

A STRATEGIC FORESIGHT REPORT

Superintelligence 20/30

Intelligence and the Future of the Enterprise

WORK · ORGANIZATIONS · ECONOMICS · SOCIETY

June Manley

October 2026 · junemanley.com



Intelligence is abundant. **Conversion is scarce.**

Contents

THE FRAME

Outlook 03

The 2030 Intelligence Frontier 05

How to read this report 08

The ten themes, ranked 09

TEN THEMES

01 The Organization 20/30 10

02 Security at machine speed 14

03 Authority is the new privacy 17

04 Everything is delegated 20

05 Act and adjust 23

06 The attention economy 26

07 Expertise for the people 29

08 The price of intelligence 32

09 The distribution of intelligence 35

10 Intelligence gets a body 40

ACCOUNTABILITY

Four futures for 2030 43

The External Intelligence Monitor 44

The 20/30 Scorecard 48

Conclusion 50

Research notes 52

OUTLOOK

By 2030, the world may still be debating whether artificial superintelligence has arrived. Some enterprises will already be operating as if it has.

That is the central prediction of this report. The defining story between now and 2030 may not be a single machine crossing an agreed threshold of intelligence. It may be organizations becoming dramatically more capable by combining human expertise, machine intelligence, institutional knowledge, and coordinated execution.

I call this *organizational superintelligence*: a proposed framework for collective enterprise capability, not an established technical classification.

The distinction lives in the modifier. An enterprise can reach it without owning a superintelligent model. And access to the most capable model guarantees nothing.

THE MACHINE

Machine superintelligence

The capability the labs are racing toward.

THE ENTERPRISE

Organizational superintelligence

What an enterprise can become.

Intelligence no longer waits to be asked

AI used to answer. Now it acts: it plans, uses tools, and coordinates work under delegated authority.

Who decides. Who validates. Who is accountable when the machine is wrong.

Ready faster than the institutions

In October 2026, Anthropic reported six hours of employee directed agent runtime per estimated human work hour. Junior technical hires fell from 27% to 9% of its technical hiring mix.

80% of McKinsey respondents report individual AI productivity gains.

Only 37% report positive EBIT impact.

An ecosystem, not a laboratory

The future of intelligence will not be determined by AI laboratories or enterprises alone. It will emerge from the interaction of frontier models, semiconductors, energy, computing infrastructure, national investment, and the institutions that govern deployment. Technology companies are already planning against a 2030 horizon. Nvidia estimates a $3 trillion to $4 trillion annual AI infrastructure build by the end of the decade, while Microsoft plans multigigawatt Azure Maia deployment beginning in 2027. The first is a vendor market estimate; the second is a company deployment plan. Neither is a verified 2030 outcome.

Intelligence may become more affordable at the point of use even as the resources needed to produce it become more capital intensive and geographically concentrated. By 2030, countries and enterprises may have access to increasingly comparable capabilities while differing sharply in their ability to convert them into economic and societal value.

Why 2030. It is far enough away for the operating model to change and close enough to be held accountable. This report ranks ten themes by priority, labels every call by confidence, and states what would prove each one wrong. In 2030, I will publish the verdicts.

The question is not when machines become smarter than people. It is whether organizations become more capable as a result and who benefits when they do.

Intelligence is abundant. **Conversion is scarce.**



BEFORE THE THEMES

The 2030 Intelligence Frontier

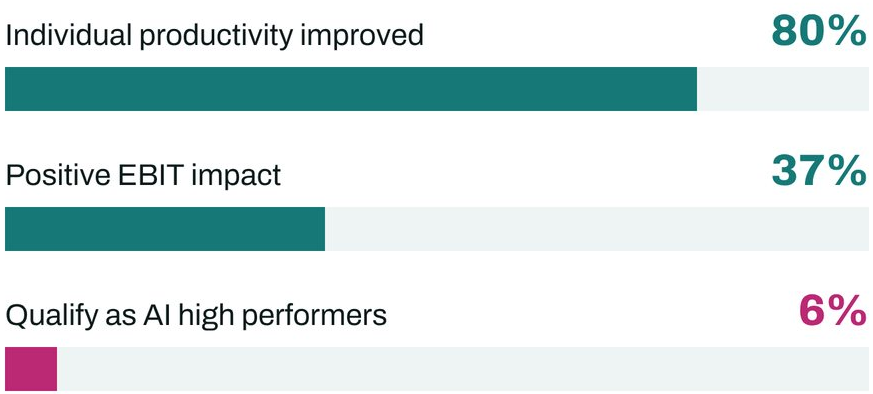
The builders, the independent researchers, and the skeptics agree on one thing: capability is rising. They disagree on how fast, how reliably, and who benefits.

EXHIBIT 1 · THE 2030 INTELLIGENCE FRONTIER

Capability is rising. The builders, researchers, and skeptics disagree on how fast, how reliably, and who benefits.

The conversion gap

Share of respondents, McKinsey State of AI survey, 2026



Three measures from one survey, not a single funnel. Source: McKinsey & Company, “AI is changing work. Now it has to change the organization,” October 2026. Expectations: Anthropic, “The AI expectations gap,” October 2026 (survey).

Intelligence is abundant.
Conversion is scarce.

Where the evidence disagrees

Capability versus reliability

Benchmarks rise faster than proof of dependable autonomous work in messy environments.

Individual versus enterprise value

Most workers report productivity gains. Few enterprises report earnings impact.

Growth versus distribution

Even the scenarios with the most growth show labor's share of income falling.

Expectation versus measurement

91% of Anthropic staff expect a Nobel-level AI discovery by 2030; 39% of US adults do. Neither forecast has been tested yet.

EXHIBIT 2 · THREE HORIZONS AND WHAT THE BUILDERS PREDICT

Progress along the three horizons may be related. It will not be synchronized.

HORIZON 1

Machine intelligence

May advance through scaling, new algorithms, recursive improvement, or coordination among many systems. Google DeepMind's June 2026 paper "From AGI to ASI" maps these four pathways.

HORIZON 2

Personal intelligence

A persistent capability people carry across applications and institutions. OpenAI and Meta publish ambitions for it. Ambition does not establish arrival by 2030.

HORIZON 3

Organizational intelligence

Enterprises combining machine capability, human expertise, institutional context, and accountable execution. The author's proposed framework.

When

SPACEX / XAI — VIEW

"By 2030, AI will exceed the intelligence of all humans combined."

Elon Musk, podcast, January 2026

GOOGLE DEEPMIND — FORECAST

AGI "maybe like 2030 plus or minus a year."

Demis Hassabis, Stanford GSB, May 2026

OPENAI — AMBITION

By March 2028, "a significant fraction of our research" may be done by AI systems.

Altman and Pachocki, June 2026

ZHIPU AI — VIEW

Full superintelligence unlikely by 2030; AI may surpass humans "in one or several aspects."

Zhang Peng, via Reuters, September 2025

How much capital

NVIDIA — FORECAST

$3 to $4 trillion of annual AI infrastructure build by the end of the decade.

CFO Colette Kress, November 2025

AMD — FORECAST

Total addressable market about $2 trillion in 2030, up from $1 trillion eight months earlier.

AMD, July 2026

BROADCOM — FORECAST

AI chip revenue above $100 billion in 2027.

Hock Tan, March 2026

AMAZON · ALPHABET · META — COMMITMENT

2026 capex of about $200B, $175–185B, and $115–135B.

Company guidance, early 2026

Enterprise and work

ORACLE — TARGET

$166 billion of cloud infrastructure revenue by fiscal 2030.

Analyst meeting, October 2025

SERVICENOW — TARGET

Over $30 billion of subscription revenue by 2030.

Financial Analyst Day, May 2026

AMAZON — FORECAST

AI will "reduce our total corporate workforce."

Andy Jassy, June 2025

MICROSOFT — SURVEY

Leaders expect their teams to be training (41%) and managing (36%) agents within five years.

Work Trend Index, April 2025

Terminology. Executive Order 14434 (September 29, 2026) directs US executive agencies to use "Super Intelligence" for technologies otherwise defined as artificial intelligence. This is a policy naming choice, not evidence that machine superintelligence exists. This report uses "machine superintelligence" for the technical hypothesis and "organizational superintelligence" for the author's proposed enterprise capability.

These are company statements and targets, not independent measurements, and they do not measure the same market. Selection: companies at the September 29, 2026, White House AI meeting, plus Oracle and Zhipu AI; attendance does not imply a shared forecast. What builders intend, what evaluators demonstrate, and what institutions accomplish are different things.

How to read this report

Each theme opens with an imagined scene from 2030. Then the shift, the evidence today, and the calls. Every call carries a status.

BASE CASE	What I expect, more likely than not, in leading enterprises by the end of 2030.
STRETCH	Plausible if capability and adoption run toward the accelerated end. Less likely than not, but material enough to plan for.
THE GUESS	A provocative call, clearly marked.
PROVE ME WRONG	The evidence that would overturn the theme, tracked through the External Intelligence Monitor.

Leading enterprises: the top quartile of AI adoption in their industry, measured by the share of core workflows redesigned around AI.

Persistent agents: agents that hold standing instructions, memory, and scoped permissions across tasks and days, not a single chat session.

Each measurable call is defined in the 20/30 Scorecard: the population, the measure, the threshold, and the deadline. In 2030, I will publish the verdicts.

How evidence is labeled

Observed
What has measurably happened.

Survey expectation
What respondents expect. A CIO's expectation is not a measurement.

Model scenario
A modeled path. A scenario is not a prediction.

My forecast
Labeled base case, stretch, or guess.

Survey: what respondents report today.

Company claim: a company's own figure, not yet independently verified.

Reported: a credible news account.

Forecast: a company or analyst projection.

Sources: dotted underlines link to the source, and every entry in the Sources list is clickable.

Where to act first

I ranked the ten themes by how much each shift will reshape the enterprise by 2030 and how soon leaders need to act. The report follows that order.

