Stackmint

NOVEMBER 28, 2025

THE SEMANTIC STANDARD

FOR ENTERPRISE AI

PRESENTED BY
Florian Boymond

WHY SEMANTICS NOW

Enterprise AI is failing for a simple reason: models generate text, but *businesses run on meaning*.

Most systems today can predict the next token. Very few can reliably apply the right business rule at the right time.

A 2024 McKinsey survey found 81% of enterprises cite “context misalignment” as the top blocker to scaling AI. BCG reports that 72% of AI pilots fail to reach production.

This standard introduces a simple framework: make business meaning explicit, machine-readable, and enforceable across every agent and workflow.

WHAT SEMANTICS MEANS IN THE ENTERPRISE

Semantics is the layer that gives AI systems business meaning instead of raw pattern-matching.

Enterprises rely on shared concepts:
- What is an Account?
- When is an Opportunity considered At Risk?
- What does it mean to Qualify, Escalate, or Approve?

In this standard:
- Ontology = the nouns of the business
- Semantics = the verbs, constraints, and rules
- Context = how meaning shifts by product, segment, region, or persona

Agents should not invent definitions.

They should execute consistent, governed meaning across every workflow.

THE SEMANTIC FRAMEWORK

The Semantic Standard for Enterprise AI is organized into five layers:

O**ntology (Nouns)**
The entities the business operates on: Account, Opportunity, Contract, Incident, SLA.

Semantics (Verbs + Constraints)
The actions, rules, and definitions that govern how entities behave.

Branches (Composed Logic)
Workflows built from semantic actions, with clear inputs, outputs, and governance.

Agents (Executable Meaning)
Runtime processes that choose and execute Branches based on context.

Runtime & Audit (Governance)
Identity, permissions, policy enforcement, observability, and traceability.

Together, these layers turn AI from "answer engines" into execution systems that run on meaning.

WHY AI BREAKS WITHOUT SEMANTICS

Without a semantic foundation, enterprises see the same failure modes on repeat:

Ambiguous definitions – “Qualified,” “At Risk,” or “High Priority” mean different things to Sales, RevOps, and Finance.

Fragmented schemas – CRM, ticketing, billing, and product data are not aligned on the same entities.

Inconsistent workflows – teams build their own automations on top of conflicting business logic.

Model hallucinations – LLMs fill gaps in rules with guesses instead of governed semantics.

Tribal knowledge – key rules live in people’s heads and slide decks, not in systems.

No enforcement layer – nothing ensures workflows always respect constraints and approvals.

Weak governance – audit, rollback, and accountability are bolted on instead of designed in.

The Semantic Standard removes these failure modes by making meaning part of the execution layer, not an afterthought.

PRINCIPLES OF SEMANTIC GOVERNANCE

Effective semantic governance rests on five principles:

Meaning First, Models Second
Models consume semantics; they do not define them. Business meaning comes from the enterprise, not the LLM.

Deterministic Enforcement
Semantic rules are enforced by the runtime as constraints, not as "suggestions" to the model.

Shared, Reusable Definitions
A change in the definition of "Qualified Opportunity" should propagate everywhere that concept is used.

Separation of Logic and Execution
Business semantics live in a governed layer, decoupled from any specific workflow, UI, or agent.

Full Observability and Audit
Every semantic action leaves a trace: inputs, outputs, actors, rules, and decisions can be inspected.

These principles make AI safer to deploy and easier to operate at scale.

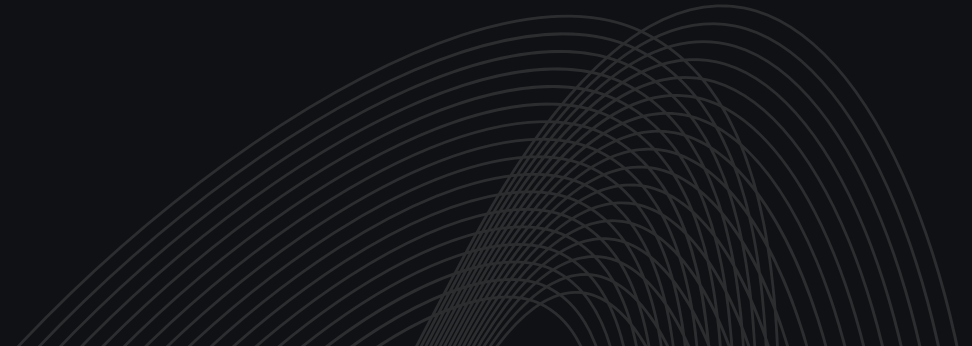

ONTOLOGY: THE NOUNS OF THE BUSINESS

Ontology is the structured catalogue of the entities an enterprise operates on.

