

Three products, one instrument, one ecosystem

Each product answers a different question, sells to a different buyer, and could stand alone. What makes them one company is that they all read the same recording. Ethos captures behaviour and internals together. Panacea reads those internals for concepts. Epoch measures how much of the recording was information at all. Each measurement tells the next one what to do.

What each measures

Product	Measures	Emits	Buyer
Ethos	Behaviour under perturbation, with internals captured on the same run, compared across substrates	Drift detection, attestation	Labs running biological or agentic workflows
Panacea	Internal representations — structure the output cannot surface	Hidden concepts, located	Anyone with a black box and no interpretation of it
Epoch	How much a corpus still has left to teach a model	Allocation decision: what to buy next	Anyone pretraining or acquiring data under a limit

Behaviour, representation, information — **three readings of one recording.**

How they hand off

Ethos captures Runs a controlled perturbation and records both the behavioural and the internal response from the same run. Produces the capture everything downstream reads.	Panacea interprets Probes those internals for structure the behavioural readout missed — concepts present in the representation but absent from the output. Produces candidate concepts and where they sit.	Epoch meters Measures how much of the capture was new information, and how much more is needed before another probe is worth fitting at all. Produces the size and shape of the next capture.
Ethos → Panacea: the internals, and the perturbation that caused them	Panacea → Epoch: the target worth metering against	Epoch → Ethos: how many runs, subjects or models the next round needs

The loop closes: Epoch's answer determines what Ethos captures next, which is why one compute budget funds three product lines.

The shared instrument

A common measurement procedure, applied to every substrate, that returns readings which can be compared from one substrate to the next. The procedure was built for behavioural drift in Ethos. It turns out to supply what Epoch needs to meter information content, and what Panacea needs to tell a concept-bearing signal from a redundant one.

More than a convenience

One procedure, three products. Nothing had to be rewritten to move between them, which is the actual reason a small team can hold three product lines at once, and the reason a result in one substrate transfers to the others instead of staying local.

Methods: available under NDA.

Experiment map

Experiment	What varies	Serves	Uniquely gives	Status
LLM arm	Controlled workflow conditions	Ethos, Panacea, Epoch	The only substrate with a language channel	Running
Evo 2 regulatory probe	Single-base variant	Panacea, Ethos	Panacea's proving ground, with labels that were measured rather than predicted	Pre-registered
Brain imaging	Task and recording conditions	Epoch, Panacea, Ethos	Human cortex — the black box interpretability tooling was never built for	Next
Genome corpora	Corpus composition	Epoch	A finite data source, so the limit is real	Next
Model internals	Capture scale and model count	Epoch, Panacea	A system we control, so the method can be checked against a known answer	Next

Each row is a different substrate read by the same instrument.

Panacea's proving ground

The Evo 2 regulatory probe is the cleanest test of the white-box premise. On identical labels, in a domain where existing tools are thinnest, we compare what can be read from a model's internals against what can be read from its output.

The claim under test is that **reading the representation beats reading the output**. If it does, the representation is holding regulatory signal the output cannot surface — which is the entire proposition Panacea sells, stated in a form that can fail.

What it returns, and where it stops

The output is a set of DNA positions where the model predicts an effect on brain gene activity that nothing else explains. It locates signal; it does not name mechanism. Naming mechanism needs a regulatory genomics collaborator, and we say so rather than implying we close that gap ourselves.

Two anchors

Does the instrument see signal?

Answered by the Evo 2 arm, where the answer was measured rather than judged. If the instrument separates those cases it is detecting something real — and that verdict carries to every other substrate, because the procedure is the same.

Does the metering return the right number?

Answered by the model internals arm, where we control the system and can check the method against a known answer. That is why this work validates the genome and imaging tracks rather than competing with them.