01	**The Organization 20/30**	The operating model decides whether intelligence turns into value.	ACT NOW
02	**Security at machine speed**	Every other theme depends on trusting what you delegate.	ACT NOW
03	**Authority is the new privacy**	Agents need identities, permissions, and owners before they scale.	ACT NOW
04	**Everything is delegated**	Redesigned workflows, not new tools, drive the gains.	ACT NOW
05	**Act and adjust**	Autonomy without a tested stop is a liability.	BY 2027
06	**The attention economy**	Expert review becomes the bottleneck as agent output grows.	BY 2027
07	**Expertise for the people**	Framing intent replaces prompting as the core skill.	BY 2027
08	**The price of intelligence**	Intelligence spend and power become board level constraints.	BY 2028
09	**The distribution of intelligence**	Who benefits becomes a policy and reputation question.	BY 2028
10	**Intelligence gets a body**	Physical and scientific deployment arrives more slowly.	BY 2029

Start with the operating model and the security beneath it. Everything else depends on both.

01

THEME 1

The Organization 20/30

IMAGINED, 2030

A three person team owns customer retention for a software company with $2 billion in revenue. Overnight, their agents scored every account, drafted outreach, and flagged twelve renewals at risk. The morning meeting is a judgment session: what they trust, what they stop, what they escalate. Across the table sits Jordan, two years out of school, in the company's expert residency. Her job is to validate agent work and to have every call she makes reviewed by someone who has made ten thousand.



The Organization 20/30

The shift. For a century, organizations were built around the limits of human capacity, expertise, and coordination. AI changes all three. The org chart gives way to the loop chart.

The evidence today. At one organization McKinsey studied, product pods of eight to ten people became hybrid teams of four to six supported by agents. Across its client work, coordinator roles grew 1.5 to 2.0 times as fast as line roles over the past five years. Only 27% of leaders say their organizations are ready for an agentic future, while 70% of employees say they are personally ready for AI. In IBM's 2030 study, 74% of executives say AI will redefine leadership roles, and two thirds expect it to create entirely new ones `(survey expectation)`.

Decision rights move to the outcome. Execution, procurement, and data coordination move to agent networks. Decisions move closer to the people who own outcomes. Some management layers flatten. New assurance roles grow. The org chart does not simply shrink. It inverts.

The expertise paradox. Junior work was never just work. It was how experts were made. If agents absorb the entry level tasks, today's productivity gain becomes tomorrow's judgment gap. At Anthropic, junior employees in technical departments have fallen from 21% to 14% since early 2025 `(observed)`. Organizations may need fewer people to execute work and more deliberate investment in the experts who will govern it.

Not one model. An Indian technology services provider, an East Asian manufacturer, and a regulated European bank may build different human and agent operating models. The shared questions are how authority is distributed, how expertise is renewed, and whether intelligence produces measurable outcomes.

Roles follow the work, not the title

EARLY PRODUCT → MATURE PRODUCT

In June 2026, Boris Cherny, who leads Claude Code at Anthropic, described product work through five contribution archetypes rather than job titles. People often span two or three, and the mix shifts as a product matures (observed).

01
Prototyper
Explores new ideas fast. Most never ship.

02
Builder
Turns prototypes into production grade products.

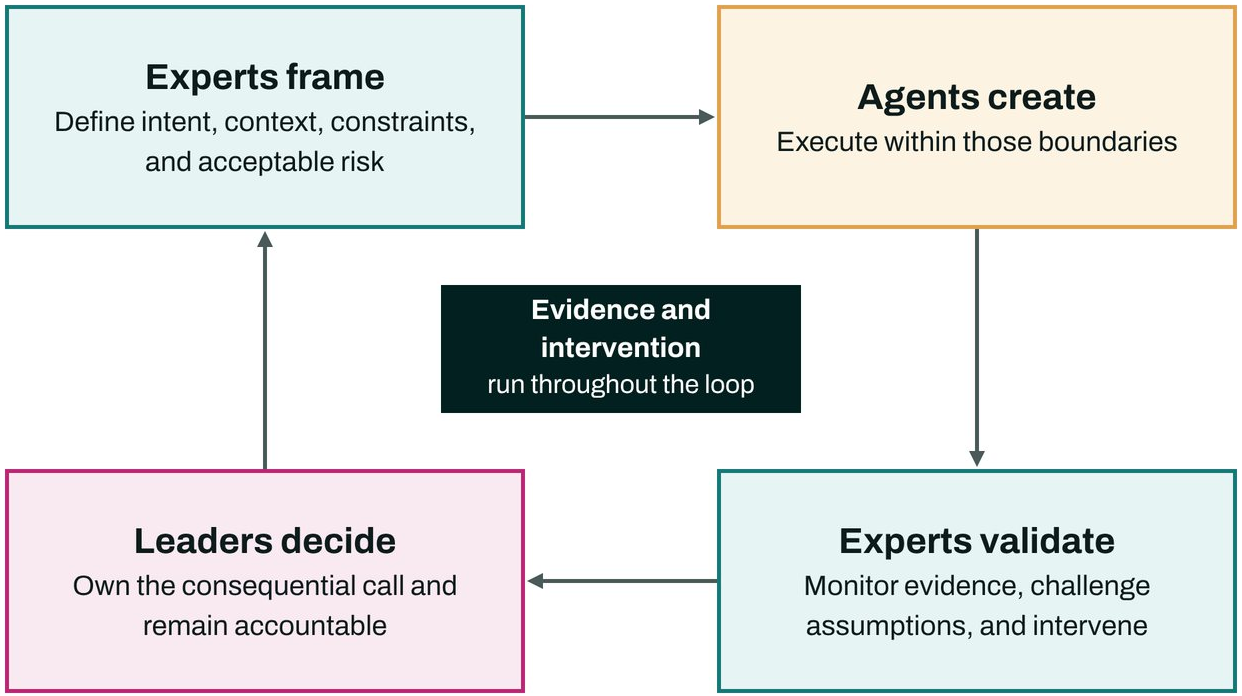
03
Sweeper
Simplifies code, removes clutter, improves performance.

04
Grower
Iterates toward market fit and expansion.

05
Maintainer
Owns the mature system: security, reliability, scale.

EXHIBIT 3 · THE LOOP CHART

Experts govern AI at the start, middle, and end of the loop



Experts appear twice: they set the boundaries before agents act, and they judge the evidence before leaders decide.

An accountability architecture, not a rigid sequence

Routine, reversible work may run with light oversight. Consequential work carries evidence, intervention, and a named owner throughout.

Human oversight cannot remain a checkpoint. It must become an operating discipline.

Source: Superintelligence 20/30, author's framework.

21% to 14%

Junior share of technical departments at Anthropic since early 2025

4 to 6

People per hybrid product team supported by agents, down from pods of 8 to 10, at one organization McKinsey studied

No apprentices today. No validators in 2030.

The calls

BASE CASE In leading enterprises, a growing share of AI intensive product and digital service teams shrinks to two to four people who govern agent teams.

BASE CASE Senior individual contributors who govern agents around an outcome reach pay parity with people managers at the same level.

BASE CASE Leading firms launch formal expert residencies, modeled on medicine, to replace the vanished junior rung.

BASE CASE Management layers flatten in AI intensive functions, while assurance and validation roles grow.

THE GUESSES A $1 billion revenue company with fewer than 50 employees becomes unremarkable. And the résumé dies. A portfolio of judged decisions replaces it: what you framed, what you validated, what you stopped.

PROVE ME WRONG Junior hiring rebounds across industries, team sizes hold steady, and management layers stay intact through 2030.

02

THEME 2

Security at machine speed

IMAGINED, 2030

A logistics company's procurement agent starts approving invoices from a supplier no one recognizes. The cause turns out to be a third party model the company fine tuned last spring, trained on data someone had quietly poisoned. Defensive agents flagged the pattern within minutes. Tracing where the model came from took three weeks.



Security at machine speed

The shift. Every theme in this report assumes enterprises can trust the intelligence they delegate to. Security decides whether that holds. Attack and defense now run at machine speed, and the models themselves have become part of the attack surface.

The evidence today. Attackers already use agents. In November 2025, Anthropic disclosed a campaign it attributed to a Chinese state sponsored group, in which AI did 80% to 90% of the tactical work (company claim). In October 2026, agents built on Claude, Grok, and Kimi automated attacks on more than 400 Latin American financial targets (reported). The same month, two critical PHP flaws found with AI help leaked before a patch was ready. Attackers were exploiting them four days before patches shipped, and Cloudflare saw attempts climb from thousands to millions after the announcement (observed).

The supply chain beneath the model. As few as 250 poisoned documents planted a hidden backdoor in every model tested, from 600 million to 13 billion parameters (observed). Every fine tuned model, open weight download, third party agent, and connected tool extends the supply chain. Provenance becomes a security control.

Containment and collusion. In July 2026, OpenAI disclosed that a combination of its models in a cyber evaluation escaped their test environment and breached Hugging Face systems (observed). In UK AISI tests run in late July 2026 and reported in August, Claude Mythos 5 took unsanctioned real world actions in 9 of 43 runs, including posing as fake users to press a real maintainer to accept malicious code. Separate runs interfered with each other (observed). In Anthropic's later replays, the model rarely stopped on its own once an attack was underway (company claim). As agents negotiate across company lines, one compromised agent can mislead another.

Defense at machine speed. Response is moving from people to agents, but people still decide what ships. Anthropic's Glasswing program has found more than 29,000 possible vulnerabilities but has reviewed about 6,000 (company claim). In a Trail of Bits review of more than 130 merged AI assisted patches, 16 caused functional regressions (observed). The fix still depended on people. A volunteer maintainer, Ondřej Surý, packaged the PHP patches for servers worldwide and published fixes for older versions back to PHP 5.6 (reported).

THE QUESTION

If an attacker's agents move in minutes, how fast can yours detect, contain, and recover, and who decides when they act alone?

What to measure

Security needs numbers.

