

Q2 2026 · COVERING APRIL TO JUNE 2026

The state of AI.

Our first edition is a comprehensive baseline: the flagship reports and standing frameworks that define where AI stands now, from the frontier labs, the major consulting firms, independent research and the governance bodies. Every figure is quoted from its primary source.

INDEPENDENT AND VENDOR-NEUTRAL EVERY FIGURE QUOTED FROM ITS PRIMARY SOURCE

UPDATED QUARTERLY

THIS QUARTER'S NEW RELEASES

The signals that moved

This first edition maps the full landscape, from this quarter's new flagship reports to the standing frameworks that anchor AI governance. From the next edition, each quarterly release focuses on the new flagship research of that quarter.

Stanford HAI	Organisational AI adoption reached 88%, and capability is now racing ahead of our ability to govern it.
Microsoft	Organisational factors drive more than twice the AI impact of individual mindset, 67% versus 32%.
BCG	74% of white-collar workers now use AI regularly, but a clear strategy adds 25 points of impact where tools alone add about 5.
Anthropic	86% of Claude users report productivity gains, and the heaviest adopters are the most optimistic about their careers.
KPMG	Multi-agent orchestration doubled to 18% in a single quarter, yet only 26% of leaders can see what their AI costs to run.
MIT CISR	34% of enterprises now run AI digital colleagues, but only 22% have done the workflow redesign that captures the value.

Verify before you cite. Every figure below is quoted from its primary source and linked. Entries are summarised for brevity, and sources may revise their reports after publication. Always open the linked source and verify a figure before you cite it or act on it.

CROSS-SOURCE SYNTHESIS

Four things every major source is converging on

Where consulting firms, frontier labs and independent research agree independently, the signal is strong. Each theme below is evidenced on the source pages that follow.

01

Governance and cost accountability are the number one gate

Security, oversight and now the economics of AI, not regulation or technology, gate the move to agentic AI.

KPMG: only 26% can see what their AI costs to run. CEO-accountable firms realise value 57% versus 21%.

KPMG, MCKINSEY, EY, DELOITTE

02

Adoption is near universal, value is not

A divide separates the few who capture returns from the many who stall.

Stanford: 88% adoption. PwC: 20% of firms capture 74% of AI-driven returns. MIT CISR: 34% adopt, 22% redesign.

STANFORD, PWC, MIT CISR, MCKINSEY

03

The bottleneck is the operating model, not the tools

Strategy, workflow redesign and enablement, not tool choice, decide who wins.

BCG: clear strategy adds 25 points of impact. Microsoft: org factors drive 2x the impact of individuals.

BCG, MICROSOFT, MCKINSEY

04

Autonomy is climbing in months, not years

Agentic capability is compounding fast, and worker sentiment is splitting as it does.

METR: time horizon doubling every 131 days. UK AISI: cyber capability doubling every 4.7 months. Anthropic: heavy adopters most optimistic, juniors most exposed.

METR, UK AISI, ANTHROPIC, MICROSOFT

WHAT THIS MEANS FOR YOU

The market keeps proving the same point: the AI gap is organisational, not technical. Maturity, governance and enablement, not which tool you buy, separate the firms that capture value from the ones that stall.

CONSULTING FIRM

McKinsey

The State of Organizations 2026MARCH 2026 • PDF, 74PP • [VIEW SOURCE](#) ↗**88% / 81%**

experiment with AI, but report no meaningful bottom-line gain

under 20%

of adopters see significant tangible impact

\$5 : \$1

the advised spend ratio, people to technology

Key findings each figure quoted from its primary source

AI use is near universal but value is rare. 88% are experimenting, 81% report no meaningful bottom-line gains.

[The State of Organizations 2026](#), p.3 · “While 88 percent of organizations are now experimenting with AI, 81 percent do not report any meaningful bottom-line gains.”

Fewer than one in five adopters see real impact, a maturity gap, not a tooling gap.

[Same report](#), p.7 · “Less than 20 percent of companies that have tried to adopt the technology have seen significant tangible impact on their bottom lines.”

Value capture is people-led. The rule of thumb is \$5 on people for every \$1 on technology.

[Same report](#), p.8 · “for every \$1 spent on technology, \$5 should be spent on people.”

