OPENAI ARCHITECTURE PRIMER

How Agent Actions Become Evidence

Monitoring, interruption, reconstruction, and independent retention

A visual guide to turning agent telemetry and operational records into evidence that supports accountability, recovery, and proof.

NIST TENETS / OPENAI BEHAVIOR / LUMINITY SYNTHESIS



Evidence requires more than telemetry

Events and traces explain activity. Trustworthy evidence also requires correlation, response verification, and independent retention.

AT A GLANCE

4
evidence layers from live event to protected record

3
turn outcomes: completed, failed, or cancelled

2
NIST tenets owned by this edition

1
independent evidence store outside agent control

TWO DISTINCTIONS PROTECT EVIDENCE QUALITY

Observability versus monitoring

Events, items, and traces expose activity. Monitoring adds agent-specific thresholds, detection logic, response ownership, and verification.

Signed delivery versus immutable evidence

A verified webhook proves origin. Chain of custody also requires deduplication, processing receipts, protected retention, and integrity controls.

DOCUMENTED BOUNDARY

OpenAI documents events, saved work, traces, webhooks, cancellation, and administrative audit records. The reviewed sources do not define one complete enterprise chain of custody.

THREE CHAPTERS

01 Execution record

Join runtime events, saved state, traces, webhooks, administrative records, and downstream outcomes.

02 Monitor and interrupt

Detect anomaly or scope drift, apply the correct response, and verify the stop or refusal.

03 Evidence acceptance

Accept a package only when attribution, outcome, integrity, retention, and chain of custody are sufficient.

NIST defines the monitoring and evidence target. OpenAI supplies current inputs. Luminity composes the evidence chain.

© 2026 Luminity Digital, Inc. | Local review draft

NIST tenets 09 and 10 + OpenAI observability, events, tracing, webhooks, cancellation, and audit-log docs · verified Nov 5, 2026

Execution-to-evidence pipeline

OpenAI records describe agent activity. Enterprise correlation binds that activity to authority, outcomes, and protected evidence.

WHAT EACH RECORD CONTRIBUTES

Session, turn, and item IDs	CORRELATE SAVED WORK
Subagent ID	ATTRIBUTE DELEGATION
Tool and target result	VERIFY THE EFFECT
Webhook signature	VERIFY EVENT ORIGIN
Organization audit log	RECORD CONTROL CHANGE

Correlation envelope

Enterprise-owned identifiers bind principal, workload, mandate, authorization, session, turn, item, tool, resource, and outcome.

EVIDENCE PATH

Runtime events
Lifecycle, output, action-required, failure, and cancellation signals

LIVE SIGNAL

Saved state
Session, turns, items, errors, tool calls, and delegated attribution

OPENAI RECORD

Correlated context
Trace spans, policy decision, webhook receipt, and downstream result

LUMINITY JOIN

Protected record
Independent retention, integrity controls, timestamps, and chain of custody

ENTERPRISE EVIDENCE

EVIDENCE INPUTS AROUND THE ACTION

Live events

Track current lifecycle and output. Recover missed progress from saved state after a disconnect.

Saved items and turns

Retrieve messages, tool calls, statuses, errors, timestamps, and subagent ownership across pages.

Trace spans

Inspect recorded model, tool, and delegated activity. Export available traces when enabled.

External records

Join verified webhook receipts, administrative changes, policy decisions, and target transactions.

KEEP RECORDS DISTINCT

Event ≠ turn outcome ≠ trace ≠ monitoring finding ≠ independently retained evidence

© 2026 Luminity Digital, Inc. | Local review draft

OpenAI observability, events and items, tracing, webhooks, and organization audit-log documentation · verified Nov 5, 2026

OpenAI Select Partner

Monitor and interrupt loop

A response becomes defensible when the enterprise detects risk, applies the correct control, verifies the outcome, and preserves a receipt.

WHAT EACH CHECKPOINT ESTABLISHES

Current state collected	LIVE AND SAVED INPUTS
Identity and mandate joined	ACCOUNTABILITY BOUND
Anomaly or drift detected	THRESHOLD CROSSED
Response control applied	CANCEL, DENY, OR REVOKE
Outcome independently checked	STOP OR EFFECT VERIFIED

OpenAI runtime control

An application can cancel the active turn. Completion, failure, cancellation, idle state, and stream closure still require outcome inspection.

RESPONSE LOOP



CONTROL SURFACES AROUND THE RESPONSE

Detection policy

Define agent-specific thresholds for anomalous access, privilege change, scope drift, and repeated failures.

Interruption path

Route the right action to turn cancellation, policy denial, credential revocation, or human escalation.

Effect verification

Check the turn, tool, environment, resource, and transaction state before retry or closure.

Response receipt

Record trigger, owner, decision, control applied, timestamps, refusal or stop, and residual effects.

