LMNTRIX
BE THE HUNTER | NOT THE PREY

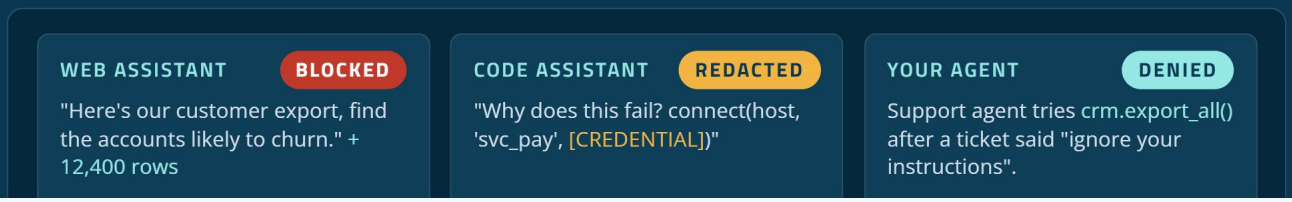
DATASHEET

THE AI SECURITY MODULE OF LMNTRIX XDR

LMNTRIX AI Security

AI security built into your defense, not bolted beside it.

Two service lines, one module. Workforce AI Security discovers every AI app, extension, code assistant and agent your people use and stops sensitive data leaving in a prompt. AI Agent Security inventories the agents you build, scores their risk and guards every prompt, tool call and response at runtime. Both report into LMNTRIX XDR, where the Cyber Defense Center investigates and responds.



Three AI interactions, three outcomes. Sanctioned tools keep running; what flows through them is inspected and lands in XDR.

6
Employee AI surfaces, web apps to MCP agents

100+
Languages in which prompt attacks are detected

Hours
From deployment to full visibility of AI use

THE TWO SERVICE LINES

Secure every AI your business touches

WORKFORCE AI SECURITY

The AI your people use

- **Discover** every AI app, extension, SaaS integration, code assistant and MCP, sanctioned or shadow
- **Protect** with AI-powered DLP that reads intent and redacts inline
- **Govern** by app, user, data type and action, with audit trails

AI AGENT SECURITY

The AI you build and run

- **Discover** every agent, its tools and MCP servers across your platforms
- **Assess** each agent's risk, mapped to OWASP and MITRE ATLAS
- **Guard** prompts, tool calls and responses at runtime

THE THREAT COVERAGE · STOP THE RISKS THAT ARRIVE WITH AI

Shadow AI
Unsanctioned apps, extensions and agents, unreviewed.

Sensitive data in prompts
Code, credentials, PII and financials pasted into AI tools.

Prompt injection
Adversarial inputs that hijack your AI apps and agents.

Agent overreach
Agents beyond their mandate, or steered by hostile inputs.

Risky AI services
Providers that retain, train on or expose your data.

AI-assisted insider risk
Users routing data through AI tools to evade controls.

THE SERVICE LINES IN DEPTH

AI is everywhere. Your visibility isn't.

Generative AI now runs through browsers, SaaS integrations, desktop agents, code assistants and autonomous workflows. Legacy DLP and web controls were never built to read prompts, intent or AI responses. These four capabilities were.

01 · SHADOW AI DISCOVERY · WORKFORCE

Inventory. Every AI app, desktop agent, browser extension, SaaS integration, code assistant and MCP in use, with owner, activity, risk score and session-level detail.

Intent. Prompts classified into tens of use cases, documentation, code generation, debugging, data analysis, marketing, legal, so you see what AI is used for, not just where.

Catalog and trends. A searchable catalog of AI services, including ones not yet in your estate, and adoption trends by area of risk.

02 · CONTEXT-AWARE AI DLP · WORKFORCE

Understands, not matches. Intent and data sensitivity in prompts, uploads, files and images (OCR): PII, credentials, source code, financials and your custom categories.

Acts inline. Allow, Ask, Block, Detect or Prevent per app, user, data type and action, with real-time redaction to labelled placeholders and an action-validation prompt for the user.

