TAKTIK TECHNOLOGIES LLC
AI Governance & Cybersecurity Advisory

NIST AI RISK MANAGEMENT FRAMEWORK

Compliance Summary & AI Governance Framework

Prepared by TakTik Technologies LLC | June 2026

> *This framework provides a practical guide to the NIST AI Risk Management Framework (AI RMF 1.0) and a structured, actionable governance approach for SMB and mid-market organizations developing, deploying, or using AI systems. Whether you are subject to the EU AI Act, U.S. federal AI expectations, or sector-specific regulation, the NIST AI RMF provides the governance foundation you need.*

1. What Is the NIST AI Risk Management Framework?

The NIST AI Risk Management Framework (AI RMF 1.0) was published by the U.S. National Institute of Standards and Technology in January 2023. It is the most widely adopted voluntary AI governance framework in the world, providing organizations of all sizes and sectors with a structured, evidence-driven approach to designing, developing, deploying, and monitoring AI systems responsibly.

Unlike prescriptive regulations such as the EU AI Act, the NIST AI RMF is voluntary and technology-agnostic. It does not mandate specific technical implementations but instead provides a flexible, scalable framework that organizations can tailor to their risk tolerance, sector, and operational context. This makes it equally applicable to a 50-person healthcare practice and a 5,000-person logistics company.

The framework was expanded significantly in 2024–2025 with companion materials including:

- The AI RMF Playbook — detailed suggested actions for each subcategory
- The Generative AI Profile (NIST AI 600-1, July 2024) — addressing 12 risk categories specific to large language models and foundation models, including confabulation, prompt injection, and data privacy
- Sector-specific profiles for healthcare, financial services, and critical infrastructure
- Closer alignment crosswalks with the EU AI Act, ISO 42001, and HIPAA

> *Key Distinction: The NIST AI RMF is not a compliance checklist. It is a governance operating model. Organizations that treat it as infrastructure — embedding it into existing risk management, IT governance, and procurement processes — gain both risk protection and competitive advantage.*

2. The Seven Characteristics of Trustworthy AI

The NIST AI RMF is anchored in seven characteristics that define what it means for an AI system to be trustworthy. These characteristics form the evaluation criteria for the MEASURE function and inform risk identification in MAP. Every AI system your organization uses should be assessed against all seven.

Characteristic	What It Means in Practice
Valid & Reliable	AI outputs are accurate, consistent, and perform as intended across diverse conditions and populations.
Safe	The AI system does not cause unacceptable physical, psychological, financial, or societal harm.
Secure & Resilient	The system withstands adversarial attacks, data poisoning, prompt injections, and unexpected failures.
Accountable & Transparent	Clear ownership of AI decisions exists; stakeholders can access meaningful information about system operations.
Explainable & Interpretable	Outputs can be understood and traced by users, affected individuals, and auditors in context-appropriate depth.
Privacy-Enhanced	The system minimizes personal data use, mitigates re-identification risk, and respects data subject rights.
Fair / Bias Managed	The system does not produce discriminatory outcomes or amplify harmful biases across demographic groups.

3. The Four Core Functions

The AI RMF Core is organized into four functions that work together as a continuous cycle — not a linear sequence. GOVERN is cross-cutting and infuses all other functions. MAP, MEASURE, and MANAGE apply iteratively across the AI system lifecycle.

GOVERN	MAP	MEASURE	MANAGE
Establishes organizational culture, accountability structures, policies, and oversight mechanisms for AI risk management. Cross-cutting — infused throughout MAP, MEASURE, and MANAGE.	Contextualizes AI risk by identifying the system's purpose, capabilities, stakeholders, and potential harms. Scopes what needs to be measured and managed.	Analyzes and tracks risks using quantitative and qualitative methods. Evaluates trustworthiness characteristics and determines whether controls are working.	Allocates resources to treat prioritized risks, defines response and recovery plans, handles residual risk, and drives continuous improvement.
6 categories | 19 subcategories	**5 categories | 23 subcategories**	**4 categories | 14 subcategories**	**4 categories | 16 subcategories**

3.1 GOVERN — Build the Foundation

GOVERN is the backbone of the entire framework. It establishes the organizational culture, structures, and accountability needed to make AI risk management repeatable and scalable. Six categories define its scope:

GOVERN 1 — Policies, Processes & Procedures

Organizational policies for AI risk management are established, communicated, and enforced. This includes defining what AI systems are permissible, who owns AI risk, and how decisions are escalated.

