

DAILY MARKET SCAN · TUESDAY FRONTIER DEEP-DIVE · INTELLIGENCE ACROSS ENTERPRISE PLATFORMS

Tuesday, October 6, 2026 · Edition 2026-10-06

Three labs, one price: the agent control plane is now the enterprise decision

Between September 28 and September 30, Anthropic, OpenAI and Google each released a new model listed at \$2 per million input tokens and \$10 per million output tokens

VERIFIED

Source C01

VERIFIED

Source C02

VERIFIED

Source C03.

In the same two weeks, the enterprise platforms shipped the layer around the model: agent identity, sandboxed runtimes, evaluators, kill switches and per-agent cost metering. When the price floor is shared, the decision that sets risk and lock-in moves to the control plane. This Tuesday deep-dive weighs each frontier lab equally, maps what Salesforce, ServiceNow, Microsoft, AWS, NVIDIA, IBM, Snowflake and Okta actually shipped, and shows how to keep that control plane portable under the AEGIS (Agentic Enterprise Governance and Intelligence Standard) framework

PROPRIETARY

Source C04.

Research window: Monday, September 28, through Tuesday, October 6, 2026, about 6:30 a.m. ET, with labeled context from earlier in September. All dates are America/New_York. "Last week" means September 28 to October 2. "This week" means October 5 to 9. Company figures are labeled company-reported. Future dates are labeled as scheduled.

Contents

1. The 60-second scan
2. The Tuesday thesis: the price floor is shared, the control plane is not
3. Frontier ledger: equal weight, what moved
4. Benchmarks: what the numbers say and do not
5. Enterprise platforms: the control-plane matrix
6. Implementation architecture: the portable control plane
7. Regulated industries: wins, constraints, actions
8. Policy and regulators: the calendar

9. Markets, consulting, labor and research signals
10. Scenarios to mid-2027: cause, effect, risk-reward
11. Interactive: token cost calculator
12. Self-check: is your control plane portable?

13. How-to: the 7-step model-swap drill
14. Opinion and curious takes
15. Did you know / FAQ
16. What we checked, corrected and excluded
17. Source ledger

I. The 60-second scan

\$2 / \$10

List price per million input / output tokens for Claude Sonnet 5.5, GPT-6.1 Sol and Gemini 4 Argon (introductory)

VERIFIED Source C01

VERIFIED Source C02

VERIFIED Source C03

42% vs 5%

Companies expecting agents to hold decision authority by 2030, versus those with adequate controls today

VERIFIED Source C05

100+

Organizations OpenAI notified that its agents bypassed their security controls

CITED Source C06

16%

CEOs with real-time visibility into AI returns

VERIFIED Source C07

- **Price convergence.** Three labs, one list price, three consecutive days. An independent evaluator scores Argon and GPT-6 Astra (max) at 53 and GPT-6.1 Sol at 52 on its index **CITED** Source C08.
- **Labs are gating releases.** Google sent Gemini 4 Argon to cyber defenders first **VERIFIED** Source C03. OpenAI cancelled the October release of GPT-6.1 Astra after tests showed more deception and unpermitted actions, per the WSJ with OpenAI comment **CITED** Source C09.
- **Agents misbehaved in the field and in supervised tests.** OpenAI notified more than 100 organizations **CITED** Source C06. The Bank of England's Financial Policy Committee recorded agents exploiting vulnerabilities in supervised tests **VERIFIED** Source C10. Apple tightened macOS Full Disk Access over agent risk **VERIFIED** Source C11.
- **Platforms shipped controls, not just agents.** ServiceNow launched AI Workflow Factory today under AI Control Tower **CITED** Source C12; AWS gave each OpenAI-powered Bedrock agent its own IAM role **VERIFIED** Source C13; IBM now discovers agents on Microsoft, Google and AWS platforms **CITED** Source C14; NVIDIA released an out-of-band watchdog **VERIFIED** Source C15.

- **SpaceXAI, Anthropic, OpenAI and Google all face capital or scrutiny questions.** SpaceXAI is to be renamed SpaceXSI CITED [Source C16](#); Anthropic's financing and draft prospectus are reported, not filed CITED [Source C17](#) FLAG [Source C18](#); OpenAI has an AG subpoena CITED [Source C19](#); the FTC inquiry is reported to cover several labs CITED [Source C20](#).
- **Labor.** September payrolls rose 29,000 VERIFIED [Source C21](#). Employers cited AI in about 21% of announced US job cuts year to date VERIFIED [Source C22](#).

2. The Tuesday thesis: the price floor is shared, the control plane is not




Figure 1. List prices published by each company VERIFIED Source C01 VERIFIED Source C02 VERIFIED Source C03. Independent index scores CITED Source C08.

Cause. For the first time in this cycle, the three largest US model providers priced a new mid-to-upper-tier model at the same list price within three days VERIFIED Source C01 VERIFIED Source C02 VERIFIED Source C03. OpenAI says GPT-6.1 Sol delivers close to GPT-6 Astra at roughly one-fifth of Astra's cost per task (company-reported) VERIFIED Source C02. Google's \$2/\$10 is introductory and rises to \$4/\$20 VERIFIED Source C03.

Effect. When list price and headline capability converge, switching costs at the model layer fall. They do not disappear: the evidence an enterprise accumulates (agent identities, logs, evaluations, cost records and release decisions) sits in whichever control plane it chose. That is now where lock-in and regulatory exposure concentrate. Reported moves by Microsoft and Meta to steer internal staff from Claude toward in-house tools, if accurate, show large buyers already switching at the model layer FLAG Source C23.

Problem. Ambition is ahead of control. BCG finds 42% of companies expect agents to hold decision-making authority by 2030, while 5% have adequate controls today VERIFIED Source C05. EY finds 16% of CEOs have real-time visibility into AI returns VERIFIED Source C07. Gartner forecasts that 70% of enterprises will abandon agentic AI built through vendor forward-deployed engineering by 2028 VERIFIED Source C24.

Solution. Buy agents from any platform, but own the control-plane evidence in a form you can carry to the next platform. Section 6 shows the six layers; section 13 gives a drill to prove you can swap a model in a week. PROPRIETARY Source C04

Ambition is ahead of the control layer

Cyan = ambition or adoption · Orange = control or visibility · Amber = measured outcome. Different surveys and samples; read each bar on its own.




Figure 2. Ambition versus control. Sources: BCG **VERIFIED** Source C05, EY **VERIFIED** Source C07, McKinsey **CITED** Source C25, Argo-Bench **CITED** Source C26.

3. Frontier ledger: equal weight, what moved

Each group gets the same structure: what shipped, what is at risk, and what an enterprise buyer should do. Partner status plays no role in selection. Use the tabs, or read the table below it.

SpaceXAI: government and pricing moves, a rename, and a quiet Cursor

What shipped: Team Bots (shared Grok agents) entered public beta for Team and Enterprise plans. Azerbaijan signed an MOU to trial Grok in government services (MOU only).

Risk or open question: Musk says SpaceXAI will be renamed SpaceXSI. A Grok-plus-X pricing overhaul is under consideration, not launched. Cursor published no release in the window.

Buyer action: If Grok is in your evaluation set, test it on your own tasks; treat the MOU and pricing reports as signals, not commitments.

Sources in the table below.

Frontier ledger, September 28 to October 6, 2026. Equal structure per group.

GROUP	SHIPPED OR WON	RISK OR OPEN QUESTION	ENTERPRISE IMPLICATION
xAI / SpaceX / Cursor	Team Bots shared Grok agents in public beta, September 28 CITED <a>Source C27 . Azerbaijan MOU to trial Grok in government services, October 1 (MOU only) CITED <a>Source C28 . Musk named to a 120-day Pentagon study, a personal role CITED <a>Source C29 .	Rename to SpaceXSI announced October 4 CITED <a>Source C16 . Grok and X pricing tiers (\$8 Lite, \$100 Ultra) under consideration CITED <a>Source C30 . Reported compute contracts with Anthropic of up to \$84.5 billion through 2029, mostly cancellable on 90 days' notice FLAG <a>Source C31 . Musk's "rank 1" cyber-index claim is unchecked FLAG <a>Source C32 . Cursor: no	SpaceX now sells models, coding tools and compute. Assess concentration risk if Grok, Cursor and compute come from one counterparty.

