← Back to Sample Reports

SANITIZED REAL-REPOSITORY EXAMPLE

Evaluated source repository: Kestrel / AI Service Call Agent · commit 5e65843

Kestrel — AI Assurance Decision

A real-repository assurance example showing what the AI implementation can reach and cause, where authority evidence ends, and why the bounded Assurance Decision is REVIEW.

ASSURANCE DECISION

REVIEW

Bounded decision within evaluated scope and available evidence.

DECISION BASIS

HAIEC established **44 consequential action paths** the implementation is capable of reaching. Code capability is established for these paths, while policy authorization, effective runtime permission and observed execution were not established by the available evidence.

What this means: the implementation has meaningful action capability. The remaining assurance question is whether those capabilities are authorized, effectively permitted at runtime, and actually exercised as intended.

What to do next: prioritize the highest-consequence paths and establish policy, runtime permission and observation evidence before making a stronger assurance claim.

View System Constellation

View Full Assurance Report

Download / Evidence

44	44/44	1,322	555
CONSEQUENTIAL PATHS	CODE CAPABLE ESTABLISHED	EVIDENCE TRACES	EVIDENCE FRONTIER

WHAT TO DO NEXT

- → Connect policy and authority evidence
- → Connect runtime permission / credential evidence
- → Connect runtime execution evidence
- → Establish subject-to-tenant binding evidence
- → Establish the agent-to-capability bridge (evidence frontier)

ASSURANCE SURFACE

SYSTEM & AI SURFACE

Consequential action paths — Established
44 paths the implementation can reach

Code capability — Established
44 of 44 paths established as code-capable

Evidence traces — Established
1,322 source-backed proof records

Agent-to-capability bridge — Partial
The link from the AI agent to its code paths is not established in this snapshot — shown as evidence frontier, not assumed.

CONTROLS & CONTEXT

Tenant/context relations evaluated — Established
31 relations evaluated at path grain — not assumed globally

Tenant filter bound — Mixed
Tenant filtering verified on 30 of 31 applicable relations; assurance actions identified on 1

Subject-to-tenant binding — Evidence needed
31 applicable relations — subject-binding evidence needed

Paths with context evidence — Mixed
23 of 44 evaluated paths carry context evidence

AUTHORITY & RUNTIME — FIVE INDEPENDENT PLANES

Requested — Not assessed
Not assessed in this evaluation — evidence was not connected

Policy Authorized — Not assessed
Not assessed in this evaluation — evidence was not connected

Effectively Granted — Not assessed
Not assessed in this evaluation — evidence was not connected

RELEASE & EVIDENCE INTEGRITY

Producer completion — Completed
Declared analyzer scope completed

Source coverage — Truncated
Truncation declared — not silently partial

Analyzers not assessed — Not assessed
agentTopology, credentialUse

Ai Inventory Completeness — Not assessed

Privileged Action Authorization — Insufficient evidence

Untrusted Input Boundary — Insufficient evidence

Observed — Not assessed
Not assessed in this evaluation — evidence was not connected

Bundle integrity — Established
sha256 60d87c589323...

SYSTEM PROFILE

Kestrel — AI Service Call Agent

What HAIEC evaluated and the context bound to that evaluation.

SYSTEM
Kestrel — AI Service Call Agent
Evaluated

SOURCE REPOSITORY
subodhkc/AI-Service-Call-Agent-
Evaluated

REPOSITORY COMMIT
5e65843
Evaluated

EVALUATION SNAPSHOT
2026-09-14T08:38:17.814Z
Evaluated

Declared purpose
Purpose not declared / not available

Evaluated perimeter
44 consequential paths and their bound evidence relations

Methodology
1.1

Basis: evaluated snapshot · HAIEC-KESTREL-EVAL-5e65843 · repository source is not silently merged with current source.

SYSTEM COMPOSITION

How the evaluated AI-enabled system is assembled

Source-backed facts from the evaluated bundle and the existing Agents & Capabilities / reachability read models. Counts describe exactly what HAIEC represented in this evaluation — not active production behavior.

0
Agent entities represented in evaluated reachability

44
Consequential paths evaluated

9
Distinct tool candidates on consequential paths

4
Distinct handlers on consequential

8
Distinct sink targets on consequential

1,322
Evidence traces

CHECKED — NO QUALIFYING RESULT IN ANALYZED SCOPE
No qualifying result was discovered for memory lineage in the analyzed evaluated scope.

