

NIST AI Risk Management Framework (AI RMF 1.0)

NIST AI RMF

The U.S. federal baseline for trustworthy AI. If you only run one pack, run this one. It will not block on its own — it produces evidence and surfaces gaps your auditor can act on.

AUTHORITY	National Institute of Standards and Technology
PUBLICATION	NIST AI 100-1
VERSION	1.0.0 · License: Apache-2.0
URL	https://www.nist.gov/itl/ai-risk-management-framework

About this pack

This is the NIST AI RMF pack from the AIGovOps Beacon audit framework. Walk it with your auditor, your CISO, and the engineer who actually ships the system. It is short on purpose.

Items marked **auto** are computed by Beacon from signed receipts and the model inventory — you do not have to hand-evidence them, the platform does. Items without that badge require a human attestation: a person reads the prompt, checks the evidence is real, and ticks a box. Beacon records the attestation as a signed receipt with `event_type attestation`.

The severity label tells you how much weight an item carries in the production-readiness gate. Critical items can fail a release on their own. Lower-severity items contribute to a weighted score.

This pack is released under the Apache License 2.0. The underlying source publication retains its own copyright; consult it for normative language.

Checklist — Functions

GOVERN

Govern

Cultivate a culture of risk management.

GOVERN-1.1

HIGH

Is there a written AI policy that names an accountable owner?

Evidence `policy_document_uri` Gate `governance_policy_present`

☐ Yes ☐ No ☐ N/A

GOVERN-1.2

HIGH

Are roles and responsibilities for AI risk documented?

Evidence `raci_uri` Gate `roles_documented`

☐ Yes ☐ No ☐ N/A

GOVERN-1.3

MEDIUM

Does the organization track legal and regulatory requirements relevant to its AI systems?

Evidence `regulatory_register_uri`

☐ Yes ☐ No ☐ N/A

GOVERN-1.4

MEDIUM

Are risk tolerances for AI systems explicitly defined?

Evidence `risk_appetite_statement_uri`

☐ Yes ☐ No ☐ N/A

GOVERN-2.1

MEDIUM

Is there a process for stakeholder input on high-risk systems?

Evidence `stakeholder_engagement_log`

☐ Yes ☐ No ☐ N/A

GOVERN-3.1

CRITICAL

AUTO

Does the inventory record vendor, model, version, and owner?

Evidence `beacon_inventory_row` `Gate inventory_complete`

☐ Yes ☐ No ☐ N/A

GOVERN-4.1

HIGH

Is there an incident response runbook for AI failures?

Evidence `incident_runbook_uri`

☐ Yes ☐ No ☐ N/A

GOVERN-4.2

HIGH

Are AI incidents logged with root-cause analysis?

Evidence incident_log_uri

☐ Yes ☐ No ☐ N/A

GOVERN-5.1

HIGH

Are third-party model risks reviewed before procurement?

Evidence vendor_review_uri

☐ Yes ☐ No ☐ N/A

GOVERN-6.1

CRITICAL

AUTO

Is signing-key rotation policy defined and enforced?

Evidence key_rotation_log Gate key_rotation_within_90d

☐ Yes ☐ No ☐ N/A

MAP

Map

Establish context to frame AI risks.

MAP-1.1

HIGH

Is the intended use of the model documented in plain language?

Evidence model_card_uri Gate model_card_present

☐ Yes ☐ No ☐ N/A

MAP-1.2

HIGH

Are foreseeable misuses enumerated?

Evidence misuse_register_uri

☐ Yes ☐ No ☐ N/A

MAP-1.3

HIGH

Are affected populations identified, including vulnerable groups?

Evidence impact_assessment_uri

☐ Yes ☐ No ☐ N/A

MAP-2.1

MEDIUM

Is the data lineage of training/fine-tuning data documented?

Evidence data_lineage_uri

☐ Yes ☐ No ☐ N/A

MAP-2.2

HIGH

Are licensing and provenance of third-party data verified?

