INTERNATIONAL RESEARCH INSTITUTE — FLAGSHIP REPORT

Artificial Intelligence in the Enterprise: Adoption, Governance and Market Trajectory

Artificial intelligence · 2026-03-09 · Licence CC BY 4.0

Executive summary

The corporate adoption of artificial intelligence has moved, in roughly three years, from a specialist research concern to a line item on most large organisations' operating agendas. The release of capable general-purpose language models in late 2022 compressed a long-running trend — machine learning embedded quietly in search, advertising and logistics — into a visible, board-level question. The result is an unusual configuration: a technology that is simultaneously over-hyped in public discourse and under-deployed in practice. Adoption, measured as any reported use, is now common among large firms in advanced economies. Value capture, measured as durable margin or productivity gains attributable to AI, remains concentrated in a minority of adopters and a narrow set of use cases.

This report treats that gap between experimentation and realised return as the central analytical fact of the current market. It is not primarily a story about model capability, which continues to improve rapidly, but about the organisational, data and governance conditions that determine whether capability becomes advantage. The firms capturing measurable returns are disproportionately those that already possessed clean data pipelines, clear process ownership and the appetite to redesign workflows rather than bolt a model onto an unchanged one. For most others, the constraint is not the availability of a model but the readiness of the organisation around it.

Two structural forces will shape the market to 2030. The first is a cost paradox: the price of a unit of model inference has fallen sharply as competition and efficiency gains compound, even as the capital expenditure required to build and run frontier-scale infrastructure has risen into the hundreds of billions of dollars annually across the largest providers. This pulls the market in two directions at once — commoditising basic access while concentrating supply among a handful of well-capitalised platforms. The second is the shift of governance from voluntary principle to enforceable obligation, led in Europe by a phased, risk-tiered legal regime and internationally by management-system standards that make responsible deployment an auditable, fixed-cost activity rather than a reputational nicety.

Our central judgement is that enterprise AI is a durable, material market rather than a passing enthusiasm, but that near-term returns are widely overstated and their distribution will be uneven. We estimate the broad addressable market at roughly USD 210–320 billion in 2025 and sketch a 2030 range of USD 750 billion to USD 1.3 trillion under explicitly stated assumptions. The wide band reflects genuine uncertainty about definitional boundaries, the pace of organisational absorption, and whether infrastructure investment is validated by revenue. Decision-makers should plan for a market that grows substantially but rewards patience, data discipline and governance maturity over speed of procurement.

By the numbers

INDICATOR	VALUE
Addressable enterprise AI market, 2025 — broad definition; estimate	USD 210–320bn
Sketched 2030 range — scenario band, not a forecast	USD 750bn–1.3tn
Per-unit inference cost fall, 2023–2025 — order of magnitude or more	≥10×
Implied CAGR to 2030 — under stated assumptions	High-20s–30s%

Headline findings

- **Adoption is broad but shallow.** A large majority of surveyed enterprises in advanced economies now report using AI in at least one function, but only a minority report material, measured impact at scale. The modal outcome is a portfolio of pilots, of which a small share reach production. The distance between "using AI" and "benefiting from AI" defines the 2024–2026 market.
- **The cost curve and the capital curve are diverging.** Per-unit inference costs for comparable capability have fallen by roughly an order of magnitude or more over 2023–2025, while combined capital expenditure by the largest cloud and model providers has climbed steeply. Cheap access sits atop expensive, concentrated supply — a tension that will govern pricing power and margins through the decade.

- **Governance is becoming a fixed cost of doing business.** The EU AI Act, phased in from 2025, alongside the NIST AI Risk Management Framework and the ISO/IEC 42001 management-system standard, is converting "responsible AI" from principle into documented, auditable obligation, raising the entry cost of compliant deployment and advantaging larger firms with dedicated assurance functions.
- **The binding constraint is organisational, not technical.** Data quality, process redesign, workforce capability and change management explain most of the variance in outcomes. Model choice is increasingly a secondary decision; the primary ones concern data readiness and workflow ownership.
- **Value is migrating up the stack, unevenly.** Foundation-model access is commoditising at the base, while defensible value accrues to proprietary data, distribution, workflow integration and domain-specific applications. The application and services layers hold most enterprise spend and most durable differentiation.
- **Measurement remains weak.** Few organisations can attribute productivity or revenue effects to AI with methodological confidence. Absent credible measurement, capital allocation is driven partly by competitive anxiety, raising the risk of misallocation and a later correction in expectations.
- **The market is durable, but near-term returns are overstated.** We find no evidence that enterprise AI is a transient phenomenon, and substantial evidence that its near-term financial impact is being anticipated faster than most organisations can absorb it.

