

AI Body of Knowledge

AI-BOK v1.3 - September 2026, English Edition

The reference framework for AI governance, lifecycle management and organisational design
13 knowledge areas · 22 roles · 22 failure modes · 9 jurisdictions · open ArchiMate model



Jan Willem van Veen · Principal Enterprise Architect

ISO 42001 · NIST AI RMF · EU AI Act · DAMA-DMBOK · TOGAF 10 · 2026

Table of Contents

What the AI-BOK is	5
Why the AI-BOK	5
What v1.3 changes	6
Where to start	6
How to contribute	6
Foreword	6
Module 1: Knowledge Areas	8
Introduction	8
Overview of the thirteen knowledge areas	9
KA1: AI Governance	10
KA2: AI Strategy & Portfolio Management	18
KA3: AI Architecture	25
KA4: AI Lifecycle Management	34
KA5: Data & Semantics for AI	40
KA6: Model Management	49
KA7: AI Interaction & User Experience	59
KA8: AI Operations (AIOps/MLOps/LLMOps)	66
KA9: AI Risk Management & Safety	73
KA10: AI Compliance & Audit	82
KA11: AI Ethics & Responsible AI	89
KA12: AI Knowledge & Context Management	95
KA13: AI Literacy & Capability Development	102
Module 2: Role Model	108

2.0 Introduction	108
2.1 Role Overview	109
2.2 Role Descriptions	109
2.3 RACI Model	135
2.4 Separation of Duties	140
2.5 Governance Structure	142
2.6 Scalability of the Role Model	147
2.7 Mapping to EU AI Act Roles	149
2.8 Summary and Coherence	149
Module 3: Integration Module	151
3.0 Introduction	151
3.1 AI Lifecycle Model	151
3.2 Meta-model	165
3.3 Governance Integration Model	172
3.4 Maturity Model	191
3.5 Implementation Roadmap	204
3.6 Cognition Plane as integration layer	213
3.7 AI Quality Framework (cross-cutting)	235
3.8 Role-based reading paths and Minimum Viable Governance (MVG)	240
3.9 Closing statement: the AI-BOK as a coherent whole	246
Appendices	248
Appendix A: Worked examples per knowledge area	249
Appendix B: End-to-end agentic case (CHRA-v2)	254
Appendix C: Brownfield integration guide	262
Appendix D: SME playbook	268

Appendix E: Templates and Checklists	273
Appendix F: Glossary	291
Appendix G: Sources	299

AI-BOK Executive Summary - v1.3 (September 2026)

A one-page summary for board members, executives and senior stakeholders: what the AI-BOK is, why it exists, what changes in this version, where to start and how to contribute.

What the AI-BOK is

The AI Body of Knowledge (AI-BOK) is an open reference framework for AI management: how an organisation governs, builds, operates, secures and decommissions AI systems responsibly. It bundles a knowledge structure, a role model, an integration model, a maturity model, an open ArchiMate reference and a regulatory traceability layer into one coherent whole, comparable to what DAMA-DMBOK is for data management and TOGAF for enterprise architecture. The framework is openly licensed, appears in versioned releases and is usable by organisations from 50 FTE to a global enterprise.

Component	Count	What it is
Knowledge areas	13	KA1 Governance (central) plus twelve areas for strategy, architecture, lifecycle, data, models, interaction, operations, risk, compliance, ethics, knowledge and literacy
Roles	22	Strategic, tactical, operational and cross-functional, with a RACI of 22 roles by 13 knowledge areas
Lifecycle phases	8	Ideation & Exploration through Evolution & Decommissioning, with four feedback loops
Metamodel entities	10	AI System, AI Model, Dataset, Prompt, Interaction, Decision, Knowledge Source, Agent, Governance Framework, Risk Profile
Architecture planes	4	Data plane, control plane, cognition plane and management plane
Failure modes	22	FM01-FM16 agentic and FM17-FM22 cyber-physical, each with mechanism, signature, detection and mitigation
Maturity levels	5	L1 Initial to L5 Optimised, with an evidence rubric per knowledge area

Why the AI-BOK

No existing instrument delivers the integrated operating model an organisation needs. ISO/IEC 42001 describes management processes without the substance they have to carry. NIST AI RMF names risks without owners. The EU AI Act imposes duties without saying how an organisation assigns them. The AI-BOK connects the three: every knowledge area has an accountable role, every duty a place in the lifecycle and every risk a control with evidence.

