OPERATOR DESK REFERENCE

BUILD ONCE. COMPLY TWICE.

INTEGRATED AI GOVERNANCE

One management system. One evidence spine.
Aligned to ISO/IEC 27001 and ISO/IEC 42001.

OPERATING SPINE

AI System	AI Inventory	**AISIA**
Model Card	**Reuse / Extend / New**	Unified Risk Register
Evidence Index	Unified Evidence Library	**Technology & Risk Committee**

Raf Rafaqut

Build Once. Comply Twice.

Operator Desk Reference for Integrated AI Governance Aligned to ISO/IEC 27001 and ISO/IEC 42001

Raf Rafaqut

Copyright page

Published by AIBI Systems | Turn Complexity into Clarity™
Capital Office, 124-128 City Road, London, EC1V 2NX, United Kingdom

For permissions and enquiries:
hello@aibisystems.co.uk
aibisystems.co.uk

Disclaimer

This book is for informational purposes only. It is not legal, regulatory, security, audit, or professional advice. You are responsible for how you interpret and apply the material, including any decisions and outcomes. No audit outcomes, certification outcomes, regulatory compliance, or risk reduction results are guaranteed. Seek competent professional advice where needed.

Standards and trademarks

References to ISO standards are used for alignment and intent framing only. ISO publications are copyrighted by ISO. The author and publisher are not affiliated with ISO. All trademarks are the property of their respective owners.

First edition: 2026

Preface

This book is for implementers who need AI governance to work under operational pressure. It is written for governance, risk, and compliance (GRC) and InfoSec teams, AI governance leads, internal audit, and accountable owners who need a model they can run, defend, and maintain.

The problem is rarely intent. It is fragmentation. A supplier enables a feature. A local team starts using it. Procurement holds one part of the story. Risk holds another. Evidence sits somewhere else again. By the time someone asks what was approved, by whom, on what basis, and whether the current use still fits that decision, the organisation has activity but not control.

The weakness AI exposes is not ignorance of governance. Most organisations already have management systems, policies, risk registers, controls, and evidence practices built around information security, data protection, or broader operational risk. The problem is that AI arrives outside those systems. It arrives through supplier contracts, platform activations, and commercial renewals rather than planned technology projects with formal governance gates. By the time a governance process formally notices a system, the dependency is already formed and the leverage to shape how it is governed has narrowed. That is the operating reality this book is designed for.

This book gives you one design rule: Build Once. Comply Twice. Use the Information Security Management System (ISMS) where it already answers the governance question. Extend it where AI changes the decision, the trigger, or the evidence burden. Create something new only where no credible route already exists.

The emphasis is operational. For most organisations, the urgent issue is supplier AI, embedded features, third-party services, quiet scope expansion, and weak ownership. The real

question is simple: where the organisation relies on AI, can it explain and defend that reliance?

This book does not replace the standards. It does not offer legal, regulatory, certification, or audit advice.

Judge the book by whether you can open a live AI system and follow the current decision without reconstruction. If you can do that without explanation or recovery work, the model is doing exactly what it was built for.

How to use this book

This book is a desk reference. You do not need to read it like a textbook unless you are building the full operating model from the ground up.

Use it against live AI systems, live ownership questions, live approval decisions, and live evidence gaps. Read with implementation in mind. The aim is control without building a second governance system. If the route survives challenge on a real AI system, keep it. If it breaks, fix the route before you widen coverage.

Standards and guidance will evolve. Keep the book's operating spine - scope, assessment, risk, evidence, governance, and review - as the stable layer. Draft guidance such as ISO/IEC 42005 (DIS) can add context. They do not replace the management-system base. If you operate in or sell into the EU, remember this: ISO/IEC 42001 can support your operating model, but it does not by itself create a presumption of conformity with the EU AI Act.

Reading paths

Start with the pressure you are under, not chapter order.

If you are building deliberately, read in sequence. The early chapters define the problem, scope what is in play, and set the assessment route. The middle chapters establish governance authority, integration logic, evidence, and risk. The later chapters deal with responsible use, supplier AI, and implementation.

