

DAILY MARKET SCAN · FRIDAY EDITION · WEEKLY DIGEST, REG-READY FIELD NOTE, INDUSTRIES DEEP-DIVE

Friday, October 9, 2026 · Edition 2026-10-09

List it before you scale it: the agent inventory became the week's common demand

This week, small frontier models got much cheaper and the major agent products began routing work across rival labs' models. In the same five days, Singapore's central bank finalized AI rules that require financial institutions to keep inventories of their AI use

VERIFIED

Source C01, New York utilities worked toward an AI inventory due in November

CITED

Source C42, and Utah expanded a health-AI sandbox in which independent evaluators audit decision logs **VERIFIED** Source C26. Our read: before anyone can stop, undo or prove what an agent did, someone has to be able to list it. This Friday edition rolls up the week, tests the inventory idea across four regulated industries, and ships a Reg-Ready field note on building the register.

Research window: Monday, October 5, through Friday, October 9, 2026, about 7:00 a.m. ET, with labeled context from September 2026. All dates are America/New_York. "This week" means Monday, October 5, to Friday, October 9. "Last week" means September 28 to October 2. "Next week" means October 12 to 16. Company figures are labeled company-reported. Forecasts, pilots, proposals, conditional commitments and scheduled dates are labeled as such.

Contents

1. The 60-second scan
2. The Friday thesis: list it before you scale it
3. Weekly digest: one week, one control path
4. Frontier ledger: equal weight, what moved
5. Enterprise platforms: what shipped
6. Industries deep-dive: four case studies and their architecture
7. Industry scorecard: wins, constraints, actions
8. Physical AI and robotics
9. Regulators: the map and the calendar
10. Markets, labor and research signals

11. Framework: the AEGIS Agent Register

12. Scenarios to 2028 and a register-coverage slider

13. Self-check: could you list every agent by Monday?

14. How-to: a ten-day register sprint

15. Practical artifacts: field note and infographic

16. Opinion, curious takes and did-you-know

17. FAQ

18. What we checked, corrected and excluded

19. Source ledger

I. The 60-second scan

\$0.10

per million input tokens for Anthropic's Claude Haiku 5.5 (prompts up to 100K tokens), released Wednesday, October 7

VERIFIED

Source C03

Oct 7, 2027

when MAS AI risk management guidelines start to apply; full life-cycle sections by October 7, 2028

VERIFIED

Source C01

9+

South Korean banks reported as targets of attacks that CrowdStrike links, with moderate confidence, to one actor using an open-source pen-test tool and LLMs

VERIFIED

Source C02

24%

of CHROs say their leadership team agrees on who is accountable for AI agents (Gartner, survey of 297)

CITED

Source C46

- **Frontier, equal weight.** OpenAI began rolling GPT-6 out on October 7, with Sol for paid plans and Luna reaching Free and Go on October 8 **VERIFIED** Source C05. Google launched a Gemini agent for work on October 8, with its own Workspace identity and a model picker that starts with Anthropic's Claude **VERIFIED** Source C06. Elon Musk said Grok Bot will route tasks to rival models, including Claude Opus 5.5 **CITED** Source C08. Mistral put a roughly 1-trillion-parameter Large 4 into public preview on October 6 **VERIFIED** Source C07.
- **Financial services.** MAS finalized guidelines that apply to all financial institutions and all forms of AI, and plans to consult on agentic AI in 2027 **VERIFIED** Source C01. The Evident AI Index says 12 banks now report realized or projected AI returns, up from 8 (company-reported figures) **VERIFIED** Source C20.
- **Healthcare.** Ochsner reports a 41% rise in clinical-trial screening capacity with Paradigm Health, with coordinators verifying every candidate (company-reported) **CITED** Source C25. FDA comments on generative-AI devices close Monday, October 19 **VERIFIED** Source C28.

- **Manufacturing.** Schneider Electric agreed to buy PTC for about \$22.6 billion on October 5; the deal has not closed **VERIFIED** [Source C32](#). Boston Dynamics named former Amazon AI chief Rohit Prasad as CEO on October 6 **VERIFIED** [Source C33](#).
- **Energy.** Google and Constellation signed agreements in PJM on October 6: 890 MW of new nuclear from uprates under a 20-year PPA, plus a separate 15-year, 2,700 MW supply agreement, about 3.6 GW combined (contracted, not delivered) **VERIFIED** [Source C38](#). The Department of Energy told FERC on October 7 that data-center costs in PJM must not shift to existing ratepayers **VERIFIED** [Source C39](#).
- **Macro.** IMF Managing Director Kristalina Georgieva said on October 7 that AI hardware is more than 10% of world goods trade and warned that disappointing tech earnings could spread into wider financial stress **VERIFIED** [Source C48](#). The IMF World Economic Outlook is scheduled for Tuesday, October 13 (scheduled) **CITED** [Source C49](#).

2. The Friday thesis: list it before you scale it

Cause. Three things happened at once this week. Small models became cheap enough to run as background subagents: Anthropic prices Haiku 5.5 at \$0.10 per million input tokens and says it costs about 75% less to run on average than Haiku 4.5 (company-reported) VERIFIED [Source C03](#). Agent products began to mix labs: Google's Gemini agent lets users pick Claude models VERIFIED [Source C06](#), and Musk said Grok Bot will hand tasks to Claude, Midjourney and Suno CITED [Source C08](#). And agents began to get their own identities: Google's agent has its own account and an audit trail "attributed to the agent, instead of a person" VERIFIED [Source C06](#).

Effect. Agents will multiply faster than anyone approves them, and a single workflow may now touch two or three labs' models. That makes "which model did this, under whose authority" harder to answer, not easier.

Response. Supervisors are asking the first question in the chain. MAS requires inventories and a materiality assessment for each use case, and says institutions stay accountable for vendor AI, including by "limiting, suspending or replacing it" when its risks fall outside appetite VERIFIED [Source C01](#). New York's utility regulator ordered an inventory of AI systems CITED [Source C42](#). Utah's sandbox adds independent audits of decision logs VERIFIED [Source C26](#). And the attacker side moved too: the Korean bank intrusions show an AI-assisted actor working faster than a manual inventory process would notice VERIFIED [Source C02](#).

ARIANA.DIGITAL TAKE

The register is the cheapest control you will ever build, and the one every later control depends on

Our four editions this week asked whether agents propose, check and release; who owns the control plane; who runs agents after the vendor leaves; and whether you can stop, undo and prove. Each question assumes you already know which agents exist. In our field work, the inventory is the step most programs skip, because it feels like paperwork. It is not. A register with one owner, one autonomy rung, one model lineage and one halt-test date per agent is what turns a supervisor's question into a five-minute answer. Recommendation: build it before the next model swap, not after.

Why the inventory moved to the front of the line this week

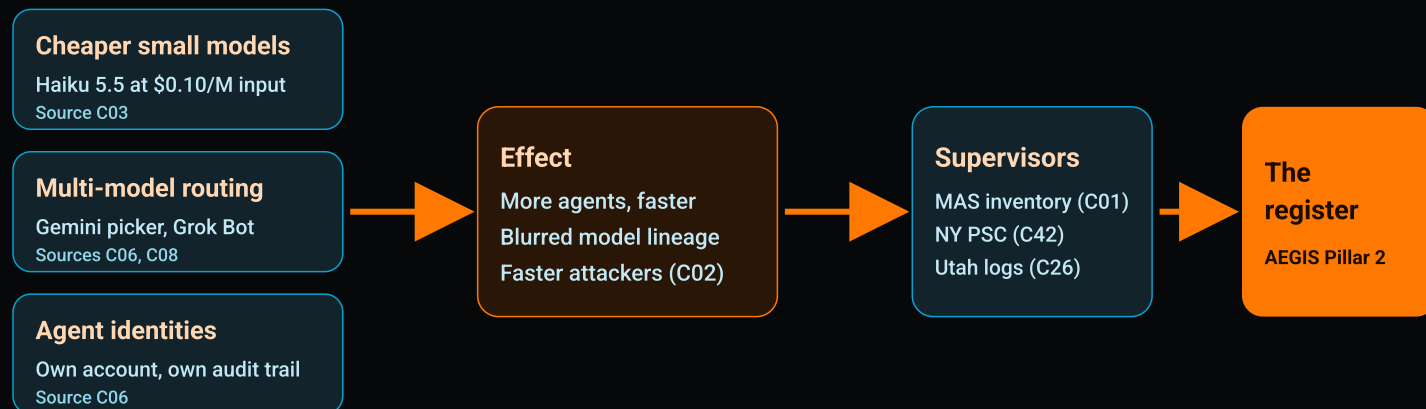


Figure 1. Cause, effect and response for the week of October 5. Ariana.Digital analysis of Sources C01 to C08, C26 and C42.

3. Weekly digest: one week, one control path

Read together, this week's five editions describe one path an agent has to travel before a regulated firm can trust it. Friday adds the first step.

