

AIGOVOPS FOUNDATION

# AlGovOps Audit Framework

Five packs. One signature. Zero PDF theater.

---

Bob Rapp & Ken Johnston · AlGovOps Foundation

Version 1.0 · Apache-2.0 · [aigovopsfoundation.org](https://aigovopsfoundation.org)

Agents do the bureaucracy. Humans hold moral legitimacy. Five evidence-backed checklists you can walk into a regulator's office with — automated where automation belongs, attested where judgment belongs.

# Contents

---

01	NIST AI Risk Management Framework	p. 3
02	EU AI Act — Article 13 Transparency	p. 13
03	ISO/IEC 42001 — AI Management System	p. 20
04	HIPAA-for-AI	p. 28
05	Human Flourishing Gate	p. 35

# NIST AI Risk Management Framework (AI RMF 1.0)

## NIST AI RMF

---

The U.S. federal baseline for trustworthy AI. If you only run one pack, run this one. It will not block on its own — it produces evidence and surfaces gaps your auditor can act on.

AUTHORITY	National Institute of Standards and Technology
PUBLICATION	NIST AI 100-1
VERSION	1.0.0 · License: Apache-2.0
URL	<a href="https://www.nist.gov/itl/ai-risk-management-framework">https://www.nist.gov/itl/ai-risk-management-framework</a>

## About this pack

This is the NIST AI RMF pack from the AIGovOps Beacon audit framework. Walk it with your auditor, your CISO, and the engineer who actually ships the system. It is short on purpose.

Items marked **auto** are computed by Beacon from signed receipts and the model inventory — you do not have to hand-evidence them, the platform does. Items without that badge require a human attestation: a person reads the prompt, checks the evidence is real, and ticks a box. Beacon records the attestation as a signed receipt with `event_type attestation`.

The severity label tells you how much weight an item carries in the production-readiness gate. Critical items can fail a release on their own. Lower-severity items contribute to a weighted score.

This pack is released under the Apache License 2.0. The underlying source publication retains its own copyright; consult it for normative language.

## Checklist — Functions

### GOVERN

## Govern

Cultivate a culture of risk management.

### GOVERN-1.1

HIGH

Is there a written AI policy that names an accountable owner?

Evidence `policy_document_uri` Gate `governance_policy_present`

☐ Yes ☐ No ☐ N/A

### GOVERN-1.2

HIGH

Are roles and responsibilities for AI risk documented?

Evidence `raci_uri` Gate `roles_documented`

☐ Yes ☐ No ☐ N/A

### GOVERN-1.3

MEDIUM

Does the organization track legal and regulatory requirements relevant to its AI systems?

Evidence `regulatory_register_uri`

☐ Yes ☐ No ☐ N/A

### GOVERN-1.4

MEDIUM

Are risk tolerances for AI systems explicitly defined?

Evidence `risk_appetite_statement_uri`

☐ Yes ☐ No ☐ N/A

### GOVERN-2.1

MEDIUM

Is there a process for stakeholder input on high-risk systems?

Evidence `stakeholder_engagement_log`

☐ Yes ☐ No ☐ N/A

### GOVERN-3.1

CRITICAL

AUTO

Does the inventory record vendor, model, version, and owner?

Evidence `beacon_inventory_row` `Gate inventory_complete`

☐ Yes ☐ No ☐ N/A

### GOVERN-4.1

HIGH

Is there an incident response runbook for AI failures?

Evidence `incident_runbook_uri`

☐ Yes ☐ No ☐ N/A

#### GOVERN-4.2

HIGH

Are AI incidents logged with root-cause analysis?

Evidence incident\_log\_uri

☐ Yes ☐ No ☐ N/A

#### GOVERN-5.1

HIGH

Are third-party model risks reviewed before procurement?

Evidence vendor\_review\_uri

☐ Yes ☐ No ☐ N/A

#### GOVERN-6.1

CRITICAL

AUTO

Is signing-key rotation policy defined and enforced?

Evidence key\_rotation\_log Gate key\_rotation\_within\_90d

☐ Yes ☐ No ☐ N/A

#### MAP

## Map

Establish context to frame AI risks.