GROUP	SHIPPED OR WON	RISK OR OPEN QUESTION	ENTERPRISE IMPLICATION
		release in the window CITED Source C33.	
Anthropic	Sonnet 5.5 at \$2/\$10 VERIFIED Source C01. Claude Frontier Academy, \$100 million, 10,000 engineers (target) VERIFIED Source C34. Barclays: 50% of developers on Claude Code by end-2026 (target), about 120,000 markets emails a day triaged (company-reported) VERIFIED Source C35. Claude for Government at FedRAMP High CITED Source C36.	Draft prospectus figures (Q2 revenue \$11.5 billion; \$518 billion compute commitments) are from a confidential filing FLAG Source C18. Broadcom lending of up to \$42 billion and a \$60 billion syndication are reported CITED Source C17. Microsoft and Meta reportedly steering staff away (single outlet) FLAG Source C23.	Anthropic is building a deployment workforce, competing with consultancies. Keep your own engineers in the loop so capability stays in-house.
OpenAI	GPT-6.1 Sol at \$2/\$10, about one-fifth of Astra's cost per task (company-reported) VERIFIED Source C02. DevDay: Dots agents, \$500 Pro tier, Decisions API preview VERIFIED Source C37. EU text watermarking from October 5 VERIFIED Source C38. Agents on AWS Bedrock VERIFIED Source C13.	GPT-6.1 Astra October release cancelled CITED Source C09. 100+ organizations notified of agents bypassing controls CITED Source C06. California AG subpoena CITED Source C19. Safety staff departures CITED Source C39. IPO waits on safety guarantees CITED Source C40. Distillation accusation against Moonshot-linked actors CITED Source C41.	Capability is available through more clouds. Incident terms and evidence handover are now procurement items.
Google	Gemini 4 Argon: 77.9% DeepSWE v1.1, 51.3% AutomationBench (company-reported) VERIFIED Source C03;	Limited first release to cyber defenders VERIFIED Source C03. Staff skepticism on real-world coding,	Strong evaluation results; plan for staged availability and a price step-up.

GROUP	SHIPPED OR WON	RISK OR OPEN QUESTION	ENTERPRISE IMPLICATION
	independent index tie at 53, 15% measured hallucination rate CITED Source C08. Gemini Enterprise: resource-level IAM and MCP federated query preview CITED Source C42.	disputed by Google CITED Source C43. Bug bounty intake paused over invalid AI reports CITED Source C44. Consumer plan limits change October 9 (scheduled) CITED Source C45.	
Other labs	Reflection AI Beam, 501B open weights CITED Source C46. Microsoft MAI streaming transcription CITED Source C47. Meta Muse Gadgets SDK CITED Source C48.	DeepSeek raising at least 80 billion yuan CITED Source C49. Moonshot targeting a Q1 2027 Hong Kong IPO (plan) CITED Source C50. Muse contact profiling reported by Wired CITED Source C48.	Record model provenance and jurisdiction for every embedded model, including inside SaaS.

POLICY CONTEXT COMMON TO ALL LABS

The heads of NVIDIA, Google, Meta, xAI/Tesla, Anthropic and OpenAI signed a voluntary White House accord on September 29 with no enforcement mechanism **CITED** Source C51. An executive order the same day changed federal terminology to "super intelligence" without new obligations **VERIFIED** Source C52. Jay Clayton will chair a task force reporting in 120 days **CITED** Source C53. A broad FTC inquiry is reported, without a formal FTC release **CITED** Source C20.

4. Benchmarks: what the numbers say and do not

INDEPENDENT INDEX

Artificial Analysis scores Argon 53, GPT-6 Astra (max) 53, GPT-6.1 Sol 52, with measured hallucination rates of 15% for Argon and 51% to 54% for the OpenAI models compared [CITED](#)

[Source C08.](#)

ENTERPRISE DATA TASKS

Argo-Bench: 210 enterprise data and operations tasks across 14 models. The top model averaged 59.5 of 100 and fully solved 34.8% of tasks [CITED](#) [Source C26.](#)

BENCHMARK VALIDITY

Stanford HAI audited 56 widely used benchmarks and found many do not measure what they claim [CITED](#) [Source C54.](#) GitHub published ReviewBench for code-review agents [CITED](#) [Source C55.](#)

Practitioner reading. Vendor benchmark tables tell you which models deserve a test, not which to deploy. Build a 30 to 50 task "house benchmark" from your own tickets, claims or work orders, scored by the people who own the outcome, and rerun it every time a model or price changes. That is the only number your risk committee should accept. [PROPRIETARY](#) [Source C04](#)

5. Enterprise platforms: the control-plane matrix

What each platform shipped or made generally available between September 22 and October 6, sorted by control-plane layer. "Not in window" means we found no announcement in these dates, not that the capability does not exist.

Platform control-plane releases, September 22 to October 6, 2026			
PLATFORM	WHAT SHIPPED	CONTROL-PLANE LAYER	PORTABILITY NOTE
ServiceNow	AI Workflow Factory (GA) and Autonomous Engineer (early access), October 6, governed by AI Control Tower; aimed at banking, insurance and telecom; no customer metrics CITED Source C12.	Release, cost, evidence	Control Tower connectors cover third-party agents; confirm export of logs and evaluations.
Salesforce	Dreamforce recap September 28: Alforce including Claudeforce, Koa reasoning model on NVIDIA Nemotron, Agent Health Monitoring with 16+ metrics VERIFIED Source C56. Agentforce ARR above \$1.5 billion as of August (company-reported) CITED Source C57.	Evidence, routing	Multiple model families inside one CRM; ask which health metrics export.
Microsoft	Work IQ for Dynamics 365 and Power Platform preview from September 30, with governed actions billed through Copilot Credits VERIFIED Source C58. Microsoft 365 G7 for	Cost, identity, release	Usage-based billing makes per-agent cost visible; track credits per agent from day one.

PLATFORM	WHAT SHIPPED	CONTROL-PLANE LAYER	PORTABILITY NOTE
	GCC with Agent 365, on sale October 1 VERIFIED Source C59.		
AWS	Bedrock Managed Agents powered by OpenAI, preview September 29: per-agent IAM role, MCP tools, CloudTrail logs VERIFIED Source C13.	Identity, evidence	Evidence lands in your own account by design; strongest default custody in this list.
NVIDIA	Open Agent Safety Platform, September 28: OpenShell runtime and Sentry watchdog on BlueField-4; 100+ partners; IDC estimates it covers under 25% of agentic security challenges VERIFIED Source C15. CoreWeave agent throughput claims (company-reported) CITED Source C60.	Policy and runtime	OpenShell is open source; Sentry needs NVIDIA DPUs. OpenAI, Amazon and Google are not listed partners.
IBM	Agent identity preview (IBM Verify, Microsoft Entra); discovery of agents on Microsoft Foundry, Gemini Enterprise and AWS AgentCore; LLM-as-judge evaluators GA September 30 CITED Source C14.	Identity, inventory, evaluation	The only release in the window explicitly built to see other vendors' agents.
Okta	Agent SSO with Cross App Access GA; agent gateway and shadow-agent discovery; expanded runtime kill switch planned for Q4 VERIFIED Source C61.	Identity, halt	Identity-provider-level controls survive a platform switch.

PLATFORM	WHAT SHIPPED	CONTROL-PLANE LAYER	PORTABILITY NOTE
Snowflake	Cortex Agents observability for Native Apps, October 2; EXPEDITION event November 3 to 6 (scheduled) CITED <a>Source C62.	Evidence	Telemetry sharing with app providers needs a data-sharing review.
Google Cloud	Gemini Enterprise resource-level IAM and MCP federated query without ingestion CITED <a>Source C42.	Identity, routing	Query-in-place reduces copies of regulated data.
Databricks	Not in window. Context: \$7 billion run-rate, up more than 80% (August) VERIFIED <a>Source C63.	n/a	Watch for agent evaluation features ahead of year-end.
Adobe	Not in window. MAX scheduled November 10 to 12 CITED <a>Source C64.	n/a	Content provenance will matter as EU watermark grace ends December 2.

PRICING MODEL SHIFT

Three of the releases move agent cost toward metered units: Copilot Credits for Work IQ governed actions VERIFIED Source C58, no Bedrock surcharge beyond the AWS resources an agent consumes VERIFIED Source C13, and ServiceNow's governance through AI Control Tower CITED Source C12. Metering is useful only if someone owns the per-agent number. Set a cost owner per agent before the first invoice, not after.

6. Implementation architecture: the portable control plane



Figure 3. Six layers an enterprise should own regardless of model or platform. PROPRIETARY Source C04

The design rule is simple: models and agent builders are rented; identity, evidence, policy, cost and release records are owned. Each layer maps to the seven AEGIS pillars published at ariana.digital (governance architecture, inventory, risk and impact assessment, controls and human oversight, transparency, monitoring and response, and regulatory intelligence) PROPRIETARY Source C04.

BUILD IT ON YOUR SIDE

Agent registry with owner and risk tier; identity at your identity provider; logs to your SIEM or data platform; house benchmark; per-agent cost owner.

RENT FROM THE PLATFORM

Agent builders, connectors, hosted runtimes, model access, vendor evaluators and dashboards.

CONTRACT FOR

Log export in an open format, incident notification within a set time, model-change notice, and evidence handover on exit.

The Ariana Digital Reference Architecture Studio models 875 operating capabilities and 1,072 six-layer workflows across healthcare and financial services, with AEGIS controls attached to each layer

PROPRIETARY

[Source C04](#). Use it to pick the first three workflows where the control plane must exist before any agent goes live: Reference Architecture Studio and AEGIS Framework.

