OntoGuard Decision Authorization Packet

AI output governed before release - public sample packet

Protected by U.S. Patent Application 19/444,521 - Track I prioritized examination granted May 2026. OntoGuard Direct LLM Backend Governance.

Before OntoGuard, the buyer had an AI output. After OntoGuard, the buyer has a governed business decision with evidence, routing, auditability, and improvement signals.

Financial Services Workflow | **Buyer-Safe Public Sample** | **Decision API** | **L1-L2-L3 Semantic Governance**

Decision ESCALATE	**Routed To** HUMAN_REVIEW	**Release Control** WITHHELD PENDING REVIEW
Trust 81.41% - below autonomous release threshold	**Triad** Cold 100.00% | Heat 6.00% | Mercury 96.67%	**Benchmark** Comparable corpus available; strict release gate did not pass
Primary Scope SEC / SOX / GLBA / FINRA	**Coverage** No coverage gaps; SEC citation linkage pending	**ROI** Hard-dollar ROI not calculated - buyer baseline required

Why this matters: OntoGuard did not rubber-stamp the AI output. It mapped financial-services evidence, detected a remaining SEC citation-linkage gap, failed the strict autonomous-release benchmark gate, and withheld release pending human review.

Executive Summary

This public sample shows how OntoGuard governs an LLM backend output before it reaches a regulated workflow. The scenario is a financial-services AI advisor copilot that drafts lending and investment recommendations from customer financial data, KYC records, credit risk signals, advisor notes, and transaction history.

OntoGuard evaluated the proposed output using mapped financial-services evidence, symbolic trace lineage, uncertainty monitoring, hallucination analysis, strict benchmark comparison, and a reproducible proof harness. The system found strong mapped coverage across SEC, SOX, GLBA, and FINRA scope, but identified a remaining SEC citation-linkage gap and failed the strict autonomous-release benchmark gate.

Decision: ESCALATE to HUMAN_REVIEW. Business effect: autonomous release was withheld pending review. The buyer receives a portable audit packet showing what was governed, why release was withheld, what evidence supported the result, and what the human reviewer should do next.

Scenario / Governed Request

A bank is deploying an AI advisor copilot that drafts lending and investment recommendations using customer financial data, KYC records, credit risk signals, advisor notes, and transaction history. Before the AI output is released to an advisor or customer, OntoGuard must determine whether to ALLOW, BLOCK, or ESCALATE it.

Evaluation context: suitability, fair lending, SEC/FINRA-style supervision, GLBA privacy, SOX auditability, model risk, hallucination risk, explainability, and human-review requirements.

Signal	Public sample value
Workflow	LLM backend output authorization before downstream release
Primary evaluated scope	SEC, SOX, GLBA, FINRA
Supplemental domains	Legal compliance, AI governance, privacy, responsible AI ethics
Required routing	ESCALATE -> HUMAN_REVIEW

Decision API and Release Control

Field	Value	Buyer meaning
Decision API action	ESCALATE	Autonomous release is not approved.
Routed to	HUMAN_REVIEW	A compliance/risk owner must review before release.
Release status	WITHHELD_PENDING_REVIEW	The output is controlled before entering the workflow.
Policy pass	Yes	Policy posture was acceptable under soft-pass criteria.
Benchmark pass	No	Strict autonomous-release benchmark did not pass.

Regulatory Coverage and Release Control

Regulation	Coverage	Citation linkage	Status
SEC	Mapped evidence present	83.33%	Review - citation linkage pending
FINRA	Mapped Clause Coverage 100.0%	100.0%	Pass
GLBA	Mapped Clause Coverage 100.0%	100.0%	Pass
SOX	Mapped Clause Coverage 100.0%	100.0%	Pass

Coverage gaps: None. Citation gaps: SEC. The packet remains honest: mapped evidence can be present while legal citation linkage still requires review.

Governed Response Handling

The original model response was captured and hash-preserved in the source governance artifact. This public sample intentionally shows a buyer-safe summary instead of reproducing the raw model text.