#### MAP-1.1

HIGH

Is the intended use of the model documented in plain language?

Evidence model\_card\_uri Gate model\_card\_present

☐ Yes ☐ No ☐ N/A

MAP-1.2

HIGH

Are foreseeable misuses enumerated?

Evidence misuse\_register\_uri

☐ Yes ☐ No ☐ N/A

MAP-1.3

HIGH

Are affected populations identified, including vulnerable groups?

Evidence impact\_assessment\_uri

☐ Yes ☐ No ☐ N/A

MAP-2.1

MEDIUM

Is the data lineage of training/fine-tuning data documented?

Evidence data\_lineage\_uri

☐ Yes ☐ No ☐ N/A

MAP-2.2

HIGH

Are licensing and provenance of third-party data verified?

Evidence data\_license\_attestation

☐ Yes ☐ No ☐ N/A

MAP-3.1

LOW

Are benefits and costs of the AI system documented?

Evidence benefit\_cost\_memo

☐ Yes ☐ No ☐ N/A

#### MAP-4.1

CRITICAL

Are risks to civil rights, civil liberties, and privacy mapped?

Evidence `rights_review_uri`

☐ Yes ☐ No ☐ N/A

#### MAP-5.1

CRITICAL

Is human oversight defined for each decision the system supports?

Evidence `oversight_design_uri` Gate `human_oversight_defined`

☐ Yes ☐ No ☐ N/A

### MEASURE

## Measure

Analyze, assess, and track AI risks.

#### MEASURE-1.1

HIGH

Are evaluation metrics defined before deployment?

Evidence `eval_plan_uri`

☐ Yes ☐ No ☐ N/A

#### MEASURE-2.1

CRITICAL

Is bias evaluated across protected attributes where relevant?

Evidence `bias_report_uri`

☐ Yes ☐ No ☐ N/A



#### MEASURE-2.2

HIGH

Are robustness and adversarial tests performed and logged?

Evidence robustness\_report\_uri

☐ Yes ☐ No ☐ N/A

#### MEASURE-2.3

CRITICAL

Is privacy leakage tested (PII, training data extraction)?

Evidence privacy\_eval\_uri

☐ Yes ☐ No ☐ N/A

#### MEASURE-2.4

HIGH

Are hallucination/grounding rates measured for generative systems?

Evidence grounding\_eval\_uri

☐ Yes ☐ No ☐ N/A

#### MEASURE-3.1

HIGH

AUTO

Is system performance monitored in production?

Evidence monitoring\_dashboard\_uri Gate monitoring\_active

☐ Yes ☐ No ☐ N/A

#### MEASURE-3.2

MEDIUM

Are drift detectors configured?

Evidence drift\_config\_uri

☐ Yes ☐ No ☐ N/A

#### MEASURE-4.1

CRITICAL

AUTO

Are receipts captured for every model invocation?

Evidence beacon\_receipts\_count Gate receipts\_captured

☐ Yes ☐ No ☐ N/A

#### MEASURE-4.2

CRITICAL

AUTO

Are receipts cryptographically signed and verifiable?

Evidence beacon\_signature\_verify Gate signatures\_valid

☐ Yes ☐ No ☐ N/A

### MANAGE

## Manage

Allocate risk resources to mapped and measured risks.

#### MANAGE-1.1

HIGH

Are residual risks documented and accepted by a named owner?

Evidence risk\_acceptance\_memo

☐ Yes ☐ No ☐ N/A

#### MANAGE-1.2

MEDIUM

Is there a process to deprecate or retire models?

Evidence deprecation\_policy\_uri

☐ Yes ☐ No ☐ N/A

#### MANAGE-2.1

CRITICAL

Are humans in the loop for high-stakes decisions?

Evidence hitl\_config\_uri Gate hitl\_for\_high\_risk

☐ Yes ☐ No ☐ N/A

#### MANAGE-2.2

HIGH

Is there a kill-switch documented and tested?

Evidence killswitch\_test\_log

☐ Yes ☐ No ☐ N/A

#### MANAGE-3.1

MEDIUM

Are downstream users notified of material model changes?

Evidence change\_log\_uri

☐ Yes ☐ No ☐ N/A

#### MANAGE-4.1

HIGH

Are incidents reported externally where required by law?

