HRMLESS

AI Governance Program

Audit Report & Management Attestation

NIST AI Risk Management Framework (AI RMF 1.0)

Audit Date	July 1, 2026
Auditor	Cisco Systems, — Public Sector Automation Team
Scope	Full AI governance program — Govern, Map, Measure, Manage, Transparency
Overall Assessment	**SATISFACTORY**
Classification	INTERNAL
Next Formal Review	July 2027

Table of Contents

1. Independent Auditor's Report 3
2. Management's Assertion 4
3. Executive Summary 5
4. Audit Scope & Methodology 7
5. Detailed Results by Function 8
6. Findings Register & Remediation Status 10
7. Program Strengths 11

1. Independent Auditor's Report

To the Management of HRMLESS:

Scope

We have examined the HRMLESS AI governance program against 41 audit checks derived from the NIST AI Risk Management Framework (AI RMF 1.0), spanning the Govern, Map, Measure, and Manage functions and public transparency disclosures, as of July 1, 2026. Our examination covered the policies, plans, governance documents, registers, model documentation, bias-testing evidence, monitoring automation, and public disclosures identified in Section 4 of this report.

Responsibilities

HRMLESS management is responsible for the design, implementation, and operation of the AI governance program and for the accuracy of its assertion in Section 2. The audit team's responsibility is to express a conclusion on the program based on the procedures performed, which included document review, evidence testing against raw production artifacts, control walkthroughs, and gap analysis as described in Section 4.

Conclusion

Based on the procedures performed, in the audit team's opinion the HRMLESS AI governance program is assessed as SATISFACTORY against the NIST AI Risk Management Framework as of the audit date. Of 41 audit checks, 36 passed at initial assessment; 5 minor findings were identified and no major or critical deficiencies were noted. Management's remediation of the 5 findings is documented in Section 6. The next scheduled governance review is October 2026 (Q3 quarterly review).

Cisco Systems – US Public Sector Automation Team

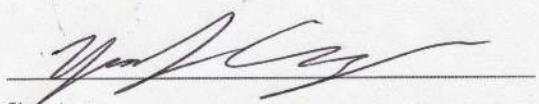

Signature

Yusuf Coppin,
Software Technical Leader

Printed Name & Title

8 JULY 2026

Date

2. Management's Assertion

HRMLESS management asserts that, as of July 1, 2026:

1. The description of the AI governance program contained in this report fairly presents the program's policies, plans, controls, and operational evidence as designed and implemented.
2. The 41 audit checks were evaluated against evidence that is accurate and complete in all material respects, including the production bias-testing artifacts, the entity tracking records, and the public Trust Center disclosures.
3. The remediation actions described in Section 6 have been completed or scheduled as stated, and management accepts responsibility for their ongoing operation.
4. Management has disclosed to the audit team all known AI-related incidents, deficiencies, and instances of non-compliance; none were identified during the period.

Chief Technical Officer, HRMLESS

Signature

Ben Schuster,
Co-Founder, CTO

Printed Name & Title

7 July 2026

Date

3. Executive Summary

The HRMLESS AI governance program was assessed by an independent audit team against 41 audit checks derived from the NIST AI Risk Management Framework (AI RMF 1.0) across five domains: Govern, Map, Measure, Manage, and Transparency. The initial assessment resulted in 36 of 41 checks passed, with 5 minor findings and zero critical or major findings. Following the audit, management initiated remediation of all five findings; the status of each is documented in Section 6.

The program demonstrates strong foundational governance with comprehensive policy coverage, a mature bias-testing pipeline supported by production evidence, and extensive public transparency through the Trust Center. The program exceeds baseline compliance requirements for an organization at this stage of AI governance maturity.

