


DEFENSELOGIX AI RISK MANAGEMENT

Who approved this AI, *and how do you know it works?*

Most organizations use more AI than leadership realizes: features switched on inside existing software, tools bought by one team, consumer apps used with business data. DefenseLogix helps regulated and trust-sensitive organizations govern that AI with the NIST AI Risk Management Framework.

The goal is practical: an inventory you can trust, rules people can follow, and documentation that answers the questions examiners, auditors and customers now ask about AI.

The framework is voluntary, but it is increasingly the common language behind those questions. It also maps into programs you already run, such as NIST SP 800-53, SOC 2, ISO 27001 and CMMC.



01 • INVENTORY

Know what you run.

You cannot govern what you have not listed.

- AI products, AI features inside existing software, and consumer tools
- Owner, data received, and any effect on decisions about people
- Vendor terms on training and data retention

02 • RULES

Set clear rules.

Governance people can apply without escalating.

- A risk-tiering rule a manager can apply in a minute
- A short list of pre-approved tools
- An approval path faster than working around it

03 • EVIDENCE

Show that it works.

Answers from documentation, not memory.

- Testing on your own use cases, not vendor claims
- Monitoring after a system goes live
- Records an auditor or examiner can follow

WHAT AN ENGAGEMENT CAN INCLUDE

AI system inventory and classification	GOVERN • MAP	Risk-tiering rule and approval process	GOVERN
AI acceptable use and governance policy	GOVERN	Third-party AI vendor review	GOVERN • MANAGE
Current and target profile, with a gap roadmap	ALL FOUR	AI incident response procedures	MANAGE
Crosswalk to NIST SP 800-53, CMMC and SOC 2	MAPPING	Audit and assessment support for AI controls	EVIDENCE

THE FRAMEWORK, APPLIED

What each function asks, and what we help you build.

FUNCTION	THE QUESTION IT ASKS	WHAT WE HELP YOU BUILD
GOVERN	<i>Who decides, and under what rules?</i>	Accountability and approval authority, an AI policy, a risk tolerance statement, the AI inventory, and oversight of third-party AI.
MAP	<i>What is each system for, and who does it affect?</i>	For each system: intended use, data sources, the people affected and potential harms, plus a risk tier for each use case.
MEASURE	<i>How do you know it works, and keeps working?</i>	A test and evaluation approach on your own use cases, bias checks where outputs affect people, and monitoring for drift after launch.
MANAGE	<i>What happens when it does not?</i>	Prioritized risk responses, deployment gates, AI incident response, vendor monitoring and decommissioning.

GENERATIVE AI IS IN SCOPE

The Generative AI Profile (NIST AI 600-1, July 2024) identifies 12 risks specific to generative AI, including confabulation, data privacy and information security, and maps suggested actions to the same four functions. We apply it wherever you use large language models.

We also look at the whole system, not only the model: the tools it can call, the data it reads, and the access controls around it.

STRAIGHT ANSWERS

The AI RMF is voluntary, and there is no AI RMF certification. What it gives you is a defensible structure and a shared vocabulary with auditors, examiners and customers.

BUILT FOR REGULATED WORK

Banking

The April 2026 interagency model risk management guidance places generative and agentic AI outside its scope. The AI RMF gives those systems a governance structure until further guidance arrives.

Healthcare

Any AI service that receives protected health information needs a business associate agreement, including AI features added to a platform you already use.

Defense contractors

AI used near Controlled Unclassified Information needs governance too, and AI RMF findings map into the CMMC requirements you already track.

Start with an inventory. *We'll help with the rest.*

Start a conversation: tell us where AI is in use today and what is driving the question, whether an examiner, a customer or a new tool, and we'll suggest a practical starting point.

Or read the AI RMF Field Guide, our five-part plain-English series on the DefenseLogix blog.

defenselogix.com/services/ai

