

EXECUTIVE WHITE PAPER

March 2026

The Real State of Enterprise AI

What the Numbers Say, What Leadership Must Do

Basil C. Puglisi, MPA

Human-AI Collaboration Strategist & Responsible AI² Governance

basilpuglisi.com

Research conducted independently. No vendor sponsorship.

Produced under the HAIA-RECCLIN Human-AI Governance Methodology. #Alassisted

EXECUTIVE SUMMARY

More than seventy percent of enterprises now use AI in at least one function, and organizations citing that number as evidence of progress are measuring the wrong thing entirely, because adoption does not compound into value without the governance architecture that converts deployment into accountability and accountability into results. The independent evidence on what those investments have produced is convergent, sourced from three separate research bodies, and materially different from the narrative circulating in vendor reports and conference keynotes.

NBER Working Paper 34836, drawing on nearly 6,000 senior executives across four countries, found that nine out of ten report no measurable impact on their firm's employment or productivity after three years of AI adoption, despite two-thirds of those executives using AI tools personally. S&P Global Market Intelligence found that forty-two percent of organizations abandoned the majority of their AI initiatives before reaching production in 2025, more than double the rate of the prior year, with the average organization scrapping nearly half its proofs-of-concept before production. Workday research found that nearly forty percent of reported AI productivity gains are consumed by rework. These findings come from independent sources using independent methodologies and they point to the same structural failure: organizations are managing AI as a technology problem when the actual constraint is governance architecture.

This paper examines why that gap exists and what closes it. It addresses five dimensions in sequence: why agentic AI behaves more like capital than software and what that means for investment governance; why returns remain elusive for the majority of organizations even when the technology performs; why most enterprises stay locked in pilot experimentation rather than scaling to production; why individual productivity gains consistently fail to reach enterprise P&L; and two risks most boards have not yet named, one with a statutory compliance deadline in August 2026 and one with deployment costs that routinely run two to three times what business cases project. Each section closes with a specific governance action and a measurable outcome, because the argument throughout is not that AI requires more investment or better technology but that it requires named human accountability at the decision points that currently have none.

The five decisions the paper identifies are available to any organization now, require no new infrastructure, and can each produce measurable governance signals within a quarter. The organizations that close the gap between AI adoption and AI value will not be the ones that moved fastest. They will be the ones that governed most deliberately, and the evidence for that claim is already in the quarterly reviews of the organizations that have not.

SECTION 1

What the Numbers Actually Say

More than seventy percent of enterprises now use AI in at least one function, but organizations citing that number as evidence of progress are mistaking the starting line for the finish line. The independent evidence on what those organizations have to show for their investment tells a materially different story than the one circulating in vendor reports and conference keynotes.

Ninety percent of senior business executives across the United States, United Kingdom, Germany, and Australia report that AI has had no measurable impact on their firm's employment or productivity over the past three years, according to NBER Working Paper 34836, published February 2026, drawing on surveys of approximately 6,000 executives at firms already using AI. Two-thirds of those executives personally use AI tools, averaging 1.5 hours per week. The data does not describe laggards or skeptics; it describes the people who approved the budgets, and nine out of ten cannot yet point to a measurable firm-level result.

In the same period, forty-two percent of organizations abandoned the majority of their AI initiatives before reaching production, up from seventeen percent the prior year, according to S&P Global Market Intelligence's survey of more than 1,000 enterprise professionals across North America and Europe. The average organization scrapped forty-six percent of its AI proofs-of-concept before production. These are not outlier findings from a single study. They are the convergent self-reporting of the organizations making the investments, and the gap between what they describe and what the dominant narrative promises is not a rounding error. It is a structural problem, and understanding what drives it requires looking past the technology.

SECTION 2

The Classification Problem

The most consequential strategic shift in how to think about AI is not about what the technology can do. It is about what the technology is. Agentic AI systems, those that plan, decide, and execute rather than simply generate content, behave more like capital than software: they can be owned, compounded over time, and governed with the discipline applied to any long-lived productive asset. Organizations that have started treating AI investment through that lens are making different decisions and getting different results than those still managing AI as a technology subscription line item, and the divergence between the two groups is accelerating.

