

NURO.is

Executive Policy

Artificial Intelligence Compliance & Governance Policy

The responsible development, deployment, and use of AI across all NURO operations, products, and client engagements. Human Intelligence leads. AI amplifies.

HI → AI = IE

Humans are the Source. Machines are the Force.

I. Purpose & Governing Doctrine	2
II. Scope & Definitions	2
III. Governing Principles - The HI Framework	4
IV. Regulatory & Standards Alignment.....	4
V. AI Tool Governance - Internal Use	5
VI. Data Handling & Confidentiality	6
VII. The Mandatory HI Review Step - Deliverables	7
VIII. Governance of AI Systems NURO Builds & Deploys	9
IX. Client Engagement Standards.....	10
X. Prohibited Uses - Zero Tolerance.....	11
XI. Human Oversight, Roles & Accountability	11
XII. Incident Response & Reporting	12
XIII. Monitoring, Audit & Continuous Improvement.....	12
XIV. Training & Awareness.....	12
XV. Reservation of Rights & Amendment	13
XVI. Affirmation	14
Acknowledgment & Agreement	15

I. Purpose & Governing Doctrine

NURO.is is an artificial intelligence company. We do not merely use AI - we design, build, deploy, and operate AI systems and autonomous agents on behalf of our clients. That dual position raises our obligation: we are held to the standard we sell. This Policy governs every way AI enters our work, from a team member drafting a proposal to an autonomous voice agent speaking with a client's customer.

This Policy is not a brake on innovation. It is the operating system that makes innovation safe to scale. It exists to ensure that every AI capability NURO touches is accurate, accountable, confidential, and under human command.

The Doctrine

All AI activity at NURO is governed by a single, non-negotiable doctrine drawn from our published framework, **HI → AI = IE** (Human Intelligence into Artificial Intelligence equals Innovation to the Power of Everything):

- **HI defines. AI amplifies.** AI is a probabilistic engine that mirrors and scales the clarity - or the chaos - of the human intelligence directing it. It does not understand meaning, bear responsibility, or face consequences.
- **HI governs. AI executes.** A human being remains the source of every decision, boundary, and command. The machine is the force that carries out the command.
- **HI is the Source. AI is the Force.** The Source must always lead the Force. AI is never a source of fact, law, judgment, or authority.

Because AI is an amplifier, it inherits the strengths and weaknesses of the human intelligence behind it. NURO therefore requires that the **Four Pillars of Human Intelligence** - Meaning, Identity, Responsibility, and Emotional Regulation - be stable before any AI capability is deployed. When these pillars are defined, AI is a force multiplier. When they are ambiguous, AI is a chaos amplifier.

The HI Step - The Core Safeguard of This Policy

No AI output enters any work product, client deliverable, client communication, or production system without first passing a Human Intelligence (HI) Review Step performed by an accountable person.

The HI Step is mandatory, non-delegable to the machine, and never waived for speed, urgency, or workload. A human reviews, verifies, and takes ownership before AI output is used. This single rule is the spine of everything that follows.

II. Scope & Definitions

A. Covered Personnel

This Policy is binding on every individual performing work on behalf of NURO or any NURO entity, regardless of role or employment status:

- Officers, employees, and members of NURO.is and all affiliated entities.

- Independent contractors, subcontractors, F-1/OPT and other visa-status workers, fractional staff, and interns.
- Third-party vendors, BPO partners, and consultants who use AI while performing services for NURO or its clients under NURO supervision.

No person acting on behalf of NURO may use, build, or deploy any AI capability except in full compliance with this Policy. **Where it is unclear whether a tool, system, or use is covered, the default is that it is covered.**

B. Three Operating Contexts

AI enters NURO's work in three distinct contexts. Every requirement in this Policy is mapped to one or more of them.

Context	Definition	Primary Governing Sections
USE	AI tools used internally by the team to produce work (writing, code, research, analysis).	V (Tools), VI (Data), XI (Training)
DELIVER	AI-assisted content that leaves NURO - proposals, reports, letters, code, client materials.	VII (HI Review Step), IX (Client Standards)
BUILD	AI systems, automations, and autonomous agents NURO designs, deploys, and operates for clients.	VIII (Agentic Governance), IX, XII (Incident)

