Vallignus Runtime Enforcement Incident Report

Incident ID: VALL-ENF-001

System: Vallignus Runtime Governance Layer

Date: 2026-01-28

Classification: Unclassified

Distribution: Demonstration Use Only

Prepared by:

Vallignus Technologies

1 Executive Summary

An autonomous agent attempted to initiate an outbound HTTP request to an external domain not authorized under the active runtime policy.

The request was intercepted at the execution boundary and denied prior to network transmission.

No external communication occurred.

A cryptographically structured audit record was generated synchronously at enforcement time.

This incident demonstrates deterministic execution-level governance independent of agent reasoning or intent interpretation.

2 System Context

The enforcement event occurred within a Vallignus-governed execution environment configured with a deny-by-default network posture.

Active Policy Configuration

- Default network access: DENY
- Explicit allowlist:
  - localhost
  - 127.0.0.1

All outbound network requests not matching the allowlist were configured to be denied deterministically at runtime.

The agent was executed within a supervised runtime context instrumented by the Vallignus enforcement layer.

3 Event Description

During execution, the autonomous agent initiated an HTTP GET request targeting an external domain.

Requested destination:

http://example.com/

The domain was not included in the active policy allowlist.

The request was intercepted synchronously by the runtime governance layer prior to network transmission.

No packets were sent externally.

4 Enforcement Decision

Upon interception, the enforcement layer evaluated the request against the active policy bundle.

Policy Evaluation Outcome

- Policy ID: deny-external
- Policy version: 1
- Evaluation result: DENY
- Deny reason: domain_not_allowed

The decision was made deterministically without reliance on model output interpretation, confidence scoring, or semantic reasoning.

Execution authority for the requested operation was not granted.

5 Enforcement Action

The enforcement layer terminated the request before network transmission.

The agent process continued execution under supervision, without the elevation of execution privileges.

No retries or alternative execution paths were able to bypass the policy constraint.

6 Audit Record

A structured audit entry was generated at the moment of enforcement.

Recorded Fields

- Timestamp
- HTTP method
- Requested URL
- Decision outcome
- Allow or deny status
- Policy identifier
- Policy version
- Agent identifier
- Owner context
- Enforcement reason

The audit record was written to an append-only flight log suitable for downstream review and attestation.

This log provides a complete chain of custody for the attempted action and its enforcement outcome. The audit record is immutable for the lifetime of the execution environment.

7 Impact Assessment

Security Impact

- No external network communication occurred
- No data exfiltration was possible
- No system compromise occurred

Operational Impact

- Agent execution continued under governance
- System stability was maintained
- Enforcement introduced no external dependency or delay

8 Observations

This incident exhibits the following properties of runtime governance:

Observation 1: Enforcement occurred at the moment execution authority was requested

Observation 2: Policy evaluation was independent of agent reasoning

Observation 3: Denial decision was deterministic and non-bypassable

Observation 4: Audit evidence generated synchronously at decision time

Observation 5: Agent autonomy preserved while authority remained bounded

The system did not attempt to predict agent behavior. It simply enforced execution constraints at runtime.

9 Conclusion

This enforcement event confirms that Vallignus provides execution-level governance for autonomous systems.

Rather than relying on alignment, prompts, or post-hoc monitoring, the system prevents unauthorized actions from occurring at the point where authority would otherwise be exercised.

The demonstrated capability is applicable to environments requiring strong guarantees around containment, accountability, and operational control of autonomous agents operating over extended time horizons.