

❄️ Content frozen ❄️

Inside the Platform: Bugcrowd's Vulnerability Trends Report

Vulnerability insights from the crowdsourced security ecosystem for CISOs and security leaders

Letter from the Editor

An Introduction from Bugcrowd CISO Nick McKenzie

Every year, Bugcrowd conducts landmark research on the global vulnerability landscape and publishes its findings for the benefit of security leaders. As an industry, we're truly on the precipice of so many changes, and the goal of this report is to arm security leaders and practitioners alike with the necessary trend information, data, and expert predictions to prepare for these changes.

I'm balancing a lot of priorities as the CISO of Bugcrowd, so believe me when I say I know how hard it is to keep a pulse on everything going on in your organization—let alone what's happening in the industry more broadly. That's why I'm thrilled to share this report, wherein you'll find that my team and I have done much of that groundwork for you. Leveraging vulnerability data from the last 12 months, this report offers critical context, insights, and opportunities for security leaders looking for new information to bolster their risk profiles.

Throughout the research process, I wasn't surprised to find that vulnerabilities are still

on the rise. When you combine an overall increase in rapid digitization (including new technologies that businesses are adding into business processes like generative AI) with more products boasting many new features, it's inevitable that you end up with an exponential increase in bugs.

Another insight from the report that I found especially telling is an increase in the trend toward favoring public crowdsourced security programs over private programs. More programs are dropping the clutch and shifting their gear to "public."

Looking ahead, we can use insights from this report in conjunction with other key learnings from the industry to predict what is coming next. Thinking holistically about risks and threats, I often look at what's happening publicly in terms of events and combine that with information on emerging technologies, people trends, and usage. Here are three predictions I have for the year ahead:

1. **Threat actors will use adversarial AI to speed up enterprise attacks.**

In general, security teams are now dealing with an increased number of events and more noise. With the use of AI increasing, I believe we'll see a higher volume of attacks, ultimately leading to more noise for those who are on defense to sift through.

2. Supply chain security, third-party risk, and inventory management will get hotter.

Coming off the back of high-profile events and breaches that occurred over the past couple of years, this topic will permeate into the future, pushing with it a stronger focus on third-party risk management practices, quality software bills of materials (SBOM), tooling integrations with various governance, risk and compliance (GRC) requirements, and IT asset management-type tools.

3. We will see a stronger focus on the human factor.

This will come in many forms, such as controlling malicious insiders and

preventing accidental or unintentional control failures, like the actions of misguided employees falling prey to social engineering or focusing on improving application security and development/coding practices. To counter the cyber talent skills gap and help their security teams scale, organizations will more broadly adopt the crowdsourcing of human intelligence to continuously weed out unique or previously unidentified vulnerabilities.

Every organization's risk and threat profile is unique. Challenges can arise as a result of anything from an organization's business model to its IT footprint, industry vertical, people (costs, scarcity, and skills shortage), security maturity, and geographical sprawl. With this in mind, Bugcrowd's goal in publishing this annual vulnerability research report is to arm security leaders with key information about trends, which they can apply to their unique challenges.

Report highlights and methodology

- The most successful programs were those that offered higher rewards (e.g., \$10,000 or more for P1 vulnerabilities).
- Programs with open scopes received 10x more P1 vulnerabilities than those with limited scopes.
- The government sector experienced a 151% increase in vulnerability submissions and a 58% increase in the number of P1s rewarded this year compared to last year.
- The financial services industry and government sector offered the highest median payouts for P1 vulnerability submissions.
- A new AI-related category was added to Bugcrowd's [Vulnerability Rating Taxonomy](#) (VRT).

Methodology

In preparing this edition of *Inside the Platform*, millions of proprietary data points and vulnerabilities were analyzed. These data were collected from across thousands of programs on the Bugcrowd Platform from January 1, 2022, to October 31, 2023. When referring to “this year,” we imply measurements taken from January 1, 2023, to October 31, 2023.

What millions of vulnerabilities tell us about the year to come

With every annual edition of this report, we analyze activity on the Bugcrowd Platform to identify larger cybersecurity trends and better understand the specific **challenges that security leaders face**. The Bugcrowd Platform is a multi-solution crowdsourced security platform that provides the scalability and adaptability needed to proactively safeguard organizations from increasingly sophisticated threat actors. It is built on **the industry’s richest repository of insights** into vulnerabilities, assets, and hacker profiles, which have been curated over the course of more than a decade.

Bugcrowd connects organizations with **trusted hackers** (aka ethical hackers, security researchers, or white hat hackers) to **proactively defend** their assets against sophisticated threat actors. Through solutions like penetration testing as a service, managed bug bounty, and vulnerability disclosure programs (VDPs), organizations can **unleash the collective ingenuity of hackers** to better mitigate risks across all their applications, systems, and infrastructure.

This report offers a glimpse into the millions of proprietary data points behind the Bugcrowd Platform, looking at hacker

vulnerability submissions from every possible angle to truly understand what vulnerability trends tell us about the future of cybersecurity. For the third year in a row, we’ve seen **growth in a number of key metrics**, driven by the security demands of hybrid workplaces and threat actors’ adoption of generative AI as a tool. This article breaks down these metrics and comments on why these trends exist.

Overall submissions, critical submissions, and payouts

At the start of the decade, global lockdowns and the associated uptick in remote work came with an unsurprising spike in vulnerability reports. The increase in submissions in 2023 (up a double digit percentage compared to the same timeframe in 2022) is testament to **the value of sustained investment in crowdsourced security**. It is also indicative of the continued expansion of work being carried out over the internet and the infrastructure around it, which continues to result in a **high number of vulnerabilities**.

Before looking at data around critical submissions vs. overall submissions, it is

helpful to understand how a rewarded submission is assigned a critical severity level. When a hacker sends in a submission, it is validated and checked to ensure it isn't a duplicate by Bugcrowd's global team of in-house application security engineers. This team adds **important context** to each hacker submission before it is triaged according to the [VRT](#), an open source framework for assessing, prioritizing, and benchmarking the severity of security vulnerabilities.

Since 2017, Bugcrowd has been the creator and maintainer of the VRT. The VRT was designed to be a simple-to-use, evolving method for assigning a severity level to a specific vulnerability class—and taking an open source approach to managing it enables us to keep our ear to the ground, ensuring that the taxonomy stays aligned with the market. Since the VRT's creation, **hundreds of thousands of vulnerability submissions** on the Bugcrowd Platform have been created, validated, triaged, and accepted by program owners who subscribe to this rubric.

We use the VRT as a **common point of reference** for setting the priority of submissions on the Bugcrowd Platform. Every vulnerability submitted is tagged with a category that has an associated technical severity. For a complete listing of categories, [visit the VRT page](#).

One key metric we track at Bugcrowd is **critical vulnerabilities**, which we call P1s.

Generally, P1 vulnerabilities offer the **highest payouts**. The incentive for hackers to report P1 vulnerabilities is more compelling than ever, as the average payout **increased by 7%** this year. This increase is one that we were expecting to see. We've observed [significant economic changes in the market](#) over the past few years, fueled by both greater investment in crowdsourced security and increased geopolitical uncertainty. Today, the most successful programs pay **\$10,000** or more for P1 vulnerabilities.

One of the main reasons for this increase in payouts is the **increased complexity and nuance** behind the findings, which is a function of hackers moving and growing beyond **"recon-style" bug hunting**. While still heavily prevalent today, low-hanging fruit isn't as easy to find as it was in years prior. In response, the most successful hackers have adapted to go deeper, as opposed to wider, knowing that the most impactful vulnerabilities still lie beneath the surface.

