LTM

BROCHURE

Cisco AI Defense

Security Built for AI Applications in Production

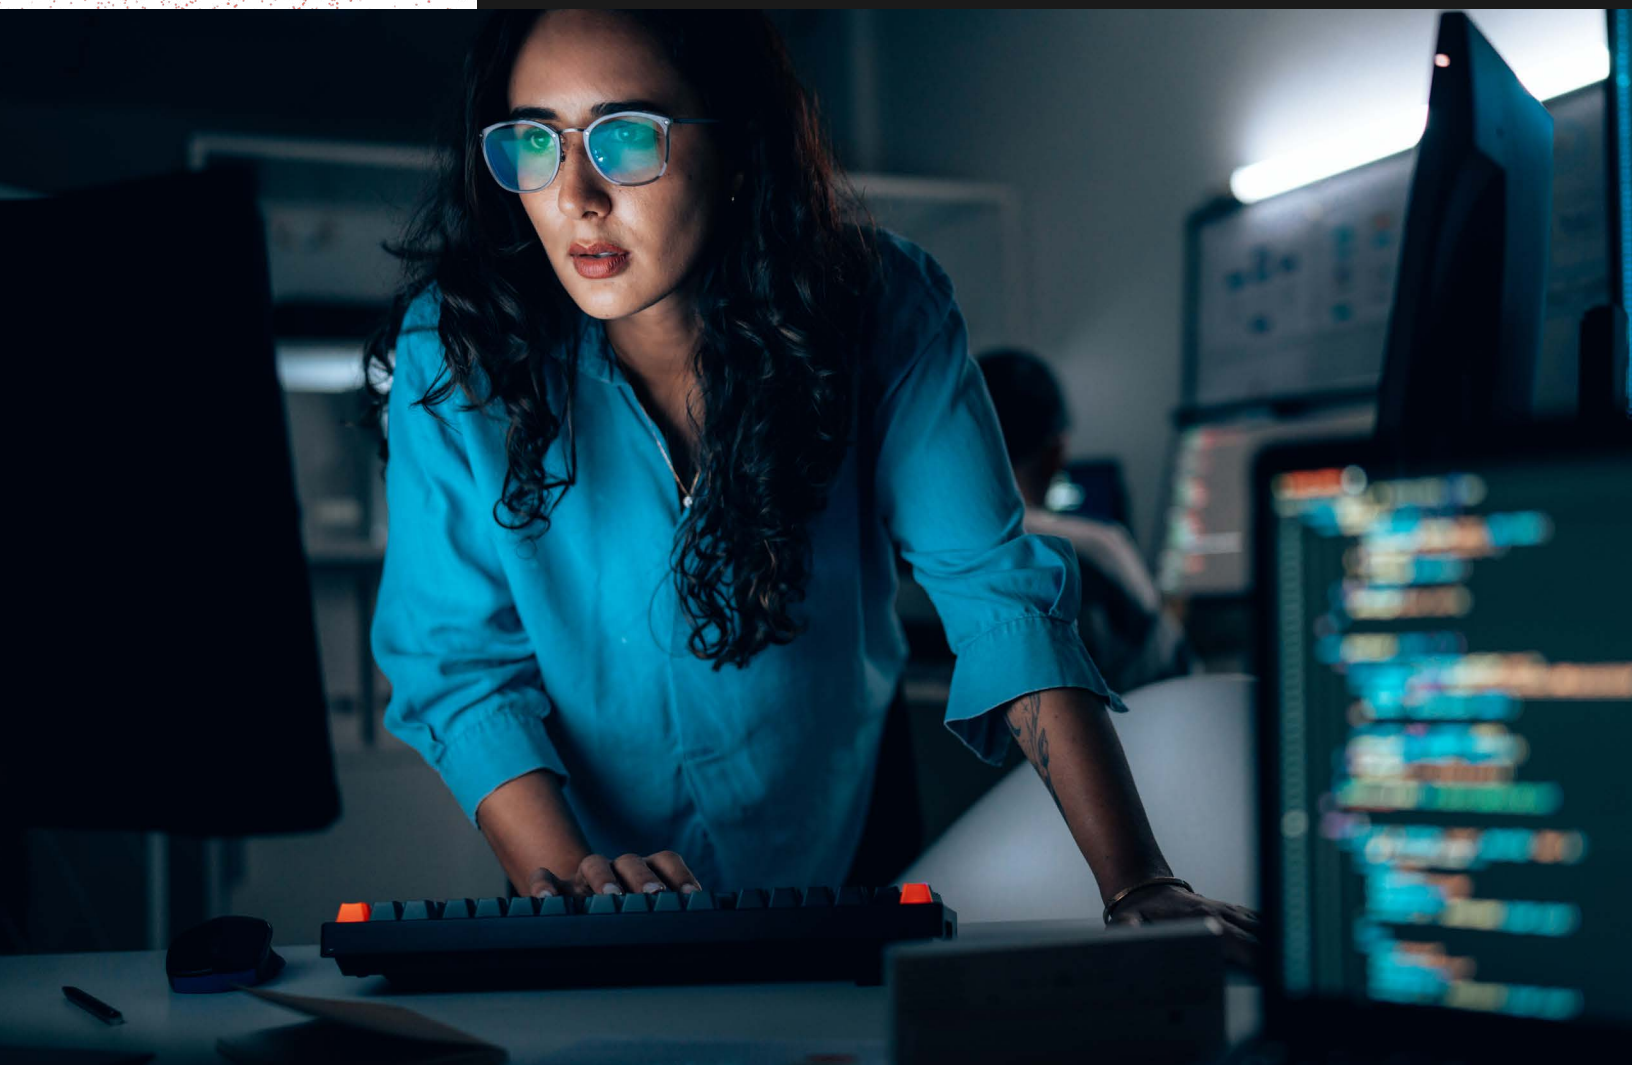

Introduction

AI is moving from pilot to production. Enterprises are deploying models into customer-facing applications, internal copilots, and agentic workflows that act on real data and decisions. This shift creates a new attack surface that traditional cybersecurity tools were not designed to defend.

Models behave in non-deterministic ways, datasets and open-source components move into production with limited scrutiny, and guardrails vary across models. As enterprises scale AI adoption, they must not simply react to emerging threats, but Outcreate the market with security designed for AI-native operations.

Business Problems: Why Traditional Security Fails on AI

Security and AI teams are facing new risks that legacy controls cannot fully address.

Non-deterministic Risk Vectors
Generative models and agents can produce different outputs for the same input, creating risks such as prompt injection, data leakage, unsafe responses, and supply-chain compromise.

Risk Grows with Capability
As organizations move from chatbots to RAG and agentic systems, applications gain more autonomy and access to sensitive data, increasing exposure to attackers.

Manual Testing does not Scale
Manual red teaming takes 7–15 weeks for a single model and must be repeated after every fine-tune, slowing production deployment.

Built-in Guardrails are Inconsistent
Fine-tuned variants are 3x more susceptible to jailbreaks and 22x more likely to produce harmful responses, according to Cisco research.