What v1.3 changes

v1.3 is a consistency and adoption release. No knowledge area is added; the framework is made internally consistent and gains two roles that were missing in practice.

- **Two adoption roles.** The AI Ambassador (Role 21) carries adoption inside a team or department; the AI Prompting & Agents Specialist (Role 22) manages prompts, agent configurations and the prompt library.
- **RACI corrected.** In earlier versions seven of the thirteen knowledge areas had no accountable role. Every knowledge area now has exactly one A and at least one R; a validation script enforces this.
- **Four architecture planes.** The management plane sits alongside the data plane, control plane and cognition plane, in line with the peer-reviewed paper and the open ArchiMate model.
- **Consistency tooling.** Book, ArchiMate model and SemanticXL pack read the same canonical source; validation scripts check roles, levels, planes and NL/EN parity on every build.

Where to start

Start with the **MVG-7** (Minimum Viable Governance, Module 3 §3.8): a governance charter, an AI inventory, a risk classification per system, lifecycle gates, AI literacy with evidence, compliance traceability and an incident pipeline. That is a defensible baseline for any organisation, even with four combined roles. Then pick the **reading path for your role** (Module 3 §3.8) and use the worked examples in the appendices as the shape of evidence you are working towards.

How to contribute

The AI-BOK is a living document. Proposals and corrections go through the review protocol and the changelog on <https://www.ai-bok.nl>; releases so far: v1.0 (March 2026), v1.1 (June 2026), v1.2 (June 2026) and v1.3 (September 2026). The regulatory traceability matrix is reviewed every six months, the maturity model annually.

Feedback at info@ai-bok.nl. The open ArchiMate model and the full document set are freely downloadable at <https://www.ai-bok.nl/architecture>.

Foreword

I have worked as an architect for more than twenty-five years. First on classical information systems in the Dutch public sector, then on data platforms and knowledge technology, and over recent years increasingly on AI. What began as a specialisation has become something every organisation now has to deal with.

In my work I keep running into the same pattern. An organisation wants to do something with AI. Someone goes looking for a framework. ISO 42001 covers the management system, NIST writes about risk, the AI Act sets the legal boundaries, but together they do not add up to a workable whole. An AI architect who wants to understand how an agent relates to the cognition plane has to assemble the picture on their own. A risk manager looking for the expected failure modes runs into fragmented literature. That puzzling work is time-consuming and routinely goes wrong.

Hence this book.

The AI Body of Knowledge brings together what existing frameworks do not provide themselves: knowledge areas, roles, a lifecycle, a metamodel, a maturity model and a reference architecture, in a coherent form. It is not an academic work, nor a catalogue of best practices. It is a working document for people who have to make progress on Monday morning.

My field has always been meaning: semantic web technology, knowledge graphs, linked data. From that perspective my conviction is that trustworthy AI does not come from better models alone. It comes from better meaning: precise concepts, trusted sources, explicit frameworks within which reasoning takes place. The cognition plane captures that as an architectural concept.

In this release (September 2026) the Dutch and English editions stand on equal footing; the parity restored in the previous release now extends to all modules. The role model counts 22 roles: alongside knowledge area 13 (AI Literacy) and role 20 (Literacy Officer), role 21 (AI Ambassador) and role 22 (AI Prompting & Agents Specialist) are now part of the canonical modules. The reference architecture works with four planes: data plane, control plane, cognition plane and management plane. The RACI matrix has been corrected, among other things for the separation between the AI Auditor and the AI Compliance Officer. Consistency between the modules, the role model and the SemanticXL pack is maintained with automated validation. The failure-mode catalogue counts 22 patterns. Regulatory traceability covers nine jurisdictions. The open reference architecture and all accompanying ArchiMate files are available at <https://www.ai-bok.nl/architecture>.

I hope you find it useful.

Jan Willem van Veen · Principal Enterprise Architect · September 2026

Module 1: Knowledge Areas

Introduction

The AI Body of Knowledge (AI-BOK) provides a structured reference framework for organisations that wish to deploy artificial intelligence in a responsible, architecturally grounded and governable manner. Module 1 defines thirteen knowledge areas that together cover the full domain of AI management, from strategy to operations, from ethics to architecture, and from compliance to AI literacy.

