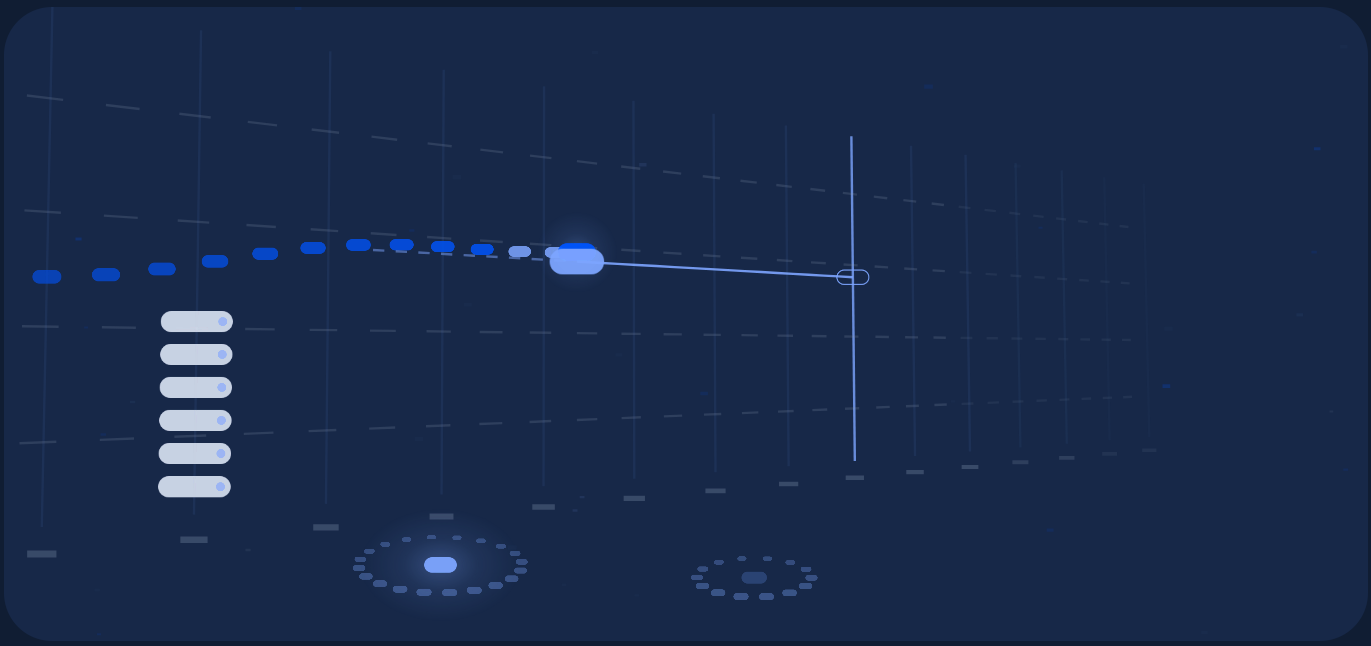
Rillor

RESEARCH REPORT

What sets GPU system lead times in 2026

Filings and earnings calls show what sets a GPU server lead time: packaging, high-bandwidth memory, substrates and testing, system builders and powered buildings.



Rillor

Contents

01 Summary

02 Background

03 Method

04 Findings

05 Implications

06 Limits of this analysis

07 Sources

Summary

- No chip maker or system builder in these sources gives a lead time for a named GPU server model or a GB300 NVL72 rack. NVIDIA reports lead times "of more than 12 months" for certain supply [1], and Super Micro says its lead times have "extended beyond twelve months for some products" [5]. Week counts for finished servers appear on a seller's page [32].
- Packaging and memory are the upstream inputs the makers describe as short. TSMC's chief executive said in July 2026 that its packaging capacity was "limiting my customers' growth" [25]. Micron has completed agreements for the vast majority of its calendar 2027 HBM bit supply [12]. NVIDIA's supply commitments rose from $119 billion to $279 billion in the quarter to 26 July 2026, primarily for memory and manufacturing facilities [2].
- NVIDIA's announcement of 21 October 2026 dates most of its new substrate, interposer, testing and optics capacity between February 2027 and the end of 2028 [4].
- System builders tie delivery timing to components, customer readiness and acceptance. Dell held a $95 billion AI server backlog at 31 July 2026 [8] and cites "the varying stages of customer readiness" [9]. HPE cites "longer customer acceptance timelines on AI-related orders" [10].
- Powered buildings set the last date. NVIDIA calls securing land, power and shell "the next critical phase" of the buildout [3]. In North America's primary markets, 80.4% of capacity under construction was preleased in the first half of 2026 [28]. The Department of Energy puts lead times for critical grid equipment at two or more years [23].
- Two disclosed delays show both ends of the chain: a late powered shell at CoreWeave in late 2025 [17], and 16 B300 servers at SANGRIX whose delivery moved from the third quarter of 2026 to the week of 3 November 2026 [18, 19].

