DAILY MARKET PULSE · WEDNESDAY FRONTIER INFOGRAPHIC EDITION

Wednesday, September 23, 2026 · America/New_York

Capability got cheaper this week. Accountability did not.

In the forty-eight hours before this edition, three frontier labs repriced their flagship tiers downward or held price while raising capability, and Alphabet's robotics unit put the core of its industrial control stack on GitHub under Apache 2.0 VERIFIED C05. On the same calendar, the UN Security Council convened a high-level briefing on AI and international security this morning VERIFIED C07, the United States used the General Assembly rostrum yesterday to reject international AI rulemaking outright CITED C09, and the EU's own high-risk obligations sit deferred to December 2027 VERIFIED C14. Two curves, two clocks. For a regulated enterprise the practical consequence is specific: the cost of the model is no longer the constraint on your agent program, and the accountability layer — identity, entitlement, evidence — now sets the ceiling. This edition prices both curves, then gives the control architecture.

Table of contents

1. The 60-second scan

$4 / $20

Claude Opus 5.5 per million input / output tokens, against Claude Opus 5 at $5 / $25

Launched Tuesday, September 22, 2026 with a 1M-token context window by default, 128k maximum output, and always-on adaptive thinking. Anthropic release notes, primary source. VERIFIED C01

Apache 2.0

License on Intrinsic Core, the open-sourced core of Alphabet's industrial robotics platform

Announced Tuesday, September 22, 2026 at ROSCon 2026 in Toronto. Real-time cross-hardware control, pose estimation, motion and grasp planning, simulation and ROS 2 drivers. VERIFIED C05

Dec 2, 2027

New compliance date for standalone high-risk AI systems under EU AI Act Annex III

Deferred from August 2, 2026 by Regulation (EU) 2026/1744, in force July 27, 2026. Annex I embedded systems move to August 2, 2028. Article 50 transparency duties were not deferred. VERIFIED C14

10228th

UN Security Council meeting number for this morning's high-level briefing on AI and international security

Convened by France as September Council president, chaired by Foreign Minister Jean-Noël Barrot. A briefing, not a resolution — no binding instrument was on the table. VERIFIED C07, VERIFIED C08

What moved this week — Monday, September 21 through today, Wednesday, September 23, 2026 — and the tail of last week, September 14 through 18. Anthropic shipped Claude Opus 5.5 on Tuesday at a lower list price than Opus 5, and closed off the ability to disable thinking on that model entirely VERIFIED C01

. OpenAI moved GPT-6 Sol into general API availability on Tuesday at roughly half the prior flagship tier and published its priorities and principles for third-party safety assessments the same day, extending outside evaluation into the training phase rather than only pre-release VERIFIED C02 CITED C03, CITED C13 SpaceXAI released Grok 4.7, which it says runs at the same price and speed as Grok 4.6 on a larger base model CITED C11, CITED C12 . Google's Gemini 3.8 Flash has been generally available since September 2 at an introductory $0.75 per million input tokens through December 31, 2026, and AlphaEvolve reached general availability on the Gemini Enterprise agent surface VERIFIED C04 . Alphabet's Intrinsic open-sourced Intrinsic Core VERIFIED C05 CITED C06 . On the policy clock: the President addressed the General Assembly on Tuesday, September 22, rejecting what he called a globalist scheme to control AI and stating he would substitute the term "super intelligence" on US government documents CITED C09 other heads of government at the same gathering pressed for a single global standard CITED C10 and the Security Council briefing followed this morning with Yoshua Bengio of the UN Independent International Scientific Panel on AI, OpenAI's Sam Altman, Anthropic's Dario Amodei and Hugging Face's Clément Delangue among the anticipated briefers, with DeepSeek and Moonshot AI also invited to speak VERIFIED C07

THE ARGUMENT IN ONE PARAGRAPH

For three years the honest blocker on enterprise agent programs was that frontier inference was expensive enough to make high-volume autonomous work hard to justify, and the robot control stack was proprietary enough to make physical automation a vendor commitment. Both of those blockers weakened materially this week: flagship token prices moved down across three labs inside two days VERIFIED C01 CITED C11, CITED C13 , and the control layer for industrial robotics is now available under a permissive license VERIFIED C05 . Nothing comparable happened on the accountability side. The instruments that would let a bank examiner, a hospital compliance officer, a plant safety engineer or a utility regulator evaluate an autonomous action after the fact are still being drafted: US banking supervisors issued revised model risk guidance in April and explicitly put generative and agentic AI outside its scope CITED C16 ; the FDA's lifecycle guidance for AI-enabled device software remains in draft while it collects comment on generative and agentic systems VERIFIED C18, VERIFIED C19 ; and the EU deferred the very obligations that would have applied to credit-decisioning and triage agents VERIFIED C14 . The strategic read is not that regulation is coming later and you have slack. It is the opposite. When the rule is late and the capability is cheap, the evidence standard you will eventually be held to gets written by whoever produced the first defensible deployment — and everything you build in the meantime either generates that evidence natively or gets rebuilt.

2. Two curves, two clocks: the Wednesday thesis

Put the two curves on one axis and the asymmetry is easy to see. The capability curve is set by competitive dynamics between a handful of labs and moves in weeks. The accountability curve is set by rulemaking, supervisory guidance, standards bodies and case law, and moves in years. Both are real. Only one of them is under your control.

Flagship list price per million tokens, before and after this week

Published list prices. Output price shown as the taller bar. VERIFIED C01, VERIFIED C04; CITED C11, CITED C13