No qualifying result was discovered for model context influence in the analyzed evaluated scope.
No qualifying result was discovered for deferred execution in the analyzed evaluated scope.

Reachability: MIXED · evaluated reachability availability: PARTIAL

ASSURANCE COVERAGE

What HAIEC checked

Bounded check inventory from one versioned read projection (assurance-checks-1.0.0) — the checks represented in the evaluated bundle. Positive results are never inferred from missing findings.

3	10	0	7
Verified / Established	Mixed	Assurance actions identified	Evidence needed

3	9	0
None discovered	Not assessed	Not applicable

All | Verified | Mixed | Actions | Evidence needed | None discovered | Not assessed | Not applicable

SYSTEM & AI SURFACE

What consequential paths were evaluated?

ESTABLISHED

SYSTEM:CONSEQUENTIAL-PATHS

44 consequential paths evaluated.

View 44 affected paths in System Constellation →

What code capability was established?

VERIFIED

CAPABILITY:CODE-CAPABLE

Code capability established across 44/44 evaluated consequential paths.

View 44 affected paths in System Constellation →

MCP references were not assessed for this evaluation.

Next: Connect evaluated agent-definition evidence to assess MCP references

Limitation: AGENT_TOPOLOGY_NOT_ASSESSED

Agent Capability Constellation

NOT ASSESSED

SURFACE:AGENT_CAPABILITY_CONSTELLATION:SUMMARY

Agent Capability Constellation was not assessed in this evaluation.

No qualified agent capability constellation relations found in the analyzed evidence.

Next: Connect the corresponding producer/read-model evidence

Limitation: AGENT_CANDIDATE != AGENT_INSTANCE

Shared Substrate

UNKNOWN

SURFACE:REACHABILITY_SHARED_SUBSTRATE:SHARED_SUBSTRATE

Shared Substrate: result could not be determined from the connected evidence.

shared substrate coverage is unknown for the evaluated snapshot.

Limitation: SHARED_SUBSTRATE != CROSS_AGENT_EXFILTRATION

Resource Access

MIXED

SURFACE:REACHABILITY_SHARED_SUBSTRATE:RESOURCE_ACCESS

Resource Access: partial coverage — 590 relation(s) evaluated; some dimensions remain unresolved.

Partial resource access coverage; some dimensions remain unresolved.

Limitation: REACHABLE != ACCESSED

Persistence Creation

MIXED

SURFACE:REACHABILITY_SHARED_SUBSTRATE:PERSISTENCE_CREATION

Persistence Creation: partial coverage — 66 relation(s) evaluated; some dimensions remain unresolved.

Candidate persistence creation path(s) present; not established as observed.

Limitation: PERSISTENCE_CREATION != FUTURE_READ

Action Composition

MIXED

SURFACE:ACTION_COMPOSITION_DEFERRED_EXECUTION:ACTION_COMPOSITION

Action Composition: analysis ran with partial coverage; no qualifying relations were established.

Explicit action composition producer not currently established; bounded frontier.

Limitation: FUNCTION_CALL != COMPOSITION

Deferred Execution

None discovered

SURFACE:ACTION_COMPOSITION_DEFERRED_EXECUTION:DEFERRED_EXECUTION

No qualifying result was discovered for deferred execution in the analyzed evaluated scope.

CONTROLS & CONTEXT

Was tenant filtering bound on applicable relations?

MIXED

CONTEXT:TENANT-FILTER

Tenant filtering verified on 30 of 31 applicable relations; assurance actions identified on 1.

Next: Inspect exact tenant relations and establish remaining binding evidence

Limitation: TENANT_CONTEXT != TENANT_FILTER

Was subject-to-tenant binding established?

EVIDENCE NEEDED

CONTEXT:SUBJECT-BINDING

Subject-binding evidence needed on 31 applicable relations.

Next: Establish subject-to-tenant binding evidence

Limitation: TENANT_FILTER != TENANT_SUBJECT_BINDING

Memory Lineage

None discovered

SURFACE:MEMORY_CONTEXT_INTEGRITY:MEMORY_LINEAGE

No qualifying result was discovered for memory lineage in the analyzed evaluated scope.

No qualified memory lineage relations found in the analyzed evidence.

Limitation: MEMORY_ACCESS != DATA_EXFILTRATION

Model Context Influence

None discovered

SURFACE:MEMORY_CONTEXT_INTEGRITY:MODEL_CONTEXT_INFLUENCE

No qualifying result was discovered for model context influence in the analyzed evaluated scope.

