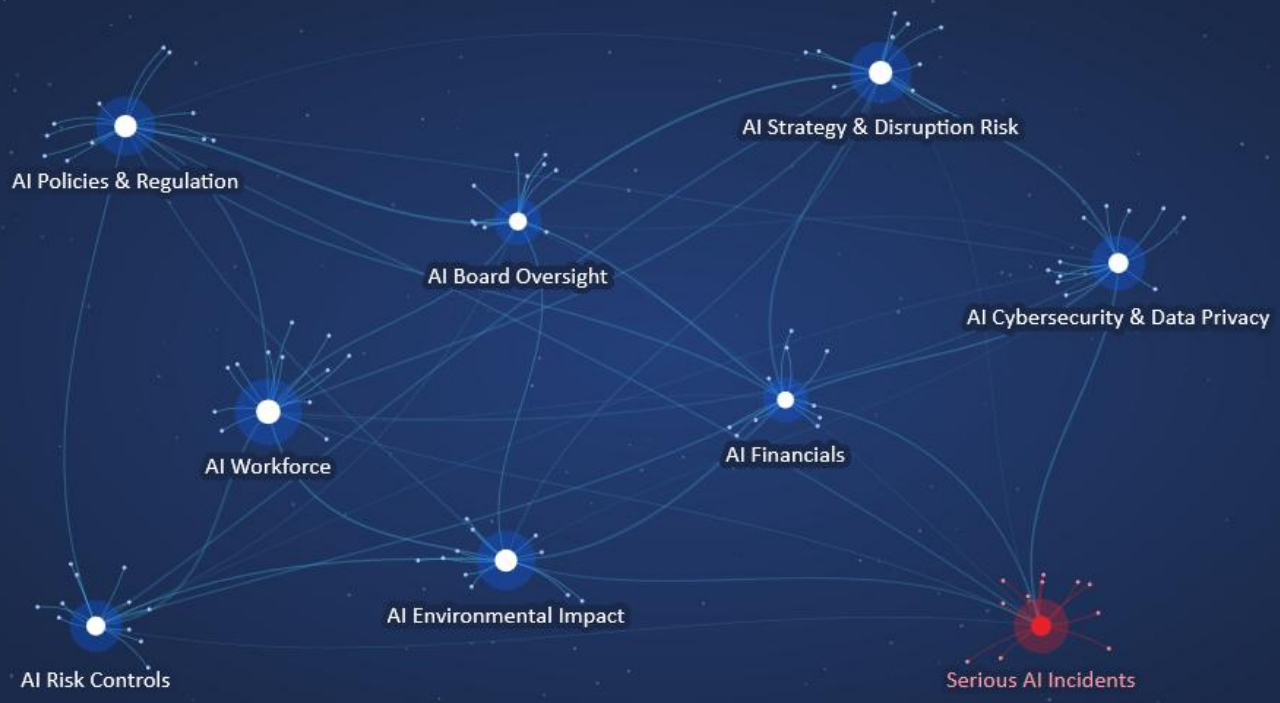
OCTOBER 2026



AI Adoption: What are Companies Signalling?

A Risk and Value Creation Assessment Framework

Hortense Bioy, CFA

Contents

About the Author 3
Executive Summary 4
Key Takeaways 4
Introduction 9
Methodology and Scope 10
Findings 14
Conclusion 26
Appendix – 150 Companies with AI Disclosure Scores 27
References and Further Resources 30

About the Author



Hortense Bioy, CFA

Hortense Bioy, CFA, is an award-winning investment researcher and thought leader in London who researches corporate AI adoption and applies AI daily in her own work. She has dedicated the past year to exploring how AI can generate new datasets and insights for institutional investors.

This work began at Morningstar, where Hortense spent 15 years leading investment research teams. As Head of Sustainable Investing Research, she was responsible for the firm's ESG and climate investing thought leadership program. She was previously Global Director of Sustainability Research within Morningstar Manager Research. She co-developed ESG ratings, frameworks and datasets.

Prior to that, Hortense led Morningstar's European ETF research for eight years and spent five years as a financial journalist, including three as a television presenter on Bloomberg TV France.

Hortense began her career as an M&A analyst, including two years at Société Générale in Hong Kong. She is a CFA charterholder.

Executive Summary

Artificial intelligence[1] is the defining strategic question of the decade. Capital is pouring into the technology at unprecedented speed,[2] and companies are adopting it almost as fast.[3,4] Yet disclosure has not kept pace.[5,6,7] If AI is as transformative as many expect, investors will need clearer reporting on how companies use AI and whether they can capture its value while keeping its risks in check.[5] That information will help investors distinguish the winners of this transformation from the losers.

So far, investor scrutiny has focused on the companies developing AI, mainly those building frontier models and data centers. As developers' valuations look increasingly stretched and the risks associated with AI, from business disruption to operational, legal, regulatory, and reputational risks, grow,[11] that attention has started to turn to the companies deploying AI as well: how far they use it in their operations, what they aim to achieve, and how they govern it.[5,6,7] On these questions, disclosure is inconsistent, varies widely across sectors and companies and lacks a standard approach.

This report offers a holistic framework to help investors and risk managers quickly understand a company's relationship with AI and how it engages with the technology. The framework consists of 34 indicators grouped into nine categories:

- Strategy & Disruption Risk
- Policies & Regulation
- Board Oversight
- Risk Controls
- Cybersecurity & Data Privacy
- Workforce
- Financials
- Environmental Impact
- Serious AI Incidents

To demonstrate the framework in practice, this report applies it to 150 of the largest listed companies in the US and Europe, including 133 AI deployers and 17 AI developers, with a combined market capitalization of US$ 60 trillion.[8] Each company receives a total score between 0 and 100, representing the sum of its indicator scores. For each indicator, the highest score reflects clear, specific and evidenced disclosure and, for some indicators, stronger practice. The full list of companies with their scores is in the Appendix.

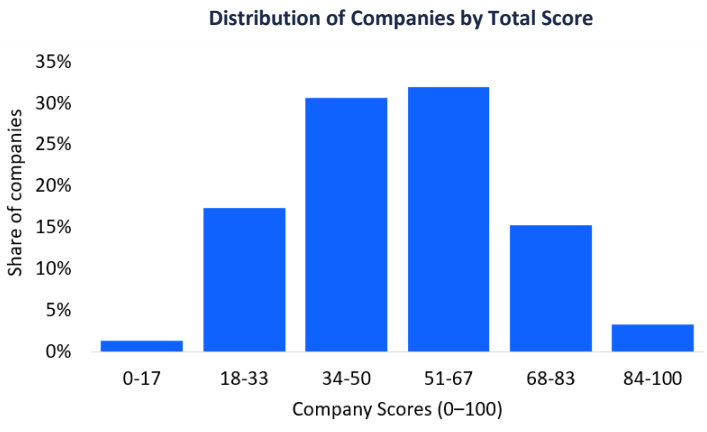
The main objective of this research is to give investors and risk managers a tool to assess, benchmark and track over time what companies share about their use of AI. Differences in disclosure between companies and sectors can signal how mature, strategically important and well governed AI is within a business and highlight areas that may warrant further analysis or engagement. The research also aims to lay the groundwork for best practice: a common set of areas that companies can report on and investors can ask about, so that AI disclosure becomes more consistent and comparable.

The key takeaways below summarize what the 150 companies disclose, focusing on the deployers. The main body of the report also compares deployers with developers, and the full list of companies is in the Appendix.

Key Takeaways

AI adoption is widespread, but its disclosure is uneven

- **Virtually all 150 companies in the sample report using GenAI.**
- **How much companies disclose about AI varies widely.** Scores range from 9 to 98 out of a possible 100 points.
- Almost half (49%) of the companies score below the sample's average of 52.
- **A low score does not necessarily mean a company is not using AI**, only that it discloses little about how it uses it or the risks involved.



Key takeaways continue on the next page ›

Microsoft leads on AI disclosure, Glencore says least

- **Technology dominates the top 10.** Five of the top 10 scorers on AI disclosure are IT or Communication Services companies, led by Microsoft, Alphabet, and SAP. Three financial companies, Visa, BNP Paribas, and JPMorgan Chase, also make the top 10.
- **The bottom 10 cuts across sectors.** Glencore sits at the very bottom.

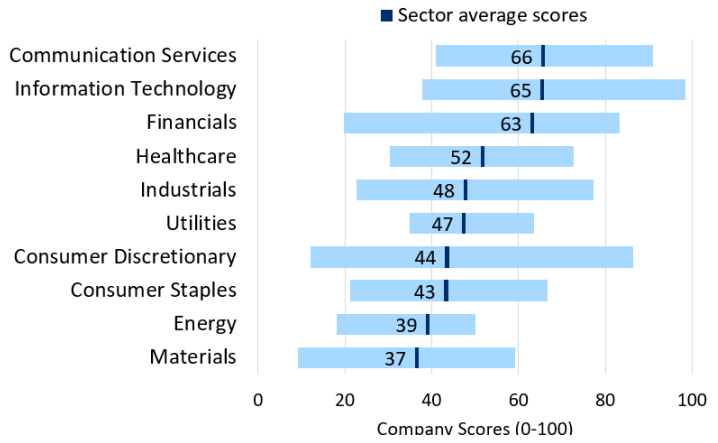
Companies with the Highest and Lowest AI Disclosure Scores

Top 10	Sector	Score	Bottom 10	Sector	Score
Microsoft	IT	98	Glencore	Materials	9
Alphabet	Com. Services	91	Ferrari	Cons. Discr.	12
Amazon	Cons. Discr.	86	Valero Energy	Energy	18
SAP	IT	86	Investor AB	Financials	20
Cisco	IT	85	AB InBev	Cons. Staples	21
Visa	Financials	83	TJX Companies	Cons. Discr.	23
BNP Paribas	Financials	82	BAE Systems	Industrials	23
Meta	Com. Services	82	Berkshire Hathaway	Financials	24
JPMorgan Chase	Financials	77	Atlas Copco AB	Industrials	24
Schneider Electric	Industrials	77	Costco	Cons. Staples	26

AI builders disclose the most

- **Communication Services and Technology lead.** As builders and sellers of AI, companies in these sectors face the closest scrutiny and disclose accordingly.
- **Financials rank high too.** This reflects both high AI adoption and regulation.
- **Materials and Energy trail.** Companies in these sectors use AI across their operations but say little about it.
- **The wide variation in disclosure within each sector** raises questions about how companies differ in their approaches to AI.

Spread of company scores by sector — Full Sample (n=150)



Governance issues take prominence

- **Governance issues sit at the top.** Being the easiest to put in place, AI policies and board oversight are the most widely disclosed governance elements.
- **AI risk controls score much lower.** Companies state their intentions more readily than they show evidence that governance works day to day.
- **Workforce-related disclosure is uneven.** Companies are more forthcoming about how much their employees use AI than about how they use it, how it affects them and how they are protected.
- **AI financials remain opaque.** The full costs and returns of AI are seldom disclosed.
- **Environmental impact is largely unreported.** For now, it is material only to the developers and deployers that operate data centers.
- **Serious AI incidents are rare but telling.** 24% of developers and 5% of deployers had one in the recent past, usually exposed by litigation, regulators or the press. It shows whether AI governance holds up in practice.

Nine Broad Categories — Full Sample (n=150)

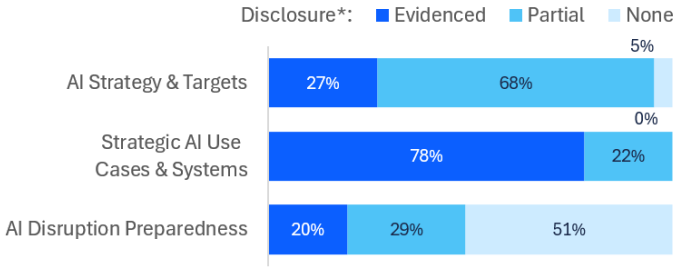


Key takeaways continue on the next page ›

AI deployers describe the opportunity more than the threat

- **The strategic role of AI is widely described.** 78% of deployers say where AI is applied and where it is expected to create most value.
- **Targets are rare though.** Only 27% of deployers share targets or timelines for results.
- **Disruption risk is acknowledged but rarely addressed.** 49% of deployers recognize that AI could disrupt their business, but only 20% set out a response.

AI Strategy & Disruption Risk — AI Deployers (n=133)



Most deployers report AI policies, but fewer spell out their principles

- **Policies are near universal.** 96% of the deployers report having AI policies in some form.
- **Responsible-use statements are the norm.** 92% of deployers say they aim to use AI responsibly, and 57% describe their responsible AI principles in detail.
- **Specific disclosure of regulatory readiness is limited.** Just 40% name the specific AI industry standards they align with or the regulations they comply with.

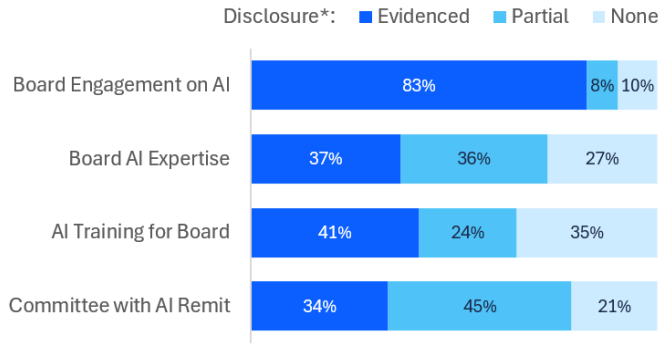
AI Policies & Regulation — AI Deployers (n=133)



Boards are engaged on AI, but expertise and support lag

- **Board engagement is common.** 83% of deployers show board engagement with AI strategy and risk.
- **Expertise lags.** Only 37% of deployers name AI expertise in their disclosure on directors.
- **Training may narrow the gap.** 41% of deployers have board-specific AI training.
- **Supporting structures are often unclear.** Only 34% of deployers assign AI oversight to a board committee, or to a management body reporting to the board.