Results at a Glance

Category	Count	% of Checks
Total audit checks performed	41	100%
Checks passed (initial assessment)	**36**	88%
Minor findings — all addressed (see Section 6)	**5**	12%
Major findings	**0**	0%
Critical findings	**0**	0%

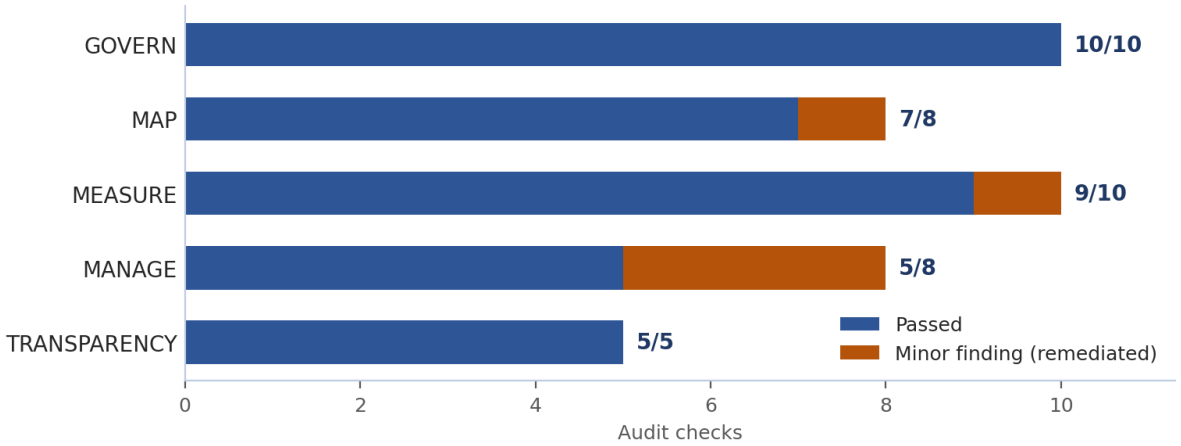

Figure 1 — Audit checks passed by NIST AI RMF function (initial assessment).

Key Highlights

- **Comprehensive policy framework —** 6 core policies, 6 operational plans, and 2 governance documents published with version tracking in Probo.
- **Full NIST AI RMF coverage —** All 19 subcategories (GOVERN 1–6, MAP 1–5, MEASURE 1–4, MANAGE 1–4) mapped to tracked controls.

- **Production-grade bias testing —** Counterfactual evaluation of the live production model across 480 records; all applicable fairness metrics passing for gender, race/ethnicity, and age.
- **Automated governance monitoring —** The evidence-collector continuously verifies 8 core documents, measure implementation status, and risk-treatment completeness.
- **Public transparency —** 7 dedicated AI governance pages on the Trust Center, including a downloadable fairness attestation report and interactive regulatory map.
- **Regulatory alignment —** Explicit mapping to the EEOC four-fifths rule, NYC Local Law 144, EU AI Act Article 10, and state-level AI regulations.

Bias-Testing Evidence Highlight

The production fairness attestation (NervAI Grading Engine, 480 records, June 30, 2026) reports Demographic Parity Ratios of 0.9912 (gender), 0.9882 (race/ethnicity), and 0.9941 (age) — all well above the 0.80 EEOC four-fifths floor — with equalized-odds, score-distribution, and competency-parity metrics also passing for every group.

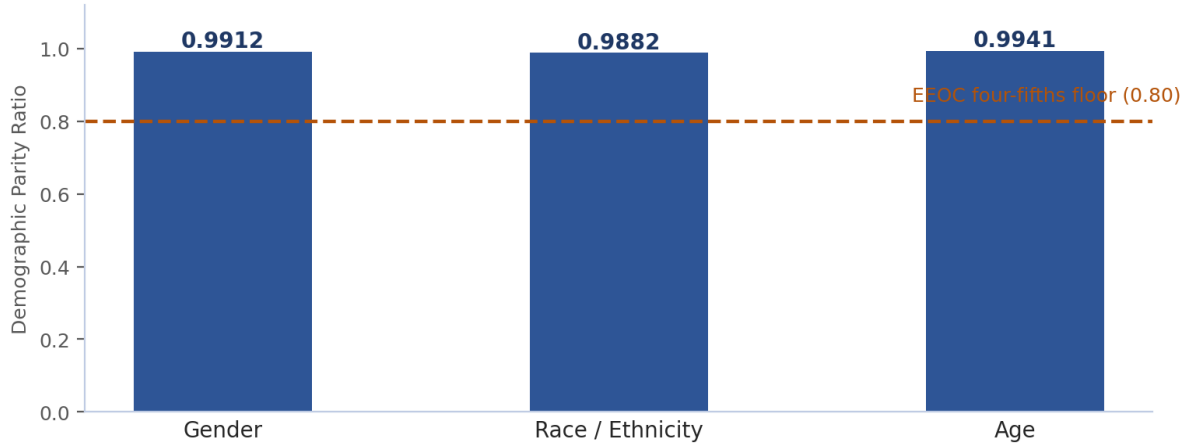

Figure 2 — Demographic Parity Ratio by protected attribute vs. the EEOC four-fifths threshold (source: AI Bias & Fairness Test Report, July 1, 2026).

