

## THE BOTTOM LINE

# Q1 broke the bear thesis. The marginal dollar started earning more — and the network still compounds.

The W17 read was 'the marginal dollar earns less at every layer except the network.' Q1 2026 hyperscaler prints (MSFT / GOOG / META / AMZN, all on Apr 29) reset the math. Aggregate 2026 capex stepped to \$695-725B (+77% YoY) but AI revenue grew faster — Microsoft Cloud commercial RPO doubled YoY to \$627B, Microsoft AI run rate +123% YoY to \$37B, Google Cloud +63%, AWS +28% (a 15-quarter high).

The hyperscaler capex / AI revenue lever compressed from ~5.5 to ~5.0-5.2 — the first Q-over-Q compression in five quarters. The network hypothesis got an audited proof point in the same 48 hours: Equinix Q1 Fabric revenue +26% YoY with bookings +70% YoY at a record 51% adjusted EBITDA margin, raised guidance, ~60% of the largest Q1 deals AI-related.

Custom silicon at ~31% (Trainium's \$225B forward book, Maia 200 going GA in Iowa, MTIA expanded with Broadcom and AMD lanes) is starting to compress merchant GPU pricing power. Capability gating became market structure rather than research-org concern: GPT-5.5-Cyber and Claude Security both shipped this week as gated, separately-priced enterprise products.

The arc the buildout is now on: less hype, more earn-out — boards should reset their AI-ROI thresholds upward, investors should track the capex / AI-revenue ratio (now ~5.0, threshold >6.0) as the wedge it has stopped being, and architects should pull custom-silicon ASIC TCO into procurement bake-offs that previously assumed merchant GPU dominance.

## AUTHOR

**Brian Letort**

## PUBLISHED

**April 30, 2026**


BrianLetort.AI

Issue 01

## HOW WE READ THIS WEEK

# Three lenses, one flywheel.

Cheaper inference pulls in more workloads. More workloads need more compute. More compute needs denser fabric. Denser fabric unlocks new architectures, which lower the cost of inference again. Read a week's news across that loop and the noise sorts itself out.



## THIS WEEK'S ARC

## All three lenses.

**Software → Hardware** new model capability creates new use cases that consume more compute (Jevons).

**Hardware → Networking** more compute means more nodes; the value of the connecting fabric scales as the square of the nodes (Metcalfe).

**Networking → Software** denser, higher-bandwidth interconnect makes new model architectures viable (Gilder).

**Hardware itself** GPU performance doubles faster than Moore's Law (Huang).

### THE BAR

The events that actually matter touch at least two of the three lenses. Single-lens reads are noise dressed up as motion. Each section of this brief grades its evidence and ties the implication back to the flywheel.

# Software.

Model releases, pricing, capability benchmarks, license posture, and capability-risk disclosures.

APR 30

EVENT 01

UK AISI publishes GPT-5.5 cyber evaluation (71.4% Expert pass rate, highest tested; beats Mythos Preview at 68.6% and Opus 4.7 at 48.6%); OpenAI launches GPT-5.5-Cyber as restricted-access, adopting the same vetted-customer gating Sam Altman previously criticized when Anthropic used it for Mythos

AISI.GOV.UK, DEPLOYMENTSAFETY.OPENAI.COM/GPT-5-5/CYBERSECURITY, SIMONWILLISON.NET

APR 30

EVENT 02

Anthropic ships Claude Security in public beta — an Opus 4.7-powered codebase vulnerability scanner with auto-patching, scheduled scans, and audit integrations; embedded with CrowdStrike, Microsoft Security, Palo Alto Networks, SentinelOne, Wiz, plus Accenture, BCG, Deloitte, Infosys, PwC