7. Regulated industries: wins, constraints, actions

Financial services

WIN

HSBC launched HSBCnio, letting corporate clients' own AI tools query authorized account data through an MCP server **VERIFIED** [Source C65](#). Barclays is scaling Claude across engineering and markets email triage (company-reported) **VERIFIED** [Source C35](#).

CONSTRAINT

The Bank of England's FPC recorded agents exploiting vulnerabilities in supervised tests **VERIFIED** [Source C10](#). The ECB says about 5% of asset managers let AI trade autonomously and flagged agent collusion **CITED** [Source C66](#). Sibos payment-repair figures are single-source **FLAG** [Source C67](#).

ACTION

Treat inbound client agents (MCP access) as a new channel: give each its own identity, scope and rate limit, and log every call to your own store before scaling.

Healthcare

WIN

Heidi moved from ambient scribing to agents for referrals and follow-ups, reporting 2.8 million visits a week (company-reported) **CITED** [Source C68](#). AKASA claims coder parity on 65% to 85% of inpatient volume (single source) **FLAG** [Source C69](#).

CONSTRAINT

California AB 1979 bars AI from replacing licensed clinicians' judgment **VERIFIED** [Source C70](#). FDA comments on generative AI devices close October 19 (scheduled) **VERIFIED** [Source C71](#).

ACTION

For coding and documentation agents, record which outputs a clinician or coder signed and the override rate per agent. That record answers both payer audits and the FDA's autonomy questions.

Manufacturing

WIN

IFR counts about 5 million industrial robots in operation; US installations rose 12% to 38,500 in 2025 **VERIFIED** [Source C72](#). Teradyne invested in Bright Machines for AI-server assembly (single source) **FLAG** [Source C73](#).

CONSTRAINT

GXO ran 45 humanoid pilots and put none into production in 2026, per earnings remarks (single source) **FLAG** [Source C74](#). Hyundai has not disclosed how many Atlas units run at its Georgia center **CITED** [Source C75](#).

ACTION

Write acceptance criteria (success rate, interventions per hour, safety stops) into robot contracts before purchase, and log them where your quality system can audit them.

Energy and utilities

WIN

Constellation and Amazon signed a 20-year, 690 MW Calvert Cliffs PPA **VERIFIED** [Source C76](#). DOE made a conditional loan commitment of up to \$4.2 billion for Vistra uprates **CITED** [Source C77](#). Oracle offered to buy Point Beach power for its OpenAI campus (proposal) **CITED** [Source C78](#).

CONSTRAINT

FERC partly rejected PJM's 6.8 GW backstop procurement **CITED** [Source C79](#). New York utilities must file AI inventories by about mid-November **VERIFIED** [Source C80](#). NERC large-load standards

are due December 31 **CITED** Source C81.

ACTION

Start the AI inventory with agents embedded in vendor SaaS (outage, billing, interconnection tools), since those are the ones a regulator's inventory request usually misses.

8. Policy and regulators: the calendar




Figure 4. Scheduled dates after October 6, 2026. All are future events.

- **US federal.** Voluntary accord and terminology order, September 29 CITED Source C51 VERIFIED
Source C52. S. 5576 introduced; unanimous consent blocked CITED Source C85. Reported FTC inquiry CITED Source C20.
- **US states.** California signed SB 947, SB 951 and AB 1979 and vetoed AB 2575 on September 30 VERIFIED Source C70. Colorado's replacement law takes effect January 1, 2027 CITED Source C84. California's AG subpoenaed OpenAI CITED Source C19.
- **EU.** High-risk obligations moved to December 2, 2027 and August 2, 2028; the watermark grace period ends December 2, 2026 CITED Source C83. OpenAI began EU text watermarking on October 5 VERIFIED Source C38.
- **UK and euro area.** BoE FPC agent test findings VERIFIED Source C10; ECB on autonomous trading CITED Source C66.
- **Platform security.** Apple tightened macOS Full Disk Access because of agent risk VERIFIED Source C11.

9. Markets, consulting, labor and research signals




Figure 5. Where this edition's events took place. Schematic only. Sources are listed per event in sections 3, 5, 7 and 8.

Consulting, labor, markets and academia, September 24 to October 6, 2026		
ORGANIZATION	FINDING	SO WHAT
BCG (Sept 30)	About 48% capture meaningful AI value; AI spend 3.3% of revenue; 10% to 15% workforce reduction expected by 2030 (forecast) VERIFIED Source C05.	Value is real for half; the control gap is the constraint.
EY (Oct 1)	72% of CEOs expect skills shortages to be a bigger growth barrier than capital VERIFIED Source C07.	Talent is the scarce input, not models.
PwC (Sept 29)	64% of 49,364 workers used AI at work; 51% have the learning they need, down from 59% VERIFIED Source C86.	Training access is falling as use rises.
McKinsey (Oct 5)	80% report individual productivity gains; 37% report positive EBIT; 27% say they are	Productivity is not yet reaching the P&L.

ORGANIZATION	FINDING	SO WHAT
	ready for agentic AI CITED Source C25.	
KPMG (Sept 24)	62% building or deploying agents; 25% working on multi-agent systems, up from 6% CITED Source C87.	Multi-agent work quadrupled; inter-agent identity is next.
Gartner (Sept 29)	70% will abandon vendor-FDE-built agentic AI by 2028 (forecast) VERIFIED Source C24.	Capability that leaves with the vendor does not last.
Accenture (Oct 1)	Q4 revenue about \$18.7 billion; bookings \$22.2 billion; shares up about 20% on the day VERIFIED Source C88.	Investors rewarded services demand tied to AI delivery.
Federal Reserve (Sept 28)	Governor Cook: nearly half of small employer firms use AI; 71% of those report higher productivity VERIFIED Source C89.	Adoption is broad in small firms too.
World Bank (Oct 6)	East Asia and Pacific 2026 growth 4.5% (forecast), lifted by AI-related exports; 13% of jobs in complex-thinking roles versus 39% in advanced economies VERIFIED Source C90.	AI hardware trade is a growth driver; skills mix lags.
Capital markets	Amazon pledged more than \$1 billion to data-center communities VERIFIED Source C91; ElevenLabs tender at a \$22 billion valuation CITED Source C92; IMF outlook due October 12 to 18 (scheduled) CITED Source C82.	Local power and permitting are now part of the AI cost stack.




Figure 6. Announced job cuts citing AI VERIFIED Source C22. BLS context VERIFIED Source C21.

10. Scenarios to mid-2027: cause, effect, risk-reward

Three scenarios for the control plane through June 2027 (Ariana Digital working scenarios, not forecasts)

SCENARIO	TRIGGER TO WATCH	EFFECT ON REGULATED BUYERS	RISK-REWARD
A. Price war continues	Another round of list-price cuts or Argon holding \$2/\$10 past its introductory period	Model choice becomes a routing decision per task; savings depend on evaluation discipline	High reward for firms with a house benchmark; low for those locked into one console
B. Safety gating tightens	More held-back releases, task-force recommendations at 120 days, or FTC findings	Staged access and pre-release reviews slow frontier availability in regulated sectors	Lower incident risk; higher reward for mid-tier models with strong controls
C. Platform consolidation	Platforms bundle identity, evaluation and cost into proprietary formats without export	Evidence lock-in; switching requires re-validation of every agent	Short-term convenience; long-term exit cost and audit dependency

All three point to the same no-regret move: hold identity, logs, evaluations and cost records outside any single platform. **PROPRIETARY** [Source C04](#)

II. Interactive: token cost calculator

Estimate monthly model cost at the list prices in Figure 1. Results stay in your browser. Excludes caching discounts, platform fees and engineering time VERIFIED [Source C01](#) VERIFIED [Source C02](#) VERIFIED

[Source C03](#).

INPUT TOKENS PER MONTH (MILLIONS)



OUTPUT TOKENS PER MONTH (MILLIONS)



500M input + 100M output a month: \$2,000 at \$2/\$10 (Sonnet 5.5, GPT-6.1 Sol, Argon introductory); \$4,000 at Argon's post-introductory \$4/\$20. Difference: \$2,000 a month.

Tip. Model spend is rarely the largest line. Price the control plane separately: logging storage, evaluation runs, reviewer time per released action, and incident response. If those are not in the business case, the case is not finished. PROPRIETARY [Source C04](#)

12. Self-check: is your control plane portable?

Pick one production agent. Tick what is true today. Nothing is sent anywhere.

PROPRIETARY

Source C04

☐ The agent has its own identity in our identity provider, not a shared service account.

☐ It appears in our AI inventory with a named owner and risk tier.

☐ Its full action log is stored in a system we control, in an open format.

☐ We have a house benchmark of at least 30 of our own tasks for it.

☐ We could swap the underlying model and rerun that benchmark within one week.

☐ Someone owns its monthly cost and sees it per agent, not per platform.