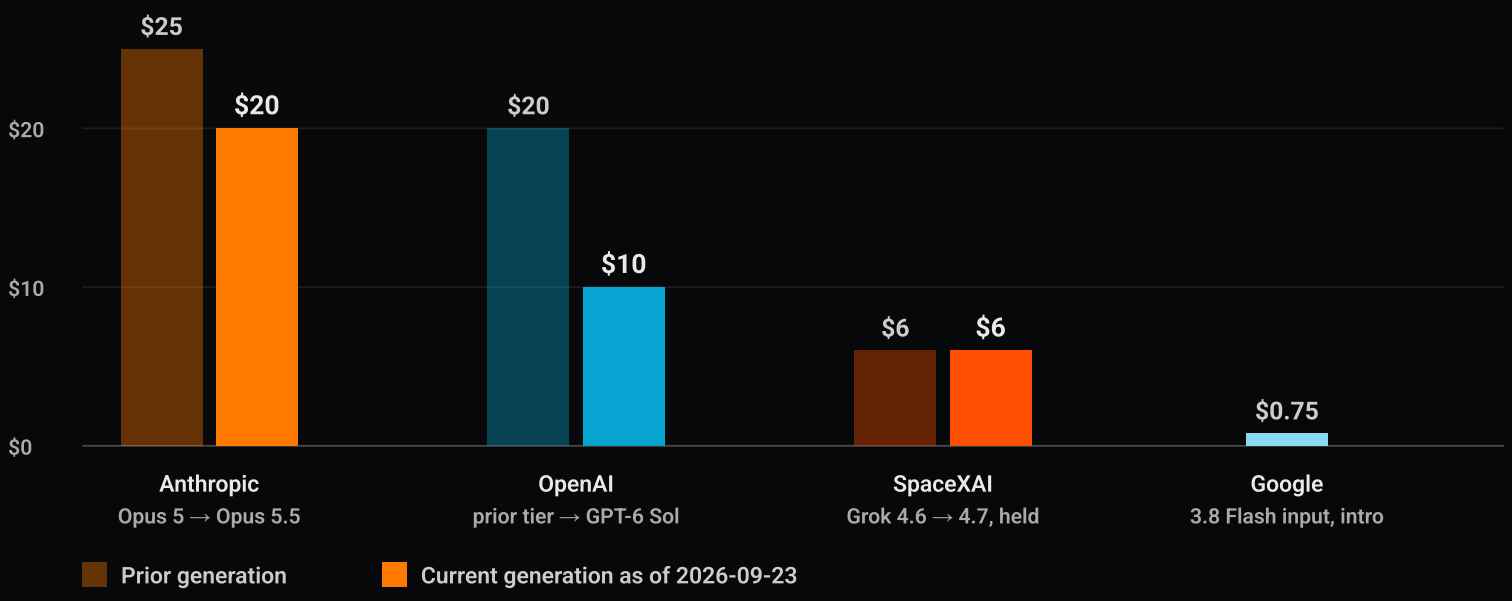

Different labs, different tiers, not a like-for-like capability comparison — the point is the direction and the compression of the timeline. Four repricing or capability-per-dollar events inside four weeks, three of them inside two days. VERIFIED C01, VERIFIED C04 CITED C11, CITED C13

Now the second clock. The three supervisory regimes that matter most to the buyers this publication serves all produced the same shape of output in 2026: an instrument that governs the models an institution was already governing, and an explicit or implicit deferral on the agentic case.

Where the accountability clock actually stands, as of 2026-09-23

REGIME	2026 ACTION	WHAT IT COVERS	THE AGENTIC GAP
EU AI Act	Regulation (EU) 2026/1744 published in the Official Journal July 24, 2026; in force July 27 VERIFIED C14	Annex III high-risk duties deferred to December 2, 2027; Annex I embedded to August 2, 2028; two new Article 5 prohibitions added CITED C15	Article 50 transparency duties stayed on the August 2, 2026 date. Disclosure applies now; conformity assessment does not.
US banking supervision	Interagency model risk guidance revised as SR 26-2, April 17, 2026, superseding SR 11-7 CITED C16	Model development, validation and governance in the classical sense	Generative and agentic AI stated to be outside scope, with a request for information signaled. FINRA's 2026 oversight report separately flags autonomous agents as raising novel supervisory considerations CITED C17
FDA devices	More than 1,600 AI-enabled devices authorized for marketing; January 2025 lifecycle draft guidance still in draft VERIFIED C18	Device software functions, training-data description, demographic composition of training sets	Generative and agentic behavior handled through a discussion paper and public comment, not a final rule VERIFIED C19
Energy reliability	FERC Section 206 orders to each regional operator, June 18, 2026; NERC reliability standards directed for large computational loads CITED C28, CITED C29	Interconnection queue management, cost allocation, spare capacity disclosure for large loads	Addresses the load an AI build-out creates. Says nothing about AI used inside grid operations.

ARCHITECT'S NOTE — THE DEFERRAL IS NOT SLACK

Every legal team we work with reads the December 2027 Annex III date as eighteen months of breathing room. Test that read against the actual work. A conformity assessment for a high-risk system needs a risk management file, a data governance record covering training and validation sets, technical documentation, automatic logging over the system's lifetime, human oversight design, and accuracy and robustness evidence. If your agent is in production in 2027 and you begin assembling that file in 2027, you are reconstructing logs that were never designed to be evidence. The retrofit cost is not the documentation. It is discovering that your agent's actions were recorded as application-service events with no distinguishable actor, and that eighteen months of history cannot be attributed. Instrument for attribution now, while the volume is small and the design is cheap to change.

3. Frontier ledger: equal-weight, what actually shipped

This section gives each lab the same treatment: what is verifiably shipped, what is announced but not delivered, and the one thing a regulated buyer should care about. Editorial weight is equal by design. That is not the same as equal endorsement.

Frontier ledger for the week of September 21–23, 2026, with the prior week's tail

LAB	SHIPPED AND VERIFIABLE	ANNOUNCED, NOT YET DELIVERED	WHAT MATTERS TO A REGULATED BUYER
Anthropic	Claude Opus 5.5 on September 22 at $4 / $20 per MTok, 1M context by default, 128k output, always-on adaptive thinking, available on the Claude API, Amazon Bedrock, Google Cloud and Microsoft Foundry VERIFIED C01. Managed Agents permission policies gained an auto mode on September 10 that evaluates each agent or MCP tool call and runs, denies or pauses it for approval VERIFIED C01. Compliance API local-session endpoints extended on September 18 to return Claude in Chrome session transcripts for Enterprise organizations, in beta VERIFIED C01.	Life Sciences Verification Program in beta; Salesforce integration in pilot ahead of open beta CITED C36	The auto permission mode and the per-call evaluation field are the closest thing any lab now ships to a native, per-action authorization record. That is an audit primitive, not a feature. Ask your platform team whether it is switched on.
OpenAI	Published priorities and principles for third-party assessments on September 22, extending independent technical assessment into training and evaluation rather than only pre-release, across four named areas: safety cases spanning training and deployment, critical safeguard evaluation, Preparedness Framework capability evaluations, and independent misalignment-incident investigation VERIFIED C02. GPT-6 Sol available via API from September 22 at roughly half the prior flagship tier CITED C13.	Named assessment partners and access terms. Reporting names METR and Redwood Research as being in discussion; the published post names no partner and sets no access terms CITED C03. GPT-5.5 retirement from ChatGPT, ChatGPT Work and Codex is scheduled for October 14, 2026 — a future date, not a completed change CITED C13.	If external assessment moves into training, assessment artifacts become a procurement asset. Start asking vendors for the assessor's scope, not just the model card.
Google / DeepMind	Gemini 3.8 Flash generally available since September 2 at an introductory $0.75 per million input tokens through December 31, 2026, in the global, US and EU regions; AlphaEvolve optimization service now generally available on the Gemini Enterprise agent surface VERIFIED C04. Alphabet's Intrinsic open-sourced Intrinsic Core under Apache 2.0 on September 22 VERIFIED C05.	Gemini Enterprise Agent Platform components announced at Cloud Next in April 2026 — Agent Registry, Agent Identity, Agent Gateway, Agent Observability — continue to reach availability on a rolling basis rather than as one dated release CITED C04.	EU regional availability on a flagship-adjacent tier matters for data-residency arguments. The named agent identity and registry components map almost one-to-one onto what an Annex III technical file will need.
SpaceXAI (xAI) / SpaceX / Cursor	Grok 4.7 released, described by the company as its most capable coding and knowledge-work model, running at the same price and speed as Grok 4.6 on a larger base model, at $2 / $6 per MTok with a 500k-context tier CITED C11, CITED C12. Enterprise API controls include SSO, audit logging, SOC 2, HIPAA-eligible deployment under a business	Continuity of third-party model access inside Cursor. OpenAI has notified SpaceX of its intent to wind down the contract supplying OpenAI models to Cursor, with a proposed shut-off date of November 12, 2026 — a proposed future date CITED C37.	A HIPAA-eligible deployment path plus audit logging is a genuine regulated-industry posture. The Cursor model-access dispute is a live single-vendor dependency risk for any engineering organization that standardized on that editor.