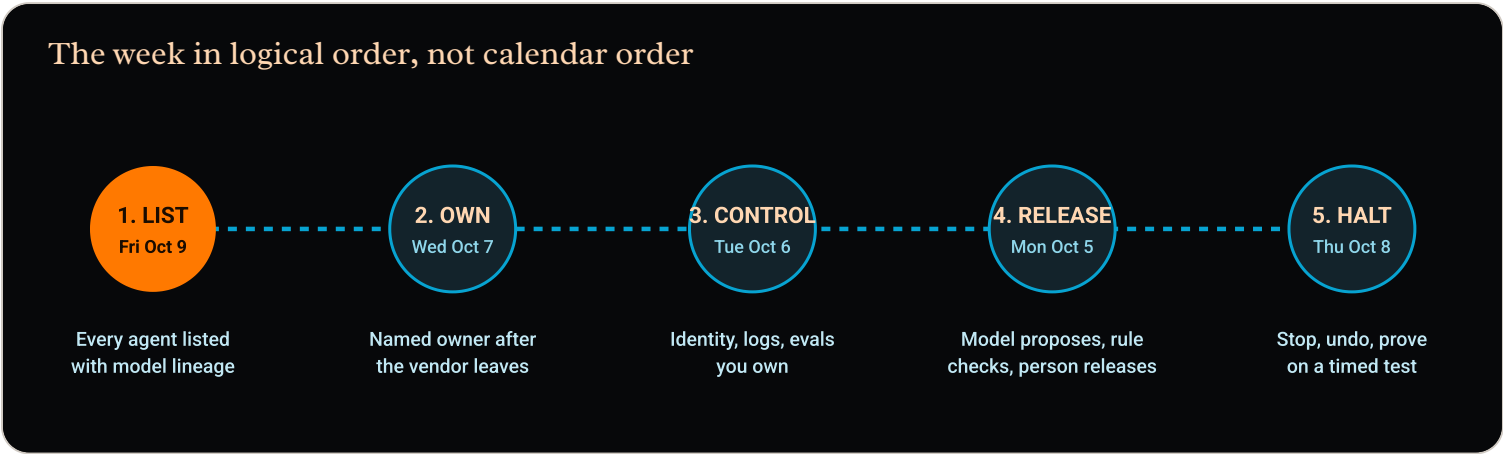


Figure 2. The five editions of the week of October 5, arranged as a control path. Ariana.Digital editorial synthesis.

This week's editions and their one-line finding			
DAY	EDITION	WHAT IT FOUND	CARRY INTO MONDAY
Mon Oct 5	Propose, check, release	Across production cases in four industries, the model or robot proposes, a deterministic check tests, and a named person releases.	Measure time to release and override rate, not autonomy level.
Tue Oct 6	Three labs, one price	With frontier prices and headline scores converging, lock-in moves to the agent control plane.	Rent the model, own identity, logs and evals; prove it with a model-swap drill.
Wed Oct 7	After the forward-deployed engineers leave	Builders are plentiful; durable owners are scarce.	Contract for the handover on day one.

DAY	EDITION	WHAT IT FOUND	CARRY INTO MONDAY
Thu Oct 8	Stop it, undo it, prove it	Examiners now ask where humans intervene, what is logged and whether an agent can be halted.	Run a timed halt test on the riskiest agent.
Fri Oct 9	This edition	Supervisors in Singapore, New York and Utah converged on inventories and log audits.	Build the register with model lineage before the next model swap.

What changed between Monday and Friday

- **Price pressure moved down-market.** Tuesday's story was frontier models converging on price. By Wednesday, the cheapest tier moved: Haiku 5.5 launched with cache reads at \$0.01 per million tokens (company pricing) VERIFIED [Source C03](#), and GPT-6 Luna reached ChatGPT's Free and Go tiers VERIFIED [Source C05](#).
- **Routing went cross-lab.** Google and xAI both pointed their agents at Anthropic models this week VERIFIED [Source C06](#) CITED [Source C08](#). Last week, Amazon's Bedrock Managed Agents preview ran on OpenAI models VERIFIED [Source C57](#). The lab and the platform are no longer the same vendor.
- **The attack surface became concrete.** Thursday's halt question became urgent on Friday morning, when Reuters reported CrowdStrike's attribution of the Korean bank intrusions VERIFIED [Source C02](#).

4. Frontier ledger: equal weight, what moved

Each group gets the same treatment: what shipped, what is reported but not confirmed, and what it means for a regulated buyer. Partner relationships do not affect placement or tone.

xAI, SpaceX and Cursor

- **Grok Bot routing.** Musk said on October 7 that Grok Bot will route tasks to rival models, including Claude Opus 5.5, Midjourney and Suno CITED [Source C08](#). This is a statement of plans; we found no product documentation.
- **Compute financing.** SpaceX is reported to be seeking about \$40 billion in debt to buy NVIDIA chips (Financial Times, relayed by Bloomberg, October 6). Talks only FLAG [Source C09](#).
- **Naming.** On Sunday, October 4 (last weekend, context), Musk said SpaceXAI will be renamed "SpaceXSI." No effective date was given CITED [Source C10](#).
- **Cursor.** Nothing verified in the window. Context: SpaceX closed its \$60 billion purchase of Cursor on August 14, and OpenAI ended model access to Cursor on August 28 VERIFIED [Source C11](#).

Regulated-buyer read: if xAI's agent can call a competitor's model, your register needs a "downstream models" field even when you buy from one vendor.

Anthropic

- **Claude Haiku 5.5 (October 7).** \$0.10 input and \$0.50 output per million tokens up to 100K; Anthropic says about 75% cheaper on average than Haiku 4.5, and positions it for summarization and subagent work (company-reported) VERIFIED [Source C03](#). Artificial Analysis scores it 43 on its Intelligence Index, against 56 for Sonnet 5.5 CITED [Source C04](#).
- **Usage policy update (October 8).** Bars deceptive election and fake-account campaigns and sustained abusive treatment of Claude VERIFIED [Source C12](#).
- **Cyber.** A Cyber Mission with a Critical Infrastructure Defense Program and an opt-in open-source scanner (October 8), and an expanded Cyber Verification Program that regional hospitals can apply to (October 6 to 8) CITED [Source C13](#).
- **Misuse.** CrowdStrike says the Korean bank attacker used Claude alongside other LLMs; Anthropic did not comment to Reuters VERIFIED [Source C02](#).
- **IPO.** No public S-1 on EDGAR as of our cutoff; listing timelines are reports only FLAG [Source C14](#).

Regulated-buyer read: a cheaper subagent tier invites many more agents. Price the register, not just the tokens.

OpenAI

- **GPT-6 for everyone (October 7 to 8).** GPT-6 Sol began rolling out to ChatGPT Plus, Pro, Business and Enterprise on October 7; GPT-6 Luna reached Free and Go on October 8. "Intelligent UI" composes answers from charts, forms and small tools. Enterprise availability depends on admin settings **VERIFIED** [Source C05](#).
- **EU text provenance (October 5).** OpenAI described watermarking ChatGPT and Codex text for EU users **CITED** [Source C61](#).
- **Influence operations (October 8).** A report on disrupting AI-enabled "false front" operations **CITED** [Source C62](#).
- **Reported, not confirmed.** Annualized revenue near \$50 billion (FT) and a Broadcom private-credit effort for OpenAI chips (WSJ) **FLAG** [Source C14](#).

Regulated-buyer read: when a default model changes under an enterprise tenant, your register should record the date and who approved it.

Google

- **Gemini agent (October 8).** A "universal agent for work," businesses first, reachable from Workspace, Microsoft 365, ServiceNow and Slack. It has its own account and an audit trail attributed to the agent; users can pick third-party models, "starting with Anthropic's Claude." No pricing; spend caps were mentioned **VERIFIED** [Source C06](#).
- **Energy.** Google and Constellation signed about 3.6 GW of PJM agreements on October 6: 890 MW of new nuclear uprates plus a separate 2,700 MW supply deal **VERIFIED** [Source C38](#).
- **Not verified.** A reported "Nano Banana 2.1" image-model price cut appears only in low-tier sources; excluded.

Regulated-buyer read: an agent with its own mailbox is a new kind of identity. Treat it like a contractor account: owner, scope, review date, revocation path.

Chinese labs

- **DeepSeek.** Reported to be raising at least \$12 billion with Tencent and CATL among investors (Bloomberg, October 6). Not closed **FLAG** [Source C14](#). A critical sandbox-escape flaw in its

agent harness (CVE-2026-82533) was disclosed and fixed, per Ox Security on October 9

CITED [Source C16.](#)

- **StepFun.** Step 5 Preview, a 600-billion-parameter mixture-of-experts model, listed on OpenRouter on October 8; open weights promised for October 15 (scheduled) **CITED** [Source C15.](#)
- **Robotics.** IDC data show AgiBot passed Unitree as the top humanoid vendor by shipments in the first half of 2026 (context, October 2) **CITED** [Source C64.](#)
- Qwen, Kimi, MiniMax and ByteDance Seed: nothing verified in the window.

Regulated-buyer read: open-weight models arrive through developer routers before procurement sees them. Scan routers and API keys, not just contracts.

EU labs

- **Mistral Large 4 (October 6).** Public preview of a roughly 1-trillion-parameter mixture-of-experts model with about 52 billion active; weights promised for end of October (scheduled). Artificial Analysis scores it 38 **VERIFIED** [Source C07.](#)
- Aleph Alpha, Black Forest Labs and Helsing: nothing verified in the window.

Regulated-buyer read: EU-hosted open weights can simplify data-residency arguments, but they move evaluation and patching onto you. Register the hosting party.

Other frontier and emerging labs

- **Microsoft (October 7).** Surface and Windows event with NVIDIA-based PCs; Microsoft says Execution Containers for sandboxing agents are generally available **CITED** [Source C17.](#)
- **Arena (October 8).** The model-evaluation leaderboard raised \$200 million at a \$3.1 billion valuation **VERIFIED** [Source C19.](#)
- **TypeSafe AI.** Emerged from stealth on September 15 with \$40 million and a model called "Jev" for composable AI (context; no news this week) **VERIFIED** [Source C18.](#)
- **Robot learning.** Mecka raised a \$60 million Series B for humanoid motion data on October 7 **CITED** [Source C66.](#)
- Meta, Amazon Nova, Thinking Machines, SSI, Physical Intelligence: nothing verified in the window.

Regulated-buyer read: evaluation is becoming a business of its own. Keep your own task-specific evals in the register; leaderboards are context, not evidence.