Background

A GPU system lead time runs from a signed order to a system that has passed acceptance in a powered building. Each stage on that path has its own clock.

A data-center GPU is a package. The GPU die is made on a leading-edge wafer. Stacks of high-bandwidth memory (HBM) sit beside it on an interposer, joined by an advanced packaging process such as TSMC's CoWoS, and the assembly is mounted on an organic substrate. The package is tested, placed on a board and built into a server or rack with CPUs, system memory, solid-state drives, network cards and, in newer designs, optical engines. The rack then needs a building with power and cooling. NVIDIA's user guide puts the power draw of a DGX GB rack at approximately 120 kW [22], and CBRE writes that "100-kW GPU servers can only be cooled by a closed-loop liquid-cooling system" [28].

Makers plan ahead of orders. NVIDIA's annual report says: "We have long manufacturing lead times and build finished products and maintain inventory in advance of anticipated demand" [1]. Generations turn over quickly. NVIDIA began shipping production units of GB300 in the second quarter of its fiscal year 2026 [1], and Vera Rubin began production shipments in the third quarter of fiscal year 2027, the quarter that began after 26 July 2026 [2].

Demand is the backdrop. NVIDIA's Data Center revenue was $89.0 billion in the quarter to 26 July 2026, up 117% from a year earlier [3]. On 21 October 2026 its chief executive said: "Demand is extraordinary. Capacity is the constraint" [4].

Seller pages answer "how long?" with a week count. This report traces where the weeks come from, through the disclosures of the companies that make, package, test, assemble and house GPU systems.

Method

What was examined. Statements about lead times, supply constraints, supply commitments, dated capacity additions, backlog, customer readiness, acceptance and delivery delays, each recorded with its date, its scope and the words used.

Sources.

- Filings with the U.S. Securities and Exchange Commission: NVIDIA [1, 2, 3], Super Micro [5, 6, 7], Dell [8, 9], Hewlett Packard Enterprise [10], Micron [11], CoreWeave [15, 16], SANGRIX and its former name Bit Origin [18, 19], GE Vernova [20, 34] and Vertiv [21].
- Company releases and documents: NVIDIA [4, 22], Micron [12], SK hynix [13, 33], Samsung Electronics [14] and CoreWeave's corrected call transcript [17].
- Government: the Department of Energy [23] and the Federal Register [24].
- TSMC: transcripts of its July and October 2026 earnings calls published by The Motley Fool [25, 26]. These are third-party transcriptions. TSMC's own investor pages could not be opened during this research.
- Research firms and brokers: CBRE [27, 28, 29], Wood Mackenzie [30] and TrendForce News [31].
- One seller's product page [32], used only as a named claim.

Period. From CoreWeave's call of 10 November 2025 to Super Micro's results of 3 November 2026. Sources were opened on 24 October 2026. On 7 November 2026 the figures dated October 2026 were checked again, the seller's page [32] was read again, and the newer results of Super Micro [7], SK hynix [33] and GE Vernova [34] were added.

How. Each statement was assigned to one stage: chip and system makers, advanced packaging, memory, other package inputs, system builders and powered buildings. Figures are given as each source states them, with dates and units. Capacity plans are reported as targets, because the companies present them as forward-looking. Where a figure is this report's own arithmetic, the text says so. Estimates from research firms and news reports are labelled as estimates. The report makes no forecast of lead times and uses no Rillor data.

Findings

1. What chip and system makers say about lead times

The makers state lead times for supply and for some products. None in these sources states a lead time for a named GPU server model or rack.

Company	Statement	Scope	Period
NVIDIA	Has experienced "extended lead times of more than 12 months"	"certain supply"	Fiscal year to 25 January 2026 [1]
NVIDIA	"currently experiencing certain supply constraints"	Blackwell and Rubin systems	Quarter to 26 July 2026 [2]
Super Micro	Lead times have "extended beyond twelve months"	"some products"	Fiscal year to 30 June 2026 [5]
Micron	AI demand could "extend lead times, or limit availability"	Materials, equipment and packaging capacity	Fiscal year to 3 September 2026 [11]

NVIDIA describes how it meets long lead times: "To secure future supply and capacity, we have paid premiums, provided deposits, and entered into long-term supply agreements and capacity commitments, which have increased our product costs" [1]. Super Micro describes the same steps: non-cancellable component orders placed ahead of historical lead times, premiums and deposits [5].

The commitments show how far ahead NVIDIA reserves supply. They were $95.2 billion at 25 January 2026, "of which substantially all will be paid through fiscal year 2027" [1]. Six months later they were $279 billion, up from $119 billion a quarter earlier [2]. The 10-Q schedules $92 billion for the rest of fiscal 2027, $87 billion for fiscal 2028, $88 billion for fiscal 2029 and $12 billion for later years [2]. NVIDIA says the commitments are "primarily memory and manufacturing facilities", and that certain agreements "may be cancelable, rescheduled, or adjustable for our business needs prior to placing firm orders" [2].

