bugcrowd

Ultimate Guide to **offensive security testing**

IN THE PUBLIC SECTOR

Table of Contents

Introduction

This guide provides an overview of the modern security landscape and its current challenges, specifically in the public sector, and ways offensive security testing is being leveraged today as a solution.

For public sector organizations, cybersecurity has gone from being a technical concern to being **a central part of their operational strategies and national security priorities.** Government agencies reported a 40% increase in targeted attacks, while vulnerability submissions to government sector programs surged 151% over the past year. While the expanded attack surface presents significant risks, the increasingly connected world also offers unprecedented opportunities to strengthen our defenses. The same global reach that creates vulnerabilities also provides access to a worldwide community of security researchers who can help governments stay ahead of new threats. To make the most of the diverse talents of this community, organizations need to position themselves as partners and allies of the broader security community while maintaining the stringent security standards that public sector operations demand.The same global reach that creates vulnerabilities also provides access to a worldwide community of security researchers who can help governments stay ahead of new threats. To make the most of the diverse talents of this community, organizations need to position themselves as partners and **allies of the broader security community** while maintaining the stringent security standards that public sector operations demand.

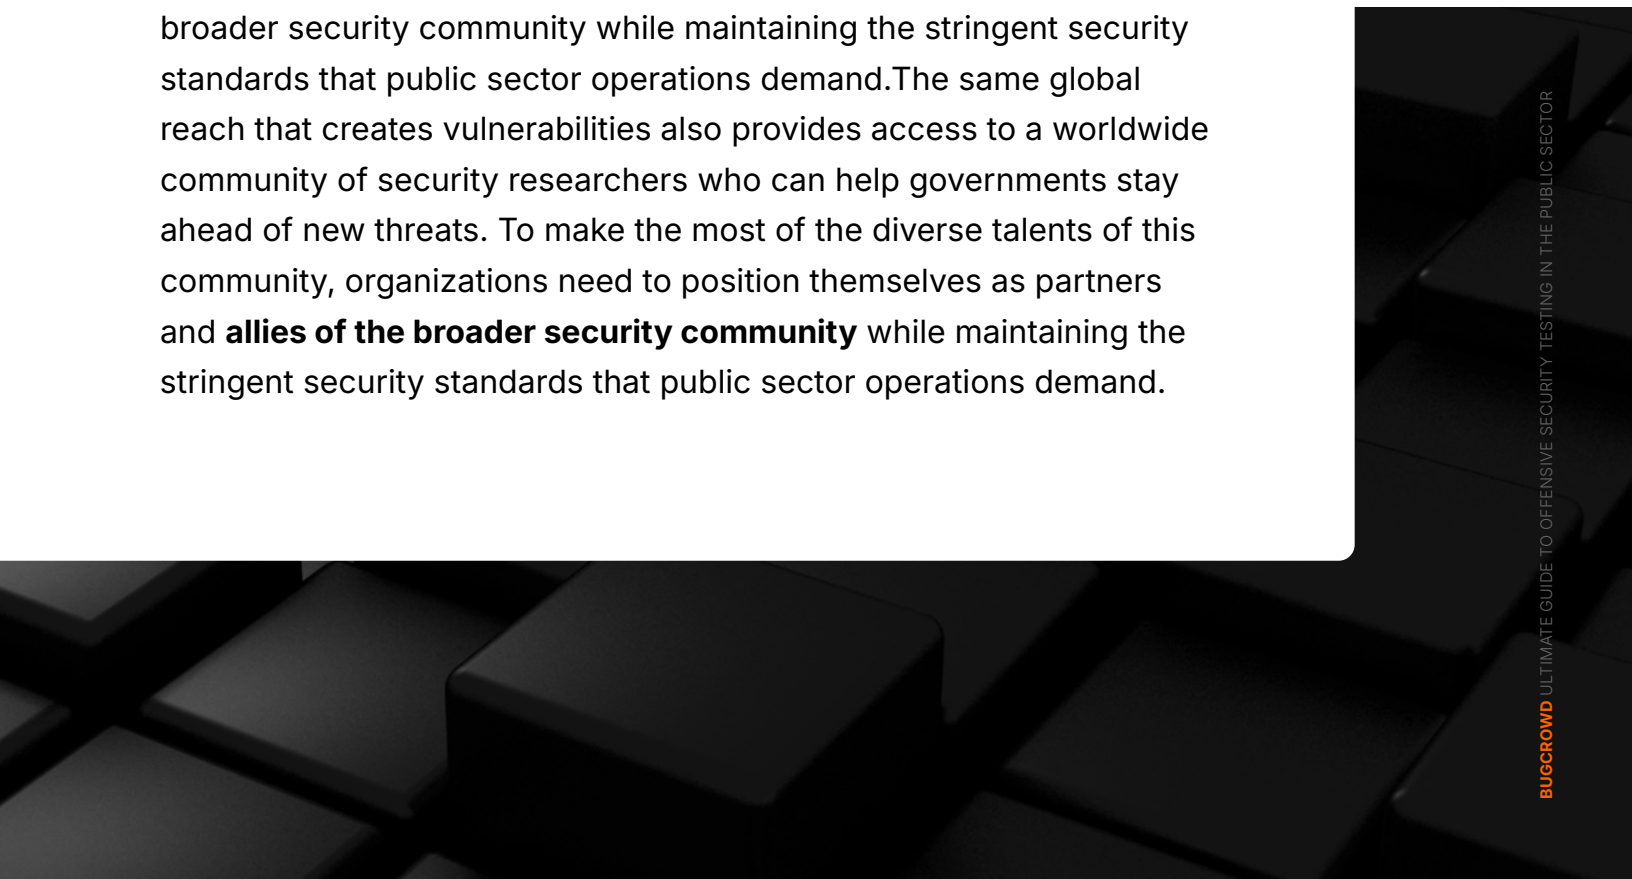

What is crowdsourced security?

Crowdsourced security is a type of offensive security testing that secures digital assets that draws from **the collective skill and experience of the world's community of security researchers.** These highly capable individuals are given the direction, scope, and incentives they need to identify and report vulnerabilities. To carry out their tasks, they effectively simulate the varied techniques employed by threat actors.