- Publish a formal AI Use Policy covering approved and prohibited use cases
- Define the organization's AI risk appetite in writing
- Establish accountability for AI risk at the C-suite or board level

GOVERN 2 — Accountability Structures

Roles and responsibilities for AI risk management are clearly assigned across the AI lifecycle — from procurement and development through deployment and retirement.

- Appoint an AI Risk Owner (maybe CISO, CTO, or COO for SMBs)

- Define who can approve deployment of new AI systems
- Establish an AI Oversight Committee with cross-functional representation

GOVERN 3 — Organizational Culture

AI risk management is embedded in organizational culture, not just policy documents. Teams are empowered to raise concerns about AI behavior without fear of reprisal.

GOVERN 4 — Workforce & Training

All staff involved in AI development, deployment, or oversight have appropriate AI literacy. Training is role-differentiated and kept current as AI capabilities evolve.

GOVERN 5 — Stakeholder Engagement

Affected communities, customers, regulators, and other stakeholders are identified and engaged. Feedback mechanisms allow external perspectives to inform AI governance decisions.

GOVERN 6 — Policies for Third-Party Risk

AI risks from vendors, partners, and open-source model providers are identified and governed. Contracts specify AI risk responsibilities. Third-party AI is included in the organizational risk register.

3.2 MAP — Establish Context & Identify Risks

MAP is the scoping function. For each AI system under consideration, MAP defines what it does, who it affects, what could go wrong, and what success looks like. Five categories structure this work:

MAP 1 — Context Establishment

Define the AI system's intended purpose, capabilities, and operational context. Document the business problem it solves, the decisions it influences, and the population it affects.

- Complete an AI System Context Card for every tool in the inventory
- Document intended users, affected populations, and use-case boundaries
- Identify potential dual-use or misuse scenarios

MAP 2 — Scientific & Technical Grounding

Evaluate whether the AI system is fit for purpose based on available evidence. Assess whether the underlying model has demonstrated validity for the intended use case and context.

MAP 3 — AI Risk Identification

Identify risks across all seven trustworthiness characteristics. Consider risks from the AI system's design, the data it uses, the context of deployment, and the humans who interact with it.

- Conduct structured risk identification workshops for each high-impact AI system
- Use NIST AI RMF Playbook risk taxonomies as a starting checklist
- Document risks to operators, end users, third parties, and society

MAP 4 — Organizational Risk Tolerance

Assess identified risks against the organization's risk appetite. Determine which risks are acceptable as-is, which require mitigation, and which would prevent deployment.

MAP 5 — Impacts Assessment

Evaluate the potential magnitude of harm — considering probability, severity, reversibility, and breadth of impact. Prioritize risks accordingly.

3.3 MEASURE — Analyze & Track Risks

MEASURE turns risk identification into evidence. It establishes how organizations will know whether their AI systems are performing as intended and whether risks are being controlled. Four categories guide this:

MEASURE 1 — Measurement Approach

Define the metrics, methods, and evaluation criteria that will be used to assess the AI system's trustworthiness characteristics. Distinguish between technical performance and real-world impact.

- Define success metrics for each of the seven trustworthiness characteristics
- Select quantitative (accuracy rates, bias metrics, drift indicators) and qualitative (user feedback, expert review) methods
- Establish baselines against which degradation will be measured

MEASURE 2 — AI System Evaluation

Conduct systematic evaluations of AI systems against defined metrics before deployment and on an ongoing basis. Document findings and remediation actions.

- Conduct pre-deployment testing across diverse data distributions
- Perform adversarial testing (red-teaming) for high-impact AI systems
- Evaluate fairness across demographic subgroups before production release

MEASURE 3 — Risk Tracking

Monitor AI system behavior in production. Track metrics over time and flag anomalies, drift, or unexpected behavior for investigation.

- Implement real-time or periodic monitoring dashboards for production AI systems
- Define alert thresholds that trigger human review
- Maintain audit logs of AI system inputs, outputs, and decisions

MEASURE 4 — Feedback & Improvement

Use measurement outputs to improve the AI system, refine risk assessments, and update governance policies. Close the loop between monitoring and governance.

