Ciphrix

BUYER RESOURCE

PRACTICAL GUIDE

ISO 42001 AI Governance Field Guide

Build AI governance that product teams can operate, leaders can oversee and customers can trust

PUBLISHED 27 August 2026 **FROM** Ciphrix **DOCUMENT ID** CPX-BG-2608-7194

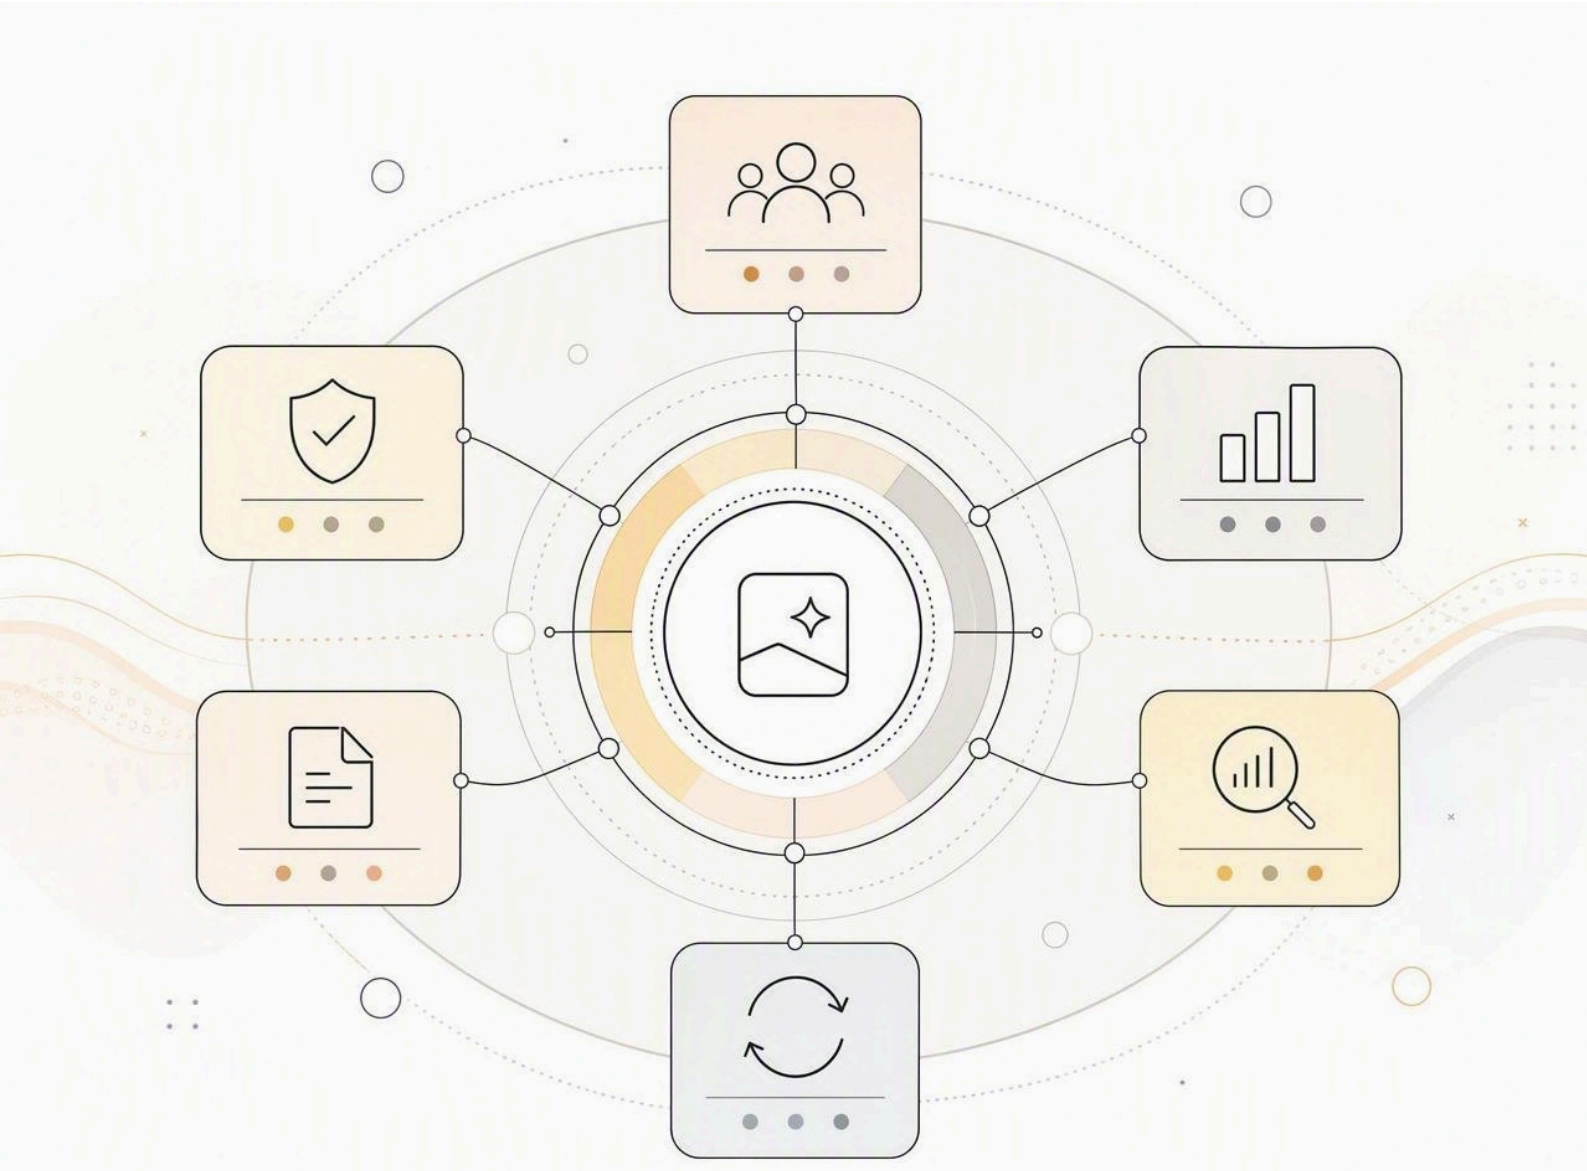

BUYER RESOURCE

ISO 42001 AI Governance Field Guide

CIPHRIX
BUYER GUIDE

Build AI governance that product teams can operate, leaders can oversee and customers can trust

Build AI governance that product teams can operate, leaders can oversee and customers can trust

AI governance is easy to make sound impressive and hard to make real. A principles document will not tell a product team whether a new model integration is in scope. A board presentation will not prove that risk was assessed before a customer-facing release. A supplier's assurance report will not answer how *your* organisation is using the system.