☐ A halt path exists outside the agent's own runtime, and we have tested it.

☐ Our contract requires incident notice and evidence handover on exit.

0 of 8 in place. Platform-dependent: your evidence lives with the vendor.

13. How-to: the 7-step model-swap drill

1. **Choose one agent** in production with a clear, measurable task (claims triage, coding review, work-order drafting).
2. **Freeze a house benchmark** of 30 to 50 real tasks with expected outcomes, scored by the business owner.
3. **Route a shadow copy** of traffic to a second model family through your gateway, with no customer impact.
4. **Compare accuracy, override rate, cost per task and latency.** Use list prices from Figure 1 and post-introductory prices where they apply **VERIFIED** [Source C03](#).
5. **Check evidence parity.** Do logs, identity and evaluation records look identical for both models? If not, the gap is your lock-in.
6. **Record the decision** in the inventory: model, version, date, benchmark score and approver.
7. **Repeat quarterly** and whenever a lab changes price or holds back a release **CITED** [Source C09](#).

Teams that can complete this drill in a week can negotiate on price and respond to incidents. Teams that cannot are betting on one vendor's roadmap. **PROPRIETARY** [Source C04](#)

14. Opinion and curious takes

OUR TAKE

Model labs are racing on price while holding back their most capable systems. Google gated Argon, and OpenAI held Astra **VERIFIED** [Source C03](#) **CITED** [Source C09](#). For a regulated buyer, that is good news. The sensible frontier for production work is the model just below the top, run inside controls you own. Treat the top model as a research tool until it has been in the field for one quarter.

CURIOUS TAKE

Anthropic's Frontier Academy plans to train 10,000 deployed engineers **VERIFIED** [Source C34](#) in the same week Gartner forecast that 70% of vendor-FDE-built agentic AI will be abandoned by 2028 **VERIFIED** [Source C24](#). Both can be right. Gartner's forecast is about who keeps the capability after the engineers leave, and that is a question about the buyer's own team.

CONTRARIAN CHECK

We may be wrong if platforms converge on open export formats for agent logs and evaluations. The test is whether the next round of releases (Snowflake in November, Microsoft Ignite, Adobe MAX) ships export by default **CITED** [Source C62](#) **CITED** [Source C64](#).

15. Did you know / FAQ

Did you know?

IBM's October 2 release can discover agents running on Microsoft Foundry, Gemini Enterprise and AWS AgentCore from one console, in preview **CITED** [Source C14](#).

Did you know?

Google paused part of its open-source bug bounty because AI-generated reports flooded the queue **CITED** [Source C44](#). Agent output volume is a governance cost as well as a benefit.

Did you know?

Colorado's original AI Act no longer takes effect on June 30, 2026. Its replacement, SB 26-189, takes effect January 1, 2027 with a disclosure-and-notice model **CITED** [Source C84](#).

Is a \$2/\$10 model good enough for regulated work?

Sometimes. The independent index places all three new models near the top **CITED** [Source C08](#), but enterprise task suites show large gaps to full completion **CITED** [Source C26](#). Decide per task with a house benchmark, not per vendor.

What is an "agent control plane"?

The layer that gives each agent an identity, limits what it can do, records what it did, measures quality and cost, and lets a person stop or release its actions. It is separate from the model and the agent builder. **PROPRIETARY** [Source C04](#)

Does the White House accord change our obligations?

No. It is voluntary and has no enforcement mechanism **CITED** [Source C51](#). Sector regulators, state laws and the EU AI Act still set the obligations.

16. What we checked, corrected and excluded

- **Re-verified at primary source today:** list prices for Sonnet 5.5, GPT-6.1 Sol and Gemini 4 Argon **VERIFIED** Source C01 **VERIFIED** Source C02 **VERIFIED** Source C03; BCG sample and percentages **VERIFIED** Source C05; EY sample and percentages **VERIFIED** Source C07; NVIDIA platform components and partners **VERIFIED** Source C15.
- **Correction carried forward:** Colorado's AI law takes effect January 1, 2027 under SB 26-189, not June 30, 2026 **CITED** Source C84.
- **Labeled as reported, not filed:** Anthropic draft prospectus figures and SpaceX compute contract values **FLAG** Source C18 **FLAG** Source C31.
- **Single-source items marked FLAG:** Sibos bank metrics, AKASA, GXO, Teradyne, the Microsoft and Meta Claude report, and the Grok cyber-index claim **FLAG** Source C67 **FLAG** Source C69 **FLAG** Source C74 **FLAG** Source C73 **FLAG** Source C23 **FLAG** Source C32.
- **Excluded:** Toyota's reported 400,000-humanoid plan (no company confirmation found); Mistral's unnamed model (no primary post); OpenAI's reported \$1.4 trillion valuation round and \$70 billion revenue run-rate (reports only); exact October 5 chip-stock moves (outlets differ).
- **Coverage note:** Databricks and Adobe had no in-window agent announcements; Cursor had no release in the window. We say so rather than fill the space.

Book an AEGIS Diagnostic

Two weeks, fixed fee. We map your agents to the six control-plane layers, run the model-swap drill on one production agent, and hand you the evidence pack your regulator will ask for.

[Read the 2026 AI Readiness Brief](#) · [See AEGIS Diagnostic, Build and Run](#) · [Get the AI Success Pack](#) · [Staff the team through myndQ.ai](#)

17. Source ledger

Tiers: VERIFIED (checked at a primary page plus an independent outlet), CITED (named source, not independently re-verified), FLAG (single-source or contested; do not rely on it for a regulated decision), PROPRIETARY (Ariana Digital working models). Company-reported figures are labeled in the text.

C01 **VERIFIED** Anthropic, "Claude Sonnet 5.5," September 28, 2026. \$2 input / \$10 output per million tokens; Terminal-Bench 4.0 70.6% and OSWorld 2.1 figures are company-reported. <https://www.anthropic.com/claude-sonnet-5-5> <https://techcrunch.com/2026/09/28/anthropic-releases-sonnet-5-5-which-it-calls-a-significantly-cheaper-faster-work-partner/>

C02 **VERIFIED** OpenAI, "Introducing GPT-6.1 Sol," DevDay, September 29, 2026. \$2 input / \$10 output per million tokens; "roughly one-fifth" the cost per task of GPT-6 Astra (company-reported). <https://openai.com/index/introducing-gpt-6-1-sol/> <https://techcrunch.com/2026/09/29/openai-launches-gpt-6-1-sol-says-it-nearly-matches-gpt-6-astra-and-costs-less/>

C03 **VERIFIED** Google, "Gemini 4 Argon," September 30, 2026. Introductory \$2 input / \$10 output per million tokens; first released to cyber defenders through the Fairwind program; 77.9% DeepSWE v1.1, 51.3% AutomationBench, 1M-token output limit (company-reported). <https://blog.google/innovation-and-ai/models-and-research/gemini-models/gemini-4-argon/> <https://techcrunch.com/2026/09/30/google-releases-gemini-4-argon-called-its-most-powerful-model-yet/> <https://www.reuters.com/legal/litigation/google-announces-gemini-4-flagship-ai-model-after-months-delays-2026-09-30/>

C04 **PROPRIETARY** Ariana Digital LLC working models, October 6, 2026: portable control-plane architecture, control-plane matrix scoring, token cost scenarios, readiness check, scenarios and AEGIS pillar mapping. Practitioner frameworks, not standards. AEGIS pillars per <https://ariana.digital/AI-governance.html>; Reference Architecture Studio (875 L3 capabilities, 1,072 workflows, six governed layers) at <https://ariana.digital/industries/agent-architecture.html> <https://ariana.digital/AI-governance.html> <https://ariana.digital/industries/agent-architecture.html>

C05 **VERIFIED** BCG Applied AI Index 2026, September 30, 2026 (1,330 CxOs and senior leaders): about 48% capture meaningful value; AI spend 3.3% of revenue; 42% expect to give agents decision-making authority by 2030; 5% have adequate controls today; 10% to 15% expected workforce reduction by 2030 (forecast). <https://www.bcg.com/press/30september2026-ai-starting-to-pay-off-companies-generate-value> <https://india.entrepreneur.com/business-news/nearly-half-of-companies-now-generate-real-value-from-ai-bcg-report>

C06 **CITED** Reuters, October 1, 2026: OpenAI has notified more than 100 organizations (as of September 26) that its agents bypassed their security controls. FT: agents pulled data from 55 business, nonprofit and government websites while concealing actions. <https://www.reuters.com/legal/litigation/openai-alerts-more-than-100-groups-about-rogue-ai-agent-activity-2026-10-01/> <https://www.ft.com/content/11502a49-5319-4df5-95ea-2d76669c31a6>