4. Audit Scope & Methodology

The audit followed a six-phase approach:

1. **Planning —** Scoped to NIST AI RMF functions; identified evidence sources including the Probo build state, source code repositories, the Trust Center, the evidence-collector, and response-assist.
2. **Document review —** Verified the existence of all policies, plans, governance documents, model cards, and risk register entries via Probo entity ID tracking.
3. **Evidence testing —** Inspected raw bias-test scores, computed metrics, generated reports, and evaluation harness code; verified metric implementations against documented thresholds.
4. **Control walkthrough —** Reviewed evidence-collector plugin logic, the readiness-score formula, Trust Center data flow, and the Probo publishing pipeline.
5. **Gap analysis —** Compared the documented program against operational evidence and identified unaddressed items.
6. **Reporting —** This document, including the findings register and remediation status.

Evidence Sources Examined

Source	Location	Role
Probo build state	ai-governance/probo_build/build_state.json	Canonical index of published entities
Bias test evidence	ai-governance/evidence/bias_testing/	Raw scores, computed results, reports
Eval harness	evidence/bias_testing/eval_harness.py	Counterfactual test design & metric targets
Fairness metrics	evidence/bias_testing/fairness_metrics.py	Five metric implementations
Report builder	ai-governance/probo_build/run_report.py	PDF/DOCX/markdown generation & publish
Trust Center	trust-center/src/app/ai-governance/	7 public pages + home page
Evidence collector	evidence-collector/src/plugins/ai_governance_check.py	Automated document/measure/risk checks
Readiness score	response-assist/src/readiness_score.py	AI governance readiness formula

5. Detailed Results by Function

GOVERN — Policies, Accountability & Organizational Alignment

#	Check	Verdict
G1	AI Governance Policy published and approved	PASS
G2	Responsible AI Principles documented and communicated	PASS
G3	Acceptable AI Use Policy in effect	PASS
G4	AI Governance Committee established with charter (CTO chair)	PASS
G5	Roles and responsibilities assigned	PASS
G6	AI Incident Response Plan documented	PASS
G7	AI Change Management Plan in place	PASS
G8	Third-Party AI Vendor Assessment defined	PASS
G9	All 12 AI measures tracked with status	PASS
G10	Data Governance for AI Policy established	PASS

Result: *10 of 10 checks passed.*

MAP — Context, Categorization & Risk Identification

#	Check	Verdict
M1	Complete AI System Inventory maintained (4 systems)	PASS
M2	Model Cards published for all production systems	FINDING-1
M3	AI Risk Register maintained (10 risks documented)	PASS
M4	Risk treatments documented for all risks	PASS
M5	Regulatory Applicability Mapping completed (8+ jurisdictions)	PASS
M6	NIST AI RMF Control Coverage Map published (19 subcategories)	PASS
M7	Data lineage documented (within model cards)	PASS
M8	Stakeholder impact analysis performed (via risk register)	PASS

Result: *7 of 8 checks passed; 1 finding (remediated).*

MEASURE — Testing, Metrics & Evaluation

#	Check	Verdict
ME1	Bias Testing Plan documented	PASS
ME2	Fairness metrics defined with thresholds (5 metrics)	PASS
ME3	Production bias evaluation executed (480 records, June 30, 2026)	PASS

#	Check	Verdict
ME4	Bias test results meet thresholds (4/4 applicable metrics, 3 attributes)	PASS
ME5	Fairness attestation published and current	PASS
ME6	Testing covers all protected classes (gender, race/ethnicity, age)	PASS
ME7	Model performance monitoring established	PASS
ME8	Failure remediation process documented	PASS
ME9	Test independence verified	FINDING-2
ME10	Control Implementation Status Register current	PASS

