

Your AI is regulated. Your governance probably is not.

Six global regimes are now enforceable. Ten sectors have overlays of their own. This is the operating map the board, the GC, and the AI Officer should be reading from. Not the press release version.

The governance gap costs more than the technology.

If you ship AI to customers, employees, or regulated decisions, you are already inside the perimeter of at least three of the regimes in this document. Most boards do not know that. Most legal teams know in the abstract but cannot point at the artifact a regulator would actually ask for.

This playbook collapses six global frameworks and ten sector overlays into the obligations, the evidence, and the questions worth asking at the next risk committee.

"The risk is not that you fail an audit. The risk is that you cannot prove you ran one."

What this document is

A field reference. Not a legal opinion. Cite primary regulator guidance for binding interpretation, and use this as the orientation layer above it.

Who should read it

Chief AI Officers, General Counsel, Chief Risk Officers, audit-committee chairs, and the operators who have to assemble the evidence file three days before an inspection.

One operating picture. Six rule books.

Each frame below carries a different enforcement model and a different evidence expectation. Treat them as overlapping, not interchangeable.

01 EU AI Act

EU

Any AI system placed on the EU market or whose output affects people in the EU. Extraterritorial. Applies whether your headquarters is in the EU or not.

CRITICAL

02 NIST AI RMF

US-FEDERAL

Voluntary framework, but the de facto standard for US federal procurement, federal-adjacent buyers, and any vendor security questionnaire that mentions AI. Increasingly cited in enterprise...

HIGH

03 ISO/IEC 42001

GLOBAL

Certifiable management system standard for organizations that develop, provide, or use AI. Parallel structure to ISO 27001. Increasingly demanded by enterprise procurement.

HIGH

04 UK AI Framework

UK

Sectoral, principles-based, regulator-led. Five cross-cutting principles enforced by existing regulators (ICO, FCA, MHRA, CMA, Ofcom). Statutory legislation expected mid-decade.

ELEVATED

05 Singapore Model AI

SINGAPORE

Voluntary but highly influential in APAC. Adopted as a reference by Hong Kong, Thailand, the Philippines. AI Verify is the world's first AI governance testing framework and toolkit.

MODERATE

06 Canada AIDA

CANADA

Federal regime focused on high-impact AI systems. OSFI Guideline E-23 (model risk management) already in force for federally regulated financial institutions and applies to AI/ML models.

ELEVATED

EU AI Act and NIST AI RMF.

EU AI Act European Union Artificial Intelligence Act EUROPEAN UNION, 27 MEMBER STATES **CRITICAL**

Scope. Any AI system placed on the EU market or whose output affects people in the EU. Extraterritorial. Applies whether your headquarters is in the EU or not.

Core obligations. Risk classification and Annex III determination. Conformity assessment before market entry for high-risk systems. Technical documentation, post-market monitoring, incident reporting within 15 days.

Evidence regulators expect. Risk management system documentation. Data governance procedures (training, validation, testing). Quality management system records.

Timeline. Phased: prohibited practices Feb 2025, GPAI obligations Aug 2025, high-risk systems Aug 2026, full enforcement Aug 2027.

Penalty exposure. Up to €35M or 7% of global annual turnover for prohibited practices. Up to €15M or 3% for high-risk violations. Up to €7.5M or 1.5% for incorrect information to authorities.

NIST AI RMF NIST AI Risk Management Framework 1.0 + AI 600-1 GenAI Profile UNITED STATES, FEDERAL GUIDANCE **HIGH**

Scope. Voluntary framework, but the de facto standard for US federal procurement, federal-adjacent buyers, and any vendor security questionnaire that mentions AI. Increasingly cited in enterprise contracts.

Core obligations. Implement the four functions: Govern, Map, Measure, Manage. Cover 19 categories and approximately 72 subcategories. For generative AI: address the 12 risk categories in AI 600-1 (hallucination, data poisoning, IP infringement, etc.).

Evidence regulators expect. Per-subcategory audit artifacts. Risk register tied to business impact. TEVV (testing, evaluation, verification, validation) procedures.

Timeline. AI RMF 1.0 published Jan 2023. AI 600-1 GenAI Profile published Jul 2024.

Penalty exposure. No direct fines. But failure to align is increasingly cited in enterprise contract rejections, board liability disputes, and post-incident litigation. NIST alignment is treated as the duty-of-care baseline.

ISO/IEC 42001 and UK AI Framework.

ISO/IEC 42001 ISO/IEC 42001:2023 AI Management System

INTERNATIONAL, CERTIFIABLE

HIGH

Scope. Certifiable management system standard for organizations that develop, provide, or use AI. Parallel structure to ISO 27001. Increasingly demanded by enterprise procurement.

Core obligations. Establish an AI Management System with leadership commitment. AI impact assessments tied to organizational context. Operational controls across the AI lifecycle.

Evidence regulators expect. AI management system manual. Statement of Applicability with selected controls. Risk treatment plan and residual risk acceptance.

