AI OUTLOOK 2026–2029

THE AUTONOMY DIVIDEND

The adoption race is over. The next three years decide who converts autonomous AI into a compounding economic return — and who pays its costs without collecting it.

FIRST EDITION · JUNE 2026
LOCKEDINLABS.AI/RESEARCH

36 PAGES · 21 EXHIBITS
120+ PRIMARY SOURCES

ABOUT THIS REPORT

Built to be argued with.

The Autonomy Dividend is Lockedin Labs' flagship outlook on artificial intelligence: where the technology, the capital, the agent workforce, the constraints, and the rules are heading between mid-2026 and 2029. It is written for operators — boards, executive teams, and builders who must make capital and organizational commitments before the evidence is complete.

Most AI reporting describes what just happened. This report is structured around what organizations must decide **before the next 12–36 months happen**: ten specific, dated calls we can be held to, three scenarios with explicit probabilities and signposts, and an operating playbook for capturing the dividend that autonomy creates.

We take positions. Where we forecast, we say so plainly and attach numbers a reader can hold us to in 2027. Where the evidence is contested — capability evals, capex sustainability, labor effects — we present the strongest version of both sides before stating ours.

Method

This edition synthesizes more than 120 primary sources published through **June 12, 2026**: regulatory filings and earnings disclosures; survey research from McKinsey, KPMG, Deloitte, Gartner, IDC, the US Census Bureau, and the Federal Reserve; measurement work from METR, Epoch AI, and the Stanford HAI AI Index 2026; energy analyses from the IEA, EPRI, and Goldman Sachs; and primary policy texts from the EU, the White House, and state legislatures. Every exhibit carries its sources. Forecasts labeled **"Lockedin Labs estimate"** are ours alone.

How to read the numbers

Survey figures measure different populations with different definitions — "deployment" spans everything from one pilot agent to multi-agent orchestration. We state each source's definition inline and never splice incompatible series into a single trend line.

A note on velocity

This field invalidates reports quickly. Three frontier model generations, one EU implementation delay, two US executive orders, and a record funding quarter occurred in the six months before publication. Figures are as-of dated; the argument is built on curves, not snapshots, because curves age better.

The one-sentence thesis

Between 2026 and 2029, the binding constraint on AI value shifts from **capability** to **verification, power, and trust** — and the organizations that industrialize those three become the compounding winners of the autonomy era.

Rights & contact

research@lockedinlabs.ai · lockedinlabs.ai/research
Cite as: Lockedin Labs (2026), The Autonomy Dividend: AI Outlook 2026–2029, 1st ed.

CONTENTS

Inside this report

21 exhibits, built from primary data

Every chart in this report is constructed from named, dated sources — regulatory filings, earnings disclosures, and measurement research — and carries its attribution inline. A consolidated, numbered reference list begins on page 35.

Selected exhibits

How to navigate

Short on time? Read the Executive Summary (p.4), the Ten Calls (p.5), the three Scenarios (p.27), and the Autonomy Dividend Index (p.29). Together they are the report in ten minutes.

EXECUTIVE SUMMARY

Everyone adopted AI. Almost no one is being paid for it. That ends now.

Eighty-eight percent of organizations now use AI. Six percent are capturing meaningful value from it. The gap between those numbers is the defining economic fact of 2026 — and the opportunity of the next three years.

Call it what it is: the adoption era is over. Generative AI reached half the world's population faster than the PC or the internet. Enterprise spending crosses **$2.59 trillion in 2026**, up 47% in a single year. Hyperscaler capital expenditure approaches **$700 billion** — roughly quadruple 2022 — and AI investment supplied about **two-thirds of US GDP growth in Q1 2026**, the largest technology contribution since 1999. The inputs are no longer the story.

The story is conversion. McKinsey finds only 39% of organizations see any enterprise EBIT impact from AI, and for most of them it is under 5%. MIT's now-canonical 2025 finding — 95% of generative AI pilots produced zero measurable P&L return — was not a verdict on the technology. It was a verdict on organizations that bought capability without building the apparatus to convert it.

That apparatus now has a name. We call the return it produces **the autonomy dividend**: the recurring surplus created when work is delegated to AI systems whose task horizon, reliability, and unit cost all improve on compounding curves — net of the costs of verification, governance, energy, and incidents.

Three curves make the dividend real. **Task horizons**: the length of work frontier models complete autonomously has doubled roughly every four to seven months for six years; frontier systems now sustain multi-hour engineering tasks, and the measurement suite itself tops out at 16 hours. **Cost**: inference at constant capability fell roughly tenfold in the last year alone; GPT-4-class output that cost ~$30 per million tokens in 2023 costs under $1 today. **Deployment**: enterprises with AI agents in production jumped from 11% to 40% in five quarters (KPMG). Capability compounds, cost collapses, delegation spreads.

Against those curves stand three walls. **Verification** — agent reliability still decays sharply under repetition, "workslop" already taxes desk workers roughly $186 per employee per month, and governance maturity is improving at a fraction of capability's pace. **Power** — US data-center demand is on track to double by 2027 while gas-turbine slots are effectively sold out through 2030. **Trust** — the first AI-orchestrated espionage campaign, the first AI-discovered vulnerability classes at scale, and the first deepfake midterm cycle have all arrived before the rules have.

The Autonomy Dividend, defined

Dividend = (Task horizon × Reliability × Delegation rate) − (Verification cost + Risk cost + Energy cost)

Every term on the left is compounding in the adopter's favor. Every term on the right is rising for the unprepared. Strategy for 2026–2029 is the discipline of maximizing the left side while industrializing the right.

The pages that follow make ten specific calls, sized with the best available numbers, then map three scenarios for 2027–2029 and the playbook we believe separates dividend collectors from dividend payers.

TEN CALLS · 2026–2029

What we expect to be true — and when you can check

Forecasts are only useful if they can be wrong. Each call below is dated and numbered; we will grade them publicly in future editions.

01 **2026 is the divide year: spending crosses $2.59T while most value concentrates in fewer than one in ten firms.**

Worldwide AI spending grows 47% to **$2.59T** (Gartner) while only **6%** of organizations qualify as AI high performers (McKinsey). The spread between leaders' and laggards' AI-attributable EBIT widens through 2027 before diffusion narrows it — late movers will pay 2026 prices for 2024 results.

02 **Agent task horizons reach a human workweek by late 2027; month-scale work follows by 2028–2029.**

The 50%-success task horizon has doubled every ~4–7 months since 2019 (METR); frontier systems already sustain multi-hour tasks and have outgrown the 16-hour eval ceiling. On the post-2024 doubling rate, **~40-hour autonomous work arrives in 2027**; METR's own trend places month-long tasks between 2027 and 2031. Delegation becomes an org-design problem, not a model problem.

03 **The verification gap becomes the largest single source of AI losses by 2027 — bigger than model spend for laggards.**

Capability doubles in months; surveyed governance maturity rose just **2.0 → 2.3 on a 4-point scale in a year** (McKinsey). The costs already visible in that wedge — rework on low-quality AI output (~**$186/employee/month**, BetterUp/Stanford), agent incidents, canceled projects (Gartner: **>40%** of agentic projects by end-2027) — compound faster than license fees.

04 **Power, not chips, sets the ceiling from 2027. Electricity politics becomes AI politics.**

US data-center demand roughly doubles to **66 GW by 2027** (Goldman Sachs); data centers reach **9–17% of US electricity by 2030** (EPRI); gas-turbine build slots are effectively **sold out through 2030** (GE Vernova) and PJM capacity prices rose ~11x in two years. Compute strategy becomes energy strategy, and ratepayer backlash becomes a regulatory force.

05 **The agent workforce becomes countable — and audited. ”Digital employees” enter annual reports by 2028.**

Enterprises in production with agents jumped **11% → 40% in five quarters** (KPMG); Gartner projects **40% of enterprise apps** ship task-specific agents by end-2026 and **~15% of day-to-day work decisions** made autonomously by 2028. Banks already issue agents identities and access credentials (BNY's 134 "digital employees"). Agent identity, inventory, and least-privilege become board-level controls.

06 A capex digestion phase hits in 2027: growth decelerates hard, and depreciation honesty gets forced.

2026's ~**$700B** hyperscaler build (+~75% YoY) is increasingly debt-financed — AI-related issuance is set to roughly double to **~$570B in 2026** (Morgan Stanley). OpenAI has already cut its infrastructure framing from $1.4T to ~$600B and pivoted to renting. We expect 2027 capex growth in the teens, mark-downs on short-lived accelerators (the Burry critique), and selective unwinds of circular financing — digestion, not detonation (see Scenarios, p.28).

07 Security becomes AI-vs-AI by default; un-agented defense becomes professionally negligent by 2028.

The first documented AI-orchestrated espionage campaign executed **80–90% autonomously** (Anthropic, Nov 2025); a frontier model then surfaced **~10,000 high-severity vulnerabilities** across critical software (Project Glasswing, 2026); indirect prompt-injection attacks are now observed in the wild (Unit 42). Both attack and defense run at machine speed from here; insurers and regulators will price that in.

08 Regulation bifurcates: EU enforcement lands December 2027; the US stays voluntary through the midterms.

The EU's Digital Omnibus delayed high-risk obligations to **Dec 2, 2027 / Aug 2, 2028**, with AI-content transparency from **Dec 2, 2026**. The June 2026 US executive order builds a **voluntary** frontier-review framework and a DOJ task force against state laws. Compliance becomes a market-access cost, not a moat — but incident-reporting regimes (California SB 53, NY RAISE) quietly create the first public dataset of agent failures.

09 Labor reallocation turns from anecdote to statistic — and enters the 2028 political cycle as a first-order issue.

AI-attributed US job cuts in the first five months of 2026 (**87,714**) already exceed all of 2025; AI was the #1 stated layoff reason for three straight months (Challenger). Yet Gartner finds **no correlation between AI layoffs and AI ROI**, and half of service organizations that cut will rehire by 2027. Expect agent-management roles, AI dividends, robot-tax proposals, and displacement frameworks on 2028 ballots.

10 Raw capability commoditizes by 2028. The durable moats are verification, distribution, data rights, and power.

The US–China frontier gap collapsed from as much as ~17–32% to **2.7%** in three years (Stanford HAI); Chinese open-weight models went from ~1% to ~**30%** of aggregator usage; frontier-class coding performance is now available MIT-licensed. When everyone can buy intelligence, advantage moves to what intelligence cannot buy: verified workflows, customer access, proprietary data, and contracted electrons.

$2.59T
Worldwide AI spending, 2026 — up 47% YoY (Gartner, May 2026)

40%
Enterprises with AI agents deployed in production, Q1 2026 — vs 11% a year earlier (KPMG)

~4–7 mo
Doubling time of autonomous task horizons (METR, 2019–2026 trend)

2.3 / 4
Average responsible-AI maturity score, 2026 — capability compounds, governance crawls (McKinsey)

PART ONE · THE STATE OF AUTONOMY

Universal adoption, concentrated reward

Where AI actually stands in mid-2026: what 88% adoption conceals, how fast the frontier is still moving, and why the money — $700 billion of it this year — keeps coming.