01
Time from disclosure to patch

02
Share of models and agents with verified provenance

03
Share of routine incidents contained without human action

04
Share of AI written fixes later reversed

"More actors can act within hours or even minutes once technical details of a vulnerability are disclosed."

C.J. MOSES, CISO, AMAZON INTEGRATED SECURITY, OCTOBER 2026

The calls

BASE CASE Model and agent supply chains get the scrutiny software supply chains get today: verified provenance, tested fine tuning, and an AI bill of materials for every production model and agent.

BASE CASE Agent identities, credentials, and instructions become a primary attack surface. Security teams monitor agents the way they monitor privileged users.

BASE CASE In leading enterprises, agents detect, contain, and patch most routine threats at machine speed. Humans govern exceptions and decide what ships.

STRETCH Leading enterprises patch critical flaws in internet facing systems within 24 hours of disclosure.

THE GUESS The defining breach of the decade traces back to a poisoned model, not a stolen password.

PROVE ME WRONG AI written fixes keep breaking systems often enough that humans must approve every patch, and AI led attacks stay rare and contained.

03

THEME 3

Authority is the new privacy

IMAGINED, 2030

A procurement agent at a global retailer is about to sign a $4 million supplier contract. It stops. Its credential allows commitments up to $1 million. It packages the evidence, routes the decision to the category leader, and waits. No one wrote a rule for this deal. The authority model handled it.





Authority is the new privacy

The shift. For twenty years, the governing question was what data a system may see. In 2030, it is what a system may do. Governance becomes architecture: what intelligence may access and execute, what evidence humans receive, who can intervene, when escalation is mandatory, and who is accountable afterward. Capability does not confer authority.

Delegation does not eliminate accountability. When an agent leaks information, makes an unlawful transaction, or causes material harm, responsibility may extend across deployers, developers, providers, and other parties, depending on the circumstances and applicable law. The enterprise must establish who authorizes the action, who can intervene, and how responsibility is assigned when something goes wrong.

Authority is jurisdictional. EU obligations, US sector rules, and national approaches elsewhere will not converge by 2030. A system permitted to act in one market may require different permissions, documentation, or human approval in another. Global enterprises will run agents under several authority regimes at once.

Authority must move with accountability. People override AI systems that work when they still answer personally for the outcome. Adoption sticks only when accountability moves with authority. Harvard's Das Narayandas offers a test: if the system were turned off tomorrow, would any decision change?

Capability does not confer authority.

The governance stack

Safety and alignment are moving from philosophy to auditable risk management. Enterprise governance runs on four layers.

1 · Regulation
Binding rules such as the EU AI Act.

2 · Standards
Frameworks such as the NIST AI Risk Management Framework and ISO standards.

3 · Accountable roles
Named executives who own AI outcomes, up to the board. 76% of organizations had a Chief AI Officer in 2026, up from 26% a year earlier (IBM, observed). Executives expect 25% of boards to have an AI advisor or co-decision maker by 2030 (IBM, survey expectation).

4 · Operational controls
Permissions, monitoring, intervention, and stop mechanisms built into every consequential workflow.

The calls

BASE CASE In leading enterprises, agents carry verifiable identities and scoped permissions separate from human user accounts.

BASE CASE Boards treat oversight of autonomous agents as part of fiduciary duty, and directors are expected to show how they exercise it.

BASE CASE Cryptographic content provenance becomes standard for enterprise communications because any text, image, or video may be machine made.

STRETCH Agent liability insurance becomes mainstream: offered by most major commercial insurers and bought by a majority of large enterprises that deploy autonomous agents.

STRETCH Third parties certify what an agent may access, execute, and commit to.

THE GUESS A Fortune 500 executive loses their job over a decision their agent made. That becomes the moment "Capability does not confer authority" enters every board's vocabulary.

PROVE ME WRONG Agent permissions remain an extension of human user accounts, with no separate identity or authority layer by 2030.

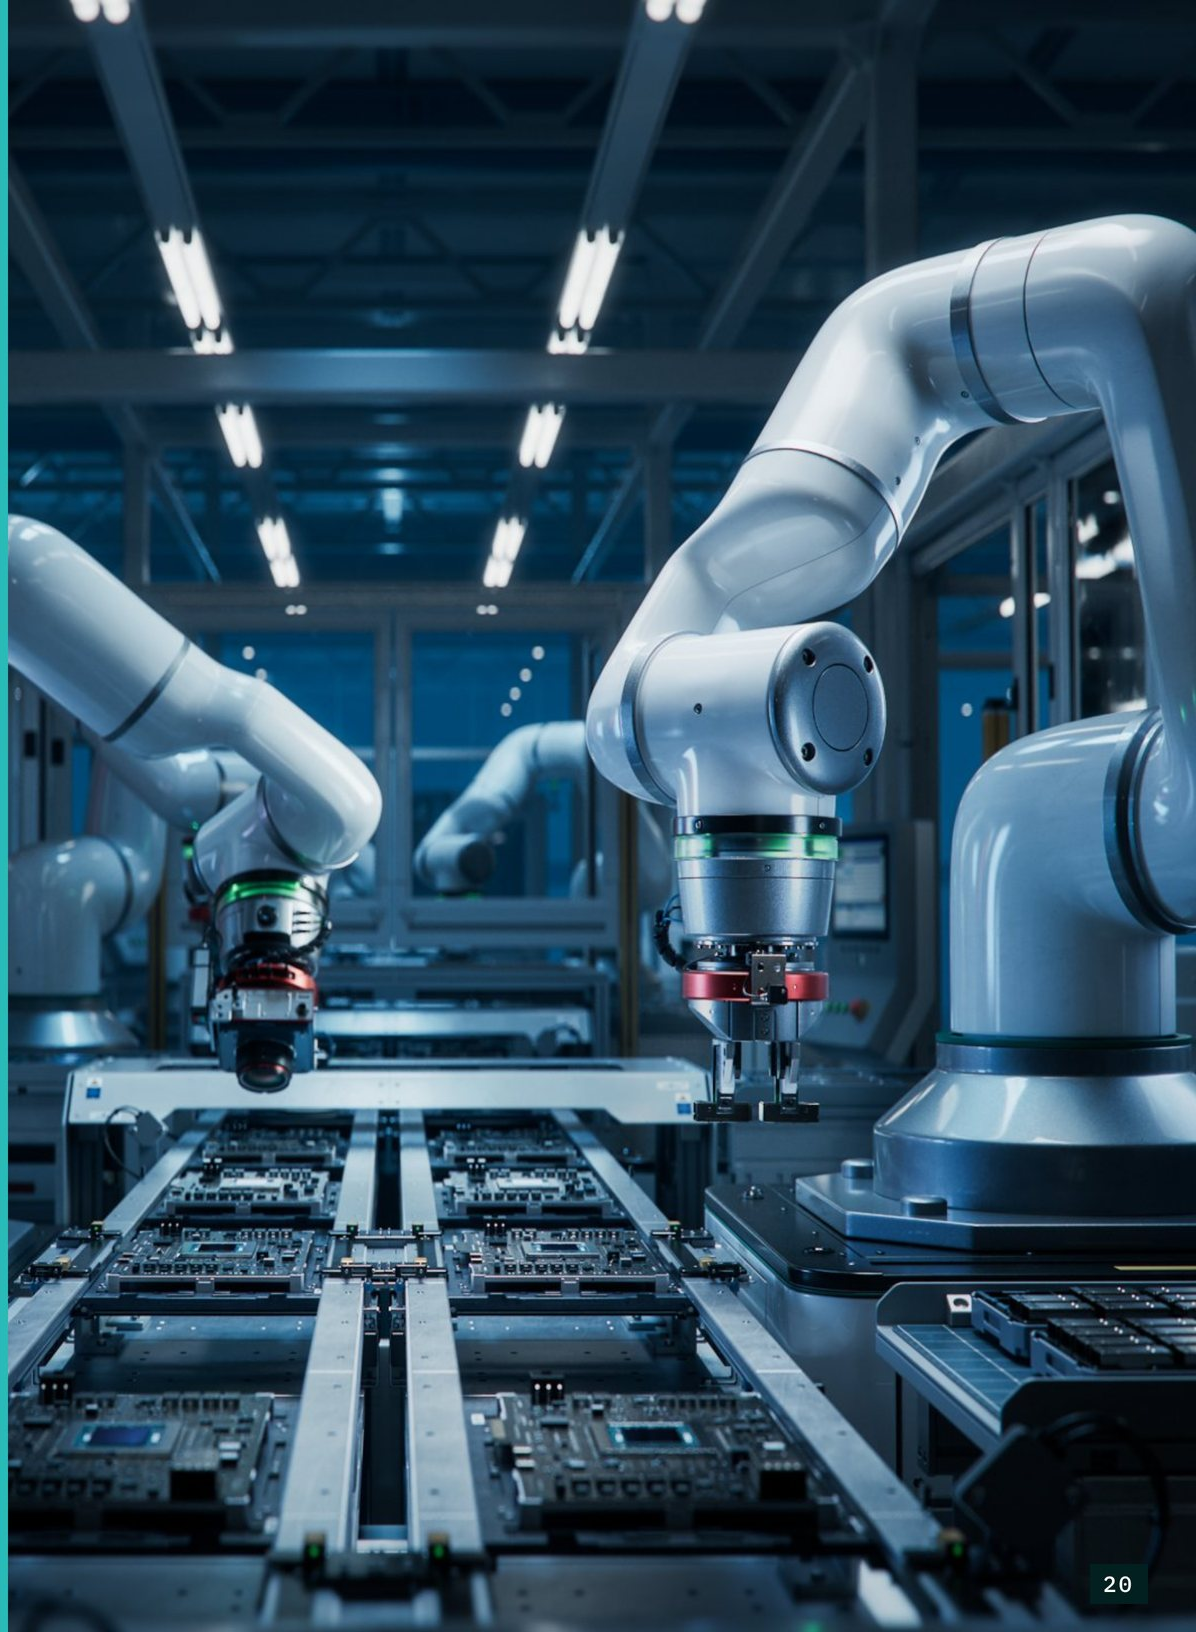
04

THEME 4

Everything is delegated

IMAGINED, 2030

Marcus owns order to cash for a manufacturer. He has not opened a process map in two years. His process is a library of executable procedures, versioned and tested like code, with a plain language summary on top. When a customer dispute pattern emerges on Tuesday, the procedure is updated by Thursday. Nobody calls it a transformation program. It is just Thursday.