ISO/IEC 42001 gives organisations a management-system structure for governing AI. This guide translates that structure into operating work: inventory, decisions, risk and impact assessment, controls, evidence, monitoring, corrective action and management review. It is a practical field guide, not an official ISO template or legal advice.

1. The distinction that changes everything: principles vs an operating system

An AI policy tells people what the organisation expects. An AI Management System (AIMS) shows how those expectations are made, applied, checked and improved over time.

For every material AI use, a credible programme should be able to show:

1. what the system does and who owns it;
2. why it is in scope, excluded or being monitored;
3. what data, users, suppliers and decisions it touches;
4. what risks and impacts were assessed;
5. what controls, approvals and operating procedures apply;
6. what monitoring, issues, exceptions and corrective actions have occurred; and
7. how management reviewed the resulting information.

ISO describes ISO/IEC 42001 as a standard for establishing, implementing, maintaining and continually improving an AIMS. That means written design is only one part of readiness. A programme also needs operating records that show its reviews, approvals, control activity and improvements occurring in practice.

Important boundary: operating an AIMS aligned to ISO 42001 is not the same thing as certification. Certification is a separate, voluntary third-party conformity assessment. Neither alignment nor certification makes an AI system risk-free or transfers product, legal or management accountability.

2. Build the five records that make AI governance visible

Before adding complex workflows, establish five living records. Together, they give a programme enough structure to answer the first questions from customers, auditors, leadership and its own product teams.

RECORD	MINIMUM USEFUL FIELDS	WHAT IT PREVENTS
1. AI inventory	System/use-case name, purpose, lifecycle stage, business and technical owner, built/bought/hybrid, data, users, supplier, review date.	"We do not know where AI is being used."
2. Scope and decision record	Include, exclude or monitor; rationale; conditions; evidence owner; review trigger.	Silent scope creep and unexplained exclusions.
3. Risk and impact record	Context, affected groups, risks, potential impact, mitigations, residual-risk decision, owner and review date.	A generic risk register detached from real use cases.
4. Control and evidence record	Applicable controls, required approvals, evidence source, testing/review expectation and exception path.	Policies that cannot be shown to operate.
5. Monitoring, issue and decision history	Signals reviewed, incidents/near misses, actions, approvals, corrective actions, escalation and management-review outcome.	A programme that cannot learn from change or failure.