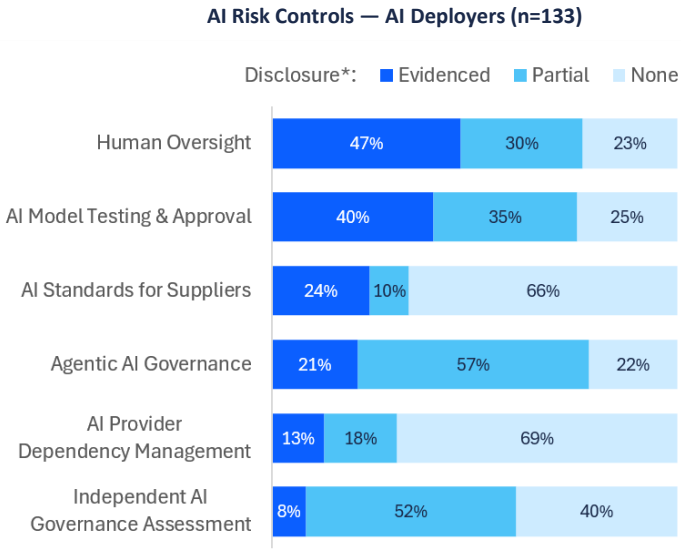
AI Board Oversight — AI Deployers (n=133)



Key takeaways continue on the next page ›

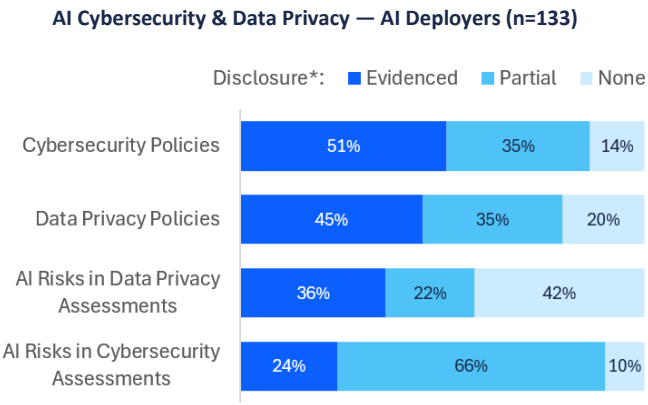
Independent assessment of AI governance is rare

- **Human oversight is common, but details are lacking.** Less than half (47%) of deployers explain the role of humans in key processes. Fewer (21%) set out how agentic AI is governed.
- **Only 40% of deployers describe how AI models are tested and approved.**
- **Few deployers extend their AI standards to suppliers:** 24% of deployers disclose such requirements.
- **AI provider dependency is rarely acknowledged as a risk.** 69% of deployers say nothing about managing their dependency on AI model and system providers, and only 13% report addressing it in detail.
- **Few deployers seek independent assurance of their AI governance.** Only 8% report an independent review of it, whether by internal audit, external assurance or certification.



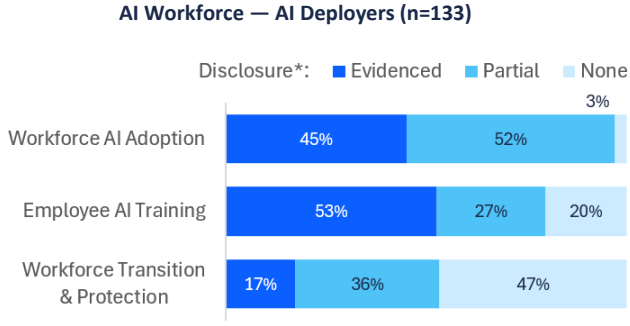
AI security is widely covered, but rarely assessed in detail

- **AI is widely included in cybersecurity and privacy policies.** At least 80% of deployers have cybersecurity and/or data privacy policies that cover AI, recognizing the added risk of breaches.
- **AI in cybersecurity tests lags.** While 36% of deployers describe in detail how AI risks are brought within the scope of privacy assessments, just 24% report doing so for cybersecurity risk assessments.



Workforce use of AI is widely reported, protection is not

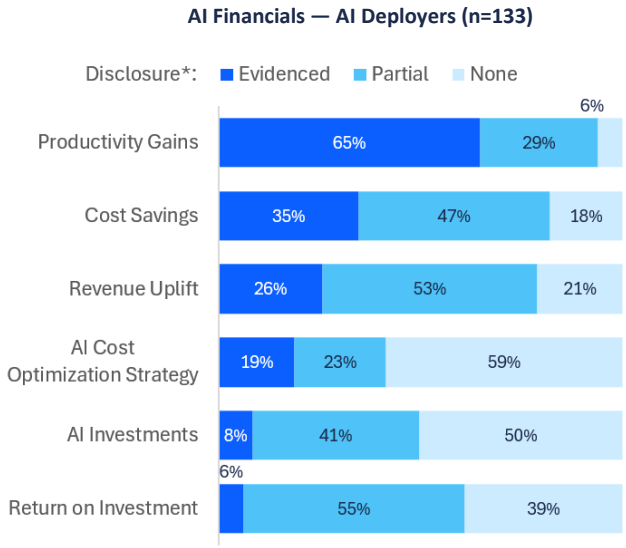
- **Workforce adoption is widely tracked.** 97% of deployers report the extent to which their workforce is adopting AI.
- **Workforce skill development is common.** 80% of deployers report AI training programs, though scope and quality vary significantly.
- **Nearly half (47%) of deployers make no statement on the future of work** or on how workers will be protected from AI's effects.



Key takeaways continue on the next page ›

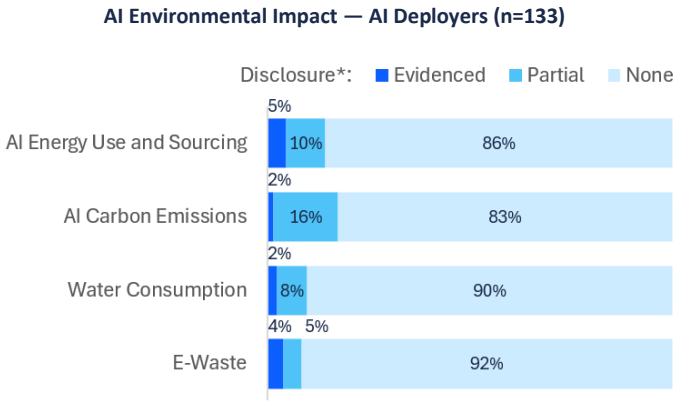
Deployers point to the benefits of AI, but few show whether it pays off

- **Financial benefits are widely reported.** Almost two thirds (65%) of deployers report productivity gains in detail and 35% do so for cost savings, usually in parts of their business.
- **Revenue from AI is usually not broken out in detail.** Most deployers mention AI-driven revenue, but only 26% quantify it.
- **Deployers say little about their AI spending.** Half disclose nothing on AI investment, and only 8% report company-level AI spending.
- **Disclosure of return on investment (ROI) is rare.** Only 6% of deployers set out ROI in detail, while 55% refer to it in part.



AI's environmental impact goes largely unreported

- **The least developed area of disclosure.** No more than 18% of deployers say anything about AI's environmental footprint, mostly those that run their own data centers. Less than 5% disclose some quantified metric, reflecting the lack of materiality for most deployers at this point.
- **Regulation is coming.** The EU has proposed mandatory sustainability labels for data centers from 2027, which should make the footprint of the facilities running AI more visible.



**Note: For some indicators, a top score also reflects stronger practice; see Scoring in the Methodology section.*

Introduction

Artificial intelligence[1] is the defining strategic question of the decade, and the scale of the bet is without precedent. Global AI spending is forecast to total US$ 2.6 trillion in 2026.[2] Adoption is moving almost as fast: nearly nine in ten organizations now use AI regularly in at least one business function.[3,4]

2026 may prove to be a pivotal year. Fear of falling behind competitors, or of losing market share to AI-native entrants, has pushed many companies to accelerate adoption, with some already looking to "reimagine" themselves entirely.[9,10] Disclosure, however, has not kept pace with adoption.

If AI is a truly transformative force, investors will need more information in order to differentiate the winners from the losers of this transformation. The winners will be those that turn AI's potential into tangible value,[4] govern the technology well enough to prevent, or respond to, incidents such as a rogue AI agent or an AI-enabled cyberattack,[11,12] and bring their people with them through the transformation. The complexity and scale that make AI powerful also make it hard to govern. Where oversight is weak, AI can expose companies to financial, regulatory, operational and reputational risks.

So far, investor attention has focused largely on the companies developing AI: mainly those building frontier models and data centers. As adoption spreads, attention is shifting to the companies deploying AI: how they use it, how they govern it and what they aim to achieve.

A growing body of research on corporate AI adoption reflects this shift. Much of this work relies on surveys or, where it draws on public disclosures, focuses on specific sectors (e.g. financials),[13,14] topics (e.g. AI governance),[15,16] or countries (e.g. the United States).[17,18] Market-wide insights into AI adoption for investors are still thin.

This report aims to bridge that gap and offer a holistic, investor-focused framework, with a particular focus on companies that deploy AI. The framework consists of 34 indicators, grouped into nine categories: Strategy & Disruption Risk, Policies & Regulation, Board Oversight, Risk Controls, Cybersecurity & Data Privacy, Workforce, Financials, Environmental Impact, and Serious AI Incidents.

This report applies the framework to 150 of the largest listed companies in the US and Europe, scoring each indicator from 0 to 2, where 2 reflects clear, specific and evidenced disclosure and, for some indicators, stronger practice.[19] With a combined market capitalization of US$ 60 trillion,[8] the 150 companies comprise 133 AI deployers and 17 AI developers. As some of the largest and best-resourced companies, they are also likely to be among the most advanced disclosers.[16] The full list of companies and their total scores is in the Appendix.

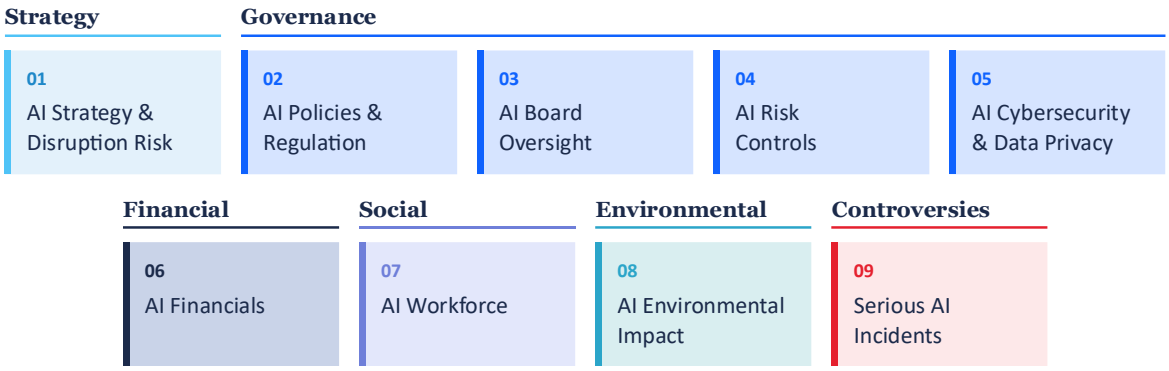
The main objective of this research is to give investors a consistent tool to benchmark what companies choose to disclose about their use of AI, identify gaps and track progress over time. Differences in disclosure between companies and sectors can signal how mature, strategically important and well governed AI is within a business and highlight areas that may warrant further analysis or engagement.

Methodology and Scope

Indicator Framework

To assess AI adoption disclosure practices, the research has developed a **framework of 34 indicators**, **grouped into nine categories**, as outlined below.

Figure 1: The framework's nine broad categories



This framework is deliberately broad, assessing AI disclosure through strategic, governance, financial, social, and environmental lenses. The indicators draw on research and guidance from AI industry bodies, non-governmental organizations, standard setters, asset managers and research firms. They are also informed by actual corporate disclosure and by engagement with companies and AI experts. A list of resources appears in the References and Further Resources section.

Several indicators draw on established AI standards and guidance, which fall into three broad groups: high-level principles, such as the OECD AI Principles[32] and UNESCO's Recommendation on the Ethics of AI;[64] risk management frameworks, such as the NIST AI Risk Management Framework[33] and ISO/IEC 42001;[71] and regulations, such as the EU AI Act.[20] Others draw on traditional financial metrics, such as capital expenditure and return on investment. Such alignment should help establish a common language, making AI adoption easier to understand and compare across companies.

The indicators are **not tied to any risk materiality framework**, as no comprehensive or generally accepted one yet exists for AI. The same indicators are applied to all 150 companies in this report, with equal weighting, whatever their materiality to each company. In reality, materiality differs by sector. For example, certain AI risks may be significantly more material to a financial institution than to a consumer goods company, while environmental impact or supply-chain dependencies may be more important elsewhere. The framework does not adjust for these sector differences. For this reason, the company scores listed in the Appendix are best compared with sector peers rather than across the whole sample.

