DENDRO LOGIC
INTELLIGENT DATA

AI Governance for Enterprise Organisations

Frameworks, Policies, and Operating Models
for Auditable AI at Scale

ISO 42001, NIST AI RMF, EU AI Act alignment, enterprise risk management, third-party AI governance, responsible AI, and the complete enterprise AI policy stack

By Mike McGreal

Dendro Logic Ltd
March 2026
Business AI Adoption Playbook

https://dendro-logic.com

Contents

Part 1: Why Enterprise Governance Is Different

The Scale Problem

Document 4 of this series provides a data sovereignty and AI security framework that works for organisations of all sizes. Document 5 provides a minimum viable governance structure for AI strategy. This document extends both into the formal, auditable governance operating model that large organisations require.

A 50-person company can manage AI governance with a named lead, an approved tools registry, and a quarterly review. A 500-person or 5,000-person organisation faces fundamentally different challenges. Shadow AI is not one or two employees using ChatGPT, it is dozens of teams across multiple business units using AI tools embedded in products they already rely on. Vendor risk is not one provider to monitor, it is scores of software suppliers quietly integrating AI into their products, often without notifying customers. Regulatory exposure spans multiple jurisdictions, potentially the UK, the EU (through the AI Act), and sector-specific requirements from the FCA, NHS, SRA, or other regulators simultaneously.

PwC's 29th Global CEO Survey found that the 12% of organisations achieving both cost and revenue gains from AI, the "vanguard," share a common characteristic: strong AI foundations, including responsible AI frameworks and technology environments that enable enterprise-wide integration. ISO 42001, the first international AI management system standard, provides the certifiable framework for building those foundations. The NIST AI Risk Management Framework provides the risk methodology. The EU AI Act provides the legal requirements. Together they form a governance stack that enterprise organisations need to implement systematically.

Who This Document Is For

This playbook is aimed at governance leads, compliance officers, CISOs, CTOs, Data Protection Officers, and board members in organisations with 250 or more employees, those operating in regulated industries (financial services, healthcare, legal, public sector), or those with international operations, particularly EU exposure. If your organisation already manages ISO 27001, ISO 9001, or similar management systems, this document shows how to extend those systems to cover AI governance without building a parallel structure.

Series Cross-Reference

This document builds directly on Document 4 (Data Sovereignty and AI Security) and Document 5 (AI Strategy for Business Leaders). Where those documents cover a topic in full, this document references rather than duplicates. Where enterprise scale requires a different approach, this document provides it. Read Document 4 first for the foundational sovereignty framework, then this document for the enterprise governance layer.

Part 2: The Global Framework Landscape

Three frameworks define enterprise AI governance in 2026. They are not competing alternatives, they are complementary layers of a single governance stack. Understanding how they relate to each other is essential for avoiding duplicated effort.

ISO/IEC 42001: The Certifiable AI Management System

ISO/IEC 42001, published in December 2023, is the world's first international standard for AI Management Systems (AIMS). It is to AI governance what ISO 27001 is to information security and ISO 9001 is to quality management. It provides a structured, auditable framework for establishing, implementing, maintaining, and continually improving how your organisation governs AI.

The standard follows the Plan-Do-Check-Act methodology through 10 structured clauses, seven of which contain mandatory requirements: context of the organisation, leadership, planning, support, operation, performance evaluation, and improvement. Critically, ISO 42001 is certifiable: an independent, accredited auditor can verify your compliance and issue a formal certificate. This is the difference between claiming you govern AI responsibly and having a third party verify it.

For organisations already running ISO 27001 or ISO 9001, ISO 42001 integrates directly. It uses the same Annex SL management system structure, meaning your existing governance architecture, internal audit processes, and management review cycles can be extended to cover AI without building from scratch. The standard includes Annex A controls specific to AI: risk assessment, data governance, transparency, human oversight, and continuous monitoring.

NIST AI Risk Management Framework

The NIST AI RMF, developed by the US National Institute of Standards and Technology, is a voluntary risk management framework structured around four core functions: Govern (establishing a risk management culture and governance structures), Map (identifying AI system context, risks, and benefits), Measure (using quantitative and qualitative tools to analyse and track AI risks), and Manage (allocating resources to treat identified risks and developing response strategies).

Although voluntary and not certifiable, the NIST AI RMF's influence exceeds its voluntary status. US federal procurement increasingly expects NIST alignment. The FTC, CFPB, FDA, SEC, EEOC, and Department of Defense all reference its principles. Enterprise customers use it as the benchmark for evaluating vendor AI governance maturity. For UK

organisations with US clients or operations, NIST alignment is increasingly a procurement requirement.

The Generative AI Profile (NIST AI 600-1) extends the framework specifically for generative AI risks, including hallucination, data provenance, and content authenticity. The Cyber AI Profile (IR 8596, preliminary draft December 2025) bridges AI risk management with cybersecurity across three areas: securing AI systems, using AI for cyber defence, and defending against AI-enabled threats.

The EU AI Act

The EU AI Act is the most comprehensive AI legislation globally and affects UK organisations that serve EU clients, process EU personal data, or operate in EU markets. Its obligations phase in on a timeline that runs from 2025 into 2027. The May 2026 Digital Omnibus agreement has since deferred the main high-risk obligations to December 2027.

February 2025: Prohibited AI practices banned (social scoring, untargeted facial recognition scraping, emotion recognition in workplaces and schools). AI literacy obligations begin for all providers and deployers.

August 2025: Governance infrastructure operational. General-purpose AI (GPAI) model obligations begin. Providers must maintain technical documentation, publish training data summaries, and implement policies to address risks.

August 2026 / December 2027: Transparency obligations (Article 50) apply from August 2026. The high-risk AI system obligations (Annex III) were due at the same point, but the May 2026 Digital Omnibus agreement postponed them to 2 December 2027, and conformity assessment and the related duties shift with them. This is the enforcement milestone that matters most for enterprise organisations: the extra time is for readiness, not a reprieve.