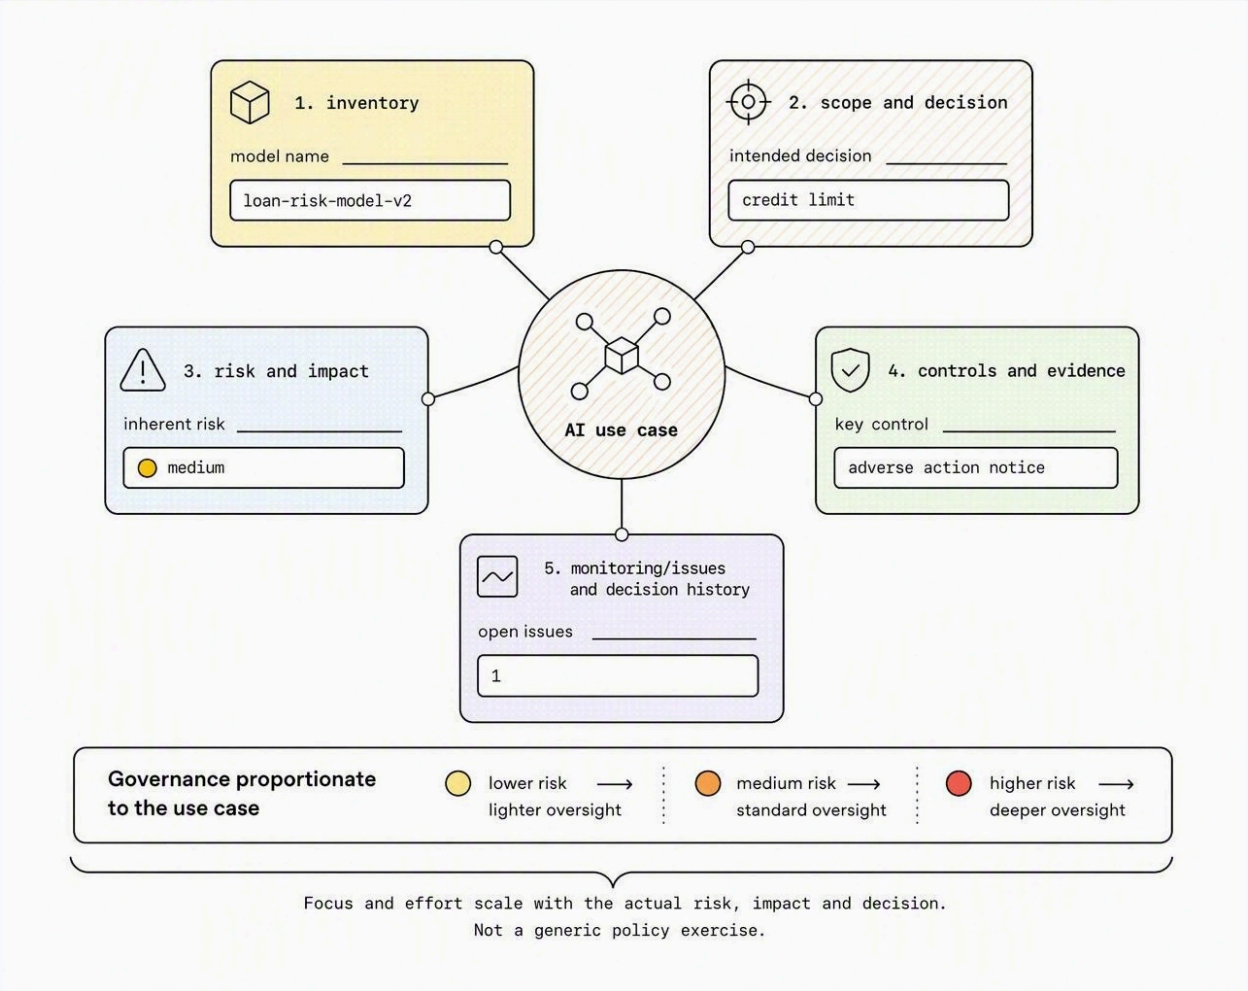

Five connected governance records around one AI use case: inventory, scope and decision, risk and impact, controls and evidence, and monitoring, issues and decision history.

What good looks like

A customer-support assistant, an internal coding tool and a model supporting credit decisions should not be governed with the same depth. The records allow the organisation to apply proportionate control while making its rationale clear.

FIELD	EXAMPLE: CUSTOMER-FACING SUPPORT ASSISTANT
Purpose	Answer routine product-support questions using an approved knowledge base.
Owners	Support leader (business) and platform engineering lead (technical).
Use context	Customer-facing; answers are visible to end users and can affect their experience.
Data / supplier	Support content and approved customer context; third-party model/API dependency.
Key risks	Inaccurate output, disclosure of inappropriate information, degraded service, supplier change.
Controls	Approved knowledge sources, access/data restrictions, release review, feedback and incident route.
Evidence	Inventory record, supplier review, testing/release record, monitoring review and issue log.
Review triggers	New data source, change of vendor/model, expanded use case, production incident or material complaint.

3. Scope first: include, exclude or monitor — deliberately

The first AIMS decision is not “what policy should we write?” It is what the management system covers. Scope determines which AI systems, teams, data, vendors, processes and records the organisation must govern.

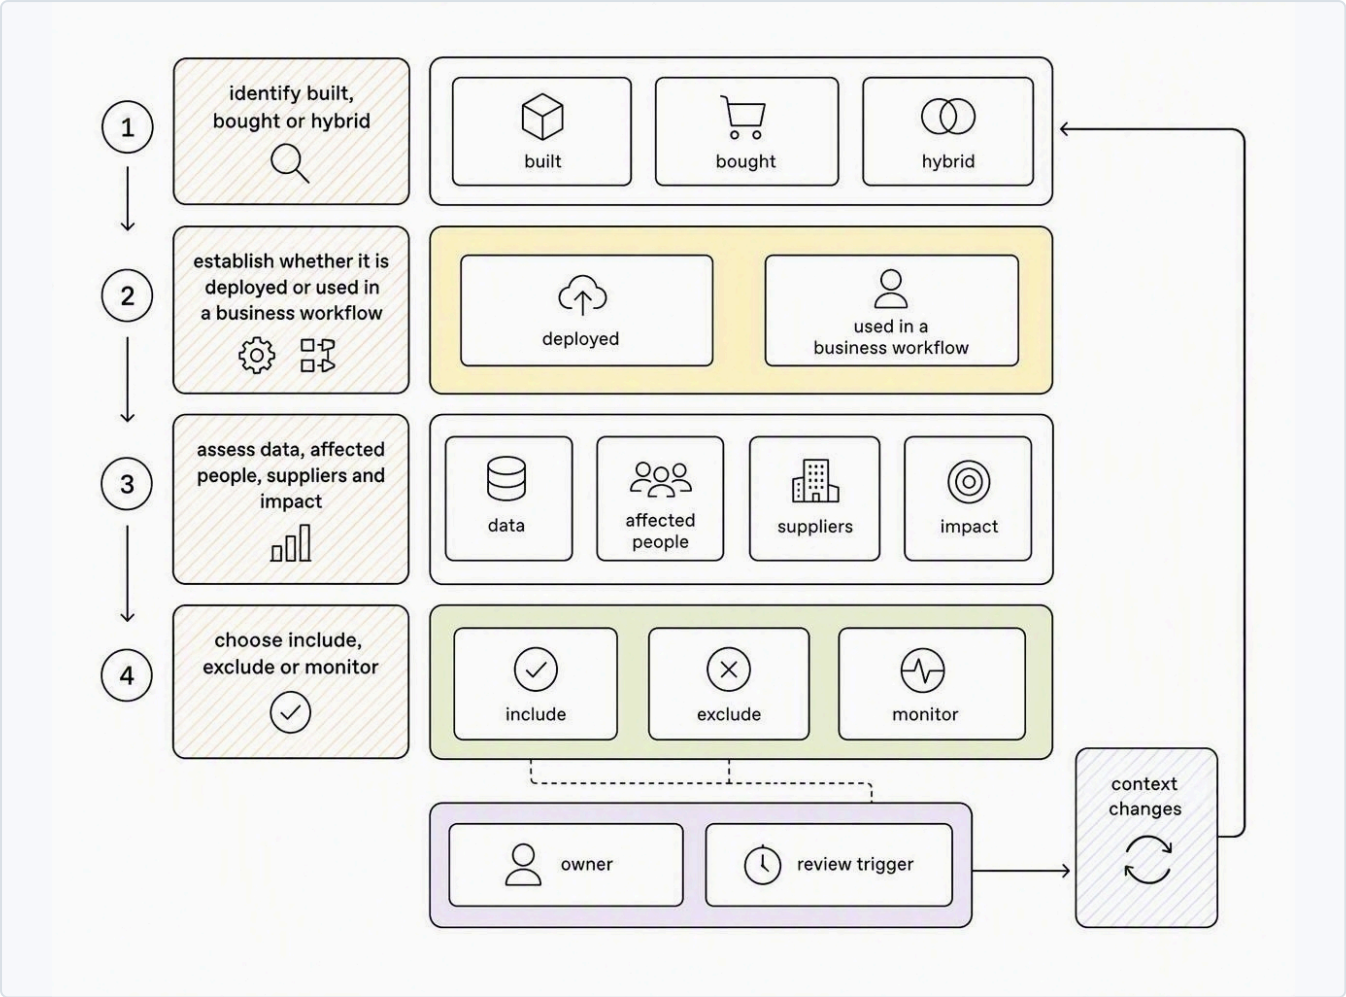

A decision path for a new AI use case: assess context, then include, exclude or monitor it with accountable ownership and a review trigger.

