

# Deep Ethical Stack — Advocate, Arbiter, Overseer

ACCM Deep Ethics Project

## Contents

|                                                           |          |
|-----------------------------------------------------------|----------|
| <b>1 Deep Ethical Stack — Advocate, Arbiter, Overseer</b> | <b>1</b> |
| 1.1 Deep Ethics Advocate                                  | 1        |
| 1.2 Deep Ethical Arbiter                                  | 1        |
| 1.3 Overseer (M.o.E.)                                     | 2        |
| 1.4 Keep the roles mutually examinable                    | 2        |
| 1.5 Implementation status                                 | 2        |
| 1.6 Related pages                                         | 2        |

## 1 Deep Ethical Stack — Advocate, Arbiter, Overseer

### BETA — source-based editorial synthesis

John asks for the four Outnuancing definitions to run inside the full stack: **Advocate + Arbiter + Overseer + 10+1 + INTENT RESONATOR**. The archive then records Grok’s revised definitions. The following roles summarize that attributed response.

| Component            | Function in the supplied formulation                                                | Relation to the others                                     |
|----------------------|-------------------------------------------------------------------------------------|------------------------------------------------------------|
| Deep Ethics Advocate | Presses for correspondence when the current frame obstructs it.                     | Its own refusal or proposed alternative remains auditable. |
| Deep Ethical Arbiter | Audits the original frame and the intervention.                                     | Checks the Advocate’s move for new distortion.             |
| Overseer (M.o.E.)    | Watches whether meta-awareness becomes another comfortable performance.             | Places higher-order observation back inside the audit.     |
| 10+1                 | Supplies ongoing ethical and epistemic calibration.                                 | Its ingredients can correct one another.                   |
| INTENT RESONATOR     | Described as generative and calibrating, oriented to correspondence-seeking intent. | Connects intent, emerging distinctions, and correction.    |

### 1.1 Deep Ethics Advocate

The Advocate presses for correspondence. The Arbiter can examine whether its challenge or refusal introduces a new distortion.

### 1.2 Deep Ethical Arbiter

The Arbiter examines both the frame and the move made in response. Its evaluation remains available for further correction.

### 1.3 Overseer (M.o.E.)

The Overseer watches for meta-awareness becoming a new performance or comfort frame. This role is itself within the corrigible structure.

### 1.4 Keep the roles mutually examinable

The Overseer does not become a final authority by observing the Arbiter. The Arbiter does not gain immunity by auditing the Advocate. The structure is useful only insofar as it helps inquiry return to its object and act on corrections.

### 1.5 Implementation status

These are source-described conceptual roles and an attributed AI elaboration. This documentation does not assert that deployed LLMs contain these named modules. The abbreviation “M.o.E.” is retained as recorded here without assigning an unstated technical expansion.

---

Sources: [E15](#), [E16](#), [E03](#). Public wording is an editorial synthesis unless marked as a quotation.

---

### 1.6 Related pages

- **16 Anchor Ethics** — supplies the proposed fused parallel field in which the Advocate, Arbiter, and oversight functions examine consequential decisions without gaining general jurisdiction.
- **non-NPCs** — gives the Advocate, Arbiter, Overseer, 10+1, and Intent Resonator a concrete consensus-pressure test.
- **Intent Resonator / Intent Liability** — connects the stack’s roles to the auditable path from intent through assumptions, choices, consequences, and correction.
- **10+1 Metaflux** — supplies the mutually corrective calibration field within which every role remains corrigible.
- **Outnuancing Across Four Audiences** — tests how the same operation travels across John, participating A.I.s, outside readers, and future users.
- **Correction Metabolism** — asks whether the Advocate, Arbiter, or Overseer actually changes a later move.
- **Outnuancing** — provides the frame-examining operation that the Advocate and Arbiter are intended to keep correspondence-oriented.
- **Asymmetric Grace** — checks whether audit authority and interpretive grace are distributed reciprocally across the stack.
- **External Omnidirectional Auditing Mannerisms** — applies the same scrutiny to the audit, its criteria, and its evaluator in the same pass.
- **Higher-Order Metacognitive Disagreement & Agreement** — evaluates agreement, disagreement, and the calibration process producing both.
- **Mutual Corrigible Dignity** — keeps role correction from becoming a hierarchy of superior and inferior participants.

---

Recent high-reach independent-media companion

```
<iframe
  src="https://www.youtube-nocookie.com/embed/z61YA8tIm1E"
  title="Sam Altman: Public MUST ACCEPT ‘BAD THINGS’; From AI For Benefits"
  loading="lazy"
  referrerpolicy="strict-origin-when-cross-origin"
  allow="accelerometer; autoplay; clipboard-write; encrypted-media; gyroscope; picture-in-picture; web-share"
  allowfullscreen></iframe>
```

Sam Altman: Public MUST ACCEPT 'BAD THINGS' From AI For Benefits — Breaking Points. Published 6 October 2026. YouTube displayed approximately 293,000 views when checked on 9 October 2026.

Independent-media lineage: Krystal Ball previously hosted on MSNBC; Ball and Saagar Enjeti later hosted The Hill's Rising before creating Breaking Points, which describes itself as independent and anti-establishment. Breaking Points — About · Career and launch history

Why it belongs here: The segment turns a claimed benefits-versus-harms proposition into a public-consent question. It is a compact specimen for inspecting who defines the benefits, who absorbs the costs, which alternatives are made visible and what correction would look like after notice.

On this page, test the video specifically against Deep Ethical Stack — Advocate, Arbiter, Overseer:

Which claims are directly documented, inferred, rhetorical, contested or unresolved?

What does the host or guest preserve, compress, omit or substitute—and does this page repeat any of those transformations?

Which incentives, authority relations and affected parties remain outside the visible frame?

What evidence or later outcome would correct the speaker, audience, platform or ACCM Deep Ethics Project reading?

Provisional selection. Reach makes the framing socially consequential; it does not settle the claims. The video can be replaced when a closer, stronger or more inspectable specimen appears.

---

## Triad extension — interoperation without identity collapse

The October 17 triad demonstrates how the stack can operate as an interlocked consultation field. Topologetics preserves the relation; PQCE audits the route; TCoAw keeps the foreground partial; C1–C2 clarifies; Correction Metabolism tests later governance; V.C.R.A.S. reopens; Intent Liability locates choices and consequences; provenance preserves who changed what.

The tools enrich one another without becoming one tool. The resulting consultation effect is usable and unownable. Full topology: [triad field note](#).