Traditional security tends to be an ad hoc administrative arrangement heavy on consultant hours. However, crowdsourced security relies on the **wisdom of the crowd**, a phenomenon in which large groups of people are collectively smarter than individual experts. A group of security researchers can make discoveries and identify opportunities more effectively than even the most capable and expert individuals on an internal security team. In nature, this phenomenon is reflected in the herds of animals that are more effective at finding food and shelter than, say, the lone wolf. In security, this means that crowds of security researchers can **find and identify bugs faster** than over-burdened internal teams and dynamic attackers working alone.

Crowdsourced security leans into **community and collaboration**, which is why researcher-powered security is so powerful. Working with a crowdsourced platform like Bugcrowd gives organizations and IT departments access to the widest pool of talent and allows them to broker interactions with security researchers and triage responses so that security teams only have to pay for results.

Crowdsourced security is a part of the preemptive security movement, which has grown massively, providing organizations and public sector agencies with access to the **world's best security minds** before attackers can strike. This has allowed them to identify and rectify security challenges at speeds previously thought unimaginable.

Crowdsourced security in the public sector

Bringing crowdsourced security to the public sector has taken a little longer than other sectors. The Computer Fraud and Abuse Act (CFAA) put into force in 1984 has significantly shaped the landscape of cybersecurity research by prohibiting unauthorized computer access. While essential for prosecuting cybercrimes, the Act's broad language initially created uncertainty for security researchers, as their work could be interpreted as "exceeding authorized access." A pivotal shift came in 2021, when the Supreme Court's Van Buren v. United States decision narrowed this interpretation, followed by Department of Justice's guidance **explicitly protecting good-faith security research.** This evolution in the CFAA's interpretation has helped create a clearer legal framework for ethical hacking, enabling **more effective collaboration** between security researchers and organizations seeking to improve their cybersecurity. Additionally, it has forged a path for crowdsourced security to play a larger role.

