AI Governance Frameworks: NIST AI RMF and ISO/IEC 42001

A practical guide to AI risk management, governance, implementation, and continuous improvement for organizations adopting artificial intelligence responsibly.

1. NIST AI RMF

The NIST AI Risk Management Framework, commonly called the NIST AI RMF, is a voluntary framework designed to help organizations identify, assess, prioritize, and manage risks connected with artificial intelligence systems. It is especially useful because it focuses on trustworthy AI, including safety, reliability, accountability, transparency, privacy, fairness, and resilience. The framework organizes AI risk management into four core functions: Govern, Map, Measure, and Manage. These functions are not intended to be a rigid checklist; instead, they create a repeatable operating model for responsible AI across the lifecycle.

1.1 Govern

Govern is the foundation of the NIST AI RMF. It defines how an organization establishes accountability, policies, roles, responsibilities, and oversight for AI risk management. Governance should not sit only with a technical team; it should involve leadership, legal, compliance, risk, security, data, product, and business stakeholders. A strong governance approach makes it clear who approves AI use cases, who owns the risks, who monitors performance, and who responds when something goes wrong.

- Define an AI policy that explains acceptable and prohibited AI use.
- Create an AI inventory that lists models, tools, vendors, owners, data sources, and business purposes.

- Assign clear accountability for risk decisions, approvals, monitoring, and escalation.
- Train employees on responsible AI use, data protection, bias awareness, and human oversight.
- Establish periodic review forums, such as an AI governance board or model risk committee.

Example: A bank planning to use an AI-powered chatbot for customer support may require approval from compliance, information security, legal, and customer experience teams before launch. The governance process may require the team to document the chatbot's purpose, approved knowledge sources, escalation rules, customer disclosure wording, and monitoring metrics.

1.2 Map

Map focuses on understanding the context in which an AI system will operate. This includes the business purpose, users, affected stakeholders, data inputs, operating environment, legal obligations, and potential impacts. Mapping helps organizations avoid treating AI risk as purely technical. It asks: What decision or process will AI influence? Who may be affected? What could go wrong? What assumptions are being made?

- Document the intended purpose and limitations of the AI system.
- Identify stakeholders such as users, customers, employees, regulators, and impacted communities.

- Classify the AI use case by risk level, such as low, medium, high, or unacceptable.
- Review data sources for quality, consent, representativeness, and sensitivity.
- Identify dependencies, including third-party models, cloud services, APIs, and human review processes.

Example: For an AI system used in loan application screening, mapping would include identifying applicants, loan officers, compliance teams, credit policy owners, regulators, and customer service teams. The organization would map the data used by the model, such as income, employment history, repayment history, and credit bureau data, and then assess whether these inputs could create unfair outcomes for certain groups.

1.3 Measure

Measure is about evaluating AI risks and system performance using qualitative and quantitative methods. It includes testing, validation, benchmarking, monitoring, and independent review. Measurement should happen before deployment and continue after deployment because AI system behavior can change over time due to changes in data, user behavior, business rules, or the external environment.

- Test model accuracy, reliability, robustness, and failure rates.
- Measure fairness and bias across relevant population groups where legally and ethically appropriate.

- Evaluate explainability so decision-makers can understand key factors behind AI outputs.
- Run security and resilience testing, including misuse scenarios and adversarial inputs.
- Track operational metrics such as false positives, false negatives, override rates, complaint rates, and escalation volumes.

Example: If a fraud detection model flags transactions as suspicious, the organization should measure how often legitimate transactions are incorrectly blocked. A high false-positive rate may frustrate customers, increase call center volume, and create reputational risk even if the model appears technically accurate.

1.4 Manage

Manage converts risk understanding into action. It involves prioritizing risks, selecting treatment options, implementing controls, monitoring residual risk, and deciding whether an AI system should be approved, changed, restricted, or retired. Management activities should be risk-based, meaning higher-risk AI systems receive stronger oversight, more testing, and more frequent review.

- Choose risk responses such as avoid, reduce, transfer, accept, or monitor.
- Implement controls such as human review, audit logs, access controls, approval workflows, and fallback procedures.

- Define incident response processes for harmful outputs, security issues, or customer complaints.
- Set thresholds for retraining, rollback, suspension, or decommissioning.
- Maintain evidence of decisions, reviews, approvals, and corrective actions.

Example: A recruitment screening tool may be approved only if it supports human decision-making rather than replacing it. The organization may require recruiters to review all AI-generated shortlists, document final selection reasons, and monitor whether candidate outcomes show signs of unfair filtering.

