Standard Intelligence

White paper

Enterprise Shift

Full workforce transformation in the agentic age

PUBLISHED

February 2026

WEB

standardintelligence.com

AUDIENCE

For boards and C-suite leading workforce change

CONTENTS

Summary 3

1. Agentic AI is a workforce question before it is a technology question 4

2. Efficiency gained without governance is fragile, which makes governance a commercial question 5

3. Binding rules differ, but employment AI already has specific, dated anchors 6

4. Three principles govern the shift, and each needs the other two 7

5. Roles, management, learning and compliance each need a structural shift 8
   - 5.1 Job descriptions give way to capability portfolios 8
   - 5.2 Hierarchical control gives way to distributed accountability 8
   - 5.3 Training events give way to continuous capability building 9
   - 5.4 Compliance obligation gives way to ethical commitment 9

6. Governance starts at the board, because the board cannot delegate judgement about AI 9

7. Governance needs a structure, named roles and a budget sized to the task 10
   - 7.1 Five layers, each with named roles 10
   - 7.2 Governance should take a larger share of the AI budget than benchmarks show 11
   - 7.3 The structure scales with the organisation 11

8. Roles, training and careers must change with the work 12
   - 8.1 Four role archetypes need different oversight 12
   - 8.2 Training runs on three pillars, at depths set by role 12
   - 8.3 Performance management and careers must credit governance work 13

9. Governance holds only when it is someone's job, with consequences attached 13

10. Employees must know how AI affects them and have a voice in its use 14

11. Trust is built in sequence, and psychological safety keeps the monitoring honest 14

12. Intellectual property and AI costs need governance of their own 15
   - 12.1 Intellectual property 15
   - 12.2 AI costs 15

13. Aim for at least Level 3 of five maturity levels, and govern at speed 16

14. Report governance performance to the board each quarter, in four domains 17

15. The Standard Intelligence platform is designed to support both sides of the shift 18

Where to start, and what it costs 18

Basis and Standing 19

Sources 20

Summary

Agentic AI is a workforce question before it is a technology question. By agentic AI we mean systems that act towards a goal within a mandate, initiating and completing steps rather than only answering a question. Such systems change who decides, who is accountable and what roles mean, yet many organisations are adopting them as they adopted earlier software, by buying the tool and measuring productivity.

Efficiency gained without governance is fragile. A recruitment system that speeds up hiring while systematically disadvantaging some groups of candidates can cost more in penalties, litigation, reputational damage and lost staff than it saves. We expect organisations that build transparency, fairness and accountability into their AI strategy to outperform those that add governance later; that is a reasoned expectation, not a measured result.

Binding rules differ, but employment AI already has specific, dated anchors. Under the AI Act, AI used for recruitment, for decisions on terms of work, promotion and termination, for task allocation and for monitoring workers is high-risk under Annex III point 4, with obligations applying from 2 December 2027 (subject to Article 111(2) for systems already on the market). Inferring workers' emotions from biometric data has been prohibited since 2 February 2025. In the UK, Articles 22A to 22D of the UK GDPR have governed significant decisions taken without meaningful human involvement since 5 February 2026, and national works-council rights, such as Germany's, may already apply to workplace AI.

Three principles govern the shift, and each needs the other two. Transparency makes AI's role in decisions visible to the people affected, fairness requires equitable outcomes monitored continuously, and accountability keeps people responsible for outcomes even where execution is delegated.

Roles, management, learning and compliance each need a structural shift. Job descriptions give way to capability portfolios, hierarchical control to distributed accountability, training events to continuous capability building, and compliance obligation to ethical commitment.

Governance starts at the board, because the board cannot delegate judgement about AI. That judgement needs fluency. Each C-suite role carries its own fluency requirement and governance accountability, and the CFO decides whether governance is funded at all.

Governance needs a structure, named roles and a budget sized to the task. We recommend five governance layers and a governance allocation of 15 to 20% of the AI budget for large enterprises, rising to 25 to 30% for smaller organisations. That is our recommended target, well above the low single-digit shares that published benchmarks report.

Roles, training and careers must change with the work. Different role archetypes need different oversight, and training should run on three pillars (tool fluency, governance fluency, and reporting and process fluency) at depths set by role. Article 4 of the AI Act

already requires providers and deployers to take measures to support the AI literacy of their staff.

Governance holds only when it is someone's job, with consequences attached. Governance that depends on goodwill will not survive commercial pressure. We recommend tying 15 to 25% of variable pay in governance-critical roles to governance outcomes, with governance gates on every bonus and a check that the incentives produce the behaviour they are meant to reward.

Employees must know how AI affects them and have a voice in its use. Every employment deployment needs a communication plan, and the people affected hold information about a system's impact that monitoring data cannot capture.

Trust is built in sequence, and people report what goes wrong only when it is safe to do so. Leadership actions taken in the right order earn trust, and where employees are punished for raising AI concerns, the monitoring data that governance depends on become unreliable.

Intellectual property and AI costs need governance of their own. Ownership, protectability and confidentiality of AI-assisted work are often unsettled in practice, and agentic workflows make AI spending harder to forecast.

Most organisations seeking initial compliance should aim for at least Level 3 of five maturity levels, and govern at speed. Risk-proportionate review lets an organisation adopt AI quickly while avoiding both governance paralysis and governance abdication.