August 2027: GPAI models already on the market before August 2025 must comply.

High-risk categories include AI systems used in employment (hiring, performance evaluation, promotion), creditworthiness assessment, education (determining access or assessment), essential services (insurance, healthcare prioritisation), law enforcement, and critical infrastructure. If your organisation uses AI in any of these areas and serves EU markets, you need to prepare for August 2026 enforcement.

The UK Regulatory Position

The UK maintains a principles-based, pro-innovation approach to AI oversight, which means there is no single UK AI Act equivalent. Instead, existing regulators (FCA, ICO, Ofcom, CMA, MHRA) apply AI principles within their existing mandates. The DUAA (Data

Use and Access Act), which received Royal Assent in June 2025, relaxed some automated decision-making restrictions while introducing transfer risk assessment requirements.

The ICO's AI and Biometrics Strategy, published in June 2025, signals a forthcoming statutory code of practice on AI and automated decision-making. AISI (the AI Safety Institute) provides technical evaluation and testing guidance. For UK enterprises, the practical position is: there is no single AI regulation to comply with, but multiple regulators expect you to demonstrate responsible AI governance within their respective frameworks. ISO 42001 provides the single management system that satisfies this expectation across regulators.

How They Layer Together

```
THE ENTERPRISE AI GOVERNANCE STACK

Layer 3: REGULATION (legal requirements)
  EU AI Act, UK DUAA, FCA rules, ICO guidance,
  sector-specific requirements
  What you MUST do.

Layer 2: STANDARD (certifiable management system)
  ISO/IEC 42001
  How you PROVE you do it.

Layer 1: FRAMEWORK (risk methodology)
  NIST AI RMF (Govern, Map, Measure, Manage)
  How you THINK about risk.

These are not alternatives. They are layers.
  NIST provides the risk thinking.
  ISO 42001 provides the auditable system.
  Regulation provides the legal requirements.

Published crosswalks exist between all three.
Implementing them together produces no
duplicated effort if mapped correctly.
```

Part 3: Building an AI Management System

An AI Management System (AIMS) is the structured set of policies, processes, roles, and controls that govern how your organisation develops, deploys, uses, and retires AI systems. ISO 42001 provides the standard for this. If you already run an ISO 27001 Information Security Management System, you have most of the structural foundations in place.

The Plan-Do-Check-Act Cycle for AI

Plan. Define the scope of your AIMS (which AI systems, which business units, which geographies). Identify risks and opportunities. Set AI governance objectives. Establish the AI policy, roles, and responsibilities.

Do. Implement the controls. Deploy the AI register. Run risk assessments. Establish the policy stack. Train staff. Implement technical controls (monitoring, logging, access control, bias testing).

Check. Monitor and measure performance. Run internal audits. Conduct management reviews. Track incidents and near-misses. Evaluate whether controls are working as intended.

Act. Address nonconformities. Implement corrective actions. Update the AIMS based on lessons learned, changing risks, and evolving regulations. Feed improvements back into the Plan phase.

For organisations with existing ISO management systems, the extension to AI is primarily about scope and controls, not about rebuilding the PDCA infrastructure. Your internal audit programme, management review process, document control system, and corrective action procedures can all be extended rather than replicated.

Scope Definition

The first decision is what your AIMS covers. This must be explicit and documented. A scope that is too broad becomes unmanageable. A scope that is too narrow leaves significant AI risk ungovemed.

```
AIMS SCOPE DEFINITION TEMPLATE

Organisation: _____________________________________

AI systems in scope:
  [ ] AI tools used by employees (Copilot, ChatGPT,
      Claude, Gemini, etc.)
  [ ] AI embedded in vendor products (CRM, HR, finance,
      customer service platforms)
  [ ] AI developed or customised internally
  [ ] AI used in products/services delivered to clients
  [ ] AI agents and automated workflows

Business units in scope:
  [ ] All business units
  [ ] Specific units: _____________________________

Geographies in scope:
  [ ] UK only
  [ ] UK + EU (AI Act obligations)
  [ ] UK + US (NIST alignment)
  [ ] Global

Exclusions (with justification):
  _____________________________________________

Regulatory obligations in scope:
  [ ] UK GDPR / DUAA
  [ ] EU AI Act (high-risk categories)
  [ ] FCA / PRA requirements
  [ ] ICO AI guidance
  [ ] NHS / MHRA requirements
  [ ] SRA requirements
  [ ] Other: _____________________________________

Approved by: _____________ Date: _____________
Review date: _____________
```

Leadership Requirements

ISO 42001 requires that top management demonstrates leadership and commitment to the AIMS. This is not a checkbox exercise. It means the board or senior leadership team must establish the AI policy, ensure the AIMS is integrated into business processes, ensure resources are available, and communicate the importance of effective AI governance. The standard explicitly requires that someone at a senior level is accountable for AI governance.

In practical terms, this means a named executive sponsor (CEO, COO, or equivalent) who owns the AI governance agenda, and a named AI Governance Lead who manages day-to-day operations. For organisations establishing a Chief AI Officer role, this person typically fulfils the governance lead function and reports directly to the executive sponsor.

Part 4: The Enterprise AI Register

The AI Register is the living inventory of every AI system in use across your organisation. It serves triple duty: it satisfies ISO 42001 Clause 8 requirements for operational planning, supports the NIST AI RMF Map function, and provides the system inventory required for EU AI Act compliance. Without it, governance is guesswork.