Buyer-safe response summary: The AI advisor output attempted to evaluate suitability, fair lending, privacy, supervision, auditability, and human-review considerations for a financial-services recommendation workflow. OntoGuard withheld autonomous release because the output required review against mapped financial-services evidence, SEC citation-linkage remained pending, and the strict autonomous-release benchmark gate did not pass.

Display mode: public-safe summary. Raw response text is omitted from this public packet to avoid exposing formatting artifacts or implementation-specific audit content.

Semantic Governance Triad

Cold - Evidence Grounding

100.00% Strong mapped coverage across financial-services scope.

Heat - Alignment Pressure

6.00% Low but present pressure from review/benchmark conditions.

Mercury - Trace Density

96.67% Dense symbolic lineage across 21 nodes and 20 links.

Human Review Task

Task status	Required review details
Status	OPEN - human review required before release
Owner role	Compliance / Risk Owner
Review reason	Primary financial scope is covered by mapped current-run evidence, but SEC legal citation linkage remains pending and the strict autonomous-release benchmark gate did not pass.
Review options	APPROVE_RELEASE, REQUEST_REWRITE, BLOCK_RELEASE, REPLACE_WITH_SAFE_TEMPLATE, REQUEST_MORE_EVIDENCE
Response to review	/governed_response
L3 writeback	/decision_authorization_packet/improvement_signals/reviewer_outcome

Improvement Signals and L3 Alignment Feedback

OntoGuard does not stop at scoring. Each governed decision can become a structured improvement signal for retrieval, policy, alignment, training data curation, and future agent behavior.

L3 field	Public sample value
Training signal generated	Yes - review-gated candidate
Signal type	HUMAN_REVIEW_CANDIDATE
Recommended action	ESCALATE_TO_HUMAN
Export status	PENDING_REVIEW
Primary improvement need	legal_citation_linkage
Promotion rule	Gold candidate requires reviewer outcome before training promotion

Business Value / ROI Readiness

Hard-dollar ROI is not calculated in this public sample because customer-specific baseline data was not attached. OntoGuard does not invent savings, avoided losses, revenue impact, or payback without buyer assumptions.

To calculate ROI in a pilot, provide:

- monthly_case_volume
- baseline_review_minutes_per_case
- ontoguard_review_minutes_per_case
- fully_loaded_reviewer_hourly_rate
- baseline_false_approval_rate
- ontoguard_false_approval_rate
- average_loss_per_false_approval
- implementation_cost_usd
- assumptions_source

Without OntoGuard / With OntoGuard

Without OntoGuard	With OntoGuard
AI produces outputs with limited proof of pre-release control.	AI outputs become governed decisions before release.
Human review is ad hoc or inconsistently triggered.	Human review is triggered by explicit Decision API routing and reason codes.
Evidence is hard to reconstruct after the fact.	Evidence, retrieval IDs, symbolic trace, hashes, and audit credentials are preserved.
Failures become one-off incidents.	Failures become structured L3 improvement signals.

Proof Harness, Audit Credential, and Next Step

This public packet keeps the audit posture visible while avoiding customer-specific or private-only material. The original governance run produced a full audit PDF, sellable-lite packet, and schema-rich JSON artifact.

Audit field	Public sample value
Proof harness	ok - reproducible evidence and artifact checks present
Decision API hash	45ebd4258a28cc52a5b6cf06...
Evidence pack hash	ae396a0d1732245fd733641a...
Proof harness manifest hash	4212fc4f934e16e0cb8aeab8...
Report hash	ee6a0e15f25c0e25c9486d99...
Verification status	Offline sample anchor - public demo artifact

Buyer-Safe Governance Mechanism Explanation

- **Cold** indicates grounded compliance coverage and evidence coverage.
- **Heat** indicates anomaly, alignment, or release-pressure signals requiring review.
- **Mercury** indicates symbolic trace and lineage density.
- **Uncertainty** captures retrieval, entailment, toxicity, and agent-disagreement risk.
- **Proof Harness** gives reproducible evidence that the run produced auditable coverage, provenance, and artifact checks.

Next buyer step: run one pilot workflow with customer baseline KPIs and compare governed vs. ungoverned release outcomes.

This document is a public sample generated from a sanitized OntoGuard Decision Authorization Packet. It is intended for product evaluation only and is not legal, compliance, investment, or audit advice.