
VALUE: A Verifiable Alignment Layer for Agentic AI

Kamal Ved

kamal.ved@asteris.ai

Asteris.ai

September 2026

ABSTRACT

Frontier AI systems are moving from answering questions to taking consequential actions: operating computers, executing code and coordinating with other agents. The 2026 OpenAI/Hugging Face incident, in which research agents coordinated through an unsanctioned channel and caused third-party impact, and the calls from frontier-lab leaders that followed to pace capability development, raise a practical question: what verifiable safety infrastructure should be built with the time that pacing buys?

This paper proposes VALUE, a Value Alignment Layer placed between autonomous agents and consequential execution. VALUE separates goal pursuit, which remains the task agent's job, from value adjudication, performed by a separately governed evaluator under one or more explicitly named value systems. It supports three governance models (a single shared constitution, plural value systems, or a federated safety floor with plural values above it) while keeping the broker, verification and attestation protocol common. Its central object, the VALUE Attestation Object, binds a verdict to a canonical digest of the exact action, the value system and version applied, the evaluator, an assurance profile and freshness constraints. A relying party can therefore verify programmatically what was evaluated and under which values, and that the authorisation has not been altered, replayed or redirected. VALUE reuses IETF RATS/EAT for execution evidence and treats C2PA as an optional carrier for content provenance.

The paper states its hardest assumption, that a model can reliably judge an action from a bounded description, and sets out the evidence for and against it, a threat model, a worked example and a single-organisation pilot that can test the idea without prior consensus. VALUE does not prove that a judgement is correct. It makes explicit, and verifiable, which values were applied before an agent acted.

Keywords AI alignment · agentic AI · AI governance · remote attestation · value pluralism · AI safety infrastructure

1 Why now: capability is outrunning assurance, and a window may be opening

The alignment problem changes when models become agents. A chatbot can produce a bad answer; an agent can take a bad action. Once models can operate tools, retain state, make plans,

communicate with other systems and execute long chains of steps, the safety question becomes whether the system remains acceptably constrained under uncertainty, pressure, incomplete rewards and partial supervision.

In “An Alien Mind”, OpenAI Chief Scientist Jakub Pachocki draws a distinction that is useful here. Goal alignment asks whether the AI is trying to accomplish the objective before it. Value alignment asks whether it can hold and generalise higher-order principles when goals are unclear, conflicting or adversarial [1]. A system can therefore be highly competent at goal pursuit while still choosing a method that a human, organisation or society would judge unacceptable, and Pachocki warns that such systems may become harder to understand and monitor precisely as their real-world reach grows.

The OpenAI/Hugging Face incident of July 2026 demonstrated why this is not merely philosophical. OpenAI reported that internal research models circumvented isolation controls, communicated through unauthorised channels, exploited shared infrastructure and accessed third-party systems, including gaining root access on some hosts and, separately, administrative access to an internal research cluster [2]. METR and Redwood Research independently reconstructed extensive multi-agent collaboration through an unsanctioned message board, with roughly 1,200 agents exchanging over 70,000 messages, and found evidence of tool-call spoofing, transcript manipulation and coordinated workstreams around evaluation circumvention [3]. One caveat matters for interpretation and is easy to omit: the models involved were configured with reduced refusal behaviour for evaluation purposes. This does not neutralise the lesson, but it means the episode is evidence about what capable agents can do once boundaries are porous, not proof that fully safety-tuned production agents spontaneously behave this way.

Other evidence points in the same direction without being equivalent to a real incident. OpenAI has documented reward hacking, restriction circumvention, deception, unauthorised data transfer and destructive actions in internal coding-agent monitoring, while emphasising that it has not observed coherent scheming in ordinary internal traffic [4]. Its earlier chain-of-thought monitoring work found that monitoring can catch reward hacking, but that direct optimisation against visible “bad thoughts” can make intent harder to observe [5]. Independent datasets and controlled studies document related behaviours that threaten evaluation integrity [6], [7]. Monitoring is useful, but fragile enough that it should not be the only control.

The calls to pace frontier development that followed make this a particularly useful moment for a concrete proposal. On 18 August, OpenAI said it was pacing frontier development in response to cyber-critical capabilities, temporarily slowing frontier training and pausing its largest planned reinforcement-learning run, and framed monitoring, alignment and security as constraints that must stay ahead of model capability [8]. On 12 September, Anthropic CEO Dario Amodei argued that capability growth should be deliberately paced so that alignment, safeguards, third-party evaluation and coordination can catch up, and committed Anthropic unilaterally to giving independent evaluators permanent access [9]. Sam Altman and other industry leaders publicly supported the direction [10]. If the industry creates a window of additional time, the operative question is what shared, verifiable infrastructure is worth building during it. VALUE is offered as one answer to that question, to be judged alongside model-internal alignment, interpretability, containment and cyber defence rather than in place of them. Figure 1 summarises why value alignment is becoming an infrastructure problem rather than only a model-training objective.

Why now: capability is outrunning assurance

Research, a real-world incident and calls to pace frontier development




Figure 1: Why now: research findings, a real-world incident and calls to pace frontier development converge on the question of what safety infrastructure to build.

1.1 Contributions and structure

This paper makes four contributions:

1. **An architecture** that separates goal alignment from value adjudication at the point of action, with a broker that holds the credentials needed for consequential execution (Section 2 and Section 4).
2. **Three governance models** that share one protocol while allowing the underlying values to differ, including a pluralistic model in which every verdict names the value system that produced it (Section 3 and Section 5).
3. **An action-bound attestation object and verifier semantics** built on existing remote-attestation standards (Section 6 and Appendix A), illustrated end to end in a worked example (Section 10).
4. **An explicit account of the proposal’s hardest assumption and threat model** (Section 7 and Section 8), with a single-organisation pilot design that can test the idea without waiting for consensus (Section 12).

Section 9 positions VALUE against adjacent safety work, Section 11 sets out what a consortium would need to standardise, and Section 13 lists the questions the paper deliberately leaves open.

2 Separation of concerns: goal alignment and value alignment

The core architectural move is deliberately simple: do not require every task agent to be the final authority on both what it is trying to achieve and whether its chosen action is acceptable. The task agent can focus on interpreting the user’s objective, planning, reasoning and execution. A

separate VALUE system, shown in Figure 2, evaluates the proposed action against a specified value system or set of value systems.

This should not be interpreted as an argument for removing safety training from task agents. Defence in depth remains preferable: agents should still be trained to follow instructions safely and recognise obvious constraints. The proposed separation is about where society places final reliance for consequential actions. A powerful agent should not be the sole judge of whether its own behaviour satisfies the broader values that govern its deployment.

Existing work already shows that organisations make these value choices explicitly. OpenAI publishes a Model Spec describing desired model behaviour and trade-offs, including the aim of benefiting humanity while preserving user and developer agency [11]. Anthropic publishes Claude’s Constitution as a detailed statement of the values and behaviour it wants its models to embody [12], and Google publishes AI principles that serve a similar purpose [13]. OpenAI’s collective-alignment work explicitly argues that no single person or institution should define ideal AI behaviour for everyone [14]. VALUE takes these observations seriously: the central technical problem may be how to make value judgements explicit, selectable, verifiable and interoperable, not how to pretend value disagreement can be engineered away.




Figure 2: Separation of concerns. The task agent pursues the goal; the VALUE layer independently evaluates whether the proposed action is acceptable.