Governance performance should be reported to the board each quarter, in four domains. An organisation that cannot report governance in figures will find it much harder to show that its governance works.

The first step is an inventory of every AI system that touches employment decisions. Named owners, board fluency and baseline training follow in the first 90 days; structure, budget, incentives and quarterly measurement come after.

This paper states the position as at 11 October 2026. Basis and Standing, at the end, sets out which propositions are binding law, which are guidance, and which are our own recommended controls.

1. Agentic AI is a workforce question before it is a technology question

AI agents can now be deployed to draft contracts, screen candidates, monitor performance, generate code and coordinate logistics with little human involvement in each step. We use "agentic AI" for systems that act towards a goal within a mandate, initiating and completing steps rather than only answering a question.

A system that drafts is a tool. A system that decides and executes raises questions a tool never did: who is accountable for what it does, what oversight fits, and how roles change when routine judgement is delegated to software.

Many organisations are adopting these systems as they adopted earlier software, by procuring the tool, training a few people and measuring productivity. That approach governs the technology and leaves the work around it ungoverned.

The work itself changes. In our view, AI agents compress management chains, because frontline employees gain direct access to analysis and context that previously had to pass through several layers of management, and the case for some of those layers weakens. Management has to be rebuilt around what AI does not do well: building trust, resolving ambiguity, exercising ethical judgement in novel situations and shaping the culture.

How organisations govern this human side of the change will decide which of them benefit from it.

2. Efficiency gained without governance is fragile, which makes governance a commercial question

Gains from ungoverned AI can be clawed back. A recruitment system that speeds up hiring but systematically disadvantages some groups of candidates, a performance tool that erodes trust through surveillance, or a customer service agent that harms vulnerable customers can each cost more, through regulatory penalties, litigation, reputational damage and staff attrition, than the efficiency it delivered.

Transparency makes failure visible while it is still small. A transparent system's failures can be seen and corrected; an opaque system can fail silently and accumulate harm that is costly to remediate. An organisation with working monitoring and escalation can deploy AI with more confidence, because problems are more likely to be found early.

Fair systems produce better data. A system that discriminates distorts the outcomes it produces, and those outcomes become the organisation's data. Recruitment AI that screens out qualified candidates on protected characteristics narrows the talent pool, and performance AI that favours some groups over others produces misleading productivity data.

The commercial case is an expectation, not a measured result. We expect organisations that build transparency, fairness and accountability into their AI strategy to outperform those that treat governance as an afterthought. Building governance into strategy also makes compliance easier to sustain, because the governance then works between audits as well as during them (section 5.4).

3. Binding rules differ, but employment AI already has specific, dated anchors

Binding AI rules have not converged. The EU has a horizontal AI Act, the United States has no comprehensive federal AI statute, the UK has no AI-specific statute and relies largely on its existing regulators, and other jurisdictions have taken their own routes.

What the main regimes and frameworks share is a set of principles. The AI Act, the NIST AI Risk Management Framework (which describes trustworthy AI as, among other things, "accountable and transparent" and "fair with harmful bias managed"), UK sectoral regulation and many frameworks elsewhere address transparency about AI use, fairness in outcomes, accountability for decisions, human oversight of consequential systems and protection of the people affected. An organisation built on these principles will find most requirements easier to meet, but it still has to identify and evidence each regime's specific obligations.

For employment AI (recruitment, performance management, workforce planning and worker monitoring), those obligations are specific and dated.

High-risk classification. Annex III point 4 of the AI Act classes as high-risk the AI systems intended to be used for the recruitment or selection of natural persons, and those intended to make decisions affecting terms of work-related relationships, promotion or termination, to allocate tasks based on individual behaviour or personal traits or characteristics, or to monitor and evaluate the performance and behaviour of workers. A system listed in Annex III can fall outside the high-risk regime under Article 6(3), but never where it performs profiling of natural persons, as many workforce tools do.

Application dates. The high-risk obligations apply from 2 December 2027 (Article 113, as amended by Regulation (EU) 2026/1744, the Digital Omnibus on AI). Under Article 111(2) as amended, a high-risk system placed on the market or put into service before then is caught only if its design changes significantly from that date, and systems intended for use by public authorities must comply by 2 August 2030.

Deployer duties for employers. From the same date, a deployer must use a high-risk system in line with its instructions for use and assign human oversight to people with "the necessary competence, training and authority, as well as the necessary support" (Article 26(1) and (2)). An employer must also inform workers' representatives and the affected workers before a high-risk system is put into service or used at the workplace (Article 26(7)).

The emotion recognition prohibition. Since 2 February 2025, Article 5(1)(f) has prohibited the use of AI systems to infer the emotions of a natural person in the workplace, except for medical or safety reasons. The Commission's guidelines on prohibited practices (C(2025) 5052 final, 29 July 2025) limit the prohibition to inferences from biometric data (para 245), so sentiment inferred from written text falls outside it.

Penalties. Breach of a prohibition can attract fines of up to EUR 35 million or 7% of total worldwide annual turnover, whichever is higher (Article 99(3)), and breach of a deployer's

Article 26 obligations up to EUR 15 million or 3% (Article 99(4)(e)). For SMEs, including start-ups, the lower of the two figures applies (Article 99(6)).