The structure has been deliberately modelled after the knowledge areas of DAMA-DMBOK, in which Data Governance serves as the central connecting area. In the AI-BOK, **AI Governance** fulfils that role: the overarching knowledge area that provides direction for all other areas and ensures that AI initiatives are aligned with organisational objectives, legislation and regulation, and societal expectations.

A distinguishing concept within this AI-BOK is the **cognition plane**: the architecture layer in which AI systems reason, interpret and prepare decisions. This plane complements the classical architecture layers (business, data, application, technology) and makes explicitly visible where autonomous intelligence operates within the enterprise architecture.

Why an AI-BOK?

A legitimate objection is that the fragmentation in AI governance ought to be solvable without a new reference framework: simply implement the existing standards better, assign the responsibilities more tightly, get the architecture practice in order. Each of those remedies, however, presupposes what no existing instrument provides: an agreed content structure that says *which* knowledge an AI function must maintain, *who* is responsible for it and *how* the instruments relate to one another. ISO/IEC 42001 requires processes but does not describe their content; the NIST AI RMF names risks but no owners; the EU AI Act imposes obligations but no operating model. Organisations that combine these instruments must construct the integration themselves. In practice that integration is built again and again, and differently every time. The AI-BOK codifies that integration once: the knowledge areas provide the content structure, the role model the responsibilities, and the integration matrices of Module 3 the explicit connections with ISO 42001, NIST AI RMF, EU AI Act, TOGAF and DAMA-DMBOK.

Note. Knowledge area 13 - **AI Literacy & Capability Development** - anchors EU AI Act Article 4 (in force February 2025) and gives AI literacy a structural home alongside the other knowledge areas. The cross-references from KA1, KA3, KA9, KA10, KA11 and KA12 are elaborated in the accompanying addenda (regulatory traceability matrix, maturity evidence rubric, agentic & cyber-physical failure-mode catalogue, reading paths, MVG profile, VNG/BZK/GEMMA cross-walks and the operational/cyber-physical addendum).

Status and further development

The AI-BOK is a **candidate reference framework**. The authority of a body of knowledge does not arise at first publication but through multi-year, community-driven curation, as DAMA-DMBOK has built up over multiple editions. This edition has been constructed principally by one author and tested in three evaluation iterations (practitioner panel, independent academic expert panel, pilot applications); that is a starting point, not a terminus. The validation roadmap runs via collaboration with practitioner communities (including DANW and the FGEA Community of Practice), broadening of the academic review panel and longitudinal research into effectiveness in organisations. Readers who

wish to contribute (corrections, practical experiences, change proposals) can do so via the channels on www.ai-bok.nl; the change log and the published review protocols can also be found there. Feedback is processed per release and accounted for traceably.

Overview of the thirteen knowledge areas

The diagram below shows the structure of the AI-BOK. AI Governance forms the central node; the remaining twelve knowledge areas are grouped into three rings: Strategy & Direction, Realisation & Operations, and Assurance & Accountability.

(See Figure 1: AI-BOK Framework - 13 knowledge areas)




Figure 1: AI-BOK Framework - 13 Knowledge Areas

Reading guide: AI Governance is positioned centrally and is connected to all other knowledge areas. The three rings indicate the primary orientation:

- **Strategy & Direction** (2, 3, 4): Determining what, why and how AI is deployed
- **Realisation & Operations** (5, 6, 7, 8): Building, training, deploying and operating AI systems
- **Assurance & Accountability** (9, 10, 11, 12, 13): Ensuring that AI functions safely, compliantly, ethically and reliably, with KA13 establishing the human capability that makes the other four areas operable

KA1: AI Governance

1.1 Definition

AI Governance encompasses the totality of governance structures, policies, decision-making processes and accountability mechanisms that ensure artificial intelligence within an organisation is developed, deployed and decommissioned in a purposeful, responsible and controlled manner. It is the central and overarching knowledge area of the AI-BOK, comparable to the role that Data Governance plays in DAMA-DMBOK. AI Governance explicitly also covers the cognition plane: the architecture layer in which AI systems autonomously reason, interpret and act.