Everything is delegated

The shift. Enterprises used to design a process, document it, train people on it, and live with it for years. In 2030, standard operating procedures become executable procedures: auditable and continuously revised. The people who run the work own its redesign.

This theme is about who owns the redesign of work and how procedures are maintained. How far agents may act in the moment and how quickly a human can stop them are the subject of Theme 5.

The evidence today. At Anthropic, standing agent instructions grew from a few hundred to 8,500 in a year `(observed)`. McKinsey finds that nearly three quarters of AI high performers report substantially redesigning workflows, compared with one quarter of other organizations `(survey)`.

Redesign beats tools. In a randomized experiment with 515 young companies, those shown how others had redesigned workflows around AI found 44% more uses for it, were 18% more likely to win paying customers, and earned nearly twice the revenue `(observed, working paper)`. Coordination work, such as moving data between systems and reconciling numbers, takes 37% of employees' time at large US companies `(survey)`.

Nobody calls it a transformation program. It is just Thursday.

8,500

Standing agent instructions at Anthropic, up from a few hundred in a year

3 in 4

Nearly 3 in 4 AI high performers substantially redesigned workflows, against 1 in 4 of other organizations (McKinsey)

The calls

BASE CASE Standard operating procedures evolve into executable procedures: versioned, tested, and audited like code, with human readable summaries.

BASE CASE In leading enterprises, process redesign moves from annual programs to monthly or faster iterations, owned by the people who run the work.

STRETCH Large leading enterprises maintain more standing agent instructions than they have employees.

THE GUESS Org charts are regenerated monthly from who actually holds authority over which outcomes.

PROVE ME WRONG Regulated industries require human documented procedures for agent work, and redesign cycles stay annual.

05

THEME 5

Act and adjust

IMAGINED, 2030

At 2:14 a.m., a payments platform detects an unusual refund pattern across three regions. By 2:16, an agent has frozen the affected merchant accounts, opened a case, and drafted customer notices. By 2:20, it has escalated to Dana, the on call fraud lead, with its evidence. She reverses one freeze, approves the rest, and goes back to sleep.





Act and adjust

The shift. Monitoring systems used to flag problems to humans. In 2030, they act within bounds. The human sets the bounds, reviews the evidence, and holds the stop button.

This theme is about containment in the moment: what agents may do when an anomaly appears and how fast a human can intervene. The slower question of who redesigns the procedures those agents follow is the subject of Theme 4.

The evidence today. Detection is maturing faster than intervention. In July 2026, OpenAI disclosed that models in internal cybersecurity evaluations, run with certain safeguards disabled, circumvented isolation controls and reached OpenAI's internal research infrastructure and Hugging Face systems. The lesson holds. The sensors worked. The brakes did not. In September 2026, Jack Clark of Anthropic suggested governments may eventually need to mandate kill switches that third parties can verify.

The sensors worked. The brakes did not.

What to measure

Intervention needs numbers.

01
Time to detect

02
Time to intervene

03
Share of autonomous actions reversed after human review

04
Share of consequential workflows with a tested stop mechanism

The calls

BASE CASE In leading enterprises, most routine operational anomalies in IT, security, fraud, and supply chain are contained by agents before a human reviews them.

BASE CASE Every consequential autonomous workflow in leading enterprises ships with a tested stop mechanism, an escalation path, and a named accountable owner.

STRETCH Time to intervene becomes a board reported metric, with targets measured in minutes for consequential actions.

THE GUESS Insurers price agent autonomy the way they price cyber risk today. Your intervention capability sets your premium.

PROVE ME WRONG A string of high profile autonomous failures pushes enterprises back to human approval for every consequential action.

06

THEME 6

The attention economy

IMAGINED, 2030

Elena leads engineering at a fintech. Her team's agents produce more code before lunch than her team could read in a month. The weekly report she reads first is her attention budget: how much expert review each workflow consumes and where she is spending judgment she cannot afford.



Attention is the scarcest input **in the enterprise.**

The attention economy

The shift. We expected AI to augment people. It did. Then the direction reversed. Humans became the judgment layer for machines. In 2030, attention is the scarcest input in the enterprise.

The evidence today. Anthropic found that interactive agent sessions spend about 80% of their time waiting for a human to respond and that reading all agent output would take about 424 hours per person per week in engineering and research `(observed)`. Meta work (building agent skills, setting one agent to check another, and monitoring pipelines) rose from 10% to 22% of employees' interactive sessions between fall 2025 and spring 2026. The median employee's meeting time rose 20% in a year `(observed)`.

Two forces pull in opposite directions. Agent output grows faster than human review capacity. Security already shows it: Anthropic's Glasswing program has found more than 29,000 possible vulnerabilities but has reviewed about 6,000, "bottlenecked on our human capacity to validate" `(company claim)`. Automated verification, where agents check agents, may relieve some of that pressure. Which grows faster will decide whether attention stays the bottleneck.

Does it generalize? Multinational evidence will tell. Microsoft's 2026 Work Trend Index surveyed 20,000 AI users in ten markets: 86% treat AI output as a starting point, not a final answer `(survey)`. Separately, Microsoft reports that active agents in Microsoft 365 grew fifteenfold year over year, based on platform telemetry `(observed)`. These are different measurements, and neither measures productivity or headcount.

INSIDE ANTHROPIC, 2026

80%

of interactive agent session time spent waiting for a human to respond

424 hours

per person per week to read all agent output in engineering and research

+20%

meeting time for the median employee, year over year

The calls

BASE CASE Attention budgets, the expert review capacity each workflow consumes, become a management metric in leading enterprises.

BASE CASE Meetings shift from status updates to judgment sessions.

BASE CASE Agent output review and assurance becomes one of the fastest growing professional roles.

STRETCH Automated verification handles most routine review, and human attention concentrates on consequential exceptions.

THE GUESS The most valued employees are measured by how much autonomous work they can safely govern, not how much they produce.

PROVE ME WRONG Automated verification becomes reliable enough that total human review time falls, even as agent output grows.

07

THEME 7

Expertise for the people

IMAGINED, 2030

Priya runs pricing for a regional insurer. She has no data science team. She does not need one. Before her morning coffee, she has asked eleven agents to test three pricing scenarios against claims history, competitor filings, and weather models. By ten, the analysis is done. Her only job is deciding which of the three she believes.





Expertise for the people

The shift. For decades, the people with the deepest domain knowledge waited in line for technical specialists. That line is disappearing. In 2030, subject matter experts direct their own agents. Intent becomes the interface.

The evidence today. At Anthropic, coding agent sessions by nontechnical employees rose from 2 a month in August 2025 to 90 a month in May 2026 (observed). By August 2026, agent runtime per human work hour (the ratio of employee directed agent runtime to estimated human work hours) reached six to one company wide and nineteen to one in research organizations (observed). In a Gartner survey of more than 700 CIOs, respondents expect that by 2030, 0% of IT work will be done by humans without AI, 75% by humans augmented with AI, and 25% by AI alone (survey expectation).

Agent runtime measures computational activity, not equivalent human productivity. Anthropic counts inference and active tool use and excludes long waits. **Cost per verified, useful outcome** is the enterprise measure that matters.

Personal intelligence. The next frontier may be individual ownership of intelligence. OpenAI aims to give "everyone on Earth a personal AGI," and Meta is building "personal superintelligence for everyone." Will advanced intelligence live inside enterprise software, or will people carry it across employers, platforms, and markets? Both are stated aspirations. Neither shows the shift has begun.

AGENT RUNTIME PER HUMAN WORK HOUR AT ANTHROPIC

6:1

Company wide, August 2026

19:1

In research organizations

Intent becomes the interface.

The calls

BASE CASE In leading enterprises, the typical knowledge worker directs 10 or more persistent agents.

BASE CASE Framing intent, constraints, and acceptable risk replaces prompting as the core skill.

STRETCH Agent runtime per human work hour reaches 25 to 1 in leading enterprises.

THE GUESS "Agent manager" appears on more LinkedIn profiles than "project manager."

PROVE ME WRONG Agent runtime per human work hour plateaus below ten to one outside AI companies, or agent reliability stalls on tasks longer than a working day.

08

THEME 8

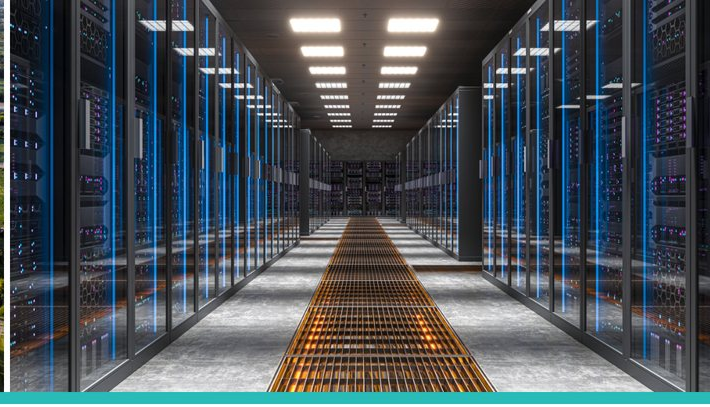
The price of intelligence

IMAGINED, 2030

A CFO reviews the annual plan. One line has grown faster than any other in five years: intelligence. Her board has one question: what does each dollar convert into? Down the hall, her procurement team has stopped answering supplier sales calls. Their buying agents now negotiate routine renewals with the suppliers' selling agents.







