SEMICONDUCTOR BRIEFING

The State of Computer Chips

Vol. 1 · July 2026

NVIDIA & THE GPU INCUMBENTS

CUSTOM SILICON & CHALLENGERS

FOUNDRY

EQUIPMENT & EUV

MARKET, PRICING & DEMAND

GEOPOLITICS & CHINA

BY **DANIEL SHANKLIN**

danielshanklin.com/stateof/chips

SECTION 01

Executive Overview

The compute supercycle is intact, the supply chain is straining, and Washington and Beijing are both tightening their grips. The numbers that define May–July 2026.

$81.6B

NVIDIA Q1 FY27 revenue, +85% YoY — data center alone $75.2B (+92%)

NVIDIA IR · MAY 20

$91.0B

NVIDIA Q2 FY27 revenue guide (±2%) — with China data-center compute excluded entirely

NVIDIA IR · GM ~75%

$40.2B

TSMC record Q2 2026 revenue, +33.7% YoY at 67.7% gross margin; net income +77%

TSMC IR · JUL 16

>40%

TSMC's raised FY2026 USD revenue-growth guidance (from >30%); 2026 capex raised to $60–64B

TSMC IR · JUL 16

€43–45B

ASML FY2026 sales guidance; EUV system sales up more than 45% YoY, ~65 low-NA systems targeted

ASML IR · JUL 15

~$140B

Lam Research's raised 2026 wafer-fab-equipment forecast (from $135B) on a record $5.84B quarter

LAM IR · APR 22

+110%

Reported DDR5 price surge as AI devours memory; AI-server memory lead times reported >40 weeks

REPORTED · NEWS FEEDS, JUL

~$20B

NVIDIA's largest deal ever: a non-exclusive license to Groq's inference tech — plus its CEO and ~90% of staff

CNBC · DEC 24, 2025

1. NVIDIA posted $81.6B in Q1 FY27 revenue (+85% YoY, +20% QoQ)[1], with data center at $75.2B (+92%) and networking nearly tripling (+199% to $14.8B). It guided Q2 to $91.0B ±2% at ~75% gross margin — and excluded China data-center compute from the outlook entirely.

2. TSMC's Q2 was a record NT$1,270.38B (US$40.2B, +33.7% YoY) at 67.7% gross margin[2]; it raised full-year USD growth guidance to "slightly more than 40%" and 2026 capex to $60–64B. HPC is now 66% of revenue, N2 booked its first commercial revenue, and N3 is sold out.

3. The memory famine is the sector's binding constraint: DDR5 prices are reported up ~110%, AI-server memory lead times reportedly exceed 40 weeks, and SK Hynix's CEO warned July 11 that **2027 will be "the worst year" of the shortage**, with constraints possibly persisting to 2030 (as reported via news feeds).

4. The custom-silicon wave broke into the open: Google split its 8th-gen TPU into training and inference parts and began direct TPU sales (reported from Cloud Next); AWS Trainium3 is reportedly sold out; Microsoft's Maia 200 is reportedly in production (per DCD); OpenAI and Broadcom unveiled a first custom inference chip on June 24; and Meta's MTIA is reportedly headed to production in September (per Reuters).

5. NVIDIA paid ~$20B — its largest deal ever — for a *non-exclusive* license to Groq's inference technology[3], hiring CEO Jonathan Ross and ~90% of staff (CNBC, Dec 24, 2025). The antitrust-shaped structure is the story; the remainder of Groq raised $650M in June to rebuild as an inference neocloud.

6. Cerebras completed the largest US IPO of 2026 on May 14, reportedly raising $5.55B at $185/share and opening near $350 (~$95B cap) — then its first earnings beat on revenue ($193.4M, +92%) but guided to negative full-year margins on giant OpenAI/AWS cloud deals, and the stock fell hard (reported).

7. Equipment is booming: ASML did €9.3B in Q2 at 54.0% gross margin and guided Q3 to €11–12B[4], with EUV capacity expanding 30% for 2027; Lam posted a record $5.84B[5] and raised its 2026 WFE forecast to ~$140B.

8. Intel's 18A entered high-volume production in January — still with no named major external customer[6] — and its Q1 filing warned 14A may be *paused* without sufficient committed demand. The US government holds a 9.9% passive stake. This is the existential foundry gate of 2026–27.

9. The China standoff hardened on every front: H200 export approval is "approved-but-broken" — Commerce told Congress July 14 that only "a small number" have actually shipped; Taiwan ran its first formal smuggling crackdown; NVIDIA slashed its Asian authorized-customer list and deployed field inspectors; and the AI OVERWATCH Act was secured in the Senate NDAA (Jul 15). Meanwhile SMIC printed $2.505B at ~93% utilization[7] and Huawei's Ascend 910D is reportedly drawing orders from ByteDance and Alibaba.

