



**From Knowledge Quality
to Neuroplastic Enterprise Intelligence**

**A pre-market category thesis for structural defect
removal in organizations and AI.**

intelligence, made immanent.

immanait Strategy Series – Pre-Market Category Thesis

Dr. Thomas R. Glück

Abstract

Organizations often fail not because relevant knowledge is absent, but because it is structurally unable to affect selection. Around that defect, markets arise that are paid to manage recurring symptoms.

**The status quo monetizes the *treatment* of organizational defects.
immanait is built to monetize their *removal*.**

immanait is a deep-tech program for organizational intelligence. It combines two complementary capabilities:

- **immanait KQ** — an empirically grounded standard for detecting structural losses of decision-relevant knowledge.
- **immanait NEO** — a patented functional graph architecture for implementing and governing Dynamic Structure Systems.

A longer-term research track, **immanait NANN**, extends the same structural principle to learning systems and the prospect of structurally controllable AI.

This is a category-creation project, not a conventional product comparison. The thesis is not that one more tool can optimize the existing integration and governance stack. It is that substantial parts of that stack become unnecessary when organizational reality is represented and controlled through a structurally superior architecture.

Contents

1	The category break	3
2	The blind spot: Passive Qualitative Disinformation	4
3	Organizational intelligence is substrate-independent	5
4	Why now	7
5	Two coupled levers	8
6	immanait NEO: radically simple	9
7	More controllable complexity, less structural overhead	11
8	Implementation without a big bang	12
9	Evidence and current status	13
10	Strategic consequence: category creation, not competition	15
11	The research horizon: structurally controllable AI	16
12	The immanait convergence	17
13	Perspective	18
14	Further material	19

1. The category break

Category creation does not begin by claiming superior performance on incumbent criteria. It begins by showing that the incumbent criteria describe the wrong problem.

immanait is therefore not intended to be:

- another organizational development framework,
- another graph database,
- another integration platform,
- another AI-governance layer.

It is a combined epistemic and technical architecture for making organizational intelligence directly assessable, designable, and operational.

The core progression is:

1. **PQD** identifies the structural loss of decision-relevant knowledge.
2. **KQ** makes that loss visible, measurable, and legitimately contestable.
3. **NEO** preserves and operationalizes relevant distinctions in a directly controllable architecture.
4. **NANN** extends the principle to adaptive artificial systems.

In condensed form:

KQ makes better selection legitimate.
NEO makes it operational.
NANN extends it to learning systems.

2. The blind spot: Passive Qualitative Disinformation

The mechanism is quieter than error or ignorance. Relevant knowledge is discounted, overwritten, or filtered out by routine long before its quality is ever tested.

This is **Passive Qualitative Disinformation (PQD)**: the systematic devaluation of potentially decision-relevant knowledge before its quality is adequately assessed. The information is present, but prevented from participating fairly in selection.

Its invisibility is what makes it durable. Once embedded in routines, hierarchies, incentives, reporting systems, or accepted management language, the defect becomes self-stabilizing. The organization keeps processing information while the quality of its selection quietly degrades.

What forms around such defects is a **compensation market**: an economy that manages recurring symptoms without removing the structural cause. It can remain perfectly rational for its participants while remaining irrational for the client, because its economics may reward continued treatment more reliably than permanent removal.

This logic is not confined to consulting. It also applies to the parts of the technology market whose demand grows with unresolved fragmentation: repeated integration, translation, reconciliation, governance, and remediation. Both sides can monetize permanence.

immanait does not claim that all advisory or integration work is compensatory. It targets the subset whose demand is generated by avoidable fragmentation, indirect control, and structurally poor knowledge quality.

A defect that pays for its own management does not get removed.
It gets an industry.

3. Organizational intelligence is substrate-independent

For this project, intelligence is defined functionally, not by substrate.

Organizations are engineered cognitive systems. They observe, retain information, select actions, learn, coordinate, and adapt across individuals and generations. In that functional sense, they can be analyzed as a form of artificial intelligence — and their potential exceeds that of any individual member through:

- cross-generational technical progress,
- accumulated institutional knowledge,
- distributed observation,
- collective decision architectures,
- coordinated action across large systems.