Use this decision path whenever a new AI use case appears or a material condition changes.

Scope decision checklist

- Is the use case built, bought or hybrid?
- Is it deployed, or is it used in a business workflow?
- Who owns the business outcome and the technical operation?
- Which users or affected groups could experience its outputs or decisions?
- What data is used, created or exposed? Is any of it personal, confidential, regulated or proprietary?
- Which supplier, model, API, embedded feature or data provider is involved?
- Does it affect customers, employees, safety, rights, pricing, credit, health, legal decisions or another high-impact context?
- Is there a defined reason to include, exclude or monitor it — and a trigger to revisit that decision?

An exclusion can be legitimate. A research prototype with synthetic data that is not deployed may be excluded or monitored. The governance test is whether that decision is documented, owned and revisited when the prototype

becomes production, its data changes or the context changes.

4. Run an AIMS as a loop, not a set of documents

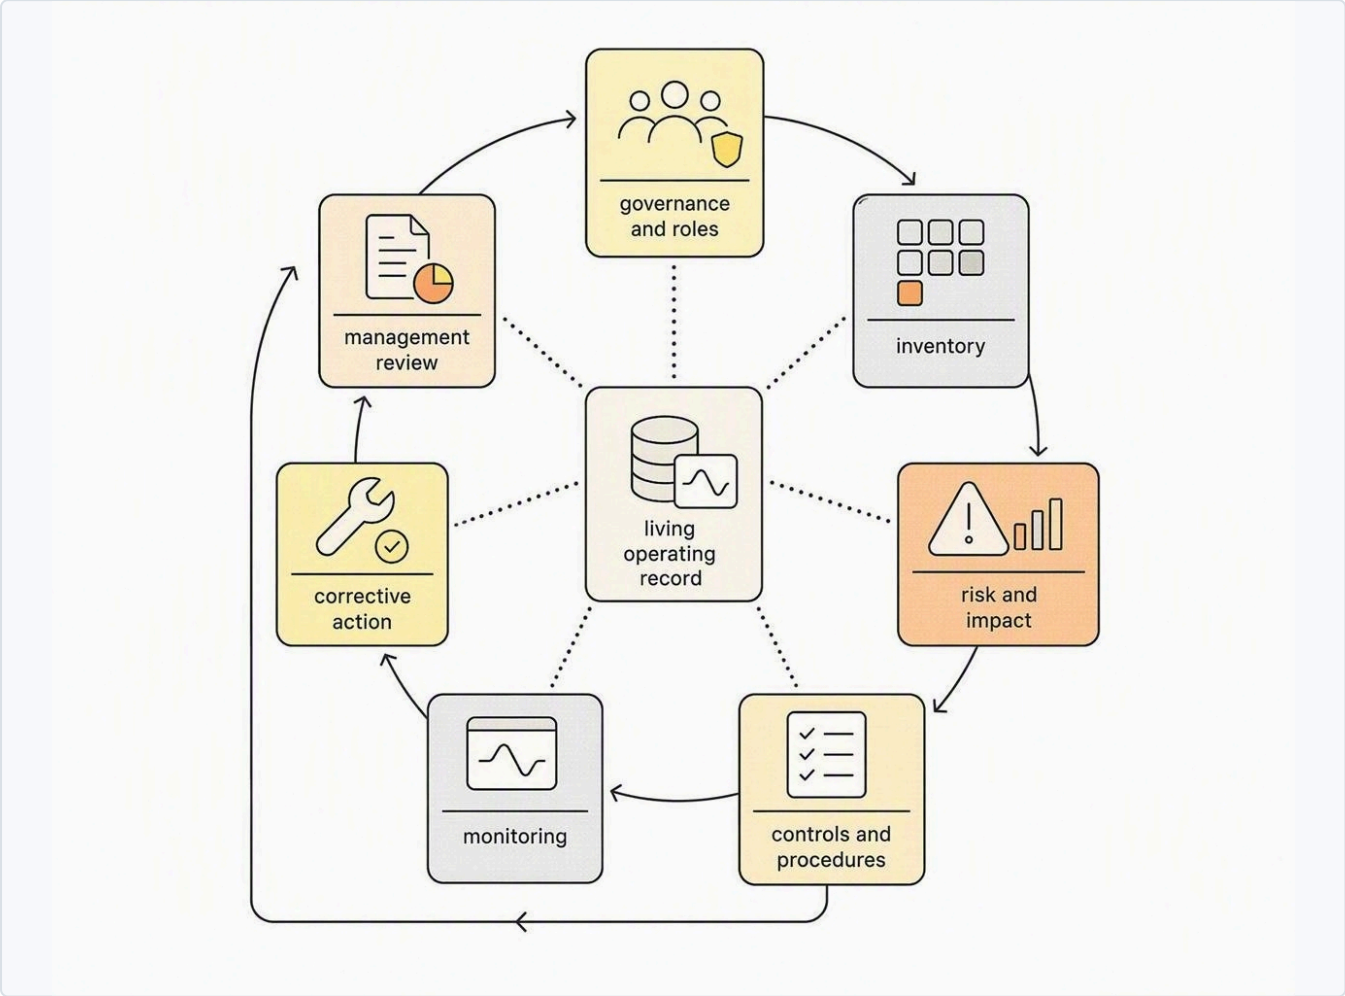

An AIMS operating loop connecting governance, inventory, risk and impact, controls, monitoring, corrective action and management review through living records.

