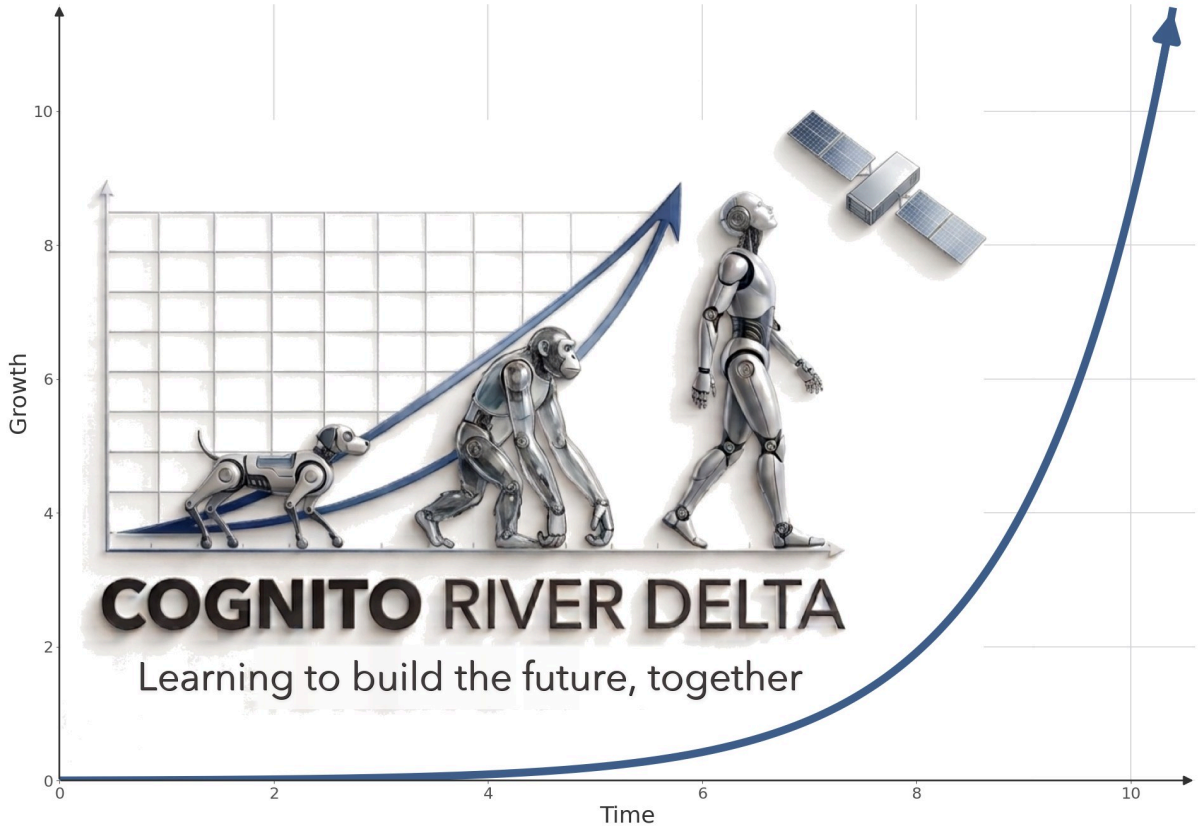


CATEGORY ROLLUP

Semiconductors & chip supply chain - 2026-W45

Week of November 2, 2026

This week, the narrative within the semiconductor and AI infrastructure sector underwent a fundamental transformation. The industry has moved past the relatively simple debate over "chip demand" and entered a far more vo

CATEGORY ROLLUP

Semiconductors & chip supply chain - 2026-W45

Week of November 2, 2026 · 19 min read

Semiconductors & chip supply chain | Policy, geopolitics & China

Semiconductors & chip supply chain · week 2026-W45: Nov 02 - Nov 08, 2026 · 3 subtopic(s) covered · 2816 words · expanded

Weekly Notebook: The Silicon Sovereignty Race

Overview

This week, the narrative within the semiconductor and AI infrastructure sector underwent a fundamental transformation. The industry has moved past the relatively simple debate over "chip demand" and entered a far more volatile era defined by the struggle over "physical and political reality." We are witnessing a collision between unprecedented technological ambition—typified by SpaceX's massive Blackwell deployments and Elon Musk's trillion-dollar vision for a revamped manufacturing paradigm—and the sudden, sharp friction of terrestrial constraints: power shortages, political moratoriums, and high-stakes national security scandals.