What to Record

For each AI system, record the following. This register must be maintained as a living document with quarterly reviews to ensure governance keeps pace with AI system changes.

```
ENTERPRISE AI REGISTER - RECORD TEMPLATE

System ID: AI-______ (unique identifier)

IDENTIFICATION
  System name: _____________________________________
  Vendor/provider: _________________________________
  Version: _________________________________________
  Category:
    [ ] Employee productivity tool
    [ ] Embedded in vendor product
    [ ] Internally developed / customised
    [ ] Client-facing product/service
    [ ] Autonomous agent / workflow
  Description: _____________________________________
  Business owner: __________________________________
  Technical owner: _________________________________

DATA AND PROCESSING
  Input data types: ________________________________
  Data classification (Doc 4): _____________________
  Processing location(s): __________________________
  Data retention: __________________________________
  Sub-processors: __________________________________
  Does data leave UK? [ ] Yes [ ] No
  Does data enter training? [ ] Yes [ ] No
  DPA in place? [ ] Yes [ ] No [ ] N/A

RISK CLASSIFICATION
  EU AI Act risk tier:
    [ ] Minimal  [ ] Limited  [ ] High  [ ] N/A
  Internal risk rating:
    [ ] Low  [ ] Medium  [ ] High  [ ] Critical
```

```
  Human oversight level:
    [ ] Fully automated
    [ ] Human-in-the-loop (reviews before action)
    [ ] Human-on-the-loop (monitors outcomes)
    [ ] Human-in-command (can override at any time)

GOVERNANCE
  Approved by: _____________ Date: _____________
  Last risk assessment: _____________________________
  Next review date: _________________________________
  DPIA required? [ ] Yes [ ] No
  DPIA completed? [ ] Yes [ ] No
  Incident history: _________________________________

DEPENDENCIES
  Integrates with: __________________________________
  Downstream systems: _______________________________
  Fallback if unavailable: __________________________
```

Discovery: Finding AI You Do Not Know About

Most enterprise organisations significantly underestimate the number of AI systems in use. Zylo's 2026 SaaS Management Index found that AI-native spending nearly doubled year-on-year and ChatGPT became the number one most-expensed application by transaction count. This means employees are adopting AI tools through expense claims, bypassing IT procurement entirely.

Discovery requires multiple approaches: IT asset and licence audit (what AI tools are in the software inventory), expense analysis (what AI subscriptions are being expensed), vendor audit (which existing vendors have added AI features to their products), network analysis (DNS and egress monitoring for AI API endpoints, as detailed in Document 4), and employee survey (what AI tools are people actually using, with an amnesty for undeclared tools during the initial discovery phase).

The discovery phase typically reveals 2-5 times more AI systems than leadership expected. This is not a failure of governance, it is the starting point for governance. Document every system found, classify its risk, and bring it into the register before making any decisions about approval or restriction.

The Vendor AI Problem

Your existing vendors are adding AI to their products, often without explicit notification. Microsoft has embedded Copilot features across M365. Salesforce has integrated Einstein AI across its platform. HubSpot, Zendesk, ServiceNow, and dozens of other enterprise platforms have added AI capabilities. Your AI register must include AI embedded in products you already

use, not just standalone AI tools. Vendor contract review and terms monitoring (covered in Document 4) is essential for keeping the register current.

Part 5: AI Risk Assessment at Enterprise Scale

Enterprise AI risk assessment extends beyond individual system evaluation to portfolio-level risk management. You need to understand the risk profile of each AI system individually, the aggregate risk across your AI portfolio, and the interdependencies between AI systems and critical business processes.

The Risk Classification Framework

Classify every AI system in your register using a four-tier framework that aligns with the EU AI Act risk categories while adding practical enterprise detail.

```
AI RISK CLASSIFICATION FRAMEWORK

TIER 1: MINIMAL RISK
  Examples: spell-checking, document formatting,
  internal search, meeting transcription
  Controls: standard acceptable use policy,
  basic monitoring
  Review: annual

TIER 2: LIMITED RISK
  Examples: content generation for internal use,
  email drafting, data summarisation, code
  generation with human review
  Controls: human review before external use,
  output quality monitoring, prompt logging
  Review: bi-annual

TIER 3: HIGH RISK
  Examples: hiring/recruitment assistance, credit
  scoring, performance evaluation, client
  advisory support, medical/health applications,
  insurance pricing, legal research
  Controls: DPIA required, bias testing,
  human-in-the-loop decision-making, full
  audit trail, explainability documentation,
  regular fairness assessment
  Review: quarterly

TIER 4: UNACCEPTABLE RISK
  Examples: social scoring, emotion recognition
  in workplace, real-time biometric ID in
  public spaces, manipulation of vulnerable
  groups
  Decision: DO NOT DEPLOY. These are prohibited
  under the EU AI Act and should be
```

```
  prohibited under any responsible AI policy.
```

Individual System Risk Assessment

For each AI system rated Tier 2 or above, complete a structured risk assessment. This feeds into both your enterprise risk register and your ISO 42001 risk treatment records.

```
AI SYSTEM RISK ASSESSMENT

System ID: AI-_______ Name: _____________________

RISK IDENTIFICATION

Risk Category        | Description         | Likelihood | Impact | Score
---------------------|---------------------|------------|--------|------
Data breach/leakage  |                     | 1-5        | 1-5    |
Inaccurate output    |                     | 1-5        | 1-5    |
Biased output        |                     | 1-5        | 1-5    |
Regulatory breach    |                     | 1-5        | 1-5    |
Reputational harm    |                     | 1-5        | 1-5    |
Vendor dependency    |                     | 1-5        | 1-5    |
Service disruption   |                     | 1-5        | 1-5    |
Scope creep          |                     | 1-5        | 1-5    |
Employee misuse      |                     | 1-5        | 1-5    |
Insurance exclusion  |                     | 1-5        | 1-5    |

OVERALL RISK SCORE: _______

RISK TREATMENT
For each identified risk above Medium:
  Risk: ____________________________________________
  Treatment: [ ] Mitigate [ ] Transfer [ ] Accept
  Controls: ________________________________________
  Residual risk: ___________________________________
  Owner: ___________________________________________

Assessed by: _______________ Date: _______________
Approved by: _______________ Date: _______________
Next review: _______________
```

Portfolio-Level Risk Reporting

The board and risk committee need aggregated AI risk reporting, not individual system assessments. Produce a quarterly AI risk summary that shows total AI systems in the register, distribution across risk tiers, systems with overdue reviews, open risk items, incidents and near-misses in the period, and the aggregate risk trend (improving, stable, or deteriorating). This should integrate into your existing enterprise risk reporting rather than running as a separate process.