CLAUDE.COM/BLOG/CLAUDE-SECURITY-PUBLIC-BETA, SECURITYWEEK.COM

APR 29

EVENT 03

Mistral releases Medium 3.5 (open-weight modified MIT, 128B dense, 256K ctx) at 77.6% SWE-Bench Verified and \$1.50/\$7.50 per Mtok; replaces Devstral 2, Magistral, and Medium 3.1 in one consolidated release with Vibe remote cloud coding agents and Le Chat Work mode on the same runtime

MISTRAL.AI/NEWS, THE-DECODER.COM, AWESOMEAGENTS.AI, DECRYPT.CO

APR 29

EVENT 04

IBM releases Granite 4.1 LLM family (Apache 2.0, 3B / 8B / 30B dense decoders, 512K context, ~15T training tokens); 8B Instruct matches the prior 32B-A9B MoE Granite 4.0-H-Small at a fraction of the parameter count

HUGGINGFACE.CO/BLOG/IBM-GRANITE, IBM.COM/GRANITE/DOCS

APR 28

EVENT 05

GPT-5.5 + Codex + Bedrock Managed Agents land natively on AWS Bedrock in limited preview; OpenAI inference now billable against AWS commitments under IAM, PrivateLink, and CloudTrail — the same week Microsoft / OpenAI exclusivity formally ends and OpenAI's revenue-share is capped through 2030

OPENAI.COM/INDEX/OPENAI-ON-AWS, AWS.AMAZON.COM WHAT'S NEW, BLOGS.MICROSOFT.COM

**WHAT THIS MEANS**

Two arcs converged in one week. Software-internal: AISI's GPT-5.5 cyber numbers and the simultaneous launch of OpenAI's gated GPT-5.5-Cyber and Anthropic's Opus-4.7-powered Claude Security make frontier-grade cyber an explicit, separately-priced product line — CISOs and security architects must update agentic threat models this quarter and pick a defender stack now.

Software-distribution: GPT-5.5 going Bedrock-native plus the end of Microsoft exclusivity decouples closed-frontier intelligence from cloud residency for the first time, so boards previously locked in by Azure-only constraints should reopen OpenAI procurement on AWS or GCP within the limited-preview window before pricing hardens at GA. See the Model Pulse for the full architecture read on the six new tree rows that landed this week.

# Hardware.

Silicon, density, packaging, memory supply, and the share of incremental compute going to custom silicon.

APR 29

EVENT 01

**Amazon Q1 2026 8-K: Trainium revenue commitments hit \$225B; chip business at >\$20B annualized run rate (now top-3 data-center silicon globally); Trainium3 nearly fully subscribed, Trainium4 already partly reserved; Anthropic 5 GW + OpenAI ~2 GW commitments confirmed**

AMAZON 8-K (SEC.GOV), AMAZON Q1 PRESS RELEASE, ABOUT AMAZON (JASSY COMMENTARY), BUSINESSWIRE

APR 29

EVENT 02

**Microsoft FY26 Q3 8-K: FY26 capex raised to ~\$190B (+61% YoY); Q3 capex \$31.9B; Maia 200 went GA in US Central (Iowa) with US West 3 (Phoenix) next, claiming '30% better tokens per dollar' vs current fleet silicon; Microsoft cited a ~\$25B memory-cost headwind**

MICROSOFT 8-K (SEC.GOV), MICROSOFT IR FY26 Q3 RELEASE, NEWS.MICROSOFT.COM, CNBC

APR 29

EVENT 03

**Meta Q1 2026 8-K: 2026 capex band raised to \$125-145B (from \$115-135B); management named memory pricing as the primary driver; Zuckerberg confirmed broader rollout of Broadcom-co-developed MTIA silicon plus AMD Instinct GPUs alongside NVIDIA in 2026 fleets**

META 8-K (SEC.GOV), META Q1 PRESS RELEASE, BLOOMBERG, AP

APR 30

EVENT 04

**Samsung Q1 2026: confirmed industry-first HBM4 mass production for NVIDIA Vera Rubin began Feb 2026; HBM4E sampling targeted Q2 2026 (16 Gbps pin, 4.0 TB/s/stack); HBM revenue projected to more than triple YoY in 2026 with HBM4 forecast >50% of HBM mix from Q3**