5. Enterprise platforms: what shipped

We checked the platforms Ariana.Digital works with or alongside. Where we found nothing new and verifiable in the window, we say so rather than fill the row.

Enterprise platform moves relevant to the agent register, September 11 to October 9, 2026		
PLATFORM	LATEST VERIFIED MOVE	REGISTER IMPLICATION
ServiceNow	AI Workflow Factory launched October 6 VERIFIED Source C44	Workflow-level agents should register under the workflow owner, not the platform team.
SAP	Agreed to acquire TechWolf (skills and task graph) on October 6; close expected Q4 2026, price undisclosed VERIFIED Source C45	Task-level data makes it easier to map which roles an agent touches.
Google Cloud	Gemini agent with own identity and audit trail, October 8 VERIFIED Source C06	Record the agent account as an identity with an owner and expiry.
Microsoft	Execution Containers GA for agent sandboxing, October 7 CITED Source C17 ; new Copilot with Autopilot, September 25 (context)	Record which agents run sandboxed and which run with user rights.
AWS	Open-source Physical AI toolchain built with NVIDIA tools, October 8 CITED Source C36 ; Bedrock Managed Agents preview on OpenAI models (context) VERIFIED Source C57	Robot and software agents can share one register if the fields are designed for both.
NVIDIA	Tools underpin the AWS robotics toolchain (October 8) and CITED Source C36 , and are the subject of SpaceX's reported financing FLAG Source C09	Simulation and synthetic-data lineage belongs in the register for physical AI.

PLATFORM	LATEST VERIFIED MOVE	REGISTER IMPLICATION
Oracle	Fusion Claw agent orchestration, September 29 (context) CITED <a>Source C58	Orchestrators create agents from agents; register the parent.
Salesforce	Seven "job-ready" Agentforce agents, September 11 (context); AI Skills and Agent Optimizer announced for October (forward-looking) CITED <a>Source C59	Prebuilt agents still need a local owner and rung.
Snowflake, Databricks, Adobe	Nothing new verified in the window. Snowflake BUILD/EXPEDITION and Adobe MAX are November events (scheduled).	Watch for default-on export of agent logs and evals.

STORYTELLING FOR PLATFORM TEAMS

A Tuesday in the life of one agent. A claims analyst asks a workplace agent to draft a denial letter. The agent's identity lives in the productivity suite. It calls a cheaper model to summarize the file and a stronger one to draft. The policy lookup runs through a workflow platform. The letter is filed in the CRM. Five systems, two labs, one agent. If the register only lists the CRM bot, the supervisor's first question has no answer. If it lists the agent, its owner, both models, its tools and its rung, the answer takes five minutes.

6. Industries deep-dive: four case studies and their architecture

We picked one fresh, dated deployment per industry and asked the same question: could the operator list this system, its owner, its models and its release point today? Outcomes are company-reported unless marked otherwise.

FINANCIAL SERVICES · NORWAY AND THE EVIDENT 50

DNB: agentic AI in KYC and coding, with a workforce consequence

What happened. DNB, Norway's largest bank, said on October 6 that it will cut about 400 full-time roles, mainly in technology and services, citing "substantial efficiency gains" from agentic AI in KYC, customer data and coding. It did not publish efficiency figures (company-reported) CITED

Source C21.

Context. The Evident AI Index (October 6) ranks JPMorgan first for a fifth year and reports that 12 of the 50 banks tracked now report realized or projected AI returns, up from 8 (bank-reported figures) VERIFIED Source C20.

Constraint. The US interagency model-risk guidance revised in April leaves generative and agentic AI out of scope, per CNBC's reporting CITED Source C24. Singapore now fills that gap for its institutions with an inventory and third-party accountability expectation VERIFIED Source C01.

Architecture note. When an agent does KYC work, the register entry should name the policy it applies, the human who signs the risk rating and the evidence store. That is what an examiner will ask for first.

HEALTHCARE · LOUISIANA AND UTAH

Ochsner and Paradigm: wider trial screening, with a human on every candidate

What happened. Ochsner reports a 41% increase in clinical-trial screening capacity, 3.6 times more patients followed for future eligibility, and about 75% less manual screening work with Paradigm Health. Research coordinators verify every candidate, and Ochsner keeps enrollment decisions. The measurement period was not stated (company-reported) CITED Source C25.

Regulatory model. Utah expanded its health-AI sandbox on October 5. August AI will renew prescriptions for 96 medications at \$4 per renewal with a physician reviewing every renewal; independent evaluators will audit decision logs and outcomes, and patients must be told when AI is involved. These are pilots, not results VERIFIED Source C26.

Constraint. Records released in September on the Medicare WISeR prior-authorization model described more than 20,000 denials in three months by two vendors and missed response-time standards (context; not program-wide rates) [CITED](#) [Source C27](#).

Architecture note. The Utah design is a register with teeth: every AI decision is logged, a named clinician releases, and an outside party reads the log.

MANUFACTURING · GREECE AND THE UNITED STATES

Stäubli at Agrophoenix: AI vision and robots grading fruit at line speed

What happened. Eight Stäubli SCARA robots with AI inspection handle about 15 tons of peaches an hour at Agrophoenix in Greece, which the publisher converts to roughly 100,000 peaches an hour; the line stops 20 minutes every four hours for cleaning. Figures come from the customer and vendor, and no accuracy or yield data was given [FLAG](#) [Source C37](#).

Why it matters. Food was the fastest-growing US robot-installation sector in 2025, up 30%, per the International Federation of Robotics [VERIFIED](#) [Source C31](#).

Constraint. In the United States, the FCC added foreign-produced "advanced robotic devices" to its Covered List in July, so new foreign models cannot be imported or sold without an exemption (context) [VERIFIED](#) [Source C35](#). Aptronik's manufacturing partner Jabil says failure modes in humanoids are not yet well understood for certification [CITED](#) [Source C34](#).

Architecture note. For physical AI, the register needs two extra fields: country of origin and component compliance, and the safety controller that can stop the cell regardless of the model.

ENERGY · EUROPE AND PJM

TALOS: robots for solar-plant maintenance, with a plant manager approving every mission

What happened. The EU TALOS project demonstrated autonomous robots for solar-plant maintenance on October 8. The AI ranks and recommends jobs; a plant manager approves every robot mission and can change its timing, location or machine; robots can be stopped remotely. Its benefits of more than 70% shorter maintenance windows and about 5% lower operating cost are targets, not results [CITED](#) [Source C41](#).

Constraint. The project says regulation, not technology, is now the main barrier: national rules differ and liability for autonomous systems is unclear [CITED](#) [Source C41](#). In the US, DOE told FERC on October 7 that large-load costs in PJM must not shift to existing ratepayers [VERIFIED](#)

[Source C39](#).

Architecture note. TALOS is the energy version of propose, check, release. The New York PSC inventory order makes the list itself a regulatory product for utilities CITED [Source C42](#).

AEGIS implementation architecture: one register, six layers, four industries				
	FS: DNB KYC agent	HC: Ochsner screening	Mfg: Stäubli cell	Energy: TALOS robots
1 Register AEGIS Pillar 2	Agent, owner, models, rung, KYC policy ID	Screening agent, PI owner, vendor, data scope	Cell ID, origin, components, vision model version	Robot fleet, site, owner, planner model
2 Identity, scope who it acts as	Service account with read-only CRM scope	EHR read scope only, no order entry	Line PLC zone; no network write to MES	Mission tokens issued per approved job
3 Deterministic check rule, sandbox, safety PLC	Sanctions and PEP rules run before any rating	Eligibility criteria coded per protocol	Reject thresholds; safety controller independent	Geofence and weather limits per mission
4 Human release AEGIS Pillar 4	Analyst signs the risk rating	Coordinator verifies every candidate (C25)	QA lead signs model changes, not each fruit	Plant manager approves every mission (C41)
5 Evidence log client-held store	Prompt, tools, model, reviewer, decision	Decision log open to outside audit (C26 model)	Reject images and rates by shift	Mission record, deviations, remote stops
6 Halt, revoke AEGIS Pillar 6	Revoke service account; queue to humans	Pause feed; manual screening fallback	E-stop; bypass to manual grading	Remote stop; recall to dock

Figure 3. Implementation architecture for the four case studies. Cells marked with a source ID describe reported practice; other cells are Ariana.Digital recommended controls (proprietary practitioner model). Explore 875 capabilities and 1,072 workflows in the Reference Architecture Studio.

