ARIANA.DIGITAL

Frontier & Industry Intelligence : Regulated Sectors - FinServices, Healthcare, Energy, Manufacturing

DAILY MARKET PULSE · Wednesday 19 August 2026

Ninety-seven percent of the world's humanoids shipped from China this year. One of the makers listed in Shanghai today.

Unitree opened up about 629% on the Shanghai STAR Market this morning and closed its first session up about 542%, raising roughly CNY 6.1 billion. Four hours away by train, the World Robot Conference opened in Beijing with more than 300 exhibitors. Neither event is the story. The story is that 19,100 humanoids shipped globally in the first half of this year, more than 97% of them from Chinese makers, and industrial and commercial buyers took over 70% of them. If you run a plant, a warehouse or a substation in a regulated sector, that is a supplier concentration question sitting in your capital plan, not a robotics headline. Today: the concentration picture, a nine-move frontier ledger that now includes Cursor launching a GitHub rival three days after SpaceX closed its USD 60 billion acquisition, and four sectors with a win, a constraint and a control apiece.

97%

Chinese share of H1 2026 global humanoid shipments, 19,100 units VERIFIED C03

542%

Unitree first-session close over its CNY 150.80 offer price, 19 Aug VERIFIED C01

9 in 9

Material frontier moves in the nine days to this edition VERIFIED C27

LEAD

The humanoid supply chain got a public price today. Most robot procurement policies still read like forklift procurement policies.

Three facts landed within nine days of each other and they belong together. First, global humanoid shipments reached about **19,100 units in the first half of 2026**, up 272% from roughly 5,100 a year earlier, with Chinese manufacturers accounting for more than 97% of the total VERIFIED C03. Second, the mix shifted: industrial and commercial applications took more than 70% of shipments, against about half a year earlier VERIFIED C03. Third, this morning one of those manufacturers became a listed company, with Unitree opening about 629% above its CNY 150.80 offer and closing its first session up about 542% VERIFIED C01. The offer was priced on 6 August at a valuation near USD 9 billion for about 10% of enlarged capital, making it the first pure-play humanoid maker on a mainland Chinese exchange VERIFIED C02. The World Robot Conference opened the same morning in Beijing E-Town and runs to 23 August, with more than 300 exhibitors, up 36% year on year, and over 150 global product debuts VERIFIED C04.

WHERE THE WORLD'S HUMANOIDS CAME FROM IN THE FIRST HALF OF 2026

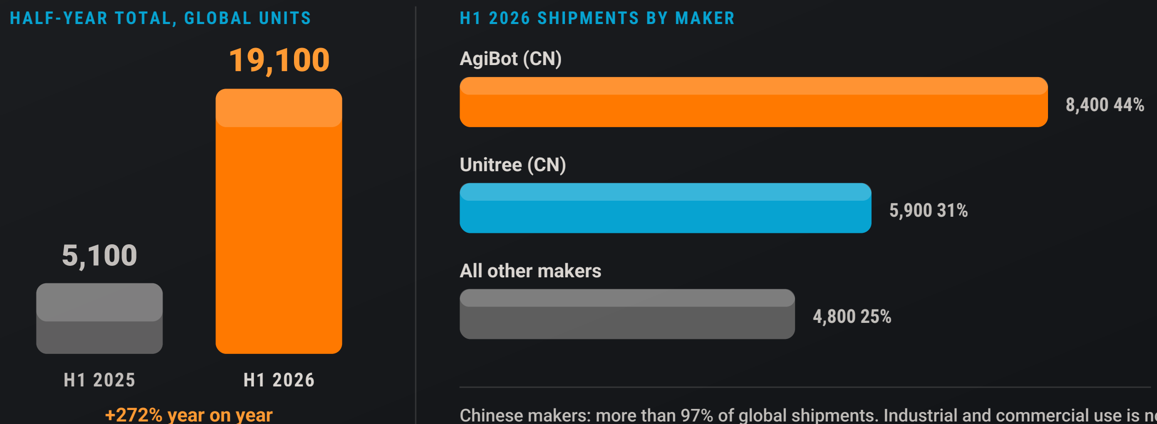

Source C03. Vendor shipment tracking, half-year totals.

Chinese makers: more than 97% of global shipments. Industrial and commercial use is now over 70%.

AgiBot at roughly 8,400 units and Unitree at roughly 5,900 account for about three quarters of first-half global shipments between them VERIFIED C03. Unitree separately reports approximately 18,000 cumulative bipedal units produced by July 2026 and a stated plan to lift its Hangzhou G1 line from about 200 to about 500 units a month by the fourth quarter CITED C06. The second figure is a capacity plan, not delivered output.

WHAT ACTUALLY CHANGED

A listing does two things a private round does not. It publishes a cost of capital, and it publishes a disclosure obligation. From today there is a quoted comparable for humanoid manufacturing capacity, which is what makes the next twelve months of capacity expansion financeable at scale.

WHAT DID NOT CHANGE

Nothing about your ability to answer four questions your auditor will eventually ask. Whose firmware is on the machine. Where does its telemetry terminate. Who can push an update to it. And what happens to the line if that supplier becomes unavailable for ninety days.

The US side of the ledger, stated plainly

Concentration is not the same as capability. Figure AI has forty Figure 03 units placed at BMW's Spartanburg plant following an eleven-month Figure 02 pilot that ran alongside production of more than 30,000 BMW X3 vehicles, and recorded its 1,000th Figure 03 build on 23 July 2026 CITED C05. That is a deeper integration into a regulated automotive production process than unit counts suggest. The honest read is that Chinese makers are winning on volume and price while Western deployments are further into the operational learning curve inside individual plants. Both statements are true at once, and only one of them shows up in a shipment table.