The threads across this category are no longer disparate; they are inextricably linked by the emergence of the "AI factory" concept. In this new paradigm, the ability to generate intelligence is becoming limited not by the microscopic design of the silicon wafer, but by the macroscopic ability to power it, cool it, and secure the human and political environments required to build it. This week's developments highlight a strange duality: while industry titans like NVIDIA and SpaceX are looking toward the heavens to bypass Earth's limits through

projects like "StarMind," they are simultaneously being pulled down by the gravity of Congressional probes, espionage allegations, and state-level regulatory crackdowns in critical hubs like Texas.

Terafab: The Quest for Hyper-Vertical Integration

The Terafab project, the ambitious chip-manufacturing arm associated with the Musk ecosystem (Tesla and SpaceX), has transitioned from a theoretical roadmap to a primary source of extreme valuation tension. The core of the Terafab storyline is a radical attempt at vertical integration designed to solve the industry's most persistent headache: the development cycle. Musk’s stated objective is to shrink the chip development cycle from the industry standard of one year to a mere single month—a move that would fundamentally break the current semiconductor cadence and force a total reconfiguration of how hardware is iterated.

The Scale of Ambition and Financial Divergence

The scale of this ambition is causing massive divergence in reporting, reflecting the uncertainty surrounding whether Terafab is a specialized production line or a total reconfiguration of the global manufacturing landscape. On one end of the spectrum, some reports characterize the project as a $16.8 billion bet aimed at producing 100,000 chips per month. However, other estimates are significantly more staggering: Creative Loafing Charlotte describes a "$119 billion AI chip factory ambition," and Elon Musk himself has suggested the cost of a new Texas-based chip production facility could reach $1 trillion.

This isn't just a difference in accounting; it represents a fundamental question of capital intensity. The project is expected to involve a research "Fab Lab" at Giga Texas and a production "tariff fab" in Grimes County. While the timeline for the first test chips is aggressively set for mid-2025 (specifically May or June), with volume production targeted for late 2028, the sheer capital required to realize a trillion-dollar manufacturing footprint represents an unprecedented bet on the future of proprietary silicon.

Strategic Prioritization and Corporate Consolidation

Current developments suggest a tactical shift in Terafab’s roadmap. While early discussions focused on the development of "AI5" chips, recent reports indicate that Terafab’s initial

production runs will prioritize memory. This is a strategic response to the critical, persistent memory shortages that are currently bottlenecking the broader AI industry. By securing in-house memory production, Musk aims to mitigate geopolitical risks and accelerate the deployment of robotics and autonomous systems.

Furthermore, Terafab is acting as a massive catalyst for corporate consolidation. The boards of Tesla and SpaceX have confirmed the formation of special committees to investigate a potential merger of the two companies, with Terafab's massive requirements cited as a primary driver for deepening the ties between the two entities.

The Partnership Landscape and Foundry Sovereignty

The partnership landscape for Terafab remains a high-stakes negotiation over operational control. Intel is reportedly joining the project, and TSMC has expressed significant interest. However, TSMC's participation comes with a critical caveat: the foundry reportedly demands the right to run the fabrication facility itself. This highlights the central tension of the project: while Musk seeks the hardware necessary to power his robots and rockets, the world's premier foundry is not willing to act as a silent partner in a Musk-led ecosystem, insisting instead on maintaining its own manufacturing standards and operational sovereignty.

Semiconductor Fabrication and Chip Manufacturing: The Shift from Silicon to Infrastructure

For the broader fabrication sector, the industry's primary "bottleneck" has undergone a regime change. For much of the last two years, the conversation focused on silicon availability. This week, the consensus shifted toward a different, more difficult set of constraints: power, permitting, and cooling.

The Terrestrial Wall: Power, Permitting, and the Texas Risk

While long-term demand for fabrication remains robust—evidenced by ASML receiving orders extending through 2030 and TSMC's construction of nearly 20 new fabs—the immediate growth of AI compute is hitting a terrestrial wall. The primary risks are no longer merely technical; they are political and physical.

