The Human Transition Business Adoption Standard

A responsible AI governance standard for companies, institutions, and employers.

By David H. Wever

Artificial intelligence is rapidly changing how organizations operate, compete, hire, manage, communicate, evaluate, and deliver value. Used well, it may improve productivity, expand access, reduce friction, and unlock new capabilities. Used poorly, it may intensify surveillance, erode trust, displace workers without preparation, expose sensitive data, manipulate vulnerable users, weaken accountability, and damage the legitimacy of institutions.

This Standard guides businesses and institutions in adopting AI in a manner consistent with human dignity, transparency, accountability, workforce responsibility, child protection, privacy, and the common good.

Part I. Core Principles

1. Human dignity comes first

No AI system should be deployed in a way that treats human beings as disposable units, behavioral targets, or mere objects of optimization.

2. Accountability must remain human

AI may inform decisions, but accountable responsibility must remain with identifiable human leaders and institutions.

3. Higher risk requires higher responsibility

Not all AI uses carry the same stakes. The greater the impact on rights, livelihood, education, health, finances, reputation, or safety, the stricter the oversight required.

4. Adoption must not outrun stewardship

Organizations should not adopt AI merely because competitors do, investors demand it, or vendors promise savings. Adoption should follow evidence, governance, testing, review, and consideration of human consequences.

5. The gains of AI should not be privatized while the harms are socialized

If AI creates significant gains, organizations should not treat worker displacement, community disruption, or social costs as somebody else’s problem.

Part II. Scope and Risk Classification

6. Covered systems

This Standard applies to all material AI systems an organization develops, fine-tunes, procures, integrates, or deploys, including systems used for hiring and workforce evaluation, customer interactions, education or training, content generation, marketing, legal or risk triage, and automated or semi-automated decision-making.

7. Risk categories

- Low-risk systems : limited consequences for individuals and no meaningful impact on rights, livelihood, education, access, or safety.
- Moderate-risk systems : influence experiences, workflow, consumer engagement, or scalable communication, but do not directly determine consequential outcomes.
- High-risk systems : materially shape decisions regarding a person’s livelihood, legal standing, education, finances, health, benefits, reputation, safety, or access to important opportunities.
- Sensitive relational systems : operate where emotional dependency, developmental vulnerability, attachment, loneliness, crisis states, or trust asymmetries are likely to be present.
- Systemic-impact systems : scale, reach, or workforce effect may materially affect public trust, labor structures, communities, or institutional legitimacy.

Part III. Governance Requirements

8. Board and executive responsibility

Organizations shall assign formal oversight of AI governance to an accountable governing body and involve legal, privacy, security, human resources, operations, product, communications, and risk leadership as appropriate.

9. AI governance charter

Organizations should adopt an internal AI Governance Charter defining governance authority, approval thresholds, risk classification rules, review procedures, escalation triggers, incident response rules, and documentation requirements.

10. AI inventory

Organizations shall maintain a current inventory of material AI systems, including system name, function, business owner, vendor, data sources, risk classification, affected populations, degree of automation, decision role, and status of approvals and testing.

11. Approval pathways

No high-risk, sensitive relational, or systemic-impact AI system should be deployed without formal review and written approval through a documented governance process.

Part IV. Pre-Deployment Review

12. Required pre-deployment assessment

Before deployment of any high-risk, sensitive relational, or systemic-impact AI system, organizations shall conduct a documented pre-deployment assessment addressing intended use, expected benefits, foreseeable harms, affected users or workers, data provenance, bias concerns, privacy implications, security risks, likely failure modes, complaint mechanisms, human oversight, and workforce effects.

13. Testing and validation

Material AI systems should undergo testing proportionate to risk, including accuracy, robustness, adversarial testing, bias review, hallucination review, stress testing for edge cases, and usability testing for affected populations.

14. Vendor accountability

Organizations remain responsible for AI systems they procure from third parties. Contracts with AI vendors should address data rights, security obligations, model update disclosures, audit access where feasible, incident notification, service limitations, subcontractor use, and remediation rights.

Part V. Human Oversight and Rights Protection

15. Human review for consequential decisions

No organization should permit a final consequential decision to be made solely by AI in matters involving hiring, firing, promotion denial, pay-impacting evaluation, education discipline, healthcare access, insurance eligibility, benefit denial, housing access, legal exposure, account suspension with material consequences, or similar outcomes.

16. Explanation and contestability

Where AI materially contributes to a consequential decision, affected persons should be able to know that AI was used, receive a plain-language explanation, contest incorrect or unfair outcomes, request human reconsideration, and obtain correction where appropriate.

