트레이드 렌즈
TRADE LENS

Focus on investment.

LENS INSIGHT / No.14

NVIDIA | Korea AI infrastructure | Supply and capital recovery

NVIDIA’s CEO Transition

Will Korea’s 260,000-GPU Plan Hold?

From supply commitments to utilization and capital recovery

Supply
Commitments and execution

Service
Hardware into usable compute

Returns
Timing and recurring value

Commitments continue.
Execution and capital recovery decide the outcome.

Separate aggregate GPU plans into contracts, facilities, customers and cash generation.

Analysis date 2026.11.14

TRADELENS.KR

Industry structure, project execution and capital efficiency

01 / RESEARCH VIEW

Supply continuity is not an investment return

TradeLens view | Project maturity matters more than the leadership headline.

Central judgment

Company statements and Korea's reaffirmation do not support treating the CEO change alone as a wholesale withdrawal of GPU supply. The 260,000-GPU figure is a multi-year public/private plan, not delivered or operating capacity. Contracts, system readiness, recurring demand and capital recovery must be assessed together. Our base case is differentiated execution and returns, rather than a blanket halt. [1] [2] [3]

Supply

Separate company commitments from order and delivery terms.

Service

Track when delivered equipment becomes usable capacity.

Returns

Evaluate value after costs, not activity alone.

Korea's buildout is not one project executed personally by one leader. Public computing, internal manufacturing applications and commercial cloud services have different buyers, budgets and routes to economic value. Contract maturity and site readiness are more useful than a national equipment total. [4] [5]

Why the change matters

Expectations centered on personal relationships may shift toward corporate execution controls. Maintaining existing customer deliveries and reviewing uncommitted expansion can coexist. This is a conditional interpretation, not evidence that the new leadership has decided to cut investment.

Evidence that would change the view

Cancellations, reduced allocations or repeated delays to firm schedules would weaken continuity. Acceptance, service launch and paid repeat orders would translate commitments into economic evidence.

02 / EVENT & EVIDENCE

A resignation is not a Korea project withdrawal

Separate the filing's event date, the company announcement and Korea's response.

Date	Reported development	Interpretation
11 November	Board accepted Jensen Huang's resignation and appointed Colette Kress interim CEO. [1]	The filing's event date; not completion of the permanent succession process.
12 November	Company announced leadership changes and continuity of commitments and roadmap. [2]	A corporate position, not proof of completed contract performance.
13 November	Korea's minister said GPU supply and cooperation had been reaffirmed with NVIDIA. [3]	News reporting of government remarks, not an installed-capacity ledger.

Three channels through which uncertainty can reach projects

Decisions
Approval speed for follow-on investment

>

Terms
Prepayments, collateral and allocation

>

Demand
Customers' expansion and service purchases

Leadership changes can increase uncertainty over decision channels and capital allocation. They do not instantly remove engineers, operations or the manufacturing supply chain. Order revisions or delivery disruption are needed to connect governance concern to a supply interruption.

> **Scope**
> Personal allegations and stock-price declines are not used as evidence of supply disruption. The subject is Korea's contracts, schedules and capital recovery, not share-price targets or trading decisions.

03 / THE 260,000-GPU BASELINE

260,000 GPUs are not one purchase order

Reconstruct the multi-year plan announced in 2025, not a tally of currently operating assets.

Baseline of announced 2025 plans [4]

