WORKING PAPER · APRIL 2026

The Last Mile of AI:

Judgment Infrastructure, Defensible Audit Logs, and the End of Information Retrieval

Lok Yek Soon

Founder & CEO, askOdin Pte. Ltd.

yeksoon@askodin.app · askodin.app

askOdin

Judgment is the last unscalable asset.

ABSTRACT

The proliferation of foundational Large Language Models has reduced the marginal cost of information retrieval and synthesis to zero. When every capital allocator operates from the same omniscient information baseline, information asymmetry ceases to generate alpha. This paper argues that the competitive frontier in capital allocation has fundamentally shifted: the premium is no longer on accessing data, but on the rigor with which logic derived from that data is stress-tested. We define this transition as the emergence of AI Judgment Infrastructure™—a dedicated architecture for evaluating the structural soundness of investment theses, rather than merely retrieving and summarizing the information they contain. Drawing on a corpus of 60,000+ Clarity Scores™ benchmarked through the RUNE Protocol (patent-pending), we identify seven empirically derived archetypes of investment thesis failure—the Grammar of Failure—and introduce the Judgment Graph™ as a proprietary data structure mapping relationships between claims, evidence, and historical failure patterns. We further propose the Defensible Audit Log as the canonical output artifact of this infrastructure: an immutable, machine-verifiable proof of analytical rigor for institutional stakeholders. The Clarity Framework™ and The Rigor Protocol are presented as the applied methodology and organizational standard required to operationalize AI Judgment Infrastructure at scale. Our central thesis: judgment is the last unscalable asset, and the infrastructure we build today will determine who commands the next decade of capital deployment.

Keywords: AI Judgment Infrastructure, Judgment Graph, Clarity Framework, Defensible Audit Log, RUNE Protocol, venture capital diligence, brittle assumptions, capital allocation, investment thesis evaluation, institutional AI

PART I

The Market Thesis: From Retrieval to Rigor

1.1 The Commoditization of Information

For the better part of a decade, competitive advantage in capital allocation was defined by information asymmetry. The firm with the fastest access to the most accurate data held the structural edge. Analysts devoted thousands of hours to scraping, aggregating, and synthesizing market landscapes. Proprietary data pipelines were treated as institutional moats.

That era ended in 2024. The proliferation of foundational Large Language Models has reduced the marginal cost of information retrieval and summarization to zero. The consequence is categorical: when every market participant operates from the same omniscient information baseline, information ceases to generate alpha.

The market has systematically misclassified the utility of LLMs. Institutions have deployed generative AI primarily as a sophisticated search engine or content-generation mechanism. This is a critical strategic error. The competitive frontier has migrated upstream—from finding the data to stress-testing the logic derived from it.

1.2 The Judgment Gap

This commoditization exposes a structural vulnerability in modern capital allocation that we define as the Judgment Gap: the growing asymmetry between our exponentially expanding capacity to gather information and the persistently static cognitive capacity of human beings to evaluate it.

Human judgment at institutional scale is inherently brittle. It is degraded by cognitive fatigue, confirmation bias, the availability heuristic, and the social dynamics of partnership structures. Investment memos, in the absence of structured adversarial infrastructure, frequently devolve into post-hoc narrative justifications of decisions implicitly made during the first ten minutes of a pitch.

> *“The void between perfect information and rigorously evaluated logic is where capital is currently being destroyed at scale.”*

Analysis of the RUNE Protocol’s 60,000+ benchmarked Clarity Scores™ reveals an average structural quality rating of 39 out of 100 across all evaluated investment theses. The Series A pass rate registers at only 20 percent. The rarest signal—a genuine Paradigm Shift warranting immediate escalation—appears in fewer than 2 percent of submissions. These are not indictments of founder quality. They are indictments of an evaluation infrastructure that has failed to keep pace with deal flow.

1.3 The LLM as Adversarial Reasoning Engine

Closing the Judgment Gap requires a fundamental inversion of how institutions deploy AI—from using AI to generate text to deploying AI to evaluate logic. When architected correctly, the LLM functions as an adversarial reasoning engine: an analytical co-pilot structurally immune to consensus thinking, social pressure, and cognitive fatigue.

This requires moving categorically beyond consumer-grade prompts and retrieval-augmented generation. It demands a dedicated, enterprise-grade architecture: AI Judgment Infrastructure™.