UK automated decisions. In the UK, Articles 22A to 22D of the UK GDPR, substituted by section 80 of the Data (Use and Access) Act 2025 and in force since 5 February 2026, require safeguards for any significant decision based solely on automated processing, meaning one taken without meaningful human involvement (Articles 22A and 22C). They restrict such decisions where special category data are processed or recognised legitimate interests are relied on (Article 22B).

Works councils. National employment law adds its own requirements. In Germany, for example, the works council has co-determination rights over the introduction and use of technical devices designed to monitor the behaviour or performance of employees (section 87(1) no. 6 of the Works Constitution Act), which can reach much workplace AI.

Governing the human side of AI with the same rigour as the technology is therefore partly a legal requirement. It is also how an organisation keeps the human capabilities that AI does not supply: judgement, creativity, ethical reasoning and the ability to earn trust.

4. Three principles govern the shift, and each needs the other two

Every structure, process and practice in this framework follows from three principles, put into practice through governance mechanisms, incentives and accountability requirements.

Principle	What it requires	What happens without it
Transparency	AI systems, their purposes, capabilities and limitations, and their roles in decisions are visible and understandable to those affected: employees, managers, regulators and individuals subject to AI decisions.	Opaque failures accumulate harm. Employees cannot contest decisions, regulators cannot verify compliance, trust erodes and adoption stalls.
Fairness	AI systems produce equitable outcomes across protected characteristics. The benefits and burdens of deployment are distributed justly, affected individuals are treated with dignity, and monitoring and adjustment are continuous.	Discriminatory outcomes can lead to regulatory penalties and litigation. Data quality degrades, talent from affected groups leaves, and the social licence to operate is lost.
Accountability	People remain responsible for AI-driven decisions and outcomes. Responsibility is assigned at every level, consequences for governance failures are proportionate and consistently applied, and incentives reward ethical AI use.	Governance exists on paper but is bypassed in practice. Failures go unaddressed, and box-ticking replaces protection.

The three reinforce one another. Transparency without accountability produces disclosure without consequence; accountability without transparency produces blame without understanding; fairness without either produces aspiration without evidence.

5. Roles, management, learning and compliance each need a structural shift

The principles require four changes in how organisations conceive roles, management, learning and compliance. Each is a structural change, and incremental improvement to existing practice will not deliver it.

5.1 Job descriptions give way to capability portfolios

A traditional job description (a static list of duties, qualifications and reporting lines) does not capture agentic work. A marketing analyst who uses AI agents to generate campaign copy, segment audiences, run A/B tests and synthesise performance data is doing a different job from the one in a specification written before those tools existed, even if the title is unchanged.

A capability portfolio is a living description of a role, kept current as the role changes. It covers:

- the human capabilities the role requires (judgement, creativity, stakeholder management, ethical reasoning);
- the AI capabilities it uses, from data analysis agents to workflow automation;
- its governance responsibilities (oversight of AI outputs, escalation triggers, compliance checkpoints);
- its ethical obligations;
- how the role is expected to evolve over the next 12 to 24 months.

For every role that involves AI, the portfolio should say what human oversight is required and why, what fairness obligations apply to the AI outputs the role uses, and what the role holder must disclose to colleagues, customers and others affected. Where the role oversees a high-risk system, the portfolio is also where the organisation can record the competence, training and authority that Article 26(2) requires (section 3).

5.2 Hierarchical control gives way to distributed accountability

Where AI agents compress management chains (section 1), distributed accountability replaces the chain with a framework saying which decisions people must make, which AI agents may make on their own and which need both, and that framework should encode the organisation's ethical commitments.

As a matter of policy, we recommend that decisions affecting individual rights, determining access to employment or services, or carrying significant consequences for vulnerable people always involve meaningful human judgement. That goes further than the law, which permits some solely automated decisions subject to safeguards (section 3).

5.3 Training events give way to continuous capability building

AI-related skills can date within months; prompting techniques learned in January may be superseded by July as models and ways of working change. Annual training programmes cannot keep pace, so learning has to be part of daily work, with the currency of skills measured continuously rather than course completion tracked once a year.

Capability building must cover ethical reasoning and governance alongside technical skill. An employee who uses AI tools well but cannot recognise when they produce unfair outcomes adds risk as well as capability.

5.4 Compliance obligation gives way to ethical commitment

Compliance-driven governance tends towards box-ticking: documentation written to satisfy a requirement and never read, bias audits to the minimum standard, oversight that exists on paper and is routinely bypassed, and disclosures that are accurate but incomprehensible. An organisation run that way can satisfy regulators during audits and still fail between them.

An ethically committed organisation designs its AI systems for equitable outcomes from the start and treats the audit as verification. Its governance runs continuously, which makes compliance easier to demonstrate, though not automatic.

6. Governance starts at the board, because the board cannot delegate judgement about AI

The board cannot delegate its understanding of AI, and still less its ethical judgement about it. Capital allocation, market strategy, risk appetite and organisational design are increasingly shaped by AI capabilities that change quickly. A board without fluency in them is likely to misjudge what proposed strategies can deliver, underestimate deployment risk and lose the ability to hold technical leadership to account for ethical outcomes.

Fluency at board level does not mean understanding transformer architectures. It means enough understanding to ask searching questions, compare competing proposals, judge risk and recognise a recommendation that is technically unsound or ethically problematic.

