The AI Controls Catalog

Enterprise AI Governance for Boards, CISOs, and Auditors

A structured controls framework for responsible AI deployment across the enterprise.

NIST AI RMF | ISO/IEC 42001:2023 & 27001:2013 | SOC 2

Version 1.0 · May 2026

READER'S GUIDE

How to Use This Catalog

This catalog is structured for three audiences. Each domain card follows the same format.

FOR THE BOARD

Use the Executive Context and Overview cards, then scan the 'If you skip this' row on each domain card to understand where liability concentrates if controls are absent.

FOR THE CISO / COMPLIANCE LEAD

Use the domain cards as a gap assessment checklist. Each Policy, Control, and Evidence row maps directly to what an auditor will request.

FOR THE AUDITOR

Framework Coverage rows map each control to NIST AI RMF, ISO/IEC 42001:2023, ISO/IEC 27001:2013, and SOC 2. Evidence rows specify the artifacts required to substantiate each control.

POLICY

The governing document that authorizes the control.

CONTROL

The specific mechanism or process that implements the policy.

EVIDENCE

The artifacts that prove the control is operating.

MAPPING

The regulatory framework clauses this control satisfies.

EXECUTIVE CONTEXT

Why This Catalog Exists

Boards are now directly accountable for AI control failures. The gap between deployed AI systems and documented, auditable controls is where regulatory enforcement and legal liability concentrate. This catalog maps seven control domains to NIST AI RMF, ISO/IEC 42001, ISO/IEC 27001:2013, and SOC 2 so the organization proves governance, not just claims it.

7	4	28+
Control Domains	**Regulatory Frameworks**	**Mapped Controls**

OVERVIEW

7 Control Domains at a Glance

Enterprise AI controls mapped to NIST AI RMF · ISO/IEC 42001:2023 · ISO/IEC 27001:2013 · SOC 2

#	Domain	Key Control	Framework Coverage
01	Governance & Accountability	RACI; Quarterly Review	NIST AI RMF GOV · ISO/IEC 42001: 5.2 · SOC 2 CC1.2
02	Data & Privacy	PIAs/DPIAs; Approved Sources	NIST AI RMF MAP/GOV · ISO/IEC 27001 A.8 · SOC 2 CC6
03	Model Development & Testing	Independent review pre-deploy	NIST AI RMF MEA/MAP · ISO/IEC 42001 8.3/8.5 · SOC 2 CC7
04	Security & Access	RBAC; MFA; VA scans	NIST AI RMF PROTECT · ISO/IEC 27001 A.9/A.12 · SOC 2 CC6
05	Human Oversight & IR	Kill switch; Runbook	NIST AI RMF GOVERN/MAP/MEA · ISO/IEC 42001 8.6 · SOC 2 CC7.4
06	Monitoring & Post-Deployment	Scheduled reviews; Alerts	NIST AI RMF MEA · ISO/IEC 42001 8.7 · SOC 2 CC7
07	Third-Party / Vendor AI	Vendor list; Contract clauses	NIST AI RMF GOVERN/MAP · ISO/IEC 27001 A.15 · SOC 2 CC3

Source: NIST AI RMF · ISO/IEC 42001:2023 · ISO/IEC 27001:2013 · AICPA SOC 2.

DOMAIN 01 OF 07 FOUNDATION DOMAIN

Governance & Accountability

Establishing accountability structures, policies, and oversight mechanisms for AI systems.

Risk Tier:	High. Absence of governance ownership is the primary trigger for regulatory enforcement action.
Policy	AI Governance Policy
Control	RACI matrix; Quarterly Board Review
Evidence	Board minutes; Risk Register; RACI documentation
Mapping	**NIST AI RMF GOV · ISO/IEC 42001: 5.2 · SOC 2 CC1.2**
If you skip this:	No named owner means no defensible answer when a regulator asks who approved the model.

