Dropbox Blocks Prompt Injection Attempts While Preserving Exceptional, Real-Time User Experience

INDUSTRY:
Technology Solutions

HEADQUARTERS:
San Francisco, CA

COMPANY SIZE:
700 million registered users

OVERVIEW

Dropbox is one place to keep life organized and keep work moving. With more than 700 million registered users across 180 countries, Dropbox is on a mission to design a more enlightened way of working.

OUR SOLUTION

AI Agent Security (Formerly Lakera Guard)

CHALLENGE

Dropbox, the global provider for cloud storage, file sharing, and collaboration, has established itself as a pioneer in deploying generative AI (GenAI) capabilities at enterprise scale. Dropbox has been actively developing GenAI services to enhance product experiences and content management capabilities for its 700 million registered users worldwide. As Dropbox accelerated the development of its GenAI services - rolling out features including intelligent search, document summarization, chat interfaces, and automated content analysis — the security team faced a critical challenge: how to comprehensively protect and monitor all GenAI applications within its self-hosted, custom-built software architecture without compromising performance.

> After evaluating several AI security vendors and internal options, Check Point AI Agent Security quickly stood out in terms of its operational performance and ultra-low latency.
>
> Adrian Wood, security engineer, Dropbox

The Dropbox team's primary security concern was emerging AI-specific threats that traditional cyber security tools cannot stop. These include prompt injection attacks — where malicious users craft inputs designed to manipulate LLM behavior and bypass safety guardrails -posed significant risks to data confidentiality and system integrity. There are also "jailbreak" attempts that threaten to override model alignment and safety protocols, potentially enabling unauthorized access to sensitive user data or the generation of harmful content.

Beyond threat protection, Dropbox also faced stringent operational requirements that further raised the bar for vendor selection. The company's architecture relies on self-hosted infrastructure with proprietary integrations, requiring on-premises deployment to maintain data privacy and regulatory compliance, without relying on external APIs. Performance was another non-negotiable constraint: Dropbox had to maintain ultra-low latency to preserve the real-time user experience that defines Dropbox's competitive positioning. Security tools that introduce hundreds of milliseconds of latency would render collaboration features unusable, forcing unacceptable trade-offs between protection and functionality.

The security team conducted comprehensive evaluations of multiple AI security vendors but struggled to find options that simultaneously delivered superior threat detection, on-premises deployment flexibility, extended context support for long documents, and millisecond-level response times. The challenge was identifying a security partner capable of protecting Dropbox's diverse GenAI applications — spanning RAG systems, document Q&A, and text summarization — while integrating seamlessly with existing machine learning infrastructure and scaling alongside the company's expanding AI portfolio.

SOLUTION

After a rigorous technical evaluation spanning multiple vendors and deployment architectures, Dropbox selected Check Point AI Agent Security as its enterprise AI security platform, which has distinguished itself through a combination of architectural flexibility, operational performance, and a collaborative partnership that precisely aligned with Dropbox's requirements. "After evaluating several AI security vendors and internal options, Check Point AI Agent Security quickly stood out in terms of its operational performance and ultra-low latency," Adrian Wood, security engineer, Dropbox, said.

Check Point AI Agent Security features three core capabilities that proved decisive: real-time, centralized protection and monitoring, ultra-low-latency processing, and flexible on-premises deployment options.

With centralized protection enabled, Check Point AI Agent Security establishes a unified security layer across all GenAI applications rather than implementing disparate security controls for discreet use cases. The platform integrates directly with Dropbox's centralized LLM library, providing security-by-default for all GenAI traffic flowing through the company's machine-learning infrastructure. This architecture delivered Dropbox immediate protection, visibility, and control, enabling its security team to monitor threat patterns across product lines while empowering individual product teams to innovate on new use cases without having to implement security measures independently.

Dropbox deployed Check Point AI Agent Security as a containerized Docker service within its Kubernetes environment, exposing it as a remote procedure call (RPC) service accessible to any LLM pipeline via standard internal service discovery.

The implementation leveraged LangChain as the orchestration framework, with Dropbox engineers creating multi-stage "security chains" that route user prompts through Check Point's AI Agent Security prompt-injection and jailbreak-detection APIs before reaching the language models, and then subject LLM responses to content-moderation analysis before returning results to users. This level of defense-in-depth provides multiple interception points for threat detection, catching adversarial inputs before they reach models and filtering harmful outputs before they reach users. Performance optimization represented a significant collaborative achievement: through close engineering partnership between Dropbox and Check Point's AI Security teams, the companies achieved a 7x latency improvement for extended-context scenarios involving prompts exceeding 8,000 characters — transforming a potential deployment barrier into a competitive advantage.