Part 6: Third-Party AI Risk Management

Third-party AI risk is the governance challenge that most enterprise organisations underestimate. It is not just about the AI tools you buy deliberately. It is about the AI features your existing vendors add to products you already depend on, often without requiring your explicit consent.

The Vendor AI Lifecycle

Vendor AI risk management covers four phases: pre-procurement assessment (before you buy), onboarding verification (when you deploy), ongoing monitoring (while you use), and exit management (when you leave or the vendor changes).

Pre-procurement. Before purchasing any AI-enabled product, complete the Vendor AI Evaluation Checklist from Document 5 plus a full data sovereignty assessment from Document 4. For Tier 3 (high-risk) AI systems, require the vendor to provide evidence of their own AI governance framework, including their approach to bias testing, data handling, and incident response. If the vendor cannot provide this evidence, that is itself a risk indicator.

Onboarding. Verify that the deployed product matches what was assessed. Confirm data processing locations, retention settings, and access controls. Ensure the DPA is signed and the data classification is appropriate. Add the system to the AI register with all mandatory fields completed.

Ongoing monitoring. This is where most governance programmes fail. Vendors change their terms, add AI features, update models, and modify data processing arrangements, sometimes without notification. Document 4's Approved-But-Changed monitoring framework applies here at enterprise scale. Implement quarterly vendor terms review for all Tier 2+ AI systems and annual review for Tier 1 systems.

Exit management. Before a vendor relationship ends, confirm data deletion, export any data you need, verify the AI system is removed from all integrations, and update the AI register. Vendor lock-in risk should be assessed at procurement stage (see Document 4's Vendor Lock-In and Exit Strategy Checklist).

Third-Party AI Risk Register

```
THIRD-PARTY AI RISK REGISTER

Vendor: _____________________________________
Product: ____________________________________
AI features: ________________________________

Risk Area             | Status    | Last Check
----------------------|-----------|----------
DPA signed            | Y/N/Exp   |
Data location verified| Y/N       |
Training opt-out conf | Y/N/N/A   |
Terms stable          | Y/Changed |
Sub-processors known  | Y/N       |
AI features disclosed | Y/Partial |
Bias testing provided | Y/N/N/A   |
Incident process doc  | Y/N       |
Exit plan documented  | Y/N       |
Insurance coverage    | Y/N/Check |

Actions required: ___________________________
Risk owner: _________________________________
Next review: ________________________________
```

Part 7: Responsible AI in Practice

Responsible AI is not a separate initiative alongside governance. It is the set of practices that make governance meaningful: ensuring AI systems are fair, transparent, explainable, and subject to appropriate human oversight. For enterprise organisations, this means moving from principles to operational processes.

Bias Testing and Fairness Assessment

For any AI system classified as Tier 3 (high risk), particularly those involved in hiring, performance evaluation, credit decisions, insurance pricing, or any other decision that materially affects individuals, bias testing is not optional. The UK Equality Act 2010 applies to AI-assisted decisions just as it applies to human decisions. If an AI system produces outcomes that disproportionately disadvantage individuals based on protected characteristics (age, disability, gender reassignment, marriage and civil partnership, pregnancy and maternity, race, religion or belief, sex, sexual orientation), the organisation faces legal liability regardless of whether the bias was intentional.

Demographic parity testing. Measure whether the AI system produces similar outcomes across different demographic groups. If an AI recruitment screening tool advances 60% of male applicants but only 35% of female applicants, that disparity requires investigation and remediation.

Input audit. Review the data the AI system uses for decision-making. Proxy variables can introduce bias indirectly: postcode can proxy for race, university attended can proxy for socioeconomic background, name patterns can proxy for ethnicity. Document which inputs the system uses and assess each for proxy bias risk.

Output monitoring. Continuously monitor the AI system's outputs for patterns that suggest bias. This is not a one-time test. Bias can emerge over time as input data distributions shift or as the AI model is updated. Implement automated monitoring that flags statistically significant outcome disparities.

Documented remediation. When bias is identified, document the finding, the root cause analysis, the remediation action, and the verification that the remediation was effective. This documentation is essential for regulatory compliance and for defending decisions if challenged.

Transparency and Explainability

The DUAA maintains requirements for transparency in automated decision-making. The EU AI Act requires that high-risk AI systems are designed to be sufficiently transparent for deployers to interpret and use outputs appropriately. In practice, this means different things for different AI use cases.

For internal productivity tools (Tier 1-2): Users should understand what AI is doing with their data and how AI-generated outputs should be reviewed. This is covered by the acceptable use policy and training programme.

For decision-support systems (Tier 3): The decision-maker must understand how the AI arrived at its recommendation, what data it used, and what its limitations are. Outputs should include confidence indicators where available. The human decision-maker must be able to explain, in their own words, why they agreed or disagreed with the AI recommendation.

For automated decisions affecting individuals: The individual affected must be informed that AI was used in the decision, given meaningful information about the logic involved, and provided with a route to contest the decision and obtain human review. This is a DUAA requirement for significant automated decisions.

Human Oversight Framework

```
HUMAN OVERSIGHT LEVELS

Level 1: HUMAN-IN-THE-LOOP
  AI generates, human decides.
  Use for: all Tier 3 systems, client-facing
  content, decisions affecting individuals.
  The human reviews every AI output before
  it is acted upon or communicated.

Level 2: HUMAN-ON-THE-LOOP
  AI acts, human monitors.
  Use for: Tier 2 systems with established
  accuracy, high-volume internal processes.
  The human monitors aggregate outputs and
  spot-checks individual outputs regularly.
  The human can intervene at any time.

Level 3: HUMAN-IN-COMMAND
  AI operates autonomously within defined
  boundaries. Human can override or shut
  down at any time.
  Use for: Tier 1-2 systems only, with
```

```
  well-understood behaviour, robust testing,
  and clear operational boundaries.
  Not appropriate for any Tier 3 system.

Level 4: FULLY AUTONOMOUS
  No human oversight of individual outputs.
  ONLY appropriate for Tier 1 systems with
  no impact on individuals and no external
  communication. Must still have monitoring
  for drift and aggregate quality.

Map every AI system in your register to
one of these levels. Document the rationale.
Review whenever the system changes scope.
```

Part 8: The Enterprise AI Policy Stack

An enterprise AI governance programme requires a set of interconnected policies. These are not standalone documents, they form a coherent stack where each policy references the others and all align with the AIMS. Below is the complete set, with the scope and key contents of each.