3.4 MANAGE — Treat Risks & Drive Continuous Improvement

MANAGE is where risk identification and measurement translate into action. It allocates resources, defines treatment strategies, and ensures the organization can respond to AI incidents. Four categories:

MANAGE 1 — Risk Treatment Prioritization

Prioritize identified risks based on likelihood, impact, and strategic importance. Assign treatment owners and timelines. Document risk acceptance decisions with appropriate authorization.

- Maintain a living AI Risk Register with treatment status and ownership
- Apply risk treatment options: accept, mitigate, transfer (insurance/contracts), or avoid (discontinuation)
- Escalate unacceptable residual risks to executive decision-makers

MANAGE 2 — Risk Response Planning

Define and document specific actions to address prioritized risks. Implement technical controls, process changes, contractual protections, and oversight mechanisms.

- Implement model cards and system cards to document known limitations for deployers
- Apply data governance controls to address bias and data quality risks
- Establish human-in-the-loop requirements for high-stakes AI decisions

MANAGE 3 — Incident Response & Recovery

Establish AI-specific incident response procedures covering AI errors, bias events, adversarial attacks, data breaches, and regulatory notifications. Test these procedures regularly.

- Define AI incident classification tiers (low / medium / high / critical)
- Build response playbooks for the most likely AI failure modes in your environment
- Conduct tabletop exercises simulating AI incidents annually
- Define communication protocols for affected users, customers, and regulators

MANAGE 4 — Third-Party & Supply Chain Risk

Continuously monitor third-party AI systems and vendor compliance. Update risk assessments as vendors release new model versions or change their terms of service.

- Require AI vendors to provide model documentation and risk disclosures
- Include AI governance requirements in vendor contracts and SLAs
- Monitor open-source AI model updates for new vulnerabilities or behavioral changes

4. The Generative AI Profile (NIST AI 600-1)

Released in July 2024, the Generative AI Profile (NIST AI 600-1) extends the AI RMF specifically for large language models (LLMs) and foundation models such as ChatGPT, Claude, Gemini, Copilot, and similar tools. Any organization using generative AI tools — even as deployers of third-party products — should apply this profile alongside the core framework.

The profile identifies 12 unique risk categories for generative AI:

Risk Category	What to Watch For
Confabulation (Hallucination)	AI generates plausible sounding, but factually incorrect outputs presented as true
Prompt Injection	Malicious inputs manipulate the model into ignoring instructions or exposing data
Data Privacy	Training data may include sensitive personal or proprietary information
Intellectual Property	Model outputs may reproduce copyrighted content or trade secrets
Cybersecurity	Models may generate malware, phishing content, or attack code
Harmful Content Generation	Models may produce illegal, dangerous, or discriminatory content
Information Security	Sensitive business data entered into AI tools may be retained or exposed
Value Chain Complexity	Downstream risk from foundation models embedded in vendor products
Bias & Discrimination	Output disparities across demographic groups may perpetuate or amplify bias
Transparency & Explainability	Black-box outputs are difficult to audit or explain to affected parties
Accountability	Unclear ownership of AI-generated decisions or errors
Environmental Impact	High computing requirements contribute to carbon emissions at scale

5. Taktik Technologies AI Governance Framework

Taktik Technologies has developed a practical 10-pillar AI Governance Framework aligned with NIST AI RMF 1.0, the EU AI Act, and ISO 42001. Designed specifically for SMB and mid-market organizations, this framework makes enterprise-grade AI governance accessible without requiring a dedicated compliance department.

Framework Pillar	Key Actions	AI RMF Mapping
1. AI Governance Charter	Define AI use policy, risk appetite, and accountability. Appoint an AI Risk Owner. Establish an AI Oversight Committee with C-suite sponsorship.	**GOVERN 1, GOVERN 2**
2. AI System Inventory	Catalog every AI system in used across the organization — vendor, custom, and embedded. Document purpose, data inputs, decision outputs, and affected populations.	**MAP 1, MAP 2**
3. Risk Context & Impact Mapping	For each AI system, define intended use, potential misuse, affected stakeholders, and downstream consequences. Map to business processes and regulatory obligations.	**MAP 3, MAP 4, MAP 5**
4. Trustworthiness Assessment	Evaluate each AI system against all seven trustworthiness characteristics. Document gaps. Prioritize remediation by risk level and business impact.	**MEASURE 1, MEASURE 2**
5. Risk Metrics & Monitoring	Define measurable indicators for AI risk (accuracy drift, bias rates, incident frequency). Establish dashboards and regular review cadences.	**MEASURE 3, MEASURE 4**
6. Risk Treatment Planning	Prioritize identified risks. Assign treatment owners. Define accept, mitigate, transfer, or avoid decisions. Document residual risk and escalation thresholds.	**MANAGE 1, MANAGE 2**
7. Incident Response & Recovery	Define AI-specific incident classifications. Build response playbooks covering AI errors, bias	**MANAGE 3**

