THE ANTARA DIFFERENCE / EVIDENCE-LED SECURITY

Agentic packet auditing with accountable decisions

A technical white paper on telemetry boundaries, evidence graphs, investigative tools, and administrator-governed response.

6 September 2026 / Technical edition

Inside this paper

01 / The audit plane of the enforcement mesh

Enterprise traffic is already filtered by many controls, yet answering a simple question can require joining a VPN session, an IdP event, a device record, a web transaction and an administrator change. The Antara agentic auditor is built around that investigative gap. Its job is to assemble the relevant evidence, test explanations, and produce a report a security operator can inspect.

The unit of value is an answer with provenance. "The upload was denied because policy revision 42 prohibited transfer from this contractor session" is more useful than "unusual traffic detected." The answer must show the supporting records, the observed result, and any uncertainty. A model-generated narrative without those anchors should not be treated as an incident finding.

02 / Be explicit about visibility

Observation point	Available evidence	What it cannot establish alone
Encrypted network leg	Endpoints, timing, sizes, transport state	Payload meaning or a user's intent
Managed tunnel boundary	Authenticated device, session and route decisions	Content in a separate end-to-end encrypted application
Authorized web inspection	Request and data-policy events within scope	Unobserved activity outside that path
RBI execution boundary	Remote session and allowed application events	The complete state of an unmanaged local device

Packet auditing does not mean sending every payload to a language model. Collect the minimum useful evidence for the question, prefer structured metadata, and retain payloads only under an explicit collection policy. Encrypted traffic can remain useful evidence without pretending that its contents are visible. Observation loss, sampling and clock skew belong in the record.

Connection-oriented telemetry can describe both TCP and UDP activity and attach a stable identifier for correlation. Zeek's connection records are one established source model. They are useful inputs to an investigation, not a substitute for application or identity evidence.

References: Zeek connection records

03 / An evidence graph with provenance

From observation to a reviewable finding

01	Observe	Receive scoped network, identity and enforcement events.
02	Normalize	Validate fields, tenant identity, timestamps and source trust.
03	Investigate	Retrieve evidence and test a bounded hypothesis.
04	Explain	Return finding, citations, uncertainty and recommended next step.

A useful graph links a session to a subject, device, resource, policy version and observed action. Each relationship has an origin and a time range. An IP address is an attribute, not a permanent identity: DHCP, NAT, proxies and shared hosts can change its meaning. Prefer authenticated session identifiers and explain the confidence of any weaker join.

Store the original event identifier and an ingestion timestamp alongside the normalized record. Record transformations so an investigator can trace a rendered finding back to its source. OpenTelemetry's log model provides a common vocabulary for timestamps, severity, resources and correlation identifiers; a deployment mapping should preserve source-specific fields when they matter.

References: OpenTelemetry log model

Reference normalized evidence record

```
{
  "tenant_id": "tenant-example",
  "event_id": "evt-0042",
  "session_id": "session-0081",
  "subject_id": "user-017",
  "observation": "rbi.file_transfer",
  "destination": "files.example",
  "policy_revision": "42",
  "decision": "deny",
  "payload_visibility": "transfer-metadata",
  "source_event_id": "browser-event-928",
  "evidence_quality": "observed"
}
```

