VERISTACK · GOVERNANCE ARCHITECTURE

The VeriStack Governance Gate Framework

One reusable control model. Every agentic build.

> **VeriStack does not govern one use case. We apply the same seven-gate control framework across every agentic system we build, tuned to each domain's rules and regulator. The architecture is reused. Only the thresholds and lists change.**

Why this matters

Most teams building agentic AI treat governance as a one-off, rebuilt from scratch for every project. That is slow, inconsistent, and impossible to audit at scale. VeriStack solved this once and reuses it. Every agent decision, in any domain, passes through the same seven independent gates before an action is allowed to execute. This is defense in depth applied to agent actions, and it is our core reusable intellectual property.

We know it is reusable because it already is. Four VeriStack builds run a version of this framework today, covering AML compliance, invoice intelligence, accounting close, and facilities maintenance. The evidence below is drawn directly from those production codebases, not from theory.

The seven gates

#	Gate	What it enforces	Portability
1	**Confidence**	Low-confidence agent output is routed to a human, not executed	**Universal, as-is**
2	**Materiality**	High-value or high-impact actions get extra review regardless of confidence	**Universal, as-is**
3	**Hard-Stop Block**	A confirmed prohibited condition auto-blocks and notifies, no override	**Pattern, configured**
4	**Human-in-the-Loop**	Defined decision types always require a human approver in the path	**Universal, as-is**
5	**Escalation Chain**	Risk-tiered routing sends critical patterns to the right authority	**Universal, as-is**
6	**Eligibility Escalation**	Identity or eligibility failures escalate rather than silently pass	**Pattern, configured**
7	**Authority Gate**	The highest-stakes decisions route by role and value limit	**Pattern, configured**

Proof of reuse, from live code across four builds

Gate 1, Confidence, appears in all four builds

- AML blocks automated action below a 0.7 agent-confidence score.
- Invoice Intelligence gates both OCR confidence (below 0.85) and agent confidence (below 0.70).
- Accounting Close escalates any journal entry below its confidence threshold.
- Facilities Maintenance escalates low-confidence work orders for verification.

Gate 2, Materiality, appears across the money-touching builds

- AML triggers additional review above NGN 50M.
- Invoice routes by three approval tiers at $5K, $25K, and $100K.
- Accounting Close routes by a three-level materiality ladder (L1, L2, L3).
- Facilities escalates any work order above its cost threshold for budget approval.

Gates 4 and 5, Human-in-the-Loop and Escalation Chain, are the spine of every build

- AML routes by risk level to the compliance officer and above.
- Facilities routes by five distinct reasons: safety, high cost, equipment replacement, low confidence, and compliance flags.
- Accounting Close routes by materiality tier to the right approver.
- The mechanism is identical everywhere. Only the trigger changes.

Gates 3, 6, and 7 are the same pattern with a domain-swapped rule

> **Gate 3** is a sanctions list in AML, a duplicate-invoice block in Invoice, and a safety-compliance flag in Facilities. The gate is a confirmed prohibited condition that hard-stops with no override. The blocklist is configuration.
>
> **Gate 6** is a KYC check in AML and a vendor-registry eligibility check in Invoice. Same shape.
>
> **Gate 7** routes onboarding decisions by authority in AML and journal entries by approval level in Close.

What is portable, and what is not

Portable as-is (the universal core). Confidence, Materiality, Human-in-the-Loop, Escalation Chain. Four of the seven gates drop into any new build unchanged.

Portable as a pattern (swap the rule). Hard-Stop Block, Eligibility Escalation, Authority Gate. The machinery is reused, the specific list or check is configured per domain.

Not portable, thin by design. The exact threshold values, the specific blocklists, and the regulatory bindings, such as AML's NGN 50M, its OFAC and UN lists, and its CBN escalation policy. These are configuration and regulatory mappings, not architecture.

Net reuse, roughly 90 percent architecture and 10 percent per-domain configuration.

What this unlocks for VeriStack

1. **A methodology, not a project.** Clients do not buy a one-off. They buy a proven control framework applied to their domain. That is what lets a firm charge premium and scale, the same way a Big 4 sells a methodology rather than hours.
2. **A delivery accelerator.** A new governed build becomes define the gate rules for this domain, not build governance from scratch. That compresses timelines and improves margin.
3. **A moat.** The seven gates are real, tested, and running in production across four domains. Competitors selling built-in governance tools as a marketing bullet cannot demonstrate this. We can.
4. **A single audit story.** One control model across every system means one consistent audit trail, one explanation for a regulator, one thing to certify.

The governing principle

> In regulated environments, trust is not a model property. **It is a proof property.**
>
> You earn it by showing control, accountability, and an audit trail on every autonomous decision. The seven gates are how VeriStack delivers that proof, once, and reuses it everywhere.

Framework derived from four live governance engines: AgenticAML, AgenticInvoiceIntelligence, AgenticAccountingClose, AgenticFacilitiesMaintenance.