ARIANA.DIGITAL TAKE, IMPLEMENTATION ARCHITECT VIEW

The mistake we see is treating robot fleets as capital equipment and agent fleets as software, with two separate governance regimes. A humanoid on a production line is an agent with a body. It perceives, plans, acts on the physical world and updates itself over a network. If your agent inventory does not have rows for the machines on the floor, it is not an agent inventory. It is an application register.

Method and correction policy

Every edition is researched fresh against sources published within the preceding seven days where the item is time-sensitive. Figures carry a chip: VERIFIED means named, dated and publicly checkable; CITED means named source, not independently re-verified; FLAG means contested and pending re-verification. Where market commentary conflicted with primary legal sources this week, notably on EU high-risk applicability, we followed the primary legal sources and said so.

FRONTIER LEDGER

Nine material moves in nine days, across four countries and two capital markets

We hold this section to one rule: equal factual treatment regardless of who we partner with. Below is what actually shipped or moved between 11 and 19 August 2026, with the implementation consequence stated next to it rather than in a separate opinion column.

FRONTIER CADENCE, 11 TO 19 AUGUST 2026

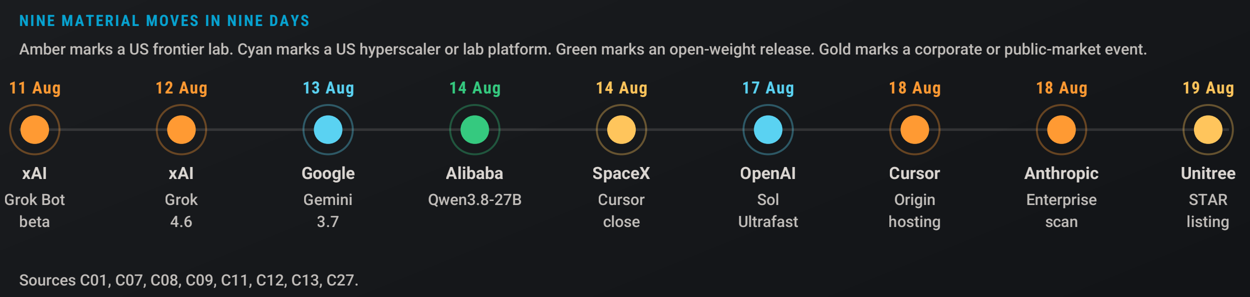

Three are capability releases, two are enterprise-control or platform releases, one is an open-weight release, one is a throughput preview, and two are corporate events: the SpaceX close of the Cursor acquisition and the Unitree listing. Capability is no longer the scarce input in this window.

LAB OR VENDOR	WHAT SHIPPED OR MOVED	IMPLEMENTATION READ
xAI	Grok 4.6 on 12 August, thirty-five days after 4.5. 500K context, USD 2 and USD 6 per million tokens held flat, new xhigh reasoning level, reported 61 on the Artificial Analysis Intelligence Index, agent tasks reported at roughly half the prior turn count. Grok Bot entered beta on 11 August. Added to GitHub Copilot's model picker. CITED C07	Turn efficiency is the buried number. Halving turns on the same task halves orchestration surface, retry logic and audit volume. Ask your vendor for turns per completed task, not tokens.
Google	Gemini 3.7 Flash on 13 August, three weeks after 3.6 Flash. Vendor-reported FrontierCode 1.1 moving from 34.4% to 43.6% and	A three-week minor-version cadence on a workhorse tier is a change-control problem before it is a capability gain. Pin to a

LAB OR VENDOR	WHAT SHIPPED OR MOVED	IMPLEMENTATION READ
	DeepSWE v1.1 from 49% to 65.3%. CITED C08	capability tier in contract language, not a version string.
OpenAI	Limited API preview of Ultrafast mode for GPT-5.6 Sol on Cerebras hardware, reported at roughly 750 output tokens per second and up to fourteen times standard throughput. Daybreak cybersecurity programme expanded with GPT-5.6-Cyber, Daybreak Blue and Daybreak Red for approved defenders. Teen-tailored ChatGPT mode shipped with age prediction. CITED C09 Separately, a confidential S-1 was reported filed with the SEC in June, with a listing window reported from September into the fourth quarter and press-reported revenue near USD 2 billion a month against a projected 2026 loss near USD 14 billion. OpenAI has not stated a target valuation, so the widely circulated USD 1 trillion figure remains speculation. CITED C10	Throughput at that level changes which workloads can sit in a synchronous user path. It also changes how fast an unbounded agent loop can burn budget. Rate-limit at the agent, not the account. A listed provider is a different vendor-risk profile than a private one, so start that assessment before it happens.
Anthropic	Three enterprise controls landed in August, and they map to three separate risks. Beta skill and plugin security scanning checks third-party uploads for malicious content, which addresses the extension supply chain. An Admin API in beta covers listing members, changing roles, removing members and managing invites and groups, which addresses joiner-mover-leaver drift on agent access. Claude in Chrome now defaults to off on Enterprise plans with admin-controlled domain allowlisting, which addresses browser agents reaching systems nobody scoped. VERIFIED C11	This is the first widely available answer to the supply-chain question for agent extensions, and it is worth more than a capability release to anyone in a regulated estate. If you allow staff-authored skills, scanning plus an allowlist is the minimum viable control pair. The Admin API is the piece most teams will underrate: without programmatic role and group management, agent access reviews stay manual, and manual access reviews are the finding that shows up in every audit. Ask whether your identity team can revoke an agent as fast as it can revoke a person.