1. Context & why it matters

Artificial intelligence is not new to the enterprise. Statistical machine learning has underpinned credit scoring, demand forecasting, fraud detection, recommendation and industrial optimisation for well over a decade. What changed after late 2022 was the arrival of general-purpose models — principally large language models, and subsequently multimodal systems — that could be applied to unstructured text, code and images with minimal task-specific training. This lowered the technical barrier to entry dramatically. A capability that once required a specialist data-science team became accessible through an application programming interface or a chat interface, which is why adoption widened so quickly and why the surrounding discourse became so heated.

The significance for decision-makers is threefold. First, the addressable surface area is unusually large. Because language is the connective tissue of white-collar work — drafting, summarising, coding, searching, classifying, advising — a general-purpose language capability touches functions that previous waves of automation left largely untouched. Second, the technology is a general-purpose one in the economic sense: like electricity or the corporate database, its returns depend heavily on complementary investment in processes, skills and organisational structure, and therefore arrive with a lag. Third, the geopolitics of compute, data and standards mean that enterprise AI is not a purely commercial question. Export controls on advanced semiconductors, national investment in compute capacity, and divergent regulatory regimes all bear directly on what firms can build, where, and at what cost.

The history of general-purpose technologies counsels against both dismissal and credulity. The productivity gains from earlier waves of information technology were real but delayed, uneven and dependent on organisational change — the so-called productivity paradox, in which measured gains lagged investment by years. A reasonable prior is that AI follows a similar path: heavy investment and disappointing aggregate statistics, followed, for firms that make the complementary changes, by durable advantage. The current market sits early in that cycle, which is precisely why the gap between adoption and value is the right thing to study now.

2. Market structure and scale

Estimating the size of the "enterprise AI market" is complicated by definitional boundaries. A narrow definition counts only spending explicitly attributed to AI software, platforms and dedicated services. A broad definition includes the underlying cloud compute consumed by AI workloads and the AI-attributable share of embedded features in existing enterprise software. These definitions differ by a factor of two or more, so any single headline number is misleading without its boundary stated. We therefore present a segmented estimate with explicit assumptions rather than a single figure.

The market can be understood as a stack of five layers, each with different economics. At the base, **compute and infrastructure** — data centres, accelerators and cloud services on which models are trained and run — is capital-intensive, highly concentrated among a few hyperscale providers and chip vendors, and currently absorbing the largest absolute investment. Above it, **foundation models and model access** provide general capability through APIs or licences, a competitive layer subject to rapid price erosion

for comparable capability. The **platform and tooling** layer — orchestration, retrieval, vector stores, evaluation and monitoring, often labelled MLOps — is where models are made operable, governed and connected to enterprise data. The **application** layer embeds AI into software end-users touch, whether standalone tools or features within existing suites. The **services** layer — integration, data engineering, change management and assurance — captures the substantial human effort required to move from pilot to production.