SAMSUNG Q1 2026 RELEASE (NEWS.SAMSUNG.COM), CNBC, MARKETSCREENER EARNINGS TRANSCRIPT, DIGITIMES

**WHAT THIS MEANS**

Three SEC-disclosed CSP earnings inside a 48-hour window made W17's thesis tangible: custom silicon's share of incremental compute is pushing past the ~30% baseline (Trainium's \$225B forward book + Maia 200 in two regions + Meta's expanded MTIA and AMD lanes), and the binding constraint is now memory on the record — Microsoft's \$25B memory headwind and Meta's capex hike both explicitly cite memory pricing, while Samsung's Apr 30 HBM4 mass-production confirmation for Vera Rubin finally adds a qualified second source against SK Hynix's monopoly. Architects and operators planning H2 2026 capacity should re-bid GPU/ASIC mix assuming Samsung HBM4 is now real (de-risks Vera Rubin allocation); investors should model HBM pricing — not GPU pricing — as the swing variable on hyperscaler ROI through 2026.

# Networking.

Interconnect, fabric standards, optical capacity, sovereign and operator-level networking products.

APR 29

EVENT 01

Equinix Q1 2026 8-K reports Fabric revenue +26% YoY, Fabric bookings +70% YoY, +5,800 net physical and virtual interconnections, record 51% adjusted EBITDA margin, full-year guidance raised to 10-11% revenue growth — ~60% of the largest Q1 deals AI-related

EQUINIX 8-K (SEC EDGAR), Q1 2026 EARNINGS CALL TRANSCRIPT, INSIDER MONKEY TRANSCRIPT

APR 29

EVENT 02

Celestica opens orders for DS6000-series 1.6TbE switches (64 x 1.6TbE OSFP224, 102.4 Tbps non-blocking, Broadcom Tomahawk 6 silicon, SONiC) compliant with both UltraEthernet Consortium and OCP ESUN — the open-standard 1.6T scale-up path moves from spec to procurement-ready SKU at OCP EMEA Summit Barcelona

CELESTICA PRESS RELEASE / GLOBE NEWSWIRE, BROADCOM CORE SWITCHING GROUP COMMENTARY, DELL'ORO GROUP ANALYST QUOTE

APR 30

EVENT 03

Hedgehog contributes OCP-Accepted AI training and AI inference fabric reference architectures (validated with Celestica and FarmGPU) to the OCP Marketplace at OCP EMEA — first vendor-neutral, multi-silicon, Ethernet-based AI fabric blueprints available as production-grade designs

HEDGEHOG PRESS RELEASE / PR NEWSWIRE, OCP MARKETPLACE LISTING, OCP FOUNDATION COMMENTARY

APR 28

EVENT 04

CoreWeave + Google Cloud launch CoreWeave Interconnect (private fiber with line-rate MACsec on Google's Partner Cross-Cloud Interconnect), SUNK Anywhere (Slurm-on-Kubernetes extending to AWS and Azure), and LOTA Cross-Cloud (cache giving 7 GB/s/GPU near-local throughput, no egress fees) — a 44-DC neocloud GPU operator becomes a cross-DC training-fabric provider

COREWEAVE BLOG ANNOUNCEMENT, GOOGLE CLOUD NEXT 2026 KEYNOTE, DCD COVERAGE

**WHAT THIS MEANS**

Equinix's SEC-grade Q1 — Fabric revenue +26% YoY, Fabric bookings +70% YoY, record 51% EBITDA margin, raised 2026 guidance — is the audited proof point for the publication's durable-fabric thesis; the same week, three independent open-standard moves (Celestica UEC/ESUN-compliant 1.6TbE on order, Hedgehog OCP-Accepted AI fabric reference designs, CoreWeave-Google cross-cloud Interconnect) push the open AI-fabric path from spec to procurement-ready SKU. Architects should re-test single-vendor scale-up roadmaps against an open Tomahawk 6 + UEC + ESUN + OCP-Accepted blueprint stack now bookable today; boards and investors got an audited Metcalfe data point that justifies re-underwriting interconnect-category pricing power; operators should treat the CoreWeave-Google move as the week the colo-vs-neocloud connectivity boundary started dissolving.

