ARCHITECTURE BRIEF

From Identity to Evidence—and Coverage

Trust boundaries, outcome states, route inventory, security repairs and customer deployment responsibilities.

Every Decision Begins With Authority.

Version 7 · 18 November 2026 · Trace Continuity Labs · tracecontinuity.com
Patent Pending — U.S. Provisional Patent Application No. 64/137,172. Not a certification.

1 · Architecture

Trace Continuity is an execution-time authority, governance and evidence layer for AI agents and enterprise systems. It answers three separate questions: who authorized the action, what actually happened, and whether the operation was inside the required enforcement coverage. It provides governed memory, exact-action enforcement and verifiable evidence without claiming visibility or control over systems that never pass through Trace.

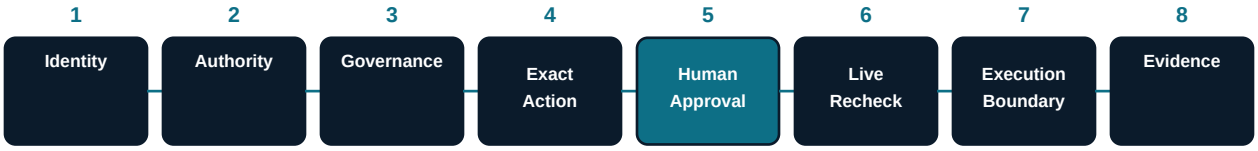

Authority is re-checked live. Approval applies only when required.

Outcome and coverage remain separate evidence dimensions after the execution boundary.

2 · Execution and evidence boundaries



Each state answers a different question. None is silently promoted to the next.

Dispatch ≠ target completion. Confirmation requires an explicit target-backed check.

Unknown or unavailable evidence remains unknown or unconfirmed.

No collapsed success state
Authorization is not execution. Consumption is not dispatch. Dispatch is not target completion. Coverage is a separate assessment, not an outcome.

3 · Trust boundaries

- Trace owns authority/governance decisions, exact-action permit issuance and redemption, governed memory, evidence and coverage-report generation.
- The customer enforcement point holds the downstream credential, recomputes the fingerprint, redeems immediately before execution, performs the exact action and reports evidence.
- The customer decides whether alternate credentials, APIs, jobs, integrations, network routes and manual operations can reach the same protected operation.
- A target system may supply a receipt or read-back. Trace records the enforcement point's attributed target check; it does not magically observe the target. When the target or a connector signs the receipt, Trace verifies that signature itself and records who signed.

- Historical evidence uses saved single-statement captures, not filters on alleged event time. Signed claims and verification audits are atomic; clean captures are independently exportable and private captures are withheld.

4 · Coverage information model

Input or signal	Meaning
Registered enforcement binding	A known enforcement-point identity can participate in the protocol; not proof of route completeness
Declared protected route	Customer-owned append-only declaration of operation, tool, target, environment, owner and credential boundary
Trace-observed operation	Authorization/consumption evidence actually recorded by Trace in the selected window
Customer-declared inventory	Systems, environments and completeness asserted by a named customer reviewer
Unregistered live key	Live Trace credential without a registered agent mapping; review signal, not proof of a bypass
Multiple credentials	More than one live key for an authority; potential alternate path requiring review
Unmapped observed operation	Observed action does not exactly match a current declared route
Verified observed coverage	Exact declared route matched observed governed evidence, with other risk fields still separate
Unknown / incomplete coverage	Missing, stale, ambiguous, out-of-scope or truncated information prevents a stronger conclusion

5 · Security boundary repairs

Repair	Current behavior and test boundary
Offboarding authority binding	The presented credential's pinned authority is preserved; a restricted authority cannot fall back to another eligible authority
Protected evidence reads	Audit, condition, lineage and action evidence require live authority and applicable scope/right; read decisions are audited
Matter-wall failures	Matter, wall, access and preservation lookup errors deny distinctly; an error is not treated as no wall or no hold
Matter-screened evidence	Sensitive rows and target-check records are screened through their underlying action before return
Protected-memory reads	Direct ordinary-client SELECT is explicitly revoked and refused on a fresh isolated database
Atomic signed evidence intake	Each signed claim and its verification audit commit together; audit failure rolls back the claim
Historical privacy and key lifecycle	Clean captures remain exportable, private details stay withheld, and intake-time key facts survive later revocation
Expanded bypass scanning	Person-suspension, coverage-declaration and approval-detail records are included in scanner inventories

6 · Evidence-read and matter model

- Audit and condition reads filter matter-bound rows; provenance screens roots, ancestors and descendants and reports excluded counts and incomplete visibility.
- External-action and approval evidence checks the matter boundary. Target-check records inherit the underlying action's screening; missing references stay hidden.
- Lookup failure is not absence. Matter, wall, access and preservation errors refuse with distinct reasons.
- Evidence-read decisions use the existing runAuthorityCheck decision point and are audited. Authorize/redeem performance paths are unchanged by these read-side checks.