10. Watch next: AMD's "Advancing AI" event lands July 22–23 — **two days after this issue publishes** — with the MI400 volume launch expected (reported), after a Q1 in which AMD's data-center revenue reportedly passed Intel's for the first time. And NVIDIA's Rubin ramp, promised at "giant amounts" for Q3, resolves this fall.

SOURCES

[1] nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-first-quarter-fiscal-2027

[2] investor.tsmc.com/english/quarterly-results

[3] www.cnbc.com

[4] www.asml.com/en/investors/financial-results/q2-2026

[5] investor.lamresearch.com/2026-04-22-Lam-Research-Corporation-Reports-Financial-Results-for-the-Quarter-Ended-March-29,-2026

[6] intc.com

[7] www.smics.com/en/site/news_read/7956

SECTION 02

NVIDIA & the GPU Incumbents

An $81.6B quarter, a $91B guide with China zeroed out, a $1T bookings claim, and a Rubin rumor cycle that whipsawed the market for three months.

NVIDIA Revenue Trajectory — FY26 into FY27

Total revenue (solid) and data-center revenue (muted) by quarter, all primary-sourced from NVIDIA IR[1]. Dashed: Q2 FY27 guidance of $91.0B ±2% — issued with China data-center compute excluded. FY26 closed at $215.9B (+65%).

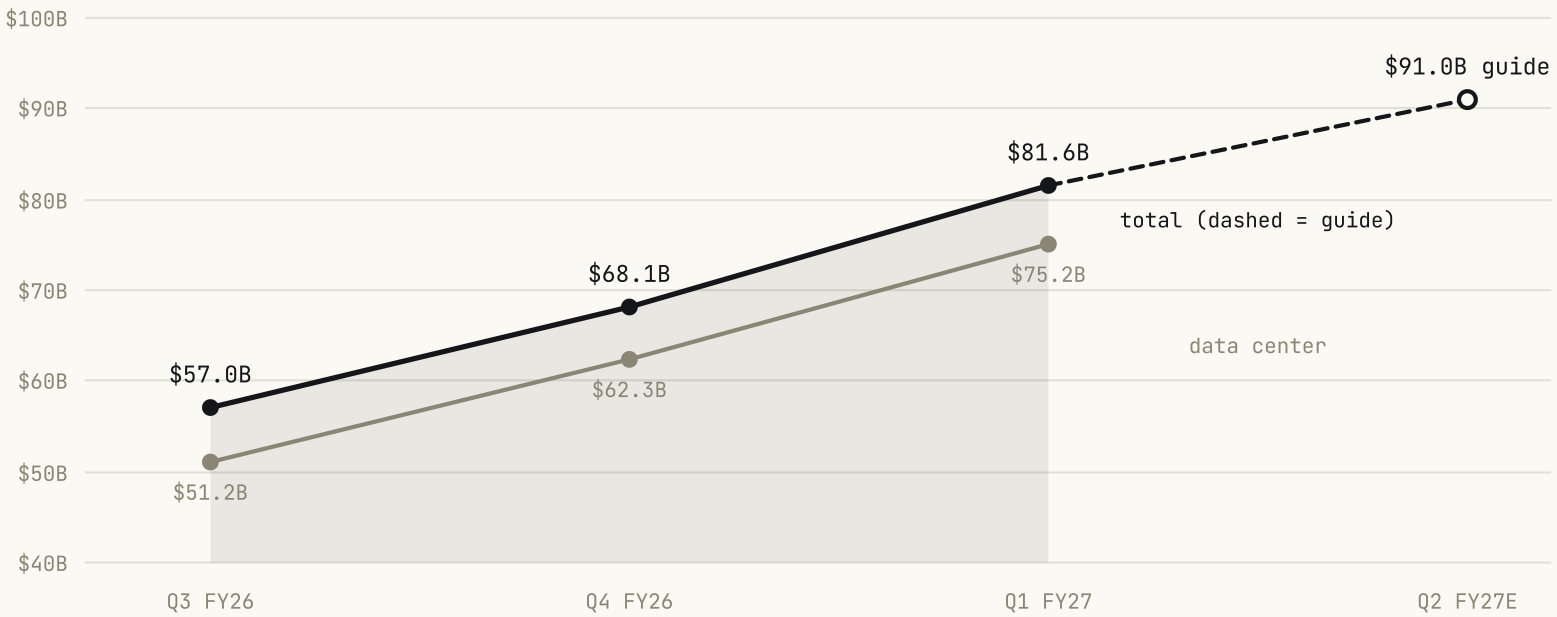

Sources: NVIDIA Q1 FY27 results[1] · NVIDIA Q4 FY26 results[2].

