

NIST AI RMF

Quick Reference Guide

Plain-language translations of every subcategory

The NIST AI Risk Management Framework organizes AI governance into four functions: GOVERN, MAP, MEASURE, and MANAGE. This guide translates each subcategory into plain language so your team can act — not just read.

GOVERN**MAP****MEASURE****MANAGE**

How to Use This Guide

The NIST AI Risk Management Framework (AI RMF 1.0) is a voluntary framework published by the National Institute of Standards and Technology in January 2023. It provides a structured approach to managing the risks of AI systems throughout their lifecycle.

This guide translates every subcategory into plain language — what it means, and what it requires your organization to actually do. Use it as a working reference, a leadership conversation starter, or a roadmap for your AI governance program.

Function	Purpose	Subcategories
GOVERN	Policies, roles, and culture	19 subcategories
MAP	Context, uses, and harm identification	18 subcategories
MEASURE	Testing, metrics, and monitoring	22 subcategories
MANAGE	Response, recovery, and improvement	13 subcategories

RiskGov Tip: Start with GOVERN. Without accountability structures in place, MAP, MEASURE, and MANAGE will not stick. Governance is not the last step — it's the foundation.

GOVERN

Put accountability in writing before AI goes live.

Establishes the policies, roles, and culture needed to manage AI risk across your organization.

GOV-1.1

Know your legal landscape

Identify all AI-relevant laws, regulations, and contractual obligations — and document how you comply with each.

GOV-1.2

Embed trustworthiness into policy

Make fairness, transparency, and safety explicit in your AI policies — not just implied.

GOV-1.3

Right-size your risk effort

Calibrate how much governance work each AI system needs based on your organization's actual risk tolerance.

GOV-1.4

Make risk decisions visible

Publish the criteria and controls that guide your AI risk decisions so stakeholders can see and trust the process.

GOV-1.5

Build in regular review

Schedule periodic reviews of your risk management process — and assign who owns them.

GOV-1.6

Inventory your AI systems

Maintain a living register of all AI systems in use, with risk priority levels assigned to each.

GOV-1.7

Plan for retirement

Define how AI systems will be safely decommissioned so shutdowns don't create new risks.

GOV-2.1

Map roles to responsibilities

Document who owns what across the AI risk lifecycle — and make sure people actually know their role.

**GOV-
2.2****Train your people**

Provide AI risk management training to everyone involved — not just technical teams.

**GOV-
3.1****Build diverse teams**

Bring varied perspectives — demographic, disciplinary, experiential — into AI risk decisions.

**GOV-
4.1****Foster safety-first culture**

Encourage critical thinking and 'speak up' norms around AI risk — not just compliance checkboxes.

**GOV-
4.3****Enable testing and incident sharing**

Create clear channels for AI testing, incident reporting, and internal knowledge sharing.

**GOV-
5.2****Act on what you hear**

Build feedback loops that feed external input back into system design — not just reports.

**GOV-
6.2****Prepare for third-party failure**

Have a contingency plan ready for when a vendor's AI system fails or causes harm.

**GOV-
2.3****Leadership owns the risk**

Executive leadership takes responsibility for decisions about AI system risks associated with development and deployment.

**GOV-
3.2****Define human oversight roles**

Specify exactly how humans supervise AI systems and what authority they have to intervene.

**GOV-
4.2****Document and communicate risk**

Teams must write down what risks their AI creates and share that knowledge broadly.

**GOV-
5.1****Collect external feedback**

Actively gather input from people outside the development team about real-world AI impacts.

**GOV-
6.1****Govern your vendors**

Extend your AI risk policies to third-party tools, data sources, and partners.

MAP

Understand the system before you deploy it.

Identifies the context, intended uses, potential harms, and risk landscape of each AI system.

MAP-1.1

Define purpose and context

Document what the AI is for, who uses it, what it could harm, and what assumptions underlie it.

MAP-1.2

Assemble the right team

Ensure those mapping AI risk include people with diverse expertise — and document who contributed.

MAP-1.3

Align to organizational mission

Confirm the AI system supports your organization's actual goals — not just a technical capability.

MAP-1.4

Validate business value

Clearly define what business problem this AI solves — and re-evaluate it for existing systems.

MAP-1.5

Document risk tolerance

Write down how much risk your organization is willing to accept for this specific AI system.

MAP-1.6

Capture stakeholder requirements

Gather system requirements from actual users and affected parties — with privacy and fairness built in.

MAP-2.1

Specify AI tasks and methods

Define exactly what the AI does: classify, generate, recommend — and how it does it.

MAP-2.2

Document system limits

Be explicit about what the AI doesn't know and how humans should oversee its outputs.

MAP-2.3

Address scientific integrity

Document data sources, selection criteria, and validation methods — especially for high-stakes decisions.

MAP-3.1

Examine intended benefits

Document the positive outcomes the AI is designed to produce — and test whether they materialize.

**MAP-
3.2****Cost what can go wrong**

Quantify or qualify the costs of AI errors — financial, reputational, and human.

**MAP-
3.3****Bound the application scope**

Define exactly where and how the AI can be used — and what's out of scope.

**MAP-
3.4****Assess operator proficiency**

Determine if the people running the AI actually understand it well enough to catch failures.

**MAP-
3.5****Define human oversight processes**

Write down how humans will monitor, intervene in, and correct the AI during operation.

**MAP-
4.1****Map third-party risks**

Identify risks from external data, APIs, or AI models baked into your system.

**MAP-
4.2****Document internal risk controls**

List the safeguards you've built into the system for each identified risk.

