OntoGuard Decision Authorization for Agentic AI

OntoGuard evaluates the exact action an AI agent is attempting to take and returns ALLOW, BLOCK, or ESCALATE before that action becomes an enterprise consequence. A score is not authorization. A prior ALLOW does not bind a mutated action.

The Decision API is the sole ALLOW / BLOCK / ESCALATE authority. Semantic interpretation informs standing. It does not grant release.

Product surfaces

OntoGuard Decision Authorization Platform is the full operating environment for authorization, human review, evidence, analysis, reporting, testing, and governance operations. OntoGuard Headless Decision Authorization Runtime 1.2.21 is the self-contained embedded service for placing the same Decision Authorization core directly in an execution path. Governance Pack v1.2.0 is a separate public artifact/schema version line.

Runtime 1.2.21 adds a bounded OpenShell-compatible supervisor-middleware adapter and makes cryptographically sealed authorization material a normative handoff. The Decision Authorization kernel itself remains the same standing, binding, review, consequence, and recovery product described for the 1.2.20 line.

The OntoGuard TRACE integration is currently listed at the AgenTrust Verified tier. Verification covers the published integration behavior against released packages; it does not assert hardware attestation, customer-production L5, or universal non-bypassability.