Automated tools are a standard part of the hacker's toolbox, but a **breadth of skills** is required to find the more interesting and high-impact vulnerabilities. On average, the level of effort required to find a critical vulnerability is much higher. But as the bar climbs, so do the incentives. Organizations that keep up with the latest market rates for vulnerabilities always **attract top talent**—there is simply no substitute when it comes to engaging the best hackers in any given area of expertise.

Table: Recommended reward ranges based on historical data available from the Bugcrowd Security Knowledge Platform at the time of publication

Severity Level per Vulnerability Rating Taxonomy (VRT)

	P1	P2	P3	P4
Low Range (Attracts Generalists) Best for: Untested web apps with basic credentialed access and no hacker restrictions (e.g., geolocation) new to crowdsourced testing—for any target with restrictions in place, rewards should default to one range higher.	\$3,500–\$4,500	\$1,500–\$2,500	\$500–\$750	\$175–\$225
Mid Range (Attracts Experienced Hackers) Best for: Well-tested web apps that have been part of longstanding crowdsourced programs, moderately tested APIs or mobile apps, and presumed-to-be-vulnerable thick clients/binaries and/or embedded devices	\$5,500–\$7,500	\$2,500–\$3,500	\$750–\$1,500	\$250–\$500
High Range (Attracts P1 Specialists) Best for: Extremely hardened and sensitive web apps, APIs, mobile apps, and moderate-to-highly secure thick clients/binaries and/or hardened embedded devices	\$11,000–\$20,000	\$3,500–\$7,500	\$1,000–\$2,500	\$300–\$600
Hardware Providers	\$5,000–\$10,000+	\$2,000–\$4,000	\$600–\$900	\$200–\$400
Cloud Providers	\$5,000–\$15,000+	\$3,000–\$5,000	\$1,000–\$2,500	\$250–\$700
Financial Services	\$8,000–\$20,000+	\$3,000–\$8,000	\$600–\$1,500	\$250–\$350
Cryptocurrency	\$50,000+	\$10,000–\$20,000	\$2,000–\$3,000	\$500–\$750

The number of P1 vulnerabilities Bugcrowd rewarded in 2023 aligns with the observations of other industry experts. According to a report by Statista, hackers discovered more than **25,000 new common IT security vulnerabilities and exposures (CVEs) this year**—the largest number reported in a single year to date.

In 2023, it was nearly impossible to ignore the news coverage and social media activity around critical bugs and hackers chaining vulnerabilities together to pull off bigger exploits. Some examples, along with descriptions from the National Vulnerability Database, include the following:

- [CVE-2023-027350](#)

This bug allows remote attackers to bypass authentication on affected installations of PaperCut NG 22.0.5 (Build 63914). Therefore, authentication is not required to exploit this vulnerability. This specific flaw exists within the SetupCompleted class, with the issue resulting from improper access control.

- [CVE-2023-34362](#)
This SQL injection vulnerability found in the MOVEit Transfer web application allows an unauthenticated attacker to gain access to MOVEit Transfer's database. Depending on the database engine being used, an attacker may be able to infer

information about the structure and contents of the database and execute SQL statements that alter or delete database elements.

- [CVE-2023-26360](#)
Adobe ColdFusion versions 2018 Update 15 (and earlier) and 2021 Update 5 (and earlier) are affected by an improper access control vulnerability that could result in arbitrary code execution in the context of the current user. Exploitation of this issue does not require user interaction.

Keeping all of this in mind, it's no surprise that [ICS2](#) recently found that **75% of security professionals** believe that the current threat landscape is the most challenging one in the past five years.

SIDE BAR

Navigating Changing Reward Ranges

Bugcrowd recently increased our suggested reward ranges to keep pace with changes in the industry. The last major update to the suggested reward ranges happened almost five years ago. A lot has changed in the past five years, including the following:

- Inflation—\$1 in 2018 is worth approximately \$1.21 in 2023.
- Increased rewards throughout the market—There is now competitive pressure from more programs offering higher rewards.

When it comes to setting reward ranges for your program, there are six guiding principles that I recommend keeping in mind:

1. Reward ranges are suggestions, not absolutes.

In some cases, a set of targets will call for higher rewards, and in other cases, lower. It all depends on your security maturity and your appetite for success.

It's important to remember that everyone is vulnerable at the right price point—perhaps there are no findings right now for a program that offers \$1,000 for a P1, but it can be reasonably guaranteed that if you offer \$1M for a P1, someone will find one.

To that end, the goal of setting a reward range is to blend your organization's desire for findings with your security maturity and willingness to pay a certain amount. A program that desperately wants findings but isn't willing to pay for them needs to review and

possibly reset expectations.

2. Reward ranges are not universally applied to the entirety of a program.

Some assets will be more secure or complex than others. It is reasonable to offer higher rewards for more difficult targets and lower rewards for easier targets, even within the same brief/program. A brief doesn't have to offer the same rewards for all targets. In fact, adding in this variability will help craft a more precise program.

3. The crowd is a free market.

In a free market, the value of any given thing is set by the people willing to pay for it. Paradoxically, even though organizations are the ones paying dollars, the real currency here is the attention of hackers (and often, more specifically, top hackers).

If a program isn't receiving the desired level of engagement or results, provided the program has a large enough sample size to be representative of the crowd as a whole, low results are typically a reflection of the fact that the rewards are inadequate incentives in driving the desired results.

It's important to note that there may also be other reasons why (e.g., inaccessible targets or creds issues), but at the end of the day, the reality is that individuals might not be willing to spend or invest their time on a target for the perceived return on their investment (their time). As covered above, at a high-enough dollar amount, just about anyone can be activated—the reward just has to be worth their time and investment.

4. More complex programs require higher rewards.

Although this sounds intuitive, it cannot be overstated—programs that apply onerous preconditions to testing will not see activity (or at least not activity by highly qualified parties) unless their rewards are attractive enough to get the attention of said parties.

5. Programs that want to attract the best talent and yield the most findings should pay the highest rewards.

Coupled with an open scope, this is a surefire way to get all of the best talent working on your program en masse. Clearly demonstrated by the data on major programs on the Bugcrowd Platform, big scopes and big rewards will bring big talent. If significant impact is what you want, this is the way to do it.

6. "High-" and "mid-level" ranges can also apply to things that are more complex.

When looking at the recommended reward table above, it's important to keep in mind that while the headings used involve terms like "high range," this guidance also applies to things that are more complex. For example, you could have an entry-level, extremely complex program that needs to pay at the "high" level to get attention. Conversely, a moderately complex entry-level program may need to pay at the mid-level, as opposed to the low range.

Notable targets and VRT categories

Every submission to the Bugcrowd platform is tagged with a **target category**. Categories include Android, Hardware, IOS,

IoT, API, Network, Web, and Other. Web continues to be, by far, the largest target category, making up **58% of all**

submissions created. Interestingly, in this year's [Inside the Mind of a Hacker](#) report, **70% of hackers** identified web applications as their area of specialty.

Compared to 2022, this year saw a **30% increase** in Web submissions created, an **18% increase** in API submissions created, a **21% increase** in Android submissions created, and a **17% increase** in IOS submissions created.

Breaking vulnerability data up by VRT category, the **top three categories** of critical submissions rewarded were broken authentication and session management, sensitive data exposure, and server-side injection.