If you are under immediate pressure, start with scope, ownership, assessment, and evidence retrieval.

If your organisation is supplier-heavy, keep that bias from the start. For many organisations, embedded supplier AI is the main exposure. The hard questions are visibility, approved use, supplier opacity, change triggers, compensating controls, and whether you can still explain the live position when the supplier moves first.

- **Executive sponsor:** Introduction, Chapter 1, Chapter 4, Chapter 5, Chapter 11, Chapter 12, and the Closing Note. This path shows the problem, the assessment and authority model, the supplier reality, and the implementation route without the full evidence architecture.
- **Implementer under pressure:** Chapters 3, 4, 5, 7, 8, 9, 10, and 11. Read in that order. The pressure appears in that order: what is in scope, how it is assessed, who decides, how risk is recorded, what use boundary is approved, what gate stops unsafe use progressing, and how supplier AI stays inside the same route.
- **Assurer or internal auditor:** Introduction, Chapter 3, Chapter 4, Chapter 7, Chapter 8, Chapter 9, Chapter 11, the Appendix, and the Closing Note. This path goes straight to scope boundary, assessment logic, evidence retrieval, live risk posture, transparent use, and supplier exposure, with the worked examples in the Appendix as the practical anchor.
- **Fast orientation:** Introduction, Chapter 1, Chapter 3, Chapter 4, Chapter 7, and Chapter 12. This takes you through the operating spine in order: the problem, scope boundary, assessment gate, evidence route, and implementation pattern. Start here if you have one hour and need the model before a meeting.

How to use the end-of-chapter block

If you are moving quickly, read the Executive Summary and the end-of-chapter block first. That gives you the problem the chapter solves, the outputs you should create, and the evidence you should be able to retrieve if the control is really operating.

Use the end-of-chapter block as a working check, not as a recap. If you cannot point to the outputs at the end of a chapter, the chapter has been read but not implemented. If you cannot retrieve the evidence, the control is not yet defensible.

After each chapter, ask three questions. What should now exist? Who owns it? Could we retrieve it tomorrow, fast and unaided?

How to avoid duplication

Before creating any new governance record or process for an AI requirement, ask whether an existing ISMS control already carries it. Apply the Reuse / Extend / New test: reuse what already works, extend where AI changes the question, and create something new only where no existing route can carry it cleanly. Chapter 6 shows how to apply that test in practice.

Toolkits and AIBI Systems

The book is complete on its own. You do not need a toolkit for the governance model to be valid, usable, or defensible in audit.

The toolkits provide reusable artefacts, workflows, and evidence structures that reduce drafting effort. They do not fill gaps in the operating model, and they are not a second method.

Before working through the book, you can run the free AIBI Governance Diagnostic at check.aibisystems.co.uk. It takes around five minutes and identifies where your current AI governance model is likely to break under challenge - across scope, ownership, evidence, risk, supplier AI, and lifecycle control. It requires no preparation and produces a readiness snapshot with one recommended next step.

Operator default

Your default reading mode should be practical and selective. Start with the chapter that matches the live problem in front of you, then trace backward or forward only as far as needed to make the route hold.

Do not try to perfect the architecture before you have tested it on a live AI system. Pick one or two live AI systems. Identify the AI Owner. Check whether the AI Inventory entry exists. Test

whether the AISIA is proportionate. Review whether the current approved-use boundary, risk position, and evidence route can be retrieved cleanly and unaided.

If the answer is no, fix that first before widening coverage

Introduction

At a glance

Purpose: Set the operating spine and explain how the desk reference works.
You leave with: A clear picture of the integrated system before the detail starts.
Common failure point: Building governance around documents rather than decisions.
System link: Every chapter that follows extends from this spine.

Executive Summary

The operating artefacts are few and each has a distinct purpose. The AI Inventory makes AI systems visible and owned. AISIA (AI System Impact Assessment) decides how deeply each system must be assessed and what approval path it follows. The Model Card records approved use, limits, and change visibility. The Unified Risk Register logs material uncertainty, treatment decisions, and acceptance. The Evidence Index points reviewers to controlled artefacts in the Unified Evidence Library. The Technology & Risk Committee takes the decisions that exceed delegated authority. Each artefact does a different job. Together they turn general intent into attributable decisions, proportionate controls, and retrievable proof.

