Publication date:
14 Oct 2025
Author(s):
Rik Turner, Chief Analyst, Cybersecurity

On the Radar: DeepKeep secures models, GenAI, and AI agents

Summary

Catalyst

DeepKeep offers a technology platform that delivers security for large language models (LLMs), computer vision, and artificial intelligence (AI) variants that rely on them, namely generative AI (GenAI) and agentic AI. The company offers its technology in two products—DeepKeep for LLM and DeepKeep for Vision—and has three identified use cases:

- Employees using GenAI apps delivered in software as a service (SaaS), such as ChatGPT or Claude, as well as developers working with Cursor via their integrated development environment (IDE).
- Organizations developing homegrown applications with LLMs built into them.
- Organizations deploying AI agents.

Omdia view

You may be seeing media and analyst reports suggesting AI enthusiasm is waning, with talk of a "trough of disillusionment" regarding GenAI. Be that as it may, the march of AI adoption, in both its generative and agentic forms, is relentless, and Omdia suspects, with or without troughs, these new and evolving flavors of AI are here to stay. Agentic, in particular, holds immense promise, given its ability to automate processes within business, operations, and security.

However, this also makes security tools for all AI variants a growing requirement across most, if not all, areas of business. LLMs already underpin SaaS-delivered GenAI services and are increasingly incorporated into apps organizations are developing. While the cut-down versions, known as small

language models (SLMs), may grow in popularity, the fundamental security challenges models represent will not diminish.

Agentic AI further complicates matters by virtue of the levels both of access and autonomy that agents enjoy within an organization's infrastructure. The nightmare scenario for the deployment of AI agents is that they go rogue, operating like the "sentinels" in *The Matrix* movie franchise (mechanical hunters zipping around an organization's infrastructure and autonomously killing, blocking, or otherwise deactivating systems). Worse still, they could start exfiltrating sensitive data to points unknown on behalf of threat actors.

Why put DeepKeep on your radar?

DeepKeep's technology sits in what it describes as the semantic layer—the application layer in the OSI networking model of functionality in application stacks. This, together with its focus on understanding the customer- and application-specific context, is what enables its model scanning and red teaming capabilities, as well as informing its AI firewalling technology. Its approach to GenAI and agentic AI has both proactive and reactive capabilities, which Omdia considers essential to a comprehensive security approach.

Market context

The speed with which the newer variants of AI (generative and agentic) are developing and being adopted across multiple areas of business is truly unprecedented. If the software as a service (SaaS) delivery model for computing took around 15 years to go mainstream and cloud-native computing around 10, AI adoption is proceeding much faster. We are barely three years into the AI revolution, yet it is already invading multiple areas of business and, more broadly, of human activity. And of course, in doing so, it expands the attack surface of the adopters significantly.

The security challenges posed by GenAI and vibe coding

GenAI interactions present a variety of risks, the most obvious being an employee might upload sensitive corporate data, either in order to have a platform manipulate data (to formulate a press release, for example) or because they were tricked into providing more data than was appropriate.

In addition, GenAI has increasingly been applied to writing application code, a trend dubbed "vibe coding" in February 2025. The security issues here are immense: DevOps was already incorporating large amounts of open-source software into apps, with all the concomitant vulnerability risks that come with it, even before the vibe coding wave. Turning much of the task over to AI systems accelerates the output of code, but also increases the potential production of insecure apps.

The threats from agentic AI

Meanwhile, the implications of agentic AI are potentially even more serious, given both the level of access and autonomy that agents enjoy. The adoption of AI agents opens the door to increasing amounts of process automation, such as systems being trusted to perform actions without human intervention or even, in some cases, oversight. As such, a compromised agent (there is already talk of calling them "double agents") would enjoy considerable autonomy of action, as well as potentially extensive access to corporate assets. Another issue from a security perspective is the non-

deterministic nature of agents—their responses/actions may differ from one time to the next, even when they are working on the same input each time.