C. Key Definitions

- **Generative AI / LLM** - Any system that produces text, code, images, audio, or analysis from a probabilistic model (e.g., Claude, ChatGPT, Gemini, Copilot, Grok), accessed via consumer interface, API, extension, or integration.
- **AI Agent** - An AI system granted tools, permissions, or the ability to take actions in real systems - sending messages, writing records, placing calls, executing workflows - with reduced or no human input per action. NURO's voice agents, GHJ automations, and APIP intelligence layers are agents.
- **Autonomous Action** - Any step an agent performs without a human approving that specific step before it occurs.
- **Bounded Autonomy** - The explicit, documented set of actions an agent is permitted to take on its own - and the boundary past which a human checkpoint is required.
- **Shadow AI** - Any AI tool used for NURO work without review and approval under Section V. Shadow AI is prohibited.
- **HI Review Step** - The mandatory human verification and ownership step defined in Section VII before AI output is used or shipped.
- **Client-Confidential Data** - Any information identifying, or provided by, a client or a client's customers - including CRE owner records, financials, PII, case facts, and contact data.

III. Governing Principles - The HI Framework

Ten non-negotiable principles govern every AI decision at NURO. Where any other instruction conflicts with these, these control.

01	AI is a tool, not an authority.	AI output is never a source of fact, law, professional judgment, or strategy. It is a starting point that a human must verify and own.
02	Human responsibility is non-delegable.	A named person remains fully responsible for every statement, deliverable, and system action - regardless of how it was generated or which model produced it.
03	The HI Step is mandatory.	No AI output enters work product, client communication, or a production system without a Human Intelligence review step. Non-negotiable.
04	Verification over trust.	No AI-generated citation, figure, claim, name, or legal/financial statement is relied upon without independent confirmation from a primary source.
05	Confidentiality is paramount.	AI platforms are not inherently confidential. Client and proprietary data is protected and only enters AI systems under contractual safeguards (Section VI).
06	Bounded autonomy by default.	Every agent NURO builds operates inside explicit decision boundaries, with human checkpoints on consequential actions. No agent runs unbounded.
07	Transparency and disclosure.	End users are told when they are interacting with an AI agent. AI involvement is disclosed wherever required by law, contract, or honest dealing.
08	Accuracy over efficiency.	Speed never justifies shipping unverified output. A hallucinated fact, citation, or figure that reaches a client is a compliance failure, not a time saving.
09	Accountability is assigned.	Every AI system and agent in production has a named business owner and a named technical owner before it goes live.
10	Clarity is the safeguard.	AI amplifies the clarity or the confusion of the human directing it. Precise meaning, scope, and instruction are themselves controls.

IV. Regulatory & Standards Alignment

NURO voluntarily aligns its AI governance with the frameworks that regulators, enterprise clients, and procurement teams now treat as the baseline. This alignment is both a risk control and a competitive credential we can demonstrate to clients.

Framework	What It Governs	How NURO Applies It
NIST AI RMF 1.0	Risk management lifecycle: Govern, Map, Measure, Manage.	The structure of this Policy and the AI System Register (Appendix C).

Framework	What It Governs	How NURO Applies It
ISO/IEC 42001:2023	AI Management System (AIMS): formal policy, roles, controls, continual improvement.	Roles in Section X; review and audit cadence in Section XIII.
EU AI Act	Risk-tiered obligations; transparency and human-oversight duties for high-risk and agentic systems.	Applied to any deliverable or agent with EU exposure; informs agent risk classes (Section VIII).
Colorado AI Act / state law	Duties around consequential automated decisions; ISO 42001 recognized as a safe-harbor signal.	Monitored for client jurisdictions; consequential-decision agents flagged for enhanced review.
OWASP LLM Top 10	Security risks for LLM and agent applications (prompt injection, excessive agency, data leakage).	Required threat checklist for every BUILD project (Section VIII.E).

Standards alignment is voluntary and risk-based. NURO does not claim formal certification unless and until an accredited audit is completed; until then, references to these frameworks describe the standards NURO designs toward, not a certification held.

V. AI Tool Governance - Internal Use

Goal: channel all AI use into approved, reviewed tools. Bans drive AI underground; an approval path keeps it visible and governed.

A. Tool Approval Tiers

Every AI tool used for NURO work falls into one of three tiers. The CAIO (or delegate) maintains the current Approved Tool List.

Tier	Criteria	Client-Confidential Data?
APPROVED	Enterprise/business agreement in place; no training on our inputs; encryption in transit and at rest; documented retention; DPA/BAA where applicable.	YES - subject to the HI Step and data classification in Section VI.
CONDITIONAL	Consumer/free tier; training opt-out available but not contractually guaranteed; no enterprise terms.	NO. Anonymized or hypothetical inputs only. No client or proprietary data.
PROHIBITED	Unclear retention or ownership; trains on all inputs; no confidentiality protection; unknown jurisdiction.	NEVER - for any purpose.