It's important to note here that the **VRT categories aren't set in stone**—they evolve just as the security industry evolves. The VRT is a reflection of the current threat environment. Changes to the VRT are often canaries in the coal mine when it comes to the opportunities associated with certain kinds of vulnerabilities. The VRT evolves dynamically, just like the security industry itself.

One key change to the VRT is the addition of **new AI-related vulnerabilities**, formalizing how AI vulnerabilities get defined, reported, and prioritized. This release reflects the **profound influence** that AI is having on the threat environment and the ways that hackers, customers, and the Bugcrowd triage team view certain vulnerability classes and their relative impacts.

SIDE BOX

Top 5 Most Commonly Reported Vulnerability Types Explained

Digging deeper into the VRT, let's look at the specific vulnerabilities submitted within various VRT categories. We asked prominent hacker **Joseph Thacker, aka rez0**, to break down the meaning of each vulnerability type.

Thacker is a security researcher who specializes in application security and AI. He's helped Fortune 500 companies find vulnerabilities by submitting and collaborating on more than 1,000 reports. Thacker currently works as an offensive security engineer at AppOmni, a SaaS security posture management company based in California.

Measured by the total number of valid submissions found over the past year, the top 5 most commonly identified vulnerability types were as follows:

1. Reflected (cross_site_scripting_xss)

Reflected XSS refers to when a hacker injects malicious code into a website's context like other XSS, but the code only executes if an unsuspecting user were to click on it. When they do, the harmful code is "reflected" back from the website to the user's browser. The browser, thinking the code is safe because it comes from the website it trusts, executes the harmful code. This can have varying impacts depending on the architecture, just like other XSS vulnerabilities. It can always change content, but

account takeover is sometimes also possible. Reflected XSS is notorious for requiring user interaction, meaning that the malicious link needs to be emailed/messaged to the victims or placed where they are likely to click it.

2. Insecure Direct Object References, aka IDOR (broken_access_control)

Insecure direct object references, or IDOR, are a type of vulnerability that occurs when an application allows an unauthorized user to reference an object (essentially any data, be it a user, file, org, or something else) they shouldn't be able to access. The key aspect of IDORs is that they are always categorized by accessing the object via an ID. Sometimes, it's numerical, but other times, it could be a username or UUID.

3. Disclosure of Secrets (sensitive_data_exposure)

Disclosure of secrets, also known as sensitive data exposure, is a vulnerability that happens when a website or application does not properly protect deployment secrets, tokens, user data, etc. Common avenues for data to be exposed are accidentally uploaded files on a webserver, hard-coded data in JavaScript files, deployment files that don't use runtime variables, or GitHub exposures. The impact varies depending on what is exposed.

4. Authentication Bypass (broken_authentication_and_session_management)

This is when a hacker gets around a website's authentication in some way. It might be as simple as accessing the API without a token or navigating directly to an admin panel without logging in. Many methods are used for authentication bypasses, but some common ones are forced browsing (browsing straight to the desired web page without logging in first), default credentials, path traversal, and unique characters in the path of a request to bypass proxy rules. If an attack can bypass the authentication requirements, it's an authentication bypass.

5. Misconfigured Domain Name System (DNS) (server_security_misconfiguration)

A misconfigured Domain Name System (DNS) occurs when a website's DNS is set up incorrectly. The DNS can be incorrect in many ways. Regarding related security impacts, the most common DNS vulnerabilities are ones that result in subdomain takeovers where a DNS record points at an IP address or domain that an attacker can take control of. An example of the former is an elastic IP address in a cloud provider. An example of the latter would be a DNS record pointing to a SaaS provider that allows you to "register" a domain for the SaaS instance. If you no longer use it, an attacker registers it, and you still have a DNS record pointing to it. The subdomain you are pointing to will direct all users to the attacker's instance. This has implications for XSS and businesses if the attacker were to host explicit content on it.

Public vs. Private Programs

There are more public programs on the Bugcrowd Platform than ever before. Although all bug bounty programs begin as private, the idea that they need to stay private is a **legacy mindset** that stems from what we call “**Crowd Fear**” here at Bugcrowd. Crowd Fear is the fear that when a program goes public, it will open the floodgates to thousands and thousands of hackers simultaneously testing the organization’s assets while reporting a significant number of findings, making the program unmanageable and overwhelming. It’s understandable for an organization to be hesitant about taking a program public.

There are **three main drivers** behind the increasing appetite for public bug bounty programs. First, there is **growing customer demand for more transparency** when it comes to security and data. Customers want to support businesses that are proactive in their security approach. It is not in an organization’s best interest to hide the fact that it is on the cutting edge of cybersecurity and doing everything it can to help secure its clients and their data.

We can also see this demand for security transparency in the **growing adoption of VDPs**. A VDP is a structured framework for

hackers to document and submit security vulnerabilities to organizations. VDPs reduce risk by enabling organizations to accept, triage, and rapidly remediate valid vulnerabilities submitted by the security community. **87%** of organizations reported receiving a critical or high-impact vulnerability through a VDP. They also **signal a public commitment** to cybersecurity best practices, which helps **improve confidence** in the organization among customers and the hacker community.

The second driver behind the increase in public bug bounty programs is **increased trust in the hacker community**. As organizations spend more time on the Bugcrowd Platform leveraging crowdsourced security and working closely with hackers, they build confidence in their programs and want to increase their scope and impact.

The third and final driver is that organizations want to take the opportunity for **more impactful outcomes** from a wider pool of talent. The more accessible programs are to hackers, the more successful outcomes organizations can expect.

SIDEBAR/BOX

3 Steps Rapyd Took to Make its Program Public

Rapyd is a cutting-edge fintech leader focused on helping businesses create great commerce experiences anywhere. It had been using crowdsourced security for years, but about a year ago, it made the switch to Bugcrowd with the goal of launching a public program, which it did six months later.

Rapyd has experienced outstanding results so far, uncovering almost **40 unique and valid vulnerabilities—15 of which were critical**. We spoke to **Achiad Aviv**, who is responsible for

application security at Rapyd, for his advice on how to successfully take a bug bounty program public.

Tip 1: Find the right hackers for your program and engage with the community.

While your program is still private, focus on finding specialized hackers for engagements so you have the right fit. By picking the right hackers for specific programs, researchers remain engaged, setting up a future public program for success. Be sure to respond quickly to hackers and engage with them to build positive relationships and a good reputation.

Tip 2: Build confidence in your security posture across the organization.

Be sure you have the right roadmap in place before launching a public program. We worked with Bugcrowd to build this. Our entire team participates in the strategy and operations of our program. We've integrated the platform with numerous DevSec tools for tracking program findings and routing to the appropriate stakeholders. By preparing our process in advance, we felt confident in going public.

Tip 3: Leverage unparalleled expertise from the Bugcrowd team.

Launching a public program is a journey, not a destination. We haven't stopped looking for ways to continuously improve our program, and we work very closely with the Bugcrowd team via email, meetings, and Slack for advice on how best to do this. I encourage you to take similar advantage of these channels.

"We quickly felt safe to take our program public with Bugcrowd. We value the way Bugcrowd finds the right hackers with the right expertise for our programs." – Achiad Avivi, applications security, Rapyd

Learn more about Rapyd's journey with crowdsourced security [here](#).

The Value of an Open-Scope Program

When organizations launch bug bounty programs, they decide what type of scope is right for their programs. A scope is the defined set of targets that have been listed by an organization as assets that are to be tested as part of a particular engagement. Hackers are incentivized to report (and get rewarded for) what is in scope, and what's out of scope is off limits, meaning no compensation is awarded for findings in those targets. Organizations choose

between **three main types of scope**—limited scope, wide scope, and open scope.

