AI Updates, Tools & Guides | Security

AI can write code faster. Can enterprises secure it

AI can shrink days of work into minutes and generate code at remarkable speed. But as enterprises hand more work to AI agents, the harder challenge is securing, validating, and tracking everything those agents do.

Ashok Pandey

29 Jul 2026 15:16 IST

Artificial Intelligence (AI) is changing how quickly enterprises can build software, investigate incidents, and automate work. But every gain in speed brings a less visible problem. The more AI generates and acts on its own, the more businesses need to validate, monitor, and secure.

In an interview with PCQuest, ***Shomiron Das Gupta, CEO and Founder of Bloo***, shared his insights on whether enterprise security practices can keep pace with what AI now makes possible.

Productivity is no longer the only question

The practical case for AI can be compelling. In one personal example shared during the interview, AI enabled roughly four times the amount of work possible earlier. A product idea that might once have remained a written thesis can now become a prototype in about three days, helping a team see the idea in action.

The impact can be even sharper in enterprise operations. When a critical application or security incident triggers a Priority 1 (P1) investigation, specialists from cloud, infrastructure, applications, and development may spend five to eight hours finding the cause and preparing a root cause analysis (RCA). Difficult cases can continue for several days.

A tool described during the interview can perform such an investigation in about three minutes at a cost of USD 5 by examining different hypotheses and expanding its analysis when it identifies a clue. The productivity opportunity is clear. The harder question is what happens when this speed spreads across enterprise technology.

AI has moved the software bottleneck

AI-assisted coding can sharply increase the amount of software produced. Security teams, however, still have to validate that code. That shifts the bottleneck from creation to verification. More code means more potential attack surface. If AI can generate thousands of lines of code far faster than before, organizations may not have enough human capacity to inspect every line for vulnerabilities.

The fact that AI-generated code works is not enough. Enterprises also need to know whether it can be exploited, whether it meets security requirements, and whether it will remain secure as it scales. Security benchmarks need to become part of the development process. Organizations can define how an application should work, but they must also define the security tests it needs to pass before code enters a repository.

Ironically, AI itself may become part of the answer. If humans cannot manually validate the growing volume of generated code, AI may need to help test and secure the code it creates.

The real risk is not knowing what went wrong

As AI moves from assisting people to taking actions within enterprise systems, observability becomes critical. The biggest concern may not be that an AI agent makes a mistake. It may be that nobody knows it made one.

When an autonomous system fails, enterprises need to trace what happened. Which agent made the decision?

What information did it access? What action did it perform? Where did the process go wrong?

This requires AI agents to have clear identities, authentication, controlled access, and observable activity. The principle is familiar. Enterprises already apply identity and access controls to people, cloud infrastructure, and other systems. AI agents will require similar treatment. An agent accessing a database, data pipeline, data fabric, or web application programming interface (API) should not operate as an invisible actor. Its identity and actions need to be traceable.

Security teams may need to watch the verbs

One practical approach is to monitor AI agents based on their actions.

Read. Write. Share. Create. Delete.

These verbs describe what an agent is actually doing inside an enterprise environment. Security controls can be placed around each action, just as access controls have traditionally governed what people and machines are allowed to do.

This approach could help security teams adapt existing principles to autonomous AI. The actor may be new, but the actions are familiar. The same logic applies to vector databases and other components of the AI stack. Authentication, authorization, and access control remain essential, particularly when AI systems rely on contextual business data. New technology does not remove the need for established security practices.

Getting from a fast prototype to production

Perhaps the bigger unresolved question is what happens after AI produces version one. A single person can now use AI to perform work that traditionally involved different developers handling front-end, back-end, database, and other tasks. That can make building an initial product fast.

Production is different.

The application still needs to scale. It needs to be secured and tested. Other engineers need to participate. Large teams must find ways to combine their knowledge while AI handles more of the coding work. This

creates a new collaboration problem. If a software project involves 100 people, how do all 100 use AI to build, test, and secure one application together?

That question remains open.

AI has accelerated the act of building software, but enterprise engineering involves more than generating working code. The next challenge is creating development and security models that can match AI's speed without losing the collective judgment that complex software demands.

Speed needs visibility

Enterprise AI security is unlikely to be solved by one product or control. Agents will need identity and authentication. Their access to data and systems must be governed. AI-generated code must be tested against defined security benchmarks. Autonomous actions need to be monitored and traced.

AI has made it easier to build something quickly and see it work right away. That is already changing expectations around productivity. But speed has moved the problem rather than removed it.

The next test for enterprises is whether they can make AI-driven systems secure, observable, and manageable at the same pace they are learning to build them.