One arm proves the instrument works. One arm proves the ruler is right. **Everything else is application.**

Brain imaging serves all three

Panacea does not care what produced the activity it reads. Pointed at a model's internals it surfaces concepts the model never states. Pointed at human cortical recordings it does the same job on a system that has never had interpretability tooling aimed at it, and where — unlike a model — nobody can inspect the weights.

That single relationship serves all three products. **Epoch** meters how much of the recording is information and how much data a probe needs before it is worth fitting at all — the difference between a probe that means something and one that fits noise. **Panacea** reads what the information contains. **Ethos** takes anything mechanistic that comes out and tests whether the same mechanism appears in the other substrates.

The honest asymmetry

One comparison available in the language substrate has no counterpart in the others, because those substrates have no channel through which the equivalent condition could be set. Any diagram of this programme should show that comparison stopping at the LLM arm rather than implying clean symmetry across every substrate. *Detail: NDA sheet.*

How each one sells

Product	Shape	Entry	Recurring
Ethos	Eval suite built against the client's workflow, then run continuously	\$25,000 build	\$6,000/mo
Epoch	Fixed-fee measurement per corpus, priced by substrate; licence for in-pipeline use	\$18,000–45,000	\$3,000–6,000/mo
Panacea — tooling	Interpretability tooling for neuroscience: read internal activity, surface concepts, no interpretability team required. Single lab up to five seats, or department-wide.	—	\$2,500–6,000/mo
Panacea — embedded	We run the probes alongside your group — an extension of your research team rather than a tool they have to learn. Three-month minimum.	—	\$12,000/mo

Epoch's range: \$18,000 for model internals, \$30,000 for brain imaging, \$45,000 for genome corpora. Embedded licences from \$60,000 a year. Panacea is priced for the neuroscience use case, which is where the first experiment already sits; the AI-model side stays unpriced until the Evo 2 probe clears.

Pricing principles

Fixed fee, never hourly.

Hourly caps the upside on work whose value has nothing to do with how long it takes, and turns every conversation into a scope argument. The market is moving the same way: deliverable-based pricing is displacing the billable hour precisely because the hours have collapsed while the value has not.

The first read is free, always.

Giving it away costs us almost nothing, proves the value before anyone signs, and is the entire top of the funnel.

Design partners pay nothing.

Two slots per substrate. A refundable commitment holds the slot and filters out tyre-kickers, but the work is free and the exchange is written down: named reference, usable quote, publication rights, and a pre-agreed price if they convert. The method is unproven on real client data. Charging half for it would be charging for our own risk.

Gainshare only where the saving is auditable.

Avoided compute can be documented against a training plan agreed in advance, so we will take a reduced fee plus a capped share of it. Everything else — bench time, recruitment, researcher hours — cannot be measured cleanly, so we do not price against it.

Where these numbers come from

Every figure above is anchored to something checkable, because the first question a serious buyer asks about a price is how it was arrived at.

Anchor	Figure	Source
GPU cost	Median H100 on-demand about \$3.40 per GPU-hour, September 2026; cheapest verified in-stock about \$1.55	Public multi-provider price trackers
A mid-scale genomic run	7B parameters on 1T tokens is roughly 70,000–75,000 H100-hours, about \$250,000 at the median rate	Our arithmetic from the standard compute estimate, shown so you can check it
A frontier genomic run	Evo 2 used over 2,000 H100s for several months — millions of dollars of compute at any rate you like	Arc Institute and NVIDIA, publicly stated
Neuroimaging cohort cost	\$16,494 per randomised participant fully loaded, \$8,247 of it imaging alone; academic scanner cores \$500–650 an hour	Published trial cost analysis and university core rate cards
Consulting comparables	Scoped boutique AI pilot \$10,000–50,000; strategy sprint \$25,000–75,000; boutique hourly \$200–500; advisory retainers \$15,000–20,000 a month	2026 published rate surveys
Interpretability tooling	No public list price exists — the category sells by negotiation	Which is why Panacea is anchored to research-software budgets instead

Where a figure is our own estimate rather than a published number, it says so, and the arithmetic is available.