	events, data breaches, and model manipulation. Test annually.	
8. Third-Party AI Risk (TPRM)	Assess all AI vendors and open-source model providers. Require contractual accountability. Map vendor AI into organizational risk register.	**GOVERN 6, MANAGE 4**
9. AI Literacy & Training	Train all staff by role: awareness for all employees, deeper governance training for AI owners, technical training for developers and data teams.	**GOVERN 4, GOVERN 5**
10. Continuous Improvement	Conduct annual AI risk reviews. Update the AI inventory as new tools are adopted. Align with evolving NIST profiles (Generative AI, Critical Infrastructure).	**GOVERN 3, MANAGE 4**

6. AI Governance Maturity Model

Understanding your current AI governance maturity is the starting point for any improvement program. Taktik Technologies uses the following five-level maturity model — aligned with the NIST AI RMF — to baseline organizations and track progress over time.

Level	Stage	Characteristics
Level 1	**Initial**	Ad hoc AI use with no formal governance. Risks are unidentified. No AI inventory. No policies or oversight mechanisms in place.
Level 2	**Developing**	Basic AI inventory exists. Some risk awareness. Governance discussions underway but policies are informal and inconsistently applied.
Level 3	**Defined**	Formal AI governance policy in place. Risk register established. Key roles assigned. AI literacy training launched.
Level 4	**Managed**	Quantitative risk metrics in use. Regular monitoring dashboards. Incident response procedures tested. Vendor AI risk assessed.
Level 5	**Optimizing**	Continuous improvement cycle embedded. AI governance integrated with enterprise risk management. External audits conducted. Board-level AI risk reporting in place.

Most SMBs assessed by Taktik Technologies begin at Level 1 or Level 2. Our AI Governance Readiness Assessment will identify your current maturity level, priority gaps, and a targeted roadmap to reach Level 3 or higher within 12 months — at no cost to qualify.

7. Implementation Roadmap

Taktik Technologies recommends a phased approach aligned with NIST AI RMF function sequencing — beginning with governance foundations before moving to measurement and management:

Phase 1: Governance Foundation (0–30 Days)

- Appoint an AI Risk Owner and establish initial oversight accountability
- Draft and publish an AI Use Policy covering approved tools, prohibited uses, and data handling
- Conduct an AI system inventory — catalog every AI tool in use across the organization
- Launch AI literacy awareness training for all staff
- Subscribe to NIST AI RMF Playbook updates and relevant sector profiles

Phase 2: Map & Assess (30–90 Days)

- Complete AI System Context Cards for every tool in the inventory
- Conduct structured risk identification for high-impact AI systems
- Assess each system against all seven trustworthiness characteristics
- Apply the Generative AI Profile (NIST AI 600-1) to any LLM or foundation model tools
- Prioritize risks by likelihood and business impact; update AI Risk Register

Phase 3: Measure & Control (90–180 Days)

- Define success metrics and monitoring thresholds for each AI system
- Implement pre-deployment testing and adversarial evaluation for high-risk systems
- Establish production monitoring dashboards and anomaly alerting
- Review vendor AI documentation; update TPRM assessments and contracts
- Conduct bias and fairness evaluations for AI systems influencing people's decisions

Phase 4: Manage & Improve (Ongoing)

- Operationalize AI incident response playbooks; conduct annual tabletop exercises
- Run quarterly AI risk register reviews with executive stakeholders
- Conduct annual third-party AI governance audits
- Update AI governance framework as new NIST profiles and guidance are released
- Align NIST AI RMF compliance with EU AI Act obligations ahead of the August 2026 deadline

8. Framework Crosswalk: NIST AI RMF, EU AI Act & ISO 42001

The NIST AI RMF is widely recognized as the governance foundation that makes compliance with other AI regulations significantly more achievable. The table below maps NIST AI RMF functions to equivalent obligations under the EU AI Act and ISO 42001, alongside the corresponding Taktik Framework pillar.