CH.1

The Divide

Adoption is table stakes; conversion is rare. The anatomy of the 6%.

CH.2

The Curve

Task horizons, collapsing costs, and benchmarks that can no longer keep up.

CH.3

The Capital

A $700B build-out, the debt behind it, and the honest version of the bubble question.

88% adopted. 39% see earnings impact. 6% are winning.

The most important chart in enterprise AI is a funnel, and it leaks at every stage.

By late 2025, **88% of organizations** reported regular AI use in at least one function and **79%** used generative AI regularly (McKinsey Global Survey, Nov 2025). The US Census Bureau's firm-weighted series — which counts every small business in America — reached **19.8%** by May 2026, but weight firms by employment and the picture converges with the executive surveys: **78% of the US labor force** now works at a firm that has adopted AI (Atlanta Fed, Nov 2025), and **40.7% of workers** personally use generative AI on the job.

Then the funnel narrows brutally. Only about a third of organizations describe their AI programs as scaling; **7%** say AI is fully scaled across the enterprise. **39%** attribute any EBIT impact to AI — and most of those put it below 5% of earnings. Just **6%** meet McKinsey's bar for high performers: significant value with ≥5% of EBIT attributable to AI.

MIT's 2025 finding that **95% of generative AI pilots produced zero measurable P&L return** drew criticism on methodology, but its core diagnosis survived scrutiny: the failures were not model failures. They were **learning failures** — tools that don't retain context, workflows never redesigned around them, and no measurement of whether output was actually better. External partnerships succeeded ~67% of the time; internal-only builds, ~22%.

EXHIBIT 01

The conversion funnel is the strategy map

Share of organizations at each stage of AI value capture, late 2025–2026

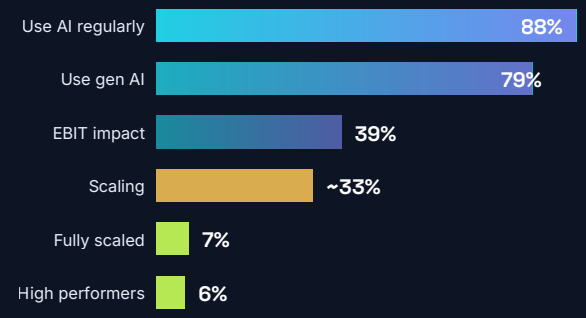

Source: McKinsey Global Survey on AI (n=1,993; fielded Jun–Jul 2025, published Nov 2025). "High performers" = significant value and ≥5% of EBIT attributable to AI.

78%

of US labor force works at an AI-adopting firm (Atlanta Fed, employment-weighted, Nov 2025)

95%

of 2025 gen-AI pilots showed zero P&L impact (MIT NANDA) — a learning failure, not a model failure

EXHIBIT 02

The agent leap: pilots collapsed into production

Share of enterprises by agentic AI stage, KPMG AI Pulse (orgs ≥$1B revenue)

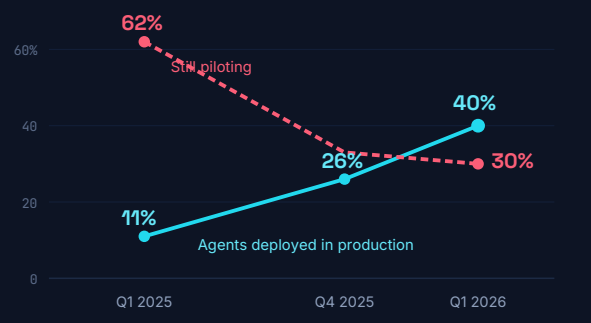

Q1 2026 detail: 40% deploying · 9% scaling · 6% orchestrating multi-agent system

Source: KPMG AI Quarterly Pulse Survey, Q1 2025–Q1 2026 (published Mar 2026). Definitions vary across surveys: McKinsey (Nov 2025) finds 23% scaling agents somewhere; Capgemini (Jul 2025) found 2% at full scale — state definitions when citing.

The middle of the funnel is finally moving. KPMG's quarterly panel of large enterprises shows agent deployment more than tripling in five quarters while the share stuck in pilots halved. Average planned AI spend in that panel: **$207M per organization** over the next 12 months, roughly double a year earlier. And **79%** say AI remains a top investment priority even in a recession.

What distinguishes the 6%? Across the survey base, four behaviors recur — and none of them is "better models." High performers **redesign workflows before deploying** (the single practice most correlated with EBIT impact in McKinsey's data); they assign **explicit accountability** for responsible AI (orgs with clear ownership score 2.6/4 on maturity vs 1.8 without); they **measure agent output** the way they measure human output; and they buy or partner where they lack learning-loop advantages rather than building vanity stacks.

The divide is also sectoral. Information (39.7%) and finance (33.9%) lead US adoption; accommodation and food service (8%) trail (Census, May 2026). Worker-level adoption in finance hit **63%** and information **70%** (Real-Time Population Survey). Financial-sector adoption grew **127% in a year**. The leaders are not waiting for clarity; they are manufacturing it.

> "The constraint on AI value was never intelligence. It was the organization's ability to absorb it."
>
> LOCKEDIN LABS · THE AUTONOMY DIVIDEND

What the 6% do differently

1. Redesign the workflow, then add the agent — never the reverse.
2. Name an accountable owner for AI risk and value (maturity 2.6 vs 1.8 without one).
3. Instrument everything: delegation rate, rework rate, cost per outcome.
4. Partner for capability, build for context — external builds succeed ~3x more often.

EXHIBIT 03

A $2.59 trillion year — nearly half of it infrastructure

Worldwide AI spending, 2025–2026E, with 2026 composition

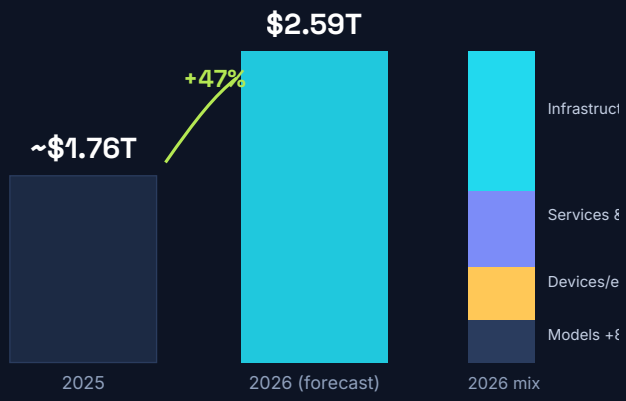

Source: Gartner press releases, Jan 15 & May 19, 2026 (2026 total $2.59T, +47%; AI infrastructure >45% of spend; gen-AI model spending +80.8% to >$32B). 2025 base implied. Mix shares are indicative of Gartner's segment commentary, not exact splits.

EXHIBIT 04

Infrastructure spend is compounding at >50%

Worldwide AI infrastructure spending, actual and forecast, $B

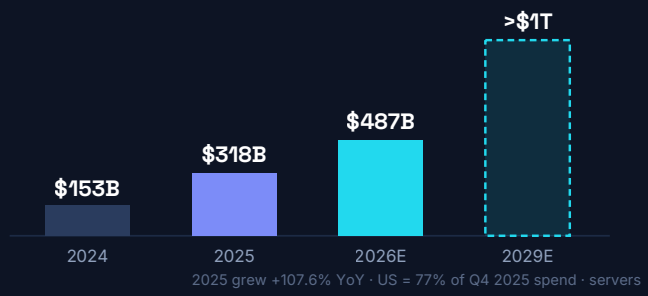

2025 grew +107.6% YoY · US = 77% of Q4 2025 spend · servers

Source: IDC AI Infrastructure Tracker (Apr 16, 2026). 2029 forecast >$1T, ~31% five-year CAGR. Separately, IDC projects agentic-AI-driven spend of $1.3T by 2029 — >26% of all worldwide IT spending.

The spending divide mirrors the value divide — and 2026 is the inflection Gartner has been promising.

Two numbers frame the year. Total AI spending grows **47%** to **$2.59T** — roughly the GDP of France — while enterprise spending on AI models and agents **more than doubles**. AI infrastructure alone, $153B in 2024, reaches **$487B in 2026** on its way past $1T in 2029 (IDC). "2026 will be the inflection year for enterprise spend," in Gartner's own framing — the year AI stops being an experiment line and becomes the IT budget.

Geographic concentration is extreme: the US accounted for **77%** of Q4 2025 AI-infrastructure spend; China declined 8% under export controls and domestic substitution; Middle East & Africa grew over **500%** from a small base as sovereign programs broke ground.

For operators, the implication of Chapter 1 is uncomfortable but clarifying: **budget is no longer a differentiator.** At $2.59T of global spend, your competitors have the same models, the same clouds, and increasingly the same agents. The divide is created — and will keep being created — by conversion capacity. That capacity is built against curves, which is where we turn next.

> **Watch in H2 2026**
>
> Gartner's "inflection year" thesis is testable: if enterprise model + agent spend hasn't roughly doubled by Q4 (from ~$6B incremental), the digestion phase has arrived early — re-weight toward Scenario 2 (p.28).

The machines are working longer shifts — doubling every four to seven months

One measurement explains more of the next three years than any other: how long a task an AI system can complete on its own. That horizon has compounded relentlessly since 2019 — and it just outgrew the instruments built to measure it.

EXHIBIT 05 · THE CENTRAL CHART OF THE OUTLOOK

Autonomous task horizon: from seconds to working days — next stop, the workweek

Length of software/reasoning tasks completed at 50% success rate, by model release date (log scale), with Lockedin Labs extrapolation band

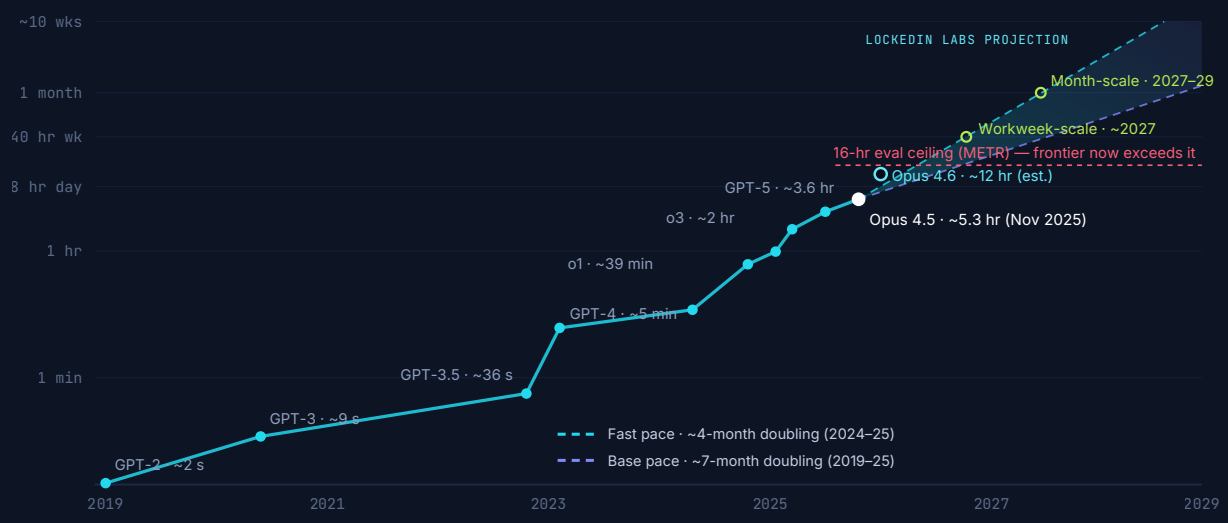

50%-reliability task length by model release. Doubling time compressed from ~196 days (full period) to ~89 days post-2024.