A major systemic risk has emerged in Texas, a critical hub for planned capacity. Following the 2026 midterms, Texas Governor-elect Gina Hinojosa has pledged a moratorium on new data center construction. Given that Texas holds between 25% and 30% of the planned U.S. data center capacity, this political shift poses a direct threat to the deployment of new hardware. This underscores a growing debate: whether the expansion of AI is being halted by a lack of chips, or by a lack of "permission" from the communities and governments that must host them.

The Bottleneck Evolution: From EUV to HBM

There is an ongoing debate regarding which technical constraint will ultimately choke the supply chain next. Some analysts, such as David Carbutt, identify EUV (Extreme Ultraviolet) lithography machines and the physical build-out of fabs as the immediate primary constraints. However, others, including Peter Diamandis, predict a shift toward a different kind of scarcity within the next 2-3 years: a bottleneck in CPUs and High-Bandwidth Memory (HBM).

This suggests that even if the "silicon" problem is solved via massive fab expansion, the industry may immediately run into a "memory and architecture" problem. This volatility is already being felt; for example, nearly 600,000 H200 chips are reportedly currently stranded in Hong Kong due to Chinese government import delays, demonstrating how quickly logistics and geopolitics can disrupt the flow of high-end hardware.

The Threat of Vertical Compression

As vertical integration becomes the stated goal for giants like Tesla (via Terafab), the traditional semiconductor model is bracing for a potential compression of profit margins. If the largest consumers of chips—the companies building the AI factories—become the manufacturers themselves, the traditional foundry model faces an existential challenge. The shift toward "Physical AI" and in-house silicon could transform the most profitable customers into the industry's most formidable competitors.

NVIDIA and Jensen Huang: A Leadership and Geopolitical Crisis

NVIDIA experienced one of its most turbulent weeks in history, characterized by a "perfect storm" of personal scandal, geopolitical friction, and unprecedented infrastructure scaling.

The Leadership Crisis and Espionage Allegations

The most explosive development is the investigation into CEO Jensen Huang. Reports from multiple outlets, including the *Financial Times* and *The New York Times*, indicate that the NVIDIA board is investigating Huang following allegations of a romantic relationship with Alice Huang, a Megaspeed executive. Crucially, U.S. intelligence reports suggest that Alice Huang is suspected of being an agent for China's Ministry of State Security (MSS).

This has escalated from a matter of corporate governance to a high-level national security issue. Huang has been summoned to brief Congress to address these potential ties. The personal and professional risk to the company is quantifiable, with analysts like Jeffrey Sonnenfeld estimating a "50-50" chance of Huang remaining in his position. While both NVIDIA and the Chinese Embassy have rejected these claims, the reputational and regulatory fallout remains a significant headline risk.

Geopolitical Decoupling and the China Pivot

Simultaneously, NVIDIA is navigating a complex geopolitical minefield as it attempts to balance U.S. export controls with the need to maintain market access in China. To satisfy massive, multi-billion-dollar orders from Chinese firms like ByteDance and Alibaba—reportedly totaling $35 billion for 3 million RTX Pro 5500 units—NVIDIA is reportedly shifting its high-end GB202 dies away from consumer products (like the RTX 5090) and toward high-margin professional chips.

This pivot is a strategic attempt to navigate the "decoupling" of the tech sectors. However, the friction is visible: NVIDIA has reportedly instructed cloud partners to sell 51% of services outside of China, and strict ID requirements have been implemented for U.S. consumers purchasing RTX 5090s to prevent unauthorized diversion to the mainland.

Unprecedented Scale: Blackwell and the "AI Factory"

Despite these crises, NVIDIA's hardware deployment remains in an era of massive, almost incomprehensible scale. The company is the primary engine behind SpaceX's infrastructure build-out, specifically the "Colossus 2" project. This project targets the deployment of 1.2 million Blackwell GPUs by the end of 2026, utilizing massive clusters like the "Gun Range" (220,000 GB300s) and "Spam" (200,000 GB300s).

To combat terrestrial constraints, NVIDIA is innovating on how hardware is delivered. The introduction of "DSX Pods"—factory-built, integrated GPU racks—is designed to slash data

center setup times from the traditional six months to just 14 days. Furthermore, NVIDIA is pushing "AI-RAN" (Radio Access Network) technology, a strategic effort to transition telecommunications companies from mere connectivity providers into AI service providers, embedding AI into the global 6G infrastructure.