PART II

The Empirical Foundation: The Grammar of Failure

2.1 The Foundry and the Field

Theoretical frameworks fail upon contact with actual capital deployment. The architecture of AI Judgment Infrastructure™ cannot be constructed in academic abstraction. It must be forged in what practitioners recognize as scar tissue: the accumulated operational experience of evaluating real founders, untangling complex market narratives, and deploying real capital under conditions of genuine uncertainty.

The empirical foundation of the RUNE Protocol was engineered directly from the operational realities of active capital allocation—not from simulated deal flow, curated case studies, or academic proxies.

2.2 The Seven Archetypes: A Grammar of Failure

Analysis of 60,000+ Clarity Scores™ benchmarked through the RUNE Protocol has enabled the identification and codification of seven empirically recurring archetypes of investment thesis failure. The fundamental insight: failure is rarely born from bad data. It is born from faulty logic applied to good data.

#	Archetype	Frequency	Structural Signal
I	The Service Trap	~31%	Platform multiples claimed on service economics; revenue scales linearly with headcount
II	The Hardware Denial Curve	~9%	Manufacturing capitalization under-modeled by 10x+; mass-production milestones with sub-threshold capital asks
III	Super-App Indigestion	~14%	3+ distinct business lines at pre-seed; complexity positioned as a moat
IV	The Structural Kill Shot	~4%	Unlicensed securities activity, fabricated pipeline, commingled custody, guaranteed return mechanics
V	The Dangerous Asset Class	~11%	High Presentation Score masking low Clarity Score; polished narrative over broken unit economics
VI	The Regulatory Arbitrage Illusion	~8%	Business model predicated on a regulatory vacuum that closes before scale is achievable
VII	The Paradigm Shift	~2%	Genuine structural novelty; the rarest signal and the only verdict warranting immediate escalation

Table 1: The Seven Archetypes of Investment Thesis Failure. Frequencies derived from 60,000+ Clarity Scores™ benchmarked through the RUNE Protocol (patent-pending). Corpus compiled in less than four months, representing 12+ months ahead of original projection trajectory.

2.3 Judgment as the Last Unscalable Asset

Venture capital operates on the golden gut mythology: the belief that elite investor intuition is mystical and irreducible. Examined rigorously, intuition is rapid, subconscious pattern recognition. The problem is not that this intuition is invalid. The problem is that it is entirely unscalable. AI Judgment Infrastructure™ is the first instantiation of infrastructure designed to scale this asset.

PART III

The Architecture of AI Judgment Infrastructure™

3.1 The Clarity Framework™

At the base layer of AI Judgment Infrastructure™ is the Clarity Framework™—a methodological protocol that operates as a standardized syntax for the evaluation of investment claims. The result is a Clarity Score™: a quantified, reproducible measure of structural investment merit that exists independently of the persuasiveness of the pitch.

3.2 The Judgment Graph™

The Judgment Graph™ stores the relationships between claims, evidence, and historical failure patterns, structured as a dynamic, densifying knowledge graph. Each additional thesis evaluated increases the graph's density and the precision of its pattern-recognition across all future evaluations. It is not a database of deals—it is a knowledge graph of how deals fail.

3.3 The Crucible

The Crucible is the application layer where investment narratives collide with mechanical reality. Its critical design principle is the elimination of narrative momentum: the Crucible evaluates business mechanics in isolation from the quality of the presentation. The most dangerous investment in any portfolio is the one with a high Presentation Score and a low Clarity Score™.

PART IV

Operationalizing Judgment: Outcomes and Artifacts

4.1 Isolating Brittle Assumptions

The immediate operational benefit of AI Judgment Infrastructure™ is the systematic identification of brittle assumptions—the one or two critical dependencies in every investment thesis that, if invalidated, destroy the entire model. The infrastructure relentlessly surfaces them before capital deployment rather than after.

4.2 The Defensible Audit Log

The canonical output artifact of AI Judgment Infrastructure™ is the Defensible Audit Log. A traditional investment memo is a static, frequently self-justifying document constructed after a decision has been implicitly made. A Defensible Audit Log is a machine-verifiable proof of analytical process: an immutable, transparent record of exactly which claims were evaluated, which assumptions were identified as brittle, and which logical chain produced the final verdict.

For Limited Partners, the Defensible Audit Log shifts the allocator's value proposition from 'trust our instincts' to 'verify our rigor.'

