ENFORCEAUTH

GUIDE

How to map AI agent authorization decision logs to NIST AI RMF and HIPAA controls

Each authorization decision about an AI agent's action records who acted, for whom, what was attempted, whether it was allowed, and under which policy version. That record is evidence for HIPAA's access control, audit control and activity review requirements, and for the NIST AI RMF outcomes on accountability, human oversight and post-deployment monitoring. Map each control to the fields that support it, mask protected health information (PHI) before the log leaves OPA, and keep the records where they cannot be altered.

Decision logs support these controls. They do not satisfy them alone.

HIPAA Security Rule

Requirement (45 CFR)	What decision logs show
164.312(a)(1) Access control	Each agent action on ePHI was checked against policy, with the result
164.312(b) Audit controls	A record of activity in systems that contain or use ePHI, including agent tool calls
164.312(c)(1) Integrity	Writes to records were allowed under a named policy version. Pair with tamper-evident logs so the log itself can be trusted
164.312(d) Person or entity authentication	The verified user and agent identity on each decision
164.308(a)(1)(ii)(D) Information system activity review	Searchable records to review agent activity, denials and unusual volume

Source: 45 CFR 164.312 and 164.308.

NIST AI Risk Management Framework 1.0

The AI RMF has four functions: GOVERN, MAP, MEASURE and MANAGE. (NIST AI 100-1)

Subcategory	What decision logs and policy show
GOVERN 1.4: "The risk management process and its outcomes are established through transparent policies, procedures, and other controls"	Agent rules written as reviewed, versioned policy
GOVERN 2.1: roles, responsibilities and lines of communication are documented	Each agent has a named owner in policy data and in every record
MAP 3.5: "Processes for human oversight are defined, assessed, and documented"	Which calls need approval or notification, and the approvals given
MEASURE 2.8: "Risks associated with transparency and accountability ... are examined and documented"	Per-decision records with reasons, searchable by agent, user and tool
MANAGE 2.4: mechanisms "to supersede, disengage, or deactivate AI systems"	A policy change that denies an agent takes effect on the next call, and the denials are logged
MANAGE 4.1: "Post-deployment AI system monitoring plans are implemented"	Ongoing records of allows, denies and approvals per agent

Step by step

1. **Decide every agent action through policy**, so every action has a record.
2. **Record the fields auditors need:** user, agent, owner, tool, resource, result, reasons, approval, policy version and time.
3. **Mask PHI before it leaves OPA.** Use `system.log.mask` to remove or replace sensitive input fields. The decision log is itself a record that can hold PHI, so restrict who can read it.
4. **Keep records tamper-evident** and retained for your required period.
5. **Write the mapping down** in your compliance documentation, control by control, and agree it with your auditor.

Example: mask PHI in decision logs

Tested with OPA v1.22.4 (2 of 2 tests passed).

```
REGO
package system.log

import rego.v1

# Runs inside OPA before a decision log event leaves the process.
# Remove values you must not store; keep the fields auditors need.
mask contains "/input/arguments/patient_dob"

mask contains "/input/arguments/ssn"

mask contains {"op": "upsert", "path": "/input/arguments/note", "value": "[redacted]"} if {
          input.input.arguments.note
}
```

Where EnforceAuth fits

EnforceAuth is the OPA/Rego authorization control plane, in the category we call Authorization Context Intelligence. Your own OPA runtimes make each runtime decision. EnforceAuth does not make that decision. You author, test and ship the Rego and bundles in Writ. Engines send decisions to Writ, and Writ feeds EvidenceIQ, an EnforceAuth capability that keeps a durable, searchable record of authorization decisions. EvidenceIQ records decisions and never makes or changes them. Herald MCP, EnforceAuth's hosted MCP server, lets you operate EnforceAuth from Claude, Cursor or ChatGPT. Your own OPA makes each runtime decision. Free tier: the first 1 million decisions a month are free, then $89 per additional million. (llms.txt, Pricing)

EnforceAuth deploys in the cloud region you choose, or self-hosted.

FAQ

How do I map AI agent decisions to the NIST AI RMF?

Use per-decision records as evidence for subcategories such as GOVERN 1.4 (transparent policies and controls), MAP 3.5 (human oversight processes), MEASURE 2.8 (transparency and accountability), MANAGE 2.4 (disengaging AI systems) and MANAGE 4.1 (post-deployment monitoring). Document the mapping control by control.

What HIPAA controls apply to AI agents accessing patient data?

The Security Rule's access control, audit controls, integrity, and person or entity authentication standards in 45 CFR 164.312, and information system activity review in 164.308. A per-call authorization decision log is evidence for each of

them.

Do decision logs make an AI agent HIPAA compliant?

No. They are evidence that supports specific controls. Compliance depends on your full program, including risk analysis, policies, training and agreements. Agree the evidence with your auditor.

Can authorization decision logs contain PHI?

Yes, if the input includes patient data. Mask sensitive fields before the event leaves OPA, for example with system.log.mask, and restrict access to the logs.