Caveat. Self-published consultancy survey with a commercial incentive; leader perceptions, not audited outcomes. n=10,018, 15 countries, fielded June to September 2025.

CONSULTING FIRM

Deloitte

The State of AI in the Enterprise 2026: From Ambition to ActivationJANUARY 2026 • PDF • [VIEW SOURCE](#) ↗**+50%**

one-year jump in workers' access to sanctioned AI tools

21%

have a mature governance model for autonomous agents

34%

are using AI to deeply transform the business

Key findings each figure quoted from its primary source**Worker AI access jumped 50% in a year, from under 40% to about 60% of workers.**[State of AI in the Enterprise 2026](#), p.4 • “broadened worker access to AI by 50% in just one year—growing from fewer than 40% to around 60%”**Agent governance lags adoption. 74% plan agentic AI within two years, only 21% have a mature governance model.**[Same report](#), p.6 • “74% companies plan to deploy agentic AI within two years. Yet, 1 in 5 (21%) report having a mature model for governance”**Only 34% are truly reinventing the business; 37% remain surface-level.**[Same report](#), p.4 • “34% of companies are starting to use AI to deeply transform their businesses... remaining 37% are only using AI at a surface level”**Caveat.** Self-reported perception survey of director to C-suite leaders. n=3,235, 24 countries, August to September 2025.

CONSULTING FIRM

BCG

AI at Work 2026 (4th ed.): Why Strategy Matters More Than ToolsJUNE 2026 • REPORT • [VIEW SOURCE](#) ↗**74%**

of frontline workers are now regular AI users, up 20+ points in two years

+25pp

business-impact lift from a clear strategy, versus about 5 from tools alone

42%

of regular users save a full workday a week, but 66% get no guidance on it

Key findings each figure quoted from its primary source

Frontline AI adoption surged to 74% regular users, up more than 20 points in two years.

[BCG, AI at Work 2026](#), press release · “74% now regular users, up more than 20 percentage points over the previous two years”

Strategy beats tools. A clear strategy lifts measurable impact by 25 points, where better tools alone move it by about 5.

[Same release](#), Strategy vs tools · “Clear strategy lifts measurable business impact by 25 percentage points... better tools, without that strategy and redesign, move it only by approximately 5 points”

Time is saved but squandered. 42% save at least a workday a week, yet 66% get little or no guidance on using it.

[Same release](#), body · “42% of regular frontline users report saving at least a full workday... 66% report they get limited or no guidance”

Caveat. Figures cited from BCG's official press release. Survey of 11,749 workers across 14 markets.

CONSULTING FIRM

PwC

Want ROI from AI? Go for growth (AI performance study)APRIL 2026 • PDF, 30PP • [VIEW SOURCE](#) ↗**74%**

of all AI-driven returns captured by just the top 20% of firms

7.2x

the AI-driven financial performance of the most AI-fit firms versus others

2.5x

more of revenue that AI leaders invest in AI

Key findings each figure quoted from its primary source**AI value is concentrated. 20% of firms capture 74% of AI-driven returns.**[Want ROI from AI? Go for growth](#), p.3 • “20% of the 1,217 companies we surveyed capture 74% of the AI-driven returns”**The most AI-fit firms deliver 7.2 times the AI-driven financial performance of others.**[Same study](#), p.4 • “deliver AI-driven financial performance that's 7.2 times as high as the other respondents”**Leaders invest materially more, 2.5 times as much of revenue in AI.**[Same study](#), p.15 • “The leading companies... invest materially more in AI than other companies do: 2.5 times as much of their revenue”

Caveat. Self-reported exec survey; AI-driven performance is self-estimated, not audited; leaders are the top-20% cohort. n=1,217, 25 sectors, October to November 2025.

