AI Governance Readiness Checklist

Truth Before It Costs Millions™

A structured governance instrument designed to surface hidden AI risk and clarify executive and board-level accountability for AI systems and decision authority.

Who It's For

This checklist is for leaders accountable for AI outcomes — not merely those implementing the technology.

Specifically:

- Boards and Board members responsible for AI oversight
- CEOs and executive leaders approving AI deployment
- CIOs, CTOs, and transformation leaders accountable for implementation
- Risk, compliance, and audit leaders evaluating governance exposure

A structured governance instrument for executive and board-level oversight of AI systems and decision authority.

How to Complete This Checklist

A checked box affirms that the control:

- Formally exists
- Is documented
- Is operational
- Can be evidenced upon request

If documentation, ownership, or operational clarity is uncertain, treat it as a governance gap and unsurfaced risk.

If it cannot be evidenced, it is not established and must remain unchecked to accurately reflect exposure.

SECTION 1 — Inventory & Visibility

If you cannot see it, you cannot govern it.

☐ We maintain a current inventory of all AI systems in use (including shadow AI).
☐ Each AI system is classified by business impact (Low / Moderate / Material).
☐ Each system is classified by regulatory exposure (None / Moderate / High).
☐ Each AI system is mapped to the workflows it influences
☐ Data ingress points into AI systems are documented.
☐ Third-party vendors powering each AI capability are identified and documented.

☐ Each AI system is mapped to the workflows it influences.

Red Flag:
AI tools are purchased by departments without centralized visibility.

Board-Level Question:
If regulators asked tomorrow, could we identify every AI system influencing our business?

SECTION 2 — Named Accountability

Authority must be explicit. Not assumed.

☐ Every material AI system has a named executive owner.
☐ Accountability is formally assigned beyond IT implementation.
☐ Authority boundaries are defined (what the AI can and cannot decide).
☐ A stop-authority exists and is documented.
☐ Escalation paths are documented and tested.

Red Flag:
Accountability for AI systems is distributed across teams but owned by no one.

Board-Level Question:
If harm occurred tomorrow, who signs their name first?

SECTION 3 — Decision Authority & Control Boundaries

Capability does not equal authority.

☐ AI decision scope is formally defined and approved.
☐ AI output classification (advisory, assisted, autonomous) is documented and enforced.
☐ Human review requirements are documented.
☐ High-impact decisions require dual validation.
☐ Autonomous execution (if any) has defined containment controls.

Red Flag:
AI can initiate actions without clearly defined containment limits.

Board-Level Question:

Who has the authority to override or halt an AI-driven decision — and is that authority exercised?

SECTION 4 — First-, Second-, Third-Order Risk Cascade

The deeper the layer, the less reversible the risk.

☐ First-order risks (model failure, bias, hallucination) are documented.
☐ Second-order risks (ownership diffusion, workflow reshaping, authority drift) are identified.
☐ Third-order risks (institutional knowledge erosion, regulatory exposure, scale harm) are evaluated.
☐ Risk reporting cadence reaches executive leadership.
☐ Material risks reach the board.
☐ Governance roles were formally re-evaluated when AI altered workflow boundaries.

Board-Level Question:
Where is our risk deepest — and where is that risk least reversible?

Red Flag:
Risk analysis stops at model failure and does not examine

authority drift, workflow reshaping, or irreversible institutional impact.

SECTION 5 — Lifecycle Discipline (AISLC™ Alignment)

Lifecycle controls require accountable human oversight. Monitoring may be automated, but governance responsibility is not.

☐ Formal approval is required before deployment.
☐ Performance metrics are defined, active, and subject to human review.
☐ Drift detection mechanisms are in place, monitored, and escalated to accountable leadership.
☐ An incident response protocol is documented and owned by a named executive.
☐ A periodic governance review cadence is formally established and conducted.
☐ Decommissioning criteria are defined and enforceable.
☐ Model updates require governance review and executive approval.

Red Flag:
AI systems remain active indefinitely without review.

SECTION 6 — Documentation & Evidence

If it isn't documented, it isn't defensible.

☐ Deployment decisions are recorded.
☐ Model version history is tracked.
☐ Monitoring evidence is archived.
☐ Governance committee decisions are formally documented and retained.
☐ Audit trail exists for material decisions.

Red Flag:
Material AI decisions cannot be reconstructed with documented evidence.

Board-Level Question:
Can we reconstruct how this system made decisions six months ago?

SECTION 7 — Workflow Integrity (JTBD Alignment)

☐ AI use is aligned to a defined Job To Be Done.
☐ Workflow redesign is formally approved.
☐ Downstream impact analysis is performed before scale.
☐ AI is not merely accelerating broken processes.
☐ Executive communications are safeguarded from AI summarization distortion.

Red Flag:
AI is summarizing work instead of transforming it.

Board-Level Question:
Are we accelerating efficiency — or accelerating flawed process design?

SECTION 8 — Procurement & Tool Proliferation

☐ AI subscriptions require governance review.
☐ Autonomous or agentic systems require centralized approval.
☐ Vendor claims are independently validated.
☐ Data-sharing terms are reviewed by legal.
☐ AI-generated integrations, plugins, or agentic extensions are subject to the same approval and oversight controls as primary systems.
☐ Tool redundancy is evaluated annually.

Red Flag:
AI systems initiate or connect to additional tools, services, or integrations without centralized governance review.

Board-Level Question:
How many AI systems do we pay for that no one governs?

Interpretation Standard

Unchecked controls represent governance gaps and unsurfaced risk.
Multiple gaps compound structural exposure.

Governance Maturity Indicator

- 0–25% Checked → High Structural Risk
- 26–50% → Governance Immature

- 51–75% → Operational but Vulnerable
- 76–90% → Structured but Actively Monitored
- 91–100% → Governance Embedded and Evidenced

AI scales what already exists.

If governance is unclear, AI scales ambiguity.
If accountability is weak, AI scales diffusion.
If documentation is thin, AI scales liability.

Governance precedes scale.

Truth Before It Costs Millions™

Prepared by:
Tom Staskiewicz, MBA, CISSP
AI Governance & Structural Risk Advisor
Former IBM Systems Engineer