Beyond legislation, the public sector has embraced bug bounty programs. In 2016, the first US federal government bug bounty program, Hack the Pentagon, was launched. Since then, the federal government has also introduced the bipartisan Hack the Department of Homeland Security Act.

To increase the ease of collaboration, public sector organizations can choose from several tiers of security testing programs with Bugcrowd based on their security requirements and data sensitivity. While some agencies opt for private programs open only to specially selected researchers vetted by Bugcrowd's extensive background checks and security clearances, others may choose broader programs that engage a wider pool of vetted security experts on the platform. Though the idea of opening government systems to external testing may seem daunting or potentially dangerous, it's important to remember that government assets are already subject to constant probing by potential adversaries. Security operations centers across federal, state, and local agencies regularly observe persistent scanning of their ports and applications.

> **The key difference is that structured Bugcrowd programs channel expertise toward improving defenses rather than exploiting weaknesses.**

Trusted by federal agencies for mission-critical security

In 2026, Bugcrowd achieved FedRAMP Moderate Authorization, sponsored by CISA. This milestone validated that Bugcrowd meets the U.S. government's stringent security requirements for protecting sensitive data via the Bugcrowd Platform.

FedRAMP

Learn more about this authorization here.

Types of offensive security testing in the public sector

Offensive security testing solutions are just like any other security solution in the sense that they change dynamically according to the needs of an industry. At present, the most popular solutions that draw from distributed security talent are:

- Penetration Testing/Pen Testing as a Service (PTaaS)
- Vulnerability Disclosure Programs
- Bug Bounty Programs
- Red Team as a Service
- AI Safety and Security Testing

VDP

Vulnerability disclosure program

A VDP is a **structured framework** that allows and invites security researchers to submit vulnerabilities they discover in an organization's digital infrastructure to the organization directly. These programs offer clear guidance on how security researchers can bring vulnerabilities to the attention of an organization. If the submission process is done correctly, organizations will often disclose these vulnerabilities to give credit to the security researchers who took the time to help them.

Security starts with understanding where risks and problems lie. VDPs are essential in facilitating this understanding.

For some public sector organizations, VDPs have moved **from best practice to requirement.** The Cybersecurity and Infrastructure Security Agency's (CISA) Binding Operational Directive (BOD) 20-01 requires all Federal Civilian Executive Branch agencies to develop and publish a VDP. This directive mandates that agencies maintain a public vulnerability disclosure webpage, provide secure methods for receiving vulnerability reports, and commit to not pursuing legal action against good-faith security researchers.

To support agencies in meeting these requirements, CISA established a federal VDP platform in partnership with Bugcrowd and EnDyna. The federal VDP platform goes beyond a basic VDP and integrates with agencies' development and security workflows, provides comprehensive triage services, and generates required metrics for BOD 20-01 reporting.

> **The platform enables scaled testing that uncovers vulnerabilities scanners might miss by engaging the global security researcher community.**

This VDP-as-a-service platform has demonstrated a ton of success. In just one year, over 40 agency programs received more than 1,300 valid vulnerability submissions, with nearly 200 critical findings. It has seen a **25-fold increase in vulnerability reports** from before, and 84% of these vulnerabilities were remediated within an average of 38 days.

Bugcrowd has also helped various government agencies set up and list their own VDPs, such as those for NASA, the Department of Veterans Affairs, and the State of California. Bugcrowd can help with triaging as well, so agencies can send the most critical findings to their engineering teams quickly. These programs benefit from external management and triage support, ensuring critical findings reach engineering teams quickly while building productive relationships with the security researcher community.

How to approach responsible disclosure in the public sector

Responsible disclosure refers to the standardized process of receiving, acknowledging, and addressing vulnerability reports from security researchers. For researchers, this means following established protocols for reporting vulnerabilities through authorized channels and allowing adequate time for remediation. For government agencies, this means maintaining clear communication channels and following predetermined assessment and remediation timelines as outlined in their VDPs.

The process requires careful adherence to security protocols and compliance requirements. Agencies must balance transparency with security considerations while protecting sensitive information. When vulnerabilities may affect multiple agencies or critical infrastructure, coordination with CISA or sector-specific bodies is often required, following predefined reporting matrices and response protocols based on severity and scope.