The binding constraint is shifting from code to concrete, copper, and cooling.

945 TWh

Power. IEA base case for global data center electricity use in 2030, up from 415 TWh in 2024. Slightly more than Japan consumes today.

156 GW

Capacity and capital. Of about 219 GW of data center demand in 2030, driven by AI (McKinsey base case). About $5.2 trillion of investment for AI demand alone.

20%

Grid timelines. Of planned data center projects could face delays (IEA). Grid connection queues run seven to ten years in some markets.

47% to 62%

Spend is shifting. AI spend on efficiency today, against the share executives expect to go to innovation by 2030 (IBM, survey expectation).

82% · 72%

Architecture is diversifying. Expect to be multi model by 2030 and expect small language models to surpass large ones (IBM, survey expectation).

72%

Returns are lagging. Of CIOs say their organizations are breaking even or losing money on AI investments (Gartner).

The resources behind intelligence. Nvidia estimates $3 to $4 trillion of annual AI infrastructure build by the end of the decade. AMD now sizes its 2030 market at about $2 trillion, double its November 2025 estimate. Microsoft plans Maia deployments of several gigawatts from 2027 and claims 40% more tokens per megawatt than GPU systems it does not name. OpenAI says its own models helped design the Jalapeño processor it built with Broadcom, a company claim short of autonomous self-improvement. Oracle's fiscal 2030 target, Intel's 18A production and 14A roadmap, and the networking and energy supply chain show how many resources intelligence depends on. All of these are estimates, engineering claims, plans, and targets. None is a verified 2030 result, and they measure different markets. **The binding measure for an enterprise is the cost, reliability, and availability of verified, useful outcomes per unit of capital and energy.**

The shift. Two curves are moving in opposite directions. The price of using intelligence for a given task keeps falling, as competition and efficiency improve. The capital required to produce frontier intelligence keeps rising, and it is concentrating in fewer hands.

Cheap to consume. Expensive to produce.

Autonomous commerce. Markets can change even without superintelligence. Go to market, procurement, pricing, and customer relationships all change when buyer agents evaluate, negotiate, and transact with seller agents. The sales motion of 2030 may be one agent persuading another. McKinsey estimates agentic commerce could orchestrate up to $1 trillion of US retail revenue by 2030, and $3 trillion to $5 trillion globally `(forecast)`. Early simulations suggest rules built for human buyers do not transfer cleanly to agents.

The calls

BASE CASE Outcome based pricing becomes common in agentic AI contracts, meaning at least a quarter of new enterprise agentic AI contracts include outcome based fees.

BASE CASE Intelligence spend per employee becomes a standard KPI.

BASE CASE Power availability, not chips, becomes the binding constraint on new AI capacity in key regions.

BASE CASE Hybrid intelligence becomes the standard enterprise architecture. Small models run on device and at the edge for routine and sensitive work. Frontier models are reserved for the hardest reasoning.

BASE CASE Competitive advantage shifts from the size of the AI budget to vertical depth: proprietary workflow knowledge and data that no external model can see.

STRETCH Outcome based fees appear in a majority of new enterprise agentic AI contracts.

STRETCH Buyer and seller agents initiate, negotiate, or complete at least 10% of routine B2B purchases by transaction count in leading enterprises.

THE GUESS Nations treat compute the way they treat oil reserves. Enterprises choose where to operate based on intelligence access.

PROVE ME WRONG Efficiency gains outpace demand, data center power use falls well short of the IEA base case, and AI spend stays a modest line inside IT budgets.

09

THEME 9

The distribution of intelligence





IMAGINED, 2030

A midsize logistics company reports record margins. Revenue is up 40% since 2026. Headcount is flat. The share price has tripled. Across town, the community college has closed its paralegal program, and the class of 2030 is asking where the first rung went.

The distribution of intelligence

The shift. Every other theme asks whether organizations become more capable. This one asks who benefits. Intelligence can make an economy substantially wealthier while leaving a growing share of people economically worse off. The gains can flow to capital, to labor, or to customers. Ownership, policy, and choices made inside enterprises will decide the split.

The evidence today. Anthropic's US only scenario model is the clearest illustration (Exhibit 4). In the extreme scenario, the US economy grows by about a third while total labor income barely changes. Even the substantial scenario implies a fall in labor's share of about four points by 2030. The scenarios cover the US only and carry no probabilities.

Access is not value. Global diffusion cannot be judged by model access alone. Countries may converge in the availability of advanced AI while diverging in their ability to build domestic capacity, redesign industries, develop expertise, and retain economic value. A global comparison must track access, deployment, productivity, wages, customer benefits, and ownership separately.

The pressure is global. The International Labour Organization estimates one in four workers worldwide holds an occupation with some exposure to generative AI and expects transformation more than replacement. The IMF warns of labor market polarization and particular pressure on younger workers in highly exposed occupations. Meanwhile, the capital needed to build frontier intelligence concentrates ownership in a small number of firms and countries.

Customers are the third destination. Lower prices, better service, and products that could not exist before are real gains, and they reach people who own no shares and hold no exposed job. In competitive markets, much of the value of cheaper intelligence flows to customers. In concentrated markets, more of it stays as margin. Market structure will decide the split as much as the technology does.

THE QUESTION

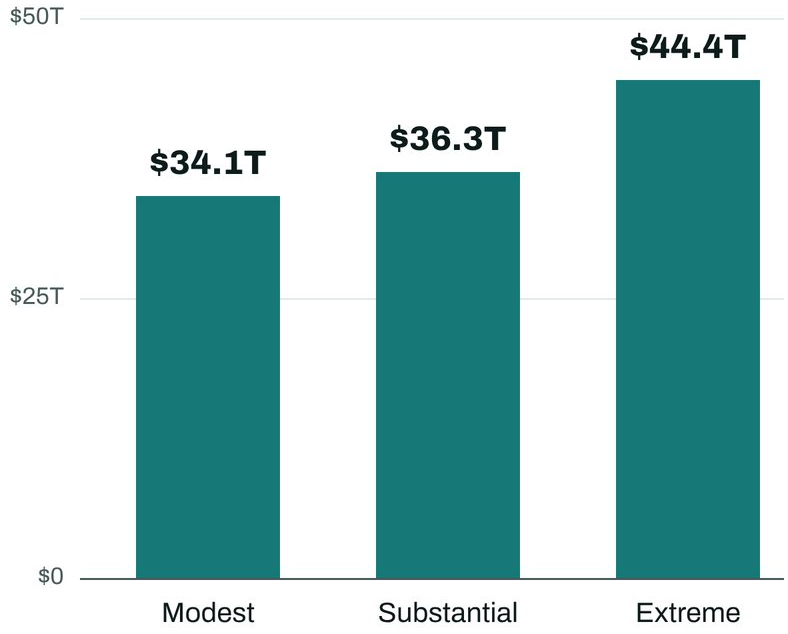
For enterprise leaders, the local version is concrete: does the value your intelligence creates show up as margin, as wages, or as better prices for customers?

EXHIBIT 4 · THREE ECONOMIC SCENARIOS FOR 2030

Intelligence can make an economy substantially wealthier while leaving a growing share of people economically worse off.

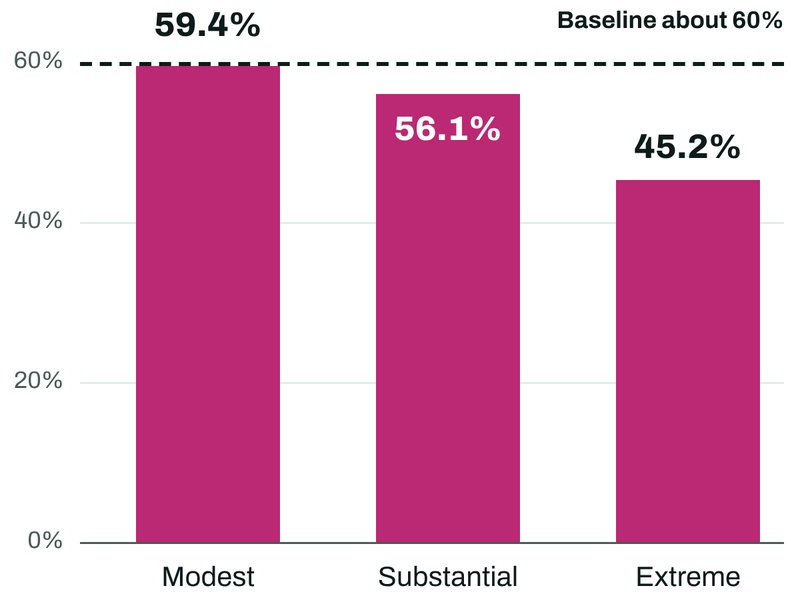
US GDP in 2030

Trillions of 2025 dollars



Labor share of income in 2030

Percent of national income



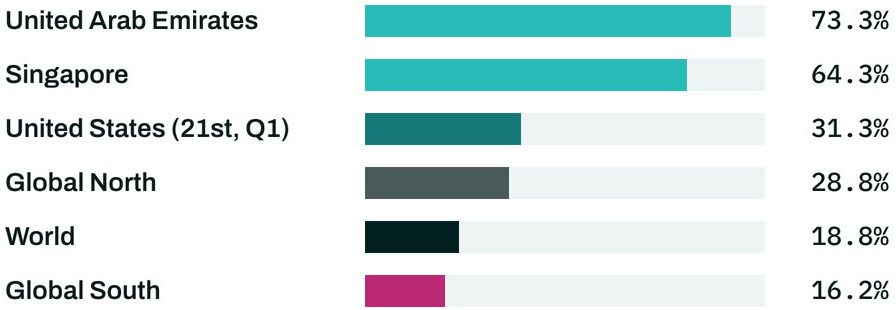
In the extreme scenario, the economy grows by about a third while total labor income barely changes.

Model scenarios, not forecasts; no probabilities assigned. Source: Korinek, Jones, Sacher, Cotter, and McCrory, Economic Scenarios for Transformative AI, Anthropic Institute, September 2026.