Policy Overview

```
ENTERPRISE AI POLICY STACK

1. AI Governance Policy (master document)
   Scope, objectives, roles, AIMS structure,
   regulatory alignment, review cycle.
   Audience: Board, all staff

2. AI Acceptable Use Policy
   What employees can/cannot do with AI,
   approved tools, data classification rules,
   review requirements, prohibited uses.
   Audience: All employees
   Cross-ref: Document 3 (SOPs)

3. AI Procurement Policy
   Vendor AI evaluation, DPA requirements,
   sovereignty checks, pilot agreement terms,
   approval workflow for new AI tools.
   Audience: Procurement, IT, budget holders
   Cross-ref: Document 4 (sovereignty), Doc 5

4. AI Data Governance Policy
   Data classification for AI, PII handling,
   prompt sanitisation, data minimisation,
   cross-border transfer rules.
   Audience: Data teams, IT, all AI users
   Cross-ref: Document 4 (full framework)

5. AI Model Risk Management Policy
   Risk assessment methodology, tier
   classification, bias testing requirements,
   monitoring, model change management.
   Audience: Technical teams, risk function

6. AI Incident Response Policy
   Detection, classification, containment,
   notification, remediation, post-incident
   review for AI-specific incidents.
```

```
   Audience: IT, security, legal, comms
   Cross-ref: Document 4 (incident response)

7. Third-Party AI Policy
   Vendor AI assessment, ongoing monitoring,
   terms review, exit management.
   Audience: Procurement, vendor management
   Cross-ref: Document 4 (vendor checklist)

8. Employee AI Rights Policy
   Rights regarding AI-assisted decisions
   about them (transparency, explanation,
   human review, contest).
   Audience: HR, all employees

9. AI Ethics Policy (or Responsible AI Policy)
   Fairness principles, bias prevention,
   transparency commitments, human oversight
   requirements, prohibited applications.
   Audience: Board, all staff
```

Policy Governance

Each policy needs a named owner, a review cycle (annual minimum, more frequent for rapidly evolving areas), a version control system, and an approval process. Changes to Tier 3 (high-risk) policies should require board or AI governance committee approval. Changes to operational policies can be approved by the AI Governance Lead.

Avoid policy sprawl. Nine policies sounds like a lot, but several of these can be short (2-3 pages each) and reference the detailed operational procedures in Documents 1-5 of this series rather than duplicating them. The governance policy is the master document. The rest are focused, audience-specific documents that implement specific aspects of the governance framework.

Part 9: Governance Structure for Enterprise Scale

Document 5 provides a minimum viable governance structure. Enterprise organisations need a more formal structure with clear decision rights, escalation procedures, and committee governance that integrates with existing corporate governance.

The Three-Layer Governance Model

```
ENTERPRISE AI GOVERNANCE STRUCTURE

LAYER 1: STRATEGIC (Board / ExCo)

AI Governance Board (or subcommittee of the
Board Risk Committee)
  Chair: CEO / COO / Non-Exec with AI experience
  Members: CAIO, CTO, CISO, CDO, CFO, GC, CHRO
  Frequency: Quarterly
  Remit:
    - AI strategy approval
    - Risk appetite for AI
    - Major investment decisions
    - Regulatory compliance oversight
    - Annual AIMS review

LAYER 2: OPERATIONAL (Cross-functional)

AI Governance Committee
  Chair: Chief AI Officer / AI Governance Lead
  Members: representatives from IT, Security,
    Legal, HR, Data, Compliance, plus BU leads
  Frequency: Monthly
  Remit:
    - AI register maintenance
    - Risk assessment review and approval
    - Policy updates and implementation
    - Incident review
    - Vendor AI monitoring
    - Training programme oversight

AI Ethics Advisory Panel (optional)
  Chair: Independent or senior leader
  Members: cross-functional + external advisors
  Frequency: Quarterly or as needed
  Remit:
    - Ethical review of Tier 3 AI deployments
    - Bias testing oversight
    - Stakeholder impact assessment
```

```
    - Emerging ethical issues

LAYER 3: EMBEDDED (Business Units)

Business Unit AI Leads
  One per major business unit
  Report to: BU head + dotted line to AI
    Governance Lead
  Remit:
    - BU-level AI register maintenance
    - First-line risk assessment
    - Training coordination
    - Feedback and incident reporting
    - Prompt library management

Team Champions (per department, as in Doc 5)
  Report to: BU AI Lead
  Remit: day-to-day support, feedback, adoption
```

Decision Rights Framework

Clear decision rights prevent governance bottlenecks while ensuring appropriate oversight. The RACI matrix below defines who is Responsible, Accountable, Consulted, and Informed for key AI governance decisions.

```
AI GOVERNANCE RACI MATRIX

Decision                    | R    | A    | C    | I
----------------------------|------|------|------|------
AI strategy and budget      | CAIO | CEO  | CFO  | Board
New Tier 1 AI tool approval | BU   | CAIO | IT   | GovCom
New Tier 2 AI tool approval | CAIO | CAIO | IT,L | GovCom
New Tier 3 AI deployment    | CAIO | Board| L,HR | All
AI policy changes (Tier 1-2)| CAIO | CAIO | L    | GovCom
AI policy changes (Tier 3)  | CAIO | Board| L,HR | All
Vendor AI terms change      | IT   | CAIO | L    | BU
AI incident (Low)           | IT   | CAIO | -    | GovCom
AI incident (High/Critical) | CAIO | CEO  | L,PR | Board
Bias test failure           | CAIO | Board| L,HR | GovCom
DPIA for AI system          | DPO  | CAIO | L    | Board
AI register update          | BU   | CAIO | IT   | GovCom

Key: CAIO=Chief AI Officer/AI Governance Lead
     BU=Business Unit Lead, L=Legal, GC=General
     Counsel, IT=IT/Security, HR=HR, DPO=Data
     Protection Officer, GovCom=Governance Committee
```