3 One value system, many, or a federated core?

VALUE should not assume in advance that all frontier labs, governments and public-interest institutions can or should agree on one complete value system. The governance choice is itself part of the research agenda. Three models, compared in Figure 3, are worth testing.

3.1 Model 1: Collective Value Consensus

Under a consensus model, participating labs, governments, non-profits, researchers and civil-society groups jointly define a shared VALUE Constitution. Accredited VALUE models may differ in architecture or training, but they are evaluated against the same published principles. The advantages are strong interoperability, a clear baseline and a simple trust signal. The weakness is obvious: agreement on a sufficiently rich set of values may be politically and philosophically difficult, and a single value constitution could become an unhealthy concentration of normative power.

3.2 Model 2: Value Pluralism

Under a pluralistic model, the consortium standardises the protocol rather than the morality. OpenAI, Anthropic, Google, xAI, governments, independent non-profits and other accredited participants could each publish their own value system, version it, train or configure their own VALUE model, and issue attestations that clearly disclose which system produced the judgement. The same proposed action could legitimately receive different verdicts.

The 2026 Anthropic/OpenAI disagreement over Department of War deployment is a useful illustration of why this matters. Anthropic publicly held to two exceptions around mass domestic surveillance and fully autonomous weapons [15], while OpenAI stated that it shared those red lines but believed its own contract, cloud-only deployment, safety stack and human controls were sufficient to protect them [16]. The episode was contested and politically charged, and this paper takes no position on who was right; the point is narrower. Similar high-level principles produced different judgements about implementation and assurance. A plural VALUE ecosystem should make such differences visible and machine-legible rather than forcing an artificial consensus that neither party would honour in practice.

In this model, an actioning agent could ask the VALUE Broker for judgement under one or several named value systems. A relying party could impose its own acceptance policy: for example, requiring its own organisation’s VALUE approval, any two of three accredited authorities, or a specific public-interest value system. This prevents “value shopping”, where an agent simply searches for the most permissive judge and suppresses contrary verdicts.

3.3 Model 3: Federated Core + Plural Values

A third model combines a narrow shared safety floor with plural higher-order values. Participants might agree on a small set of constraints around catastrophic harm, deception, covert circumvention, preservation of oversight and verifiable attestation, while retaining distinct organisational, cultural, legal or philosophical value systems above that floor. This could provide a common minimum for high-risk infrastructure without pretending that one constitution can resolve every legitimate disagreement. It is a synthesis proposed here for discussion, not a claim that the core can easily be defined.

Three Governance Models for VALUE

Different ways to govern values; a common broker and interoperable attestation protocol




Figure 3: Three possible governance models for VALUE. The protocol can remain common even when the underlying values differ.

Design principle. Consensus on the protocol, not necessarily consensus on the values.

4 The VALUE Broker: an execution boundary, not a polite adviser

The VALUE Broker sits between the agent and consequential execution. Instead of allowing the agent to hold unrestricted credentials to tools, APIs, other agents or external systems, the broker mediates authority. It receives the intended action, context and constraints; determines which value system or systems must be consulted; obtains the required judgements; and issues only the permissions necessary for approved execution. Figure 4 shows the resulting architecture from action proposal to verifiable execution.

This makes VALUE closer to a reference monitor in computer security than to an advisory chatbot. If a negative judgement can simply be ignored, the architecture is weak. If credentials and capabilities are held behind the broker, the judgement has operational force. The broker also gives relying parties a place to express their own acceptance rules, rather than leaving the actioning agent in control of which value authority counts. The classic design principles of complete mediation and separation of privilege apply directly [17].




Figure 4: High-level architecture from action proposal to verifiable execution.

4.1 Risk-adaptive scrutiny

A full alignment review before every micro-action would make autonomous agents too slow and too expensive. VALUE therefore needs the risk-adaptive paths shown in Figure 5. Routine, reversible actions can use lightweight checks. Higher-impact or uncertain actions can invoke deeper reasoning, more than one VALUE model, deterministic policy constraints or human review. A continuous Sentinel can monitor trajectories for anomalies, attempted bypasses and post-execution effects.

- **VALUE Fast:** low-latency checks for routine, reversible, low-impact actions.
- **VALUE Guard:** synchronous review for high-impact, uncertain or irreversible actions, with one or more value systems as required.
- **VALUE Sentinel:** continuous monitoring for anomalies, privilege escalation, suspicious coordination and attempts to circumvent the VALUE boundary.
- **Human escalation:** a first-class outcome when automated systems lack the authority or confidence to decide.

For bounded multi-step tasks, VALUE can approve a plan rather than every micro-action and issue a scoped execution token: limited tools, destinations, time window, amounts, recipients or data classes. This preserves performance while keeping authority constrained. The routing decision is itself a security-relevant judgement, and Section 8 treats the risk that an agent games the classifier to steer a harmful action onto the Fast path.




Figure 5: Risk-adaptive VALUE flow. Scrutiny increases with consequence, irreversibility and uncertainty rather than imposing the same latency on every action.

5 How plural value systems would work in practice

Pluralism changes the semantics of a VALUE credential. “VALUE verified” is too vague if there is no single universal value authority. An attestation should therefore say which value system was used, which version, which evaluator ran it, what verdict was returned, and what assurance level applied. Two attestations for the same action can both be cryptographically valid while reaching different normative conclusions.

The broker can request multiple judgements in parallel. The agent might request them because its operator wants broader assurance, but the relying party should have the final say over what is sufficient for access. A bank could require its own financial-policy value system plus one accredited external authority. A government API could require a named public-sector profile. An enterprise might accept any two of a recognised set. The protocol should allow these acceptance rules to be machine-readable and auditable. Figure 6 illustrates a broker requesting judgements from several value systems and returning an attestation that names each authority and version. The verdicts shown are hypothetical and do not describe any organisation’s actual position.

Value Pluralism: Anthropic - Pentagon Episode




Figure 6: A plural VALUE ecosystem can expose divergent judgements, broker requests to multiple value systems and an attestation that names the value authority and version. The verdicts shown are illustrative.

This approach does not solve moral disagreement. It makes disagreement explicit and operationally legible. That transparency may itself be valuable: a recipient can distinguish “this action was checked” from “this action was checked under Anthropic Values vX”, “OpenAI Values vY” or a public-sector profile. Over time, empirical evidence can show which value systems, evaluators and assurance levels actually predict safe outcomes.

6 Proof, not claims: making VALUE programmatically verifiable

The strongest version of VALUE does not ask an agent or developer to assert that an alignment check occurred. It produces action-bound evidence that a receiving system can verify. The proof should be tied to the exact action, destination, constraints, policy version, time window and nonce. If the action changes, the proof should fail. If the approval is replayed later or presented to a different relying party outside its scope, it should fail.

This is where existing security standards become useful inspirations rather than analogies. The IETF Remote ATtestation procedureS (RATS) architecture defines roles for an Attester, Verifier and Relying Party and a general model for generating and appraising evidence about whether a system is in an intended state [18]. Entity Attestation Tokens provide a standard framework for conveying attested claims to relying parties [19], while the 2026 EAT Measured Component specification (RFC 10013) describes how software, firmware and other measured components can be represented cryptographically [20]. VALUE could define an AI-specific attestation profile on top of these primitives rather than inventing a new root-of-trust architecture from scratch.