LAB OR VENDOR	WHAT SHIPPED OR MOVED	IMPLEMENTATION READ
Alibaba, DeepSeek, Moonshot	Qwen3.8-27B released 14 August. DeepSeek-V4-Flash on 31 July. Kimi K3 open weights on 27 July. Open-weight cadence is now roughly weekly across the three. CITED C12	Open weights are the only honest path to an air-gapped agent in a sanctioned or classified environment. They are also the fastest way to acquire a provenance question you cannot answer. Decide deliberately.
Cursor, now SpaceX	SpaceX closed its USD 60 billion all-stock acquisition of Anysphere on 14 August, issuing about 391 million Class A shares. Cursor is now a wholly owned subsidiary inside the SpaceXAI team with access to the Colossus cluster in Memphis. Widely reported as the largest startup acquisition on record. VERIFIED C13	The change of control is done, so the vendor-risk window has already closed. If Cursor is in your SDLC, the DPA and source-code-egress assessment need refreshing now, against a new parent, not at renewal.
Cursor Origin	Three days after the close, Cursor launched Origin in beta for paid users: repositories, browsing, editing and pull requests hosted inside the editor, with GitHub sync. Cursor positions it as code hosting built for agent scale. It shipped during a GitHub outage where error rates hit about 20% on web and API traffic and up to 50% on repository downloads. Independent tracking counted 257 GitHub incidents between May 2025 and April 2026, 48 of them major. VERIFIED C27	This is the first credible challenge to GitHub as the default system of record for source code, and it arrives owned by a launch company. For a regulated buyer the question is not editor preference. It is whether your source of truth, your audit trail and your access reviews are about to sit with a vendor your third-party risk process has never assessed. Decide before an engineer does it for you.
NVIDIA	Company-reported first-quarter fiscal 2027 revenue of USD 81.6 billion, up 85% year on year, data centre up 92%. Vera Rubin in full-scale production, with Foxconn indicating rack shipments from the fourth quarter of calendar 2026. Second-quarter results scheduled for 26 August 2026. CITED C14	Rack-scale generational change lands in the same window as your 2027 capacity plan. If you are signing multi-year colocation now, price the power envelope for the next generation, not this one.

THE PATTERN WORTH NAMING

In the first half of 2025 the frontier news flow was almost entirely capability. In this nine-day window, four of the nine moves are not model releases at all. They are governance features, a code-hosting platform, an acquisition close and a listing. Capability releases are arriving on a three to five week cadence, which means the binding constraint for a regulated buyer has shifted from what the model can do to how fast the organisation can re-approve what changed, and to who owns the infrastructure underneath. Both are process problems, and process problems are solvable without waiting for the next release.

Method and correction policy

Every edition is researched fresh against sources published within the preceding seven days where the item is time-sensitive. Figures carry a chip: VERIFIED means named, dated and publicly checkable; CITED means named source, not independently re-verified; FLAG means contested and pending re-verification. Where market commentary conflicted with primary legal sources this week, notably on EU high-risk applicability, we followed the primary legal sources and said so.

SECTOR, 1 OF 2

Financial services and healthcare: where the agent is allowed to finish the sentence

FINANCIAL SERVICES · THE WIN

JPMorgan Chase reported in June that it runs more than 400 production AI use cases on an enterprise machine-learning platform, and attributed a 20% increase in private banking gross sales in part to AI, while stating an intent to deploy agents that work autonomously over longer horizons CITED C16. Treat the 20% as company-reported and unaudited. Treat the 400 as the more useful number, because it tells you the organisation solved the deployment pipeline problem, which is the part most institutions are stuck on.

FINANCIAL SERVICES · THE CONSTRAINT

Roughly 23% of organisations are scaling an agentic system and about 62% are at least experimenting, but fewer than 10% have scaled agents to tangible value, and 12% of chief executives report both revenue gain and cost reduction CITED C25. In regulated banking the gap between experimenting and scaling is almost never model quality. It is that no one has written down which decisions the agent may complete without a named human, and which it may only prepare.

FINANCIAL SERVICES · THE CONTROL WE WOULD PUT IN FIRST

A two-column decision-authority schedule per agent: **prepare** versus **commit**. Prepare means the agent may assemble, draft, screen, reconcile and recommend. Commit means it may file, send, book, approve or release funds. Every agent in your inventory gets one row and one named owner for the commit column. Most institutions we work with find that no agent currently needs commit authority in year one, which makes the risk conversation with the second line considerably shorter.

HEALTHCARE · THE WIN

The regulatory path for agentic clinical AI is being built rather than argued about. ARPA-H is soliciting two agentic clinical assistants, including ADVOCATE for cardiovascular care, with a stated aim of producing the first FDA-authorised agentic AI able to provide round-the-clock specialty care on an approximately three-year timeline CITED C17. Separately the FDA has an established precedent for a protocolised clinical workflow using conversational AI around a deterministic core, cleared as 510(k) K253281 on 23 December 2025 CITED C18.

HEALTHCARE · THE CONSTRAINT

Both of those are narrow by design. The ARPA-H item is a funding solicitation and a stated goal, not an authorisation, and it should not be cited as one. The cleared device is not an autonomous language model making clinical decisions; it is conversation wrapped around deterministic logic. The FDA has separately asked for public feedback on how to regulate generative AI-enabled devices, including foundation models and agentic systems, which tells you the framework is genuinely unsettled CITED C18.

HEALTHCARE · THE CONTROL WE WOULD PUT IN FIRST

Separate the deterministic core from the conversational surface in your architecture diagram, on purpose, before your quality team asks you to. The cleared precedent works because the part that decides is testable and the part that talks is bounded. If your design has the model both deciding and speaking, you have a harder regulatory conversation and a harder incident review.

Method and correction policy

Every edition is researched fresh against sources published within the preceding seven days where the item is time-sensitive. Figures carry a chip: VERIFIED means named, dated and publicly checkable; CITED means named source, not independently re-verified; FLAG means contested and pending re-verification. Where market commentary conflicted with primary legal sources this week, notably on EU high-risk applicability, we followed the primary legal sources and said so.