The Quarter in Numbers

QUARTER	TOTAL REVENUE	DATA CENTER	NOTES
Q3 FY26	$57.0B	$51.2B	Blackwell ramp
Q4 FY26	$68.1B	$62.3B	FY26 closes at $215.9B, +65%
Q1 FY27	$81.6B	$75.2B	Compute $60.4B; networking $14.8B, +199% YoY
Q2 FY27 (guide)	$91.0B ±2%	—	GM ~75%; China DC compute excluded

Alongside earnings, NVIDIA announced the **Vera Rubin platform**, the **Vera CPU** ("first processor purpose-built for agentic AI"), **BlueField-4**, and **Dynamo 1.0**, which the company claims delivers up to 7x Blackwell performance on some workloads (NVIDIA IR[1]). Separately, Jensen Huang claimed in March more than **$1T in cumulative Blackwell + Vera Rubin bookings/visibility through 2027** — up from $500B claimed in October 2025 — per Yahoo Finance reporting (Mar 18). That figure was *not* restated at the May earnings call; treat it as claimed March 2026 and not updated since.

The Rubin Whiplash — a three-month rumor cycle (all as reported)

May: design-issue reports surface (per Wccftech) → cleared within weeks, with DigiTimes reporting a 3Q26 ramp (May 15) → NVIDIA declares "full production" (May 31) → Korean media report production cuts (Jun 23) → Huang publicly vows "giant amounts" and "the roadmap is intact" (per Tom's Hardware, Jul 15–16). Meanwhile Bloomberg reported all three HBM4 suppliers certified for Vera Rubin (Jun 4), and clouds have reportedly booked Rubin supply through 2027. Net: the bearish reports have been repeatedly overrun by the announced ramp — but the fall ramp is now a public, gradeable claim.

AMD & Intel: the other two

AMD

ASCENDING

Q1 2026 revenue reportedly $10.3B (+38% YoY) with data center at $5.8B (+57%) — passing Intel's data-center revenue for the first time (reported).

- MI450 sampling since May 5 (reported)
- MI500 engaged with customers for 2H27 (reported)
- "Advancing AI" launch event Jul 22–23 — MI400 volume launch imminent, 2 days after this issue publishes

Intel (products)

Crescent Island (Xe3P) inference GPU: 480GB LPDDR5X at ~1.5TB/s, 350W air-cooled — an explicit no-HBM play against the memory shortage, targeting year-end (per Tom's Hardware, from Computex).

- Bets on capacity-per-dollar, not peak bandwidth
- Sidesteps the HBM bottleneck entirely

SOURCES

[1] nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-first-quarter-fiscal-2027

[2] nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-fourth-quarter-and-fiscal-2026

SECTION 03

Custom Silicon & Challengers

Every hyperscaler now has a serious in-house chip program, and the challenger tier just produced the year's biggest IPO, its strangest acquisition-shaped license, and a $20B-valuation stealth exit.

Hyperscaler Silicon Scoreboard

Program status as of mid-July 2026. Most entries below reached us via news-feed reporting rather than primary IR — treat statuses as reported.

OWNER	CHIP	STATUS (REPORTED)	BUILD PARTNER	SIGNAL
GOOGLE	8th-gen TPU — "8t" train / "8i" infer	Announced Cloud Next (Apr 22); direct TPU sales begun	Broadcom (8t) / MediaTek (8i) — naming per a single feed	~$25B Google+Blackstone TPU-cloud vehicle (May 18); courting NVIDIA neoclouds to host TPUs (Jul 15)
AWS	Trainium3	Reportedly sold out; Q3'26 ASIC server targets raised up to 30%	—	Pinterest ($4B thru 2031), Uber (~50%-cheaper claim), Anthropic ($100B agreement incl. Trainium); Trainium sales beyond AWS in negotiation
MICROSOFT	Maia 200	"Already in production" (DCD, Jun 3)	—	Anthropic in talks to run Claude on Maia after $5B Microsoft investment (CNBC, May 21)

OWNER	CHIP	STATUS (REPORTED)	BUILD PARTNER	SIGNAL
OPENAI	"Jalapeño" inference chip	Unveiled Jun 24; ~9-month design cycle; codename per a single feed	Broadcom	Claimed ~50% cost savings; ~$18B program financing reported unresolved (The Information, May)
META	"Iris" MTIA	Production Sept 2026 (Reuters, Jul 9–12)	Broadcom (pact through 2029, >1GW)	Meta targeting 14GW of compute by 2027

Broadcom is the quiet winner of this table: it reportedly builds for Google, OpenAI, and Meta simultaneously, its AI revenue doubled YoY, and its CEO points to **$100B in AI revenue in 2027** (reported). That claim goes on the forecast ledger below.