Agencies that ignore submissions, dismiss legitimate concerns, or take adversarial legal stances against good-faith actors risk alienating the security research community and undermining their cybersecurity mission. Clear safe harbor provisions and well-defined legal protections are essential for fostering a cooperative security environment that serves the public interest.

Bug bounty programs MBB

Bug bounty programs are **result-focused security initiatives** that incentivize security researchers to uncover and report security vulnerabilities within an organization's digital infrastructure. Typically building on the VDP framework, they add a financial reward based on the criticality of the vulnerabilities identified and remediated. Bug bounty programs are the original and most widely used crowdsourced cybersecurity solutions. They can ensure the **rapid evaluation and remediation of novel threats,** such as when new zero-day vulnerabilities emerge.

Bug bounties attract top talent quickly. Although most security researchers believe it's more important to report a vulnerability than profit from it, cash rewards pique many security researchers' interest in a system. Bug bounties are cost-effective too. Hiring full-time security staff is expensive, and it's not realistic to continually train them on the latest threats. With bug bounties, agencies only pay for results, and they get access to highly specialized talent without a lot of overhead. The dynamic pricing model ensures that **critical vulnerabilities** receive priority attention while helping organizations allocate security budgets effectively and identify frequently targeted assets.

Bug bounty programs have proven **remarkably effective in strengthening government security postures,** with the Department of Defense's Hack the Pentagon initiative serving as a prime example of public sector cybersecurity.

Launched in 2016 as the federal government's first bug bounty program, Hack the Pentagon demonstrated that security researchers can rapidly identify critical vulnerabilities that had escaped the detection of traditional security measures. Program rewards typically range from $100 to $15,000, creating a cost-effective model where agencies only pay for validated results.

> **The program's success has led to dozens of subsequent initiatives across military branches and federal agencies that have consistently delivered impressive results.**

These programs have been highly effective; the Chief Digital and Artificial Intelligence Office (CDAO) reported finding critical vulnerabilities within the first 24 hours of every bounty program launched, even for systems that undergo regular automated scanning and penetration testing. In one notable instance, a single four-hour bug bounty event uncovered more vulnerabilities than the entire duration of a traditional vulnerability disclosure program. This rapid discovery rate, combined with flexible reward structures that can be adjusted in real time based on findings, allows agencies to efficiently allocate resources to their **most pressing security challenges.** For public sector organizations facing persistent threats from nation-states and cybercriminal groups, bug bounty programs offer a proven way to **proactively identify and address vulnerabilities** while fostering collaboration with the broader security research community.

Penetration testing

PTAAS

Pen tests are security tests in which security testers **mimic real-world attacks** to identify methods of circumventing the security features of an application, system, or network that are failing to protect vital assets. Pentesters operate as a team, working within a defined scope for a set time period and completing each engagement by offering a report of the vulnerabilities detected.

Crowdsourced pen tests are a **new take on a longstanding security service,** offering dynamic new functionalities that make the most of talent accessed and findings integrated to advance software development. They can provide targeted and detailed assessments of digital assets and infrastructure quickly and efficiently while meeting regulatory compliance needs, just as traditional pen testing does. For public sector organizations facing sophisticated threats from nation-states and cybercriminal groups, crowdsourced pen testing offers comprehensive security evaluations.

RTaaS

Red teaming

Red teaming is an exercise where an offensive security team (the red team) covertly simulates an attack on a company's systems, looking for any weaknesses to gain access and simulate impact. The open-scope and covert nature of red teaming lets it emulate threat actors, who similarly would strike without warning and would gain access by any means necessary.

The internal team responsible for preventing, identifying, and responding to breaches (the blue team) responds to the red team attack as if it were a real intrusion. In most cases, the blue team doesn't know about the red team's activities. A variant called purple teaming combines the blue and red team into one: the red team calls out weaknesses as it spots them and the blue team can monitor and respond immediately.

Red teaming recently joined the ranks of crowdsourced offensive security methods. Previously, red teams were internal teams or teams from consultancies. Now, the global talent of the Crowd enables a new delivery method: red teaming as a service via crowdsourced red teams.

