NIST AI RMF Profile

ClearPath Pharmacy: AI Governance Portfolio

Document information	Details
Organization	ClearPath Pharmacy: Specialty Pharmacy Services
Document title	NIST AI RMF Profile: Patient Adherence and Outreach Prioritization Tool
Document ID	AI-RMF-001
Classification	Confidential: Internal Use Only
Version	1.0
Effective date	May 18, 2026
Next review date	May 18, 2027
Prepared by	Information Security & Compliance Department
Approved by	Information Security Steering Committee
AI system in scope	Patient Adherence and Outreach Prioritization (AOP) tool

AI use case in scope

System: The Patient Adherence and Outreach Prioritization (AOP) tool, a third-party predictive analytics system supplied by MedSignal Analytics under a Business Associate Agreement. The system ingests patient data, predicts which patients are at risk of missing or stopping therapy, and ranks them so care staff know whom to contact first.

Why this use case matters: It processes protected health information, it influences decisions that affect patients (who receives outreach first), and adherence is a core clinical and business concern for a specialty pharmacy. That places it squarely at the intersection of HIPAA and AI governance.

The HIPAA / AI overlap (the thread through every document): Wherever the AOP tool processes PHI, HIPAA Privacy and Security Rule obligations apply alongside the AI governance work. Minimum necessary, a signed BAA, and access controls are AI governance controls and HIPAA controls at the same time.

1. Purpose and Methodology

This profile applies the NIST AI Risk Management Framework (AI RMF 1.0) to one concrete system: the AOP tool. A profile works through the framework's four core functions, Govern, Map, Measure, and Manage, for a single use case, identifies the gap between current and target state, and recommends actions. It builds directly on ClearPath's existing GRC work: the HIPAA Security Risk Assessment (NIST SP 800-30), the NIST CSF 2.0 Gap Analysis, the Security Policy Package, and the AI Use Policy (POL-AI-001).

2. System and Use-Case Description

The AOP tool is a supervised predictive model operated by MedSignal Analytics and integrated with ClearPath's pharmacy management system. It scores active specialty-therapy patients on their risk of non-adherence and produces a ranked outreach list for care coordinators. Outputs inform, but do not automate, patient outreach. No dispensing, clinical, or coverage decision is made by the tool.

3. Stakeholders and Affected Parties

- Patients: most affected; outreach priority can influence the support they receive.
- Care coordinators and pharmacists: primary users who act on the ranked list.
- Compliance and privacy officer: accountable for HIPAA and AI governance alignment.
- Vendor (MedSignal Analytics): builds, hosts, and maintains the model as a Business Associate.
- Leadership / Information Security Steering Committee: owns AI risk acceptance.

4. Function: Govern

Govern is cross-cutting: the policies, accountability, roles, and culture for managing AI risk. It asks who owns the tool, what oversight exists, and how AI risk decisions get made.

Current state	Gap	Recommended action
AI Use Policy (POL-AI-001) exists and classifies the AOP tool as a Tier 2 use; ISMS	No named owner for the AOP system specifically; no standing AI governance committee; no	Name an AOP system owner; charter an AI governance committee

Current state	Gap	Recommended action
and HIPAA program are established.	maintained AI system inventory.	under the existing Steering Committee; stand up an AI system inventory.
Vendor is under a signed BAA.	AI-specific vendor obligations (fairness reporting, drift monitoring, model documentation) are not yet in the contract.	Add AI-specific clauses to the vendor agreement at next renewal; require model documentation.

5. Function: Map

Map establishes context and identifies risks: what the system does, who is affected, the data it uses, and what could go wrong.

Current state	Gap	Recommended action
Use case, data flows, and PHI elements are generally understood by the team.	No formal AI impact assessment; data lineage and the specific PHI fields used by the model are not documented.	Complete an AI system impact assessment; document data sources, PHI elements, and the model's intended and prohibited uses.
Obvious risks (PHI exposure) are recognized.	Fairness, over-reliance, and drift risks are not formally identified or owned.	Adopt the risk register in Section 9; assign an owner to each risk.