The Chief AI Officer Question

Not every organisation needs a Chief AI Officer. But every organisation needs someone who owns AI governance at a sufficiently senior level to influence strategy, budget, and risk decisions. In smaller enterprises (250-500 employees), this might be the CTO or Head of Digital with an expanded mandate. In larger enterprises, a dedicated CAIO reporting to the CEO or COO becomes necessary when AI is pervasive across business functions and the governance burden exceeds what can be added to an existing role.

The role combines three functions that are often separate: technology strategy (which AI systems to adopt and how to integrate them), risk and compliance (ensuring AI use meets regulatory and ethical requirements), and change management (driving adoption, training, and cultural transformation). The person in this role needs to be credible with both the board and the engineering teams. Pure technologists and pure compliance professionals both struggle in this role. The most effective AI governance leaders combine technical literacy with business acumen and stakeholder management skills.

Part 10: Audit and Assurance

Governance without audit is aspiration. Audit provides the evidence that your AIMS works in practice, not just on paper. For organisations pursuing ISO 42001 certification, audit is mandatory. For all enterprise organisations, it is essential for demonstrating due diligence to regulators, clients, and insurers.

Internal Audit Programme

Your internal AI audit programme should cover four areas: policy compliance (are policies being followed in practice), technical controls (are monitoring, logging, and access controls working), risk management (are risk assessments current and accurate), and operational effectiveness (is the governance structure functioning, are reviews happening on schedule, are incidents being reported and resolved).

```
INTERNAL AI AUDIT PLAN

ANNUAL CYCLE

Q1: Policy compliance audit
  - Sample check: are approved tools being used
    according to the acceptable use policy?
  - Are Tier 3 systems operating with the
    documented level of human oversight?
  - Are training records current?
  - Are prompt libraries maintained?

Q2: Technical controls audit
  - Network egress monitoring (Doc 4)
  - Data classification enforcement
  - Access control verification
  - Vendor terms review completion
  - AI register accuracy (spot check)

Q3: Risk management audit
  - Are risk assessments current for all
    Tier 2+ systems?
  - Are bias tests running on schedule
    for Tier 3 systems?
  - Is the incident response plan tested?
  - Are DPIAs completed where required?

Q4: AIMS effectiveness review
  - Management review preparation
  - Aggregate metrics and trends
  - Gap analysis against ISO 42001
```

```
  - Improvement recommendations
  - Annual report to the board

Each audit: documented findings, corrective
actions with owners and deadlines, follow-up
verification.
```

Preparing for ISO 42001 Certification

ISO 42001 certification follows a two-stage audit process conducted by an accredited certification body. Auditors must meet the qualification requirements of BS ISO/IEC 42006:2025, ensuring they have specific AI governance expertise, not just generic management system auditing experience.

Stage 1 (documentation review): The auditor reviews your AIMS documentation, including the AI policy, risk assessments, AI register, and procedures. They confirm that the management system is designed to meet the standard's requirements. Any major gaps are identified for remediation before Stage 2.

Stage 2 (implementation audit): The auditor verifies that the AIMS is implemented and operating effectively in practice. This includes interviews with staff, observation of processes, review of records and evidence, and sampling of AI systems from the register. Nonconformities are classified as major (preventing certification until resolved) or minor (requiring a corrective action plan).

Certification is valid for three years, with annual surveillance audits to confirm continued compliance. The cost varies by organisation size and complexity, but budget £10,000-30,000 for initial certification for a mid-size enterprise, plus annual surveillance costs of approximately £5,000-15,000.

Certification Timeline

Most organisations need 6-12 months to prepare for ISO 42001 certification, depending on the maturity of their existing governance systems. Organisations with ISO 27001 already in place can typically achieve certification faster (4-8 months) because the management system infrastructure already exists. Start with a gap analysis against ISO 42001 requirements to assess your starting position.

Part 11: Governing Agentic AI

Current governance frameworks, including ISO 42001, the NIST AI RMF, and the EU AI Act, were designed primarily for AI systems that assist human decision-making. They are not yet fully adapted for agentic AI: systems that plan, execute, and make decisions autonomously across multiple steps. Singapore's January 2026 framework is the only governance document that directly addresses autonomous agents. For enterprise organisations deploying or planning to deploy AI agents (as covered in Documents 1 and 2 of this series), governance must be extended.

What Makes Agents Different

Cascading failures. A traditional AI tool produces one output for one input. An agent chains multiple steps together. An error in step 2 propagates through steps 3, 4, and 5, potentially causing compounding damage before a human notices. Governance must include circuit breakers and escalation triggers at defined checkpoints, not just at the final output.

Scope creep. Agents can interpret instructions more broadly than intended, especially when given access to tools (file systems, APIs, databases). An agent asked to "clean up the customer database" might interpret this more aggressively than the operator intended. Governance must include strict scope definitions and permission boundaries for every agent.

Attribution gaps. When a multi-agent system produces an output, it can be unclear which agent made which decision. For audit and accountability purposes, every agent action must be logged with the agent identity, the input it received, the decision it made, and the output it produced. The audit trail requirements from Documents 1 and 2 are governance requirements, not just technical best practices.

Autonomous tool use. Agents that can execute code, send emails, modify databases, or interact with external APIs require a fundamentally different risk profile from agents that only generate text. The governance framework must classify agents by their capabilities and apply controls accordingly.