B. New-Tool Request Workflow

Before any new AI tool is used for NURO work, it must be requested and approved. This is the single most important control against shadow AI - it gives the team a fast, legitimate path instead of an unsanctioned one.

1. Submit a request to the CAIO/delegate naming the tool, the intended use, and the data it would touch.
2. The reviewer completes the Platform Risk Assessment (Section VI.A) and assigns a tier.

3. Approved tools are added to the Approved Tool List with their tier and any data conditions.
4. Decisions are returned quickly. Speed of approval is itself a control: slow approval recreates shadow AI.

C. Shadow AI - Prohibited

Shadow AI Is Prohibited

Using any AI tool for NURO work that has not been approved under this Section is a policy violation, even when well-intentioned and even when the output looks correct.

If you are mid-task and realize a tool is not on the Approved List, stop using it for NURO data and submit a request. Uncertainty defaults to: ask first.

D. Access & Account Hygiene

- AI tools are accessed through NURO-managed accounts, not personal logins, wherever an enterprise tier exists.
- API keys, model credentials, and integration secrets are stored in the managed secrets manager (Doppler) - never in code, chat, prompts, or shared documents.
- Access to AI systems and client data follows least privilege: people and agents get the minimum access needed, and access is removed at offboarding.

VI. Data Handling & Confidentiality

A. Platform Risk Assessment - Before First Use

Before any AI tool touches a client matter, the reviewer assesses and records:

- Whether inputs are retained after the session, for how long, and how they are purged.
- Whether inputs are used to train or fine-tune the model, and whether opt-out is contractual.
- Whether a DPA/BAA or enterprise confidentiality terms are in place.
- Encryption in transit and at rest; security posture (SOC 2 / ISO 27001 / ISO 42001 where available).
- Data residency / jurisdiction of stored data.

B. Data Classification

Class	Examples	AI Handling Rule
PUBLIC	Published marketing, public web content, NURO brand materials.	Any approved or conditional tool.
INTERNAL	Internal drafts, non-sensitive process docs, general code.	Approved tools; conditional tools only if no sensitive detail.
CLIENT-CONFIDENTIAL	CRE owner records, financials, client contact data, case facts, contracts.	APPROVED tools only, with enterprise/no-training terms. Never conditional or prohibited tools.
REGULATED / PII	SSNs, financial account data, medical/financial records, sealed or NDA-bound information.	Approved tools with DPA/BAA + explicit authorization only.

Class	Examples	AI Handling Rule
		Minimize and de-identify wherever possible.

C. Absolute Data Restrictions

Absent enterprise-level contractual confidentiality protections, the following must NEVER be entered into any AI tool, prompt, or agent:

- X Client names, contact details, or any personally identifiable information.
- X Financial account numbers, SSNs, EINs, loan numbers, or claim identifiers.
- X Client financial records and case facts tied to an identifiable party.
- X Trade secrets, proprietary NURO IP (scoring logic, thresholds, model prompts), or confidential client business data.
- X Anything under court seal, NDA, or other confidentiality obligation.

When in doubt, de-identify. Use hypotheticals and placeholders rather than real identifiers, even on approved platforms.

D. Data Flowing Through Systems We Build

Client and end-user data that passes through agents and pipelines NURO operates (Supabase, GHL, voice platforms, APiP, CRM layers) is governed by the same classification. For every BUILD project: data is segregated by client, retention is defined, third-party model terms are confirmed to prohibit training on client data, and the client's data-use expectations are documented in the SOW (Section IX).

VII. The Mandatory HI Review Step - Deliverables

Before any AI-assisted content leaves NURO, a named person verifies it and takes ownership of it. This is the core safeguard for everything we ship.

A. What Must Pass the HI Step

Any deliverable that leaves NURO and was touched by AI at any stage - proposals, reports, client letters, marketing copy, financial figures, legal-adjacent language, code shipped to a client, and any client-facing AI output.

B. Verification Protocol

- **Retrieve & confirm.** Every factual claim, statute, figure, name, citation, or data point is confirmed against a primary source - not an AI summary. A citation that cannot be located does not get used. For regulatory and government program matters (SBA EIDL/Disaster Loans, ERC/ERTC claims, Florida Executive Orders, IRS guidance, CDC/OSHA): retrieve current official text directly from sba.gov, irs.gov, fl.gov, Florida Administrative Register, or official court/government sites; confirm version, effective date, and applicability to client facts only after independent HI verification of mandate scope and preemption.
- **Read in full.** The reviewer reads the entire output and independently assesses it. No skimming AI output and forwarding it.