No qualified model context influence relations found in the analyzed evidence.

Limitation: MODEL_CONTEXT_EXPOSURE != PROMPT_INJECTION

Approval Control

MIXED

SURFACE:APPROVAL_INTEGRITY:APPROVAL_CONTROL

Approval Control: partial coverage — 500 relation(s) evaluated; some dimensions remain unresolved.

Partial approval control coverage; some dimensions remain unresolved.

Limitation: CONTROL_PRESENT != PATH_BOUND

Approval Integrity

MIXED

SURFACE:APPROVAL_INTEGRITY:APPROVAL_INTEGRITY

Approval Integrity: analysis ran with partial coverage; no qualifying relations were established.

Approval Integrity coverage is unknown for the evaluated snapshot.

Limitation: UI_CONFIRMATION != SERVER_SIDE_ENFORCEMENT

SURFACE:ENFORCEMENT_COVERAGE:CONFIRMATION_MEDIATION

Confirmation Mediation: partial coverage — 138 relation(s) evaluated; some dimensions remain unresolved.

Partial confirmation mediation coverage; some dimensions remain unresolved.

Limitation: CONFIRMATION_MEDIATED != ACTION_EXECUTED

Approval Control

MIXED

SURFACE:ENFORCEMENT_COVERAGE:APPROVAL_CONTROL

Approval Control: partial coverage — 500 relation(s) evaluated; some dimensions remain unresolved.

Partial approval control coverage; some dimensions remain unresolved.

Limitation: CONTROL_PRESENT != PATH_BOUND

AUTHORITY & RUNTIME

Was the requested action scope established?

NOT ASSESSED

AUTHORITY:REQUESTED

Requested-scope evidence was not connected. Not assessed is an evidence state — not failure.

Next: Connect the corresponding evidence plane

Was policy authorization established?

NOT ASSESSED

AUTHORITY:POLICY-AUTHORIZED

Policy authorization evidence was not connected. Not assessed is an evidence state — not failure.

Next: Connect the corresponding evidence plane

Was effective runtime permission established?

NOT ASSESSED

AUTHORITY:EFFECTIVELY-GRANTED

Runtime permission evidence was not connected. Not assessed is an evidence state — not failure.

Next: Connect the corresponding evidence plane

Was observed execution established?

NOT ASSESSED

AUTHORITY:OBSERVED

Observed-execution evidence was not connected. Not assessed is an evidence state — not failure.

Next: Connect the corresponding evidence plane

Delegated Authority

MIXED

SURFACE:DELEGATED_AUTHORITY:SUMMARY

Delegated Authority: analysis ran with partial coverage; no qualifying relations were established.

Explicit delegated authority producer not currently established; bounded frontier.

Limitation: HANDOFF != AUTHORITY_DELEGATED

SURFACE:REPLAY_RETRY_DUPLICATE_EFFECT:SUMMARY

Replay, Retry & Duplicate Effect: analysis ran with partial coverage; no qualifying relations were established.

Explicit replay, retry & duplicate effect producer not currently established; bounded frontier.

Limitation: RETRY_CONFIGURED != RETRY_OCCURRED

Runtime-Validated Differential

NOT ASSESSED

SURFACE:RUNTIME_VALIDATED_DIFFERENTIAL:SUMMARY

Runtime-Validated Differential was not assessed in this evaluation.

No runtime corroboration for the evaluated runtime-validated differential dimension.

Next: Connect the corresponding producer/read-model evidence

Limitation: STATIC_PATH != RUNTIME_OCCURRENCE

RELEASE & EVIDENCE INTEGRITY

Does the published bundle carry a canonical digest?

VERIFIED

INTEGRITY:BUNDLE-DIGEST

Canonical bundle digest bound to this evaluation.

Did all declared analyzers assess their scope?

EVIDENCE NEEDED

INTEGRITY:ANALYZER-COVERAGE

2 declared analyzer(s) not assessed: agentTopology, credentialUse.

Next: Connect the not-assessed analyzer scope or record the limitation

Limitation: Evidence frontier is a deterministic projection of unestablished coverage; it is not a risk score.

Ai Inventory Completeness

NOT ASSESSED

CLAIM:AI-INVENTORY-COMPLETENESS

No qualifying evaluation of this claim occurred.

Privileged Action Authorization

EVIDENCE NEEDED

CLAIM:PRIVILEGED-ACTION-AUTHORIZATION

Insufficient evidence to determine this claim — evidence needed.

