PUBLIC PROJECT DOCUMENT / AI RMF 1.0 CROSSWALK

NIST AI RMF prototype profile

Governed AI Service Desk / version 0.1.0

HOW TO READ THIS PROFILE

A portfolio mapping organized by the voluntary AI RMF 1.0 functions. This is not a NIST assessment, certification or formal compliance finding. Local tests are distinguished from documented plans and unverified Azure controls.

Scope: design and local prototype; synthetic data only; no Azure deployment or cloud test. Evidence status follows the project ledger last verified 15 July 2026.

GOVERN

Accountability and risk work across the project lifecycle.

PROJECT LENS	CURRENT PROJECT EVIDENCE	REMAINING EVIDENCE
Accountability	Owner, intended use, prohibited use, oversight and review triggers are documented.	Add independent approvers and a formal responsibility matrix.
Roles	The governance pack defines needed owner, content, identity, security, AI-risk and incident functions.	Assign accountable people in an organization.
Lifecycle	Risk register, threat model, tests, change triggers and incident plan accompany the code.	Connect to enterprise risk and change processes.
Obligations	Consequential decisions and Restricted data are excluded; framework themes are mapped.	Complete jurisdiction-, sector- and data-specific reviews.
Supply chain	Azure, model-provider, package and container dependencies are identified.	Add SBOM, scanning, vendor assessment and contract review.
Safe use	Transparency and escalation expectations are documented.	Run training and usability studies.

These are project-specific summaries, not verbatim NIST outcome statements. Framework reference: NIST AI RMF 1.0 (AI 100-1).

AI RMF PROFILE / FUNCTION 2 OF 4

Map

Use context, stakeholders, data flows and risk boundaries.

MAP

PROJECT LENS	CURRENT PROJECT EVIDENCE	REMAINING EVIDENCE
Purpose	Six categories of synthetic, low-risk IT guidance define the task.	Interview service-desk users before operational use.
System boundary	The API provides guidance, without executing tools or making consequential decisions.	Inventory real interfaces, dependencies and handoffs.
Data flows	Inputs, classification, redaction, retrieval, output and audit fields are documented.	Complete a privacy assessment before real data.
Benefits and costs	Faster access to approved guidance is a hypothesis; cloud cost and oversight are noted.	Measure time saved, errors and operational burden.
Stakeholders and harms	Employees, service desk, security, privacy and audit are considered; exposure and over-reliance are risks.	Assess accessibility, language and disparate impact.
Risk tolerance	Restricted data and automated actions are prohibited; a local Restricted-classification block is tested.	Obtain approval for any expanded use.

AI RMF PROFILE / FUNCTION 3 OF 4

Measure

Local tests and planned assessment of quality, safety and privacy.

MEASURE

PROJECT LENS	CURRENT PROJECT EVIDENCE	REMAINING EVIDENCE
Test methods	The July 2026 ledger records 10 passing local tests for policy, role, grounding and audit behavior.	Expand abuse cases, performance baselines and regression data.
Answer reliability	Tests cover cited guidance and escalation where no approved article matches.	Build a labeled answer set and approve thresholds.
Safety and security	Selected prompt-risk tests and a threat-model catalog exist.	Run red-team and independent security exercises.
Privacy	Tests check selected PII redaction and absence of raw prompt text in app audit events.	Validate provider diagnostics, retention and a formal privacy assessment.
Transparency	API responses carry citations, policy action, risk labels and request ID.	Test user understanding of warnings and escalation.
Monitoring	Blocks, provider failures, missing evidence and cost are listed as target signals.	Implement dashboards, alert thresholds and feedback.

AI RMF PROFILE / FUNCTION 4 OF 4

Manage

Risk treatment, operational controls and residual risk.

MANAGE

PROJECT LENS	CURRENT PROJECT EVIDENCE	REMAINING EVIDENCE
Prioritization	A project risk register scores likelihood and impact.	Validate scoring with cross-functional reviewers.
Treatment	The register records mitigation, lab acceptance, avoidance and deferral.	Assign risk owners, deadlines and formal acceptance.
Controls	Local gateway controls have test evidence; Entra, managed identity and Azure resources are design only.	Deploy an isolated lab and collect cloud evidence.
Incident response	A plan covers severity, containment, evidence, communication and recovery.	Integrate ticketing, paging and notification processes.
Provider failure	The application handles provider failure and keeps a human service-desk path.	Define service levels and fallback approval.
Residual risk	Limitations and production prerequisites are disclosed.	Obtain independent review before organizational use.

Release boundary

A synthetic-data mock demonstration fits the current portfolio scope. Real organizational use requires accountable approval; privacy, security, legal, accessibility and vendor review; validated cloud identity and audit controls; broader answer and abuse evaluation; and an agreed human escalation process.

EVIDENCE LIMIT

The July 2026 ledger records 10 passing local tests and a clean Ruff check. No Azure deployment, independent review, penetration test, AI red-team exercise or formal audit has occurred.

Source basis: NIST_AI_RMF_PROFILE.md, EVIDENCE_STATUS.md and AI_SYSTEM_CARD.md. This public edition separates locally tested application behavior from the source profile's Azure design language.