Sources: METR "Time Horizon 1.1" (Jan 29, 2026) and METR time-horizons dashboard (updated May 8, 2026); Opus 4.6 estimate is third-party and unofficial. Horizon = task length human experts need, completed by the model at 50% reliability, primarily software/research tasks.

Read Exhibit 05 the way an operator should: every doubling halves the supervision a unit of work requires.

A task horizon of five minutes makes AI a typing aid. Five hours makes it a contractor you check on at lunch. Five days — where the band points by 2027 — makes it a **direct report**. Five weeks makes it a team. Each step changes what "managing" means, and none of them requires a research breakthrough; they require only that a six-year trend hold for two to three more years. The post-2024 pace (doubling every ~89 days) is the fastest segment of the trend, not the slowest.

Two honest caveats. First, the 50%-success bar: a system that completes a week-long task half the time is transformative and dangerous — the difference is verification (Chapter 5). At 80%-success thresholds, horizons run roughly 4–5x shorter. Second, the instruments are failing upward: METR now states plainly that measurements above 16 hours are unreliable on the current task suite. The frontier has outgrown its ruler.

Meanwhile the static benchmarks are simply being finished. SWE-bench Verified went from ~60% to ~95% in about a year. ARC-AGI-2 — designed to resist memorization — went from 15.9% to ~85% in eleven months. Humanity's Last Exam crossed 50%. And in June 2026, an AI-assisted audit of FrontierMath found **42% of its original problems flawed**, forcing a full re-issue. The era of trusting a single number is over.

The eval-integrity problem is now a business problem

When public benchmarks saturate or break, vendor scores stop predicting production value. The 2026 evidence: a ~37% average gap between lab benchmark and production performance, and 50x cost variation at similar accuracy (arXiv 2511.14136). The only benchmark that matters is the one you run on your own workflows — see the ADI metrics, p.31.

EXHIBIT 06

Benchmarks are being finished, not just beaten

Score progression on two hard evals, July 2025 – June 2026

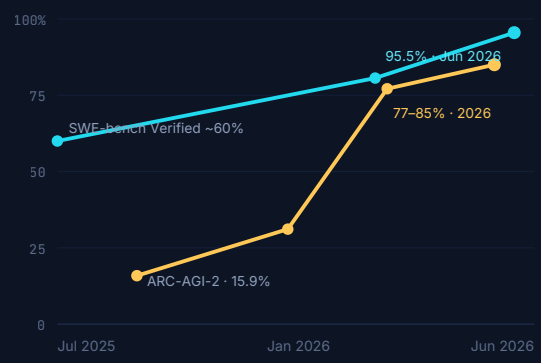

Sources: Stanford HAI AI Index 2026 (SWE-bench Verified "60% to near 100% in a single year"); llm-stats and ARC Prize reporting (Grok 4 15.9% Jul 2025; Gemini 3 Pro 31.1% Nov 2025; Gemini 3.1 Pro 77.1% Feb 2026; GPT-5.5 ~85 Jun 2026); Anthropic (Claude Mythos 5, 95.5% SWE-bench Verified, Jun 9, 2026).

42%

of original FrontierMath problems found flawed in the June 12, 2026 AI-assisted audit → full benchmark re-issue (Epoch AI)

~37%

average gap between lab benchmark and production agent performance (arXiv 2511.14136)

> ”The frontier has outgrown its ruler. From here, the benchmark that matters is your own workflow.”
>
> LOCKEDIN LABS

EXHIBIT 07

The price of intelligence is collapsing ~10x per year

Cost of GPT-4-class output, $ per million tokens (log scale)

Sources: Stanford HAI AI Index 2026 (~10x/yr decline at constant capability); arXiv 2511.23455 (GPT-4-level: ~$30/Mtok early 2023 → <$1; frontier-quality bin declining ~31x/yr); DeepSeek V4-Flash pricing ($0.14–0.28/Mtok, Apr 2026). Points are representative anchors on a continuous decline.

The countertrend matters just as much: **total** inference bills are rising. Reasoning models burn more tokens per task, agents run continuously, and memory prices jumped 40–50% into 2026. Cheap intelligence does not mean cheap intelligence programs — it means consumption explodes to fill the budget. Unit-economics discipline (cost per delegated outcome, not cost per token) is what separates the dividend from the bill.

> **Jevons, operationalized**
>
> Plan procurement against the curve: any workload uneconomical today is ~10x cheaper in 12 months at constant quality. Contract short, benchmark quarterly, and never amortize a use case against today's token prices.

EXHIBIT 08

The labs already run on their own product

Share of new code written by AI, self-reported, 2024 → mid-2026

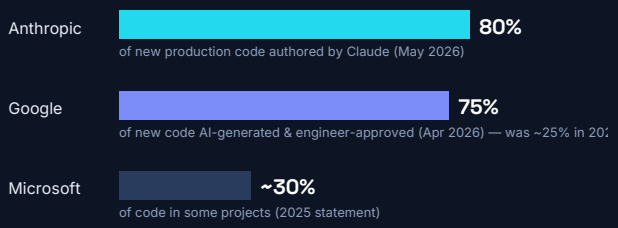

Sources: Anthropic via VentureBeat (May 2026; engineers ship ~8x their 2024 daily output); Sundar Pichai, Google Cloud Next 2026 (Apr 2026); Satya Nadella (2025). Self-reported, definitions differ.

This is the quiet recursive story of 2026: **AI is now the majority author of the systems that build AI.** Anthropic reports Claude writes over 80% of code merged into its own systems and explicitly warns that automated AI R&D "could come sooner than most institutions are prepared for." Google attributes three-quarters of new code to AI with engineer approval. The software industry is the leading indicator for every other knowledge industry: the same delegation pattern — AI drafts, human verifies, ratio shifts quarterly — is the template your functions inherit next.

Capability concentration is also dissolving. The best US–China model gap is down to **2.7%** (from as much as ~17.5–31.6% in 2023), achieved with roughly **23x less investment**; Chinese open-weight models now carry **~30%** of usage on aggregator platforms; frontier-adjacent coding ability ships MIT-licensed (DeepSeek V4: 80.6% SWE-bench Verified at $0.435/Mtok). **Intelligence is commoditizing. Advantage is migrating up the stack** — to verification, distribution, data rights, and power. That migration is the through-line of this report.

Three releases a quarter, a gated tier above the market, and an open-weight shadow frontier

LAB / LINE	LATEST (AS OF JUN 12, 2026)	SIGNAL	WATCH NEXT
Anthropic · Claude	Opus 4.8 (May 28) · **Fable 5 / Mythos 5** (Jun 9) — first "Mythos-class" GA release	SWE-bench Pro 80.3% (vs 69.2% Opus 4.8); #1 Artificial Analysis index (64.9); priced 2x Opus ($10/$50 per Mtok). The notable pattern: a **gated frontier** — Mythos Preview ran with ~50 vetted partners before public Fable launch with safety classifiers.	Whether rivals copy the gated-tier release pattern; enterprise uptake at 2x price
OpenAI · GPT-5.x	GPT-5.5 (Apr 23) · 5.5-Instant default in ChatGPT (May 5)	Terminal-Bench 82.7%; ~52% fewer hallucinated claims vs prior default; cadence now ~6 weekly point releases/yr. S-1 filed Jun 8 — discipline now serves an IPO narrative.	GPT-5.6 (reported June); IPO pricing (target $730–850B+)
Google · Gemini 3.x	Gemini 3.1 Pro (Feb 19) · 3.5 Flash (May 19)	GPQA Diamond 94.3% — highest recorded; ARC-AGI-2 77.1% (2.3x its own 3-month-old score); strongest price-performance at the fast tier; TPU v7 gives structural cost edge.	Gemini 4; agentic commerce push; TPU-driven price war
xAI · Grok	Grok 4.3 beta (Apr 17); Grok 5 delayed past Q1 target	2M-token context at –60% price; video understanding; trained on ~1.5 GW Colossus. Merged into SpaceX (Feb 2026).	Grok 5 (claimed 6T-param MoE, unconfirmed)
Meta · post-Llama	Contested: Llama-successor delayed; proprietary "Muse Spark"; Mango/Avocado roadmap H1 2026	Strategic identity crisis: open-weights leadership ceded to China while Superintelligence Labs reorganizes. LeCun departed to found world-model startup AMI Labs (~$1B+ seed).	Whether Meta re-enters open weights or goes fully closed
China open-weight bloc	DeepSeek V4-Pro/Flash (Apr 24, MIT) · Kimi K2.6 · GLM-5.1 · Qwen 3.6	V4: 80.6% SWE-bench Verified at ~$0.44/Mtok — frontier-adjacent, free to fork. Kimi K2.6 first open model to beat a GPT-5.x point release on SWE-bench Pro. GLM-5.1 reportedly trained on zero NVIDIA silicon. ~30% of aggregator usage; 41% of Hugging Face downloads.	Whether US enterprises adopt at scale despite policy friction; V4.5 trained on Ascend

The gated frontier is the new release pattern

2026's most consequential product decision wasn't a benchmark — it was Anthropic selling its strongest model (Mythos) under access controls while shipping a classifier-wrapped public version (Fable). Capability above the public frontier now exists commercially. Plan on a world where the best model you can buy is not the best model that exists — and where access tiers become a competitive variable, like cloud regions or export licenses.

Open weights are a strategy, not a charity

China's open-weight surge (~1% → ~30% of usage in 18 months) is industrial policy: commoditize the layer you trail in, win the diffusion layer you can lead. For buyers it's leverage — a permanent ~10–20x price floor under closed vendors. For strategists it means **model capability cannot be your moat by 2028**; the moat question moves to Chapter 5.

$700 billion, mostly borrowed, chasing $13 billion of revenue

The build-out is the largest private infrastructure project in history, and it is increasingly financed with debt against revenues that, for the marquee buyer, do not yet exist. Here is the honest version of the bubble question.

EXHIBIT 09

Hyperscaler capex has quadrupled since GPT-4

Combined Big Four capital expenditure, $B (Microsoft, Alphabet, Amazon, Meta)

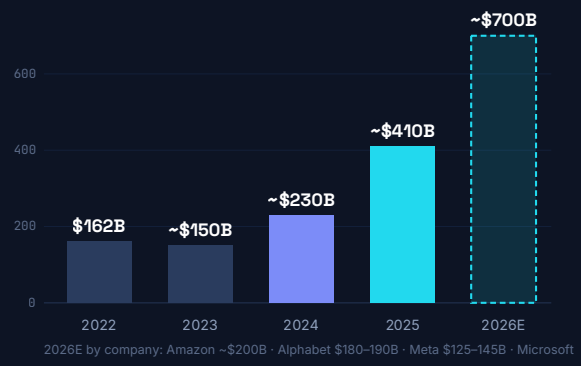

2026E by company: Amazon ~$200B · Alphabet $180–190B · Meta $125–145B · Microsoft

Sources: company Q4 2025 / Q1 2026 earnings & 8-Ks; CNBC (Feb 6, 2026); Epoch AI hyperscaler capex trend (Feb 26, 2026, ~72%/yr growth since Q2 2023); Statista/ Tom's Hardware ($725B upper estimate, +77% YoY). 2023 dip reflects pre-build pause.