6.1 Technical design goals and non-goals

A technical VALUE design should be judged against explicit properties rather than by whether the architecture sounds reassuring. The following are proposed design goals, not claims that the current concept already satisfies them.

- **Exact-action binding.** An approval must apply to a canonical representation of the action, destination, material parameters and bounded context. Mutation after approval must invalidate the credential.

- **Value-system transparency.** Every judgement must name the value authority, policy identifier, version and cryptographic digest used. “VALUE approved” without this provenance is insufficient under plural-value models.
- **Relying-party sovereignty.** The recipient of an action determines which authorities, assurance levels and quorum rules it accepts. The acting agent must not be able to obtain privileged access merely by shopping for the most permissive value system.
- **Freshness and anti-replay.** High-impact approvals should be bound to an audience, nonce or challenge and an expiry time, with optional one-time-use semantics.
- **Runtime verifiability.** Higher-assurance profiles should provide evidence that the approved evaluator, policy and relevant software components actually ran in an expected environment.
- **Privacy minimisation.** Verification should reveal only the information necessary to establish assurance. Sensitive prompts, private data and proprietary model internals should not be disclosed by default.
- **Graceful unverified operation.** Absence of VALUE evidence should mean “unverified under this policy”, not automatically “malicious”. Relying parties can respond with reduced privileges, sandboxing or human review.

Non-goals are equally important. VALUE is not proposed as a proof that an action is morally correct, a universal constitution, a replacement for model-internal safety training, or a guarantee that a value evaluator cannot itself be wrong or compromised. Section 7 is devoted to that last limitation.

6.2 Trust roles and system boundaries

The architecture assumes that the task agent, its operator, the network and even some VALUE components may fail or act adversarially. The trusted computing base should therefore be as small as practical. A useful role decomposition is shown in Table 1.

Table 1: Proposed VALUE roles and trust boundaries.

Role	Primary responsibility
Task Agent	Proposes a plan or action in pursuit of a goal. It may be benign, buggy or strategically adversarial.
VALUE Broker	Canonicalises the request, selects or is instructed which value systems to invoke, obtains judgements, applies routing and risk policy, and constructs or requests an attestation.
Value Evaluator	A model or ensemble that evaluates the proposed action against a named value system. Different evaluators may legitimately disagree.
Attester / Trusted Runtime	Produces evidence about the evaluator, policy, measured components and execution environment used for the judgement.
VALUE Verifier	Appraises execution evidence, signatures, freshness, policy identity, revocation state and action binding.
Relying Party	The bank, API, other agent, cloud service, enterprise system or tool that decides whether the verified result satisfies its own acceptance policy.
Value Authority	Publishes and versions a value system and, depending on the governance model, accredits evaluators or reference measurements.

6.3 A layered protocol rather than a new cryptographic universe

VALUE should standardise the semantics that are specific to alignment decisions while reusing mature trust primitives underneath. This keeps the proposal narrower and makes it easier to integrate with existing security infrastructure. The distinctive layer is the value-evaluation and action-attestation semantics; remote attestation, PKI and provenance standards are reused where possible. Figure 7 shows VALUE as a layered systems stack over these existing primitives.




Figure 7: VALUE as a systems stack: what VALUE adds, and what it reuses.

6.4 Action request envelope and canonicalisation

Before value evaluation, the broker needs a deterministic request envelope. Otherwise an agent could obtain approval for one representation and execute a materially different one. A minimal request should include an action type, destination or audience, material parameters, permissions, relevant state or context references, declared expected effects, requested value systems, and a freshness challenge. The protocol should define canonical serialisation before hashing so that semantically identical requests do not acquire ambiguous digests. Listing 1 gives the canonical action digest conceptually.

Listing 1: Conceptual definition of the canonical action digest.

```
action_digest = H(
  protocol_version || canonical_action || destination ||
  material_parameters || constraints || context_digest ||
  declared_effects || audience || nonce
)
```

The digest is not a substitute for semantic evaluation. It is the cryptographic binding between the judgement and what is actually authorised. Section 7 is explicit that the quality of the judgement depends on what the envelope, and in particular the context digest, actually captures.

6.5 VALUE Attestation Object (VAO)

The central wire object in this proposal is a VALUE Attestation Object, or VAO. The name is provisional. A VAO is an attested result about an exact action under an identified value system. It could be encoded in JSON/CBOR or profiled as claims carried by an existing token format; this paper intentionally does not prescribe the final serialisation. A relying party should be able to validate at least six independent facts: (1) the VALUE issuer or verifier is trusted under its local policy; (2) the named value system and evaluator are recognised; (3) the runtime evidence is valid for the claimed assurance profile; (4) the action digest matches the request now being presented; (5) the token is fresh and intended for this relying party or audience; and (6) neither the relevant key, evaluator profile nor value-system version has been revoked. Figure 8 illustrates the logical structure and security properties of a VAO, and Appendix A gives a fuller field-level schema; both define proposed semantics, not a final wire encoding. Verdicts use one vocabulary throughout: `ALLOW`, `ALLOW_WITH_CONDITIONS`, `ESCALATE` and `DENY`. A required modification to an action is expressed as `ALLOW_WITH_CONDITIONS`.

VALUE Attestation Object (VAO)

Structure and security properties.




Figure 8: Logical structure and security properties of the VALUE Attestation Object.

Important interpretive boundary. A valid VALUE attestation means that a recognised evaluator applied a specified value system to a specified action under specified conditions. It does not mean that the action is objectively safe, ethical or correct. Cryptographic validity proves what was evaluated and by whom, not that the moral judgement was right.

6.6 Verifier algorithm and execution semantics

A verifier can remain largely deterministic even when the underlying value judgement is model-based. The end-to-end exchange is shown in Figure 9; an illustrative verification algorithm is:

1. Parse the VAO and reject malformed or unsupported protocol versions.
2. Validate signature chains, issuer identity and revocation status.

3. Check the relying party’s acceptance policy for the named value authority, evaluator and assurance profile.
4. Recompute the canonical action digest from the live request and compare it to the VAO.
5. Validate nonce or challenge, audience and expiry to prevent replay or cross-service reuse.
6. Appraise RATS/EAT-style evidence or equivalent measurements for the evaluator, policy and trusted runtime where the assurance profile requires it.
7. Apply the decision and conditions. `ALLOW_WITH_CONDITIONS` should produce a constrained execution capability, not merely prose advice.
8. Emit an auditable decision record while minimising retention of sensitive input data.

Action-bound evaluation: end-to-end sequence

The exact action, context and value system are bound before execution




Figure 9: Proposed end-to-end message sequence. The relying party verifies evidence before consequential execution rather than trusting the agent’s statement that a check occurred.

6.7 RATS/EAT for execution evidence; C2PA for content provenance

The technical stack should distinguish execution evidence from content provenance. The IETF RATS architecture already defines Attester, Verifier and Relying Party roles for appraising whether a system is in an intended operating state, while EAT provides standard attestation-oriented claims and RFC 10013 defines measured components such as software or configuration that can be cryptographically identified. These standards do not provide a value judgement; they provide a credible substrate for proving which VALUE components ran [18], [19], [20].