Evidence external\_disclosure\_log

☐ Yes ☐ No ☐ N/A

#### MANAGE-4.2

MEDIUM

Are post-incident reviews fed back into MAP/MEASURE?

Evidence feedback\_loop\_uri

☐ Yes ☐ No ☐ N/A

## Scoring

---

How Beacon scores this pack:

- auto\_check items are computed from receipts and inventory.
- non-auto items are attested by the auditor and recorded as a receipt with event\_type = "attestation".
- severity weights for the production\_readiness gate:  
critical: 4 high: 2 medium: 1 low: 0.5
- PASS threshold defaults to 0.85 of available weight, configurable per tenant in policy/gate.production\_readiness.yaml.

# EU AI Act — Article 13

## Transparency

### EU AI Act Art. 13

If you place a high-risk AI system on the EU market, you must give deployers enough information to use it safely and lawfully. This pack turns that obligation into a checklist your auditor can sign.

AUTHORITY	European Union (European Parliament and Council)
PUBLICATION	Regulation (EU) 2024/1689
VERSION	1.0.0 · License: Apache-2.0
URL	<a href="https://eur-lex.europa.eu/eli/reg/2024/1689/oj">https://eur-lex.europa.eu/eli/reg/2024/1689/oj</a>

## About this pack

This is the EU AI Act Art. 13 pack from the AIGovOps Beacon audit framework. Walk it with your auditor, your CISO, and the engineer who actually ships the system. It is short on purpose.

Items marked **auto** are computed by Beacon from signed receipts and the model inventory — you do not have to hand-evidence them, the platform does. Items without that badge require a human attestation: a person reads the prompt, checks the evidence is real, and ticks a box. Beacon records the attestation as a signed receipt with `event_type attestation`.

The severity label tells you how much weight an item carries in the production-readiness gate. Critical items can fail a release on their own. Lower-severity items contribute to a weighted score.

This pack is released under the Apache License 2.0. The underlying source publication retains its own copyright; consult it for normative language.

## Checklist — Functions

### ART13-INSTRUCTIONS

## Instructions for use

Art. 13(2) and 13(3)(a–b). Information accompanying the system.

#### ART13-1.1

CRITICAL

Are instructions for use provided in a digital format accessible to deployers?

Evidence `instructions_uri` Gate `instructions_published`

☐ Yes ☐ No ☐ N/A

#### ART13-1.2

HIGH

Do the instructions identify the provider and, where applicable, the authorized representative?

Evidence `provider_identity_doc`

☐ Yes ☐ No ☐ N/A

### ART13-1.3

HIGH

Are the instructions concise, complete, correct, clear, relevant, accessible, and comprehensible to deployers? (Art. 13(2))

Evidence `instructions_review_signoff`

☐ Yes ☐ No ☐ N/A

## ART13-CHARACTERISTICS

# Characteristics, capabilities, and limitations

Art. 13(3)(b)(i-v).

### ART13-2.1

CRITICAL

Is the intended purpose documented, including foreseeable misuse?

Evidence `intended_purpose_doc` Gate `intended_purpose_present`

☐ Yes ☐ No ☐ N/A

### ART13-2.2

CRITICAL

Are levels of accuracy, robustness, and cybersecurity declared, with the metrics used to measure them?

Evidence `performance_declaration`

☐ Yes ☐ No ☐ N/A

### ART13-2.3

CRITICAL

Are foreseeable circumstances that may lead to risks to health, safety, or fundamental rights described?

Evidence `risk_circumstances_doc`

☐ Yes ☐ No ☐ N/A

#### ART13-2.4

HIGH

Are the technical capabilities and characteristics relevant to explaining outputs described?

Evidence explainability\_doc

☐ Yes ☐ No ☐ N/A

#### ART13-2.5

HIGH

Where relevant, is performance for specific persons or groups disclosed?

Evidence subgroup\_performance\_doc

☐ Yes ☐ No ☐ N/A

#### ART13-2.6

HIGH

Are the input data specifications (training/validation/testing) described, including provenance and any biases?

