Humanostic:
A Term for the Enforcement Posture That Binds Humans and AI Identically

Jordan Newell
Independent researcher

November 8, 2026

Abstract

As AI agents become operators — running workflows, merging code, spending budgets, taking actions that once required a person — governance needs vocabulary for a property it has never had to name: a rule enforced such that it applies identically whether the actor is human or machine. The literature reaches for "human-agnostic," in three incompatible senses, none of them enforcement-scoped; security practice reaches for "mandatory access control" and "technical controls," which name machine enforcement but never the human/AI symmetry. We coin **humanostic** (human + agnostic): of a guardrail, enforced below the conversational layer such that it binds human and AI operators identically, with no conversational override path. We give the term a contrast class (advisory, actor-differentiating, humanostic — a three-posture taxonomy of enforcement), trace its ancestry (rule-of-law constitutionalism; nondiscretionary access control), state precisely what is new in it (the symmetry claim: humans and AI agents as indifferently bound equals), demonstrate its retro-fit value on published arguments it compresses, and document that the coinage space is clear: no prior usage, no registry collisions, no squatters as of November 8, 2026.

Keywords: humanostic; enforcement; AI governance; guardrails; access control; rule of law; terminology

1 Definition

> **humanostic** (adj.) — of a guardrail, constraint, or gate: *enforced at the machine layer such that it binds human and AI operators identically, with no conversational override path.* Overrides, where they exist, are machine-mediated and leave audit trails.

Etymology: a portmanteau of *human* + *agnostic*, in the engineering lineage of *platform-agnostic* and *device-agnostic* — where *agnostic* means "valid regardless of the variable in question." Here the variable abstracted away is the operator's species. (The word deliberately echoes *anthropic* — Greek *anthropos*, "human" — in order to invert it: where *anthropic* centers the human, *humanostic* refuses to distinguish one.)

First use: Jordan Newell, live working session, November 8, 2026, describing a work-in-progress limit that "trips the same regardless of who's driving." This note timestamp-priorities the coinage (Zenodo DOI).

2 The contrast class: three postures of enforcement

The term earns its keep by naming one cell of a taxonomy that governance practice already uses but had never laid out:

Posture	Enforcement	Canonical examples
Advisory	prompt, policy document, or norm; binds only whoever complies	behavioral guidelines in system prompts; corporate AI-use policies; any "oversight" requirement with no machinery behind it
Actor-differentiating	machine-enforced, but binds differently by identity or content origin	access rules with named-actor bypass lists; CI rules that treat bot-suffixed accounts differently; permission systems that gate agents while waving humans through
Humanostic	machine-enforced; binds identically regardless of species	branch protection with `enforce_admins` (declines the repository owner's own direct push); server-side required checks; protocol-level rate limits

Two properties of the taxonomy matter. First, **enforcement does not entail symmetry**: the machine-enforced cell splits into the differentiated and the indifferent, so "humanostic" is not a synonym for "technical control" — it names a deliberate design choice within that category. Second, **the boundary is a trust boundary, not an absolute**: every humanostic gate is configured by a human upstream, and boundary-transcending overrides (config changes, vendor support channels) always exist. A humanostic gate is therefore one whose *operators' conversational channel* offers no override, and whose boundary-transcending overrides are themselves gated. This is the honest form of the property; the crisp binary fails against everyday escape hatches.

3 Ancestry, and what is actually new

Honesty requires citing the ancestors, because they are giants:

- **Rule-of-law constitutionalism.** "A government of laws and not of men" (John Adams, Massachusetts Constitution, 1780) states the shape — rules enforced by structure, indifferent to who holds power, overrides requiring formal process — 246 years early. *Humanostic* is not rule-of-law; it is its machinery translation, for the first governance domain where the posture is cheap to implement at every gate.
- **Nondiscretionary access control.** NIST defines mandatory access control (MAC) as policy decided by a central authority such that users — and, per the security-engineering literature, administrators — cannot override it. MAC already names machine enforcement that binds the powerful. It was, however, designed *against users*: the question "which species is bound?" never arose.
- **Policy-as-code / technical controls.** The AI-agent governance literature converges on the finding that advisory guidance alone cannot constrain agents at machine speed and must be backed by deterministic, externally enforced controls. This literature names the enforcement layer but never the symmetry property.

