AI SECURITY ENGINEERING FIELD GUIDE · 2026

Chapter 10 · AI Governance, Risk, and Compliance

Standalone learning module for LMS delivery and required reading.

FORMAT	USE	SCOPE	AUDIENCE
Standalone PDF	LMS reading module	Single chapter	Learners

CHAPTER 10

AI Governance, Risk, and Compliance

Evidence

GOVERNANCE EVIDENCE

Map policy statements to controls, owners, tests, telemetry, and public-safe evidence.

Chapter 10 · AI Governance, Risk, and Compliance AI Security Engineering Field Guide 2026

> Governance becomes real when policy maps to controls, owners, telemetry, and evidence.
>
> FIELD GUIDE

> **FIELD GUIDE: GOVERNANCE EVIDENCE REVIEW**
>
> Inspect whether AI risk language can be traced to engineering controls, evidence records, approvals, exceptions, and remediation.

This domain covers AI governance, risk, compliance, audit readiness, policy-to-control mapping, framework translation, risk acceptance, executive reporting, customer assurance, and evidence management. It matters when governance teams need proof that AI risks are controlled by engineering work rather than policy language alone.

The useful governance review asks whether a public or internal statement can survive an evidence request. A policy, framework row, or committee decision is not enough unless it points to a control owner, system, test, log, exception path, and current artifact.

> **FIELD USE**
>
> Start with one external AI security statement or one governance requirement. Trace it to a control, owner, system, test, evidence record, exception status, and remediation path.

