

NARRIEL White Paper (EN)

*Semantically Stable, Auditable, and Reproducible Prompt
Architecture for Productive AI Systems*

Version: 2026.05-R2

Date: May 20, 2026

Copyright © Andrey Stepanenko. All rights reserved.

www.narriel.com

1. Executive Summary

NARRIEL addresses a fundamental problem in productive GenAI systems: semantic drift. This refers to the gradual shift of meaning even when inputs, prompts, or workflows appear formally unchanged. NARRIEL approaches this challenge not as a stylistic prompt issue, but as an infrastructure and governance problem.

The architecture separates stable invariants, evidence ingestion, and goal-oriented rendering into independent layers. This enables more reproducible, traceable, and auditable AI workflows.

2. The Problem Space: Semantic Drift

In productive AI environments, operational failures rarely originate from poor language quality alone. More often, they emerge from inconsistent meaning. The same input may produce different interpretations depending on model behavior, context, pipeline order, or organizational conditions.

This instability increases debugging costs, governance risks, inconsistent communication, and reduced operational reliability.

3. Methodological Core

NARRIEL uses a layered semantic architecture:

- Canon: stable invariants, semantic boundaries, terminology, and structural constraints
- Ingest: evidence, source material, and contextual candidate space
- Prompt/Goal: goal-oriented selection and rendering

Processing is intentionally non-commutative. Sequence itself is treated as a semantic contract.

4. Quantum Semantic Modeling

Within the NARRIEL context, the term 'quantum semantic' does not describe a physics claim. It represents a structural model for ambiguity, contextual selection, and observer-dependent interpretation.

Meaning is therefore not treated as a fixed entity, but as a context-dependent realization within a semantic probability space.

5. Technical Architecture

The project structure demonstrates a modular API and pipeline architecture featuring deterministic control mechanisms, traceability, shadow execution, hash logic, and observability layers.

Core mechanisms include:

- Idempotency and fingerprint systems
- Trace and audit logs
- KPI and policy gates
- Scope and guardrail controls
- Shadow comparison strategies

6. Governance and Reproducibility

NARRIEL does not claim absolute determinism. Instead, it aims for controlled reproducibility under explicitly defined conditions.

This creates a practical foundation for governance, quality assurance, and semantically traceable AI operations.

7. Failure Taxonomy

A central characteristic of the architecture is the explicit modeling of semantic failure classes.

Examples include:

- Canon Drift
- Layer Blending
- Evidence Break
- Processing Order Drift
- Certainty Overstatement

These mechanisms make semantic instability observable, measurable, and operationally manageable.

8. Operational Use Cases

NARRIEL is particularly suitable for:

- Enterprise governance
- regulated AI environments
- multi-agent and multi-tenant systems
- brand and communication consistency

- documentation-critical operations
- semantically stable assistant systems

9. Limitations and Risks

Even with stable architecture, AI systems remain dependent on changes in models, runtime environments, and infrastructure. Semantic stability can be improved, but never fully guaranteed.

Poor or contradictory source material also limits the achievable output quality.

10. Strategic Positioning

NARRIEL treats meaning not primarily as a stylistic problem, but as an infrastructure problem. The central question is therefore not: 'How do we generate better text?', but rather: 'How do we prevent semantic entropy in productive AI systems?'

11. Conclusion

The architecture combines semantic governance, technical traceability, and structural layer separation into a coherent operational model for productive AI systems.

NARRIEL positions itself less as a conventional prompt framework and more as a semantic orchestration and control layer for reproducible AI workflows.

Legal Notice

This document is provided solely for informational and analytical purposes. All concepts, texts, models, structural approaches, and architectural principles remain the intellectual property of Andrey Stepanenko unless otherwise stated.

Any reproduction, redistribution, publication, commercial use, modification, or derivative work based on this document or its components requires prior written permission from the copyright holder.

NARRIEL, its semantic architecture principles, and associated modeling approaches are protected intellectual property.

© Andrey Stepanenko – All rights reserved.