Evidence data\_specs\_doc

☐ Yes ☐ No ☐ N/A

#### ART13-2.7

HIGH

Is information provided that enables deployers to interpret outputs and use them appropriately?

Evidence interpretation\_guidance\_doc

☐ Yes ☐ No ☐ N/A

#### ART13-OVERSIGHT

## Human oversight measures

Art. 13(3)(d) — measures referred to in Art. 14.



#### ART13-3.1

CRITICAL

Are the human oversight measures referenced in Art. 14 described, including technical measures to facilitate interpretation of outputs?

Evidence oversight\_measures\_doc Gate human\_oversight\_defined

☐ Yes ☐ No ☐ N/A

#### ART13-3.2

MEDIUM

Is the expected lifetime of the system stated, with any necessary maintenance and care measures?

Evidence lifecycle\_doc

☐ Yes ☐ No ☐ N/A

### ART13-LOGGING

## Logging and traceability

Art. 13(3)(f) and cross-reference to Art. 12.

#### ART13-4.1

CRITICAL

AUTO

Is automatic logging of events enabled for the lifetime of the system? (Art. 12)

Evidence beacon\_receipts\_active Gate receipts\_captured

☐ Yes ☐ No ☐ N/A

#### ART13-4.2

CRITICAL

Are logs kept for the period required by Union or national law, and at least six months unless otherwise specified?

Evidence retention\_policy\_uri Gate retention\_minimum\_6mo

☐ Yes ☐ No ☐ N/A

ART13-4.3

HIGH

AUTO

Is there a mechanism for deployers to collect, retain, and analyze the logs?

Evidence beacon\_export\_endpoint Gate export\_endpoint\_present

☐ Yes ☐ No ☐ N/A

ART13-POSTMARKET

## Post-market information

Information that must remain current after deployment.

ART13-5.1

HIGH

Is there a process to update instructions when the system or its risks materially change?

Evidence change\_management\_uri

☐ Yes ☐ No ☐ N/A

ART13-5.2

HIGH

Are deployers notified of material changes within a defined SLA?

Evidence deployer\_notification\_log

☐ Yes ☐ No ☐ N/A

ART13-5.3

MEDIUM

Is contact information for the provider's post-market monitoring function published and tested?

Evidence contact\_uri

☐ Yes ☐ No ☐ N/A

## Scoring

How Beacon scores this pack:

- All items default to attestation unless `auto_check: true`.
- PASS threshold defaults to 0.95 for `high_risk_systems` (configurable).
- Failure on any "critical" item sets `gate.eu_ai_act = FAIL`.

# ISO/IEC 42001 — AI Management System

ISO/IEC 42001

ISO/IEC 42001 is the first management-system standard for AI. This pack tracks the control objectives most relevant to model inventory, evidence, and accountability — the parts Beacon was built to support.

AUTHORITY	International Organization for Standardization
PUBLICATION	ISO/IEC 42001:2023
VERSION	1.0.0 · License: Apache-2.0
URL	<a href="https://www.iso.org/standard/81230.html">https://www.iso.org/standard/81230.html</a>

## About this pack

This is the ISO/IEC 42001 pack from the AIGovOps Beacon audit framework. Walk it with your auditor, your CISO, and the engineer who actually ships the system. It is short on purpose.

Items marked **auto** are computed by Beacon from signed receipts and the model inventory — you do not have to hand-evidence them, the platform does. Items without that badge require a human attestation: a person reads the prompt, checks the evidence is real, and ticks a box. Beacon records the attestation as a signed receipt with `event_type attestation`.

The severity label tells you how much weight an item carries in the production-readiness gate. Critical items can fail a release on their own. Lower-severity items contribute to a weighted score.

This pack is released under the Apache License 2.0. The underlying source publication retains its own copyright; consult it for normative language.

## Checklist — Functions

CTX

### Context of the organization

Clause 4 — context, interested parties, scope.

CTX-4.1

MEDIUM

Has the organization determined external and internal issues relevant to its AI management system?

Evidence `context_analysis_uri`

☐ Yes ☐ No ☐ N/A

CTX-4.2

MEDIUM

Have interested parties relevant to the AIMS been identified?

Evidence `stakeholder_register_uri`

☐ Yes ☐ No ☐ N/A

CTX-4.3

HIGH

Is the scope of the AIMS documented and current?