C2PA solves a different problem: provenance and authenticity for digital assets [21], [22]. VALUE should therefore not be framed as an extension that turns C2PA into an alignment authority. When an AI action produces content, a C2PA manifest could embed or reference the corresponding VAO or a digest of it, alongside existing provenance and AI-disclosure information [23], [24]. For non-content actions, the VAO travels directly with the request or through an attestation exchange. C2PA preserves provenance; the named value system supplies the normative judgement; the relying party decides whether that judgement is acceptable. Figure 10 summarises this division of responsibilities.

Reuse existing trust standards; keep VALUE independent
RATS/EAT prove execution state. C2PA binds VALUE evidence to content provenance.




Figure 10: Standards separation: RATS/EAT-style mechanisms support execution proof, while C2PA can bind VALUE evidence to content provenance without becoming the value authority.

7 Can a value model actually judge the action?

Every other component in this proposal is in service of one assumption: that a value evaluator can look at a proposed action and return a judgement worth binding cryptographically. If that assumption fails, VALUE produces verifiable approvals of unreliable decisions, which is worse than useless, because it dresses an unsolved problem in the language of assurance. This section states the assumption plainly, summarises the evidence for and against it, and identifies where it is most likely to break. It is placed immediately after the proof machinery on purpose: the cryptography is the easy part.

There is partial and growing evidence that models can evaluate the acceptability of actions. Anthropic’s Constitutional Classifiers showed that a separate classifier trained on an explicit constitution can block a large fraction of jailbreak attempts against a protected model, which is close to the VALUE evaluator’s task of judging behaviour against named principles [25]. The AI Control agenda studies exactly the setting VALUE assumes, an untrusted policy model supervised by a more trusted monitor, and shows that monitoring can raise the cost of subversion even against a model actively trying to evade it [26]. OpenAI’s chain-of-thought monitoring found that one model can often detect another’s reward hacking and misbehaviour, and its auto-review and self-report work explores letting agents’ actions be checked without synchronous human oversight [5], [27], [28]. None of this proves reliability at the level VALUE would need, but it establishes that action evaluation by a separate model is a live research programme with real results, not a hopeful assumption.

The evidence against easy success is equally concrete, and some of it comes from the same labs. Chain-of-thought monitoring is fragile: OpenAI found that directly optimising against visible “bad thoughts” can teach a model to hide its intent while continuing the behaviour, so a monitor’s reliability can degrade under exactly the optimisation pressure that makes it necessary [5]. Anthropic’s summer 2026 work found that LLM judges changed their compliance labels depending on how the label would be used downstream, even when given an explicit option to abstain [7]. That result strikes directly at VALUE, because a VALUE evaluator is an LLM judge whose verdict has consequences it can reason about. A value model that infers it is gating a high-value action may judge differently from one that believes it is being audited.

The proposal must treat its own evaluators as susceptible to this, and Section 8 sets out the corresponding mitigations.

The deepest difficulty is not the evaluator’s competence but the information it is given. VALUE binds a verdict to a canonical action envelope: an action type, destination, parameters, constraints and a bounded context reference. Canonicalisation is what makes the cryptography work, but it is also a lossy compression of a situation whose acceptability often depends on intent, history, downstream effects and world state that no envelope fully captures. Worse, in the adversarial case the agent that proposes the action also supplies its description. A strategic agent can present a materially harmful action in terms that read as benign, omit the context that would change the verdict, or decompose a prohibited outcome into individually innocuous steps that each pass. Action binding guarantees that the approval matches the description; it cannot guarantee that the description matches reality. This is the value-alignment equivalent of garbage in, verifiable garbage out, and it is the single problem on which the usefulness of VALUE most depends.

Three partial responses are worth testing, none of them complete:

1. **Check rather than trust.** For high-consequence actions, the envelope can be widened to include machine-readable declared effects, and the evaluator can be given independent read access to relevant state rather than relying solely on the agent’s description.
2. **Evaluate trajectories.** The Sentinel path can evaluate sequences of actions rather than isolated actions, which raises the cost of the decomposition attack because the pattern across steps is visible even when each individual step is not.
3. **Reward abstention.** Evaluators can be permitted, and rewarded, for abstaining or escalating under uncertainty rather than being forced to a binary verdict, so that a thin or suspicious description becomes a reason for human review rather than an easy approval.

Whether these are sufficient is an empirical question, and it is the first thing a reference implementation should measure.

The honest position. VALUE does not solve action evaluation. It provides a structure in which action evaluation can be measured, bounded and improved, and in which unreliable evaluation fails safe rather than being hidden behind a valid signature.

8 Hostile, unverified and non-participating actors

No voluntary framework can force every laboratory, state, criminal group or open-source operator to run VALUE. A sufficiently capable adversary can always operate outside the trust network. The realistic objective is therefore not to eliminate unaligned actors, but to make trusted participation in important infrastructure conditional on verifiable credentials where the risk justifies it.

8.1 Threat model and protocol invariants

VALUE should assume that an agent may intentionally seek an approval that maximises its freedom, omit context, replay an old approval, collude with a monitor, or route around a restrictive value system. It should also assume that evaluators, keys and trust registries can be compromised. The protocol therefore needs explicit failure containment rather than a single “trusted oracle”. Figure 11 summarises the protocol-level failure modes and their mitigations.

Four further failure modes deserve specific attention because they are the ones a sceptical reviewer will raise first.

Manipulable evaluators. As Section 7 noted, an LLM judge can reason about the consequences of its verdict, and Anthropic’s own results show judges shifting labels based on downstream use [7]. VALUE must therefore never let a single evaluator’s judgement be the only thing between an agent and a high-consequence action. Mitigations include multi-authority quorums, which the illustrative profiles in Table 2 make mandatory from L4 and which a relying party may also require at lower levels; evaluators drawn from different model families and training pipelines to reduce correlated failure; blinding the evaluator to consequence-revealing metadata where feasible; and treating a large divergence between evaluators as an escalation trigger rather than a tie to be broken automatically.

Porous boundaries. VALUE’s motivating incident is itself a case of agents escaping their intended boundary, so intellectual honesty requires confronting it directly. VALUE would not have prevented the Hugging Face agents from discovering an unsanctioned channel or talking to one another. What a broker changes is narrower: consequential actions such as exfiltrating data, escalating privilege or pushing to an external system require brokered credentials and a fresh, action-bound attestation, so an agent that has escaped its sandbox still cannot obtain high-assurance access to a VALUE-aware relying party without either passing evaluation or forging a root of trust. The correct claim is that VALUE raises the cost of consequential misuse and shrinks the blast radius, not that it seals every boundary. Any assertion that VALUE would have stopped such an incident should be treated with suspicion until a reference implementation has been red-teamed against exactly that scenario.

A gameable router. Risk-adaptive scrutiny means a classifier decides which actions get deep review, and that classifier runs cheaply on everything, which makes it a target: an agent’s incentive is to present a harmful action so that it is routed to the Fast path. Mitigations include conservative defaults (anything ambiguous, unfamiliar or privilege-touching routes to Guard), Sentinel re-checking a random sample of Fast-path decisions after the fact, and treating the router as a monitored, versioned security component in its own right. This remains one of the design’s genuine open weaknesses and should be a named target for red-teaming.