The most useful definition comes from W. Ross Ashby:

intelligence := power of appropriate selection

Intelligence is therefore not binary. It is a variable capability: measurable, improvable, and — decisively — capable of degeneration.

A second principle concerns **compression**. An intelligent system must preserve the distinctions required for good selection while shedding representational and coordination overhead that carries no decision value. This is **not** a call to reduce relevant complexity. It means:

- increasing **controllable variety** — the variety the system can actually regulate,
- reducing redundant and inconsistent representations,
- preserving decision-relevant distinctions,
- eliminating avoidable translation and meta-layers.

In short:

More controllable variety. Less representational redundancy and inconsistency.

Modern foundation models offer one illustration of how substantial capability can be encoded in distributed representations that compress regularities across large bodies of data. immanait pursues an analogous architectural principle for enterprise intelligence: more decision-relevant variety carried by less representational redundancy.

4. Why now

Bad entropy is free; good organization is expensive.

Ashby's law of requisite variety sets the benchmark: effective regulation requires enough controllable variety to absorb the relevant variety of disturbances and system states.

In practice, the variety an organization can actually control often remains far below the variety of the environments, processes, technologies, and dependencies it is expected to govern. The gap is papered over through simplification, local "optimization", manual coordination, duplicated data, and expanding governance layers.

The AI revolution is the ultimate stress test for this gap. It both enables and forces a substantial increase in complexity competence.

A local AI assistant can succeed inside one bounded workflow. Scaling hundreds of agents across inconsistent data, shifting dependencies, permissions, controls, and audit requirements is a different problem — an architectural coordination problem. That is why many AI pilots work locally but fail to scale organizationally: fragmented infrastructures, indirect integration, slow change cycles, and proliferating control layers turn every extension into another project.

meta is beta: every compensatory meta-layer adds translation and coordination cost without repairing the underlying representation.

Organizations that cannot raise their controllable variety may gain isolated AI productivity while losing system-level control.

5. Two coupled levers

5.1. immanait KQ: problem-definition and epistemic authority

immanait KQ is an empirically and epistemologically grounded knowledge-quality standard.

It identifies when relevant knowledge is being qualitatively devalued and provides an evidence-based foundation for challenging incumbent system rationality. It does not legitimize arbitrary disruption; it legitimizes change when the existing selection architecture can be shown to be inferior.

KQ therefore supplies:

- **problem-definition authority** — the ability to identify structural knowledge-quality defects;
- **epistemic legitimacy** — the basis for distinguishing justified change from mere preference;
- **simulation capability** — the ability to examine how changes in input quality affect organizational intelligence in either direction;
- **auditability** — a traceable basis for evaluation and intervention.

Without legitimacy, even trivial improvements can fail against entrenched system rationality. Criticism without a demonstrably better alternative risks becoming merely destructive.

5.2. immanait NEO: implementation capability

immanait NEO provides the technical architecture for implementing justified redesign.

It translates epistemically grounded improvement into a versioned, directly controllable operational structure. NEO is therefore not simply a data layer. It is an integration and intelligence kernel that coordinates measurement, control, dependencies, change, and artificial intelligence in a single functional architecture.

KQ determines which distinctions deserve to govern decisions. NEO provides the architecture in which those distinctions are represented, propagated, versioned, and audited.

KQ and NEO can be used independently. Their combination is the stronger category proposition.

6. immanait NEO: radically simple

immanait NEO is built around three components.

1. Patented functional graph architecture

Organizational states, dependencies, measurements, controls, and changes are represented in a dynamic functional graph.

A single versioned structure replaces uncontrolled mutation and eliminates the architectural need for independently maintained copies of the same functional state. State transitions become traceable and auditable, while prior represented states remain recoverable.

2. Recursion-free processing protocol

The processing core propagates changes through affected dependency structures without recursion-limited failure.