The regulatory environment is beginning to reflect this shift, even if incompletely. FASB issued Accounting Standards Update 2025-06 in September 2025, modernizing the framework for internally developed software costs and moving disclosure requirements into alignment with Property, Plant and Equipment rules. The standard stops well short of creating a formal AI capital class; FASB stated explicitly that capitalization practice will not significantly change for most software under the new rules, and mandatory compliance does not begin until December 2027.

The IASB has an active intangible assets project underway and in 2025 staff concluded that AI was not sufficiently different from other intangible assets to warrant a separate test case, meaning no new AI capital class is under active development. Academic literature is moving faster than the standards bodies, increasingly treating AI as a distinct factor of production alongside capital and labor, with characteristics that conventional accounting frameworks were not designed to capture: near-zero marginal replication cost, capacity for self-improvement, and scalability that does not require proportional labor input.

The legal codification is incomplete. The economic argument does not wait for it. Organizations that apply capital governance logic to AI now, which means naming owners, assigning binding accountability, building depreciation thinking into investment decisions, and treating AI capability as an asset to compound rather than a subscription to renew, are developing institutional muscle that regulatory requirements will eventually demand of everyone. Requiring every AI initiative to be classified explicitly as either an expense or a capital investment, with corresponding governance obligations tracked quarterly, is a decision available to any organization today and worth making before the regulatory mandate arrives, not after.

SECTION 3

Why Returns Remain Elusive

The mainstream narrative describes a two-to-four-year payback horizon for AI investment, which is already substantially longer than the seven-to-twelve months enterprises expect from conventional technology. Independent research suggests the two-to-four-year frame describes the successful minority, not the median case, and three structural factors explain why the majority remain outside it.

MIT's Project NANDA studied more than 300 publicly disclosed enterprise AI deployments alongside executive interviews and leadership surveys conducted between January and June 2025 and found that ninety-five percent of integrated AI pilots deliver no measurable P&L impact, with only five percent extracting significant financial value. The methodology is directional rather than definitive; the study sample is relatively small, the six-month measurement window may undercount returns that take longer to materialize, and independent analysts have challenged whether the underlying exhibit data clearly supports the 95% headline figure. Treat it as a directional signal consistent with the S&P Global and NBER data, not as a precise market measurement, but worth taking seriously as a description of the median case. The 95% figure frames the problem; the NBER and S&P Global data measure it.

The first structural factor is the rework tax. Workday research found that nearly forty percent of reported productivity gains are consumed by correcting or verifying AI-generated output. Productivity metrics that count output volume rather than net time saved systematically overstate AI's actual contribution to enterprise performance, and most organizations are measuring the wrong thing without knowing it.

The second is measurement misalignment. Organizations typically measure AI ROI at the use-case level, tracking a specific process, a particular workflow, a defined team, while the NBER data measures at the firm level. The gap between those two vantage points is exactly where AI value disappears, because a process can run faster without the enterprise P&L moving if the time saved is reabsorbed elsewhere rather than converted into revenue growth or cost reduction. Without a measurement framework that connects use-case efficiency to firm-level outcomes through a named person with accountability for the translation, organizations will continue generating impressive pilot reports and unimpressive income statements.

The third is investment misdirection. Research on enterprise AI deployments consistently finds that budgets concentrate in high-visibility applications, sales tools, customer-facing products, marketing automation, while the highest and most consistent ROI comes from back-office automation: eliminating outsourced processes, cutting external agency spend, streamlining operations where the cost base is controllable and the gains are directly measurable. The forcing question for any AI budget review is not which application is most visible to leadership. It is which workflow, if automated, would most directly reduce a controllable cost or compound a proven revenue stream, and who is named to own the measurement of whether it did.