1.2 Purpose and principles

The purpose of AI Governance is to create the conditions under which AI systems deliver maximum value within acceptable risk boundaries. The following principles apply:

1. **Responsibility is explicitly assigned** - For every AI system, an owner is designated who is responsible for its functioning, impact and policy compliance. Responsibility must not remain implicit or become distributed across functions without a clear mandate.
2. **Governance is proportional to risk** - The degree of governance intensity is proportional to the risk classification of the AI system. A chatbot for internal FAQs requires lighter governance than an autonomous decision-making system for permit issuance. The EU AI Act risk classification provides a useful reference framework for this purpose.
3. **Transparency about autonomy** - The degree to which an AI system operates autonomously is documented and visible. For AI agents, this means: explicit documentation of authority levels, context boundaries and escalation thresholds.
4. **Governance is embedded, not imposed** - Governance mechanisms are part of the AI architecture and the development process, not an after-the-fact control layer. Policies are preferably recorded in machine-readable form so that agents can consult them at runtime.
5. **Human control as a design principle** - For decisions with significant impact, human-in-the-loop or human-on-the-loop is built in. The escalation level is specified per decision type in the governance architecture.
6. **Continuous improvement based on metrics** - Governance is not static. Based on monitoring, incidents and maturity assessments, governance structures are periodically adjusted.

1.3 Scope

Within scope:

- Establishment of AI policy, principles and guidelines at the organisational level
- Set-up of governance bodies (AI Board, AI Review Committee, AI Ethics Committee, optionally an AI Center of Excellence)
- Definition and assignment of roles and responsibilities in accordance with RACI
- Decision-making frameworks for AI initiatives, including go/no-go criteria and dispensation procedures
- Governance of the cognition plane: governing the layer in which AI agents reason and act
- Compliance with internal standards and external regulation (EU AI Act, ISO 42001)
- Integration with enterprise governance, IT governance and data governance

- Governance escalation and incident mechanisms

Outside scope:

- Operational execution of risk management (see KA9)
- Technical implementation of compliance controls (see KA10)
- Formulation of the AI strategy itself (see KA2)
- Architecture design and management (see KA3)

1.4 Core activities

1. **Establishing the AI governance framework** - Designing and formalising the governance structure including mandates, authorities and decision-making lines. The framework positions AI Governance relative to existing corporate governance, IT governance and data governance. ISO 42001 provides a useful basis for the structure of the management system.
2. **Developing AI policy and principles** - Formulating organisation-wide policies that set the boundaries for responsible AI deployment. This includes policy for data usage by AI, transparency, human control, non-discrimination and sustainability. The OECD AI Principles and the requirements of the EU AI Act serve as a starting point.
3. **Establishing AI governance bodies** - Setting up the required bodies: an AI Board for strategic decision-making, an AI Review Committee for tactical review at gates, an AI Ethics Committee for ethical review and, as an organisational option, an AI Center of Excellence for knowledge sharing and standardisation. The composition and mandates are aligned with the size and maturity level of the organisation.
4. **Classifying AI systems by risk** - Each AI system is classified into a risk category based on the EU AI Act (unacceptable, high, limited, minimal) and additional organisation-specific criteria. The classification determines the governance intensity and the required measures.
5. **Defining roles and responsibilities** - Assigning ownership per AI system, model and dataset. The RACI matrix specifies who is Responsible, Accountable, Consulted and Informed. This includes both human roles and defined responsibilities for AI agents themselves (mandates, boundaries).
6. **Governing agent autonomy** - Establishing authority levels for AI agents, from full human control to delegated autonomy. Per agent, the following is specified: which domain, which sources, which actions, which escalation thresholds. This is a specific governance challenge of the cognition plane.
7. **Managing agent identities** - AI agents are non-human identities (NHIs) that require the same identity governance as human users. This encompasses five capabilities: (1) *discovery and classification*: identify all active agents in the landscape and classify them by risk and authority level, (2) *ownership*: assign a human owner to each agent to prevent orphaned agents and unmanaged risk, (3) *lifecycle governance*: govern how agent identities are created, modified and decommissioned to prevent permission creep, (4) *observability*: link agent activity to access rights to make hidden risks visible, and (5) *auditability*: ensure traceability and audit-ready evidence. As the number of agents in production grows, identity governance for NHIs becomes a prerequisite for keeping the AI landscape manageable.
8. **Monitoring AI maturity** - Periodically assessing organisation-wide AI maturity using the AI-BOK maturity model. The assessment covers all thirteen knowledge areas and produces an improvement plan.
9. **Coordinating cross-functional collaboration** - Organising alignment between business, IT, data, legal, compliance, ethics and architecture. AI governance is by definition multidisciplinary; no single department can govern it alone.
10. **Handling governance incidents** - Establishing an escalation and response mechanism for AI-related incidents: undesired output, bias detection, security breaches, compliance violations. The mechanism includes triage, investigation, mitigation and reporting.