Check Point AI Agent Security detection capabilities proved strong across Dropbox's diverse use cases and threat scenarios. The platform analyzes prompts in real time using machine-learning models trained on adversarial AI attack patterns, providing granular confidence scores that enable Dropbox product teams to tune sensitivity levels to their risk tolerance and operational requirements. Check Point AI Agent Security achieved detection rates above 98 percent for prompt injection and jailbreak attempts while maintaining a false positive rate below 0.5 percent -critical specifications that ensure security effectiveness without degrading the user experience through excessive blocking of legitimate queries. The on-premises deployment model met Dropbox's data privacy requirements, ensuring that customer data never leaves the company's infrastructure while security analysis is performed locally within milliseconds.

OUTCOME

The deployment of Check Point AI Agent Security enabled Dropbox to establish a scalable security infrastructure that protects GenAI applications without compromising user experience or hindering innovation. "Check Point's AI Agent Security team has accelerated our GenAI journey, allowing us to create secure GenAI experiences at scale. They are great partners, and the team exhibits leadership in the space," Wood said.

The centralized security architecture provides protection-by-default for GenAI traffic flowing through Dropbox's machine learning systems, enabling product teams to rapidly develop and deploy new AI-powered features with confidence that security guardrails are applied. This eliminated the previous bottleneck where security reviews and custom protection implementations delayed feature launches, accelerating Dropbox's competitive positioning in the AI-enhanced productivity software market.

System integrity and security posture improved dramatically through comprehensive threat detection across all GenAI touchpoints. Check Point AI Agent Security successfully intercepted most of the prompt injection and jailbreak attempts during testing and production operations, providing security teams with centralized visibility into attack patterns and emerging threats. The platform's confidence scoring APIs enabled Dropbox to implement nuanced response strategies: immediately blocking high-confidence threats, flagging medium-confidence suspicious inputs for human review and logging, and allowing low-risk queries to proceed without friction. This granular control balanced security effectiveness with operational flexibility, ensuring that legitimate edge-case user behavior wasn't unnecessarily restricted while genuine threats were rapidly neutralized.

"We worked closely with the Check Point AI Agent Security team to minimize added latency as much as possible, and our current average latency is now a 7x improvement for prompts with more than 8,000 characters," Wood said.

The result of that collaboration? User experience remained uncompromised despite the addition of comprehensive security layers. Check Point AI Agent Security's ultra-low-latency processing — maintaining sub-50ms response times even for extended-context prompts after optimization — ensured that security analysis occurred transparently, without perceptible delays in interactive features.

The platform's remarkably low false-positive rate, tuned through collaborative engineering between Dropbox and Check Point AI Agent Security for the company's specific use cases, ensured that legitimate user queries proceeded smoothly without security interruptions. This achievement validated the principle that security and performance represent complementary rather than competing objectives when solutions are properly architected and optimized for production deployment.

"Check Point's AI Agent Security successfully protects our GenAI applications from most of the prompt injection and jailbreak attempts while maintaining incredibly low false positive rates. This means our users get the security they need without friction in their experience," Wood added.

"Check Point's AI Agent Security has given us the security foundation we need to scale our GenAI confidently across all product lines. As we continue expanding our AI capabilities, we know we have centralized protection that won't become a bottleneck," Wood said.

> The Check Point AI Security team has accelerated our GenAI journey, allowing us to create secure GenAI experiences at scale. They are great partners, and the team exhibits leadership in the space.
>
> Adrian Wood, security engineer, Dropbox

About Check Point Software Technologies Ltd.

Check Point Software Technologies Ltd. is a global cyber security leader protecting more than 100,000 organizations worldwide. Its mission is to secure enterprises' AI transformation. With a prevention-first approach and an open ecosystem architecture, Check Point helps organizations block advanced threats, prioritize exposures, and automate security operations across complex digital environments. The unified architecture simplifies protection across hybrid networks, multi-cloud environments, digital workspaces, and AI systems. Structured around four strategic pillars, Hybrid Mesh Network Security, Workspace Security, Exposure Management, and AI Security, Check Point delivers consistent protection and visibility across multivendor environments, enabling organizations to reduce risk, improve efficiency, and accelerate innovation without increasing complexity.

LEARN MORE

Worldwide Headquarters
5 Shlomo Kaplan Street, Tel Aviv 6789159, Israel | Tel: +972-3-753-4599
U.S. Headquarters
100 Oracle Parkway, Suite 800, Redwood City, CA 94065 | Tel: 1-800-429-4391
www.checkpoint.com

June 2026