RATIUM.AI / STRATEGIC DEVELOPMENT

LoopGuard-AI Partner Brief

Non-Confidential Discussion Version

Turn AI evaluation into governed action.

Document field	Current position
Prepared by	Benny Dunavich / RATIUM.AI
Version and date	V1.0 / 28 July 2026
Purpose	Partner qualification and bounded development discussion
Current stage	Concept and reference-architecture stage
Prepared for	Qualified technology, investment, validation, and pilot counterparties
Contact	contact@ratium.ai

> **Decision requested:** Determine whether to fund a bounded 90-day discovery and reference-build phase that can produce evidence for a rational decision to continue, restrict, redesign, or terminate development.

This brief does not present a completed product, a securities offer, or a request for passive sponsorship. It presents a disciplined first transaction: funded founder participation combined with engineering, falsification, and commercial definition.

ratium.ai/partner | contact@ratium.ai

DECISION BRIEF

1. Executive proposition

AI systems can produce scores, logs, benchmark results, policy classifications, red-team findings, and observability traces. These outputs become governance only when they can change what happens next. LoopGuard-AI is a proposed AI governance and evaluation control layer intended to connect signals, evidence, policy, authority, reversibility, drift, and decision stability to explicit operational gates: SHIP, RESTRICT, HOLD, and ROLLBACK.

RATIUM.AI is not asking a counterparty to accept an unvalidated architecture as a finished product. The proposed first commitment is a paid, bounded program that translates the public reference architecture into formal requirements, a minimal executable gate flow, a falsification plan, and a decision on whether deeper investment is justified.

> **Commercial object:** The first commercial object is not a license to a proven product. It is a funded decision process designed to determine whether a credible product path exists.

The decision a serious partner should make

- **Problem fit:** identify one bounded AI event, release process, agent action, or workflow transition where evaluation currently lacks a decision-bearing consequence.
- **Execution fit:** confirm access to engineering, security, evaluation, product, and responsible human authority.
- **Funding fit:** define a founder-transition retainer, technical discovery budget, and any independent evaluation or pilot budget.
- **Evidence fit:** agree in advance what evidence would justify continuation, restriction, redesign, or termination.

Why the founder matters to execution

The public corpus allows independent inspection. An official relationship adds interpretive continuity across CEP, LoopGuard-AI, and the wider RATIUM.AI corpus; direct clarification of architectural dependencies; coordinated development of future implementation assets; disciplined claim control; and a canonical relationship to the source.

Founder availability is therefore an execution input. A viable first engagement must fund enough defined professional participation to reduce the constraint imposed by unrelated full-time industrial employment. This is not a charitable add-on; it is a resource requirement.

ARCHITECTURE

2. What LoopGuard-AI proposes

LoopGuard-AI is designed to sit above, around, or alongside AI applications, language models, agents, and AI-enabled workflows. It treats an AI output, agent action, tool call, release candidate, evaluator conflict, drift signal, or human override request as a governance event that may require an operational consequence.

> **Reference pipeline:** AI event -> ingestion -> signals and metrics -> policy and evidence -> stability assessment -> gate decision -> audit, escalation, monitoring, and feedback.

Gate	Operational meaning	Illustrative consequence
SHIP	Proceed under the current governance conditions.	Allow the output, action, release, or workflow transition.
RESTRICT	Proceed only under defined controls.	Reduce authority, limit scope, require approval, or increase monitoring.
HOLD	Pause until a blocking condition is resolved.	Request evidence, clarify authority, resolve policy conflict, or obtain human review.
ROLLBACK	Reverse or downgrade a prior state.	Disable access, revert a configuration, restore a prior version, or withdraw a release.

Category boundary

LoopGuard-AI is not a foundation model, chatbot, prompt wrapper, content filter, observability dashboard, legal compliance guarantee, or replacement for human governance in high-impact domains. Its proposed category is an AI governance and evaluation control layer for LLMs and agentic AI systems.

Relation to CEP

The Central Equilibrium Problem (CEP) supplies a theoretical framework for examining systems that stabilize around inferior or difficult-to-correct decision regimes. LoopGuard-AI attempts to translate that concern into decision-stability assessment, gate logic, and audit-ready governance records. This is a conceptual and architectural bridge, not a claim that CEP has already been empirically validated inside production AI systems.

EVIDENCE CONTROL

3. Current stage and maturity boundary

The current public work defines the problem, architecture, gate semantics, candidate component topology, Policy Pack logic, metric interfaces, evidence and decision packages, replay, audit, rollback, and staged validation direction. Architecture-level coherence is not implementation proof.