SECTION 4

Why Most Organizations Stay Stuck

Pilots create slides. Transformation builds operating systems. That distinction is not about the quality of the technology in the pilot or the talent of the team running it. It is about whether a named human holds binding accountability for the decision to advance from experimentation to production, with documented rationale, defined success criteria, and consequences for the outcome either way.

Almost every enterprise is experimenting with AI and very few are scaling it. The gap between those two states is most precisely described as a governance architecture gap rather than a technology gap, a talent gap, or a data gap.

A pilot runs on clean data in a controlled environment, with a small team, a short timeline, and a forgiving success threshold. Nobody owns the outcome in a binding way, which is precisely why the pilot can succeed on its own terms and still produce nothing of enterprise value.

Production requires integration with systems that were never designed with AI in mind, security reviews, compliance audits, ongoing monitoring, and maintenance arrangements that continue long after the team that built the pilot has moved to the next experiment. A use case estimated at three months in pilot can stretch to eighteen months when real integration complexity surfaces. Failures that were learning opportunities in the pilot become business-critical incidents in production.

BCG's framework, built on work across hundreds of enterprise transformations, weights AI implementation success as ten percent dependent on algorithms, twenty percent on data and technology, and seventy percent on people, processes, and culture, a ratio that is nearly the inverse of how most AI investment decisions are structured. Organizations spend on the technology first and discover the organizational change requirement after the spending is committed, at which point the path of least resistance is to declare the experiment complete and launch another one. The result is a self-reinforcing cycle in which, without clear production-level ROI, the next pilot always feels safer than the next scale-up, cost and risk accumulate, and value does not.

Breaking that cycle requires a governance gate at the point of proposal rather than at the point of scale. Any initiative that cannot name a production owner, document an integration architecture, establish a data quality baseline, and show a workflow redesign plan does not advance past planning. The number worth targeting is not how many pilots are running but how many of last year's pilots became production systems.

SECTION 5

The Productivity Paradox Explained

The evidence that AI delivers real productivity gains at the individual worker level is credible and consistent. Workers complete more tasks, access capabilities they previously could not, and expand into work that used to require specialists. These gains are real, documented, and not in dispute, which makes the NBER finding more striking, not less: those individual gains are not reaching the enterprise P&L.

The mechanism that explains why is straightforward once named. When individual workers gain efficiency from AI, some of that efficiency is absorbed privately, as workers use AI to improve their own work quality or reduce personal cognitive load rather than converting the saved time into additional measured output. This is rational individual behavior entirely invisible to enterprise metrics.

The gains that do flow into measured output create a second, more structural problem: AI-enabled workers produce more, more code, more documents, more analyses, more drafts, and the downstream processes responsible for reviewing, approving, and integrating that output were designed for the prior volume. Research analyzing telemetry from over 10,000 developers found that teams with high AI adoption completed twenty-one percent more tasks and generated nearly double the pull requests, while review time for those pull requests increased ninety-one percent, meaning the individual gain was real and the system absorbed it because the system had not been redesigned to handle the new volume. This research comes from Faros AI, a vendor with a commercial interest in identifying AI productivity bottlenecks, and should be treated as directional support for the NBER finding rather than as independent confirmation.

What the NBER evidence establishes on its own terms is that ninety percent of firms report no firm-level impact after three years of AI adoption. The path from individual efficiency gains to enterprise P&L movement runs through a decision that most organizations have not made: who holds the authority to redesign the workflows that receive AI-generated output, and what is the checkpoint that ensures the redesign happens before additional AI investment is approved upstream. Organizations that assign explicit workflow redesign authority and bind it to investment approval will see individual productivity gains compound into enterprise performance. Organizations that layer AI onto existing workflows without redesigning what comes after will see those gains neutralized by the same bottlenecks they already had, now running faster toward the same constraints.

SECTION 6

The Shadow AI Economy

The MIT NANDA research surfaces a finding that most governance conversations skip. While only forty percent of organizations hold official AI subscriptions, ninety percent of workers surveyed reported daily personal use of AI tools, and the gap between those two numbers is not a compliance failure waiting to be closed. It is a signal about human motivation that leadership misreads at significant cost when it treats shadow AI as purely a security problem.