Having an open scope is quite possibly the single most effective thing an organization can do to help secure its external attack surface. An open-scope bug bounty program is one that imposes **no limitations** on what hackers can or cannot test, so long as the target or asset belongs to the

organization. Open scopes generally look something like “any externally facing asset belonging to Example Organization,” where nothing is excluded.

Most organizations and bug bounty programs tend to follow a general progression as they grow their security postures over time, starting with a limited scope, expanding to a wide scope, and eventually ending with an open scope.

There’s nothing wrong with running a bug bounty program with a limited scope, but there’s almost always **an opportunity to do more**. There are two main reasons why having an open scope is so valuable for identifying flaws before they are exploited in the wild. The first is that **bad actors don’t have to play by any set scope or rules**. They go wherever they want to find the path of least resistance. If the goal of a bug bounty is to harden and secure assets by finding issues before bad actors, then both

sides need to operate from the same perspective. To give your organization a shot at defending against attackers, it’s critical to give the good guys as much opportunity as possible to find the issues before the bad actors do. Otherwise, it’s a lopsided race out of the gate.

The second reason is that **there is always more than one way in**. The reality is that while you may have a bank vault for a front door, you may have a wide-open window in the back. It’s often far easier to find a way around via a less secure vector versus attacking things head-on where defenses are the strongest.

In 2023, programs with open scopes received **10x more P1 vulnerabilities** than those with limited scopes. This supports the idea that bad actors aren’t asking for permission to test everywhere, and by limiting where the good actors can test, organizations only further disadvantage themselves.

Making the Internet a Safer Place to Hack

Now that you know more about vulnerability trends and recent crowdsourced security insights, let's tease apart **the policy that makes submissions from hackers possible**. There is a **deep societal misunderstanding** of the hacking community, which is reflected in **outdated laws** that hinder their creativity at best and hold them **criminally liable** for ethical disclosures at worst. Although progress has been made, there is still a lot of work to be done.

Bugcrowd founder and Chief Strategy Officer Casey Ellis is on the advisory committee of the **Hacking Policy Council**, an organization that strives to:

- Create a more favorable legal environment for vulnerability disclosure and management, bug bounties, independent repair for security, good faith security research, and pen testing.
- Foster collaboration among the security, business, and policymaking communities.
- Prevent new legal restrictions on security research, pen testing, or vulnerability disclosure and management.
- Strengthen organizations' resilience through the effective adoption of vulnerability disclosure policies and security researcher engagement.

The Hacking Policy Council works on many initiatives throughout the year, but here are three highlights from 2023:

NIST SP 800-171 Rev. 3

In July, the National Institute of Standards and Technology (NIST) updated draft guidelines for NIST Special Publication 800-171—*Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations*.

Ellis, in partnership with the Hacking Policy Council, recommended the addition of VDPs to these guidelines. VDPs help organizations **mitigate risk** by supporting and enabling the disclosure and remediation of vulnerabilities before hackers exploit them. VDPs usually contain a program scope, safe harbor clause, and remediation method. VDPs generally cover all publicly accessible, internet-facing assets.

The addition of VDPs into these guidelines would prevent hackers from having to face the **legal consequences** of good-faith participation in VDPs.

State Charging Policies for Good-Faith Security Researchers

In August, Ellis worked with the Hacking Policy Council to submit a letter encouraging state attorney generals to support the advancement of independent cybersecurity research and the security community for the benefit of all.

Their recommendation encourages state attorney generals to establish policies that clarify and **protect the rights of hackers** conducting security research in good faith.

Hacker Workshops

At the **DEF CON 31** event in Las Vegas last August, the Hacker Policy Council ran a workshop to show hackers how to engage with their governments to **influence local hacking regulations**.

The DEF CON workshops highlighted the process of submitting official comments to hacking regulations and legislation. They covered the process of using regulations.gov and congress.gov as a way to find open opportunities to influence regulations, along

with how to form an **advocacy strategy** to amplify its impact. By attending these workshops, hackers can become active participants in crucial conversations around hacking policy on a government level.

In addition to his work with the Hacking Policy Council, Ellis is also a founding member of **The disclose.io Project**. The goal of this project is to make vulnerability disclosure safe, simple, and standardized for everyone.

Customer Spotlight

The Cybersecurity Skills Gap in a Changing Threat Landscape

In a recent conversation, we had the opportunity to speak with the **Director of Cybersecurity** at a leading data networking organization. Our discussion provided insights into their experience launching a Vulnerability Disclosure Program and shed light on the current threat landscape.

This director is a winner of *Cyber Defense Magazine's* 2023 Top Global CISO Award with over **19 years of experience** in the IT security field. This Bugcrowd customer is a data networking hardware leader that strives to build the world's most reliable, innovative, future-ready wireless technologies that securely connect every person and everything, effortlessly.

We spoke to their director, who is seeing major changes to the threat landscape since the beginning of the pandemic, specifically in the complexity of security threats and the **prevalence of automated and AI-based cyber threats**. "As vulnerabilities and threats continue to increase and become more complex in the wake of AI-based cyber threats, security professionals need to upskill themselves in automation and AI-based technologies to tackle such threats," they said.

They also cite the **cybersecurity skills gap** as one of the top security risks impacting organizations at the moment. "We are at the cusp of an ever-changing technology landscape and evolving, sophisticated security threats. Without adequate cybersecurity skills, organizations are at a significant risk of getting compromised."

They predict the cybersecurity skills gap will continue to increase in the short term, but they aren't without hope. "The security industry has been a **pioneer in hiring people from diverse technology and education backgrounds**, helping them train in cybersecurity skills to fix the hiring gap," they said. They recommend that organizations focus on **investing resources** in their cybersecurity personnel to ensure they stay updated on the latest and greatest with respect to cybersecurity skills. "**Cybersecurity upskilling** should be one of the top business priorities in the organization."

Crowdsourced security is another way to address the cybersecurity skills gap, helping organizations connect with thousands of security experts around the world. The data networking organization decided to partner with Bugcrowd to establish a Vulnerability Disclosure Program (VDP) in order to help manage the vulnerabilities reported by the hacker community. With the help of Bugcrowd, they were able to put a structure to vulnerability submissions, helping them comply with security, compliance, and regulatory requirements.

Beyond compliance requirements, they explored adopting a VDP because they wanted to do everything possible to **proactively reduce risk exposure**, innovating in security instead of just checking boxes. "We want to visibly demonstrate our commitment to security, building productive relationships with the

hacker community. We want security testing and remediation to keep up with the pace of innovation,” they said.

They chose to partner with Bugcrowd because it offers a **multi-solution platform**, extensive experience and a track record of

fast triage response times, reporting and analytics capabilities, adoption and integration, and emphasis on long-term success. Looking forward, they intend to complement their VDP with other Bugcrowd products, including launching a Managed Bug Bounty program and utilizing Pen Testing as a Service.

Casey Speaks: 2024 Cybersecurity Predictions

Prediction #1: The chaotic threat actor returns.

The last time defenders had their attention focused squarely on “asymmetric” or “chaotic” threat actors was Lulzsec and Anonymous in 2013. In 2023, Lapsu\$ demonstrated that defenders have focused on financially and state-motivated attackers, leaving open doors for those whose goals might seem “irrational.” The increasing array of reasons for hacktivists to use hacking as a protest tool puts chaotic threat actors at the top of my list for 2024.