**MAP-
5.1****Assess impact likelihood and magnitude**

Estimate how likely each harm is and how severe — using real data where possible.

**MAP-
5.2****Engage stakeholders regularly**

Create formal mechanisms for ongoing external feedback and impact monitoring.

MEASURE

Test it. Track it. Trust only what you can verify.

Analyzes and quantifies AI risks using metrics, evaluations, and ongoing monitoring.

MEA-1.1

Select the right metrics

Choose measurements that match your highest-priority risks — and document what you're choosing not to measure.

MEA-1.2

Review metrics regularly

Reassess whether your measures still capture real risk — including impacts on affected communities.

MEA-1.3

Use independent reviewers

Bring in internal experts not on the build team — and external assessors where needed.

MEA-2.1

Document your test setup

Record every test set, metric, and tool used during evaluation — so others can reproduce and audit.

MEA-2.2

Protect human subjects

Evaluations involving people must meet ethical and legal requirements and be representative of real users.

MEA-2.3

Measure actual performance

Test the system under conditions that match real deployment — not just ideal lab settings.

MEA-2.4

Monitor in production

Track system behavior after deployment — not just before launch.

MEA-2.5

Validate reliability

Demonstrate the system performs consistently and document where its generalizability breaks down.

**MEA-
2.6****Evaluate safety risks**

Actively test for safety failures — and confirm the system can fail safely when pushed beyond its limits.

**MEA-
2.8****Examine transparency and accountability**

Measure whether stakeholders can understand and contest AI decisions.

**MEA-
2.10****Assess privacy risk**

Evaluate whether the AI exposes personal data or enables re-identification.

**MEA-
2.12****Track environmental impact**

Assess the energy and environmental cost of training and running the AI model.

**MEA-
3.1****Track emerging risks**

Continuously identify new risks that weren't visible at deployment — and log them.

**MEA-
3.3****Build user feedback channels**

Give end users a way to report problems and appeal AI decisions — and feed that back into evaluation.

**MEA-
2.7****Assess security and resilience**

Test for adversarial attacks, data poisoning, and system brittleness under stress.

**MEA-
2.9****Explain model outputs**

Document how the model works and interpret outputs in the context they'll actually be used.

**MEA-
2.11****Test for fairness and bias**

Measure disparate performance across demographic groups — and document findings even if unflattering.

**MEA-
2.13****Evaluate your evaluation**

Review whether your measurement approach itself is working — and update it when it isn't.

**MEA-
3.2****Account for hard-to-measure risks**

Acknowledge risks you can't yet quantify — and document your monitoring approach for them.

**MEA-
4.1****Ground measures in deployment context**

Ensure your metrics reflect how the system actually performs for real users in real settings.

**MEA-
4.2****Validate with domain experts**

Have subject-matter experts and affected communities confirm the system is working as intended.

**MEA-
4.3****Track performance changes**

Document measurable improvements or regressions over time — with input from real users.

MANAGE

When something goes wrong — and it will — be ready.

Prioritizes and addresses AI risks through response plans, monitoring, and continual improvement.

MAN-1.1

Make the go/no-go call

Formally decide whether the AI system should proceed — based on whether it achieves its intended purpose.

MAN-1.2

Prioritize your risks

Rank documented risks by impact and likelihood — and allocate resources accordingly.

MAN-1.3

Plan your risk responses

For high-priority risks, document your response: mitigate, transfer, avoid, or accept — and why.

MAN-1.4

Own residual risk

Document the risks that remain after mitigation — and disclose them to downstream users.

MAN-2.1

Budget for risk management

Allocate real resources to AI risk controls — and consider non-AI alternatives where appropriate.

MAN-2.2

Sustain deployed systems

Maintain processes that keep AI systems performing as intended over time.

MAN-2.3

Respond to unknown risks

Have procedures ready for risks that emerge unexpectedly — before they occur.

MAN-2.4

Be able to shut it down

Establish clear authority and mechanisms to override, disengage, or deactivate any AI system.

**MAN-
3.1****Monitor third-party AI**

Continuously watch for risks from AI components you didn't build — and apply controls.

**MAN-
3.2****Track pre-trained models**

Include foundation models and pre-trained components in your ongoing monitoring — they drift too.

**MAN-
4.1****Implement monitoring plans**

Execute post-deployment monitoring with user feedback, override mechanisms, and incident response.

**MAN-
4.2****Improve continuously**

Make measurable improvements part of every system update — with regular stakeholder input.

**MAN-
4.3****Communicate incidents**

Report errors and incidents to all affected parties — and follow documented recovery processes.

Ready to build your AI governance program?

RiskGov helps organizations translate frameworks like the NIST AI RMF into working governance programs — not binders that sit on a shelf. We bring federal compliance depth, human-centered design, and senior expertise to every engagement.

AI Data Exposure Assessment	Map where your AI systems touch sensitive data and where governance gaps exist.
AI Governance Program Build	Design and implement a governance framework aligned to NIST AI RMF, EU AI Act, and your sector.
Fractional Chief AI Risk Officer	Senior AI governance leadership without the full-time hire.
GRC Advisory	Integrate AI risk into your existing Governance, Risk, and Compliance program.
Proprietary GRC Platform	Operationalize your governance framework with our platform — built to federal compliance standards.

riskgov.ai

info@riskgov.ai

Reston, Virginia

This guide is based on NIST AI RMF 1.0 (January 2023). It is provided for informational and governance planning purposes only and does not constitute legal advice. Subcategory descriptions are RiskGov plain-language translations — refer to the official NIST publication for authoritative text.