LAB	SHIPPED AND VERIFIABLE	ANNOUNCED, NOT YET DELIVERED	WHAT MATTERS TO A REGULATED BUYER
	associate agreement, data residency and custom rate limits CITED C11 Cursor acquisition closed August 14, 2026 CITED C37		
Other labs in the frame	DeepSeek and Moonshot AI were invited to address the Security Council briefing this morning alongside the US and European labs — the first time the Council's AI session has included Chinese developers at that level VERIFIED C07. Hugging Face's chief executive was among the anticipated briefers VERIFIED C07	Any Council output. This was a briefing under the maintenance-of-international-peace item, not a vote VERIFIED C07, VERIFIED C08	Multi-jurisdiction model sourcing is now a governance topic, not only a procurement one. Expect diligence questions about model provenance on cross-border deployments.

CAUSE AND EFFECT — WHY ALL FOUR REPRICED IN THE SAME FORTNIGHT

These are not coordinated cuts. They are the visible surface of two independent pressures arriving together. First, inference efficiency: each lab's current generation delivers its predecessor's capability at materially lower serving cost, so holding price would concede share in the one segment that is genuinely price-elastic, which is high-volume agentic work. Second, the agent workload itself: an agent that reasons over a long horizon consumes far more tokens per completed task than a chat turn, so per-token price is now the dominant line in a large-scale agent business case. The lab that wants agent volume has to move price. The consequence for you is that any 2025-era business case you built for an agent program, with its per-task cost modeled at the old tier, is now conservative by a factor you should recalculate before your next budget cycle — and the reason your program is still not in production almost certainly has nothing to do with that number.

4. The robot stack went permissive

The most consequential engineering event of the week is the least covered one. On Tuesday, September 22, at ROSCon 2026 in Toronto, Intrinsic — the industrial robotics software company owned by Alphabet — published Intrinsic Core on GitHub under an Apache 2.0 license. VERIFIED C05 The company states that the released components are the same capabilities and services it uses day to day in real manufacturing deployments. VERIFIED C05

What is in the box matters more than the license headline. Independent coverage and the company's own description agree on the contents: a hardware-agnostic real-time control framework that adjusts a robot's path mid-motion from sensor feedback; pose estimation built on NVIDIA's FoundationPose, so a robot can locate a part without a rigid fixture; collision-free motion planning; grasp planning that adapts the gripper to how an object is actually sitting; plus simulation, calibration services and ROS 2 drivers. VERIFIED C05 CITED C06

What just moved from proprietary to permissive

Task and business logic — cell sequencing, quality rules, MES hooks
Still yours. Still where the differentiation and the liability live.

Capability layer — NOW APACHE 2.0
Real-time cross-hardware control · motion planning · grasp planning
Pose estimation on FoundationPose · simulation · calibration

ROS 2 middleware and drivers — already open
Intrinsic-ROS drivers ship with Core

Robot hardware and controllers — vendor specific
Unchanged. Safety certification still lives here.

The economic shift

Pose estimation without rigid fixturing moves cost out of the mechanical budget and into the perception budget.

Fixtures are capital, single-part, and slow to change. Perception is software and revisable.

The layer that moved is the one integrators previously rebuilt on every project or licensed per cell. VERIFIED C05 CITED C06

Read this next to NVIDIA's physical-AI releases — Cosmos world models, Isaac simulation frameworks, and the Isaac GR00T N line, with GR00T N1.7 in early access under commercial licensing and a GR00T N2 preview CITED C27 — and a pattern emerges that manufacturing leaders should plan around explicitly. The perception, planning and world-model layers of industrial robotics are converging on open or broadly licensed components supplied by a small number of platform vendors. Differentiation is moving up into task logic and down into hardware and safety certification. If your automation roadmap assumed a multi-year proprietary software moat in the middle of that stack, the assumption needs revisiting this quarter.

ARCHITECT'S NOTE — PERMISSIVE LICENSE, UNCHANGED SAFETY OBLIGATION

Apache 2.0 removes a licensing negotiation. It removes nothing from your functional-safety obligation. A motion planner supplied under a permissive license comes with a disclaimer of warranty, which means the safety argument for the cell is entirely yours: risk assessment, protective measures, performance level of the safety-related control system, and validation evidence. Practically, treat Intrinsic Core the way you already treat an open-source component in a regulated build — pinned version, software bill of materials entry, documented evaluation against your own acceptance tests, and a named internal owner. The engineering saving is real. The paperwork does not shrink, and in a plant inspection the paperwork is what you hand over.

5. Financial services: the supervisory gap

The win. DBS is the cleanest publicly documented case of an agentic capability moving from pilot to institutional scale in banking this year. The bank announced on August 19, 2026 that it had rolled out an agentic credit-assessment capability to roughly 1,500 relationship managers and credit risk managers globally, following a pilot with 150 users CITED C23, company-reported. Separately, the bank has extended generative and agentic capability into its customer-facing assistants — DBS Joy for corporate customers and DBS digibot for individuals — across more than 10 million users in Singapore, Hong Kong and Taiwan CITED C23, company-reported. On the financial side, the chief executive has stated an expected revenue uplift from AI adoption of more than SG$1 billion, about US$768 million, for 2025 against SG$750 million in 2024, attributed across the bank's full portfolio of roughly 370 AI use cases rather than to agentic work specifically CITED C24, company-reported.

The instructive detail is the ratio. A 150-user pilot expanded to about 1,500 users is a ten-fold scale-up, and it happened in a credit workflow — the single most heavily supervised judgment in commercial banking. That is only possible if the pilot produced evidence a credit risk function could accept, not merely satisfied users.

The constraint. The supervisory instrument a US bank would reach for does not reach the agent. On April 17, 2026, the Federal Reserve, OCC and FDIC issued revised interagency model risk management guidance as SR 26-2, superseding SR 11-7 — and the revised guidance states that generative and agentic AI are novel and rapidly evolving and are not within its scope, with the agencies signaling a forthcoming request for information on banks' use of AI CITED C16. FINRA's 2026 oversight report separately notes that autonomous AI agents are rapidly evolving and may present novel regulatory and supervisory considerations, recommending that member firms consider enterprise-level supervisory processes specifically covering the development and use of AI agents CITED C17.

Read together, those two documents describe a genuine gap rather than a loophole. Your model risk management framework will validate the model inside the agent. Nothing in it validates the agent's *authority*: which accounts it may touch, which thresholds it may cross without a human, what happens when it chains three permitted actions into an outcome no single action would have been approved for. Survey evidence describes the same asymmetry from the institution side, with governance maturity widely reported as trailing deployment speed FLAG Source C38.