Employees are not using consumer AI tools solely because their employer's approved options fall short. Many are using AI at home, on personal devices, on their own time, to build capabilities they believe will matter for their careers. It shows up in managers drafting sharper analyses, developers learning new frameworks outside working hours, and analysts building domain knowledge through tools the organization has not yet sanctioned. These are acts of professional development happening outside the governance perimeter because the governance perimeter has not created space for them inside it. That distinction matters for how organizations respond.

The risk is genuine and specific. When employees apply personal AI tools to enterprise content, even with good intentions, data sovereignty obligations, confidentiality agreements, and regulatory requirements do not pause. The liability lands on the organization regardless of whether the tool was sanctioned, which means the governance exposure from shadow AI does not require a data breach to become real.

The response that treats all shadow AI use as a threat to prohibit misreads what is driving it, and organizations that restrict access without addressing the underlying motivation will find the behavior continues in less visible ways. The governance action with the highest return is not prohibition but channel design: building structured pathways for employees to develop AI capability using approved tools, with organizational support rather than personal risk. The metric that tells leadership whether this is working is not a shadow AI incident count. It is the percentage of employees with an active, sanctioned AI development pathway, tracked quarterly against the production conversion rate of the initiatives those employees own.

SECTION 7

Two Risks Most Boards Have Not Named

Sovereign AI: The Decision Most Organizations Have Not Made

Sovereign AI is the practice of designing, training, and deploying AI under a country's own laws, on infrastructure that country controls, using locally governed data. It has moved from a public sector compliance topic to a boardroom decision with direct commercial consequences on a statutory timeline. More than eighty percent of companies now view sovereign AI as at least moderately important to their strategic planning. Seventy-seven percent factor the country of origin of AI solutions into vendor selection decisions, and fifty-eight percent build their AI stacks primarily with local vendors. The shift is not philosophical; it is driven by concrete regulatory pressure with a specific date attached.

Ireland published the General Scheme of its Regulation of Artificial Intelligence Bill in February 2026, establishing the AI Office of Ireland with a statutory operational deadline of August 1, 2026. The EU AI Act's high-risk obligations go live August 2, 2026, for standalone high-risk AI systems, applying to any company operating AI systems in the EU market regardless of where that company is headquartered. AI embedded as safety components in regulated products such as medical devices and industrial machinery carries an extended deadline of August 2, 2027. Penalties for prohibited AI practices reach up to seven percent of worldwide annual turnover, a penalty structure that exceeds GDPR maximums.

The US CLOUD Act compounds the picture for US multinationals operating in Europe: US law permits the US government to compel US companies to produce data held anywhere in the world, EU data sovereignty requirements rest on a conflicting assumption, and that tension remains unresolved at the treaty level as of early 2026, though documented enforcement against EU-stored data by US authorities is rare in practice, making this a legal exposure and planning risk rather than an active operational threat for most enterprises today. The minimum governance act available to any organization is to name who decides where each AI workload runs, not who recommends but who decides, and to document that decision with the rationale. Organizations that map every AI workload against the EU AI Act's high-risk categories before August 2026 have four months to remediate whatever that mapping surfaces. Organizations that do not begin now will be remediating under enforcement pressure at a cost that may reach seven percent of worldwide annual turnover.

Physical AI: The Governance Gap Behind the Adoption Numbers

Physical AI, encompassing robotics, autonomous vehicles, drones, and AI-directed physical systems, has crossed from emerging to embedded in enterprise operations, with fifty-eight percent of companies currently using physical AI to some extent and adoption projected to reach eighty percent within two years. The investment cases being built to support that adoption are consistently underestimating two structural costs that, when they surface, convert promising deployments into expensive stalls.

