

On the Radar: NuGuard AI delivers pre-production security for agentic environments

Summary

Catalyst

NuGuard AI offers a platform that delivers what it describes as behavioral validation for agentic AI. It discovers an organization's AI assets and generates an AI software bill of materials (AI-SBOM), which can then serve as both a definitive inventory and a list of each asset's expected behavior. It goes on to validate the behavior of each agent in a sandbox before it goes into the production environment.

Omdia view

The rate at which agentic AI technology is being adopted across multiple verticals is unprecedented, even when compared with other waves of technological innovation such as the internet or cloud computing. Omdia is already hearing accounts of banks and retailers operating agent infrastructures that number in the hundreds of thousands. Yet, as is so often the case, the drive for speed continues to outpace concern for security.

As such, there is a clear market opportunity for technology that provides governance capabilities around the development and deployment of AI agents, focusing particularly on their levels of autonomy and access. NuGuard's technology, combining the generation of AI-SBOMs with red teaming and behavioral validation, looks well placed to address many of the security requirements created by the boom in agentic AI.

Why put NuGuard AI on your radar?

NuGuard offers an open source version of its platform with all the basic functionality required to provide governance or oversight for agentic environments, which enables organizations to gain familiarity with how it works. From there, they can choose the Enterprise version of the product, gaining features such as risk scoring and audit reporting to guide remedial action.

Market context

It is clear that the business world is in the middle of a full-scale revolution thanks to agentic AI. Ever more agents are incorporated into customer-facing and internal applications across multiple verticals, accelerating and streamlining human activity in some areas while they are replacing employees altogether in others. Studies indicate that some of the more forward-thinking organizations in finance, retail, and other verticals are already running tens and even hundreds of thousands of agents within their IT infrastructure.

As is often the case, speed of adoption increases risk

The rush to adopt agentic AI, however, brings with it a potentially significant expansion in an organization's attack surface. Given the levels of autonomy and access that can be granted to agents, oversight and control of what they can and cannot do are becoming vital capabilities.

While the notion of guardrails has been around in platform engineering since the late 2010s, it has really come to the fore in the AI era. Indeed, the need to provide security for AI (rather than harnessing AI to deliver security, a trend Omdia describes as security with or by AI) that underpins NuGuard's technology offering.

The agentic AI behavior gap

The vendor speaks of a "behavior gap" in agentic AI, describing it as "the difference between intent and real behavior." In other words, agents are interacting with classified data and acting autonomously, mostly without a well-defined behavior policy and without behavior being validated prior to deployment. NuGuard argues that a full agentic stack requires four layers, namely:

- A foundation model for reasoning
- A tools layer, consisting of APIs, databases, and execution environments
- An orchestration layer for task decomposition and multi-agent coordination
- A trust and policy layer.

One could see the fourth layer here (trust and policy layer) as a sort of control plane, in which trust is established and policies are defined, though NuGuard notes that in most deployments today, it is missing entirely. The vendor cites statistics from Cloud Security Alliance studies suggesting that 53% of organizations have had AI agents exceed their intended permissions, while 65% reported an AI agent-related security incident in the past year with real business impact. Most worryingly, it also has data to suggest that as many as 79% of organizations lack written policies for governing AI agents, despite having already deployed them into production.

The AI-SBOM concept

One of the cornerstones of NuGuard's product offering is the generation of AI-SBOMs. An AI-SBOM can serve as the reference for what software components an agentic infrastructure contains, cataloging the datasets, base models, software dependencies, algorithms, and prompt pipelines that make up the AI application. This helps organizations track supply chain risks, detect bias, and ensure regulatory and internal policy compliance.

The AI-SBOM concept is an evolution of the SBOMs already mandatory in certain market segments (for example, for vendors supplying the US federal government). Although it has not yet reached the status of a legal requirement, it makes sense as a baseline for an AI governance strategy. It should enable organizations not only to set the necessary guardrails but also to test for drift, monitor ongoing behavior, and take corrective action when needed.

Product/service overview

The workflow of the NuGuard AI Trust Platform is the classic one of “discover, define, validate, and govern” that determines how most proactive security products work. The platform:

- Scans an organization's codebase
- Maps its AI asset (models, MCP servers, agents, tools, etc.) before they have gone into the production environment (i.e., at the build stage and prior to the push request)
- Generates an AI-SBOM based on what it has found
- Defines the assets' roles and entitlements (access rights), which will determine what constitutes acceptable behavior
- Validates these entitlements in the AI-SBOM against the developer's original intent
- Carries out automated red team exploits against the assets in a sandbox to stress-test their resilience
- Simulates jailbreak scenarios to see how susceptible they might be to malicious manipulation
- Governs the assets, continually checking that they are not drifting beyond their defined roles and behaviors.

When it detects drift, the platform also provides suggestions on how to fix it, offering the code snippets and policy adjustments required to bring the agentic system back into line with the behavior sanctioned by the organization.

Differentiators

NuGuard highlights what it considers to be four key differentiators of its product offering. They are:

- The fact that it operates in the pre-production environment (i.e., not after an incident has taken place). This is the essence of proactive security, enabling organizations to catch behavioral issues before code ships. It notes that pre-production fixes cost anywhere between 10 and 100 times less than post-incident remediation.

- The open source and auditable nature of the platform, which provides visible logic, community-contributed patterns and, crucially, no vendor lock-in. Security teams can inspect every test, which is particularly critical in regulated industries.
- The Redteam Innovation Engine, which generates in excess of 1,000 tests on-the-fly for the customer's specific AI stack and behavior policy and tests agents, tools, APIs, and guardrails simultaneously.
- The platform's delivery of audit-ready evidence, with every run producing an AI-SBOM, supply-chain analysis, compliance mapping (including the NIST AI RMF, EU AI Act, PCI, and HIPAA), and sandbox validation, all in a format that can be audited by regulators and insurers.