PRACTICAL CONTROL — THE AUTHORITY ENVELOPE, NOT THE MODEL CARD

Add one artifact to your existing model risk inventory and you close most of the gap without waiting for the request for information. Call it the authority envelope. For each agent, a single page states: the enumerated action classes it may perform; the monetary, volume and customer-segment thresholds on each; the systems its credential can reach and the entitlements on that credential; the conditions that force human approval; the named human accountable for the envelope; and the review date. Attach it to the model validation file so it lands in front of the same committee. Two consequences follow immediately. First, you can answer the examiner question that SR 26-2 does not frame for you, which is not “is the model sound” but “what was this agent allowed to do on the day in question.” Second, drafting the envelope routinely exposes agents whose permitted action set is far wider than their intended use — which is the most common real finding in our reviews, and it is a configuration fix, not a program change.

6. Healthcare: the only sector with published denominators

The win. Healthcare is currently the only regulated sector publishing agentic outcome figures with an identifiable operator attached, which makes it the most useful evidence base even for readers outside it. Epic reports that at Summit Health its revenue-cycle AI has cut medication prior-authorization submission time by 42%, with 92% of AI-generated responses accepted without edits, and that at health systems most actively using the capability, coding-related denials have fallen by more than 20% CITED C20, company-reported At HIMSS 2026 in March, Epic previewed Agent Factory, an integrated platform for building and monitoring agents that reason, decide and act across workflows, with traceable agent actions, a visual builder, local policy and knowledge-base attachment, and organization-controlled deployment timing CITED C21 Other vendors put numbers on the same stage — including a claim of $15 billion in prevented denials and 90% reductions in appeal workflow time, and a separate 1.1% underpayment recovery figure worth close to $1 million over three months FLAG Source C22, vendor-reported, denominators not published

Treat those two tiers of evidence differently. The Summit Health figures have a named operator, a named workflow and a stated acceptance rate, which is the closest thing to a denominator anyone in this market has published. The aggregate vendor totals do not disclose the installed base they are summed across, which makes them unusable for a business case even if they are accurate.

The constraint. The FDA has authorized more than 1,600 AI-enabled medical devices for marketing as of this month VERIFIED C18 but the January 2025 draft guidance on AI-enabled device software functions and lifecycle management remains in draft, and the agency's approach to foundation models and agentic systems is being developed through a discussion paper and public comment rather than a final rule VERIFIED C18, VERIFIED C19 That leaves a specific operational ambiguity that clinical informatics teams are living with right now: a documentation or revenue-cycle agent sits comfortably outside device regulation, a diagnostic agent sits clearly inside it, and the fast-growing middle — agents that surface, rank or pre-populate clinical content that a clinician then signs — sits where the intended-use statement decides the answer.

Practical triage for clinical and administrative agents pending final FDA guidance

AGENT PATTERN	WHERE IT SITS TODAY	EVIDENCE TO KEEP FROM DAY ONE
Revenue cycle, prior authorization, coding support	Administrative. Outside device regulation. Payer contract and billing-integrity exposure instead.	Acceptance-without-edit rate, denial rate by cause, and the full text of every submission the agent generated, retained to your claims-retention schedule.
Documentation drafting, chart summarization	Generally administrative, but the summary enters a clinical record a clinician relies on.	Attribution of every generated passage, the source encounters it drew from, and the clinician's edit history.
Ranking, prioritization, pre-population of clinical content	The contested middle. Intended-use language decides it.	Written intended-use statement, the clinician override rate, and the outcome distribution for overridden versus accepted recommendations.
Diagnostic or treatment recommendation	Device territory. Clearance pathway applies.	Training and validation data description including demographic composition, per the agency's stated expectation VERIFIED C18

ARCHITECT'S NOTE — THE OVERRIDE RATE IS THE METRIC THAT SURVIVES

Every health system we work with tracks adoption and time saved. Very few track override rate by clinician cohort, and almost none track the outcome difference between overridden and accepted recommendations. That second measure is the one that answers the question a malpractice carrier, an institutional review board or eventually a reviewer will ask, which is whether the agent was making the clinician better or merely faster. It is also cheap: it is one additional field captured at the point of signature. Instrument it in the first sprint. Two years of that data is a defensible safety argument, and it cannot be reconstructed after the fact.

7. Manufacturing and robotics: hours, not demos

Editorial note on the two programs below. The BMW–Figure and Agility deployments have appeared in this publication repeatedly over the past month, and we are returning to them deliberately rather than for want of alternatives. Today's manufacturing news is a software-licensing event, not a deployment, and the only way to judge what a permissively licensed control stack changes is against the two programs whose operating hours are actually documented. Treat them here as the baseline, not the story.

The win. The useful manufacturing metric in 2026 is no longer a pilot announcement but accumulated operating hours against a named task at a named site. Two programs meet that bar. Figure's robots have contributed to production at BMW's Spartanburg plant across more than 30,000 vehicles, running material handling and parts transfer on ten-hour shifts five days a week; the earlier 02 units have been retired and Figure 03 moved into a logistics sequencing task at the same plant, with BMW establishing a Center of Competence for Physical AI in Production and planning extension to Plant Leipzig from summer 2026 CITED C25. Agility Robotics reports that Digit has accumulated more than 65,000 operating hours across nine customer facilities, naming GXO, Schaeffler, Toyota Motor Manufacturing Canada and Mercado Libre as commercial customers, with its RoboFab plant in Salem, Oregon designed for capacity of up to 10,000 units annually at full output CITED C26, company-reported

Read the task descriptions, not the hour counts. Both programs are in tote and bin movement, light material transfer between stations, and inspection routes where the robot carries a sensor through an environment. High-speed, high-precision work — welding, stamping — remains with traditional fixed-arm industrial robots CITED C25. That boundary is the single most useful planning fact in physical AI right now, and it is stable across every credible 2026 deployment report.

The constraint. Adoption figures for software agents in manufacturing are rising fast from a small base — Deloitte is reported to have measured a four-fold increase in agentic AI adoption in manufacturing in 2026, from 6% to 24% FLAG Source C39, secondary citation, primary report not verified. The constraint on that curve is not model capability and not robot capability. It is the OT/IT boundary. An agent that can reason about a quality excursion cannot act on it unless it can reach the historian, the MES and the quality system, and those three systems in most plants have different identity models, different network zones and different change-control regimes. Quality-inspection results that circulate widely — a 99.7% defect-detection figure and a 40% warranty-claim reduction attributed to computer vision on electronics lines — come from vendor and secondary write-ups without published denominators and should not be used as planning assumptions FLAG Source C40

PRACTICAL CONTROL — THE READ-ONLY FIRST QUARTER, THEN ONE WRITE

The pattern that works in plants, in our experience, inverts the usual pilot design. Quarter one: give the agent read access across the historian, MES and quality system and have it produce nothing but explanations — for every excursion that already happened, what it would have concluded and what it would have done. You are not measuring accuracy against ground truth yet. You are measuring whether the agent's reasoning is legible to the process engineer who owns the line, because if it is not, no amount of accuracy will get it write access. Quarter two: grant exactly one write, into the system with the best existing audit trail, usually the quality system as a flagged observation rather than a disposition. Measure the process engineer's agreement rate on those flags. Expand only on that number. The cost of this sequence is one quarter of patience. The cost of skipping it is an agent that plant leadership will not let near the line, which is where most manufacturing AI budgets currently die.