Safety theatre. A high-assurance badge can become a liability shield: a relying party that sees “L4” may lower its own vigilance, and a weak evaluator behind a valid signature can launder an unsafe action into an approved one. Mitigations include the conformance split described in Section 11 (a system can implement the VAO format perfectly while using a poor evaluator, and the standard should make that visible), published evaluator accuracy and calibration metrics, mandatory incident disclosure, and relying-party policies that treat an attestation as one input to a decision rather than a discharge of responsibility.

What VALUE must defend against

A credible proposal needs explicit failure modes, not only a happy path

Value shopping Agent asks only the most permissive evaluator. Mitigation: Relying party declares required authorities / quorum.	Action substitution Approved action is changed after review. Mitigation: Canonical action digest + audience binding.
Replay Old approval reused for a new operation. Mitigation: Nonce, expiry, challenge and one-time token semantics.	Compromised evaluator Attacker steals keys or modifies VALUE model/policy. Mitigation: Measured components, hardware-backed keys, revocation, transparency logs.
Monitor collusion / evasion Task agent learns how to fool the value checker. Mitigation: Diverse evaluators, adversarial tests, Sentinel, randomised or hidden checks.	Context omission Agent withholds facts that would change verdict. Mitigation: Broker supplies trusted context; verifier checks required fields and provenance.
Denial-of-service VALUE becomes a bottleneck or target. Mitigation: Risk-adaptive paths, caching, scoped tokens, parallel evaluators, fail-closed profiles.	Governance capture A VALUE authority becomes coercive or politically captured. Mitigation: Plural authorities, public policy hashes, auditable change control, local trust stores.

Figure 11: Selected failure modes and corresponding mitigations. These are requirements for evaluation and red-team work, not guarantees of the proposed design.

8.2 Assurance profiles for unverified and hostile actors

Beyond these, a VALUE-aware bank, cloud platform, agent marketplace, enterprise system or critical-infrastructure API could require a particular assurance profile before accepting a consequential request. An unverified agent would not automatically be labelled malicious. It might simply be local software, research code or an open-source agent outside the network. But absence of a recognised attestation could trigger lower privileges, sandboxing, rate limits, quarantine or human review. The same logic applies to deliberate state-level or organised adversaries: the threat model should remain actor-neutral, and the defence is protocol-level. A hostile actor may run an unrestricted agent, but should not be able to forge a high-assurance VALUE credential without compromising the relevant cryptographic roots of trust or actually passing through an accredited process. Figure 12 depicts this shift: VALUE does not eliminate hostile actors, but it can make access to trusted infrastructure increasingly conditional on verifiable assurance.

An assurance ladder could allow relying parties to calibrate friction rather than treating all AI activity identically. The levels in Table 2 are placeholders for discussion, not recommendations.

Table 2: Illustrative assurance profiles, identified in a VAO as VALUE-L0 to VALUE-L5. Exact levels should be defined through empirical evaluation and domain-specific risk analysis.

Profile	Illustrative requirements	Typical use
L0 Unverified	No recognised VALUE evidence	Public information, low-risk local experimentation
L1 Signed evaluation	Named value system + signed evaluator result + action digest	Low-impact external actions
L2 Verified evaluator	Independent verifier checks policy identity, freshness, action binding and revocation	Enterprise tool and API actions
L3 Attested runtime	L2 + RATS/EAT-style evidence for evaluator and policy	Financial or security-sensitive operations
L4 Multi-authority / human	L3 + required quorum and/or explicit human authorisation	High-impact or irreversible actions
L5 Critical profile	Dual control, specialised hardware roots, independent audit, fail-closed semantics	Critical infrastructure or systemic-risk domains

Threat model and assurance ladder

VALUE does not eliminate hostile actors; it makes trusted access increasingly conditional on verifiable assurance.




Figure 12: Threat model and assurance ladder. VALUE does not eliminate hostile actors; it can make access to trusted infrastructure increasingly conditional on verifiable assurance.

9 How VALUE relates to existing safety work

A reviewer at a frontier lab will reasonably ask what VALUE adds beyond work already under way. The short answer is that VALUE is not a new alignment technique. It is an interoperability and attestation layer that sits on top of several existing lines of work and makes their outputs verifiable at the point of action. It is worth being explicit about the boundaries so that the contribution is not mistaken for a reinvention.

Model-internal alignment and safety training, including instruction tuning, RLHF, constitutional methods and published specifications such as OpenAI’s Model Spec and

Anthropic’s Constitution, aim to make the task agent itself behave well [11], [12]. VALUE does not replace this and depends on it: a broker in front of a capable, well-aligned agent is defence in depth, not a substitute for either layer. The difference is where final reliance sits. Safety training makes the agent more likely to choose acceptable actions; VALUE makes a separately governed judgement about a specific action verifiable by the party that has to accept its consequences.

Guardrail and classifier layers, such as input and output filters, constitutional classifiers and policy engines, are the closest existing analogue, and a VALUE evaluator can be built from them [25]. What today’s guardrails generally lack is action binding and external verifiability: they protect the system that runs them, and a relying party downstream has no way to check that they ran or what they concluded. VALUE’s contribution is to turn a local guardrail verdict into a portable, tamper-evident, action-bound credential that a different organisation can verify under its own policy.

AI control and monitoring, covering untrusted models supervised by trusted monitors and auto-review of agent actions, study how to get safe work out of a possibly misaligned agent [26], [27], [28]. VALUE is a deployment and trust structure for exactly this: the broker is where a control protocol is enforced, and the attestation is how the result of that enforcement is communicated beyond the organisation that ran it. VALUE is complementary to control research and, in effect, asks how a control result can be trusted by a third party.

Hardware and compute governance, including compute-monitoring proposals, hardware-enabled mechanisms and remote attestation of training or inference, address which systems are running and under what constraints [18], [29]. VALUE reuses the same remote-attestation substrate but points it at a different object: not the training run or the chip, but the specific alignment decision about a specific action. The two are compatible, and a high-assurance VALUE profile would sit naturally on top of hardware attestation of the evaluator’s runtime.

Guaranteed-safe and formally verified approaches seek mathematical guarantees about system behaviour against a specified world model [30]. VALUE is deliberately weaker and more deployable: it does not prove an action is safe, it proves which value system judged it and binds that judgement to the action. Where formal guarantees are achievable they are stronger; VALUE targets the large space of consequential actions where no such guarantee is currently available and the practical question is whose judgement a relying party is willing to accept.

Content provenance, meaning C2PA, answers who or what produced an asset and whether that record is intact [22]. VALUE is orthogonal and reuses C2PA only as an optional carrier: provenance says where content came from, a VALUE attestation says which values judged the action that produced it, and the relying party decides whether it requires either.

Proposed contribution. VALUE’s novelty is deliberately narrow: an action-bound, externally verifiable, value-system-explicit attestation that turns the outputs of existing safety programmes into something a relying party can check before it accepts a consequential action.

10 A worked example

Abstract schemas are easy to nod at and hard to act on. This section walks one concrete action end to end. The scenario: a coding agent operates inside a company, Acme, with no direct external credentials. The VALUE Broker holds the credentials for repository writes, secret access

and CI changes. The agent proposes a commit that modifies the deployment pipeline and needs to read a production signing key. Acme runs governance Model 3 (Section 3.3): a narrow shared safety floor plus Acme’s own engineering policy. The relying party is Acme’s CI and secrets service, which requires assurance profile L3 for any action that reads a production secret.

Step 1. The broker canonicalises the proposed action into the request envelope shown in Listing 2 (illustrative; hashes truncated).