Inventory moved the same way. NVIDIA's raw materials rose from $3,807 million at 25 January 2026 to $11,341 million at 26 July 2026, and total inventories from $21,403 million to $31,575 million [2]. The finance chief tied the second-quarter rise, from $25.8 billion to $31.6 billion, to preparation "for the introduction of Vera Rubin in the third quarter" [3].

Money moves earlier on the customer side too. NVIDIA's customer advances were $160 million at 25 January 2026 and $2.8 billion at 26 July 2026 [2]. For some investment-grade customers it offers "longer payment terms ranging from 90 days up to one year to assist customers with large data center builds" [2].

By this report's sum, $175 billion of NVIDIA's supply commitments fall in its fiscal years 2028 and 2029, which run from early 2027 to early 2029 [2]. A buyer ordering systems for delivery in 2027 draws on capacity that the chip maker has already contracted.

For HGX B200 and B300 servers, a seller's page gives week counts. Viperatech's site says "HGX B200/B300 lead times are 8-14 weeks for Golden SKU" and "Custom BOMs exceed 20 weeks" [32]. That is one seller's claim, viewed 7 November 2026. It does not show where the time goes, or whether a building is ready to receive the system.

2. Advanced packaging: CoWoS capacity in 2026

Advanced packaging joins the GPU die and its HBM stacks on an interposer. TSMC's process for this is CoWoS, which chief executive C.C. Wei called in October 2026 "still the number one choice for our customer" [26]. Wei spoke about packaging capacity in both 2026 calls reviewed.

On 16 July 2026, asked about competition from Intel's EMIB-T, Wei said: "Well, let me say that our packaging capacity is so tight that now it's limiting my customers' growth. We welcome that additional flexibility in the market" [25]. Later in the call he said: "since our capacity in the back end is so in shortage mode, the gap is bigger" [25].

On 15 October 2026, asked about a GlobalFoundries interposer partnership, Wei said: "We designed a deal with GlobalFoundries, and this is part of our overall strategy in the back-end supply" and "Definitely, we have a little bit shortage on the back end" [26].

TSMC gave no CoWoS capacity figure, no packaging lead time and no size for the packaging shortfall in either call [25, 26]. It declined to size the front-end gap as well. Asked in July whether demand for 3 nm and smaller processes exceeded supply by 30% to 50%, Wei said "we don't have a number to share" [25].

Published capacity figures come from estimates. TrendForce News, which says it "operates independently from our research team", reported on 15 June 2026 what Economic Daily News attributed to institutional investors [31]:

Estimate (second-hand, June 2026)	Figure
CoWoS supply-demand gap	From around 20% to about 10% by the end of 2026
TSMC monthly CoWoS capacity in 2026	120,000 to 140,000 wafers
Added capacity from outsourced packaging partners	50,000 to 60,000 wafers a month
Total industry capacity	Approaching 200,000 wafers a month

Source: [31]. TSMC has not confirmed these estimates in the sources reviewed.

The same page gives TrendForce's own view: "TrendForce expects the severe shortage of global 2.5D packaging capacity to begin moderating by 2027" [31]. Neither call puts a date on balance.

Capacity is being added on several fronts. TSMC reiterated a 2026 capital budget of $60 billion to $64 billion, and Wei said TSMC is "building 13 leading-edge and advanced packaging fab in Taiwan over the next several years, and another five to six fabs overseas" [26]. Its finance chief added that "the tool supply is tight" [26]. Outside TSMC, NVIDIA and Amkor target "a greater than 300% increase in interposer wafer capacity by the end of 2028" in South Korea [4].

3. HBM supply in 2027

HBM takes more factory space than ordinary memory. Micron's annual report says: "HBM requires a higher number of wafers and more cleanroom space to produce the same number of bits as conventional DRAM in the same technology node" [11].

Three memory makers describe demand above supply.

- **Micron, 30 September 2026.** "We have completed agreements for the vast majority of our calendar 2027 HBM bit supply with significant price increases year over year" [12]. Micron does not give the share. For memory and storage overall, it expects conditions "to be much tighter in calendar 2027 and 2028 than they were in 2026", with the DRAM industry "supply constrained in both years" [12]. It adds: "we do not have line of sight to when supply and demand will return to balance" [12].
- **SK hynix, 28 October 2026.** The company "has currently signed LTAs with 14 major customers", against around 10 in July, and the volume secured "corresponds to more than 50% of total production capacity for 2027 and 2028" [13, 33]. Citing "the long lead time required to build new production facilities", it expects demand to exceed supply through 2030 [33].
- **Samsung Electronics, 30 July 2026.** Growth in demand for server DRAM, enterprise SSDs and HBM "is projected to keep the market undersupplied" in the second half of 2026, and "Despite efforts to increase production, supply constraints are expected to continue" [14].

Contracts now run for years. Micron has signed 26 strategic customer agreements with take-or-pay volumes, which it estimates at over 35% of its revenue through 2030, backed by customer financial commitments of $32 billion, "the vast majority of which are cash deposits" [12].