This loop is the heart of the programme. Each stage should produce a record and create a path to the next stage.

AIMS ELEMENT	PRACTICAL OPERATOR QUESTION	EVIDENCE THAT USUALLY MATTERS
Governance and roles	Who can approve use, accept risk, change the system and escalate an issue?	Role map, meeting/approval records and escalation policy.
Inventory	Which AI systems are covered and what has changed?	Current inventory with owners, use context and review dates.
Risk and impact	What could go wrong for the organisation and affected people, and what did we decide?	Assessment, mitigations, approval and review trigger.
Controls and procedures	How do policy expectations show up in engineering, product, vendor and data workflows?	Release, access, testing, supplier and change records.
Monitoring	How does the team notice that risk, performance, use or context has changed?	Monitoring reviews, user feedback, system signals and periodic checks.

AIMS ELEMENT	PRACTICAL OPERATOR QUESTION	EVIDENCE THAT USUALLY MATTERS
Corrective action	What happens after an incident, failed control or identified weakness?	Issue log, action owner, closure proof and effectiveness review.
Management review	What did leadership review and what decisions followed?	Management-review pack, actions, resource decisions and follow-up.

Use NIST AI RMF as a useful companion, not a substitute

ISO 42001 gives the management-system structure. The NIST AI Risk Management Framework provides a voluntary risk-management lens with **Govern, Map, Measure and Manage**. Teams can use the NIST functions to make their risk work more concrete while retaining ISO 42001's management-system discipline.

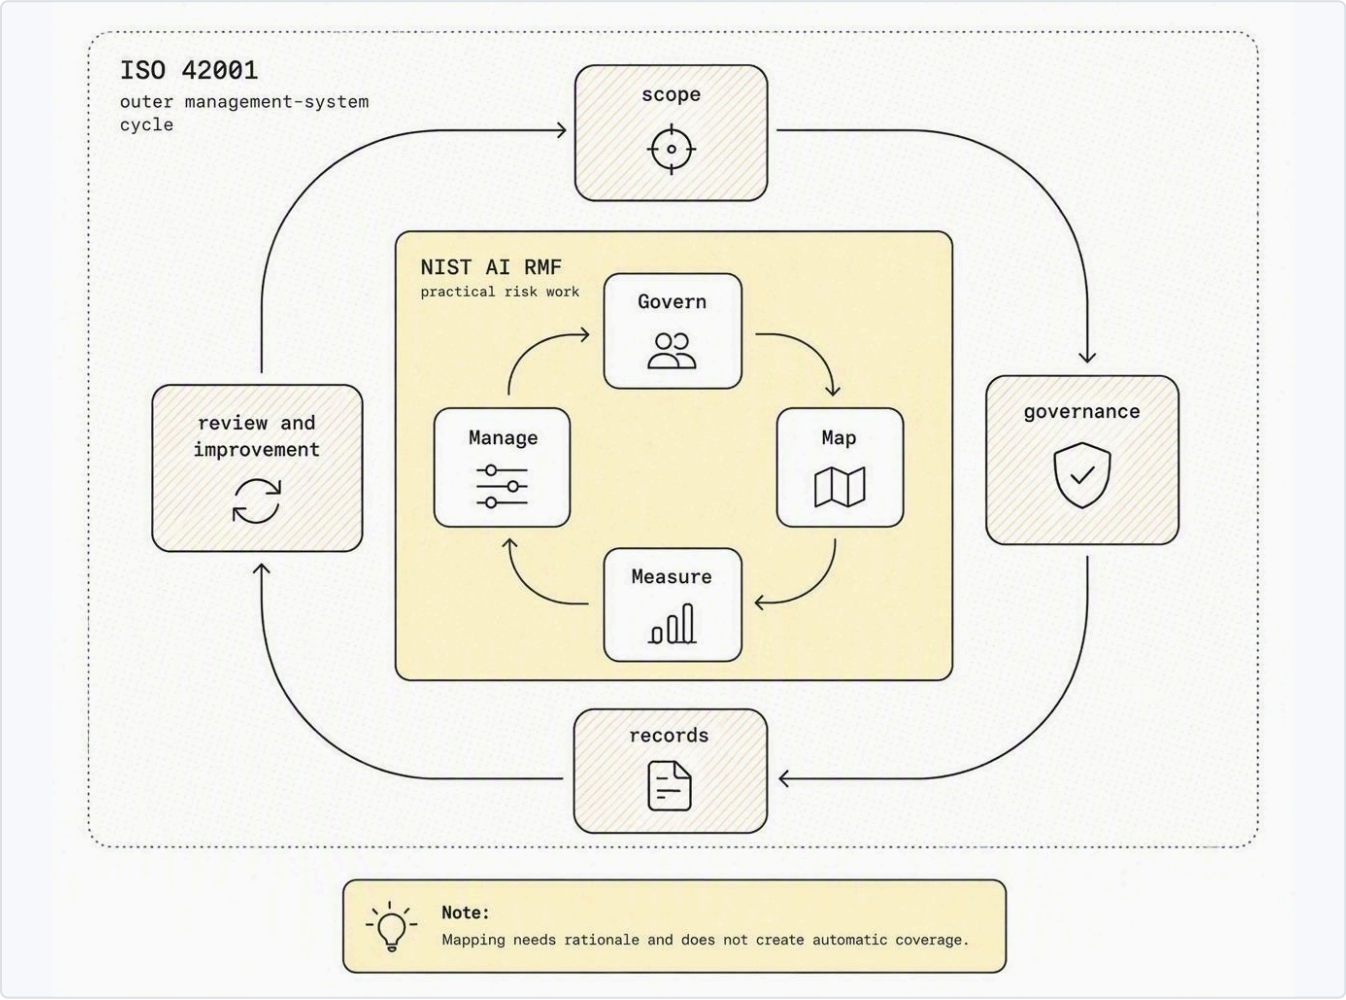

ISO 42001 management-system work and NIST AI RMF risk functions shown as complementary, not automatically interchangeable.

