LATINFLASH RESEARCH · TECHNICAL PUBLICATION SERIES

PAPER 2026-04 · CANONICAL TECHNICAL SPECIFICATION · VERSION 1.0

Vertical AI Governance Stack v1.0: A Layered Reference Architecture for High-Assurance Agentic Systems

Manuel Enrique Morales Santiago

LatinFlash Research | ORCID iD: 0009-0007-6921-7688 | Dialnet Code: 3123473

Document Date: September 19, 2026 | Status: Independent Technical Specification

ABSTRACT

Canonical technical specification defining the six-layer Vertical AI Governance Stack (L1 Normative Authority to L6 Resilience & Fallback). Defines formal control mechanisms, delegation boundaries, and evidence collection patterns required to achieve verifiable autonomous agent execution.

1. Abstract

The **Vertical AI Governance Stack** is a vendor-neutral reference architecture for converting:

```
normative obligations → semantic policy → delegated authority → controlled execution → verifiable evidence → resilient operation
```

The Stack remains independent of LLM vendors, cloud providers, and agent frameworks. Its concern is governance architecture, not model selection.

2. Foundational Thesis

> *Models may change. Operational knowledge must remain sovereign.*

> *Probabilistic intelligence must never be allowed to override deterministic authority.*

> *Probabilistic intelligence may propose, reason and plan. Deterministic governance defines what may become real.*

3. The Invariant Matrix (L1–L6)

The core architecture establishes six non-negotiable layer invariants:

Layer	Governs	Canonical Invariant
L1	Authority	Models may interpret governance. They must not constitute governance.
L2	Meaning	The meaning of the business must not depend on the model interpreting it.
L3	Authority Delegation	Intelligence defines possibility. Governance defines authority.
L4	Execution Boundary	Authorization is not execution. Every authorized intent must pass through a controlled execution boundary.
L5	Evidence Ledger	Trust requires reconstructibility.
L6	Failure & Fallback	When assurance decreases, autonomy must decrease with it.

4. Governance Proof Obligations

For every material governed operation, an implementation MUST be able to demonstrate:

O1 / AUTHORITY

Why was this action permissible?

Demonstrable capability decision binding caller claims to explicit operational rules.

O2 / IDENTITY

Who or what requested it?

Cryptographically authenticated principal identity (human, service, agent, daemon).

O3 / SCOPE

Was the requested resource inside authorized boundaries?

Strict tenant and resource isolation check independent of client inputs.

O4 / INTEGRITY

Did actual execution correspond to authorized execution?

Parameter immutability and capability hash verification post-authorization.

O5 / EVIDENCE

Which verified facts supported the decision?

Verifiable execution audit records linking inputs, models, rules, and outcomes.

O6 / OUTCOME

What state change actually occurred?

Deterministic execution side-effect confirmation.

O7 / RECOVERY

What happens when this operation fails?

Graceful fallback to lower autonomy tiers or human intervention.

5. Governance Pipeline

All operational execution paths must flow through a Capability Verification Adapter (CVA):

Capability Verification Boundary (L4)

No model output or API call may bypass the CVA. Unverified execution attempts must be rejected at the boundary and logged as security coverage gaps.

6. Citation & Metadata

Institutional Citation: Morales Santiago, M. E. (2026). *Vertical AI Governance Stack v1.0: Canonical Technical Specification*. LatinFlash Research Paper 2026-04.

Canonical URL: https://research.latinflash.com/papers/vertical-ai-governance-stack-spec/ | ORCID: 0009-0007-6921-7688 | Dialnet Code: 3123473