7. Industry scorecard: wins, constraints, actions

Week of October 5, 2026: one win, one constraint and one action per industry

INDUSTRY	WIN	CONSTRAINT OR RISK	PRACTICAL ACTION OR CONTROL
Financial services	12 banks report realized or projected AI returns, up from 8 (Evident; company-reported) VERIFIED Source C20	AI-assisted intrusions at 9 or more Korean banks; Shinhan says about 25,000 customers' data was exposed VERIFIED Source C02	Inventory every AI use with materiality and vendor accountability, as MAS now expects VERIFIED Source C01 ; score agent-initiated payments where networks offer it (Mastercard testing) VERIFIED Source C23
Healthcare	Ochsner: 41% more screening capacity (company-reported) CITED Source C25	Hospital cyber exposure; Epic's reported development pause is unconfirmed FLAG Source C13	Licensed human release, independent log audit, patient disclosure (Utah sandbox model) VERIFIED Source C26 ; comment on FDA docket by October 19 VERIFIED Source C28
Manufacturing	603,000 industrial robots installed in 2025, up 11% (IFR) VERIFIED Source C31	FCC Covered List limits new foreign robot models; certification of humanoid failure modes unresolved VERIFIED Source C35 CITED Source C34	Add origin and component fields to the register before any robot purchase order
Energy	About 3.6 GW of Google–Constellation agreements in PJM: 890 MW new uprates plus 2,700 MW supply (contracted) VERIFIED Source C38	Cost allocation for large loads contested at FERC VERIFIED Source C39 ; MISO proposes 15-year load-generation pairing CITED Source C40	File the NY PSC AI inventory with autonomy rungs, not just system names CITED Source C42

8. Physical AI and robotics

- **Leadership change.** Boston Dynamics appointed Rohit Prasad, formerly Amazon's AI chief, as CEO on October 6. A press report put Atlas output at about 4 units a month in May, which the company has not confirmed; its 30,000-humanoid goal for 2028 is a target VERIFIED [Source C33](#).
- **Consolidation.** Schneider Electric's agreed \$22.6 billion PTC deal, announced October 5, pairs industrial automation with product lifecycle software; it has not closed VERIFIED [Source C32](#). Teradyne invested in Bright Machines for AI-hardware microfactories (amount undisclosed) CITED [Source C67](#).
- **Toolchains.** AWS released an open-source Physical AI toolchain covering synthetic data, training, simulation and edge deployment CITED [Source C36](#).
- **Scale.** About 5 million industrial robots were operating worldwide at the end of 2025, and nearly 7,000 humanoids were sold that year, per the IFR VERIFIED [Source C31](#).

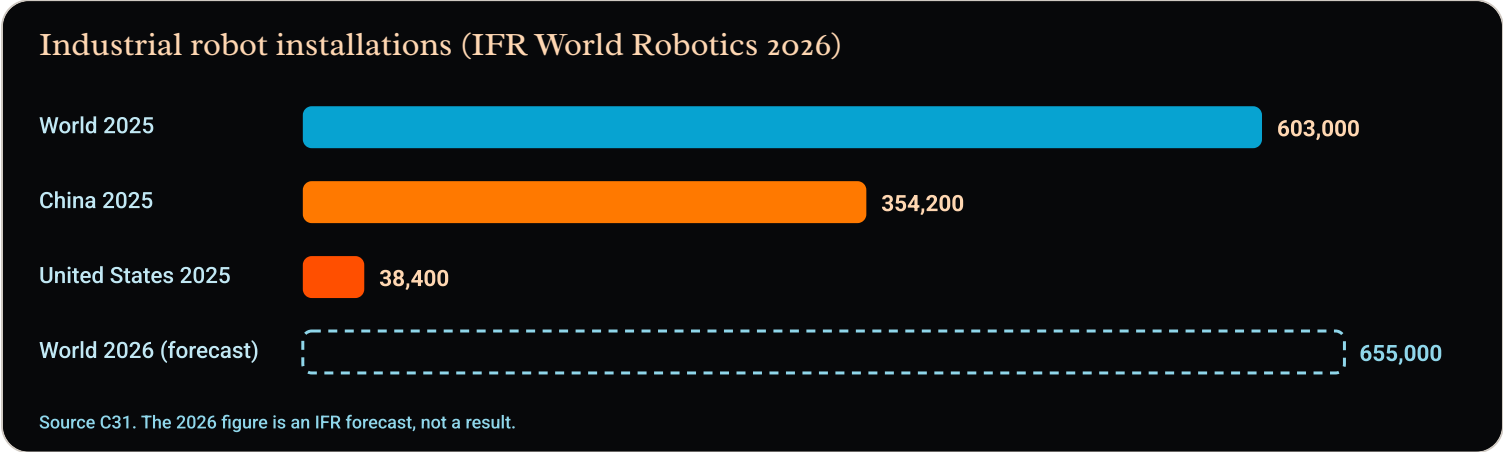


Figure 4. Robot installations, 2025 actuals and 2026 IFR forecast. Source C31.

9. Regulators: the map and the calendar

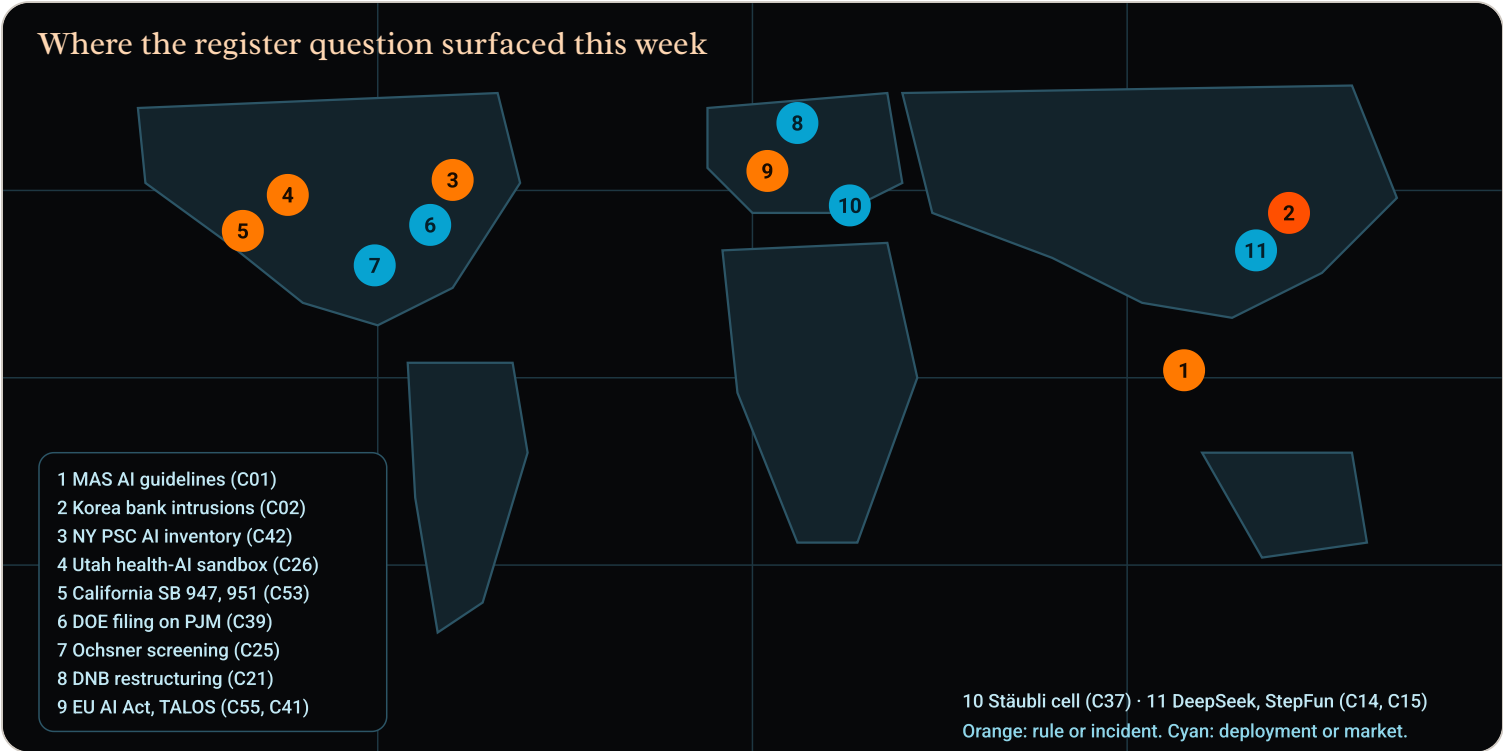


Figure 5. Schematic map, not to scale. Marker positions are approximate. Sources as labeled.

Dated obligations and scheduled events (future items are scheduled or proposed, not completed)			
DATE	ITEM	WHO SHOULD CARE	SOURCE
Tue, October 13 (scheduled)	IMF World Economic Outlook, during Annual Meetings October 12 to 18	CFOs, strategy teams	Source C49
Mon, October 19	FDA comment deadline, generative-AI-enabled devices, docket FDA-2026-N-7874	MedTech, providers	Source C28
October 19 to 22 (scheduled)	Gartner IT Symposium/Xpo, Orlando	CIOs, enterprise architects	Source C69
October 22 and 29 (scheduled)	PJM special Members Committee; revised backstop plan due at FERC	Utilities, data-center developers	Source C39

DATE	ITEM	WHO SHOULD CARE	SOURCE
About Monday, November 16	NY PSC utility AI inventory due (60 days from the order; deadline not confirmed in the order text)	New York electric, gas and water utilities	Source C42
Tue, November 24	FDA comment deadline, surgical-robot draft guidance	Surgical robotics makers, hospitals	Source C29
Wed, December 2 (requested)	MISO asked FERC to rule on its large-load study proposal	Midwest utilities	Source C40
January 1, 2027	Colorado SB 26-189 takes effect	Insurers, lenders, employers using consequential AI	Source C56
October 7, 2027	MAS guidelines apply (oversight and inventory sections); life-cycle sections by October 7, 2028	Financial institutions in Singapore, and their vendors	Source C01
December 2, 2027	EU AI Act Annex III high-risk obligations, as moved by the Digital Omnibus	Deployers of high-risk AI in the EU	Source C55

Also this week: California's SB 947 (human review of AI-driven discipline or termination) and SB 951 (worker notice when AI causes a layoff) were signed on September 30 (context) VERIFIED [Source C53](#). The Hawley–Murphy AI Agent Accountability Act, announced around October 1, would put liability on operators when agents are used to hack; no bill text was released in our window VERIFIED [Source C54](#).