Its scaling behavior is designed to follow the affected subgraph rather than require repeated global recomputation. Preliminary dense-graph tests show processing cost per edge decreasing as graph size grows across the tested range.

3. AI-native intelligence kernel

The NEO Core acts as an operational intelligence layer for the organization. It can coordinate measurement and control systems, propagate change, preserve lineage, support simulation, and provide a structural basis for adaptive AI integration.

Traditional fragmented architectures generate translation, consistency, and change problems as relational complexity grows. Graph-native representations are better suited to highly interdependent domains — but NEO goes beyond using a graph database.

Its category claim lies in combining, in a single architecture:

- functional dependency structures,
- versioned state,
- controlled propagation,
- organizational coordination,

- auditability,
- native coordination of AI agents and control systems.

The objective is a common functional model spanning planning, operations, simulation, AI coordination, compliance, and other highly interdependent domains.

7. More controllable complexity, less structural overhead

NEO is not built on eliminating complexity indiscriminately. Its design principle is to increase the amount of relevant complexity that can be controlled while reducing the representational and bureaucratic overhead created by fragmented systems.

Legacy and many current architectures do not compress organizational reality. They duplicate it across applications, translate it through interfaces, reconcile it through governance processes, and stabilize it with additional meta-layers.

NEO instead represents organizational dependencies in a unified, living variable graph. The intended result is:

- fewer redundant mutable copies,
- fewer translation layers,
- less reconciliation effort,
- faster structural change,
- direct lineage and governance,
- greater controllable variety.

Enterprise intelligence can then scale not only through additional resources, but through better representation and lower structural entropy.

8. Implementation without a big bang

NEO emerged from repeated attempts to generalize custom systems built to overcome the constraints of standard software. Each implementation moved more functionality into a reusable base architecture; the current system is the result of that cumulative path.

Standard software usually requires the organization to adapt to the assumptions embedded in the product — at high implementation cost, with collateral process damage and provider lock-in. NEO inverts that relationship:

system follows organization

The architecture adapts to organizational requirements rather than forcing organizational reality into a fixed application model.

8.1. Migration principles

- **Context-adaptive modeling:** a domain-independent architecture represents organization-specific structures and dependencies across different domains and scales.
- **Semantic preservation:** existing structures and processes are migrated with their operational meaning preserved and explicitly validated.
- **Parallel validation:** migrated components can run alongside incumbent systems, so correctness, performance, and fit are tested before cutover.
- **AI-assisted migration:** analysis of schemas, source code, interfaces, and process documentation can substantially reduce reverse-engineering and migration effort.
- **Structural governance:** lineage, versioning, documentation, and control become architectural properties rather than downstream layers.

The objective is not a risky, organization-wide replacement, but a sequence of controlled migrations in which each transferred component is benchmarked against the incumbent implementation.

After migration, changes that once required projects can increasingly be implemented as controlled modifications to the common functional structure.

9. Evidence and current status

immanait is in its pre-market phase. The conceptual and technical foundations exist; the category proposition is still undergoing validation.

Area	Current status
KQ / PQD	Empirical and theoretical foundations documented; operationalization published
KQ Explorer	Interactive implementation and standalone demonstration available
Dynamic Structure Systems	Technical category and architecture documented
NEO architecture	Patented and implemented in an initial proof of concept
Processing core	Formalized; preliminary internal dense-graph benchmarks completed
Migration	Initial implementation refined and migrated to an industry-standard platform
Benchmarking	Reproducible external documentation in preparation
NANN	Confidential research program active; public IP-strategic technical disclosure released; implementation remains a future R&D objective

9.1. Patents

Patent protection has been granted in Japan, Canada, India, and — most recently — the United States in May 2026. The European application is currently pending.

The claims are designed to cover the core functional design space. The architectural simplicity narrows the space of obvious, functionally equivalent design-arounds, though claim scope and enforceability remain jurisdiction-specific.

9.2. Proof of concept

The initial proof of concept demonstrated:

- correct propagation under intertwined multi-edits,
- termination in the tested dense and cyclic dependency structures,

- recursion-free processing.