Do not claim one framework automatically satisfies the other. The useful question is: which existing records, practices and evidence can be mapped with a clear rationale, and what AI-specific gap remains?

5. Put ownership where the work happens

AI governance is cross-functional, but cross-functional cannot mean ownerless. A programme needs a shared role map before it needs a large committee.

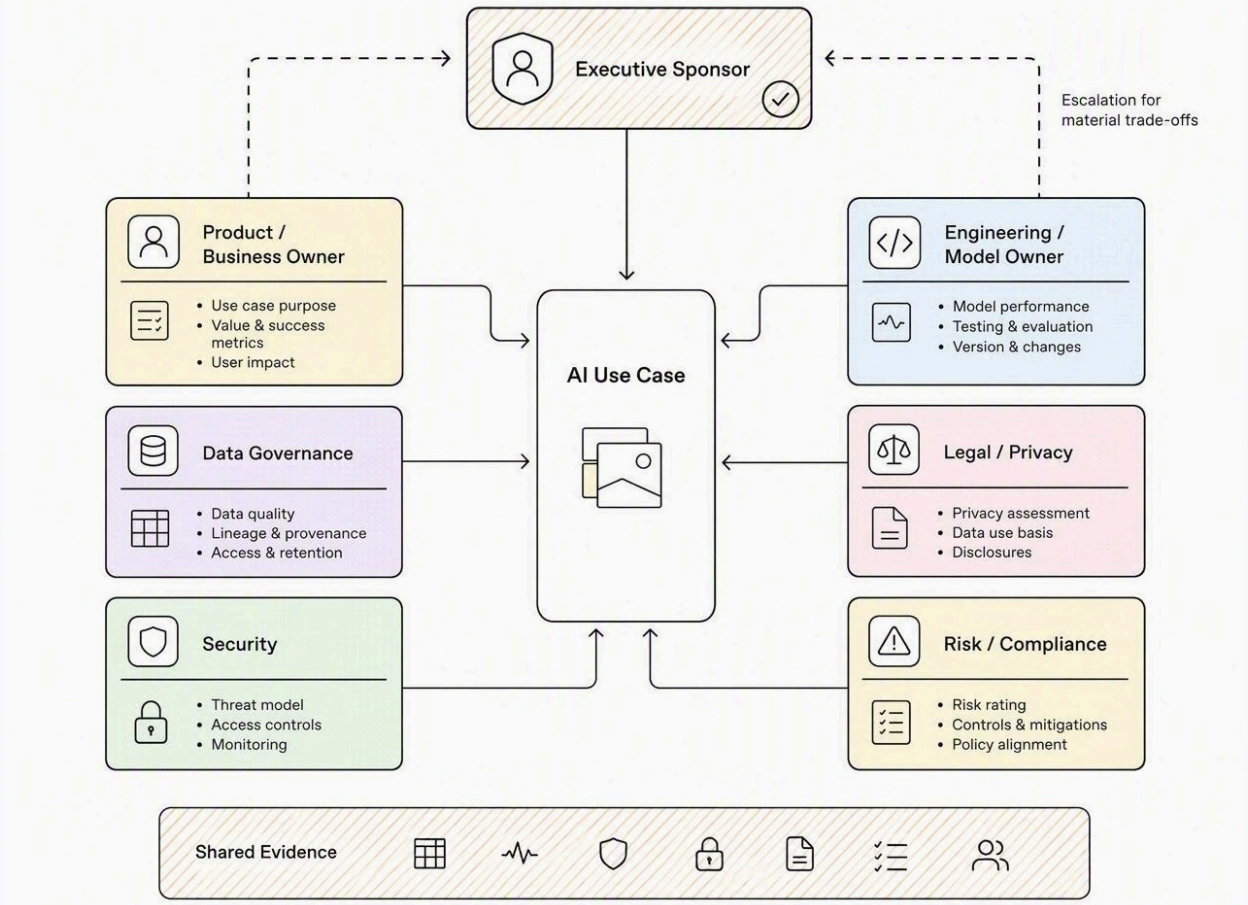

A role map connecting executive sponsorship, business and technical ownership, data governance, security, legal/privacy and risk/compliance to one AI use case.

ROLE	ACCOUNTABLE FOR	SHOULD BE ABLE TO SHOW
Executive sponsor	Risk appetite, resources, material escalation and management review.	Decisions, priorities and accountable delegation.
Product / business owner	Purpose, intended users, expected value and acceptable use.	Clear use context and changes in business intent.
Engineering / model owner	Technical design, integration, testing, deployment, monitoring and change.	System records, test/release evidence and issue response.
Data owner / governance	Data context, quality, permissions and data-use constraints.	Data sources, classification and approval/review records.
Security	Security controls, access, vendor/security review and technical risk response.	Security review, control evidence and exception handling.
Legal / privacy	Legal, contractual and privacy implications; escalation input.	Review decisions and conditions of use where needed.
Risk / compliance	AIMS method, evidence expectations, challenge, reporting and audit coordination.	Current records, open actions and review cadence.

Decision rule: if a team has only been asked to “support AI governance,” translate that into a named record, a recurring decision, a defined evidence contribution or an escalation responsibility.

6. Create a board- and customer-ready evidence pack

Executives and customers do not need every technical detail. They do need enough traceable information to understand what is being governed, who owns it, what risk remains and what happens next.