The Groq Deal — an acquisition wearing a license

NVIDIA paid ~$20B — its largest deal ever — for a *non-exclusive* license to Groq's inference technology, and hired CEO Jonathan Ross plus ~90% of the staff (CNBC[1], Dec 24, 2025). The structure is antitrust-shaped: no merger to block, but the team and the technology move anyway. The husk raised $650M in June to rebuild as an inference neocloud. Expect this template to recur.

Challenger Tier: IPOs, Rounds & Exits (all as reported)

Cerebras IPO (CBRS)

MAY 14

Largest US IPO of 2026

$185/share, $5.55B raised; opened ~$350 (~$95B cap). First earnings (Jun 23): Q1 revenue $193.4M (+92% YoY, a beat) — but negative full-year margin guidance on OpenAI (~$10B) / AWS (~$20B combined) cloud deals. Stock fell 10%, then 22–28%.

Etched

JUN 30

Out of stealth

Working chip, $800M raised, >$1B in contracts, and reportedly in talks at a ~$20B valuation.

SambaNova

JUL 8

$1B Series F

First close at an $11B valuation, led by General Atlantic.

Tenstorrent

JUN 17

Qualcomm acquisition talks

Reported at $8–10B (Reuters).

SOURCES

[1] www.cnbc.com

SECTION 04

Foundry

TSMC is printing records and raising every number it publishes. Everyone else is fighting for the scraps of leading-edge relevance — and Intel's fight is existential.

TSMC: the record machine

Q2 2026 (reported Jul 16)[1]: record revenue of **NT$1,270.38B (US$40.2B, +33.7% YoY)**, gross margin **67.7%**, net income **+77% YoY**. Q3 guided to **$44.6–45.8B**. Full-year USD growth guidance raised to **"slightly more than 40%"** (from >30%); 2026 capex raised to **$60–64B**. HPC is **66% of revenue**; **N2 booked its first commercial revenue**; **N3 is sold out**.

In the US, TSMC added $100B to its commitment — an Arizona pipeline now at **$265B across six fabs**, with AZ Fab 2 (3nm) equipment move-in reported for Q3–Oct 2026 and HVM in 2027 (EE Times, as reported[2]). On the roadmap (Apr 22 symposium[3]): **A16 (1.6nm, backside power) slips to 2027**; A12, A13 and N2U announced; N2+A16 capacity planned to grow at ~70% CAGR through 2028.

Foundry Market Share — Q1 2026

Per TrendForce, as reported via aggregators — treat as indicative, not confirmed. TSMC rose from 70.4% the prior quarter.

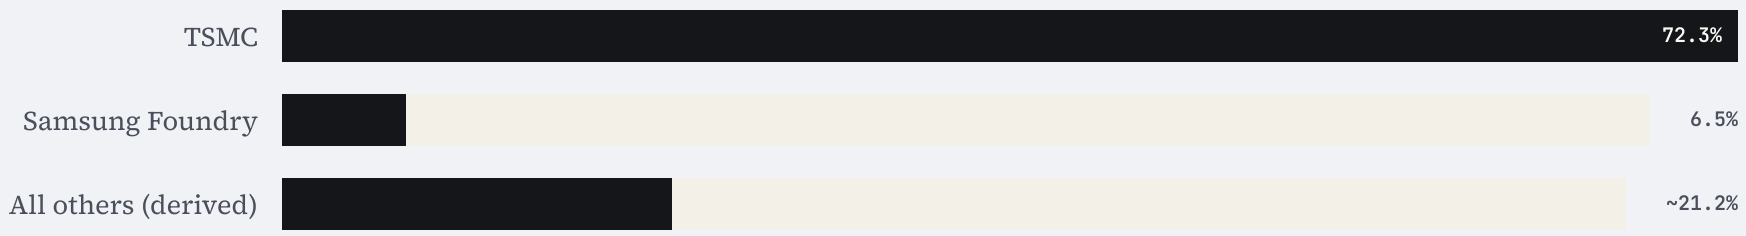

CoWoS: the other monopoly

TSMC advanced-packaging capacity in wafers/month — per Silicon Analysts[4], an aggregator; treat figures as indicative. Reported fully booked against ~1.0M wafers of 2026 demand.