SECTOR, 2 OF 2

Manufacturing and energy: the two sectors where an agent's mistake has mass

MANUFACTURING · THE WIN

Predictive maintenance is the one applied-AI category in industry with a durable, repeatedly reported outcome range: unplanned downtime reduced 20% to 40% and maintenance costs reduced 25% to 40% across documented production deployments CITED C20. Siemens continues to build this into the Industrial Edge and Industrial AI Suite stack, and published on 13 August on applying it to semiconductor fab uptime CITED C19. GE Vernova extended the pattern in August 2025 by feeding autonomous robotic inspection data from ANYbotics units into its Asset Performance Management platform CITED C20.

MANUFACTURING · THE CONSTRAINT

Those ranges are vendor-reported across mixed deployments. They are a reason to run the business case, not the business case itself. The sharper constraint today is the one from page two: if more than 97% of humanoid units ship from a single jurisdiction VERIFIED C03, then fleet firmware, telemetry endpoints and update authority concentrate too. A dual-source policy for robots is not a geopolitical statement. It is the same continuity logic you already apply to a single-source bearing.

MANUFACTURING · THE CONTROL WE WOULD PUT IN FIRST

A robot and edge-device provenance register with five columns: make and model, firmware version and signing authority, network segment, telemetry destination, and update approval owner. Populate it for every autonomous or semi-autonomous machine on the floor, including AMRs and inspection units. Two weeks of work in most plants. It is also the artefact that makes an OT security review possible at all.

ENERGY: WHAT THE COMPUTE BUILD IS ACTUALLY ASKING THE GRID FOR

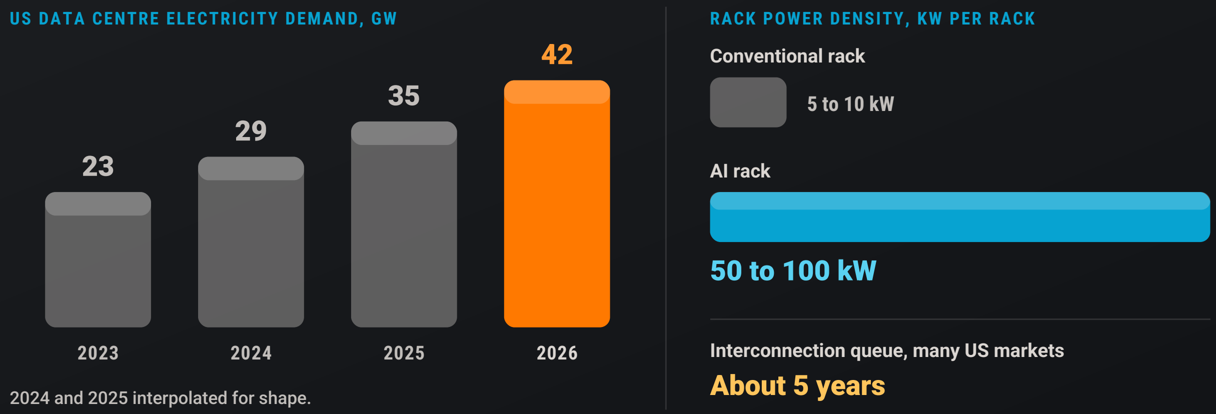

2024 and 2025 interpolated for shape.

Sources C21, C22. Power availability, not land or permitting, is now the leading cause of construction delay.

The rack-density line is the one that breaks existing facility assumptions. A hall designed for 5 to 10 kW racks cannot be retrofitted to 50 to 100 kW by adding capacity; it is a different building CITED C21.

ENERGY · THE WIN

Procurement got creative and it worked. Hyperscalers have collectively committed to close to 10 GW of nuclear and small modular reactor capacity since 2023, and the first direct power purchase agreements between large technology buyers and reactor operators are closing this year CITED C22. That is a genuine structural change in how baseload gets financed.

ENERGY · THE CONSTRAINT

Committed is not energised. Interconnection queues average about five years in many US markets, and power availability rather than land or permitting is now the leading cause of construction delay CITED C21. A signed PPA does not move a queue position.

ENERGY · THE CONTROL WE WOULD PUT IN FIRST

Put an energy line in the agent business case. Not a sustainability line, a unit-economics line: watts per completed task, and the marginal cost of one additional agent run at your contracted rate. Teams that do this stop deploying agents on a per-seat licence assumption and start designing for turn efficiency, which is exactly the number the Grok 4.6 release put on the table CITED C07. Cost discipline and carbon discipline turn out to be the same engineering decision.

Method and correction policy

Every edition is researched fresh against sources published within the preceding seven days where the item is time-sensitive. Figures carry a chip: VERIFIED means named, dated and publicly checkable; CITED means named source, not independently re-verified; FLAG means contested and pending re-verification. Where market commentary conflicted with primary legal sources this week, notably on EU high-risk applicability, we followed the primary legal sources and said so.

WORKFORCE

Employers are not cutting to pay for AI. They are re-specifying who they hire and retraining who they have.

The Federal Reserve Bank of New York published in August that most firms expecting to use AI in the next six months plan to retrain their workforces, with far fewer reporting adjustments to planned headcount, and concluded that the measured labour-market effect runs through changing skill requirements rather than job elimination VERIFIED C23. That is a central-bank survey finding, published this month, and it is the most reliable current read available.