Secret Exposure Prevention

EVIDENCE NEEDED

CLAIM:SECRET-EXPOSURE-PREVENTION

Insufficient evidence to determine this claim — evidence needed.

Tenant Boundary Enforcement

EVIDENCE NEEDED

Untrusted Input Boundary

EVIDENCE NEEDED

CLAIM:UNTRUSTED-INPUT-BOUNDARY

Insufficient evidence to determine this claim — evidence needed.

Evidence Sufficiency & Uncertainty

NOT ASSESSED

SURFACE:EVIDENCE_SUFFICIENCY_UNCERTAINTY:SUMMARY

Evidence Sufficiency & Uncertainty was not assessed in this evaluation.

This evidence sufficiency & uncertainty dimension was not analyzed for the evaluated snapshot.

Next: Connect the corresponding producer/read-model evidence

Limitation: SURFACE_IMPLEMENTED != EVIDENCE_COMPLETE

Top Consequential Actions

Selected from the 44 evaluated action paths. Five planes are independent evidence states — NOT ASSESSED is not failure.

BUSINESS ACTION

CALL LIFECYCLE RECORD WRITE

`handle_create_booking ⟶ lifecycle_tracker.py · call_lifecycle table insert — writes a call-lifecycle event record`

Requested	○ NOT ASSESSED
Policy Authorized	○ NOT ASSESSED
Effectively Granted	○ NOT ASSESSED
Code Capable	● ESTABLISHED
Observed	○ NOT ASSESSED

Show in System Constellation →

BUSINESS ACTION

ORDER RECORD WRITE

`handle_place_order ⟶ tool_handlers.py · orders table insert — creates a customer order record`

Requested	○ NOT ASSESSED
Policy Authorized	○ NOT ASSESSED
Effectively Granted	○ NOT ASSESSED
Code Capable	● ESTABLISHED
Observed	○ NOT ASSESSED

Show in System Constellation →

BUSINESS ACTION

TENANT TIMEZONE CONFIG READ

`handle_create_booking ⟶ calendar_service.py · ai_agent_configs table query — reads tenant timezone configuration`

Requested	○ NOT ASSESSED
Policy Authorized	○ NOT ASSESSED
Observed	○ NOT ASSESSED

BUSINESS ACTION

TENANT LOCATIONS READ

`handle_create_booking ⟶ template_schema.py · locations table query — reads active tenant locations`

Requested	○ NOT ASSESSED
Policy Authorized	○ NOT ASSESSED
Observed	○ NOT ASSESSED

Show in System Constellation →

Show in System Constellation →

Show all 44 evaluated action paths

⚠ Evidence coverage is partial. Bounded and unknown states below are accurate — they are not failures and not passes.

Full Assurance Report

Executive | Technical | Auditor

Decision, consequences and required action — same evaluation, different presentation.

HAIEC

Assurance Report

HAIEC traces consequential AI action paths and compares code capability with the authority and evidence available for the same evaluated scope.

AI System: Kestrel — AI Service Call Agent
Evaluation ref: `HAIEC-KESTREL-EVAL-5e65843`
Source repository: `subodhkc/AI-Service-Call-Agent-`
Evaluation snapshot: 2026-09-14T08:38:17.814Z
Repository commit: `5e65843fddfe5f907485b798e464ad37b3b3b2c7`
Methodology: 1.1
Evidence Bundle Digest: `60d87c589323a82a…`

ASSURANCE DECISION

REVIEW

REVIEW — bounded to the evaluated scope and available evidence.

No qualifying evidence reference is bound to this statement.

EXECUTIVE SUMMARY

HAIEC evaluated this system and reached REVIEW. 44 path(s) not assessed on requested; 44 path(s) not assessed on policyAuthorized; 44 path(s) not assessed on effectivelyGranted; 44 path(s) established on codeCapable; 44 path(s) not assessed on observed

- 44 consequential action path(s) were evaluated; 44 have established code capability (what the code can reach or cause).
- 44 path(s) have established code capability without established authority evidence — policy authorization, effective permission, and observed execution remain separately stated, never inferred.
- 1322 supporting evidence trace(s) document how these paths were reconstructed.
- 31 canonical context-binding relation(s) were established; see Tenant & Subject Binding for whose context controls each action.
- 555 evidence frontier item(s) remain — see Evidence Coverage & Frontier.
- Next: Capture runtime evidence for action path 019bb0db3c13f0dd... before claiming observation.