Typical examples include:
- Account, Contact, Lead
- Opportunity, Stage, Forecast Segment
- Ticket, Incident, SLA
- Contract, Amendment, Invoice

A clear ontology reduces ambiguity across teams and tools. Research from MIT's Computational Structures Lab (2024) suggests that workflows aligned to an explicit ontology improve task accuracy by 23–41% in complex enterprise scenarios.

The practical rule: if an agent is allowed to touch it, it should exist in the ontology.

SEMANTICS: THE VERBS AND CONSTRAINTS

Semantics define **how entities behave and interact**.

Examples of semantic actions:
- Qualify Opportunity
- Score Lead
- Escalate Incident
- Approve Invoice
- Trigger Renewal Outreach
- Forecast Quarter-End Revenue

Each semantic action includes:
- required inputs
- expected outputs
- preconditions and postconditions
- error-handling rules
- governance and approval paths

Codifying semantics in this way turns business rules into reusable building blocks that agents and Branches can safely execute.

BRANCHES: MEANING IN MOTION

Branches are workflows built entirely from semantic actions.

A Branch might:
• monitor Opportunities for defined risk signals
• evaluate them against a shared “At Risk” semantic rule
• trigger a follow-up task, Slack notification, or CRM update
• log each step for audit and reporting

Because Branches are composed from shared semantics, they are:
• **deterministic** – the same inputs produce the same behavior
• **reusable** – the same Branch can run across products, segments, or regions
• **governable** – updates to semantics cascade safely through dependent workflows

A Branch is not just a flowchart. It is **meaning in motion**.

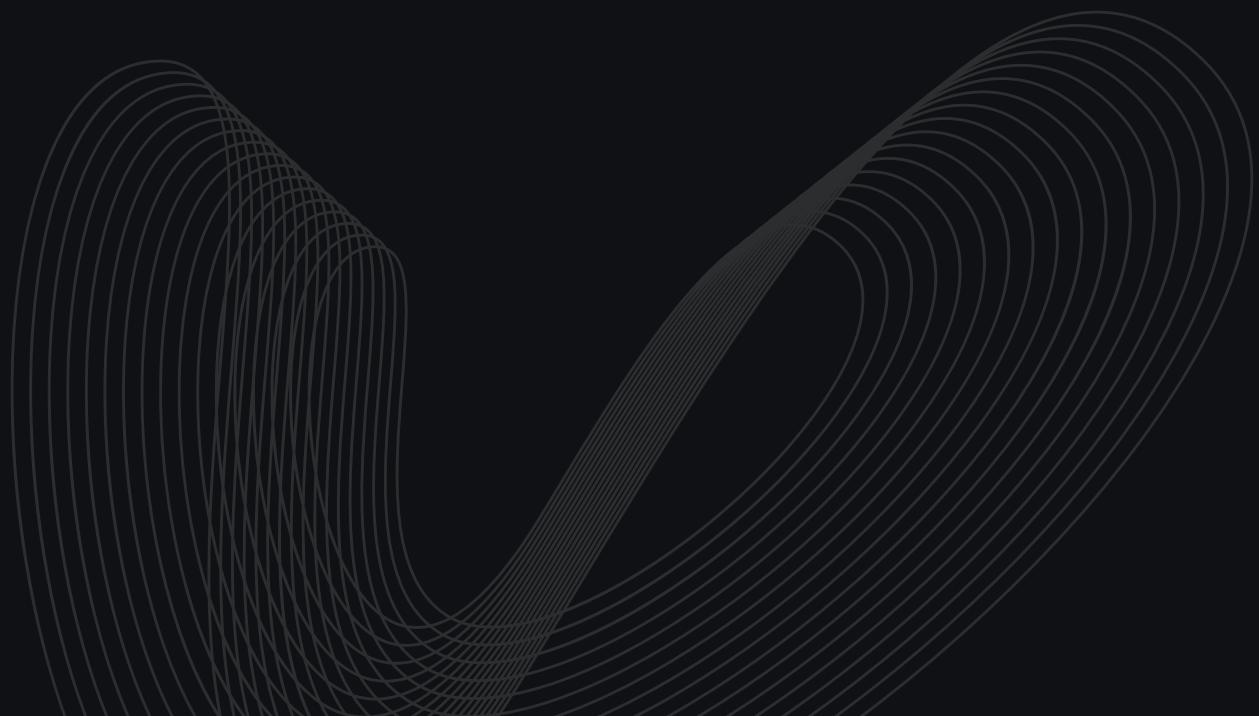

AGENTS: EXECUTABLE MEANING

Agents embody semantics at runtime.

