SECURE THE AI YOUR BUSINESS NOW DEPENDS ON

Your AI is already deployed. *Is it secure?*

Illumant helps CISOs and engineering leaders secure the AI systems their organizations have quietly come to depend on — from shadow Copilot deployments to customer-facing chatbots to autonomous agents with production access.

THE ADOPTION GAP

90 / 5

of organizations deploy AI, but only 5% feel confident in their AI security.

Industry analysis, 2025

ACCESS CONTROLS MISSING

97%

of AI-related breaches involved systems lacking proper access controls.

IBM Cost of a Data Breach, 2025

SHADOW AI COST

$670K

higher breach cost when shadow AI is involved — 1 in 5 orgs affected.

IBM Cost of a Data Breach, 2025

Five practices. *Full coverage.*

01 · SERVICES

01

AI Governance & Risk Advisory

Top-down strategy work: risk frameworks, policies, inventories, and regulatory alignment your board and auditors expect.

`Risk Assessment` `Policy Development` `EU AI Act` `Vendor Review`

02

AI Implementation Assessment

Evaluate how securely AI is deployed inside your environment — Copilot oversharing, over-permissioned agents, shadow AI.

`Oversharing Assessment` `Access Review` `Shadow AI Discovery`

03

AI Product Security Testing

Adversarial testing of AI-powered products — chatbots, assistants, and RAG pipelines. Structured around the OWASP LLM Top 10 (2025).

`OWASP LLM Top 10` `Prompt Injection` `RAG Testing`

04

Agentic AI Security Testing

Testing for AI that takes action — tool-using agents, autonomous systems, MCP integrations. Built on the OWASP Agentic Top 10 (2026).

`OWASP ASI Top 10` `Goal Hijack` `Tool Misuse` `MCP Security`

05

AI Red Team

Full-scope adversary simulation targeting AI systems, pipelines, and the humans who trust them.

`Attack Simulation` `AI-Assisted SE` `Multi-Step Chains`

FRAMEWORK

Built around *OWASP* — LLM Top 10 & Agentic Top 10.

Every engagement produces a deliverable your auditors, boards, and regulators recognize. LLM testing follows OWASP's LLM Top 10 (2025); agentic testing follows the new Agentic Applications Top 10 (2026).

LLM01
Prompt Injection
Direct and indirect instruction override.

LLM02
Sensitive Information Disclosure
Model leaks PII, credentials, or proprietary data.

LLM03
Supply Chain Vulnerabilities
Compromised models, datasets, or components.

LLM04
Data and Model Poisoning
Malicious training data manipulating behavior.

LLM05
Improper Output Handling
Unvalidated output executed downstream — XSS, SSRF, RCE.

LLM06
Excessive Agency
Agents with too much functionality or autonomy.

LLM07 NEW
System Prompt Leakage
Exposure of business logic or credentials in prompts.

LLM08 NEW
Vector & Embedding Weaknesses
RAG and vector database attacks.

LLM09
Misinformation
Hallucinations and overreliance on unverified output.

LLM10 NEW
Unbounded Consumption
Resource exhaustion, wallet attacks, model theft.

Why *Illumant.*

03 · DIFFERENTIATION

01

Operators, not academics.

Led by practitioners with deep offensive backgrounds — the people who break networks also break AI.

02

Standards-mapped.

OWASP LLM Top 10, NIST AI RMF, EU AI Act, ISO 42001 — deliverables your auditors and board recognize.

03

One firm, full coverage.

From strategy to red team in a single engagement — no vendor sprawl, one consistent voice.

Start with a scoping call.

A 30-minute conversation with one of our AI security leads. We'll map your AI footprint and identify where the meaningful risks live — no generic sales deck, no commitment.

BOOK A CONSULTATION

info@illumant.com