DEMAND SIGNAL VERSUS DELIVERY CAPABILITY

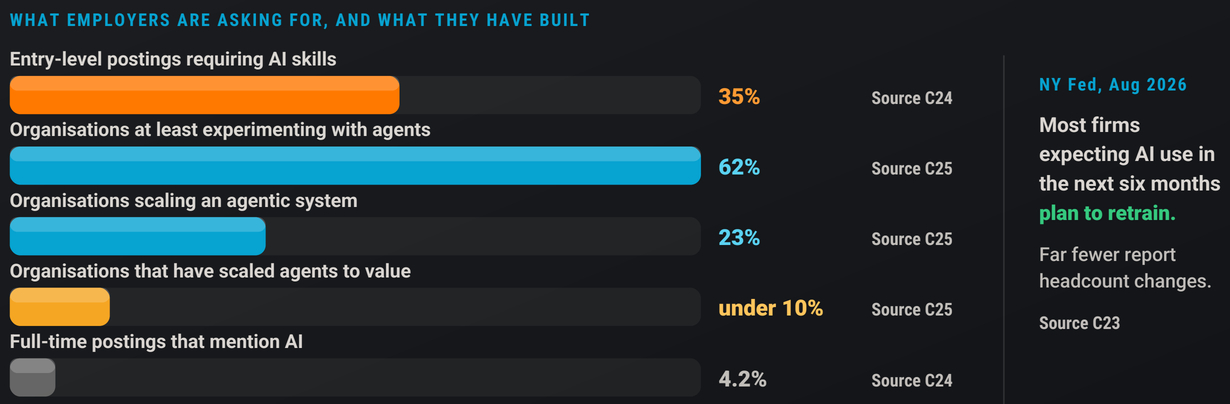

Bars share one scale to 62%. Adoption survey bases differ, so read the gap between rows two, three and four rather than the absolute levels.

About 35% of entry-level postings now require AI skills and the share of full-time postings mentioning AI has roughly doubled year on year to about 4.2% CITED C24. Set that against the delivery picture: 62% experimenting, 23% scaling, under 10% at tangible value CITED C25.

CAUSE AND EFFECT

Postings are a demand signal that leads capability by twelve to eighteen months. The gap between the 35% asking for AI skills at entry level and the under 10% who have scaled agents to value is not hypocrisy. It is the ordinary lag between deciding to build a capability and having one.

RISK AND REWARD

The risk in that lag is hiring for tool fluency and getting no change in outcomes, because the constraint was never tool access. The reward, for the organisations that get it right, is that they are buying judgment about where an agent is allowed to act, which is scarce and does not arrive with a certification.

A NOTE ON THE SURVEY EVIDENCE, BECAUSE IT MATTERS HERE

The most cited enterprise AI survey in circulation, Deloitte's State of AI in the Enterprise 2026, was fielded across 3,235 leaders between August and September 2025 CITED C26. That is roughly a year before this edition, and it predates the agentic release cadence documented on page three. It remains useful for direction, particularly its finding that agentic impact is expected highest in customer support with supply chain, R&D, knowledge management and cybersecurity close behind. It is not a current measurement, and we would not use it to size a 2027 budget. Where we quote it, we date it.

What we would tell a chief people officer this week

Stop screening for AI tool names. Screen for three things instead. One, can this person describe a workflow precisely enough that it could be handed to an agent, including its failure modes. Two, can they say what should never be automated in their function and defend it. Three, have they ever caught a confident wrong answer from a system and been able to explain how. Those three are teachable, testable and they transfer across every tool release in the nine-day window on page three.

ARIANA.DIGITAL TAKE

Retraining budgets are being approved right now, in the same planning cycle as agent pilots, by different people, against different success measures. The single highest-return intervention we make with clients is putting those two plans on one page. The agent inventory tells you exactly which roles change and in what order. Without it, retraining is a generic literacy programme and the pilot has no receiving organisation.

Method and correction policy

Every edition is researched fresh against sources published within the preceding seven days where the item is time-sensitive. Figures carry a chip: VERIFIED means named, dated and publicly checkable; CITED means named source, not independently re-verified; FLAG means contested and pending re-verification. Where market commentary conflicted with primary legal sources this week, notably on EU high-risk applicability, we followed the primary legal sources and said so.

PRACTITIONER DESK

Questions we were asked this week, answered the way we would answer them on a call

We deploy AMRs and are evaluating humanoids. Does today's listing change our decision?

Not the decision. It changes the diligence. A listed supplier publishes financials and carries disclosure obligations, which is genuinely useful for a continuity assessment. What it does not do is answer where firmware is signed, where telemetry terminates, and who can push an update. Ask those three in the RFP regardless of who wins, and ask them of the incumbent too.

We shipped an AI feature before 2 August. Are we already non-compliant in the EU?

Article 50 transparency has applied since 2 August 2026 and was not deferred by the Digital Omnibus VERIFIED C15. There is one narrow grace period: for systems placed on the market before 2 August 2026, the Article 50(2) marking and machine-readable detection duty applies from **2 December 2026**. That is about fifteen weeks from today. The rest of Article 50, including telling people they are interacting with an AI system, is live now.

Our agents run on three different model providers. Is that a governance problem or a hedge?

Both, and which one depends on a single artefact. If your agent definitions name model versions, three providers is three change-control streams on independent release cadences, and this month alone that would have meant re-approving against Grok 4.6, Gemini 3.7 Flash and Qwen3.8 CITED C07 CITED C08 CITED C12, while Anthropic changed enterprise defaults underneath a fourth VERIFIED C11. If your definitions name capability tiers with an evaluation gate between the tier and the model, three providers is a real hedge. The abstraction layer is the whole difference.

Our staff are writing their own agent skills. Is that a problem?

It is the extension supply chain, and it is the least governed part of most agent estates. A staff-authored skill or plugin is executable code that runs with the agent's permissions, and until recently the only control most platforms offered was a policy telling people not to. Anthropic's August release notes put a beta scanner in front of third-party uploads on Claude Enterprise plans, and defaulted Claude in Chrome to off with admin-controlled domain allowlisting VERIFIED C11. Treat that as the shape of the control rather than a product recommendation: a scanner at upload, an allowlist for what the agent may reach, and a named owner for the exception list. If your platform does not offer those three, the compensating control is a review gate before a skill is shared beyond its author.