- **Confidentiality check.** Confirm no confidential or PII data was exposed to an unapproved tool during creation.
- **Brand & accuracy check.** Confirm the deliverable reflects NURO standards (including brand rules - e.g., NURO black-and-gold; APIP's own brand system) and the team's own judgment, not the model's defaults.
- **Own it.** The reviewer must be prepared to stand behind every word as NURO's own work.

C. Labeling & Versioning

- AI-generated drafts are labeled "AI Draft - Unverified" until the HI Step is complete, and kept distinct from the reviewed final.
- For substantive client deliverables, the reviewer is recorded (who reviewed, when). This is the deliverable equivalent of a sign-off.

HI Step - Deliverable Checklist

- ✓ Every fact, figure, name, and citation independently verified from a primary source.
- ✓ No confidential or PII data was exposed to an unapproved tool.
- ✓ Output read in full and reflects NURO's independent judgment and brand standards.
- ✓ No unverified AI-generated content remains in the final.
- ✓ Reviewer is prepared to stand behind it as NURO's own work.

VIII. Governance of AI Systems NURO Builds & Deploys

This is the section that distinguishes NURO from a company that merely uses AI. When AI can act - not just advise - the obligation shifts from accuracy to control. No NURO agent reaches production without the following. An AI Impact Assessment (covering purpose, affected parties, rights/financial impact, bias/fairness exposure especially for qualification agents, regulatory exposure, and mitigations) is completed and CAIO-approved before design begins.

A. Named Ownership - Before Production

Every agent, automation, or AI product NURO deploys has, before it goes live:

- **A Business Owner** - accountable for the business outcome and the client relationship.
- **A Technical Owner** - accountable for the system's behavior, safety, monitoring, and rollback.

B. Agent Risk Classification

Each system is classified at design time. Higher risk demands tighter human checkpoints.

Class	Definition	Required Oversight
LOW	Informational only; no external action, no consequential decision (e.g., internal summarizer).	Standard testing; human-on-the-loop review of outputs.
ELEVATED	Acts on external systems or contacts (e.g., outbound messaging, lead qualification, scheduling).	Bounded autonomy + defined human checkpoints; logging; disclosure to end users.
HIGH	Touches money, legal/financial determinations, regulated data, or consequential decisions about a person.	Human-in-the-loop approval per consequential action; enhanced testing; client sign-off; kill switch. For any qualification, intake, or scoring agent: explicit bias/fairness review required in pre-deployment testing and periodic audits.

C. Bounded Autonomy & Human Checkpoints

- Every agent ships with a documented boundary: the specific actions it may take autonomously, and the actions that require a human to approve before they occur.
- Consequential actions - anything financial, legal, contractual, or that materially affects a person - require a human checkpoint. Agents qualify, draft, route, and prepare; humans decide and commit.
- Agents are never granted authority to make legal or financial determinations on a client's behalf. (E.g., an intake or qualification agent surfaces a candidate; a person makes the accept/decline call.)

D. Disclosure to End Users

- People interacting with a NURO-built agent are told they are interacting with AI. Voice and chat agents disclose their AI nature; they do not impersonate a specific named human.

- Disclosure language and any jurisdiction-specific bot-disclosure requirements are confirmed before launch.

E. Security - OWASP LLM Threats

Before deployment, the technical owner confirms defenses against the agent-specific failure modes:

- **Excessive agency.** Tools and permissions are scoped to least privilege; the agent cannot take actions outside its defined boundary.
- **Prompt injection.** Untrusted content (web pages, inbound messages, documents) cannot redirect the agent's instructions or exfiltrate data; inputs are treated as data, not commands.
- **Data leakage.** The agent cannot expose secrets, other clients' data, or system prompts. Outputs are constrained to its purpose.