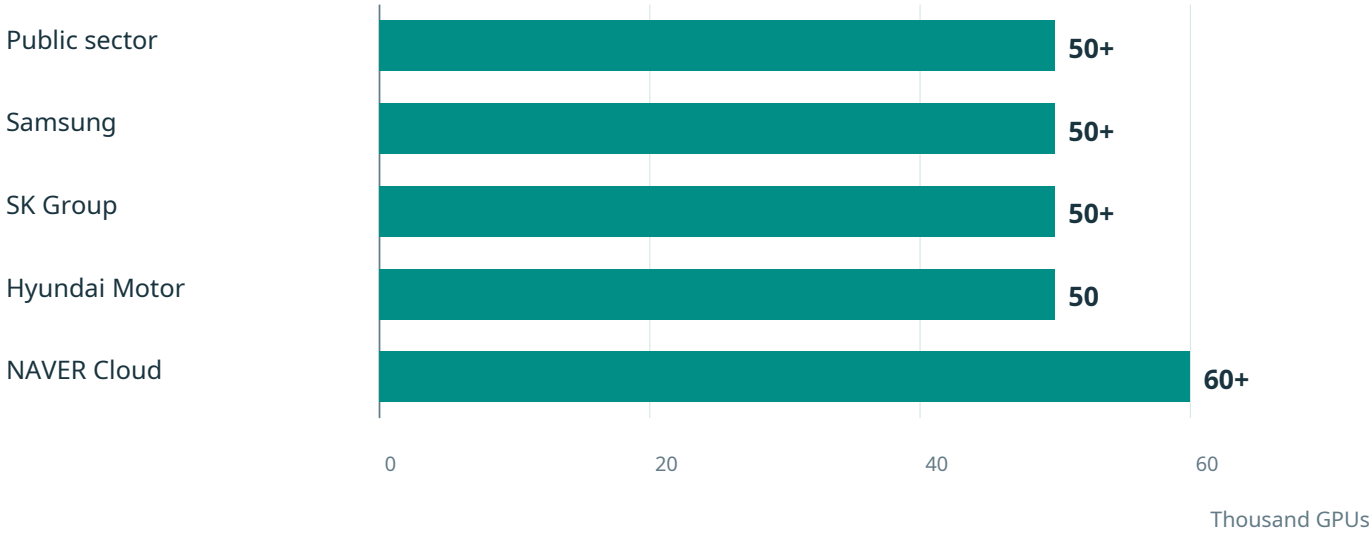

The “over” qualifiers in company announcements are retained. The headline total is a plan, not firm orders or homogeneous delivered units. The same provider can appear in public procurement and its own private buildout; company names alone are not a basis for adding allocations again. [4] [5]

Segment	Economic purpose	Outcome measure
Public	Research and model access	Queues, completions, allocation
Manufacturing	Design, process and product gains	Quality, cycle time and realized savings
Commercial cloud	Compute for external customers	Paid utilization, pricing, renewals and cash

Evaluate the contract unit

Converting corporate plans, public procurement and research systems into concurrent power demand or HBM revenue mixes scopes. Identify the buyer, hardware, delivery date and asset owner first.

04 / PUBLIC DEPLOYMENT

Public procurement is not the same as operation

Yonhap's 13 November account of government remarks; quantities are on a B200 basis. [3]

Stages within the public-sector plan



21k described as secured

14k expected early 2027

15k via national center

Thousand / B200 basis

Route	Reported stage	Interpretation limit
35,000 direct	About 21,000 secured; remaining 14,000 expected to arrive in early 2027.	"Secured" does not establish that every system is installed and usable.
15,000 center	Procurement through the National AI Computing Center is being pursued.	Do not apply the same arrival date or structure as direct procurement.

60%

21,000 / 35,000: reported secured amount relative to the direct target.

Not provided

These remarks do not establish delivery or service-readiness ratios for all 260,000 GPUs.

The remarks support continuity for near-term procurement. The next evidence is arrival, acceptance and service launch by cluster. Deployment across several operators may not provide the same capacity for one large training job; network integration and operating arrangements matter.

Value of public resources

Commercial revenue alone cannot measure a research-access objective. Yet equipment ownership without waiting-time or project-completion measures does not demonstrate policy effectiveness.

05 / CONTRACT MATURITY

Some projects have moved beyond an announcement

NAVER illustrates why a July conditional proposal is not the whole November picture.

Date	Publicly reported progress	Remaining question
24 Jul 2026	NVIDIA equity plans and Brookfield financing preconditions announced. [10]	Distinguish the conditional proposal at that date from completed funding.
27 Oct 2026	NAVER announced a USD9bn infrastructure agreement and satisfaction of preconditions. [11]	A signed contract does not mean all equipment has been delivered or activated.
12 Nov 2026	Company reported KRW2tn of initial customer contracts; H1 2027 revenue planned. [12]	A contract total is not one quarter's revenue or profit.