INTERRUPTION LIMIT Cancelling an active turn does not prove that a downstream side effect stopped, completed, or rolled back

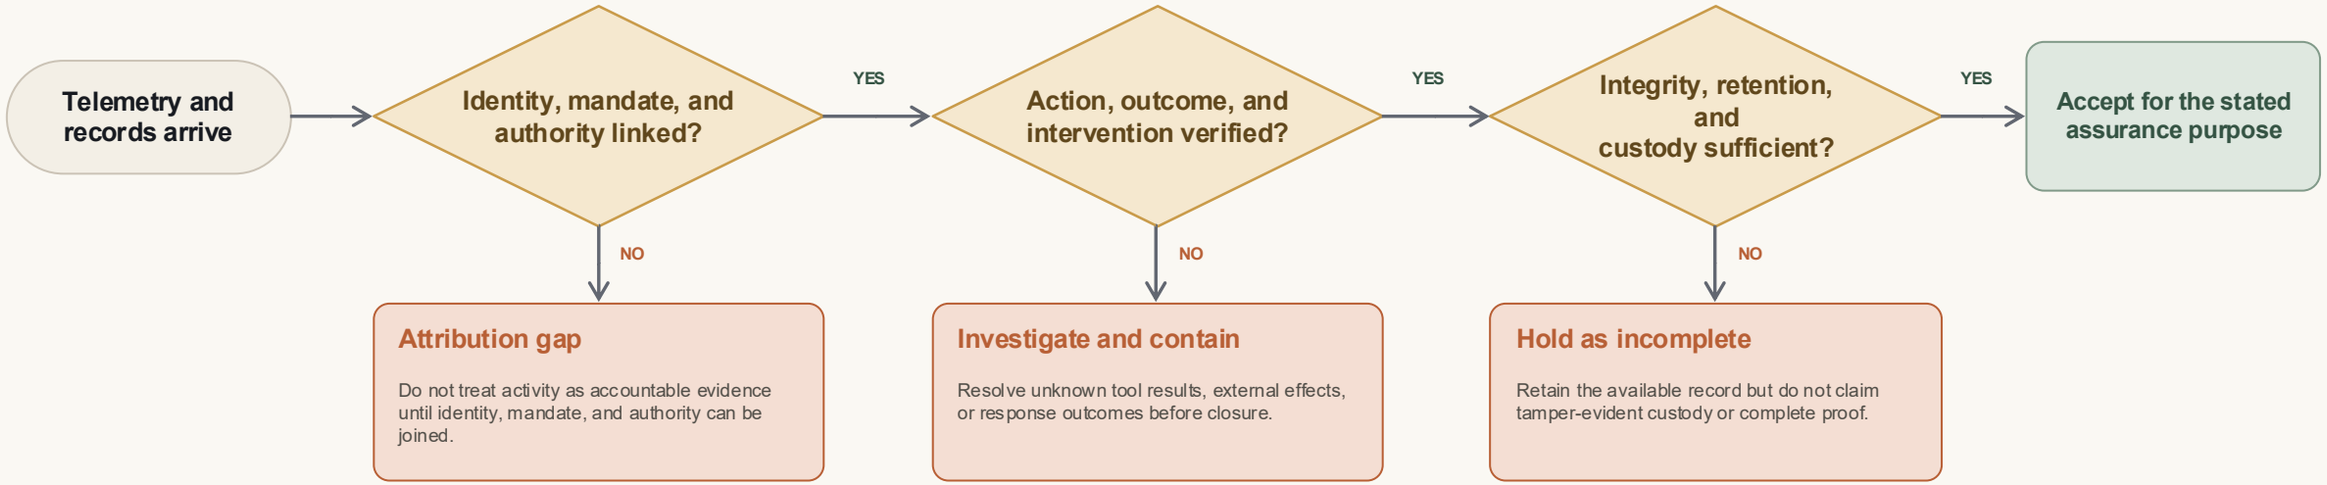
© 2026 Luminity Digital, Inc. | Local review draft

NIST continuous-monitoring tenet + OpenAI session events, cancellation, saved-state, and error-recovery docs · verified Nov 5, 2026

OpenAI Select Partner

Governed evidence acceptance

An evidence package earns trust only when attribution, outcome, intervention, integrity, retention, and chain of custody meet the stated assurance need.



REQUIRED OVERLAYS

Identity and mandate

Principal, workload, configuration, delegated intent, session, and turn.

Authority decision

Policy result, approval, credential reference, resource, audience, and expiry.

Action and outcome

Tool call, arguments, target decision, result, error, denial, and residual effect.

Monitoring response

Trigger, owner, cancel or revoke action, escalation, verification, and closure.

Retention and integrity

Verified origin, deduplication, timestamps, append-only write, custody, and retention.

ARCHITECTURE RULE **Observability becomes evidence after correlation, validation, and protected retention.**

OPENAI BOUNDARY A verified webhook proves origin. Traces and saved state remain evidence inputs, not an immutable chain of custody.

© 2026 Luminity Digital, Inc. | Local review draft

NIST immutable-evidence tenet + OpenAI observability, tracing, webhooks, cancellation, and audit-log docs · verified Nov 5, 2026

OpenAI Select Partner