PRIMUS PARTNERS®

Quote by Devroop Dhar, Co-founder & CEO, Primus Partners

Published in COMPUTERWORLD
September 17, 2025

Check Point acquires Lakera to create a unified AI security stack

The acquisition, which brings runtime protection, continuous red teams, and multilingual defenses to Check Point's Infinity platform, comes at a time when companies are facing risks arising from LLMs, agents, and generative AI workloads.

.

Read on: https://www.computerworld.es/article/4058766/check-point-adquiere-lakera-para-crear-una-pila-de-seguridad-de-ia-unificada.html

Article Content- The acquisition, which brings runtime protection, continuous red teaming, and multilingual defenses to Check Point's Infinity platform, comes at a time when companies are facing risks stemming from LLMs, agents, and generative AI workloads.

Check Point has signed an agreement to acquire Lakera, an AI-native security platform specialized in agent-based AI applications. The deal, whose amount has not been disclosed, is expected to close in the fourth quarter of 2025, and is also expected to boost Check Point's AI security stack with the purpose of strengthening corporate defenses as companies accelerate AI adoption.

Over the past 18 months, companies worldwide have rushed to incorporate large language models, generative AI, and autonomous agents into their core workflows. This change is driving innovation but has also expanded the attack surface.

In statements to CSO, Nataly Kremer, Chief Product Officer at Check Point Software Technologies, explained that "customers are already reporting risks such as command

injection attacks that manipulate outputs, confidential data leaks through LLMs and agents, model manipulation and poisoning, and new vulnerabilities introduced by collaboration between multiple agents and autonomous decision-making."
Kremer argued that traditional cyber defenses were not designed to address these model-specific, real-time risks. Therefore, "Lakera addresses this gap with AI-native runtime protection, continuous red teaming, and multilingual defenses. Proven at scale, its technology already protects advanced enterprise AI deployments using its adversarial engine Gandalf, which leverages more than 80 million attack patterns."
According to Check Point, Lakera has detection rates above 98%, latency below 50ms, and false positives under 0.5% to protect AI workloads without affecting speed or accuracy.

Integration of Lakera into the Infinity platform

Check Point's approach to this change in offerings includes generative AI security protection, SaaS and API security, advanced data loss prevention, and machine learning–based defenses across applications, cloud, and endpoints.
What the company is doing is integrating Lakera's technology directly into its Check Point Infinity architecture. Consequently, the first integrations will appear in Check Point CloudGuard WAF, which protects AI-enabled applications, and Check Point GenAI Protect, which does the same for user traffic to GenAI applications.
This led Kremer to state that "customers will immediately see real-time protection for LLMs and agents against attacks and leaks, continuous red teaming intelligence applied across all their AI deployments, high detection accuracy with minimal false positives, and faster protection times, leveraging the same Infinity Portal they already use for firewall, endpoint, cloud, and email security. Over time, these AI-native defenses will extend across the entire Infinity portfolio."
Current Check Point Infinity customers will perceive AI security as an additional capability within their existing deployments, while new customers will be able to adopt the unified end-to-end AI security stack directly through Infinity, as Check Point has confirmed. Being API-based, Lakera's platform is delivered via the cloud (with on-premises options) and also ensures that protection time is almost immediate.

Closing a critical gap

Experts have considered this acquisition significant, as it does not merely add another tool to the stack. Amit Jaju, Senior Managing Director at Ankura Consulting, explained that "this acquisition closes a real gap by adding AI-native runtime protection barriers and continuous red teaming to Check Point's stack," adding: "Customers can now protect LLMs and agents alongside their existing network, cloud, and endpoint controls."
As an immediate advantage, all of them have highlighted the reduction of integration friction and the unification of policies and telemetry for AI use cases that are already spreading across enterprises—especially when agents have access to tools and handle sensitive data. To do this, companies must immediately treat AI applications and agents as first-class assets, implement runtime protection barriers, apply continuous red teaming, and integrate AI telemetry with existing policies.
While dedicated AI security is still emerging, it is believed that demand will grow enormously as more companies adopt large language models, agents, RAG (retrieval-augmented generation) systems, etc.
According to Devroop Dhar, co-founder and managing director of Primus Partners, "demand exists, especially in sectors with regulated data (finance, healthcare), technology companies creating AI-based products, and large-scale enterprises that cannot afford unknown unknowns.

The greatest demand comes from early adopters of this technology, such as cloud providers, AI SaaS companies, large enterprises with mature security programs, and organizations developing AI in production (not just in pilot phase)."
Check Point also believes there is strong demand in the manufacturing and logistics sectors, where agent-based AI is used for automation; and in the public sector and critical infrastructure, where trust and regulatory compliance are fundamental.

Vendors compete to secure AI

Security companies are rapidly reshaping their portfolios to address the risks posed by artificial intelligence. Jaju acknowledged the emergence of two clear themes: the platform and depth in LLM-specific risk.

In his view, "buyers want AI security integrated into existing suites for unified visibility and response, while vendors are doubling down on controls tailored to LLMs and agent workflows that cover prompt injection, jailbreaks, unsafe tool use, data exfiltration, and supply chain risks."

That's why vendors are going beyond traditional network and endpoint protection to secure AI models, agents, and applications. Consequently, many companies are choosing to acquire firms that help them fill these gaps instead of building solutions from scratch.

As Dhar added: "SentinelOne acquired Prompt Security, Cato Networks bought Aim, and other companies are also heavily investing in startups that protect AI models, APIs, and training data. The generative AI cybersecurity market is expected to grow rapidly, given how quickly adoption is increasing and how new attack vectors will continue to emerge."