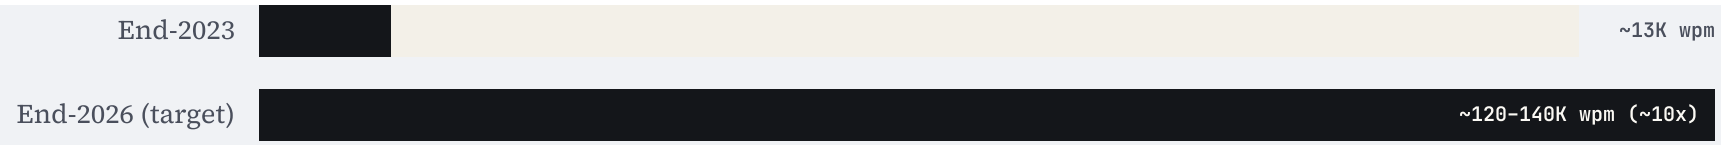

The Chasing Pack

Samsung Foundry — REVIVING

SF2P (2nm) yield reportedly ~70% as of January, up from ~20% in late 2025 (aggregator-sourced — treat as indicative). Expects its first quarterly operating profit since 2022 in Q3 2026.

- Tesla AI5 taped out at Samsung 2nm, dual-sourced with TSMC (Tom's Hardware, Jul 13[5])
- 2nm wafers ~$20K vs TSMC's ~$30K — competing on price

Intel Foundry — AT THE GATE

18A entered high-volume production in January — but with **no named major external customer**. The CFO now says 18A is positioned for external customers (a reversal), and the Q1 filing warns 14A may be **paused** "if we are unable to secure sufficient committed demand" (Intel[6]).

- US government holds a 9.9% passive stake
- 14A go/no-go is the existential foundry gate

Rapidus — PILOT

2nm GAA verified on 300mm wafers; pilot test production began April 2026; PDKs delivered to early customers; mass production targeted ~2027 (Rapidus/NEDO[7]).

SOURCES

[1] investor.tsmc.com/english/quarterly-results

[2] www.eetimes.com/tsmc-boosts-2026-expansion-budget-adds-100b-to-u-s-investment

[3] www.tomshardware.com/tech-industry/semiconductors/tsmc-unveils-process-technology-roadmap-through-2029-a12-a13-n2u-announced-a16-slips-to-2027

[4] siliconanalysts.com/market-data/cowos-capacity

[5] www.tomshardware.com/tech-industry/artificial-intelligence/teslas-ai5-with-2nm-class-node

[6] intc.com

[7] www.rapidus.inc/en/news_topics/information/nedo-approves-rapidus-fy2026-plan-and-budget-for-2nm-semiconductor-projects

SECTION 05

Equipment & EUV

The toolmakers are the purest read on the buildout — and every reading is up and to the right, including the ones flowing to China through side doors.

€9.3B

ASML Q2 2026 net sales at 54.0% GM; net income €2.9B; Q3 guided to €11.0–12.0B

ASML IR · JUL 15

+45%

ASML EUV system sales growth YoY (more than); ~65 low-NA EUV systems targeted for 2026

ASML IR · JUL 15

$5.84B

Lam Research record March quarter, +24% YoY; June quarter guided to ~$6.6B

LAM IR · APR 22

~$19B

Combined AMAT + Lam + KLA China revenue across FY2025 periods — while Chinese fabs import record US equipment via Singapore/Malaysia

TOM'S HARDWARE

ASML[1] is expanding EUV capacity **+30% for 2027** and studying another +30% for 2028. **High-NA EUV reached high-volume logic production** — with Intel as first adopter. One transparency note for the ledger: ASML **no longer discloses quarterly bookings**, removing the sector's favorite forward indicator. Lam[2] raised its 2026 wafer-fab-equipment forecast to **~$140B** (from $135B) — that forecast goes on the watch list below. And the China equipment channel persists: Chinese fabs are importing record volumes of US equipment via Singapore and Malaysia[3] — an evasion channel running in parallel with export controls.

SOURCES

[1] www.asml.com/en/investors/financial-results/q2-2026

[2] investor.lamresearch.com/2026-04-22-Lam-Research-Corporation-Reports-Financial-Results-for-the-Quarter-Ended-March-29,-2026

[3] www.tomshardware.com/tech-industry/chinese-chip-tool-makers-posted-record-2025-revenues-while-margins-slipped

SECTION 06

Market, Pricing & Demand

Compute prices are whipsawing, memory is the famine everyone shares, and the deal sheet reads like sovereign-wealth statements. Most items here arrived via news feeds — treat as reported.