Did you know

Turn count is becoming a governance metric, not just a cost metric. xAI reported Grok 4.6 completing agent tasks in roughly half the turns of the previous release CITED C07. Every turn an agent takes is a place it can call a tool, touch a record, or be prompt-injected. Halving turns halves the audit surface for the same completed work. Ask vendors for turns per completed task alongside tokens.

What is the smallest useful first step, this week?

One page, four columns, every agent and every autonomous machine: what it touches, who owns it, what it is allowed to decide alone, and who can stop it. Include the robots. Most organisations can complete it in a fortnight and are surprised by the row count. Every framework we work with, AEGIS included, needs this artefact before anything else in it becomes usable.

Scenario, if you want one to plan against

Assume shipment concentration holds through 2027 and one major supplier becomes commercially or politically unavailable for a quarter. The question is not whether you can buy a replacement machine. It is whether your line control software, safety interlocks and inspection models are portable to a different chassis, or whether you have written twelve months of integration work into a single vendor's API. That is answerable today, on paper, for the cost of an afternoon.

Method and correction policy

Every edition is researched fresh against sources published within the preceding seven days where the item is time-sensitive. Figures carry a chip: VERIFIED means named, dated and publicly checkable; CITED means named source, not independently re-verified; FLAG means contested and pending re-verification. Where market commentary conflicted with primary legal sources this week, notably on EU high-risk applicability, we followed the primary legal sources and said so.

RESEARCH BASE

Sources

Every figure in this edition maps to an entry below. Chips mark verification tier. Company-reported results are labelled as such in the body. Solicitations, announced transactions, capacity plans and scheduled events are not treated as completed facts.

- **C01** VERIFIED C01 Unitree Robotics STAR Market debut, Wednesday 19 August 2026, ticker 688836. Shares opened up about 629% against the CNY 150.80 offer price and closed the first session up about 542%. The offering raised roughly CNY 6.1 billion, about USD 905 million. Intraday market value moved from roughly USD 9 billion at the offer to a reported peak near USD 66 billion before paring. https://www.cnbc.com/2026/08/19/china-backflipping-robot-maker-unitree-jumps-shanghai-ipo.html https://finance.yahoo.com/markets/stocks/articles/chinese-robot-maker-unitree-jumps-071833994.html
- **C02** VERIFIED C02 Unitree IPO terms. Priced at CNY 150.80 per share on 6 August 2026, valuing the company near USD 9 billion. 40.45 million new shares, about 10% of enlarged capital. First pure-play humanoid robot maker to list on a mainland Chinese exchange. Retail subscription was oversubscribed several thousand times. https://kraneshares.com/a-complete-guide-to-unitree-robotics-2026-ipo-why-it-matters-for-star-market-etf-kstr-humanoid-robotics-etf-koid/ https://qz.com/unitree-robotics-ipo-oversubscribed-shanghai-081126
- **C03** VERIFIED C03 Global humanoid robot shipments of about 19,100 units in the first half of 2026, up 272% from roughly 5,100 units a year earlier. Chinese manufacturers accounted for more than 97% of the total. AgiBot shipped about 8,400 units, roughly 44% of the global total; Unitree about 5,900. Industrial and commercial applications were more than 70% of shipments, against about 50% a year earlier. https://www.bloomberg.com/news/articles/2026-08-10/china-humanoid-makers-hold-97-of-global-shipments-report-says https://technode.com/2026/08/11/chinese-makers-accounted-for-more-than-97-of-global-humanoid-robot-shipments-in-h1-2026/ https://www.techrepublic.com/article/news-humanoid-robot-shipments-chinese-vendors/
- **C04** VERIFIED C04 World Robot Conference 2026, Beijing E-Town, 19 to 23 August 2026. More than 300 exhibitors, up 36% year on year, over 2,000 exhibits and more than 150 global product debuts. 19 August is Release Day; 20 August is Procurement Day. https://english.beijing.gov.cn/latest/news/202608/t20260819_4828602.html https://www.globaltimes.cn/page/202608/1368527.shtml
- **C05** CITED C05 Figure AI deployment reporting, July and August 2026. Forty Figure 03 units placed at the BMW Spartanburg plant following an eleven-month Figure 02 pilot that ran alongside production of more than 30,000 BMW X3 vehicles. BotQ recorded its 1,000th Figure 03 build on 23 July 2026. Company-reported through deployment trackers, not independently audited. https://www.technology.org/2026/07/18/humanoid-robots-in-2026-what-is-actually-deployed/ https://www.solidmarketresearch.com/post/humanoid-robots-cross-the-pilot-threshold-where-factory-deployment-actually-stands-in-2026
- **C06** CITED C06 Unitree production capacity reporting. Approximately 18,000 cumulative bipedal humanoids produced by July 2026. Hangzhou G1 line running at about 200 units per month with a stated plan to reach 500 per month by the fourth quarter of 2026. Company-stated capacity plan, not delivered output. https://www.humanoidsdaily.com/news/unitree-reclaims-the-spotlight-with-18-000-cumulative-bipedal-humanoids-produced
- **C07** CITED C07 xAI Grok 4.6, released 12 August 2026, thirty-five days after Grok 4.5. Retains a 500K context window and USD 2 per million input and USD 6 per million output token pricing, adds an xhigh reasoning level,

and is reported at 61 on the Artificial Analysis Intelligence Index with agent tasks completed in roughly half the turns of the prior release. Also added to GitHub Copilot's model picker. Composite benchmark, not an audited evaluation.
https://x.ai/news https://codersera.com/blog/grok-4-6-launch-guide-2026/