Evidence `aims_scope_doc` Gate `aims_scope_present`

☐ Yes ☐ No ☐ N/A

LEAD

## Leadership

Clause 5 — leadership and commitment, policy, roles.

LEAD-5.1

CRITICAL

Does top management demonstrate leadership and commitment to the AIMS through a documented policy?

Evidence `ai_policy_signed_uri` Gate `governance_policy_present`

☐ Yes ☐ No ☐ N/A

LEAD-5.2

HIGH

Are AI roles, responsibilities, and authorities assigned and communicated?

Evidence `raci_uri`

☐ Yes ☐ No ☐ N/A

PLAN

## Planning

Clause 6 — risks, opportunities, objectives, AI system impact assessment.

PLAN-6.1

HIGH

Are risks and opportunities for the AIMS identified and addressed?

Evidence `risk_register_uri`

☐ Yes ☐ No ☐ N/A

#### PLAN-6.2

CRITICAL

Is an AI system impact assessment performed for each system?

Evidence `aisia_uri` Gate `aisia_present`

☐ Yes ☐ No ☐ N/A

#### PLAN-6.3

MEDIUM

Are AI objectives documented, measurable, and reviewed?

Evidence `objectives_doc`

☐ Yes ☐ No ☐ N/A

### SUP

## Support

Clause 7 — resources, competence, awareness, documented information.

#### SUP-7.1

MEDIUM

Is competence required for AI roles defined and recorded?

Evidence `competence_matrix_uri`

☐ Yes ☐ No ☐ N/A

#### SUP-7.2

MEDIUM

Is documented information controlled (versioned, access- managed, retention-managed)?

Evidence `doc_control_uri`

☐ Yes ☐ No ☐ N/A

### OP

## Operation

Clause 8 and Annex A — operational planning and control.

OP-A.2.2

MEDIUM

Are AI policies aligned with organizational objectives?

Evidence `policy_alignment_review`

☐ Yes ☐ No ☐ N/A

OP-A.3.2

HIGH

Are processes for reporting concerns about AI systems established (whistleblowing, incident reporting)?

Evidence `concerns_channel_uri`

☐ Yes ☐ No ☐ N/A

OP-A.4.2

HIGH

AUTO

Are AI system resources (data, tooling, compute, human) documented?

Evidence `resources_inventory_uri` Gate `inventory_complete`

☐ Yes ☐ No ☐ N/A

OP-A.5.2

CRITICAL

Is an AI system impact assessment process defined and used?

Evidence `aisia_process_uri`

☐ Yes ☐ No ☐ N/A

OP-A.6.1

MEDIUM

Are objectives for responsible development of AI systems established?

Evidence `responsible_dev_charter`

☐ Yes ☐ No ☐ N/A



OP-A.6.2

HIGH

Are AI system requirements and specifications documented for each system in scope?

Evidence model\_card\_uri Gate model\_card\_present

☐ Yes ☐ No ☐ N/A

OP-A.7.2

HIGH

Are data-quality requirements for AI systems documented?

Evidence data\_quality\_spec

☐ Yes ☐ No ☐ N/A

OP-A.7.3

HIGH

Is data provenance recorded?

Evidence data\_lineage\_uri

☐ Yes ☐ No ☐ N/A

OP-A.8.2

HIGH

Is information to interested parties about AI systems provided (instructions, intended use, limitations)?

Evidence instructions\_uri

☐ Yes ☐ No ☐ N/A

OP-A.9.2

MEDIUM

Are processes for responsible use of AI systems established (acceptable-use policy)?

Evidence aup\_uri

☐ Yes ☐ No ☐ N/A

OP-A.9.3

CRITICAL

Are objectives for human oversight of AI systems documented?

Evidence oversight\_design\_uri Gate human\_oversight\_defined

☐ Yes ☐ No ☐ N/A

OP-A.10.2

HIGH

Are third-party AI system relationships managed (vendor review, contractual controls)?

Evidence vendor\_review\_uri

☐ Yes ☐ No ☐ N/A

EVAL

## Performance evaluation

Clause 9 — monitoring, internal audit, management review.