The Memory Famine

+110%

Reported DDR5 price increase as AI devours memory supply

REPORTED · NEWS FEEDS

>40wk

Reported AI-server memory lead times

REPORTED · NEWS FEEDS

$4,300+

RTX 5090 street price vs $1,999 MSRP; vendors re-releasing 2020-era RTX 3060/3050

REPORTED · RTX 50 SUPER SLIPPING TO 2027

2027

SK Hynix CEO: 2027 will be "the worst year" of the shortage — could persist to 2030

REPORTED · JUL 11

GPU Rental Rollercoaster

Per SemiAnalysis data reported via MSN (Jul 19): Blackwell rentals hit **$4.08/hr in mid-May** (a +48% surge), then fell **−30% in three weeks** by June 22 — while H100 rates are still up **~40% net over six months**. B200 rental prices are reported "set to double," with orders pushed to Q2 2027. Read: spot compute is now a genuinely volatile commodity market, and the forward book is sold out anyway.

The Deal Sheet (May–July, as reported)

DATE	DEAL	SIZE	WHY IT MATTERS
Jun 9	Anthropic capacity deal w/ Apollo, Blackstone, Broadcom	$35B	Private credit is now funding compute
Jun 13	AMD–Meta multi-year supply	~$100B	Second source at hyperscale, reported
Jun 10	OpenAI weighing Ohio data center	10GW	Single-site gigawatt-class ambition
Jul 3	SoftBank US compute rental plan	10GW	Compute-as-a-commodity thesis
Jul 7	SK Hynix Nasdaq listing	$28B	Record listing; memory is the new oil

Longer-arc forecasts entering the ledger: Bloomberg Intelligence sees the AI accelerator market passing **$600B by 2033**, and IDC sees total semiconductors passing **$1T** (both reported).

SECTION 07

Geopolitics & China

The export-control regime produced its strangest artifact yet: an approval that ships almost nothing. Both sides are now enforcing scarcity on the same chips.

H200-to-China: approved, but broken

December 2025 brought conditional approval (per-shipment licenses plus a 25% revenue share); April brought easing to ~75K-unit per-firm caps — but Beijing blocked imports through buy-local directives. On July 14, Commerce told Congress that only **"a small number"** of H200s have actually shipped, with Alibaba and ByteDance permitted limited quantities (CNBC, Jul 14[1]; detail per tech-insider.org[2], an aggregator — hedge accordingly). The H20 is discontinued. Net: a legal China channel exists on paper and barely exists in silicon.

Enforcement Hardens on Every Side

Taiwan ran its first formal crackdown on NVIDIA AI-chip smuggling: 12 locations raided, 3 arrests, Japan transshipment routes and falsified documents[3] (May 21–27) — and is weighing a criminal ban on AI-chip exports to all of China (Jun 9). **NVIDIA** slashed its Asian authorized-customer list and deployed field inspectors (Jul 14). **Super Micro** faces a shareholder securities-fraud suit over allegedly concealed illicit China sales (Mar 26). In **Washington**, the AI OVERWATCH Act was secured in the Senate FY27 NDAA (Jul 15)[4], with the Chip Security Act — on-chip location verification — advancing in the same cycle. A Section 232 action has reportedly placed a 25% tariff on a narrow range of advanced computing chips (H200, MI325X) with buildout exemptions — wiki-sourced, treat as unconfirmed. And the CHIPS Act apparatus is being remade: Natcast shut down, quantum grants restructured as federal equity stakes.

China's Domestic Push

SMIC — PRIMARY

Q1 2026 revenue $2.505B at 20.1% GM, ~93% utilization; Q2 guided +14–16% QoQ[5]. Separately, AEI estimates SMIC 5nm DUV yields ~20% (34 patterning steps vs 9 with EUV) — aggregator-sourced, unverified. China reportedly targets ~1.6M advanced AI dies in 2026.

Huawei Ascend — SCALING

910D customer testing reportedly went well; ByteDance and Alibaba are planning orders[6]. Plans: ~100K 910D + ~600K 910C units in 2026 — an output claim graded against a skeptical HBM-capacity counterview in the forecast ledger.

RISC-V — EMERGING

Alibaba's XuanTie C950: a 3.2GHz server-class RISC-V core on TSMC 5nm (Mar 26). China state guidance now explicitly promotes RISC-V — the architectural hedge against both x86 and Arm.