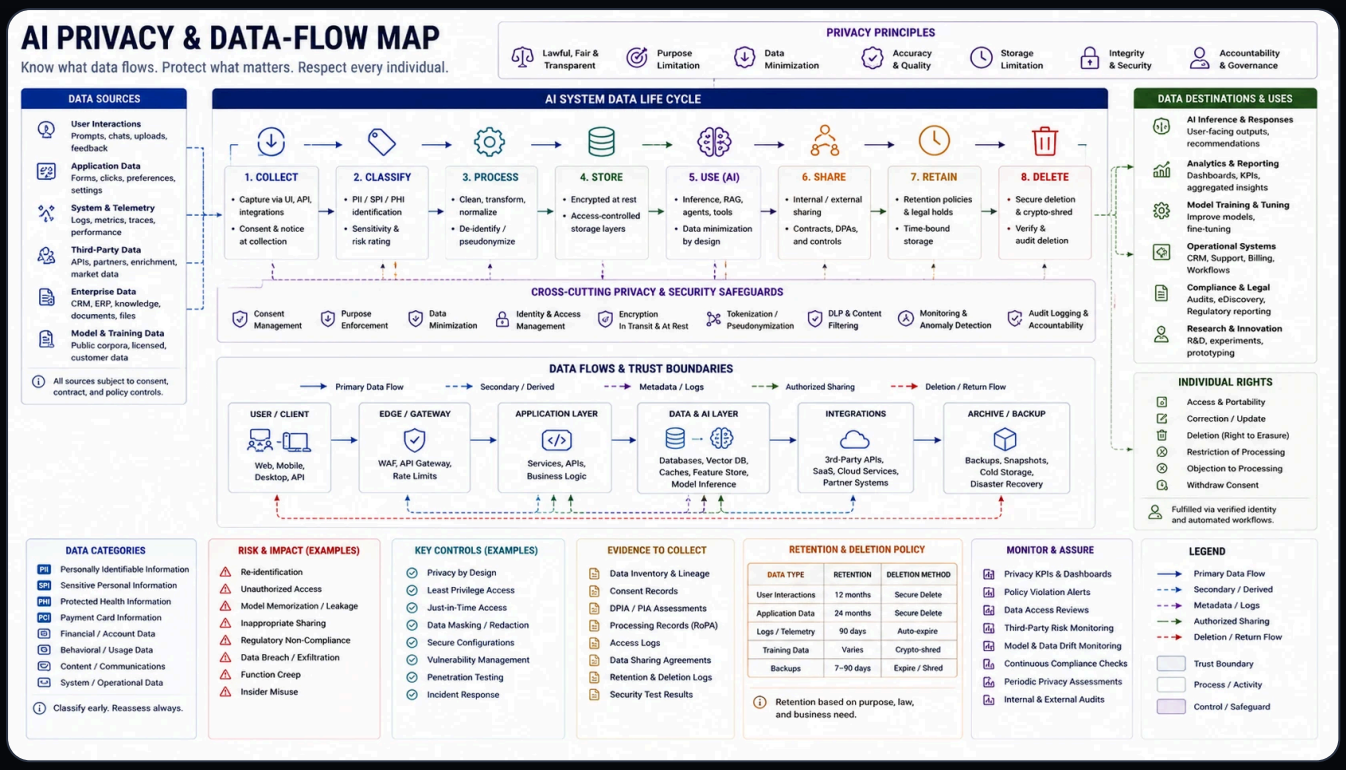

FIGURE 1: FIGURE 11: GOVERNANCE EVIDENCE CHAIN LINKING POLICY TO CONTROL OWNER, TEST ARTIFACT, CLAIM-READINESS STATUS, AND REMEDIATION BACKLOG

WHAT BREAKS HERE?

- Policy says controls exist, but systems do not produce evidence.
- Framework mappings are broad, stale, or disconnected from owners.
- AI risks have no severity model, expiry, or remediation path.
- External claims exceed available evidence.
- Governance meetings do not change backlog, gates, telemetry, or approvals.
- Customer questionnaire answers drift from the evidence the product can actually show.

WHAT TO INSPECT

- AI policies, standards, risk registers, framework mappings, and control owners.
- Evidence repositories, control tests, approval records, and exception registers.
- Trust center claims, sales answers, customer questionnaire responses, and audit narratives.
- Board or executive AI risk summaries.
- Backlog items linked to governance findings.
- Evidence freshness, artifact owners, review dates, and public-safety caveats.

WHAT TO ASK

- Which AI controls can be proven today?
- Who owns each control and each evidence record?
- What happens when an AI control fails?
- Which external claims are `public_claim_ready`, `public_claim_with_caveat`, `internal_or_teaser_only`, or `do_not_claim`?
- How do governance findings become engineering work?
- What evidence would satisfy a customer, auditor, or regulator?
- Which claims need caveats because evidence is directional, partial, or aggregate?

WHAT TO TEST

- Trace one policy requirement to a control, owner, log, test, and artifact.
- Review customer-facing AI claims against available evidence.
- Sample risk acceptances for owner, rationale, expiry, and compensating controls.
- Check whether framework mappings cite real artifacts.
- Verify governance findings appear in remediation backlog with retest dates.
- Compare questionnaire answers against trust-center language and evidence packs.

CONTROLS AND GUARDRAILS

- Control register with owner, system, evidence, test cadence, and failure response.
- Claim-readiness workflow for external AI and security statements.
- Risk acceptance workflow with expiry and compensating controls.
- Framework crosswalk tied to artifacts, not generic policy text.
- Governance-to-backlog process with severity, owner, due date, and retest evidence.
- Evidence freshness review and public-safe caveat rules.

EVIDENCE TO COLLECT

- AI control crosswalk.
- Evidence matrix.
- Claim-readiness ledger.
- Risk register and exception records.
- Audit and customer assurance artifacts.
- Governance remediation backlog.
- Evidence owner, review date, and claim-readiness status.

OUTPUT ARTIFACTS

- AI governance evidence matrix.
- Control mapping summary.
- Buyer-ready evidence pack.
- Risk acceptance review.
- Governance backlog.

RELATED SERVICES AND WORKBENCH TOOLS

TYPE	RELATED PATHS
Workbench	AI Control Crosswalk, Trust Scanner, Runtime Proxy
Services	Enterprise AI Security Readiness, AI Security Operating Model, AI Product Security Assessment
Handbook	Governance Evidence and Customer Trust, Logging and Telemetry, Evaluation and Regression Testing

DECISION - CONDITIONAL: CLAIM-READINESS DECISION

Do not publish an AI security claim unless the claim maps to a control, an owner, a test, and evidence that can be shown safely.