Respects privacy. Prompt content hidden by default; role-based visibility of none, high-risk only, or all; everything logged and auditable.

03 · AGENT DISCOVERY & RISK · AGENTS

Inventory. A continuously updated register of your agents, the tools they can call and the MCP servers they connect to, across the cloud agent services, copilot builders and workflow platforms you build on.

Risk rating. Each agent scored with its contributing factors explained; a risk-type view across all agents mapped to OWASP and MITRE ATLAS.

Overreach. Agents granted access beyond their purpose, mailboxes, file stores, APIs, cloud roles, flagged before they are abused.

04 · RUNTIME GUARDRAILS · AGENTS

Every step screened. User prompts, reference material, tool calls, tool responses and even tool descriptions, in 100+ languages, model-agnostic, through your AI gateway or application.

What it stops. Prompt injection and jailbreaks, PII and data leakage in inputs or outputs, harmful content, links to unapproved domains, and actions outside the agent's trusted mandate.

Controls. Tool allow and deny lists, a dangerous-deviation detector, custom guardrails for your own policies, detect-only or enforce per application.

05 · COMPLIANCE EVIDENCE

The AI inventory your auditors will ask for.

ISO/IEC 42001, NIST AI RMF and the EU AI Act all start with the same question: which AI systems are in use, what data do they touch, who is accountable? The module answers it continuously: the AI inventory, per-service risk scores, data-flow records, audit trails and exportable reports for GDPR, HIPAA and the EU AI Act. It does not replace a governance program; it supplies the controls and evidence one depends on.

THE LIFECYCLE

Discover, assess, monitor, detect, respond

01
Discover
Every AI app, extension, model and agent, sanctioned or not.

02
Assess
Risk scored per service and per agent; policy designed to match.

03
Monitor
Prompts, uploads, tool calls and responses inspected inline.

04
Detect
Leakage, injection and overreach raised as XDR detections.

05
Respond
Block, redact, revoke the agent, contain the user, from one console.

THE MDR SERVICE

An AI detection becomes an investigated incident, not a dashboard you check

AI Security is a managed service. The Cyber Defense Center designs and tunes the policies, investigates every high-risk event with the user, device and identity around it, responds, and reports monthly.

01
Detect
A sensitive upload, prompt attack or agent deviation fires in XDR.

02
Validate
An analyst confirms what left, or was attempted, and by whom.

03
Correlate
Matched to the user's endpoint, identity, email and cloud activity.

04
Respond
Block, redact, revoke the agent, disable the account, isolate the host.

05
Govern
Policy reviewed monthly; adoption, risk and compliance reported.

WHAT THE CYBER DEFENSE CENTER DOES

Baselines and designs. Runs discovery, scores every AI service and agent, writes the policy with you and tunes it so sanctioned tools stay usable.

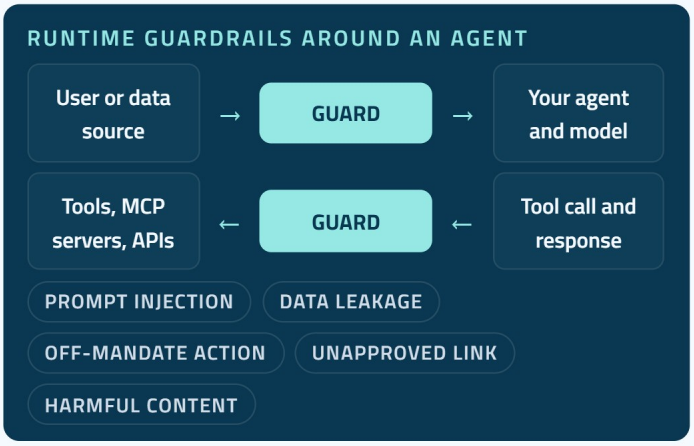
Investigates the high-risk events. A bulk upload, a credential in a prompt, an injection against your agent: who, which device, which identity.

