IDC Innovators: Autonomous Penetration Testing for DevSecOps, 2026

Katie Norton

Senior Research Manager,
DevSecOps and Software Supply Chain Security

July 2026

This IDC Excerpt Features: Xint.io

Synopsis

The content for this excerpt was taken directly from “IDC Innovators: Autonomous Penetration Testing for DevSecOps, 2026” (July 2026, Doc #US54175326)

This IDC Innovators presentation highlights emerging vendors in AI agent–driven autonomous penetration testing (pentesting) for DevSecOps. As applications become more distributed and frequently updated, security teams need offensive testing that can run closer to development and production changes.

This presentation profiles seven vendors: Corgea, Equixly, Novee, Prime Security, Ridge Security, XBOW, and Xint.io. The vendors in this report take different approaches, including source code–informed validation, product context modeling, proprietary offensive models, model orchestration, and multi-agent exploit validation. What they share is a common goal: making pentesting more continuous and aligned with the speed of software delivery.

"AI agents are changing the operating model for pentesting by combining adversarial reasoning with more continuous execution. This gives security teams a way to validate exploitable risk closer to the pace of software change."

Katie Norton

Senior Research Manager,
DevSecOps and Software Supply Chain Security

Contents

Executive summary

Application security teams are under pressure to validate risk at the pace of software change. Traditional pentesting remains valuable for expert-led assessment, but periodic engagements are increasingly difficult to align with fast release cycles, API-heavy architectures, AI-generated code, and application environments that may change daily. At the same time, automated testing continues to generate findings that often require additional validation before security teams can determine whether they represent real exploitable risk.

Autonomous pentesting is emerging to address this gap by using AI agents to conduct offensive security testing. These solutions are designed to discover or ingest target context, map target behavior and attack surfaces, generate exploit hypotheses, execute controlled tests, validate exploitability, and return evidence-backed findings with remediation guidance. The category is still early, but its direction is clear: Pentesting can now be delivered as continuous, software-driven validation.

This IDC Innovators presentation features seven companies operating in autonomous pentesting for DevSecOps: Corgea, Equixly, Novee, Prime Security, Ridge Security, XBOW, and Xint.io. The vendors selected for this research focus on testing applications, including web applications, APIs, and AI-enabled applications, rather than broader network or infrastructure pentesting. They share a focus on exploit validation, application context, agent-led testing, and developer-consumable findings, but differ in approach. Some incorporate source code or product design context, some emphasize proprietary offensive models, and others differentiate through orchestration, safety controls, or model flexibility.

For technology buyers, evaluation should center on whether the platform can safely test within approved boundaries, understand enough application context to find meaningful paths, prove exploitability, and deliver findings that developers can reproduce and fix. The value of autonomous pentesting will depend less on finding volume and more on whether the output is validated, prioritized, and usable in existing DevSecOps workflows.

Market assessment

- Current usage data underscores a structural gap between the cadence of traditional pentesting and pace of change in modern environments. IDC's March 2025 *Exposure Management Survey* found 76% of organizations conduct external pentesting twice a year or less, including 11.5% that never engage an external tester, whereas only 5.5% test monthly or more. This cadence is difficult to reconcile with applications, APIs, and AI-enabled systems that can change continuously as new code, functionality, and workflows are introduced. Autonomous pentesting is designed to close this gap by replacing periodic audits with continuous, machine-speed validation.
- IDC survey data shows a consistent gap between demand signals for automated testing and traditional pentesting. Application security testing tools, such as dynamic application security testing (DAST), rank as a top 3 investment priority for 33.9% of organizations, compared with 19.8% for pentesting, red teaming, or bug bounty programs. Further, the top reason organizations most recently purchased a new application security tool was the adoption of AI/GenAI-powered security features (30.5%, according to IDC's July 2025 *DevSecOps and Software Supply Chain Survey*. Autonomous pentesting is likely to draw on this demand.
- Trust rankings for agentic AI tasks placed vulnerability scanning third among the activities organizations trust AI agents to complete (22.4%), behind testing code (46.5%) and writing code (31.3%), according to IDC research. Testing and code generation are already established, high-adoption use cases for agentic AI, and vulnerability scanning's proximity to the top tier suggests autonomous pentesting is positioned to follow the same adoption curve.
- Autonomous pentesting currently occupies its own distinct category, but its boundaries with several adjacent testing approaches are already thin. Those boundaries are likely to erode further as the market matures toward a single AI-driven offensive security testing category:
  - DAST already automates testing the running application at scale, running predefined payloads against live applications on a continuous basis. Autonomous pentesting adds a layer of reasoning that DAST lacks, interpreting application logic, state changes, and role-based workflows to chain multiple weaknesses into a single exploit path rather than flagging isolated findings.
  - Manual testing remains the category most associated with adaptive reasoning, but it depends on scheduled, resource-intensive human engagements. Autonomous pentesting performs a comparable analytical function continuously and at machine speed.
  - Red teaming is an objective-based, human-led exercise typically run over an extended period. Pentesting as a service (PTaaS) is a delivery model that combines human testers with cloud-based platforms rather than a distinct testing methodology. Both intersect with autonomous pentesting without fully overlapping with it since autonomous agents can sustain a red team–like adversarial posture continuously and increasingly serve as the engine inside PTaaS delivery models.