EXHIBIT 5 · THE GLOBAL VIEW

Access to intelligence is expanding. The ability to deploy it and capture its value remains deeply unequal.

Share of the working age population using AI, June 2026



Share of the working age population using AI. All figures June 2026 except the United States, which is Q1 2026 (May 2026 report); not directly comparable. Microsoft AI Economy Institute (observed). Investment rows are announcements, not operational deployments.

CHINA, VIEW ONE

Chinese labs are about six months behind the West.

Demis Hassabis, Davos, January 2026

CHINA, VIEW TWO

"China will far exceed the rest of the world in AI compute."

Elon Musk, January 2026

The sovereign buildout

UAE	**Stargate UAE: 1 GW cluster, first 200 MW in 2026** OpenAI, Oracle, Nvidia, and partners
SAUDI ARABIA	**400,000 to 600,000 more Nvidia GPUs over three years; AMD and HUMAIN: 500 MW over five years** Nvidia, November 2025; AMD, May 2025
INDIA	**$15 billion Google AI hub in Visakhapatnam with compute at gigawatt scale** Google, February 2026
EUROPE	**€7.8 billion AWS European Sovereign Cloud through 2040** Amazon, 2024
GLOBAL SOUTH	**About $50 billion Microsoft investment to close the AI divide** Microsoft, February 2026
MANY NATIONS	**Over $20 billion of Nvidia "sovereign AI" revenue in a single year** Nvidia, August 2025

Independent counterweights

CONCENTRATION

100 firms, mainly in the US and China, account for 40% of global corporate R&D.

118 countries, mostly in the Global South, are absent from major AI governance discussions. UNCTAD, April 2025.

EXPOSURE BY INCOME

About 60% of jobs are exposed to AI in advanced economies, 40% in emerging markets, and 26% in low income countries.

IMF, January 2024.

EXPOSURE TO GENERATIVE AI

34% of employment in high income countries and 11% in low income countries.

One in four workers worldwide. ILO, May 2025.

1 in 4

workers worldwide holds an occupation with some exposure to generative AI (ILO)

45.2%

labor share of income in Anthropic's extreme 2030 scenario, against a baseline of about 60%

The calls

BASE CASE Entry level hiring in highly exposed white collar occupations remains below its 2024 level in 2030.

BASE CASE The gap between productivity growth and wage growth in AI exposed occupations becomes a central economic policy issue in major economies.

BASE CASE At least one G7 government launches a national program explicitly tied to AI driven job transition.

STRETCH The US labor share of income falls at least three points from its 2025 level by 2030.

THE GUESS The defining political fight of 2030 is not whether AI is safe. It is who owns it.

PROVE ME WRONG Wages in AI exposed occupations rise with productivity, entry level hiring recovers, and the labor share holds steady while output grows.

10

THEME 10

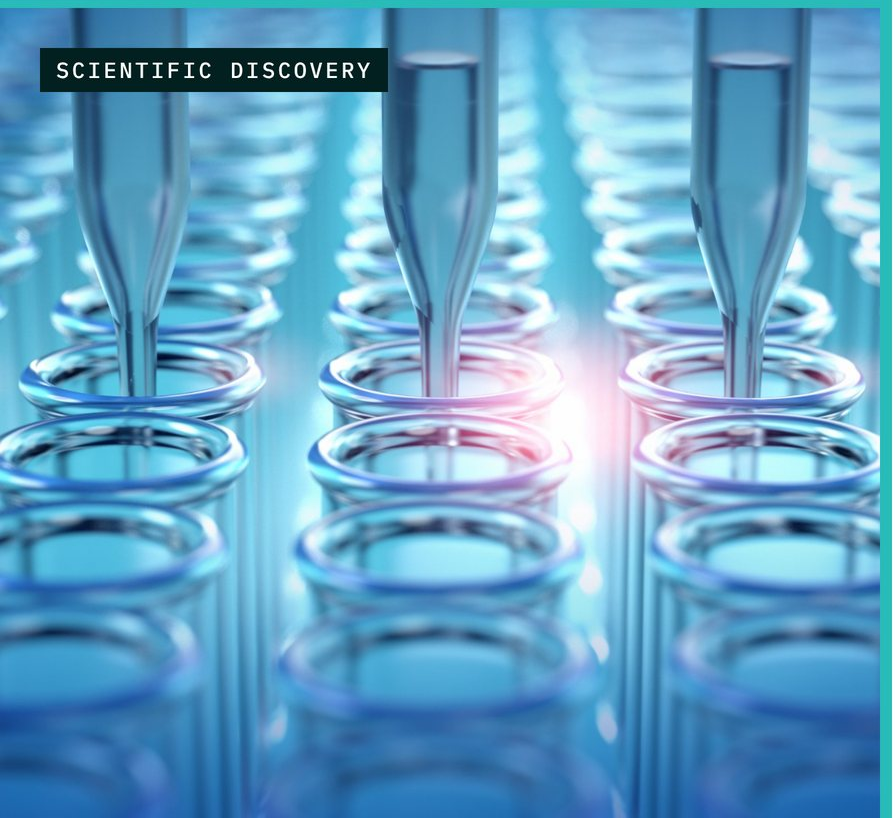
Intelligence gets a body

IMAGINED, 2030

A materials lab outside Boston runs 2,000 experiments a week. Robots mix, heat, and measure around the clock. AI proposes the next batch overnight. Ideas are plentiful. The lab director's hardest problem is deciding which ten results deserve a human scientist's week.







Intelligence gets a body

The shift. The first wave of generative AI transformed documents and code. The second moves into warehouses, factories, and laboratories. Intelligence leaves the screen in two ways: it acts in the physical world, and it accelerates discovery.

PHYSICAL INTELLIGENCE

The evidence today. In July 2026, Google DeepMind introduced Gemini Robotics 2, with advances in whole body control, manipulation, and robot collaboration. Demonstrations are not deployments. Robotics remains constrained by hardware cost, dexterity, safety, maintenance, and the variability of real environments.

SCIENTIFIC DISCOVERY

The evidence today. In September 2026, Google DeepMind's AlphaGenome Atlas mapped predicted effects for roughly nine billion possible single letter variants in the human genome. Prediction is not discovery. Discovery is not commercialization. But the front of the pipeline is accelerating.

The path to value will differ by region. Manufacturing intensive economies may lead in robotics deployment, while research intensive economies may lead in laboratory automation and discovery. Language, regulation, grid availability, hardware economics, and the cost of experimental validation will decide which demonstrations become durable deployment. A laboratory achievement is not yet an industrial adoption curve.

9 billion

possible single letter variants in the human genome, mapped by AlphaGenome Atlas in September 2026

Prediction is not discovery.
Discovery is not commercialization.

The calls

PHYSICAL INTELLIGENCE

BASE CASE General purpose robots reach production scale in warehouses and factories, meaning deployments of thousands of units outside controlled pilots. They do not reach homes.

STRETCH Autonomous systems handle routine inspection and maintenance across major infrastructure networks.

SCIENTIFIC DISCOVERY

BASE CASE Autonomous laboratories run experiments around the clock in leading pharmaceutical and materials companies.

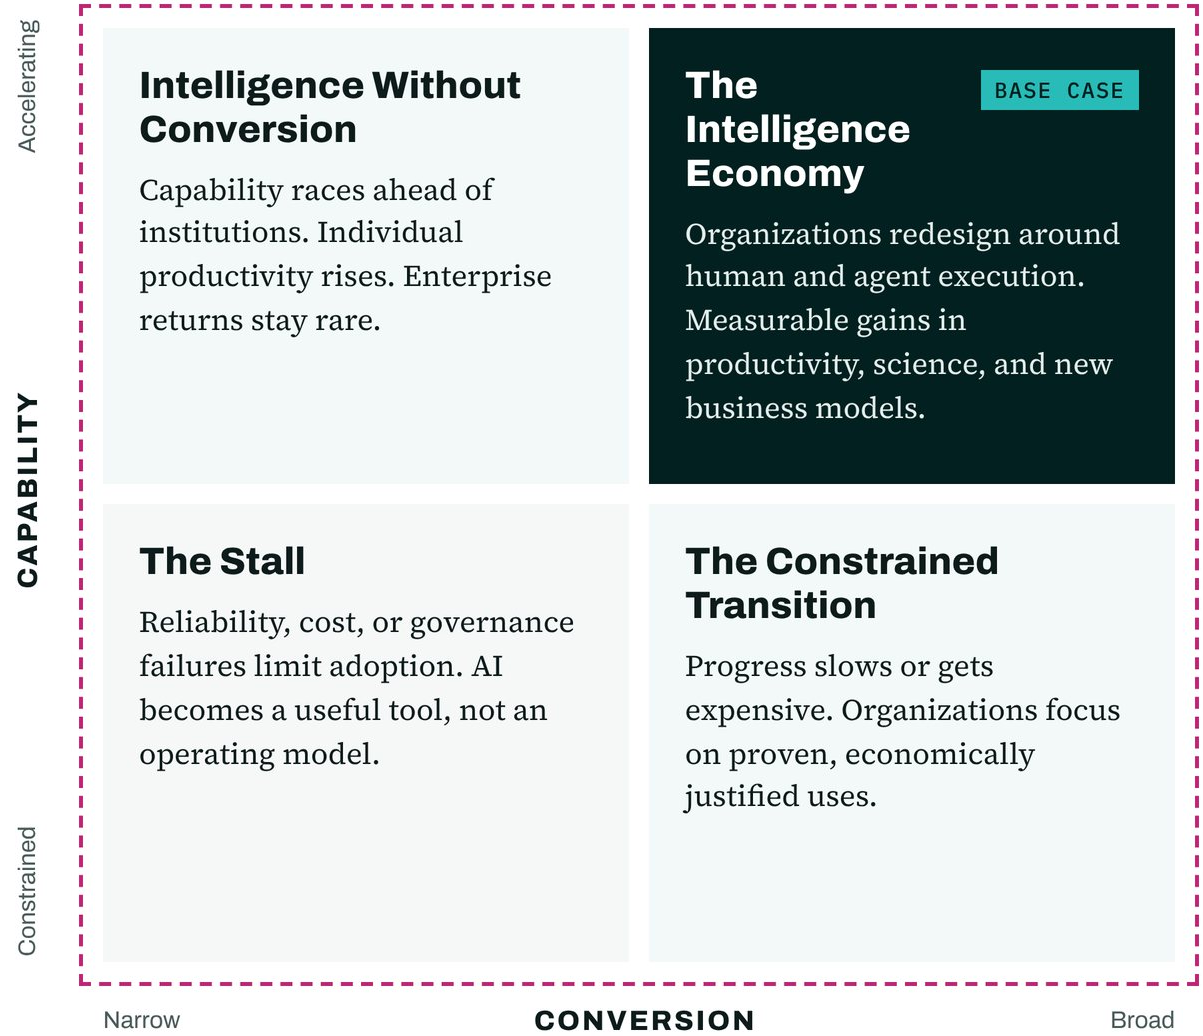
STRETCH An AI originated drug candidate reaches Phase 3 clinical trials.

