

What this means for clients

AI agents now act inside client systems. The first question every client asks is whether to trust what the agent did. The NIA AI Harness is designed so that the answer rests on evidence the client can inspect.

00 IN SHORT

Clients get verified outcomes, set their own limits and can prove both on request

Verified outcomes	Limits the client sets	Proof on request
Clients receive results that have been checked by an independent evaluator, not an agent's own report that the job is done.	Permissions, approvals and budgets are defined per workflow in a profile and enforced in code, not left to the model.	Evidence and replay give the board, the auditor and the regulator something concrete to inspect.

01 BUYER QUESTIONS

EXHIBIT 1

The harness answers the six questions buyers ask before they trust an agent

THE CLIENT ASKS	HOW THE HARNESS ANSWERS	THE CLIENT RECEIVES
Can we trust the result?	An independent evaluator checks each task against an agreed outcome contract. Missing proof produces INSUFFICIENT, never VERIFIED.	A verdict with reason codes and evidence references for every task
Can we limit what the agent can do?	Permissions are denied by default and enforced in code at the moment of action. Material actions wait for an authenticated human approval.	Scoped permissions, approval records and spending caps for each agent
Can we prove it to a board, auditor or regulator?	Evidence is hashed, linked and timestamped. Any run can be replayed from its frozen manifest without repeating its side effects.	An evidence manifest and, where issued, a signed AgenticAssure attestation
What does a result really cost?	Usage is attributed by tenant, product, agent, run, tool and model. Retries are never counted as extra successes.	Cost per verified outcome, not cost per token
Are we tied to one model vendor?	Model adapters keep the contracts constant. Governance applies whichever model runs, and behaviour is re-evaluated before a new model is promoted.	Freedom to change provider without rebuilding controls
Will our data stay separate from other clients?	Tenant isolation applies across memory, retrieval, tools, queues, logs and exports.	Tenant-scoped data and evidence

Source: AI Harness Build Blueprint v1.0. Design principles, evaluator verdicts, observability and the security model.

02 WHAT CLIENTS RECEIVE

EXHIBIT 2

Every governed task leaves a receipt the client can check

An outcome receipt states the verdict, which criteria passed and on what evidence, who approved what, and which versions ran. The example follows a customer ticket workflow.

OUTCOME RECEIPT

ILLUSTRATIVE EXAMPLE

Ticket T-123 resolved for the correct customer

OUTCOME CONTRACT
ticket-resolution-v1

EVALUATOR
assure-evaluator@2.1.0

VERDICT
VERIFIED

HUMAN APPROVAL
refund.issue approved; approver identity, time and evidence viewed recorded

CRITERION	TYPE	RESULT	EVIDENCE
correct_customer	Mandatory	pass	crm.customer_record, ticket.identity_match
policy_compliance	Mandatory	pass	policy_snapshot, decision_trace
recorded_action	Mandatory	pass	ticket.audit_log
customer_message	Optional	pass	message.delivery_receipt

policyHash sha256:... · evidenceManifestHash sha256:... · signature ...

Values are examples drawn from the reference workflow.

03 AUTHORITY

EXHIBIT 3

Authority is set per workflow: what the agent may do, what needs approval, what is never allowed

The reference customer-resolution profile shows the pattern. Each run also carries hard limits, so cost and time stay bounded.

May act	Needs human approval	Never permitted
Within scope, no approval needed	Held until an authenticated approver decides	Denied before any side effect
- knowledge.search - crm.customer.read - ticket.comment	refund.issue	- identity.modify - credential.export

LIMITS PER RUN	6	180 s	USD 0.50	12
	iterations	wall time	cost cap	tool calls

04 COMMITMENTS AND BOUNDARIES

What NIA commits to, and what it does not claim

NIA commits to

- Evidence comes before verdict. Confidence never replaces proof.
- The agent that does the work is never the only judge of it.
- Missing or conflicting evidence fails closed.
- Approvals are binding, authenticated and expire.
- Every consequential action can be traced and replayed.

NIA does not claim

- Zero human oversight or universal autonomy.
- Certification or regulatory compliance without an independent assessment. The harness maps controls and collects evidence.
- Identical outputs from different models. Behaviour is evaluated per model and version.
- That an attestation exists until a signed one has been received and verified.
- Readiness beyond what has been tested. Status is reported from repository evidence.

05 ONE FOUNDATION

The same controls, evidence and verdicts apply across every NIA product

A client who adds a second NIA product does not add a second governance model. The harness is the shared foundation beneath each of them. IntelliStudio is the product where AI employees run.

IntelliStudio

AI employee management

AgenticAssure

Independent testing and assurance

Future solutions

Onboarded through a profile, on the same controls

NEXT STEP

Discuss how the harness applies to your workflows

Contact the NIA team to review which of your agent workflows need approvals, evidence and independent verification.

contact@nia.sg

This document describes the design of the NIA Universal AI Harness and the commitments it is built to meet. It is not a certification, a compliance attestation or a statement of the implementation status of any specific deployment.