WHY IT MATTERS

Without a named owner and a documented review cycle, no one is accountable when a model causes harm. Regulators do not accept 'we were unaware' as a defense.

KEY ARTIFACTS

- AI Governance Policy document
- RACI matrix (approved)
- Quarterly review minutes
- Enterprise risk register

DOMAIN 02 OF 07 FOUNDATION DOMAIN

Data & Privacy

Governing data quality, lineage, and privacy protections across the AI data lifecycle.

Risk Tier:	High. Unlawful data processing exposes the organization to regulatory fines and invalidates model outputs as evidence.
Policy	Data Governance & Privacy-by-Design
Control	PIAs/DPIAs; Approved Data Sources list
Evidence	DPIA reports; Data access logs; Source approval records
Mapping	**NIST AI RMF MAP/GOV · ISO/IEC 27001 A.8 · SOC 2 CC6**
If you skip this:	The organization cannot demonstrate that personal data used in training was lawfully processed or that privacy risks were assessed before use.

WHY IT MATTERS

AI models are only as trustworthy as the data they consume. This domain enforces privacy-by-design principles, mandates impact assessments before data use, and restricts training data to approved, auditable sources.

KEY ARTIFACTS

- Data Governance Policy
- DPIA / PIA reports
- Approved data source registry
- Access control logs

DOMAIN 03 OF 07 FOUNDATION DOMAIN

Model Development & Testing

Ensuring AI models are rigorously validated, documented, and independently reviewed before deployment.

Risk Tier:	High. Unvalidated models in production are the most common source of AI-related audit findings and liability claims.
Policy	Model Risk Management Policy
Control	Independent validation review pre-deployment; Model cards required
Evidence	Testing logs; Model cards; Validation sign-off records
Mapping	**NIST AI RMF MEA/MAP · ISO/IEC 42001 8.3/8.5 · SOC 2 CC7**
If you skip this:	The organization cannot prove a model was independently validated before deployment, which is the first question any regulator or auditor will ask.

WHY IT MATTERS

A model that has not been independently validated before deployment is a liability, not an asset. If it fails, the organization cannot show it exercised reasonable care. That is the standard regulators and plaintiffs apply.

KEY ARTIFACTS

- Model Risk Management Policy
- Model cards (per model)
- Independent validation reports
- Pre-deployment test logs

TRANSITION

From Standards to Operations

Foundation domains set the policy and model standards that define what responsible AI looks like inside the organization; operational domains prove those standards hold under real conditions of access, oversight, monitoring, and vendor risk. Together the seven domains form one auditable control stack, not seven separate checklists.

FOUNDATION DOMAINS

Governance & Accountability · Data & Privacy · Model Development & Testing

OPERATIONAL DOMAINS

Security & Access · Human Oversight & IR · Monitoring & Post-Deployment · Third-Party & Vendor AI

DOMAIN 04 OF 07 OPERATIONAL DOMAIN

Security & Access

Protecting AI systems and data through robust access controls, secure development, and vulnerability management.

Risk Tier:	High. Unauthorized access to AI systems and training data is a direct SOC 2 and ISO/IEC 27001 audit failure point.
Policy	Secure Development & Access Control Policy
Control	RBAC; MFA enforcement; Vulnerability assessment scans
Evidence	Access review reports; VA scan results; Penetration test records
Mapping	**NIST AI RMF PROTECT · ISO/IEC 27001 A.9/A.12 · SOC 2 CC6**
If you skip this:	The organization cannot demonstrate that access to AI systems and training data was restricted, logged, and reviewed on a defined schedule.

WHY IT MATTERS

AI systems are high-value targets. A compromised model or training dataset can be manipulated without detection. This domain ensures that who can access what is defined, logged, and reviewed on a fixed schedule. Not discovered after a breach.

KEY ARTIFACTS

- Access control policy & RBAC matrix
- MFA enforcement logs
- Quarterly VA scan reports
- Penetration test results

DOMAIN 05 OF 07 OPERATIONAL DOMAIN

Human Oversight & Incident Response

Maintaining human control over AI decisions and ensuring rapid, structured response to AI-related incidents.