The desk reference works in two ways. Read end to end, it takes you from first principles to implementation. Read selectively, it answers live questions that need a defensible response: what counts as an AI system in scope, who owns it, how deeply it needs to be assessed, and what evidence should already exist.

Why this exists

Figure 1 shows the integrated operating spine. Start here. Every chapter that follows develops one segment of this route, but the system only works when the segments connect.

This Introduction exists to stop you reading the later chapters as separate topics. The route matters more than the individual artefacts. If the spine breaks, governance breaks with it.

Figure 1: Integrated AI governance operating model - One management system, one evidence spine

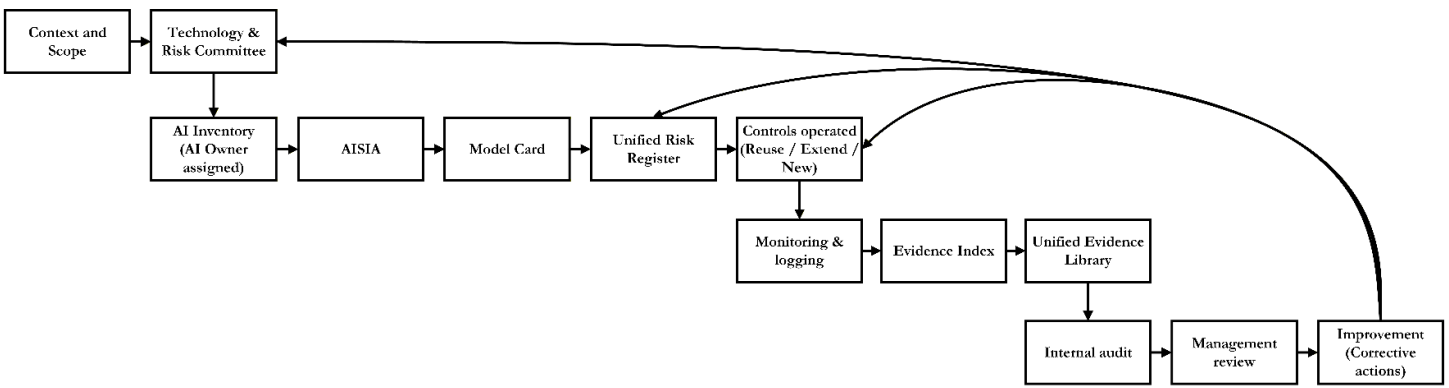

The operating reality

The spine runs left to right. Context and scope feed the Technology & Risk Committee. The AI Inventory, with an AI Owner assigned, triggers the AISIA. The AISIA output drives the Model Card and the Unified Risk Register. Controls operate under the Reuse / Extend / New classification. Monitoring and logging feed the Evidence Index, which links to the Unified Evidence Library. Internal audit draws from the library. Management review closes the loop. Corrective actions feed back into improvement, and material changes return to the Technology & Risk Committee.

The model is deliberately small. One AI system becomes visible in the AI Inventory, moves through AISIA, resolves into a Model Card and a Unified Risk Register entry, and proves itself through the Evidence Index and the Unified Evidence Library.

Nothing is built twice. Keep this diagram visible during implementation. It is the quickest way to test whether an artefact belongs on the main route or is creating duplication.

Regulatory grounding

This book is built on ISO/IEC 27001 and ISO/IEC 42001, but most organisations do not operate in a standards-only environment. They also face live legal, supervisory, and sector pressure.

In Europe, the EU AI Act is being applied in stages. In the UK, existing data protection, conduct, and sector rules already shape how AI use is judged in practice. The point of this book is not to summarise those regimes. It is to give you one operating model that still holds when external scrutiny arrives.