Prediction #2: Questionable election security will continue.

Has it been four years already? Despite progress in election system security, a deepening distrust in election integrity in North America will once again bring the subject of vulnerabilities, hacking in good faith, and the place of security research into public discourse.

Prediction #3: AI will accelerate everything in security.

Since the mainstream adoption of generative AI, brought on by the release of ChatGPT, the potential of AI has captured imaginations everywhere, including those of adversaries. The cat-and-mouse game between attack and defense is as old as time, but the general availability of powerful AI tooling is poised to speed things up.

Prediction #4: Risk management will get noisier.

The first three predictions beget this one—with a higher volume, greater volatility, and a wider range of “baddies” to think about—the importance of efficient prioritization will never be more obvious, and the core role that risk plays in priority assessments will breathe fresh interest into risk management and calculations.


Vulnerabilities Reported by Industry Breakdown

Now that we've seen an overview of vulnerability trends over the past few years, let's dive deeper into what is happening in specific industries to uncover more of the story. There is a **misconception** that only software and technology companies leverage crowdsourced security; however, our data show that this isn't accurate. Although crowdsourced security is heavily used in these spaces, organizations from a wide variety of industries worked with hackers on the Bugcrowd Platform in 2023.

For this report, we narrowed our case study down to **six key industries**—computer software, computer hardware, corporate services, financial services, government, and retail.

Across the board in five of the six industries, the number of submissions increased in 2023 compared to 2022. The biggest increase was in government, which experienced a **151% increase in submissions**. Federal directives that require organizations to develop VDPs may contribute to increased reporting in these industries. In 2020, the Cybersecurity and Infrastructure Security Agency issued [Binding Operational Directive 20-01](#). This directive requires federal, executive departments and agencies to implement their own VDPs and maintain handling procedures—challenging the perception that less regulated industries like technology embrace crowdsourced security the most.

(DESIGN NOTE - This is how H1 formats this section in their report and I really like it)



Computer software

- 12% increase in submissions

Computer hardware

- 2% decrease in submissions

Corporate services

- 20% increase in submissions

Financial services

- 11% increase in submissions

Government

- 151% increase in submissions

Retail

- 34% increase in submissions

Average Payouts

Payouts for P1s are increasing in all industries. The graph below shows the average, median, and 90th percentile bounties paid for P1 submissions in 2023.

Computer software

- Average: \$7,278
- Median: \$2,500
- 90th percentile: \$10,000

Computer hardware

- Average: \$3,058
- Median: \$2,500
- 90th percentile: \$5,200

Corporate services

- Average: \$2,708
- Median: \$2,500
- 90th percentile: \$3,000

Financial services

- Average: \$10,247
- Median: \$10,000
- 90th percentile: \$20,000

Government

- Average: \$5,000
- Median: \$5,000
- 90th percentile: \$8,000

Retail

- Average: \$1,066
- Median: \$500

- 90th percentile: \$2,500

Yet again, we see that the financial services sector offers the **highest average payouts** for critical vulnerabilities. The financial sector has experienced continuous growth in crowdsourced security adoption over the past decade. Financial services institutions were one of the first industries to adopt crowdsourced security.

One reason for this is the regular occurrence of **mergers and acquisitions** in the financial sector. Unfortunately, companies involved in mergers and acquisitions have become **prime targets for ransomware** and other kinds of cyberattacks. To minimize the chances of a breach in the midst of a deal—or after its closing—many organizations in the financial sector find value in crowdsourced security as a way to assess risk and help protect IT infrastructure, applications, and assets during the merger and acquisition process.

Digital transformation is another core driver of crowdsourced security popularity in this sector. The speed with which financial institutions have adopted new technologies, moved to adopt the cloud, and adopted new collaborative tools and technologies has likely contributed to an increase in vulnerabilities. Additionally, security environments are often siloed and fragmented, leaving more blind spots that attackers can exploit. As such, many financial organizations use up to [50 security tools](#), sometimes more.

Although we only presented data on six key industries in this section, it's important to remember that all industries are currently adapting to today's security environment. A

recent [ICS2 study](#) found that those in the healthcare, military, energy/power/utilities, government, and manufacturing industries

believe that they are **more sensitive to threats** than other industries in the modern threat landscape.

Customer Spotlight

Meet Martin Choluj, VP of Security at ClickHouse

We recently had the privilege of speaking with **Martin Choluj**, the vice president of security at ClickHouse. Our discussion yielded valuable insights into his experience collaborating with Bugcrowd and the critical role that crowdsourced security plays in safeguarding a brand's intellectual property.

Choluj is a seasoned security professional with an impressive **15-year track record** in the field. He is currently VP of security at ClickHouse, a company renowned for its efficient open source database solutions.

Before stepping into this role, Choluj spent nearly six years as CISO at Campaign Monitor and held various security leadership roles in international financial institutions. Bolstering his practical experience, he holds a Master's Degree in Security and Forensic Computing and a Bachelor's Degree in Information Technology.

At its core, ClickHouse champions the principles of **trust and risk reduction**, and it's this ethos that led it to explore a bug bounty program. Choluj highlighted that the company's aim is not simply compliance but fostering innovation in security and building constructive relationships with the hacker community.

Choluj's partnership with Bugcrowd started in 2016 in a previous role, which led ClickHouse to choose our platform over

others. With Bugcrowd, ClickHouse was able to tap into a global community of hackers to identify and address hidden, high-impact vulnerabilities.

According to Choluj, a proactive approach is essential to any large-scale assurance program. He underscored the importance of crowdsourced security by saying, "Interacting with the hacker community is vital for our assurance program to operate on a large scale effectively."

He praised Bugcrowd's triage response time and commitment to long-term customer success, both underpinned by a solid track record of experience. The primary challenge for ClickHouse was **anticipating attack vectors and attacker ingenuity**—an area where Bugcrowd's expertise has proven invaluable.

Choluj also acknowledged a **skill gap in cybersecurity**, particularly when bridging the divide between security and engineering. He sees the Bugcrowd platform as a viable solution to this challenge, enabling organizations to augment their internal teams by tapping into the collective creativity of hackers. This approach **effectively bridges the workforce gap**, fostering stronger synergy between different domains of expertise.

A wave of **digital revolution** has prompted organizations to rethink their security

strategies. Old-school methods centered on safeguarding known environments and networks no longer suffice. Choluj asserted that the shift to remote work, amplified by the pandemic, requires a new focus on securing systems and users, regardless of location.

Choluj's experience highlights the importance of treating cybersecurity as an **ongoing strategic endeavor** rather than as

a one-off project. His partnership with Bugcrowd exemplifies how a platform-driven approach to crowdsourced security can strengthen an organization's defenses, turning potential vulnerabilities into fortified security measures.

Embracing crowdsourced security is more than a wise business decision in today's intricate digital landscape; it's a necessary step toward a secure digital tomorrow.

Thought Piece: How Different Hacker Roles Contribute to Crowdsourced Security with Bugcrowd

VP of Security Operations and Hacker Success

Michael Skelton

We can't say this often enough: Adopters of crowdsourced security are only as successful as the hackers/security researchers with whom they collaborate, whether it's in a crowdsourced penetration test, bug bounty, or something else. A major ingredient in that success is the ability to match and activate the right hackers and/or pentesters for the task at hand—and quite often, the *types of hacker roles* involved also make a big difference in the results.