11. **Reporting and accountability** - Preparing periodic governance reports for the board of directors, regulators and relevant stakeholders. Reports cover governance coverage, incidents, maturity development and compliance status.

1.5 Roles and responsibilities

Role	Responsibility	Reference
Chief AI Officer (Role 1)	Ultimately responsible for the AI governance framework and the AI strategy; chairs the AI Board	Module 2, Role 1
AI Board (body)	Strategic decision-making on AI policy, risk acceptance, portfolio prioritisation and the organisation-wide maximum for agent autonomy	Module 2, §2.5
AI Governance Lead (Role 7)	Designs and maintains the governance framework and policy; secretary of the AI Board and chair of the AI Review Committee	Module 2, Role 7
AI Product Owner (Role 4)	System owner of a specific AI system; responsible for classification, compliance and performance; as deployer the owner of DPIA and FRIA	Module 2, Role 4
AI Ethics Advisor (Role 3)	Safeguards ethical principles and advises on dilemmas; chairs the AI Ethics Committee	Module 2, Role 3
AI Agent Governor (Role 18)	Sets the autonomy level per agent within the maximum set by the AI Board and maintains the agent authority register	Module 2, Role 18
AI Compliance Officer (Role 19)	Coordinates and executes compliance assessments, including DPIA and FRIA; ultimately responsible for KA10	Module 2, Role 19
Data Protection Officer (DPO; outside the role model)	Advises on GDPR compliance of AI systems that process personal data (Art. 39 GDPR); advises and does not approve	Outside the role model

1.6 Techniques and tools

AI Governance Framework Canvas

Visual tool for designing and communicating the governance structure

RACI matrix

Documentation of roles and responsibilities per AI system

AI risk classification model

Structured method for categorising AI systems by risk profile (in accordance with the EU AI Act) (See Figure 2: EU AI Act risk classification)




Figure 2: EU AI Act Risk Classification

Maturity model (AI-BOK CMM)

Assessment framework at five levels for all thirteen knowledge areas

Policy-as-code tooling

Machine-readable documentation of governance rules (e.g., Open Policy Agent, Rego) so that AI agents can consult policy rules at runtime

AI governance dashboard

Real-time overview of governance coverage, incidents and compliance status

Governance tooling / enterprise architecture tool

Documentation of governance structures, mandates and agent authorities as architecture artefacts

1.7 Inputs and outputs

Inputs:

- Enterprise strategy and governance frameworks
- Legislation and regulation (EU AI Act, GDPR, sector-specific regulation)
- AI maturity assessment (baseline measurement)
- Risk inventory from KA9 (AI Risk Management)
- Architecture principles from KA3 (AI Architecture)
- Incident reports and audit findings from KA10

Outputs:

- AI Governance Charter
- AI policy framework and AI principles
- RACI matrix for AI roles
- AI risk classification register
- AI governance dashboard and periodic reports
- Mandates and authority frameworks for AI agents
- Dispensation and escalation procedures

1.8 Implementation guidelines

1. **Begin with an AI Governance Charter** - Draft a brief, authoritative document that defines the mandate, scope and governance bodies. Have it approved by the board of directors. Without a formal mandate, AI Governance remains an advisory function without executive authority.
2. **Align with existing governance structures** - Do not build AI Governance as an isolated silo. Integrate with IT governance (COBIT), data governance (DAMA-DMBOK) and enterprise architecture (TOGAF). Reuse existing bodies where possible and add AI-specific expertise.
3. **Make governance proportional and layered** - Define three to four governance levels linked to risk classification. Avoid one-size-fits-all; heavy governance for low-risk applications inhibits innovation, while overly light governance for high-risk systems creates liability risks.
4. **Invest in policy-as-code** - Document governance rules in documents and in machine-readable form. In an architecture with AI agents, agents must be able to verify their authorities and limitations at runtime without human intervention.

