



PRE-SEED INVESTMENT PRESENTATION

TauGuard Limited

Governance infrastructure for increasingly capable AI

TauDIL GOS™ · ARCHON GOS™ · TauGuard Lab · Continuity Science


€1.2M Pre-Seed Round

AI is moving from recommending actions to taking them.

TauGuard governs what AI is allowed to make happen — deterministically and structurally, before the action occurs.

Capability is advancing faster than authority. TauGuard builds the layer that keeps the two aligned — not a better model, but the infrastructure around one.

From recommendation to action



- Copilots have become agents; agents are gaining tool access, payment rails, and system permissions.
- Each step forward removes a human checkpoint that used to double as a safety control.
- The industry has optimized for what AI can do. Almost nothing governs what it is allowed to do.0

The governance gap

Capability curve

- AI systems now execute code, move funds, and modify records with little or no human in the loop.
- New tool integrations ship weekly; permission scopes widen by default.
- Vendors compete on what their agents can reach, not on what they refuse to reach.

Oversight curve

- Most safeguards remain advisory — prompts, guidelines, and human review applied after output.
- Controls describe intended behaviour; they do not enforce it structurally.
- When something goes wrong, teams explain it after the fact instead of preventing it before.

Post-hoc explanation is not governance.

Capability $\neq$ authority

CAPABILITY

What a system is technically able to do

VS

AUTHORITY

What a system is permitted to make happen

- Most AI safety today conflates the two: if a model can call a tool, it is treated as authorized to call it.
- TauGuard separates them structurally — authority is granted, scoped, and enforced independently of capability.
- An action outside the authority boundary is not flagged after execution. It is unreachable before execution.



The commercial Governance Operating System

1

Deterministic

Rules are executable, not advisory

2

Structural

Invalid actions are unreachable,
not just flagged


3

Deployable

Ships today as an API, SDK, or on-
prem layer

TauDIL sits between AI systems and the actions they attempt to take — approving, scoping, or structurally refusing each transition before it executes.

How TauDIL works



- Every candidate action is mapped onto a governed topology of permitted states and transitions.
- The gate evaluates the transition against explicit rules — not a trained model's judgment call.
- Outcomes are binary and logged: authorized, or structurally refused. Nothing executes in between.


Deterministic governance

- Governance rules are executable code, not guidance text a model may or may not follow.
- The same input state and the same rule set always produce the same governance outcome.
- Enforcement happens outside the model's own reasoning — there is no probabilistic interpretation of policy at the decision point.
- Auditability follows directly: every decision traces to a specific rule, not a statistical inference.

Reproducible by design

- **Same state + same rules = same outcome, every time**
- **Decisions log to a specific rule, not a confidence score**
- **No prompt-level negotiation with the governance layer**

Structural refusal



- Traditional AI safety relies on the system recognizing a bad action and declining it in the moment.
- TauDIL GOS removes that reliance: invalid transitions are not represented as reachable states in the topology at all.
- There is no prompt to defeat, because there is no path to defeat.

Governed Cognitive Topology — inside TauDIL

- The Governed Cognitive Topology is deployed as a native component of TauDIL GOS™.
- It is not a dependency on ARCHON or any external reasoning layer — TauDIL governs on its own stack.
- The topology encodes permitted states, transitions, and the conditions under which each transition is authorized.
- This keeps TauDIL deployable as a standalone governance layer in front of any AI system a customer already runs.

Clean architectural line

- TauDIL: governed action layer, self-contained
- ARCHON: separate architecture, separate stack
- No shared runtime dependency between the two

Semantic drift

Semantic drift: the gradual divergence between what a system's outputs mean and what its governing rules assume they mean.

- Left ungoverned, drift quietly erodes the reliability of every downstream control built on top of it.
- TauDIL GOS treats semantic drift as a Governed Cognitive Topology problem: meaning is checked against the topology at each transition, not assumed.
- This is distinct from constitutional identity drift, addressed separately by Continuity Science's identity mathematics (slide 24) — a different failure mode with a different fix.

Real-world use case: agentic payments

1

A finance team's AI assistant is asked to reconcile and pay a batch of vendor invoices.

2

The assistant proposes a batch of payment actions across multiple vendors and amounts.

3

TauDIL evaluates each payment transition against scoped authority — vendor allow-list, per-payment limits, dual-control thresholds.

4

In-scope payments execute automatically; out-of-scope ones are structurally refused before submission — not flagged after the fact.

