AI Governance

Minimum Viable Control Model - Public Reference Edition

Author: Hypotenuse Consulting Ltd

Status: v1.0

Classification: Public - Reference Edition

Date: 2026-01-01

This document is released publicly for reference purposes only and does not constitute legal, contractual, or professional advice.

Contents

1 Executive Summary 2

2 Scope and Applicability 2

2.1 AI Systems in Scope 2

2.2 Exclusions 2

3 Standards and Alignment Context 2

4 Governance Control Objectives 3

5 Governance Control Activities 3

6 Evidence and Assurance Expectations 3

6.1 Minimum Evidence Requirements 3

6.2 Evidence Purpose 4

7 Roles and Responsibilities 4

8 Key Risks and Governance Gaps 4

9 Recommendations 4

10 Next Steps and Decisions Required 5

1 Executive Summary

This document defines a minimum viable control model for the governance of artificial intelligence systems at board and executive level. It is intended to support oversight, accountability, and defensible decision-making where AI capabilities introduce material and often poorly understood operational, legal, or reputational risk.

The model focuses on governance intent and control objectives rather than technical implementation, development practices, or tooling.

2 Scope and Applicability

2.1 AI Systems in Scope

This model applies to the governance of AI systems that:

- Are developed, procured, or operated by the organization
- Influence decisions, outcomes, or processes with material impact
- Introduce regulatory, ethical, safety, or reputational considerations

2.2 Exclusions

This reference does not define:

- Technical controls or model architectures
- Development or deployment standards
- Model training, testing, or monitoring procedures

Those aspects are delegated to management and technical functions.

3 Standards and Alignment Context

This control model is intended to align with:

- Emerging AI regulatory regimes and guidance
- ISO 27001 and related governance standards
- Enterprise risk management practices
- Data protection and accountability principles

Alignment provides structure but does not replace organization-specific judgement or accountability.

4 Governance Control Objectives

The primary governance objectives are to ensure that:

- AI use is intentional, approved, and traceable
- Accountability for AI-driven outcomes is explicitly assigned
- Risks introduced by AI are identified and consciously accepted
- AI systems can be challenged, suspended, or retired when required
- The organization can demonstrate responsible oversight to regulators and stakeholders

These objectives apply regardless of whether AI capability is developed internally or procured externally.

5 Governance Control Activities

At governance level, control activities typically include:

- Approval of AI use cases based on defined risk criteria
- Classification of AI systems by impact and criticality
- Assignment of accountable owners for AI outcomes
- Periodic review of AI system performance and assumptions
- Escalation processes for AI-related incidents or concerns

Operational implementation and technical execution are delegated to management.

6 Evidence and Assurance Expectations

6.1 Minimum Evidence Requirements

Minimum evidence to support governance and assurance includes:

- A register of approved AI use cases and classifications
- Documented accountability for each AI system
- Records of risk assessment and risk acceptance decisions
- Outcomes of periodic governance reviews
- Decisions to suspend, modify, or retire AI systems

6.2 Evidence Purpose

Evidence should support defensible decision-making. It should not be optimized for exhaustive compliance or volume.

7 Roles and Responsibilities

Clear separation of responsibilities is required:

- Board: sets risk appetite and approves high-impact AI use
- Executive management: ensures governance is implemented and effective
- AI or technology leadership: advises on capabilities, limitations, and risks
- Assurance functions: independently review governance effectiveness

Accountability for AI outcomes must not default to technical teams alone.

8 Key Risks and Governance Gaps

Common governance failures include:

- Deployment of AI without explicit approval
- Unclear ownership of AI-driven outcomes
- Over-reliance on vendor assurances or certifications
- Inability to explain or justify AI decisions
- Lack of mechanisms to pause, withdraw, or decommission AI systems

This model is designed to surface and mitigate these risks.

9 Recommendations

Organizations adopting this model should:

- Define clear criteria for AI approval and escalation
- Limit governance metrics to decision-relevant indicators
- Separate governance oversight from delivery and experimentation
- Regularly test governance assumptions through scenario discussion
- Treat AI withdrawal and suspension as planned capabilities

10 Next Steps and Decisions Required

Boards using this reference should determine:

- Which AI uses require explicit board approval
- Acceptable levels of AI-related risk and uncertainty
- Evidence required to support continued AI operation
- Conditions under which AI systems must be paused, modified, or retired

This model should be adapted to organizational context and reviewed as AI capabilities and regulatory expectations evolve.