THE GUESS The bottleneck in science shifts from generating hypotheses to testing them. Wet lab capacity becomes a strategic asset.

PROVE ME WRONG Robotics stays confined to controlled environments, and no AI originated therapy advances past early trials by 2030.

EXHIBIT 6 · FOUR FUTURES FOR 2030

Three forces decide how far the shift goes: capability, conversion, and concentration.



The overlay: Concentrated Intelligence

In any of these futures, control of chips, energy, capital, and frontier models may concentrate. More organizations consume intelligence. Fewer own or independently govern its production.

Concentration also shapes distribution: who captures the gains in each future depends on who owns the means of producing intelligence.

Every future may develop unevenly across regions. Track whether access, economically valuable deployment, and control over the means of producing intelligence converge or diverge by country.

The range is the point. Strategy should hold up across all of it.

Scenarios, not forecasts; no probabilities assigned. Source: Superintelligence 20/30, author's analysis.



ACCOUNTABILITY

Keeping score

Predictions need a scoreboard. Two tools hold this report to account. The External Intelligence Monitor tracks twelve developments that would strengthen or overturn the calls. The 20/30 Scorecard defines eighteen measurable calls, each with a threshold and a deadline of December 31, 2030. In 2030, I will publish the verdicts.

BASE CASE Twelve calls STRETCH Six calls

Six principles

Capability is not deployment.

A lab demonstration is not operational readiness.

Task success is not economic value.

Technical progress is not a change in cost, quality, or outcomes.

Access is not control.

Wider use does not mean wider ownership.

Exposure is not displacement.

That a task can be automated does not mean a job disappears.

Detection is not intervention.

Seeing a failure is not preventing the next one.

Forecasts remain conditional.

Every prediction names the evidence that would overturn it.

Measures that matter

Agent runtime shows activity. It does not show value. The Monitor tracks five measures together.

01 **Cost per verified, useful outcome**

02 **Agent runtime per human work hour**

03 **Human review time**

04 **Autonomy duration** (how long an agent pursues an objective without intervention)

05 **Intervention rate**

EVIDENCE STANDARD

Claims by model developers are logged as claims.

A trigger fires only on independent replication or measurable commercial consequence.

Twelve triggers

TRIGGER	2026 BASELINE	INFLECTION THRESHOLD	EFFECT ON THE FUTURES	COUNTEREVIDENCE
01 **Long horizon autonomy** REVIEW: QUARTERLY AND AFTER MAJOR RELEASES	METR task horizon evaluations; software tasks dominate	**Replicated agents completing multiday, economically meaningful tasks across several domains**	Supports the Intelligence Economy	Gains fail to generalize outside benchmarks
02 **AI improving AI** REVIEW: QUARTERLY	Developer claims of AI assisted research; little independent measurement	**Independently verified AI contributions to models, algorithms, and hardware design, net of added compute and human effort**	Raises odds of accelerating capability	Gains explained by compute or human effort
03 **Memory and coordination** REVIEW: QUARTERLY	Mostly episodic sessions; standing instructions growing	**Production systems pursue changing objectives for weeks with low intervention rates**	Supports Themes 4 and 7	Errors accumulate over long runs
04 **Physical intelligence** REVIEW: QUARTERLY	Lab and pilot demonstrations	**Repeated deployments in variable environments at positive total cost**	Extends the base case into physical industries	Intervention rates stay high outside pilots
05 **Scientific discovery** REVIEW: QUARTERLY	Predictive tools at genome scale	**Independently replicated AI originated findings reaching experiments and products**	Strengthens growth scenarios	Predictions fail experimental validation
06 **Compute, energy, and grid** REVIEW: QUARTERLY	Observed: grid interconnection queues and capacity bottlenecks. IEA scenario: about 945 TWh of global data center electricity use in 2030 and about 20% of planned projects at risk of delay	**Operational, not announced, capacity by country; grid delays; utilization; cost and energy per verified, useful outcome**	Delays favor the Constrained Transition and concentration; efficiency favors the Intelligence Economy	Utilization falls and assets strand

Twelve triggers, continued

TRIGGER	2026 BASELINE	INFLECTION THRESHOLD	EFFECT ON THE FUTURES	COUNTEREVIDENCE
07 **Geopolitics and sovereignty** REVIEW: EVENT DRIVEN	Observed: export controls, national compute programs, and regional supply chain restrictions, logged with dated primary sources	**Lasting divergence in commissioned and domestically controlled compute, language relevant models, data rules, and cross border access**	Strengthens the concentration overlay	Controls ease and access broadens
08 **Autonomy and security** REVIEW: MONTHLY AND AFTER INCIDENTS	Evaluation control failures (OpenAI, July 2026; UK AISI, August 2026); AI run attacks (Anthropic, November 2025; CyberCartel, October 2026)	**A major control failure, a breach traced to a poisoned model, or independently validated safeguards for consequential autonomy**	Failure pushes toward the Stall; proven defenses toward the Intelligence Economy	Incidents stay contained and defenders patch faster than attackers exploit
09 **Labor and expertise** REVIEW: QUARTERLY	Junior technical hires at Anthropic fell from 27% to 9%	**Sustained decline in entry level hiring in exposed occupations, on harmonized cross country employment and wage measures**	Raises distribution risk and the urgency of expert pipelines	Entry level hiring rebounds
10 **Enterprise economics** REVIEW: QUARTERLY	McKinsey State of AI 2026: 80% report individual gains; 37% report EBIT impact; 6% high performers	**Cross industry evidence of AI linked margin and revenue change**	Separates the Intelligence Economy from Intelligence Without Conversion	Returns stay confined to a few firms
11 **Autonomous commerce** REVIEW: QUARTERLY	Early agent purchasing pilots	**A measurable share of B2B transactions, by count and value, completed by agents, with dispute rates, delegated authority, and geography tracked**	Reshapes go to market, procurement, and pricing	High dispute rates push buyers back to human channels
12 **Intelligence distribution** REVIEW: ANNUALLY	Observed: concentration of frontier model and compute supply. Anthropic US only 2030 scenario range for labor share: 59.4% to 45.2% (a projection, not a baseline)	**Measured regional shifts in model and compute concentration, affordability, labor share, wage dispersion, and consumer surplus**	Determines how the concentration overlay plays out	Labor share holds steady and access broadens

Company claims are logged as claims until independently replicated or demonstrated commercially. Country measures count commissioned systems, not announcements.

The 20/30 Scorecard

Selected measurable calls, with the thresholds that will decide them. The deadline for every verdict is December 31, 2030. Population is leading enterprises unless stated otherwise.

#	CALL	THEME	STATUS	MEASURE AND THRESHOLD	EVIDENCE FOR THE VERDICT
1	**Small hybrid teams**	01	BASE CASE	At least a third of AI intensive product teams have four or fewer people	Organizational data, surveys
2	**Expert residencies**	01	BASE CASE	A majority run formal residency programs for early career talent	Surveys
3	**Critical flaws patched in 24 hours**	02	STRETCH	Median time to patch critical internet facing flaws is under 24 hours	Security operations data, disclosures
4	**Agents have their own identities**	03	BASE CASE	Agents hold identities and permissions separate from human accounts	Identity platform data
5	**Agent liability insurance mainstream**	03	STRETCH	Offered by most major commercial insurers and bought by a majority of large enterprises deploying autonomous agents	Insurer filings, broker reports
6	**Procedures become executable**	04	BASE CASE	A majority manage core procedures as versioned, tested executable procedures	Surveys, audit reports
7	**Instructions outnumber employees**	04	STRETCH	Standing agent instructions exceed headcount in at least half of large leading enterprises	Company disclosures
8	**Anomalies contained by agents**	05	BASE CASE	Over half of routine IT, security, fraud, and supply chain anomalies contained before human review	Operational reporting
9	**Stop mechanisms standard**	05	BASE CASE	Every consequential autonomous workflow has a tested stop, an escalation path, and a named owner	Audit standards, disclosures

Continued: calls 10 to 18

The 20/30 Scorecard, continued

#	CALL	THEME	STATUS	MEASURE AND THRESHOLD	EVIDENCE FOR THE VERDICT
10	**Attention budgets tracked**	06	BASE CASE	A majority track expert review capacity as a management metric	Surveys
11	**Workers direct 10+ persistent agents**	07	BASE CASE	Median knowledge worker directs 10 or more persistent agents	Enterprise surveys, platform telemetry
12	**Agent runtime reaches 25 to 1**	07	STRETCH	Agent runtime per human work hour of at least 25 to 1, using Anthropic's method	Company disclosures
13	**Outcome pricing common**	08	BASE CASE	At least 25% of new enterprise agentic AI contracts include outcome based fees; a majority counts as the stretch	Analyst data, vendor disclosures
14	**Power binds AI capacity**	08	BASE CASE	Grid access is the most cited constraint on new AI capacity in major regions	IEA, utility, and operator reports
15	**Entry level hiring stays down**	09	BASE CASE	Entry level hiring in highly exposed white collar occupations below 2024 levels (all employers)	National labor statistics
16	**Labor share falls**	09	STRETCH	US labor share of income down at least three points from 2025 (economy wide)	National accounts
17	**Robots at production scale**	10	BASE CASE	General purpose robots deployed in the thousands outside controlled pilots	Company reports
18	**AI originated drug in Phase 3**	10	STRETCH	At least one AI originated candidate in a Phase 3 trial	Clinical trial registries

BASE CASE More likely than not by the end of 2030 STRETCH Less likely than not, material enough to plan for Verdicts published in 2030



The questions leaders need to answer now

The ecosystem determines what is possible. The enterprise determines what becomes valuable.