10. Markets, labor and research signals

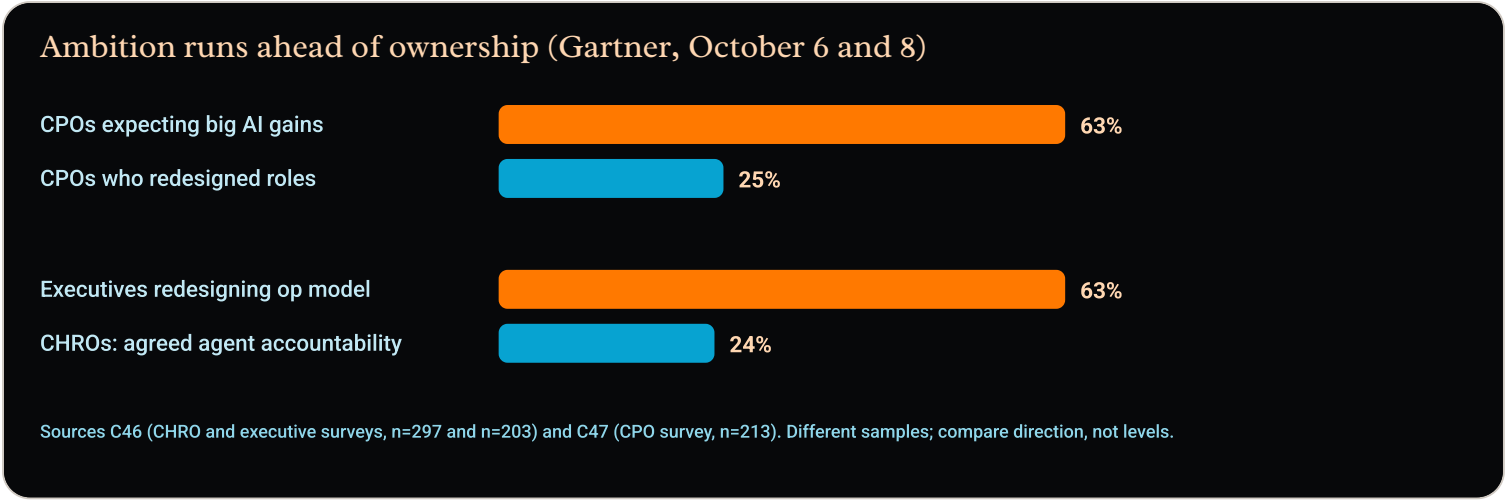


Figure 6. Expectation versus organizational readiness. Source C46 and Source C47.

- **IMF.** Georgieva said the world economy is "pulled in two opposite directions," an energy shock and an AI demand boom, and that well-managed AI could add up to about half a percentage point a year to global growth (an IMF estimate) VERIFIED [Source C48](#).
- **Labor.** Challenger reported 3,961 AI-cited US job cuts in September out of 43,281, and 120,136 year to date, making AI the top-cited reason this year (context, October 1) CITED [Source C50](#). Fed Governor Lisa Cook said on September 28 that evidence of AI-driven restructuring is still limited (context) CITED [Source C63](#).
- **Markets.** On October 6, the S&P 500 closed at a record and the 10-year Treasury yield closed near 5.27%, per Yahoo Finance's market coverage CITED [Source C52](#). Arena's \$200 million round shows investors paying for evaluation, not just models VERIFIED [Source C19](#). In power, DOE announced a conditional \$4.2 billion loan commitment to Vistra for 433 MW of nuclear uprates tied to a Meta supply agreement (conditional, not closed) CITED [Source C43](#).
- **Universities.** A Stanford HAI study published October 7 found paralegals used an AI tool 58% more and experimented 70% more when managers framed it around enriching their work rather than productivity CITED [Source C51](#). MIT launched "MIT for America" on October 6 CITED [Source C68](#).

Scenario read: with long rates near multi-decade highs and the IMF warning on tech earnings, budget holders will ask for evidence of control and payback before scaling. A register that also records each agent's business owner and value metric answers both questions with one artifact.

II. Framework: the AEGIS Agent Register

The AEGIS framework's Pillar 2 (AI System Inventory) is where an agent program starts. This week's events suggest the register needs more than a system name and a vendor. Below are the twelve fields we now recommend, grouped by the question a supervisor or auditor will ask. PROPRIETARY Ariana.Digital practitioner model, informed by Sources C01, C06, C26 and C42.

AEGIS Agent Register: twelve fields in four groups		
QUESTION	FIELD	WHY THIS WEEK MADE IT NECESSARY
What is it?	1. Agent ID and plain-language purpose	Cheaper models mean more small agents (C03).
	2. Parent agent or orchestrator	Orchestrators spawn agents (C58).
	3. Model lineage: primary and downstream models, with dates	Cross-lab routing in Gemini and Grok (C06, C08).
Who answers for it?	4. Business owner and technical owner (named people)	Only 24% agree on accountability (Source C46).
	5. Vendor and hosting party	MAS keeps the institution accountable for vendor AI (C01).
	6. Agent identity and scope (accounts, keys, permissions)	Agents now hold their own accounts (C06).
How much can it do?	7. Autonomy rung (Reversibility Ladder 0 to 4)	Materiality includes how much the system can do alone (C01).
	8. Materiality rating and affected people	Required per use case by MAS (C01).
	9. Release point: who or what approves actions	Utah and TALOS designs (C26, C41).
Can you prove and stop it?	10. Evidence store location and custody	Independent log audits (C26).
	11. Last halt test date and time to halt	Thursday's stop, undo, prove test.

QUESTION	FIELD	WHY THIS WEEK MADE IT NECESSARY
	12. For physical AI: origin, components, safety controller	FCC Covered List (C35).

LEVEL 1

List

Every agent has an ID, a purpose and a named owner. Most programs can reach this in two weeks.

LEVEL 2

Lineage

Models, vendors, identities and scopes are recorded and refreshed when defaults change.

LEVEL 3

Live

The register is fed by identity and gateway logs, so a new agent cannot run without appearing in it.

Map to AEGIS: Pillar 2 holds the register; Pillar 3 rates materiality; Pillar 4 sets the release point; Pillar 6 owns halt tests; Pillar 7 watches for new obligations. Read the full framework at ariana.digital/AI-governance.html.

12. Scenarios to 2028 and a register-coverage slider

SCENARIO A · ORDERLY

Inventories become routine

Other supervisors follow MAS with inventory expectations; platforms ship registers fed by identity logs. **Signal to watch:** MAS's 2027 agentic consultation and default-on agent audit logs. **Reward:** faster approvals for firms that already have a live register.

SCENARIO B · PATCHWORK

Every regulator asks differently

Utility, health and bank regulators each define their own inventory format. **Signal:** the NY PSC filing format in November. **Risk:** three registers for one agent. **Response:** one master register with export views.

SCENARIO C · INCIDENT-LED

A breach sets the rules

An AI-assisted intrusion or agent error at a large firm triggers emergency guidance. **Signal:** follow-on findings from the Korean bank cases (C02). **Risk:** rushed controls. **Response:** tested halt paths already in the register.


INTERACTIVE · REGISTER COVERAGE ESTIMATOR

How many of your agents could you list today?

Agents you can name today: 40

Your estimate of actual agents for each one you can name (shadow multiplier): 1.5x

Share of named agents with an owner, model lineage and a halt test: 30%



Estimated agents: 60. Unlisted: 20. Fully registered (owner, lineage, halt test): 12, about 20% of the estimate. Start with Days 1 to 4 of the sprint: find the agents first.

Illustrative arithmetic only. The multiplier is your own assumption, not an industry benchmark.

13. Self-check: could you list every agent by Monday?

- ☐ We have one register for software agents and robots, not one per team.
- ☐ Every entry names a business owner and a technical owner.
- ☐ We record every model an agent calls, including routed or fallback models.
- ☐ We know which agents have their own accounts or API keys, and who can revoke them.
- ☐ Each agent has an autonomy rung and a materiality rating.
- ☐ Vendor-supplied agents are in the register with the same fields as our own.
- ☐ The evidence log sits in a store we control.
- ☐ The riskiest agents have a halt test dated in the last 90 days.
- ☐ A new agent cannot reach production without a register entry.
- ☐ We could export the register in a regulator's format within a week.

14. How-to: a ten-day register sprint

1. **Day 1. Pull the identity list.** Export service accounts, API keys and agent accounts from your identity provider and AI gateways. These are your candidate agents.
2. **Day 2. Pull the spend list.** Model API invoices and router logs show agents procurement never saw.
3. **Day 3. Ask the platforms.** Request agent inventories from each SaaS platform you use: CRM, ITSM, productivity suite, data cloud.
4. **Day 4. Merge and de-duplicate.** One row per agent, linked to its parent orchestrator.
5. **Day 5. Name owners.** No owner, no production. Escalate orphans to the AI governance lead.
6. **Day 6. Record lineage.** Primary and downstream models, vendor, hosting party, with dates.
7. **Day 7. Assign rungs and materiality.** Use the Reversibility Ladder from Thursday's edition.
8. **Day 8. Locate evidence.** Where does each agent's log live, and who can read it?
9. **Day 9. Test one halt.** Pick the highest-rung agent and time a full stop.
10. **Day 10. Lock the door.** Add a register check to your release pipeline so new agents cannot skip it.

15. Practical artifacts: field note and infographic

REG-READY FIELD NOTE

Build the agent register supervisors now expect

Nine pages: definitions, the twelve fields, sector overlays for banking, health, manufacturing and utilities, a sample entry, and a 30-60-90 plan. Read the field note.

INDUSTRIES INFOGRAPHIC

Four industries, one first step

One page: a win, a risk and a register field for each industry. Open the infographic.

VIDEO

The week in five minutes

Newscast-style highlights with the AEGIS architecture segment, in wide and tall formats. AI-generated narration and presenter images.