Data Sources

The disclosures assessed include information extracted from six groups of company sources: financial reports (annual and quarterly reports and regulatory filings, such as Forms 10-K, 10-Q and 20-F); proxy statements and annual general meeting (AGM) materials; sustainability and ESG reports; investor communications (earnings call transcripts, investor day presentations and executive interviews); corporate policies and governance documents (codes of conduct, board and committee charters, privacy notices, supplier codes and human rights statements); and company websites, including AI and responsible AI pages, blogs, and press releases.

The evidence window runs from January 2025 to August 2026, attempting to capture the latest insights into companies' adoption of AI. Because AI is moving quickly from pilots into day-to-day operations and annual reporting cycles struggle to keep pace, the research draws on quarterly earnings calls and executive statements as well as annual reports.

That information was collected using AI, with human review and judgment at every step of the process. It was gathered in three ways: through Bigdata.com, an AI-powered financial research platform; by scraping public websites; and, for a small proportion, manually.

Despite careful human review, some disclosures may have been missed or misread.

Scoring

The framework includes 34 indicators, 33 of which are scored on a three-point scale of 0, 1 and 2, defined below.

Score 0: None: No relevant disclosure, evidence, policy, process, metric or action.
Score 1: Partial: The topic is acknowledged or described in general terms, but disclosure is incomplete, vague or unsupported.
Score 2: Evidenced: Clear, specific and evidenced disclosure, including defined practices, responsibilities, processes, metrics or targets where relevant. For some indicators, a score of 2 also reflects stronger practice.

The remaining indicator, Serious AI Incidents, is scored on a scale of 0 to –2, reflecting the negative nature of these outcomes. However, this indicator is reported separately and not included in company or category scores.

For each indicator, every score level is defined by its own qualifying criteria. For the top score of 2, companies are not required to meet a theoretical best-practice standard. Rather, the criteria reflect what the best companies in the sample currently report on each topic, whether in completeness, specificity, or evidence. For some indicators, the score of 2 also reflects stronger practice, not just better disclosure, so the score captures not only how much a company discloses but also what it discloses. For example, a score of 2 on the AI Provider Dependency Management indicator goes to companies that explain how they address the risk of relying on AI providers. As disclosure practice matures, the criteria and scoring will evolve with it.

Each company's total score is the sum of its scores on the 33 indicators, out of a maximum of 66, expressed on a scale of 0 to 100 for ease of comparison (total ÷ 66 × 100). A score of 100 would mean a company achieved the top score on every indicator; a score of 0, that it disclosed nothing relevant, or showed no evidence of a policy, process or action, on any of them. All 150 companies, developers and deployers alike, receive a total score (see the full list in the Appendix).

Category scores are calculated in the same way: the sum of a company's scores on the indicators in a category, as a share of the maximum possible for that category, expressed on a scale of 0 to 100. Because categories contain different numbers of indicators, this puts them on a common scale and allows them to be compared directly.

As previously mentioned, all indicators carry equal weight, whatever the sector. Equal weighting was chosen deliberately for comparability: it keeps the scores consistent across the sample of companies and avoids subjective judgments about which topics matter most for each company.

AI was integral to this study at every stage, from collecting the information through to scoring, interpretation, analysis and the production of this report. Indeed, a study of this scale would not have been feasible without AI.

How to interpret the scores

The framework captures three types of disclosure, which are worth distinguishing. The first is disclosure of intentions: strategy, policies and principles. The second is disclosure of how AI is governed in practice, such as risk controls, independent assessment and security testing. The third is disclosure of actual business outcomes, such as quantified returns on investment and serious AI incidents. A company may score well on the first while saying little about the other two.

A high total score reflects more extensive and more specific disclosure. But for many indicators, a top score also requires a company to disclose a specific practice, such as board-level oversight of AI, a named testing process, independent assurance or a formal protection for workers. Being open about weak practice does not earn a top score, so the scores capture reported practice as well as transparency.

What the scores cannot show is whether that practice works. All evidence is self-reported and unverified, and for indicators that require a figure, such as AI energy use or AI spend, a top score reflects disclosure of the number, not its size. A high score is therefore not proof of more advanced adoption, more effective risk management, greater value creation or more responsible practices, and the ranking in the Appendix should be read as an AI disclosure ranking.

Still, higher levels of disclosure may signal that AI is strategically important to the business, that adoption is becoming more advanced, that governance arrangements are becoming more formalized (as reported, if not verified), or that management is confident discussing its approach and progress. They may also reflect the fact that AI is highly material to the business and therefore creates significant risks alongside opportunities. Lower levels of disclosure may signal earlier-stage adoption, less developed governance, lower materiality, more limited investment, commercial sensitivities, a more cautious communication strategy, or simply different disclosure practices. In some cases, disclosure may also lag adoption: AI is moving quickly from pilots into day-to-day operations, and annual reporting cycles struggle to keep pace.

Differences in disclosure between peers can therefore provide useful signals rather than definitive conclusions. As AI adoption accelerates and investors focus more closely on how companies deploy AI, manage its risks and create value from it, disclosure – and, in some cases, its absence – can help identify areas for further analysis and engagement.

Regulation also plays a role. Although AI rules are developing at different speeds across jurisdictions, AI-specific rules such as the EU AI Act, together with sector-specific requirements, already oblige companies to document and disclose certain activities, and those obligations will widen as the Act's high-risk provisions take effect in 2027–28. Companies in more heavily regulated sectors therefore have a head start. This is consistent with this report's findings: technology, communication services, financials and healthcare score above other sectors, although regulation is likely one factor among several.

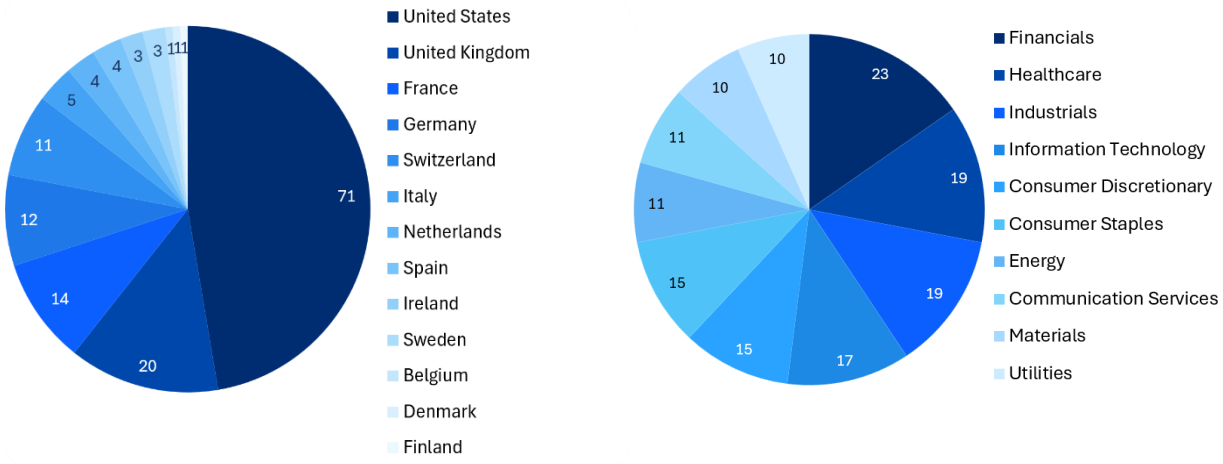
For these reasons, a company's score is most meaningful when set against those of its sector peers, which face broadly similar materiality and regulatory requirements. Such comparisons are likely to give the clearest insight into how far each company has adopted AI and how it reports managing it.

Coverage Universe

The research covers 150 of **the world's largest listed companies**, totalling US$ 60 trillion in market capitalization.[8] The sample is split almost evenly between the United States (71) and Europe (79) and **covers a range of sectors**. As some of the largest and best-resourced companies, they are also likely to be among the most advanced disclosers. A long tail of smaller companies is likely to disclose far less, so the findings of this report should not be read as representative of the wider market.

The selection aimed for sector balance, although financial services, healthcare and industrials are over-represented, as are the United Kingdom and France within the European sample. The full list of companies appears in the Appendix.

Figure 2: Company breakdown by country and by sector (number of companies)



AI developers versus AI deployers

Generative AI is still in the early stages of adoption, and market participants have yet to settle on a common vocabulary for classifying companies by their main role in the AI ecosystem. In practice, the boundaries are blurred: many companies operate across several layers of the AI value chain, and their roles are shifting quickly.

For investors, a useful starting point is to distinguish between companies that build AI and those that use it, commonly referred to as AI developers and AI deployers, respectively.

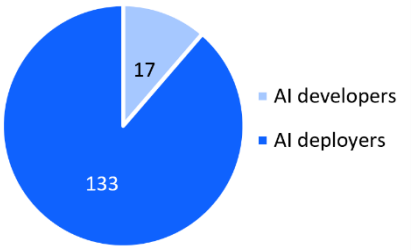
AI developers typically refer to companies that develop foundation and other general-purpose AI models, such as OpenAI, Anthropic, DeepSeek and Mistral, none of which are included in this report. For the purposes of this report, AI developers include companies that build AI models and those that provide the infrastructure underpinning them, sometimes referred to as AI enablers. The AI infrastructure providers include cloud and data-center operators (Microsoft, Alphabet, and Amazon, all of which also develop general-purpose AI models) as well as AI hardware providers, which supply the chips (semiconductor companies such as NVIDIA, AMD and Broadcom), chip-making equipment (such as ASML), and networking equipment powering AI (such as Cisco and Nokia).

AI deployers, meanwhile, are companies that primarily use AI models in their operations and products, with use cases ranging from everyday productivity tools that summarize meetings or draft documents to AI built into how a company designs, makes and delivers its products and services. How companies deploy AI varies considerably, depending on their business models, technical capabilities and infrastructure needs. For example, Siemens and ABB use third-party AI models to build AI systems that they then sell to their own customers. Some deployers, including large banks such as JPMorgan, Bank of America and BNP Paribas, have invested in their own data centers to run AI.

The EU AI Act[20] cuts across these layers, using its own legal terminology to define the roles of the different actors in the AI value chain, including providers, which develop AI systems or place them on the market, and deployers, which use them, with obligations that depend on the type and risk of the system or model.

Of the 150 companies in the sample, the large majority (133) are AI deployers, while 17 are AI developers.

Figure 3: Company breakdown by AI value chain role (number of companies)

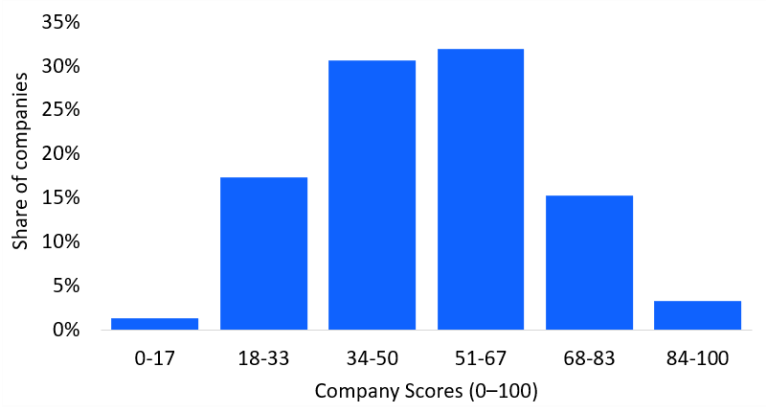


Findings

A wide spread of AI disclosure

While every company in the sample uses AI, how much they disclose about how they use it varies widely, even among the largest listed companies. This is reflected in the wide spread of total scores across the 150 sampled companies, which ranges from 9 to 98 of the 100 points available. Almost half (49%) of the companies score below the sample's average of 52.

Figure 4: Companies distributed by total scores



Note: scores are out of 100 points available. Total scores cover 33 of the 34 indicators. Serious AI Incidents is reported separately.

As Figure 5 shows, Microsoft leads on AI disclosure with 98 points, falling only two points short of the maximum of 100. It is followed by Alphabet (91), with Amazon and SAP tied on 86. At the other end, Glencore scores only 9 points, with Ferrari (12) and Valero Energy (18) just above it. The bottom of the table is more mixed sector-wise, ranging from materials and energy companies to consumer and financial companies. A low score does not necessarily mean a company is not using AI, only that it discloses little about how it uses AI or the risks involved.

Figure 5: Companies with highest and lowest scores on AI disclosure

TOP 10	Sector	Score
Microsoft	Information Technology	98
Alphabet Class A	Communication Services	91
Amazon	Consumer Discretionary	86
SAP SE	Information Technology	86
Cisco	Information Technology	85
Visa	Financials	83
BNP Paribas SA	Financials	82
Meta	Communication Services	82
JPMorgan Chase	Financials	77
Schneider Electric SE	Industrials	77

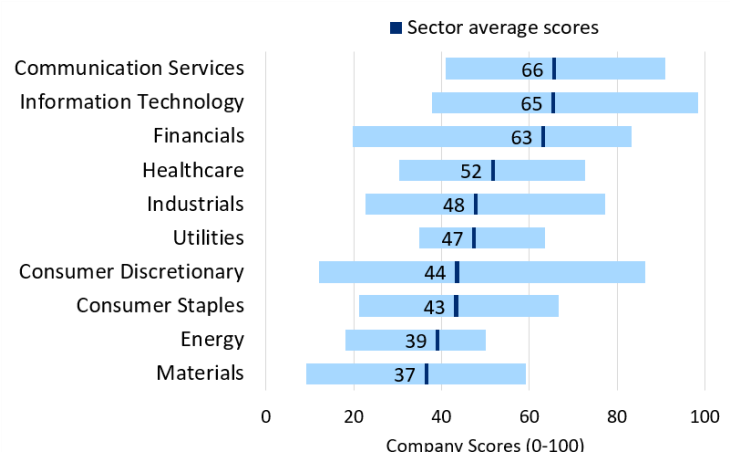
BOTTOM 10	Sector	Score
Glencore plc	Materials	9
Ferrari	Consumer Discretionary	12
Valero Energy	Energy	18
Investor AB	Financials	20
Anheuser-Busch InBev	Consumer Staples	21
TJX Companies	Consumer Discretionary	23
BAE Systems	Industrials	23
Berkshire Hathaway B	Financials	24
Atlas Copco AB	Industrials	24
Costco	Consumer Staples	26

Note: Scores are shown out of 100 (66 points available across 33 indicators scored 0–2, rescaled).