The Orbital Escape Hatch: Project StarMind

Perhaps most ambitiously, NVIDIA is co-designing "StarMind" with SpaceX. This project aims to move AI compute into orbit, potentially bypassing the terrestrial limitations of power, water, and political opposition. If successful, this would represent a total pivot for the industry: moving from terrestrial data centers to an orbital, space-based compute economy. This move is intended to leverage the declining launch costs of Starship to create a compute environment that is physically and politically insulated from Earth's constraints.

Cross-cutting themes

The unifying theme of this week's developments is the **systemic attempt to circumvent terrestrial limits.** We are observing three distinct, massive-scale strategies designed to bypass the traditional bottlenecks of the semiconductor industry:

1. **Vertical Integration (The Terafab/Tesla Model):** Bypassing the slow, third-party foundry cycle and the standard one-year development cadence by building in-house manufacturing to collapse development time to one month.
2. **Infrastructure Innovation (The NVIDIA/DSX Model):** Bypassing the logistical and physical delays of data center construction through rapid-deployment "pods" and moving compute to the edge via the "AI-RAN" telecommunications pivot.
3. **Orbital Migration (The StarMind/SpaceX Model):** Bypassing the physical (power, cooling) and political (permitting, community opposition) constraints of Earth by moving the entire compute economy into space.

These threads demonstrate that the semiconductor industry is no longer just a "tech" sector; it has become a "geopolitical and physical infrastructure" sector. The central tension is whether these massive, capital-intensive workarounds—such as $1 trillion fabs or orbital satellite constellations—will arrive in time to prevent a period of AI scarcity caused by the "final constraint" of political and resource opposition on Earth.

Where sources agree

- **The Infrastructure Bottleneck:** There is a clear consensus that the primary constraints on AI growth have shifted from mere chip availability to the availability of power, cooling, and political permits.
- **SpaceX's Massive Scale:** Multiple sources confirm that SpaceX is executing an unprecedented hardware deployment, specifically targeting massive clusters of NVIDIA Blackwell chips to power its autonomous products.
- **The Memory Crisis:** Analysts and industry leaders agree that a critical and persistent shortage of memory is a systemic risk that will continue to bottleneck the industry.
- **NVIDIA’s Market Position:** Despite the leadership and geopolitical crises, there is general agreement that NVIDIA remains the primary engine of the AI scaling era.

Where sources disagree

- **The Financial Sustainability of NVIDIA:** A fundamental disagreement exists regarding NVIDIA's financial health. Some analysts, such as Scott Galloway, view the company's participation in debt offerings for its own customers (like Meta) as "classic bubble behavior" reminiscent of a circular financing model. Others, like Randy Kirk, view the company's massive earnings and hardware dominance as a fundamental, stable market reality.
- **The Feasibility of Orbital Compute:** A significant scientific debate exists over the "StarMind" concept. Critics and astrophysicists argue that the vacuum of space creates a "thermos trap" that makes cooling impossible, while proponents argue that the plummeting cost of Starship launches and the efficiency of radiative cooling make orbital compute the only long-term solution.
- **The Magnitude of Terafab:** There is a massive discrepancy in reported costs, with estimates ranging from a $16.8 billion "bet" to an astronomical $1 trillion requirement for a Texas-based facility.
- **Political Impact:** There is disagreement over whether state-level political opposition (like the Texas moratorium) will actually succeed in halting data center expansion, or if the economic incentives provided by tax revenue and jobs will ultimately force politicians to relent.

Numbers and claims to verify

- **$1 trillion cost** for the Terafab chip plant (Elon Musk).
- **$895 billion** computing spending commitment by Anthropic (Prof G Markets).
- **1.2 million Blackwell GPUs** to be deployed by the end of 2026 (Farzad Mesbahi/Herbert Ong).
- **$35 billion in orders** from Chinese firms for RTX Pro 5500s (Gamers Nexus).
- **14-day setup time** for NVIDIA's DSX Pods (David Carbutt).
- **600,000 H200 chips** stranded in Hong Kong due to import delays (Data Center Dynamics).
- **$40 billion principal** for the NVIDIA Verarubin GPU finance lease (Farzad Mesbahi).
- **80% gross margins** maintained by NVIDIA (Peter Diamandis).