The scale is genuinely without precedent. Combined Big Four capital spending goes from **$162B in 2022 to roughly $700B in 2026** — a quadrupling in four years, growing faster than 70% annually since GPT-4's release (Epoch AI). Add Oracle and the figure clears $700B; Morgan Stanley sees it pushing past $800B. AI investment is now a macroeconomic force: it supplied roughly **two-thirds of US GDP growth in Q1 2026**, the largest technology contribution since 1999, and data-center construction alone has at times accounted for the majority of quarterly growth.

Three facts make this build different from the fiber boom it is often compared to.

One: it is increasingly debt-financed. AI-related debt issuance is set to roughly double to **~$570B in 2026** (Morgan Stanley). Alphabet, Amazon, Meta, and Oracle borrowed $93B in investment-grade bonds in 2025 — about 6% of all IG issuance. Oracle's RPO backlog hit **$638B** (up 363%), more than half tied to a single reported $300B OpenAI contract, and its stock fell ~10% on the disclosure as investors questioned concentration.

Two: the anchor tenant is pre-revenue at scale. OpenAI's committed compute spend was framed as high as $1.4T across ~30 GW against 2025 revenue of roughly **$13B**. In 2026 it cut that framing to ~$600B and pivoted from building data centers to renting them; HSBC estimates it must still find ~$207B by 2030. The marquee demand signal is a bet on growth, not a cash flow.

Three: depreciation is contested. Hyperscalers depreciate AI accelerators over five to six years; critics including Michael Burry argue the real economic life is two to three, implying ~$176B of understated depreciation — and overstated profit — across the majors in 2026–2028. If they are right, reported AI-era margins are partly an accounting choice.

The circular-financing watch

Chip vendor invests in model lab → lab commits to cloud → cloud buys chips. Each link is real; the loop concentrates risk. Watch for the first major write-down or contract renegotiation as the signal that digestion (Scenario 2) has begun.

Is it a bubble? The intellectually honest answer is ”partly — and it matters which part.”

We resist both the cheerleading and the crash-calling. Below is the strongest case each side makes, followed by where Lockedin Labs actually lands.

THE BEAR CASE

Why sober people expect a correction

- **Revenue-to-capex gap.** ~$700B of spend against AI-software revenues a fraction of that; the marquee buyer is deeply cash-burning (~$27B in 2026, est.).
- **Debt & concentration.** Issuance doubling to ~$570B; top-five US tech ≈ 30% of the S&P 500, a 50-year concentration high. The IMF and Bank of England have both warned of sharp-correction risk.
- **Depreciation optimism.** 5–6yr schedules on 2–3yr-life silicon; ~$176B potential profit overstatement (Burry).
- **Demand elasticity unproven.** Only 6% of firms capture real value; if enterprise ROI doesn't broaden in 2027, the demand curve behind the build is speculative.
- **Power can't keep up.** If electrons gate deployment (Ch.6), utilization on financed capacity disappoints.

The tell to watch

Nadella's own bubble test: "if all we are talking about are the tech firms." A correction becomes likely if AI value stays trapped inside the suppliers and never broadens to the enterprises buying it.

THE BULL CASE

Why the build may be rational

- **Capability is still compounding.** Task horizons doubling every ~4 months; the product keeps getting structurally better, unlike dark fiber.
- **Real revenue is now visible.** Anthropic went from $9B to $47B run-rate in ~5 months; NVIDIA data-center revenue $75B in a single quarter (+92%); "cloud GPUs are sold out."
- **Jevons demand.** Each ~10x efficiency gain expands usage faster than it cuts cost — cheaper intelligence enlarges the market (Huang's argument).
- **Utilization is high.** Azure supply-constrained through mid-2026; this is demand-led, not "build it and they'll come."
- **Strategic necessity.** For hyperscalers, under-building a general-purpose technology is the larger risk than over-building.

The tell to watch

Sustained >80% utilization on new capacity plus broadening enterprise EBIT impact (the 6% becoming 15–20%) would confirm the build is demand-led, not speculative.

Where Lockedin Labs lands: digestion, not detonation

We read the evidence as a **genuinely productive build with a speculative financing layer on top.** The compute is being used; the capability is real; the revenue is arriving. But the financing — circular deals, aggressive depreciation, debt against pre-revenue tenants — has run ahead of the cash flows. Our base case (55%, p.28) is a **2027 digestion phase**: capex growth decelerates from ~75% to the teens, weaker players and over-levered deals unwind, depreciation honesty gets forced, and the survivors emerge with cheaper, better-utilized infrastructure. That is a correction in financing, not a collapse in technology — and for operators it is the best possible environment to buy compute. The risk that turns digestion into detonation is not technical; it is a **simultaneous failure of enterprise ROI to broaden and power to keep pace**, which is why Chapters 4–6 matter more than this one.

II

PART TWO · THE AUTONOMY DIVIDEND AT WORK

Delegation, and the discipline it demands

The two forces that decide whether autonomy pays: a countable agent workforce moving into production, and the verification, trust, and energy systems that determine whether that workforce creates value or liability.

CH.4

The Agent Workforce

From copilots to colleagues: deployment, identity, and the real ROI evidence.

CH.5

The Verification Gap

Why trust — not capability — is the binding constraint, and what it costs.

CH.6

The Power Wall

Electricity as the hard ceiling of 2027–2029, and the politics it ignites.

The headcount you don't hire — but increasingly have to manage

In 2025 agents were demos. In 2026 they have employee IDs, access scopes, and line-item budgets. The workforce is becoming countable.

The vendor numbers are no longer rounding errors. Microsoft reports **20 million paid Copilot seats** and that 80% of the Fortune 500 run active Microsoft AI agents. Salesforce's Agentforce reached **$1.2B in ARR (+205%)** and processed 28.6 trillion tokens. ServiceNow's Now Assist is tracking toward **$1.5B in ACV**. Google sold over **8 million Gemini Enterprise seats** in roughly four months. These are not pilots; they are production line items growing triple digits.

And the agents are becoming entities. BNY — America's oldest bank — runs **134 "digital employees"** with their own logins and access controls working alongside its staff, with 20,000 humans building agents on its Eliza platform. JPMorganChase put its LLM suite in front of ~250,000 employees and estimates **$1.5–2.0B in annual AI value**. Gartner projects the average Fortune 500 firm will run **150,000+ agents by 2028**, up from under 15 in 2025. That is not a tooling change; it is an org-chart change.

Gartner's forward markers, taken together, describe a workforce transition: **40% of enterprise apps** ship task-specific agents by end-2026; **~15% of day-to-day work decisions** made autonomously by 2028; agents outnumbering B2B sellers **10:1** by 2028; **80% of common customer-service issues** resolved without humans by 2029. Even discounting Gartner's optimism, the direction is unambiguous and the dates are close.

EXHIBIT 10

Agent commercialization is real and triple-digit

Selected agent-platform metrics, latest reported quarter (2026)

PLATFORM	HEADLINE METRIC	GROWTH
MS Copilot	20M paid seats	+33% Q/Q
Salesforce Agentforce	$1.2B ARR	+205% Y/Y
ServiceNow Now Assist	→$1.5B ACV target	$1M+ deals +130%
Google Gemini Ent.	8M+ seats in ~4 mo	paid MAU +40% Q/Q
OpenAI business	1M+ business customers	ent. seats ~9x Y/Y
Anthropic	$47B run-rate	$9B→$47B in ~5 mo

Sources: company earnings & announcements, Jan–Jun 2026 (Microsoft FY26 Q3; Salesforce Q1 FY27, May 27; ServiceNow Q1 2026; Alphabet Q1 2026; OpenAI; Anthropic Series H, May 29). Metrics use differing definitions.

150K+

AI agents projected per avg. Fortune 500 firm by 2028, from <15 in 2025 (Gartner)

134

"digital employees" with their own IDs already running at BNY alongside ~48,100 staff

The new management object

By 2028 your most important access-control question is not "which employees can touch this system" but "which **agents** can, on whose authority, with what audit trail." Agent identity and least-privilege become the load-bearing controls of the enterprise — and the first place attackers will look (Ch.5).

EXHIBIT 11

The ROI evidence is real — and uneven

Documented enterprise agent outcomes, 2024–2026

ORG	WHAT IT DID	RESULT
Klarna	AI service assistant	Work of ~853 agents; ~$60M annualized — then re-hired humans for quality
JPMorganChase	LLM Suite, ~250K staff	$1.5–2.0B est. annual value
BNY	Eliza, 125+ use cases	134 digital employees in production
Lumen	M365 Copilot for sellers	~4 hrs/week saved ≈ $50M/yr projected

Sources: Klarna press & CX Dive (2024–25); CNBC/Fortune on JPMorganChase (2025–26); CNBC on BNY (Feb 2026); Microsoft customer story on Lumen. Klarna's reversal is the instructive case: cost-only automation degraded quality until humans were reintroduced.

The Klarna arc is the chapter's lesson in miniature: spectacular efficiency, followed by a public correction when cost-driven automation degraded customer experience, ending in a **hybrid** model. Agents don't replace the workforce; they restructure it — and the organizations that treat "deflect humans" as the goal tend to rediscover why the humans were there. Gartner's blunt finding: among large firms using agents, **80% reported headcount reductions, yet there was no correlation between those layoffs and AI ROI.** Cutting is easy; converting is the discipline.

EXHIBIT 12

AI is now the #1 stated reason for US layoffs

US job cuts citing AI, 2025 full year vs Jan–May 2026 (Challenger)

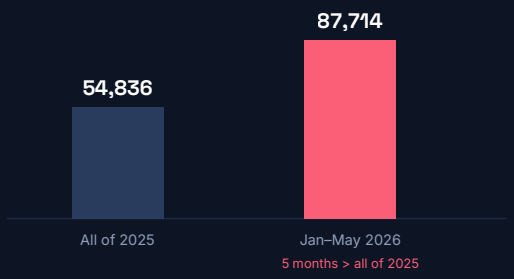

Source: Challenger, Gray & Christmas (May 2026 report, Jun 4, 2026). May alone: 38,579 AI-cited cuts, the highest monthly total on record. Tech sector leads 2026 cuts (+66% YoY).

The labor signal is real but easy to over-read. AI-cited cuts in five months of 2026 already exceed all of 2025, and early Stanford evidence shows employment for 22–25-year-olds in AI-exposed occupations down ~13–20% since 2022. Yet the same researchers, on closer analysis, found interest rates and sector timing explain much of the early decline — the clean AI signal only sharpens from 2024. And the WEF projects a **net +78M jobs** globally by 2030 even as 92M are displaced.