Timeline. Published December 2023. First certifications issued through 2024 to 2025.

Penalty exposure. No fines. But absence increasingly blocks enterprise deals, particularly in financial services, healthcare, government, and any sector where the procurement team is mature.

UK AI Framework UK Pro-Innovation AI Regulation Framework UNITED KINGDOM

ELEVATED

Scope. Sectoral, principles-based, regulator-led. Five cross-cutting principles enforced by existing regulators (ICO, FCA, MHRA, CMA, Ofcom). Statutory legislation expected mid-decade.

Core obligations. Apply the five principles: safety/security/robustness, transparency/explainability, fairness, accountability/governance, contestability/redress. Comply with sector regulator AI guidance. Consider AI Safety Institute model evaluation participation for frontier models.

Evidence regulators expect. Sector-specific compliance documentation. Algorithmic transparency records. Data protection impact assessments where automated decisions affect individuals.

Timeline. White Paper March 2023. Government response Feb 2024. AI Safety Institute operational. Statutory regime under consultation 2025 to 2026.

Penalty exposure. Penalties levied through existing regulators. ICO up to £17.5M or 4% of global turnover. FCA unlimited fines. MHRA criminal penalties for medical devices including AI as SaMD.

Singapore Model AI and Canada AIDA.

Singapore Model AI Singapore Model AI Governance Framework + SINGAPORE

AI Verify

MODERATE

Scope. Voluntary but highly influential in APAC. Adopted as a reference by Hong Kong, Thailand, the Philippines. AI Verify is the world's first AI governance testing framework and toolkit.

Core obligations. Apply the four-area framework: internal governance, decision making, operations management, stakeholder communication. For generative AI: address the nine dimensions in the Generative AI Framework (accountability, data, trusted development, incident reporting, testing, security, content provenance, safety/alignment, AI for public good). Optionally run AI Verify tests for fairness, robustness, explainability.

Evidence regulators expect. Governance structure documentation. Risk assessment and treatment records. AI Verify test reports where applicable.

Timeline. Model Framework v2 in 2020. Generative AI Framework in 2024. AI Verify open-source toolkit operational.

Penalty exposure. No direct fines. But MAS (financial sector), PDPC (data protection), and sectoral regulators can act through existing regimes. Penalties up to S\$1M under PDPA.

Canada AIDA Canada Artificial Intelligence and Data Act (AIDA, CANADA, FEDERAL

Bill C-27)

ELEVATED

Scope. Federal regime focused on high-impact AI systems. OSFI Guideline E-23 (model risk management) already in force for federally regulated financial institutions and applies to AI/ML models.

Core obligations. Anticipate AIDA-style obligations: identify high-impact systems, mitigate harms, monitor compliance, publish plain-language descriptions. For financial institutions: OSFI E-23 model risk governance, model inventory, validation, monitoring. PIPEDA transparency requirements for automated decisions.

Evidence regulators expect. Model inventory with risk tiering. Validation and monitoring reports. Plain-language descriptions of high-impact systems.

Timeline. Bill C-27 introduced June 2022. Bill died on order paper at 2025 election dissolution. Likely re-introduction in revised form. OSFI Guideline E-23 on model risk in effect for federally regulated financial institutions.

Penalty exposure. Under proposed AIDA: up to CAD\$25M or 5% of global revenue for serious violations. OSFI can impose administrative monetary penalties and supervisory actions.

When the rule books start grading on a curve.

Most regulators publish well before they enforce. The phase from publication to first inspection is the window most companies waste. Use the timeline below to back-plan governance milestones.

PHASED: PROHIBITED PRACTICES FEB 2025, GPAI OBLIGATIONS AUG 2025, HIGH-RISK SYSTEMS AUG 2026, FULL ENFORCEMENT AUG 2027.

EU AI Act

Any AI system placed on the EU market or whose output affects people in the EU. Extraterritorial. Applies whether your headquarters is in the EU or not.

AI RMF 1.0 PUBLISHED JAN 2023. AI 600-1 GENAI PROFILE PUBLISHED JUL 2024.

NIST AI RMF

Voluntary framework, but the de facto standard for US federal procurement, federal-adjacent buyers, and any vendor security questionnaire that mentions AI. Inc...

PUBLISHED DECEMBER 2023. FIRST CERTIFICATIONS ISSUED THROUGH 2024 TO 2025.

ISO/IEC 42001

Certifiable management system standard for organizations that develop, provide, or use AI. Parallel structure to ISO 27001. Increasingly demanded by enterprise...

WHITE PAPER MARCH 2023. GOVERNMENT RESPONSE FEB 2024. AI SAFETY INSTITUTE OPERATIONAL. STATUTORY REGIME UNDER CONSULTATION 2025 TO 2026.

UK AI Framework

Sectoral, principles-based, regulator-led. Five cross-cutting principles enforced by existing regulators (ICO, FCA, MHRA, CMA, Ofcom). Statutory legislation ex...

