Vault *for LLM* Security

Pre-deployment security that runs *inside your perimeter.* Zero gaps. Zero compromise.

Full-stack pre-deployment auditing of RAG pipelines, system prompts, tool permissions, and output validation rules — before a single request reaches production. Built for teams who ship AI at scale.

● 100% local processing · zero data egress

0–100	2	0 B
SECURE DEPLOYMENT SCORE PER SCAN	AUDIT MODULES · RAG PIPELINE + GUARDRAIL CHECKER	DATA TRANSMITTED EXTERNALLY

AUDIT COVERAGE

Two modules. Complete pre-deployment coverage.

Vault decomposes your deployment stack into a structured audit graph — catching configuration gaps that manual reviews miss before they reach production.

MODULE 01 · RAG PIPELINE AUDIT

DATA POISONING RISKS
Identifies adversarially crafted or contaminated documents in your retrieval corpus that could manipulate model outputs at query time.

CROSS-TENANT LEAKAGE
Validates namespace isolation and retrieval boundaries to prevent one tenant's data surfacing in another tenant's context window.

MISSING ACCESS CONTROLS
Surfaces unprotected retrieval endpoints, absent authentication gates, and over-permissive index access policies across the pipeline.

MODULE 02 · GUARDRAIL CHECKER

SYSTEM PROMPT HARDENING
Validates that system prompts enforce role boundaries, resist override attempts, and include structural resistance to injection sequences.

TOOL PERMISSIONS
Audits connected tool definitions for implicit over-grants, missing scope restrictions, and unrestricted execution paths to external APIs.

OUTPUT VALIDATION RULES
Checks that output filters, format constraints, and PII redaction layers are correctly configured and cannot be bypassed by adversarial inputs.

SECURE DEPLOYMENT SCORE

One number. Full deployment readiness.

Every Vault scan produces a 0–100 Secure Deployment Score with specific, node-level remediation guidance — not just a pass/fail verdict.

SAMPLE · PRE-REMEDIATION

Data Poisoning Resistance	91/100
Tenant Isolation Integrity	78/100
Access Control Coverage	54/100
System Prompt Hardening	82/100
Tool Permission Scope	41/100
Output Validation Rules	87/100

ANALYSIS PIPELINE

Four stages. Fully on-prem.

Every audit runs inside your infrastructure boundary — no external call, no data residency risk.

01

Ingest & Normalize

RAG configs, system prompts, tool definitions, and output rules ingested and normalized into a canonical audit schema.

02

Graph Construction

Dependency graph built from data flows, access scopes, retrieval namespaces, and instruction boundaries.

03

Policy Analysis

1,800+ security policies matched syntactically and semantically across both RAG pipeline and guardrail surfaces.

04

Score & Remediate

Secure Deployment Score generated with severity-ranked findings and node-level fix templates for each gap.

DEPLOYMENT OPTIONS

Your environment. Your rules.

No architecture changes. No migration. Vault fits the stack you already run.

DEVELOPER · CI GATE

CLI & Pipeline Integration

- — Pre-commit hook support
- — GitHub Actions marketplace action
- — Exit code contracts for CI enforcement
- — JSON score output for custom tooling
- — Monorepo & multi-config support

PLATFORM · ENTERPRISE

Dashboard & API Server

- — Role-based access control (RBAC)
- — SSO via SAML 2.0 and OIDC
- — Webhook alerts — Slack, PagerDuty
- — Audit log trail for compliance
- — Custom policy authoring interface

COMPLIANCE ALIGNMENT | SOC 2 Type II | ISO 27001 | NIST AI RMF | EU AI Act | OWASP LLM Top 10 | GDPR

"Most teams treat go-live as the start of security review. Vault inverts that — every RAG pipeline, guardrail, and tool permission is audited before the first production token is generated. The Secure Deployment Score gives engineering leadership a single, defensible readiness signal."

CORELAYER · SECURITY RESEARCH TEAM · 2026

Request an enterprise evaluation.

A solutions engineer will scope Vault to your environment, integrate it into your CI/CD pipeline, and deliver a full Secure Deployment Score for your prompt infrastructure within five business days.

REQUEST A DEMO →

READ THE DOCS