Security for, with, by, and against AI

These existing and emerging scenarios raise the profile of technology that Omdia refers to as "**security for AI**", referring to technology that enables AI to be used in a more secure and trustworthy fashion, regardless of which type of application it is underpinning. It entails various types of action, such as:

- Making sure input data, whether for training or actual inference, has not been "poisoned" by a malicious actor.
- Checking that an LLM has not been compromised.
- Verifying an agentic action is legitimate and its behavior is in keeping with what it was designed to do.

We prefer this description of the technology to the more generic catch-all of "AI security" since that term can also refer to:

- "**Security with/by AI,**" which is where AI, whether generative (the "with") or agentic (the "by"), is harnessed to enable better cybersecurity outcomes.
- "**Security against AI,**" where cybersecurity tools are used to detect and block attacks that have been crafted using AI, such as deepfakes.

DeepKeep operates in the market of security for AI.

Product/service overview

Both DeepKeep for LLM and DeepKeep for Vision comprise a series of modules that are deployed on a common underlying platform. These are:

- **Model Scanning** is designed to guarantee the model, whether an LLM, a vision-language model (VLM) in the computer vision world, or another type of model, is safe to incorporate in a customer's application. This module looks for the presence of trojans, signs of tampering, or license issues.
- **Automated AI Red Teaming** aims to identify and assess vulnerabilities, then suggest mitigations for a customer's AI applications. A key differentiator here, compared to other automated red teaming products, is the module's focus on understanding the customer-specific context of the application and adjusting its actions accordingly. Contextual information is gathered via the deployment of software agents within the customer's environment.
- **AI Firewall** provides the reactive security dimension of detecting and responding to attacks once an AI app is in production. The module monitors, filters, and blocks threats seeking to interact with AI models and ships with some 60+ built-in guardrails for this purpose.
- **Agentic AI security** delivers proactive security for agentic environments. It enables customers to define policies to permit or block specific Model Context Protocol (MCP)

servers, aiming to ensure usage is aligned with security requirements. It also looks at what privileges an agent has with a view to rightsizing them.

- **AI lens** provides governance around GenAI usage by a customer's employees, monitoring and controlling it across the organization.

Company information

Background

DeepKeep was founded in 2021 by CEO Rony Ohayon and CTO Yossi Altevet. Ohayon previously co-founded two other tech start-ups, namely DriveU, a developer of a connectivity platform for autonomous vehicles, where he was CEO, and LiveU, a vendor of video transmission and streaming technology, where he was CTO. Altevet spent time as a product manager at Cisco before working in the CTO's office at GlobalLogic, DriveU, and HCL Enterprise.

DeepKeep emerged in May 2024 and announced a $10m seed round led by Awz Ventures.

The vendor was founded before the current explosion of interest in, and adoption of, GenAI, agentic AI, and all things LLM-related, and its initial focus was on another area of AI—computer vision, where its technology beta'd in late 2022. This is a market that continues to be served by DeepKeep for Vision.

However, it decided the skills it brought to bear, and the tech platform it had developed, were relevant to the broader AI market taking shape since OpenAI's launch of ChatGPT in November 2022. This enabled the vendor to expand, rather than pivot, to address the security requirements of widespread Gen and agentic AI adoption. DeepKeep for LLM started beta in late 2023. The Agentic AI modules launched in 2025.

Current position

DeepKeep's target customers are primarily organizations in the large enterprise segment, with financial service institutions and telcos as leading adopters. Most of its business is currently in North America, though there are already some customers in the European and Asia & Oceania regions. Its go-to-market is a mixture of direct sales and channel partnerships.

The vendor offers its technology both for on-premises deployment and as a SaaS delivered from the cloud. While this may seem unusual for a start-up vendor at such an early stage of development, the company's first customers were in the automotive sector, developing autonomous vehicles and so requiring computer vision technology. All such customers required DeepKeep to deliver its product in an on-premises and airgapped manner. All its products are software, with no hardware component, and the on-premises version is available as microservices.

The charging mechanism for DeepKeep's technology varies between modules:

- For **Model Scanning** and **AI Red Teaming,** there is a fee per execution.
- The **AI Firewall** fee is charged on the prompts that the module inspects.
- For **Agentic AI,** the charge is per token, which can be thought of as per execution.
- **AI Lens** is per employee/developer.