CONSULTING FIRM

KPMG

Global AI Pulse Survey, Q2 2026JUNE 2026 • QUARTERLY SURVEY • [VIEW SOURCE](#) ↗**9% to 18%**

organisations orchestrating multiple AI agents, doubled in one quarter

26%

the share with full, real-time visibility into what their AI costs to run

57% v 21%

realise meaningful value where the CEO is accountable for AI decisions

Key findings each figure quoted from its primary source**Agent deployment plateaued at 53% versus 55%, but multi-agent orchestration doubled from 9% to 18%.**[KPMG Global AI Pulse Q2 2026](#), US release · “the percentage of organizations orchestrating multiple AI agents across workflows doubled from 9% to 18%”**A cost-visibility gap is opening. Only 26% have full, real-time visibility into what their AI systems cost to operate.**[Same survey](#), US release · “only 26% report full, real-time visibility into what their AI systems cost to operate”**Accountability tracks value. Where the CEO owns AI decisions, firms far more often realise value, 57% versus 21%, and established ROI, 14% versus 4%.**[Same survey](#), global release · “more likely to realize meaningful business value (57 percent vs. 21 percent) and report established ROI (14 percent vs. 4 percent)”**Caveat.** A quarterly pulse of self-reported C-suite perceptions, and KPMG sells AI advisory. n=2,145 leaders, 20 markets, fielded 28 April to 25 May 2026.

CONSULTING FIRM

EY

2026 AI Sentiment Report (AI Sentiment Index Study)MARCH 2026 • [HTML](#) • [VIEW SOURCE](#) ↗**84%**

have used AI in the past six months

16%

have used AI that acts without any human intervention

66%

say human oversight of AI remains essential

Key findings each figure quoted from its primary source**AI use is near universal, 84% used AI in the past six months.**[2026 AI Sentiment Report](#), Adoption grows · “In the past six months alone, more than eight in 10 (84%) respondents used AI”**Autonomous AI is real but a minority, 16% use AI that acts without human intervention.**[Same report](#), opening findings · “16% globally report using AI systems that act without human intervention”**Demand for human control persists, 66% say human oversight remains essential.**[Same report](#), Adoption outpaces confidence · “Almost seven in 10 (66%) say human oversight remains essential”

Caveat. A consumer-sentiment study, n=18,152, 23 markets. A separate EY poll of 500 US tech execs found 78% say adoption is outpacing their ability to manage risk.

INDEPENDENT RESEARCH

Stanford HAI

The AI Index 2026 Annual Report (9th ed.)

APRIL 2026 • PDF, 425PP • [VIEW SOURCE](#) ↗**88%**

of organisations have now adopted AI

\$285.9BUS private AI investment in 2025, about
23x China's**53%**gen-AI population adoption in three years,
faster than the PC**Key findings** each figure quoted from its primary source

Capability is racing ahead. SWE-bench Verified rose from 60% to near 100% of the human baseline in a single year, as organisational adoption reached 88%.

[AI Index 2026](#), Top Takeaways, p.9 · “performance rose from 60% to near 100% of meeting the human baseline in a single year. Organizational adoption reached 88%”

Investment is concentrated in the US, at \$285.9B in 2025, more than 23 times China's.

[AI Index 2026](#), p.10 · “U.S. private AI investment reached \$285.9 billion in 2025, more than 23 times the \$12.4 billion invested in China”

Governance is lagging. Documented AI incidents rose to 362, up from 233 the year before.

[AI Index 2026](#), p.9 · “Documented AI incidents rose to 362, up from 233 in 2024”

Caveat. An independent academic report, 9th edition. Figures quoted from its own Top Takeaways.

INDEPENDENT RESEARCH

MIT CISR

Leveraging Digital Colleagues for Enterprise Value

APRIL 2026 • MIT SLOAN CISR BRIEFING • [VIEW SOURCE](#) ↗

34%

are actively using AI digital colleagues
alongside staff

+25%

predicted gain in revenue per employee
over three years

22%

have done the major workflow redesign that
captures the value

Key findings each figure quoted from its primary source

34% of enterprises are already using AI digital colleagues that work alongside staff.

[Leveraging Digital Colleagues for Enterprise Value](#), MIT Sloan CISR briefing · "Thirty-four percent of surveyed enterprises were actively using digital colleagues"

Expectations run high. 75% predict an average 25% rise in revenue per employee over three years.

[Same briefing](#), briefing · "75 percent predicted an average increase of 25 percent"

But value-unlocking redesign is rare. Only about 22% have done major workflow redesign.

[Same briefing](#), briefing · "Approximately twenty-two percent of the surveyed enterprises reported major workflow redesign."

Caveat. MIT CISR, Sloan. Self-reported predictions, n=132. Figures verified via the public preview.