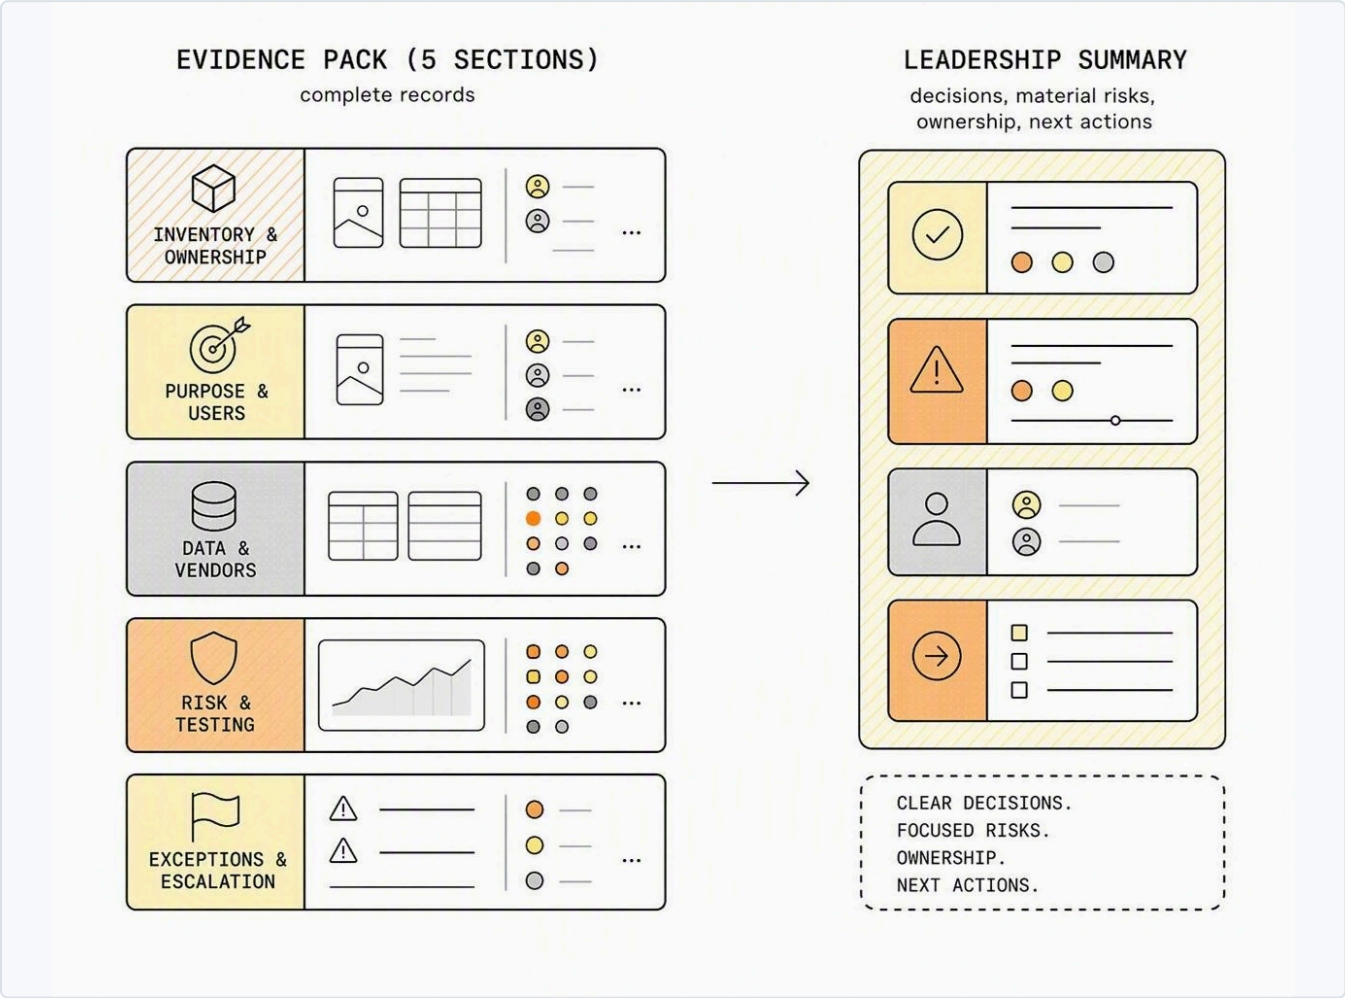

A board- and customer-ready AI evidence pack: five record sections distilled into a leadership summary of decisions, risks, ownership and next actions.

The five sections

1. **Inventory and ownership** — current AI systems, lifecycle stage, business owner, technical owner and criticality.
2. **Purpose and users** — intended outcome, user groups, decision context and scope.
3. **Data and vendors** — key data sources, sensitivity, third-party dependencies, contractual or assurance inputs.
4. **Risk and testing** — risk/impact summary, controls, testing/validation status and monitoring approach.
5. **Exceptions and escalation** — known issues, policy exceptions, incidents, mitigation actions, owners, dates and decision path.

What belongs in a leadership dashboard

Keep it small and decision-led:

- inventory coverage and material new use cases;
- high-risk or high-impact systems and their approval/testing status;
- systems involving sensitive or confidential data;
- material vendor AI dependencies and pending reviews;
- open exceptions, incidents and overdue remediation;
- decisions requested from the relevant executive, committee or board.

Volume is not the point. A growing number of use cases can be positive if inventory coverage, ownership and controls are keeping pace. A small number of high-risk systems may be unacceptable if approval, testing or exception ownership is unclear.

7. Start in 90 days — one use case at a time

Days 1–30: establish visibility and decision rights

- Name the executive sponsor and core operating roles.
- Build the first AI inventory, including internal tools, embedded vendor features and shadow-use visibility where proportionate.
- Define AIMS scope, inclusion/exclusion rationale and review triggers.
- Select one representative in-scope system and trace its product, data, supplier, risk and release path.
- Establish a simple risk/impact assessment and decision-record format.

Deliverable: an owned inventory, a visible scope decision and one end-to-end use-case trace.

Days 31–60: make controls and evidence operate

- Define required approvals, controls and evidence for the pilot system.
- Connect use-case changes to existing product, engineering, security, privacy and vendor workflows.
- Capture testing, supplier, data and release records where relevant.
- Establish monitoring/review cadence and the issue/exception route.
- Turn the pilot learnings into a proportionate approach for other risk tiers.

Deliverable: a usable control-and-evidence model, not a generic policy collection.

Days 61–90: test the evidence and management loop

- Conduct a review as if a customer, auditor or board committee asked for the evidence pack.

