



Datadog Expands AI Security Capabilities to Enable Comprehensive Protection from Critical AI Risks

June 10, 2025 at 4:00 PM EDT

Launch of Code Security and new security capabilities strengthen posture across the AI stack, from data and AI models to applications

NEW YORK – JUNE 10, 2025 – [Datadog](#), Inc. (NASDAQ: DDOG), the monitoring and security platform for cloud applications, today announced new capabilities to detect and remediate critical security risks across customers' AI environments—from development to production—as the company further invests to secure its customers' cloud and AI applications.

AI has created a new security frontier in which organizations need to rethink existing threat models as AI workloads foster new attack surfaces. Every microservice can now spin up autonomous agents that can mint secrets, ship code and call external APIs without any human intervention. This means one mistake could trigger a cascading breach across the entire tech stack. The latest innovations to Datadog's Security Platform, presented at [DASH](#), aim to deliver a comprehensive solution to secure agentic AI workloads.

"AI has exponentially increased the ever-expanding backlog of security risks and vulnerabilities organizations deal with. This is because AI-native apps are not deterministic; they're more of a black box and have an increased surface area that leaves them open to vulnerabilities like prompt or code injection," said Prashant Prahlad, VP of Products, Security at Datadog. "The latest additions to Datadog's Security Platform provide preventative and responsive measures—powered by continuous runtime visibility—to strengthen the security posture of AI workloads, from development to production."

Securing AI Development

Developers increasingly rely on third-party code repositories which expose them to poisoned code and hidden vulnerabilities, including those that stem from AI or LLM models, that are difficult to detect with traditional static analysis tools.

To address this problem, Datadog [Code Security](#), now Generally Available, empowers developer and security teams to detect and prioritize vulnerabilities in their custom code and open-source libraries, and uses AI to drive remediation of complex issues in both AI and traditional applications—from development to production. It also prioritizes risks based on runtime threat activity and business impact, empowering teams to focus on what matters most. Deep integrations with developer tools, such as IDEs and GitHub, allow developers to remediate vulnerabilities without disrupting development pipelines.

Hardening Security Posture of AI Applications

AI-native applications act autonomously in non-deterministic ways, which makes them inherently vulnerable to new types of attacks that attempt to alter their behavior such as prompt injection. To mitigate these threats, organizations need stronger security controls—such as separation of privileges, authorization bounds, and data classification across their AI applications and the underlying infrastructure.

Datadog [LLM Observability](#), now Generally Available, monitors the integrity of AI models and performs toxicity checks that look for harmful behavior across prompts and responses within an organization's AI applications. In addition, with Datadog [Cloud Security](#), organizations are able to meet AI security standards such as the NIST AI framework out-of-the-box. Cloud Security detects and remediates risks such as misconfigurations, unpatched vulnerabilities, and unauthorized access to data, apps, and infrastructure. And with Sensitive Data Scanner (SDS), organizations can prevent sensitive data—such as personally identifiable information (PII)—from leaking into LLM training or inference data-sets, with [support for AWS S3 and RDS instances](#) now available in Preview.

Securing AI at Runtime

The evolving complexity of AI applications is making it even harder for security analysts to triage alerts, recognize threats from noise and respond on-time. AI apps are particularly vulnerable to unbound consumption attacks that lead to system degradation or substantial economic losses.

The [Bits AI Security Analyst](#), a new AI agent integrated directly into Datadog [Cloud SIEM](#), autonomously triages security signals—starting with those generated by AWS CloudTrail—and performs in-depth investigations of potential threats. It provides context-rich, actionable recommendations to help teams mitigate risks more quickly and accurately. It also helps organizations save time and costs by providing preliminary investigations and guiding Security Operations Centers to focus on the threats that truly matter.

Finally, Datadog's [Workload Protection](#) helps customers continuously monitor the interaction between LLMs and their host environment. With new LLM Isolation capabilities, available in preview, it detects and blocks the exploitation of vulnerabilities, and enforces guardrails to keep production AI models secure.

To learn more about Datadog's latest AI Security capabilities, please visit: <https://docs.datadoghq.com/security/>.

Code Security, new tools in Cloud Security, Sensitive Data Scanner, Cloud SIEM, Workload and App Protection, as well as new security capabilities in LLM Observability were announced during the keynote at DASH, Datadog's annual conference. The replay of the keynote is available [here](#). During DASH, Datadog also announced launches in [AI Observability](#), [Applied AI](#), [Log Management](#) and released its [Internal Developer Portal](#).

About Datadog

Datadog is the observability and security platform for cloud applications. Our SaaS platform integrates and automates infrastructure monitoring,

application performance monitoring, log management, user experience monitoring, cloud security and many other capabilities to provide unified, real-time observability and security for our customers' entire technology stack. Datadog is used by organizations of all sizes and across a wide range of industries to enable digital transformation and cloud migration, drive collaboration among development, operations, security and business teams, accelerate time to market for applications, reduce time to problem resolution, secure applications and infrastructure, understand user behavior and track key business metrics.

Forward-Looking Statements

This press release may include certain "forward-looking statements" within the meaning of Section 27A of the Securities Act of 1933, as amended, or the Securities Act, and Section 21E of the Securities Exchange Act of 1934, as amended including statements on the benefits of new products and features. These forward-looking statements reflect our current views about our plans, intentions, expectations, strategies and prospects, which are based on the information currently available to us and on assumptions we have made. Actual results may differ materially from those described in the forward-looking statements and are subject to a variety of assumptions, uncertainties, risks and factors that are beyond our control, including those risks detailed under the caption "Risk Factors" and elsewhere in our Securities and Exchange Commission filings and reports, including the Annual Report on Form 10-K filed with the Securities and Exchange Commission on May 6, 2025, as well as future filings and reports by us. Except as required by law, we undertake no duty or obligation to update any forward-looking statements contained in this release as a result of new information, future events, changes in expectations or otherwise.