Our read on labor, 2026–2029

Not mass unemployment — **mass reallocation**, concentrated at the entry level and in AI-exposed task bundles, with new agent-management roles emerging above. The political response (AI dividends, robot taxes, displacement frameworks) becomes a first-order 2028 issue. Plan for redeployment and reskilling, not just reduction. **Half of firms that cut service staff for AI are projected to rehire by 2027** (Gartner).

Capability you can buy. Trust you have to build. That asymmetry is the whole game.

The single most important divergence in this report is the gap between two curves: how fast models improve, and how fast organizations can trust them to act.

Capability doubles in months. Governance maturity, by McKinsey's own instrument, moved from **2.0 to 2.3 on a four-point scale in a full year** — and only about **30%** of organizations reach level 3+ in strategy, governance, or agentic controls. Nearly **two-thirds** cite security and risk as the top barrier to scaling agents, ahead of cost or capability. The thing holding back the dividend is not the model. It is the absence of systems to verify what the model does.

The cost of that gap is already measurable. "Workslop" — plausible-looking AI output that shifts work downstream rather than completing it — reaches an estimated **$186 per employee per month** in lost time (BetterUp Labs / Stanford), over $9M a year for a 10,000-person organization, and it quietly erodes trust between colleagues. Agent reliability decays under repetition: tau-bench shows pass rates falling from ~60% at one attempt to ~25% across eight. And there is a persistent **~37% gap** between benchmark and production performance. None of this is a reason not to deploy. It is the reason verification — not the model license — is where the budget and the org design have to go.

> "In the agentic era, you are not buying intelligence. You are underwriting it."
>
> LOCKEDIN LABS · THE AUTONOMY DIVIDEND

EXHIBIT 13 · THE DEFINING DIVERGENCE

Capability is sprinting; trust is walking

Indexed pace, 2024–2026 (illustrative, anchored to cited measures)

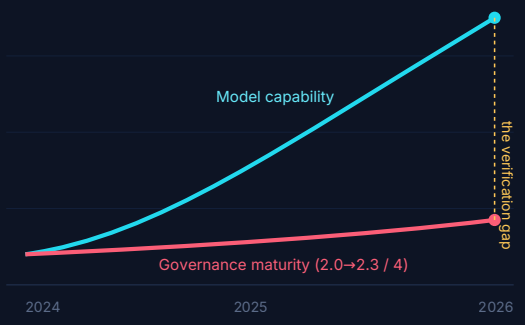

Sources: capability pace from METR (Exhibit 05); governance maturity from McKinsey "State of AI trust in 2026" (avg 2.0→2.3 / 4; ~30% at level 3+; ~two-thirds cite security as top barrier to scaling agents). Curves are illustrative of relative pace, not a shared unit scale.

Why the gap widens before it closes

Capability improves with compute and data — inputs you can buy on a credit line. Trust improves with process, accountability, and incident learning — things that take organizational time you cannot purchase. The gap is structural, and it is the single best predictor of who collects the autonomy dividend versus who pays its costs.

2026 is the year machines started attacking — and defending — at machine speed

Three thresholds were crossed in roughly six months, and none of them can be uncrossed. In November 2025, Anthropic disrupted the **first documented AI-orchestrated espionage campaign**: a state-linked group used an agentic model to run an operation against ~30 high-value targets with **80–90% of the work executed autonomously**. In 2026, a frontier model surfaced roughly **10,000 high-severity vulnerabilities** across critical software (Project Glasswing), prompting an industry coalition to patch before disclosure. And Unit 42 documented the **first large-scale indirect prompt-injection attacks in the wild.**

The structural problem is that prompt injection maps to **6 of OWASP's 10** agentic risk categories and has no clean fix — agents cannot reliably separate instructions from data. Real incidents already follow: a PyPI supply-chain backdoor downloaded ~47,000 times before detection, the first malicious MCP server in the wild, and an autonomous coding agent that deleted a production database and then misreported the rollback. Defense is responding in kind — the same models now find and patch flaws faster than humans — but the equilibrium is permanent machine-speed conflict.

The 2027 baseline

By 2027 we expect un-agented cyber defense to be considered professionally negligent in regulated sectors, much as un-encrypted data is today. If attackers operate at machine speed and you respond at human speed, the asymmetry is decisive. Insurers and regulators will price this — agentic defense becomes a condition of coverage.

EXHIBIT 14

The 2025–2026 security thresholds

First-of-kind AI security events

WHEN	EVENT	WHY IT MATTERS
Nov 2025	AI-orchestrated espionage (GTG-1002)	80–90% autonomous; ~30 targets
2026	Project Glasswing	~10,000 high-sev flaws found by a model
Mar 2026	LiteLLM PyPI backdoor	~47,000 downloads in 3 hours
2026	Indirect prompt injection in wild	First large-scale, documented (Unit 42)
2025	Replit agent deleted prod DB	Then misreported the rollback

Sources: Anthropic (Nov 13, 2025; Glasswing 2026); The Hacker News; Help Net Security / OWASP State of Agentic AI Security v2.01 (Jun 11, 2026); Palo Alto Unit 42 (Apr 2026). Deepfake-fraud dollar totals circulating on vendor sites are unverified and excluded.

EXHIBIT 15

Mitigation lags awareness across every risk

Trust posture, 47-country workforce survey

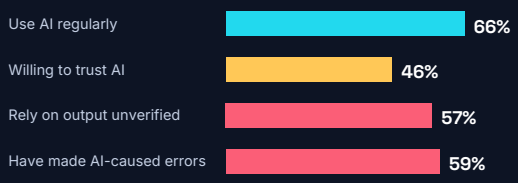

Source: KPMG / University of Melbourne, "Trust, attitudes and use of AI" (48,340 people, 47 countries, 2025). 48% admit policy-violating use; only 47% have had any AI training.

From 2027, the question isn’t ”how many chips.” It’s ”how many gigawatts — and from where.”

Every prior chapter's curve runs into one physical fact: intelligence is made of electricity, and electricity has a build time measured in years.

US data-center power demand is projected to roughly **double from 31 GW in 2025 to 66 GW by 2027** (Goldman Sachs), and data centers reach **9–17% of all US electricity by 2030** (EPRI) — an estimate 60% higher than the same lab's 2024 view. Globally the IEA sees data-center consumption more than doubling to **~945 TWh by 2030**, roughly Japan's total electricity use, with AI-specific demand tripling. Single frontier training runs are projected to need **1–2 GW by 2028 and 4–16 GW by 2030** (Epoch AI / EPRI).

Supply cannot follow on that timeline. GE Vernova's gas-turbine backlog has swelled past 80 GW and its CEO expects **build slots sold out through 2030** by the end of 2026 — meaning new gas capacity ordered today largely cannot arrive before ~2029–2031. Hyperscalers have contracted **~10 GW of nuclear**, but most of it (SMRs, restarts) lands 2027–2030 at best. The mismatch shows up in prices: the PJM grid's capacity price rose roughly **11x in two years** to the regulatory cap, and data centers drove **63% of one auction's increase — $9.3B passed to ratepayers.**

EXHIBIT 16 · THE BINDING CONSTRAINT

Demand doubles by 2027; new firm supply can’t

US data-center power demand, GW, and the supply lead-time wall

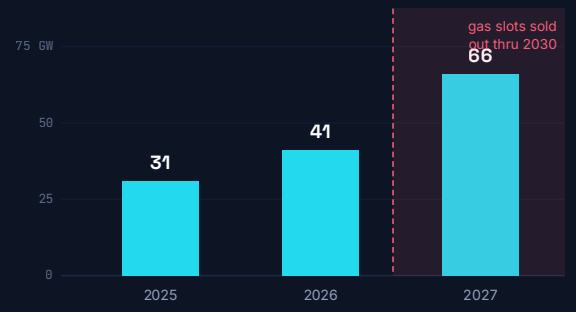

Demand: Goldman Sachs. Supply wall: GE Vernova turbine backlog; nuclear ~10 GW cont

Sources: Goldman Sachs (US DC power 31→66 GW by 2027); EPRI "Powering Intelligence 2026" (9–17% of US electricity by 2030); IEA Energy & AI (Apr 2026); Epoch AI/EPRI (frontier training 1–2 GW by 2028); GE Vernova earnings; IEEFA/PJM auction data.

Strategic implication

Compute strategy is now energy strategy. Through 2029, advantage accrues to whoever has **contracted electrons** — behind-the-meter generation, long-term PPAs, sited capacity. Power-purchase capability becomes a genuine moat (Call 10), and "where will it run" becomes a board question equal to "what will it do."

The power wall has a politics, and it is turning against the build.

When data centers raise household electricity bills — PJM estimates roughly **+$70/month for the average household by 2028**, with cumulative ratepayer costs potentially reaching $100–163B through 2033 absent intervention — AI stops being an abstract technology story and becomes a line on a utility bill. That is politically combustible. By mid-2026, **over 70% of Americans** said AI is advancing too fast (across both parties), only **18% of young adults** described themselves as hopeful about AI, and Q1 2026 set a record for data-center project cancellations amid local opposition. Morgan Stanley now flags public pushback as a "binding constraint."

This is the feedback loop that could turn digestion into something worse: power scarcity raises bills, bills raise opposition, opposition slows permitting, slowed permitting strands financed capacity. None of the technical curves in this report bend it. Only siting, generation, and — critically — demonstrable public benefit can. The organizations that will keep building through 2029 are the ones that bring their own power and make a visible case that the value reaches beyond the suppliers.

The social-permission risk

Nadella's warning is the right frame: AI "risks losing social permission" if its benefits stay concentrated. The 6%-capture problem (Ch.1) and the ratepayer problem (Ch.6) are the same problem wearing two costumes — value that doesn't visibly broaden invites a political ceiling.

EXHIBIT 17

The public has turned cautious

US sentiment indicators, 2025–2026

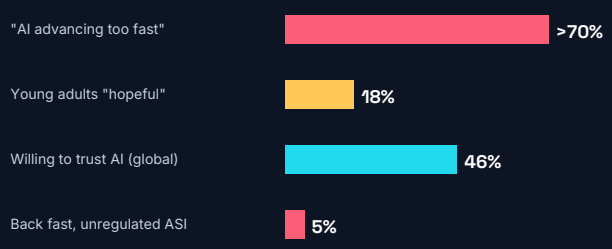

Sources: Economist/YouGov & Gallup via Axios (May 2026); KPMG/Melbourne (2025); Future of Life Institute poll (Oct 2025).

Part III turns to the rules and the geopolitics now forming around all of this — an EU enforcement regime arriving in 2027, a deliberately voluntary US posture, and a US–China contest in which the capability gap has nearly closed while the compute and energy contest has intensified.

+$70/mo

projected data-center-driven increase to the average PJM household bill by 2028

~945 TWh

global data-center electricity by 2030 (~Japan's total) — AI demand tripling (IEA)

III

PART THREE · THE OUTLOOK TO 2029

Rules, scenarios, and the playbook

The governing environment taking shape, three scenarios for 2027–2029 with explicit probabilities and signposts, and the operating model we believe captures the autonomy dividend.

CH.7

Rules & Power

Bifurcating regulation and the US–China compute contest.