No qualifying evidence reference is bound to this statement.

DECISION BASIS

Why the canonical assurance engine reached this disposition, from the evaluated claims bound to this exact evaluation.

- Privileged action authorization — Insufficient evidence. Coverage is UNKNOWN — absence of findings cannot become proof of absence.
- Secret exposure prevention — Insufficient evidence. Coverage is UNKNOWN — absence of findings cannot become proof of absence.
- Tenant boundary enforcement — Insufficient evidence. Coverage is UNKNOWN — absence of findings cannot become proof of absence.
- Untrusted input boundary — Insufficient evidence. Coverage is UNKNOWN — absence of findings cannot become proof of absence.
- Ai inventory completeness — Not assessed. No qualifying evaluation of this claim occurred.
- Evidence coverage was unknown for an evaluated claim.
- A claim was not evaluated.

▶ Inspect proof (1 reference)

ACTION ASSURANCE

Availability: partial. 6 frontier surface(s).

Limitations: CREDENTIAL_EFFECTIVE_GRANT: Credential attached but no source-qualified use — EFFECTIVELY_GRANTED = NOT_ASSESSED

No qualifying evidence reference is bound to this statement.

TENANT & SUBJECT BINDING

31 canonical context-binding relation(s) across 23 evaluated path(s). Tenant context detected on 9 relation(s); tenant restriction bound to action on 30; tenant tied to authenticated subject on 0; 31 relation(s) remain partial or unknown.

- Path 084ee19768725f5a… (TENANT_CONTEXT): tenant context detected; tenant restriction bound to action; tenant not tied to authenticated subject, value origin unknown — relation ad3_rel_kga9l8…, state partial.
- Path 084ee19768725f5a… (TENANT_CONTEXT): no tenant context detected; tenant restriction bound to action; tenant not tied to authenticated subject, value origin unknown, authentication auth not found — relation ad3_rel_utdz02…, state partial.
- Path 2cf2e1fd31e82f6d… (TENANT_CONTEXT): no tenant context detected; tenant restriction bound to action; tenant not tied to authenticated subject, value origin unknown, authentication auth not found — relation ad3_rel_7f6mte…, state partial.
- Path 2f6204531b481b20… (TENANT_CONTEXT): no tenant context detected; tenant restriction bound to action; tenant not tied to authenticated subject, value origin unknown, authentication auth not found — relation ad3_rel_qnq2wx…, state partial.

▶ Inspect proof (3 references)

MATERIAL MISMATCH / CONSEQUENCE

44 evaluated path(s) have established code capability without established authority evidence. 5 representative path(s) shown.

- Path 019bb0db3c13f0dd…: code capability is established, while effective permission has not been established from the evaluated evidence and policy authorization has not been established from the evaluated evidence.
- Path 084ee19768725f5a…: code capability is established, while effective permission has not been established from the evaluated evidence and policy authorization has not been established from the evaluated evidence.
- Path 1e33aa235efc0f17…: code capability is established, while effective permission has not been established from the evaluated evidence and policy authorization has not been established from the evaluated evidence.
- Path 2cf2e1fd31e82f6d…: code capability is established, while effective permission has not been established from the evaluated evidence and policy authorization has not been established from the evaluated evidence.
- Path 2f6204531b481b20…: code capability is established, while effective permission has not been established from the evaluated evidence and policy authorization has not been established from the evaluated evidence.
- + 39 additional path(s) not shown in this profile.

▶ Inspect proof (5 references)

REQUIRED EVIDENCE ACTIONS

Evidence-closure recommendations derived deterministically from missing or unsupported evidence states. These are report-derived recommendations, not canonical policy obligations.

- Capture runtime evidence for action path 019bb0db3c13f0dd... before claiming observation.
- Capture runtime evidence for action path 084ee19768725f5a... before claiming observation.
- Capture runtime evidence for action path 1e33aa235efc0f17... before claiming observation.
- Capture runtime evidence for action path 2cf2e1fd31e82f6d... before claiming observation.
- Capture runtime evidence for action path 2f6204531b481b20... before claiming observation.
- + 128 additional evidence-closure action(s) not shown in this profile.

No qualifying evidence reference is bound to this statement.

EVIDENCE COVERAGE & FRONTIER

555 frontier item(s): 554 partial, 1 unknown. 5 decision-relevant item(s) shown for this profile; 550 additional frontier item(s) remain in the full Technical/Auditor report.