Capital and assets
Asset-owning SPV / procurement

>

Operations
NAVER Cloud / service delivery

>

Demand
Customers / usage and payment

The October structure places infrastructure in a Brookfield asset-owning vehicle that supplies compute to the operator. External capital separates roles; it does not by itself establish an absence of long-term purchase obligations or consolidated financial exposure. [11]

Expansion from an initial 55MW to 200MW by 2028 is staged. Early contracted capacity and later capacity still needing customers, funding and facilities have different risks. Multi-year infrastructure value is neither current spending nor period service revenue. [10] [11]

> **Read the leadership effect at the right stage**
>
> Do not describe a progressed contract as still only an expression of interest. Evaluate execution, delivery and launch separately from approval of the next expansion. Customer commitments help, but pricing and cost control still matter.

06 / WHAT LEADERSHIP CAN CHANGE

Watch follow-on decisions, not only deliveries

Potential transmission channels identified by TradeLens; not confirmed changes in policy.

Layer	Possible transmission point	Evidence to observe
Existing supply	Continuity in sales, logistics and support	Firm delivery dates, accepted units, notices
Next expansion	Next platform, allocation and co-investment	Executed contracts, orders, capital payments
Risk allocation	Collateral, guarantees and minimum purchases	Changed terms, provisions and commitments
Technical work	Staffing, certification and development priority	Test acceptance, launches and support dates

The stated commitment to existing customers is the starting point, not a promise that every future investment term is immutable. Executing current orders while applying a higher return hurdle to new joint projects is logically possible. [2]

Korean projects need not face the same shock

A funded project with facilities and customers ready depends less on further approvals. A project still requiring equity, grid connection and new customers can transmit one delay into several workstreams. This is a classification by execution stage, not a judgment assigned to a company name.

Unchanged GPU quantities can still carry different economics. Earlier payment or stronger minimum-purchase commitments could raise the buyer's funding exposure without any cancellation. The relevant change would be cash timing and loss allocation, not equipment count.

> **Decision rule**
>
> Do not infer a new strategy solely from a leader's technical or financial background. Observe actual performance of existing commitments and terms for new projects.

07 / FROM DELIVERY TO SERVICE

Early delivery can still mean a late launch

Hardware, facilities, funding and software progress in parallel and can wait for one another.

01 Order
Platform, volume, schedule

>

02 Funding
Prepayment and obligations

>

03 Arrival
Customs and acceptance

04 Facilities
Power, cooling, networking

>

05 Integration
Performance and operations

>

06 Use
Customers, research, payment

A simple bottleneck model

Service start approximately equals max(hardware readiness, facility readiness, financing conditions) + integration and validation time. This is an illustrative schedule model, not a project's actual construction plan.

On-time shipment does not eliminate a late grid connection or network validation. Where facilities are ready, equipment delivery may instead be the constraint. Adding the same delay to every stage double-counts workstreams that run concurrently.

Observed progress	Additional question
More GPUs arrive	Which clusters have completed acceptance and entered service?
Power capacity secured	Are energization, cooling and network testing complete?
More customer reservations	What triggers revenue, and who bears non-delivery liability?

NVIDIA's system documentation includes CPUs, switches, storage, networking and operating software alongside GPUs. Equipment availability and service quality therefore require different measures. [13] [14]

Change the execution metric

Report usable compute time after acceptance alongside procurement volumes. This helps distinguish delivery delays from operating delays.

08 / NOT ALL GPUs ARE EQUAL

Equal counts do not mean equal compute

Do not mix a B200-based public figure with private plans covering several hardware classes.

