

AI LIFECYCLE SECURITY

CORELAYER

SecureAI

A unified AI security platform that defends every phase of your AI lifecycle from the first line of code to the end user's keyboard.

PLATFORM

CoreLayer SecureAI

PHASES

Five (Build → End User)

DEPLOYMENT

Cloud / On-Premise

01

T H E P R O B L E M

AI systems are under attack

*"Traditional WAFs and SIEMs were not built for LLM attack surfaces.
Your AI needs purpose-built defense."*

Prompt injection, jailbreaks, data poisoning, cross-tenant leakage these threats emerge at every phase of the AI lifecycle. Point-in-time tools and perimeter defenses leave entire attack surfaces unmonitored.

74%

of enterprises report AI-specific security incidents

3.2×

faster attacker innovation vs. AI defense adoption

\$4.8M

average cost of an AI data breach

0

legacy tools purpose-built for LLM attack surfaces

02

THE PLATFORM

One platform.

Five phases.

Full lifecycle coverage - intelligence shared bi-directionally across every phase.

CORELAYER · 5 PHASES · 1 UNIFIED INTELLIGENCE LAYER · CONTINUOUS DEFENSE

01

BUILD

CoreLayer Radar

AST prompt scanning before
code ships

02

TEST

CoreLayer Striker

2,000+ adversarial payloads in
CI/CD

03

VALIDATE

CoreLayer Vault

Pre-deployment RAG audit &
Secure Score

04

RUNTIME

CoreLayer Shield

Real-time behavioral defense,
<10ms

05

END USER

CoreLayer SecureAgent

PII & credential masking at the
edge

03

PHASE 01 — BUILD

CoreLayer Radar

Stop vulnerabilities before a single line of code ships.

AST-Style Prompt Parsing

Deep structural analysis detects injection surfaces manual review misses entirely

Role Override Detection

Identifies unsafe instruction patterns, privilege escalation paths, and confusion vectors

Tool Configuration Audit

Scans permissions, parameter handling, and execution chains for exploitable configs

100% Local Execution

Zero cloud upload — all analysis on-premise, your prompts never leave your environment

KEY OUTCOMES

100%

LOCAL ANALYSIS

Zero cloud upload of proprietary prompts

0

CLOUD UPLOADS

All scanning runs entirely on-premise

<30s

SCAN TIME

Full prompt template scan per run

04

PHASE 02 TEST

CoreLayer Striker

Adversarial pressure-testing before attackers get the chance.

2,000+

REAL-WORLD ATTACK PAYLOADS

Continuously updated threat library

15

ATTACK CATEGORIES

Across the full kill chain

CI/CD NATIVE

- GitHub Actions
- GitLab CI
- Jenkins
- CircleCI
- Azure DevOps

CoreLayer Vault

No AI goes live without a passing score.

RAG PIPELINE	Data Poisoning Audit	Detects manipulated documents injected into your knowledge base
ISOLATION	Cross-Tenant Leakage	Verifies tenant data boundaries and isolation controls
ACCESS	Missing Access Controls	Identifies permission gaps and privilege escalation paths
HARDENING	System Prompt Validation	Guardrail strength check and injection resistance scoring
PERMISSIONS	Tool Permission Review	Least-privilege enforcement across all tool integrations
OUTPUT	Output Validation Rules	Completeness check on response filters and content policies

SECURE DEPLOYMENT SCORE

0 100
FAILING PASSING

Every scan produces a 0-to-100 score with specific remediation guidance for each gap found — so your team knows exactly what to fix before deployment.

Integrates with: AWS Config · Azure Security Center · GCP SCC

CoreLayer Shield

Three coordinated engines. Continuous defense. Zero blind spots.

CORELAYER SHIELD · 3 ENGINES · SHARED REAL-TIME INTELLIGENCE · <10ms POLICY EVALUATION

LCAC

Layered Context Access Control

Controls precisely what context the model can access. Enforces information boundaries, prevents context bleeding, and applies dynamic access policies based on user role and request sensitivity.

LBF

Live Behavioral Fingerprinting

Monitors model behavior in real time to detect anomalies and zero-day jailbreaks. Builds a behavioral baseline and flags deviations that signature-based systems would completely miss.

CBE

Constrained Behavior Enforcement

Enforces hard limits on what the model can do. Caps tool chaining depth, sets execution ceilings, and applies kill-switches the moment a constraint boundary is crossed.

CoreLayer SecureAgent

Protecting the people interacting with your AI.

D A T A I N T E R C E P T E D A N D M A S K E D

- API Keys & Tokens
- Passwords & Credentials
- PII Names, Addresses
- Aadhaar & PAN Numbers
- Credit Card Numbers
- UPI IDs
- Email & Phone Numbers

P R I V A C Y G U A R A N T E E S

- ✓ **Local masking only**
Processing happens on the user's device before any data is transmitted to an LLM
- ✓ **Zero data collection**
No sensitive data is ever collected, stored, or logged by CoreLayer systems
- ✓ **Zero transmission**
Only masked prompts reach the model — raw sensitive content is never sent
- ✓ **No third-party exposure**
Complete privacy chain from user input through AI response and back

Every attack makes the platform smarter.

Bi-directional intelligence flows between all five phases a closed-loop defense system.

Build → Test

Vulnerability found in prompt scan automatically generates a Test attack case in Striker's library

Test → Runtime

Exploit confirmed in CI/CD pipeline triggers a new Runtime enforcement rule in Shield

Runtime → Build

Anomaly detected in production feeds back into Build scanning rules for future deployments

Validate → Runtime

Pre-deployment score gap configures Runtime enforcement threshold automatically

Runtime → SecureAgent

Behavioral pattern flagged at runtime updates SecureAgent interception rules at the edge

Purpose-built for the LLM threat landscape.

ASPECT	CORELAYER SECUREAI	LEGACY / TRADITIONAL TOOLS
Coverage	✓ Full lifecycle — Build through End User	X Point-in-time or perimeter only
AI Awareness	✓ LLM-native attack detection and response	X Signature-based, not LLM-aware
Feedback Loop	✓ Bi-directional phase intelligence sharing	X Siloed tools, no cross-signal sharing
Privacy	✓ 100% local option — zero cloud upload	X Data uploaded to vendor cloud
CI/CD Integration	✓ Native — blocks deployment on failure	X Manual or post-deployment only

CORELAYER

Continuous Intelligence Platform

What's the next move?

CoreLayer SecureAI is available for enterprise pilots now. The immediate priority is onboarding 3–5 design partners to validate the full five-phase deployment model.

[REQUEST A SECURITY BRIEFING →](#)

WEBSITE

corelayersecurity.ai

ENTERPRISE SALES

sales@corelayersecurity.ai

SECURITY BRIEFING

security@corelayersecurity.ai