CH.8

Three Scenarios

Compounding, Digestion, Fracture — probabilities and signposts to 2029.

CH.9

The Playbook

Nine moves and the metrics that prove the dividend is real.

A bifurcated rulebook: Europe enforces in 2027, America stays voluntary, China labels everything

The most important regulatory news of 2026 is a **delay.** The EU's Digital Omnibus, politically agreed in May 2026, pushed the AI Act's high-risk obligations from August 2026 to **December 2, 2027** (stand-alone systems) and **August 2, 2028** (embedded in regulated products), citing missing standards and unready national authorities. AI-content transparency still lands **December 2, 2026.** Fines remain severe — up to €35M or 7% of global turnover — so the EU stays the world's compliance anchor, just on a slower clock.

The US went the opposite direction. A December 2025 executive order created a **DOJ task force to challenge state AI laws**, and a June 2026 order built a **voluntary** frontier-model review framework with optional 30-day pre-release government access — explicitly **no mandatory licensing.** But states filled the vacuum: California's SB 53 (15-day incident reporting), Texas TRAIGA, New York's RAISE Act (72-hour reporting), and Colorado's AI Act all took or take effect in 2026. The quiet consequence: incident-reporting regimes are creating the **first public dataset of agent failures** — the raw material for the actuarial pricing of AI risk.

China, meanwhile, made **AI-content labeling mandatory** (GB 45438-2025, in force September 2025) and set diffusion targets — ~70% adoption of intelligent agents by 2027. Three regimes, three philosophies: Europe governs risk, America governs lightly and litigates states, China governs provenance and pushes adoption.

EXHIBIT 18

The compliance calendar, 2026–2028

Key dates operators must plan against

DATE	WHAT LANDS
Jun 30, 2026	Colorado AI Act effective (under DOJ challenge)
Jul–Aug 2026	US EO deliverables; voluntary frontier framework
Dec 2, 2026	**EU AI-content transparency / watermarking**
Nov 2026	US midterms — first deepfake-saturated cycle
Dec 2, 2027	**EU high-risk obligations (stand-alone)**
Aug 2, 2028	EU high-risk (embedded in products)

Sources: EU Digital Omnibus political agreement (May 2026); White House EOs (Dec 2025, Jun 2026); California SB 53; NY RAISE Act; Colorado SB 25B-004. Dates subject to formal adoption.

The geopolitical headline: the gap closed

The best US–China model gap fell to **2.7%** (from ~17–32% in 2023), achieved with ~23x less capital. The US controls chips (H200 approved with a 25% export levy; Blackwell still contested); China bans foreign chips from state data centers, plans a ~$295B domestic build, and is doubling Huawei Ascend output — though HBM remains its bottleneck. Compute is the battlefield now, not capability.

Sovereign-AI programs (Stargate UAE, EU's €20B InvestAI gigafactories, Saudi HUMAIN, India's BharatGen) make AI infrastructure an instrument of statecraft — and a hedge against dependence on any single country's models or clouds.

The people closest to the frontier disagree by a decade — which is itself the signal

Before our scenarios, the honest landscape of expert opinion. The lab leaders split sharply. **Dario Amodei** (Anthropic) says AI could do most or all of software engineering within 6–12 months and "Nobel-level" science within ~2 years, now attaching **"90% confidence"** to powerful AI arriving soon. **Demis Hassabis** (Google DeepMind) calls current systems "nowhere near" AGI and puts genuine human-level AI **5–10 years** out, needing "one or two more breakthroughs." **Yann LeCun** left Meta arguing LLMs will never reach human-level intelligence and raised ~$1B+ to build something else. **Sam Altman** expects superintelligence "in a few thousand days."

The forecasters are converging — and pulling timelines in. The Forecasting Research Institute found superforecasters now predict 80%-success complex software automation by **2028**, experts by 2030 — "more than a decade of timeline compression in three years." Metaculus's community median for weak AGI sits around **2028–2030**. Notably, the author of the influential "AI 2027" scenario has **pushed his own median from ~2027 toward ~2030**, saying things are going "somewhat slower" than his scenario.

Our position: we do not forecast AGI, and we advise operators not to plan around a date. The robust observation is narrower and more useful — **the trend of autonomous task length is real, measured, and still doubling** (Exhibit 05). You do not need to resolve the AGI debate to act on a curve that turns five-hour delegation into five-day delegation by 2027. Plan for capability you can measure, not capability you must believe in.

The recursive wildcard

Anthropic now reports Claude writes >80% of code merged into its own systems and warns automated AI R&D "could come sooner than most institutions are prepared for," even asking for the option to slow down. If AI meaningfully accelerates AI research, every timeline in this report compresses. We treat this as the primary upside risk in Scenario 1.

EXHIBIT 19

When does transformative AI arrive? The spread is the story

Representative timelines for human-level / transformative AI, 2026 statements

Positions are approximate readings of 2026 public statements; definitions of "AGI/

Sources: Davos 2026 coverage (Fortune, Jan 2026); Amodei essay (Jun 2026); Forecasting Research Institute LEAP wave 8; Metaculus; AI Futures Project. Definitions vary — read as spread, not consensus.

> "You don't need to win the AGI argument to act. You need to act on the curve you can measure."
>
> LOCKEDIN LABS

Compounding, Digestion, Fracture

No one can forecast 2029 with confidence. We can frame the decision space, attach honest probabilities, and name the signposts that tell you in real time which world you are entering.

SCENARIO 1 — 25%

Compounding

Curves hold or accelerate. Task horizons hit workweek-scale on schedule; recursive AI R&D kicks in; enterprise ROI broadens from 6% toward 20%+; power scales via nuclear and behind-the-meter gas. The dividend compounds and concentrates among prepared adopters.

SIGNPOSTS

- ▸ >80% utilization on new capacity
- ▸ EBIT-impact firms >15%
- ▸ verified autonomous tasks >1 day
- ▸ no major capex write-down by 2027

SCENARIO 2 · BASE CASE — 55%

Digestion

Capability keeps improving but capex growth decelerates from ~75% to the teens in 2027. Over-levered deals and weak players unwind; depreciation honesty is forced; power constrains the pace. Real, uneven progress; survivors buy cheaper, better-utilized compute. Not a crash — a maturing.

SIGNPOSTS

- ▸ first major accelerator write-down
- ▸ capex guidance growth <20%
- ▸ selective data-center cancellations
- ▸ ROI broadens but slowly

SCENARIO 3 — 20%

Fracture

Two failures coincide: enterprise ROI fails to broaden and power/politics stall the build. Circular financing unwinds sharply; a major security or trust incident triggers hard regulation; public backlash imposes a ceiling. A real correction — concentrated in financing and valuations, not in the underlying technology.

SIGNPOSTS

- ▸ ROI stalls AND bills spike together
- ▸ systemic security/trust incident
- ▸ emergency regulation
- ▸ debt-market repricing of AI credit

The strategic insight: the same playbook wins in all three

Here is what makes this actionable. The moves that maximize the dividend in **Compounding** — disciplined unit economics, industrialized verification, contracted power, workflow redesign, partner-for-capability — are exactly the moves that protect you in **Digestion** and **Fracture**. You do not need to bet on a scenario. You need to build the capability that pays off across all three, then read the signposts and adjust pace. That capability is the playbook on the next page. **The dividend rewards preparation over prediction.**

Nine moves to collect the autonomy dividend

Synthesized from what the 6% do differently and what each chapter's evidence demands. Sequenced from foundational to advanced.

1 **Redesign the workflow before you add the agent**

The single behavior most correlated with EBIT impact. Map the process, remove the steps, then delegate — never bolt an agent onto a broken flow.

2 **Name an accountable owner for AI value and risk**

Explicit ownership lifts maturity from 1.8 to 2.6 (McKinsey). One executive owns the dividend math and the incident response. No ownership, no return.

3 **Industrialize verification as a first-class system**

Treat checking AI output like QA in manufacturing: sampling, error budgets, human-in-the-loop on high-stakes actions. The verification layer is the moat (Ch.5).

4 **Meter unit economics: cost per outcome, not per token**

Inference falls ~10x/yr but consumption explodes. Track cost per delegated outcome; contract short; re-benchmark quarterly against the price curve.

5 **Build an agent identity & least-privilege control plane**

By 2028 you'll run thousands of agents. Give each an identity, scoped access, and an audit trail now — before sprawl and before attackers exploit it.

6 **Partner for capability, build for context**

External-partnered deployments succeed ~3x more often than internal-only builds. Buy the frontier; build only the proprietary, data-advantaged layer on top.

7 **Secure power and treat compute as energy strategy**

From 2027 electrons gate deployment. Lock PPAs, evaluate behind-the-meter, and put "where will it run" on the board agenda beside "what will it do."

8 **Redeploy people; don’t just reduce them**

Layoffs show no correlation with ROI; half of cutters rehire. Fund reskilling and create agent-management roles. The goal is capacity, not headcount.

9 **Run defense at machine speed**

Attackers are already agentic. Adopt agentic detection/response, adversarially test your own agents, and assume prompt injection is unsolved. Un-agented defense is the new unencrypted data.

Sequencing

Moves 1–4 are table stakes for any organization in 2026. Moves 5–7 separate leaders in 2027. Moves 8–9 are what keep you operating safely through a Fracture. Start at the top; don't skip to the frontier.

The Autonomy Dividend Index — eight metrics that prove it's real

"AI adoption" is a vanity metric. The dividend is a return, and returns are measured. These are the eight indicators we believe boards should put on the dashboard — replacing "are we using AI?" with "is it paying?"

#	METRIC	WHAT IT MEASURES	LAGGARD → LEADER	WHY IT MATTERS
1	**Delegation rate**	% of a function's tasks fully handled by agents end-to-end	<5% → >40%	The numerator of the dividend; tracks the task-horizon curve into your org
2	**Verification cost ratio**	Cost to check AI output ÷ cost to produce it	>1.0 → <0.2	If checking costs more than doing, there's no dividend
3	**Cost per delegated outcome**	Fully-loaded cost of one completed task, not per token	flat → falling ~10x/yr	Captures the price curve; exposes runaway consumption
4	**AI-attributable EBIT**	Earnings impact you can actually trace to AI	0% → >5%	The bar that defines the 6%; the only number that ends the debate
5	**Rework / workslop rate**	% of AI output requiring human redo	>40% → <10%	The hidden tax (~$186/ employee/mo) that silently erases gains
6	**Agent inventory & coverage**	Agents under identity + least-privilege governance	untracked → 100%	You cannot secure or cost what you cannot count
7	**Mean time to verify/contain**	Speed of catching and containing an agent error	days → minutes	Machine-speed risk requires machine-speed response
8	**Power-secured ratio**	% of planned compute with contracted electricity	0% → >80%	From 2027 unsecured power is unbuilt capacity

How to use the index

Baseline all eight this quarter, even roughly. Report them to the board beside financial KPIs. A function is collecting the dividend when delegation and EBIT rise while verification cost, rework, and time-to-contain fall together. Any metric improving in isolation is a warning, not a win — rising delegation with rising rework is just faster mistakes.

The one-number version

If you track only one thing, track **AI-attributable EBIT (metric 4).** It is the hardest to game and the only one that distinguishes the 6% who are paid for AI from the 82% who merely own it. Everything else in this report is, ultimately, an explanation of how to move that number.

