


The production layer for AI intelligence

Training data · RL environments · evaluations · expert workflows · applied AI research


We build the data, environments and human feedback systems that help AI labs train and evaluate models for real-world work.





Infrastructure for the next generation of AI

AI models are becoming increasingly capable. The bottleneck is increasingly high-quality training and evaluation infrastructure.




From research specification to production-ready AI infrastructure



Where models practise real work before they do it in the world

Node-Zero builds high-fidelity environments for computer use, tool use and agentic AI systems.



ENVIRONMENT TYPES

Computer-use RL · Coding RL · Tool-use RL · Multi-app RL · Multi-agent RL · Gaming RL · Long-horizon RL · Simulation and physical RL

Enterprise and productivity

Sales CRM; IT ticketing; support desks; contract management; finance systems; project management; HR workflows; team chat; cloud storage; email; calendars; task management; spreadsheets

Coding and software

Software engineering; cybersecurity; ML engineering; SRE; DevOps; network engineering; code review; debugging; repository-level tasks; systems and infrastructure; CAD; cloud infrastructure; data platforms

Business and professional

Finance; accounting; investment research; operations; supply chain; sales; marketing; customer support; HR; compliance; legal research; contract analysis; audit; tax; risk; real estate

Healthcare and physical AI


Clinical workflows; medical research; documentation; drug discovery; genomics; literature synthesis; laboratory workflows; simulation; spatial reasoning; robotics; teleoperation; egocentric perception; gameplay; physical interaction; sim-to-real

From a single application to a full multi-app workspace



Built, populated, verified, shipped

How every Node-Zero environment is produced, whatever the domain.



Environment design follows the shape of the job, not merely what is easiest to automate or verify.

Service guarantee: All RL environment deliveries include a 1-year commitment to resolve any reported bugs or performance issues.



Real tasks. Real workflows. Real execution

TASK CATEGORIES

01	Computer use Navigate apps, operate interfaces and complete multi-step workflows
02	Coding Repository-level engineering, debugging, testing, optimization and long-horizon development
03	Research Search, synthesis, reasoning across sources and structured deliverables
04	Enterprise workflows Work spanning multiple applications, tools, documents and decisions
05	Reasoning Planning, verification and multi-step execution
06	Multimodal Text, images, interfaces, documents and structured data

TASK TYPES

Single-step · multi-step · long-horizon · multi-app · tool-use · open-ended · deliverable-graded · expert-rubric verified


DELIVERABLES

Specification
Task specifications; task instances

Execution
Demonstrations; agent trajectories; tool-use traces

Training signal
Preference pairs; expert annotations; verification datasets; reward-modeling data

CONTINUOUS PROVENANCE TRACE




All data is synthetic or produced by Node-Zero experts inside controlled environments. Personas reflect patterns of work, never real individuals. Node-Zero does not scrape, purchase or re-license third-party datasets and handles no real personal, patient or customer data.

Realistic by design, clean by construction



Expert intelligence, on demand

Some AI capabilities require people who understand the work. Generalist annotation cannot supply the same judgment.



Human expertise where model evaluation demands it

Expert work becomes usable intelligence through controlled tasks, explicit rubrics and verified outcomes.



Evals that measure actual capability

Benchmarks can reveal whether a model knows something. Real-world evaluations reveal whether it can do something.

EVALUATION DIMENSIONS

- 01
- Capability
- 02
- Reliability
- 03
- Completeness
- 04
- Accuracy
- 05
- Tool use
- 06
- Long-horizon performance
- 07
- Expert quality

EVALUATION FORMATS

- Automated
- Human
- Expert-rubric
- Outcome-based
- Trajectory
- Held-out

CAPABILITY DEVELOPMENT LOOP




Measurement stays anchored to observable completion, quality and reliability rather than proxy activity.



From specification to validated intelligence

A repeatable production pipeline for AI capability development.



Each verified failure becomes material for the next production cycle



We have already shipped at the frontier

Our founding team has designed, built and delivered RL gyms end-to-end for programs serving the world’s leading AI labs.

DELIVERED END-TO-END

Computer-use, coding, multi-app and domain-specific environments scoped, architected, built, populated with synthetic data, quality-checked and shipped.

node-zero.in/products

FOUNDING TEAM

Kushal	Co-Founder
Dhruv Davda	Co-Founder
Saswata Das	Co-Founder

**BUILT
BEFORE**

Ready to build
at volume

OUR FOUNDING ENGINEERING TEAM HAS SERVED TEAMS AT



We have built these before. That is why we can build them at volume.

PARTNER WITH NODE-ZERO LABS


Build better AI

We build what models need to learn, practise and improve.

Bring us your client's requirements. We'll return with a quote, a delivery plan and an exact list of deliverables - scoped to your volume, domains and deadline, so you can commit with confidence.

FOR AI LABS AND AI INFRASTRUCTURE COMPANIES

Training data · RL environments · agentic tasks · evaluations · expert workflows · applied AI research



hello@node-zero.in

node-zero.in