MODEL FRAMEWORK V2 IN 2020. GENERATIVE AI FRAMEWORK IN 2024. AI VERIFY OPEN-SOURCE TOOLKIT OPERATIONAL.

Singapore Model AI

Voluntary but highly influential in APAC. Adopted as a reference by Hong Kong, Thailand, the Philippines. AI Verify is the world's first AI governance testing...

BILL C-27 INTRODUCED JUNE 2022. BILL DIED ON ORDER PAPER AT 2025 ELECTION DISSOLUTION. LIKELY RE-INTRODUCTION IN REVISED FORM. OSFI GUIDELINE E-23 ON MODEL RISK IN EFFECT FOR FEDERALLY REGULATED FINANCIAL INSTITUTIONS.

Canada AIDA

Federal regime focused on high-impact AI systems. OSFI Guideline E-23 (model risk management) already in force for federally regulated financial institutions a...

Three questions worth asking before any regulator asks them for you.

1

Can you produce a current inventory of every AI system in production, including shadow pilots and vendor agents, in under 48 hours?

2

For your three highest-risk AI use cases, can you name the human accountable for each decision the model influences?

3

If a regulator asked for the impact assessment behind your most consequential AI deployment, would it exist, and would it survive cross-examination?

The eight-artifact evidence file.

If you can produce these eight artifacts on demand, you are ahead of 80% of mid-market peers across every regime in this playbook.

01 AI system inventory

Every model, agent, and vendor system with owner and risk tier.

02 Impact assessment library

Pre-deployment assessments for any consequential decision.

03 Model and data documentation

Model cards, datasheets, provenance, evaluation results.

04 Human oversight protocols

Who can intercept, override, escalate. Documented and tested.

05 Incident response runbook

Detection, containment, notification, and post-mortem procedures.

06 Bias and fairness evaluations

Pre-deployment and recurring, with mitigation actions.

Ten industries. Different overlays. Same operating math.

Every sector below carries its own enforcement story on top of the six global regimes. Companion industry playbooks treat each sector in depth.

Healthcare & Life Sciences

AI in healthcare is regulated as a medical device, a privacy risk, and a clinical-safety question. All three at once.

OVERLAYS

FDA SAMD

HIPAA AI

EU MDR/IVDR

US STATE AI HEALTH

Financial Services & Banking

Model risk is not a new idea in banking. AI is forcing every existing framework to extend.

OVERLAYS

SR 11-7

NYDFS AI

CFPB AI CREDIT

BASEL + ECB AI

Manufacturing & Industrial

Operational AI on the factory floor is regulated as a safety system, a product, and a workforce decision.

OVERLAYS

ISO/IEC 23894

EU MACHINERY REG

OSHA AI

SUPPLY CHAIN AI

Retail & E-commerce

Personalization, dynamic pricing, and recommendation engines are now regulated as consumer-protection issues.

OVERLAYS

EU DSA

FTC AI ENFORCEMENT

CPRA ADM

DYNAMIC PRICING

Transportation & Logistics

Autonomy, route AI, and driver monitoring are regulated as safety, labor, and infrastructure decisions all at once.

OVERLAYS

NHTSA AV

FMCSA AI

AVIATION AI

IMO MARITIME AI

Energy & Utilities

AI on the grid is regulated as critical infrastructure. The threat model includes nation-state actors.

OVERLAYS

NERC CIP

FERC AI

CSRD AI

NRC AI

Technology & Software

Software vendors are the supply chain. Your customers are downloading your AI risk

Agriculture & AgTech

Precision agronomy and autonomous equipment AI is regulated where it

What to do before the next risk committee meeting.

DAYS 1 TO 30

Diagnose

MAP THE PERIMETER

- Inventory every AI system in production, including shadow pilots
- Rank by consequence, not by sophistication
- Identify regime exposure (EU AI Act, NIST RMF, sector overlays)
- Brief the audit committee on the gap

DAYS 31 TO 60

Govern

BUILD THE OPERATING LAYER

- Stand up the eight-artifact evidence file
- Assign human accountability to every consequential decision
- Adopt impact-assessment template for new deployments
- Run a tabletop incident response on one production system

DAYS 61 TO 90

Operate

MAKE IT ROUTINE

- Quarterly inventory refresh becomes business as usual
- Vendor AI questionnaire enforced at procurement
- Red-team or fairness review on top-three risk tier
- Publish internal governance posture for sales enablement

Take this to your next board meeting.

If you want to pressure-test your governance posture against the six regimes and your sector overlay, the Diagnostic Sprint compresses the entire exercise into two weeks with a board-ready output.

ENGAGE

sophizo.net/checkout/diagnostic-sprint

EMAIL

john@sophizo.net

Sources. European Commission AI Act Implementation Office; NIST AI Risk Management Framework 1.0; ISO/IEC 42001:2023; UK DSIT AI Regulation Policy Paper; OECD AI Policy Observatory; US State legislative trackers (Colorado SB 24-205, California AB 2013, Texas TRAIGA, NYC LL144). Sophizo Research, 2026.