Risk Tier:	High. Absence of a tested override mechanism is cited in AI incident investigations as evidence of negligent deployment.
Policy	Human-in-the-Loop Policy; Incident Response Plan
Control	Kill switch / override mechanism; IR runbook; Escalation thresholds
Evidence	IR tickets; Post-mortem reports; Override activation logs
Mapping	**NIST AI RMF GOVERN/MAP/MEA · ISO/IEC 42001 8.6 · SOC 2 CC7.4**
If you skip this:	The organization cannot show it had a tested mechanism to stop a failing AI system or a documented process for responding when one caused harm.

WHY IT MATTERS

AI systems fail. The question regulators ask is not whether a failure occurred but whether the organization had a tested mechanism to stop it and a documented process to respond. This domain provides both.

KEY ARTIFACTS

- Human-in-the-Loop Policy
- IR runbook (current version)
- Kill switch test records
- Post-mortem reports

DOMAIN 06 OF 07 OPERATIONAL DOMAIN

Monitoring & Post-Deployment

Continuously tracking AI model performance, drift, and behavioral changes after production deployment.

Risk Tier:	High. Regulators treat post-deployment monitoring failures as evidence that governance was performative, not operational.
Policy	Continuous Monitoring & Model Drift Policy
Control	Scheduled performance reviews; Automated drift alerts; Threshold-based triggers
Evidence	Monitoring dashboards; Alert tickets; Drift detection reports
Mapping	**NIST AI RMF MEA · ISO/IEC 42001 8.7 · SOC 2 CC7**
If you skip this:	The organization cannot prove it detected or acted on model degradation, bias drift, or behavioral change after a system went live.

WHY IT MATTERS

A model approved at deployment is not the same model six months later. Data shifts, user behavior changes, and outputs drift. Without scheduled reviews and automated alerts, degradation goes undetected until it becomes a regulatory finding or a public failure.

KEY ARTIFACTS

- Monitoring & Drift Policy
- Live performance dashboards
- Drift detection alert logs
- Scheduled review records

DOMAIN 07 OF 07 OPERATIONAL DOMAIN

Third-Party & Vendor AI

Managing AI risk introduced through external vendors, APIs, and third-party model integrations.

Risk Tier:	High. Most enterprise AI vendors process personal data or make automated decisions. Contractual controls are necessary but insufficient without active assessment.
Policy	Vendor AI Risk Management Policy
Control	Approved vendor list; AI-specific contract clauses; Annual vendor assessments
Evidence	Vendor risk assessments; DPA agreements; SLA documentation
Mapping	**NIST AI RMF GOVERN/MAP · ISO/IEC 27001 A.15 · SOC 2 CC3**
If you skip this:	The organization cannot demonstrate that third-party AI systems were assessed for risk before use or that vendors are contractually bound to the same governance standards.

WHY IT MATTERS

When a vendor's AI system causes harm, the contracting organization shares liability. Signing a DPA is not due diligence. This domain requires active assessment before onboarding and annual reviews to confirm the vendor's controls remain in place.

KEY ARTIFACTS

- Vendor AI Risk Management Policy
- Approved vendor registry
- Data Processing Agreements (DPAs)
- Annual vendor assessment reports

NEXT STEP

Benchmark Your Control Stack

Ethical Tech Matters (ETM) assesses an organization's seven-domain control coverage against NIST AI RMF, ISO/IEC 42001, ISO/IEC 27001:2013, and SOC 2 and delivers a board-ready gap roadmap with prioritized remediation actions.

Assessment led by an IEEE CertifAIEd Lead Assessor

MODEL YOUR EXPOSURE

Quantify your AI governance risk before the assessment.

roi.ethicaltechmatters.com

SCHEDULE ASSESSMENT

Book a scoping call with our lead assessor.

Book a Scoping Call