INTELLIGENT SYSTEMS PRESS™

EXECUTIVE ASSESSMENT SERIES

EA-004: Responsible AI Governance Assessment™

Can Your Enterprise Govern AI With Accountability, Evidence, and Trust?

Purpose

A practical executive assessment for evaluating whether AI governance is embedded across accountability, policy, risk classification, data, human oversight, security, monitoring, assurance, and continuous improvement.

Educational Executive Self-Assessment | Version 1.0

www.intelligentsystemspress.com

About This Assessment

Responsible AI governance is the enterprise capability to direct, control, and continuously improve AI use so that value, accountability, risk, and stakeholder protection remain connected throughout the lifecycle. This assessment helps leaders identify whether governance is operational or exists mainly as policy intent.

How to Complete It

1. Rate each statement Use the 1-5 scale based on current, observable enterprise practice - not future intent, isolated pilots, or vendor assurances.

2. Score each domain Add the five ratings in each domain. Every domain score ranges from 5 to 25.

3. Calculate the total Add all eight domain scores. The overall governance score ranges from 40 to 200.

4. Interpret the pattern Review the total and the distribution. Any domain below 15 represents a material governance dependency.

5. Select 90-day actions Choose a limited number of executive actions that strengthen accountability, controls, evidence, or assurance.

Rating Scale

1 Strongly Disagree	Practice is absent or contradicted by current evidence.
2 Disagree	Practice is inconsistent, informal, or limited to isolated areas.
3 Partly Agree	Practice exists in several areas but is not yet enterprise-wide.
4 Agree	Practice is generally reliable, measured, and consistently applied.
5 Strongly Agree	Practice is integrated, evidence-driven, and continuously improved.

Recommended use: Complete individually first, then compare results as an executive team. Differences in ratings are evidence of governance misalignment and should be discussed before averaging scores.

Domain 1: Governance Mandate & Accountability

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

Assessment statement	1	2	3	4	5
Executive leadership has approved a clear mandate for responsible AI governance that applies across business units, functions, and delivery models.	○	○	○	○	○
Named executives and governance bodies have defined authority for approving, escalating, suspending, and retiring AI uses.	○	○	○	○	○
Business, technology, data, legal, privacy, security, risk, quality, and workforce responsibilities are explicitly assigned and understood.	○	○	○	○	○
Governance accountability continues throughout the AI lifecycle rather than ending at initial approval or deployment.	○	○	○	○	○
Leaders receive timely evidence about AI value, performance, incidents, exceptions, and unresolved risks.	○	○	○	○	○

Domain score: _____ / 25

Evidence or observations:

Domain 2: Policy, Standards & Decision Rights

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

Assessment statement	1	2	3	4	5
Enterprise AI policies define permitted, restricted, and prohibited uses according to impact, consequence, and regulatory exposure.	○	○	○	○	○
Decision rights and approval thresholds are proportionate to the sensitivity of data, autonomy of the system, and severity of potential harm.	○	○	○	○	○
AI standards are translated into practical requirements, controls, templates, and acceptance criteria for delivery teams.	○	○	○	○	○
Policy exceptions require documented rationale, accountable approval, compensating controls, and expiration or review dates.	○	○	○	○	○
Policies and standards are reviewed as laws, technologies, business models, and organizational risk tolerance change.	○	○	○	○	○

Domain score: _____ / 25

Evidence or observations:

Domain 3: Use-Case Intake & Risk Classification

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

Assessment statement	1	2	3	4	5
Every proposed AI use begins with a documented problem, intended outcome, affected users, and accountable business owner.	○	○	○	○	○
AI uses are classified using consistent criteria for impact, data sensitivity, human dependency, reversibility, and external exposure.	○	○	○	○	○
High-impact or high-uncertainty uses receive enhanced review before experimentation, procurement, deployment, or material change.	○	○	○	○	○
The intake process identifies whether AI is necessary or whether a simpler process, rule, or technology would produce better value with less risk.	○	○	○	○	○
Approval records preserve the assumptions, evidence, conditions, and controls that justified each governance decision.	○	○	○	○	○

Domain score: _____ / 25

Evidence or observations:

Domain 4: Data, Privacy & Information Governance

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

Assessment statement	1	2	3	4	5
Data used to train, configure, test, operate, or evaluate AI has defined ownership, provenance, permitted use, quality expectations, and retention rules.	○	○	○	○	○
Privacy requirements are assessed before data is collected, combined, inferred, transmitted, or exposed through AI-enabled workflows.	○	○	○	○	○
Sensitive, confidential, regulated, copyrighted, and third-party data are protected through enforceable access and usage controls.	○	○	○	○	○
Data quality, representativeness, lineage, and known limitations are evaluated in relation to the intended decision or business process.	○	○	○	○	○
Governance controls address prompts, embeddings, retrieval sources, generated content, logs, feedback data, and other AI-specific information assets.	○	○	○	○	○

Domain score: _____ / 25

Evidence or observations:

Domain 5: Human Oversight, Ethics & Stakeholder Protection

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

Assessment statement	1	2	3	4	5
Human oversight is designed around the consequence and complexity of AI-assisted decisions rather than added as a generic approval step.	○	○	○	○	○
Users understand when they are interacting with AI, when human review is available, and how to challenge or appeal significant outcomes.	○	○	○	○	○
The enterprise evaluates potential unfairness, exclusion, accessibility barriers, manipulation, and disproportionate stakeholder impact.	○	○	○	○	○
Human reviewers have sufficient authority, time, context, and competence to question, override, or stop AI-supported decisions.	○	○	○	○	○
Stakeholder feedback and lived experience inform governance decisions for uses that materially affect people, rights, services, or opportunities.	○	○	○	○	○

