2021

ULTIMATE GUIDE TO Vulnerability Disclosure

bugcrowd

TABLE OF CONTENTS

INTRODUCTION

The digital world comprises billions of lines of code written by human programmers—unique in critical skills, experience, and education. To err is human, and variation and errors are inevitable as it continues to expand at a rate that far outpaces any one organization's ability to keep up. Vulnerabilities are, quite simply, a fact of software development. Software deployed across millions of devices will be recycled and reused, multiplying attack surface and risk. To prosper in this environment, organizations must adopt a model of humility, transparency, collaboration, and action to maintain the trust of customers, partners, and the security community at large.

Organizations need a mechanism for identifying and remediating vulnerabilities discovered outside the typical software development lifecycle. Vulnerability Disclosure Programs provide a means to align these considerations efficiently and economically, while building a stronger security brand.

EXECUTIVE SUMMARY

THIS REPORT EXAMINES...

- The strategic, legal, and social nuances associated with vulnerabilities discovered “in the wild”
- The basics of VDPs, including key benefits
- Why the NIST Cybersecurity Framework lists vulnerability disclosure as a requirement for every organization
- Best practices for implementing and managing a VDP
- How to combine a VDP with Bug Bounty programs or penetration testing

87%

of organizations receive critical or high priority findings with a VDP

79%

of organizations award monetary payments for impactful findings

99%

of organizations consider pairing a bug bounty program with their VDP

78%

of organizations expand their VDP with multiple programs

VULNERABILITY 101

What Are Vulnerabilities?

Components of code that can be exploited to negatively impact the security of data, systems, people, or IP.

What Is The Cause Of Vulnerabilities?

Vulnerabilities can be the result of erroneous scripting or can arise from changes in the deployment environment, or from several seemingly intentional commands combined in unintentional ways.

How Common Are Vulnerabilities?

The average software application reportedly has 15-50 "bugs" per thousand lines of code.

15-50

How Are Vulnerabilities Surfaced?

Most internally developed software progresses through similar development lifecycles, which include several phases of targeted testing prior to, and throughout production. Unfortunately, it's impossible to simulate every possible use case, permutation or potential interaction in such controlled settings. Additionally, software is always evolving—expanding and contracting like a living organism to adapt to new operating environments and an ever-growing list of connected tools and services.

WHO FINDS VULNERABILITIES? →

- Internal Software Developers
- End-users
- Hackers/Security Researchers

Vulnerabilities are inevitable and are not a sign of weakness.
It's all about how an organization responds to these vulnerabilities.

THE BASICS OF VULNERABILITY DISCLOSURE PROGRAMS

WHAT IS A VULNERABILITY DISCLOSURE PROGRAM (VDP)?

Vulnerability disclosure programs, or VDPs, may be best described as the **internet's "neighborhood watch."** Neighborhood watches leverage a formal system run on voluntary effort to report suspicious activity. While the city does much to protect inhabitants through routine police patrols, and emergency response services, neighborhood watches help "fill in the gaps," for **24/7 community-lead protection.** These communications are incentivized by an altruistic desire to make the neighborhood safer, as well as build relationships that persist even when neighbors move away.

Just like neighborhood watches, VDPs encourage anyone that uses your corner of the internet, to take care of it, **for the benefit of all.** VDPs provide a framework to encourage and facilitate the secure reporting of vulnerabilities discovered outside of typical testing cycles. And as they usually cover all publicly-accessible, internet-facing assets, anyone with an internet connection can participate. Additionally, just as the simple presence of neighborhood watch signs tend to **deter nefarious activity**, publicly posted VDPs indicate that the organization is unlikely to be an easy target.

Presence of a VDP in an organization is an acknowledgment that they understand the inevitability of vulnerabilities and are **committed to security transparency.** It's worth noting that VDPs are different from bug bounty programs and penetration testing, which are focused on actively engaging with the security community. We'll cover this more later on in the guide.

The method for managing VDPs differs by organization and is often dependent on **goals, resources, and bandwidth.** Some choose self-management, while others rely on third parties like Bugcrowd to monitor intake channels, triage findings, and provide feedback to the submitting party. Rewards for valid vulnerabilities also differ by program and management type, and while "kudos points," are the standard method for showing appreciation, some programs offer ad-hoc payments for findings with significant impact. While this may seem similar to a bug bounty program, there is an important distinction—bug bounty programs *incentivize* submissions, VDPs selectively *reward* them.

87% of organizations report receiving a critical vulnerability through their VDP

A VDP reduces risk, while publicly showcasing a company's commitment to security in a way that is both easily understood, and easily verified

Vulnerability disclosure programs may be best described as the internet's "neighborhood watch"

It is this last point that tends to cause the most consternation in both communities. While many organizations derive great value from highly active vulnerability disclosure programs, **the purpose of a VDP is first and foremost to provide a secure channel for well-intentioned, externally-sourced security feedback.** VDPs do not replace Bug Bounties, and equally, Bug Bounties do not replace VDPs.

By allowing for the communication of vulnerabilities found in the routine use or testing of externally-facing products and services, organizations can **greatly expand their risk-reduction with minimal disruption to existing security and production lifecycles.** By offering recognition to well-intentioned hackers that abide by a defined process, VDPs simultaneously **build and enhance an organization's reputation for security.** Brian Adeloye, Principal Product Security Engineer at Atlassian, states that, "a VDP is a reciprocation of the good faith shown by hackers who identify and share vulnerabilities of their own volition. This provides an opportunity for organizations to give and get respect within the security community."

KEY BENEFITS OF A VDP

Countless vulnerabilities are being written into new and existing software every day, and organizations need to maximize their catchment area for discovering these. Yet a large majority of security professionals claim that at some time or another, they have been unable to report a vulnerability that they discovered. In fact, less than 9% of Fortune 500s have a VDP in place today. All organizations are well practiced in paying lip service to “taking security seriously” whenever they make the news, but the actual prevalence and maturity of VDPs say otherwise.

