

HAPI OPEN STANDARD AND CERTIFICATION MODEL

Preserving Human Agency Without Owning the Gate

A Human Agency Preservation Infrastructure Working Thesis Paper

Author: Michael Bower

Project: Human Agency Preservation Infrastructure (HAPI)

Version: v0.1

Status: Working research, standards, and certification design manuscript, not legal advice

Core thesis: Human agency preservation should remain a public-interest standard, not a private monopoly. HAPI should keep its doctrine teachable and broadly adoptable while protecting the trust layer: its name, certification marks, audit quality, implementation discipline, and institutional credibility.

Table of Contents

1. Introduction
 2. The Problem of Owning the Gate
 3. Open Doctrine and Protected Trust
 4. The HAPI Open Standard
 5. What Must Remain Public
 6. What May Be Protected
 7. Certification as Proof, Not Permission
 8. Certification Levels
 9. Evidence Requirements
 10. Avoiding Certification Theater
 11. Independent Audit Governance
 12. Defensive IP Without Mission Capture
 13. Compatibility With Commercial Implementation
 14. Public Standard Roadmap
 15. Risks and Failure Modes
 16. The HAPI Open Standard Statement
 17. Conclusion
- Appendix A: Certification Principles
- Appendix B: Sample Certification Claims
- Appendix C: Open Standard Checklist

1. Introduction

Human Agency Preservation Infrastructure, or HAPI, is built on a boundary claim: systems should preserve and restore the human capacity to understand, choose, refuse, participate, repair, and remain meaningfully responsible. Because that claim is bigger than any single product, the question of ownership matters.

A framework created to preserve agency can be corrupted if it becomes a gate that others must pass through without legitimate reason. A company can own software. A foundation can protect a name. A certification body can control the use of a mark. But the moral function of preserving human agency should not be reduced to a proprietary toll booth.

This paper develops a model for HAPI as an open standard and certification system. The principles remain open. The official trust layer is protected. Commercial implementation is allowed. Certification is earned by evidence, not purchased by reputation. The goal is to prevent agency preservation from becoming agency capture.

The central distinction is simple: HAPI should not own the gate as domination. HAPI should steward the gate as responsibility.

2. The Problem of Owning the Gate

A gate is a boundary that determines what may pass. In the HAPI framework, a true gate preserves agency by governing passage according to rightful authority, evidence, context, and accountability. A false gate captures agency by blocking what it has no authority to block or permitting what it has no authority to permit.

If a company tries to monopolize the entire idea of agency-preserving governance, it risks becoming the very pattern HAPI is meant to diagnose. The problem is not ownership itself. The problem is ownership becoming control over a public-interest function.

The danger can be stated plainly:

- If HAPI claims that no one may preserve agency without permission, HAPI becomes a false gate.
- If HAPI protects nothing, the mission can be copied, diluted, misrepresented, or turned into theater.
- If HAPI protects only the brand and implementation quality, while keeping the principles open, the mission can scale without becoming capture.

This paper argues for the third path: open principles, protected trust, commercial implementation, and evidence-based certification.

3. Open Doctrine and Protected Trust

HAPI should separate doctrine from trust. Doctrine is the public teaching: what agency is, how agency is lost, how systems capture agency, what false participation looks like, and how governance restores agency. Trust is the institutional layer: who is allowed to claim official certification, official audit status, official HAPI compliance, or official use of the HAPI mark.

Open doctrine allows people, institutions, researchers, churches, schools, companies, hospitals, and public systems to learn the framework without asking permission. Protected trust prevents low-quality or bad-faith actors from claiming HAPI approval without evidence.

The distinction is:

- Doctrine should be open enough to teach and use.
- Standards should be public enough to inspect and criticize.
- Certification should be protected enough to mean something.
- Implementation should be commercial enough to sustain the work.
- The mission should be independent enough to refuse capture.

This structure lets HAPI become useful without becoming possessive. It also lets the company and foundation earn revenue through real work, not through ownership of the moral boundary itself.

4. The HAPI Open Standard

The HAPI Open Standard should define the minimum conditions under which an institution, product, workflow, or AI deployment can claim to preserve human agency. The standard should not require every organization to use HAPI software. It should require organizations to demonstrate agency-preserving behavior through evidence.

The open standard should answer five questions:

1. Is human agency being preserved, restored, captured, simulated, or stripped away?
2. Can affected humans understand what is happening and why?
3. Can affected humans refuse, contest, revise, or escalate before consequence binds?
4. Does the system preserve memory, accountability, and proof?
5. Does governance reduce agency loss over time, or does it only document compliance?