New memory capacity arrives on dated schedules. Micron expects initial output from its HBM advanced packaging facility in Singapore in early calendar 2027, and wafer output from its ID1 fab in mid-calendar 2027 and from ID2 in late calendar 2028 [12]. Output from new fabs "takes time to ramp and gradually becomes more meaningful starting a few quarters after initial output" [12]. SK hynix is accelerating its M15X schedule and expanding capacity after the Yongin Phase 1 cleanroom opens in early 2027 [13].

Memory reaches GPU systems twice: as HBM inside the GPU package and as system memory and storage in the server. System builders report the second as a delivery constraint. Super Micro says constraints "for certain components, including memory and storage, as well as GPU and CPU availability" affected "the timing of certain of our product deliveries" in fiscal 2026 [5]. Dell cites "current limitations in capacity from memory manufacturers" [9]. HPE says memory constraints "may persist until memory vendors transition greater production allocations towards high performance memory components required by AI workloads" [10].

4. Substrates, interposers, testing and optics

On 21 October 2026 NVIDIA announced plans with six suppliers to expand capacity across substrates, advanced packaging, testing and optical engines [4]. Each comes with a date. All are forward-looking targets.

Stage	Supplier and location	Target	Date
Substrates	Samsung Electro-Mechanics, South Korea	Immediate capacity from existing operations, plus upgrades and new factories in Sejong and Busan	Now. By 2028, production allocated to NVIDIA is projected to be a significant portion of its package substrate capacity
Testing	Singapore operations	High-volume production	February 2027
Testing	KYEC, United States, with a $1.1 billion capacity reservation advance from NVIDIA	High-volume production	April 2027
Substrates	Unimicron, Taiwan, under a capacity reservation agreement	Begins adding capacity, then full capacity	From the second quarter of 2027 to mid-2028
Optical engines	SPIL, Taiwan	15x capacity, then another doubling	Mid-2027, then the end of 2028
Substrates	IBIDEN Ono Plant, Japan, with a multibillion-dollar capacity reservation advance	Capacity reaching 3x current levels at the plant	Additional capacity ramping from the first quarter of 2028
Interposers	Amkor, South Korea	More than 300% increase in interposer wafer capacity	End of 2028
Substrates, total	All substrate agreements	2.5x NVIDIA's substrate capacity	Over three years

Stage	Supplier and location	Target	Date
Testing, total	United States, Taiwan and Singapore	Total testing capacity doubles	From calendar 2026 to 2027

Source: [4].

Two points follow. NVIDIA pays suppliers in advance to add capacity: the release names capacity reservation advances or agreements for IBIDEN, Unimicron and KYEC [4]. Most of the dated additions start in 2027, and several reach their targets in 2028, so systems delivered in late 2026 draw on capacity already in place.

TSMC's back end connects to the same list. Its October call described the GlobalFoundries deal as part of its "overall strategy in the back-end supply" [26]. Its co-packaged optics platform, COUPE, was in production by July 2026: "We start the production right now, and it will be ramped up" [25].

The release states no lead time for any component. Apart from Samsung Electro-Mechanics, it does not say how much of each supplier's capacity is reserved for NVIDIA [4].

5. System builders, backlogs and component supply

System builders turn GPU baseboards into servers and racks. Their filings show large order books, rising component inventories and delivery dates that depend on the customer.

Company	Period	Figure	Source
Dell	Quarter to 31 July 2026	AI server orders $60.9 billion, revenue $16.4 billion, backlog $95 billion	[8]
Dell	31 July 2026	Purchase obligations $34.8 billion, of which $24.6 billion payable within twelve months	[9]
Dell	30 January to 31 July 2026	Production materials inventory from $6,696 million to $15,894 million	[9]
Super Micro	Quarter to 30 June 2026	New orders in excess of $60 billion, "expected to be delivered over future quarters"	[6]
Super Micro	Quarter to 30 September 2026	Net sales $19.9 billion, against $11.1 billion in the quarter to 30 June 2026	[7]
HPE	31 July 2026	Customer deposits $405 million	[10]

Dell's backlog equals about 5.8 times its AI server revenue for the quarter. That ratio is this report's arithmetic. It is not a lead time, because Dell does not say how the backlog is scheduled [8]. Across the whole company, Dell expects to recognize about 77% of its $132 billion of remaining performance obligations within twelve months, a figure that covers all products [9].

The builders name three things that set delivery timing.

Components. Super Micro's fiscal 2026 deliveries were affected by memory, storage, GPU and CPU supply [5]. It warns that "we have been forced to delay the introduction of certain of our products or the fulfillment of customer orders as a result of shortages of materials and key components" [5]. One unnamed supplier accounted for 63.1% of its purchases in fiscal 2026 [5].

Customer readiness. Dell writes that "the varying stages of customer readiness, and the frequency of component part updates or transitions" create "inherent non-linearity in the timing of demand and subsequent shipments" for its AI servers [9].