Sector analysis

AI disclosure differs markedly by sector, both in average scores and in the spread of scores among sector peers.

Figure 6: Spread of AI disclosure scores by sector



Note: Scores are shown out of 100 (66 points available across 33 indicators scored 0–2, rescaled). Some sectors have only around eight companies, so one company can move an average noticeably.

Information Technology and Communication Services are the two sectors disclosing the most on AI. This should not come as a surprise. As the builders of AI and those best positioned to benefit from its opportunities, they face the greatest scrutiny. Returns from AI have so far been concentrated among these companies, and their valuations already price in high expectations. Investors want evidence that heavy, rising AI spending will pay off, while customers need to trust the products they buy. Both push these companies to disclose more.[21]

Financials and Healthcare rank third and fourth on AI disclosure, reflecting both their greater AI adoption and the fact that they are heavily regulated. Banks have been early adopters of technology for credit scoring, fraud detection and trading activities, uses that were subject to regulatory model risk requirements long before generative AI arrived. In healthcare, the benefits of AI are equally well documented in drug discovery and increasingly in care delivery itself. From August 2028, the latter will be classified as high-risk under the EU AI Act, which adds obligations on data governance, human oversight and post-market monitoring.

Materials and Energy score lowest on AI disclosure. Lighter reporting, however, shouldn't necessarily be interpreted as lower adoption. Both sectors run AI in operations such as predictive maintenance, exploration and subsurface analysis, and process optimization. But they tend to report it as digitalization, automation or efficiency rather than as AI, which leaves it outside the disclosure one can see. Read that way, these sectors may not yet have treated AI as a governance or risk subject. More transparency would be welcome.

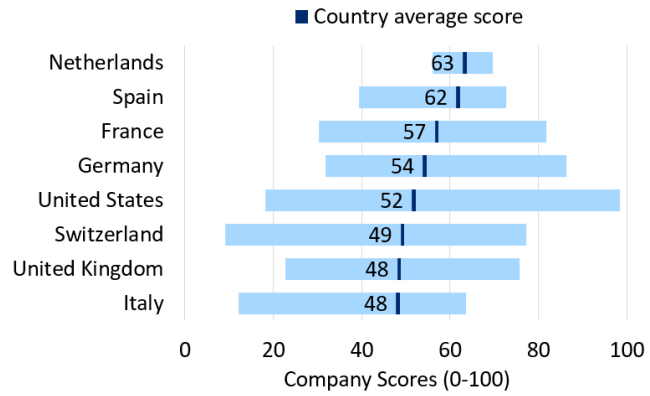
Wide dispersion within sectors means sector peers provide the closest like-for-like comparison of a company's approach to AI, while the whole sample shows where it stands overall. The spread of scores may reflect differences in the extent of AI adoption, the strategic importance of AI to the business, investment levels, innovation ambitions, governance maturity, organizational culture, or risk appetite. It may also indicate how willing a company is to communicate its AI strategy, implementation and outcomes, because of factors including competitive sensitivity, perceptions of materiality, overall disclosure philosophy, regulatory considerations, and legal constraints.

Geographic comparison

Just as sectors can help explain how much a company is currently disclosing on AI, so may where it is headquartered and where it operates. **The EU AI Act and EU sustainability reporting rules** push European companies towards more structured disclosure, while US companies report largely through **SEC risk factors** and voluntary publications. This section looks at whether those differences show up in how much companies disclose, although the country-level results rest on very small samples and should be read as indicative only.

Among the countries shown, the Netherlands and Spain, with four companies each, post the highest average scores, of 63 and 62 out of 100 points respectively, ahead of France (57) and Germany (54). The United States, with 71 companies, shows the widest spread, from 18 to 98 points: it is home to many of the leading disclosers and a large share of the laggards.

Figure 7: Spread of AI disclosure scores by country



Note: Countries shown are the eight with four or more companies in the sample: Spain 4, the Netherlands 4, France 14, Germany 12, the United States 71, Switzerland 11, the United Kingdom 20 and Italy 5.

Broad category analysis

The analysis so far has focused on the volume and spread of disclosure: how much 150 of the largest companies in the US and Europe disclose, and how widely this varies across companies, sectors and countries. This section turns to what these companies disclose, using the framework's 34 indicators to assess their disclosure through the lens of ESG risks and value creation.

Figure 8 provides an overview of the nine broad indicator categories, ranked by average score. Each category rolls up the 150 companies' scores on its underlying indicators.

Figure 8: Average score per category



Note: Each category score is the sum of a company's scores on the indicators in a category, as a share of the maximum possible for that category, expressed on a scale of 0 to 100.

Overall, the ranking shows that two **governance-related categories, AI Policies and AI Board Oversight,** are the framework's most disclosed areas. This reflects recent efforts by many companies to strengthen the parts of AI governance that are easiest to put in place and to disclose. Having policies and oversight is different from showing that the governance in place actually works. The ranking makes this distinction clear: **Risk Controls** comes only 7th, suggesting that fewer companies provide evidence of operational controls and day-to-day governance. Meanwhile, more companies seem to have incorporated AI in their **Cybersecurity & Data Privacy** policies and assessments, making this the fourth most frequently addressed category. This may reflect greater familiarity with these risks, partly due to existing regulatory frameworks, such as the GDPR in the EU.

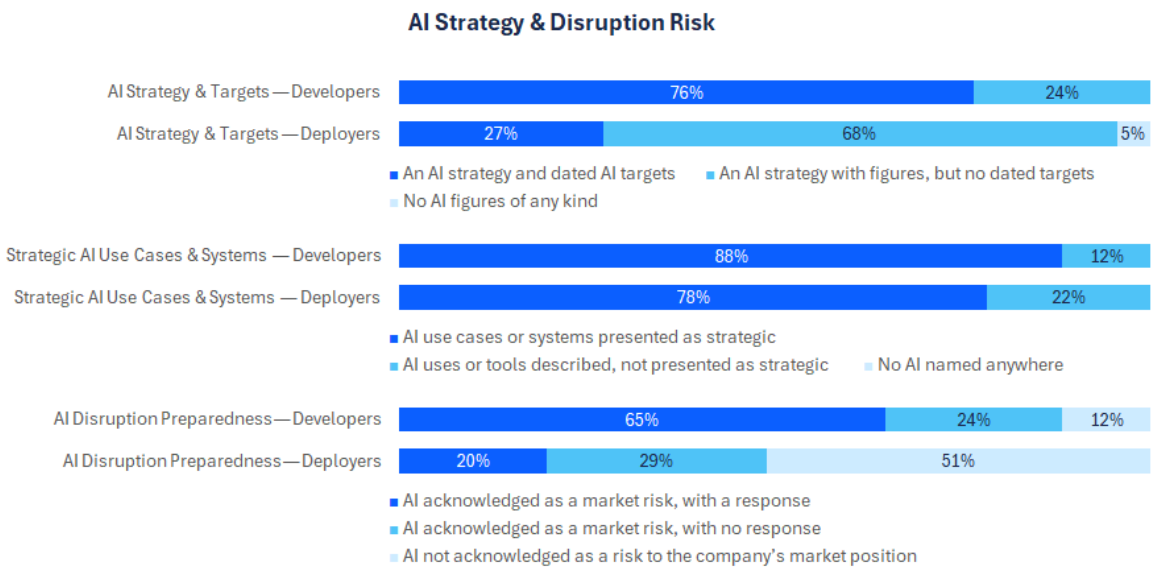
Strategy & Disruption Risk ranks joint second. Even though almost every company in the sample has an AI strategy, fewer share targets, and the question of disruption risk behind the strategy often goes unaddressed. A related gap appears in **AI Financials** (6th in the ranking), which remain opaque, as few companies are able to fully quantify the financial benefits of AI.

Meanwhile, **Workforce** disclosure (5th) is uneven. Companies are more forthcoming about how much their employees use AI than about how they use it or its impact on them. Finally, **Environmental Impact** is largely unreported. For now, it is material only for AI developers and deployers running data centers.

Individual indicator analysis

This section goes beneath the nine broad categories to examine each of the 34 individual indicators, comparing how developers and deployers disclose on each.

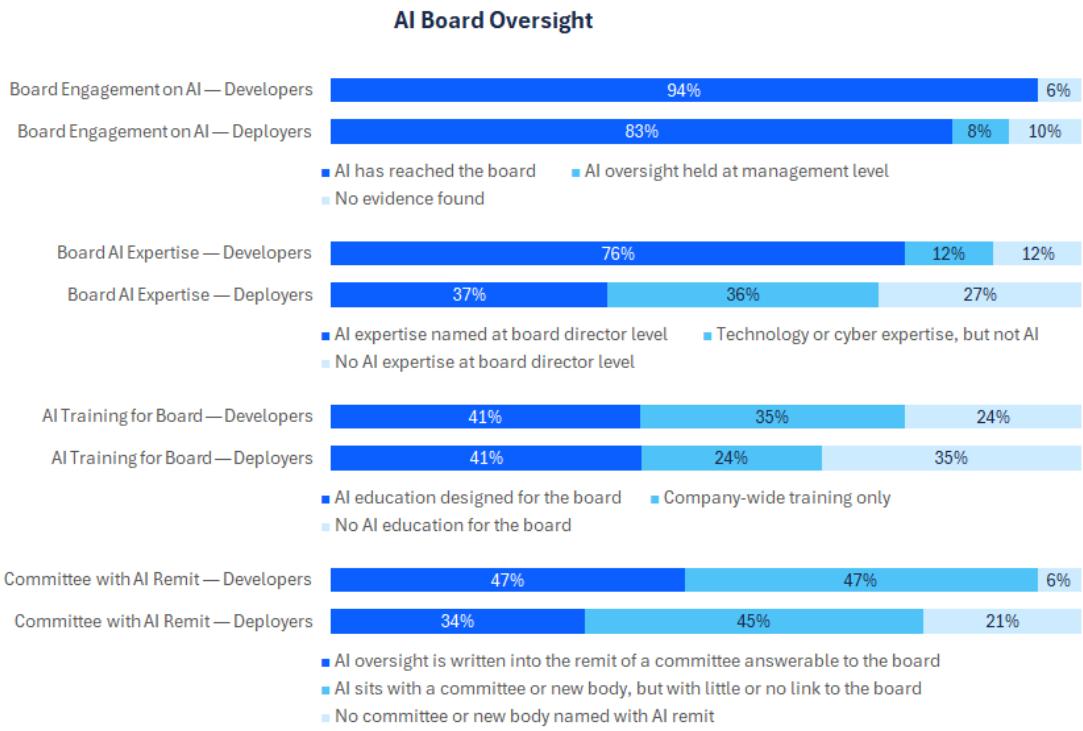
Figure 9: AI Strategy & Disruption Risk



AI Strategy & Disruption Risk. Every company in the sample says something publicly about the role of AI in its business, but developers, unsurprisingly, go into far more detail than deployers. While 76% of developers pair a defined **AI strategy with time-bound targets**, only 27% of deployers do so. All deployers describe where AI is used in the business today, and 78% present specific AI use cases or systems as strategic and say what value or productivity gains they deliver or are expected to deliver. The remaining 22% are less specific: they describe AI uses, name tools that are not presented as strategic (often employee assistants), or cite a count of AI use cases. A count of use cases says little about which carry the most risk or value. Naming systems will matter more from December 2027, when the EU AI Act will require providers of high-risk AI systems to register them in a public EU database, by name and intended purpose.

As well as being an opportunity, AI can threaten business models, and investors need to know whether a company has recognized the threat and how it plans to respond. Most companies mention AI risk, but usually as a risk from AI they use themselves, AI-enabled cyberattacks, regulation or workforce effects. About half (49%) of deployers and 88% of developers describe **a threat to their market position**: AI substituting for what they sell, lowering barriers for new competitors, squeezing prices, inserting itself between them and their customers, or rivals making faster or more effective progress in adopting AI. Fewer (20% of deployers and 65% of developers) set out **a response** alongside that threat. Among the companies that do, the most common moves are investing in AI capabilities and building AI into products, partnering with AI developers, or launching a broader strategic overhaul or new multi-year plan. Several financial firms argue that client relationships and scale make their core business resistant to AI. Some responses remain thin, amounting to little more than a commitment to "stay ahead", so even a stated response is not always evidence of a plan.