The first is full deployment cost. Decision-makers routinely account for AI models and software while underestimating facility retrofits, safety infrastructure, integration with existing operational systems, maintenance contracts, spare parts, and downtime during implementation and troubleshooting, costs that can collectively exceed the software investment by multiples. Organizations that do not build those costs into business cases before commitment risk abandoning deployments midway through, which is more expensive than not starting.

The second is governance lag. Humanoid robots do not yet have a mature dedicated global safety standard for unrestricted human collaboration. Industrial collaborative robots operate under certified ISO standards, but humanoid systems deployed in enterprise settings lack equivalent certification, and physical deployment costs run two to three times stated estimates when facility and safety infrastructure are included, according to OECD analysis of physical AI deployment costs. Organizations deploying humanoid systems at scale are operating in a governance vacuum that creates liability exposure, and the accountability for decisions made in that vacuum lands on specific people whether those people have been named in advance or not. The governance measure before any physical AI commitment is a full-cost business case that includes facility, safety, and integration costs alongside the AI investment itself, with a named human accountable for validating that the total is reflected in the approved budget.

SECTION 8

What Leadership Has to Do

The evidence across every dimension points to the same five decisions, each available now, each requiring no new technology, each capable of producing measurable governance signals within a quarter of implementation.

Reclassify AI as a capital investment.

Ninety percent of firms report no firm-level productivity impact after three years of AI adoption, and the primary structural reason is that AI investment is governed as an operating expense with no ownership, no depreciation logic, and no compounding strategy. Organizations that reclassify AI as a capital investment category before the next budget cycle, building named ownership and quarterly asset tracking into the classification itself, create the accountability architecture that connects investment decisions to P&L outcomes. The recommended target: every AI initiative above a defined materiality threshold carries a named capital owner with documented accountability, reaching one hundred percent within the first budget cycle.

Name a production owner at the point of proposal.

Forty-two percent of organizations abandoned the majority of their AI initiatives before reaching production in 2025, and the governance record consistently traces that abandonment to the absence of a named human with binding authority over the advance-to-production decision. Requiring every initiative to name a production owner at the point of proposal, not at the point of scale-up, converts the governance gate from a downstream obstacle into an upfront filter that eliminates programs that were never resourced to reach production. The number worth reporting to the board is the percentage of last year's pilots now running in production, with forty percent or above as a recommended target within twelve months of implementing the gate.

Measure net productivity, not output volume.

Workday research finds that nearly forty percent of AI productivity gains are consumed by rework, which means organizations measuring gross output are systematically overstating AI's contribution to enterprise performance. Establishing measurement that tracks net time saved at the function level, connected to firm-level P&L movement through a named person accountable for the translation, replaces the metrics that make pilots look successful with the metrics that reveal whether value is reaching the income statement. Rework hours per function per quarter, measured against a baseline set before the next investment cycle begins, is the number that tells leaders whether their measurement discipline is closing the gap or papering over it.

Redesign the workflows that receive AI-generated output.

The Faros AI telemetry finding, supported by the NBER firm-level data, establishes that individual productivity gains do not compound into enterprise performance when the downstream processes receiving AI-generated output are inherited unchanged from the pre-AI workflow. Requiring the redesign of those receiving workflows as a precondition for approving additional AI investment upstream, with a named human accountable for each redesign decision, converts the bottleneck from a structural constant into a managed variable. Downstream approval latency, measured in hours per cycle against the pre-redesign baseline, is the operational signal that redesign is holding rather than reverting.

Map sovereign AI exposure before August 2026.

The EU AI Act's high-risk obligations for standalone AI systems go live August 2, 2026, with penalties reaching seven percent of worldwide annual turnover for prohibited practices. Organizations that have not mapped every AI workload against the Annex III high-risk categories, documented the sovereign AI status of every vendor in the stack, and named the human accountable for each compliance decision are carrying an unquantified liability that compounds in cost the closer enforcement arrives. Every AI workload with a documented sovereign AI classification and named decision owner before August 1, 2026 is a workload that does not require emergency remediation under enforcement pressure.