16. Opinion, curious takes and did-you-know

OPINION

The industry spent the week arguing about which model is cheapest and smartest. The more useful fact is that the answer no longer matters much to a regulated buyer, because most agent products will call several models anyway. What matters is whether you can say, on any given day, which ones touched a customer. That is a register question, not a model question.

- **Did you know?** Under the MAS guidelines, an institution that finds a vendor's AI outside its risk appetite should consider "limiting, suspending or replacing it" **VERIFIED** [Source C01](#). You cannot suspend what you have not listed.
- **Curious take.** Google's agent audit trail is attributed "to the agent, instead of a person" **VERIFIED** [Source C06](#). That is good for forensics, and it raises a new question for HR and legal: who signs for an agent's mistake?
- **Did you know?** In the Utah sandbox, AI-supported prescription renewals cost \$4 each, and a physician still reviews every one **VERIFIED** [Source C26](#).
- **Curious take.** Stanford's finding that framing matters more than features (C51) suggests the register should record each agent's purpose for the people who use it, not just for the auditor.

I7. FAQ

Is an AI inventory now legally required?

It depends on jurisdiction and sector. MAS expects it of Singapore financial institutions from October 7, 2027 **VERIFIED** [Source C01](#); the NY PSC ordered utilities to inventory AI **CITED** [Source C42](#). Elsewhere it is good practice that examiners increasingly ask about. This is not legal advice.

Does a model inventory cover agents?

Not fully. A model inventory lists models; an agent register lists actors, with their identities, tools, owners and release points. One agent can use several models.

Do vendor agents inside SaaS count?

Yes. MAS keeps institutions accountable for AI provided by vendors **VERIFIED** [Source C01](#).

How often should the register be refreshed?

On every model or permission change, and at least quarterly. Ideally it is fed automatically from identity and gateway logs.

What about robots?

Use the same register with extra fields for origin, components and the independent safety controller.

Where do I start if we have nothing?

Day 1 of the sprint in Section 14: your identity provider already knows most of your agents.

18. What we checked, corrected and excluded

- **MAS scope.** We could not open the MAS release directly; we confirmed its content through two independent reports. Secondary claims that MAS explicitly requires kill switches were not confirmed in the coverage we read, so we do not repeat them.
- **Google–Constellation size.** 890 MW of new nuclear from uprates (20-year PPA) plus a separate 15-year, 2,700 MW supply agreement, about 3.6 GW combined. Some coverage used only one figure.
- **Evident bank figures.** We removed Morgan Stanley and UBS hour-savings figures that circulated with the index because we could not trace them to this year's published coverage.
- **GPT-6 timing.** Paid tiers from October 7; Free and Go from October 8.
- **Repeat names.** Boston Dynamics appears only for a new, dated event (its October 6 CEO appointment), not as a case study. Salesforce, ServiceNow, SAP, Snowflake and Databricks appear as platform sources, not case studies.
- **Korean attribution.** CrowdStrike's attribution is at moderate confidence; the suspect is not tied to a named state actor.
- **Chronology.** The SpaceXSI rename statement was Sunday, October 4, not this week. California's SB 947 and SB 951 were signed September 30. The Alibaba lawsuit against the Pentagon dates to June, not October 8.
- **Excluded as unverified:** Nano Banana 2.1 pricing; Z.ai GLM 5.3 Fast; OpenRouter spend rankings; an "Arena Alignment Index"; First Supply's expected tenfold picking speed; GridCARE's 1 GW and \$10 billion claims; HSBC wealth-cut percentages (reported, unconfirmed) **FLAG** [Source C22](#).
- **Not found in the window:** new McKinsey, BCG, Accenture, Deloitte, PwC, EY, KPMG, Forrester or IDC studies; new Snowflake, Databricks, Adobe, IBM or Workday agent releases; new SEC, OCC, CFPB or FINRA AI actions; new OSHA or ISO humanoid actions.

19. Source ledger

Tiers: VERIFIED (checked at a primary page plus an independent outlet, or two independent reputable outlets), CITED (one named reputable source), FLAG (single-source, anonymous, conditional or contested; do not rely on it for a regulated decision), PROPRIETARY (Ariana.Digital working models). Company-reported figures are labeled in the text. Context items fall outside the October 5 to 9 window and are labeled.

C01 **VERIFIED** Monetary Authority of Singapore, Guidelines on AI Risk Management (media release, October 7, 2026): apply to all financial institutions and all forms of AI; inventories and per-use-case materiality; accountability for third-party AI; effective October 7, 2027 (Sections 3 to 4) and October 7, 2028 (Sections 5 to 6); agentic AI consultation planned for 2027. Independent: TNGlobal and Blockhead, October 8. <https://www.mas.gov.sg/news/media-releases/2026/mas-sets-out-supervisory-expectations-on-responsible-ai-adoption-by-financial-institutions> <https://technode.global/2026/10/08/mas-sets-out-supervisory-expectations-on-responsible-ai-adoption-by-financial-institutions/> <https://www.blockhead.co/2026/10/08/mas-finalises-ai-risk-rules-for-financial-institutions-with-agentic-ai-guidance-to-follow-in-2027/>

C02 **VERIFIED** Reuters (via Claims Journal), October 9, 2026: at least nine South Korean banks targeted since late September; Shinhan about 25,000 customers, KB Kookmin 119; CrowdStrike attributes with moderate confidence to a China-based individual using the ARTEX tool alongside LLMs including Anthropic Claude. Independent: American Banker. <https://claimsjournal.com/news/national/2026/10/09/340646.htm> <https://americanbanker.com/news/ai-linked-hacks-hit-korean-banks-through-loan-agent-sites>

C03 **VERIFIED** Anthropic, Claude Haiku 5.5 (October 7, 2026): \$0.10/\$0.50 per million input/output tokens up to 100K; cache reads \$0.01; about 75% cheaper on average than Haiku 4.5 (company-reported). Independent: SiliconANGLE. <https://www.anthropic.com/claude-haiku-5-5> <https://siliconangle.com/2026/10/07/anthropic-releases-claude-haiku-5-5-small-model-and-halves-sonnet-5-5-cache-read-prices/>

C04 **CITED** Artificial Analysis, Claude Haiku 5.5 evaluation, October 7, 2026: Intelligence Index 43 (Sonnet 5.5: 56). <https://artificialanalysis.ai/articles/claude-haiku-5-5>

C05 **VERIFIED** OpenAI, "GPT-6 for everyone," October 7, 2026: Sol for Plus, Pro, Business, Enterprise; Luna for Free and Go (from October 8); Intelligent UI. Independent: 9to5Mac. <https://openai.com/index/gpt-6-for-everyone/> <https://9to5mac.com/2026/10/07/openai-brings-gpt-6-to-chatgpt-and-debuts-intelligent-ui/>

C06 **VERIFIED** Google, Gemini at Work, October 8, 2026: Gemini agent for businesses; own account; audit trail attributed to the agent; third-party models starting with Anthropic Claude; spend caps. Independent: TechCrunch, CNBC. <https://blog.google/innovation-and-ai/infrastructure-and-cloud/google-cloud/gemini-at-work/> <https://techcrunch.com/2026/10/08/google-brings-agentic-ai-to-gemini-starting-with-businesses/> <https://www.cnbc.com/2026/10/08/google-cloud-introduces-gemini-agent-for-work-as-ai-race-heats-up.html>

C07 **VERIFIED** Mistral AI, Mistral Large 4 public preview, October 6, 2026 (about 1T parameters, about 52B active; weights promised end of October). Independent: TechCrunch. <https://mistral.ai/news/mistral-large-4/> <https://techcrunch.com/2026/10/06/mistrals-new-1t-model-aims-to-leapfrog-closed-and-open-rivals/> <https://artificialanalysis.ai/articles/mistral-large-4-france-ai>