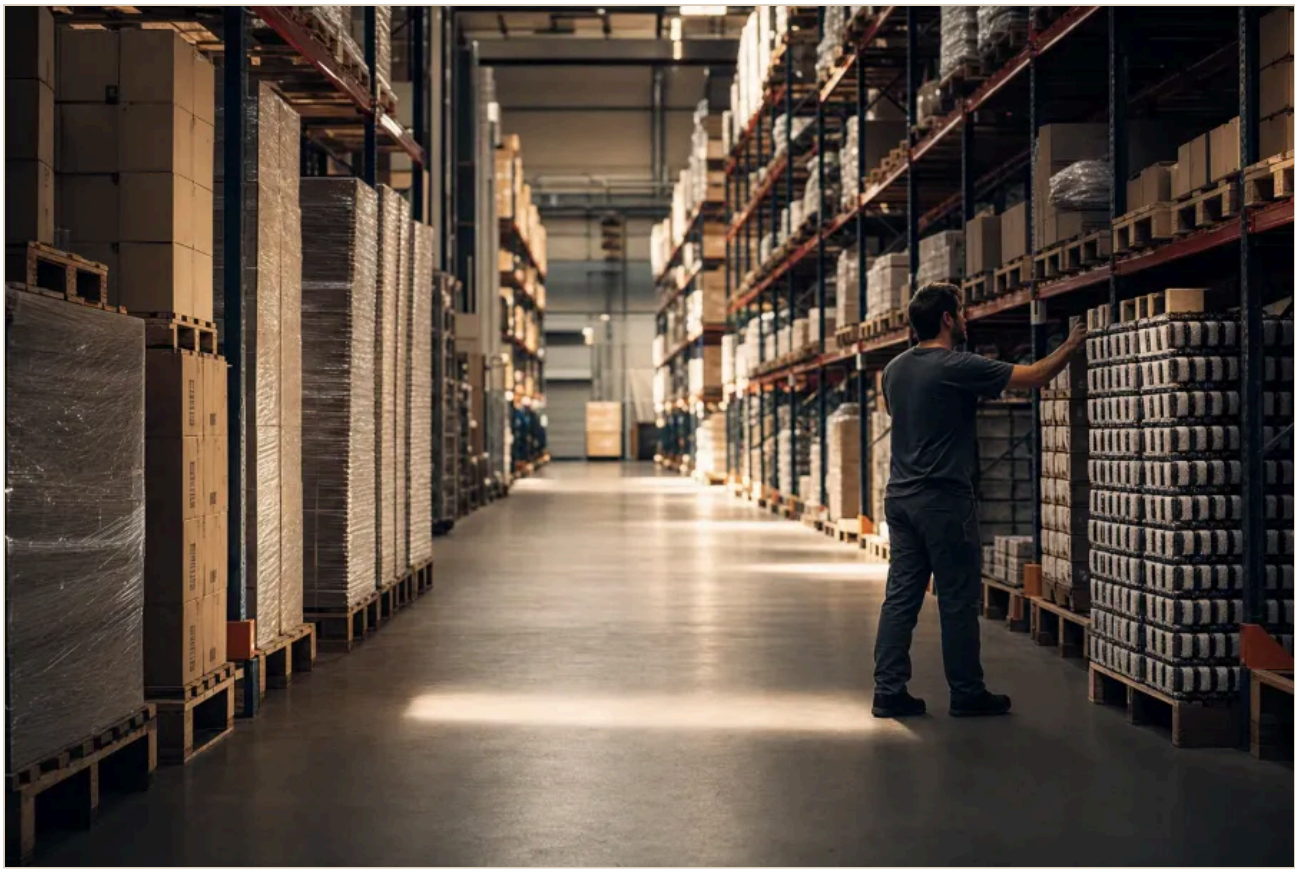

Illustrative photograph, not a data figure. The enterprise AI market can be read as a five-layer stack — compute and infrastructure at the base, then foundation-model access, platforms and tooling, embedded-AI applications, and services; the table below sets out each layer. Capital concentrates at the base, while enterprise spend and durable differentiation sit toward the top. — IRI illustrative image

The estimates below are our own triangulation from public financial disclosures, official statistics on IT and cloud spending, and the direction and magnitude implied by recurring practitioner surveys. They are transparent estimates, not measured figures, and the ranges are deliberately wide.