Asset Class	Risk Assessment	Audit Standard	Judgment Infrastructure
Credit	Credit scoring	Underwriting	Moody's / S&P;
Insurance	Actuarial modeling	Reserving standards	Actuarial institutes
Public equity	Fundamental analysis	GAAP / IFRS audit	PCAOB standards
Venture capital	Partner intuition	Investment memo	AI Judgment Infrastructure™ (askOdin)

Table 2: Venture capital remains the only major asset class without a formalized judgment infrastructure layer. AI Judgment Infrastructure™ is its first instantiation.

4.3 Eliminating Cognitive Debt

By delegating structural stress-testing to the infrastructure layer, AI Judgment Infrastructure™ eliminates cognitive debt. The General Partner's cognitive capacity is redirected toward activities software cannot replicate: relationship assessment, board governance, and the contextual judgment that only accrues from decades of direct market engagement.

PART V

The Empirical Validation: Corpus Metrics

5.1 Scale and Trajectory

The following metrics are drawn from live infrastructure deployment, not from projected or modeled data.

Metric	Value
Clarity Scores benchmarked	60,000+ (as of April 2026)
Corpus compilation period	Less than 4 months
Projection variance	12+ months ahead of original trajectory
Average Clarity Score	39 / 100
Structural pass rate (overall)	68%
Series A pass rate	20%
PRIORITY verdicts issued	1 in 75 evaluated
Paradigm Shift detections	2 per 134 audited decks
Peak single-day throughput	7,064 Clarity Scores (March 29, 2026)

Table 3: RUNE Protocol operational corpus metrics as of April 2026. Peak single-day throughput of 7,064 Clarity Scores™ (March 29, 2026) demonstrates infrastructure capacity at institutional scale.

5.2 What the Corpus Reveals

The 39-point average Clarity Score™ means that 61 percent of investment theses contain brittle assumptions, logical gaps, or architectural contradictions that would be material to an investment decision if identified before capital deployment rather than after. The 20 percent Series A pass rate means the current institutional evaluation infrastructure is not catching the failure modes the Clarity Framework™ systematically identifies.

The peak single-day throughput of 7,064 Clarity Scores™ demonstrates AI Judgment Infrastructure™ operating in a regime where human-only evaluation is not theoretically possible.

PART VI

The Rigor Protocol: An Organizational Standard

6.1 Internalizing the Standard

The Rigor Protocol is a commitment to a different epistemic standard: decisions are made by whoever presents the most structurally sound defense against the infrastructure's adversarial models, not by whoever constructs the most compelling narrative.

6.2 The Infrastructure Imperative

The first wave of AI leveled the information playing field. The second wave will aggressively separate the market based on the rigor with which information is evaluated. Firms operating with analog judgment mechanisms will find themselves structurally outmaneuvered not by better-funded competitors, but by better-infrastructured ones.

> *"Implementing AI Judgment Infrastructure™ is no longer a forward-looking experiment. It is the minimum viable standard for institutional capital allocation in the current era."*

The Judgment Graph™ densifies with each evaluation. The Clarity Framework™ improves with each identified archetype. The advantage of moving first is not incremental. It is architectural.

CONCLUSION

Judgment Is the Last Unscalable Asset

The central claim of this paper is not that AI will replace the General Partner. It is that the General Partner without AI Judgment Infrastructure™ is operating with a structural disadvantage that will compound over every fund cycle that passes without correction.

Credit has underwriting. Insurance has actuarial science. Public equity has GAAP and PCAOB. Venture capital has had partner intuition. That era is ending.

AI Judgment Infrastructure™ does not make judgment easier. It makes rigor obligatory. It makes the brittle assumption visible before it destroys the portfolio company. It makes judgment—finally, structurally, at scale—auditable.

Judgment is the last unscalable asset. The infrastructure described in this paper is the first credible attempt to scale it.

DISCLOSURE

Author Note & Contact

Lok Yek Soon is Founder & CEO of askOdin Pte. Ltd. (UEN: 202531656N), incorporated in Singapore. The RUNE Protocol, Clarity Framework™, Judgment Graph™, AI Judgment Infrastructure™, and Defensible Audit Log are proprietary methodologies and terms of askOdin Pte. Ltd. The RUNE Protocol is patent-pending (U.S. Provisional Application No. 63/948,559). Corpus metrics are drawn from live operational deployment. This working paper is published simultaneously on SSRN, arXiv, and askodin.app/research. The author has operational and financial interest in the subject matter of this paper.

Contact: Lok Yek Soon · yeksoon@askodin.app · askodin.app · Crucible (free): crucible.askodin.app