The implementation has subsequently been migrated to an industry-standard platform.

Preliminary benchmarks indicate that processing cost follows the affected dependency structure. In the tested dense-graph range, cost per edge decreased as graph size increased. The test run reached an infrastructure limit before any algorithmic limit was observed.

These results are promising but preliminary. A reproducible benchmark protocol and full documentation are in preparation.

10. Strategic consequence: category creation, not competition

Rather than competing on incumbent terms — more tools, more integration, more governance layers — immanait NEO changes the value curve itself. Its intended strategic effect is:

- lower integration effort,
- lower governance overhead,
- reduced provider dependency,
- simpler architecture,
- faster change cycles,
- greater scalability of AI integration,
- substantially greater organizational neuroplasticity.

This does not optimize the existing trade-offs; it aims to dissolve substantial parts of them.

If organizational data, dependencies, controls, and change logic can be represented in a directly governed, dynamic functional structure, substantial parts of the compensatory integration stack lose their function — including parts of ETL, iPaaS, API management, master-data management, downstream data governance, and reconciliation-heavy integration services.

The claim is conditional and testable: wherever NEO replaces fragmented representations and indirect coordination with a unified functional structure, the associated compensation effort should decline materially.

The economic opportunity is therefore proportional to the integration, governance, maintenance, and change costs that can be removed structurally. Quantifying that opportunity for initial target markets is part of the current pre-market work.

AI is already eroding the economics of labor-intensive integration and advisory models. That creates a transition path for capable incumbents: from selling compensatory effort to implementing structural improvement.

11. The research horizon: structurally controllable AI

immanait starts from the hypothesis that behavioral rules, prompting, monitoring, and post-hoc governance are necessary but insufficient for reliable AI control. These approaches remain relevant, but they operate primarily above the architecture that generates behavior.

NANN — Neuroplastic Artificial Neural Networks — is the research hypothesis that learning systems should be able to modify not only their weights, but also their topology, under explicit structural constraints, versioning, and continuous auditability.

The dynamic data structure underlying NEO is already claimed in the patent as a potential base architecture for artificial neural systems.

NANN would not replace alignment, governance, or institutional control. It would add a structural control layer beneath them, addressing limitations that policies and behavioral correction alone cannot resolve.

The accompanying public NANN paper is not a disclosure of the confidential research program. It is an IP-strategic technical disclosure and position paper that establishes authorship, publicly stakes out the broader topology-adaptive design space, supports market positioning, and creates a citable disclosure that may constitute relevant prior art in relation to later third-party patent claims. The concrete research architectures, experimental priorities, and implementation path remain confidential.

NANN is a research-and-development horizon, not a completed product claim.

12. The immanait convergence

The three parts form one progression:

KQ makes organizational knowledge more intelligent by improving the quality and legitimacy of selection.

NEO makes the more intelligent organization technically operable, reconfigurable, and auditable.

NANN extends the principle toward structurally controllable learning systems.

Together, they support a category thesis built on:

- problem-definition and epistemic authority through KQ,
- implementation capability through NEO,
- a research option for structural AI control through NANN.

13. Perspective

immanait's category ambition rests on four moves:

- identifying problems that incumbent categories cannot adequately name,
- legitimizing departures from inferior system rationality,
- implementing improvement in a directly controllable architecture,
- setting new standards where existing trade-offs can be dissolved rather than optimized.

The prevailing economics reward repeated treatment.

immanait is built to reward structural removal.

No competition within existing categories — the setting of new dominant standards.

This is not yet a claim of achieved market dominance. It is a testable research, architecture, and implementation program. The next stage is validation through partnerships, benchmarked applications, technical due diligence, and carefully selected implementations.

14. Further material

A current list of the author's public publications is available through the author's ORCID record: [ORCID 0009-0009-8514-8592](https://orcid.org/0009-0009-8514-8592)

For research partnerships, pilot applications, and technical due-diligence inquiries, see immanait.ai



For supporting materials and technical background, see immanait.ai

© Dr. Thomas R. Glück. All rights reserved.