Investment and strategic implications

- **Key Person and Regulatory Risk:** The investigation into Jensen Huang introduces unprecedented leadership instability and potential national security regulatory risk for NVIDIA. Any outcome of the Congressional briefing could fundamentally shift NVIDIA's valuation.
- **The Verticalization Threat:** The emergence of projects like Terafab suggests that the "customers" of the semiconductor industry are increasingly becoming its most direct competitors, threatening the long-term profit margins of traditional foundries.
- **Shift from Compute to Efficiency:** As power and political constraints tighten, investment value may shift from companies that provide the most raw compute to those that provide the highest "tokens per megawatt." NVIDIA’s focus on power efficiency and rapid-deployment DSX Pods is a direct strategic play for this shift.
- **The Memory Upside:** The confirmed, persistent memory shortage provides a massive strategic tailwind for manufacturers like Micron, which currently appears undervalued relative to the growth ceiling of the AI sector.
- **The "Space-Compute" Frontier:** If the "StarMind" project successfully moves from a concept to a viable orbital data center model, it would represent a massive, currently unpriced strategic hedge against terrestrial regulatory and environmental risks.

What to watch next week

- **Congressional Testimony:** Any updates regarding the mandatory briefing of Jensen Huang to U.S. lawmakers regarding the espionage allegations.
- **NVIDIA Board Updates:** Any official communications regarding the internal investigation into the CEO and the potential for leadership succession.
- **Texas Regulatory Status:** News on whether the pledged moratorium on data center construction in Texas is formally codified into law.
- **TSMC/Tesla Negotiations:** Any confirmation regarding the specific terms of TSMC's potential participation in the Terafab project, particularly regarding fab management.
- **Blackwell Deployment Milestones:** Updates on the first large-scale Blackwell clusters being activated in SpaceX facilities.

Appendix: Individual perspectives

We see further because we stand on the shoulders of giants. The people in this appendix are some of those giants: their public work, followed and studied every day and every week, is how CognitoRiverDelta brings the world's AI knowledge and concerns together, so that, in time, our expert panel can help shape real solutions and paths forward for all of us. We learn by observing them, and we think of them as kindred spirits: people helping humanity and AI educate each other and contribute to solutions. Together they form a much larger community of minds, a worldwide panel; ours is a small copy of it. Most people don't want to become expert kayakers, reading rocks and rapids every day; yet that is how this fast-flowing river of change feels to many of us, and the evidence that could ease those fears, the grounded ones and the imagined ones, is still being gathered. That is the work. We credit their ideas here with gratitude; they are not affiliated with us and do not endorse this report.

- **Ben Thompson:** Argues that NVIDIA's competitive advantage is increasingly defined by power efficiency (tokens per megawatt) as political and power constraints become the primary bottlenecks.
- **Bruce Schneier:** Criticizes NVIDIA for using complex marketing of its security hardware to boost stock prices while deploying unsafe products.
- **David Carbutt:** Views NVIDIA’s transition toward "SI factories" and "DSX Pods" as a move to provide a complete, rapid-deployment computing stack via a "reference

design" model.

- **Ed Zitron:** Warns that NVIDIA's growth is threatened by a "circular" financing model and increasing political/municipal resistance to data centers.
- **Herbert Ong:** Highlights the role of Blackwell chips as the critical hardware foundation for SpaceX's autonomous product infrastructure.
- **Jason Calacanis:** Claims NVIDIA has a massive economic multiplier effect, where every $1 spent at NVIDIA translates into significantly more revenue for model providers like Anthropic.
- **Jeff Lutz:** Characterizes Jensen Huang as a high-velocity leader who uses hands-on execution to bypass corporate bureaucracy.
- **Jensen Huang:** Views the era as a shift toward "AI factories" that transform energy and data into intelligence.
- **Peter Diamandis:** Sees NVIDIA chips as the essential hardware driving the rapid scaling of global AI capacity.
- **Scott Galloway:** Views NVIDIA as a bubble risk due to its participation in debt offerings for its own customers (circular financing).
- **Steve Burke:** Compares NVIDIA's ecosystem to a closed, "Apple-like" system that uses political influence and proprietary hardware to protect its market.
- **Steven Mark Ryan:** Questions whether the allegations against Jensen Huang are a major security breach or a temporary distraction.
- **Tony Seba:** Points to the massive improvement rates predicted by Huang as a driver of imminent systemic disruption.