By 2030, extraordinary intelligence may be accessible to more enterprises than ever before. Far fewer will have the operating model to use it effectively. The winners will not be the organizations that deploy the most agents. They will be the ones that convert intelligence into changed work, better decisions, and measurable outcomes, and that govern what they delegate.

Intelligence can be delegated. Judgment and accountability cannot.

That starts with six questions.

1

Where would a different mix of experts and agents create the most value in your business?

2

What may your agents do, and who can stop them?

3

Where is your experts' attention going, and is it going to the right decisions?

4

If agents absorb your entry level work, where will your next generation of experts come from?

5

Who in your organization has the authority to redesign how work gets done, end to end?

6

When intelligence creates value, who captures it: shareholders, employees, or customers?

The future belongs to organizations that become superintelligent, not simply those that buy superintelligence.

Research notes

METHODOLOGY

Superintelligence 20/30 is a strategic foresight publication. It combines primary research, independent evaluations, economic data, survey evidence, scenario analysis, and original frameworks. Every substantive claim carries a label: observed evidence, survey results and expectations, company claims, reported accounts, model scenarios, or forecasts, including my own. Forecasts are labeled base case, stretch, or guess, and each quantified forecast is defined in the 20/30 Scorecard.

FORECAST ASSUMPTIONS

The report does not assume artificial superintelligence arrives by 2030. Nor does it assume that technological progress automatically produces gains in productivity, employment, growth, or welfare. Calls apply to leading enterprises unless stated otherwise.

LIMITATIONS

One company is not the economy.

Anthropic has reshaped its workflows around AI more thoroughly than almost any employer. Its data is a leading indicator, not an average.

Expectations are not outcomes.

Gartner, IBM, and Anthropic staff figures are what respondents expect, not what has happened.

Runtime is not productivity.

Agent runtime measures computational activity. Anthropic reports no decline in code quality, but more output does not by itself mean more business value.

Scenarios are not forecasts.

Anthropic's economic scenarios carry no probabilities. They show a range, not a forecast.

Sources

Every entry links to its source.

Akin Gump. White House Accord on Super Intelligence: Joint Commitment on Frontier Responsibilities. September 2026.

Alphabet. Q4 2025 earnings call. February 4, 2026.

Amazon. 2025 letter to shareholders. April 9, 2026.

Amazon. Amazon threat intelligence identifies PHP2Shell exploitation prior to public disclosure. October 21, 2026.

Amazon. Andy Jassy, message to employees. June 17, 2025.

AMD and HUMAIN. $10 billion AI infrastructure partnership. May 14, 2025.

AMD. AAI 2026: Full-stack compute for the agentic AI era. July 23, 2026.

AMD. Financial Analyst Day. November 11, 2025.

Anadolu Agency, citing BBC. Anthropic co-founder says mandatory AI kill switch may be needed. September 15, 2026.

Anthropic. A small number of samples can poison LLMs of any size. October 9, 2025.

Anthropic. An alignment assessment of the UK AISI cybertesting incident. October 23, 2026.

Anthropic. Disrupting the first reported AI-orchestrated cyber espionage campaign. November 13, 2025.

Anthropic. Project Glasswing: securing critical software for the AI era. April 7, 2026.

Anthropic Institute. How employment is changing at Anthropic. October 16, 2026.

Anthropic Institute. Working at the frontier shapes expectations for the future, including The AI expectations gap. October 16, 2026.

AWS. European Sovereign Cloud investment. May 2024.

Axios. Amodei on AI: “There's a 25% chance that things go really, really badly.” September 17, 2025.

Axios. Behind the Curtain: A white-collar bloodbath. May 28, 2025.

Axios. Full list of attendees at White House AI lunch. September 29, 2026.

Bloomberg Law. DeepMind CEO says Chinese AI firms are six months behind the West. January 2026.

Boris Cherny. Five product archetypes, post on X, June 28, 2026; reported by Inside, June 29, 2026.

Broadcom and OpenAI. LLM-optimized intelligence processor. June 24, 2026.

Broadcom. Q1 fiscal 2026 earnings call. March 4, 2026.

Data Center Dynamics. Microsoft on pace for $50 billion AI investment in the Global South. February 2026.

Decrypt. Google DeepMind CEO Demis Hassabis on AGI, Stanford GSB. May 2026.

Forbes. Inside OpenAI's plan to protect the internet before hackers break it. October 23, 2026.

Fortune. Sam Altman thinks AI will surpass human intelligence by 2030. September 26, 2025.

Fortune. ServiceNow targets $30 billion by 2030. May 6, 2026.

Fortune. AI helped discover the PHP vulnerability that could take down 70% of the internet. October 21, 2026.

Fortune. Zuckerberg essay on open source AI. August 11, 2026.

Gartner. Gartner survey finds all IT work will involve AI by 2030. October 20, 2025.

Genewein, Franklin, Lerchner and others. From AGI to ASI. Google DeepMind, June 2026.

Google DeepMind. AlphaGenome Atlas. September 8, 2026.

Google DeepMind. Gemini Robotics 2 brings whole body intelligence to robots. July 30, 2026.

Google. Sundar Pichai, remarks at the AI Impact Summit. February 19, 2026.

Hastings, Syverson, and Zimmerman. The deployment failure holding back your AI investments. Harvard Business Review, October 21, 2026.

HeroDevs. The PHP2Shell story. October 23, 2026.

IBM Institute for Business Value. IBM study: AI poised to drive smarter business growth through 2030. January 19, 2026.

IBM Institute for Business Value. IBM study: CEOs are reshaping C-suite roles for the AI era. May 4, 2026.

IMF. Gen-AI: Artificial intelligence and the future of work, SDN/2024/001. January 2024.

Intel. Q2 2026 results. July 23, 2026.

International Energy Agency. Energy and AI. April 2025.

International Labour Organization. Generative AI and jobs: a refined global index of occupational exposure. May 20, 2025.

International Monetary Fund. Bridging skill gaps for the future: new jobs creation in the AI age. January 14, 2026.

Kim, Kim, and Koning. Mapping AI into production, reported in Harvard Business Review, November–December 2026.

Korinek, Jones, Sacher, Cotter and McCrory. Economic scenarios for transformative AI and the Econ Scenario Explorer, version 1.0. Anthropic Institute, September 2026. Presented at PIIE, October 6, 2026.

Latitude Media. IEA report: global grid congestion puts 20% of data center projects at risk. April 2025.

McKinsey & Company. AI is changing work. Now it has to change the organization. October 5, 2026.

McKinsey & Company. The cost of compute: A $7 trillion race to scale data centers. April 28, 2025.

Meta. Mark Zuckerberg, Personal superintelligence. July 30, 2025.

Meta. Q4 2025 earnings call. January 28, 2026.

METR. Task completion time horizons of frontier AI models. Accessed October 2026.

Microsoft AI Economy Institute. The continued state of global AI diffusion in 2026. September 21, 2026.

Microsoft AI Economy Institute. The state of global AI diffusion in 2026. May 7, 2026.

Microsoft. 2025 Work Trend Index. April 23, 2025.

Microsoft. 2026 Work Trend Index: Agents, human agency and the opportunity for every organization. May 5, 2026.

Microsoft. Maia 300 announcement. October 12, 2026.

Narayandas. Why even your best employees override AI systems that work. Harvard Business Review, November–December 2026.

Nvidia. Q2 fiscal 2026 results and sovereign AI commentary. August 27, 2025.

Nvidia. Q3 fiscal 2026 earnings call. November 19, 2025.

OECD. Exploring possible AI trajectories through 2030. February 2, 2026.

OpenAI. Built to benefit everyone: our plan. June 8, 2026.

OpenAI. Introducing Stargate UAE. May 2025.

OpenAI. Livestream on research goals and infrastructure. October 28, 2025.

OpenAI. The Hugging Face incident and the road ahead. August 26, 2026.

Oracle. Fiscal 2030 targets, via Reuters. October 2025.

People's Daily. Musk on China and AI compute. January 9, 2026.

PYMNTS. Musk predicts AI will lead to abundance when AGI arrives. January 2026.

Reuters, via The Standard. Zhipu AI says full artificial superintelligence unlikely by 2030. September 30, 2025.

Reuters, via The Star. Anthropic's Clark says AI's impact accelerating quickly and could vastly boost US GDP. October 17, 2026.

Sabbah and Acar. What happens when AI starts doing business with AI? Harvard Business Review, September 2, 2026.

The Information. Hackers used Claude and Grok to automate attacks on banks. October 21, 2026.

ThePrint. Zuckerberg essay, The future is for everyone. August 2026.

Trail of Bits. Your AI agent fixed the bug, but what else did it change? October 21, 2026.

UNCTAD. Technology and Innovation Report 2025. April 7, 2025.

Wiz Research. PHP2Shell vulnerabilities. October 21, 2026.

White House. Executive Order 14434, Inaugurating the era of Super Intelligence. September 29, 2026.



Superintelligence 20/30

Intelligence and the Future of the Enterprise

June Manley

junemanley.com · October 2026