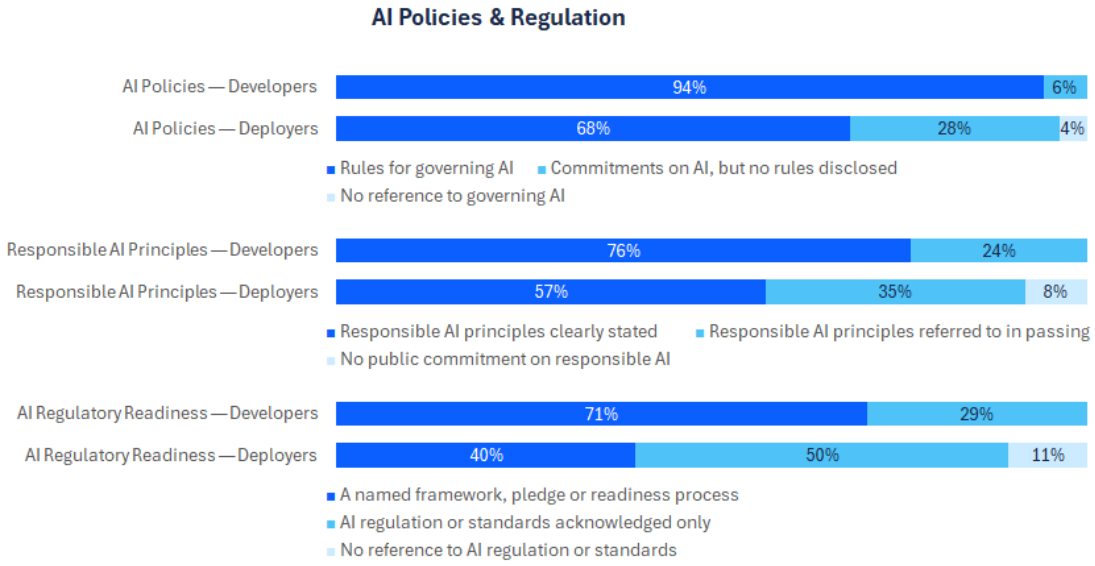
Figure 10: AI Board Oversight



AI Board Oversight. AI has reached the board of almost every company in the sample: 94% of developers and 83% of the deployers show **board engagement with AI strategy and risk**, reflecting broad recognition that AI has moved from a management concern to a board-level governance priority. Expertise lags, though, especially among deployers: 76% of the developers but only 37% of the deployers name **AI expertise** in their disclosure on directors. Broader technology, digital and cybersecurity skills remain more common on boards. Some companies are seeking to narrow the gap through **board-specific AI training**, disclosed by 41% of both developers and deployers. The rapid evolution of AI means that even current expertise can become outdated quickly. For investors, the question is therefore not only whether boards have AI expertise today, but whether they have embedded mechanisms for continuously developing and refreshing the board's expertise as the technology, its regulation and its applications evolve.

Meanwhile, among deployers, the structures to support board oversight are often unclear: only 34% of deployers clearly disclose **a committee, whether existing or new, with an AI remit** that reports to the board. This can be a board committee whose charter or stated responsibilities include AI risk, ethics, or governance, or a management-level AI forum with a disclosed reporting line to the board. For 45% of the deployers, AI sits with a committee or body, but with little or no link to the board. The remaining 21% don't report having a committee with an AI remit. Reporting lines and oversight structures matter. Rather than taking claims about AI principles and oversight at face value, investors should establish which committee holds the AI remit and how often it reports to the board.

Figure 11: AI Policies & Regulation



AI Policies & Regulation. Almost every company in the sample says something about governing AI: every developer and almost all (96%) of the deployers disclose principles, policies or standards related to AI. The developers tend to go into more detail on their commitments and how they are applied, compared to the deployers.

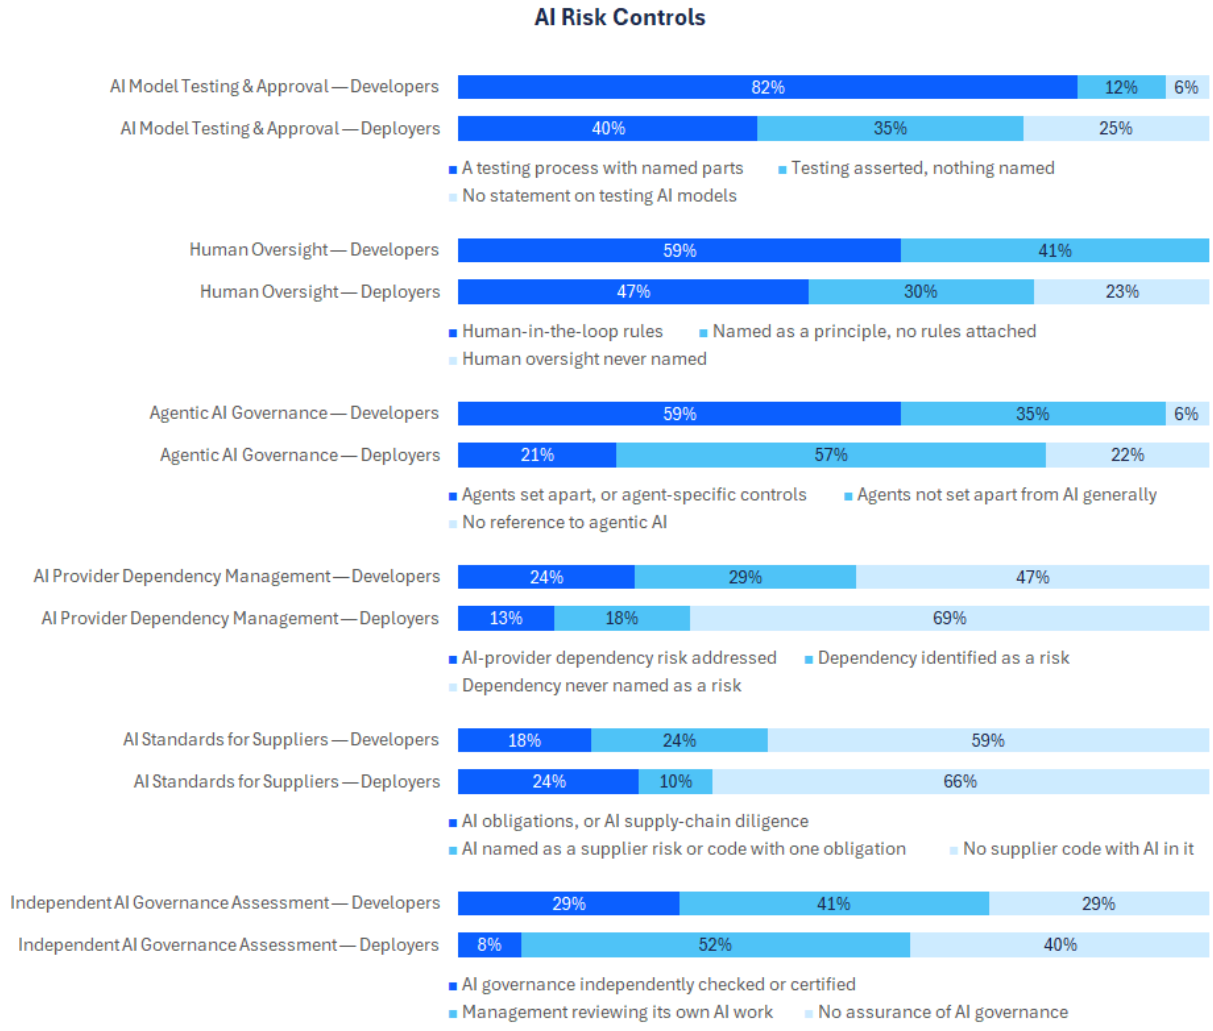
Three quarters (76%) of developers and 57% of deployers explain their **Responsible AI principles** in detail either in a standalone report, on a dedicated webpage, or in a section of a sustainability report. A minority (24% of developers and 35% of deployers) mention in passing that they have responsible AI principles, or simply say they aim to use AI responsibly. AI principles are high-level statements about how AI should be used safely and responsibly. They reflect the organization's core values and provide aspirational guidance for employees to follow and for public awareness, and they also guide the wider governance program. By sector, Technology and Communication Services set out the most detailed Responsible AI principles, a way to build trust with their users. Detailed Responsible AI principles seem to matter less to companies in other sectors, including financial institutions, which emphasize their risk management systems instead.

AI-related policies are the governance machinery: they set out how a company's AI principles are put into practice. They typically cover the roles and responsibilities, the processes (including risk management and assessment processes), the regulations that apply, the related policies (such as data privacy and cybersecurity) and the consequences of not following the rules. Here the sector order changes. Financials lead because AI policy extends what they already have: supervisory model-risk rules such as SR 11-7 and the PRA's SS1/23. Energy and Materials typically refer to AI in general terms without naming a policy, which points to a gap in disclosure rather than necessarily in adoption.

AI regulatory readiness: Two fifths (40%) of the deployers in the sample show a readiness process for an AI law (typically the EU AI Act), have adopted a well-established framework (such as the NIST AI RMF, the OECD AI Principles, UNESCO's Recommendation on the Ethics of AI, and ISO/IEC 42001), or signed an AI pledge (such as the EU AI Pact and the White House Voluntary AI Commitments). These offer shared reference points for transparency, human oversight, risk management and accountability. Among deployers, financial and healthcare companies lead on this indicator, reflecting their exposure to both existing and future regulatory requirements. Banks and insurers using AI for credit scoring or for pricing life and health insurance will fall within the EU AI Act's high-risk framework from December 2027, adding to the model-risk rules they already follow. In healthcare, AI used in medical devices, including clinical decision support, will fall within the same framework from August 2028. Materials, Energy and Consumer Staples trail. Their companies tend to cite AI regulation as a risk without naming a law or standard they apply. This may partly reflect lighter exposure: their main AI uses, such as predictive maintenance, process optimization and demand forecasting, mostly fall outside the high-risk framework. The main exception is AI used in hiring and managing staff.

Generally speaking, regulatory readiness matters because AI rules are evolving quickly but unevenly across jurisdictions. This raises compliance costs for global companies and exposes those that fall behind to fines, litigation and reputational damage. Companies that engage openly with regulators are better placed to anticipate new rules and build trust, and investors may want to see that engagement reflected in their disclosures.

Figure 12: AI Risk Controls



AI Risk Controls. Most companies in the sample say they assess the AI foundation models and systems they use and keep humans in the loop, but the detail varies sharply. Developers are twice as likely to explain **how testing works**: 82% report setting out the framework they test against and who signs off the results, against only 40% of deployers. Testing an AI model means checking, before and after it goes live, that it performs as intended and does not produce biased or harmful outcomes. The strongest disclosures combine a risk or impact assessment, bias and fairness testing, approval by a review body before release, and monitoring once in use.

All developers and 77% of deployers use some form of **human-in-the-loop** language, but only 59% and 47%, respectively, provide examples of the roles these individuals play. When it comes to human oversight, the real question is whether a person merely signs off at the end of a process or genuinely shapes the decision, a difference that grows as AI takes over more of the analysis that used to inform the decision.

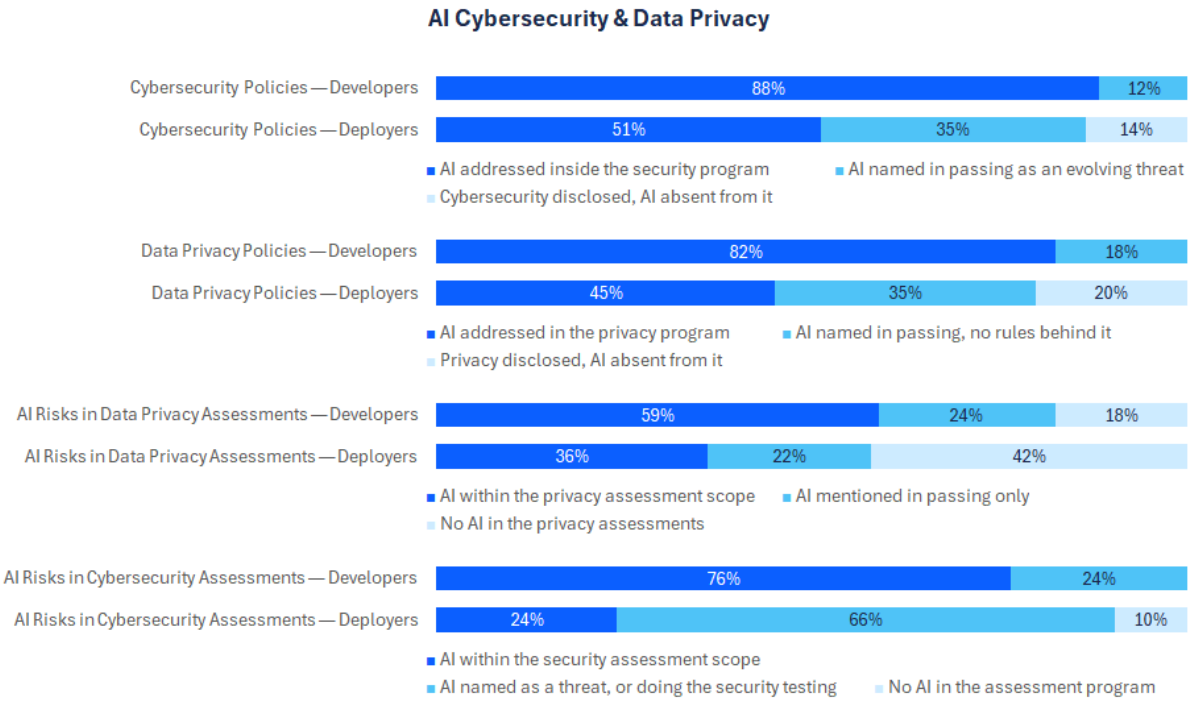
Autonomous AI agents raise questions a general AI policy does not answer, such as how far an agent may act on its own and who can stop it, yet only 21% of deployers report governing agents as something distinct, against 59% of developers. Agents bring risks of their own: they can act beyond their intended limits, pursue goals that drift from policy, and, when working together, pass one agent's error through the others before a person notices.[22,23] Human oversight remains essential, although reviewing every step of an agent's work may not be practical at scale. For investors, the gap in agent-specific governance matters because regulators are starting to expect it. Singapore's Model AI Governance Framework for Agentic AI, published in January 2026 as the first government guidance written specifically for agentic AI, asks organizations to limit what an agent may do and to make a person accountable for it, while the EU AI Act requires human oversight of high-risk AI systems to include the ability to intervene and stop them.