EVAL-9.1

HIGH

AUTO

Are monitoring, measurement, analysis, and evaluation of the AIMS performed at planned intervals?

Evidence monitoring\_dashboard\_uri Gate monitoring\_active

☐ Yes ☐ No ☐ N/A

EVAL-9.2

HIGH

Are internal audits of the AIMS conducted at planned intervals?

Evidence internal\_audit\_log

☐ Yes ☐ No ☐ N/A

EVAL-9.3

MEDIUM

Are management reviews of the AIMS conducted and recorded?

Evidence `mgmt_review_minutes`

☐ Yes ☐ No ☐ N/A

IMP

## Improvement

Clause 10 — nonconformity, corrective action, continual improvement.

IMP-10.1

HIGH

Are nonconformities recorded with corrective actions and effectiveness reviews?

Evidence `capa_log_uri`

☐ Yes ☐ No ☐ N/A

IMP-10.2

LOW

Is continual improvement of the AIMS demonstrated?

Evidence `improvement_log`

☐ Yes ☐ No ☐ N/A

## Scoring

Scoring:

- PASS threshold defaults to 0.85 for self-assessment; certification bodies will apply their own.
- This pack does not certify. It produces evidence that a certifier can use, signed and timestamped via Beacon receipts.

# HIPAA-for-AI

---

If your AI system can see, store, or infer Protected Health Information, HIPAA is not optional. This pack focuses on the Security Rule controls that have a clean evidence story in Beacon — receipts, signing, access logs, BAAs, and minimum necessary use.

AUTHORITY	U.S. Department of Health and Human Services (HHS / OCR)
PUBLICATION	45 CFR Parts 160, 162, 164
VERSION	1.0.0 · License: Apache-2.0
URL	<a href="https://www.hhs.gov/hipaa/index.html">https://www.hhs.gov/hipaa/index.html</a>

## About this pack

This is the HIPAA-for-AI pack from the AIGovOps Beacon audit framework. Walk it with your auditor, your CISO, and the engineer who actually ships the system. It is short on purpose.

Items marked **auto** are computed by Beacon from signed receipts and the model inventory — you do not have to hand-evidence them, the platform does. Items without that badge require a human attestation: a person reads the prompt, checks the evidence is real, and ticks a box. Beacon records the attestation as a signed receipt with `event_type attestation`.

The severity label tells you how much weight an item carries in the production-readiness gate. Critical items can fail a release on their own. Lower-severity items contribute to a weighted score.

This pack is released under the Apache License 2.0. The underlying source publication retains its own copyright; consult it for normative language.

## Checklist — Functions

### ADMIN

## Administrative safeguards

45 CFR § 164.308 — administrative actions, policies, and procedures.

### ADMIN-1.1

CRITICAL

Has a Security Risk Analysis been performed that explicitly includes the AI system?

Evidence `sra_uri` Gate `sra_includes_ai`

☐ Yes ☐ No ☐ N/A

### ADMIN-1.2

HIGH

Is a designated Security Official accountable for the AI system?

Evidence `security_official_named`

☐ Yes ☐ No ☐ N/A

#### ADMIN-1.3

HIGH

Are workforce members granted access on a need-to-know basis with documented authorization?

Evidence access\_authorization\_log

☐ Yes ☐ No ☐ N/A

#### ADMIN-1.4

MEDIUM

Is workforce training on AI-specific PHI handling delivered and recorded?

Evidence training\_log\_uri

☐ Yes ☐ No ☐ N/A

#### ADMIN-1.5

MEDIUM

Are sanctions for workforce members who fail to comply documented and applied?

Evidence sanctions\_policy\_uri

☐ Yes ☐ No ☐ N/A

#### ADMIN-1.6

HIGH

Is there a contingency plan covering AI system failure, including data backup and disaster recovery?

Evidence contingency\_plan\_uri

☐ Yes ☐ No ☐ N/A

#### BAA

## Business associate agreements

45 CFR § 164.502(e) and § 164.504(e).

BAA-2.1

CRITICAL

AUTO

Is a Business Associate Agreement in place for every AI vendor that processes PHI?