7 · Deployment responsibilities

Area	Owner / requirement
Exclusive downstream credential custody	Customer IAM, secrets manager and workload identity
Complete execution routing	Customer APIs, integrations, jobs, SDKs, MCP servers and manual paths
Network egress and alternate-route closure	Customer infrastructure
Exact-action construction and fail-closed adapter	Customer enforcement point
Target receipt, read-back or reconciliation	Customer integration where actual effect matters
Coverage inventory ownership and review	Named customer owner; review before due date
Checkpoint anchor retention	Customer storage or escrow outside Trace
KMS/HSM, SSO/SCIM, SIEM, residency and formal service objectives	Separate deployment work; not claimed current

8 · Verification and external evidence

Reverified 18 November 2026: 633 ordinary tests passed with 208 environment-dependent tests skipped; 838 isolated-database tests passed with three port-gated affiliate tests skipped, and those three passed separately. Twenty-nine named historical-evidence checks passed, including atomic rollback, privacy, concurrency, replay, tamper and later-key-revocation cases. Live production synthetic memory write/read, cross-tenant refusal and revoked-key refusal passed. The full hosted signed external-action flow remains unestablished because no external-action authority or action records are provisioned. These are not certification results.

The separate 4 November internal staging benchmark remains a staging result: approximately 1.5 seconds per complete authorize → redeem → outcome → evidence workflow, approximately 3.17 workflows/second, concurrency 10, 25 and 50, and 160/160 successful workflows. Replay, cross-tenant and mid-action suspension attempts were refused; that run reported 714/714 security tests. It was not rerun for this brief and does not establish production capacity, sustained throughput or support for 1,000+ agents.

The commissioned October external evaluation remains 14 PASS / 3 FAIL / 1 BLOCKED plus 2 PASS / 1 NOT TESTED. F-01 was repaired and separately retested on staging. F-02/F-03/F-04 are reported staging repairs tested internally, without independent retest or production parity confirmation. Historical results are preserved, not rewritten.

9 · Stated limits

- Trace cannot detect or prevent an action occurring entirely outside its integration and network visibility. Customers must control credentials, IAM, network egress, jobs, alternate APIs and manual paths.
- A registered enforcement point or an observed governed request does not prove that all customer routes are governed. Inventory completeness and exclusive credential custody remain customer declarations until independently established.

- Target confirmation is attributed to the bound enforcement point. An unsigned target-evidence hash fingerprints reported content only. A signed receipt is verified by Trace against a registered key and labeled target-system-signed or connector-signed; it cannot prove the signer's key was never stolen.
- Signed checkpoints do not stop a database owner from changing data. Externally retained anchors make later alteration or deletion of committed history detectable; later activity is covered once a subsequent checkpoint is retained.
- Signing keys are held in managed secret storage, not an HSM. Tenant vault keys remain in the application database. No independent timestamping or public-ledger anchoring is claimed.
- Trace evaluates reported state and reported lineage. It does not discover hidden derivations, terminate arbitrary external processes or recall data/actions already released.
- Historical reconstruction requires a previously saved capture. Append-only rows, alleged event dates or an arbitrary as_of time cannot establish what was visible in the past. Signed statements attribute claims; they do not prove physical effects or external wall-clock time.
- Sensitive-data recognition is pattern/context-based and not exhaustive identification of arbitrary prose. No compliance certification or universal security guarantee is claimed.

10 · References

Current references: docs/CURRENT-BRIEF-VERIFICATION.md, ENTERPRISE-GOVERNANCE.md, API.md, INTEGRATION-CURRENT.md, EXTERNAL-ACTION-AUTHORIZATION.md, ENFORCEMENT-COVERAGE.md and ENTERPRISE-PILOT-FAQ-SECURITY.md.