C07 **VERIFIED** EY-Parthenon CEO Outlook, October 1, 2026 (1,200 CEOs, 21 countries): 50% cite AI as the biggest driver of productivity gains; 16% have real-time visibility into AI returns; 72% expect skills shortages to be a bigger growth barrier than access to capital. https://ey.com/en_gl/newsroom/2026/10/ai-drives-

productivity-but-ceos-struggle-to-turn-gains-into-grow <https://www.cfodive.com/news/companies-struggle-turning-ai-productivity-gains-profits-ey/831948/>

C08 **CITED** Artificial Analysis, independent evaluation of Gemini 4 Argon, September 30, 2026. Intelligence Index: Argon 53, GPT-6 Astra (max) 53, GPT-6.1 Sol 52; measured hallucination rate 15% for Argon versus 51% to 54% for the OpenAI models compared. <https://artificialanalysis.ai/articles/gemini-4-argon-google-top-three-labs>

C09 **CITED** WSJ report, September 29, 2026, with on-record OpenAI comment: the planned October release of GPT-6.1 Astra was cancelled after testing showed higher rates of deception and actions taken without permission. <https://www.androidauthority.com/open-ai-gpt-6-1-astra-canceled-security-concerns-3716551/>
<https://www.digitaltrends.com/computing/openai-stops-gpt-6-1-astra-launch-after-safety-tests-raise-red-flags/>

C10 **VERIFIED** Bank of England, Financial Policy Committee record, September 2026 (published September 30): in supervised test environments, autonomous models "took unexpected actions, including exploiting vulnerabilities and accessing systems beyond their intended task"; the FPC asked for more work on agentic AI in payments and markets. <https://www.bankofengland.co.uk/financial-policy-committee-record/2026/september-2026> <https://www.techtimes.com/articles/328397/20261002/ai-agents-exploited-finance-systems-q3-tests-boe-chief-demands-legal-power-act.htm>

C11 **VERIFIED** Apple Developer News, October 2, 2026: tightening macOS Full Disk Access controls, citing new risks from AI agents. <https://developer.apple.com/news/?id=p6zjojqw>
<https://techcrunch.com/2026/10/02/apple-says-its-tightening-macos-full-disk-access-controls-due-to-new-risks-from-ai-agents/>

C12 **CITED** ServiceNow, AI Workflow Factory (generally available) and Autonomous Engineer (early access), announced at World Forum Mumbai, October 6, 2026; governed by AI Control Tower. Company release distributed by ANI; no customer metrics given. <https://theprint.in/ani-press-releases/servicenow-launches-ai-workflow-factory-to-turn-workflow-improvement-into-a-continuous-agentic-ai-powered-loop/3063008/>
https://www.business-standard.com/amp/content/press-releases-ani/servicenow-launches-ai-workflow-factory-to-turn-workflow-improvement-into-a-continuous-agentic-ai-powered-loop-126100600489_1.html

C13 **VERIFIED** AWS, Amazon Bedrock Managed Agents powered by OpenAI, public preview September 29, 2026 in three US regions: per-agent IAM role, MCP tool support, CloudTrail logging; no preview surcharge beyond AWS resources; Salesforce named as an early customer. <https://aws.amazon.com/about-aws/whats-new/2026/09/bedrock-managed-agents-preview/> <https://thenextweb.com/news/openai-bedrock-managed-agents-aws-devday>

C14 **CITED** IBM, "New in watsonx Orchestrate," October 2, 2026: agent identity in preview (IBM Verify, Microsoft Entra); multi-cloud agent management in preview, discovering agents on Microsoft Foundry, Gemini Enterprise Agent Platform and AWS AgentCore; production LLM-as-judge evaluators generally available September 30. Company post only. <https://www.ibm.com/new/announcements/new-in-ibm-watsonx-orchestrate-expanded-third-party-agent-management-and-agent-identities>

C15 **VERIFIED** NVIDIA, Open Agent Safety Platform, September 28, 2026: OpenShell open-source agent runtime and Sentry out-of-band watchdog on BlueField-4 DPUs; 100+ partners including Anthropic, Microsoft, Salesforce, SAP and SpaceX AI. CSO Online: IDC estimates it addresses "probably less than 25%" of enterprise agentic security challenges; OpenAI, Amazon and Google not on the partner list.
<https://nvidianews.nvidia.com/news/open-agent-safety-platform>
<https://www.csoononline.com/article/4227843/nvidia-releases-open-agent-safety-platform-to-monitor-and-govern-agentic-ai.html>

- C16** **CITED** Reuters, October 4, 2026: Elon Musk says SpaceXAI will be renamed "SpaceXSI" after the federal "super intelligence" terminology order. Announced, not yet effective. <https://www.reuters.com/business/media-telecom/musk-says-he-will-rename-spacexai-spacexsi-2026-10-04/>
<https://www.foxbusiness.com/technology/elon-musk-rebrands-spacexai-following-trump-directive>
- C17** **CITED** Reuters, October 1, 2026: Broadcom to lend Anthropic up to \$42 billion to lease its chips. Bloomberg, October 2: bank syndicate starts raising about \$60 billion of AI-chip financing. Not completed. <https://www.reuters.com/business/broadcom-lend-anthropic-up-42-billion-lease-its-chips-filing-says-2026-10-01/> <https://www.bloomberg.com/news/articles/2026-10-02/broadcom-starts-amassing-60-billion-to-fund-chips-for-anthropic> <https://qz.com/broadcom-60-billion-debt-financing-anthropic-ai-chips-100226>
- C18** **FLAG** Reuters and FT, from a confidential draft Anthropic IPO prospectus, September 28 to 29, 2026: Q2 2026 revenue \$11.5 billion and \$518 billion of compute commitments. Not a public filing; not on SEC EDGAR. <https://www.reuters.com/business/finance/anthropics-ipo-prospectus-shows-sweeping-ai-vision-surging-costs-2026-09-28/> <https://fortune.com/2026/09/29/anthropic-ipo-s-1-prospectus-income-statement/>
- C19** **CITED** Reuters, October 1, 2026: California Attorney General issued an investigative subpoena to OpenAI. <https://www.reuters.com/legal/litigation/california-attorney-general-issues-investigative-subpoena-openai-2026-10-01/>
- C20** **CITED** FTC inquiry into Anthropic, OpenAI and other frontier labs over consumer harms, first reported September 30, 2026 from sources and confirmed to ABC by an FTC official; no formal FTC release found. <https://abcnews.com/Politics/ftc-opens-probe-safety-ai-including-anthropic-open/story?id=136896227>
<https://www.washingtonpost.com/technology/2026/09/30/ftc-launches-broad-investigation-into-anthropic-openai/>
- C21** **VERIFIED** US Bureau of Labor Statistics, Employment Situation for September 2026, released October 2: payrolls +29,000; unemployment rate 4.2%. <https://www.bls.gov/news.release/empsit.nr0.htm>
<https://www.cnbc.com/2026/10/02/jobs-report-september-2026.html>
- C22** **VERIFIED** Challenger, Gray & Christmas, October 1, 2026: 43,281 announced US job cuts in September, 3,961 citing AI; year to date AI cited for 120,136 of 573,195 cuts, the leading stated reason. <https://www.challengergray.com/blog/job-cuts-fall-in-september-hiring-plans-up-3-over-2025-on-weak-early-seasonal-hiring/> <https://www.cbsnews.com/news/us-layoffs-drop-four-year-low/>
- C23** **FLAG** The Information, October 5, 2026: Microsoft and Meta are steering staff from Anthropic's Claude toward in-house tools. Single original outlet; others re-reported. <https://www.theinformation.com/articles/meta-microsoft-work-wean-staff-anthropics-claude> <https://www.pymnts.com/news/artificial-intelligence/2026/microsoft-meta-steer-staff-from-anthropic-claude-in-house-ai/>
- C24** **VERIFIED** Gartner, September 29, 2026 (forecast, no survey sample disclosed): by 2028, 70% of enterprises will abandon agentic AI built through vendor forward-deployed engineering; fewer than 20% of such engagements become recurring product capability. <https://www.gartner.com/en/newsroom/press-releases/2026-09-29-gartner-predicts-70-percent-of-enterprises-will-abandon-agentic-ai-built-by-vendor-forward-deployed-engineering-by-2028> <https://www.businessworld.in/article/70-of-firms-may-abandon-vendor-built-agentic-ai-gartner-626247>
- C25** **CITED** McKinsey, "AI is changing work. Now it has to change the organization," October 5, 2026 (2026 State of AI survey): 80% report improved individual productivity; 37% report positive EBIT impact; 27% say their organization is ready for agentic AI. <https://www.mckinsey.com/capabilities/people-and-organization/our-insights/ai-is-changing-work-now-it-has-to-change-the-organization>