SOURCES

[1] www.cnbc.com

[2] tech-insider.org/nvidia-h200-chip-sales-china-2026

[3] www.tomshardware.com/desktops/servers/taiwan-raids-12-locations-in-its-first-formal-crackdown-on-nvidia-ai-chip-smuggling-hunts-three-fugitives-for-document-forgery-fraudulent-declarations-in-super-micro-smuggling-case

[4] www.wbiw.com/2026/07/15/sen-banks-secures-ai-overwatch-act-in-senate-ndaa

[5] www.smics.com/en/site/news_read/7956

[6] www.cnbc.com/2026/03/27/bytedance-alibaba-planning-to-order-huaweis-new-ai-chip-reuters.html

SECTION 08

Sector Timeline

The running milestone history of the series — regime changes, firsts, and records. Vol. 1 seeds the ledger; every future issue appends.

2025-12-24

NVIDIA–Groq: a ~$20B license instead of an acquisition

NVIDIA's largest deal ever buys a non-exclusive license to Groq's inference tech — plus the CEO and ~90% of staff (CNBC). A new antitrust-shaped M&A template.

2026-01

Intel 18A enters high-volume production

America's leading-edge node ships in volume — but with no named major external foundry customer.

2026-03-18

Huang claims >$1T cumulative bookings through 2027

Blackwell + Vera Rubin bookings/visibility, doubled from the $500B claimed Oct 2025 (reported via Yahoo Finance; not restated at May earnings).

2026-04-22

Google splits the TPU and starts selling it directly

8th-gen TPU divides into training and inference parts (reported); direct TPU sales begin — the first hyperscaler chip offered as a merchant product.

2026-05-14

Cerebras IPO – largest US IPO of 2026

Reportedly $185/share, $5.55B raised, opening near $350 (~$95B cap) — the first pure-play AI-chip challenger at mega-cap scale.

2026-05-20

NVIDIA posts an $81.6B quarter

+85% YoY with data center at $75.2B; guides Q2 to $91.0B with China data-center compute excluded (NVIDIA IR).

2026-05-21

Taiwan's first formal chip-smuggling crackdown

12 locations raided, 3 arrests, Japan transshipment and forged documents (Tom's Hardware) — enforcement goes kinetic at the source.

2026-05-31

NVIDIA declares Vera Rubin in "full production"

Weeks after design-issue reports surfaced and were cleared (reported); all three HBM4 suppliers certified per Bloomberg (Jun 4).

2026-06-24

OpenAI + Broadcom unveil a first custom inference chip

Reported ~9-month design cycle and claimed ~50% cost savings ("Jalapeño" codename per a single feed); ~$18B program financing reported unresolved.

2026-07-11

SK Hynix CEO: 2027 will be the worst year of the memory shortage

Constraints could persist to 2030 (reported). The famine gets an official long horizon.

2026-07-14

NVIDIA slashes its Asian authorized-customer list, adds field inspectors

Vendor-side export enforcement — the same day Commerce told Congress only "a small number" of H200s have shipped to China.

2026-07-15

AI OVERWATCH Act secured in the Senate FY27 NDAA

With the Chip Security Act (on-chip location verification) advancing in the same cycle (WBIW) — chip tracking heads toward statute.

2026-07-16

TSMC record quarter; FY26 guidance raised past 40% growth

US$40.2B revenue at 67.7% GM; capex raised to $60–64B; N2 books first commercial revenue; N3 sold out (TSMC IR).

SECTION 09

Forecasts & Scorecard

The calibration instrument: our falsifiable calls, the sector's open external claims, and grades for claims already resolved. Our own Vol. 1 forecasts get graded starting Vol. 2.

Our Forecasts — Vol. 1

Each claim is concrete, dated, and gradeable from public data on the resolve date.

SOCC-2607-01 — 80%

NVIDIA reports Q2 FY27 revenue of $89.2B or higher (i.e., at or above the bottom of its $91.0B ±2% guide). Resolves by 2026-09-01.

- Basis: guide issued with China already excluded[1]; Q1 beat its own guide

SOCC-2607-02 — 75%

NVIDIA confirms a Rubin volume ramp on its Q3 FY27 earnings call (Rubin revenue or volume shipments explicitly cited), despite the June production-cut reports. Resolves by 2026-12-01.

- Basis: "full production" declared May 31; Huang's July "giant amounts" vow (reported); HBM4 suppliers certified (reported)

SOCC-2607-03 — 85%

TSMC finishes FY2026 with USD revenue growth above 38%. Resolves by 2027-01-31 (Q4 report).

- Basis: guidance raised to "slightly more than 40%"[2]; Q3 guided $44.6–45.8B; N3 sold out

SOCC-2607-04 70%

AMD launches MI400 in volume at "Advancing AI" (Jul 22–23) with at least one named hyperscale customer. Resolves by 2026-08-01.