- **C08** CITED C08 Google Gemini 3.7 Flash, released 13 August 2026, three weeks after Gemini 3.6 Flash. Vendor-reported movement on FrontierCode 1.1 from 34.4% to 43.6% and DeepSWE v1.1 from 49% to 65.3%. Vendor-reported benchmark figures.
https://releasebot.io/updates/google/gemini https://www.digitalapplied.com/blog/ai-model-releases-august-2026-tracker
- **C09** CITED C09 OpenAI product activity in the week to 19 August 2026: limited API preview of Ultrafast mode for GPT-5.6 Sol running on Cerebras hardware at roughly 750 output tokens per second and up to 14 times standard throughput; expansion of the Daybreak cybersecurity programme with GPT-5.6-Cyber, Daybreak Blue and Daybreak Red for approved defenders; and a teen-tailored ChatGPT mode for users aged 13 to 17 using age prediction to route minors automatically.
https://releasebot.io/updates/openai https://techstartups.com/2026/08/17/top-tech-news-today-august-17-2026-ge-microsoft-nvidia-open-stripe-unitree-more/
- **C10** CITED C10 OpenAI listing preparation. A confidential S-1 registration statement was reported filed with the SEC in June 2026, with a listing window reported from September 2026 into the fourth quarter and Goldman Sachs and Morgan Stanley engaged. Reported revenue of roughly USD 2 billion a month, about USD 25 billion annualised, against a projected 2026 loss near USD 14 billion. These are press-reported figures. OpenAI has not stated a target valuation, and the widely circulated USD 1 trillion figure should be treated as speculation until it appears in an offering document.
https://www.techi.com/openai-ipo/ https://www.buildfastwithai.com/blogs/ai-news-today-august-17-2026
- **C11** VERIFIED C11 Anthropic Claude enterprise controls, August 2026 release notes: beta skill and plugin security scanning on Enterprise plans to check third-party uploads for malicious content; Admin API in beta for listing members, changing roles, removing members, managing invites and groups; Claude in Chrome off by default on Enterprise with admin-controlled domain allowlisting.
https://platform.claude.com/docs/en/release-notes/overview https://www.anthropic.com/news
- **C12** CITED C12 Open-weight release cadence around this edition: Alibaba Qwen3.8-27B on 14 August 2026, DeepSeek-V4-Flash on 31 July 2026, and Moonshot AI Kimi K3 open weights on 27 July 2026. Release-tracker aggregation.
https://aireleasetracker.com/latest https://llmgateway.io/timeline
- **C13** VERIFIED C13 SpaceX closed its USD 60 billion all-stock acquisition of Anysphere, maker of Cursor, on 14 August 2026. Approximately 391 million SpaceX Class A shares were issued. Cursor becomes a wholly owned subsidiary folded into the SpaceXAI team, with access to the Colossus cluster in Memphis, reported at roughly 200,000 Nvidia GPUs. Widely reported as the largest startup acquisition on record. Anysphere had reached roughly USD 4 billion annualised revenue by June 2026, about USD 2.6 billion of it enterprise. Completed transaction, corrected from the pre-close position reported earlier.
https://finance.yahoo.com/technology/ai/articles/spacex-completes-record-60-billion-131311785.html https://satnews.com/2026/08/13/spacex-finalizes-regulatory-procedures-to-close-60-billion-acquisition-of-ai-platform-cursor/ https://thenextweb.com/news/spacex-cursor-acquisition-completed-gpu-fleet
- **C27** VERIFIED C27 Cursor Origin, a code hosting platform built directly into the Cursor editor, launched in beta for paid users on 17 and 18 August 2026, three days after the SpaceX close. It covers repositories, browsing, editing and pull requests, syncs with GitHub, and is positioned by Cursor as code hosting built for agent scale rather than as a GitHub clone. It shipped during a multi-hour GitHub outage in which error rates reached about 20% on web and API traffic and up to 50% on repository downloads and raw content. Independent outage tracking recorded 257 GitHub incidents between May 2025 and April 2026, 48 of them classified as major.
https://techcrunch.com/2026/08/18/cursor-capitalizes-on-github-frustration-launches-rival-hosting-platform/ https://venturebeat.com/infrastructure/cursor-launches-origin-code-hosting-platform-as-github-outage-exposes-opening-in-ai-coding-race https://techstartups.com/2026/08/17/cursor-launches-origin-a-github-rival-built-for-ai-coding-agents/