Company profile

Xint.io

Company profile: Xint.io

IDC Innovator | Xint.io

Founded
2025

Number of employees
20–50

Headquarters
Austin, Texas, United States

Founders
Brian Pak (CEO) and Andrew Wesie (CTO) are Carnegie Mellon University alumni who founded Theori, an offensive security research and consulting firm, before spinning out Xint to translate its offensive security depth into an automated platform.

Product/service description
Xint.io is an AI-native application security platform, with Xint Web providing live application testing for exploitable vulnerabilities and business logic flaws and Xint Code providing source code analysis.

Funding
The company is privately funded through Theori, the parent company, which has operated as an independent offensive security firm since 2016.

Why Xint.io was chosen as an IDC Innovator

Black-box live application testing can miss security-relevant paths that are difficult to discover from the outside alone. Modern applications often contain client-side routes, undocumented API calls, hidden parameters, and workflow logic that may not be visible through crawling or API documentation. Broader attack surface discovery is critical because exploitability often depends on reaching the right application path, role, or state before testing can determine whether a weakness is real.

Xint Web tests running applications while drawing on complementary source-derived context from Xint Code, in which it can improve discovery and targeting. The product reflects Xint.io's origin as a spinout of Theori, an offensive security firm with experience running red team and pentesting engagements for global enterprises and governments. This background brings practitioner-led offensive security expertise to an AI-driven approach for live application testing, with an emphasis on the reasoning, validation, and evidence expectations associated with pentesting.

Company profile: Xint.io

IDC Innovator assessment

1 Application discovery builds the testing map.

Xint Web can ingest OpenAPI or Swagger specifications to enumerate endpoints, parameters, authentication schemes, and expected request/response structures, accelerating initial coverage. When such documentation is missing or incomplete, Xint Code analyzes client-side JavaScript and other artifacts to extract hardcoded routes, API calls, and hidden parameters that may not be externally visible. These inputs are normalized into a structured map that captures endpoint relationships, parameter types, state transitions, and application paths, providing a detailed foundation for targeted testing.

2 Exploit hypotheses guide live investigations.

The testing engine uses the application map to generate exploit hypotheses, which are reasoned assumptions about how an attacker might combine inputs, states, and application behaviors to achieve an unintended outcome. Rather than testing each endpoint in isolation, the system considers how different parts of the application interact, forming scenarios, such as bypassing authentication through workflow manipulation, exploiting authorization gaps across roles, or chaining state-dependent actions to reach restricted functionality. Xint Web then investigates these hypotheses using LLM-based reasoning, focusing on whether the predicted attack paths can be realized in practice.

3 Validation turns issues into proof-backed findings.

Investigation and validation focus on executing candidate scenarios against the running application and packaging confirmed issues with proof of vulnerability evidence. This process emphasizes confirming exploitability under realistic conditions and reducing false positives through repeatable testing. Findings can include reproduction steps, proof-of-concept scripts, HTTP request and response evidence, impact, rationale, and related threat scenario detail when the issue is part of a multistep attack chain.

Company profile: Xint.io

Key differentiators

→ Behavioral testing focuses on application logic and workflows.

- Many high-impact application vulnerabilities emerge from valid actions performed in the wrong order, by the wrong role, or under the wrong state. Xint Web's focus on workflow sequence, state, roles, and business logic helps move testing beyond endpoint coverage and toward abuse path validation. For security teams, this creates a better fit for finding issues, such as authorization gaps, privilege escalation paths, and multistep business logic flaws, that conventional dynamic testing can miss.

→ The testing workflow is not tied to one model.

