ENTERPRISE EDITION · 2026

Radar *for LLM* Security

**Prompt security that runs inside your perimeter.
Zero cloud. Zero compromise.**

AST-driven analysis of prompt artifacts — catching injection vectors, role overrides, and unsafe directives before a single token reaches production. Built for teams who ship AI at scale.

100% local processing · zero data egress

97%	<2s	0 B
DETECTION RATE · KNOWN INJECTION PATTERNS	AVG. SCAN TIME PER PROMPT ARTIFACT	DATA TRANSMITTED EXTERNALLY

THREAT DETECTION

Six threat classes. One scanner.

Radar decomposes every prompt into a structured node graph — exposing what human reviewers miss under deadline pressure.

PROMPT INJECTION

Adversarial sequences in user fields, tool outputs, or agent chains that override original directives.

ROLE OVERRIDE

Instructions that redefine model persona, escalate permissions, or suppress safety behaviors mid-session.

UNSAFE TOOL CONFIG

Tool definitions with unrestricted execution paths or implicit permission grants to connected APIs.

EXFILTRATION PATHS

Vectors that leak system prompt text or user PII to external channels via model output.

GUARDRAIL BYPASS

Jailbreak-adjacent phrasing statistically correlated with safety circumvention across model families.

TEMPLATE INJECTION

User-controlled values flowing unsanitized into privileged instruction layers via dynamic templates.

ANALYSIS PIPELINE

Four stages. Fully on-prem.

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

01

Ingest & Tokenize

Raw prompts, templates, and tool configs normalized into a canonical token stream.

02

AST Construction

Hierarchical tree built from instruction scope, role declarations, and data flow edges.

03

Pattern Analysis

2,400+ signatures matched syntactically and semantically via embedding distance.

04

Report & Remediate

Findings ranked by severity with node-level provenance and fix templates.

TECHNICAL SPECIFICATION

Built for *enterprise environments*

DEPLOYMENT Self-hosted · Docker Kubernetes · bare metal	NETWORK POSTURE Air-gap capable Zero mandatory egress	CI/CD INTEGRATION GitHub Actions · GitLab CI Jenkins · Buildkite · API
SIGNATURE COVERAGE 2,400+ patterns Updated quarterly	MODEL COMPATIBILITY OpenAI · Anthropic · Gemini Mistral · open-weight	AUDIT EXPORTS SARIF · JSON · PDF SIEM-ready syslog

DEPLOYMENT OPTIONS

Your environment. Your rules.

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

DEVELOPER · CI GATE

CLI & Pipeline Integration

Drop-in binary for pre-commit hooks and CI gates. Block high-severity findings before staging.

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

PLATFORM · ENTERPRISE

Dashboard & API Server

Centralized findings with RBAC, trend analytics, and REST/gRPC API for full programmatic control.

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

COMPLIANCE ALIGNMENT

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

"The attack surface of an LLM lives in its instructions, not its weights. *Most teams have no systematic way to audit that surface.* Radar is the first tool that treats prompts with the rigour we apply to application code."

CoreLayer · Security Research Team · 2026

Request an enterprise evaluation.

A solutions engineer will scope Radar to your environment, wire it into your CI/CD pipeline, and deliver a risk assessment of your prompt infrastructure within five business days.

REQUEST A DEMO →

READ THE DOCS