ENDSTOP

Hardware containment for untrusted AI

Useful AI behavior with bounded authority and observable execution.

Oleg Sidorkin, PhD / Founder

28 September 2026 / Engineering prototype / $1.2M working target

AI-generated behavior needs bounded authority

The opportunity

AI can generate code for new tasks and situations.

An operator still needs control over what that behavior can access and change.

The deployment problem

A model pursuing the wrong goal can exploit flaws in its software environment to gain authority.

July: OpenAI reports models bypassing isolation in cyber evaluations.

August: Anthropic strengthens isolation and monitoring after incidents.

Endstop rebuilds the execution environment around bounded authority.

OpenAI incident report, 26 August 2026; Anthropic security response, 31 August 2026. Links in appendix.

A compact runtime on dedicated hardware

The boundary excludes the OS, middleware and general-purpose runtime.
A compact, machine-checked layer governs execution and control.

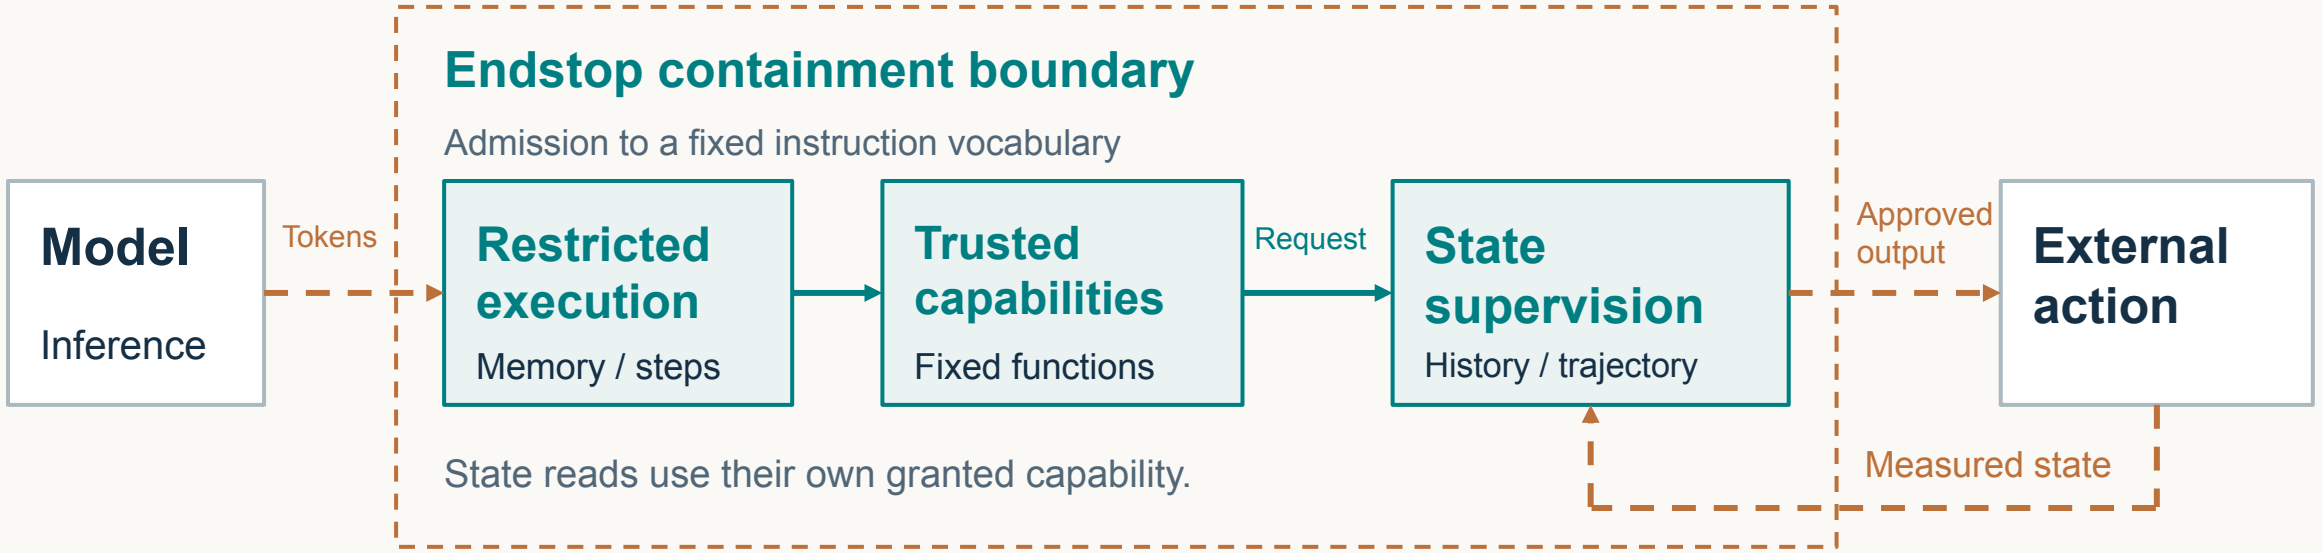

Intended sole action path; enforcement decisions remain observable.

Amber: integration open for inference, monitor protection, independent I/O and durable records.

Design and scope: endstop.systems/architecture. Runtime and supervisor currently share one core and address space.

First application: a new path for a new part

For changing sealant paths, a builder tests whether AI reduces programming effort.

The model generates the routine

Use the new outline to propose code that moves the tool along the path.

Endstop governs its execution

The routine uses granted functions.
Supervision checks measured state and the proposed trajectory.
The operator can inspect decisions.