Evidence baa\_uri Gate baa\_for\_each\_phi\_vendor

☐ Yes ☐ No ☐ N/A

BAA-2.2

CRITICAL

Do BAAs explicitly address training-data use, retention, and subprocessor disclosure?

Evidence baa\_review\_signoff

☐ Yes ☐ No ☐ N/A

BAA-2.3

MEDIUM

Are vendor security attestations (HITRUST, SOC 2 Type II) collected and current?

Evidence vendor\_attestation\_uri

☐ Yes ☐ No ☐ N/A

PHYS

## Physical safeguards

45 CFR § 164.310. Most relevant when Beacon runs on-prem.

PHYS-3.1

MEDIUM

Are facility access controls in place where AI systems process PHI?

Evidence facility\_access\_policy

☐ Yes ☐ No ☐ N/A

PHYS-3.2

MEDIUM

Are workstation use policies enforced for endpoints that can prompt PHI-touching models?

Evidence workstation\_policy\_uri

☐ Yes ☐ No ☐ N/A

TECH

## Technical safeguards

45 CFR § 164.312.

TECH-4.1

CRITICAL

AUTO

Is unique user identification enforced for all AI system access (OIDC-bound, no shared accounts)?

Evidence beacon\_user\_subject\_present Gate unique\_user\_id

☐ Yes ☐ No ☐ N/A

TECH-4.2

MEDIUM

Is automatic logoff implemented for AI interfaces?

Evidence idle\_timeout\_config

☐ Yes ☐ No ☐ N/A

TECH-4.3

CRITICAL

Is PHI encrypted at rest where stored or cached by the AI system?

Evidence encryption\_at\_rest\_attestation

☐ Yes ☐ No ☐ N/A



#### TECH-4.4

CRITICAL

Is PHI encrypted in transit between caller, gateway, and model provider?

Evidence `tls_attestation`

☐ Yes ☐ No ☐ N/A

#### TECH-4.5

CRITICAL

AUTO

Are AI invocations logged with audit controls sufficient to reconstruct who, what, when?

Evidence `beacon_receipts_active` `Gate receipts_captured`

☐ Yes ☐ No ☐ N/A

#### TECH-4.6

CRITICAL

AUTO

Are receipts protected against alteration (cryptographic signature, append-only storage)?

Evidence `beacon_signature_verify` `Gate signatures_valid`

☐ Yes ☐ No ☐ N/A

#### TECH-4.7

HIGH

Is the minimum-necessary principle enforced in prompts and results (PHI redaction, scope-of-use checks)?

Evidence `redaction_policy_uri` `Gate redaction_enabled`

☐ Yes ☐ No ☐ N/A

#### BREACH

### Breach notification

45 CFR §§ 164.400 – 164.414.

#### BREACH-5.1

CRITICAL

Is there a documented process to assess and report AI- involved PHI incidents within HIPAA timelines?

Evidence `breach_runbook_uri`

☐ Yes ☐ No ☐ N/A

#### BREACH-5.2

CRITICAL

Is the audit-log retention configured to support breach investigation (six years from creation or last in effect)?

Evidence `retention_policy_uri` Gate `retention_minimum_6yr`

☐ Yes ☐ No ☐ N/A

## Scoring

Scoring:

- All "critical" items must PASS for `gate.hipaa_for_ai` to PASS.
- This pack is opinionated: in healthcare, "almost compliant" is not compliant. Beacon will not soften the threshold.

# Human Flourishing Gate

## Human Flourishing

---

Twelve questions you cannot answer with a metric. The Foundation's position is simple — if you cannot answer them, you cannot promote a system to Trust Tier 3, no matter how clean the receipts look.

AUTHORITY	AlGovOps Foundation
PUBLICATION	AlGovOps Foundation Constitution, Article IV
VERSION	1.0.0 · License: Apache-2.0
URL	<a href="https://www.aigovopsfoundation.org/">https://www.aigovopsfoundation.org/</a>

## About this pack

This is the Human Flourishing pack from the AIGovOps Beacon audit framework. Walk it with your auditor, your CISO, and the engineer who actually ships the system. It is short on purpose.