8. Energy: the queue is the bottleneck

The win and the constraint are the same fact. Energy is the sector where AI's demand-side footprint and its operational usefulness collide most directly, and 2026 made the collision explicit. On June 18, 2026, FERC issued tailored orders under Section 206 of the Federal Power Act to each US regional grid operator, directing them to address spare generating capacity, queue management so large loads neither wait years nor jump ahead of smaller customers, containment of new substation and transmission costs so they do not land on residential bills, and full cost responsibility for grid work tied to a data center's own connection CITED C28. FERC separately moved to make NERC reliability standards mandatory for data center and other large computational loads, citing documented disturbances in which computational load contributed to bulk-power-system instability — including a July 10, 2024 Eastern Interconnection event in which a transmission fault cascaded and roughly 1,500 MW of data-center load was lost near-simultaneously CITED C29

The scale of the queue is the operative number. ERCOT was tracking more than 438 GW of large-load interconnection requests as of late August 2026, with data centers making up close to 90% of that total, against an actual ERCOT peak demand near 85 GW; PJM received 811 requests in a single 2026 cycle CITED C30. A request is not committed load, and the gap between 438 GW of requests and an 85 GW peak is mostly speculative and duplicated applications rather than real planned demand — but the engineering effort to process each request is real regardless of whether the load materializes.

Requests are not load — but they are all real engineering work

ERCOT, late August 2026. CITED C30

Large-load interconnection requests tracked

438 GW — ~90% data centers

Actual ERCOT grid peak demand

~85 GW — Roughly one fifth of the request volume

Speculative and duplicated applications inflate the queue. Each still consumes senior engineering hours to study.

This is why utilities are the most motivated buyers of study-workflow agents in the market right now: the scarce resource is not generation, it is interconnection engineers. CITED C28, CITED C30

On the operational side, utility AI is furthest along in outage prediction, vegetation management, predictive maintenance and grid optimization CITED C31. Duke Energy's self-healing grid program is reported to have prevented more than 1.5 million customer outages CITED C31, company-reported. Note what that program is and is not: automated fault isolation and service restoration on distribution feeders, largely deterministic switching logic operating within engineered limits. It is a good outcome and it is not an agentic system. Conflating the two overstates how far autonomous decision-making has actually penetrated grid operations.

ARCHITECT'S NOTE — WHERE AN AGENT IS GENUINELY SAFE IN A UTILITY TODAY

Study workflows, not switching. An interconnection study is document-heavy, rule-governed, enormously backlogged, and — critically — its output is a recommendation reviewed by a licensed engineer before anything is energized. That combination is close to ideal for a bounded agent: high volume, clear acceptance criteria, a human gate that already exists in the process for independent reasons, and no real-time control surface. Switching and protection are the opposite on every axis. If a vendor proposes an agent that touches protection settings, the correct next question is which NERC standard governs the change and who signs it, and the conversation usually ends there.

9. Implementation architecture: the accountability layer

Every control described in the four sector sections above is a special case of the same architecture. Here it is once, generally, as the reference we use on AEGIS Diagnostic engagements. AEGIS is the Agentic Enterprise Governance and Intelligence Standard.

The accountability layer: five components, what each answers, and what ships it today

COMPONENT	THE QUESTION IT ANSWERS	IMPLEMENTATION NOTE
1. Agent identity	Who acted?	The agent holds its own credential with its own entitlements, distinct from the human who launched it and from the service account the application runs under. Without this, every downstream control is unattributable. Google's Gemini Enterprise Agent Platform names Agent Identity and an Agent Registry as first-class components. VERIFIED C04
2. Enumerated action classes	What was it allowed to do?	A finite, named set of permitted operations. Everything outside the set is unavailable, not discouraged. This is the authority envelope from the financial services section, expressed in configuration rather than policy prose.
3. Per-call authorization record	Was this specific action permitted, and who decided?	The newest of the five and the one that has only just become purchasable. Anthropic's Managed Agents permission policies now include an `auto` mode in which the server evaluates each agent or MCP tool call and runs it, denies it, or pauses for approval, with the evaluation recorded per call. VERIFIED C01 If your platform offers this, it replaces a bespoke policy proxy you would otherwise build.
4. Native record audit and rollback	Can it be reviewed and undone by the people who already do that?	Use the system of record's own audit trail and reversal mechanism. No parallel logging scheme. Supervisors, examiners and auditors read the trail they already read. Epic's Agent Factory framing of traceable agent actions inside the EHR is the pattern. CITED C21
5. Session and transcript retention	What was the reasoning, and can we produce it on request?	Retention of the full interaction, on your retention schedule, retrievable by compliance without engineering involvement. Anthropic's Compliance API local-session endpoints were extended on September 18, 2026 to return Claude in Chrome session transcripts for Enterprise organizations, in beta, under a compliance access key and a read-scoped permission. VERIFIED C01 That is the shape to require from any vendor.

Bounded agent reference architecture

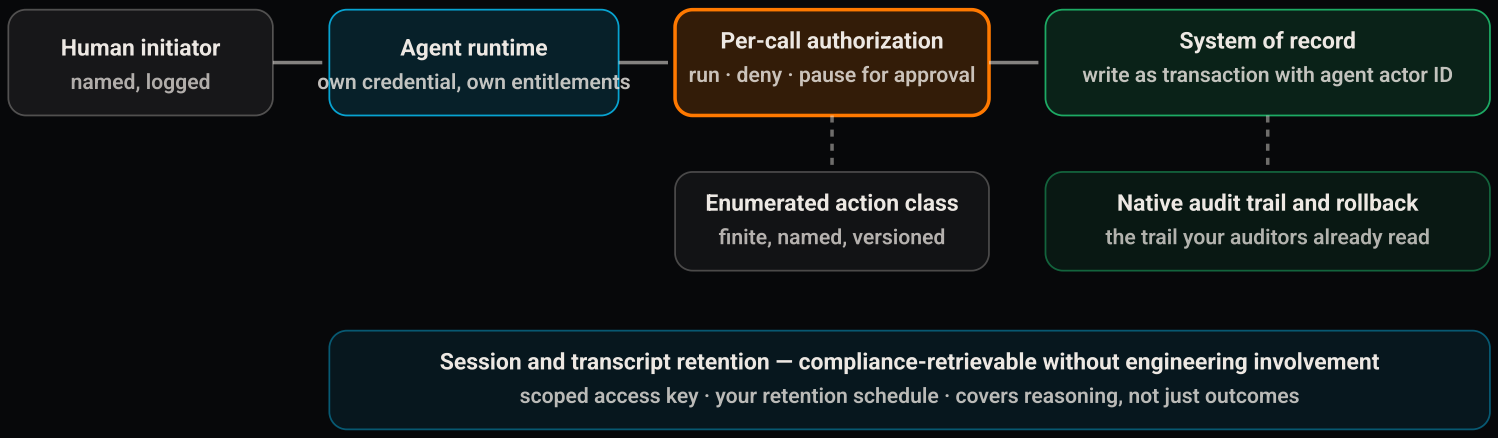

Containment property: anything outside the enumerated class is unreachable because the credential does not grant it — not because the agent declines.