The standard should be public, versioned, and criticizable. It should evolve as HAPI learns from audits, case studies, implementation results, and public feedback.

5. What Must Remain Public

Some parts of HAPI should remain public-interest doctrine. Keeping them public strengthens the movement, establishes prior art, prevents monopoly over the moral frame, and invites adoption across sectors.

The public layer should include:

- Core definitions of agency, agency loss, agency capture, agency theater, false gates, true gates, and restorative governance.
- General HAPI principles for preserving understanding, refusal, participation, accountability, memory, and repair.
- Public education materials explaining how people and institutions lose agency.
- Open checklists that help individuals identify agency-stripping systems.
- Public standards for what agency-preserving governance should demonstrate.
- Research papers that develop the theory and invite critique.
- High-level audit categories and evaluation questions.

- Ethical statements about avoiding mission capture and false gate behavior.

The public layer allows HAPI to become a shared language. The point is not to make everyone dependent on one organization. The point is to help civilization recognize and restore human agency more clearly.

6. What May Be Protected

Keeping doctrine open does not mean leaving the institution defenseless. HAPI can protect the parts that preserve trust, quality, continuity, and accountability. Without protection, the name could be misused, certification could be counterfeited, and weak implementations could damage the mission.

The protected layer may include:

- The HAPI name, logo, certification marks, and official labels.
- Specific software products, dashboards, audit tooling, report generators, and implementation code.
- Official audit methodologies that require training, quality control, and reviewer discipline.
- Certification review processes and authorized certification marks.
- Templates, report formats, training materials, and enterprise implementation packages.
- Quality-control procedures for auditors, assessors, partners, and certified implementers.
- Narrow technical inventions or defensive IP where protection prevents capture by bad actors.

The guiding rule is that HAPI should protect the trust layer, not privatize the moral layer. The brand can be protected. The certification can be protected. The software can be commercial. The principle of preserving human agency should remain broadly teachable.

7. Certification as Proof, Not Permission

Certification should not mean that an organization has permission to preserve agency. Any organization can attempt to preserve agency. Certification means that a specific organization, system, workflow, or deployment has produced evidence that it meets a defined HAPI standard.

This distinction prevents certification from becoming a false gate. HAPI does not say: “You may not preserve agency unless we approve you.” HAPI says: “You may not claim official HAPI certification unless you have demonstrated the required evidence.”

Certification should therefore function as proof, not permission. It should indicate that an assessment occurred, evidence was reviewed, standards were applied, limitations were disclosed, and claims were bounded.

A certification mark should never be a decorative symbol. It should be a compact public claim backed by reviewable evidence.

8. Certification Levels

HAPI certification can become more useful if it recognizes different levels of maturity. Not every organization will meet the highest standard immediately. A level system allows organizations to enter the process honestly without pretending to be fully mature.

Suggested certification levels:

6. HAPI Aligned Intent: the organization has adopted agency-preservation principles and begun mapping agency risks, but evidence is early.
7. HAPI Assessed: the organization has completed an agency audit and received documented findings, limitations, and restoration recommendations.
8. HAPI Remediation in Progress: the organization has accepted findings and is implementing specific agency-restoring changes.
9. HAPI Certified Workflow: a specific workflow has demonstrated agency-preserving controls, evidence, and review practices.
10. HAPI Certified Deployment: a specific AI, institutional, or operational deployment has demonstrated agency preservation under defined scope.
11. HAPI Mature Governance: the organization shows repeated evidence of restored agency, accountable governance, continuous improvement, and reduced theater over time.

The levels should be bounded by scope. An organization should not be able to certify one workflow and imply that the whole institution is agency-preserving. Scope clarity is essential to prevent reputation laundering.

9. Evidence Requirements

Certification must rest on evidence. If the evidence is weak, the certification becomes theater. HAPI should require artifacts that show not only policy language, but actual agency-preserving behavior.

Evidence may include:

- Authority maps showing who can approve, refuse, revise, or escalate consequential decisions.
- Workflow maps showing where humans remain meaningful participants rather than symbolic approvers.
- User or affected-person documentation explaining rights, choices, contestability, and limits.
- Logs, receipts, or audit trails showing what was proposed, approved, refused, changed, and executed.
- Review packets showing that humans had enough context to make meaningful decisions.
- Examples of refusals, reversals, corrections, and escalations actually changing outcomes.
- Evidence that human review happens before consequence, not after commitment.
- Metrics showing whether agency improved, not merely whether compliance was documented.
- Remediation plans for unresolved agency risks.
- Limitations statements that prevent overclaiming.