Sources

- "Congress will investigate": Nvidia CEO Jensen Huang faces probe over alleged Chinese spy — Axios (via Google News), Nov 08
- "Jensen Huang's Eldest Daughter" Madison Huang Abruptly Cancels South Korea Visit... "Fami — 아시아경제 (via Google News), Nov 08
- 'A Very Complex Issue': Trump Leaves Door Open to Dropping Nvidia CEO as Tech Adviser Afte — International Business Times UK (via Google News), Nov 02
- A Retiree Who Bought 1,000 Nvidia Shares at $122 Just Watched Chips Have Their Worst Day S — 24/7 Wall St. (via Google News), Nov 07
- AI Watch: Anthropic's IPO plans hit hurdles, Nvidia under scrutiny, Google ramps up invest — CNBC TV18 (via Google News), Nov 06

- AI’s Stock Market Fallout, Google’s Profit Problem, and Altman’s “Hold My Beer” Moment | # — Peter H. Diamandis, Oct 31
- AMD & NVIDIA Are About to Go Bonkers — David Carbutt, Oct 31
- AMD's AI Data Center Guarantees Reach $8.5 Billion as It Challenges Nvidia — Yahoo Finance (via Google News), Nov 04
- ASML Rebounds 5.6% as AI Regulation Fears Ease and Musk Flags Faster Terafab Deliveries - — Google News, Nov 07
- Alleged Jensen Huang spy plot highlights Nvidia export concerns ... — eeNews Europe (via Google News), Nov 02
- Ask HN: Are Tenstorrent cards better than Nvidia? — Hacker News, Nov 03
- Breakingviews - COMMENTARY: Nvidia crisis may hurt China more than it helps — Reuters (via Google News), Nov 09
- CNBC's The China Connection newsletter: Nvidia's humanoid push relies on China-made robot — CNBC (via Google News), Nov 02
- Chamath Reveals Elon's Devastating Power Move — Steven Mark Ryan, Nov 08
- China Newspaper Dismisses Jensen Huang Spy Claim, Warns on Ties — Hacker News, Nov 09
- China espionage probe: Why US lawmakers want Nvidia CEO Jensen Huang to testify — Firstpost (via Google News), Nov 06
- China’s top spy agency recruited woman who targeted Nvidia’s Jensen Huang — The Japan Times (via Google News), Nov 07
- Chinese Intelligence Targeted NVIDIA CEO Jensen Huang in AI Chip Espionage Effort: Report — Security Boulevard (via Google News), Nov 07
- Chinese Spy Who Targeted Jensen Huang Worked Directly for China's Top Spy Agency: Report — Gizmodo (via Google News), Nov 07
- Crypto’s biggest heist ever, the spy who loved Nvidia and Hacktober’s unsolved mysteries : — NPR (via Google News), Nov 04
- Crypto’s biggest heist ever, the spy who loved Nvidia and Hacktober’s unsolved mysteries : — NPR (via Google News), Nov 04
- Does SpaceX’s $2Trillion Bet have a Physics Problem? — Jo Bhakdi, Nov 07
- Elon Goes Back On His Word For Tesla | It's Not Happening Anymore — Ryan Shaw, Oct 22
- Elon Musk Claims His Terafab Will Produce Its First Chip By Mid-2027 — Gizmodo (via Google News), Oct 22
- Elon Musk Details SpaceX Tesla Merger Talks and 2028 Terafab Production Target — Technetbook (via Google News), Oct 22