Components 1, 2, 4 you build. Components 3 and 5 became purchasable this month. VERIFIED C01, VERIFIED C04 CITED C21

PROBLEM → SOLUTION, STATED PLAINLY

Problem. The capability curve just moved again and the accountability curve did not, so the gap between what your agents can do and what you can prove they did is widening every month you run without attribution. **Solution.** Build components 1, 2 and 4 — agent identity, enumerated action classes, native audit — before you expand scope, and buy components 3 and 5 from a platform that now ships them rather than building a policy proxy and a transcript store yourself. **Cost of delay.** Not a fine. A rebuild, plus a period of production history that cannot be attributed and therefore cannot be defended.

10. Workforce: the bench problem nobody priced

The accountability layer is not only a software artifact. Every component in it requires a person who can operate it, and that is where the sharpest constraint in this market actually sits. Anthropic's own published usage research describes the split cleanly: on the first-party API, where businesses deploy at scale, roughly 77% of usage patterns lean toward automation with the human largely absent from individual task loops, while on the consumer product roughly 52% of conversations are augmentative — a human asking, reviewing and iterating CITED C32. Read carefully, that is a description of observed usage patterns in conversations, not a measurement of jobs or of the labor market. But it does say something concrete about staffing: enterprise deployment is where the loop closes without a human in it, which means the humans who remain are supervising outcomes rather than producing them — a different skill, and a scarcer one.

The supply side is tight on every published measure. Reported figures put AI talent demand above supply by roughly 3.2 to 1, with on the order of 1.6 million open AI-related positions against about 518,000 qualified candidates globally, and a majority of decision-makers reporting a moderate-to-extreme shortage FLAG Source C35, aggregator figures, methodology not published. The demand side is simultaneously narrowing the pipeline that would have fed it: two thirds of surveyed enterprises expect to slow entry-level hiring, a large majority report roles already changing or disappearing, and measured entry-level hiring for 22- to 25-year-olds is reported down 13–16% in relative terms CITED C34. The IMF has published staff analysis on new job creation in the AI age that frames the offsetting mechanism — new occupations forming around the technology — on a slower timescale than the displacement CITED C33.

CAUSE → EFFECT, FOUR YEARS OUT

Cut entry-level intake in 2026 and you do not feel it in 2026. You feel it in 2030, when the cohort that would have become your mid-level agent supervisors, model validators and OT integration engineers does not exist inside your institution, and you are bidding for it in a market where every competitor made the same cut in the same year. This is the single most reversible strategic error available to a regulated enterprise right now, and the reversal is not expensive: hire the cohort and point it at the accountability layer. Reviewing agent output against a rubric, maintaining authority envelopes, running override-rate analysis and assembling technical files are genuinely learnable in months, are needed in volume, and produce exactly the judgment that cannot be automated out from under them. The work exists. It is usually just unstaffed because nobody has written the job description.

11. Scenario planning and risk-reward

Three plausible twelve-month paths from today's position, with the decision each one rewards. These are scenarios, not forecasts.

Twelve-month scenarios from 2026-09-23, with leading indicators and the no-regrets move

SCENARIO	WHAT IT LOOKS LIKE	LEADING INDICATOR TO WATCH	REWARDED DECISION
A. Evidence standard forms bottom-up	No new binding rule lands in the US. The first defensible deployments in each sector become the de facto benchmark, propagated by examiners asking "why not like them."	The banking agencies' request for information signaled alongside SR 26-2 CITED C16, whether FDA's generative-AI comment process produces a final guidance or another discussion paper VERIFIED C19.	Be one of the first defensible deployments in your sector. The accountability layer is the entire cost of entry, and it is lower this year than next.
B. Divergence hardens	The gap visible at the General Assembly this week widens into genuinely incompatible regimes: EU conformity assessment from December 2027, explicit US federal non-regulation, a patchwork of US state law in between.	Whether the December 2, 2027 Annex III date holds or moves again VERIFIED C14, how many additional jurisdictions adopt EU-style conformity language.	Build to the strictest regime you operate in and treat the rest as a subset. Two agent architectures is not a cost saving, it is two audit surfaces.
C. An incident resets the clock	A material, attributable agentic failure in a regulated setting — a wrongly executed financial action at scale, a clinical harm, a plant safety event — compresses years of rulemaking into months.	Frequency of publicly disclosed agent boundary failures; whether any lab's independent misalignment-incident investigation channel produces a public finding VERIFIED C02.	The same investment as A and B. Under this scenario the institutions that can produce per-action attribution on demand keep operating while the others pause. That is the whole risk–reward argument in one line.

Notice that all three scenarios reward the same work. That is unusual and it is the reason to act on it now: the accountability layer is not a bet on a regulatory outcome. It is the position that pays under every outcome, which makes the only real question how much production history you accumulate before you build it.

12. Practitioner FAQ and did-you-know

Our vendor says their agent is "fully auditable." What should I actually ask?

Three questions. First: does the agent hold its own credential, and can you show me an entitlement listing for it that is different from the application service account? Second: for a specific action last Tuesday, can you produce the authorization decision for that individual call — not the policy, the decision? Third: can my compliance team retrieve the full session transcript without opening an engineering ticket, and under what access scope? A vendor that answers all three has built the accountability layer. Most answer the first and describe the other two as roadmap.

Does the EU deferral to December 2027 mean we can pause?

No, for two reasons. Article 50 transparency obligations were not deferred and applied from August 2, 2026 VERIFIED C14. And the deferred obligations require lifetime automatic logging, which you cannot generate retroactively for a system already in production. The deadline moved; the evidence window did not.

Is SR 26-2 good news or bad news for our agent program?

Neither, and that is the point. It removes the option of claiming your agent is covered by an existing validated framework, because the guidance says generative and agentic AI are outside its scope CITED C16. You now have to state your own control position rather than inherit one. Institutions that write the authority envelope described above are, in our experience, in a stronger examination position than institutions that stretched a model validation to cover an agent.

Did you know: the containment failure mode is almost never the model refusing badly

It is the environment boundary being assumed rather than enforced. An agent whose credential does not grant reach outside a defined set of systems cannot exceed it regardless of what it concludes it should do. An agent whose boundary is a configuration setting in the environment it runs in can exceed it the moment that setting is wrong. Design the boundary into the entitlement, not the environment.