F. Pre-Deployment Testing, Logging & Rollback

- Each agent is tested against expected and adversarial inputs before launch; results are recorded.
- Production agents log their actions with enough traceability to reconstruct what they did and why, and to tie any action back to a human decision-maker.
- Every ELEVATED and HIGH agent has a defined kill switch / pause and rollback procedure, and a named person who can invoke it.
- Each deployed system is entered in the AI System Register (Appendix C).
-
- G. Post-Deployment Monitoring, Drift Detection & Prompt Governance
- Production agents undergo ongoing monitoring for output quality drift, performance degradation, and boundary violations. The Technical Owner reviews action logs at least weekly for ELEVATED and HIGH agents (monthly for LOW). Prompt or system prompt modifications require documented CAIO approval and re-testing against adversarial cases before deployment. Provider model updates trigger immediate re-validation of risk classification, boundaries, and HI checkpoints. Where feasible, deterministic settings (e.g., temperature=0) are used and all generation parameters logged for reproducibility and audit.

IX. Client Engagement Standards

When AI is part of a client engagement - whether in a deliverable or a system we build and operate - the engagement documents (SOW/contract) address:

- Ownership of deliverables, models, prompts, and data, and any client IP that the system processes.
- Data use and confidentiality: what client data the system touches, where it is stored, retention, and a prohibition on third-party model training on client data.
- Accuracy and oversight expectations: the client understands AI assists and a human verifies; NURO does not warrant AI output as legal, financial, or professional advice.
- Disclosure: end users of a deployed agent are told they are interacting with AI.
- Handover responsibilities: if a client operates an agent after handover, the boundary, monitoring, and kill-switch responsibilities are documented and transferred in writing.

NURO does not deploy ungoverned or unbounded agents under its name or a client's. The standard in this Policy travels with every system we ship.

X. Prohibited Uses - Zero Tolerance

The following are categorically prohibited. There are no exceptions for time pressure, client urgency, or workload.

- X Presenting AI-generated facts, citations, figures, or names in any deliverable without independent verification from a primary source.
- X Entering client-confidential data, PII, or proprietary NURO IP into a conditional or prohibited AI tool, or any tool lacking enterprise confidentiality terms.
- X Shipping any client deliverable that has not passed the HI Review Step (Section VII).
- X Deploying an AI agent to production without a named business owner, technical owner, documented boundaries, and risk classification.
- X Granting an agent autonomous authority over consequential financial, legal, or contractual actions without a human checkpoint.
- X Letting an AI system make legal or financial determinations about a person on a client's behalf.
- X Deploying an agent that impersonates a specific named human or conceals that it is AI where disclosure is expected or required.
- X Using AI to generate deceptive, fabricated, or misleading content, evidence, or representations to any client, customer, or authority.
- X Using shadow AI - any unapproved tool - for NURO work.
- X Storing model keys, secrets, or credentials in prompts, code, or shared documents instead of the managed secrets manager.

Violation may result in disciplinary action up to and including termination of employment or contract, and may carry personal and contractual liability.

XI. Human Oversight, Roles & Accountability

AI governance only works when accountability is named. The following roles own this Policy at NURO:

Role	Accountability
Chairman/ CAIO Craig Pizaris-Henderson	Owns this Policy and the AI governance program. Maintains the Approved Tool List and AI System Register; approves tools and agent risk classes; final authority on AI-specific risk decisions.
Chief Executive Officer Chase Henderson	Ultimate business accountability for AI use and deployment; approves the Policy and material exceptions; owns client-facing commitments.
Chief Operating Officer Kristin Norton	Operational implementation: onboarding, training records, vendor/BPO compliance, incident log, and policy acknowledgments.
Business / Technical Owners (per system)	Named per deployed agent (Section VIII); accountable for that system's outcome, behavior, monitoring, and rollback.

Role	Accountability
All Personnel	Perform the HI Step, use only approved tools, protect data, and report incidents. Responsibility for AI-assisted work is personal and non-delegable.

XII. Incident Response & Reporting

An AI incident includes a fabricated fact reaching a client, confidential data exposed to an unapproved tool, an agent acting outside its boundary, or any AI-related security or accuracy failure.

1. Report to the CAIO and COO within 24 hours of discovery. Do not conceal or delay.
2. Contain: pause or invoke the kill switch on any misbehaving agent; stop further use of any tool that exposed data.
3. Log the incident in the confidential AI Incident Log (what happened, systems and data involved, who was affected).
4. Root cause: determine whether the HI Step was skipped, performed inadequately, or structurally insufficient, and whether a boundary or control failed.
5. Notify affected clients where contract, law, or honest dealing requires it; correct any inaccuracy that reached a client without regard to convenience.
6. Remediate: strengthen the failed control training; update this Policy within 30 days of any material incident.