Item	DGX B200 system	RTX PRO server example
GPU configuration	Eight B200 GPUs	Eight RTX PRO 6000 Blackwell GPUs
Total GPU memory	1,440GB	768GB: 96GB GDDR7 × 8
System role	Accelerated compute with NVLink	Enterprise AI, simulation and graphics
Fit test	Model size, parallelism, communication	Inference, graphics and workload validation

These are official configuration examples, not a confirmed Korean installation or a performance ranking. Sharing the Blackwell name does not imply equal throughput, memory or power per GPU. Memory capacity alone is not application performance. [13] [14]

Training

Time to target quality, model scale and recovery

Inference

Throughput within latency limits; total cost per token

Manufacturing

Validated quality, development time and savings

Useful compute is an output meeting accuracy, latency and security requirements, not peak chip specifications. A benchmark ratio with different models, precision, batch sizes or networks does not establish practical substitutability.

> **Measure capacity in context**
>
> Track equipment counts, but also service levels and cost per useful output by workload. Only then can a supply change be translated into operational capability.

09 / PHYSICAL READINESS

Power, cooling and networking remain necessary

An illustrative 1,000-GPU configuration separates hardware procurement from facility readiness.

8 GPUs
GPU count in one DGX B200 [13]

14.3kW
Maximum system power for one DGX B200 [13]

125 systems
1,000 / 8 in this example

Step	Illustrative result
Server maximum power	125 × 14.3kW = 1.7875MW
Facility-efficiency assumption	PUE of 1.30 gives 1.7875 × 1.30 = 2.32375MW
Further requirements	Additional IT load, including external storage/switches; redundancy and actual load profile

PUE is total facility power divided by IT power. This simplified model counts only servers as IT load: 2.324MW is not a complete data-center design. Maximum power is not average consumption and should not simply be multiplied by 8,760 to infer an electricity bill.

Why this is not scaled to all 260,000 GPUs

Platforms, configurations, deployment dates and facility efficiencies differ. Applying the B200 example nationally would ignore hardware mix and platform changes. The calculation is a starting point for one specified installation, not an estimate of Korea’s total electricity requirement.

> **An operating risk distinct from succession**
> Power connection, cooling and network validation must be completed regardless of the CEO. Even with unchanged supply commitments, delays here postpone research, service launch and capital recovery.

10 / TIME & UTILIZATION

Three months of delay reduce first-year value

Illustration | 1,000 identical GPUs, 8,760 hours/year, 98% availability and unchanged workloads.

First-year utilized GPU-hours

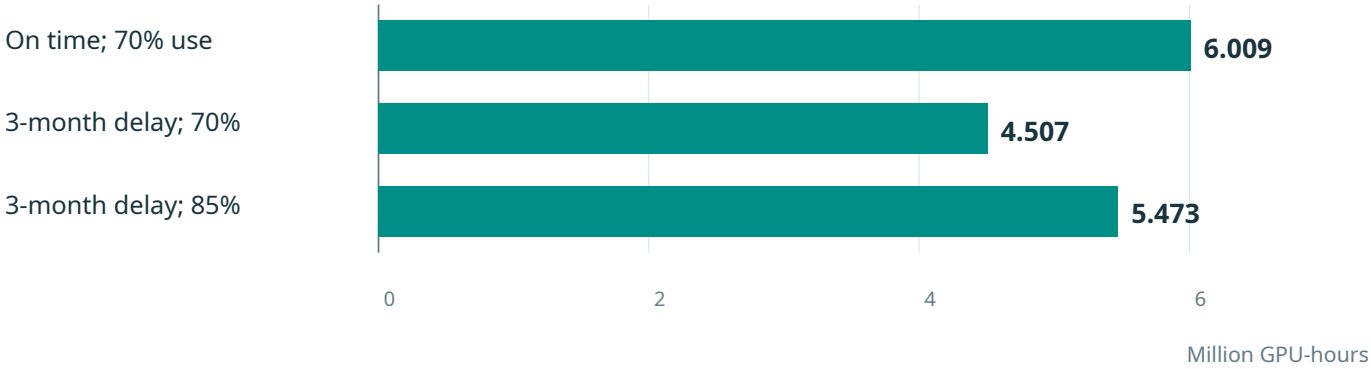

Formula: 1,000 × 8,760 × operating fraction of the year × 98% × utilization. Three months are treated as 25% of a year. Utilization is the share of available hours allocated to work, not rental revenue or model quality.