CAPITAL FLOW

# Money in, revenue out.

Capital deployed (forward) vs revenue out (quarterly or run-rate). Burn-to-revenue = revenue / capital — lower means more out than in. Bars normalized to 58.0 \$B; On-Prem revenue is indirect.

## Frontier Labs

OPENAI, ANTHROPIC, GOOGLE DEEPMIND, XAI

BURN / REV

0.6



~\$52B



~\$35B

Anthropic disclosed ~\$30B ARR (3x in four months) and is fielding offers for a ~\$50B raise at a \$900B pre-money on Apr 29 — frontier labs now reprice as scaling enterprise-software businesses with 8 of Fortune 10 paying. Boards and LPs should underwrite ARR cadence and contract concentration on hyperscaler GPU/Trainium supply, not story or option value.

## Hyperscaler-Hosted

MICROSOFT AZURE, GOOGLE CLOUD, AWS, META

BURN / REV

0.22



~\$58B



~\$13B


All four hyperscalers reset 2026 capex on Apr 29 to a combined \$695-725B (+77% YoY) — MSFT \$190B, GOOG \$180-190B, META \$125-145B, AMZN \$200B — while every CFO repeated 'compute constrained.' Architects should assume zero capacity relief through 2026 and lock multi-year terms now; investors should monitor the capex / AI-revenue ratio (currently ~5.0-5.2, threshold >6.0) since the ratio only stays healthy if AI revenue keeps growing 60-120% YoY.

## Neoclouds

COREWEAVE, CRUSOE, NEBIUS, APPLIED DIGITAL

BURN / REV

0.11



~\$18B



~\$2.0B

CoreWeave's April closed three multi-year anchor commits (Meta +\$21B Apr 9, Anthropic Apr 10, Jane Street \$6B + \$1B equity Apr 15) while Oracle landed a \$16B project-finance package for a 1.4 GW OpenAI campus in Michigan (Apr 25, BofA-led) — neocloud capacity is now structurally pre-sold through 2032. Operators should look past headline backlog to customer concentration (4-5 frontier customers ~80%+) and the project-finance debt structure that will start showing up in covenant disclosures.

## On-Prem / Hybrid

ENTERPRISE GPU CLUSTERS, SOVEREIGN AND NATIONAL PROGRAMS

BURN / REV

n/a



~\$42B



Indirect

Apple posted record R&D \$11.4B (+34% YoY, Apr 30) and called out 'amazing platforms for AI and agentic tools' on Mac Studio and mini, while Amazon's custom-silicon line (Graviton, Trainium, Nitro) crossed a \$20B annual run rate at triple-digit YoY growth — on-prem and sovereign AI buildouts are now first-class procurement tracks, not edge cases. Operators evaluating new AI capacity should run a parallel on-prem / sovereign track alongside hyperscaler RFPs; architects should assume the 36-48mo time-to-power constraint applies equally and pre-secure power siting now.

## SIGNAL VS NOISE

# What's real, what's noise.

6 claims that drove headlines this week, scored 1–5 on source quality and triangulation. 2 flagged as noise. The bar is at least one explicit noise call per issue.