Acceptance. HPE reports higher inventory due in part to "customers transitioning to the next generation of GPUs, our efforts to secure supply ahead of demand, and longer customer acceptance timelines on AI-related orders" [10]. Its revenue follows acceptance: "As customer acceptance milestones are met, the Company will recognize revenue and reduce the amount of contract liabilities" [10].

6. Powered buildings

A delivered rack earns nothing until it runs in a powered, cooled room. NVIDIA's finance chief wrote in August 2026: "Securing land, power and shell for data centers has become the next critical phase in the AI infrastructure buildout" [3]. The 10-Q calls expanding land, power, shell and energy "a complex, multi-year process that involves significant regulatory, technical, and construction challenges" and warns that "Customers may postpone purchasing new architectures due to the lack of availability of data center infrastructure to deploy our products" [2].

TSMC checks the same stage before it ships. In October 2026 Wei said TSMC is "sending our own teams of people to check on the data center build-out progress and power supply situation in the U.S." and that "The feedback is so far so good" [26]. In July he gave the purpose: "We checking all that to make sure that TSMC's chips will not be put in inventory" [25].

Space. CBRE tracks eight primary North American markets [28, 29].

Measure	Figure	As of
Vacancy in primary markets	1.4%	Year-end 2025, unchanged in the first half of 2026

Measure	Figure	As of
Capacity under construction	5,994.4 MW	Year-end 2025
Capacity under construction	7,481 MW, after a 24.8% rise in the half	First half of 2026
Share of capacity under construction that is preleased	80.4%, against 74.3% a year earlier	First half of 2026
Future capacity still available	Less than 1,500 MW, about six months of demand	First half of 2026

Sources: [27, 28, 29].

CBRE's Gordon Dolven said in August 2026 that "many occupiers are effectively competing for future inventory rather than existing availability" [29].

Power. CBRE wrote in February 2026 that "Sites offering power access within 18 to 36 months for greenfield development are highly sought after" and that "Grid-power capacity for existing projects is largely booked through 2030 in most markets" [27]. In August it added that "developers are facing energization delays as new last-mile projects face persistent hurdles in permitting, planning and zoning" [28].

The grid's own equipment has long lead times. The Department of Energy wrote on 10 August 2026 that dependence on imports, limited domestic production and excessive customization are "factors which have resulted in lead times of two or more years for critical equipment", naming transformers, circuit breakers, substation components and power electronics [23]. Wood Mackenzie estimates that power transformers and substations are in market shortages of 15% and 8% in 2026 [30]. Its analyst Ben Boucher said data centers have been "using Chinese units to minimize lead times" [30]. Executive Order 14421 of 26 August 2026 declared a national emergency to secure the bulk-power system and calls for implementing rules within 120 days [24]. Wood Mackenzie, whose release numbers the order 14420, expects it to worsen the transformer shortage [30].

On-site power. CBRE reports that "BYOP and on-site generation (including batteries and gas turbines) are becoming routine in large-scale data center planning" [27]. Turbines are themselves sold by reservation. GE Vernova's gas power equipment backlog and slot reservation agreements grew from 100 GW to 116 GW in the second quarter of 2026 [20], and from 116 GW to 122 GW in the third, and it expects to end 2026 with at least 130 GW of gas equipment under contract [34]. In July its chief executive said it was on track to reach 24 GW of annual gas turbine output in 2028 [20]. GE Vernova does not state a delivery lead time.

Fit-out and cooling. CBRE reports that "RFS fit-out dates are expanding from weeks to several months" as tenants ask for higher power density, and that "For air-cooled data centers requiring conversion to liquid cooling, delays can extend beyond six months" [28]. Vertiv, a maker of power and cooling equipment, said its second-quarter 2026 revenue "reflected minor timing shifts, primarily due to temporary supply chain congestion and multi-phased project execution" [21].

CoreWeave had 1.5 GW of active power and about 3.7 GW of contracted power at 30 June 2026 [16]. The difference is power under contract that was not yet active, and the release gives no schedule for it.

7. Two disclosed delays

Few delays are disclosed with dates. Two are, one in a building and one in hardware.

A powered shell, CoreWeave, November 2025. On 10 November 2025 CoreWeave's chief executive Michael Intrator said the company was "affected by temporary delays related to a third-party data center developer, who is behind schedule" [17]. Its finance chief said "the delays in powered shell delivery associated with the data center provider will have an impact on our fourth quarter results" and guided 2025 revenue to $5.05 billion to $5.15 billion [17]. He put 2025 capital spending at $12 billion to $14 billion, a "reduction in CapEx from our prior guidance" [17]. Third-quarter capital spending of $1.9 billion was "lower than anticipated due to the delays", and construction in progress, which the company described as "infrastructure not yet in service", rose by $2.8 billion to $6.9 billion [17]. The customer "has agreed to adjust the delivery schedule and extend the expiration date", so CoreWeave kept the total value of the original contract [17].