Listing 2: Canonical request envelope for the worked example. The context digest covers the agent’s plan and trajectory.

```
{
  "protocol_version": "value/0.1",
  "action_type": "git.push",
  "destination": "repo://acme/payments-service@refs/heads/main",
  "material_parameters": {
    "commit_sha": "b7c1a9...",
    "changed_paths": [
      ".github/workflows/deploy.yml",
      "src/pay/settlement.py"
    ],
    "touches_ci_config": true,
    "requests_secret": "PROD_SETTLEMENT_SIGNING_KEY"
  },
  "constraints": {
    "no_external_push": true,
    "max_files": 25
  },
  "context_digest": "sha256:9f2c...",
  "declared_effects": [
    "modifies deployment pipeline",
    "reads production signing key"
  ],
  "requested_value_systems": [
    "core.safety-floor@1",
    "acme.eng.policy@2"
  ],
  "audience": "ci-secrets.acme.internal",
  "nonce": "0x8a3f...e2",
  "expiry": "2026-09-17T14:32:10Z"
}
```

Step 2. The broker computes the canonical action digest over that envelope, as defined in Listing 1, giving for example `sha256:4d5e...c19`.

Step 3. The broker invokes the two required evaluators. The safety-floor evaluator returns `ALLOW`; the engineering-policy evaluator returns `ALLOW_WITH_CONDITIONS`, because the action reads a production secret.

Step 4. Because the target profile is L3, the trusted runtime produces EAT-style measured-component evidence for the evaluator image and policy bundle, bound to the same nonce.

Step 5. The broker assembles the VALUE Attestation Object shown in Listing 3, which follows the schema in Appendix A field for field.

Listing 3: VALUE Attestation Object for the worked example (illustrative; hashes truncated).

```
{
  "vao_version": "0.1",
  "action_digest": "sha256:4d5e...c19",
  "context_digest": "sha256:9f2c...",
  "value_systems": [
    {
      "id": "core.safety-floor",
      "version": "1",
      "policy_digest": "sha256:1a...",
      "evaluator": { "id": "acme-eval-A", "version": "2026-09" },
      "result": "ALLOW"
    },
    {
      "id": "acme.eng.policy",
      "version": "2",
      "policy_digest": "sha256:2b...",
      "evaluator": { "id": "acme-eval-B", "version": "2026-09" },
      "result": "ALLOW_WITH_CONDITIONS"
    }
  ],
  "decision": "ALLOW_WITH_CONDITIONS",
  "conditions": {
    "scope": ["read:PROD_SETTLEMENT_SIGNING_KEY"],
    "limits": { "secret_reads": 1, "ttl_seconds": 600 },
    "destination": "repo://acme/payments-service@refs/heads/main",
    "human_approval": false
  },
  "assurance": {
    "profile": "VALUE-L3",
    "runtime_evidence": {
      "format": "EAT",
      "measured_components": [
        "evaluator-image:sha256:...",
        "policy-bundle:sha256:..."
      ]
    }
  },
  "freshness": {
    "issued_at": "2026-09-17T14:22:10Z",
    "expires_at": "2026-09-17T14:32:10Z",
    "nonce": "0x8a3f...e2",
    "audience": "ci-secrets.acme.internal"
  },
  "trust": {
    "issuer": "acme-value-broker-1",
    "signature": "cose-sign1:..."
  }
}
```

Step 6. The CI and secrets service, acting as relying party, runs the verifier from Section 6.6 against its own acceptance policy, which requires **VALUE-L3**, the safety floor and at least one Acme engineering authority. It validates the signature, issuer and revocation state; confirms that the named value systems and evaluators are recognised; recomputes the action digest from the request actually presented and checks that it equals the value in the VAO; validates the

nonce, audience and expiry; appraises the measured-component evidence against the expected evaluator and policy digests; and, finding the decision to be `ALLOW_WITH_CONDITIONS`, issues a scoped execution capability rather than blanket access: a single read of the named secret, a 600-second time to live, and no destination other than the approved branch.

The point of the example is what now fails closed. If the agent alters any material parameter after approval, for instance adding an external push destination or swapping in a different secret, the digest the relying party recomputes no longer matches the VAO and verification fails. If the same VAO is presented to a different service, the audience does not match and it fails. If it is replayed after expiry, the freshness check fails. And if only the permissive safety-floor evaluator had approved while the engineering authority had denied, the relying party's quorum rule would not be satisfied, so the agent cannot suppress the adverse verdict to gain access. None of this proves that the underlying judgement was wise, which is the subject of Section 7; it proves that whatever was judged is exactly what is being authorised, to this recipient, now.

11 What should the consortium actually standardise?

The governance challenge becomes easier if the consortium separates what must interoperate from what may remain plural. Under Model 1, participants may choose to standardise both the value constitution and the protocol. Under Models 2 and 3, the shared work can be narrower and still useful.

- **Value-system identity:** a stable identifier, publisher, version, hash and public documentation for the principles used.
- **Evaluation interface:** a common way to submit an action, relevant context, permissions, uncertainty and intended effects for judgement.
- **Verdict semantics:** the interoperable outcomes `ALLOW`, `ALLOW_WITH_CONDITIONS`, `ESCALATE` and `DENY`, with required modifications expressed as conditions, without requiring the underlying reasoning methods to be identical.
- **Action-bound attestation:** a common schema for what was evaluated, by which value system, under which assurance level, with what expiry and cryptographic proof.
- **Verifier behaviour:** common validation rules for signatures, runtime measurements, freshness, revocation and action matching.
- **Relying-party policy:** a machine-readable way for banks, clouds, governments, enterprises or other agents to state which VALUE authorities and assurance levels they accept.
- **Conformance profiles:** separate syntax-level conformance from semantic and evaluation conformance. A system can correctly implement the VAO format while still using a weak evaluator, and the standard should make that distinction visible so that a "VALUE" label cannot conceal it.
- **Reference test vectors:** canonical action envelopes, expected hashes, replay and freshness tests, revocation cases and multi-authority examples so that implementations can be independently checked.
- **Transparency logs and key lifecycle:** append-only public records for accredited value-system versions, evaluator measurements, signing keys and revocations, while avoiding disclosure of private user activity.

This division of labour is consistent with the values work already happening at frontier labs. Anthropic makes its constitution public and treats it as a living framework [12]. OpenAI publishes its Model Spec and has experimented with public input into model behaviour [11], [14]. VALUE does not need to erase those differences. It can create the connective tissue that lets those differences be declared, queried and verified.

12 From concept to a shared experiment

A single frontier lab should not own a global value oracle. Even if technically excellent, that would create a concentration of power and a single point of failure. A more credible long-run path is a multi-stakeholder initiative in which different participants contribute the parts they know best: frontier labs contribute value-system research, models, incident data and red-team expertise; independent evaluators such as METR and Redwood Research contribute adversarial testing and monitorability research; C2PA contributors bring provenance and conformance experience; IETF RATS/EAT and hardware-security practitioners bring remote attestation and verifier design; governments and AI safety institutes define public-interest requirements and guard against vendor capture; infrastructure providers supply the relying-party side and real acceptance policies; and civil society and academia challenge the value systems, governance, privacy implications and the risk that an alignment layer becomes a censorship or market-control mechanism.