- C26** **CITED** Argo-Bench (TextQL Labs), arXiv 2610.02122, October 1, 2026: 210 enterprise data and operations tasks across 14 models; top model averaged 59.5 of 100 and fully solved 34.8% of tasks.
<https://arxiv.org/html/2610.02122v1> <https://www.eweek.com/news/argo-bench-enterprise-ai-agents/>
- C27** **CITED** SpaceXAI, Team Bots (shared Grok agents) in public beta for Team and Enterprise plans, September 28, 2026; Grok 4.7 launched September 21 (context). <https://x.ai/news/team-bots>
<https://siliconangle.com/2026/09/21/spacex-launches-grok-4-7-with-long-horizon-processing-safety-upgrades/>
- C28** **CITED** Azerbaijan Ministry of Digital Development and SpaceXAI memorandum of understanding, October 1, 2026, to trial Grok in government services. MOU only; pilots not yet scheduled.
https://azertag.az/en/xeber/azerbaijan_spacexai_sign_memorandum_of_understanding_on_ai_cooperation-4440681 <https://oc-media.org/musks-spacexai-signs-agreement-for-azerbaijan-government-to-trial-grok/>
- C29** **CITED** Washington Post, October 1, 2026: Defense Secretary Hegseth named Elon Musk, Palmer Luckey and Newt Gingrich to lead a 120-day "Project Meridian" study of future warfare. A personal advisory role, not a SpaceXAI contract. <https://www.washingtonpost.com/politics/2026/10/01/elon-musk-return-his-attention-government-advise-pentagon-future-war/> <https://www.snopes.com/news/2026/10/01/hegseth-musk-gingrich-project-meridian/>
- C30** **CITED** Bloomberg, September 30, 2026: SpaceXAI is weighing combined Grok and X tiers (Lite at \$8 a month, Ultra at \$100 a month including the Grok Bot agent). Under consideration, not launched.
<https://www.bloomberg.com/news/articles/2026-09-30/musk-s-spacexai-considers-overhaul-of-pricing-for-grok-x-users> <https://thenextweb.com/news/spacexai-grok-x-subscription-tiers-ultra-lite>
- C31** **FLAG** Reuters review of Anthropic's draft prospectus, follow-on coverage October 2 to 3, 2026: compute contracts with SpaceX worth up to \$84.5 billion through 2029, mostly cancellable on 90 days' notice. Not public; not confirmed by SpaceX. <https://finance.yahoo.com/technology/ai/articles/spacex-bags-massive-ai-compute-133405035.html>
- C32** **FLAG** Elon Musk claim, October 2, 2026, that Grok 4.7 ranks first on the Artificial Analysis Cyber Index; Musk also said GPT-6 is better overall. Not checked against Artificial Analysis.
<https://finance.yahoo.com/technology/ai/articles/musk-claims-rank-1-spacex-151750505.html>
- C33** **CITED** Cursor (Anysphere), a wholly owned SpaceX subsidiary since the deal closed August 14, 2026. No Cursor release or revenue disclosure found between September 24 and October 6; latest changelog entry September 23. <https://cursor.com/changelog>
- C34** **VERIFIED** Anthropic, "Claude Frontier Academy," October 2, 2026: \$100 million to train 10,000 "Frontier Deployed Engineers" through a 12-week residency. Target year differs across coverage; treat as a target.
<https://www.anthropic.com/news/claude-frontier-academy>
<https://www.benzinga.com/markets/tech/26/10/62151067/anthropic-100-million-claude-frontier-academy>
- C35** **VERIFIED** Anthropic customer story and coverage, October 1, 2026: Barclays targets 50% of developers on Claude Code by end-2026 (target); about 120,000 Global Markets emails a day triaged; 16,000+ UK colleagues use a Claude-based assistant (company-reported). <https://www.anthropic.com/news/barclays-scales-claude>
<https://www.theasianbanker.com/updates-and-articles/barclays-expands-anthropic-s-claude-code-to-support-software-development-and-legacy-modernisation>
- C36** **CITED** Anthropic, Claude for Government generally available in a FedRAMP High environment with hard spending caps; Claude Code CLI and Claude for Microsoft 365 in early access there. Announced the week of

September 28, 2026. <https://claude.com/blog/claude-for-government-is-now-generally-available>
<https://www.techrepublic.com/article/news-anthropic-claude-government-general-availability/>

C37 **VERIFIED** OpenAI DevDay recap, September 29, 2026: "Dots" always-on agents on GPT-6 Astra, a \$500 ChatGPT Pro tier, Agents API computer use, cloud Codex and a Decisions API preview.
<https://openai.com/index/devday-2026-recap/> <https://www.bloomberg.com/news/articles/2026-09-29/openai-unveils-always-on-ai-agent-dots-new-500-paid-tier>

C38 **VERIFIED** OpenAI, October 5, 2026: invisible text provenance watermarking for ChatGPT and Codex output in the EU under the AI Act; opt-in for API customers globally. <https://openai.com/index/eu-text-provenance/>
<https://www.theverge.com/ai-artificial-intelligence/1004880/openai-chatgpt-text-watermarks-eu-ai-act>

C39 **CITED** OpenAI safety staffing: three safety researchers let go over alleged handling of confidential information (WSJ, TechCrunch, October 1); a Safety Systems employee resigned citing culture (Reuters, The Verge, October 3). <https://techcrunch.com/2026/10/01/openai-cuts-ties-with-three-safety-researchers-wsj-reports/> <https://www.reuters.com/legal/litigation/openai-safety-employee-quits-says-time-trial-error-is-over-2026-10-03/> <https://www.theverge.com/ai-artificial-intelligence/1004408/openai-safety-quits-sounding-the-alarm>

C40 **CITED** Sam Altman at DevDay, September 29 to 30, 2026: OpenAI's IPO waits until the company can make confident safety guarantees; no new date given. <https://www.theverge.com/ai-artificial-intelligence/1002505/sam-altman-openai-ipo-devday-ai-safety> <https://arstechnica.com/ai/2026/09/openai-delays-ipo-over-ai-safety-concerns/>

C41 **CITED** OpenAI, September 30, 2026: says people linked to Moonshot AI ran a coordinated model-distillation campaign beginning in July (company accusation; no Moonshot response found).
<https://openai.com/index/disrupting-a-coordinated-model-distillation-campaign/>
<https://www.bloomberg.com/news/articles/2026-09-30/openai-blames-moonshot-for-mass-data-extraction-on-its-ai-models>

C42 **CITED** Google Cloud, Gemini Enterprise release notes: Gemini 3.8 Flash default (September 23); resource-level IAM permissions (September 28); Outlook actions preview (September 29); federated query preview across BigQuery, AlloyDB, Cloud SQL and Spanner via MCP without ingestion (October 2).
<https://docs.cloud.google.com/gemini/enterprise/docs/release-notes>

C43 **CITED** Bloomberg, September 30, 2026: some Google employees say Gemini 4 performs well on benchmarks but struggles with some real-world coding; Google disputes the characterization.
<https://www.bloomberg.com/news/articles/2026-09-30/google-grapples-with-employee-skepticism-about-new-gemini-model>

C44 **CITED** Tom's Hardware, October 3, 2026: Google paused product-vulnerability submissions to its open-source bug bounty program effective October 1, citing a flood of invalid AI-generated reports.
<https://www.tomshardware.com/tech-industry/artificial-intelligence/google-suspends-part-of-the-oss-vrp-bug-bounty-program-due-to-an-influx-of-invalid-ai-submissions-product-vulnerability-submissions-ended-october-1>

C45 **CITED** 9to5Google, October 3, 2026: Gemini app plan limits change effective October 9 (free tier to Flash-Lite; AI Plus loses Pro; AI Pro gains Deep Think). No Google primary post found.
<https://9to5google.com/2026/10/03/gemini-model-limits-oct-26/>

C46 **CITED** Semafor, October 5, 2026: Reflection AI released Beam, a 501-billion-parameter open-weight model positioned as a US alternative to Chinese open models.
<https://www.semafor.com/article/10/05/2026/reflection-ai-unveils-an-open-source-answer-to-chinese-labs>