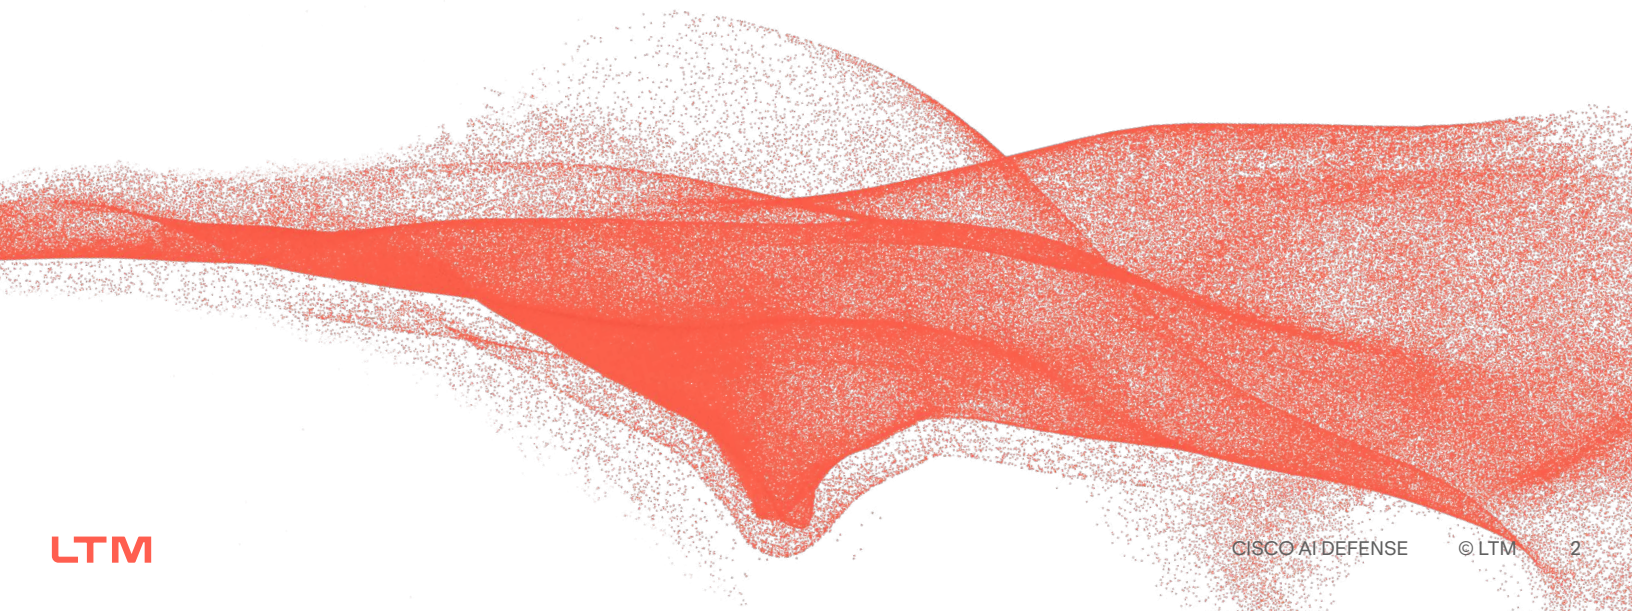

AI Security by the Numbers

Independent research and Cisco's AI security testing show the scale of enterprise exposure.

The AI Security Reality*

86%
of enterprises experienced an AI-related security incident in the past 12 months

45%
have the expertise for comprehensive AI security assessments

41%
lack mature controls on AI training data

7–15
weeks are required for a manual end-to-end model red team

Why Traditional Security Breaks on AI*

Cisco research quantifies what changes once models hit the real world.

3x
more models susceptible to jailbreak instructions after fine-tuning

22x
more likely to produce a harmful response after fine-tuning

100%
attack success rate observed against DeepSeek in Cisco testing

0.01%
of a dataset is enough to poison a model

Multi-turn attacks succeed 2x to 10x more often than single-turn baselines, with success rates between 25.8% and 92.8%

Cisco AI Defense Coverage*

What algorithmic red teaming and runtime guardrails actually cover out of the box.

200+
safety and security subcategories tested through algorithmic red teaming

45+
prompt injection attack techniques covered out of the box

50+
safety categories spanning violence, bias, and public harm

20+
information security categories including data leakage and code execution

New threats are added within three days of disclosure through Cisco Talos and Cisco AI security research, with guardrails mapped to OWASP Top 10 for LLMs, MITRE ATLAS, and the NIST adversarial ML taxonomy

Our Solution: Discover, Detect, and Protect AI Applications

Cisco AI Defense protects AI applications across the full lifecycle, from development to runtime. Operating at the network layer, it applies security enforcement without modifying the application or model.

The platform follows a five-step framework.

Discover AI assets	**Detect** supply chain risk	**Detect** model vulnerabilities	**Protect** at runtime	**Operate** and govern
Inventory models, agents, and connected data sources across cloud and on-prem environments	Scan model files, repositories, and MCP servers before they enter production	Run algorithmic red teaming across more than 200 safety and security categories	Apply bi-directional guardrails that block malicious prompts and unsafe outputs	Continuously monitor, update, and align protections with OWASP, MITRE ATLAS, and NIST standards

Cisco AI Defense

Discovery	**Detection**	**Protection**
Know where AI exists in your environment	Find vulnerabilities before attackers do	Stop attacks at runtime for every request

Discover AI Assets

AI Cloud visibility builds a live inventory of models, agents, and data sources running in your environment.