Each C-suite role has its own fluency requirement and governance accountability.

C-suite role	Core fluency requirement	Key governance accountability
CEO	Understanding of AI capabilities; investment and risk judgement; signalling through visible AI use and AI-informed questions.	Signals that AI fluency is a leadership requirement; holds executives accountable for AI-driven outcomes.
CFO	Valuing AI investments with uncertain payoffs; quantifying AI-specific financial risks, including exposure to AI Act fines (section 3); ensuring governance functions are adequately funded.	Executive sponsor for AI cost governance (section 12.2); ensures every AI business case includes explicit governance cost lines.

C-suite role	Core fluency requirement	Key governance accountability
CPO / CHRO	Capability portfolio design; distributed accountability frameworks; continuous capability building; cultural governance.	Architect of the new employment relationship; presents the governance performance dashboard to the board each quarter (section 14).
General Counsel	Interpreting AI regulation across jurisdictions; negotiating AI procurement contracts; advising on intellectual property in AI-generated work; evolving AI liability exposure.	Executive responsibility for the intellectual property framework (section 12.1); managing unsettled law across several jurisdictions at once.
CRO	Model risk (drift, bias, adversarial robustness, explainability); regulatory risk across jurisdictions; concentration risk in AI supply chains.	Integrates AI-specific risk categories into enterprise risk management; translates technical assessments into board-level business language.
CMO	Generative AI capabilities and limitations; personalisation and privacy trade-offs; attribution modelling; the line between effective marketing and manipulation.	Establishes an ethical marketing framework defining approved personalisation techniques and targeting boundaries.

Collective board fluency matters as much as individual fluency. We recommend structured board education, an annual AI strategy offsite, cross-functional governance simulations, and AI fluency as a criterion in board composition and succession planning.

The CFO decides whether governance is funded. Organisations tend to fund AI capability generously and governance sparingly, treating oversight as overhead rather than as part of responsible deployment. Governance spending protects the value of capability spending, and governance cut under budget pressure creates deferred risk that can cost more than it saved.

7. Governance needs a structure, named roles and a budget sized to the task

7.1 Five layers, each with named roles

We recommend that workforce AI governance operate across five layers:

- executive sponsorship, providing strategic direction and authority over resources;
- a steering committee, making cross-functional governance decisions;
- operational oversight, managing day-to-day compliance and monitoring;
- reporting, keeping governance performance visible across the organisation;
- documentation, keeping the audit trail that supports assurance and regulatory defence.

Each layer needs defined roles, decision authority, escalation paths and accountability. We propose four roles:

- a Head of Workforce Architecture, who maintains the overall framework and drives its evolution;

- an AI Employment Ethics Officer, who manages the organisation's ethical commitments, oversees bias monitoring and ensures fairness obligations are met;
- Human-AI Collaboration Coaches, who support employees through changing roles;
- Workforce Compliance Specialists, who maintain documentation standards and meet regulatory obligations across employment-related AI systems.

7.2 Governance should take a larger share of the AI budget than benchmarks show

In our view, an organisation spending 95% of its AI budget on capability and 5% on governance is likely to be undergoverned. The few published benchmarks point to low single-digit shares, though none measures governance as a share of the total AI budget.

A Black Book Research survey of 182 US hospital leaders (November 2025) put the median share of 2026 IT and quality and safety budgets allocated to AI governance and safety at 4.2%. Forrester expects off-the-shelf AI governance software to reach $15.8 billion, about 7% of overall AI software spending, by 2030.

We recommend a deliberately higher allocation: 15 to 20% of the total AI budget for large enterprises, rising to 25 to 30% for smaller organisations, which have fewer economies of scale in governance. These figures are our recommended target for the share of the AI budget allocated to governance, not an observed market norm.

The target also measures something different from the cost of reaching a basic governance baseline, which our paper *Ungoverned AI Systems* estimates at around 0.8% to 1.4% of annual AI-related spend, so the two figures should not be compared directly.

7.3 The structure scales with the organisation

Organisation size	Recommended structure	First 90-day priority actions
Large enterprise (500+ employees)	Federated model: a central governance office sets standards and provides specialist expertise, business units implement within those standards, and HR business partners coordinate. A dedicated governance office of 8 to 15 people.	Appoint the Head of Workforce Architecture and the AI Employment Ethics Officer. Launch a board AI literacy programme. Commission a complete inventory of employment-related AI systems, including shadow AI. Establish the steering committee. Deploy baseline regulatory awareness training across the organisation.
Mid-size (50 to 500 employees)	Combined committee structures, with named individuals holding governance responsibility alongside their main roles. External advisory support for specialist requirements. Governance committee meetings at least monthly.	Assign governance responsibility to named individuals with explicit time allocated. Inventory all employment-related AI systems. Establish a combined governance committee. Schedule executive AI literacy sessions.
Small (under 50 employees)	Governance as a standing item on the leadership team's agenda. External advisory support for the initial compliance assessment and ongoing specialist input.	Add workforce AI governance to leadership meetings as a standing item. Engage external support for an initial compliance assessment. Build AI fluency through executives' hands-on use of the tools deployed.

Shadow AI here means AI tools adopted by staff or teams outside approved procurement.