- Xint.io treats models as interchangeable components within a broader testing system. The surrounding orchestration determines how models are selected, configured, and applied across discovery, reasoning, and validation tasks. The approach may be more durable as model performance changes or new models are introduced in the ecosystem. Xint.io can adapt the underlying model mix while preserving the testing workflow, validation discipline, and evidence structure on which buyers rely.

→ Source context helps expand the live testing surface.

- Xint Code can complement Xint Web by identifying routes, API calls, and client-side references that may not appear in OpenAPI documentation or standard crawling. Bringing code and running application context into the same product give the testing engine a stronger basis for validation. Source analysis can point to paths worth investigating, whereas live testing determines whether those paths are reachable and exploitable in the deployed environment.

Company profile: Xint.io

Challenges

→ Cross-category value can complicate budget ownership.

- Xint.io's code-plus-web model sits between SAST, DAST, and pentesting budgets. This positioning is differentiated, but it may require clear packaging and value mapping so buyers can understand which existing workflows, tools, or services the platform augments or replaces.

→ Static-to-dynamic correlation is not yet available.

- Today, source analysis can support live testing by improving endpoint discovery, but the platform does not yet close the loop from code-level finding to live exploitability and fix validation. As Xint.io positions code and runtime testing together, buyers will expect a unified findings experience that connects related issues across source analysis, live testing, and remediation.

→ Developer workflow fit is still evolving.

- Xint Web is currently oriented toward deeper security team testing cycles, including release assurance, staging or production validation, audits, and post-incident review. As Xint.io extends toward more incremental and developer-connected workflows, it will need to balance the depth of live exploit validation with the speed and frequency expected in PR and CI/CD processes.

Technology definition

Autonomous pentesting for DevSecOps refers to AI-driven offensive security platforms that use autonomous agents, model orchestration, and security tooling to execute scoped pentesting against web applications, APIs, and AI-enabled application surfaces. These platforms discover or ingest target context, map attack surfaces and workflows, generate and pursue exploit hypotheses, adapt based on observed behavior, validate exploitability through safe and reproducible proof, and return findings with evidence, reproduction detail, remediation guidance, and retesting support. The technology is designed to make adversarial testing more continuous, repeatable, and integrated with software delivery.

Core capabilities include:

- Autonomous discovery, planning, execution, validation, and reporting within approved boundaries
- Context-aware testing using live application behavior, API relationships, source code, design/cloud context, credentials, roles, or combinations of these inputs
- Evidence-backed findings that prioritize confirmed exploitable risk over theoretical vulnerability flags
- DevSecOps workflow integration, including ticketing, CI/CD, PRs, coding agent handoff, reporting, or retesting

IDC Innovators inclusion criteria

An IDC Innovators document recognizes emerging companies chosen by an IDC analyst because they offer an innovative new technology, a groundbreaking business model, or both and were approved by the IDC Innovators Review Process. It is not an exhaustive evaluation of all companies in a segment or a comparative ranking of the companies.

An IDC Innovators document highlights companies that meet the following criteria:

- In IDC's opinion, the company exhibits innovative technology or a new business model.
- The company's annual revenue is under $100 million at the time of selection.
- Customers are currently using the company's products and services (i.e., the products and services are not conceptual or in the process of being released).
- The product, service, or business model must solve or help alleviate an IT buyer challenge.

In addition, companies in the process of being acquired by a larger company may be included provided the acquisition has not been finalized at the time of publication of the document. Companies funded by venture capital firms may also be included even if the venture capital firm has a financial stake in the company.

Related research

Document title	Document number	Publication date
IDC Market Glance: DevSecOps, 2Q26	US52088124	June 2026
IDC Survey: Goals, Benefits, and Barriers to Agentic AI Adoption in Software Development — Insights from IDC's Agentic Application Development and DevOps Survey	US54283526	March 2026
Exposure Management Views, 2026	US54199826	January 2026
DevSecOps in Practice: What the Data Tells Us About Implementation and Impact	US53869825	October 2025

About IDC

International Data Corporation (IDC) is the premier global market intelligence, data, and events provider for the information technology, telecommunications, and consumer technology markets. With more than 1,300 analysts worldwide, IDC offers global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries. IDC's analysis and insight help IT professionals, business executives, and the investment community make fact-based technology decisions and achieve their key business objectives.

Global headquarters

One Beacon Street
Suite 33100
Boston, MA 02108
508-872-8200
x.com/idc
blogs.idc.com
www.idc.com

Copyright and trademark notice

Thank you