- Elon Musk Touts Terafab's Ambitious Goal to Cut Chip Development Time From a Year to a Mon — Benzinga (via Google News), Nov 07
- Elon Musk claims Terafab's output value "in a single building" will place it in "top 20 of — primetimer.com (via Google News), Oct 22
- Elon Musk's Terafab chip plant could cost $1 tr... — Pluang (via Google News), Nov 07
- Elon Musk: Terafab Austin factory to test chip mid 2027 — NewsBytes (via Google News), Oct 22
- Every Tesla Investor Needs To Watch This — Brighter with Herbert, Nov 04
- FT: US government believes Nvidia CEO Jensen Huang had romantic relationship with suspecte — 헤럴드경제 (via Google News), Nov 07
- Formalizing Nvidia PTX in Lean to Prove GPU Kernel Correctness — Hacker News, Nov 03
- From Chips to Skills, NVIDIA and Korea Expand Sovereign AI Collaboration — NVIDIA Blog, Nov 09
- FuriosaAI raises $620M Series D, 4Xs valuation to $2.9B in Nvidia challenge — Hacker News, Nov 05
- Here's what smart people are saying about Nvidia's investigation of CEO Jensen Huang — Business Insider (via Google News), Nov 09
- House Committee Chairs Ask NVIDIA CEO Jensen Huang for Closed-Door Briefing on Alleged Chi — XenoSpectrum (via Google News), Nov 07
- House Republicans Seek Closed-Door Briefing With Nvidia CEO Over China Spy Allegations — Businesskorea (via Google News), Nov 06
- How Nvidia came to investigate its CEO's affair with suspected Chinese spy — ABC News & Headlines – Australian Broadcasting Corporation (via Google News), Nov 09
- How Open Models Win Against OpenAI and Anthropic — Nathan Lambert, Nov 02
- How Will Markets React to Elections. — Randy Kirk, Nov 04
- Hwang Madison cancels Korea trip, derails meetings as Nvidia AI Day proceeds — Chosunbiz (via Google News), Nov 08
- Inside Colossus 2: How Elon Is Wiring 1.2 Million GPUs — Farzad Mesbahi, Nov 04
- Inside Mar-a-Lago's Jungle Gala: Donald Trump Blasts Alina Habba Over Reported $50M Pardon — Istituto Impariamo (via Google News), Nov 02
- Intel (INTC) Joins Musk's Terafab AI Chip Venture In High Profile Alliance — Simply Wall Street (via Google News), Oct 22
- Is Jensen Huang’s job on the line? Yale’s Jeffrey Sonnenfeld estimates it’s “about a 50-50 — TechCrunch (via Google News), Nov 09

- Is Jensen Huang’s job on the line? Yale’s Jeffrey Sonnenfeld estimates it’s “about a 50-50 — TechCrunch AI, Nov 09
- Jensen Huang Defends Nvidia After Spy Investigation Report. Why I’d Buy the Dip. — Barchart.com (via Google News), Nov 04
- Jensen Huang Is Bigger Than Nvidia — And It’s a Problem — Bloomberg.com (via Google News), Nov 09
- Jensen Huang Loses Top 10 Richest Spot as Nvidia Shares Plunge — Tempo.co English (via Google News), Nov 06
- Jensen Huang Rumored to Have Had Romantic Relationship with Chinese Spy... Controversy Gro — SBS 뉴스 (via Google News), Nov 08
- Jensen Huang Thinks A.I. Alarmism Has Gone Too Far | The Ezra Klein Show — finance.biggo.com (via Google News), Nov 02
- Jensen Huang and Nvidia's Only "Offense"? Doing Business — RealClearMarkets (via Google News), Nov 03
- Jensen Huang says Nvidia sees a place for both open and closed AI models — Mezha.net (via Google News), Nov 03
- Jensen Huang to Appear Before Congress to Explain Ties to Woman Suspected of Spying for Ch — 파이낸셜뉴스 (via Google News), Nov 06
- Jensen Huang's Eldest Daughter Cancels Korea Visit Amid Suspicions of Chinese Spy Ties? — SBS 뉴스 (via Google News), Nov 08
- Jensen Huang's Robot Strategy Embeds China Supply Chain: Both Nvidia Training Systems Use — BigGo Finance (via Google News), Nov 03
- Jensen Huang: Silicon Valley superstar engulfed by personal scandal — BBC (via Google News), Nov 09
- Key facts: SPCX Rebounds; NASA Adds $946M; Terafab Advances — TradingView (via Google News), Nov 07
- Korea Advances Physical AI Across Industries With NVIDIA — NVIDIA Blog, Nov 09
- Korea’s AI Ecosystem Advances and Secures Agentic AI With NVIDIA — NVIDIA Blog, Nov 09
- MSI XpertStation WS300 Review Nvidia DGX Station GB300 — Hacker News, Nov 04
- Marvell Stock Has Plunged 46% Since Jensen Huang Called It the Next $1 Trillion Company. H — Yahoo Finance (via Google News), Nov 04
- Marvell to Buy Lumentum for $115 Billion in Stock Deal; Nvidia CEO Huang Backs Merger — bloomingbit (via Google News), Nov 03
- Megaspeed exec was recruited to spy on Nvidia CEO, according to US intelligence memo — TechCrunch (via Google News), Nov 07