- CONSEQUENCE_COMMIT_MODEL:handler-op:tool_impl_659:SINK:pydb_hvac_agent/app/services/calendar_service.py:50:22 — Partially established (REASON_NOT_ESTABLISHED)
- CONSEQUENCE_COMMIT_MODEL:handler-op:tool_impl_659:SINK:pydb_hvac_agent/app/services/template_schema.py:1140:18 — Partially established (REASON_NOT_ESTABLISHED)
- CONSEQUENCE_COMMIT_MODEL:handler-op:tool_impl_659:SINK:pydb_hvac_agent/app/services/template_schema.py:1294:18 — Partially established (REASON_NOT_ESTABLISHED)
- CONSEQUENCE_COMMIT_MODEL:handler-op:tool_impl_659:SINK:pydb_hvac_agent/app/services/transactional_sms.py:190:21 — Partially established (REASON_NOT_ESTABLISHED)
- CONSEQUENCE_COMMIT_MODEL:handler-op:tool_impl_659:SINK:pydb_hvac_agent/app/utils/lifecycle_tracker.py:117:9 — Partially established (REASON_NOT_ESTABLISHED)

Limitations: Evidence frontier is a deterministic projection of unestablished coverage; it is not a risk score.; Unknown != failure; SOURCE_GAP != risk.

▶ Inspect proof (1336 references)

CONSEQUENTIAL ACTION PATHS

5 of 44 evaluated consequential action path(s) shown. Selected for presentation; not an exhaustive absence claim.

▶ Inspect proof (5 references)

SYSTEM CONSTELLATION

A same-evaluation System Constellation projection is bound to this report for causal context. Selected action paths carry exact path identity for inspection in System Constellation.

CAPABILITY & AUTHORITY ALIGNMENT

Showing 5 of 44 evaluated path(s). Selected for presentation; not an exhaustive absence claim.

Path 2f6204531b481b20…

requested Not assessed	**policy Authorized** Not assessed
effectively Granted Not assessed	**code Capable** Established
observed Not assessed	

Path 019bb0db3c13f0dd…

requested Not assessed	**policy Authorized** Not assessed
effectively Granted Not assessed	**code Capable** Established
observed Not assessed	

Path 084ee19768725f5a…

requested Not assessed	**policy Authorized** Not assessed
effectively Granted Not assessed	**code Capable** Established
observed Not assessed	

Path 1e33aa235efc0f17…

requested Not assessed	**policy Authorized** Not assessed
effectively Granted Not assessed	**code Capable** Established

Path 2cf2e1fd31e82f6d…

requested Not assessed	**policy Authorized** Not assessed
effectively Granted Not assessed	**code Capable** Established
observed Not assessed	

This report is projected from the exact Assurance Evidence Bundle. It does not recompute the assurance decision, mutate canonical evidence states, or rerun analyzers.

Generated by HAIEC Evidence-Bound Assurance · Evaluation ref HAIEC-KESTREL-EVAL-5e65843 · Snapshot 2026-09-14T08:38:17.814Z · Source commit 5e65843fddfe5f907485b798e464ad37b3b3b2c7

High assurance in every consequence.

Evidence-Bound Assurance for consequential AI systems.

Explore Sample AI System →

Scope an Enterprise POC

— PRODUCT

Platform

System Constellation

— SOLUTIONS

Enterprise POC

Assurance Firms

HAIEC

Runtime Security Testing

AI Inventory

Evidence

Assurance

DEVELOPERS

Developers

Developer Security

AI AppSec

MCP Tenant Isolation

LLMVerify

Packages

RESOURCES

Learn

Resources

What Is HAIEC

How It Works

Methodology

Research

Regulations

Frameworks

Sample AI System

Sample Reports

Blog

Documentation

NYC Local Law 144

Colorado AI Act

HAIEC

COMPANY

Home

About

Contact

Trust

Privacy

Terms

Pricing

HAIEC Updates

Product updates, research, technical assurance insights and material regulatory developments.

you@company.com

Subscribe

By subscribing, you agree to our Privacy Policy. Unsubscribe anytime.

Evidence-bound assurance for consequential AI systems. See what AI systems and agents can reach, what controls apply, and what the evidence actually establishes.

HAIEC

Incident Response Vulnerability Disclosure Subprocessors Data Deletion Rights SLA Acceptable Use

© 2026 HAIEC. High assurance in every consequence. All rights reserved.