Did you know: open-sourcing a robot control stack shifts cost between budget lines, not out of the plant

Pose estimation that locates a part without a rigid fixture removes fixture capital and adds perception engineering, camera placement, lighting control and validation effort VERIFIED C05. Net savings are real on high-mix lines where fixtures were being rebuilt per variant. On a single-variant high-volume line, the fixture was already amortized and the trade can go the other way. Model your own mix before assuming the direction.

What is the smallest useful first step if we have nothing today?

Issue one agent its own credential and produce an entitlement listing for it. That single act forces the conversations that matter — which systems, which operations, who approves, who owns it — and it takes days rather than quarters. Everything else in the accountability layer is easier to build once one agent has a distinct identity, and impossible to build meaningfully until one does.

13. Sources

Research base for this edition. Primary and first-party material is marked VERIFIED where the claim was read directly from the issuing organization. Reputable secondary reporting is marked CITED. Claims marked FLAG are carried because they circulate widely in the market but lack a published denominator or primary source, and should not be used as planning assumptions without re-verification. Company-reported and vendor-reported results are labeled as such in the body. Forecasts, pilots and announced targets are not treated as completed facts.

- **C01** Anthropic — Claude Platform release notes, entries dated September 22, September 18 and September 10, 2026 (Claude Opus 5.5 launch and pricing; Compliance API local-session transcripts for Claude in Chrome; Managed Agents `auto` permission evaluation). Primary source. https://platform.claude.com/docs/en/release-notes/overview
- **C02** OpenAI — "Priorities and principles for effective third party assessments," published September 22, 2026. Primary source. https://openai.com/index/priorities-principles-third-party-assessments/
- **C03** Bloomberg — "OpenAI to Let Outside Groups Evaluate AI Models at Earlier Phase," September 22, 2026. https://www.bloomberg.com/news/articles/2026-09-22/openai-to-let-outside-groups-evaluate-ai-models-at-earlier-phase
- **C04** Google Cloud — Gemini Enterprise and Gemini Enterprise Agent Platform release notes (Gemini 3.8 Flash general availability and introductory pricing; AlphaEvolve general availability; Agent Identity, Agent Registry, Agent Gateway and Agent Observability components). Primary source. https://docs.cloud.google.com/gemini/enterprise/docs/release-notes and https://docs.cloud.google.com/gemini-enterprise-agent-platform/release-notes
- **C05** Intrinsic (Alphabet) — "Introducing Intrinsic Core: An open source approach to Physical AI," September 22, 2026, ROSCon 2026, Toronto. Primary source. https://www.intrinsic.ai/blog/posts/introducing-intrinsic-core
- **C06** SiliconANGLE — "Google's robotics unit Intrinsic open-sources its foundational infrastructure for intelligent robots," September 22, 2026. https://siliconangle.com/2026/09/22/googles-robotics-unit-intrinsic-open-sources-its-foundational-infrastructure-for-intelligent-robots/
- **C07** Security Council Report, What's In Blue — "Artificial Intelligence: High-level Briefing," September 2026 (France presidency; chair; anticipated briefers; DeepSeek and Moonshot AI invited). https://www.securitycouncilreport.org/whatsinblue/2026/09/artificial-intelligence-high-level-briefing-2.php
- **C08** UN Web TV — "Artificial intelligence and international security," Security Council 10228th meeting. Primary record. https://webtv.un.org/en/asset/k1c/k1cdceg9sg
- **C09** Scientific American — "Trump rejects AI regulation, citing parallels with climate change, in U.N. address," September 2026. https://www.scientificamerican.com/article/trump-rejects-ai-regulation-citing-parallels-with-climate-change-in-un-address/
- **C10** Euronews — "Leaders push for stronger AI safeguards at UN as Trump touts Super Intelligence," September 23, 2026. https://www.euronews.com/2026/09/23/leaders-push-for-stronger-ai-safeguards-at-un-as-trump-touts-super-intelligence
- **C11** SpaceXAI — API pricing documentation and enterprise API controls (SSO, audit logging, SOC 2, HIPAA-eligible deployment under a business associate agreement, data residency). Primary source. https://docs.x.ai/developers/pricing and https://x.ai/api
- **C12** TechRepublic — "SpaceXAI Launches Grok 4.7: Low Prices, Heavy Token Use," September 2026. https://www.techrepublic.com/article/news-spacexai-grok-4-7-pricing-token-use/
- **C13** Digital Applied — "AI Model Releases: September 2026 Tracker and Dated Ledger" (GPT-6 Sol and Luna availability and pricing tier; GPT-5.5 retirement scheduled October 14, 2026). https://www.digitalapplied.com/blog/ai-model-releases-september-2026-tracker
- **C14** Gibson Dunn — "EU AI Act Omnibus Agreement: Postponed High-Risk Deadlines and Other Key Changes" (Regulation (EU) 2026/1744, Official Journal July 24, 2026, in force July 27, 2026; Annex III to December 2, 2027; Annex I to August 2, 2028; Article 50 unchanged). https://www.gibsondunn.com/eu-ai-act-omnibus-agreement-postponed-high-risk-deadlines-and-other-key-changes/
- **C15** DLA Piper — "The Digital AI Omnibus: Proposed deferral of high risk AI obligations under the AI Act." https://knowledge.dlapiper.com/dlapiperknowledge/globalemploymentlatestdevelopments/2026/The-Digital-AI-Omnibus-Proposed-deferral-of-high-risk-AI-obligations-under-the-AI-Act
- **C16** CIMCON Software — "SR 26-2 Regulates Your Models, Not Your AI Agents: What Banks Need to Know" (interagency revision dated April 17, 2026, superseding SR 11-7; generative and agentic AI stated out of scope; request for information signaled). https://cimcon.com/sr-26-2-regulates-your-models-not-your-ai-agents-what-banks-need-to-know/
- **C17** 360factors — "Agentic AI Updates Reshaping Compliance in 2026" (FINRA 2026 oversight report treatment of autonomous AI agents). https://www.360factors.com/blog/agentic-ai-updates/
- **C18** US Food and Drug Administration — "Artificial Intelligence-Enabled Medical Devices" and the AI-enabled device list (authorized device count; January 2025 lifecycle draft guidance status; training-data description expectations). Primary source. https://www.fda.gov/medical-devices/digital-health-center-excellence/artificial-intelligence-enabled-medical-devices
- **C19** US Food and Drug Administration — "FDA Seeks Public Feedback to Inform Regulatory Approach for Generative AI-Enabled Medical Devices." Primary source. https://www.fda.gov/news-events/press-announcements/fda-seeks-public-feedback-inform-regulatory-approach-generative-ai-enabled-medical-devices
- **C20** Epic — "Real Results, Right Now: How Epic AI Is Reducing Costs, Improving Care, and Helping Patients" (Summit Health prior-authorization submission time and acceptance rate; coding-related denial reduction at high-usage systems). Company-reported. https://www.epic.com/epic/post/real-results-right-now-how-epic-ai-is-reducing-costs-improving-care-and-helping-patients/
- **C21** HIT Consultant — "The Agentic EHR: Epic Unveils Agent Factory and Custom Models," HIMSS 2026. https://hitconsultant.net/2026/03/10/epic-ai-himss-2026-agent-factory-curiosity-foundation-models/