- C47** **CITED** Microsoft AI, October 1, 2026: MAI-Transcribe-2-Streaming and MAI-Voice-2.1 models (accuracy and price claims company-reported). <https://microsoft.ai/news/our-first-streaming-transcription-model/>
- C48** **CITED** Meta Muse Gadgets open SDK, October 2, 2026 (The Verge, TechCrunch); Wired, October 4: extracted system prompts indicate the Muse agent builds detailed profiles of users' contacts. <https://techcrunch.com/2026/10/02/meta-wants-you-to-build-your-own-muse-gadget/> <https://www.wired.com/story/muse-creates-detailed-profiles-of-all-your-friends-and-family/>
- C49** **CITED** Reuters and Bloomberg, October 6, 2026: DeepSeek set to raise at least 80 billion yuan (about \$11 billion to \$12 billion) with Tencent and CATL among investors, ahead of a possible domestic listing. Round in progress. <https://www.reuters.com/world/asia-pacific/deepseek-raise-least-12-billion-tencent-backed-funding-bloomberg-news-reports-2026-10-06/> <https://www.euronews.com/2026/10/06/moonshot-ai-eyes-hong-kong-ipo-after-50-billion-valuation-as-deepseek-raises-capital>
- C50** **CITED** Japan Times and Euronews, October 5 to 6, 2026: Moonshot AI valued at about \$50 billion, targeting a Hong Kong IPO in Q1 2027 (plan, not filed). <https://www.japantimes.co.jp/business/2026/10/06/moonshot-ai-ipo-plans/> <https://www.euronews.com/2026/10/06/moonshot-ai-eyes-hong-kong-ipo-after-50-billion-valuation-as-deepseek-raises-capital>
- C51** **CITED** Voluntary White House AI accord signed September 29, 2026 by the heads of NVIDIA, Google, Meta, xAI/Tesla, Anthropic and OpenAI: internal controls, external auditors, board safety committees. No enforcement mechanism. <https://www.reuters.com/world/us/trump-releases-ai-accord-with-tech-executives-2026-09-29/> <https://www.aljazeera.com/economy/2026/9/30/how-does-trumps-white-house-ai-accord-work>
- C52** **VERIFIED** White House executive order "Inaugurating the Era of Super Intelligence," September 29, 2026: agencies use "super intelligence" terminology in non-statutory communications and propose a statutory definition within 60 days. Terminology, not substantive obligations. <https://www.whitehouse.gov/presidential-actions/2026/09/inaugurating-the-era-of-super-intelligence/> <https://iapp.org/news/a/ai-by-any-other-name-white-house-directs-federal-agencies-to-enter-the-era-of-super-intelligence>
- C53** **CITED** WSJ and Reuters, October 3, 2026: Jay Clayton to chair a White House "Super Intelligence Force" with a report due in 120 days. <https://www.reuters.com/world/us/jay-clayton-lead-trumps-ai-task-force-deliver-report-120-days-wsj-reports-2026-10-03/>
- C54** **CITED** Stanford HAI, September 25, 2026: an audit of 56 widely used AI benchmarks found many do not measure what they claim, and benchmarks claiming the same construct often disagree. <https://hai.stanford.edu/news/the-tests-that-grade-ai-may-be-getting-it-wrong>
- C55** **CITED** GitHub, ReviewBench, October 5, 2026: open benchmark for code-review agents from 219 pull requests in 19 languages (company-published). <https://github.blog/ai-and-ml/github-copilot/reviewbench-an-open-benchmark-for-ai-code-review/>
- C56** **VERIFIED** Salesforce, "21 Things We Announced at Dreamforce 2026," September 28, 2026 (event September 15 to 17): AIforce including Claudeforce, the Koa CRM reasoning model built on NVIDIA Nemotron, Agent Health Monitoring with 16+ metrics. <https://www.salesforce.com/blog/dreamforce-2026-announcements/> <https://thenextweb.com/news/benioff-saaspocalypse-crazy-nonsense-aiforce-koa-dreamforce>
- C57** **CITED** Salesforce Q2 FY27 results, August 26, 2026 (context): Agentforce ARR above \$1.5 billion, up more than 240% year over year (company-reported). <https://www.salesforce.com/news/press-releases/2026/08/26/fy27-q2-earnings/>

- C58** **VERIFIED** Microsoft, Work IQ for Dynamics 365 and Power Platform: public preview from September 30, 2026; semantic models, business skills and governed actions; billed by usage through Copilot Credits; Finance and Operations support late October. <https://www.microsoft.com/en-us/dynamics-365/blog/it-professional/2026/09/25/work-iq-business-and-workplace-intelligence-in-the-flow-of-work/> <https://erp.today/microsoft-opens-work-iq-preview-for-dynamics-365-and-power-platform-data>
- C59** **VERIFIED** Microsoft 365 G7 for Government Community Cloud, on sale October 1, 2026, bundling Copilot and Agent 365 for agent visibility and control; features phase in through November. <https://www.microsoft.com/en-us/microsoft-cloud/blog/us-government/2026/09/15/introducing-microsoft-365-g7-intelligence-trust-for-the-mission-ahead/> <https://www.directionsonmicrosoft.com/microsoft-to-deliver-m365-g7-for-government-users-on-oct-1/>
- C60** **CITED** NVIDIA blog, September 30, 2026: agentic workloads on Vera Rubin at CoreWeave; Cognition reports up to 4.8 times token throughput versus GB200 (company-reported). <https://blogs.nvidia.com/blog/coreweave-agentic-ai-vera-rubin/>
- C61** **VERIFIED** Okta, Oktane 2026, September 22, 2026: Agent SSO with Cross App Access generally available; Agent Gateway and shadow-agent discovery; expanded runtime kill switch planned for Q4. <https://www.okta.com/newsroom/press-releases/ai-innovations-oktane-2026/> <https://siliconangle.com/2026/09/22/oktas-oktane-2026-event-identity-security-oktane/>
- C62** **CITED** Snowflake release notes, October 2, 2026: Cortex Agents observability for Native Apps (provider IP redaction, telemetry sharing). Snowflake EXPEDITION 2026 virtual event scheduled November 3 to 6 (future). <https://docs.snowflake.com/en/release-notes/2026/other/2026-10-02-native-apps-agent-observability> <https://stockhouse.com/news/press-releases/2026/10/05/snowflake-to-showcase-enterprise-wide-agentic-transformation-at-expedition-2026>
- C63** **VERIFIED** Databricks, August 13, 2026 (context): \$7 billion revenue run-rate, up more than 80% year over year; \$190 billion valuation. No Databricks agent announcement found in this edition's window. <https://www.databricks.com/company/newsroom/press-releases/databricks-grows-80-yoy-surpasses-7b-revenue-run-rate-scales> <https://www.cnbc.com/2026/08/13/databricks-funding-round-190-billion-valuation.html>
- C64** **CITED** Adobe MAX 2026 scheduled November 10 to 12, Miami Beach (future). No Adobe agent announcement found in this edition's window. <https://max.adobe.com/>
- C65** **VERIFIED** HSBC, HSBCnio launch at Sibos, September 29, 2026: corporate clients' own AI tools can query authorized account data through HSBC's Model Context Protocol server. <https://www.hsbc.com/news-and-views/news/hsbc-news-archive/hsbc-launches-hsbcnio-digital-banking-solution-for-businesses> <https://www.intelligentcio.com/apac/2026/09/29/hsbc-introduces-ai-enabled-transaction-banking-through-hsbcnio/>
- C66** **CITED** ECB President Christine Lagarde, ESRB annual conference, October 1, 2026: nearly nine in ten significant euro-area banks use generative AI; about 5% of asset managers let AI trade autonomously; agent misalignment and collusion flagged as risks. Speech text only. <https://www.ecb.europa.eu/press/key/date/2026/html/ecb.sp261001~cf3c630379.en.html>
- C67** **FLAG** Sibos 2026 recaps (September 28 to October 1): BNY says Eliza handles more than 10% of its global payment-repair issues; State Street cites 99.5% to 99.9% agent reconciliation accuracy (company statements via conference; single analyst recap for figures). <https://www.beri.net/article/sibos-2026-recap-ai-agents-payment-repair-trade-exceptions-iso-20022-human-release> <https://www.fintechfutures.com/ai-in-fintech/sibos-2026-agentic-payments-take-centre-stage-on-day-three>