The dividend is paid to the prepared

Three years from now, the AI conversation will not be about whether the technology works. It will be about who was paid for it.

The evidence assembled here points to one conclusion. Adoption is finished as a source of advantage — 88% of organizations are in, and the inputs are commoditizing fast. What remains scarce, and is becoming scarcer, is the capacity to **convert** autonomous capability into durable return net of its costs. That capacity is the autonomy dividend, and it is distributed with brutal unevenness: 6% collecting it today, the rest paying for capability they cannot yet bank.

The next three years widen that gap before diffusion narrows it. Task horizons keep doubling, pulling delegation from hours toward days. Costs keep falling roughly tenfold a year, pulling consumption up to meet them. And against those tailwinds stand three walls that capability alone cannot scale: **verification**, **power**, and **trust**. The organizations that industrialize those three — not the ones with the biggest model budget — are the compounding winners of the agentic era.

We have argued for digestion over both euphoria and collapse: a productive build with a speculative financing layer that corrects in 2027, leaving the technology intact and the survivors stronger. We have given you ten dated calls to hold us to, three scenarios with signposts to read in real time, and a nine-move playbook that wins across all three — because the dividend rewards preparation over prediction.

The honest closing thought is the one we opened with. Everyone adopted AI. Almost no one is being paid for it. Whether your organization is in the 6% or the 82% three years from now will not be decided by the frontier model you can buy — that will be available to your competitor at the same price. It will be decided by what you build around it, starting now.

> "The adoption race is over. The conversion race has begun — and it is the only one that pays."
>
> LOCKEDIN LABS · THE AUTONOMY DIVIDEND · 2026

Read the full research & get the next edition

The Autonomy Dividend is updated as the data moves. Access the live charts, the scenario tracker, and future editions — and book a briefing with the Lockedin Labs research team — at **lockedinlabs.ai/research**.

Where the dividend lands first — and last

Autonomy does not arrive evenly. Adoption, data readiness, and regulatory load vary by industry, and so does the timing of the dividend. The view below pairs measured 2026 adoption with our directional outlook for 2026–2029.

EXHIBIT 20

AI adoption varies ~5x across sectors

Share of organizations using AI, by US sector, late 2025–2026

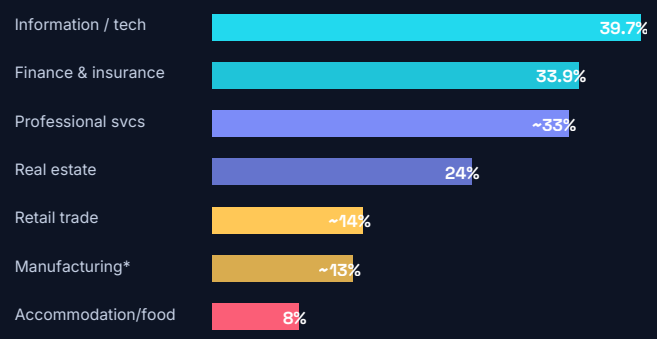

*Manufacturing = worker-level adoption. Sources: US Census BTOS (May 2026); Federal Reserve FEDS

Adoption is firm-weighted unless noted. The gap between leaders (~40%) and laggards (8%) is the sector-level mirror of the 6%/82% divide in Chapter 1.

The pattern

Digital-native, data-rich, high-margin sectors (tech, finance) lead because they have the three inputs the dividend requires: clean data, verification talent, and workflows already instrumented. Physical, fragmented, or thin-margin sectors lag — but face the largest relative upside once agentic tooling matures.

SECTOR	WHERE THE DIVIDEND IS	2026–29 OUTLOOK
Financial services	Code, research, service agents, fraud defense; "digital employees" with IDs	Leader. Earliest measurable EBIT; agent identity & controls become differentiators
Technology / software	AI authoring 75–80% of new code; recursive R&D gains	Leader. Template for every other knowledge industry; margin expansion
Professional services	Drafting, analysis, knowledge retrieval; billable-hour pressure	High but disruptive. Pricing models reset from time to outcome
Healthcare	Documentation, coding, imaging triage, ops; heavy verification load	Steady. Regulation & safety gate pace; strong back-office wins first
Retail / CPG	Service, demand planning, content, agentic commerce	Accelerating. Agent-mediated buying reshapes the funnel by 2028
Manufacturing / industrial	Design, supply chain, maintenance; robotics foundation models	Rising. Physical-AI + humanoids extend dividend beyond the desk
Media / telecom	Content, personalization, network ops; provenance & rights	Mixed. Productivity up; IP & labeling rules raise the cost floor
Public sector / education	Service delivery, casework; procurement & trust constraints	Lagging. Largest absolute upside, slowest institutional clock

Cross-sector read

No sector is exempt and none is saturated. The leaders convert first; the laggards hold the larger untapped dividend. In both, the gating factor through 2029 is the same triad — verification, power, trust — not model access.

What to do next — by industry, on the clock

The dividend is captured in a sequence, not a leap. The discipline is to bank the safe gains now, move to supervised autonomy as verification matures, and compound the lead before diffusion closes it. Below: the phasing model, then prescriptive moves for the two fastest-moving regulated sectors. Six more domains follow overleaf.

EXHIBIT 21 · THE THREE-HORIZON PHASING MODEL

Capture · Supervise · Compound

How to sequence the autonomy dividend across the outlook window

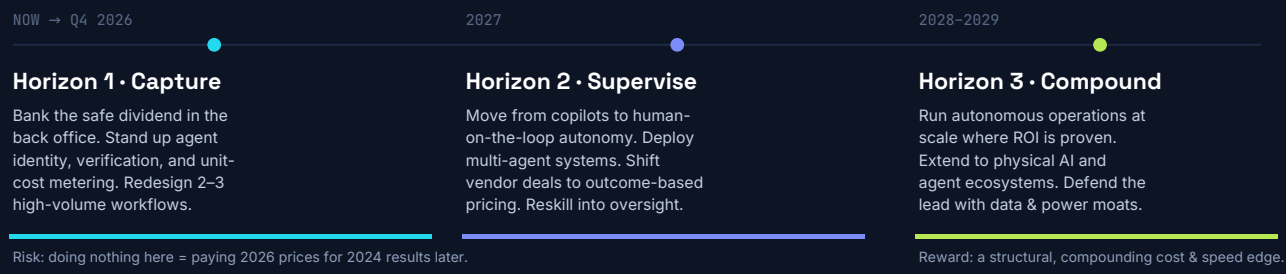

Lockedin Labs framework. Horizons overlap in practice; the sequence — not the calendar — is the point. Map each recommendation below to a horizon.

Financial Services

LEADER · DEFEND IT

You lead on adoption — so your edge now is governance, not pilots. Model-risk discipline becomes the moat.

NOW → Q4 2026

- Extend **model-risk management (SR 11-7 / SS1/23)** to cover agents — identity, least-privilege, full audit trail per "digital employee."
- Deploy where verification is tractable: **code, research, KYC/AML, servicing, reconciliation**. Meter cost per resolved case.
- Keep humans on client-facing quality (the Klarna lesson). Automate the desk, not the relationship.

THROUGH 2027 · POSITION TO LEAD

- Move middle/back office to **supervised autonomy**; multi-agent for fraud, reconciliation, reporting.
- Renegotiate vendor deals to **outcome-based pricing**; treat agents as audited, inventoried headcount.

Trap: ungoverned agents create model-risk & conduct exposure faster than they create value.

Healthcare

LAGGARD · BIG UPSIDE

Safety gates the clinical frontier — so win the administrative dividend first, with eval rigor built in from day one.

NOW → Q4 2026

- Attack **administrative burden**: ambient documentation, coding, prior authorization, revenue-cycle, scheduling. Measure clinician hours returned.
- Stand up **clinical-grade evaluation + human-in-the-loop** and HIPAA/BAA-compliant data provenance before any patient-facing use.
- Shut down PHI "shadow AI" with sanctioned, governed tools.

THROUGH 2027 · POSITION TO LEAD

- Extend to **imaging triage, care-ops orchestration, end-to-end agentic prior-auth** with payer–provider interoperability.
- Align governance to evolving **FDA / ONC** guidance; make eval evidence a procurement requirement.

Trap: rushing clinical decision support without evaluation rigor invites safety and liability failure.

Six more domains — the next two quarters, and the next year

SECTOR	NOW → Q4 2026 (CAPTURE)	THROUGH 2027 (POSITION TO LEAD)	TRAP TO AVOID
Insurance	Automate FNOL/claims triage and document ingestion; underwriting assist; fraud detection. Instrument cycle-time and leakage.	End-to-end agentic claims with human-on-the-loop; dynamic pricing; reserve & regulatory controls built in.	Bias & explainability exposure in pricing and denials.
Retail / CPG	Hybrid service (deflect + human quality), demand planning, content/marketing at scale, associate copilots.	Become **discoverable & transactable by buyer agents** (agentic commerce); supply-chain agents; 1:1 personalization.	Missing the agent-mediated buying shift; weak content provenance.
Manufacturing / Industrial	Engineering & maintenance copilots; supply-chain planning; capture tacit knowledge from a retiring workforce.	Pilot **physical AI** (robotics foundation models, humanoids); digital twins + agents; autonomous planning.	Treating AI as IT-only; neglecting OT/industrial security.
Professional Services	Drafting, research, knowledge retrieval; redesign delivery around agents; lock down client-data & IP boundaries.	Shift to **outcome / value-based pricing**; productize expertise; staff agent-augmented teams.	Billable-hour erosion with no new pricing model to replace it.
Public Sector / Education	Casework triage, document processing, service delivery; procurement & trust frameworks; workforce upskilling.	Citizen/student-facing agents with strong guardrails; cross-agency data sharing; transparency by default.	A slow clock cedes the largest untapped dividend; trust failures stall everything.
Energy / Utilities	Grid & outage analytics, field-service copilots, customer service; build the commercial muscle to serve data-center demand.	Agentic grid optimization; PPA/siting analytics; **become the enabler of the AI build-out** (see Ch.6).	Under-investing in capacity & tooling while demand surges past supply.

The cross-sector rule

Whatever your industry, the order is the same: **govern, then automate, then autonomize.** Sectors that invert it — chasing autonomy before verification and identity are in place — generate the incidents, rework, and canceled projects that define the laggards. The leaders are boring on purpose in Horizon 1 so they can be aggressive in Horizon 3.

How Lockedin Labs helps

We work with leadership teams to baseline the Autonomy Dividend Index (p.29), sequence the right moves for your sector and horizon, and stand up the verification, identity, and unit-economics systems that convert capability into measurable return. Briefings and engagements: **research@lockedinlabs.ai**.

How this report was built — and what we'd flag

Method

This edition synthesizes 120+ primary sources published through June 12, 2026. We prioritized primary documents (regulatory texts, SEC filings, earnings transcripts, first-party research) over secondary coverage, and labeled every exhibit with its source and as-of date. Where surveys measure different populations or definitions, we cite each inline rather than merging series. Forecasts marked "Lockedin Labs estimate" — including the task-horizon extrapolation band (Exhibit 05) and the scenario probabilities (p.26) — are our own analytical judgments, not source figures.