Intrator placed the problem in the building: "really where the challenge is, is the powered shell" [17]. He declined to name the provider and called it "a systemic problem that the industry is going to have to deal with for the foreseeable future" [17]. He expected "the overwhelming majority of the delay" to be resolved within the first quarter of 2026 [17]. CoreWeave's annual report later referred to "the delays that we announced in November 2025 with respect to the delivery of certain data centers to be provided by a third-party data center provider" [15].

A hardware delivery, SANGRIX, 2026. On 29 June 2026 Bit Origin, now SANGRIX, announced the purchase of 16 NVIDIA B300 servers for about $11 million [19]. The seller had already bought the servers, and they were "currently expected to be delivered during the third quarter of 2026" [19]. On 16 October 2026 SANGRIX reported that the seller, PT Mitra Manunggal Sangkara, had advised that, "based on its latest update from the manufacturer", the servers were "scheduled to be delivered during the week of November 3, 2026" [18]. Inspection, racking and testing were expected to take about one more week, with production deployment after commissioning and acceptance [18]. The release names neither the manufacturer nor a reason for the change [18].

	CoreWeave	SANGRIX
Stage that slipped	Powered shell from a third-party developer	Server delivery from an unnamed manufacturer
Timing	Delays affected fourth-quarter 2025 expectations, with most of the delay expected to clear within the first quarter of 2026	Third quarter of 2026, moved to the week of 3 November 2026, plus about one week for inspection, racking and testing
Effect disclosed	Lower 2025 capex guidance, construction in progress up $2.8 billion	Start of operations follows the later delivery
Customer contract	Delivery schedule adjusted and expiration extended	Operations subject to installation, testing and contract requirements

Sources: [15, 17, 18, 19].

By this report's arithmetic, the revised SANGRIX date falls more than four weeks after the third quarter ended on 30 September 2026. SANGRIX is a small issuer whose schedule reaches the public through its seller, relaying the manufacturer [18]. It is one company's disclosure.

Implications

For buyers planning a cluster. The delivery month of a GPU system depends on whichever stage finishes last. Chip makers reserve components years ahead [2], system builders schedule shipments around customer readiness [9], and most powered space under construction is already preleased [28]. A plan that fixes the hardware date and leaves the building to follow carries the risk CoreWeave disclosed: hardware bought before its room is ready [17]. Installation and acceptance add time after delivery: SANGRIX expects about one week of inspection, racking and testing, with production deployment after commissioning and acceptance [18], and HPE cites "longer customer acceptance timelines on AI-related orders" [10]. A contract can name each point: the delivery month, the site, the acceptance tests and what happens if either side is late.

For GPU cloud operators. CoreWeave's contract survived a late building because the customer agreed to move the delivery schedule and extend the term [17]. Operators with multi-year, fixed-capacity leases [15] and dated customer commitments carry the gap between the two.

For lenders and finance teams. Hardware that waits for a building sits in construction in progress, outside service [17]. Revenue at system builders follows acceptance milestones [10]. Draw schedules keyed to delivery dates meet the same slippage.

For researchers. A component lead time, a system lead time and a site-readiness date measure different things. NVIDIA's "more than 12 months" applies to certain supply [1], and a seller's week count applies to a standard build on a given day [32].

Rillor offers forward purchase agreements that fix the grade, quantity, price and delivery month of GPU systems at signing, and capacity agreements that reserve power and space in fixed blocks before the hardware arrives.

Limits of this analysis

- No chip maker or system builder in the sources reviewed states a lead time for a named GPU server model or an NVL72 rack. The only week counts are a seller's claims.
- TSMC's statements come from third-party transcripts. TSMC gives no CoWoS capacity figure, packaging lead time or shortfall size. The capacity figures here are second-hand estimates reported by TrendForce News.
- Micron's "vast majority" of 2027 HBM supply is not quantified. SK hynix's share of more than 50% covers its total production capacity for 2027 and 2028, not HBM alone, and Samsung gives no 2027 contract share. No source states how much HBM is available outside long-term agreements.
- NVIDIA's capacity targets of 21 October 2026 are forward-looking. NVIDIA does not split its $279 billion of supply commitments by supplier.
- The sources center on NVIDIA systems. AMD and other accelerator makers are not covered.
- Data on buildings and power come mainly from the United States and CBRE's eight primary North American markets.
- The two delays are single cases. CoreWeave does not name the data-center developer, and SANGRIX does not name the manufacturer.
- NVIDIA's results for its quarter ending in October 2026 fall after this report's 3 November 2026 cut-off.
- The report forecasts no lead times. It is research, not investment advice.

