Tracelet 1.1 + EDG

Sovereign Microkernel Architecture for Constitutional AI Execution

ETHRAEON Systems Papers | 79 provisional applications, all to be refiled; 79 technical specifications; ten families; no filing stated

Abstract

Tracelet 1.1 represents the next evolution of ETHRAEON's sovereign microkernel architecture, combining governance-first execution with the Ethical Decision Graph (EDG). This paper documents the complete architecture, performance characteristics, and constitutional enforcement mechanisms that enable 93-97% multi-engine convergence while maintaining sub-200ms audit trail generation.

79 provisional applications, all to be refiled; 79 technical specifications; ten families; no filing stated (Schedule A+ Protected).

1. Architectural Overview

1.1 Tracelet 1.1 Microkernel

Tracelet is a turn-metered, substrate-agnostic middleware that sits between an orchestrating authority and the underlying execution substrate. Unlike token-metered wrappers that bill and bound work by model token consumption, Tracelet meters by turn: a turn is a discrete, bounded unit of work that begins with an authorized request and ends with a sealed outcome. This turn boundary is the unit of governance, the unit of evidence, and the unit of metering simultaneously.

The microkernel keeps the trusted base small by design. Only three responsibilities live inside the kernel: admission (deciding whether a turn is permitted to begin), sealing

(binding a turn's inputs, decision path, and outcome into an append-only record), and arbitration (routing the turn to the substrate selected for it). Everything else, including the model engines themselves, the storage backends, and the transport layer, is treated as an untrusted outer ring. Because the kernel is substrate-agnostic, the same governance guarantees hold whether a turn is served by a frontier model, a local model, or a deterministic rule engine, and Tracelet is designed to operate down to a minimum viable device.

Isolation is structural rather than advisory. A turn cannot reach across its own boundary to mutate prior state, because state is append-only: the kernel exposes no in-place write path. A turn that attempts an action for which it holds no evidence-backed authorization does not proceed in a degraded mode; it halts. This fail-closed posture is what makes Tracelet a governance substrate and not merely an observability layer.

1.2 Ethical Decision Graph (EDG)

The Ethical Decision Graph (EDG) is the evidence-and-decision structure produced by every turn. Each node in the graph records a decision point: the constraint evaluated, the evidence consulted, and the disposition reached (permit, halt, or escalate). Edges record the ordered dependency between decisions, so that the graph is a directed, acyclic trace of how a turn arrived at its outcome.

Three node classes structure the graph. Constraint nodes evaluate a constitutional rule against the turn's request. Evidence nodes anchor a decision to a concrete artifact, a prior sealed turn, a state record, or an external attestation, so that no authorization rests on unevidenced assertion. Validation nodes confirm that the constitutional constraints relevant to the turn were all evaluated before the outcome was sealed; a turn whose validation node cannot close fails closed. Because the EDG is emitted as a first-class artifact and content-addressed, any party holding the reference can verify the decision path independently of the runtime that produced it, which is the basis for the canonical verification surface referenced in this paper's binding layer.

1.3 Multi-Engine Convergence

Tracelet does not assume a single engine is authoritative. A turn classified as high-stakes may be dispatched to multiple engines, and their dispositions are reconciled through the EDG rather than by naive majority vote. Convergence is measured as the rate at which independently dispatched engines reach the same governed disposition for the same turn under the same constraint set. Where engines diverge, the divergence

itself is recorded as an EDG node and routed to escalation rather than silently resolved, so that disagreement becomes evidence rather than noise.

The coordination protocol is deterministic at the governance layer even when the underlying engines are stochastic: the kernel fixes the constraint set, the evidence set, and the validation requirement before dispatch, so that what converges is the governed decision, not the raw generation. The convergence figures cited in Section 2 are reported as observed operating ranges from production telemetry; the mechanism described here is the architecture that produces and bounds them.