The integrated operating model is designed to hold under that external pressure without becoming a law guide. The operating question is the same under any regime: can the organisation show what is in scope, who owns it, what use is approved, what risk position was accepted, what change reopens review, and what evidence proves the route ran? If those answers exist, the organisation is in a stronger position to respond to standards-based audits and live regulatory challenge alike. Chapters 10, 11, and 12 carry the deeper treatment of how that operating model should hold under live regulatory, supplier, and implementation pressure.

Audit posture

Audit readiness begins with a clean trace from one AI system to its current decision basis and supporting evidence.

Audit test:

- Sample one AI system from the AI Inventory to its AISIA decision
- Verify the AI Owner is named and current
- Confirm the Evidence Index resolves to controlled evidence locations

Minimum evidence:

- AI Inventory entry with owner and change history
- AISIA outcome and approval record
- Evidence Index export pointing to three retrievable artefacts

Key Insights and Recommended Next Steps

Build the first pass around one visible AI system, one named AI Owner, one proportionate assessment route, one linked risk record, and one evidence path that works under challenge. Once that spine holds, the later chapters can deepen it without duplicating records.

A) Outputs

- One live AI system traceable through a named AI Owner, a proportionate AISIA, and a retrievable evidence path
- AISIA workflow defined with approval roles
- Unified Evidence Library location established
- Evidence Index skeleton published
- Unified Risk Register fields confirmed for AI risks

B) Evidence to retain

- AI Inventory export and owner assignment rule
- AISIA workflow and approval-role definition
- Unified Risk Register field set
- Evidence Index versioned copy

Chapter 1 - The Duplication Trap

At a glance

Purpose: Show what duplication costs before it becomes visible - and make one operating route non-negotiable.
You leave with: A clear test for whether your current AI governance has a shadow route running alongside it.
Common failure point: Keeping a temporary tracker because the existing route feels slower.
System link: Nothing in Chapters 3 through 12 holds if the single-route discipline breaks here.

Executive Summary

Duplication does not feel like a governance failure when it starts. It feels like pragmatism. A team runs a quick pilot and keeps a tracker because the formal route is slower. Procurement sends a supplier questionnaire that starts carrying approval weight because nobody challenged it. A manager enables a feature and records a decision in a project tool because that is where the work already lives. Each choice is reasonable in isolation.

The problem arrives later, and it arrives fast. An incident forces review. A supplier changes how a model works. An auditor asks who approved the current use and on what basis. The organisation then discovers that no single record answers the question. It has activity distributed across teams, formats, and filing locations. Multiple answers exist to basic questions: what is in scope, who approved it, what risk was accepted, where is the proof?

Reconstruction is expensive. It consumes time that should be spent answering the real question. It exposes gaps that are hard to defend. It produces a picture that looks controlled only in retrospect. For an internal auditor, a regulator, or a sceptical senior reviewer, reconstruction is not governance. It is the absence of it.

The operating answer is one route of record. If an AI system is in use, it exists in the AI Inventory, has a named AI Owner, and follows the same decision route as any material technology change. Working inputs - pilot notes, questionnaire responses, team spreadsheets - can still exist. None of them becomes the point of record. That discipline, held at entry, is what prevents the later reconstruction problem.

Everything that follows depends on this discipline holding.

Why this exists

Duplication is the earliest governance failure and the hardest to recover from. It is also the most preventable. This chapter exists to make the defect visible before it hardens and to establish the single-record discipline that makes every later control viable.

The operating reality

Duplication forms through momentum, not negligence. A working document stays live because it is accurate for now and the formal route feels like overhead. That is the entry point. Once a working document starts carrying live decisions - approvals in meeting notes, risk positions in a team wiki, supplier assessments in an email thread - the shadow route has formed.

Four failures then compound each other.

Scope becomes negotiable. When ownership is split across trackers, teams interpret what is in scope differently. The AI system one team treats as a minor tool is treated by another as a governed deployment. Nobody agrees on the boundary because nobody owns the boundary.

Risk becomes non-comparable. When risk is recorded in more than one place using different criteria and different owners, the organisation cannot produce a current risk position for a given AI system. It can produce multiple positions and hope they are consistent.