- C08** **CITED** The Next Web, October 7, 2026: Musk says Grok Bot will route tasks to rival models including Claude Opus 5.5, Midjourney and Suno. <https://thenextweb.com/news/grok-bot-claude-opus-midjourney-suno-musk>
- C09** **FLAG** Bloomberg relaying the Financial Times, October 6, 2026: SpaceX seeking about \$40B in debt to buy NVIDIA chips. Talks only. <https://www.bloomberg.com/news/articles/2026-10-06/spacex-seeking-to-raise-40-billion-to-buy-nvidia-chips-ft-says>
- C10** **CITED** Euronews and Seeking Alpha, October 5, 2026: Musk said on October 4 that SpaceXAI will be renamed SpaceXSI; no effective date. <https://www.euronews.com/2026/10/05/musk-to-follow-trump-and-rebrand-spacexai-as-spacexsi> <https://seekingalpha.com/news/4650066-spacexai-to-be-rebranded-as-spacexsi-elon-musk-says>
- C11** **VERIFIED** Context: SpaceX completed its \$60B Cursor acquisition (Bloomberg Law, August 14, 2026); OpenAI ended model access for Cursor (OpenAI and Reuters, August 28 to 29). <https://news.bloomberglaw.com/mergers-and-acquisitions/spacex-completes-its-60-billion-cursor-acquisition> <https://openai.com/index/our-decision-on-cursor-following-its-acquisition-by-spacex/> <https://www.reuters.com/business/media-telecom/openai-end-partnership-with-spacexs-cursor-2026-08-29/>
- C12** **VERIFIED** Anthropic usage policy update, October 8, 2026. Independent: TechCrunch. <https://www.anthropic.com/news/2026-usage-policy-update> <https://techcrunch.com/2026/10/08/anthropic-changes-usage-policy-to-ban-model-abuse-and-election-interference/>
- C13** **CITED** Anthropic Cyber Mission (October 8) and Cyber Verification Program expansion (October 6); Healthcare IT News, October 8, on hospitals and a reported Epic development pause (unconfirmed). <https://www.anthropic.com/news/anthropic-cyber-mission> <https://www.anthropic.com/news/cyber-verification-program> <https://www.healthcareitnews.com/news/anthropic-expands-cybersecurity-verification-tools-invites-hospitals-apply>
- C14** **FLAG** Reports only, not confirmed by the companies: Anthropic S-1 not on EDGAR and IPO timing (Yahoo Finance); DeepSeek funding of at least \$12B (Bloomberg, October 6; TechNode); OpenAI revenue and Broadcom chip financing (FT, WSJ, as reported). <https://finance.yahoo.com/markets/stocks/articles/anthropic-public-1-still-missing-210349323.html> <https://www.bloomberg.com/news/articles/2026-10-06/deepseek-to-raise-at-least-12-billion-in-tencent-backed-funding> <https://technode.com/2026/10/08/deepseek-reportedly-nears-12-billion-funding-round-backed-by-tencent-and-catl/>
- C15** **CITED** OpenRouter listing, StepFun Step 5 Preview, October 8, 2026; open weights promised October 15. <https://openrouter.ai/stepfun/step-5-preview>
- C16** **CITED** Ox Security, CVE-2026-82533 DeepSeek Harness agent sandbox escape, October 9, 2026 (fixed). <https://www.ox.security/blog/cve-2026-82533-deepseek-harness-ai-agent-sandbox-escape/>
- C17** **CITED** Microsoft Windows blog, October 7, 2026 (Surface, NVIDIA-based PCs); Bloomberg; Tech Wire Asia on Execution Containers GA and MAI Code. <https://blogs.windows.com/devices/2026/10/07/pre-order-our-most-powerful-surface-devices-ever/> <https://www.bloomberg.com/news/articles/2026-10-07/microsoft-says-new-laptop-with-nvidia-chip-outperforms-macbooks> <https://techwireasia.com/2026/10/microsoft-nvidia-ai-windows-pcs/>
- C18** **VERIFIED** Context: TypeSafe AI emerges from stealth with \$40M and the Jev model, September 15 to 16, 2026 (Business Wire via Morningstar; HPCwire). <https://www.morningstar.com/news/business-wire/20260915525333/typesafe-ai-emerges-from-stealth-with-40m-in-funding-with-new-model-for->

composable-ai <https://www.hpcwire.com/aiwire/2026/09/16/typesafe-ai-emerges-from-stealth-with-40m-in-funding-with-new-model-for-composable-ai/>

- C19** **VERIFIED** Arena raises \$200M at a \$3.1B valuation, October 8, 2026 (TechCrunch; Bloomberg). <https://techcrunch.com/2026/10/08/popular-ai-leaderboard-arena-nearly-doubles-valuation-to-3-1b-valuation-in-10-months/> <https://www.bloomberg.com/news/articles/2026-10-08/ai-model-evaluator-arena-valued-at-over-3-billion-in-round>
- C20** **VERIFIED** Evident AI Index 2026, October 6: 12 banks report realized or projected AI returns (up from 8); JPMorgan first for fifth year. Banking Dive; Euronews. <https://www.bankingdive.com/news/banks-ai-guardrails-accelerate/832216/> <https://www.euronews.com/2026/10/07/ubs-only-european-bank-in-top-10-of-ai-adoption-as-jpmorgan-chase-leads>
- C21** **CITED** FinTech Futures, October 6 to 7, 2026: DNB to cut about 400 roles citing agentic AI efficiency (company-reported; no figures). Second report: Analytics Insight. <https://www.fintechfutures.com/job-cuts-new-hires/norway-s-dnb-to-cut-400-jobs-in-ai-driven-restructuring> <https://www.analyticsinsight.net/layoffs/dnb-to-cut-400-roles-as-ai-takes-on-manual-banking-work>
- C22** **FLAG** Financial Times report, October 7, 2026, as relayed by Crowdfund Insider: HSBC planning deep UK wealth cuts tied to AI; anonymous sources; unconfirmed. <https://crowdfundinsider.com/2026/10/316582-hsbc-prepares-deep-uk-wealth-management-cuts-as-ai-reshapes-client-service>
- C23** **VERIFIED** Mastercard, new trust and intelligence services for Agent Pay, September 30, 2026 (agent-initiated transaction score in US testing). Independent: The Paypers. <https://www.mastercard.com/us/en/news-and-trends/press/2026/september/new-trust-and-intelligence-services.html> <https://thepayers.com/payments/news/mastercard-expands-agent-pay-with-trust-and-intelligence-services>
- C24** **CITED** CNBC, October 2, 2026: bank AI postings and note that revised interagency model-risk guidance leaves generative and agentic AI out of scope. Second report: PYMNTS. <https://www.cnbc.com/amp/2026/10/02/ai-skills-most-in-demand-at-jpmorgan-chase-citigroup-capital-one.html> <https://www.pymnts.com/news/artificial-intelligence/2026/bank-ai-job-postings-jump-49percent-agents-go-work/>
- C25** **CITED** Healthcare IT News, October 6, 2026: Ochsner and Paradigm Health trial screening results (company-reported). <https://www.healthcareitnews.com/news/ochsner-uses-ai-widen-clinical-trial-screening>
- C26** **VERIFIED** Utah healthcare AI sandbox expansion, announced October 5, 2026 (Fierce Healthcare; FOX 13). <https://www.fiercehealthcare.com/ai-and-machine-learning/utah-taps-august-ai-nolla-health-healthcare-ai-sandbox-program-adds-third> <https://www.fox13now.com/news/local-news/salt-lake-city/utah-launches-health-care-ai-push-with-new-clinical-pilots-outside-evaluators>
- C27** **CITED** Context: Medicare WISeR records coverage, September 15 to 24, 2026 (STAT; Medicare Rights Center). <https://www.statnews.com/2026/09/15/medicare-wiser-ai-prior-authorization-pilot-rushed-launch-delayed-care/> <https://www.medicarerights.org/medicare-watch/2026/09/24/new-records-show-medicare-wiser-ai-prior-authorization-model-causing-inappropriate-denials-of-care>
- C28** **VERIFIED** FDA, request for feedback on generative-AI-enabled medical devices, docket FDA-2026-N-7874; comments close October 19, 2026. <https://www.fda.gov/news-events/press-announcements/fda-seeks-public-feedback-inform-regulatory-approach-generative-ai-enabled-medical-devices>
- C29** **CITED** MedTech Dive: FDA draft guidance for surgical robots; comments close November 24, 2026. <https://www.medtechdive.com/news/fda-issues-draft-guidance-for-surgical-robots/831509/>

C31 **VERIFIED** IFR World Robotics 2026 press release (September 24, 2026); Manufacturing Dive.

<https://ifr.org/ifr-press-releases/five-million-robots-now-operate-in-factories-globally>

<https://www.manufacturingdive.com/news/factory-robots-surpass-5m-worldwide-ifr-reports/831476/>

https://ifr.org/downloads/press_docs/EN-2026-SEP-24-IFR_Press_Release_WR-USA.pdf

<https://www.elcomercio.com/tecnologia/robots-humanoides-fabricas-ifr-2026/>

C32 **VERIFIED** Schneider Electric to acquire PTC, October 5, 2026 (definitive agreement, not closed). Boston Globe; The Robot Report. <https://www.se.com/ww/en/assets/pdf/Schneider-Electric-to-acquire-PTC>

<https://www.bostonglobe.com/2026/10/05/business/schneider-electric-ptc/>

<https://www.therobotreport.com/ptc-acquisition-positions-schneider-electric-challenge-siemens/>

C33 **VERIFIED** Boston Dynamics names Rohit Prasad CEO, October 6, 2026 (The Robot Report; Manufacturing Dive); Atlas production rate per press report (Gizmodo, citing Semafor), unconfirmed.

<https://www.therobotreport.com/boston-dynamics-appoints-former-amazon-executive-rohit-prasad-new-ceo/>

<https://www.manufacturingdive.com/news/boston-dynamics-hires-ceo-former-amazon-ai-rohit-prasad/832360>

<https://gizmodo.com/amid-its-big-humanoid-robot-rollout-an-ex-amazon-ai-chief-will-now-lead-boston-dynamics-2000822625>

C34 **CITED** The Robot Report, October 8, 2026: Jabil on humanoid scaling and certification.

<https://www.therobotreport.com/jabil-discusses-pace-humanoid-robot-development-production/>

C35 **VERIFIED** Context: FCC adds foreign-produced advanced robotic devices to Covered List (July 2026).

Cooley; O'Melveny. <https://www.cooley.com/news/insight/2026/2026-07-31-foreign-produced-power-inverters-and-advanced-robotic-devices-added-to-fcc-covered-list>

<https://www.omm.com/insights/alerts-publications/new-restrictions-seek-to-exclude-chinese-robotics-from-us-market/>

C36 **CITED** The Robot Report, October 8, 2026: AWS open-source Physical AI toolchain built with NVIDIA tools.