NIST AI RMF	EU AI Act	ISO 42001	Taktik Framework
GOVERN 1–2 (Policies & Accountability)	Art. 9 (Risk Mgmt), Art. 4 (AI Literacy)	Clause 5 (Leadership), Clause 6 (Planning)	**AI Governance Charter**
MAP 1–3 (Context & Risk Identification)	Art. 11 (Technical Docs), Annex IV	Clause 8 (Operations), Annex A	**AI Inventory & Risk Mapping**
MEASURE 1–2 (Analysis & Evaluation)	Art. 10 (Data Governance)	Clause 9 (Performance)	**Trustworthiness Assessment**
MANAGE 1–2 (Treatment & Response)	Art. 14 (Human Oversight), Art. 13	Clause 10 (Improvement)	**Risk Treatment & Incident Response**
GOVERN 6 (Third-Party Risk)	Art. 28–29 (Deployer obligations)	Clause 8.4 (Externally provided processes)	**Vendor AI Risk (TPRM)**

> *Organizations that implement the Taktik Technologies AI Governance Framework aligned with NIST AI RMF will be well-positioned to demonstrate conformance with the EU AI Act's risk management, documentation, and human oversight requirements — reducing duplicated compliance effort by an estimated 40–60%.*

9. Cybersecurity & AI Risk Integration

The NIST AI RMF explicitly treats cybersecurity as a trustworthiness characteristic — not a separate workstream. Organizations must integrate AI risk management with existing cybersecurity programs. Key integration points include:

- Zero Trust Architecture: apply least-privilege access to AI model endpoints, APIs, and training data pipelines
- Threat modeling: extend existing threat models to cover AI-specific attack vectors (data poisoning, adversarial inputs, model inversion, membership inference)
- Vulnerability management: monitor AI libraries, frameworks, and open-source models for CVEs and behavioral regressions
- Data security: protect training data, inference inputs, and model outputs with appropriate encryption and access controls
- Incident response: add AI failure modes to existing security incident playbooks — including model manipulation, prompt injection, and hallucination-driven fraud
- Software Asset Management (SAM): track AI tools alongside traditional software in the asset inventory for licensing, patching, and risk visibility

> *Taktik Technologies integrates NIST AI RMF implementation with your broader NIST Cybersecurity Framework (CSF) program — ensuring AI governance reinforces rather than duplicates your existing security investments.*

10. Key Definitions

Term	Definition
AI RMF 1.0	NIST Artificial Intelligence Risk Management Framework, published January 2023. Voluntary, technology-agnostic guidance organized into GOVERN, MAP, MEASURE, and MANAGE functions.
AI Actor	Any entity that plays an active role in the AI lifecycle — including developers, deployers, operators, users, and evaluators.
Confabulation	The tendency of AI systems to generate plausible-sounding but factually incorrect information — commonly called hallucination.

GPAI / Foundation Model	A general-purpose AI model trained on large datasets that can perform a wide range of tasks. Covered by the EU AI Act and the NIST Generative AI Profile (NIST AI 600-1).
Model Card	A standardized document describing an AI model's intended use, performance characteristics, limitations, and ethical considerations.
Prompt Injection	An adversarial attack in which malicious input manipulates an AI system into ignoring its instructions or disclosing sensitive information.
Residual Risk	The risk that remains after controls have been applied. Must be documented, accepted by an appropriate authority, and monitored over time.
TPRM	Third-Party Risk Management — the process of identifying, assessing, and managing risks arising from vendors, partners, and other external AI system providers.
Trustworthy AI	AI that is valid, reliable, safe, secure, resilient, accountable, transparent, explainable, privacy-enhanced, and fair — as defined by the NIST AI RMF.

Disclaimer: This document is provided for informational and educational purposes only. It does not constitute legal advice. The NIST AI RMF is voluntary guidance; specific regulatory obligations vary by jurisdiction, sector, and AI system type. Organizations should consult qualified legal and compliance counsel to assess their specific obligations. NIST guidance continues to evolve; consult nist.gov/itl/ai-risk-management-framework for current publications.

Ready to Implement the NIST AI RMF?

Taktik Technologies delivers end-to-end AI Governance advisory services aligned with NIST AI RMF 1.0, the EU AI Act, and ISO 42001 — purpose-built for SMB and mid-market organizations.

Contact us today for a no-cost AI Governance Readiness Assessment

www.taktiktechnologies.com | sales@taktiktech.com | (678) 662-5700