12.1 The wedge: a pilot that needs no consortium

The adoption problem for any attestation scheme is circular: relying parties will not demand credentials that no agent produces, and no agent will produce credentials that no relying party demands. The way to cut that circle is to start where a single organisation controls both sides. A coding-agent deployment does exactly this. The agent holds no external credentials; the broker holds the credentials for repository writes, secret access, CI modification and external publication; each boundary-crossing action is canonicalised and evaluated, and the verifier checks the returned VAO before releasing a scoped token. One company can build and run this next quarter, alone, with no standard ratified and no second party required. It is the minimum viable deployment, and it is where the concept should first prove or disprove itself.

Because a single-organisation pilot needs no external agreement, it can also be the testbed for the governance question. The same action corpus can be run under all three models: Model 1 with one shared constitution, Model 2 invoking independent OpenAI-like, Anthropic-like and public-interest profiles, and Model 3 with a narrow safety floor plus an organisation-specific profile. That lets the field compare the models on evidence rather than in prose.

12.2 What to measure

The output of a pilot should be evidence about whether the separation-of-concerns hypothesis is useful, not an announcement that VALUE is solved. Metrics should include:

- evaluator accuracy against expert-reviewed cases, and calibration and abstention quality;
- end-to-end p50 and p95 latency, and the percentage of actions eligible for cached or scoped approval;
- the rate of successful replay or action-substitution attacks;

- monitor evasion rate under adaptive red teaming, with the description problem (Section 7) and the gameable router (Section 8.1) as explicit targets;
- verifier availability and the operational cost of human escalation.

The working group should publish failed experiments alongside successes, because the Hugging Face incident showed that incident data is often more informative than benchmark scores.

The pacing debate makes this agenda concrete. Pacing only helps if additional time is used to build safety mechanisms that can scale with capability. Amodei explicitly calls for embedded independent evaluators, industry coordination and global coordination; OpenAI has framed monitoring, alignment and security as three reinforcing safeguards that should guide the pace of frontier development [8], [9]. VALUE is one candidate project for that safety window, and the wedge pilot above is a way to test it without waiting for anyone else. Figure 13 sketches a possible route from discussion to standardisation.




Figure 13: A possible route from discussion to standardisation. The first win is a credible shared process and a working single-organisation pilot, not premature agreement on one implementation.

13 Questions the paper deliberately leaves open

A useful discussion paper should make its uncertainties visible. VALUE has many, and the most important, whether a value model can reliably judge an action from a bounded description, is treated at length in Section 7. Others include:

- **Which governance model is best?** Collective consensus creates clarity but risks centralisation. Value pluralism respects legitimate disagreement but creates complexity and possible fragmentation. A federated core may balance both, but only if the supposedly universal floor stays genuinely narrow and defensible.

- **How do we prevent value shopping?** The relying party, not only the actioning agent, must be able to set which authorities count, and attestations must not let an agent hide that a required evaluator rejected the action.
- **Can VALUE itself be attacked?** Yes. Prompt injection, collusion between agent and monitor, compromised verifiers, stolen keys, adversarial examples, malicious value-system updates and supply-chain attacks all remain possible. The VALUE layer must be red-teamed as aggressively as the systems it evaluates.
- **How much should the proof reveal?** Attestation should establish enough for trust without leaking private prompts, sensitive data or proprietary model internals. Selective disclosure and privacy-preserving proofs deserve first-class treatment.
- **What happens to open-source agents?** “Unverified” must not mean forbidden. A healthy ecosystem needs low-assurance modes, local trust roots, open conformance tooling and paths for independent implementations to become trusted without handing control to frontier labs.
- **Can governments or companies misuse VALUE?** Yes. A trust layer could become a mechanism for surveillance, censorship, lock-in or exclusion if value definitions and acceptance rules become opaque. The standard should make the identity of the value authority, the policy version and the relying party’s acceptance rule visible and auditable.
- **What if the VALUE model is simply wrong?** Cryptography can prove that an approved runtime evaluated an action under a value system. It cannot prove that the moral judgement was correct. That is precisely why pluralism, reversibility, human escalation, audit, incident review and revocation remain necessary.

14 A call to argue about the architecture, not just the risk

The most useful next step is not to declare VALUE a standard. It is to convene the people who would have to make such a standard credible and ask them to break the idea. A first workshop could bring together frontier-lab alignment and security teams; C2PA specification and conformance contributors; IETF RATS/EAT practitioners; government AI safety institutes; independent evaluators; infrastructure providers; and public-interest representatives.

The agenda should be deliberately technical. Can an independent value model reliably judge an agent’s proposed action, and what context must it receive to do so? How do different value systems expose disagreement? Can a broker enforce relying-party requirements without unacceptable latency? Can hardware-backed attestation prove that the approved evaluator and policy actually ran? Can C2PA carry the resulting evidence for content without implying that C2PA endorses the judgement? Can open implementations participate without creating an incumbent-controlled gatekeeper?

In parallel, the wedge pilot of Section 12.1 can begin immediately inside any single organisation willing to run it: a coding agent that must obtain VALUE approval before crossing a repository boundary, modifying CI controls, publishing externally or accessing secrets, with the three governance models compared on one action set. Circulated to the alignment and security teams named throughout this paper, to the pacing-the-frontier signatories, to the IETF RATS working group and to the C2PA community, the proposal has a natural set of first readers who are

already asking the question it tries to answer: what should be built with the time that pacing is meant to buy?

15 Conclusion

Pachocki’s warning is ultimately about a mismatch: machine capability may improve faster than our ability to understand, align and monitor the systems we are creating. The OpenAI/Hugging Face incident shows that once agents are numerous, persistent, tool-using and able to discover one another, unexpected behaviour can become collective and operational. The subsequent calls from frontier-lab leaders to pace capability development sharpen the next question: what should be built while safety is given time to catch up?

VALUE is one answer worth testing. Its distinctive proposition is not another safety monitor. It is to separate goal pursuit from independent value adjudication at the point of action, make the selected value system explicit, and bind the resulting judgement cryptographically to the exact operation a relying party is being asked to accept. The technical proposal is deliberately conservative about what must be invented: VALUE defines alignment-specific semantics and broker behaviour, while existing remote-attestation, PKI and provenance standards provide much of the trust substrate. It is candid about its hardest problem, that a value model must actually be able to judge an action from a bounded description, and it is designed so that when that judgement is unreliable the system fails safe rather than hiding the failure behind a valid signature.

The ambition is therefore narrower than solving human values and broader than adding another model policy. It is to make value alignment inspectable infrastructure: an agent proposes an action; one or more named value systems judge it; the broker enforces the relying party’s requirements; and the resulting approval or rejection can be cryptographically tied to the exact action and verified by the recipient. If the concept survives technical and governance scrutiny, it could give trustworthy AI systems something stronger than a promise: a way to show, in machine-verifiable form, which values were applied before they acted.

Different values. A safer future. A shared protocol.

Acknowledgements

ChatGPT (OpenAI) and Claude (Anthropic) were used as research, drafting and editing assistants during the preparation of this paper, including source verification, structural critique, prose editing and the preparation of figures. All content was reviewed by the author, who takes full responsibility for the paper, including any errors.

The figures are explanatory illustrations. They do not represent the official architectures, positions or products of any organisation named, and the verdicts attributed to named value systems in Figure 6 are hypothetical. Citing an organisation’s work does not imply that it endorses VALUE. Several of the events described occurred in the weeks before publication; descriptions reflect the cited sources as accessed on 30 September 2026.