When evaluating the value of crowdsourced security, many people focus on the number of hackers who will focus on their targets. While this is a logical approach, it's just as important to consider the **diversity of perspectives** that a "crowd" can provide. For example, in a traditional penetration test, the findings usually reflect the perspective of a single "type" of tester (more on that below), which produces results aligned with that one perspective, albeit these results conform to a methodology. In contrast, a genuinely crowdsourced pen test

(not a **"crowd-washed"** one) inherits value from the **full range of thoughts**, approaches, and styles that only a crowd can provide—and that enables more comprehensive, intensive testing to find more diverse types of bugs. Furthermore, it's a strong signal that "pay for effort" (typical of an industry-standard pen test) and "pay for impact" (typical of a bug bounty) testing models are **highly complementary**.

At Bugcrowd, we think of hackers/pentesters as occupying one of five distinct roles: *Beginners*, *Recon Hackers*, *Deep Divers*, *Generalists*, and *Specialists*. (It's also important to keep in mind that over time, hackers/pentesters can and will journey from one role to another.) Each role plays an important part in a given program, and these roles are relevant to how the Bugcrowd Platform's CrowdMatch technology matches the right crowd to a customer's needs, at the right time, across hundreds of dimensions.

Next, let's take a look at each type of role in more detail.

The Beginner

Beginners on the Bugcrowd Platform refer to those who are **new to the concept of crowdsourced security** in general rather than just being new to the platform specifically. When assessing a hacker's level of experience, we may consider factors such as their participation on other platforms or their published research and tools. However, if such information is not available, we may assume that the hacker is a beginner in the ecosystem, at least initially (although this may not always be the case).

It's important to note that a Beginner is not necessarily unskilled, even if they're only submitting P3/P4 issues. For example, they may be working through a course to broaden their skill set, or they may have limited public presence but already work as a pentester and want to further develop their skills. Typically, this type of hacker covers vulnerability classes that others may not focus on as much, including P4 issues related to authentication and authorization, as well as simpler infrastructure issues (such as DMARC).

Beginners add value in terms of **coverage and consistency**. Their participation in a program ensures, for example, vulnerabilities that would typically be found in a penetration test are also identified in a bug bounty program. The last thing we want is for a customer to follow a penetration test with an overlapping bug bounty and only then learn about a bunch of lower-priority items!

The Recon Hacker

Recon Hackers focus on identifying issues across the **largest scope possible**, so these individuals often discover P2/P3 issues that would not typically be found in a penetration test.

Over the past few years, Recon Hackers have dominated every provider's leaderboard due to the proliferation of subdomain takeovers, particularly in ROUTE53 and EC2 takeovers. While these takeovers are now largely patched, the leaderboards are now askew, and thus, the highest-rated hackers may not always bring the maximum level of impact.

It's important to note that many recon-based hackers are highly skilled. However, many who take a recon-first approach have found a **lucrative niche** and thus tend to focus on refining their toolkit to further exploit only that niche.

The Deep Diver

Deep Divers are the most valuable hackers for Bugcrowd to identify, engage, retain, and uplift. These hackers tend to **focus on a particular program**, learn as much as they can about it, and provide unique and distinct value. A Deep Diver can uncover vulnerabilities that nobody else can due to their persistence and long-term knowledge of how a program operates.

Identifying these hackers is best done by analyzing the **content of their submissions**—rather than just looking at the spread of vulnerabilities across a program—due to the unique nature of these findings.

The Generalist

Generalists take a **multifaceted approach**: They have a solid foundation in reconnaissance and utilize it to cover attack surfaces thoroughly, without relying solely on large-scale monitoring and tooling. Generalists also apply a deep-diving approach to evaluating assets. While they may not spend as much time on a particular program as deep divers do, they invest considerable time across a variety of programs. Due to their dual proficiency in recon and deep diving, Generalists quickly gain a reputation on the Bugcrowd Platform and are **highly valued**.

The Specialist

Specialists are a **rare breed** who require specific sourcing for an engagement. They possess unique and rare skill sets and typically have years of experience in a particular technology (e.g., APIs, AI, IoT, and Web3) or a specific [Bugcrowd VRT](#) category.

As you read in the introduction, one of the greatest strengths of the Bugcrowd Platform is its ability to **source and activate specialists** to meet a program's specific skill set needs. Due to their specialized knowledge, Specialists can uncover issues

that other hackers may miss, and they often provide invaluable, unique solutions to a problem.

An Engineered Approach

To maximize the contributions of each hacker role, Bugcrowd is strategic in its approach to sourcing and engaging with hackers. For example, adding Beginners to a program that has been running for three months may lead to frustration and a high number of duplicates, while adding Generalists too early dilutes the ability of Beginners to up-level themselves through their findings. Therefore, **program maturity** is an important input for our platform's CrowdMatch technology when it comes to sourcing the appropriate roles.

To summarize, different hacker roles contribute to crowdsourced security programs in different ways, and it's important to deeply understand a program's needs to make the most of these contributions. To respect this process, unlike other providers that rely on leaderboards or coarse-grained methods, Bugcrowd's **engineered approach** intelligently sources and activates the right role types and skills for your programs, at the right time.

Why a Crowdsourced Security Platform is the Best Early Warning System for Vulnerabilities

By providing access to the **collective expertise of thousands of trusted ethical hackers**, a crowdsourced security platform can serve as an early warning system that enables organizations to discover and remediate vulnerabilities before attackers can exploit them.

As organizations look to mature their security strategies, many have found that a crowdsourced security platform best addresses their needs. Whether you are just getting started with crowdsourced security or improving upon an existing program, implementing a **platform-based solution** is a good place to start. The ideal solution will:

- Connect your organization with trusted ethical hackers who have the skills to meet your specific requirements.
- Offer fast triage and prioritization of vulnerability submissions.
- Provide contextual intelligence without the noise of traditional scanning solutions.
- Streamline workflows with integration across the platform and DevOps tools you rely on.

To better understand the benefits of a platform-based approach to crowdsourced security, let's take a look at some of the **key benefits** of the Bugcrowd Platform.

The Bugcrowd Platform

Our platform is designed to apply over a decade of expertise and context to an organization's cybersecurity program. Its **massive knowledge graph** of historical hacker, vulnerability, interaction, asset, and remediation data can inform workflows.

Using the platform, organizations can easily create and incentivize bug bounty, vulnerability disclosure, and pen testing

programs. To ensure organizations are connected with the "right" hackers, the platform uses **AI models** and data from our vast knowledge base to match hacker skill sets, interests, and availability with an organization's specific needs.

For successful crowdsourcing to be effective, **rapid triage and prioritization** are required. Thus, our platform enables rapid vulnerability triage at any scale with the industry's best signal-to-noise ratio. Our global team of security engineers adds **critical context** to hacker submissions by rapidly validating and triaging bugs (with the most critical ones handled within hours). Severity levels are based on a rich, open source VRT developed over a decade and our deep connections to the hacker community.

Finally, a good crowdsourced security strategy does not exist in a vacuum. It must be part of a **broader workflow** that extends across DevOps tools and the software development lifecycle (SDLC). That's why the platform includes pre-built connectors, webhooks, and rich APIs to flow findings into your DevOps tools and life cycle in real time.

PLATFORM GRAPHIC HERE

Social Media Spotlight

What was the most common bug of 2023?

Now that we've seen the data around the most common bugs of 2023, we wanted to get some anecdotal feedback from the hacking community via X to understand if what they are seeing aligns with the data. [Here](#) are some of their thoughts.