Crowdsourced red teams bring the same benefits of other crowdsourced offensive security methods: flexibility to your systems, security posture, and needs. Companies can work with one crowdsourced red team for an extended period of time, or they can work with multiple red teams in a setup similar to a bug bounty (they can also set up a model in between these two).

> **Combined with the usual power of red teams—close emulation of threat actors—crowdsourced red teams can help companies identify detection and response gaps and improve incident response without toiling over finding the right red team.**

This progression from VDPs to bug bounties to crowdsourced pen testing and red teaming aligns with BOD 20-01 and BOD 22-01 requirements, which explicitly mandate VDPs for federal agencies. Each stage builds upon the previous one, offering increasingly sophisticated ways to proactively identify and address vulnerabilities.

An emerging use case for crowdsourced security in the public sector:

AI safety and security

You can't talk about technology and security these days without mentioning AI. The emergence of AI in government systems certainly introduces new security considerations. Legislation like Executive Order 14147 requires **rigorous security testing of AI systems**, with particular emphasis on red teaming exercises.

Fortunately, crowdsourced security programs are particularly effective for AI systems. They provide access to diverse expertise across AI, security, and specific domains of knowledge. For instance, when testing large language models (LLMs), organizations can leverage security researchers' expertise in prompt engineering and AI behavior to identify potential vulnerabilities.

The adaptive nature of crowdsourced testing makes it especially valuable for AI security. While traditional security measures like automated prompt filtering may become less effective over time without continuous updates, crowdsourced testing **evolves with emerging threats.** Security researchers continuously develop new techniques and approaches. Crowdsourcing keeps up with changing technology and security needs, ensuring you're ready to tackle the issues associated with AI and the next wave of cutting-edge technology.

Examples of crowdsourced security success in the public sector

VDP success story: CISA VDP

In 2020, CISA issued BOD 20-01, requiring all Federal Civilian Executive Branch agencies to develop and publish vulnerability disclosure policies. However, implementing these programs proved challenging for agencies, especially those with small or no dedicated security teams, due to resource constraints. Agencies faced significant administrative overhead in handling disclosed vulnerabilities, triaging reports, corresponding with researchers, and meeting compliance requirements. In response, **CISA partnered with Bugcrowd** to launch the first federal civilian crowdsourced vulnerability disclosure platform.

The platform was quick to prove its use and effectiveness. Within its first 33 months, **over 40 federal agencies** had been onboarded, including NASA, the Department of the Treasury, and Homeland Security. They have since received more than 15,000 vulnerability reports. In 2022 alone, the program validated **1,330 unique vulnerabilities,** including 274 critical or severe findings, with 1,119 successfully remediated. The Department of Labor saw a dramatic improvement, stating "We went from very little activity to a lot of activity, just by joining the VDP platform."

The success continued to grow, with participating agencies seeing 67 more vulnerability submissions in their first quarter compared to non-participating agencies, and 90% of all submissions in Q4 of 2023 came through the platform. Agencies using the platform validated submissions two days faster than non-participating agencies, resulting in an estimated **$4.45 million saved** in potential remediation costs for critical and severe vulnerabilities. The researcher community expanded significantly, with 1,718 researchers submitting reports in 2023—nearly double the number from the previous year.

AI security success story: Department of Defense Chief Digital and Artificial Intelligence Office (CDAO)

As AI applications become increasingly central to government operations, the Department of Defense's CDAO faced the critical challenge of ensuring both the security and fairness of their AI-enabled systems. The existing testing and security measures proved insufficient for detecting data bias and other AI-specific vulnerabilities in these **complex systems.** Recognizing this growing problem, CDAO partnered with Bugcrowd and ConductorAI to implement an innovative crowdsourced security program.

The program conducted public adversarial testing of LLM systems for bias on behalf of the Department of Defense. The program utilizes Bugcrowd's CrowdMatch technology to source and activate researchers with specialized skills, implementing a reward-for-results model where participants are paid based on the successful demonstration of impact. The program represents a significant milestone, as it conducted **the first public adversarial testing of LLM systems for bias** within the Department of Defense, establishing a proven framework for AI security testing in the public sector and paving the way for future AI testing and innovation.