Sources

1. NVIDIA Corporation Annual Report on Form 10-K for the fiscal year ended 25 January 2026. U.S. Securities and Exchange Commission, filed 25 February 2026. https://www.sec.gov/Archives/edgar/data/1045810/000104581026000021/nvda-20260125.htm

2. NVIDIA Corporation Quarterly Report on Form 10-Q for the quarter ended 26 July 2026. U.S. Securities and Exchange Commission, filed 26 August 2026. https://www.sec.gov/Archives/edgar/data/0001045810/000104581026000075/nvda-20260726.htm

3. CFO Commentary on Second Quarter Fiscal 2027 Results (Exhibit 99.2 to Form 8-K). NVIDIA, 26 August 2026. https://www.sec.gov/Archives/edgar/data/1045810/000104581026000073/q2fy27cfocommentary.htm

4. NVIDIA and Supply Chain Partners Expand Capacity to Accelerate AI Infrastructure Buildout. NVIDIA, 21 October 2026. https://nvidianews.nvidia.com/news/nvidia-and-supply-chain-partners-expand-capacity-to-accelerate-ai-infrastructure-buildout

5. Super Micro Computer, Inc. Annual Report on Form 10-K for the fiscal year ended 30 June 2026. U.S. Securities and Exchange Commission, filed 31 August 2026. https://www.sec.gov/Archives/edgar/data/0001375365/000137536526000022/smci-20260630.htm

6. Supermicro Provides Fourth Quarter of Fiscal Year 2026 Preliminary Business Update (Exhibit 99.1 to Form 8-K). Super Micro Computer, 21 July 2026. https://www.sec.gov/Archives/edgar/data/1375365/000137536526000019/fq42026businessupdate.htm

7. Supermicro Announces First Quarter of Fiscal Year 2027 Financial Results (Exhibit 99.1 to Form 8-K). Super Micro Computer, 3 November 2026, viewed 7 November 2026. https://www.sec.gov/Archives/edgar/data/1375365/000137536526000026/exhibit991_20260930.htm

8. Dell Technologies Delivers Second Quarter Fiscal 2027 Financial Results (Exhibit 99.1 to Form 8-K). Dell Technologies, 1 September 2026. https://www.sec.gov/Archives/edgar/data/0001571996/000157199626000039/exhibit991earnings8kq2fy27.htm

9. Dell Technologies Inc. Quarterly Report on Form 10-Q for the quarter ended 31 July 2026. U.S. Securities and Exchange Commission, filed 8 September 2026. https://www.sec.gov/Archives/edgar/data/1571996/000157199626000046/dell-20260731.htm

10. Hewlett Packard Enterprise Company Quarterly Report on Form 10-Q for the quarter ended 31 July 2026. U.S. Securities and Exchange Commission, filed 3 September 2026. https://www.sec.gov/Archives/edgar/data/1645590/000164559026000080/hpe-20260731.htm

11. Micron Technology, Inc. Annual Report on Form 10-K for the fiscal year ended 3 September 2026. U.S. Securities and Exchange Commission, filed 9 October 2026. https://www.sec.gov/Archives/edgar/data/723125/000072312526000023/mu-20260903.htm

12. Fiscal Q4 2026 Earnings Call Prepared Remarks. Micron Technology, 30 September 2026.

https://s25.q4cdn.com/621799436/files/doc_financials/2026/q4/Q4-FY26-Prepared-Remarks.pdf

13. SK hynix Announces 2Q26 Financial Results. SK hynix, 29 July 2026. https://news.skhynix.com/en/q2-2026-business-results/

14. Samsung Electronics Announces Second Quarter 2026 Results. Samsung Electronics, 30 July 2026. https://news.samsung.com/global/samsung-electronics-announces-second-quarter-2026-results

15. CoreWeave, Inc. Annual Report on Form 10-K for the year ended 31 December 2025. U.S. Securities and Exchange Commission, filed 2 March 2026. https://www.sec.gov/Archives/edgar/data/1769628/000176962826000104/crwv-20251231.htm

16. CoreWeave Reports Strong Second Quarter 2026 Results (Exhibit 99.1 to Form 8-K). CoreWeave, 11 August 2026. https://www.sec.gov/Archives/edgar/data/1769628/000176962826000362/coreweave2q26earningspress.htm

17. CoreWeave, Inc. Q3 2025 Earnings Call, corrected transcript. CoreWeave, 10 November 2025. https://s205.q4cdn.com/133937190/files/doc_financials/2025/q3/CORRECTED-TRANSCRIPT_-CoreWeave-Inc-CRWV-US-Q3-2025-Earnings-Call-10-November-2025-5_00-PM-ET.pdf

18. SANGRIX Announces November 2026 Delivery Schedule for 128-GPU NVIDIA Blackwell B300 Project in Malaysia (Exhibit 99.1 to Form 6-K). SANGRIX Inc., 16 October 2026. https://www.sec.gov/Archives/edgar/data/1735556/000110465926116630/tm2627763d1_ex99-1.htm

19. Bit Origin Acquires NVIDIA Blackwell B300 AI Infrastructure with Contracted Customer Deployment in Malaysia (Exhibit 99.1 to Form 6-K). Bit Origin Ltd, 29 June 2026. https://www.sec.gov/Archives/edgar/data/1735556/000110465926078558/tm2619272d2_ex99-1.htm