LAYER	2025 ENTERPRISE SPEND (ESTIMATE, USD)	BASIS / ASSUMPTIONS	CONCENTRATION	2025–2030 GROWTH CHARACTER
Compute & infrastructure (AI-attributable)	90–150bn	AI share of surging cloud and data-centre capex; excludes non-AI workloads	Very high (few hyperscalers, few chip vendors)	High absolute growth; margin depends on utilisation
Foundation models & model access	20–40bn	API and licensing revenue across leading providers; rapid price/unit erosion	High but contestable	Volume up sharply, price per unit down
AI platforms & tooling (MLOps, retrieval, monitoring)	25–45bn	Data/ML platform spend plus new orchestration and governance tools	Moderate, fragmenting	Steady, governance-driven expansion
Enterprise applications with embedded AI	45–70bn	AI-attributable share of application software and new AI-native apps	Moderate	Largest end-user spend; premium pricing tested
Services & integration	30–55bn	Systems integration, data engineering, change and assurance	Low, competitive	Front-loaded; grows with production deployment
Broad total (with overlaps removed)	**~210–320bn**	Sum less double-counting of embedded compute	—	Central 2030 range: **~750bn–1.3tn**

Plotting each layer's stated low and high bounds makes the shape of the market visible: infrastructure is the largest single band, while the application, platform and services layers together dominate the spend the typical enterprise actually controls.

2025 enterprise AI spend by layer — low–high estimate range

CATEGORY	LOW ESTIMATE (USD BN)	HIGH ESTIMATE (USD BN)
Compute & infra	90	150
Foundation models	20	40
Platforms & tooling	25	45
Embedded-AI apps	45	70
Services & integration	30	55

Section 2 estimate ranges, shown before removing double-counting. The broad total (~USD 210–320bn) nets out overlapping embedded compute. Estimates, not measured values.

Two features of this structure deserve emphasis. First, the **base of the stack is where the capital is, and the top is where the enterprise spend is**. Infrastructure absorbs the largest investment but is largely intermediated; most enterprises buy access, applications and services rather than building infrastructure themselves. This means that for the typical adopting firm, the relevant budget lines are applications, platforms and integration services, not chips and data centres. Second, **concentration decreases as one moves up the stack**. Infrastructure and frontier models are dominated by a small number of very large firms; platforms, applications and services are more fragmented and contestable, which is where most competitive entry and differentiation occur.

For context, global IT spending sits in the region of USD 5 trillion annually and enterprise software alone is on the order of USD 1 trillion, so on our broad definition enterprise AI in 2025 is still a modest single-digit percentage of total IT spend. Its significance lies less in current absolute size than in its growth rate and its potential to reallocate spending within existing IT and labour budgets. The 2030 range we sketch — roughly USD 750 billion to USD 1.3 trillion — assumes compound annual growth in the high-20s to high-30s percent, sustained production deployment rather than perpetual piloting, and no severe pullback in infrastructure investment. It is a scenario band, not a forecast, and Section 6 sets out the conditions under which the lower or upper bound becomes more likely.

3. Adoption dynamics: the pilot-to-production gap

The most robust finding across recurring enterprise surveys is not the headline adoption rate but the shape of the funnel behind it. A large and growing majority of surveyed large firms report using AI somewhere; a much smaller share report that a majority of their pilots have reached production; and a smaller share still can point to measured, material financial impact. The market is, in effect, wide at the top and narrow at the bottom.

Four mechanisms explain the attrition. **Data readiness** comes first: general-purpose models are only as useful as the enterprise data they are grounded in, and most organisations discover during pilots that their data is fragmented, poorly labelled, inconsistently governed or legally encumbered. Retrieval-augmented approaches shift the bottleneck from the model to the data pipeline, frequently the weakest part of the estate. **Workflow integration** comes second: a model that produces a plausible draft is valuable only when embedded in the actual sequence of work, with clear handoffs, review steps and accountability. Bolting a chat interface onto an unchanged process yields novelty rather than throughput.

Reliability and evaluation is third. Generative systems are probabilistic; they produce fluent output that can be wrong in ways hard to detect at a glance. Production requires evaluation harnesses, guardrails, human review where stakes are high, and monitoring for drift — an engineering discipline many organisations underestimated. **Change management and trust** is fourth: front-line staff adopt tools they trust and route around those that feel imposed or unreliable, so the adopters that reach production invest as heavily in training, incentives and process redesign as in the technology itself.