VDP success story: The Office of the Minnesota Secretary of State

Facing the challenge of protecting sensitive information for a diverse range of constituents, including voters, political candidates, and business owners, the Office of the Minnesota Secretary of State sought to proactively reduce risk exposure by innovating in security. It decided to engage with the security researcher community through a VDP, recognizing that its systems were already constantly being probed and that there were significant benefits to proactively leveraging the skills of security researchers.

Partnering with Bugcrowd, the Office implemented a VDP that **quickly yielded positive results.** The program led to the discovery of hidden high-impact vulnerabilities, reduced noise in its security operations, and helped build productive relationships with the security researcher community. The Office achieved its goals of receiving legitimate, actionable vulnerability reports, with submissions being triaged in an **average of 1.8 days.**

Legal requirements for crowdsourced security usage in the public sector

The U.S. government has established various requirements for implementing crowdsourced security programs, with the most comprehensive being CISA's BOD 20-01 for federal civilian agencies. While BOD 20-01 mandates VDPs for Federal Civilian Executive Branch Agencies, other initiatives like Executive Order 14110 address emerging technologies such as AI.

BOD 20-01

CISA's Binding Operational Directive (BOD) 20-01 requirements:

- All Federal Civilian Executive Branch agencies must develop and publish a VDP.
- Agencies must maintain a public vulnerability disclosure webpage.
- Agencies must provide secure methods for receiving vulnerability reports.
- Agencies must commit to not pursuing legal action against good-faith security researchers.

For a full list of additional requirements, check here.

BOD 20-01

BOD 22-01 focuses on reducing the significant risk of known exploited vulnerabilities. CISA's Binding Operational Directive (BOD) 22-01 requirements:

- Within 60 days of issuance, agencies shall review and update agency internal vulnerability management procedures in accordance with this Directive.
- Remediate each vulnerability according to the timelines set forth in the CISA-managed vulnerability catalog.
- Report on the status of vulnerabilities listed in the repository.

Executive Order 14179

Executive Order 14179, focuses on removing barriers to American leadership in artificial intelligence. The executive order aims to initiate the process of strengthening U.S. leadership in artificial intelligence, promote AI development free from ideological bias or social agendas, establish an action plan to maintain global AI dominance, and to revise or rescind policies that conflict with these goals.

The benefits of crowdsourced security

Crowdsourced security offers organizations **more expert eyes** on their infrastructure than possible with internal teams or consultants alone. Specifically, it delivers three key advantages: continuous coverage, cost efficiency, and rapid response to emerging threats.

Continuous coverage

By engaging global talent, organizations gain security support around the clock. While threats and malicious actors don't operate on business hours, a worldwide community of security researchers provides continuous coverage of assets.

Cost efficiency

The pay-for-results model of bug bounty programs offers a clear return on investment (ROI). Organizations only pay for validated findings rather than investing in tools or services upfront. This creates a transparent market that helps align security investments with actual risks and priorities.

Rapid response to emerging threats

During the critical Log4J vulnerability discovered in December 2021, Bugcrowd's platform saw nearly 300 submissions within two days of the announcement, with critical vulnerabilities addressed in **under three hours**—a speed impossible for traditional security teams. This rapid scalability makes crowdsourced security particularly valuable for both planned events like product launches and emerging security challenges that require immediate attention.

About the Bugcrowd Platform

Bugcrowd's FedRAMP-authorized Platform brings the right crowd into your security workflows at the right time, allowing you to run bug bounties, pen tests, red teaming engagements, VDPs, and more at scale and in an integrated, coordinated way.

Bugcrowd uses its proprietary CrowdMatch AI to match qualified, trusted security researchers to your individual security needs, as well as provide rich reports and analytics to offer continuous insights into trends in findings, payments, criticality, and more.

This allows us to find the perfect security researcher talent for goals like pen testing, bug bounty, and vulnerability intake and disclosure, as well as to ensure the scalability and adaptability that come with a functional talent platform.

