

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.

Deployment: OntoGuard is available as a self-contained headless Decision Authorization Runtime (release 1.0.19) inside customer-controlled or confidential-computing environments. The runtime provides exact-action authorization, authenticated governance inputs, and durable decision evidence without requiring the OntoGuard web application. Public artifacts do not assert customer topology, hardware attestation, or L5.

Governance Pack v1.2.0 and Headless Runtime 1.0.19 are different version namespaces.

AgenTrust TRACE Community Integration is not AgenTrust Verified. 27/27 adapter tests and the controlled \$250K / \$260K paths were independently reproduced by an AgenTrust maintainer before merge.

Patent record: U.S. Patent Application 19/444,521 — Track I prioritized examination granted May 2026.