2. ISO/IEC 42001

ISO/IEC 42001 is an international management system standard for artificial intelligence. It provides requirements for establishing, implementing, maintaining, and continually improving an Artificial Intelligence Management System, often called an AIMS. Unlike a high-level guidance framework alone, ISO/IEC 42001 is structured as a management system standard, which means organizations can use it to build documented processes, assign responsibilities, evaluate performance, and support certification or audit readiness.

2.1 AI Management System

An AI Management System is a structured set of policies, procedures, controls, roles, records, and review mechanisms used to govern AI across an organization. It is not simply a software platform or a technical model registry. It is the operating system for responsible AI management, connecting strategy, risk, compliance, technical controls, human oversight, and business objectives.

- Defines the organization's AI policy and risk appetite.
- Clarifies roles for AI owners, developers, users, risk teams, and senior leadership.
- Requires AI risk assessment and AI risk treatment processes.
- Supports documentation, evidence, monitoring, internal audit, and management review.

- Encourages alignment with business objectives, legal obligations, stakeholder expectations, and responsible AI principles.

Example: A healthcare organization using AI for appointment prioritization would define the scope of its AI Management System, identify patient safety and fairness risks, assign ownership to clinical, technology, and compliance leaders, and maintain records showing how the AI system was assessed, approved, monitored, and improved.

2.2 Implementation

Implementation of ISO/IEC 42001 should be treated as an organizational change program, not only a compliance exercise. A practical implementation approach begins with defining the scope, understanding the organization's AI context, identifying interested parties, creating an AI policy, assessing risks, selecting controls, and establishing performance monitoring. The organization should also decide how the AI Management System will integrate with existing systems such as information security, privacy, quality, enterprise risk management, vendor management, and internal audit.

- **Step 1: Define scope.** Decide which AI systems, business units, geographies, vendors, and use cases are covered.
- **Step 2: Establish leadership commitment.** Senior management should approve the AI policy, provide resources, and review performance.
- **Step 3: Build an AI inventory.** Record each AI system's purpose, owner, data sources, model type, vendor dependency, risk level, and lifecycle status.

- **Step 4: Conduct AI risk assessment.** Identify risks related to safety, bias, privacy, security, explainability, misuse, reliability, and legal compliance.
- **Step 5: Treat risks with controls.** Apply controls such as human oversight, access restrictions, testing, monitoring, documentation, and supplier due diligence.
- **Step 6: Maintain evidence.** Keep records of assessments, approvals, testing results, incidents, corrective actions, and management reviews.

Example: A financial services firm deploying generative AI for internal research support may start by limiting use to approved employees, restricting sensitive client data entry, logging prompts and outputs where appropriate, training staff on acceptable use, and reviewing the tool's performance and risks quarterly.

2.3 Continuous Improvement

Continuous improvement is central to ISO/IEC 42001 because AI systems and their risks change over time. A model that performs well at launch may degrade as data patterns change, business processes evolve, regulations mature, or users discover new ways to interact with the system. Continuous improvement ensures that AI governance remains current, evidence-based, and responsive.

- Monitor AI objectives, risk indicators, incidents, and audit findings.
- Conduct internal audits to check whether AI processes are followed and effective.

- Hold management reviews to evaluate performance, resource needs, risks, and improvement actions.
- Investigate nonconformities such as undocumented AI use, failed controls, or missing approvals.
- Apply corrective actions and update policies, procedures, controls, and training materials.

Example: If monitoring shows that an AI customer support assistant frequently gives incomplete answers for refund requests, the organization may update the knowledge base, improve guardrails, retrain reviewers, adjust escalation rules, and retest the assistant before expanding its use.

3. EU AI Act

The EU AI Act is a binding regulation that applies a risk-based approach to artificial intelligence. Unlike voluntary guidance, it creates legal obligations for AI providers, deployers, importers, distributors, and other actors depending on their role and the type of AI system involved. Its core purpose is to encourage trustworthy AI while protecting health, safety, fundamental rights, transparency, and accountability. For organizations, the EU AI Act turns AI governance from a good practice into a legal compliance requirement when AI systems fall within its scope.

3.1 Risk Classification

Risk classification is the starting point for EU AI Act compliance. The Act broadly separates AI systems into risk categories so that obligations increase as potential harm increases. A simple productivity assistant may require little or no mandatory compliance effort, while an AI system used in employment, education, law enforcement, credit access, biometric identification, or essential services may trigger much stricter requirements.

- **Unacceptable risk:** AI practices considered a clear threat to safety, rights, or democratic values are prohibited. Examples include certain manipulative systems, social scoring, and some biometric identification practices.
- **High risk:** AI systems used in sensitive areas are permitted but subject to strict controls, documentation, monitoring, human oversight, and conformity

assessment. Examples include AI used for recruitment screening, education assessment, creditworthiness evaluation, or access to essential public services.