The strongest evidence is behavioral evidence: cases where a person could understand, refuse, contest, revise, repair, or redirect the system before meaningful consequence occurred.

10. Avoiding Certification Theater

Certification theater happens when a symbol of trust replaces the work that should justify trust. In the HAPI context, this is especially dangerous because the mission is to detect agency theater. The certification itself must not become a performance layer.

HAPI should prevent certification theater by requiring:

12. Scope-limited claims: certification applies only to the assessed system, workflow, or deployment.
13. Evidence summaries: public-facing claims should describe what was actually demonstrated.
14. Limitations: unresolved risks must remain visible.
15. Expiration: certification should expire unless renewed through reassessment.
16. Revocation: HAPI must be able to suspend or revoke certification when evidence no longer supports the claim.
17. No pay-to-pass: payment may fund assessment, but it cannot purchase the conclusion.
18. Conflict disclosure: assessors should disclose conflicts that could compromise judgment.
19. Outcome review: certification should look at whether agency is preserved in practice, not only whether policy exists.

The rule is direct: certification must preserve public trust by remaining harder to obtain than it is to market.

11. Independent Audit Governance

If HAPI becomes both a software company and a certification authority, it must manage conflicts carefully. The same organization may build tools, sell implementation support, and operate a certification standard, but certification decisions need independence from sales pressure.

Suggested governance controls include:

- Separate commercial implementation teams from certification review teams where practical.
- Use written evidence criteria that reduce reviewer discretion and sales influence.
- Create an appeals process for disputed findings without weakening the standard.
- Maintain a certification registry that lists scope, status, expiration, and limitations.
- Require reviewer training in agency preservation, false gates, agency theater, and evidence quality.
- Use periodic external review or advisory board oversight as the standard matures.
- Publish high-level methodology so the public can understand what certification means.

The purpose is not to create bureaucracy. The purpose is to keep certification from becoming a sales asset detached from reality.

12. Defensive IP Without Mission Capture

HAPI does not need to reject intellectual property entirely. The ethical question is how IP is used. IP can protect the mission from misuse, counterfeit certification, low-quality copying, or hostile capture. IP becomes harmful when it tries to monopolize the basic human right to preserve agency.

A defensive IP posture would say:

HAPI protects specific names, marks, tools, methods, and implementations to preserve trust and prevent misrepresentation. HAPI does not seek to own the general principle that human agency should be preserved.

This posture allows HAPI to protect itself without contradicting itself. It can defend the official mark. It can commercialize software. It can publish standards. It can refuse unauthorized certification claims. But it should avoid presenting itself as the owner of human agency preservation itself.

The moral boundary is: protect stewardship, not domination.

13. Compatibility With Commercial Implementation

An open standard does not prevent a serious business. In many fields, open standards strengthen commercial ecosystems. Organizations still need experts, tools, training, integrations, dashboards, audits, reporting, remediation plans, and ongoing support.

HAPI can earn revenue through:

- Agency audits and implementation consulting.
- Enterprise governance software and dashboards.
- Certification assessment and renewal fees.
- Training programs for HAPI auditors and internal agency officers.
- Templates, reports, workshops, and executive briefings.
- Agentic AI governance infrastructure built on AAG, PGDL, Runtime Binding, and receipts.
- Public-sector and nonprofit agency restoration programs.
- Custom remediation plans for institutions with agency loss.

The business value comes from trusted execution, not from hoarding the idea. HAPI can be commercially strong because it is credible, useful, disciplined, and hard to fake.

14. Public Standard Roadmap

The open standard should mature gradually. A premature standard can become brittle, vague, or performative. HAPI should begin with a clear public doctrine, then turn repeated audit patterns into formal criteria.

A practical roadmap:

20. Publish v0.1 definitions: agency, agency loss, agency capture, agency theater, true gates, false gates, restoration, and governance.
21. Publish v0.1 HAPI principles: understanding, refusal, participation, accountability, memory, repair, proportionality, and anti-capture.
22. Create a public checklist for agency-preserving systems.
23. Develop a pilot audit methodology for AI deployments, healthcare, work, education, religious institutions, and public systems.
24. Run case studies and refine the standard from observed patterns.
25. Create an early certification rubric with scope limits and evidence requirements.

26. Pilot certification on friendly systems before public claims become broad.
27. Create a public registry and certification-mark policy.
28. Invite expert critique from governance, ethics, law, healthcare, education, AI safety, labor, and public policy communities.
29. Publish versioned standards with changelogs and sunset dates for outdated claims.