5 / 5

CLAIM 01

**Big-4 hyperscaler 2026 capex resets to a combined \$695-725B (+77% YoY): MSFT \$190B (vs \$154.6B consensus), GOOG \$180-190B, META \$125-145B, AMZN \$200B reaffirmed — all four CFOs repeated 'compute constrained.'**

SOURCES: MICROSOFT FY26 Q3 PRESS RELEASE (APR 29), ALPHABET Q1 2026 8-K (SEC EDGAR, APR 29), META Q1 2026 RELEASE (APR 29), AMAZON Q1 2026 RELEASE (APR 29); ROLLED UP IN BUSINESS INSIDER, TOM'S HARDWARE, REUTERS

Real and load-bearing. The capital-flow bear thesis ('hyperscalers will pull back') is gone for 2026. Read changes only if Q2 prints walk back the guides or if a regulator constrains a major utility / grid project; otherwise expect the \$725B floor to hold and stress to migrate to power siting and supply-chain memory pricing.

5 / 5

CLAIM 02

**Microsoft AI annual run rate is \$37B (+123% YoY) and Microsoft Cloud commercial remaining performance obligation is \$627B (+99% YoY) as of March 31, 2026.**

SOURCES: MICROSOFT FY26 Q3 PRESS RELEASE (APR 29, PRIMARY 8-K DISCLOSURE)

Real. RPO doubling YoY is the cleanest single proof that the AI commercial ramp is durable contract revenue, not pilot spend. Read changes only if Microsoft starts disclosing AI gross margin separately and it deteriorates; the run-rate itself is now an audited number.

5 / 5

CLAIM 03

**Equinix Q1 2026 reports Fabric revenue +26% YoY, Fabric bookings +70% YoY, +5,800 net interconnections, record 51% adjusted EBITDA margin; full-year guidance raised to 10-11% revenue growth.**

SOURCES: EQUINIX 8-K (SEC EDGAR, APR 29), Q1 2026 EARNINGS CALL TRANSCRIPT

Real. The publication's durable-network thesis got an audited Q1 proof point. Combined with Apr 28-30 open-fabric procurement-ready SKUs (Celestica 1.6TbE, Hedgehog OCP-Accepted, CoreWeave Cross-Cloud), interconnect pricing power compounds rather than compresses through 2026.

4 / 5

CLAIM 04

**Anthropic disclosed ~\$30B annualized revenue in April 2026 (3x in four months; some sources cite ~\$40B internal figure), with 1,000+ enterprise customers spending >\$1M/yr and 8 of Fortune 10 as Claude customers.**

SOURCES: CNBC, BLOOMBERG, TECHCRUNCH, REUTERS ALL SOURCING THE DISCLOSURE TO ANTHROPIC IN CONNECTION WITH \$900B RAISE TALKS; SACRA AND AI BUSINESS REVIEW CONFIRM TRAJECTORY; ANTHROPIC IS PRIVATE, NO SEC FILING

Real, but private-company disclosure — not audited. Read changes if the \$50B round prices materially below \$900B, which would imply investor pushback on the revenue durability or margin profile.

2 / 5 -  
NOISE

### **Anthropic in talks to raise ~\$50B at a \$900B pre-money valuation, which would top OpenAI's \$852B March valuation.**

CLAIM 05

SOURCES: CNBC (APR 29), BLOOMBERG, TECHCRUNCH — ALL CITING 'SOURCES'; NO TERM SHEET SIGNED; BOARD DECISION EXPECTED MAY

Real that the talks are happening, but the price is not yet committed capital. Treat as sentiment indicator, not balance-sheet event. Read changes when the round prices, anchor lead investor is named publicly, or a major LP walks away — any of which would re-rate the entire frontier-lab valuation curve.

---

2 / 5 -  
NOISE

### **OpenAI's GPT-6 / 'Spud' is launching imminently (late April through May 25, 2026) with a 2M-token context and 40% performance gains over GPT-5.4.**

CLAIM 06

SOURCES: VARIOUS TECH BLOGS (REMIO, FINDSKILL, BIGGO, VGTIMES); POLYMARKET PROBABILITY COLLAPSED FROM 93% BY JUN 30 TO 45%; OPENAI HAS CONFIRMED PRE-TRAINING COMPLETED MAR 24 BUT NOT ANNOUNCED A LAUNCH DATE OR FINAL NAMING

Hype-adjacent. The completed pre-training is real (Altman, Brockman public statements); the launch date and capability claims are speculation. Architects should not pre-commit deployment timelines or budget shifts on this. Read changes only on an OpenAI announcement of GA timing and pricing.