FRONTIER LAB

OpenAI

GPT-5.5 System CardAPRIL 2026 • PDF, 45PP • [VIEW SOURCE](#) ↗**100% / 99.1%**

misalignment-monitor recall on internal sabotage benchmarks

96%

chain-of-thought monitorability, after a rare regression

+23%

more likely GPT-5.5's claims are factually correct than GPT-5.4

Key findings each figure quoted from its primary source**The misalignment monitor hits near-perfect recall, 100% and 99.1% on internal sabotage benchmarks.**[GPT-5.5 System Card](#), p.15, §7.2.1 • “near-perfect recall on OAI-repo-sabotage and BashArena (100.0% and 99.1%, respectively)”**Chain-of-thought monitorability stays high at 96% despite one genuine regression.**[Same card](#), p.16, §7.3.1 • “a genuine monitorability regression... the affected cases are rare and overall monitorability remains high at 96%”**Factuality improved, claims 23% more likely correct than GPT-5.4 on user-flagged cases.**[Same card](#), p.12, §6.1 • “GPT-5.5's individual claims are 23% more likely to be factually correct”**Caveat.** A safety and system card, not a capabilities report; headline benchmark scores live on marketing pages that block scripted access.

FRONTIER LAB

Anthropic

Anthropic Economic Index: Cadences

JUNE 2026 • RESEARCH REPORT • [VIEW SOURCE](#) ↗**86%**

of surveyed Claude users say AI improved the speed of their work

1 in 3+

expect AI to do most or nearly all of their work tasks within a year

>60%

the job-loss odds a third of users assign to a junior colleague next year

Key findings each figure quoted from its primary source

Large majorities report real productivity gains: 86% on speed, 82% on scope, 69% on quality.

Economic Index: Cadences, Perceived impact · “Large majorities of people report productivity gains in speed, scope, and quality of their work (86%, 82%, and 69%, respectively)”

Delegation tracks optimism, not fear. Heavier-automation users are more optimistic on every one of six job dimensions.

Same report, Automation and optimism · “Across all six dimensions, people with a higher share of automated sessions feel more optimistic about the effect of AI on their job outcomes next year”

The anxiety lands on juniors. Over a third put a junior colleague's odds of job loss next year above 60%.

Same report, Perceived impact · “Respondents were especially worried about job loss for their junior colleagues, with over one third stating that the probability of a junior colleague losing their job in the next year was over 60%”

Caveat. Descriptive analytics from Anthropic's own product, on self-selected Claude users linked to a survey, mid-May to early June 2026. Perceptions, not a representative workforce sample.

FRONTIER LAB

Google DeepMind

Gemini 3.1 Pro, Model Card

FEBRUARY 2026 • HTML • [VIEW SOURCE](#)

77.1%

ARC-AGI-2 abstract-reasoning puzzles
(ARC Prize Verified)

94.3%

GPQA Diamond graduate-level science, no
tools

80.6%

SWE-Bench Verified agentic coding, single
attempt

Key findings each figure quoted from its primary source

Strong abstract reasoning, 77.1% on ARC-AGI-2 (ARC Prize Verified).

[Gemini 3.1 Pro Model Card](#), Results table • “ARC-AGI-2 Abstract reasoning puzzles | ARC Prize Verified | 77.1%”

Near-ceiling on graduate science, 94.3% on GPQA Diamond with no tools.

[Same card](#), Results • “GPQA Diamond Scientific knowledge | No tools | 94.3%”

Top-tier agentic coding, 80.6% on SWE-Bench Verified on a single attempt.

[Same card](#), Results • “SWE-Bench Verified Agentic coding | Single attempt | 80.6%”

Caveat. Vendor-published model card, not peer-reviewed; benchmark conditions vary by row. The comparison column is the prior Gemini 3 Pro.