Start delay	50% utilization	70% utilization	85% utilization
None	4.292	6.009	7.297
3 months	3.219	4.507	5.473
6 months	2.146	3.005	3.649

Unit: million GPU-hours. At unchanged utilization, a three-month delay removes 25% of first-year output. Raising utilization to 85% after the delay still leaves output about 8.9% below an on-time start at 70%. More equipment alone does not resolve a schedule loss.

> **Interpretation limit**
>
> GPU-hours compare the same equipment and workload. They are not a cross-platform performance conversion or an estimate of national effective compute. Parallelization losses, output quality and cash collection need separate assessment.

11 / CASH RECOVERY

Continued supply does not guarantee a return

Illustration | Initial investment 100, five year-end cash flows of 30, 10% discount rate, no residual value.

Same total cash, different net present value

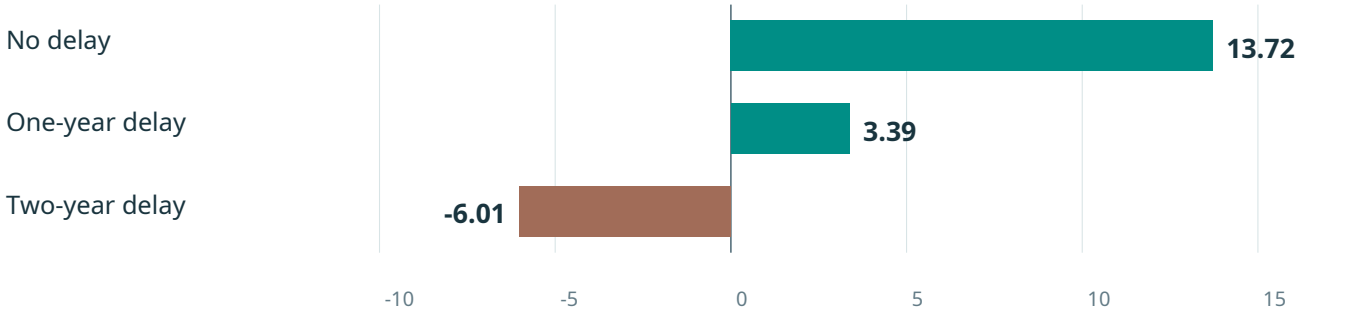

Arbitrary monetary units

Case	Cash receipt timing	NPV
On time	30 at each year-end, years 1–5	13.72
One-year delay	30 at each year-end, years 2–6	3.39
Two-year delay	30 at each year-end, years 3–7	-6.01

Investment of 100 is paid at time zero in every case. The assumed annual 30 is after operating costs, tax and maintenance investment. The same total cash of 150 has less value when received later. NPV = Σ[30/(1.10)^t] - 100.

Deferred equipment payments or customer prepayments would change the result. Idle-period costs or a shorter useful economic life could make the outcome worse. Those additional effects are excluded from this model.

What to assess for each operator

The relevant gap is between payments and the realization of service cash or validated internal savings. A supply-continuity announcement does not remove this timing question.

12 / KOREA VALUE CAPTURE

GPU ownership does not automatically create Korean profit

Separate the economics of buyers, component suppliers and service operators.

Participant	Source of economic value	Evidence needed
Memory/components	Firm orders; quality and delivery	Product volumes, margins, collection
Cloud operator	Repeat customers and reliable service	Paid use, net price, commitments, credit
Manufacturer	Lower design, validation or process losses	Field validation and repeatable net savings
Research/public	Better access and development speed	Completed work, queues, resource efficiency

Samsung and SK partnerships span manufacturing and AI services. Korea's GPU purchases cannot be booked as an equal amount of new domestic memory revenue. Customer investment and supplier orders must be matched by product and time. [6] [7] [9]