Maturity level	Current status	Evidence position
Concept	Defined	Problem model, governance purpose, and gate logic are articulated.
Architecture	Defined for review	Components, flows, interfaces, and control logic are publicly specified.
Prototype	Not yet established	Requires funded executable implementation and tests.
Evaluation	Not yet established	Requires fixed cases, error measures, baselines, and reproducible results.
Validation	Not yet established	Requires a relevant controlled or quasi-operational environment.
Production	Not claimed	Requires deployment, monitoring, incident handling, and operating evidence.
Certification	Not claimed	Requires assessment by an external formal body for a defined scope.

Available for review now

- Canonical conceptual corpus and public reference architecture.
- Evaluation-to-gate logic for SHIP, RESTRICT, HOLD, and ROLLBACK.
- Candidate schemas, metric interfaces, policy logic, evidence structures, replay, and audit concepts.
- Governance simulation procedure and staged falsification orientation.

Still requires funded execution

- Formal engineering requirements, implementation boundary, and threat model.
- Reference implementation, APIs, tests, deployment model, storage, and security controls.
- Datasets, calibration, baseline comparison, error measurement, and reproducibility.
- Controlled pilot, compliance mapping, product design, support, distribution, and independent review.

> **Claim-control rule:** No claim about LoopGuard-AI should exceed the evidence available for that claim.

BOUNDED ENTRY PROGRAM

4. Proposed 90-day discovery and reference-build phase

The recommended entry phase is deliberately narrow. It should reduce uncertainty before either side commits to a broad product promise, exclusive relationship, or large capitalization decision.

Timing	Workstream	Candidate output
Days 1-15	Architecture translation	Requirements, interfaces, asset inventory, implementation boundary, threat-model outline, and one selected event class.
Days 16-40	Minimal gate flow	Executable event-to-gate path using a canonical schema, selected metrics, a versioned policy profile, decision trace, and evidence record.
Days 41-65	Falsification and evaluation	Fixed scenarios, simpler baselines, repeatability tests, error measures, adversarial cases, and explicit stop criteria.
Days 66-90	Pilot and commercial definition	Pilot scope, integration path, operating responsibilities, budget, rights issues, and an evidence-bounded next-stage decision.

Program controls

- One bounded use case or event class; no attempt to govern every AI action in the first build.
- A primarily deterministic, rule-based gate engine at the reference stage.
- Versioned inputs, metric definitions, policy packs, decisions, evidence, and overrides.
- Comparison with a simpler baseline, not only internal demonstrations.
- Named error measures and known hard-blocker cases, including false-SHIP prevention.
- Human decision authority, escalation, and override trace defined before testing.
- A negative result treated as useful when it prevents larger misallocation.

> **End-of-phase decision:** Proceed to a controlled pilot, continue under restrictions, redesign the architecture or use case, or stop the program.

REFERENCE BUILD

5. First-build target and acceptance criteria

The first engineering milestone should prioritize a deterministic pre-release or batch decision system. Full real-time enforcement for every agent action is a later-stage possibility, not the first build target.

A credible first version should be able to

- **Ingestion:** Open a run and ingest artifacts or signals.
- **Schema:** Validate input against a canonical schema.
- **Metrics:** Compute selected versioned metrics and expose failure or incompleteness behavior.
- **Policy:** Apply a versioned Policy Pack with explicit precedence and hard blockers.
- **Gate:** Return SHIP, RESTRICT, HOLD, or ROLLBACK deterministically under defined conditions.
- **Evidence:** Provide a decision trace and generate an evidence bundle.
- **Authority:** Support documented human override with actor, authority, rationale, and timestamp.
- **Integration:** Integrate with one simple evaluation, CI/CD, or workflow boundary.
- **Replay:** Reproduce the same decision under the same inputs, metric versions, and policy version.
- **Control:** Prevent false SHIP in known hard-blocker cases and preserve audit records.

Restriction or stop conditions

- The selected metrics cannot be formalized, versioned, explained, or reproduced.
- The gate logic is materially less reliable than a simpler baseline.
- The system cannot preserve a reviewable decision and evidence trail.
- Known hard blockers still produce unacceptable false-SHIP outcomes.
- The selected use case lacks responsible human authority or a credible pilot environment.
- The value of the architecture cannot justify its operational complexity.

EXECUTION CAPACITY

6. Partner contribution and founder transition

The opportunity depends on complementary inputs. RATIUM.AI contributes the canonical source and continuing architectural interpretation; the partner must contribute the organizational capacities required to convert that source into testable work.