 **Joseph Thacker (rez0)** ✓
@rez0__

Improper auth/access control

12:14 PM · Oct 23, 2023 · **526** Views

<https://twitter.com/rez0/status/1716488336016134373>

Some accounts you follow often like this account

 **Nikhil Rajput**
@Swaggy_Singh_R

Broken Access Control.

12:04 PM · Oct 23, 2023 · **265** Views

https://twitter.com/Swaggy_Singh_R/status/1716485660058771961

 **V1k1ng**
@v1k1ng_h4ck3r

IDOR Exploits
Exposed Sensitive Information in Publicly Accessed URLs/Directories

I would have to say those 2 for sure are at least in the top 3!

1:13 PM · Oct 23, 2023 · **34** Views

https://twitter.com/v1k1ng_h4ck3r/status/1716503191099998403

Some accounts you follow often like this account

 **B_K_S**
@srb1mal

Improper Access Control and Specially auth bypasses

1:36 PM · Oct 23, 2023 · **65** Views

<https://twitter.com/srb1mal/status/1716508875728519237>



<https://twitter.com/TheLittleH4ck3r/status/1716492276262912308>

Conclusion

That's a wrap on this year's edition of *Inside the Platform*. This year, we found that vulnerability trends from past years have continued to hold true—more and more organizations across all industries are using crowdsourced security, so we continue to see a rise in the number of high-impact and valid vulnerabilities.

These data highlight many trends in the industry, but our biggest takeaway is that **crowdsourced security is both an art and a science**—and Bugcrowd makes both parts scalable. On the one hand, crowdsourced security is not a free for all. There is a rationale behind every decision, and organizations can follow these **predictors of success** to maximize their programs. However, this isn't all an exact science. Organizations must be thoughtful about how they design their programs, briefs, and incentives. In addition, organizations must think about the ways in which they interact with hackers to foster mutually beneficial relationships.

The crowdsourced security industry has matured over the course of the last decade,

and even though many still view it as a new part of the security technology stack, there is no denying that the industry is evolving. It is no longer enough to just have a bug bounty program—crowdsourced security is not a “one-and-done” exercise. Organizations must think about building **dynamic programs** with continuous improvement in mind.

This may seem daunting for many organizations, but they don't have to go it alone. Bugcrowd has a decade of experience. We know what “good” looks like, we know the **different levers** organizations call pull at different times to grow their programs, and we know the red flags to watch out for. We also help organizations **prioritize continuous improvement** through analytics, powered by a rich Security Knowledge Graph of vulnerabilities, assets, environments, and skill sets based on thousands of customer experiences. Critical insights from that data help organizations continuously improve their security posture, constantly raising the bar for excellence.

Content recommendations

Data Sheet

[CrowdMatch](#)

Unleash hacker ingenuity with AI-powered matching and activation

eBook

[Expanding Risk Reduction with a Crowdsourced Security Platform](#)

Ways the Bugcrowd Platform goes beyond crowdsourced security

Guide

[What's a Vulnerability Worth?](#)

Building a rewards model for your bug bounty program

Interactive Tour

[Bugcrowd Platform Tour](#)

A 5-minute overview of how the Bugcrowd Platform works

[Request a Demo](#)

Glossary

- **Adversary:** An individual, group, or organization that actively seeks to compromise the security of a system or network.
- **Adversarial AI:** When threat actors target the data sets, algorithms, or models that an ML system uses to deceive and manipulate their calculations, steal data appearing in training sets, compromise their operation, and render them ineffective.
- **Ally:** A person or entity that supports and cooperates with another to protect the security of a system or network.
- **Application programming interface (API):** An API is a way for two or more computer programs to communicate with each other. It is a type of software interface that offers a service to other pieces of software.
- **AI:** The simulation of human intelligence processes by machines, particularly computer systems, to execute tasks akin to learning and decision-making found in humans. Subsets of AI include expert systems, neural networks, deep learning, natural language processing, speech recognition, and machine vision. In cybersecurity, AI is applied in attack surface management, automated detection and response, and intelligent authentication and fraud prevention.
- **Asset:** Any data, device, or environmental component that supports information-related activities. Assets generally include hardware, software, and confidential information.
- **Asymmetric Intent:** Cyberwarfare that seeks to inflict a proportionally large amount of damage compared to the resources used by targeting the victim's most vulnerable security measure.
- **Attack Surface:** The sum of the different points in a software environment where an unauthorized user can enter or extract data. Minimizing the attack surface is a basic security measure.
- **Attacker:** An individual or group that performs malicious activities to destroy, expose, alter, disable, steal, or gain unauthorized access to or make unauthorized use of an asset.
- **Bad Actor:** Also called a malicious actor or threat actor, an entity that is partially or wholly responsible for an incident that impacts or has the potential to impact an organization's security.
- **Beginner Hacker:** Hackers who are new to the concept of crowdsourced security in general.

- **Bounty:** Monetary rewards offered in exchange for a vulnerability finding, discovery, or report.
- **Bounty Hunter:** A highly skilled hacker who receives recognition and compensation in exchange for reporting bugs, especially those pertaining to security exploits and vulnerabilities.
- **Breach:** A cyberattack in which sensitive, confidential, or otherwise protected data have been accessed or disclosed in an unauthorized manner.
- **Bug:** A software defect that can be exploited to gain unauthorized access or privileges to a computer system.
- **Bug Bounty:** Bug bounty programs allow independent security researchers to report bugs to an organization and receive rewards or compensation.
- **Bug Hunter:** A highly skilled hacker who receives recognition and compensation in exchange for reporting bugs, especially those pertaining to security exploits and vulnerabilities.
- **Chief Information Security Officer (CISO):** The senior-level executive within an organization responsible for establishing and maintaining the enterprise's vision, strategy, and program to ensure assets and technologies are adequately protected.
- **Critical Vulnerabilities and Exposures (CVE):** A list of publicly disclosed security flaws.
- **Crowd Fear:** The fear of an unmanageable and overwhelming number of submissions once a program goes public.
- **Crowd Washing:** The purposeful and sometimes deceptive attempt by a vendor to make their offerings sound more modern and impactful than they really are.
- **CrowdMatch:** Bugcrowd's proprietary AI technology that matches precisely the right trusted hackers to a specific program's needs across hundreds of dimensions, producing tighter engagement and better results.
- **Crowdsourced Security:** An organized security approach wherein ethical hackers are incentivized to search for and report vulnerabilities in the assets of a given organization. The power of crowdsourced security is derived from the proportion of active testers per asset/ecosystem versus more traditional testing methods.
- **Customer:** Organizations that leverage the Bugcrowd platform or its associated services.
- **Cybercriminal:** An individual or group that commits malicious activities on a system or network with the intention of stealing sensitive information or personal data to generate profit.
- **Cybersecurity Skill Gap:** The mismatch between the skills employers require in cybersecurity professions and the qualifications possessed by potential candidates.
- **Deep Diver Hacker:** Hackers who tend to focus on a particular program, learn as much as they can about it, and provide unique and distinct value.
- **DEF CON:** A hacker convention held annually in Las Vegas, Nevada.
- **DevOps:** A methodology in the software development and IT industry that integrates and automates the work of software development and IT operations as a means for improving and shortening the SDLC.
- **Digital Transformation:** The process of fundamentally changing an organization through technology and culture to improve/replace what existed before.
- **Disclosure:** The practice of reporting security flaws in computer software or hardware.
- **Engagement:** Measurable indicators of the level of interest, involvement, and influence that a crowdsourced security program generates among ethical hackers or custom-designed pen testing solutions tailored to an organization's unique needs.
- **Ethical Hacker:** A person who hacks into a computer network to test/evaluate its security rather than to carry out an act of malice.
- **Ethical Hacking:** An authorized attempt to gain unauthorized access to a computer system, application, or data.
- **Exposure:** All vulnerabilities and risks associated with an organization's networks, systems, applications, and data. Exposures encompass the potential weaknesses and threats that cybercriminals may exploit, resulting in security breaches, data loss, or other adverse consequences for an organization.
- **Generalist Hacker:** Hackers with a solid foundation in reconnaissance who utilize it to cover attack surfaces thoroughly, without relying solely on large-scale monitoring and tooling. They also apply a deep-diving approach to evaluating assets.
- **Generative AI:** A type of AI technology that can produce various types of content, including text, imagery, audio, and synthetic data in response to prompts. Generative AI models learn the patterns and structures of their input training data and then generate new data that have similar characteristics.
- **GRC:** Governance, risk (management), and compliance.
- **Hacker:** Someone who uses technical knowledge to achieve a goal or overcome an obstacle