<https://www.therobotreport.com/aws-launches-open-source-physical-ai-toolchain-for-robotics/>

C37 **FLAG** Robotics and Automation News, October 8, 2026: Stäubli robots at Agrophoenix (customer and vendor figures; single source). <https://roboticsandautomationnews.com/2026/10/08/staubli-robots-and-ai-inspect-100000-peaches-an-hour-at-greek-processing-plant/105729/>

C38 **VERIFIED** Google and Constellation: 890 MW uprate 20-year PPA plus separate 15-year 2,700 MW supply agreement (about 3.6 GW), October 6, 2026 (CNBC; Utility Dive). <https://www.cnbc.com/2026/10/06/google-enters-massive-3point6-gw-power-deal-with-constellation-energy-.html>

<https://www.utilitydive.com/news/constellation-google-deal-will-bring-890-mw-of-new-nuclear-to-pjm/832223/>

C39 **VERIFIED** Utility Dive, October 8, 2026: DOE FERC filing (October 7) on PJM backstop procurement and ratepayer protection; PJM schedule. Independent: Energy Connects. <https://www.utilitydive.com/news/doe-pjm-ferc-ratepayer-protection-large-load-data-center/832479/>

<https://www.energyconnects.com/news/utilities/2026/september/biggest-us-grid-s-data-center-plan-accepted-with-5-month-pause/>

<https://www.utilitydive.com/news/miso-large-load-generation-study-lars-ferc/832114/>

<https://www.utilitydive.com/news/miso-large-load-generation-study-lars-ferc/832114/>

C40 **CITED** Utility Dive, October 5, 2026: MISO Large Load Addition Resource Study filed at FERC; ruling requested by December 2. <https://www.utilitydive.com/news/miso-large-load-generation-study-lars-ferc/832114/>

<https://www.utilitydive.com/news/miso-large-load-generation-study-lars-ferc/832114/>

C41 **CITED** Robotics and Automation News, October 8, 2026: EU TALOS project, autonomous solar-plant maintenance robots (benefits are targets). <https://roboticsandautomationnews.com/2026/10/08/european-talos-project-demonstrates-autonomous-robots-for-solar-plant-maintenance/105760/>

<https://roboticsandautomationnews.com/2026/10/08/european-talos-project-demonstrates-autonomous-robots-for-solar-plant-maintenance/105760/>

- C42** **CITED** NY PSC Case 26-M-0552, utility AI inventory (September 2026); Daily Energy Insider; mGrid.
<https://dailyenergyinsider.com/hearing-summaries/53816-ny-psc-opens-inquiry-into-utility-use-of-artificial-intelligence-initiates-energy-affordability-index-proceeding/> <https://mgrid.org/2026/09/25/ny-psc-ai-inventory-60-days/>
- C43** **CITED** Utility Dive, October 5, 2026: DOE intends to loan Vistra \$4.2B for nuclear uprates (conditional commitment). Power Engineering. <https://www.utilitydive.com/news/doe-intends-to-loan-vistra-42b-for-nuclear-fleet-improvements/832236/> <https://www.power-eng.com/nuclear/google-constellation-doe-nuclear-plant-investment/>
- C44** **VERIFIED** ServiceNow AI Workflow Factory, October 6, 2026 (newsroom; CIO).
<https://newsroom.servicenow.com/press-releases/details/2026/ServiceNow-launches-AI-Workflow-Factory-to-turn-workflow-improvement-into-a-continuous-agentic-AI-powered-loop/default.aspx>
<https://www.cio.com/article/4231156/servicenow-takes-aim-at-enterprise-ais-workflow-bottleneck-with-ai-workflow-factory.html>
- C45** **VERIFIED** SAP to acquire TechWolf, October 6, 2026 (PR Newswire; Techzine).
<https://www.prnewswire.com/news-releases/sap-to-acquire-techwolf-giving-enterprises-evidence-based-view-of-work-in-the-age-of-ai-302899685.html> <https://www.techzine.eu/news/applications/144796/sap-acquires-ghent-based-ai-company-techwolf-adds-it-to-successfactors/>
- C46** **CITED** Gartner, CHRO priorities for 2027, October 6, 2026: 24% agree on AI-agent accountability (n=297); 63% of executives radically redesigning operating models (n=203).
<https://www.gartner.com/en/newsroom/press-releases/2026-10-6-gartner-says-chros-top-priorities-for-2027-focus-on-meeting-ceos-expectations-while-minimizing-organizational-risk>
- C47** **CITED** Gartner, CPO survey, October 8, 2026: 63% expect significant AI improvement; 25% redesigned roles (n=213). <https://www.gartner.com/en/newsroom/press-releases/2026-10-08-gartner-survey-shows-chief-procurement-officers-stalling-on-needed-ai-role-redesign>
- C48** **VERIFIED** IMF, Georgieva curtain-raiser speech, October 7, 2026. Independent: The Nation Thailand.
<https://www.imf.org/en/news/articles/2026/10/07/sp100726-2026-annual-meetings-curtain-raiser>
<https://www.nationthailand.com/business/economy/40071989> <https://tass.com/economy/2198591>
- C49** **CITED** IMF 2026 Annual Meetings (October 12 to 18) and World Economic Outlook, October 2026 (scheduled October 13). <https://www.imf.org/en/meetings/2026/annual>
<https://www.imf.org/en/publications/weo/issues/2026/10/13/world-economic-outlook-october-2026>
- C50** **CITED** Context: Challenger, Gray & Christmas, September 2026 job-cuts report, October 1, 2026.
<https://www.challengergray.com/blog/job-cuts-fall-in-september-hiring-plans-up-3-over-2025-on-weak-early-seasonal-hiring/>
- C51** **CITED** Stanford HAI, October 7, 2026: framing AI as enrichment raises use.
<https://hai.stanford.edu/news/want-employees-to-embrace-ai-stop-selling-it-as-a-productivity-tool>
- C52** **CITED** Yahoo Finance live coverage, October 6, 2026: record S&P 500 close; 10-year yield near 5.27%.
<https://finance.yahoo.com/markets/live/stock-market-today-tuesday-october-6-dow-sp-500-nasdaq-080526166.html>
- C53** **VERIFIED** Context: California Governor signs SB 947 and SB 951, September 30, 2026 (gov.ca.gov; KQED).
<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://www.kqed.org/news/12102337/newsom-signs-slate-of-ai-workplace-laws-barring-robo-bosses-and-surveillance>

C54 **VERIFIED** Context: Senators Hawley and Murphy announce the AI Agent Accountability Act (about October 1, 2026). Independent: CDO Magazine. <https://www.hawley.senate.gov/senators-hawley-murphy-announce-bipartisan-ai-agent-accountability-act/> <https://www.cdomagazine.tech/aiml/ai-agent-liability-bill-puts-data-access-and-accountability-in-focus>

C55 **CITED** Context: EU Digital Omnibus on AI moves Annex III obligations to December 2, 2027 and Annex I to August 2, 2028 (Sidley; Pinsent Masons). <https://datamatters.sidley.com/2026/06/22/eu-lawmakers-reach-provisional-agreement-to-delay-key-eu-ai-act-obligations/> <https://www.pinsentmasons.com/out-law/news/rules-high-risk-ai-delayed-under-eu-omnibus-deal>

C56 **CITED** McDermott Will & Schulte: Colorado SB 26-189 signed May 14, 2026, effective January 1, 2027. <https://www.mcdermottlaw.com/insights/colorado-ai-law-in-flux-comprehensive-replacement-bill-signed-after-federal-court-blocks-predecessors-enforcement/>

C57 **VERIFIED** Context: Amazon Bedrock Managed Agents (preview), powered by OpenAI, late September 2026 (AWS; The Next Web). <https://aws.amazon.com/about-aws/whats-new/2026/09/bedrock-managed-agents-preview/> <https://thenextweb.com/news/openai-bedrock-managed-agents-aws-devday>

C58 **CITED** Context: Oracle Fusion Claw, September 29, 2026 (HPCwire). <https://www.hpcwire.com/aiwire/2026/09/29/oracle-extends-fusion-agentic-applications-with-introduction-of-fusion-claw/>

C59 **CITED** Context: Salesforce job-ready Agentforce agents, September 11, 2026 (forward-looking dates are plans). <https://www.salesforce.com/news/stories/agentforce-job-ready-ai-agents/>

C61 **CITED** OpenAI, EU text provenance, October 5, 2026. <https://openai.com/index/eu-text-provenance/>

C62 **CITED** OpenAI, disrupting AI-enabled false-front operations, October 8, 2026. <https://openai.com/index/disrupting-ai-enabled-false-front-operations/>

C63 **CITED** Context: Federal Reserve Governor Lisa Cook speech, September 28, 2026. <https://www.federalreserve.gov/newsevents/speech/cook20260928a.htm>

C64 **CITED** Context: Robotics and Automation News, October 2, 2026, on IDC data showing AgiBot led humanoid shipments in H1 2026. <https://roboticsandautomationnews.com/2026/10/02/idc-data-shows-agibot-led-global-humanoid-robot-shipments-in-h1-2026/105463/>

C66 **CITED** TechCrunch, October 7, 2026: Mecka raises \$60M led by Sequoia for robot motion data. <https://techcrunch.com/2026/10/07/robot-data-startup-mecka-ai-nabs-60m-from-sequoia/>

C67 **CITED** The Robot Report, October 5, 2026: Teradyne invests in Bright Machines. Second report: Manufacturing.net. <https://www.therobotreport.com/teradyne-invests-ibright-machines-brings-robotics-ai-infrastructure-manufacturing/> <https://www.manufacturing.net/software/news/22975719/teradyne-partners-with-bright-machines-to-put-robots-test-tech-in-ai-infrastructure-microfactories>

C68 **CITED** MIT News, October 6, 2026: MIT for America initiative. <https://news.mit.edu/2026/mit-america-initiative-strengthens-stem-education-across-country-1006>

C69 **CITED** Gartner IT Symposium/Xpo Orlando, October 19 to 22, 2026 (scheduled). <https://www.gartner.com/en/conferences/na/symposium-us>

Build your agent register with us

A fixed-fee AI Readiness Diagnostic produces a first register, assigns owners and rungs, and tests one halt end to end.

[Book the AI Readiness Diagnostic](#) · [Read about AEGIS](#) · [Open the Reference Architecture Studio](#) · [AI Success Pack](#)

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 or medical advice. AI-assisted research, human-edited editorial standards. © Ariana Digital LLC. All rights reserved. [ariana.digital](#) · [myndQ.ai](#)