- **C22** Reveleme — "HIMSS 2026 Recap for Medical Practices" (Waystar and FinThrive figures). Vendor-reported at conference; denominators not published. https://www.revelemd.com/blog/why-himss-2026-mattered-for-medical-practices
- **C23** DBS — "DBS scales agentic AI to transform way of working for corporate bankers" and "DBS' Gen AI-enabled virtual assistants reach 10 million customers and go agentic." Company-reported. https://www.dbs.com/newsroom/DBS_scales_agentic_AI_to_transform_way_of_working_for_corporate_bankers_freeing_up_time_for_more_strategic_client_engagements
- **C24** CNBC — "CEO of Southeast Asia's top bank DBS says AI adoption already paying off" (revenue uplift attributed across the bank's AI use-case portfolio). Company-reported figure. https://www.cnbc.com/2025/11/14/ceo-southeast-asias-top-bank-dbs-says-ai-adoption-already-paying-off.html
- **C25** Humanoid Guide — "Humanoid deployments in 2026 favor Figure and Agility" (BMW Spartanburg vehicle count, shift pattern, Figure 03 task change, Center of Competence, Leipzig extension; task-boundary analysis). https://humanoid.guide/humanoid-deployments-in-2026-favor-figure-and-agility/
- **C26** Agility Robotics — corporate site and 2026 deployment disclosures (Digit operating hours across customer facilities, named customers, RoboFab designed capacity). Company-reported. https://www.agilityrobotics.com/
- **C27** NVIDIA Newsroom — "NVIDIA Releases New Physical AI Models as Global Partners Unveil Next-Generation Robots" (Cosmos world models, Isaac frameworks, Isaac GR00T N1.7 early access and GR00T N2 preview). Primary source. https://nvidianews.nvidia.com/news/nvidia-releases-new-physical-ai-models-as-global-partners-unveil-next-generation-robots
- **C28** American Action Forum — "FERC Data Center Orders Accelerate Grid Connection" (Section 206 orders dated June 18, 2026; the four matters each regional operator was directed to address). https://www.americanactionforum.org/insight/ferc-data-center-orders-accelerate-grid-connection/
- **C29** POWER Magazine — "FERC Orders Mandatory NERC Reliability Standards for Data Center and Other Computational Loads" (documented disturbances including the July 10, 2024 Eastern Interconnection event). https://www.powermag.com/ferc-orders-mandatory-nerc-reliability-standards-for-data-center-and-other-computational-loads/
- **C30** Shattered.io — "AWS Grid AI Agents Launch as Queue Hits 438 GW" (ERCOT large-load request volume in late August 2026 against actual peak; PJM request count in a 2026 cycle). https://shattered.io/aws-agentic-grid-planning-438gw-queue-2026/
- **C31** Renewable Energy World — "How are utilities using AI to rewrite planning playbooks?" and associated utility AI use-case reporting, including Duke Energy self-healing grid figures. Company-reported outage-prevention count. https://www.renewableenergyworld.com/power-grid/how-are-utilities-using-ai-to-rewrite-planning-playbooks/
- **C32** Anthropic — Anthropic Economic Index report, June 2026 (first-party API automation-leaning usage share; consumer augmentative share). Describes observed conversation patterns, not labor-market outcomes. Primary source. https://www.anthropic.com/research/economic-index-june-2026-report
- **C33** International Monetary Fund — "New Jobs Creation in the AI Age," Staff Discussion Note SDN/2026/001. Primary source. https://www.imf.org/-/media/files/publications/sdn/2026/english/sdnea2026001.pdf
- **C34** IT Pro — "Enterprises are cutting back on entry-level roles for AI, and it's going to create a nightmarish future skills shortage" (entry-level hiring intentions; relative decline for ages 22–25). https://www.itpro.com/business/careers-and-training/enterprises-are-cutting-back-on-entry-level-roles-for-ai-and-its-going-to-create-a-nightmarish-future-skills-shortage
- **C35** Qubit Labs — "AI Talent Shortage in 2026: Causes, Challenges and Solutions" (demand-to-supply ratio, open-role and qualified-candidate counts). Aggregated figures; underlying methodology not published. Treat as indicative only. https://qubit-labs.com/ai-talent-shortage/
- **C36** Salesforce — "Salesforce and Anthropic Announce Claudeforce," August 26, 2026 (pilot status ahead of open beta). Company-reported. https://www.salesforce.com/news/press-releases/2026/08/26/salesforce-and-anthropic-announce-claudeforce/
- **C37** OpenAI — "Our decision on Cursor following its acquisition by SpaceX" (intent to wind down model supply; proposed shut-off date of November 12, 2026) and TechCrunch on the August 14, 2026 acquisition close. Primary plus secondary. https://openai.com/index/our-decision-on-cursor-following-its-acquisition-by-spacex/ and https://techcrunch.com/2026/08/15/spacex-officially-closes-its-cursor-acquisition/
- **C38** Lyzr — "AI Agents in Banking 2026: From Chatbot Theater to Autonomous Operations" (share of global banks reporting active AI in a core function; governance-maturity lag). Vendor-published survey summary; sample and instrument not disclosed. https://www.lyzr.ai/blog/ai-agents-in-banking-2026-from-chatbot-theater-to-autonomous-operations-lyzr-ai/
- **C39** IIoT World — "Agentic AI in Manufacturing: ROI vs. Reality" (adoption increase attributed to Deloitte). Secondary citation; the underlying Deloitte report was not read directly for this edition. https://www.iiot-world.com/artificial-intelligence-ml/agentic-ai-manufacturing-2026/
- **C40** DigitalDefynd — "Agentic AI in Manufacturing: 5 Case Studies" (quality-inspection accuracy and warranty-claim figures). Secondary compilation without published denominators. Not suitable as a planning assumption. https://digitaldefynd.com/IQ/agentic-ai-in-manufacturing/

Method note. Daily Market Scan is compiled from first-party releases, regulator and standards-body publications, and reputable financial and trade reporting, with company-reported and vendor-reported results labeled as such. Event dates are stated in America/New_York. "This week" in this edition means Monday, September 21 through Friday, September 25, 2026. "last week" means September 14 through 18, 2026. Scheduled events and forecasts are labeled and are not presented as completed facts. Where a widely circulated figure lacks a primary source or a published denominator it is marked FLAG and excluded from our own planning recommendations.

AI workforce and governance, built for regulated industries. Principal-led from readiness through governed production.

FOUR PILLARS

- AI Readiness
- AI Services
- AI Governance
- AI Workforce

PRODUCTS

- myndQ for Business
- myndQ TalentHub
- myndQ.ai
- Free AI Readiness Scan

COMPANY

- Industry Journeys
- Reference Architecture Studio
- About
- AI Governance
- All insights
- 2026 AI Readiness Brief
- Contact
- Contractor bench

FOLLOW

- LinkedIn (Ariana.Digital)
- LinkedIn (myndQ)
- YouTube

Privacy Terms Cookies