The practical implication is that the pilot-to-production gap is not evidence that the technology does not work; it is evidence that deployment is an organisational project. Firms that treat AI as a procurement decision tend to accumulate pilots; firms that treat it as an operating-model change tend to reach production. This distinction, more than any difference in model choice, separates the adopters capturing returns from those reporting activity without impact.

> *"Firms that treat AI as a procurement decision tend to accumulate pilots; firms that treat it as an operating-model change tend to reach production."*
>
> *— §3, Adoption dynamics*

4. Economics: the cost paradox and where value accrues

The most consequential economic fact of the current market is the divergence between the falling cost of using a model and the rising cost of building the capacity to serve one. On the demand side, the price of a unit of inference for a given capability has fallen sharply — by roughly an order of magnitude or more over 2023–2025 — driven by competition, architectural efficiency, better hardware utilisation and the arrival of smaller models that match older frontier performance at a fraction of the cost. For the buyer, capable AI is getting cheaper per unit, quickly. On the supply side, the capital expenditure to build and operate frontier-scale infrastructure has

risen steeply, with combined annual capex among the largest providers climbing into the hundreds of billions of dollars. The base of the stack is becoming more expensive and more concentrated even as access at the top becomes cheaper.

> ***The cost paradox*** *— Per-unit inference cost for comparable capability has fallen by roughly an order of magnitude or more over 2023–2025, even as combined annual capital expenditure among the largest providers has climbed into the hundreds of billions of dollars — cheap access sitting atop expensive, concentrated supply.*

This paradox has clear consequences for where value accrues. Because basic model access is commoditising, it is a poor place to build durable advantage; capability parity between leading providers is narrow and temporary, and switching costs at the raw-access level are modest. Defensible value is migrating to assets that are harder to replicate: **proprietary data** that grounds and specialises a model; **distribution and installed base**, which determine who can put AI in front of users at scale; **workflow integration**, which embeds AI so deeply in a process that removing it is costly; and **domain-specific applications**, where understanding a regulated or specialised field is itself the moat. For most enterprises, this means the return on a proprietary data strategy and on deep workflow integration will exceed the return on chasing the newest model.

The pricing model itself remains unsettled. Per-seat premiums for AI features face resistance where value is unproven; consumption pricing exposes buyers to unpredictable bills as usage scales; outcome-based pricing is attractive in principle but hard to measure and attribute. We expect continued churn until measurement improves. There is also a real, unresolved question about whether current infrastructure investment will be validated by revenue. If enterprise willingness to pay does not rise to meet the capital deployed, the market could see margin compression at the infrastructure layer and a rationalisation of investment — a risk we treat explicitly in the scenarios.

5. Governance, regulation and the comparative regional lens

Governance has moved from the margins of the enterprise AI conversation to its centre, and it is doing so through enforceable instruments rather than voluntary principles. The most consequential is the European Union's AI Act, which entered into force in 2024 and applies in phases: prohibitions on a defined set of unacceptable-risk practices from early 2025, obligations on general-purpose model providers later in 2025, and the substantive obligations on high-risk systems phased in over the following years. The Act is risk-tiered — the heaviest obligations fall on high-risk uses in areas such as employment, credit, essential services and safety components — and its extraterritorial reach means that non-European firms serving European customers are within scope. For enterprises, the practical effect is that classification, documentation, human oversight, data governance and post-market monitoring become legal requirements for certain uses rather than optional good practice.

Alongside binding law, a layer of **standards and frameworks** is professionalising governance. The NIST AI Risk Management Framework, published in the United States in 2023, offers a voluntary but widely referenced structure for identifying and managing AI risks. The ISO/IEC 42001 management-system standard, published at the end of 2023, provides a certifiable framework for an AI management system, analogous to established standards for information security. Together these are converting "responsible AI" into an auditable, resourced function with owners, documentation and review cycles. The direct consequence is a higher fixed cost of compliant deployment, which advantages larger organisations able to amortise dedicated assurance teams and disadvantages smaller entrants — a dynamic worth watching for its effect on competition.