- Megaspeed exec was recruited to spy on Nvidia CEO, according to US intelligence memo — TechCrunch AI, Nov 07
- Megaspeed exec with close ties to Nvidia CEO Jensen Huang suspected of spying for China — Data Center Dynamics (via Google News), Nov 02
- Megaspeed exec with close ties to Nvidia CEO Jensen Huang suspected of spying for China — Data Center Dynamics, Nov 02
- Microsoft opens GitHub Copilot agent harness to developers through Foundry — TechGig (via Google News), Oct 31
- Musk Says His Terafab Chip Plant Could Cost $1 Trillion — Yahoo Finance (via Google News), Nov 07
- Musk Says His Terafab Chip Plant Could Cost $1 Trillion — 24/7 Wall St. (via Google News), Nov 07
- Musk announced the timeline for the release of Terafab chips and Optimus robots — Українські Національні Новини (УНН) (via Google News), Oct 22
- Musk targets Terafab mass chip production by Q3 2028 — IDNFinancials (via Google News), Oct 22
- NEW INTERVIEW: Elon Just Put Odds on $1 Trillion SpaceX Revenue — Brighter with Herbert, Nov 07
- NVIDIA CEO Jensen Huang's Vision for the Future | Cleo Abram: HUGE Conversations — finance.biggo.com (via Google News), Nov 02
- NVIDIA CEO and Sweden's Technology Leaders Discuss Compute's Next Chapter — digitimes (via Google News), Nov 03
- NVIDIA Dynamo Recipes Bring Korea's Leading Open Models to Production — NVIDIA Blog, Nov 09
- NVIDIA Enters Technical Bear Market; Jensen Huang Drops Out of Top 10 Global Billionaires — 富途牛牛 (via Google News), Nov 06
- NVIDIA GTC Berlin 2026 recap, keynote and announcements — Plain Concepts (via Google News), Nov 03
- NVIDIA GTC Berlin | NVIDIA's Jensen Huang on the 'Marvel' of AI Datacenters — Yahoo (via Google News), Nov 03
- NVIDIA's CEO Concerns Impacting AI Firms like OpenAI and Anthropic — GuruFocus (via Google News), Nov 04

🪶 ORBIT LAB · THE ROAD TO ORBIT, HANDS-ON

Launch a rocket, link satellites with lasers, beam power down and fly to the sun →

- Why one Starship launch of V3 satellites adds 20× or more the bandwidth of a Falcon 9 launch - about 10× per satellite, about 3× as many per launch
- Cooling in space: how a computer in orbit radiates its heat away - set the radiator temperature yourself
- **New, iteration 2:** See it from the Moon - a rail gun fires compute satellites into a ring around the sun (our forecast, not to scale)
- **Iteration 3:** Real scale - Earth and the Moon to scale, and the round-spheres fix a real-phone tester found
- How much of the sun do we use? - and why AI compute follows the sun

An interactive planet and solar system, built with an AI coding partner in a morning. How we built it

Related reports

COMBINED REPORT

Transition to Physical and Orbital AI

How AI is crossing from software into the physical world and into orbit - robots, factories, energy and space-based compute, read together.

September 25, 2026 · Datacenters & cloud · Semiconductors & chip supply chain

WEEKLY DIGEST

Weekly video script - 2026-W40

The AI industry has hit a definitive inflection point, moving from text-processing chatbots to agentic systems with physical arms and legs. We're also breaking down Nvidia's expansion into a sovereign infrastructure empi

Week of September 28, 2026 · Frontier labs & models · Semiconductors & chip supply chain

CATEGORY ROLLUP

Semiconductors & chip supply chain - 2026-W40

This week marked a profound evolutionary leap for NVIDIA, transitioning the company from a dominant provider of AI hardware into

what can only be described as a sovereign, multi-domain infrastructure empire. The central

Week of September 28, 2026 · Semiconductors & chip supply chain

Everything here is free, and reader-supported. Join free on Patreon to get the weekly digest by email, or become a member to get each morning's briefing by email before it reaches the site, request custom analysis, and get early access to new research tools. What members get

♥ Join or support on Patreon

Informational analysis synthesized by AI from sourced, dated material, curated by a human. Treat specific claims as unverified until checked. Not financial advice.

About · How this is made · Corrections