The novel claim, then, is narrow and should be stated narrowly: **humanostic names the symmetry** — that the bound actors include humans and AI agents *as equals*, indifferently. No established term encodes it, because no prior governance domain needed to; the species question is new.

4 Relation to incumbent "human-agnostic" usage

The existing phrase is used in three incompatible senses: (1) human-*independent*, for metrics and methods valid without human judgment; (2) human-*lacking*, for AI devoid of relational capacity; (3) human-*decentering*, for frameworks that drop the anthropocentric axiom. *Humanostic* claims only the first, sharpened from dispensing with humans to refusing to distinguish them — and substitution tests show it does not collapse into the incumbent: applied to sense-(3) usage it breaks the meaning, which is the disambiguation working.

The term's retro-fit value is strongest on the oversight-failure literature. Green's widely cited critique of human-oversight policies argues that nominal "human in the loop" requirements fail in practice; the failure compresses to one adjective: *the oversight mechanisms were advisory, not humanostic*. Practitioner writing converges on the same shape — "prompts alone cannot enforce," guardrail decisions "must be enforced by the surrounding workflow" — each phrase a paragraph where one word now stands.

5 Validation summary (November 8, 2026)

- **Novelty:** no established usage of "humanostic" located (web, literature, GitHub, package registries).
- **Retro-fit corpus:** 7 published/in-context sentences tested; 4 clean substitutions, 2 partial (scope-widening, disclosed), 1 principled misfit demonstrating disambiguation from the incumbent.
- **Necessity:** enforced-but-differentiating gates exist (named-actor bypass lists in forge rulesets; agent-gated permission prompts), so the indifferent posture is a real contrast class, not a byproduct.
- **Field proof:** branch protection with `enforce_admins` declining the repository owner's own direct push — a humanostic gate binding its author. Counter-example: an identical WIP limit enforced only by skill text, overridden in one conversational sentence.
- **Collision sweep:** npm, PyPI, GitHub (repos, users, orgs), and four TLDs — all clear, November 8, 2026.

References

1. Massachusetts Constitution, Part I, Art. XXX (drafted by John Adams, 1780): "to the end it may be a government of laws and not of men."
2. NIST Computer Security Resource Center glossary, "mandatory access control": access policy decisions made by a central authority; users cannot change access policy. `https://csrc.nist.gov/glossary/term/mandatory_access_control`
3. Green, B., "The flaws of policies requiring human oversight in computer algorithm decision-making," *Computer Law & Security Review* (2022).
4. Kiškis, M., "Legal framework for the coexistence of humans and conscious AI," *Frontiers in Artificial Intelligence* 6:1205465 (2023) — nearest-neighbor "human agnostic" usage, decentering sense.
5. Zilberman, J., "AGI imagined: social AGI and the co-formation of human-AGI society," *Frontiers in Artificial Intelligence* (2026) — "human-agnostic measure of comprehension applicable to natural and artificial agents alike."

6. AI-agent policy-enforcement literature: ValidMind, “AI governance tools for autonomous agents” (policy-as-code as the only governance model operating at agent speed); Sweet Security, agent policy-enforcement guidance (behavioral guidance is advisory and bypassable); Jackson (2025), multi-layer policy-as-code architecture for autonomous-agent risk and compliance.

7. GitHub Docs, “Managing a branch protection rule” (“restrict who can push” actor exemptions) and “About rulesets” (bypass lists for named users, teams, deploy keys); GitHub changelog, “Bypass branch protections with a new permission,” Aug 18, 2022 — evidence that enforcement does not entail species-indifference.

8. Companion paper: Newell, J., “The Missing Contributor Class: A Fourth Contributor Taxonomy for AI-Assisted Commits” (2026) — defines the disclosure layer this note’s enforcement taxonomy is orthogonal to; Zenodo concept DOI 10.5281/zenodo.24031381.