2. Performance Specifications

- Multi-Engine Convergence: 93-97% agreement rate
- Audit Trail Generation: <200ms (ISO/NIST/EU-AIA compliant)
- Constitutional Compliance: T5-rigidity enforcement at 98%+
- Context Preservation: 99.5% via Cipher integration
- Simple Checks: 28-44ms latency
- Complex Compliance: 50-120ms latency

3. Constitutional Enforcement

Tracelet operates as part of the constitutional runtime and evidence control layer. Constitutional enforcement is not a post-hoc audit applied to completed work; it is the admission condition for work beginning at all. The runtime enforces a rigid constraint posture (T5-rigidity) in which five properties hold without exception: state mutation is append-only, per-command execution is fail-closed, action authorization is evidence-required, truth handling admits no placeholders, and authority delegation is non-recursive. A turn that would violate any of these does not produce a degraded result; it halts and the halt is itself sealed into the EDG.

Constraints are sourced from the constitutional rule set rather than hardcoded into the kernel, so policy can be updated by additive anchored revision without rebuilding the trusted base. The Atlas-C constellation supplies the higher-order engines that consume

and reason over these governed turns, naming systems, governance, benchmarking, intelligence synthesis, agent handoff, scope navigation, and the creative and world-building engines, each of which receives only turns that have already passed constitutional admission. The result is that downstream engines never have to re-establish trust in their inputs: a turn that reaches an Atlas-C engine is, by construction, a turn that already carries a sealed, verifiable decision path.

4. Implementation Architecture

Tracelet is deployed as a thin middleware layer that an orchestrator calls per turn. The common deployment pattern places Tracelet inline on the path between the authority and the substrate, so that no turn reaches an engine except through admission. Because the kernel is substrate-agnostic and small, it can run co-located with a frontier-model gateway, on a local edge device, or as a standalone governance service, and the same EDG artifacts are produced in every placement.

The integration surface is deliberately narrow. An orchestrator submits a turn request carrying the intended action and the evidence references that authorize it; the kernel returns either a sealed outcome with its EDG reference or a sealed halt with the failed constraint. The two structural integration points are the constraint source (the constitutional rule set the kernel evaluates against) and the evidence store (the append-only record the kernel reads prior decisions from and writes new seals to). The canonical verification endpoint exposes EDG references for independent confirmation, which is the same mechanism used by the binding layer at the foot of this document.

The API contract follows the same fail-closed discipline as the runtime: an ill-formed request, a missing evidence reference, or an unreachable constraint source all resolve to a halt rather than to a permissive default. This keeps the integration honest; a caller cannot accidentally obtain an ungoverned result by degrading its own inputs.

5. Production Evidence

Tracelet 1.1 runs in ETHRAEON's own production estate, where it generates the audit-trail and multi-engine convergence records described in Sections 2 through 4. Named third-party deployment case studies, ninety-day operating results, and ROI analysis are

not published here: those figures are added only when measured operating data from a named deployment exists, in keeping with the no-placeholder posture that a measured claim is asserted only when measured evidence supports it. The architectural claims in this paper rest on the mechanism, not on an unpublished case study.

6. Future Evolution

The Tracelet 2.0 direction extends the turn-metered substrate along three axes while preserving the small-kernel and fail-closed guarantees that define 1.1. First, EDG enhancement: richer evidence-node typing and cross-turn graph linking, so that the decision history of a long-running agent can be traversed as a single verifiable braid rather than as a set of independent per-turn traces. Second, portability of the governance seal: making the EDG reference verifiable against an external anchor so that a sealed decision retains its proof even when moved off the originating runtime. Third, scaling the arbitration path so that multi-engine convergence can be coordinated across more engines and more substrates without enlarging the trusted base.

These directions are stated as architectural intent. Specific performance targets, convergence figures, and deployment scale numbers for 2.0 are deliberately not asserted here; they will be reported from measurement when a 2.0 build is in operation, consistent with the no-placeholder posture applied throughout this paper.

References

- SE28: Governance-First Execution Layers (paper-28-corix-measurement-framework.html)
- SE30a: Governance-First Execution Layers - EUDS
- SE30b: Governance-First Execution Layers - FULL
- Cipher Memory Architecture
- ETHRAEON 2.0.1 Complete Enhancement Suite