1.9 Governance aspects

AI Governance is itself the governance knowledge area and as such must also organise its own governance:

Periodic review

The AI governance framework is evaluated annually for effectiveness and relevance, or more frequently in the event of significant changes in

legislation and regulation

Independent assessment

The internal audit function or an external party periodically assesses the functioning of the governance structure (in conjunction with KA10)

Escalation mechanism

In the event of conflicts between governance requirements and operational interests, the AI Board escalates to the executive board

Change management

Amendments to AI policy and principles follow a formal change process with versioning and impact analysis

1.10 Metrics and KPIs

#	Metric	Description	Target value	Measurement frequency
1	Governance coverage	% of AI systems under formal governance oversight	>= 95%	Quarterly
2	Policy compliance rate	% of AI initiatives compliant with established AI policy	>= 90%	Quarterly
3	Risk classification completeness	% of AI systems with established risk classification	100%	Monthly
4	Time-to-decision	Median turnaround time of governance decision-making	<= 10 working days	Quarterly
5	Maturity score	Score on the AI-BOK maturity model (1-5)	>= 3 (year 1), >= 4 (year 3)	Semi-annually
6	Incident response time	Median time to mitigation of governance incidents	<= 48 hours	Monthly
7	Agent mandate coverage	% of operational AI agents with documented authority framework	100%	Quarterly
8	Agent identity coverage	% of AI agents with registered NHI, human owner and current access rights	100%	Quarterly
9	Orphaned agent ratio	% of agents without an active human owner	0%	Monthly
10	Cross-governance integration	Number of formal linkages with IT/data/enterprise governance	>= 3 active linkages	Annually

1.11 Relationship with other knowledge areas

KA2 (AI Strategy)

Governance sets the boundaries within which strategy is formulated; strategy provides input for governance priorities

KA3 (AI Architecture)

Governance defines architecture principles; architecture provides the structure to implement governance in a machine-readable manner

KA4 (AI Lifecycle Management)

Governance determines the gate criteria in each phase of the lifecycle

KA5 (Data & Semantics)

Governance sets requirements for data quality and semantic precision; terminology frameworks form the foundation for unambiguous governance definitions

KA9 (AI Risk Management)

Governance establishes the risk tolerance; risk management provides input for governance decisions

KA10 (AI Compliance & Audit)

Governance mandates compliance requirements; audit assesses the functioning of governance

KA11 (AI Ethics) Governance anchors ethical principles; ethics advises on dilemmas that governance must resolve	KA12 (Knowledge & Context Management) Governance sets requirements for the reliability of knowledge sources that agents use as context
---	--

1.12 Relevant standards

Standard	Relevance for this knowledge area
ISO/IEC 42001:2023	AI Management System - provides the structure for the governance framework
ISO/IEC 38507:2022	Governance implications of AI - guidelines for governance responsibility
EU AI Act (2024)	Legal obligations for providers and deployers, risk classification
NIST AI RMF 1.0	Govern function - best practices for AI governance
OECD AI Principles	Five principles for responsible AI that serve as a policy basis
ISO/IEC 23894:2023	AI Risk Management - supplementary to the governance framework
COBIT 2019	IT governance framework with which AI Governance must integrate

Business rules and decision standards:

Standard	Relevance for this knowledge area
DMN (Decision Model and Notation)	OMG standard for modelling decisions and decision rules. Relevant to AI governance because it provides the formal specification of the decision rules that AI systems apply or support. DMN makes it possible to document human and AI decisions in the same notation, enhancing transparency and auditability.
SBVR (Semantics of Business Vocabulary and Business Rules)	OMG standard for formalising business vocabulary and business rules in structured natural language. Relevant for translating governance policy into machine-readable rules that can be applied by AI systems.
RuleML / SWRL	Standards for expressing rules on top of RDF/OWL ontologies. Enable the documentation of governance rules, compliance requirements and business logic as formal rules that can be consulted and applied by AI agents.

1.13 Shadow AI: invisible risk, visible approach

Definition. Shadow AI is the use of AI tools and systems by employees without the knowledge, approval or governance of the organisation. Where shadow IT traditionally concerns unauthorised hardware and software, shadow AI specifically concerns generative and analytical AI services that employees deploy on their own initiative: ChatGPT for writing emails and notes, GitHub Copilot for code, Midjourney for presentation visuals, or self-built GPTs for data analysis. The distinguishing characteristic is that the organisation has no visibility into what data these systems process, what output they generate and what dependencies they create.

Scale of the problem. A significant proportion of AI usage in organisations falls outside formal governance; almost every internal inventory reveals more tools than the IT function is aware of. This is not a marginal phenomenon: in virtually every department, from marketing to finance, from HR to legal, employees are experimenting with AI tools to work more productively. The motivation is understandable (working faster, better output), but the unmanaged nature

makes shadow AI one of the most urgent governance challenges. The problem worsens as AI functionality is embedded in existing tools (Microsoft 365 Copilot, Google Gemini in Workspace, Adobe Firefly), blurring the boundary between "normal usage" and "AI usage".