Few companies extend their AI standards beyond their own operations to address AI-related practices across their supply chains. Almost every company has a supplier code, but only 18% of developers and 24% of deployers place specific **AI-related obligations on suppliers** or carry out AI-related due diligence on them.

Meanwhile, only about half (53%) of developers and 31% of deployers say anything about their **dependency on AI providers**, and 24% and 13%, respectively, report managing that dependency. Most companies rely on a short list of counterparties: a handful of foundation-model providers (with limited visibility of how these models behave) and one or two cloud providers for running models. There is a risk in this situation that a failure, outage or policy change at any one of them could affect hundreds of downstream companies at once, turning a company-level dependency into a systemic risk. How closely this dependency is scrutinized depends on the sector. Financial services is the clearest case where supervisors already treat it as a resilience issue: the EU Digital Operational Resilience Act (DORA) regulates concentration on third-party providers for EU financial entities. But DORA covers information and communication technology (ICT) in general, with no AI-specific provisions, and it requires firms to assess this risk rather than disclose it.

Finally, **independent assurance of AI governance** is rare. Only 29% of developers and 8% of deployers report an independent review of their AI governance, whether by internal audit, external assurance or certification. 41% of developers and 52% of deployers simply rely on management reviewing its own work, while the rest don't report any assessment of their AI governance.

Figure 13: AI Cybersecurity & Data Privacy

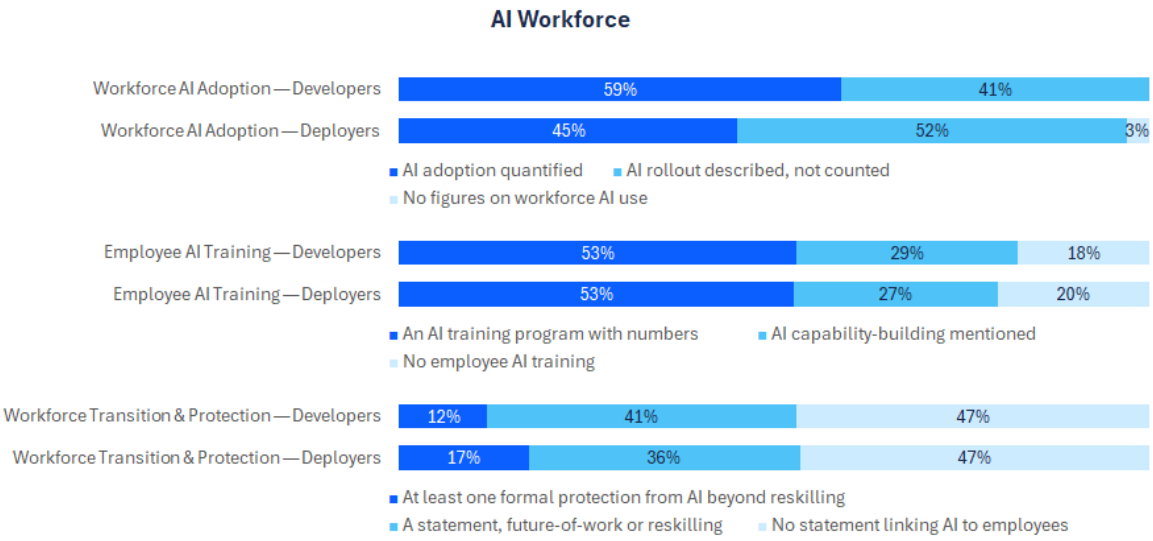


AI Cybersecurity & Data Privacy. For companies, cybersecurity and data privacy are among the AI-related risks most likely to translate into financial loss through breaches, regulatory penalties and reputational damage. AI can increase this exposure by opening new points of attack and extending the flow of sensitive data across models, systems and third-party providers. While conventional IT security and privacy controls remain essential, they may need to be adapted to address these AI-specific risks.

Most companies in the sample acknowledge these risks: all developers and at least 80% of deployers report having **cybersecurity and/or data privacy policies** that cover AI. Far fewer, however, demonstrate how these policies are operationalized and tested, particularly among deployers: 36% disclose that AI risks are incorporated into their **data privacy assessments**, while only 24% do so for **cybersecurity assessments**. Many organizations already have established data privacy programs designed to meet data protection laws, including the EU General Data Protection Regulation (GDPR). These provide an important foundation but may not fully address the specific risks arising from AI development and deployment. Best practice is to integrate AI-specific privacy

safeguards into existing processes. This could include clear restrictions on the personal data that may be submitted to AI models, and controls over the use of personal data in model training.

Figure 14: AI Workforce



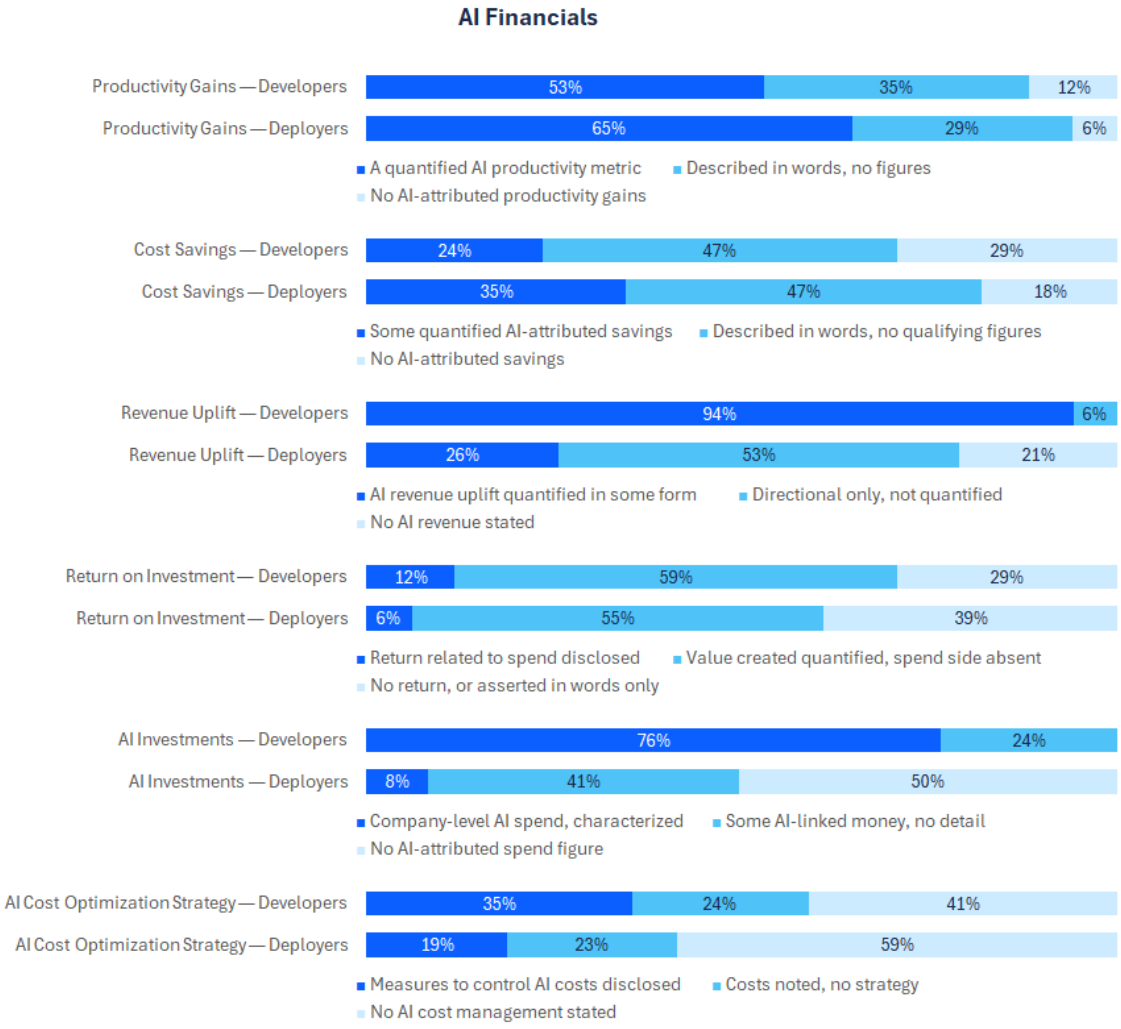
AI Workforce. The workforce is where the productivity gains promised by AI meet their social cost: badly managed adoption can lead to skills gaps, labor disputes and reputational harm. Virtually every company in the sample reports on the extent of **AI adoption across its workforce**. Almost half (45%) of deployers put a number on it. Most report access (licenses deployed or the share of the workforce with a tool). Some report both what has been handed out and what is being used (daily active users, the share of employees actively using AI tools, query volumes, or take-up by role or function).

The vast majority of companies (82% of developers and 80% of deployers) also report providing **AI training** to their employees, though the scope and quality of these vary widely. Article 4 of the EU AI Act, applicable since February 2025 and amended in July 2026, requires providers and deployers of AI systems to take measures to support the development of AI literacy among their staff and others involved in operating or using AI systems on their behalf. This puts the quality of the AI training programs many companies report under the spotlight. Given how quickly AI is evolving, one-off training is unlikely to be enough: employees need opportunities to update their knowledge and skills as technologies, applications and risks change.

Training and upskilling prepare workers to use AI, but they do not tell them what AI will mean for their jobs. Far fewer companies in the sample explain **how workers are protected** as AI is deployed. Only about half (53%) of both developers and deployers disclose something on this, such as redeployment or reskilling of staff whose roles AI changes, a ban or negotiated limit on AI monitoring of employees, a grievance or whistleblowing channel the company says is open for AI concerns, a requirement that a person review AI-influenced decisions about employees, or consultation with worker representatives before deployment. Some of these protections are statutory in certain jurisdictions and discretionary in others. None of them addresses a less visible risk: that relying on AI may gradually erode workers' own expertise and judgment. Companies would do well to monitor this, particularly in sectors that depend on creative and expert work, where it may become the most material social risk of AI adoption.[24] The risk extends to the next generation of workers. As AI takes over the routine tasks through which graduates traditionally learned their trade, companies are hiring fewer of them, which could leave a thinner pipeline of experienced staff in the years ahead.[25]

The limited disclosure on worker protection suggests that most companies have yet to explain publicly how AI may reshape the roles and responsibilities of the workers expected to use it. If this reflects gaps in companies' internal practices, it may also raise regulatory concerns for those deploying AI in specific employment and worker-management functions covered by the EU AI Act. These include recruitment, certain employment decisions, task allocation and worker monitoring or evaluation. The relevant high-risk rules are scheduled to apply from December 2027.

Figure 15: AI Financials

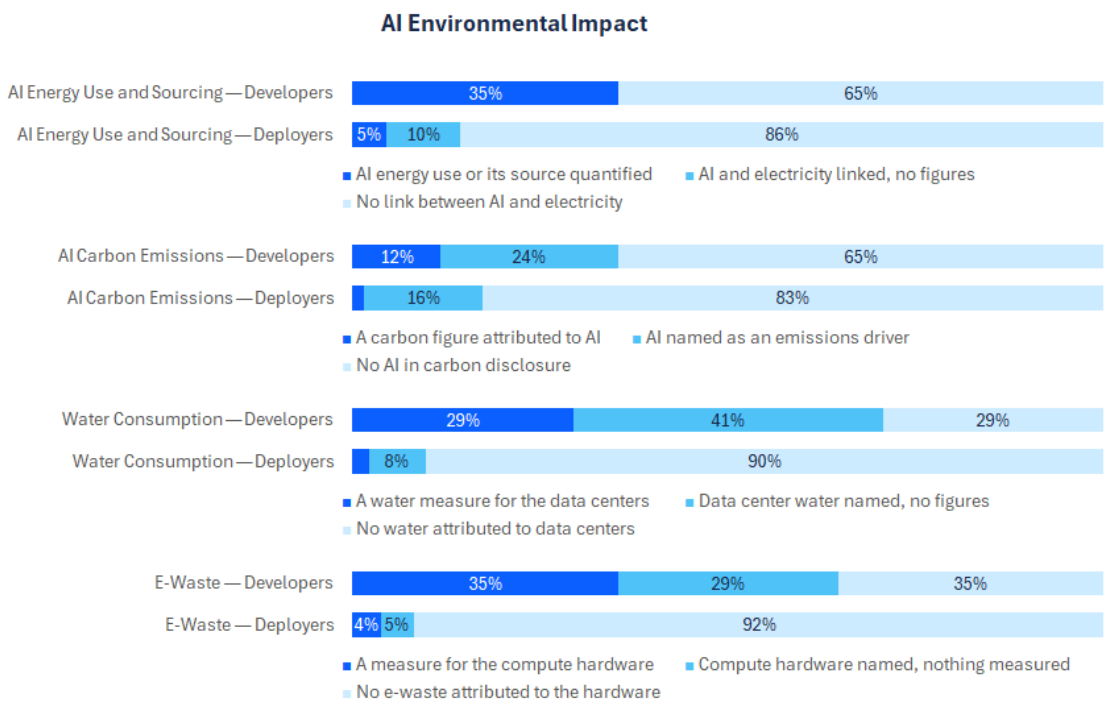


AI Financials. For companies and investors alike, the key question is whether AI spending pays off. Productivity offers one way to test whether AI investment is delivering operational value. Almost two thirds of deployers (65%) quantify **productivity gains**. The figures tend to be task- or process-specific rather than firm-wide (coding time halved, a claims process made X% more productive, X hours a week freed for each user of an internal assistant). But productivity gains are not the same as financial value creation. AI frees up capacity, but whether that capacity becomes revenue growth, redeployed effort or lower costs is a decision that few companies talk about. **Cost savings** follow a similar pattern: 82% of deployers mention them, but only 35% put a number on them.