8. Roles, training and careers must change with the work

8.1 Four role archetypes need different oversight

Roles change along a spectrum, and each point on it carries different governance, training and monitoring requirements:

- **AI-Assisted.** AI handles discrete tasks under close human supervision; the governance overhead is modest and the human judgement content stays high.
- **AI-Augmented.** AI handles a growing share of analytical and process work, and people focus on judgement, relationships and escalation.
- **AI-Collaborative.** People work in continuing dialogue with AI, shaping, challenging and refining its outputs in real time.
- **AI-Native.** The role is conceived from the outset around human-AI collaboration, so governance must be built into the role's design rather than retrofitted after deployment.

Each archetype needs its own capability portfolio, training investment and oversight mechanism.

8.2 Training runs on three pillars, at depths set by role

Article 4 of the AI Act has applied since 2 February 2025 and, as amended by the Digital Omnibus, requires providers and deployers to "take measures to support the development of AI literacy" of their staff and others operating AI systems on their behalf, taking account of their knowledge and the context of use. It does not require any specific level of literacy, so a programme tiered by role is a practical way to show that the measures fit the people they reach.

We recommend three pillars, developed together and differentiated by role tier (executives, managers, frontline staff, technical staff and governance specialists). Tool skills without governance context create compliance risk, and governance knowledge without reporting skills leaves nothing to show for it.

Pillar	Content	Depth by role tier
AI tool and agent fluency	Practical proficiency with the specific AI tools and agents each role uses, scaled to the complexity of the human-AI collaboration the role requires.	Executives: strategic overview and hands-on use. Managers: collaboration tools and oversight mechanisms. Frontline: the specific tools in their daily work. Technical: model characteristics and limitations.
Governance fluency	Regulatory literacy, ethical frameworks, bias detection, oversight responsibilities, escalation protocols.	Executives and governance specialists: multi-week certification. Managers: a structured module. Frontline: two hours' awareness at minimum. Every level must know when to escalate and to whom.
Reporting and process fluency	Documentation standards, interpreting monitoring, incident reporting, governance communication.	All employees: baseline incident recognition and reporting. Governance specialists: full audit trail management and regulatory correspondence.

8.3 Performance management and careers must credit governance work

Performance management has to change in two ways at once. It must assess human contribution where outputs are intertwined with AI, crediting the quality of judgement, the ethical choices made and the governance responsibilities discharged rather than the volume produced with AI assistance. It must also give governance contributions real weight: an employee who detects and escalates a bias issue has added value that their evaluation should reflect.

Career architecture needs new tracks for the roles agentic work creates. Human-AI collaboration specialists, AI governance professionals and AI monitoring and assurance analysts need structured development pathways, not ad hoc assignment.

9. Governance holds only when it is someone's job, with consequences attached

Governance that depends on goodwill alone will not survive commercial pressure. Accountability has to be structured, incentivised and measured.

Accountability sits at three levels. At board level, AI governance belongs in the existing risk and audit committee structures, with AI risk reported alongside financial and operational risk. At executive level, governance performance should count in performance assessment and variable pay, and at operational level, governance responsibilities should be explicit in capability portfolios and assessed in performance reviews.

Incentives should be calibrated to governance exposure. For governance-critical roles, we recommend tying 15 to 25% of variable pay to governance outcomes, measured through bias monitoring results, incident reporting rates, training completion, oversight quality assessments and transparency scores from several groups of stakeholders. For

all roles, we recommend a governance gate on bonuses (a condition that must be met before any bonus is paid), so that strong commercial performance cannot mask a governance failure.

Any of these measures can be gamed once it becomes a target. A reward for low incident counts, for example, discourages reporting. Incentive-behaviour correlation analysis tests whether the incentives produce the behaviour they are meant to reward. Run regularly, it can surface gaps before they become failures.

10. Employees must know how AI affects them and have a voice in its use

Employees should understand how AI influences their work, their performance assessment and employment decisions about them. They should know when AI outputs inform decisions about them, how to contest those decisions and who is responsible for the systems affecting their working lives.

Every deployment of AI in employment (recruitment, performance management, scheduling, workforce planning, disciplinary processes) needs a communication plan. It should cover what the system does and does not do, what safeguards exist, how employees can raise concerns and who is accountable if something goes wrong. For a high-risk system, the Article 26(7) duty to inform workers and their representatives sets a legal floor for part of this from 2 December 2027 (subject to Article 111(2) for systems already on the market), and works-council rights may already apply (section 3).

We treat employee voice as a governance requirement, not a consultation courtesy. People affected by an AI deployment hold information about its impact that monitoring data alone cannot capture. A performance system that looks fair in aggregate may be experienced as unfair by particular teams or groups in ways only they can describe, and governance that excludes them can be expected to miss the risks most likely to cause harm.

11. Trust is built in sequence, and psychological safety keeps the monitoring honest

Trust in governance is earned through leadership actions taken in the right order. Publishing a policy does not earn it. Governance announced after deployment signals that it is reactive, and governance announced alongside deployment signals that it was designed in.

During active change, we recommend briefings at least weekly, because silence creates anxiety that speculation fills. Leaders must also model the behaviour the framework asks

for: a CEO who visibly uses AI tools, asks governance questions and holds executives to account for governance outcomes sends a signal that no policy document can.