Outcome

- **No manual review queue for in-scope payments**
- **Zero unauthorized transfers by construction**
- **Full audit trail tied to a specific rule per decision**

FinTech beachhead

- Highest concentration of agentic-AI pilots that touch money movement, KYC, and account actions.
- Regulatory direction — PSD3, DORA, and emerging EU AI Act obligations — is pushing toward provable, pre-execution controls.
- Existing compliance stacks are rules engines bolted onto legacy rails; none were built for AI-initiated actions.
- A well-governed deployment produces an auditable reference case that generalizes to healthcare, insurance, and critical infrastructure.

Why fintech first

- **Money movement = clearest, highest-stakes authority boundary**
- **Regulators already demand auditable controls**
- **Beachhead, not ceiling — the pattern generalizes**

Product & deployment model

1

API Gateway

A drop-in governance layer in front of an AI system's tool calls. Fastest path to production.

2

Embedded SDK

Governance compiled directly into the customer's application, for tighter latency and control.

3

On-prem & VPC

For regulated environments where data must never leave the customer's perimeter.

Market

TAM

AI governance & agentic-
AI risk management —
estimated multi-billion €
market by the
early 2030s

SAM

Regulated-sector agentic
deployments in Europe:
finance, healthcare,
infrastructure

SOM

Initial addressable
accounts: fintech
agentic-payments and
compliance
workflows

*Figures are directional
estimates for an early,
fast-moving category
— refined as design
partners are
onboarded.*

- Every category adopting agentic AI eventually needs a governance layer — this is infrastructure, not a point solution.
- Regulation is converging on the same requirement across jurisdictions: provable, pre-execution control.

Why now

Agentic adoption

Agentic AI is compounding across every SaaS category, widening system access by default.


Regulatory pressure

The EU AI Act and DORA are moving toward requiring provable, pre-execution controls.

Visible incidents

Early agentic-AI incidents are making post-hoc governance visibly insufficient to buyers.

Competitive landscape



Moat

1

Mathematical foundation

Built on Continuity Science's identity mathematics — not a prompt technique competitors can copy.

2

Patent-pending architecture

The structural-refusal mechanism and Governed Cognitive Topology are patent-pending.

3

Compounding deployment data

Each governed deployment sharpens the topology for the next one.

4

Category definition

First mover defining “deterministic, pre-execution AI governance” as a category.

TauGuard Lab

Experimental and research environment

- Independent from the commercial TauDIL roadmap — a space to test mechanisms before they are hardened into product.
- Explores computation-energy governance, simulator-based research, and next-generation governance mechanisms.
- Findings graduate into TauDIL, ARCHON, or Continuity Science IP only once validated — the Lab does not ship directly to customers.

The computation-energy opportunity

- Ungoverned AI systems frequently compute paths they will never be authorized to act on — wasted inference, wasted energy.
- TauGuard Lab researches whether governance constraints, applied earlier in the reasoning process, reduce unnecessary computation.
- Early hypothesis: pruning inadmissible transitions before full reasoning completes lowers compute cost per governed decision.

Status

- **Lab-stage research, not a shipped TauDIL feature**
- **Feeds the Governance Simulator work (next slide)**
- **Long-run upside: lower cost-to-serve for governed AI**

Governance Simulator — early research

Priority Weighting

Ranks candidate actions by governed priority before full reasoning runs.

Attention Economy

Models compute as a scarce, budgeted resource allocated under governance.

Coherence Governance

Tests whether a reasoning path stays coherent with the governed topology.

Risk / Governance

Simulates risk-weighted outcomes of alternative governance rule sets.

Tau-Time Urgency

Explores how time-criticality should legitimately alter governed priority.

Results are from TauGuard Lab's simulator environment and inform future research — they are not current TauDIL GOS product claims.

ARCHON GOS™ — a separate product

- ARCHON is a distinct governed-intelligence architecture: separate product, separate stack, separate interface, separate solutions from TauDIL.
- Where TauDIL governs actions an AI system attempts to take, ARCHON addresses a different class of governed-intelligence problems.
- The two products share no runtime dependency — TauDIL does not require ARCHON to function, and vice versa.
- ARCHON is positioned for a later go-to-market phase, built on shared research from Continuity Science.

Continuity Science

Independent research and IP programme

- Develops the mathematical and architectural foundations underlying both TauDIL and ARCHON — identity mathematics, constitutional order, and continuity theory.
- Functions as the company's long-horizon research engine, feeding validated mechanisms into product roadmaps rather than shipping directly to customers.
- Owns the IP that underwrites TauGuard's long-term moat, independent of any single product's commercial timeline.

Constitutional identity & identity drift

Constitutional identity drift: divergence between a system's declared identity and invariants and its actual behaviour over time.

