

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	AIGovOps Foundation
PUBLICATION	AIGovOps Foundation Constitution, Article IV
VERSION	1.0.0 · License: Apache-2.0
URL	https://www.aigovopsfoundation.org/

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.