Contributor	Required contribution
RATIUM.AI / Benny Dunavich	Canonical corpus; CEP-to-LoopGuard dependency continuity; architectural clarification; requirements translation; continuing conceptual development; claim control; provenance; participation in falsification and design review.
Technology or development partner	Software architecture; backend; data; evaluation; security; product engineering; testing; integration; deployment; operations; and technical delivery.
Investor or sponsor	Founder-transition funding; defined discovery budget; governance for capital allocation; access to an operating team; and a credible path to follow-on execution.
Pilot or validation organization	Bounded use case; representative cases; responsible decision authority; review access; domain expertise; and controlled validation conditions.

Minimum funding architecture

- **Founder-transition retainer:** funds defined professional availability and reduces dependence on unrelated full-time employment.
- **Technical discovery and build budget:** funds the engineers, infrastructure, security, evaluation, product work, and program management required for the 90-day scope.
- **Independent evaluation or pilot budget:** funds external review, representative cases, measurement, and any controlled operating environment.

This V1.0 brief intentionally does not invent fees. Before signature of a statement of work, the parties must define the founder's committed availability, funding floor, payment schedule, expenses, tax and contracting route, and the partner team's committed resources. Payment instructions are exchanged only after normal counterparty verification and signed engagement documents.

> **Commercial principle:** Founder funding is part of the execution budget. It should not be framed as a donation, prize, or retrospective expression of gratitude.

COMMERCIAL STRUCTURE

7. Relationship paths and rights boundary

Path	Use when	First commitment
Paid discovery or sponsored R&D	The architecture must be translated, challenged, and tested before a larger decision.	Defined scope, team, founder availability, budget, outputs, evidence, and stop criteria.
Strategic development partnership	A capable organization wants to build, validate, integrate, and commercialize the system.	Staged development agreement followed by negotiated commercial rights after evidence and fit.
Dedicated venture	Capital and operating leadership exist for a focused company or joint vehicle.	Diligence, capitalization plan, founder and team roles, development plan, and governance.

Non-grant boundary

Receipt of this brief or access to public RATIUM.AI materials does not by itself create an official affiliation, endorsement, partnership, license, assignment, exclusivity, brand authorization, confidentiality obligation, commitment to deliver, or right to non-public materials.

Items that require a separate written agreement

- Defined deliverables, acceptance procedure, payment, expenses, and termination.
- Before any non-public disclosure: confidentiality scope, exclusions, duration, permitted use, access controls, and return or destruction.
- Background works and rights; independently created materials; and treatment of public material.
- Ownership or licensing of source code, schemas, datasets, metric definitions and formulas, thresholds, calibration materials, scenario libraries, Policy Packs, documentation, models, prompts, and other implementation assets.
- Foreground work, improvements, derivative implementations, attribution, publication, and portfolio use.
- Exclusivity, field, territory, customer class, sublicensing, assignment, change of control, and post-termination rights.
- Brand and name use, official-source claims, validation claims, and external communications.

> **Legal boundary:** This brief records a proposed negotiation structure; it is not legal advice and does not replace lawyer-reviewed engagement, confidentiality, investment, license, equity, or development agreements.

DECISION PATH

8. Qualification, diligence, and next step

A useful first inquiry should identify

- The organization, sender, decision-making role, and relevant operating mandate.
- Relevant engineering, investment, validation, product, or market capability.
- The first proposed use case, AI event, or claim to test.
- The team, budget, pilot access, and other resources available during the first 90 days.
- The requested next step and expected internal decision process.

Recommended progression

1. **Qualified inquiry.** The sender supplies enough information to assess whether execution capacity may exist.
2. **Fit call.** A 30-45 minute conversation identifies the use case, authority, resources, and likely blockers.
3. **Scoped technical review.** The parties choose one event class and identify the minimum evidence needed for a rational entry decision.
4. **Confidentiality and diligence, if necessary.** Only non-public materials needed for the next decision move into a separately agreed process.
5. **90-day statement of work.** The document defines founder availability, team, scope, budget, deliverables, acceptance criteria, stop conditions, rights, and governance.
6. **Decision.** Proceed, proceed under restriction, redesign, or stop.

Canonical public sources

LoopGuard-AI Technical Source Dossier - canonical architecture-level and engineering-reference description.

LoopGuard-AI Governance Source Dossier - foundational governance and decision-control source structure.

LoopGuard-AI Governance Simulation Protocol - proposed simulation, review, and falsification pathway.

Strategic Development Partnership - public partnership position, claim boundary, and contact route.

Strategic contact

Email: contact@ratium.ai | **LinkedIn:** Benny Dunavich / RATIUM.AI

> **Requested next action:** Reply with the organization, responsible decision-maker, candidate use case, available execution capacity, and whether a bounded 90-day program can be funded.