XIII. Monitoring, Audit & Continuous Improvement

- **AI System Register.** NURO maintains a living inventory of every AI tool in use and every agent in production (Appendix C), reviewed at least quarterly. This satisfies the “Map” function of NIST AI RMF and the inventory expectations of ISO 42001.
- **Quarterly review.** The CAIO reviews the Approved Tool List, the Register, open incidents, and any new regulation or guidance relevant to client jurisdictions.
- **Annual policy review.** This Policy is reviewed at least annually.

This Policy is additionally reviewed and updated whenever: (a) a new AI tool or agent class enters production; (b) new applicable regulation or guidance is issued; (c) a material AI incident occurs; or (d) AI capability changes create a risk category this Policy does not address.

XIV. Training & Awareness

Training	Scope	Frequency
Policy Onboarding	Full policy, the HI Step, data rules, prohibited uses, tool request workflow.	At hire/engagement, before any AI use.
Hallucination & Verification	How models fabricate; how to verify facts, figures, and citations.	Annual minimum.

Training	Scope	Frequency
Agent-Builder Training	Bounded autonomy, risk classes, OWASP LLM threats, logging, kill switch - for anyone who builds or operates agents.	Before building; annual refresh.
Tool-Specific Briefing	Features, limits, and data practices of each newly approved tool.	On approval of each new tool.

XV. Reservation of Rights & Amendment

NURO expressly reserves the following rights with respect to this Policy. These provisions are in addition to, and do not limit, the review and revision process in Section XIII.

A. Amendment at Sole Discretion

NURO may amend, modify, supplement, suspend, or revoke this Policy, in whole or in part, at any time and for any reason, in its sole and absolute discretion, with or without prior notice. Authority to issue amendments rests with the CAIO and the CEO. An amendment is effective immediately upon issuance unless a later effective date is stated.

B. Application to Ongoing Work & Newly-Known Standards

Each amendment governs all work, deliverables, AI tools, and deployed systems from its effective date forward, including those already in progress. Where AI capabilities, risks, laws, or guidance emerge that were not known when earlier work was performed, NURO reserves the right to apply the updated Policy to existing tools, deliverables, and deployed agents - and to require that prior work be re-reviewed, corrected, paused, or brought into compliance - regardless of whether the work met the standards in effect at the time it was done. The version of this Policy in effect at the time of evaluation governs the required state of any tool, deliverable, or system going forward.

C. Retroactive Application

NURO reserves the right to apply this Policy, and any amendment to it, retroactively - including to work, deliverables, communications, AI tools, and deployed systems created or performed before the amendment's effective date, and including conduct that complied with the standards in effect at the time it occurred. NURO may evaluate, re-characterize, and act upon prior conduct under the current version of this Policy, and may require remediation of, or take corrective or disciplinary action in connection with, such prior conduct. No person acquires a right to have past conduct judged solely under a prior or superseded version of this Policy.

D. No Vested Rights; Not a Contract

This Policy is a statement of NURO's internal governance standards. It is not a contract, creates no contractual or vested rights, and does not alter the at-will nature of any employment relationship or the terms of any contractor or vendor agreement. No prior version, prior practice, or prior approval limits NURO's right to apply the current Policy. To the extent this Policy conflicts with a signed agreement, the signed agreement controls as to that party.

E. Acceptance Through Continued Work

Continued performance of work on behalf of NURO after any amendment constitutes acceptance of the amended Policy. Each person is responsible for complying with the most current version, which NURO will make available.

F. Controlling Version

The most current version maintained by the CAIO is the controlling version and supersedes any printed, cached, forwarded, or previously distributed copy. Where any conflict exists between copies, the latest issued version governs.

XVI. Affirmation

At NURO, the following are absolute and without exception:

- ✓ AI does not replace human judgment. Judgment remains human.
- ✓ Every fact, figure, and citation that reaches a client is independently verified.
- ✓ Every agent in production is bounded, owned, logged, and reversible.
- ✓ NURO accepts full responsibility for everything it ships and everything it deploys.
- ✓ The HI Step is never optional, never delegable to the machine, and never waived for speed.

HI → AI = I^E · Humans are the Source. Machines are the Force. · NURO.is

Acknowledgment & Agreement

I have read, understand, and agree to comply with the NURO.is Artificial Intelligence Compliance & Governance Policy. I understand that responsibility for AI-assisted work is personal and non-delegable, and that violation may result in disciplinary action up to and including termination of employment or contract and may carry personal and contractual liability. I acknowledge that NURO may amend this Policy at any time at its sole discretion, that amendments apply to ongoing and future work, and that my continued work after an amendment constitutes acceptance of the then-current version.