ON THE WORK BEHIND THIS PAPER

The governance gaps this paper documents have been examined from the outside, without vendor sponsorship and without a commercial stake in any particular conclusion, and that independence shapes what gets asked and what gets measured.

What the evidence consistently surfaces is a distinction that most enterprise AI conversations avoid making explicit. There are two fundamentally different modes of operating with AI. The first treats AI as an automation engine, deploying it at speed, measuring output volume, and accepting that accountability is diffuse because the system moves faster than any named human can govern it. The second treats AI as a governed asset, where a named human holds decision authority at defined checkpoints, accountability survives audit, and the measure of success is not how much the system produced but whether the human who owned the decision can answer for the outcome. Both modes have a place in enterprise AI. The question every organization is answering right now, whether they have named the question or not, is where they draw the line between them.

That line defines the product. Automation at scale without governance architecture produces factory-quality outputs: consistent, fast, and indistinguishable from every other organization running the same models on the same workflows. Governed AI with named human authority at the checkpoints that matter produces something different: outputs that carry an accountable author, decisions that can be explained and defended, and a governance record that survives the regulatory scrutiny that August 2026 will begin to apply in earnest. The organizations that understand they are making a product positioning decision, not just a technology decision, will govern accordingly.

What follows is not promotional. The frameworks and works referenced below are included because the argument this paper makes, that the governance gap is real, that it can be closed, and that closing it requires deliberate architecture rather than good intentions, is stronger when the author can point to work that has actually been built, tested, and published rather than proposed. These are not presented as the only path or the definitive architecture. They are examples of how one practitioner worked through the same problems this paper describes, from diagnosis through design through submission to policy. Other paths exist. The point is that paths exist, that the work is reproducible, and that governance at the level this moment requires is not waiting on a breakthrough. It is waiting on the decision to build.

The frameworks that emerged from that body of work, including Checkpoint-Based Governance, the Human Enhancement Quotient and its composite Augmented Intelligence Score, and the Cross AI Platform Review methodology, appear here as documented, testable examples of how the governance gap can be addressed, not as prescriptions. They are open-source, published at basilpuglisi.com and github.com/basilpuglisi/HAIA, and submitted to the 119th Congress as part of a broader AI governance infrastructure package because the argument for building governance architecture is stronger when the architecture itself is visible and auditable.

That question is what *Governing AI: When Capability Exceeds Control*, published November 2025, examines at the infrastructure level, and what the forthcoming *The Minds That Bend the Machine* explores through the work of twenty-four of the most consequential AI thought leaders across five waves of the field's development. What studying those leaders reveals is that the gap between AI capability and AI governance is not an accident of timing. It is a predictable consequence of a field that systematically prioritized speed over structure. The organizations closing it will do so the way the leaders who got it right did: deliberately, with named humans, binding checkpoints, and governance that produces evidence rather than just activity.

Basil C. Puglisi, MPA

Human-AI Collaboration Strategist & Responsible AI² Governance

basilpuglisi.com | github.com/basilpuglisi/HAIA

Research conducted independently. No vendor sponsorship. Produced under the HAIA-RECCLIN Human-AI Governance Methodology. #AIassisted

SOURCES

Yotzov, I., Berson, E., Bloom, N., Davis, S. J., & Smietanka, P. (2026, February). *What do workers think about artificial intelligence?* Evidence from linked employer-employee data (NBER Working Paper 34836). National Bureau of Economic Research. <https://nber.org/papers/w34836>

Primary source for the finding that nine out of ten senior executives report no measurable firm-level productivity or employment impact after three years of AI adoption.

S&P Global Market Intelligence. (2025). *Voice of the enterprise: AI and machine learning 2025*. S&P Global. <https://spglobal.com/market-intelligence/en/news-insights/research/ai-experiences-rapid-adoption-but-with-mixed-outcomes>

Source for the 42% AI initiative abandonment rate (up from 17%) and the 46% average proof-of-concept scrap rate across more than 1,000 enterprise professionals.

