



NIST AI Risk Management Framework for Small Businesses

A plain language summary for owners and teams using AI in everyday work

How the framework applies to your business

If your business uses AI to write marketing copy, summarize customer messages, screen applicants or answer customer questions, this framework is relevant even when you did not build the technology. It helps you decide where AI is useful, what could go wrong and what checks are proportionate to the consequences.

NIST designed the AI Risk Management Framework for organizations of all sizes. A small business can start with one accountable owner and a short record of its AI tools. You do not need to adopt every action at once. Spend more attention on uses that affect people, sensitive information or important business decisions.

For example, a five person agency uses AI to draft a client newsletter. The owner approves the tool, checks its data settings, prohibits confidential client inputs and assigns someone to verify claims before publication. An incorrect draft is corrected and logged. Those ordinary controls put the framework into practice.

What this guide is: An independent introduction to AI RMF 1.0, with PBW examples and suggested starting actions. It is not a NIST publication, a certification, legal advice or a complete implementation of every framework outcome. Using it does not establish legal compliance.

The four functions in everyday work

Function	Practical question	A useful first record
GOVERN	Who approves AI use and sets the rules?	An owner and a short AI use policy.
MAP	Where is AI used and who could be affected?	A tool list showing purposes, data and affected people.
MEASURE	Does the tool work well enough for this use?	Recorded tests, errors and review findings.
MANAGE	What will we change, pause or stop?	Actions, escalation contacts and a fallback process.

Governance supports the other functions throughout AI use. Mapping, measuring and managing are repeated as tools and circumstances change; they are not a one time checklist. The records above are PBW starting suggestions, not mandatory NIST forms.



What trustworthy AI means in practice

Trustworthiness depends on the purpose and setting. A tool that produces acceptable draft slogans may be unsuitable for hiring decisions. Consider the following characteristics together; improving one does not resolve every other concern.

Valid and reliable: Check whether it performs consistently on work like yours, including difficult cases and changed conditions.

Safe: Consider how incorrect outputs or unexpected behavior could harm people or operations.

Secure and resilient: Check access controls, misuse risks and how the service recovers from disruption.

Accountable and transparent: Identify the responsible person and what information users need about AI use and its limitations. Transparency does not mean the internal decision process is fully understandable.

Explainable and interpretable: Ask whether an output can be explained and understood well enough for the intended decision.

Privacy enhanced: Limit personal information and review collection, storage, access and reuse of data.

Fair with harmful bias managed: Look for harmful differences in outcomes, including effects on people with protected characteristics.

Examples you may recognize

Use	Risk to consider	A practical safeguard
Marketing copy	Invented claims or misleading statistics.	Verify factual claims and review the final copy before publication.
Customer chatbot	Incorrect refund or pricing information.	Test common questions and offer an easy route to a person.
Applicant screening	Qualified people unfairly excluded.	Evaluate outcomes and require appropriate human review before decisions. Get specialist advice on applicable rules.
Pricing or fraud scoring	Unfair treatment or incorrect restrictions.	Review disputed outcomes, record overrides and provide a review route.

Human review needs to be meaningful: the reviewer must have enough information, time and authority to challenge the result. Simply placing a person at the end of an automated process may not adequately control its risk.



Managing AI tools supplied by vendors

Many small businesses use AI through software subscriptions, plugins and embedded features. A vendor may control the model, but your business still needs to manage how it uses the tool. Do not treat a vendor assurance as evidence that every use is appropriate.

Questions to ask before buying or renewing

1. Where is our data stored, who can access it, and is it used for model training or shared with others? What settings or contract terms control this?
2. What testing supports the specific task we intend to use? What limitations, error patterns and performance differences are known?
3. What information can you provide about training data and safeguards? If details are unavailable, what assurance can we independently assess?
4. How will we learn about material changes, outages or security incidents? Can we disable the AI feature while an issue is investigated?
5. Can a person review or override the output? Can we obtain useful logs and export or delete our data if we leave?
6. What do the contract terms say about responsibilities, support and liability? Have relevant terms reviewed when the stakes warrant it?

Keep the answers and note unresolved gaps. An error rate without its test conditions may tell you little about your task. A tool can change after launch, so recheck important uses after material updates or repeated errors.

A short AI use policy

A first policy can be brief. Cover these decisions, then obtain appropriate review for your business and industry:

Purpose and approved uses: Identify allowed tools and tasks. Explain which uses require additional approval or are prohibited.

Data restrictions: Specify what must not be entered into external tools and how approved sensitive data must be protected.

Human review: State who checks outputs and which decisions cannot proceed without that review.

Ownership and escalation: Name the approver, the contact for errors or concerns, and who can pause use.

Review and records: Set a review schedule and retain approvals, test findings and incident actions. Revisit the policy after material changes.



A manageable starting plan

This is a PBW starting suggestion, not a NIST deadline. Adjust it to your capacity. Investigate serious concerns immediately rather than waiting for a scheduled review.

When	Action	What to keep
Week 1	List the AI tools and embedded features actually used.	Purpose, owner, data involved and affected people.
Week 2	Prioritize uses with sensitive data or significant effects on people.	Reasons for priority and known gaps.
Week 3	Approve a short policy and define meaningful human checks.	Approval, restrictions and escalation contacts.
Week 4	Test priority uses and review vendor answers.	Findings, corrective actions and a fallback plan.
Ongoing	Review at a risk appropriate frequency and after changes or incidents.	Updated decisions and evidence of follow through.

Terms used in the framework

AI actor: A person or organization involved in the AI lifecycle, including deploying or operating a tool. **Residual risk:** Risk remaining after safeguards. **TEVV:** Testing, evaluation, verification and validation. **Profile:** A tailored application of framework outcomes to a particular setting or use.

Where legal and contractual requirements fit

The framework is voluntary. Separate laws, contracts and industry requirements may apply, depending on your location and AI use. Obtain qualified advice for higher impact uses such as employment, lending or health decisions. This summary does not promise a legal defense, insurance coverage or compliance status.

Sources and scope

Primary source: NIST AI 100-1, Artificial Intelligence Risk Management Framework (AI RMF 1.0), January 2023. [Read the complete framework](#). Companion resource: [NIST AI RMF Playbook](#). NIST also provides a [Generative AI Profile](#).

Checked October 3, 2026. NIST reports that AI RMF 1.0 is being revised. This guide summarizes the 2023 framework and does not cover every outcome or the complete Generative AI Profile. Check the [NIST framework page](#) for updates. PBW is not affiliated with or endorsed by NIST.

When you are ready to put it into practice

For structured implementation, explore the [AI Governance Toolkit Professional Edition](#), or begin with our free [AI checklist article](#). Neither guarantees compliance.