Evidence fragments. When the approval is in one document, the risk assessment in another, and the supplier review in a third, retrieval requires someone who knows where to look. That person is usually unavailable when the question arrives under pressure.

Accountability fractures. When ownership is distributed, escalation breaks. Nobody can call a decision final because no single person owns the outcome.

Common failure point: a temporary AI tracker stays live because nobody made a formal decision to retire it.

The cost made visible

A governance and assurance team received an internal audit finding on AI oversight with ten working days to respond. The audit had asked three questions:

what AI systems are currently in use, who approved each one, and what risk position applies?

The team had answers. The problem was that those answers came from four different sources. The AI Inventory held twelve systems. A procurement register held six more, three of which appeared in the Inventory under different names. A second-line risk team held a spreadsheet with nine AI-related risk entries, four of which referred to systems not in either register. Two approval records existed only in meeting minutes that had not been formally retained.

Six days of consolidation work produced the combined view the audit required. Four of those days involved locating record owners who were not expecting the request and could not always confirm which version of a document was current. The audit report acknowledged the effort. It also noted that the governance model depended on manual reconciliation rather than a maintained control structure.

That finding went into the executive report. The Technology & Risk Committee received it. Three senior managers were asked to explain the gap. The answer - that each team had moved quickly and kept their own records because the formal route was slower - was accurate and unconvincing. The committee required a formal remediation plan. The remediation plan required the same consolidation work to be completed again as a controlled exercise, with documented evidence that it had been done correctly. That work took three weeks and consumed more senior time than the original six days.

The cost was not the audit finding. The cost was the follow-up: committee attention, senior management time, a formal remediation exercise nobody had planned for, and the continued uncertainty about whether the consolidated picture was actually complete.

What good looks like

Good looks like one sampled AI system traceable through one current record chain without reconciliation. The AI Inventory entry is current. The latest decision is recorded once. The linked risk sits in the Unified Risk Register. The evidence route resolves cleanly through the Evidence Index.

The practical test is this: ask three teams about the same live AI system. If they point to the same scope record, the same decision route, and the same evidence path, duplication is under control. If they point to different artefacts, the duplication is still live.

Operator default: no AI system enters production or material use without an AI Inventory entry and a named AI Owner.

Audit posture

The audit test for duplication is simple and fast. Sample one AI system. If scope, decision, risk, and proof resolve through different record chains, duplication is the finding - even if each individual record is accurate.

Audit test:

- Sample one AI system and confirm it appears once in the AI Inventory
- Confirm the named AI Owner is current
- Confirm material risk sits in the Unified Risk Register
- Confirm the latest material decision followed the defined approval route
- Confirm evidence resolves through one controlled path

Each of these answers should exist as a retrieval step, not an explanation.

Once the route is singular, the next job is to turn requirement wording into live control rather than parallel paperwork.

Key Insights and Recommended Next Steps

Treat duplication as a defect, not a maturity stage. A shadow route creates extra maintenance, and the cleanup cost usually outweighs the convenience that created it.

A) Outputs

- Duplication map of active parallel artefacts
- Authoritative-route decision for each artefact class
- Named consolidation owner
- 30-day freeze, migrate, retire, redirect plan
- One completed sampled trace for a live AI system

B) Evidence to retain

- Duplication map and consolidation decision record
- Redirect log for retired artefacts
- Sampled AI Inventory and Unified Risk Register extracts
- Evidence Index export with version history
- Governance record approving consolidation

END OF PREVIEW

This is where the preview ends.

The full book continues across twelve chapters, worked examples, appendix templates, a complete back-of-book index, and two implementation bonus tools included with direct purchase

To purchase the full book

https://aibisystems.co.uk/build-once-comply-twice

Direct eBook - £17.99 | Includes 30-Day Quick-Start Card and Regulatory & Enforcement Reference Sheet

Instant PDF download. No DRM. Also available on Amazon Kindle (£9.99) and Paperback (£19.99).

AIBI Systems | Build Once. Comply Twice. | aibisystems.co.uk