Four employee anxieties need specific answers. Fear of displacement needs honest, concrete communication about how roles will evolve, not reassurance that jobs are safe, and concern about surveillance needs clear limits on what is monitored, how the data are used and how performance data are protected from misuse. Uncertainty about future value needs reskilling that is visible, funded and led by senior leaders. Frustration with rapid tool change needs a genuine employee voice in tool selection and deployment timing.

Psychological safety is part of the governance infrastructure. It is the confidence that raising a concern, reporting a governance failure or challenging an AI output will not be punished. Where it is missing, monitoring data become unreliable, incidents go unreported and failures accumulate unseen until a regulator finds them.

12. Intellectual property and AI costs need governance of their own

12.1 Intellectual property

In our experience, many organisations have yet to address three overlapping intellectual property questions systematically:

- **Ownership.** When an employee creates work with an AI tool, who owns the output: the employee, the organisation or the AI provider?
- **Protectability.** In many jurisdictions copyright depends on human authorship, so wholly AI-generated content may not be protected. The UK protects computer-generated works under section 9(3) of the Copyright, Designs and Patents Act 1988, but in its March 2026 report on copyright and AI the Government proposed that this protection be removed, while keeping protection for works created with AI assistance.
- **Confidentiality.** Employees often share commercially sensitive information with AI tools, and many organisations have no policy governing it.

An intellectual property framework should cover updated employment and contractor terms, a data classification policy setting out what may and may not be shared with AI platforms, management of open-source licensing risk in AI-generated code, and joint working with procurement and legal.

12.2 AI costs

AI costs behave differently from traditional IT spending. Foundation model API calls are charged by use, and an agentic workflow can turn a single request into dozens of chained calls, each with its own cost, so traditional IT budgeting methods are likely to forecast AI spending poorly.

Financial governance of AI (often called AI FinOps) should provide cost visibility and attribution, unit economics and cost per outcome, detection of shadow AI spending,

management of vendor concentration risk, and reporting to the board. Governance itself should be explicitly costed, since an organisation that cannot account for what it spends on governance cannot judge whether it is spending enough (section 7.2).

13. Aim for at least Level 3 of five maturity levels, and govern at speed

Governance maturity in the agentic workforce progresses across five levels, assessed against seven dimensions. The levels are this paper's own model, not a regulatory or industry standard. Two of the names (Structured and Leading) also appear in the TRACE maturity model set out in our paper *From Obligation to Evidence*, which runs from Ad Hoc to Leading, but the levels are defined differently. Neither model corresponds to the percentage governance-maturity measure used in *Ungoverned AI Systems*.

Level	Designation	Defining characteristic	Primary governance risk
1	Reactive	No formal governance. AI deployed without systematic oversight. Compliance addressed only when a regulator makes contact.	Critical: enforcement exposure, undetected discrimination, erosion of employee trust.
2	Foundational	Basic governance structures exist but are applied inconsistently. Documentation satisfies auditors but is not embedded in operations.	High: governance exists on paper and is bypassed under commercial pressure.
3	Structured	Governance is consistently implemented. Monitoring produces actionable data, training reaches the required populations and escalation works.	Moderate: governance is compliance-driven and optimises for the letter of the rules rather than their purpose.
4	Quantitatively Managed	Governance performance is measured and improving. Predictive indicators identify emerging issues, and employee voice is genuinely incorporated.	Low: governance is strong but may not yet survive changes of leadership.
5	Leading	Governance is culturally embedded, compliance is easier to sustain, and the organisation contributes to standards development and leads on transparency.	Minimal: residual risk is managed within normal operations.

The seven dimensions are:

- governance structures and data governance;
- training and capability development;
- monitoring and assurance;
- workforce transformation and career architecture;
- accountability, incentives and transparency;
- intellectual property and financial governance;

- culture, change management and employee voice.

Level 3 is the minimum we recommend. In our framework, most organisations seeking initial compliance should aim for at least Level 3, because Level 2 leaves real regulatory exposure, particularly as enforcement develops. Beyond compliance, Levels 4 and 5 can bring advantages in attracting talent, in relationships with regulators and in deploying AI at pace with confidence.

Governing at speed avoids both paralysis and abdication. The tension between governance rigour and the pace of AI change is real. Risk-proportionate review pathways, pre-approved deployment patterns for lower-risk uses and tiered governance processes allow fast adoption within guardrails. They avoid governance paralysis, where every AI output needs human approval and the productivity gain is lost, and governance abdication, where AI operates without meaningful oversight and ethical and regulatory exposure accumulates.

14. Report governance performance to the board each quarter, in four domains

Governance performance should be measured and reported to the board alongside traditional HR metrics, with the CPO presenting a governance dashboard at each quarterly review.

Domain	Key metrics	Reporting frequency
Workforce readiness	Time from deployment decision to workforce readiness; employee confidence in working with AI, by demographic group; retention of employees in AI-transformed roles.	Quarterly to the board; monthly to the governance committee.
Compliance	System-level compliance status, tracked continuously; training completion and certification rates across all three pillars; frequency of regulatory incidents and time to resolve them.	Tracked continuously; monthly summary; quarterly board report.
Fairness and equity	Demographic equity in AI tool access, adoption and career outcomes; quality of human-AI collaboration, assessed through governance audits; bias monitoring results by system.	Quarterly to the board; monthly to the AI Employment Ethics Officer.
Governance performance	Executive AI fluency scores; governance behaviour survey scores; incentive-behaviour correlation metrics; transparency assessment scores from several groups of stakeholders.	Half-yearly to the nominations committee; quarterly to the risk committee.