6. Function: Measure

Measure analyzes, assesses, and tracks the identified risks: how accuracy, fairness, and drift are measured and monitored.

Current state	Gap	Recommended action
Vendor reports overall model accuracy at onboarding.	No subgroup fairness testing; no drift monitoring; no agreed metrics or thresholds; no independent validation.	Define accuracy, fairness (subgroup), and drift metrics with thresholds; require periodic vendor reporting; review at the AI governance committee.
PHI access is logged through existing controls.	AI-specific monitoring (who acted on rankings, outcome tracking) is absent.	Add outcome tracking so the program can tell whether

Current state	Gap	Recommended action
		prioritization improved adherence equitably.

7. Function: Manage

Manage prioritizes and acts on risks: mitigations, human oversight, escalation, and ongoing monitoring.

Current state	Gap	Recommended action
Care staff currently apply judgment when using the list.	Human-in-the-loop is informal and not required by procedure; no escalation path when the tool looks wrong.	Mandate documented human review before outreach decisions; define an escalation and override path; log overrides.
Incident Response Policy (CP-POL-003) exists.	No AI-specific incident handling (e.g., discovered bias, drift breach, vendor data incident).	Extend the IR plan with AI incident triggers; tie PHI-involving AI incidents to the breach-notification branch.

8. Trustworthy AI Characteristics Check

The NIST AI RMF defines seven characteristics of trustworthy AI. Each is assessed for the AOP tool below.

Characteristic	Status	Assessment and action
Valid and reliable	Partial	Accuracy stated by vendor but not independently validated or monitored for drift. Action: define metrics and periodic validation.
Safe	Adequate	Tool informs outreach only; it cannot make clinical or dispensing decisions. Maintain that boundary in policy.
Secure and resilient	Adequate	Covered by existing ISMS controls and the BAA. Confirm vendor security posture at reassessment.
Accountable and transparent	Gap	No named AI owner and limited documentation. Action: assign ownership; maintain model and decision documentation.

Characteristic	Status	Assessment and action
Explainable and interpretable	Gap	Staff cannot explain individual rankings. Action: require vendor to provide reason codes or key factors per score.
Privacy-enhanced	Partial	BAA in place; minimum-necessary and de-identification not yet verified for the model's inputs. Action: confirm minimum-necessary data set.
Fair, with harmful bias managed	Gap	No subgroup fairness testing. Action: require fairness testing and ongoing subgroup monitoring.

9. Risk Summary

The top AI risks for the AOP tool, carried consistently across this portfolio.

ID	Risk	Description	Severity
AR1	Biased or inequitable prioritization	The model under-ranks certain patient groups (by language, geography, payer, or proxies for protected attributes), so some patients systematically receive less outreach.	High
AR2	PHI exposure through the AI vendor	PHI is processed by a third-party model and infrastructure, raising privacy, minimum-necessary, and BAA concerns.	High
AR3	Automation bias / over-reliance	Care staff defer to the ranking instead of applying clinical judgment, missing patients the model scored low.	Medium
AR4	Model drift	Accuracy degrades as the patient population, therapies, or payer mix change, without anyone noticing.	Medium
AR5	Lack of transparency and explainability	Staff cannot explain why a patient was prioritized, undermining trust and the ability to challenge outputs.	Medium
AR6	Governance and accountability gaps	No clear owner of AI risk decisions or escalation path when the tool behaves unexpectedly.	Medium

10. Prioritized Recommendations

1. Assign an AOP system owner and charter the AI governance committee (Govern).
2. Complete an AI impact assessment and document data lineage (Map).
3. Require fairness testing and drift monitoring with thresholds (Measure).
4. Mandate and document human review before outreach decisions (Manage).
5. Add AI-specific clauses and reporting to the vendor agreement (Govern).