Governance: from voluntary principle to enforceable obligation

WHEN	MILESTONE	DETAIL
2023	NIST AI RMF and ISO/IEC 42001 published	The US NIST AI Risk Management Framework (2023) and the ISO/IEC 42001 AI management-system standard (end of 2023) begin professionalising governance into an auditable, resourced function.
2024	EU AI Act enters into force	The EU's risk-tiered, extraterritorial regime begins its phased application.
Early 2025	Prohibitions apply	Bans on a defined set of unacceptable-risk practices take effect.
Later 2025	General-purpose model obligations	Obligations on providers of general-purpose models begin to apply.
2026 onward	High-risk obligations phased in	Substantive obligations on high-risk systems — employment, credit, essential services, safety components — phase in over the following years.

The regional picture is genuinely divergent, and the divergence is strategic. The **European Union** has led with comprehensive, rights-based, horizontal regulation, accepting some cost to innovation in exchange for legal certainty and fundamental-rights protection. The **United States** has pursued a more sectoral, lighter-touch and shifting posture, with federal direction changing across administrations and much of the concrete activity at agency and state level, tilted toward maintaining competitive and infrastructure leadership. **China** combines state-supported build-out of compute and models with targeted content and algorithm regulation aligned to its own governance priorities. Others — the United Kingdom, several Gulf states investing heavily in compute, and a range of Asian economies — position between these poles. For multinationals, a single global compliance posture is no longer sufficient; governance must be built to the highest applicable standard and configured by jurisdiction. This raises cost but rewards firms that treat governance as a capability rather than a constraint.

6. Risks, uncertainties and open questions

The enterprise AI market carries a distinctive risk profile. The most immediate is an **expectations correction**: investment is running ahead of demonstrated returns, and capital allocation is driven partly by competitive anxiety rather than evidence. If a critical mass of firms concludes that near-term returns were overstated, the market could see a sharp reset in sentiment and spending even as the technology keeps improving — a repricing of expectations, not a refutation of the technology, but painful for those exposed to the infrastructure and frontier-model layers.

Several operational risks compound this. **Measurement risk** means the market is allocating capital on weak evidence, raising the probability of misallocation and blurring the line between genuine value and activity. **Reliability risk** is intrinsic: probabilistic systems fail in ways that are hard to anticipate, and the cost of a fluent, confident, wrong output scales with the stakes of the use case. **Concentration and dependency risk** follows from reliance on a small number of infrastructure and model providers, exposing enterprises to pricing power, supply constraints, policy shocks such as export controls, and single points of failure. **Security and data risk** introduces new attack surfaces — prompt injection, data exfiltration through model interfaces — and the governance challenge of sensitive data flowing to third-party models.

Two further uncertainties are structural. **Organisational capability** remains scarce; the constraint on absorption is as much about people who can redesign processes as about specialists. And the **workforce and social dimension** — the distribution of gains, the effect on particular occupations, and public trust — will shape the regulatory environment in which enterprises operate, feeding directly back into the pace and permissibility of deployment.

7. Three scenarios to 2030

We present three scenarios rather than a single forecast, because the principal uncertainties — the pace of organisational absorption, the validation of infrastructure investment by revenue, and the trajectory of governance — are genuinely unresolved. Each scenario is

internally consistent and tied to observable signposts.

Scenarios & outlook to 2030 (broad-definition market)

B — Expectations correction — Downside

Investment outruns demonstrated returns; a sentiment reset compresses spending and margins at the infrastructure and frontier-model layers, and consolidation removes weaker entrants — a repricing of expectations, not a collapse of the technology.

METRIC	VALUE
Market 2030	USD 750–850bn

A — Measured absorption — Central case

Technology keeps improving and cheapening, governance settles into a workable if costly steady state, and organisations gradually build the data and process foundations for production. Strong growth without a discontinuity.

METRIC	VALUE
Market 2030	USD 900bn–1.1tn

C — Compounding returns — Upside

Organisational absorption proves faster than expected, measured productivity gains become credible and visible, and a broad base of firms reaches production across functions — validating infrastructure investment.

