

Agent Evidence Contract

A one-page operating record for AI agents that can change enterprise state.

Rule: a completed run names its authority, target, outcome proof, and recovery path.

Store references and redacted summaries by default. Do not copy sensitive raw content without a stated need.

1 Record this

Intent

Measurable task, expected result, requester reference, workflow, and risk class.

Authority

Draft, approval-required, or write. Capture policy version and approval reference.

Target

Validated customer, order, account, asset, or environment reference.

Context

Source IDs, versions, freshness, and permissions. Use references, hashes, or redacted summaries.

2 Prove this

Decision

Agent, model, prompt or policy version, selected tool, and deterministic validation result.

Action

Tool contract version, transaction or idempotency key, system of record, and before/after state reference.

Outcome

Independent system-of-record event or defined human check. An API 200 is not enough.

Recovery

Rollback, cancellation, manual remediation, incident label, and retention class.

WRITE-MODE RELEASE GATE

Do not let the agent write until all three are true.

1 Authority

Named action tier and approval reference

2 Target

Validated entity and policy result

3 Proof

System-of-record outcome and recovery handle

WHEN A RUN FAILS

1. Contain

Downgrade to approval mode and preserve the trace.

2. Reconstruct

Find the first wrong assumption, not only the last error.

3. Test

Sanitize the case and define the expected tool, target, and outcome.

4. Replay

Use dry-run tools. Restore write mode only after the full contract passes.