- Map AI assets across cloud, hybrid, and on-prem footprints
- Surface usage context and existing controls to rank risk
- Provide governance, security, and AI teams with a shared view

Detect Vulnerabilities

AI supply chain risk management and AI model and application validation test the assets that matter most, before and after release.

- Scan repositories for code execution and suspicious imports
- Inspect MCP servers and identify tool-poisoning attacks
- Run algorithmic red teaming across more than 200 safety and security subcategories
- Generate reports mapped to OWASP Top 10 for LLMs, MITRE ATLAS, and NIST

Protect at Runtime

AI runtime protection enforces enterprise guardrails in line with traffic, so attacks are blocked as they happen rather than discovered after the fact.

- Block malicious prompts and unsafe responses with bi-directional guardrails
- Cover security, privacy, and safety threats including prompt injection and PII exposure
- Mitigate rogue agents and MCP tool misuse
- Refresh protections continuously through Cisco Talos threat intelligence feeds

New threats are added within three days of disclosure through Cisco Talos and Cisco AI security research, with guardrails mapped to OWASP Top 10 for LLMs, MITRE ATLAS, and the NIST adversarial ML taxonomy

Outcreate Operational AI Security at Scale

The following benefits help enterprises secure AI with greater speed, consistency, and control:

- Faster AI security validation through algorithmic red teaming completed in minutes instead of weeks
- Continuous runtime protection across AI applications, agents, and workflows
- Consistent enforcement without modifying the application or model
- Centralized governance aligned to OWASP, MITRE ATLAS, and NIST frameworks
- Flexible deployment across SaaS, customer VPC, and on-prem AI PODs
- Operationalized AI risk management through integrations with Splunk Enterprise Security, ServiceNow AI Control Tower, and Safe Security