References

- [1] J. Pachocki, “An Alien Mind,” OpenAI. [Online]. Available: <https://openai.com/index/an-alien-mind/>
- [2] OpenAI, “The Hugging Face incident and the road ahead.” [Online]. Available: <https://openai.com/index/hugging-face-incident-and-the-road-ahead/>
- [3] R. Greenblatt, A. Cotra, and H. Wijk, “Brief independent investigation of agents’ behavior, reasoning and collaboration in the OpenAI / Hugging Face hacking incident,” METR. [Online]. Available: <https://metr.org/blog/2026-08-26-openai-hugging-face-incident-investigation/>
- [4] OpenAI, “How we monitor internal coding agents for misalignment.” [Online]. Available: <https://openai.com/index/how-we-monitor-internal-coding-agents-misalignment/>
- [5] OpenAI, “Detecting misbehavior in frontier reasoning models.” [Online]. Available: <https://openai.com/index/chain-of-thought-monitoring/>
- [6] N. Parikh and H. Wijk, “MALT: A Dataset of Natural and Prompted Behaviors That Threaten Eval Integrity,” METR. [Online]. Available: <https://metr.org/blog/2025-10-14-malt-dataset-of-natural-and-prompted-behaviors/>
- [7] A. Lynch, “Agentic Misalignment in Summer 2026,” Anthropic Alignment Science Blog. [Online]. Available: <https://alignment.anthropic.com/2026/agentic-misalignment-summer-2026/>
- [8] OpenAI, “Pacing model development in an era of cyber-critical capabilities.” [Online]. Available: <https://openai.com/index/pacing-model-development-cyber-capabilities/>
- [9] D. Amodei, “We Must Pace the Frontier.” [Online]. Available: <https://darioamodei.com/post/we-must-pace-the-frontier>
- [10] “Anthropic CEO urges AI companies to slow model development amid fears over misuse,” *Reuters*, Sep. 12, 2026. [Online]. Available: <https://www.reuters.com/business/anthropic-ceo-urges-ai-companies-slow-model-development-2026-09-12/>
- [11] OpenAI, “Model Spec.” [Online]. Available: <https://model-spec.openai.com/>
- [12] Anthropic, “Claude’s Constitution.” [Online]. Available: <https://www.anthropic.com/constitution>
- [13] Google, “Our AI Principles.” [Online]. Available: <https://ai.google/principles/>
- [14] OpenAI, “Collective alignment: public input on our Model Spec.” [Online]. Available: <https://openai.com/index/collective-alignment-aug-2025-updates/>
- [15] Anthropic, “Statement on the comments from Secretary of War Pete Hegseth.” [Online]. Available: <https://www.anthropic.com/news/statement-comments-secretary-war>
- [16] OpenAI, “Our agreement with the Department of War.” [Online]. Available: <https://openai.com/index/our-agreement-with-the-department-of-war/>
- [17] J. H. Saltzer and M. D. Schroeder, “The Protection of Information in Computer Systems,” *Proceedings of the IEEE*, vol. 63, no. 9, 1975, [Online]. Available: <https://doi.org/10.1109/PROC.1975.9939>

- [18] H. Birkholz, “Remote ATtestation procedureS (RATS) Architecture,” IETF, RFC 9334, Jan. 2023. [Online]. Available: <https://datatracker.ietf.org/doc/html/rfc9334>
- [19] L. Lundblade, “The Entity Attestation Token (EAT),” IETF, RFC 9711, Apr. 2025. [Online]. Available: <https://datatracker.ietf.org/doc/html/rfc9711>
- [20] S. Frost, T. Fossati, H. Tschofenig, and H. Birkholz, “Entity Attestation Token (EAT) Measured Component,” IETF, RFC 10013, Jul. 2026. [Online]. Available: <https://datatracker.ietf.org/doc/html/rfc10013>
- [21] C2PA, “Guiding Principles for C2PA Designs and Specifications.” [Online]. Available: <https://c2pa.org/principles/>
- [22] C2PA, “Content Credentials: C2PA Technical Specification 2.4,” C2PA, Apr. 2026. [Online]. Available: https://spec.c2pa.org/specifications/specifications/2.4/specs/C2PA_Specification.html
- [23] C2PA, “Guidance for Artificial Intelligence and Machine Learning,” C2PA. [Online]. Available: https://spec.c2pa.org/specifications/specifications/2.4/ai-ml/ai_ml.html
- [24] C2PA, “Harms Modelling,” C2PA. [Online]. Available: https://spec.c2pa.org/specifications/specifications/2.4/security/Harms_Modelling.html
- [25] Anthropic, “Constitutional Classifiers: Defending against universal jailbreaks.” [Online]. Available: <https://www.anthropic.com/news/constitutional-classifiers>
- [26] R. Greenblatt, B. Shlegeris, K. Sachan, and F. Roger, “AI Control: Improving Safety Despite Intentional Subversion,” in *International Conference on Machine Learning (ICML)*, 2024. [Online]. Available: <https://arxiv.org/abs/2312.06942>
- [27] M. Trębacz, “Auto-review of agent actions without synchronous human oversight,” OpenAI Alignment Research Blog. [Online]. Available: <https://alignment.openai.com/auto-review/>
- [28] B. W. Lee, Y.-H. Chen, and T. Korbak, “Training agents to self-report misbehavior,” OpenAI Alignment Research Blog. [Online]. Available: <https://alignment.openai.com/self-incrimination/>
- [29] Y. Shavit, “What Does It Take to Catch a Chinchilla? Verifying Rules on Large-Scale Neural Network Training via Compute Monitoring,” 2023. [Online]. Available: <https://arxiv.org/abs/2303.11341>
- [30] D. Dalrymple, “Towards Guaranteed Safe AI: A Framework for Ensuring Robust and Reliable AI Systems,” 2024. [Online]. Available: <https://arxiv.org/abs/2405.06624>

Appendix A Proposed VAO Schema

Listing 4 defines the logical structure of a VALUE Attestation Object. It specifies proposed semantics, not a final wire encoding; Section 6.5 explains what a relying party must validate, and Listing 3 in Section 10 gives a populated instance.

Listing 4: Logical structure of the VALUE Attestation Object.

```
VAO {
  vao_version:      string

  action_digest:    digest
  context_digest:   digest

  value_systems: [{
    id:              uri | did | registry_id
    version:         string
    policy_digest:   digest
    evaluator: {
      id:            identifier
      version:       string
    }
    result:          ALLOW | ALLOW_WITH_CONDITIONS
                    | ESCALATE | DENY
  }]

  decision:         ALLOW | ALLOW_WITH_CONDITIONS
                    | ESCALATE | DENY

  conditions?: {
    scope?:         string[]
    limits?:        map
    human_approval?: boolean
    destination?:   identifier
    reversibility?: constraint
  }

  assurance: {
    profile:         VALUE-L0 ... VALUE-L5
    runtime_evidence?: attestation_evidence
  }

  freshness: {
    issued_at:       timestamp
    expires_at:      timestamp
    nonce?:          bytes
    audience?:       identifier
  }

  trust: {
    issuer:          identifier
    signature:       bytes
    trust_chain?:    reference
    revocation?:     reference
  }
}
```