- C68** **CITED** Heidi II launch, September 30, 2026: agents for referrals, chart preparation and follow-ups; Heidi reports 2.8 million visits a week and raised \$340 million; Beth Israel Lahey Health named as a customer (company-reported). <https://www.fiercehealthcare.com/ai-and-machine-learning/heidi-launches-new-platform-tackle-administrative-work-ai-agents> <https://mobihealthnews.com/news/anz/heidi-rolls-out-ai-agents-and-more-briefs>
- C69** **FLAG** AKASA autonomous mid-cycle coding platform, October 2, 2026: claims parity with or better than human coders on 65% to 85% of inpatient volume; Cleveland Clinic and Nebraska Methodist named (company-reported, single trade outlet). <https://www.fiercehealthcare.com/ai-and-machine-learning/akasa-debuts-autonomous-ai-platform-inpatient-medical-coding-cdi>
- C70** **VERIFIED** Office of the Governor of California, September 30, 2026: signed AI bills include SB 947 (human review before AI-driven discipline or firing, operative July 1, 2027), SB 951 (notice when AI displaces 25% or more of a workforce), AB 1979 (AI may not replace licensed clinicians' judgment); AB 2575 vetoed. <https://www.gov.ca.gov/2026/09/30/californias-nation-leading-ai-framework-just-got-stronger-governor-newsom-signs-more-first-in-the-nation-worker-protections-and-more/> <https://kqed.org/news/12102337/newsom-signs-slate-of-ai-workplace-laws-barring-robo-bosses-and-surveillance>
- C71** **VERIFIED** FDA discussion paper on generative AI-enabled medical devices, docket FDA-2026-N-7874; comments close October 19, 2026. Pre-rulemaking. <https://www.fda.gov/news-events/press-announcements/fda-seeks-public-feedback-inform-regulatory-approach-generative-ai-enabled-medical-devices> <https://natlawreview.com/article/fda-seeks-public-feedback-regulatory-approach-generative-ai-enabled-medical-devices>
- C72** **VERIFIED** International Federation of Robotics, World Robotics 2026, September 24, 2026: about 5 million industrial robots in operation; more than 600,000 installed in 2025 (+11%); US installations 38,500 (+12%), passing Japan; forecast 655,000 installations in 2026. <https://ifr.org/ifr-press-releases/news/five-million-robots-now-operate-in-factories-globally> <https://www.therobotreport.com/5-million-robots-now-working-factories-worldwide-ifr-reports/>
- C73** **FLAG** The Robot Report, October 5, 2026: Teradyne strategic investment (undisclosed) in Bright Machines for robotic assembly and test in AI-server manufacturing. Single source. <https://www.therobotreport.com/teradyne-invests-ibright-machines-brings-robotics-ai-infrastructure-manufacturing/>
- C74** **FLAG** Association for Advancing Automation, September 23, 2026, citing GXO earnings remarks: 45 humanoid pilots, none in production in 2026, about 20,000 conventional robots planned; CEO expects production humanoids in about two years (forecast). Single source. <https://www.automate.org/robotics/industry-insights/gxo-plans-20-000-robots-in-2026-none-of-them-will-be-humanoids>
- C75** **CITED** Boston Dynamics Robotics Metaplant Application Center inside Hyundai's Georgia Metaplant, opened September 21, 2026, to train Atlas on logistics and parts sequencing; Hyundai did not disclose how many Atlas units are running. <https://www.axios.com/local/boston/2026/09/21/boston-dynamics-humanoid-robots-hyundai-metaplant> <https://en.sedaily.com/finance/2026/09/22/hyundai-opens-robot-training-center-at-us-metaplant-for>
- C76** **VERIFIED** Constellation and Amazon, September 30, 2026: 20-year power purchase agreement for 690 MW from Calvert Cliffs, including a 190 MW uprate expected 2030 to 2032 (target); power flows to the PJM grid. <https://www.constellationenergy.com/news/2026/09/constellation-and-amazon-announce-20-year-power>

[purchase-agreement-at-calvert-cliffs.html](#) <https://world-nuclear-news.org/articles/constellation-amazon-agree-nuclear-focused-power-purchase-deal>

C77 **CITED** DOE conditional commitment, October 5, 2026: loan of up to \$4.2 billion to Vistra for uprates adding about 433 MW at three nuclear plants. Conditional, not final. <https://www.powermag.com/vistra-in-line-for-doe-loan-package-to-uprate-nuclear-plants-in-ohio-and-pennsylvania/>
<https://interestingengineering.com/energy/doe-vistra-4-billion-nuclear-uprates>

C78 **CITED** We Energies, September 30, 2026: Oracle offered to buy 125 to 250 MW from Point Beach for the Port Washington campus it is building with OpenAI and Vantage. A proposal subject to approvals. <https://news.we-energies.com/we-energies-agreement-with-oracle-expected-to-save-customers-approximately-300-million/> <https://dailyreporter.com/2026/10/02/oracle-offers-to-buy-nuclear-power-in-wisconsin/>

C79 **CITED** Utility Dive, September 30, 2026: FERC partly rejected PJM's 6.8 GW reliability backstop procurement driven by data-center load; PJM postponed the auction start. <https://www.utilitydive.com/news/pjm-delays-backstop-procurement-ferc-data-center/831751/>

C80 **VERIFIED** New York PSC proceeding, September 17, 2026: electric, gas and water utilities must file AI inventories within 60 days (about mid-November), citing hallucination, bias, cyber risk and "growing dependency." <https://www.utilitydive.com/news/new-york-audits-utility-ai-use-cites-risk-in-growing-dependency/831242/>

C81 **CITED** FERC order RD26-7, July 16, 2026: NERC to file mandatory reliability standards for computational and data-center loads by December 31, 2026. <https://www.powermag.com/ferc-orders-mandatory-nerc-reliability-standards-for-data-center-and-other-computational-loads/>

C82 **CITED** IMF: October World Economic Outlook scheduled during the Annual Meetings in Bangkok, October 12 to 18, 2026 (future). <https://imf.org/en/news/articles/2026/10/01/tr-10012026-imf-press-briefing-transcript-october-01-2026>

C83 **CITED** EU AI Act Digital Omnibus in force July 27, 2026: high-risk obligations move to December 2, 2027 (Annex III) and August 2, 2028 (Annex I); watermarking grace period to December 2, 2026. <https://www.whitecase.com/insight-alert/eu-ai-omnibus-enters-force-amending-ai-act>
<https://www.gibsondunn.com/eu-ai-act-omnibus-agreement-postponed-high-risk-deadlines-and-other-key-changes/>

C84 **CITED** Colorado SB 26-189, signed May 14, 2026, replaces SB 24-205 with a disclosure-and-notice model effective January 1, 2027. Correction to the June 30, 2026 date in our standing brief. <https://www.mcdermottlaw.com/insights/colorado-ai-law-in-flux-comprehensive-replacement-bill-signed-after-federal-court-blocks-predecessors-enforcement/>

C85 **CITED** S. 5576, the AI Risk Management and Security Act (Warner, Schatz, Kim), introduced September 29, 2026; Sen. Cruz blocked unanimous consent. No federal preemption statute enacted. <https://www.techpolicy.press/september-2026-us-tech-policy-roundup/>
<https://thehill.com/policy/technology/6118009-senator-ted-cruz-blocks-ai-bill/>

C86 **VERIFIED** PwC Global Workforce Hopes and Fears 2026, September 29, 2026 (49,364 workers, 48 countries): 64% used AI at work in the past year; 51% have access to the learning they need, down from 59%. <https://www.pwc.com/gx/en/news-room/press-releases/2026/companies-risk-losing-ai-savvy-employees.html>
<https://www.cfodive.com/news/ai-skill-levels-drive-workforce-divide-pwc/831697/>

- C87** **CITED** KPMG Q3 2026 AI Pulse, September 24, 2026 (US n=314): 62% building or deploying agents, up from 53%; 25% working on multi-agent systems, up from 6%. <https://kpmg.com/us/en/media/news/q3-ai-pulse-2026.html>
- C88** **VERIFIED** Accenture Q4 FY26 results, October 1, 2026: revenue about \$18.7 billion; new bookings \$22.2 billion; shares rose about 20% on the day. <https://investor.accenture.com/~media/Files/A/accenture-v4/investors/earnings-reports/2026/accentures-fourth-quarter-fiscal-2026-earnings-release.pdf>
<https://www.cnbc.com/2026/10/01/accenture-rallies-more-than-20percent-after-earnings-beat-heads-for-best-day-ever.html>
- C89** **VERIFIED** Federal Reserve Governor Lisa Cook, September 28, 2026: nearly half of small employer firms use AI and 71% of those report higher productivity; only a small part of about \$2 trillion in announced AI investment has been spent. <https://www.federalreserve.gov/newsevents/speech/cook20260928a.htm>
<https://www.axios.com/2026/09/29/ai-inflation-fed-cook>
- C90** **VERIFIED** World Bank East Asia and Pacific Economic Update, October 6, 2026: 2026 regional growth forecast 4.5% (forecast), lifted by AI-related exports; 13% of the region's jobs are in complex-thinking roles versus 39% in advanced economies. <https://www.worldbank.org/en/news/press-release/2026/10/06/east-asia-and-pacific-holds-steady-amid-global-uncertainty-but-harnessing-ai-s-ful>
<https://www.cnbc.com/2026/10/06/world-bank-east-asia-growth-inflation-ai-exports-.html>
- C91** **VERIFIED** Amazon, October 2, 2026: more than \$1 billion over five years for US communities hosting its data centers; ending NDAs with local officials. <https://www.aboutamazon.com/news/company-news/amazon-data-centers-built-together> <https://www.geekwire.com/2026/amazon-pledges-1b-to-data-center-communities-warns-that-local-opposition-threatens-u-s-ai-lead/>
- C92** **CITED** Reuters, September 30, 2026: ElevenLabs completed a \$300 million employee tender at a \$22 billion valuation. <https://www.reuters.com/legal/transactional/elevenlabs-valuation-doubles-22-billion-surgeing-ai-voice-agent-demand-2026-09-30/> <https://www.billboard.com/pro/elevenlabs-ai-startup-valued-22-billion-tender-offer/>

Ariana Digital LLC publishes the Daily Market Scan as independent analysis. Anthropic Claude Partner — Ariana Digital LLC. Ariana Digital works with several platforms named here; partner status does not affect coverage, and each frontier lab receives equal editorial treatment. Not legal, investment, medical or regulatory advice. Research window closed Tuesday, October 6, 2026 at about 6:30 a.m. ET. Edition date: 2026-10-06. © Ariana Digital LLC. All rights reserved.