Company information

Background

NuGuard AI was founded by CEO Ranjan Goel in 2025. Goel previously held executive positions at a number of tech firms, including (most recently) observability vendor LogicMonitor, where he was VP of Product Management. Before that he was Director of Product Management at network monitoring vendor Gigamon, holding the same title earlier at networking vendors Extreme Networks and Cisco.

The vendor has yet to raise any outside funding.

Current position

NuGuard's go-to-market strategy began with the release of an open-source edition of its product, which continuously tests, validates, and red-teams agentic AI applications at the pre-production stage. The aim is to help developers close the behavior gap between an agent's intended actions and what it actually does in deployment. Its core capabilities are:

- **AI-SBOM generation:** mapping models, tools, APIs, and frameworks in an organization's stack into an AI software bill of materials (AI-SBOM)
- **Behavior definition:** standardizing the AI application's intended behavior (e.g., topics it should or should not respond to, human-in-the-loop conditions, etc.)
- **Product validation through automated testing:** scanning for structural risks, including prompt injection, supply chain attacks, data exfiltration, and tool misuse. It then proposes remediation to the specific components of the AI application
- **Governance:** providing risk scores and audit reporting to ensure compliance before production.

The vendor released the open source edition of the product to generate interest and awareness, and has now released the paid-for Enterprise version, which adds risk scoring for compliance, as well as enabling customers to identify gaps they need to fix in order to meet the requirements of their chosen AI security framework. From a compliance perspective, the risk scoring can be adjusted to reflect an organization's vertical, such as how well a healthcare company complies with HIPAA or a retailer with PCI-DSS.

The charging mechanism for the Enterprise version of the NuGuard AI Trust Platform is currently based on the number of lines of code the product scans. The expectation is that customers will purchase roughly 50 scans of 10,000 lines of code per month. Each dimension of the platform—red teaming, compliance checking, and behavioral-policy validation—consumes credits, allowing customers to choose whether they want to use one, a combination, or all of these capabilities.

NuGuard takes an enterprise-friendly approach, offering unlimited users and repositories with no feature-tier licensing complexity.

Future plans

In the coming months, NuGuard plans to partner with managed security service providers (MSSPs) to reach small and midsize businesses (SMBs), expand support for additional regulatory requirements to serve more verticals, and grow its presence in Europe, the Middle East, and Africa (EMEA)—with an initial focus on the UK, France, and Germany.

Key facts

Table 1: Data sheet: NuGuard AI

Product/service name	NuGuard AI Trust Platform	Product classification	Security for AI
Version number	N/A	Release date	An open source version is already available; Enterprise release due in July 2026
Industries covered	Fintech, healthcare, retail and manufacturing	Geographies covered	US initially
Relevant company sizes	Initially SMBs, expanding to all	Licensing options	Open source CLI is free. Enterprise: Lines of code scanned
URL	www.nuguard.ai	Routes to market	Product-led growth. Open source (free download from GitHub) Enterprise: direct, initially up to 10 customers. Then MSPs/resellers and cloud provider marketplaces
Company headquarters	San Jose, CA, US	Number of employees	8

Source: Omdia

Analyst comment

The launch of the Enterprise edition of NuGuard's tech platform is a timely one, in that there is clearly a market requirement for such functionality at the moment, given the corporate rush to introduce agentic capabilities into countless business processes and the apps that underpin them. Moreover, the fact that a lot of the product's basic capability is available in the open source version makes a compelling argument for organizations to try it out, enabling them to get acquainted with the CLI and test the platform before committing any budget to it.

One of the realities NuGuard will face as it takes its technology to market is the sheer number of vendors—large and small—claiming to offer security solutions for agentic AI. Cutting through that noise will be challenging, particularly as several of the players in this space are well-established companies with deep marketing budgets and strong brand recognition. Another dimension to be considered here is the need to articulate clearly what the platform does and where it fits within a broader infrastructure of security tooling for agentic AI. NuGuard should highlight the proactive nature of its technology—for example, that it sits in the pre-production part of the lifecycle of an agentic system (i.e., the LLM, MCP servers, and the agents), integrating into the CI/CD pipeline at the build stage to provide behavior validation oversight ahead of the push request.

As such, it is not a provider of reactive security, detecting and reacting to incoming attacks on the organization. Indeed, it may even be able to integrate with, learn from, and inform those security tools.

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

The Omdia research, data and information referenced herein (the "Omdia Materials") are the copyrighted property of TechTarget, Inc. and its subsidiaries or affiliates (together "Informa TechTarget") or its third party data providers and represent data, research, opinions, or viewpoints published by Informa TechTarget, and are not representations of fact.

The Omdia Materials reflect information and opinions from the original publication date and not from the date of this document. The information and opinions expressed in the Omdia Materials are subject to change without notice and Informa TechTarget does not have any duty or responsibility to update the Omdia Materials or this publication as a result.

Omdia Materials are delivered on an "as-is" and "as-available" basis. No representation or warranty, express or implied, is made as to the fairness, accuracy, completeness, or correctness of the information, opinions, and conclusions contained in Omdia Materials.

To the maximum extent permitted by law, Informa TechTarget and its affiliates, officers, directors, employees, agents, and third party data providers disclaim any liability (including, without limitation, any liability arising from fault or negligence) as to the accuracy or completeness or use of the Omdia Materials. Informa TechTarget will not, under any circumstance whatsoever, be liable for any trading, investment, commercial, or other decisions based on or made in reliance of the Omdia Materials.

CONTACT US

[omdia.com](https://www.omdia.com)

askananalyst@omdia.com