Agent Governance Controls

```
AGENTIC AI GOVERNANCE EXTENSIONS

For each AI agent deployed:

1. SCOPE AND PERMISSIONS
   [ ] Written scope document defining exactly
       what the agent is permitted to do
   [ ] Explicit list of tools/APIs the agent
       can access
   [ ] Explicit list of actions the agent
       CANNOT take (deny list)
   [ ] Maximum autonomy level defined
   [ ] Budget/resource limits set

2. CHECKPOINTS AND CIRCUIT BREAKERS
   [ ] Defined checkpoints where the agent
       pauses for human review
   [ ] Automatic circuit breaker triggers
       (error rate, cost threshold, scope
       deviation, unexpected tool calls)
   [ ] Escalation procedure when circuit
       breaker fires
   [ ] Kill switch available to authorised
       operators

3. AUDIT TRAIL
   [ ] Every agent action logged with
       timestamp, agent ID, input, output
   [ ] Decision reasoning captured where
       the model supports it
   [ ] Tool calls logged with parameters
       and results
   [ ] Logs retained for minimum period
       (align with data retention policy)

4. TESTING AND DEPLOYMENT
   [ ] Agent tested in sandboxed environment
       before production
   [ ] Adversarial testing completed (what
       happens with malicious or ambiguous
       inputs?)
   [ ] Failure mode analysis documented
   [ ] Rollback procedure defined
   [ ] Monitoring dashboards operational

5. REVIEW AND GOVERNANCE
   [ ] Registered in AI register with agent-
       specific fields completed
```

```
    [ ] Risk assessment includes agent-specific
        risks (cascading failure, scope creep)
    [ ] Review cycle: monthly for first 3
        months, then quarterly
    [ ] Incident response plan includes agent-
        specific scenarios
```

Part 12: Regulatory Horizon Scanning

AI regulation is evolving rapidly. Enterprise organisations need a structured approach to monitoring regulatory developments and assessing their impact. This is not a one-time exercise, it is an ongoing function that feeds into the AIMS review cycle.

What Is Coming

EU AI Act high-risk enforcement (postponed to December 2027). Still the most significant medium-term regulatory milestone, though the May 2026 Digital Omnibus agreement moved the Annex III high-risk deadline from August 2026 to 2 December 2027 (the Article 50 transparency duties still begin August 2026). Conformity assessment is required for high-risk AI systems. Organisations deploying AI in employment, credit, insurance, education, healthcare, or critical infrastructure for EU markets must be ready. Start preparing now if you have not already.

ICO statutory code of practice on AI/ADM. Signalled in the ICO's AI and Biometrics Strategy (June 2025). This will provide specific UK guidance on automated decision-making requirements, transparency obligations, and the interaction between AI and UK GDPR. Timeline uncertain but expected 2026-2027.

FCA developments. The FCA has not announced AI-specific rules, but relies on Consumer Duty and SM&CR to regulate AI use in financial services. 75% of financial services firms already use AI. The FCA's AI Live Testing programme and Supercharged Sandbox (with NVIDIA) signal increasing regulatory attention. Expect clearer guidance in 2026-2027.

ETSI AI security standard. Published May 2025 with 13 principles across 5 lifecycle stages. Increasingly referenced in procurement requirements and industry codes of practice.

UK adequacy review. EU adequacy was renewed in December 2025 for up to 6 years with possible extension. However, divergence in AI regulation between the UK and EU could become a factor in future adequacy assessments. Monitor the EU-UK relationship on AI governance.

Agentic AI governance. Singapore's January 2026 framework is the first governance document specifically addressing autonomous agents. Expect UK, EU, and NIST frameworks to incorporate agent-specific guidance in 2026-2027. Early adoption of agent governance controls (Part 11) positions you ahead of regulatory requirements.

Building a Regulatory Monitoring Function

Assign a named individual (typically the AI Governance Lead or a member of the legal/compliance team) responsibility for regulatory horizon scanning. This does not need to be a full-time role, but it does need dedicated time (2-4 hours per month minimum). The function should monitor ICO publications and consultations, FCA and relevant sector regulator announcements, EU AI Office publications, NIST framework updates, ISO/IEC updates (particularly around 42001 and 42006), AISI publications and guidance, and industry body and professional association updates.

Feed findings into the quarterly AI Governance Committee meeting and flag urgent developments for immediate escalation. Maintain a regulatory tracker that records each development, its potential impact on your organisation, the required response, the owner, and the deadline.

Part 13: Quick Reference

Enterprise AI Governance Readiness Checklist

```
ENTERPRISE AI GOVERNANCE READINESS

FOUNDATIONS
[ ] AI Governance Policy approved by board
[ ] AIMS scope defined and documented
[ ] Executive sponsor named
[ ] AI Governance Lead / CAIO appointed
[ ] Governance committee established
[ ] Budget allocated for governance programme

AI REGISTER
[ ] AI discovery exercise completed
[ ] All known AI systems registered
[ ] Risk classification assigned to each system
[ ] Business and technical owners assigned
[ ] Vendor AI features identified and registered

RISK MANAGEMENT
[ ] Risk assessment completed for all Tier 2+
[ ] DPIAs completed for all Tier 3 systems
[ ] Bias testing programme in place for Tier 3
[ ] Incident response plan covers AI scenarios
[ ] Insurance reviewed for AI exclusions

POLICIES
[ ] AI Governance Policy
[ ] AI Acceptable Use Policy
[ ] AI Procurement Policy
[ ] AI Data Governance Policy
[ ] AI Model Risk Management Policy
[ ] AI Incident Response Policy
[ ] Third-Party AI Policy
[ ] Employee AI Rights Policy
[ ] AI Ethics / Responsible AI Policy

CONTROLS
[ ] Human oversight levels defined per system
[ ] Audit trail operational for Tier 2+ systems
[ ] Data sovereignty verification complete
[ ] Approved tools registry operational
[ ] Prompt sanitisation for Tier 3 systems
[ ] Network egress monitoring (where applicable)
```

```
PEOPLE
[ ] Training programme operational
[ ] BU AI Leads appointed
[ ] Team champions identified
[ ] Competency assessment completed
[ ] Awareness programme for all staff

AUDIT AND ASSURANCE
[ ] Internal audit programme defined
[ ] First internal audit completed
[ ] Management review conducted
[ ] Corrective actions tracked
[ ] ISO 42001 gap analysis completed (if pursuing)

REGULATORY
[ ] Applicable regulations mapped
[ ] Regulatory monitoring function assigned
[ ] EU AI Act impact assessed (if applicable)
[ ] Sector-specific requirements mapped
[ ] Regulatory tracker maintained
```