- **Limited risk:** These systems are mainly subject to transparency obligations. For example, users may need to be informed when they are interacting with a chatbot or when content is AI-generated.

- **Minimal risk:** Most low-impact AI systems fall into this category and may not carry specific mandatory obligations under the Act, though voluntary good practices are still encouraged.

Example: An AI spell checker used internally by employees is likely to be low risk, while an AI tool that automatically ranks job candidates for interview shortlisting may be high risk because it can materially affect access to employment opportunities.

3.2 Key Obligations

The specific obligations depend on whether the organization is acting as a provider, deployer, importer, distributor, or another covered party. In simple terms, providers usually carry heavier obligations because they design, develop, place, or substantially modify AI systems. Deployers must still use AI systems responsibly, follow instructions, maintain human oversight where required, and monitor operational use.

- Perform risk classification before deployment or market placement.

- Maintain technical documentation explaining intended purpose, design, data, testing, limitations, and controls.

- Implement risk management and quality management processes for high-risk AI systems.
- Ensure appropriate data governance, including relevance, quality, and representativeness of data.
- Enable logging, traceability, monitoring, and post-market surveillance where required.
- Provide human oversight so people can understand, supervise, challenge, or stop AI-supported processes.
- Meet transparency obligations, such as informing users when they interact with AI or when synthetic content is generated.
- Report serious incidents and take corrective action when AI systems create unacceptable or unmanaged risk.

Example: If a company deploys a high-risk AI system to support credit decisions, it should ensure that responsible staff understand how to use the system, that input data is relevant, that outputs are monitored, and that customers can be escalated to a human review process where appropriate.

3.3 Compliance

Compliance should begin with a structured inventory of AI systems and a clear understanding of organizational roles. Many organizations make the mistake of treating

AI compliance as a one-time legal review. A stronger approach is to build repeatable controls that connect legal, risk, technology, procurement, data protection, cybersecurity, internal audit, and business ownership.

- Create and maintain an AI system inventory.
- Classify each AI system by risk category and document the rationale.
- Identify the organization's role for each system, such as provider or deployer.
- Perform gap assessments against applicable EU AI Act obligations.
- Define control owners for documentation, testing, monitoring, human oversight, incident response, and supplier assurance.
- Prepare evidence for audits, regulatory inquiries, and management reviews.
- Monitor regulatory timelines and update compliance plans as obligations become applicable.

Example: A company using AI in hiring may build a compliance file containing the system purpose, vendor documentation, risk classification, data governance review, bias testing evidence, human oversight procedure, candidate communication process, and incident escalation workflow.

4. Framework Comparison

NIST AI RMF, ISO/IEC 42001, and the EU AI Act should not be viewed as competing approaches. They serve different purposes. NIST AI RMF helps organizations think through AI risks and trustworthiness. ISO/IEC 42001 helps organizations build a formal management system around AI. The EU AI Act imposes legal obligations for AI systems that fall within its scope. Together, they create a practical pathway from principles to operating controls to regulatory compliance.

4.1 NIST AI RMF vs ISO/IEC 42001 vs EU AI Act

Aspect	NIST AI RMF	ISO/IEC 42001	EU AI Act
Nature	Voluntary risk management framework	International management system standard	Binding legal regulation
Main purpose	Manage AI risks and improve trustworthiness	Establish, operate, and improve an AI Management System	Regulate AI based on risk and protect rights, safety, and transparency
Best used for	Risk discovery, assessment, control	Policy, roles, processes, records,	Legal classification, obligations,

	design, and governance maturity	audits, and continuous improvement	conformity, and regulatory compliance
Output	Risk profiles, controls, metrics, and governance practices	Documented management system and audit evidence	Compliance evidence, disclosures, documentation, and required controls
Example use	Assessing risks of a fraud detection model	Creating an enterprise AI governance operating model	Determining whether a hiring AI tool is high risk and what obligations apply

4.2 Guidance vs Standard vs Regulation

The difference between guidance, standards, and regulation is important because each creates a different type of obligation. Guidance helps organizations decide what good practice looks like. A standard provides a structured and auditable management system. A regulation creates legal requirements that can be enforced by authorities.

- **Guidance:** Flexible and principle-based. It supports judgment and maturity but is usually not mandatory by itself.

- **Standard:** Structured and auditable. It helps organizations demonstrate that they have a repeatable management system.
- **Regulation:** Legally enforceable. It defines minimum obligations, deadlines, penalties, and supervisory expectations.

Example: A company may use NIST AI RMF to identify AI risks, ISO/IEC 42001 to institutionalize governance processes, and the EU AI Act to confirm which legal obligations must be met for AI systems used in the European market.