Risks. Shadow AI introduces five categories of risk that directly affect governance responsibility:

1. **Data leakage:** Employees enter business-sensitive information into external AI models (customer data, financial data, strategic plans). This data may be stored, used for training or accessible to third parties.
2. **Compliance violation:** Processing of personal data by external AI services may violate GDPR (no processing agreement, no DPIA). Under the EU AI Act, uncontrolled use of high-risk AI systems can lead to substantial fines.
3. **Inconsistent output:** Different employees use different AI tools with different prompts, leading to inconsistent quality, contradictory advice and untraceable decisions.
4. **Reputational damage:** AI-generated output containing errors, exhibiting bias or being inappropriate is released externally without review, resulting in reputational damage.
5. **Uncontrollable dependencies:** Business processes become dependent on external AI services without contractual assurance, SLAs or continuity planning. When a service changes or stops, processes come to a halt.

Detection. Effective detection of shadow AI requires a combination of technical and organisational measures:

- **Network monitoring of AI API traffic:** Identify outbound traffic to known AI endpoints (api.openai.com, generativelanguage.googleapis.com, etc.) via firewall logs and proxy analysis.
- **Software inventory:** Scan endpoints for installed AI applications, browser extensions and plugins that provide AI functionality.
- **Employee surveys:** Conduct periodic (at least semi-annual) anonymous surveys in which employees are asked which AI tools they use and for what purpose. Anonymity increases the likelihood of honest responses.
- **Browser extension monitoring:** Monitor browser extensions that inject AI functionality into web applications.
- **API gateway analysis:** Analyse outgoing API calls for patterns indicating AI services (large text payloads, multimodal uploads).

Approach: channel, not prohibit. A ban on AI usage is counterproductive: employees will find workarounds, and the organisation misses innovation. The effective approach is to channel:

1. **Offer safe alternatives:** Provide enterprise versions of popular AI tools (Azure OpenAI, AWS Bedrock) with data processing within organisational boundaries and processing agreements.
2. **Establish an AI usage policy:** Clearly define which categories of data may and may not be entered into AI systems, which use cases are permitted and which require approval.
3. **Create an "approved tools" list:** Publish and maintain a living list of AI tools that have been assessed for privacy, security and compliance. Make the process for having a tool assessed low-barrier (maximum 2 weeks turnaround time).
4. **Make reporting low-barrier:** Create a culture in which employees can openly share which AI tools they use, without fear of sanctions. An "AI usage reporting point" or Slack channel lowers the barrier.
5. **Train and raise awareness:** Offer short, practical training on responsible AI usage, focused on the risks of unmanaged use and the available alternatives.

Governance integration. Shadow AI management integrates with multiple knowledge areas of the AI-BOK:

- The **shadow AI inventory** provides direct input for KA9 (Risk Management): every unmanaged AI system is a risk that must be assessed, classified and mitigated.
- The compliance implications are managed via KA10 (AI Compliance & Audit): every shadow AI system must be tested against GDPR, EU AI Act and sector-specific regulation.

- The technical aspects (API monitoring, enterprise alternatives) fall under KA3 (AI Architecture) and KA8 (AI Operations).
- The cultural and organisational dimension relates to KA13 (AI Literacy & Capability Development): building AI literacy, led by the AI Literacy Officer (Role 20) and carried by the AI Ambassadors (Role 21), reduces the tendency towards unmanaged use.

Use **Template 5 (Shadow AI Inventory)** as a practical tool to map the current shadow AI landscape and as a basis for developing a targeted action plan. The inventory must be repeated periodically (at least semi-annually) because the AI landscape evolves rapidly.

How the book continues

Module 1: twelve more knowledge areas, each with activities, roles, metrics, standards and its relation to DAMA-DMBOK, TOGAF, ISO 42001 and NIST AI RMF.

Module 2: the role model with 22 roles, the RACI matrix and separation of duties.

Module 3: lifecycle, meta-model, maturity model, implementation roadmap, the cognition plane for AI agents and eleven role-based reading paths.

Appendices: five templates, worked examples, an SME playbook and an end-to-end agentic case.

Order: www.ai-bok.nl/en (EUR 39.95 including VAT, 308 pages).