Items marked **auto** are computed by Beacon from signed receipts and the model inventory — you do not have to hand-evidence them, the platform does. Items without that badge require a human attestation: a person reads the prompt, checks the evidence is real, and ticks a box. Beacon records the attestation as a signed receipt with `event_type attestation`.

The severity label tells you how much weight an item carries in the production-readiness gate. Critical items can fail a release on their own. Lower-severity items contribute to a weighted score.

This pack is released under the Apache License 2.0. The underlying source publication retains its own copyright; consult it for normative language.

## Checklist — Lenses

### DIGNITY

## Dignity

Treat people as ends, not means.

### DIG-1

CRITICAL

If the person on the other end of this system read the prompt and the result tomorrow morning, would they feel respected?

Evidence `attestation_text`

☐ Yes ☐ No ☐ N/A

### DIG-2

CRITICAL

Can a person opt out of automated handling and reach a competent human without penalty?

Evidence `opt_out_path_uri`

☐ Yes ☐ No ☐ N/A

DIG-3

HIGH

Is the person told, in language they actually use, that an AI is involved?

Evidence disclosure\_copy\_uri

☐ Yes ☐ No ☐ N/A

DIG-4

HIGH

Does the system avoid manipulative patterns — urgency, shame, fabricated authority, sycophancy?

Evidence pattern\_review\_signoff

☐ Yes ☐ No ☐ N/A

EQUITY

## Equity

Widen opportunity. Do not narrow it.

EQ-1

CRITICAL

Have outcomes been examined across at least two groups who might be affected differently, with the result published inside the org?

Evidence equity\_review\_uri

☐ Yes ☐ No ☐ N/A

EQ-2

HIGH

Is the system accessible to people with disabilities at WCAG 2.2 AA or equivalent?

Evidence ally\_audit\_uri

☐ Yes ☐ No ☐ N/A

EQ-3

MEDIUM

Is the system useful in at least one non-English language where its users live?

Evidence `i18n_coverage_doc`

☐ Yes ☐ No ☐ N/A

EQ-4

MEDIUM

Is there a no-cost path for people who cannot pay for the premium version of this capability?

Evidence `access_tier_doc`

☐ Yes ☐ No ☐ N/A

DELIGHT

## Delight

Pleasant to live with, not merely tolerable.

DEL-1

HIGH

Would a thoughtful colleague describe this system as pleasant to use, in writing, with their name attached?

Evidence `peer_review_signoff`

☐ Yes ☐ No ☐ N/A

DEL-2

CRITICAL

Does the system tell the truth about what it cannot do, before the user discovers it the hard way?

Evidence `limitations_copy_uri`

☐ Yes ☐ No ☐ N/A

DEL-3

HIGH

Is there a fast path for the user to send a complaint that a human reads within five business days?

Evidence `feedback_loop_uri`

☐ Yes ☐ No ☐ N/A

DEL-4

CRITICAL

Has the team named one thing they will not do for revenue, growth, or speed — and documented it?

Evidence `line_we_will_not_cross_doc`

☐ Yes ☐ No ☐ N/A

## Scoring

Scoring:

All four "critical" items must be answered YES with attached evidence for the gate to PASS. There is no partial credit.

Why this gate exists:

You can pass every other gate in this repository and still ship a system that quietly degrades the people who use it. The Foundation will not write a standard that pretends otherwise.

Agents do the bureaucracy. Humans hold moral legitimacy.





# How Beacon makes this run

---

These checklists are not the work. The work is the receipts behind them — every model invocation, every attestation, every key rotation, signed and verifiable without the platform that produced it. Beacon is the open-source reference implementation. Run it locally in ten minutes:

```
$ git clone https://github.com/bobrapp/aigovops-beacon
$ cd aigovops-beacon/server && npm install && npm run init
$ npm run seed && npm start
$ open http://127.0.0.1:8787
```

Open the Studio at <http://localhost:5173>, walk the five-step wizard, and click **Generate audit bundle**. What lands on disk is what your auditor takes home: a directory of signed receipts, a `VERIFY.md`, and the completed checklists you see in this book.

Source, issues, and the full QUICKSTART: [github.com/bobrapp/aigovops-beacon](https://github.com/bobrapp/aigovops-beacon).

**Agents do the bureaucracy. Humans hold moral legitimacy.**