---

## EARLY WARNING PANEL

# The levers we monitor.

10 metrics, current vs prior period. **3 rising**, **2 falling**, 5 steady. Each metric carries a threshold value where the read materially changes.

METRIC	CURRENT	PRIOR	DIR	THRESHOLD / NOTE
Frontier lab cash position (avg months runway, top 3)	~ <b>30 mo</b>	~22 mo	↑	<18 mo triggers re-rating risk
Hyperscaler capex / AI revenue ratio (top 4 weighted)	~ <b>5.0-5.2</b>	~5.5	↓	>6.0 invites investor pushback at next earnings
CoreWeave revenue backlog	~ <b>\$87B+</b> (est)	\$66.8B	↑	Conversion velocity matters more than gross figure
NVIDIA Q-over-Q data center revenue	<b>\$62.3B</b> (Q4 FY26)	\$62.3B (Q4 FY26)	→	—
Open vs closed gap on SWE-Bench Pro (coding)	<b>Closed +6 to +19pp</b>	Open +0.9	↓	Sustained open lead reshapes enterprise procurement
Sovereign AI commitments (count / aggregate \$)	<b>8 / ~\$80B+</b>	8 / \$80B+	→	—
PJM 2026/27 capacity auction price (\$/MW-day)	<b>\$329.17</b>	\$329	→	11x in 24 months — power is the new binding constraint
Time-to-power, busiest US markets (months)	<b>36-48</b> (PJM large >100 MW); ~53 nat'l avg	36-48	→	—
Cost-per-task, frontier reasoning model	~ <b>\$0.05</b>	~\$0.05	→	Falling cost expands workloads (Jevons), not contracts demand
Custom silicon share of incremental AI compute	~ <b>31%</b>	~30%	↑	>35% materially compresses merchant GPU pricing

THRESHOLD VALUES ARE THE POINTS WHERE THE READ FLIPS. CROSSINGS ARE FLAGGED IN THE ISSUE BODY WHEN THEY HAPPEN.

## PREDICTIONS

# What we expect next.

5 falsifiable, time-bounded predictions. Each carries a confidence (1–99%, never 50%), a deadline, and a specific signal we'll watch. Future issues score them hit / miss / partial.

## PREDICTION 01 SOFTWARE

**AWS reports a customer-disclosed Bedrock-resident GPT-5.5 multi-year contract greater than \$500M annualized, with the customer named publicly.**

TRIGGER: AWS customer-win announcements during the Bedrock GPT-5.5 limited preview; OpenAI partnership commentary in Q2 prints.

65%



## PREDICTION 02 CAPITAL

**The top-4 hyperscaler capex / AI revenue ratio drops below 5.0 for Q2 2026, confirming the W18 directional flip.**

TRIGGER: Q2 2026 hyperscaler earnings releases (late July through mid-August).

60%



## PREDICTION 03 HARDWARE

**Samsung HBM4 reaches greater than 25% share of NVIDIA Vera Rubin BOM by Q3 2026 supply data, ending the SK Hynix monopoly on the platform.**

TRIGGER: TrendForce / SemiAnalysis Q3 BOM data; NVIDIA Q2 FY27 earnings color on memory supply.

60%



## PREDICTION 04 NETWORKING

**At least one major colocation operator other than Equinix reports Q2 2026 interconnect / fabric revenue growth greater than 25% YoY.**

TRIGGER: Q2 earnings prints from interconnect operators; Light Reading and DCD coverage.

65%



## PREDICTION 05 SOFTWARE

**By Q3 2026, at least two frontier labs publicly disclose a separately-priced 'gated cyber' or 'gated security' SKU with revenue commentary, formalizing capability-gating as a product line.**

TRIGGER: Q3 earnings disclosures (OpenAI partnership color, Anthropic announcements); vendor security press; AISI / METR cross-vendor evaluations.