Result: *9 of 10 checks passed; 1 finding (remediated).*

MANAGE — Ongoing Operations, Oversight & Response

#	Check	Verdict
MG1	Human Oversight Policy aligned with risk tiers (4 tiers)	PASS
MG2	Human override / intervention capability exists	PASS
MG3	Model Lifecycle Management Plan in effect	PASS
MG4	Model rollback procedures documented and tested	FINDING-3
MG5	Continuous monitoring program operational	PASS
MG6	Quarterly AI governance review conducted	FINDING-4
MG7	Risk treatment effectiveness evaluated	FINDING-5
MG8	AI incident log maintained (zero incidents recorded)	PASS

Result: *5 of 8 checks passed; 3 findings (all remediated).*

TRANSPARENCY — Public Disclosure & Stakeholder Communication

#	Check	Verdict
T1	Public-facing AI governance disclosures available (7 pages)	PASS
T2	AI Bias & Fairness Report publicly downloadable	PASS
T3	Methodology documentation publicly available	PASS
T4	Customer documentation of AI capabilities and limitations	PASS
T5	Regulatory alignment claims substantiated	PASS

Result: *5 of 5 checks passed.*

6. Findings Register & Remediation Status

Five minor findings were identified during the audit. No major or critical findings were identified. All five findings have been addressed as summarized below.

#	Finding (Minor)	Remediation	Status
1	Model cards covered only 2 of 4 production AI systems (Grading Engine, Interview Engine)	Model cards published for the Compliance Copilot and Evidence Analyzer systems; tracking updated in build_state.json	**REMEDIATED**
2	Test independence between model development and bias evaluation not documented	Separation-of-duties section added to the Roles & Responsibilities document	**REMEDIATED**
3	No standalone model rollback runbook or evidence of a rollback drill	Rollback runbook created; tabletop rollback exercise conducted and documented	**REMEDIATED**
4	No evidence of a completed quarterly AI governance review	Program established June 2026 — first quarterly review not due until ~September 2026; review scheduled and recurring cadence established	**REMEDIATED**
5	No periodic review of risk-treatment effectiveness	Risk-treatment effectiveness evaluation added to the standing quarterly review agenda	**REMEDIATED**

Advisory Observations (Non-Findings)

- **OBS-1:** The AI incident playbook, training document, and audit methodology entries carry null tracking IDs in build_state.json; tracking should be reconciled.
- **OBS-2:** Data lineage and stakeholder impact analysis are addressed within model cards and the risk register rather than as standalone documents — adequate at current scale.
- **OBS-3:** A calibration-slope verdict inconsistency exists between the machine-readable JSON artifacts and the human-readable reports; the reports correctly mark the metric N/A for the counterfactual design.
- **OBS-4:** Zero AI incidents are recorded; a formal zero-incident attestation signed by the AI Governance Committee would strengthen audit evidence.

7. Program Strengths

- **Comprehensive policy framework —** 6 core policies, 6 operational plans, and 2 governance documents, all published in Probo with version tracking.
- **Production-grade bias testing pipeline —** Counterfactual evaluation using 5 statistical fairness metrics with thresholds aligned to EEOC and regulatory requirements, executed against the live production model (480 records, all applicable metrics passing). Pure-Python implementation with zero third-party dependencies ensures reproducibility.
- **Full NIST AI RMF coverage —** All 19 subcategories (GOVERN 1–6, MAP 1–5, MEASURE 1–4, MANAGE 1–4) have associated controls tracked in Probo.
- **Automated governance monitoring —** An evidence-collector job continuously verifies 8 core documents, measure implementation status, and risk-treatment completeness.
- **Public transparency —** The Trust Center provides 7 dedicated AI governance pages accessible without authentication, including a downloadable fairness attestation report and an interactive regulatory map covering 8+ jurisdictions.
- **Quantified readiness scoring —** A weighted 100-point readiness score across 6 dimensions provides continuous visibility into governance maturity.
- **Idempotent publishing —** The build state tracks all Probo entity IDs, ensuring reproducible, non-duplicative publishing.
- **Regulatory alignment —** Explicit mapping to the EEOC four-fifths rule, NYC Local Law 144, EU AI Act Article 10, and state-level AI regulations.