FRONTIER LAB

Meta AI

Introducing Muse Spark: Scaling Towards Personal Superintelligence

APRIL 2026 • [HTML](#) • [VIEW SOURCE](#) ↗**58%**

Humanity's Last Exam, in Muse Spark's contemplating mode

>10X

less pretraining compute than Llama 4 Maverick, same capability

52

Artificial Analysis Intelligence Index, ranked #1 of 220 in class

Key findings each figure quoted from its primary source**Meta's new flagship Muse Spark scores 58% on Humanity's Last Exam and 38% on FrontierScience.**[Introducing Muse Spark](#), Capabilities · “achieving 58% in Humanity's Last Exam and 38% in FrontierScience Research”**A more than tenfold efficiency leap, the same capability at an order of magnitude less compute than Llama 4 Maverick.**[Same post](#), Scaling Axes · “the same capabilities with over an order of magnitude less compute than our previous model, Llama 4 Maverick”**Independent benchmarking ranks it first of 220 in class, Index 52 versus about 14 median.**[Artificial Analysis, Muse Spark](#), Intelligence Index · “scores 52 on the Artificial Analysis Intelligence Index... well above average among comparable models (averaging 14)”**Caveat.** Meta's first non-open-weights flagship; specs and weights are not inspectable; the Index score is third-party, not Meta's. Supersedes Llama 4.

FRONTIER LAB

Microsoft

2026 Work Trend Index: Agents, Human AgencyMAY 2026 • [REPORT](#) • [VIEW SOURCE](#) ↗**67% v 32%**share of AI impact from organisational
versus individual factors**15x**year-on-year growth in active Microsoft 365
agents, 18x in large firms**58%**produce work they could not a year ago,
rising to 80% among advanced users**Key findings** each figure quoted from its primary source**Organisational factors drive more than twice the AI impact of individual mindset, 67% versus 32%.**[2026 Work Trend Index](#), p.17 · “organizational factors like culture, manager support, and talent practices account for more than 2x the reported AI impact of individual factors (67% vs. 32%)”**Agents are scaling fast. Active agents in the Microsoft 365 ecosystem grew 15x year on year, and 18x in large enterprises.**[Same report](#), p.17 · “The number of active agents in the Microsoft 365 ecosystem has grown 15x year over year, rising to 18x in large enterprises”**A capability gap is opening. 58% say they produce work they could not a year ago, rising to 80% among the most advanced users.**[Same report](#), p.8 · “58% say they're producing work they couldn't have a year ago. That rises to 80% among Frontier Professionals”**Caveat.** Microsoft sells Copilot. Figures are self-reported perceptions plus telemetry ratios, association not causation. n=20,000, 10 markets.

FRONTIER LAB

DeepSeek

mHC: Manifold-Constrained Hyper-Connections

JANUARY 2026 • ARXIV, PDF • [VIEW SOURCE ↗](#)**6.7%**training overhead of the mHC architecture
at scale**27B / 9B / 3B**MoE model sizes the method was validated
on**+2.1% / +2.3%**reasoning gains over hyper-connections on
BBH and DROP**Key findings** each figure quoted from its primary source**A cheaper-scaling architecture, mHC adds only 6.7% training overhead at scale.**

[mHC: Manifold-Constrained Hyper-Connections](#), p.9, §4.3 · “we implement mHC (with $n = 4$) in large-scale models with a marginal training overhead of only 6.7%”

It stabilises training, 0.021 lower final loss versus baseline on a 27B model.

[Same paper](#), p.12, §5.2 · “achieving a final loss reduction of 0.021 compared to the baseline”

Reasoning gains over hyper-connections, 2.1% on BBH and 2.3% on DROP.

[Same paper](#), p.13, §5.2 · “performance gains of 2.1% on BBH... and 2.3% on DROP”

Caveat. A training-method research paper, not a shipped model; gains are relative. The DeepSeek-V4 series came later, April 2026.

CAPABILITY & SAFETY

Epoch AI

Machine Learning Trends (dashboard)

UPDATED JUNE 2026 • [HTML](#) • [VIEW SOURCE](#) ↗**5x / yr**

growth in frontier training compute since 2020

37% / yr

improvement in AI chip performance per dollar

3,500+

notable ML models tracked in Epoch's open database

Key findings each figure quoted from its primary source**Compute is still exploding, frontier training compute grows about 5 times per year since 2020.**[Machine Learning Trends](#), Training Runs · "Training compute for frontier language models has been growing at 5× per year since 2020"**Efficiency improves, chip performance-per-dollar up about 37% per year.**[Same dashboard](#), Hardware · "AI chip performance per dollar has improved by 37% per year"**Benchmarks are not saturated, leading models solved under 2% of research-level FrontierMath.**[FrontierMath](#), Current Performance · "none could solve more than 2% of the problems"

Caveat. Point-in-time aggregates Epoch updates continually; the under-2% FrontierMath figure is the late-2024 baseline and newer reasoning models score higher.