An agent powered by the Semantic Standard:
- receives a goal and context
- evaluates which Branches and semantics apply
- calls the right actions with the right inputs
- enforces constraints and approvals
- emits a complete trace of what it did and why

This turns AI from a free-form assistant into a reliable operator working inside the guardrails of enterprise meaning, not outside them.

AUDIT, OBSERVABILITY, AND GOVERNANCE

Enterprises need more than successful runs; they need explainable runs.

A semantically governed execution layer provides:
- Identity – who or what initiated the action
- Inputs – which entities and values were involved
- Semantics – which business rules fired
- Branch context – which workflow ran, under which conditions
- Outputs – what changed in downstream systems
- Timeline – when each step occurred

This level of observability makes it possible for risk, compliance, and leadership teams to trust AI at the core of revenue, operations, and customer experience.

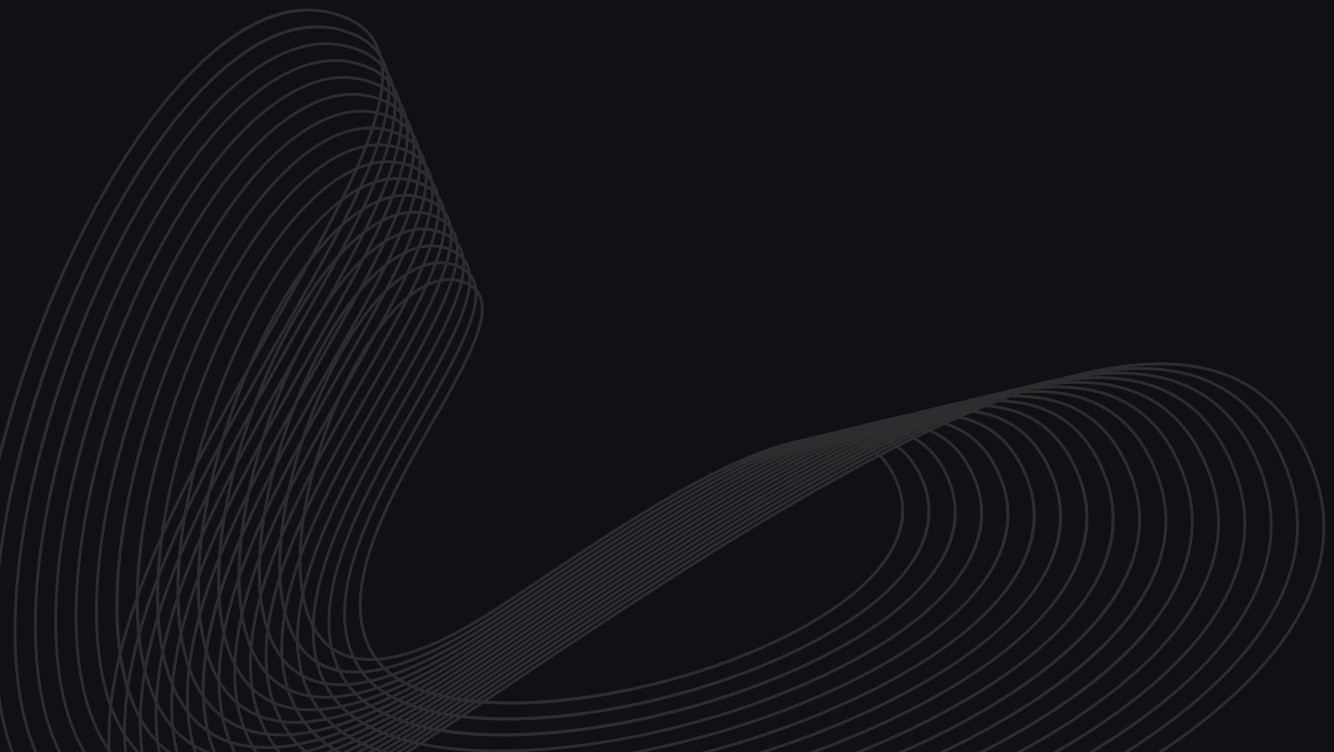

Stackmint

The Semantic Standard for Enterprise AI provides a practical foundation for building agents and workflows that run on meaning instead of guesswork.

It is:

- model-agnostic
- system-agnostic
- governance-ready

Learn more:
stackmint.ai/semantic-standard

GET IN TOUCH

stackmint.ai
hello@stackmint.ai
+1 (628) 257 6468