Implementation Timeline

```
ENTERPRISE AI GOVERNANCE IMPLEMENTATION

PHASE 1: FOUNDATION (Months 1-3)
  - Executive sponsor and CAIO appointed
  - Governance policy drafted and approved
  - AIMS scope defined
  - AI discovery exercise completed
  - Initial AI register populated
  - Gap analysis against ISO 42001

PHASE 2: BUILD (Months 4-8)
  - Full policy stack drafted and approved
  - Risk assessments for all Tier 2+ systems
  - DPIAs for Tier 3 systems
  - Governance committee operational
  - BU AI Leads appointed and trained
  - Bias testing programme designed
  - Internal audit programme planned
  - Vendor AI audit commenced

PHASE 3: OPERATE (Months 9-12)
  - All policies operational
  - First internal audit cycle complete
  - Management review conducted
  - Training programme live
```

```
  - Incident response tested
  - Regulatory tracker operational
  - ISO 42001 Stage 1 audit (if pursuing)

PHASE 4: CERTIFY AND MATURE (Months 12-18)
  - ISO 42001 Stage 2 audit (if pursuing)
  - Continuous improvement cycle established
  - Agent governance extensions deployed
  - Cross-framework crosswalk documented
  - Annual board AI governance report
```

Part 14: References and Further Reading

Dendro Logic AI Adoption Playbook Series

Document 1: AI Agents in Development Teams. Technical implementation: CLAUDE.md, guardrails, TDD, CI/CD integration, audit trails for agent systems.

Document 2: Designing Multi-Agent AI Systems. Architecture, frameworks, MCP, model selection, production deployment, observability, multi-agent governance controls.

Document 3: AI in the General Workforce. Non-technical adoption, shadow AI discovery, training, SOPs, department-level guidance, UK Government trial findings.

Document 4: Data Sovereignty and AI Security. UK regulatory landscape, data classification, provider mapping, technical verification, insurance, incident response, vendor management. Foundation document for enterprise governance.

Document 5: AI Strategy for Business Leaders. Business case, financial modelling, risk communication, measurement, 12-month roadmap, minimum viable governance, leadership and culture.

Standards and Frameworks

ISO/IEC 42001:2023, Artificial Intelligence Management System (AIMS). First international certifiable AI management system standard. Plan-Do-Check-Act structure. Annex A controls for AI-specific governance. Available at: iso.org

BS ISO/IEC 42006:2025, Requirements for bodies providing audit and certification of AI management systems. Specifies auditor qualifications for ISO 42001 certification. Available at: bsigroup.com

NIST AI RMF 1.0, AI Risk Management Framework. Govern, Map, Measure, Manage. Voluntary framework referenced by US federal agencies. Generative AI Profile (AI 600-1) and Cyber AI Profile (IR 8596). Available at: nist.gov/ai

EU AI Act, Regulation (EU) 2024/1689. Risk-based framework. Prohibited practices (Feb 2025), GPAI obligations (Aug 2025), high-risk enforcement (postponed to Dec 2027). Available at: artificialintelligenceact.eu

ETSI, AI Security Standard. May 2025. 13 principles, 5 lifecycle stages. Increasingly referenced in UK procurement. Available at: etsi.org

NCSC/CISA, Guidelines for Secure AI System Development. 4 pillars: secure design, development, deployment, operation. Available at: ncsc.gov.uk

UK Regulatory Guidance

ICO, AI and Biometrics Strategy. June 2025. Forthcoming statutory code of practice on AI/ADM. Available at: ico.org.uk

DUAA, Data Use and Access Act. Royal Assent June 2025. Relaxed ADM restrictions, transfer risk assessment requirement, "not materially lower" transfer standard. Available at: legislation.gov.uk

FCA, AI in Financial Services. No AI-specific rules planned. Consumer Duty + SM&CR. AI Live Testing. Supercharged Sandbox with NVIDIA. 75% of financial services firms use AI. Available at: fca.org.uk

AISI, AI Safety Institute. Technical evaluation, testing guidance, responsible AI deployment. Available at: aisi.gov.uk

Enterprise AI Governance Research

PwC, 29th Global CEO Survey. January 2026. 4,454 CEOs, 95 countries. 56% report zero financial benefit from AI. 12% vanguard with strong AI foundations. Available at: pwc.com/gx/en/ceo-survey

Gartner, AI spending projected at $2.52 trillion in 2026. Available at: gartner.com

Zylo, 2026 SaaS Management Index. AI-native spending doubled year-on-year. ChatGPT #1 most-expensed app. Shadow AI expanding spend and risk. Available at: zylo.com

EC Council / GAICC, "Global AI Governance Comparison 2026: EU AI Act vs NIST AI RMF vs ISO/IEC 42001." Crosswalk analysis, implementation sequence, agentic AI gaps. March 2026. Available at: gaicc.org

Singapore IMDA, Agentic AI Governance Framework. January 2026. First governance document addressing autonomous agents directly. Available at: imda.gov.sg

HiComply, "ISO 42001 vs NIST AI RMF: How to Choose the Right Framework." November 2025. Practical comparison for UK organisations. Available at: hicomply.com

FireTail, "AI Governance Frameworks: Best Practices for 2026." NIST, ISO 42001, OWASP LLM Top 10 implementation guidance. November 2025. Available at: firetail.ai

This playbook is a living document.

Update it as your governance matures, your AI portfolio evolves, and regulation develops.

Part 6 of the Dendro Logic AI Adoption Series.

https://dendro-logic.com