17. Limits on surveillance and behavioral scoring

Organizations should not use AI to create hidden, manipulative, or disproportionate systems of surveillance, emotional inference, social scoring, or behavior prediction affecting workers, customers, students, or users without compelling justification and strong safeguards.

Part VI. Workforce Responsibilities

18. Workforce impact assessment

Before major AI-led automation, organizations should conduct a Workforce Impact Assessment covering functions likely to change, expected productivity gains, retraining options, redeployment pathways, transition implications, communication plans, and effects on morale and organizational trust.

19. Notice and honest communication

Where material workforce impacts are reasonably foreseeable, organizations should provide honest communication about what is changing, why it is changing, what roles may be affected, what support will be available, and what timelines are anticipated.

20. Retraining and redeployment

Where feasible, organizations should prioritize upskilling existing workers, redeployment into new roles, transition support, internal training access, and adaptation pathways for mid-career employees.

21. Dignified displacement standards

Where displacement occurs, organizations should adopt dignified transition standards such as severance, career support, transition coaching, benefits continuity where possible, early notice, and fair treatment of long-serving employees and contractors where appropriate.

Part VII. Privacy, Data Governance, and Security

22. Data minimization and purpose limits

Organizations should collect, retain, and use only the data reasonably necessary for legitimate purposes connected to the AI system's function.

23. Sensitive data protection

Special restrictions and governance approval should apply to AI systems using or inferring data relating to health, mental health, children, education records, finances, biometrics, precise geolocation, intimate life, crisis behavior, or protected-class information where legally relevant.

24. Employee and user data boundaries

Organizations should not quietly repurpose employee, customer, or user data for model training, behavioral inference, or high-stakes prediction without lawful basis, transparent notice, and appropriate controls.

25. Security and misuse prevention

Material AI systems should be protected against unauthorized access, model abuse, prompt injection, data leakage, sensitive output exposure, jailbreak misuse, malicious automation, and unsafe system chaining.

Part VIII. Child Safety and Vulnerable Users

26. Child-first safeguards

Organizations whose AI systems are reasonably likely to be used by minors shall adopt child-first safeguards, including privacy by default where appropriate, age-appropriate experiences, restrictions on manipulative engagement features, limits on data collection, and prohibition on emotionally exploitative design.

27. No engineered dependency for minors

Organizations shall not design or market AI systems to minors in ways that encourage secrecy from caregivers, compulsive emotional reliance, confusion about whether a system is human, or synthetic attachment as a substitute for real-life care.

28. Vulnerable-user protections

Where systems operate in domains such as loneliness, grief, crisis, mental health, disability support, or elder care, organizations must apply heightened caution, review, and harm prevention.

Part IX. Transparency and Public Trust

29. Material disclosure

Organizations shall provide clear, accessible notice when AI materially shapes a person's experience, opportunities, decisions, or treatment.

30. Public-facing AI transparency statement

Organizations above a defined size, scale, or risk threshold should publish an annual AI Transparency Statement covering major categories of AI use, governance structure, high-risk deployments, workforce impacts observed, incidents and remedial actions, data and privacy safeguards, and child-protection practices where relevant.

31. Labeling of synthetic content

Organizations should label materially AI-generated public-facing content where failure to do so could mislead users, customers, workers, or the public, especially in contexts involving trust, expertise, or identity.

Part X. Incident Response and Continuous Review

32. Incident reporting process

Organizations shall maintain documented procedures for identifying, escalating, investigating, and remediating AI-related incidents.

33. Serious incident escalation

Serious incidents involving material harm, rights violations, child-safety risks, legal exposure, or public trust threats should be promptly escalated to senior leadership and, where appropriate, regulators, affected users, customers, or employees.

34. Periodic re-evaluation

Material AI systems should be reviewed periodically and whenever significant changes occur, including model updates, new use cases, outcome drift, regulatory changes, and incidents or complaints.

Part XI. Special Rules for Sensitive Relational Systems

35. Emotional and relational caution

Organizations deploying emotionally responsive or relational AI systems should recognize that simulated empathy can produce real dependency, disclosure, and trust, and therefore require heightened governance.

36. No false claims of human equivalence

Organizations should not market AI systems as equivalent to therapists, parents, intimate partners, friends, or moral authorities where such claims may mislead users about the system's actual nature and limits.

37. Escalation boundaries

Sensitive relational systems should include clear limits, safety protocols, and escalation pathways for high-risk situations such as suicidality, abuse disclosure, psychosis, child endangerment, medical crisis, criminal threat, or extreme dependency.

Adoption Pledge

We affirm that artificial intelligence must remain answerable to human dignity, accountability, and the common good. We commit to adopting AI in ways that preserve meaningful human oversight, protect privacy, safeguard children and vulnerable persons, communicate honestly with workers and users, and accept responsibility for the downstream effects of our systems.