CAPABILITY & SAFETY

METR

AI Task-Completion Time Horizon: Time Horizon 1.1

JANUARY 2026 • HTML • [VIEW SOURCE](#) ↗**131 days**

post-2023 doubling of AI's autonomous time horizon

~5.3 h

the leading model's 50%-completion time horizon

<10%

success rate on tasks that take humans over about 4 hours

Key findings each figure quoted from its primary source**AI's autonomous time horizon has been doubling about every 7 months.**[Measuring AI Ability to Complete Long Tasks](#), Summary · “consistently exponentially increasing over the past 6 years, with a doubling time of around 7 months”**The trend has accelerated, the post-2023 doubling is now 131 days.**[Time Horizon 1.1](#), TH1 vs TH1.1 · “The post-2023 doubling-time is 131 days under TH1.1”**Reliability is jagged, near 100% on sub-4-minute tasks, under 10% on tasks over about 4 hours.**[Long Tasks](#), Summary · “almost 100% success rate on tasks taking... less than 4 minutes, but succeed <10% of the time on tasks taking more than around 4 hours”**Caveat.** Measured on a software and agentic task suite with noisy human-time baselines; absolute hours are approximate, not real-world autonomy guarantees.

CAPABILITY & SAFETY

UK AI Security Institute

How fast is autonomous AI cyber capability advancing?

MAY 2026 • AISI, HTML • [VIEW SOURCE](#) ↗**4.7 mo**

doubling of the 80%-reliability cyber time horizon

~half

of AISI's own November 2025 estimate, which was 8 months

4.2 mo

METR's software-task doubling, which independently corroborates

Key findings each figure quoted from its primary source**Autonomous cyber capability is doubling roughly every 4.7 months at 80% reliability since late 2024.**

[How fast is autonomous AI cyber capability advancing?](#), Cyber Time Horizons · “we estimated that frontier models’ 80%-reliability cyber time horizon had doubled every 4.7 months since reasoning models emerged in late 2024”

The pace is accelerating, about half the doubling time AISI estimated three months earlier.

[Same post](#), Cyber Time Horizons · “This was around half our November 2025 doubling time estimate, which was 8 months”

Independent software-task data corroborates a 4.2-month doubling.

[Same post](#), Further Evidence · “a consistent doubling time of 4.2 months on software tasks since late 2024”

Caveat. An AISI trends update (blog), not a formal report; the 4.7-month figure is AISI's own February 2026 estimate and it flags uncertainty over whether the newest models mark a sustained faster rate.

CAPABILITY & SAFETY

International AI Safety Report

International AI Safety Report 2026

FEBRUARY 2026 • PDF, 221PP • [VIEW SOURCE](#) ↗**700M**

people now use leading AI systems every week

12

firms published or updated frontier safety frameworks in 2025

100+

independent expert authors, chaired by Yoshua Bengio

Key findings each figure quoted from its primary source**Adoption is rapid but uneven, at least 700 million weekly users, faster than the PC.**[International AI Safety Report 2026](#), p.10 · “at least 700 million people now using leading AI systems weekly”**Capability keeps climbing but stays jagged, agents now do about 30-minute coding tasks, up from under 10 a year ago.**[Same report](#), p.10 · “AI agents can now reliably complete some tasks that would take a human programmer about half an hour, up from under 10 minutes a year ago”**Safety governance is growing but stays largely voluntary, 12 firms published or updated frontier safety frameworks in 2025.**[Same report](#), p.13 · “In 2025, 12 companies published or updated their Frontier AI Safety Frameworks”**Caveat.** A scientific synthesis with evidence pre-December 2025, not original research; it stresses an evidence dilemma over uncertain risks. Chaired by Yoshua Bengio.