› Basis: event is 2 days out; MI450 sampling and reported ~$100B Meta deal suggest a customer name is ready

SOCC-2607-05 45%

Intel announces a named major external customer for 18A or 14A by December 31, 2026. Resolves by 2026-12-31.

› Basis: CFO repositioning 18A for external customers[3]; but the 14A pause warning shows demand hasn't committed yet

SOCC-2607-06 70%

DDR5 prices end 2026 above their July 2026 level (per major trackers, e.g. TrendForce/DRAMeXchange) — the memory famine does not break this year. Resolves by 2027-01-15.

› Basis: SK Hynix CEO's "2027 is the worst year" warning (reported); >40-week lead times (reported)

SOCC-2607-07 55%

Samsung Foundry posts its first quarterly operating profit since 2022 in Q3 2026, as the company expects. Resolves by 2026-11-15 (Q3 earnings).

› Basis: reported SF2P yield recovery to ~70%; Tesla AI5 tape-out; price advantage vs TSMC — but yield figures are aggregator-sourced

SOCC-2607-08 40%

By mid-2027, at least one third-party cloud/neocloud publicly offers Google TPUs (a non-Google operator hosting TPU capacity for sale). Resolves by 2027-07-31.

› Basis: Google began direct TPU sales and is reportedly courting NVIDIA neoclouds to host TPUs (Jul 15); ~$25B TPU-cloud vehicle reported

Forecast Watch — External

Open claims by CEOs, agencies, and analysts now on the ledger. Statuses reflect digest evidence as of Jul 20, 2026.

FORECASTER	MADE	CLAIM	RESOLVES BY	STATUS
Jensen Huang (NVIDIA)	2026-03	>$1T cumulative Blackwell+Rubin bookings/visibility through 2027 (reported)	FY28 results	TRENDING HIT
NVIDIA	2026-07	Rubin ships in "giant amounts" in the Q3 2026 ramp	2026-12	OPEN
TSMC	2026-07	FY2026 USD revenue growth "slightly more than 40%"	2027-01	TRENDING HIT
Lam Research	2026-04	2026 wafer-fab equipment spend ~$140B	2027-Q1	OPEN
SK Hynix CEO	2026-07	2027 is the "worst year" of the memory shortage; constraints to 2030 (reported)	2027–2030	OPEN
Huawei	2026-03	~600K 910C + ~100K 910D units shipped in 2026 (reported)	2026-12	OPEN — OPPOSED
SemiAnalysis (counterview)	2026	HBM capacity caps Huawei at ≤300K 910C-equivalents in 2026 (reported)	2026-12	OPEN — OPPOSED
Broadcom CEO	2026	$100B AI revenue in 2027 (reported)	2028-Q1	OPEN
Bloomberg Intelligence	2026	AI accelerator market >$600B by 2033 (reported)	2033	OPEN

FORECASTER	MADE	CLAIM	RESOLVES BY	STATUS
Intel	2026-Q1	14A proceeds only with "sufficient committed demand" — pause decision gate	2026–27	OPEN
Meta	2026-07	"Iris" MTIA in production Sept 2026; 14GW compute by 2027 (reported)	2026-10 / 2027-12	OPEN

Scorecard — Already Resolved

External claims from the coverage window that can already be graded.

FORECASTER	CLAIM	VERDICT	EVIDENCE
Jensen Huang (Oct 2025)	$500B cumulative Blackwell bookings across 2025–26 (reported)	PARTIAL	Superseded — raised to >$1T through 2027 by Mar 2026 (reported); FY26 actual revenue $215.9B[4]. Trending hit, but the original claim was replaced before it could resolve cleanly.
NVIDIA (own guidance)	Q1 FY27 revenue guide of $78B	HIT	Actual: $81.6B[1] — beat its own guidance.
May Rubin design-issue reports	Vera Rubin has design problems threatening the ramp (reported, Wccftech)	MISS	Cleared within weeks: 3Q26 ramp reported (DigiTimes, May 15), "full production" declared May 31, and all three HBM4 suppliers certified (Bloomberg, Jun 4 — reported).

SOURCES

[1] nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-first-quarter-fiscal-2027

[2] investor.tsmc.com/english/quarterly-results

[3] intc.com

[4] nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-fourth-quarter-and-fiscal-2026

THE STATE OF COMPUTER CHIPS

NEXT: VOL. 2 — AUGUST 2026

Primary sources: NVIDIA IR · TSMC IR · ASML IR · Lam Research IR · SMIC · Rapidus · Intel

Reported via: Tom's Hardware · CNBC / Reuters · EE Times · WBIW · Silicon Analysts