

Signature matching is a report

on an attack that already happened.

Multi-stage attacks orchestrated by AI don't present a signature until late. Each stage looks individually unremarkable; the pattern only resolves once the payload executes. Detection built on matching known indicators is structurally behind.

A world model learns the embedding-space trajectory of normal network and endpoint behaviour – the same self-supervised approach behind the track's embodied-agent research. An attack surfaces as divergence from a predicted trajectory, not as a match against a library.

Model system state, not just signatures

The model learns what normal looks like as a continuous trajectory across network and endpoint behaviour.

Flag deviation pre-exploit

Attacks are detected as predicted-trajectory divergence, often before payload execution rather than after a signature match.

Licensable IP for two buyer groups

Positioned as defence IP for cybersecurity vendors to embed, and for banks to deploy directly over critical infrastructure.

Predict before breach

A shift from reactive signature-matching to trajectory-level anomaly forecasting.

Illustrative target application for the JEPA track's Year-1 industry-contract layer – not yet a shipped result.

A research lab that owns what it builds.

IIDA is a non-profit research lab working on applied AI problems that are commercially real in three to five years – too long for a startup, too applied for a university. A Focused Research Organization takes on bottlenecks that are too engineering-heavy for academia and not immediately profitable enough for venture capital.

Time-bound

3-year horizons with pre-committed exit conditions

Self-sustaining

Target of $\geq 30\%$ revenue from industry by Year 2

Owns assets

IP and equipment sit with IIDA, not brokered

Kill criteria, day one

The track narrows or closes if replication doesn't reach published baselines within two quarters

WHAT A SECURITY PARTNER GETS

01 A named research pod (6–10 contributors) on trajectory-level anomaly forecasting

02 Defence IP you can embed in your own product rather than license per-seat from a model vendor

03 Open architecture and position papers from the underlying track – the method is auditable, which matters when the buyer is a regulated institution

04 A shared research base with the BFSI track, so critical-infrastructure deployment and vendor embedding are developed against the same model

ASK A design partner with production telemetry to work against, and a research lead for the track. The lead position is currently open.

What happens after you say yes.

01

WEEKS 1-4 · CONTACT & SCOPE

An NDA and an exploratory conversation. The pod maps the anomaly-forecasting track against your telemetry and environment, and names the problem statement.

02

MONTHS 2-5 · PILOT AGAINST THE KILL CRITERION

The named pod (6-10 contributors) runs the pilot against the criterion agreed at the start – reaching published replication baselines within two quarters.

03

MONTH 6 · FORMALISE OR CLOSE

The engagement becomes a licensing and collaboration agreement, or – if the criterion isn't met – it closes on schedule rather than being extended.

INTELLECTUAL PROPERTY – THE QUESTION YOU WILL ASK FIRST

Ownership follows the funding source, and is fixed in writing at project start rather than argued about at the end. An Intellectual Property Council – drawn from research, legal, industry and finance, plus a nominee of the governing board – approves filings and determines ownership. Funders hold no editorial control over research output.