within a computer system by non-standard means.

- **Hacking Policy Council:** An organization that strives to create a more favorable legal environment for vulnerability disclosure and management, bug bounties, independent repair for security, good-faith hacking, and pen testing.
- **Human Element:** The role people play in the design, implementation, and operation of technology systems, as well as their potential to introduce vulnerabilities or mitigate risks.
- **Inflation:** The measure of how much more expensive a set of goods and services has become over a certain period, usually a year.
- **Internet of Things (IoT):** Any device (often called a smart or connected device) that connects to and exchanges information over the internet.
- **Incident Response:** A term used to describe the process by which an organization handles a data breach or cyberattack, including the way an organization attempts to manage the consequences of an incident so that damage, recovery time, and costs are limited, and collateral damage, such as brand reputation, is kept to a minimum.
- **IT Asset Management:** The process of cataloging, tracking, and maintaining an organization's technology assets.
- **Lapsu\$:** An international extortion-focused hacker group known for its various cyberattacks against companies and government agencies.
- **Limited Scope:** A bug bounty program that includes only a single or specific target(s).
- **Malicious Hacker:** Someone who is actively working to disable security systems with the intent of

either taking down a system or stealing information.

- **Mergers and Acquisitions:** Business transactions in which the ownership of companies, business organizations, or their operating units are transferred to or consolidated with another company or business organization.
- **Model:** A program that analyzes mathematical representations of relationships between variables to make predictions or decisions in AI systems.
- **National Institute of Standards and Technology (NIST):** An agency of the United States Department of Commerce whose mission is to promote American innovation and industrial competitiveness.
- **Open Scope:** A bounty program with no limitations on what hackers can or cannot test, so long as the target/asset belongs to the specified organization.
- **Payout:** The money paid to a researcher once their vulnerability submission has been validated.
- **P1—Critical:** Vulnerabilities that cause a privilege escalation from unprivileged to admin or allow for remote code execution, financial theft, etc.
- **P2—High:** Vulnerabilities that affect the security of the software and the processes it supports.
- **Penetration Testing/Pen Testing:** A simulated cyberattack done by authorized hackers who test and evaluate the security vulnerabilities of the target organization's computer systems, networks, and application infrastructure.
- **Platform/SaaS Platform:** Bugcrowd is an all-in-one SaaS platform that combines actionable, contextual intelligence with the skills and experience of the world's most elite hackers to help leading organizations solve security challenges, protect customers, and

make the digitally connected world a safer place.

- **Point-in-Time Assessment/Security Testing:** A point-in-time review of a company's technology, people, and processes to identify problems. Such assessments can find vulnerabilities existing at a single moment but fail to monitor activity between assessments.
- **Program:** A program—which can be public or private—permits independent researchers to discover and report security issues that affect the confidentiality, integrity, or availability of customer or company information and rewards them for being the first to discover a bug.
- **Program Brief:** A single-page researcher-facing document that contains all relevant information regarding a bounty program (what is in/out of scope, rewards, how submissions will be rated, instructions for accessing or testing the application, etc.). This is drafted with the Bugcrowd team after the initial kickoff call.
- **Ransomware:** A type of malware designed to extort money from its victims, who are blocked or prevented from accessing data on their systems.
- **Recon Hacker:** Hackers who focus on identifying issues across the largest scope possible, so these individuals often discover P2/P3 issues that would not typically be found in a pen test.
- **Risk:** The potential for loss, damage, or negative consequences resulting from threats to the confidentiality, integrity, or availability of information or systems.
- **Software as a Service (SaaS):** A software licensing and delivery model in which software is licensed on a subscription basis and is centrally hosted.

- **Scope:** Outlines the rules of engagement for a bounty program. This includes a clearly defined testing parameter to inform researchers what they can and cannot test, as well as the payout range for accepted vulnerabilities.
- **Security Landscape:** The entirety of potential and identified cyber risks affecting a particular sector, group of users, time period, etc.
- **Security Research:** The study of technology, algorithms, and systems that protect the security and integrity of computer systems, the information they store, and the people who use them.
- **Security Researcher:** The diverse group of skilled participants who hunt for vulnerabilities using the Bugcrowd platform. These trusted experts are sometimes referred to as *white hats* or *ethical hackers*.
- **Software Bill of Materials:** A list of all the open source and third-party components present in a codebase.
- **Software Development Lifecycle (SDLC):** A structured process that enables the production of high-quality, low-cost software in the shortest possible time.
- **Specialist Hacker:** A hacker with unique and rare skill sets and who typically has years of experience in a particular technology (e.g., APIs, AI, IoT, and Web3) or a specific Bugcrowd VRT category.
- **Submission:** The report a researcher submits to Bugcrowd describing the vulnerability or bug they found.
- **Target:** A web or mobile application, hardware, or API that the Crowd tests for vulnerabilities.
- **The Crowd:** The global community of white hat hackers on the Bugcrowd platform who compete to find vulnerabilities in bug bounty programs.
- **The Disclose.io Project:** A collaborative, open source and vendor-agnostic project to standardize best practices for providing a safe harbor for security researchers within bug bounty and VDPs.
- **Threat Actor:** An individual, group, or organization that poses a potential risk to the security of information or systems through malicious activities.
- **Triage:** The process of validating a vulnerability submission from a raw submission to a valid, easily digestible report.
- **Valid:** The state of a vulnerability that has been tested and confirmed as real.
- **Vulnerability Rating Taxonomy (VRT):** The official standard used by Bugcrowd for assessing, prioritizing, and benchmarking the severity of security vulnerabilities.
- **Vulnerability:** A security flaw or weakness found in software or in an operating system that can lead to security concerns.
- **Vulnerability Disclosure Program (VDP):** Clear guidelines for researchers to submit security vulnerabilities to organizations while also helping organizations mitigate risk by supporting and enabling the disclosure and remediation of vulnerabilities before they are exploited. VDPs usually contain a program scope, safe harbor clause, and method of remediation.
- **Wide Scope:** A bounty program that includes a wildcard in the in-scope targets.
- **White Hat Hacker:** A computer security expert who uses pen testing skills to help secure an organization's networks and information system assets. A white hat hacker is also known as an ethical hacker. White hat hackers work with information technology and network operations teams to fix vulnerabilities before black hat hackers discover them. White hat hackers operate with the permission of the organization and within the set boundaries.