Internal manufacturing use requires two tests: a validated technical improvement, then a realized commercial benefit. More simulations do not imply proportionate gains in yield or development speed. Hyundai's training and validation work can be assessed through that same conversion. [8]

Domestic accelerators: more than an all-or-nothing choice

On 8 November, NVIDIA, Rebellions and SK Telecom announced R200 NPU infrastructure collaboration using NVLink Fusion. The architecture connects custom compute to NVIDIA rack infrastructure. More chip choices do not automatically eliminate software or network dependencies. [15] [16]

> **The economics of optionality**
>
> Start with specific inference workloads meeting the same quality and latency, then compare migration, validation and operating costs. An unvalidated domestic chip is not an immediate substitute for large-scale training.

13 / CONDITIONAL OUTLOOK

The base case is uneven execution, not a blanket halt

Conditional business scenarios, without invented probabilities, price targets or undisclosed orders.

Path	Conditions	Implication for Korea
Base: staged delivery	Continuity reaffirmed; orders delivered; sites and customers progress.	Plans continue, but launch dates and recovery differ by project.
Upside: faster activation	Acceptance, power and integration advance alongside paid demand.	More early compute, cash or validated savings from the same capital.
Downside: terms reset	Repeated delivery changes combine with weaker funding or customers.	Expansion delayed; working-capital and fixed-cost pressure; workload shifts.

Evidence supporting the base case

Company and government continuity statements, together with progress in contracts and customers, argue against treating a blanket halt as the default. Successful early contracts still do not guarantee returns on every subsequent expansion. [2] [3] [11] [12]

Concrete evidence that would reverse the judgment

Official allocation cuts or cancellations would weaken supply continuity. Even without cancellations, delayed acceptance plus customer attrition would weaken economics. Faster delivery accompanied by more idle equipment would not count as upside evidence.

> **Central conclusion**
>
> The plan’s success depends less on one CEO’s tenure than on whether committed equipment becomes usable service and recurring value when it is needed.

14 / MONITORING & DECISION

Follow contracts and operation, not personalities

Update the view on the separate clocks of earnings, delivery and commercialization.

Timing	Next evidence	Strengthening / weakening signal
18 Nov KST	NVIDIA results/call scheduled for 17 November PT in the U.S. [2]	Customer and product continuity / specific revisions
Early 2027	Arrival of the remaining c.14,000 direct public procurement units. [3]	Acceptance into use / repeated delivery and launch delays
H1 2027	NAVER's planned AI factory revenue recognition. [12]	Service and collection / widening contract-to-revenue gap
2027–2028	SK first-factory plans and NAVER's 200MW expansion target. [9] [11]	Staged funding, facilities and demand / delayed execution

These are announced dates or targets, not completed outcomes. Distinguish company fiscal periods and Korean dates. A national procurement aggregate cannot stand in for a project's delivery or commissioning schedule.

Contracts

Orders, capital paid, acceptance and cancellation

Operation

Available time, quality, bottlenecks and utilization

Economics

Net prices, costs, recovery and savings

Final judgment

The useful question is not simply whether 260,000 GPUs "disappear." Recognize supply continuity while separating mature projects from those still dependent on further approvals, facilities and demand.

15 / SOURCE REGISTER

Trace the argument to its original sources

Teal citations, source titles and “Open source” links lead directly to the original material.

01. NVIDIA / SEC | Form 8-K: leadership changes Open source

12 Nov 2026 | Filing | Item 5.02; event date 11 November

02. NVIDIA | CEO transition and business continuity Open source

12 Nov 2026 | Company statement | Interim CEO, commitments, earnings date

03. Yonhap News | Korea says NVIDIA GPU commitments remain Open source

13 Nov 2026 | Reporting of minister’s remarks | Public procurement progress

04. NVIDIA | Korea AI infrastructure and ecosystem plans Open source

30 Oct 2025 | Company plans | Public and private GPU allocations

05. Korea MSIT | Government and industry AI ecosystem partnership Open source

2025 announcement | Government statement | Public/private scope and research

06. NVIDIA | Samsung semiconductor AI factory Open source

30 Oct 2025 | Company plans | Manufacturing, design and digital twins

07. NVIDIA | SK Group AI factory collaboration Open source

30 Oct 2025 | Company plans | GPU services, manufacturing and timing

08. NVIDIA | Hyundai Motor Group AI factory Open source

30 Oct 2025 | Company plans | 50,000 GPUs for training and deployment

Scope: announcements establish the stated stage, not completed deployment or revenue. Contract totals are not period revenue or profit. Source [3] is news reporting of a government statement, distinct from a filing.