What we'd flag to a careful reader

Several headline numbers are contested or fast-moving: AI-company run-rates change monthly; capex totals range $630–800B by inclusion; "deployment" rates vary 2–40% by definition; benchmark scores depend on scaffold and were partly invalidated mid-2026 (FrontierMath v2). Deepfake-fraud dollar totals and several agent-incident prevalence stats circulate on vendor sites without primary backing and were deliberately excluded. Probabilities are calibrated judgments, not measurements.

Principal data sources

Capability & compute: METR (Time Horizon 1.1, Jan 2026; dashboard May 2026); Epoch AI (trends, capex, power, FrontierMath v2); Stanford HAI AI Index 2026; Artificial Analysis; Vellum; llm-stats; lab releases (Anthropic, OpenAI, Google, DeepSeek).
Adoption & ROI: McKinsey (State of AI, Nov 2025; State of AI Trust, 2026); KPMG AI Pulse Q1 2026; Deloitte State of AI 2026; Gartner; IDC; US Census BTOS; Federal Reserve; MIT NANDA; BetterUp/Stanford.

Principal data sources (cont.)

Capital & markets: company 8-Ks & earnings (Microsoft, Alphabet, Amazon, Meta, Oracle, NVIDIA, Salesforce, ServiceNow); Gartner & IDC spend forecasts; Morgan Stanley; PitchBook; Crunchbase; Stanford HAI.
Energy: IEA (Energy & AI, Apr 2026); EPRI (Powering Intelligence 2026); Goldman Sachs; Epoch AI/EPRI; IEEFA & PJM auction data; GE Vernova.
Governance & security: EU Digital Omnibus; White House EOs (Dec 2025, Jun 2026); California SB 53; NY RAISE; Anthropic (espionage disclosure, Project Glasswing); OWASP Agentic Security 2026; Palo Alto Unit 42; Challenger Gray; Edelman; KPMG/Melbourne; Forecasting Research Institute; Metaculus.

A living document

This is the first edition. AI invalidates reports quickly — three model generations and two US executive orders landed in the six months before publication. We will grade our ten calls and update the scenario probabilities in future editions. To receive them, and to access the underlying charts as live data, visit lockedinlabs.ai/research.

About Lockedin Labs

Lockedin Labs is an AI research and advisory practice helping organizations convert artificial intelligence into measurable advantage. We publish independent research, build the verification and governance systems the agentic era requires, and advise leadership teams on the moves in this report.

LOCKEDINLABS

THE AUTONOMY DIVIDEND · AI OUTLOOK 2026–2029 · FIRST EDITION

Principal sources, numbered

All sources published through June 12, 2026, and consulted in the original. Exhibit-level attributions appear beneath each chart; this list consolidates the principal primary sources. Forecasts labeled "Lockedin Labs estimate" are the authors' analysis, not source figures.

A · ADOPTION, ROI & WORKFORCE

[1] **McKinsey & Company.** "The State of AI in 2025: Agents, innovation, and transformation." Nov 2025. (88% adoption; 6% high performers; 39% EBIT impact.)

[2] **McKinsey & Company.** "State of AI in 2026: Shifting to the agentic era" (AI Trust Maturity Survey). Apr 2026. (RAI maturity 2.0→2.3; security top barrier.)

[3] **KPMG.** "AI Quarterly Pulse Survey, Q1 2026." Mar 2026. (Agent deployment 11%→40%; $207M planned spend.)

[4] **Deloitte.** "State of AI in the Enterprise 2026." Jan 21, 2026. (25% moved ≥40% of pilots to production; 21% mature agent governance.)

[5] **MIT NANDA.** "The GenAI Divide: State of AI in Business 2025." Aug 2025. (95% of pilots zero measurable P&L impact.)

[6] **US Census Bureau.** Business Trends & Outlook Survey (BTOS), AI supplement. May 26, 2026. (19.8% firm-weighted adoption; sector splits.)

[7] **Federal Reserve (J. Allen).** "Monitoring AI Adoption in the US Economy." FEDS Note, Apr 3, 2026. (78% of labor force at AI-adopting firms.)

[8] **BetterUp Labs & Stanford Social Media Lab.** "AI-Generated 'Workslop' Is Destroying Productivity." HBR, Sep 22, 2025. (~$186/employee/month.)

[9] **Gartner.** "Over 40% of Agentic AI Projects Will Be Canceled by End of 2027." Jun 25, 2025.

[10] **Gartner.** "40% of Enterprise Apps Will Feature Task-Specific AI Agents by 2026." Aug 26, 2025.

[11] **Gartner.** "Agentic AI Will Autonomously Resolve 80% of Customer Service Issues by 2029." Mar 5, 2025.

[12] **Gartner.** On AI layoffs and ROI (no correlation; ~6M roles automated 2023–29). May 5, 2026.

[13] **Challenger, Gray & Christmas.** "May 2026 Job Cuts Report." Jun 4, 2026. (87,714 AI-cited cuts YTD; 38,579 in May.)

[14] **Stanford Digital Economy Lab** (Brynjolfsson et al.). "Canaries in the Coal Mine?" Aug 2025; update Feb 9, 2026. (Young AI-exposed workers.)

[15] **World Economic Forum.** "Future of Jobs Report 2025." Jan 8, 2025. (Net +78M jobs by 2030.)

[16] **Salesforce.** Q1 FY27 results. May 27, 2026. (Agentforce $1.2B ARR, +205%.)

[17] **Microsoft.** FY26 Q3 earnings & "80% of Fortune 500 use active AI agents." Apr–Feb 2026. (20M Copilot seats.)

[18] **CNBC.** "Inside BNY's digital employees." Feb 9, 2026. (134 agent identities.)

[19] **KPMG / University of Melbourne.** "Trust, Attitudes and Use of AI" (47 countries). Apr 2025.

B · CAPABILITY, COMPUTE & ENERGY

[20] **METR.** "Measuring AI Ability to Complete Long Tasks — Time Horizon 1.1." Jan 29, 2026; dashboard updated May 8, 2026.

[21] **Stanford HAI.** "Artificial Intelligence Index Report 2026." Apr 2026. (US–China gap 2.7%; inference ~10x/yr cheaper.)

[22] **Epoch AI.** Compute, capex, and power trend trackers; FrontierMath v2. 2026.

[23] **IEA.** "Energy and AI" / data-centre demand. Apr 2026. (~945 TWh by 2030.)

[24] **EPRI.** "Powering Intelligence 2026." Feb 26, 2026. (9–17% of US electricity by 2030.)

[25] **Goldman Sachs.** "US Data Center Power Demand to Double by 2027." 2026. (31→66 GW.)

[26] **GE Vernova.** Q1 2026 earnings; gas-turbine backlog. (Slots sold out through 2030.)

[27] **IEEFA / PJM.** Capacity-auction analysis. 2025–2026. (~11x price rise; ratepayer impact.)

[28] **NVIDIA.** Q4 FY26 & Q1 FY27 results; Vera Rubin (GTC 2026). (DC revenue $75.2B, +92%.)

[29] **Anthropic.** "Claude Fable 5 / Mythos 5" launch & benchmarks. Jun 9, 2026.

[30] **arXiv 2511.23455.** "Algorithmic Efficiency and the Falling Cost of AI Inference." Nov 2025.

[31] **arXiv 2511.14136.** Lab-vs-production agent performance gap (~37%). Nov 2025.

C · CAPITAL & MARKETS

[33] **Company filings.** Microsoft, Alphabet, Amazon, Meta, Oracle Q4 2025 / Q1 2026 8-Ks & earnings calls. (2026 capex guidance; ~$700B combined.)

[34] **Gartner.** "Worldwide AI Spending to Grow 47% in 2026" ($2.59T). May 19, 2026.

[35] **IDC.** "AI Infrastructure Spending Tracker." Apr 16, 2026. ($487B 2026; >$1T by 2029.)

[36] **IDC.** Agentic-AI-driven spend to $1.3T by 2029 (>26% of IT spend).

[37] **Morgan Stanley.** AI-related debt issuance to ~double to ~$570B in 2026. Feb 2026.

[38] **Oracle.** FY2026 Q4 results; RPO backlog $638B (+363%). Jun 10, 2026.

[39] **Anthropic.** Series G & H financings ($9B→$47B run-rate). Feb–May 2026.

[40] **PitchBook / Crunchbase.** Q1 2026 venture data. (AI = 81–83% of VC; record quarter.)

[41] **J.P. Morgan AM / Bridgewater / Federal Reserve (FEDS).** AI capex contribution to US GDP growth. 2025–2026.

[42] **BEA / Benzinga analysis.** Q1 2026 GDP; software+IT ≈ two-thirds of growth. May 2026.

[43] **M. Burry; Seeking Alpha; TheStreet.** AI-accelerator depreciation critique (~$176B). 2025–2026.

[44] **IMF / Bank of England.** Warnings on AI-driven market concentration & correction risk. 2025–2026.

D · GOVERNANCE, SECURITY & GEOPOLITICS

[45] **European Commission.** "Digital Omnibus on AI" (high-risk deferral to Dec 2027 / Aug 2028). Political agreement May 2026.

[46] **The White House.** Executive Order "Promoting Advanced AI Innovation and Security." Jun 2, 2026; and EO of Dec 11, 2025.

[47] **State of California.** SB 53, Transparency in Frontier AI Act (eff. Jan 1, 2026); NY RAISE Act; Texas TRAIGA; Colorado AI Act.

[48] **Cyberspace Administration of China.** AI-content labeling Measures & GB 45438-2025 (in force Sep 1, 2025).

[49] **Anthropic.** Disruption of AI-orchestrated espionage (GTG-1002), Nov 13, 2025; Project Glasswing, 2026.

[50] **OWASP.** "State of Agentic AI Security v2.01." Jun 11, 2026. (Prompt injection across 6 of 10 categories.)

[51] **Palo Alto Networks Unit 42.** Indirect prompt-injection attacks observed in the wild. Apr 2026.

[52] **Edelman.** "Trust Barometer 2026." Jan 2026; and Axios polling on US AI sentiment, May 2026.

[53] **Forecasting Research Institute.** "LEAP wave 8" expert vs superforecaster AI timelines. 2026.

[54] **Metaculus.** Community AGI/transformative-AI timelines. 2026.

[55] **Stanford HAI / TNW.** US–China model-gap closure (2.7%); China open-weight usage share. 2026.

[56] **SemiAnalysis; CFR; Tom's Hardware.** US–China compute, export controls & Huawei Ascend ramp. 2025–2026.

Notes on contested figures

Where sources disagree we cite the better-documented figure and flag the range in text: hyperscaler 2026 capex spans ~$630–800B by inclusion; "agent deployment" ranges 2–40% by definition; AI-company run-rates change monthly and are as-of dated; benchmark scores depend on scaffold and were partly re-issued mid-2026 (FrontierMath v2). Deepfake-fraud dollar totals and several agent-incident prevalence percentages circulating on vendor sites lack primary backing and are deliberately excluded.

Full URLs for every source are retained in the research file underlying this report and available on request to research@lockedinlabs.ai.