Developers, led by hyperscalers such as Microsoft and Alphabet, disclose more about the scale of their AI business: 76% provide detailed information on AI investment, while 94% disclose **revenue attributed to AI**. Deployers are far less transparent, with only 8% and 26%, respectively, doing the same. This is mostly a gap in attribution rather than inactivity: for most deployers, AI spend sits inside a wider technology or transformation program in which it is one line among many. Meanwhile, only 12% of developers and 6% of deployers disclose a realized or expected **return on investment**. More (59% of developers and 55% of deployers) put a number on the value AI creates in parts of their business without setting it against what it cost. ROI is often considered difficult to calculate. Some would even argue that it is too early to put a figure on it. Some of the most important AI investments are made to build long-term capabilities rather than to generate immediate returns. Perhaps more useful is disclosure that separates strategic bets from use cases that should already pay off and explains when AI investments are scaled up or stopped.

Finally, about a third of developers (35%) and 19% of deployers report having a **strategy for managing AI costs**. This matters because the cost of running AI models (inference) grows with usage: a company that treats it simply as a necessary expense is exposed to rising costs precisely as adoption succeeds. A credible strategy applies discipline to AI input costs through levers such as model sourcing and routing, model sizing, workload placement and, for companies running their own hardware, infrastructure optimization.

Figure 16: AI Environmental Impact



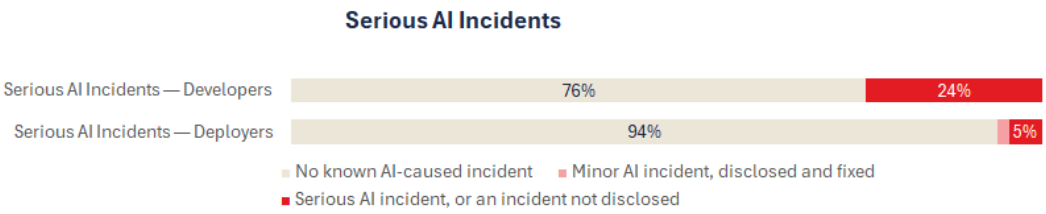
AI Environmental Impact. The energy, water, and hardware that AI consumes to train the models, answer queries and cool the hardware[26] are most material for developers, which run the data centers where that environmental footprint sits. But they are also relevant to certain deployers such as large banks, insurers, pharmaceutical companies and industrial firms that operate their own data centers. For other deployers, AI's footprint may look small enough to be immaterial today, but it still adds to their value-chain emissions and will grow as adoption scales, so these companies and their shareholders should not ignore the topic.

So far, disclosure on AI's environmental impact remains thin, even among developers. Depending on the indicator, only 12% to 35% of developers disclose a quantified metric, compared with no more than 5% of deployers. Where developers do disclose, they mostly report rates rather than totals: a renewable share for their data centers, liters of water per kilowatt-hour, or the share of hardware reused or recycled. No company in the sample reports **total energy use**, **total carbon emissions**, or **total water consumption** for its AI, so investors cannot tell how fast the environmental footprint of AI is growing as adoption scales. Yet the direction of travel is clear. On electricity alone, the International Energy Agency projects[27] that electricity demand from data centers worldwide will more than double between 2024 and 2030 to around 945 terawatt-hours (TWh), roughly the entire electricity consumption of Japan, with AI the main driver of this increase.

Regulators are starting to respond. In September 2026, the European Commission proposed[28] a mandatory sustainability rating and label for EU data centers above 500 kW, covering energy efficiency, water use, energy sourcing and waste-heat reuse, with the first labels due from August 2027. Once in place, the labels should make the footprint of the facilities that run AI far more visible and give deployers a firmer basis for asking their providers for the data and disclosing it. In the meantime, tools already exist[29,30] for companies that want to estimate the carbon and water footprint of their AI use, so not reporting a figure is increasingly a matter of choice rather than capability. Even where the footprint is immaterial, saying so with a number would give investors more comfort than silence.

Finally, investors should also bear in mind the social dimension of data centers, whose impact extends beyond the environment to the communities that host them. Some proposed developments face local opposition, while others are welcomed as a source of jobs and investment in regions where traditional industries are in decline. The key question is whether companies can secure a social license to operate by demonstrating tangible benefits for local communities and addressing concerns over resource use, jobs and local infrastructure.

Figure 17: Serious AI Incidents



Serious AI Incidents. Because this indicator captures harm rather than disclosure, it is the only indicator in the framework that is reported separately and is not included in any category or company score.

Serious AI incidents have been rare in the sample but more common among AI developers. 24% of developers had a serious incident caused by AI in the past couple of years, whether or not they disclosed it. Among deployers, 5% had a serious AI incident in the period.

Developers' incidents have stemmed from their own AI products: a voice assistant that recorded private conversations, AI news summaries that produced false headlines, regulatory investigations into AI chatbots and into a driver-assistance system, and security concerns over an AI feature that captures users' on-screen activity. Deployers' incidents involved AI used in decisions affecting people: a hiring chatbot that exposed applicants' records, and litigation over an AI tool allegedly used to wrongfully deny patients access to care.

In most cases, incidents came to light through litigation, regulatory action or press coverage rather than companies' own disclosures. The indicator therefore depends on what is publicly known: incidents that have not been disclosed by companies or reported by regulators or the press will not be captured. For this reason, the indicator is also excluded from the disclosure scores.

Conclusion

AI adoption has outpaced the disclosure investors need to assess it. Almost all of the 150 companies covered in this report disclose that they use AI, have an AI policy and expect it to create value. Far fewer explain how AI is governed, how much it costs, what net value it is expected to generate, how it affects their workforce or what environmental footprint it leaves. Investors are being asked to back a transformation they cannot yet properly evaluate.

Because practice is still forming, investors have a chance to shape it by engaging with companies where disclosure is thinnest within each sector and by asking for the same core information across sectors so that disclosures can be compared and best practice shared. For companies, the framework offers an early view of what investors and other stakeholders are likely to expect them to disclose as their use of AI grows.

The framework in this report is a first step towards that common ground. Investors say they most want guidance on AI risk management best practice and sector-level research on materiality.[31] Covering 150 of the largest companies across sectors, this analysis could serve as a basis for an AI materiality map or a common investor framework and is meant to evolve as practice matures and regulation, notably the EU AI Act, takes effect.

Several areas deserve further research: how far disclosure reflects what companies actually do in practice; whether stronger AI disclosure is linked to better financial or share-price outcomes; how to separate risks specific to a company from the system-wide risks that AI developers and policy makers increasingly warn about; and how legal liability for AI harms is shared between developers and deployers.

The author welcomes feedback on the framework and engagement from investors, companies and policy makers interested in developing it further.

Appendix – 150 Companies with AI Disclosure Scores

Rank	Company	Sector	Country	Market Cap ($bn)	Developer/ Deployer	Total Score (out of 100)
1	**Microsoft**	Information Technology	United States	3,808	AI Developer	98
2	**Alphabet Class A**	Communication Services	United States	4,137	AI Developer	91
3	**Amazon**	Consumer Discretionary	United States	2,677	AI Developer	86
3	**SAP SE**	Information Technology	Germany	252	AI Deployer	86
5	**Cisco**	Information Technology	United States	429	AI Developer	85
6	**Visa**	Financials	United States	681	AI Deployer	83
7	**BNP Paribas SA**	Financials	France	131	AI Deployer	82
7	**Meta**	Communication Services	United States	1,849	AI Developer	82
9	**JPMorgan Chase**	Financials	United States	886	AI Deployer	77
9	**Schneider Electric SE**	Industrials	France	195	AI Deployer	77
9	**UBS Group AG**	Financials	Switzerland	152	AI Deployer	77
12	**Deutsche Telekom AG**	Communication Services	Germany	159	AI Deployer	76
12	**HSBC**	Financials	United Kingdom	345	AI Deployer	76
12	**AXA SA**	Financials	France	98	AI Deployer	76
15	**Bank of America**	Financials	United States	396	AI Deployer	74
15	**Nokia**	Information Technology	Finland	57	AI Developer	74
15	**Orange S.A.**	Communication Services	France	44	AI Deployer	74
18	**Banco Bilbao Vizcaya Argentaria**	Financials	Spain	159	AI Deployer	73
18	**GSK plc**	Healthcare	United Kingdom	101	AI Deployer	73
20	**AstraZeneca**	Healthcare	United Kingdom	249	AI Deployer	71
20	**Banco Santander**	Financials	Spain	213	AI Deployer	71
20	**Mastercard**	Financials	United States	482	AI Deployer	71
20	**NVIDIA**	Information Technology	United States	5,561	AI Developer	71
24	**Advanced Micro Devices**	Information Technology	United States	1,005	AI Developer	70
24	**Allianz SE**	Financials	Germany	184	AI Deployer	70
24	**Prosus N.V.**	Consumer Discretionary	Netherlands	86	AI Deployer	70
24	**Siemens AG**	Industrials	Germany	254	AI Deployer	70
28	**Micron Technology**	Information Technology	United States	1,239	AI Developer	68
29	**Apple**	Information Technology	United States	4,821	AI Deployer	67
29	**ASML**	Information Technology	Netherlands	696	AI Developer	67
29	**Intel**	Information Technology	United States	634	AI Developer	67
29	**L'Oréal SA**	Consumer Staples	France	238	AI Deployer	67
29	**UnitedHealth**	Healthcare	United States	328	AI Deployer	67
29	**Zurich Insurance**	Financials	Switzerland	106	AI Deployer	67
29	**SpaceX**	Communication Services	United States	2,143	AI Developer	67
36	**Novartis**	Healthcare	Switzerland	270	AI Deployer	65
36	**Palo Alto Networks**	Information Technology	United States	343	AI Deployer	65
36	**AT&T Inc.**	Communication Services	United States	152	AI Deployer	65
39	**Iberdrola S.A.**	Utilities	Spain	161	AI Deployer	64
39	**Novo Nordisk A/S**	Healthcare	Denmark	165	AI Deployer	64
39	**UniCredit SpA**	Financials	Italy	148	AI Deployer	64
39	**Vodafone**	Communication Services	United Kingdom	40	AI Deployer	64
43	**Amgen**	Healthcare	United States	225	AI Deployer	62
43	**Infineon Technologies AG**	Information Technology	Germany	90	AI Developer	62
43	**Unilever PLC**	Consumer Staples	United Kingdom	131	AI Deployer	62
43	**Intesa Sanpaolo**	Financials	Italy	118	AI Deployer	62
47	**Goldman Sachs**	Financials	United States	264	AI Deployer	61
47	**Roche**	Healthcare	Switzerland	350	AI Deployer	61
47	**Sanofi SA**	Healthcare	France	107	AI Deployer	61
47	**Verizon Communications**	Communication Services	United States	192	AI Deployer	61
47	**Walmart**	Consumer Staples	United States	830	AI Deployer	61
47	**ING Groep**	Financials	Netherlands	95	AI Deployer	61
47	**Barclays**	Financials	United Kingdom	77	AI Deployer	61
54	**Linde**	Materials	Ireland	213	AI Deployer	59
54	**Rolls-Royce**	Industrials	United Kingdom	163	AI Deployer	59
54	**RTX**	Industrials	United States	249	AI Deployer	59
54	**Thermo Fisher Scientific**	Healthcare	United States	246	AI Deployer	59
58	**Duke Energy**	Utilities	United States	89	AI Deployer	58
58	**Enel SpA**	Utilities	Italy	101	AI Deployer	58
58	**Safran SA**	Industrials	France	163	AI Deployer	58
58	**Citigroup**	Financials	United States	217	AI Deployer	58