Workday, Inc. (2026, January). *New Workday research: Companies are leaving AI gains on the table*. Workday Investor Relations. <https://investor.workday.com/news-and-events/press-releases/news-details/2026/New-Workday-Research>

Source for the rework tax finding: nearly 40% of reported AI productivity gains consumed by correcting or verifying AI-generated output. Vendor-sponsored; applied as directional support.

Challapally, A., Pease, C., Raskar, R., & Chari, P. (2025, July). *The GenAI divide: State of AI in business 2025*. MIT Project NANDA (Networked Agents and Decentralized Architecture), MIT Media Lab. <https://github.com/nandapapers>

Source for the directional finding that 95% of integrated AI pilots deliver no measurable P&L impact. Methodology challenged by Futuriom (August 2025) and Professor Kevin Werbach, Wharton School; treated as pattern amplification, not primary evidence.

Faros AI. (2025). *The developer productivity paradox*. Faros AI. <https://faros.ai>

Source for the productivity bottleneck data: 21% more tasks and 98% more pull requests in high AI adoption teams alongside a 91% increase in code review time. Vendor-sponsored telemetry; applied as directional support for the NBER firm-level finding.

Boston Consulting Group. (2026, January). *From potential to profit: Closing the AI impact gap*. BCG. <https://bcg.com/publications/2025/closing-the-ai-impact-gap>

Source for the 10-20-70 framework: top-performing organizations allocate 10% of AI transformation effort to algorithms, 20% to data and technology, and 70% to people, processes, and culture.

Financial Accounting Standards Board. (2025, September). *Accounting standards update 2025-06: Intangibles — goodwill and other — internal-use software (Subtopic 350-40)*. FASB. <https://fasb.org>

Source for the software cost accounting modernization discussion; confirms capitalization practice will not significantly change for most software and no new AI capital class was created. Effective December 2027.

International Accounting Standards Board. (2025, November). *IASB update: Intangible assets project*. IFRS Foundation. <https://ifrs.org/projects/work-plan/intangible-assets>

Source for the IASB staff conclusion that AI is not sufficiently different from other intangible assets to warrant a separate test case; no new AI capital class is under active development.

Department of Enterprise, Trade and Employment, Ireland. (2026, February). *General scheme of the regulation of artificial intelligence bill 2026*. Government of Ireland. <https://enterprise.gov.ie>

Source for the AI Office of Ireland statutory establishment and August 1, 2026 operational deadline referenced in the sovereign AI compliance discussion.

European Commission. (2024). *Regulation (EU) 2024/1689 of the European Parliament and of the Council: Artificial intelligence act*. Official Journal of the European Union. <https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai>

Source for EU AI Act implementation timeline: August 2, 2026 for Annex III standalone high-risk AI; August 2, 2027 for Annex I systems embedded in regulated products; 7% of worldwide annual turnover penalty for prohibited practices.

Deloitte. (2026, January). *State of AI in the enterprise 2026*. Deloitte Insights. <https://deloitte.com/cz-sk/en/services/consulting/research/the-state-of-ai-in-the-enterprise.html>

Source for sovereign AI statistics (83% view as strategically important, 77% factor country of origin into vendor selection, 58% build with local vendors) and physical AI adoption projections (58% current, 80% within two years).

Organisation for Economic Co-operation and Development. (2025). *Physical AI deployment cost analysis*. OECD. <https://oecd.org> [provisional – pinpoint URL pending confirmation]

Source for the finding that full physical AI deployment costs, including facility retrofits and safety infrastructure, run two to three times stated software estimates. Provisional pending pinpoint citation confirmation.

International Organization for Standardization. (2025). *ISO 10218-1:2025 and ISO 10218-2:2025 – robots and robotic devices: Safety requirements for industrial robots* (incorporating former ISO/TS 15066:2016 collaborative robot safety requirements). ISO. <https://iso.org> | <https://osha.gov/robotics/standards>

Source confirming that industrial collaborative robots operate under certified international safety standards without physical separation requirements; used to narrow the physical AI governance gap claim to humanoid systems specifically.