

Adaptive autonomy. Fixed authority.

A beautiful, inspectable machine brain: it learns from failure without retraining neural weights, runs on the machine, and stays inside explicit human-set limits.

US\$5-10M


seed round

24 mo

runway, with written
gates

5-7

core team, small on
purpose



Autonomy is scaling. The hard part is the brain.

The opportunity is not more models. It is a fieldable intelligence layer that runs locally, remembers what failed, replays why it acted, and connects to the systems already in service.

Autonomy is already a priority

DoD's Replicator effort made large-scale autonomous systems explicit. Autonomous systems are moving from boutique projects toward fleet-scale deployment.

The environment is contested


Fielded systems cannot assume a perfect cloud link, a GPU rack, or constant reach-back. The autonomy layer has to keep working when bandwidth is poor, delayed, degraded or denied.

The installed base is real

Programs still depend on old controllers, buses and legacy software. Rip-and-replace is not the answer. A brain that plugs into what already exists is.

A machine brain that learns without living in the weights.

Models can perceive. Peel stores operational knowledge explicitly: state, failures, rules, evidence and provenance. The Floor decides what is actually authorised.



Weightless learning

Operational learning does not require neural-weight updates. New lessons are kept as explicit rules, conditions and evidence.

Fail-first evolution

Failures become experiments. Peel proposes bounded changes, tests them against a control, rejects what loses and promotes only what clears the gate.

One persistent brain

Bodies, sensors and models can change. The operational knowledge persists, with adapters translating each platform into the same decision grammar.

Failure becomes evidence. Evidence changes the next run.

Not prompt engineering. Not hidden fine-tuning. A replayable progression from failure → candidate rule → experiment → promotion or rejection.



Proof, not vibes.

Measured runs. Wins and losses included. The red column stays in on purpose.


TEST	WITH THE BRAIN	WHAT WENT AGAINST US
Hostile driving controller	64,952 / 64,952 bad commands overridden; 0 collisions	a boundary, not learning
Drone course	10 / 10 clean laps vs 0 / 10 fresh brain	crashes at 2–3 m/s chase
Tank brain exported to a new game	+17%; +66% after pruning	a second export only broke even
BattleZone	27,000 vs random 3,000	a dodge rule rejected
DOOM E1M1	exit at 52 s	Nightmare not cleared
3D tank arena	−1.17 → −0.26 net kills/min, hand-written to evolved	still loses to the game's own AI

What this means

- The architecture already produces measurable effects in games and simulation.
- The same runs show the boundaries: transfer is imperfect, some environments stay hard, some pilots still lose.
- Signal and honesty together: what technical diligence looks for.

Would you let a chatbot drive your tank?

General models are not useless. They are the wrong authority layer for high-consequence actuation.



What this shows

- Peel works from explicit state, mission authority and bounded adaptation.
- A general model answers to a vendor's policy and can confidently improvise the wrong thing.
- For critical systems, proposal and execution should not be the same layer.

8 / 8

tank situations, Peel rules

6 / 8

Claude, written exam

5 / 8

DeepSeek, written exam

One brain. Many endpoints. Even the ones with no API.

Not a rip-and-replace thesis. A plug-in thesis: old software, controllers, buses and systems become endpoints of the same explicit decision architecture.

Old software, no API

Driven through the interface it already has: buttons, fields, screens and receipts. This is the legacy-integration DNA behind Perslis.

Servers and mainframes

Reached through the ways in they already have. An answer is proven or refused in the system it already trusts.

Your failures, your rules

Failure telemetry becomes explicit rules for the next system: inspectable, versioned and owned, not hidden in a rented black box.

Next adapters: military buses, fielded controllers, and installed systems that will never expose a clean modern API.

Aligned to the commander. Not to the model.

The Floor is the independent authority layer: deterministic, machine-readable, replayable.
Learning changes tactics, never authority.

People

Set mission authority, rules of engagement, safety limits and human judgment over force. The command structure decides what is allowed.

The Floor

A deterministic mechanism enforcing machine-readable constraints derived from authorised policy, safety limits, mission rules and applicable ROE. Nothing Peel learns can widen it.

- Every decision leaves a written reason: same facts in, same decision out.
- Built toward appropriate human judgment over the use of force, not free-roaming model improvisation.
- The answer to 'why not just use a bigger model?'

Where it plugs in, and how it pays.

The near-term business is not selling a humanoid. It is licensing the brain, building the adapter, and being paid for evaluation and integration.

Who uses it

- Primes and integrators building uncrewed or semi-autonomous systems.
- Program offices and test & evaluation teams that need replayable decision logic.
- Sustainment programs sitting on old systems with no clean API.

How it pays (plan)

- A brain licence per platform family or deployment type.
- Integration: adapters to machines, controllers and legacy endpoints.
- Evaluation contracts: shadow-mode studies on recorded or simulated missions.
- Entry through SBIR/STTR, OTAs, prime teaming and partner pilots.

Where we are: no revenue and no partner integration yet. The first contracts are shadow-mode evaluations.

24 months. One gate at a time.

Each stage has a written gate. Nothing moves forward on an unmeasured result.

0-6 mo

Real flight controller in the loop

first partner in shadow mode

6-12 mo

Low-cost ground and air robots

first partner integration on real sensors

12-18 mo

Humanoid on the bench

military-bus adapters and expanded HIL

18-24 mo

Sanctioned robot-vs-robot ring

independent test protocol handed over

Beyond 24 months: combat-capable humanoid platforms under human command, a long-term goal, not a present capability.

A small team. A focused ask.

The leverage: a small, high-agency team building a reusable autonomy architecture, not a foundation-model lab.

Founder / architect

Autonomy engineer

Robotics & HIL

Legacy integration

Test & evaluation


Security & compliance

Programs & partnerships

Range framing

US\$5M buys the core team, HIL hardware, low-cost robots and a first partner in shadow mode. US\$10M adds the humanoid track, bus adapters, a second partner and deeper independent validation.

US\$5-10M
for about 24 months



The risks, and how each one is retired.

Naming the failure modes up front, with the path that retires each one.

RISK	RETIREMENT PATH
Transfer is not automatic	Every export gets its own gate. A lesson transfers as a candidate, not as a truth. Nothing advances on an unmeasured result.
Danger and the goal share one road	Humans set the limits up front. The brain learns tactics, not authority. The Floor holds the line.
The eye's margin is thin	Abstain instead of guessing. Add better sensors and richer perception only where the data demands it.
Nothing fielded yet	Shadow mode first, then HIL, then physical integration, then independent testing. Nothing critical is touched until the gate says so.
Policy and trust	Appropriate human judgment over force, deterministic constraints, and a replayable trace for every decision.

Why a prime should be in this round.

Not another agent wrapper. An inspectable autonomy architecture that lives inside the platforms, controllers and robots already moving toward the field.

- One brain architecture across air, ground and legacy endpoints, with only the adapter changing at the boundary.
- Operational knowledge that stays explicit: state, failures, rules, evidence and provenance, not weights you rent.
- A replayable decision record for test and evaluation, plus an independent Floor your program office can read.
- A small team, a gated 24-month plan, and a habit of publishing the losses alongside the wins.

Perslis Defense
Adaptive autonomy. Fixed authority.

The bet



Architecture	the reusable machine brain
Validation	the measured, gated path
Integration	the adapter into real systems