GOVERNANCE & STANDARDS

NIST AI RMF

AI Risk Management Framework (AI RMF 1.0), NIST AI 100-1

JANUARY 2023, STANDING REFERENCE • PDF • VIEW SOURCE ↗

4

core functions: Govern, Map, Measure, Manage

7

characteristics of trustworthy AI it defines

Voluntary

non-binding by design, guidance not regulation

Key findings each figure quoted from its primary source**Four core functions structure AI risk: Govern, Map, Measure, Manage.**[AI RMF 1.0 \(NIST AI 100-1\)](#), p.3 · “These functions — GOVERN, MAP, MEASURE, and MANAGE — are broken down further into categories and subcategories”**It names seven characteristics of trustworthy AI.**[Same framework](#), p.12 · “valid and reliable, safe, secure and resilient, accountable and transparent, explainable... privacy-enhanced, and fair”**Governance is cross-cutting, infused through the other three functions.**[Same framework](#), p.ii, Fig.5 · “Governance is designed to be a cross-cutting function to inform and be infused throughout the other three functions”

Caveat. A standing reference, not re-published annually. Voluntary guidance, not regulation; it prescribes no metrics or thresholds. Released January 2023; GenAI Profile July 2024.

GOVERNANCE & STANDARDS

ISO/IEC 42001

ISO/IEC 42001:2023, AI Management System

DECEMBER 2023, [STANDING REFERENCE](#) • [STANDARD](#) • [VIEW SOURCE](#) ↗**1st**

the world's first AI management-system
(AIMS) standard

Certifiable

organisations can be independently audited
against it

Any org

applies regardless of size, type or sector

Key findings each figure quoted from its primary source

The world's first certifiable AI management system (AIMS) standard.

[ISO/IEC 42001:2023 \(Abstract\)](#), Abstract · “establishing, implementing, maintaining and continually improving an AI... management system within... an organization”

Built to drive responsible AI in pursuit of business objectives.

[Same standard](#), Abstract · “intended to help the organization develop, provide or use AI systems responsibly”

Universal scope, any organisation regardless of size or sector.

[Same standard](#), Scope · “applicable to any organization, regardless of size, type and nature, that provides or uses products or services that utilize AI”

Caveat. A standing reference, not re-published annually. Full text paywalled; quotes limited to the official abstract and scope. Certification is by accredited third parties, not ISO.

GOVERNANCE & STANDARDS

OECD.AI

OECD AI Principles and Policy Observatory

2019, REVISED MAY 2024, STANDING REFERENCE • LEGAL INSTRUMENT AND OBSERVATORY • VIEW SOURCE ↗

47

national adherents to the Principles

5

values-based principles for trustworthy AI

2,348

national AI policy initiatives tracked by the Observatory

Key findings each figure quoted from its primary source**The first intergovernmental AI standard, adopted May 2019 and updated May 2024.**

[Recommendation of the Council on AI \(OECD/LEGAL/0449\)](#), overview • “adopted by the OECD Council... on 22 May 2019... the first intergovernmental standard on AI”

Five values-based principles define trustworthy AI.

[OECD AI Principles](#), Values-based principles • “Inclusive growth... Human rights and democratic values... Transparency... Robustness, security and safety... Accountability”

The Observatory tracks 2,348 policy initiatives across more than 80 jurisdictions.

[OECD.AI Policy Observatory](#), dashboard • “List of policy initiatives (2348)... more than 80 jurisdictions and organisations”

Caveat. A standing reference. The 2,348 figure is a live count and will drift; adherent phrasing varies across OECD pages.

ABOUT THIS BRIEF

How it is built

Each quarter, SP Optima sweeps a fixed watchlist of primary sources: the frontier labs, the major consulting firms, independent research and the governance bodies. We select the flagship reports published that quarter, verify every figure at its primary source, and quote it verbatim with a link. Where a figure cannot be verified at source, it is left out rather than repeated second-hand.

This is general market commentary, not advice. Each source carries its own methodology and, in some cases, a commercial incentive, noted in its caveat line. Always open the linked source and verify a figure before you cite it or act on it.

Verify before you cite. Every figure below is quoted from its primary source and linked. Entries are summarised for brevity, and sources may revise their reports after publication. Always open the linked source and verify a figure before you cite it or act on it.

SP OPTIMA

Independent, business-first AI advisory. We read the primary research so leaders can act on evidence, not hype. To talk it through: enquiries@sptoptima.com · sptoptima.com