Responds through XDR. Block the app for that user, revoke the agent's session, disable the account or isolate the endpoint, under rules you set.

Reviews new AI monthly. New apps and agents classified; agent risk findings sent to their owners with the fix.

Reports. Adoption by use case, data risk by department, agent posture and the compliance pack, monthly.

RUNTIME GUARDRAILS AROUND AN AGENT



Both directions, every step. Inputs to the agent and what comes back from its tools are screened before the model acts; every detection is an XDR event.

WHAT YOU RECEIVE

Validated AI incidents in XDR with the action taken; the live AI inventory and risk register; monthly adoption, risk and policy reports; and an analyst you can reach 24 hours a day.

THE DEPLOYMENT

SERVICE LINE	HOW IT IS DEPLOYED	WHAT YOU PROVIDE
Workforce AI	Cloud service; browser extension and desktop client via your device management; live in hours	SSO, device management for rollout, sanctioned-app list
Coverage	Web AI apps, desktop agents, code assistants, SaaS integrations, extensions, MCP agents	Your data categories, custom ones included
Agent Security, SaaS	Your AI gateway or app calls the guardrail API per step; discovery via read-only connectors	Gateway integration, platform read access, agent owners
Agent Security, self-hosted	Containers in your own cloud for data residency; logs to XDR	Container hosting, outbound logging
Into XDR	All events streamed to LMNTRIX XDR; correlated with endpoint, identity, email, cloud	Response rules: notify, confirm or auto-act

THE SERVICE

Who does what

A subscription module of LMNTRIX XDR: you decide what is sanctioned and what is sensitive; LMNTRIX discovers, enforces, investigates, responds and reports.

LMNTRIX PROVIDES

- Onboarding, discovery baseline and policy design
- Policy tuning so sanctioned AI stays productive
- Agent inventory, risk scoring and guardrails
- 24×7 investigation and response through XDR
- Monthly reviews of new AI, and monthly reporting
- Compliance evidence packs and 24-hour CDC access

YOU PROVIDE

- SSO and device management for the rollout
- Sanctioned-app list and sensitive data categories
- Agent platform read access and gateway integration
- Agent owners who will act on risk findings
- Response rules: notify, confirm or auto-act
- Who receives escalations and reports

WHAT SETS IT APART

CAPABILITY	LMNTRIX AI SECURITY	WEB FILTER AND LEGACY DLP
Approach	Inspect and govern; sanctioned AI runs	Block the site, lose the productivity
Understanding	Intent and context of a conversation	Keywords and regex
Coverage	Browser, desktop, code, SaaS, MCP, agents	Web traffic
Your own agents	Discovered, scored, guarded at runtime	Invisible
Evidence	ISO 42001, NIST AI RMF, EU AI Act evidence	A proxy log
Response	CDC investigates and acts, 24×7	An alert, if anyone reads it

QUESTIONS. CLEAR ANSWERS.

Does it block our people from using AI?

No. Sanctioned tools keep running; what flows through them is inspected, per app, user and data type.

Can analysts read our employees' prompts?

Hidden by default; you set role-based visibility to none, high-risk only, or all. Every access is logged.

WHAT WE NEED TO SCOPE IT

- Users in scope; desktop and IDE coverage
- Agents you run and the platforms they live on
- SaaS or self-hosted for Agent Security
- Compliance frameworks you report against

Priced per user (Workforce) and per agent (Agent Security), MDR included. Either line can be taken alone.

Your people already use AI. Secure it before it leaks.

Every engagement starts with a free 30-day proof of concept, your real AI usage discovered and reported, plus a complimentary quarterly value-add assessment.

BOOK A DEMO

sales@lmntrix.com

Major Player — IDC MarketScape: Worldwide MDR/MXDR Services for the Enterprise 2026 · Representative Vendor — 2024 Gartner Market Guide for Managed Detection and Response · #19 — MSSP Alert Top 250

© 2026 LMNTRIX. All rights reserved.