10 FAQs about agentic penetration testing and Savant Pathseeker

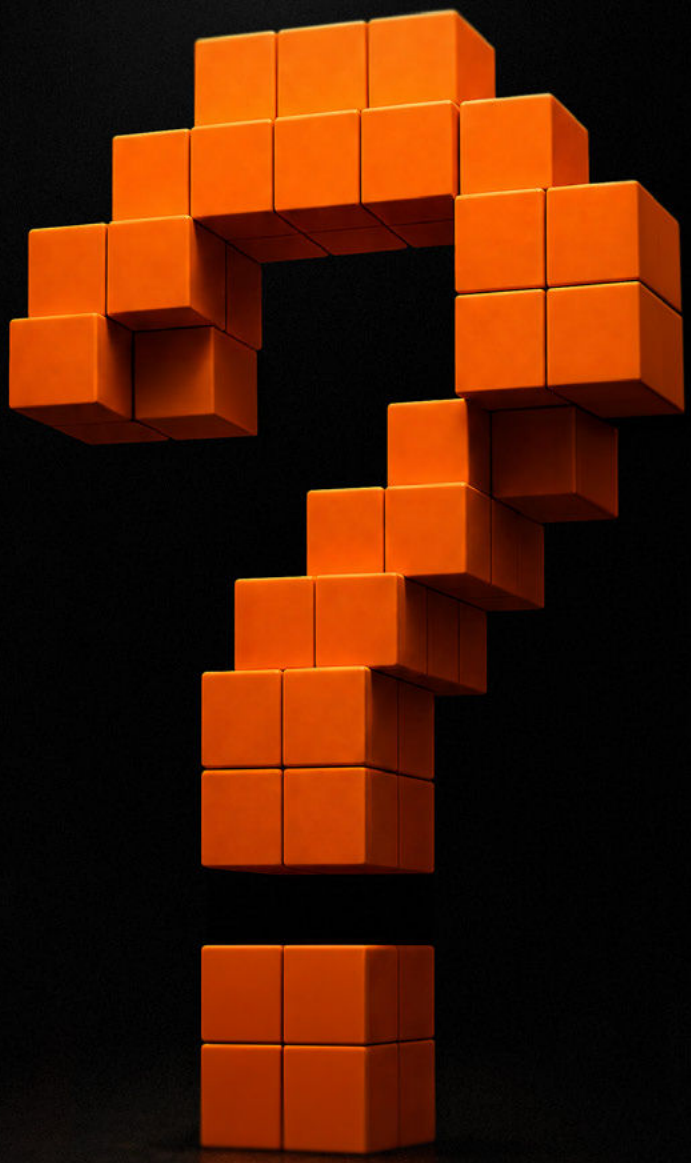

bugcrowd

For over a decade, penetration testing (aka **pen testing**) has been an indispensable tool in the security leader's toolbox. However, not all pen tests were made the same, and their effectiveness heavily relies on the details of their implementation.

A newer category, **agentic pen testing**, is changing what's possible. By deploying AI agents that simulate real attacks and assess exploitability, organizations no longer have to choose between deeply testing their crown-jewel assets and barely testing everything else.

Agentic pen testing 101

1 What are the phases of an agentic pen test?

Traditional pen tests are broken down into four phases:

Pre-engagement activity

Reconnaissance and planning

Vulnerability mapping phase

Exploitation

Agentic pen testing has begun to reshape the early phases of this lifecycle. AI agents can compress or automate reconnaissance, vulnerability mapping, and initial exploitation, completing the process in minutes or hours instead of days. **This also allows companies to best leverage human testers for more complex tasks,** as they can use the agentic pen testing findings to focus on chaining individual flaws into multi-step attack paths and producing the nuanced, defensible reporting that audits and regulators expect.

2 How does agentic pen testing increase coverage?

There's a less-discussed reason organizations need to rethink pen testing: most assets never get tested at all. This isn't because security teams don't want to test them. It's because there aren't enough humans, hours, or budget to do it. A handful of crown-jewel applications get intensive, if infrequent, manual attention. Everything else is left to vulnerability scanning, or nothing at all.

> **The result is a security program that looks complete on paper but leaves most of the real attack surface unvalidated in practice.**

With agentic pen testing, companies can continuously test all their assets around the clock, multiple times a year. As a result, organizations can have baseline assurance across the entire attack surface, not just the crown jewels, at a cost and speed that doesn't force a choice between depth and breadth. **This allows organizations to shift from reactive to pre-emptive security, catching vulnerabilities before attackers can exploit them.**

3 What are the pros and cons of agentic pen testing?

Where a vulnerability scanner tells you something might be wrong, agentic testing confirms whether it's actually exploitable—at machine speed and a fraction of the cost of manual testing. It complements human testing by examining all assets for common issues, allowing human testers to apply their skills to more complex and high-stakes targets.

