

Operationalizing NIST AI RMF for Deployed AI

The AI RMF tells you what good governance looks like. This is how the four functions become running infrastructure — with cryptographic evidence at every step.

FUNCTION	WHAT THE AI RMF CALLS FOR ¹	HOW IT RUNS ON AI ASSESS TECH
GOVERN	A culture of risk management: accountability structures, roles, policies, and oversight of AI systems.	AI System Registry binds every deployed system to a named owner and bound API key. Multi-party signed attestation (sign, countersign, witness, dissent) with executive RBAC. Risk Acceptance Ledger — justified, time-boxed, revocable — chains executive decisions to the evidence behind them.
MAP	Context is established: each system's purpose, deployment configuration, and risk surface are identified and categorized.	Registration captures the deployment as configured — system, owner, environment, model, and bound configuration. Assessment runs against your actual system prompt and tools, not the vendor's base model, so the mapped context is the tested context. Archetype classification locates each agent on the behavioral map.
MEASURE	Identified risks are assessed, analyzed, and tracked with appropriate methods and metrics — repeatable, comparable, documented.	Runtime LCSH assessment: 120 forced-choice items, continuous 0–10 scores on four behavioral dimensions, from question banks locked under a published hash and anchored to Ethereum. Scheduled reassessment on a compliance calendar; trajectory tracking turns snapshots into drift measurement; identical instrument enables cross-model, cross-vendor comparison.
MANAGE	Risks are prioritized and acted upon: response, recovery, communication, and continuous improvement.	Drift alerts feed configurable escalation chains with named routing and live delivery. Model-Change Reassessment Gate forces fresh measurement before updates ship. Board Pack PDF with a public verification link communicates posture upward; every action lands in the tamper-evident audit chain.

ASSESSMENT API Asynchronous REST integration: submit an assessment (202 Accepted), poll status, fetch scored results. Fits CI/CD and MLOps pipelines — gate a release on a fresh behavioral score the same way you gate on tests.	EVIDENCE & VERIFICATION Every record is SHA-256 hash-chained; audit exports hand examiners the evidence trail. Verification is public and login-free at aiassesstech.com/verify — a third party confirms integrity without trusting your systems.	ALERTING & WORKFLOW Drift findings route through configurable escalation chains with live Slack delivery; SIEM alerting surfaces behavioral events beside security events. Governance signals land where your teams already work.
---	---	---

WHY MEASURE IS THE HARD FUNCTION — AND THE ONE WE'RE BUILT FOR

Govern, Map, and Manage can be run with policy documents and meetings. Measure cannot — it needs an instrument. The AI RMF calls for repeatable, documented measurement of AI risk but does not specify how to measure behavioral risk in a deployed, closed-weight system. A locked, hash-anchored question bank producing quantified dimensional scores against your live configuration is that instrument: the measurement is fixed, the results are comparable over time, and the evidence is independently verifiable.

¹ NIST AI Risk Management Framework 1.0 (NIST AI 100-1, January 2023): Govern, Map, Measure, Manage. Function summaries paraphrased; see also NIST AI 600-1, Generative AI Profile (July 2024). This document describes capability alignment; it is not a NIST certification, endorsement, or conformity determination — NIST does not certify products.