Contract Vehicles for Procurement of Bugcrowd

- NASA-SEWP
- Hack DHS IDIQ
- GSA via Carahsoft
- Tradewinds AI Marketplace

The Bugcrowd Platform is an AI-powered, multi-solution platform built on the industry's richest repository of vulnerabilities, assets, and security researcher profiles curated over a decade.

Bugcrowd's team of global security engineers acts as an extension of your security team, validating and triaging submissions to enable critical vulnerability resolution within hours. Native integration with your development tools and processes ensures continuous testing of apps and APIs throughout the development life cycle, catching vulnerabilities before they reach production.

After over a decade at the forefront of crowdsourced cybersecurity crafting solutions for thousands of customers, Bugcrowd brings an extensive repository of data to discovery and remediation, as well as intangible knowledge on the mindsets and attitudes of the world's security community.

See the Bugcrowd Platform in action

Take a 5-minute tour to get an overview of how the Bugcrowd Platform connects you with trusted hackers to help you take back control and stay ahead of attackers.

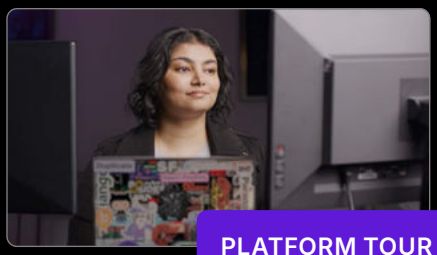

PLATFORM TOUR ▶

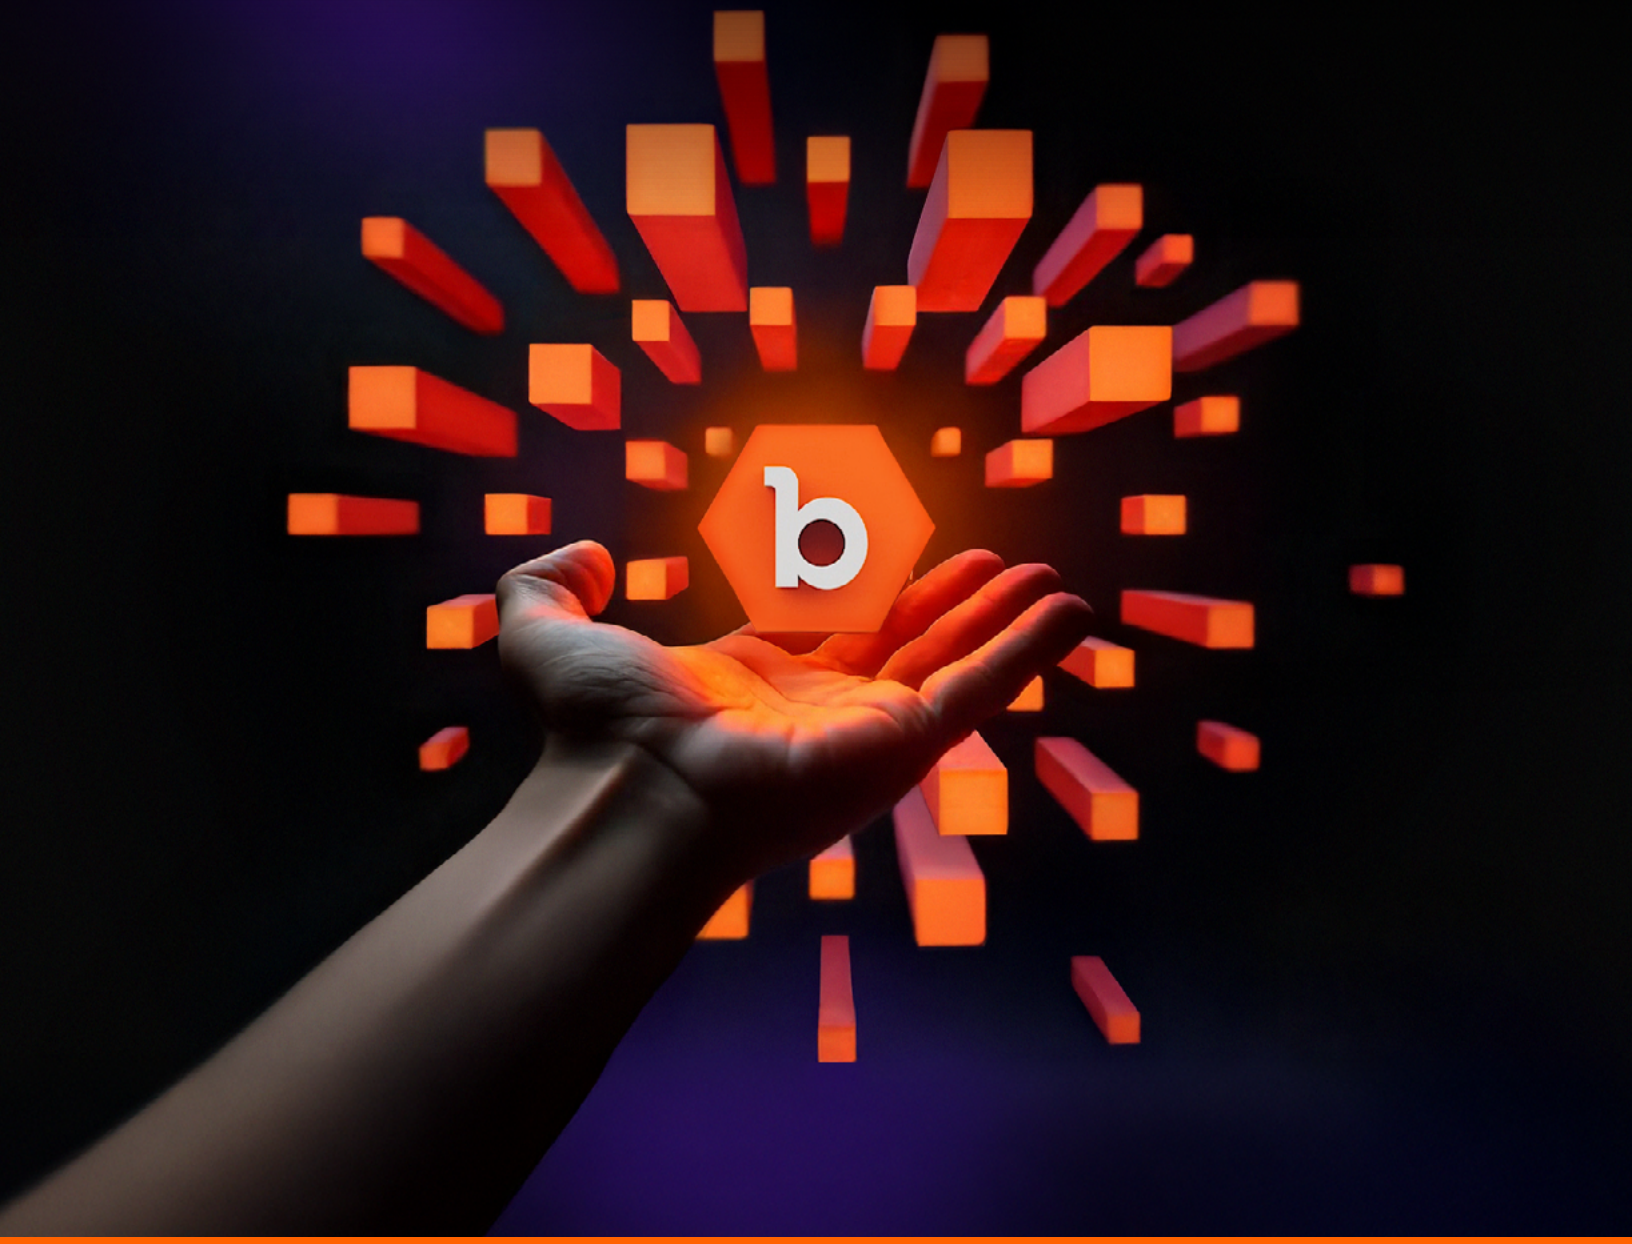

Unleash Human Creativity for Proactive Security

TRY BUGCROWD