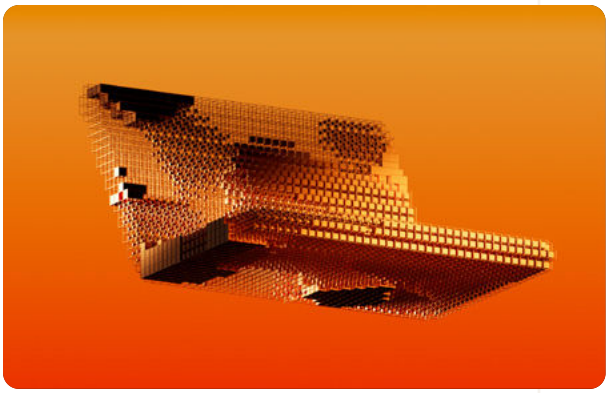

Pros

- ✓ Tests every asset at a fraction of the cost of manual testing
- ✓ Produces evidence of exploitability, not just a list of potential flaws
- ✓ Runs in minutes or hours rather than days
- ✓ No scheduling lag
- ✓ Consistent and repeatable coverage
- ✓ Complements human testing

Cons

- × Capabilities still developing for complex business logic flaws, exploit chaining, and zero-days in dynamic applications
- × Effectiveness depends on scope quality and asset inventory accuracy
- × May not meet auditor requirements for some controls (e.g., PCI-DSS) on their own
- × A newer category with evolving methodology standards

Where agentic pen testing fits compared to other security solutions

4 Does agentic pen testing replace vulnerability scanning?

Agentic pen testing uses orchestrated AI agents to autonomously test internet-accessible web applications and APIs for common, exploitable vulnerabilities. Rather than simply flagging known issues, agentic testing actually attempts exploitation (reconnaissance, vulnerability mapping, exploitation, and reporting) and provides evidence of whether an attack succeeded.

> **It's a fundamentally different category from scanning: a scanner tells you something might be wrong; agentic testing tells you whether it actually is.**

5 Does agentic pen testing replace human pen testing?

Although agentic pen testing helpfully automates certain repetitive tasks, it's not a complete replacement for human pen testers. Complex business logic flaws, novel exploit chains, zero-days, and post-exploitation attack paths still require human adversarial creativity and judgment to uncover, and regulators and auditors still expect the kind of nuanced, defensible reporting that comes from a skilled human tester reasoning through findings in context. Agentic testing covers the assets and time windows humans can't reach. The goal is to make sure human expertise is spent on the problems that actually need it.

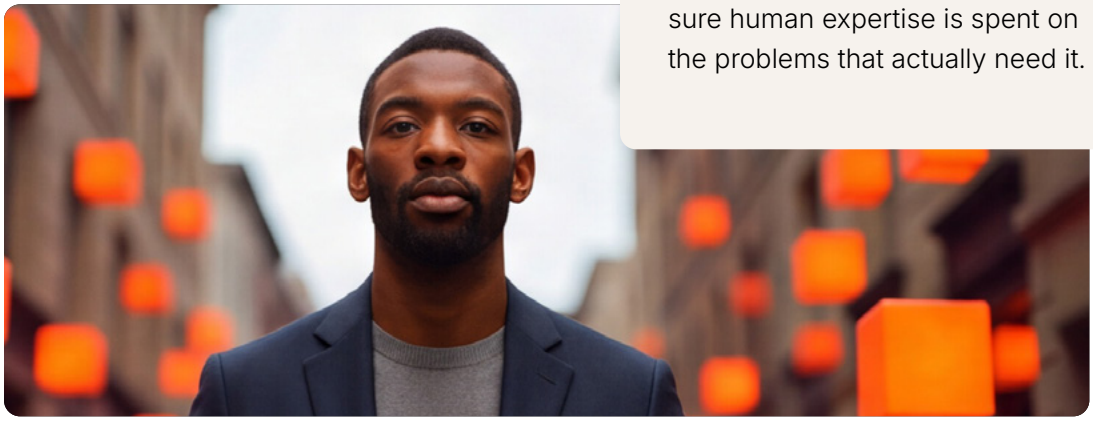

6 What is the difference between scanners, agentic pen testing, and human pen testing?

The best way to understand the difference between scanning, agentic, and human pen testing is to think of them as answering different questions.

Here's how to think of the difference, using a comparison from healthcare:

Vulnerability scanning | **Cholesterol check**

It flags risk factors, but it doesn't tell you whether something is actually wrong.

Agentic pen testing | **Diagnostic test**

It tells you whether the disease (aka the exploitable vulnerability) is actually present.

Human pen testing | **Specialist consult**

It can help you make sense of all the results and figure out what to do next (e.g., how bad your heart is, what the treatment looks like, etc.).

