TRULY AGENTIC (INDEPENDENT ENGINEERING PRACTICE — ANIL PERVAIZ)

OWASP Agentic AI (2026) v2026

Mapping of OWASP Top 10 for Agentic Applications (ASI01–ASI06) and Truly Agentic Five Runtime Controls to open-source TypeScript, PostgreSQL, and Terraform code in trulyagentic/truly-agentic-governance-kit v0.1.0.

November 24, 2026

Security Questionnaire Response

https://trulyagentic.io

security@trulyagentic.io

Response Summary

Answer	Count	Percentage
Yes	10	100.0%
Partial	0	0.0%
No	0	0.0%
N/A	0	0.0%
Total	10	100%

Goal Hijack & Context Integrity (ASI01 & ASI06)

3/3 controls satisfied

ID	Question	Answer	Explanation
ASI01-1	Does the gateway sanitize and inspect untrusted documents/emails before agent execution?	YES	scanUntrustedPayload in src/mcp-scanner.ts blocks HTML comment injections, role delimiter overrides, zero-width obfuscation characters, and instruction-hijack strings before policy evaluation.
ASI01-2	Are policy boundaries enforced deterministically outside model prompts?	YES	Language models never self-police boundaries. GovernanceGateway routes all tool invocations through deterministic TypeScript gates prior to execution.
ASI06-1	Are untrusted payload hashes captured in an immutable audit trail?	YES	Pre-execution audit entries store canonical input_sha256 digests chained via prev_hash -> entry_hash in PostgreSQL (db/migrations/001_audit_ledger.sql), preventing silent history tampering.

Excessive Agency & Privilege Control (ASI02 & ASI03)

4/4 controls satisfied

ID	Question	Answer	Explanation
ASI02-1	Are agent actions segmented across read, propose, and write authorization tiers?	YES	PolicyEngine maps tools across Five Levels of Agency (L1–L5) and enforces deterministic monetary ceilings, allowlisted currencies, IBAN country prefixes, and vendor domain checks.
ASI02-2	Do write-tier actions require named human approval and replay protection?	YES	ApprovalGate binds approval tickets to exact (agentId, toolName, inputSha256, idempotencyKey), verifies Slack HMAC-SHA256 signatures, accepts only individual user IDs (usr_*), and blocks duplicate click replays.
ASI03-1	Are agent identities locked to company-provisioned service accounts?	YES	AgentRegistry enforces registered svc-agent-* identities with explicit per-tool allowlists and rejects shared super-admin keys or personal emails.
ASI03-2	Can compromised or runaway agent leases be halted immediately?	YES	Control 4 includes pre-flight spend budget reservation (reserveBudget) and instant lease suspension (suspend), emitting an immutable kill_switch audit ledger entry.

MCP Supply Chain & Execution Safety (ASI04 & ASI05)

3/3 controls satisfied

ID	Question	Answer	Explanation
ASI04-1	Are MCP tool descriptions scanned for poisoning directives and schema drift ("rug pulls")?	YES	scanMcpToolDefinition inspects MCP tool metadata for <IMPORTANT> or cross-tool shadowing directives and verifies canonical JSON SHA-256 digests against pinned baselines (config/mcp-pins.example.json).
ASI05-1	Are tool invocations guarded against arbitrary code/shell evaluation?	YES	GovernanceGateway avoids dynamic eval/shell interpolation and strictly validates typed tool parameter schemas prior to execution.
ASI05-2	Do all automated test scenarios run cleanly without live customer data?	YES	All 14 Vitest specs and all 6 demo:attack scenarios execute cleanly in CI against synthetic lab fixtures with zero customer data or production credentials.