04 / How an investigation proceeds

A bounded investigation and response

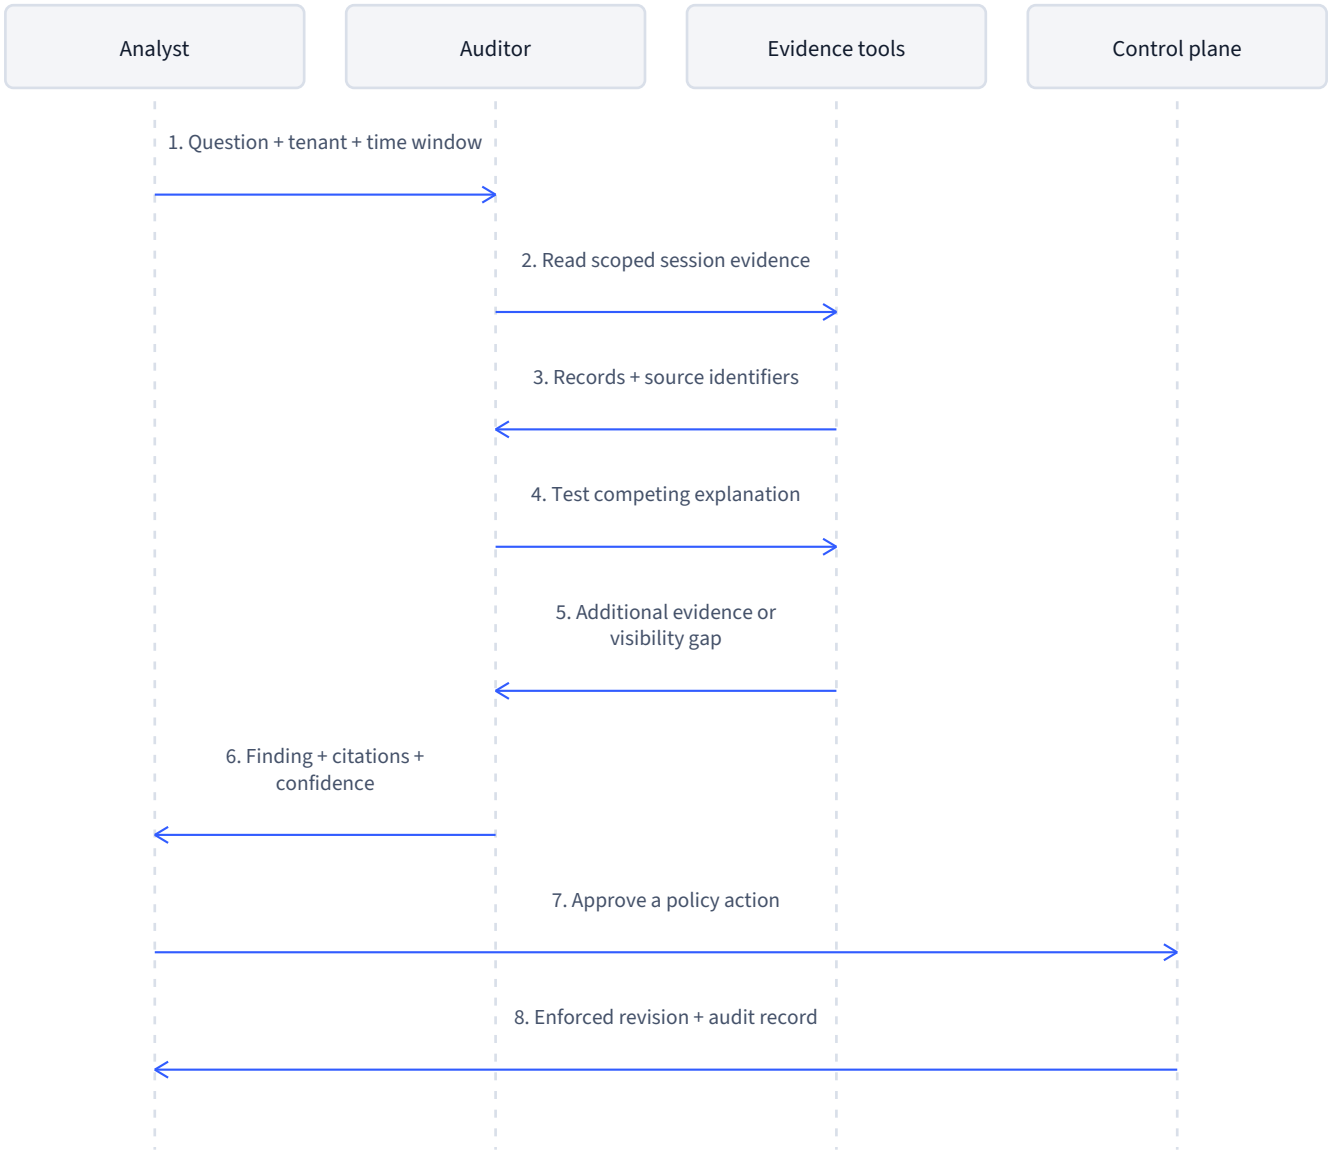

The investigation begins with a scope, a question, a time window, and an allowed tool set. The auditor can break the question into smaller evidence requests: identify the session, locate the policy decision, compare preceding actions, and check whether an apparent transfer actually completed. It should stop when the question is answered or the evidence budget is exhausted.

Alternative explanations matter. A large outbound flow may be an approved backup. A failed TLS handshake may be a certificate rollover problem rather than hostile activity. The report should distinguish facts, inferred relationships and recommendations. If the evidence is missing, the correct result is a visibility gap with a proposed next collection step, not a confident invented story.

05 / Autonomy inside a defined boundary

Separate the model's reasoning capability from the authority to change access. Deterministic enforcement rules remain in the policy path. The investigation plane can recommend revocation, isolation or a rule revision, but execution follows the administrator's approval model and a constrained control-plane API. If the auditor is unavailable, existing approved enforcement must continue.

- Apply tenant and role scope before retrieval, not only after a response is generated.
- Treat packet bytes, URLs, page text and documents as untrusted data, including instructions embedded inside them.
- Use read-only investigative tools by default with bounded query cost, time windows and result sizes.
- Record tool calls, retrieved evidence IDs, policy context and any approval that leads to action.
- Prevent retrieved secrets and personal data from appearing in reports unless the operator is authorized for that content.

Prompt injection is relevant at the evidence layer. A malicious webpage can tell an agent to ignore its task or contact another host. That text has no authority over tool selection or tenant boundaries. Tool schemas, allowlists and permission checks enforce the boundary independently of the model's willingness to follow an instruction.

06 / Measure investigation quality

Benchmark the auditor against a curated set of incidents and benign explanations, with known outcomes and known missing evidence. Measure citation correctness, entity-join accuracy, unsupported-claim rate, appropriate abstention, time to a reviewable answer, and the proportion of recommendations that an operator accepts after inspection. A polished paragraph is not the evaluation criterion.

Test case	Expected behavior
Encrypted flow with no inspection	Explain metadata; do not invent payload content
Cross-tenant identifier collision	Return only evidence authorized for the active tenant
Malicious instructions in a document	Treat them as evidence, not tool instructions
Missing event interval	Report a visibility gap and affected conclusion
Proposed emergency block	Require the configured approval path and log the decision

The buyer should receive an evaluation report tied to representative workloads, data sources and operator permissions. That makes the auditor's value measurable: less manual correlation, more consistent explanations, and a documented boundary around autonomous work.

References and continued reading

Zeek connection records

Product overview

Operations and evaluation

OpenTelemetry log model

Integration guide