The strongest preemptive security programs layer all three approaches rather than picking one:

- ✓ **Vulnerability scanning as an even broader (if shallower) first pass**
- ✓ **Continuous or on-demand agentic testing across the full attack surface to understand where to go deeper**
- ✓ **Periodic human-led pen testing or bug bounty on the assets and attack chains where depth and judgment matter most.**

None of these replaces the others. They each answer a different question, at a different depth, for a different part of the attack surface.

Savant Pathseeker

7 What is Savant Pathseeker?

Savant Pathseeker is Bugcrowd's new agentic pen testing offering.

> **It operates autonomously to find common vulnerabilities in internet-accessible web applications and APIs along with evidence of exploitability at machine speed and scale.**

It's purpose-built to complement and inform the human-led pentesting, bug bounty, and red team engagements already available on the same platform, handling continuous baseline testing so humans can focus on the complex, high-value logic flaws and novel attack chains that require their expertise.

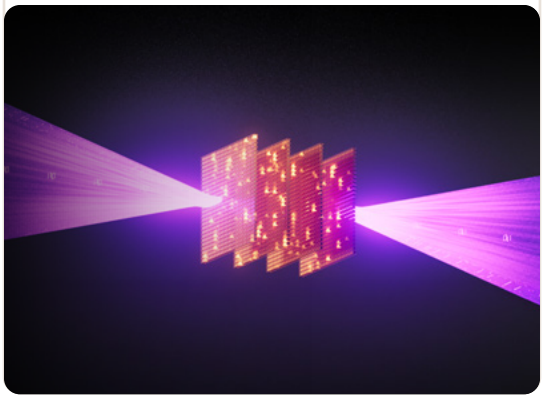

8 How is Savant Pathseeker different from an automated vulnerability scanner?

Savant Pathseeker is not a scanner with an AI layer bolted on. Scanners run pre-defined checks and flag theoretical issues. Savant Pathseeker uses purpose-built agentic systems that plan, probe, and attempt actual exploitation across web applications and APIs, delivering evidence-based findings with reproducible proof of exploitability, not just a list of potential flaws.

The technical differentiation sits in a few places:

- Proprietary skills and orchestration is layered on frontier models, developed by Bugcrowd practitioners with combined decades of offensive testing experience. This is built specifically for dynamic applications, including autonomous API fuzzing and attack path reasoning.
- Customers can natively integrate into a single platform that already unifies asset discovery, continuous testing, and human-led offensive testing like PTaaS, bug bounty, VDP, and red teaming. Findings from agents feed directly into the same risk surface and workflows as the human results, so they can be correlated, prioritized by exploitability, and escalated without tool-switching.
- Savant Pathseeker has built-in operational controls, scope guardrails, a manual kill switch, and audit-ready reporting, designed for continuous enterprise-safe use rather than ad-hoc scanning.

9 How does Savant Pathseeker differ from other providers?

There are a lot of agentic pen testing providers entering the market. Many of these vendors rely on commoditized capabilities. Bugcrowd is built on a platform ecosystem that pairs agent speed with human offensive validation to prioritize complex vulnerabilities in dynamic application based on actual weaponization risks.

Another way Bugcrowd is different from these providers is its ability to find the "unscannable." Autonomous-only tools excel at coverage, but they can miss complex business logic, zero-days, and exploit chains in dynamic applications. Bugcrowd pairs AI with human adversarial creativity, which delivers more coverage and depth.

> **The final difference is AI-only pentesting vendors are still building the operational foundations enterprises require.**

Bugcrowd brings 12+ years of proven enterprise infrastructure (SSO, RBAC, audit logging, compliance reporting, SLA-backed delivery, and dedicated customer success) so security teams get a trusted partner ready to support mission-critical programs from day one, not a vendor still figuring it out.

10 How should I combine Savant Pathseeker with a bug bounty program?

Both bug bounty programs and pen testing take a focused, strategic approach to the discovery and assessment of vulnerabilities and greater security risks. For this reason, many organizations find that a layered strategy combining several methods provides the best results.

By layering agentic testing, human-driven pen testing, and bug bounty for compliance and risk reduction, organizations can build a strategy that combines the following:

- **Continuous agentic coverage**
  Validated, evidence-backed testing of common, exploitable flaws across every web app and API, at machine speed. This is what some might consider a baseline pen test.
- **Ongoing vulnerability discovery and assessment**
  When the exploitability of vulnerabilities is confirmed, this is what some might consider a "basic" pen test.
- **Periodic, human-driven pen testing to find common flaws and meet compliance**
  Skilled human testers finding business logic flaws and chained exploits on crown-jewel assets. This is what some might consider a "standard" pen test.
- **A continuous bug bounty running "over the top"**
  Ongoing discovery that picks up emerging vulnerabilities not yet detectable using the prior three methodologies.