- Governed by Continuity Science's identity mathematics — a distinct discipline from TauDIL's Governed Cognitive Topology.
- Establishes formal invariants a system must preserve across updates, memory changes, and extended operation.
- Provides the theoretical basis eventually productized through ARCHON and future Continuity Science licensing.


Safe degradation & metabolism research

- Explores how a governed system should degrade under stress, resource limits, or partial failure — without violating its invariants.
- “Metabolism” research treats resource allocation and degradation as a governed process, not an uncontrolled fallback.
- Early-stage research, feeding directly into the Purpose → Continuity hierarchy on the vision slide.

Why it matters

- Failure without invariant violation is a governance property
- Distinct from uptime / reliability engineering
- Long-horizon research, not a near-term product claim

The research → product flywheel



Product deployments raise new governance questions, which return to Continuity Science as the next research cycle — closing the loop.


- This is what makes the four-entity structure a compounding asset rather than four separate teams.
- Every commercial deployment strengthens the research foundation it was built on.

IP strategy

Asset	Status
Structural refusal mechanism (TauDIL)	Patentable
Governed Cognitive Topology	Patentable
Identity mathematics (Continuity Science)	Proprietary architecture; filings in preparation
Governance Simulator methodology	Proprietary — trade secret

Status language is used precisely: “patent-pending,” “subject to patent protection,” and “proprietary architecture” are each applied only where that status is actually established — not as a blanket claim across every asset.

Company architecture



- Each entity has a clear mandate and a clear boundary — no product depends on another to function.
- Continuity Science research feeds all three operating arms without dictating any one product's roadmap.

Business model

Platform subscription

API / SDK access with usage-based pricing on governed decision calls.

Enterprise & on-prem

Annual licensing for regulated deployments requiring VPC or on-prem control.

Future: ARCHON & IP licensing

Post-raise roadmap —
ARCHON commercial introduction and Continuity Science licensing.

Unit economics

€40K

Average contract value
(blended, illustrative)

~80%

Gross margin at scale
(software infra)

9 months

CAC payback period, target


120%

Net revenue retention,
target

- Pricing combines a platform licence fee with usage-based pricing on governed decisions — revenue scales with customer volume, not just seat count.
- Illustrative planning assumptions based on comparable infrastructure and compliance-software benchmarks — to be validated against design-partner contracts.
- Blended ACV assumes a mix of mid-market fintech (€15–25K/yr) and early enterprise/on-prem contracts (€80K+/yr).

Financial projections


ARR trajectory (€000s, illustrative)



Reading the curve


- Year 1: design partners live, €150K ARR
- Year 2: first enterprise contracts, €900K ARR
- Year 3: gross margin ~78–80%, operating breakeven approached
- Year 4: multi-sector expansion, ARCHON early revenue contributes
- Figures are planning assumptions, not guarantees — refined with design-partner data

Go-to-market



- Land with design partners who need auditable, pre-execution control now — not a future roadmap promise.
- Convert design partners into referenceable case studies before broadening the sector mix.

€1.2M pre-seed — use of funds



Return potential for this round

Pre-seed (now)

€1.2M at an illustrative €5M post-money valuation — terms indicative, subject to negotiation.

Seed (12–18 months)

Target €4–6M raise at an illustrative €18–22M valuation — a ~3.5–4x markup on this round.


Series A (Year 3)

ARR-based valuation at ~8–10x ARR on a projected €3.2M ARR — illustrative €25–32M range.

Illustrative exit scenario: strategic acquisition by a compliance, fintech-infra, or platform acquirer in Year 5–7, at 8–12x ARR on an estimated €15–20M ARR — an indicative 25–40x range for pre-seed investors before dilution. This is a planning scenario, not a projection or guarantee.

All figures on this slide are illustrative planning assumptions for discussion, not forecasts, guarantees, or investment advice.

Milestones after the raise



- Each quarter is designed to de-risk the next — commercial proof before commercial scale.
- Continuity Science and TauGuard Lab milestones run in parallel, feeding the flywheel on slide 26.

Long-term vision

*Purpose → Identity → Invariants → Governance → Authority → Attention → Knowledge → Meaning
→ Coherence → Admissibility → Learning → Memory → Metabolism → Continuity → Time*

Continuity isn't the starting product premise — it emerges from preserving purpose through identity, invariants, and governance.

TauGuard is building the infrastructure that keeps that chain intact as AI systems keep gaining capability — today through TauDIL, and over time through ARCHON and Continuity Science.

CLOSE

TauGuard isn't building a better model.

It's building the infrastructure that governs what AI is allowed to make happen.

€1.2M pre-seed

**TauDIL GOS™
commercial traction**

**Continuity Science IP
foundation**

Kamilla Harcej · kamilla@tauguard.xyz · tauguard.xyz