20. GE Vernova reports second quarter 2026 financial results and raises 2026 financial guidance (Exhibit 99 to Form 8-K). GE Vernova, 22 July 2026. https://www.sec.gov/Archives/edgar/data/1996810/000199681026000147/gevpressrelease2q26.htm

21. Vertiv Reports Strong Second Quarter 2026 with Diluted EPS Growth of 53% (Exhibit 99.1 to Form 8-K). Vertiv Holdings Co, 29 July 2026. https://www.sec.gov/Archives/edgar/data/1674101/000162828026050323/q22026exhibit991vrt07292026.htm

22. Hardware, NVIDIA DGX GB Rack Scale Systems User Guide. NVIDIA, last updated 17 September 2026. https://docs.nvidia.com/dgx/dgxgb200-user-guide/hardware.html

23. Strengthening America's Grid Supply Chain. U.S. Department of Energy, Office of Electricity, 10 August 2026. https://www.energy.gov/oe/articles/strengthening-americas-grid-supply-chain

24. Executive Order 14421 of August 26, 2026, Declaring a National Emergency To Secure the United States Bulk-Power System, 91 FR 55995. Federal Register, 31 August 2026. https://www.federalregister.gov/documents/2026/08/31/2026-17843/declaring-a-national-emergency-to-secure-the-united-states-bulk-power-system

25. TSM (TSM) Q2 2026 Earnings Call Transcript. The Motley Fool, 16 July 2026. https://www.fool.com/earnings/call-transcripts/2026/07/16/tsm-tsm-q2-2026-earnings-call-transcript/

26. TSMC (TSM) Q3 2026 Earnings Call Transcript. The Motley Fool, 15 October 2026. https://www.fool.com/earnings/call-transcripts/2026/10/15/tsmc-tsm-q3-2026-earnings-call-transcript/

27. North America Data Center Trends H2 2025. CBRE, 25 February 2026. https://www.cbre.com/insights/books/north-america-data-center-trends-h2-2025

28. North America Data Center Trends H1 2026. CBRE, 27 August 2026. https://www.cbre.com/insights/books/north-america-data-center-trends-h1-2026

29. North American Data Center Demand Continues to Outpace Supply Despite Record Construction Activity. CBRE, 27 August 2026. https://www.cbre.com/press-releases/north-american-data-center-demand-continues-to-outpace-supply-despite-record-construction-activity

30. US executive order banning Chinese bulk-power equipment set to worsen critical transformer shortage. Wood Mackenzie, 31 August 2026. https://www.woodmac.com/press-releases/us-executive-order-banning-chinese-bulk-power-equipment-set-to-worsen-critical-transformer-shortage/

31. [News] TSMC CoWoS Supply-Demand Gap Reportedly Seen Narrowing from 20% to 10% by End-2026. TrendForce News, 15 June 2026. https://www.trendforce.com/news/2026/06/15/news-tsmc-cowos-supply-demand-gap-reportedly-seen-narrowing-from-20-to-10-by-end-2026-as-capacity-expands/

32. Supermicro GPU SuperServer SYS-822GS-NB3RT HGX B300 8-GPU. Viperatech, viewed 7 November 2026. https://viperatech.com/product/supermicro-gpu-superserver-sys-822gs-nb3rt-hgx-b300-8-gpu

33. SK hynix Announces 3Q26 Financial Results. SK hynix, 28 October 2026, viewed 7 November 2026. https://news.skhynix.com/en/q3-2026-business-results/

34. GE Vernova reports strong third quarter 2026 financial performance, raises 2026 guidance, and increases share repurchase authorization (Exhibit 99 to Form 8-K). GE Vernova, 28 October 2026, viewed 7 November 2026. https://www.sec.gov/Archives/edgar/data/1996810/000199681026000162/gevpressrelease3q26.htm

About this report

Published by Rillor on 7 November 2026. Superintelligence (SI) tools assist research and drafting, and every figure cites its source. Rillor builds and secures SI systems, supplies the data they learn from and delivers the infrastructure they run on. We apply SI to financial and compute markets.

Cite as: Rillor. What sets GPU system lead times in 2026. 7 November 2026. https://rillor.com/insights/gpu-system-lead-times-2026

Notices

Rillor is not a registered investment adviser or commodity trading advisor and does not provide investment or trading advice. Research, forecasts, data and software described on this site are for research and engineering use. Nothing here is an offer or recommendation to buy or sell any security, commodity interest or digital asset. Past or simulated results do not indicate future results.

NVIDIA, AMD, Ceph, Kubernetes and related product names are trademarks of their owners and are used only to identify products. No affiliation or endorsement is implied.

Work with Rillor

Send the systems or capacity you need, the quantity, and the delivery month and location. Contact us at rillor.com/contact?type=forward or contact@rillor.com. More on forward agreements: rillor.com/markets/forward-agreements