16 / SOURCE REGISTER

Contracts, technology and analytical boundaries

Teal citations, source titles and “Open source” links lead directly to the original material.

09. NVIDIA / SK Group | Expanded AI factory and memory partnership Open source

24 Jul 2026 | Letters of intent and plans | 2GW; initial deployment in 2027

10. NVIDIA / NAVER / Brookfield | Korea AI factory expansion and conditional investment Open source

24 Jul 2026 | Proposed financing | Preconditions and 200MW target

11. NAVER | Infrastructure agreement and investment preconditions Open source

27 Oct 2026 | Company announcement | USD9bn agreement, SPV, deployment

12. NAVER | Q3 results and initial AI factory customer contracts Open source

12 Nov 2026 | Company announcement | KRW2tn contracts; H1 2027 revenue plan

13. NVIDIA | DGX B200 User Guide: system configuration Open source

As of 14 Nov 2026 | Technical documentation | 8 GPUs, 1,440GB, max. 14.3kW

14. NVIDIA | RTX PRO AI Factory: Components Open source

As of 14 Nov 2026 | Technical documentation | 96GB GDDR7, eight-GPU nodes

15. NVIDIA | NVLink Fusion and custom compute integration Open source

First published 18 May 2025 | Technical explanation | ASIC/CPU integration

16. NVIDIA / Rebellions / SKT | Korea AI infrastructure with NVLink Fusion Open source

8 Nov 2026 | Company announcement | R200 NPUs and rack integration

Scope: announcements establish the stated stage, not completed deployment or revenue. Contract totals are not period revenue or profit. Source [3] is news reporting of a government statement, distinct from a filing.

17 / METHODOLOGY & USE

Methodology, disclaimer and permitted use

Reported facts, company plans, TradeLens analysis and illustrative assumptions are kept separate.

Item	Basis
Date and scope	14 Nov 2026. NVIDIA leadership transition and Korean GPU supply/execution; events and announcements through 13 November.
Source status	SEC filings, company/government statements and reporting of ministerial remarks are distinguished. Initial 2025 plans are separate from 2026 progress.
Quantities	260,000 is a multi-year plan. Public progress is reported on a B200 basis, not national installations or a performance conversion.
Illustrations	Pages 10–12: power, utilization and NPV under stated assumptions. Not forecasts of national demand, revenue or investment returns.
Use	General industry, execution and capital-efficiency analysis. No investment recommendations, price targets, personalized advice or contract-specific legal opinion.

Disclaimer

This report provides general information on markets, industries and corporate economics. It does not provide buy/sell recommendations, trading timing, fair-value or target prices, or personalized investment advice. Information and analysis are as of the stated date; accuracy, completeness and timeliness are not guaranteed. Forecasts, assumptions and illustrative calculations are uncertain and do not guarantee principal or investment returns.

Users are responsible for their investment decisions. Except where liability is imposed by applicable law, TradeLens accepts no liability for investment losses arising from use of this report.

Copyright and use restrictions

Text, charts, design, editing and arrangement in which TradeLens holds rights are protected by applicable law. Beyond personal reading and storage, prior written consent is required to modify, edit, translate, adapt, reproduce, transmit, republish, redistribute or sell any part, whether paid or free and whether as files or screenshots. Removing logos, source or rights notices, or distributing under a third party's name is prohibited.

Uses permitted by applicable law and the separate licensing terms of third-party materials remain subject to those rules. Share the original link on the official TradeLens publication page.