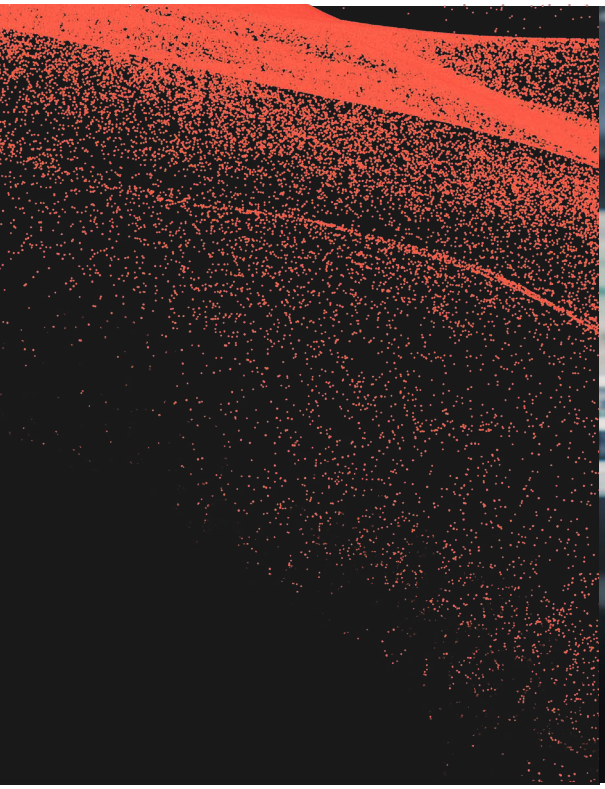

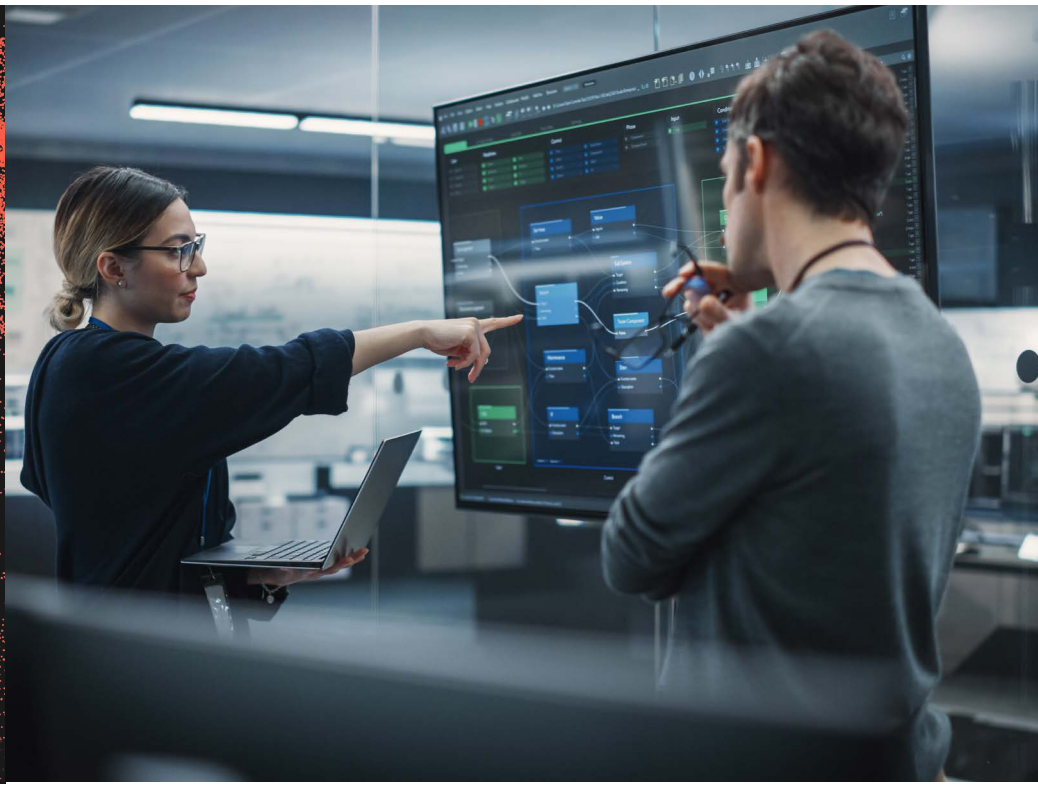

Why Cisco AI Defense with LTM

LTM follows a structured approach that takes enterprises from AI risk assessment to production rollout with confidence, speed, and control. As a Business Creativity partner, LTM brings the very best of human insights and intelligent systems together to help organizations operationalize AI security across cloud, hybrid, and on-prem environments.

Assess
Inventory AI assets, baseline controls, and identify gaps against OWASP, MITRE, and NIST

Strategize and Design
Define the target security architecture and application guardrails

Build and Deploy
Deploy Cisco AI Defense across SaaS, customer VPC, or on-prem AI PODs integrated with the SOC

Operate and Improve
Continuously validate models, tune guardrails, and align protections to new threats and standards

Each phase produces a tangible artifact including a risk register, target design, live deployment, and continuous validation plan.

As AI reshapes enterprise operations, organizations must move beyond reactive security models and build protection into every stage of AI execution. With Cisco AI Defense and LTM, enterprises can secure AI with the visibility, governance, and control needed to not just keep pace with the market, but Outcreate it.

Start building on a security layer designed for AI.

Connect with us at ***cybersecurity.pulse@ltm.com***

Reference

1. *2025 Cisco Cybersecurity Readiness Index*, CISCO, 2025: https://newsroom.cisco.com/c/r/newsroom/en/us/a/y2025/m05/cybersecurity-readiness-index-2025.html#blade_introduction

LTM is a global technology services and consulting company and the Business Creativity partner to the world's largest and most disruptive companies. We bring human insights and intelligent systems together to help enterprises across industries rewire their business models, accelerate innovation, and drive AI-centric growth. With our integrated operations, transformation, and business AI services, we design and deliver solutions that create new productivity paradigms and new roads to value. Together with 87,000 employees across 40 countries and our global network of hyperscaler partners, LTM — A Larsen & Toubro company — owns business outcomes for over 700 clients, helping them to not simply outperform the market, but to Outcreate it.

LTM/IND//DS/MAY 2026/SD06