Measuring by demographic group needs a lawful basis for the data. Where the groups are defined by characteristics such as racial or ethnic origin, health, religion or sexual orientation, the measures rely on special category data, which need a lawful

basis under Article 6 of the GDPR and a condition under Article 9(2). For bias detection and correction in AI systems, Article 4a of the AI Act has since 27 July 2026 permitted providers and deployers to process such data where strictly necessary and under strict conditions, but it is a permission, not a lawful basis in itself.

An organisation that cannot report its governance performance in figures will find it much harder to show regulators, employees or investors that its governance works. Measurement is the feedback loop that lets governance correct itself.

15. The Standard Intelligence platform is designed to support both sides of the shift

Workforce change in the agentic age is a governance problem as much as a people problem, and it depends on capability and accountability being built in rather than retrofitted. The Standard Intelligence platform supports both the human and the structural side of that change.

Platform capability	How it supports the workforce shift
AI Literacy Platform	Delivers role-based AI literacy and credentials across the workforce, evidenced and gated so that capability keeps pace with deployment (section 8.2).
Practice Workspaces	Give each function, from operators to leadership, a role-specific view of its AI responsibilities and status.
TRACE Maturity Model	Locates the organisation on the TRACE five-level path from Ad Hoc to Leading and sets the sequence of the change; its levels are defined differently from those in section 13.
Approval Workflow	Assigns a named accountable owner to each AI system, so that accountability does not diffuse as autonomy increases (section 9).

Where to start, and what it costs

The work comes in stages, and the first is an inventory.

First 90 days: find the systems and name the owners. Inventory every AI system that touches employment decisions, including shadow AI, and assign named governance responsibility at the scale section 7.3 sets out. Start the board's fluency programme (section 6) and baseline awareness training for all staff (section 8.2).

Next: build the structure and fund it. Establish the governance layers and roles (section 7.1) and set the governance budget against the target in section 7.2, which is the main cost of the framework. Write capability portfolios for roles that oversee or depend on AI (section 5.1), put governance into incentives and performance reviews (section 9), and give every employment deployment a communication plan (section 10).

Ongoing: measure, report and check. Report the four-domain dashboard to the board each quarter (section 14) and reassess maturity against the seven dimensions, aiming for at least Level 3 (section 13). Check regularly that incentives produce the behaviour they reward (section 9).

On timing. The Article 5(1)(f) prohibition and Article 4 have applied since 2 February 2025, and the UK rules on automated decisions since 5 February 2026. The Annex III obligations for employment AI apply from 2 December 2027, subject to the Article 111(2) qualification for systems already on the market, and as at 11 October 2026 no fine under the AI Act had been confirmed by any regulator or official source.

Agentic AI changes who decides, and governance has to follow the decision to the people around it. An organisation that has made the four shifts, put fluency at the top and wired governance into roles, pay and reporting can show its board, its regulators and its people who is accountable for the systems that act on its behalf.

This paper is provided for educational and informational purposes and is not legal, employment or professional advice. It describes the position as at 11 October 2026; AI Act application dates were amended by Regulation (EU) 2026/1744, in force 27 July 2026. Take qualified legal, regulatory and HR advice before acting on it.

Basis and Standing

Binding law. All AI Act references are to Regulation (EU) 2024/1689 as amended by Regulation (EU) 2026/1744 (in force 27 July 2026). Annex III point 4 and Article 6(3) on the classification of employment systems as high-risk; Article 26(1), (2) and (7) on deployers' use of high-risk systems, assignment of oversight and the duty to inform workers; Article 113 (Annex III obligations from 2 December 2027) and Article 111(2) on systems already on the market; Article 5(1)(f), applying since 2 February 2025; Article 4 on AI literacy, applying since 2 February 2025 and amended with effect from 27 July 2026; Article 4a, a permission to process special category data for bias detection and correction from 27 July 2026; Article 99(3), (4) and (6) on fines. GDPR (Regulation (EU) 2016/679) Article 6 (lawful basis), Article 9(2) (conditions for special category data) and Article 22. UK GDPR Articles 22A to 22D, substituted for Article 22 by section 80 of the Data (Use and Access) Act 2025, in force 5 February 2026 (SI 2026/82). Section 87(1) no. 6 of the German Works Constitution Act. Section 9(3) of the Copyright, Designs and Patents Act 1988.

Guidance and policy. The Commission's guidelines on prohibited AI practices, C(2025) 5052 final (29 July 2025), para 245, on the limit of Article 5(1)(f) to biometric data. The UK Government's Report on Copyright and Artificial Intelligence (March 2026), which proposes removing section 9(3) protection; it is a proposal, not law. The NIST AI Risk Management Framework 1.0 (January 2023), a voluntary framework.