The standard should grow from evidence. The work should become more precise as HAPI sees more systems.

15. Risks and Failure Modes

The open standard and certification model has its own risks. Naming them early is part of preserving agency.

Key failure modes include:

- Certification theater: the mark becomes more important than the evidence.
- Mission capture: funders, clients, or partners pressure HAPI to weaken findings.
- Standard drift: the standard becomes so broad that it no longer means anything.
- Over-control: HAPI tries to own too much and becomes a false gate.
- Under-protection: HAPI protects too little and allows counterfeit or misleading claims.
- Commercial pressure: sales goals influence audit conclusions.
- Scope laundering: organizations certify one small workflow and market themselves as fully HAPI-aligned.
- Bureaucratic load: certification becomes so heavy that it strips agency instead of preserving it.
- Ideological capture: the standard gets absorbed by a partisan, sectarian, or institutional agenda rather than staying grounded in agency preservation.

The remedy is disciplined boundaries: open principles, protected marks, evidence-based certification, limited claims, public limitations, and a willingness to refuse money when money would compromise the mission.

16. The HAPI Open Standard Statement

The following statement can serve as the foundation for a public standards page:

HAPI treats human agency preservation as a public-interest principle, not a proprietary monopoly. The core doctrine of agency preservation should remain teachable, inspectable, and broadly adoptable. HAPI protects its name, certification marks, audit quality, implementation discipline, software, and official institutional role so that public trust is not diluted or misrepresented. HAPI does not seek to own the gate as domination. HAPI seeks to steward agency-preserving infrastructure as responsibility.

This statement gives HAPI a clear ethical posture. It invites adoption while protecting credibility. It refuses monopoly over the moral boundary while defending the integrity of the institution.

17. Conclusion

HAPI can become a public standard, a certification model, a business, a foundation, and a humanitarian movement only if it handles ownership carefully. The work exists to preserve

agency. Therefore, the organization must not structure itself in a way that captures the very agency it claims to protect.

The solution is not to reject ownership entirely. The solution is to own the right things for the right reasons. Protect the mark. Protect the quality. Protect the certification. Protect the software. Protect the institution from misrepresentation. But keep the principles open enough for the world to learn and apply.

A true gate cannot be owned as domination. It can only be stewarded as responsibility.

This is the posture HAPI should take: open doctrine, protected trust, commercial implementation, evidence-based certification, and public-interest stewardship.

Appendix A: Certification Principles

30. Certification should prove evidence, not sell reputation.
31. Certification should be scoped, dated, and reviewable.
32. Certification should include limitations, not only strengths.
33. Certification should be revocable if evidence no longer supports the claim.
34. Certification should preserve agency in the certification process itself.
35. Certification should distinguish intent, assessment, remediation, workflow certification, deployment certification, and mature governance.
36. Certification should be protected by marks and quality controls, but the agency-preservation doctrine should remain publicly teachable.
37. Certification should never become a false gate that blocks legitimate agency restoration outside HAPI.

Appendix B: Sample Certification Claims

Acceptable claim examples:

- This workflow has completed a HAPI Agency Assessment under v0.1 criteria.
- This AI deployment is HAPI Certified for the documented customer-support escalation workflow through the listed expiration date.
- This organization is implementing a HAPI remediation plan for agency-theater risks identified in the assessment.
- This product is designed to support HAPI-aligned agency preservation principles but is not HAPI certified.

Unacceptable claim examples:

- Our entire company preserves agency because one workflow was certified.
- HAPI approved our system with no limitations.
- We are HAPI certified because we donated to the foundation.
- Our AI is safe because it follows HAPI values, even though no assessment occurred.
- HAPI certification means humans cannot be harmed by this system.

The purpose of claim discipline is to prevent certification from becoming marketing theater.

Appendix C: Open Standard Checklist

A HAPI open standard checklist should ask:

38. Does the system preserve the human ability to understand what is happening?
39. Does the system preserve the human ability to refuse before consequence?
40. Does the system preserve the human ability to contest or appeal after consequence?
41. Does the system preserve meaningful participation rather than symbolic presence?
42. Does the system preserve responsibility by keeping authority and accountability connected?
43. Does the system preserve memory through receipts, records, explanations, or reviewable evidence?
44. Does the system reduce dependency capture over time?
45. Does the system identify and remediate false gates?
46. Does the system avoid turning audit or certification into theater?
47. Does the system improve human agency in practice, not only in policy language?