The vendor sees customers starting from different points within its portfolio. Their development teams typically want Model Scanning and Red Teaming, whereas the production side of the house will require the AI Firewall. AI Lens may be the first port of call for a security team. Agentic adoption is generally at a much earlier stage.

DeepKeep's competitive landscape was primarily made up of other start-ups until recently, when larger vendors began buying some of them to get into security for AI. For example, Cisco bought Robust Intelligence in September 2024, and Palo Alto Networks acquired Protect AI in July 2025.

That said, the vendor says these deals have reduced the market to the point where it now sees another dedicated player, HiddenLayer, most often in competitive situations. The latter is not present in computer vision security, however.

Future plans

A development that DeepKeep is both watching and planning to address is the emergence of machine learning models capable of processing and integrating information from multiple modalities or types of data, such as text, voice, and video, a trend known as multimodal AI. Having started in computer vision and expanded to the world of generative and agentic AI, a logical next step for the vendor will be the addition of voice and other modalities into the mix.

DeepKeep also has plans for further functionality in agentic AI security, specifically additional monitoring and observability capabilities that will enable customers to know which individuals and systems within their organization are activating which agents and so on.

Finally, the company is looking to launch a marketplace so customers can acquire functionality that does not come directly from DeepKeep, but which is nonetheless complementary to and integrated with its own platform.

Key facts

Table 1: Data sheet: DeepKeep

Product/service name	DeepKeep	**Product classification**	AI security & trustworthiness
Version number	N/A	**Release date**	2024
Industries covered	All	**Geographies covered**	North America, EMEA, APAC
Relevant company sizes	Enterprise	**Licensing options**	Subscription-based (SaaS & on-premises)
URL	www.deepkeep.ai	**Routes to market**	Direct and channel
Company headquarters	Tel Aviv, Israel	**Number of employees**	45

Source: Omdia

Analyst comment

It is undeniable there is a growing addressable market for technology such as DeepKeep's. AI, particularly in its generative form, is invading pretty much every area of human activity, from business to medicine, sport, and government. Omdia's research finds most enterprises are already dabbling in agentic AI too, even though the majority of them are limiting their effort to the development of a single agent rather than going multi-agent from day one. We are currently assessing the exact size of the total addressable market in security for AI.

Part of the challenge start-ups like DeepKeep face is making themselves heard above the cacophony of vendors offering security technology for GenAI and agentic AI. Not only are there a myriad start-ups in this space, but a land grab in which larger tech players acquire specialists has already begun (see above).

While this may slow the evolution of the acquirees, the buyers have deep pockets to fund their marketing activities, potentially drowning out other specialists with which they are now in competition. DeepKeep clearly needs to redouble its efforts to raise its profile in this increasingly crowded (and noisy) space.

Omdia sees DeepKeep as well placed in this emerging market for a couple of reasons. Firstly, because it already has security for computer vision in its arsenal, as well as the more text-based LLMs and agentic AI. Expansion into security for multimodal AI looks like a logical next step and one that should prove easier to execute, given the vendor's heritage. Secondly, its portfolio has both a proactive/posture management dimension as well as a reactive/detection and response one, which is essential for the longer-term success of vendors of security for AI and their customers. Its native multilanguage support is a third important factor, in that if an organization wants to run proper AI red teaming, their datasets need to be native.

Appendix

On the Radar

On the Radar is a series of research notes about vendors bringing innovative ideas, products, or business models to their markets. On the Radar vendors bear watching for their potential impact on markets as their approach, recent developments, or strategy could prove disruptive and of interest to tech buyers and users.

Author

Rik Turner, Chief Analyst, Cybersecurity

askananalyst@omdia.com

Citation policy

Request external citation and usage of Omdia research and data via citations@omdia.com.

Omdia consulting

We hope that this analysis will help you make informed and imaginative business decisions. If you have further requirements, Omdia's consulting team may be able to help you. For more information about Omdia's consulting capabilities, please contact us directly at consulting@omdia.com.

Copyright notice and disclaimer

CONTACT US

omdia.com

askananalyst@omdia.com