Not a regulatory standard. The following are this paper's own design and are not prescribed by any regulator: the three principles as an organising framework; the four shifts and capability portfolios; the distributed accountability framework and the policy that decisions affecting individual rights always involve meaningful human judgement,

which goes further than the law (Article 22 of the GDPR and UK GDPR Articles 22A to 22D permit some solely automated decisions with safeguards); the C-suite fluency requirements and board education recommendations; the five governance layers and four roles; the budget target of 15 to 20% (large enterprises) and 25 to 30% (smaller organisations); the sizing table, including the governance office of 8 to 15 people; the four role archetypes; the three training pillars and their durations; the incentive mechanisms (15 to 25% of variable pay, governance gates, incentive-behaviour correlation analysis); weekly briefings; the five maturity levels, seven dimensions and Level 3 recommendation; and the four-domain dashboard. Recording oversight competence in capability portfolios reflects Article 26(2); the tiered training programme reflects Article 4; the communication plan reflects Article 26(7) and works-council rights.

Research and market data, and what they support. The Black Book Research survey and the Forrester forecast measure narrower things than the share of the AI budget spent on governance (a hospital sector budget share and software spending respectively); they support only the statement that published benchmarks are in the low single digits. The commercial case in section 2 is our reasoned expectation, and no study cited here measures it. The *Ungoverned AI Systems* baseline estimate is our own modelling and measures a different quantity from the budget target.

Enforcement position. As at 11 October 2026, no fine under the AI Act had been confirmed by any regulator or official source.

Sources

1. Regulation (EU) 2024/1689 (Artificial Intelligence Act), OJ L, 2024/1689, 12.7.2024, consolidated version of 27 July 2026: Articles 4, 4a, 5, 6, 26, 99, 111 and 113, and Annex III. https://eur-lex.europa.eu/eli/reg/2024/1689/2026-07-27/eng
2. Regulation (EU) 2026/1744 of 8 July 2026 amending Regulations (EU) 2024/1689, (EU) 2018/1139 and (EU) 2023/1230 as regards the simplification of the implementation of harmonised rules on artificial intelligence (Digital Omnibus on AI), OJ L, 2026/1744, 24.7.2026; in force 27 July 2026. https://eur-lex.europa.eu/eli/reg/2026/1744/oj
3. European Commission, Guidelines on prohibited artificial intelligence practices established by Regulation (EU) 2024/1689, C(2025) 5052 final, 29 July 2025, para 245. https://ai-act-service-desk.ec.europa.eu/sites/default/files/2026-01/guide-prohibited_en.pdf
4. Data (Use and Access) Act 2025, section 80, substituting Articles 22A to 22D for Article 22 of the UK GDPR; in force 5 February 2026 under the Data (Use and Access) Act 2025 (Commencement No. 6 and Transitional and Saving Provisions) Regulations 2026, SI 2026/82. https://www.legislation.gov.uk/ukpga/2025/18/section/80; UK GDPR Articles 22A to 22D. https://www.legislation.gov.uk/eur/2016/679/article/22A
5. Regulation (EU) 2016/679 (General Data Protection Regulation), Articles 6, 9 and 22. https://eur-lex.europa.eu/eli/reg/2016/679/oj
6. Works Constitution Act (Betriebsverfassungsgesetz), section 87(1) no. 6. https://www.gesetze-im-internet.de/englisch_betrvg/

7. Copyright, Designs and Patents Act 1988, section 9(3). https://www.legislation.gov.uk/ukpga/1988/48/section/9
8. HM Government, Report on Copyright and Artificial Intelligence (18 March 2026), para 47. https://assets.publishing.service.gov.uk/media/69ba692226909a14239612e4/CP2602959_-_Report_on_Copyright_and_Artificial_Intelligence_web.pdf
9. National Institute of Standards and Technology, Artificial Intelligence Risk Management Framework (AI RMF 1.0), NIST AI 100-1, January 2023. https://doi.org/10.6028/NIST.AI.100-1
10. Black Book Research, "U.S. Hospitals Underfund AI Governance as Adoption Accelerates" (ACCESS Newswire press release, 12 November 2025; 182 US hospital leaders surveyed 15 October to 8 November 2025). https://markets.financialcontent.com/chroniclejournal/article/accwirecq-2025-11-12-us-hospitals-underfund-ai-governance-as-adoption-accelerates
11. Forrester, "AI Governance Software Spend Will See 30% CAGR From 2024 To 2030" (undated blog summarising Global Commercial AI Software Governance Market Forecast, 2024 To 2030). https://www.forrester.com/blogs/ai-governance-software-spend-will-see-30-cagr-from-2024-to-2030
12. Standard Intelligence, Ungoverned AI Systems (whitepaper), on the cost of a basic governance baseline. https://standardintelligence.com/whitepapers/ungoverned-ai-enterprise-risk
13. Standard Intelligence, From Obligation to Evidence (whitepaper), on the TRACE maturity model. https://standardintelligence.com/whitepapers/trace-methodology

Standard Intelligence

GET IN TOUCH

Lead on agentic AI rather than react to it.

EMAIL

hello@standardintelligence.com

WEB

standardintelligence.com

LOCATION

Manchester, United Kingdom

AI Use: AI tools were used in research, editorial review and drafting support. All sources were retrieved and checked against the published record by the author, who holds editorial responsibility for the contents and for the conclusions drawn. This disclosure is made voluntarily. To the extent the Article 50(4) obligation in Regulation (EU) 2024/1689 applies to this paper as text published to inform the public on matters of public interest, the exception for content that has undergone human review and for which a person holds editorial responsibility applies.

STANDARD INTELLIGENCE LIMITED · REGISTERED IN ENGLAND AND WALES · COMPANY NO. 16737890