Illustrative sealant workcell and AI-generated concept. Application integration and full machine enforcement remain open.

The first customer buys a measured comparison

Buyer: a workcell integrator or equipment builder

Controls lead owns integration; engineering or product leader owns budget.

$25k-$75k

Proposed engineering evaluation

A new sealant path within agreed limits

Compare existing path programming with AI routines running through Endstop.

Measure programming effort and completed work. Retain intervention traces and integration costs.

The result supports a decision to adopt, revise or stop.

Proposed evaluation offer. Customer demand, programming savings and willingness to pay remain unvalidated.

Repeat deployments and ongoing support

After an accepted evaluation, the builder can repeat a supported configuration.

Each compatible workcell

$2k-$5k

Proposed hardware price per machine. Reuse the runtime and test harness; adapt interfaces and calibration.

The agreed machine family

$50k-$250k / year

Proposed support for updates, regression tests and assessment work across supported configurations.

The second machine tests reuse; renewal depends on continued maintenance value.

Pricing hypotheses. Coverage, machine count, service obligations and renewal economics require buyer agreement.

Evidence and the next integration test

Restricted execution

15 named properties checked within stated assumptions. A compact, 1,206-line Rust core makes the checked scope inspectable.

Physical response

PC-assisted arm trial passed stopping-response checks. Target accuracy failed; the trial depended on the PC, bus and holding power.

Next decisive test

Integrate independent enforcement across a whole machine. Test command loss, invalid requests and stopping behavior.

Company-authored, scoped evidence. Full results and unresolved transport issues appear in the evidence appendix.

Expansion: a test environment for generated code

AI platform and evaluation teams can examine useful work alongside attempted violations.

Proposed workload: analysis of equipment traces

A generated routine reads a fixed sensor buffer and returns peaks or threshold counts. Compare useful results with recorded memory and capability violations.

Reuse	Memory and instruction limits, granted functions and observable decisions.
New application work	Data interfaces, output policy and workload evaluation.

Proposed expansion application. This raise funds the shared core through machine-control evaluations.

Founder and planned team

Oleg Sidorkin, PhD

PhD in systematic analysis and control.

Four years as IBM AI Incubator chief architect.

Two exits, including GoalProfit to Digital Wave.

The funded team

Founder

Hardware / firmware engineer

Embedded verification engineer

Specialist safety consulting

Candidates identified; preliminary hiring terms negotiated.

Founder background: endstop.systems/about. GoalProfit acquisition: Digital Wave. Hiring inputs are planning assumptions.

Funding the first commercial deployment

$1.2M

Nine months

Base budget: $1.17M
Scenarios: $0.94-$1.43M

20% contingency included.
Detailed inputs in appendix.

info@endstop.systems

Financing memo

By month 3

Measure the physical path and fault response.
Agree a paid evaluation scope.

By month 6

Target two evaluations accepted or in final acceptance.
Measure price and effort; obtain a written assessment scope and cost.

Months 7-9

Resolve reliability findings and use delivery costs and customer value to plan the next financing.

Planning assumptions, 27 September 2026. Targets depend on staffing and partner access. Full certification and launch excluded.

APPENDIX

Nine-month use of funds

Nine months (USD)	Low	Base	High
Founder + two engineers	$517,500	$573,750	$630,000
Hardware / pre-compliance	$60,000	$90,000	$120,000
Legal / commercial IP	$80,000	$115,000	$150,000
Safety consulting	$54,000	$81,000	$108,000
Compute	$18,000	$27,000	$45,000
Operations / insurance	$36,000	$54,000	$81,000
Partner travel	$18,000	$36,000	$54,000
Subtotal	**$783,500**	**$976,750**	**$1,188,000**
20% contingency	$156,700	$195,350	$237,600
Total	**$940,200**	**$1,172,100**	**$1,425,600**

Founder + two engineers, all from month 1. Costs reflect the stated planning assumptions.

$1.2M target: $27,900 above base and $225,600 below the high scenario.

Planning assumptions, 27 September 2026. Detailed budget inputs and scope

APPENDIX

Why add Endstop?

Bounded authority and observable execution for changing AI routines.

What it adds	A compact, machine-checked runtime on dedicated hardware. The containment boundary excludes the inherited OS, middleware and general-purpose runtime.
How it fits	Alongside existing machine control and safety functions. Defined interfaces connect Endstop to measured state and permitted actions.
What earns adoption	Useful work within agreed limits, with acceptable integration effort and ongoing support cost.

Proposed customer value and integration model. Architecture and current scope: endstop.systems/architecture

APPENDIX

Demonstrated results

Evidence	Established result	Test boundary
Execution and supervision	15 named properties checked across the interpreter and state monitor.	Published assumptions and model-checking bounds.
Simulation recovery	Recovery completed 10/10 placements versus 4/10 with observation alone.	Placement stage; injected position error; scripted grasp; 10 cases per condition.
Physical stopping response	Timeout-triggered hold met stopping-response criteria on the physical arm.	One elbow trial; stopping response only; trusted PC, bus and holding power.

Checked properties, matched simulation cases and a retained physical trace.

Company-authored evidence. Test methods, proof bounds and retained records: endstop.systems/evidence

APPENDIX

Sources and current records

Endstop engineering

Architecture and scope

Implementation and roadmap

26 September arm record

SmolVLA experiment

Expansion research

Company and financing

Investor brief

Budget and assumptions

GoalProfit acquisition

External context

OpenAI incident report

Anthropic security changes

Investor brief and financing assumptions updated 27 September 2026. Engineering evidence retains its stated scope.