- **C14** CITED C14 NVIDIA position entering its next print. Company-reported first-quarter fiscal 2027 revenue of USD 81.6 billion, up 85% year on year, with data centre revenue up 92%. Vera Rubin described as in full-scale production, with Foxconn indicating rack shipments beginning in the fourth quarter of calendar 2026. Second-quarter fiscal 2027 results are scheduled for 26 August 2026, one week after this edition. https://www.tradingkey.com/analysis/stocks/us-stocks/262116112-nvidia-nvda-q1-81-6b-vera-rubin-production-ai-financing-openai-105b-guarantee-august-26-earnings https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-fourth-quarter-and-fiscal-2026
- **C15** VERIFIED C15 EU AI Act Article 50 transparency obligations have applied since 2 August 2026 and were not amended by the Digital Omnibus, which deferred Annex III high-risk obligations to 2 December 2027. A limited grace period applies only to the Article 50(2) marking and machine-readable detection duty, and only for systems placed on the market before 2 August 2026; those providers must comply from 2 December 2026. https://www.morganlewis.com/blogs/sourcingatmorganlewis/2026/08/eu-ai-acts-transparency-rules-what-went-into-effect-on-2-august https://www.goodwinlaw.com/en/insights/publications/2026/08/alerts-technology-dpc-eu-ai-act-transparency-obligations-now-in-force https://digital-strategy.ec.europa.eu/en/faqs/transparency-obligations-under-article-50-ai-act
- **C16** CITED C16 JPMorgan Chase agentic AI position as reported 9 June 2026: plans to deploy agents able to work autonomously over longer horizons; more than 400 production AI use cases on an enterprise machine-learning platform; company-reported 20% increase in private banking gross sales attributed in part to AI. Company-reported figures. https://www.cnbc.com/2026/06/09/jpmorgan-chase-ai-agents.html
- **C17** CITED C17 ARPA-H solicitation for two agentic AI clinical assistants, including ADVOCATE, the Agentic AI-Enabled Cardiovascular Care Transformation programme, with a stated aim of producing the first FDA-authorised agentic AI able to provide round-the-clock specialty care on an approximately three-year timeline. A funding solicitation and a stated goal, not an authorisation. https://www.fiercehealthcare.com/ai-and-machine-learning/trump-administration-creating-clinical-ai-agents-3-year-fda-approval
- **C18** CITED C18 FDA position on generative and agentic AI in devices. FDA 510(k) K253281 cleared 23 December 2025 for a protocolised clinical workflow using conversational AI around a deterministic core. FDA has separately sought public feedback on a regulatory approach for generative AI-enabled medical devices, including risk-proportionate postmarket monitoring and considerations for foundation models and agentic systems. https://www.fda.gov/news-events/press-announcements/fda-seeks-public-feedback-inform-regulatory-approach-generative-ai-enabled-medical-devices https://innolitics.com/articles/updoc-fda-cleared-ai-agent/
- **C19** CITED C19 Siemens industrial AI position, 2026. Industrial Edge ecosystem expansion announced at Hannover Messe 2026 with the Industrial AI Suite covering predictive maintenance and visual inspection. Siemens EDA published on 13 August 2026 on AI-driven predictive maintenance for semiconductor fab uptime. Vendor material. https://press.siemens.com/global/en/pressrelease/siemens-industrial-edge-ecosystem-strengthens-data-and-ai-integration https://blogs.sw.siemens.com/electronics-semiconductors/2026/08/13/ai-driven-predictive-maintenance-solving-semiconductor-fab-uptime-challenges/
- **C20** CITED C20 Predictive maintenance outcome ranges reported across documented production deployments: unplanned downtime reduced 20% to 40% and maintenance costs reduced 25% to 40%. GE Vernova extended the pattern with ANYbotics and AWS in August 2025 by feeding autonomous robotic inspection data into its Asset Performance Management platform. These are vendor-reported ranges across mixed deployments, not a single audited result. https://www.marketsandmarkets.com/ResearchInsight/ai-driven-predictive-maintenance-companies.asp https://ifactoryapp.com/predictive-maintenance/ai-predictive-maintenance-manufacturing-quality-control
- **C21** CITED C21 US data centre electricity demand rising from about 23 GW in 2023 to about 42 GW in 2026. AI racks drawing 50 to 100 kW against 5 to 10 kW for conventional racks. Grid interconnection queues averaging around five years in many US markets, with power availability rather than land or permitting now the leading cause of construction delay.

https://enkiai.com/ai-market-intelligence/ai-data-center-power-grid-limits-reshape-energy-in-2026/ https://www.morganstanley.com/insights/articles/powering-ai-energy-market-outlook-2026

- C22 CITED C22 Hyperscaler nuclear procurement. Collective commitments approaching 10 GW of nuclear and small modular reactor capacity since 2023, with the first direct power purchase agreements between large technology buyers and reactor operators closing in 2026. Contracted and announced capacity, not energised capacity. https://informedclearly.com/en/ai/53909/ai-data-centers-nuclear-power-2026 https://www.iaea.org/bulletin/data-centres-artificial-intelligence-and-cryptocurrencies-eye-advanced-nuclear-to-meet-growing-power-needs
- C23 VERIFIED C23 Federal Reserve Bank of New York, Liberty Street Economics, August 2026, AI's Impact on Labor and Hiring. Most firms expecting to use AI in the next six months plan to retrain their workforces, with far fewer reporting adjustments to planned headcount. The measured effect runs through changing skill requirements rather than job elimination, at least so far. https://libertystreeteconomics.newyorkfed.org/2026/08/ais-impact-on-labor-and-hiring/
- C24 CITED C24 US hiring signal data, 2026. The share of full-time job postings mentioning AI has roughly doubled year on year to about 4.2%, about 35% of entry-level postings require AI skills, and employers project class-of-2026 hiring up about 5.6% on last year. Job-board and survey data, not official statistics. https://www.ziprecruiter-research.org/economic-insights-research/ai-employer-report-2026 https://www.cnbc.com/select/class-of-2026-hiring-stats-and-ai-trends/
- C25 CITED C25 Agentic adoption and value capture. Roughly 23% of organisations are scaling an agentic AI system and about 62% are at least experimenting, while fewer than 10% have scaled agents to tangible value. Separately, 12% of chief executives report achieving both revenue gain and cost reduction from AI. Survey aggregation across McKinsey, Gartner and PwC sources. https://prefactor.tech/learn/ai-agent-adoption-statistics https://www.pwc.com/gx/en/1/services/ai/ai-jobs-barometer.html
- C26 CITED C26 Deloitte State of AI in the Enterprise 2026, based on 3,235 leaders surveyed between August and September 2025. Agentic AI is expected to have highest impact in customer support, with supply chain, R&D, knowledge management and cybersecurity also rated high potential. Note the survey field period predates this edition by roughly one year and is treated as directional. https://www.deloitte.com/us/en/what-we-do/capabilities/applied-artificial-intelligence/content/state-of-ai-in-the-enterprise.html

---

Method and correction policy

Every edition is researched fresh against sources published within the preceding seven days where the item is time-sensitive. Figures carry a chip: VERIFIED means named, dated and publicly checkable; CITED means named source, not independently re-verified; FLAG means contested and pending re-verification. Where market commentary conflicted with primary legal sources this week, notably on EU high-risk applicability, we followed the primary legal sources and said so.