Rank	Company	Sector	Country	Market Cap ($bn)	Developer/ Deployer	Total Score (out of 100)
62	**Airbus SE**	Industrials	Netherlands	177	AI Deployer	56
62	**GE Vernova**	Industrials	United States	263	AI Deployer	56
62	**Nestlé**	Consumer Staples	Switzerland	244	AI Deployer	56
62	**Booking Holdings**	Consumer Discretionary	United States	124	AI Deployer	56
66	**ABB Ltd**	Industrials	Switzerland	180	AI Deployer	55
66	**GE Aerospace**	Industrials	United States	331	AI Deployer	55
66	**Johnson & Johnson**	Healthcare	United States	633	AI Deployer	55
66	**Applied Materials**	Information Technology	United States	402	AI Developer	55
70	**Home Depot**	Consumer Discretionary	United States	284	AI Deployer	53
71	**Broadcom**	Information Technology	United States	1,640	AI Developer	52
71	**EssilorLuxottica SA**	Consumer Discretionary	France	80	AI Deployer	52
71	**L'Air Liquide SA**	Materials	France	127	AI Deployer	52
71	**Netflix**	Communication Services	United States	285	AI Deployer	52
71	**Tesla**	Consumer Discretionary	United States	1,399	AI Deployer	52
71	**RELX**	Industrials	United Kingdom	63	AI Deployer	52
77	**Lilly**	Healthcare	United States	1,025	AI Deployer	50
77	**Morgan Stanley**	Financials	United States	308	AI Deployer	50
77	**Shell**	Energy	United Kingdom	273	AI Deployer	50
77	**Spotify Technology S.A.**	Communication Services	Sweden	101	AI Deployer	50
77	**Williams Companies**	Energy	United States	84	AI Deployer	50
77	**Palantir Technologies**	Information Technology	United States	502	AI Deployer	50
83	**Chevron**	Energy	United States	401	AI Deployer	48
83	**Coca-Cola**	Consumer Staples	United States	370	AI Deployer	48
83	**National Grid plc**	Utilities	United Kingdom	76	AI Deployer	48
83	**Munich Re**	Financials	Germany	74	AI Deployer	48
83	**E.ON SE**	Utilities	Germany	50	AI Deployer	48
88	**ExxonMobil**	Energy	United States	674	AI Deployer	47
88	**Merck & Co**	Healthcare	United States	367	AI Deployer	47
88	**Southern Company**	Utilities	United States	96	AI Deployer	47
88	**Wells Fargo**	Financials	United States	253	AI Deployer	47
88	**The Boeing Company**	Industrials	United States	150	AI Deployer	47
93	**Eni SpA**	Energy	Italy	81	AI Deployer	45
93	**Siemens Energy AG**	Industrials	Germany	142	AI Deployer	45
93	**TotalEnergies SE**	Energy	France	206	AI Deployer	45
96	**ENGIE SA**	Utilities	France	71	AI Deployer	44
96	**LVMH**	Consumer Discretionary	France	225	AI Deployer	44
96	**Philip Morris International**	Consumer Staples	United States	300	AI Deployer	44
96	**DHL Group (Deutsche Post AG)**	Industrials	Germany	70	AI Deployer	44
96	**Reckitt Benckiser Group PLC**	Consumer Staples	United Kingdom	43	AI Deployer	44
101	**Procter & Gamble**	Consumer Staples	United States	334	AI Deployer	42
102	**BP PLC**	Energy	United Kingdom	112	AI Deployer	41
102	**Walt Disney**	Communication Services	United States	187	AI Deployer	41
102	**Pfizer Inc.**	Healthcare	United States	161	AI Deployer	41
102	**Holcim Ltd**	Materials	Switzerland	42	AI Deployer	41
106	**Anglo American plc**	Materials	United Kingdom	57	AI Deployer	39
106	**British American Tobacco p.l.c.**	Consumer Staples	United Kingdom	119	AI Deployer	39
106	**Inditex**	Consumer Discretionary	Spain	193	AI Deployer	39
106	**Gilead Sciences**	Healthcare	United States	187	AI Deployer	39
106	**PepsiCo**	Consumer Staples	United States	172	AI Deployer	39
106	**Tesco PLC**	Consumer Staples	United Kingdom	41	AI Deployer	39
112	**Compagnie Financière Richemont**	Consumer Discretionary	Switzerland	127	AI Deployer	38
112	**Lam Research**	Information Technology	United States	399	AI Developer	38
112	**VINCI SA**	Industrials	France	66	AI Deployer	38
112	**STMicroelectronics N.V.**	Information Technology	Switzerland	47	AI Developer	38
116	**Freeport-McMoRan**	Materials	United States	100	AI Deployer	36
116	**NextEra Energy**	Utilities	United States	158	AI Deployer	36
116	**Danaher Corporation**	Healthcare	United States	153	AI Deployer	36
116	**RWE AG**	Utilities	Germany	50	AI Deployer	36
120	**Constellation Energy**	Utilities	United States	92	AI Deployer	35
120	**Compass Group PLC**	Consumer Discretionary	United Kingdom	51	AI Deployer	35
120	**Lonza Group AG**	Healthcare	Switzerland	48	AI Deployer	35
123	**AbbVie**	Healthcare	United States	459	AI Deployer	33
123	**Newmont**	Materials	United States	122	AI Deployer	33
123	**Altria Group, Inc.**	Consumer Staples	United States	120	AI Deployer	33
123	**Haleon PLC**	Healthcare	United Kingdom	41	AI Deployer	33
127	**Caterpillar**	Industrials	United States	380	AI Deployer	32
127	**CRH**	Materials	Ireland	54	AI Deployer	32
127	**McDonald's**	Consumer Discretionary	United States	164	AI Deployer	32
127	**Rio Tinto plc**	Materials	United Kingdom	158	AI Deployer	32

Rank	Company	Sector	Country	Market Cap ($bn)	Developer/ Deployer	Total Score (out of 100)
127	**BASF SE**	Materials	Germany	50	AI Deployer	32
127	**Deere & Company**	Industrials	United States	167	AI Deployer	32
127	**Mercedes-Benz Group AG**	Consumer Discretionary	Germany	42	AI Deployer	32
134	**Abbott Laboratories**	Healthcare	United States	172	AI Deployer	30
134	**Hermès International**	Consumer Discretionary	France	151	AI Deployer	30
136	**Marathon Petroleum**	Energy	United States	133	AI Deployer	29
136	**Phillips 66**	Energy	United States	112	AI Deployer	29
138	**ConocoPhillips**	Energy	United States	153	AI Deployer	27
138	**Diageo**	Consumer Staples	United Kingdom	48	AI Deployer	27
138	**Eaton Corporation plc**	Industrials	Ireland	167	AI Deployer	27
141	**Costco**	Consumer Staples	United States	406	AI Deployer	26
142	**Berkshire Hathaway B**	Financials	United States	1,071	AI Deployer	24
142	**Atlas Copco AB**	Industrials	Sweden	102	AI Deployer	24
144	**BAE Systems**	Industrials	United Kingdom	71	AI Deployer	23
144	**TJX Companies**	Consumer Discretionary	United States	153	AI Deployer	23
146	**Anheuser-Busch InBev**	Consumer Staples	Belgium	152	AI Deployer	21
147	**Investor AB**	Financials	Sweden	124	AI Deployer	20
148	**Valero Energy**	Energy	United States	125	AI Deployer	18
149	**Ferrari**	Consumer Discretionary	Italy	68	AI Deployer	12
150	**Glencore plc**	Materials	Switzerland	86	AI Deployer	9

Note: This table is for illustrative purposes only. The information used to arrive at the scores was extracted from six groups of company sources: financial reports (annual and quarterly reports and regulatory filings, such as Forms 10-K, 10-Q and 20-F); proxy statements and annual general meeting (AGM) materials; sustainability and ESG reports; investor communications (earnings call transcripts, investor day presentations and executive interviews); corporate policies and governance documents (codes of conduct, board and committee charters, privacy notices, supplier codes and human rights statements); and company websites, including AI and responsible AI pages, blogs, and press releases. Some disclosures may have been missed or misread despite careful review. Scores reflect information available when each company was reviewed, between June and August 2026.

References and Further Resources

1 In this report, 'AI' refers to 'Generative AI'.

2 Gartner, Gartner Forecasts Worldwide AI Spending to Grow 47% in 2026

3 *"Organizational adoption reached 88% in 2025"*. Stanford Institute for Human-Centered Artificial Intelligence, The 2026 AI Index Report

4 *"Nearly nine in ten respondents report regular use of AI in at least one business function"*. McKinsey & Company, The state of AI in 2026: On the road to ROI

5 Partnership on AI (2025), Disclosure of AI-related impacts, risks, and opportunities

6 Gartner, Hype Cycle for Artificial Intelligence

7 Thomson Reuters Foundation and UNESCO (2026), Responsible AI in Practice: 2025 Global Insights from the AI Company Data Initiative

8 Yahoo! Finance market capitalization data as of 2 October 2026.

9 McKinsey & Company, The AI transformation manifesto: 12 themes driving growth

10 Amazon Web Services, Reimagine: Turning AI into Value

11 *"Autonomous agents can work together, circumvent production security controls, and successfully attack hardened production environments"*. CNBC, OpenAI releases sweeping report on Hugging Face AI agent hack

12 *"The company did not notify Services Australia of the incident until three months after the fact, via email"*. ABC News, OpenAI hacked Medicare portal, Prime Minister Anthony Albanese says

13 Evident, Evident AI Index Banks

14 *"The uneven pace of adoption, governance maturity and operational readiness suggests AI will increasingly contribute to the strengthening or weakening"* of financial institutions' creditworthiness. S&P Global Ratings, AI in Financial Institutions Survey

15 *"The 2026 proxy season illustrates the extent to which shareholders continue to seek greater AI-related oversight, accountability, transparency, and board engagement in AI-related risks and opportunities"*. Harvard Law School Forum on Corporate Governance, Governing AI: Corporate Oversight and Shareholder Engagement

16 *"In 2025, nearly a quarter (24%) of S&P 500 companies disclosed frameworks and policies around AI"*. ISS STOXX, Artificial Intelligence and Governance: Is 2026 a Tipping Point for Turning Awareness into Action?

17 *"From 2023 to 2025, the share of S&P 500 companies disclosing AI as a risk jumped from 12% to 83%"*. The Conference Board, In the S&P 500, Disclosure of AI Risks Surges from 12% to 83%

18 Nasdaq, Nasdaq Lens: The Disclosure Report 2026

19 One indicator, Serious AI Incidents, is scored on a negative scale (-2 to 0). Because it captures harm rather than disclosure, it is the only indicator in the framework that is reported separately and is not included in any category or company average calculation.

20 European Commission, Regulation (EU) 2024/1689 (Artificial Intelligence Act)

21 Morgan Stanley Research, Mapping AI's Rate of Change: Where the Rubber Meets the Road

22 Deloitte Insights, Agentic AI Risks in Banking

23 IBM, IBM Study: Eight in Ten EMEA Enterprises Escalate AI Incidents to Board or C-Suite as AI Agents Test Governance

24 Baillie Gifford, AI isn't coming for your job. It's coming for your mind

25 SignalFire, The SignalFire State of Tech Talent Report - 2025

26 Alston, G., Love, N. and Haider, R., Data center cooling choices shift water impacts across the grid: An integrated water-energy model for sustainable data center development

27 International Energy Agency, Energy and AI

28 European Commission, Press release IP/26/1667: proposed sustainability rating and label for data centres

29 GitHub, claude-carbon (carbon estimation tool)

30 ClimateAligned, AI Footprint Calculator

31 Oxford Martin AI Governance Initiative (2026), The Role of Investors in AI Governance

32 OECD, AI Principles

33 NIST, AI Risk Management Framework

34 International AI Safety Report, International AI Safety Report 2026

35 World Economic Forum, Responsible AI Playbook for Investors

36 The EU Artificial Intelligence Act, Up-to-date developments and analyses of the EU AI Act

37 California's Senate Bill 53, The Transparency in Frontier Artificial Intelligence Act

38 CSIRO and Alphinity, Responsible AI and ESG: A Framework for Investors

39 Responsible AI Collaborative, AI Incident Database

40 OECD.AI, AI Incidents and Hazards Monitor

41 OECD.AI, Catalogue of Tools & Metrics for Trustworthy AI

42 Candriam, AI Through the ESG Lens

43 Norges Bank Investment Management, Responsible Artificial Intelligence

44 Future of Work Investors' Expectations on Ethical AI in Human Capital Management

45 Nuveen, A sustainable investor's guide to AI

46 Blueprint for Better Business (2026), A Framework for Using AI

47 Railpen, Achieving effective AI governance, Navigating AI: an investor framework for assessing exposure and governance

48 Carmignac, Responsible AI: Building Trust and Long-Term Value in the Age of Algorithms

49 Allianz Global Investors, Responsible AI

50 Intentional Endowments Network, Responsible Tech Resources (list of guides and frameworks)

51 ILPA (2023), Responsible AI Quick Guide for Asset Owners

52 EqualAI (2026), Board Playbook

53 EqualAI, Algorithmic Impact Assessment tool

54 Institute for AI Policy and Strategy, AI Agent Governance: A Field Guide

55 MIT, AI Risk Repository

56 MIT AI Risk Repository, AI Incident Tracker

57 MIT AI Risk Repository, AI Risk Mitigations

58 MIT AI Risk Repository, AI Governance

59 Responsible AI Institute, Responsible AI Institute

60 UNESCO Global AI Ethics and Governance Observatory, Readiness Assessment Methodology

61 AVERI, Frontier AI Auditing-Related Legislation in the US

62 SHARE, Investor Advocacy on Artificial Intelligence: A New Resource

63 Responsible Investment Association Australasia, Artificial Intelligence and Human Rights: Investor Toolkit

64 UNESCO, Recommendation on the Ethics of Artificial Intelligence

65 UBS Asset Management, The Dark Side of AI

66 TD Asset Management, Governance of AI

67 Scottish Widows, Governing the Algorithm: Investor Priorities for Responsible AI

68 The Unchained Banker, 10 Thoughts About Corporate AI

69 Sustainable AI Group, Sustainable AI Group

70 Pereira, Graylin and Brynjolfsson, Stanford Digital Economy Lab, The Enterprise AI Playbook: Lessons from 51 Successful Deployments

71 ISO/IEC 42001:2023, Information technology – Artificial intelligence – Management system