Domain score: _____ / 25

Evidence or observations:

Domain 6: Security, Resilience & Third-Party Risk

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

Assessment statement	1	2	3	4	5
AI threat modeling addresses prompt injection, data leakage, model abuse, insecure integrations, adversarial inputs, and unauthorized automation.	○	○	○	○	○
AI systems are integrated with identity, access, logging, incident response, continuity, and change-control practices.	○	○	○	○	○
Third-party models, platforms, data sources, and vendors are evaluated for transparency, security, contractual protection, dependency, and exit risk.	○	○	○	○	○
Fallback procedures and safe operating modes are defined for degraded performance, vendor outages, compromised data, or unreliable outputs.	○	○	○	○	○
Security and resilience controls are tested before production use and reassessed after material changes or emerging threats.	○	○	○	○	○

Domain score: _____ / 25

Evidence or observations:

Domain 7: Validation, Monitoring & Incident Management

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

Assessment statement	1	2	3	4	5
AI solutions are validated against documented business, technical, quality, safety, and risk acceptance criteria before production use.	○	○	○	○	○
Monitoring covers outcome quality, model or workflow performance, drift, bias indicators, security events, user behavior, and operational reliability.	○	○	○	○	○
Thresholds exist for investigation, human intervention, rollback, suspension, notification, and retirement.	○	○	○	○	○
AI incidents and near misses are recorded, analyzed for systemic causes, and translated into corrective and preventive action.	○	○	○	○	○
Changes to models, prompts, data, integrations, policies, and operating context trigger proportionate revalidation.	○	○	○	○	○

Domain score: _____ / 25

Evidence or observations:

Domain 8: Transparency, Assurance & Continuous Improvement

Rate each statement from 1 (Strongly Disagree) to 5 (Strongly Agree). Base the rating on repeatable enterprise behavior and current evidence.

Assessment statement	1	2	3	4	5
The enterprise maintains an accurate inventory of AI systems, owners, purposes, data dependencies, risk classifications, and lifecycle status.	○	○	○	○	○
Documentation is sufficient for leaders, reviewers, operators, auditors, and affected stakeholders to understand how governance decisions were made.	○	○	○	○	○
Governance performance is measured using evidence such as exception trends, incident patterns, control effectiveness, adoption, and realized outcomes.	○	○	○	○	○
Independent review or challenge is used where impact, uncertainty, conflict of interest, or regulatory exposure warrants additional assurance.	○	○	○	○	○
Lessons from monitoring, incidents, stakeholder feedback, audits, and portfolio results are used to improve policies, controls, and operating practices.	○	○	○	○	○

Domain score: _____ / 25

Evidence or observations:

Scoring & Governance Interpretation

Record each domain score. Treat any domain below 15 as a priority dependency even when the overall score is higher.

Domain	Score
1. Governance Mandate & Accountability	____ / 25
2. Policy, Standards & Decision Rights	____ / 25
3. Use-Case Intake & Risk Classification	____ / 25
4. Data, Privacy & Information Governance	____ / 25
5. Human Oversight, Ethics & Stakeholder Protection	____ / 25
6. Security, Resilience & Third-Party Risk	____ / 25
7. Validation, Monitoring & Incident Management	____ / 25
8. Transparency, Assurance & Continuous Improvement	____ / 25
Total responsible AI governance score	**____ / 200**

Overall Governance Levels

180-200	**Governance Embedded**	Accountability, controls, evidence, and improvement are integrated across the AI lifecycle.
150-179	**Governed**	Strong foundations exist, with targeted gaps that should be resolved before broader scaling.
110-149	**Developing**	Several governance capabilities are established, but maturity remains uneven or dependent on local practice.
70-109	**Emerging**	Governance is fragmented, reactive, or focused mainly on policy rather than operational control.
40-69	**Exposed**	Material gaps in accountability, oversight, evidence, and risk control create significant enterprise exposure.

Important: This assessment is a discussion and prioritization tool, not a certification. Governance effectiveness should be validated using evidence, stakeholder input, legal and regulatory analysis, technical assurance, and use-case-specific risk review.

Executive Reflection & 90-Day Action Plan

Use the governance pattern to identify the few leadership actions that will most improve accountability, trust, risk control, and sustainable AI value.

Governance strengths

Which capabilities provide a reliable foundation for responsible AI oversight and decision-making?

Critical governance gaps

Which weaknesses could create legal, operational, ethical, security, or reputational exposure?

Decisions required from leadership

What ownership, authority, policy, funding, or risk-tolerance decisions cannot be delegated?

90-Day Priority Actions

Priority action	Executive owner	Success measure	Target date

Executive commitment: We will review progress by _____ and reassess affected domains using current evidence.

Continue Your Responsible AI Journey

Responsible AI governance is strongest when enterprise strategy, Process Intelligence, requirements, quality, security, data, workforce readiness, and continuous improvement operate as one connected system.

Explore the Intelligent Systems Press collection and Free Resources library

www.intelligentsystemspress.com