METRIC	VALUE
Market 2030	USD 1.2–1.3tn or above

Scenario A — Measured absorption (central case). In the most likely path, the technology continues to improve and cheapen, governance settles into a workable if costly steady state, and organisations gradually build the data and process foundations required for production deployment. Returns arrive with a lag and are unevenly distributed, concentrated among firms that make the complementary organisational investments. The market grows strongly but without a discontinuity, landing toward the middle of our 2030 range (broadly USD 900 billion to USD 1.1 trillion on the broad definition). Signposts: steady rises in the share of pilots reaching production, improving but still imperfect measurement, and no severe infrastructure pullback.

Scenario B — Expectations correction. Investment outruns demonstrated returns; a critical mass of firms concludes that near-term impact was overstated; and a sentiment reset compresses spending and margins, particularly at the infrastructure and frontier-model layers. Deployment continues but more cautiously, and consolidation removes weaker entrants. The market still grows relative to 2025 but lands toward the lower bound of our range (USD 750–850 billion). This is not a collapse of the technology but a repricing of expectations, and it may set the stage for healthier, evidence-led growth thereafter. Signposts: stalling production conversion, public retrenchment of AI budgets, and visible over-capacity in infrastructure.

Scenario C — Compounding returns. Organisational absorption proves faster than expected as tooling matures, governance is streamlined into practice, and a broad base of firms reaches production across multiple functions. Measured productivity gains become credible and visible, validating infrastructure investment and pulling spending toward the upper bound (USD 1.2–1.3 trillion or above). This scenario requires genuine progress on measurement and on the organisational bottleneck, not merely on model capability, and is therefore less about technical breakthroughs than about diffusion. Signposts: credible attribution of productivity gains, rapid production conversion across sectors, and sustained enterprise willingness to pay.

2030 market band by scenario

CATEGORY	LOWER BOUND (USD BN)	UPPER BOUND (USD BN)
Correction (B)	750	850
Central (A)	900	1100
Compounding (C)	1200	1300

Broad-definition market. Ranges as stated in §7; Scenario C is given as 'USD 1.2–1.3tn or above'. Scenario band, not a forecast.

Across all three, our central judgement holds: the market is durable and material, near-term returns are widely overstated, and their distribution will be uneven. The scenarios differ mainly in the pace of organisational absorption and the discipline of capital allocation, not in the direction of travel.

8. What this means in practice

For governments and policymakers

The policy task is to enable adoption while managing genuine risks, without entrenching the concentration that the economics already favour. Three priorities follow. First, **invest in the enabling conditions** — data infrastructure, standards, skills and compute access — that determine whether the productivity gains materialise and are broadly shared. Second, **make governance proportionate and interoperable**: risk-tiered obligation is defensible, but fragmentation across jurisdictions raises compliance cost in ways that fall hardest on smaller firms, so alignment on standards and mutual recognition is valuable. Third, **attend to competition and concentration** at the infrastructure and model layers, and to the labour-market and distributional effects that will shape public trust and the political room for continued deployment.

For business leaders and operators

Treat AI as an operating-model change, not a procurement exercise. The evidence is consistent that returns accrue to firms that redesign processes, invest in data quality, and manage change — not to those that buy the newest model. Concentrate on a small number of high-value use cases with clear ownership, credible measurement and a realistic path to production, rather than a sprawling portfolio of pilots. Build the assurance and governance function early, because it is becoming a fixed cost of compliant operation and a source of competitive advantage where trust matters. And invest in proprietary data and workflow integration, which are where durable advantage now sits, rather than in chasing model parity that erodes within months.

For investors and financial institutions

Distinguish carefully between the layers of the stack, which carry very different economics and risk. The infrastructure and frontier-model layers hold the heaviest capital intensity and the greatest exposure to an expectations correction; the application and services layers hold more of the durable enterprise spend and, often, better unit economics. Prefer businesses that can demonstrate attributable value over those trading on narrative, and stress-test exposure to a Scenario B correction, to concentration and dependency risk, and to regulatory divergence. The opportunity is real and large, but the near-term risk of overpaying for anticipated returns is equally real, and the discipline of evidence-led allocation will matter more than speed.