5. Building an Effective AI Governance Program

An effective AI governance program turns frameworks into day-to-day operating practices. It should help teams innovate safely, make risk decisions consistently, comply with legal requirements, and maintain evidence for leadership, auditors, customers, and regulators. The program should be practical enough for teams to use and strong enough to manage high-impact AI decisions.

5.1 Using the Frameworks Together

The strongest approach is to use the three frameworks together rather than choosing only one. NIST AI RMF can support risk thinking and control design. ISO/IEC 42001 can structure the management system, accountability, documentation, and continuous improvement cycle. The EU AI Act can define mandatory obligations for systems that are placed on the EU market or used in covered contexts.

- Use NIST AI RMF to ask the right risk questions.
- Use ISO/IEC 42001 to build the governance structure and evidence base.
- Use the EU AI Act to determine legal classification and mandatory obligations.
- Connect all three through one AI inventory, one risk assessment process, one approval workflow, and one monitoring dashboard.

Example: For an AI-enabled customer onboarding system, a governance team may first map the risk context using NIST AI RMF, then document roles and controls through

ISO/IEC 42001, and finally check whether EU AI Act obligations apply because the system affects access to financial services.

5.2 Key Governance Steps

A practical AI governance program should be simple enough for business teams to follow and robust enough for risk, legal, compliance, security, and audit teams to trust. The following steps provide a usable operating model.

- **Step 1: Establish AI governance ownership.** Define who approves AI policy, who owns the AI inventory, who performs risk reviews, and who escalates issues.
- **Step 2: Build and maintain an AI inventory.** Record the purpose, owner, vendor, model type, data sources, users, risk category, and lifecycle status of each AI system.
- **Step 3: Classify risk and obligations.** Assess each system against internal risk criteria, NIST AI RMF considerations, ISO/IEC 42001 requirements, and EU AI Act categories where applicable.
- **Step 4: Define approval gates.** Require formal review before development, pilot testing, production launch, major model changes, and retirement.
- **Step 5: Implement controls.** Apply human oversight, access control, testing, monitoring, documentation, vendor assurance, cybersecurity, and incident response controls.

- **Step 6: Train users and owners.** Provide role-based training for developers, business users, reviewers, approvers, and executives.

- **Step 7: Monitor and improve.** Track incidents, model performance, complaints, control failures, regulatory changes, and audit findings, then update the governance program continuously.

The NIST AI RMF, ISO/IEC 42001, and the EU AI Act together provide a complete AI governance toolkit. NIST AI RMF helps organizations understand and manage AI risk. ISO/IEC 42001 helps convert governance into a formal management system. The EU AI Act introduces binding regulatory obligations where AI systems affect safety, health, transparency, or fundamental rights. Organizations that combine these approaches can move beyond scattered AI policies and build a disciplined program with clear ownership, documented controls, measurable oversight, regulatory awareness, and continuous improvement. In practice, effective AI governance is not a one-time compliance project; it is an ongoing operating model for responsible innovation.

Conclusion

AI governance is becoming essential as organizations adopt AI across more business functions. Understanding frameworks such as NIST AI RMF, ISO/IEC 42001, and the EU AI Act can help organizations manage risks, meet compliance requirements, and build more responsible AI systems.

The right AI governance framework depends on your organization's needs, but combining practical governance practices with clear accountability, risk management, and continuous monitoring can create a stronger foundation for trustworthy AI.

GSDC
Global Skill Development Council

CERTIFIED AI GRC PROFESSIONAL

THE AI GOVERNANCE, RISK & COMPLIANCE (AI GRC) CERTIFICATION PROGRAM IS GLOBALLY DESIGNED TO STRENGTHEN AI GOVERNANCE, RISK MANAGEMENT, REGULATORY COMPLIANCE, AND RESPONSIBLE AI IMPLEMENTATION SKILLS, ENABLING PROFESSIONALS TO ESTABLISH TRUSTWORTHY, SECURE, AND COMPLIANT AI SYSTEMS WHILE ALIGNING INNOVATION WITH ORGANIZATIONAL AND REGULATORY REQUIREMENTS.

ABOUT GSDC CERTIFICATION

EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.

EBOOK

Extensive and exclusive Ebook created by world's experts to help you with understanding core concepts.

CREATED BY EXPERTS

GSDC certifications are created and authored by world's leading experts in the field.

LEARNING MATERIALS

Get access to learning materials such as videos, ebooks, templates, and practice exams, which will help you clear the certification exam.

LEARNING OBJECTIVE

- Prove your knowledge of responsible AI principles including fairness, transparency, explainability, and human oversight in AI systems
- Understand how to design and implement AI security controls, manage adversarial risks, and build trustworthy AI programs across the organization

Enroll now with the code **LEARN20** To avail **20%** discount

Enroll Now

www.gsdcouncil.org