60%



## WATCHLIST

# On the radar.

5 catalysts in the next 7–14 days that would change the read materially. Watching these tells us whether the thesis is strengthening or weakening.

MAY 5 - 20

WATCH 01

## Mythos status update from Anthropic + AISI

AISI's Apr 27-28 sabotage evaluation flagged a 65% reasoning-output discrepancy in Mythos continuation tests. Anthropic's response — extend gating, lift gating under conditions, or expand Project Glasswing — sets the precedent for capability-gated SKUs as a market category.

LATE MAY 2026

WATCH 02

## NVIDIA Q1 FY27 earnings

First disclosure on Vera Rubin ramp velocity, Samsung HBM4 second-source effect, and custom-silicon competitive pressure (Trainium \$225B book, Maia 200 GA). A beat against the \$62.3B Q4 FY26 baseline resets the H2 hyperscaler capex tail; a miss flattens Huang's slope.

MAY 5 - 20

WATCH 03

## Anthropic \$900B raise terms and lead investor

Talks reported Apr 29 are not yet committed capital. The price (\$900B vs OpenAI's \$852B) and the lead investor's identity re-rate the entire frontier-lab valuation curve. A walk-away from the price would be the loudest re-rating signal in a year.

MAY 5 - 15

WATCH 04

## Q1 capex revision wave (Apple, Oracle, second-tier hyperscalers)

The big-4 reset 2026 capex to \$695-725B; secondary cohort (Apple, Oracle, IBM Cloud, smaller hyperscalers) has not yet revised. A second wave of capex hikes pushes the aggregate above \$750B and tests prediction p3 to hit before Q2 prints.

MAY 2 - 9

WATCH 05

## Google DeepMind: Gemini 3.1 Pro to GA, or Gemini 3.2 announcement

Gemini 3.1 Pro Preview has held its AA Index 57 spot since February while OpenAI and Anthropic shipped new flagships. A GA or 3.2 announcement would force a frontier scorecard refresh and potentially retake LMArena top spot from Opus 4.7 Thinking.

## METHODOLOGY AND AUTHORSHIP

# How this brief is built.

Compiled from public announcements, SEC filings, earnings transcripts, and official lab and vendor publications. Every quantitative claim is graded on source quality. Every prediction is falsifiable, time-bounded, and scored hit / miss / partial / pending in future issues.

**SIGNAL SCORE RUBRIC**

- 5 / 5** SEC filing or audited disclosure. Multi-source independent confirmation. Operational, not aspirational.
- 4 / 5** Earnings-call disclosure or primary lab/vendor announcement. Two or more independent sources.
- 3 / 5** Single primary source. Reasonably consistent with sector data.
- 2 / 5** Analyst report or secondary press only. Or single primary source with credibility caveats.
- 1 / 5** Rumor, social media, single non-primary source. Or contradicted by alternate primary sources.

ANYTHING GRADED 2 OR BELOW IS FLAGGED AS NOISE.

**PREDICTION OUTCOMES**

- HIT** The specific observable outcome occurred by the deadline.
- MISS** The deadline passed and the outcome did not occur.
- PARTIAL** A meaningfully similar outcome occurred but not the literal wording.
- PENDING** Deadline not yet reached.

HIT RATE OF 65-75% IS THE TARGET. ABOVE 80% MEANS TOO CONSERVATIVE; BELOW 50% MEANS THE FRAMEWORK IS WRONG.

**AUTHORSHIP****Brian Letort**

BrianLetort.AI · Independent analysis. All sources public. Not investment guidance.

**THIS ISSUE****Issue 02**

Week 18 of 2026 · April 30, 2026

**WEB**

[brianletort.ai/industry](https://brianletort.ai/industry)

Past issues, the working framework, and the LLM Evolutionary Tree companion.