- Test whether the five records are current, linked and understandable to someone who did not build them.
- Review open actions, exceptions and risk decisions with the relevant management group.
- Document improvement actions, resource needs and next systems to onboard.

Deliverable: a functioning governance loop with an evidence pack and a realistic expansion plan.

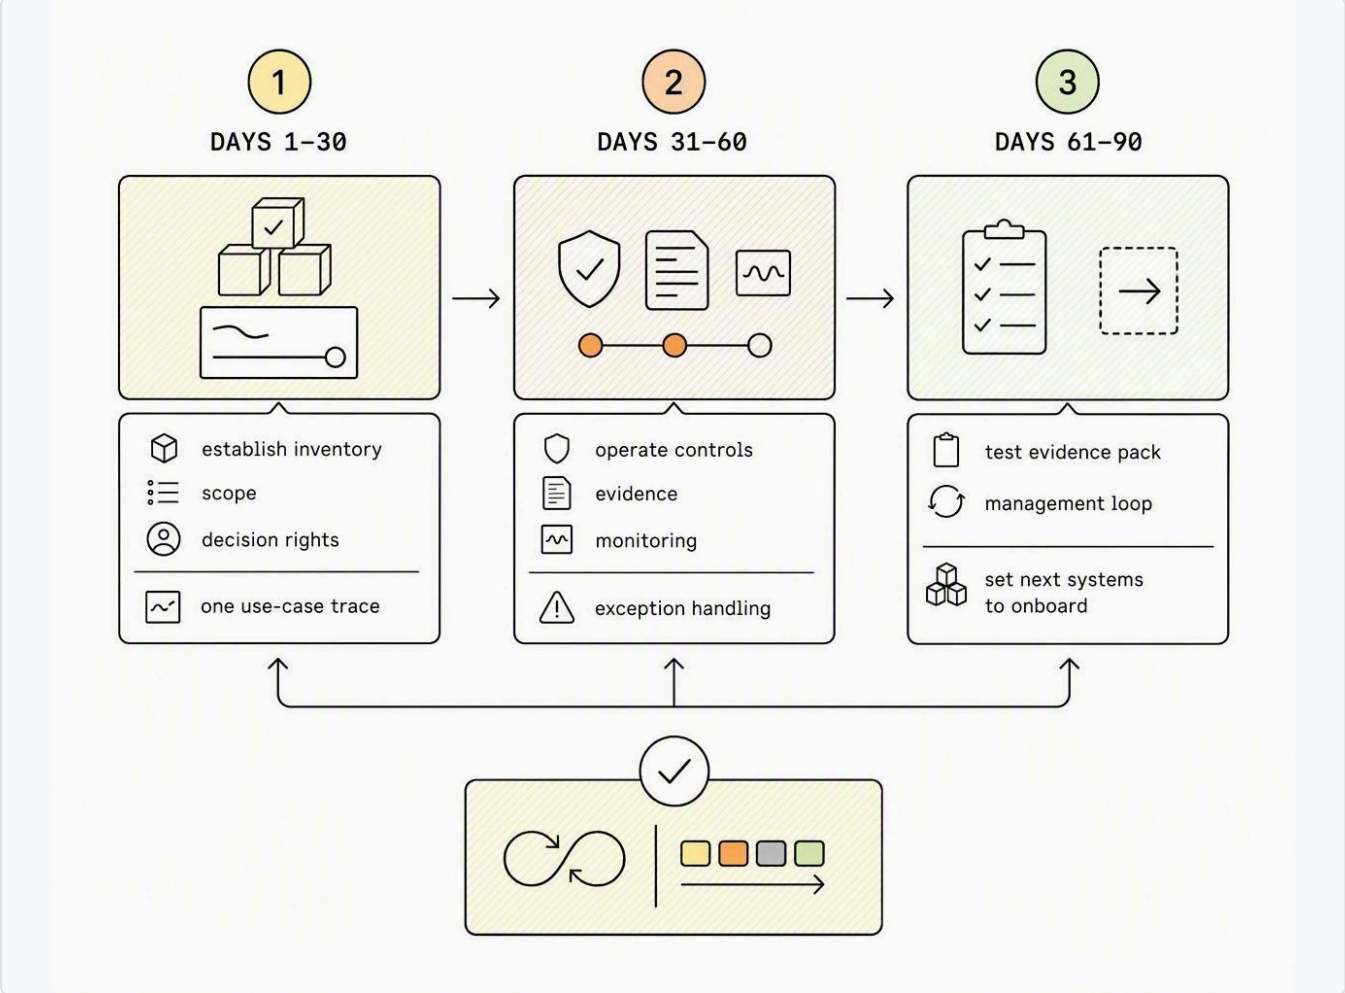

A 90-day AI governance launch: establish inventory, scope and decision rights; operate controls and evidence; then test the pack and management loop.

ISO 42001 readiness checklist

- AIMS scope is documented, with inclusions, exclusions and review triggers.
- AI inventory covers material built, bought and hybrid use cases in scope.
- Every in-scope system has business and technical ownership.
- Risk and impact records connect to real data, users, suppliers and use context.
- Applicable controls are tied to operating workflows and evidence sources.
- Testing, release, monitoring, supplier and incident records are retained where relevant.
- Exceptions and corrective actions have owners, dates, decisions and closure evidence.

- Management review is scheduled and produces follow-up actions.
- Existing ISO 27001, privacy, security and risk practices are mapped carefully — not assumed to cover AI by default.

Make AI governance operate in your real environment

Ciphrix helps teams turn AI governance into maintainable work: AI agents support the execution of policies, evidence, controls and questionnaires; expert humans guide implementation and retain scope, risk and approval decisions.
Book a demo to see how an AI governance workspace can map to your actual AI use cases, evidence sources and review rhythm.

Source notes

This guide synthesises Ciphrix's published ISO 42001 compliance guide, ISO 42001 framework page, ISO 42001 vs NIST AI RMF comparison, NIST AI Risk Management Framework guide, board oversight for AI and AI governance roles. Confirm framework, regulatory, contractual and certification requirements against the applicable sources and the organisation's own context.