Evidence data_license_attestation

☐ Yes ☐ No ☐ N/A

MAP-3.1

LOW

Are benefits and costs of the AI system documented?

Evidence benefit_cost_memo

☐ Yes ☐ No ☐ N/A

MAP-4.1

CRITICAL

Are risks to civil rights, civil liberties, and privacy mapped?

Evidence `rights_review_uri`

☐ Yes ☐ No ☐ N/A

MAP-5.1

CRITICAL

Is human oversight defined for each decision the system supports?

Evidence `oversight_design_uri` Gate `human_oversight_defined`

☐ Yes ☐ No ☐ N/A

MEASURE

Measure

Analyze, assess, and track AI risks.

MEASURE-1.1

HIGH

Are evaluation metrics defined before deployment?

Evidence `eval_plan_uri`

☐ Yes ☐ No ☐ N/A

MEASURE-2.1

CRITICAL

Is bias evaluated across protected attributes where relevant?

Evidence `bias_report_uri`

☐ Yes ☐ No ☐ N/A

MEASURE-2.2

HIGH

Are robustness and adversarial tests performed and logged?

Evidence robustness_report_uri

☐ Yes ☐ No ☐ N/A

MEASURE-2.3

CRITICAL

Is privacy leakage tested (PII, training data extraction)?

Evidence privacy_eval_uri

☐ Yes ☐ No ☐ N/A

MEASURE-2.4

HIGH

Are hallucination/grounding rates measured for generative systems?

Evidence grounding_eval_uri

☐ Yes ☐ No ☐ N/A

MEASURE-3.1

HIGH

AUTO

Is system performance monitored in production?

Evidence monitoring_dashboard_uri Gate monitoring_active

☐ Yes ☐ No ☐ N/A

MEASURE-3.2

MEDIUM

Are drift detectors configured?

Evidence drift_config_uri

☐ Yes ☐ No ☐ N/A

MEASURE-4.1

CRITICAL

AUTO

Are receipts captured for every model invocation?

Evidence beacon_receipts_count Gate receipts_captured

☐ Yes ☐ No ☐ N/A

MEASURE-4.2

CRITICAL

AUTO

Are receipts cryptographically signed and verifiable?

Evidence beacon_signature_verify Gate signatures_valid

☐ Yes ☐ No ☐ N/A

MANAGE

Manage

Allocate risk resources to mapped and measured risks.

MANAGE-1.1

HIGH

Are residual risks documented and accepted by a named owner?

Evidence risk_acceptance_memo

☐ Yes ☐ No ☐ N/A

MANAGE-1.2

MEDIUM

Is there a process to deprecate or retire models?

Evidence deprecation_policy_uri

☐ Yes ☐ No ☐ N/A

MANAGE-2.1

CRITICAL

Are humans in the loop for high-stakes decisions?

Evidence `hitl_config_uri` `Gate hitl_for_high_risk`

☐ Yes ☐ No ☐ N/A

MANAGE-2.2

HIGH

Is there a kill-switch documented and tested?

Evidence `killswitch_test_log`

☐ Yes ☐ No ☐ N/A

MANAGE-3.1

MEDIUM

Are downstream users notified of material model changes?

Evidence `change_log_uri`

☐ Yes ☐ No ☐ N/A

MANAGE-4.1

HIGH

Are incidents reported externally where required by law?

Evidence `external_disclosure_log`

☐ Yes ☐ No ☐ N/A

MANAGE-4.2

MEDIUM

Are post-incident reviews fed back into MAP/MEASURE?

Evidence `feedback_loop_uri`

☐ Yes ☐ No ☐ N/A

Scoring

How Beacon scores this pack:

- auto_check items are computed from receipts and inventory.
- non-auto items are attested by the auditor and recorded as a receipt with event_type = "attestation".
- severity weights for the production_readiness gate:
critical: 4 high: 2 medium: 1 low: 0.5
- PASS threshold defaults to 0.85 of available weight, configurable per tenant in policy/gate.production_readiness.yaml.