Methods, sources and limitations

This report is a synthesis, not a primary data collection. Its purpose is to integrate widely-established facts with transparent estimates and clearly labelled scenarios, and to separate the two throughout.

Approach. We combined four kinds of source material. First, **public financial disclosures** from major cloud, semiconductor, model and software providers, used to gauge the direction and rough magnitude of infrastructure investment and AI-related revenue. Second, **official statistics** on IT, software and cloud spending and on the wider macroeconomy, used to anchor the enterprise AI market as a share of larger, better-measured aggregates. Third, **standards and regulatory texts** — the EU AI Act, the NIST AI

Risk Management Framework and the ISO/IEC 42001 management-system standard — read directly to characterise the governance trajectory. Fourth, **recurring practitioner surveys** on enterprise AI adoption, used for the shape and direction of adoption dynamics rather than for precise point estimates.

What is established versus estimated. Established facts used here include the phased structure and entry-into-force of the EU AI Act, the publication and character of the NIST and ISO frameworks, the broad direction of falling per-unit inference costs and rising provider capital expenditure, and the order-of-magnitude scale of global IT and enterprise-software spending. These are drawn from primary or widely corroborated sources. By contrast, **every market-sizing figure and every 2030 number in this report is an estimate**, produced by triangulation and presented as a range with its assumptions stated. The layer-by-layer 2025 estimates and the 2030 scenario band in Section 2 are our own construction; they are not measured values and should not be cited as such. Adoption "shares" are described qualitatively because published survey figures vary substantially with sampling, definitions and timing, and averaging them would imply a false precision.

Limitations. Definitional boundaries make the market size sensitive to whether embedded compute and embedded features are counted, which is why we present a range rather than a single number. The field moves quickly, and figures tied to model prices, capital expenditure and regulatory milestones will date. Measurement of AI's productivity effect is genuinely weak across the economy, and we have not manufactured confidence the evidence does not support. Readers should treat the scenarios as structured reasoning about plausible futures, conditioned on stated signposts, rather than as predictions.

Independence and funding. This report was produced under the Institute's standard independence policy. It was not commissioned or sponsored by any provider, vendor or interested party, and the authors hold no positions that would benefit from a particular conclusion.

Sources and further reading

Primary and official sources were read directly. The proprietary analyst estimates that inform the market-sizing ranges are referenced by category in the methods note above rather than itemised here.

- European Parliament and Council of the European Union (2024). *Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act).* Official Journal of the European Union.
- National Institute of Standards and Technology (2023). *Artificial Intelligence Risk Management Framework (AI RMF 1.0).* NIST, US Department of Commerce, Gaithersburg, MD.
- International Organization for Standardization and International Electrotechnical Commission (2023). *ISO/IEC 42001:2023 — Information technology, Artificial intelligence, Management system.* ISO/IEC, Geneva.
- OECD (2024). *OECD Digital Economy Outlook 2024, Volume 1: Embracing the Technology Frontier.* OECD Publishing, Paris.
- OECD.AI Policy Observatory (2025). *Trends in national AI policy and compute capacity.* OECD, Paris.
- International Monetary Fund (2024). *Gen-AI: Artificial Intelligence and the Future of Work.* IMF Staff Discussion Note SDN/2024/001. Washington, DC.
- Stanford Institute for Human-Centered Artificial Intelligence (2025). *Artificial Intelligence Index Report 2025.* Stanford University, Stanford, CA.
- International Telecommunication Union (2024). *Measuring Digital Development: Facts and Figures 2024.* ITU, Geneva.
- US Bureau of Economic Analysis and Eurostat (2024–2025). *National accounts and ICT investment statistics.* Official series used to anchor information-technology and software spending as a share of larger aggregates.
- Industry analyst syntheses (2023–2026). Recurring enterprise-technology market trackers and chief-information-officer spending surveys, used in aggregate for direction and order of magnitude only.

Suggested citation

Marsh, H. and Adeyemi, R. (2026). *Artificial Intelligence in the Enterprise: Adoption, Governance and Market Trajectory.* Flagship Report. International Research Institute, Markets & Industry and Technology & Society Programmes. Licensed under CC BY 4.0.