Let’s look closer at the idea of “taking security seriously” and discuss what that actually means. This statement can usually be boiled down to a few **common goals and priorities**, such as:

- Reduce risk
- Improve security ROI
- Accelerate digital transformation
- Make better decisions on security initiatives
- Improve security transparency and customer confidence

VDPs help organizations achieve these goals in many different ways. We’ll discuss how Bugcrowd Vulnerability Disclosure Programs specifically support the above security goals later on in this guide.

WHAT DO YOU BELIEVE IS THE MAIN POINT OF VALUE FOR YOUR VDP?

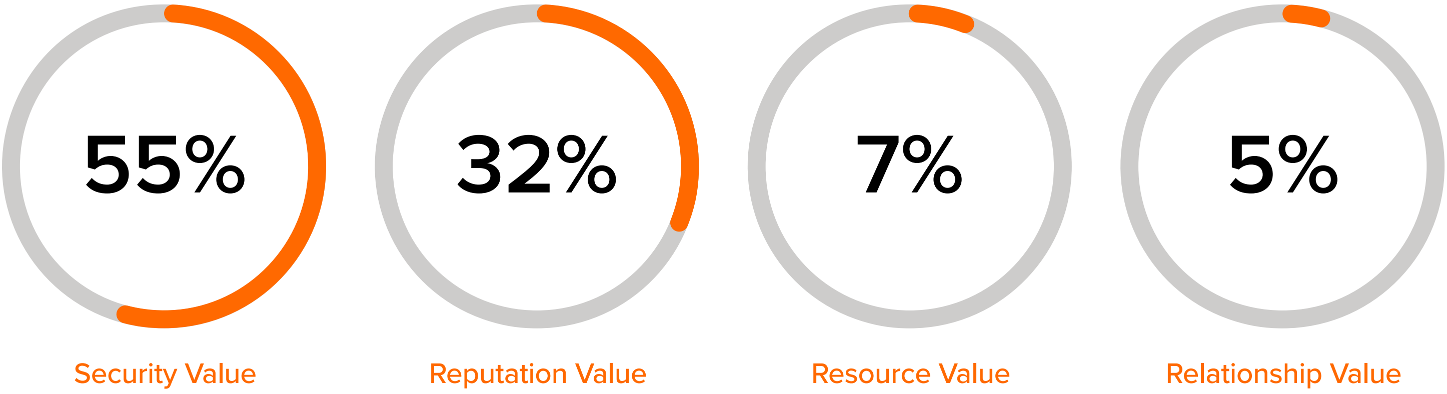

VDPs AND ELECTION SECURITY

CASEY ELLIS
FOUNDER, CHAIRMAN & CTO

As our use of the Internet expands, the "Neighborhood Watch for the Internet" plays an increasingly crucial role in creating confidence in the systems we rely on day-to-day. The most expansive and recent example of this is **the role white-hat hackers played in the 2020 US Presidential Elections.**

In 2018, the House Rules Committee invited Bugcrowd to a meeting of the minds on Capitol Hill. Our primary contribution was to call out that VDPs would play a key role in reducing cyber risk and, more importantly, that VDPs would be a transparent and effective way to build confidence in the integrity of the systems which power democracy itself.

These conversations, our input, and the input of a host of others had the desired effect:

- The primary voting machine manufacturers reversed years of tension with the security research community, **announcing VDPs and private bug bounty programs** during DEF CON
- CISA/DHS posted **Election Administrator guidance** which included VDPs with bilateral safe harbor
- Iowa, along with other states, demonstrated incredible leadership by **launching a VDP covering all election-related infrastructure**

Beyond building a stronger security posture, there are several **key benefits of a VDP** from the perspective of your customers, partners, investors, and employees, and the security researcher community.

Customers

As Bruce Schneier noted, vulnerabilities are an externality that affect end users much more than owners. Ethan Dodge, security engineer at Atlassian, agrees, "the primary goal of a VDP is to do right by your end users." This means organizations should not only prioritize the security of their users' data for their sake, but for the reputational, and ultimately financial damage the organization will incur if they fail to do so. A VDP allows companies to **reduce risk,** while publicly showcasing their commitment to security in a way that is both easily understood, and easily verified. No more lip service.

Partners, Investors, and Employees

The VDP halo extends to an organization's overall security brand, acting as a **strong indicator of security posture** for external stakeholders like prospective investors, partners, and other collaborators. These programs are public evidence of an organization's culture of remediation, recognition, respect, and commitment to rapid response. For potential security hires, the presence of a VDP often signifies the influence wielded by security leadership amongst executive peers like Marketing, Legal, and Sales. This may be best summarized by Dodge's further observation, "vulnerability disclosure is a litmus test for how advanced a company's security culture is, and it is a more accurate indicator than budget or press coverage. I have always asked about VDPs when interviewing for jobs to assess the working environment."

Security Researchers

Any discussion on the impact of VDPs would be incomplete without due attention to the finders of vulnerabilities themselves. VDPs provide emerging security researchers the opportunity to hone their skills, while established hackers can **build and extend relationships** with organizations that may also offer private, invite-only engagements like bug bounties. Moreover, both groups benefit from the knowledge that they are incrementally improving the organization's security— something that 93% of hackers cite is their primary motivation according to the 2020 "Inside the Mind of a Hacker" report.

FOR ORGANIZATIONS WITH A VDP

WHAT IS THE MAIN REASON WHY YOUR ORGANIZATION HAS IMPLEMENTED A VULNERABILITY DISCLOSURE PROGRAM?

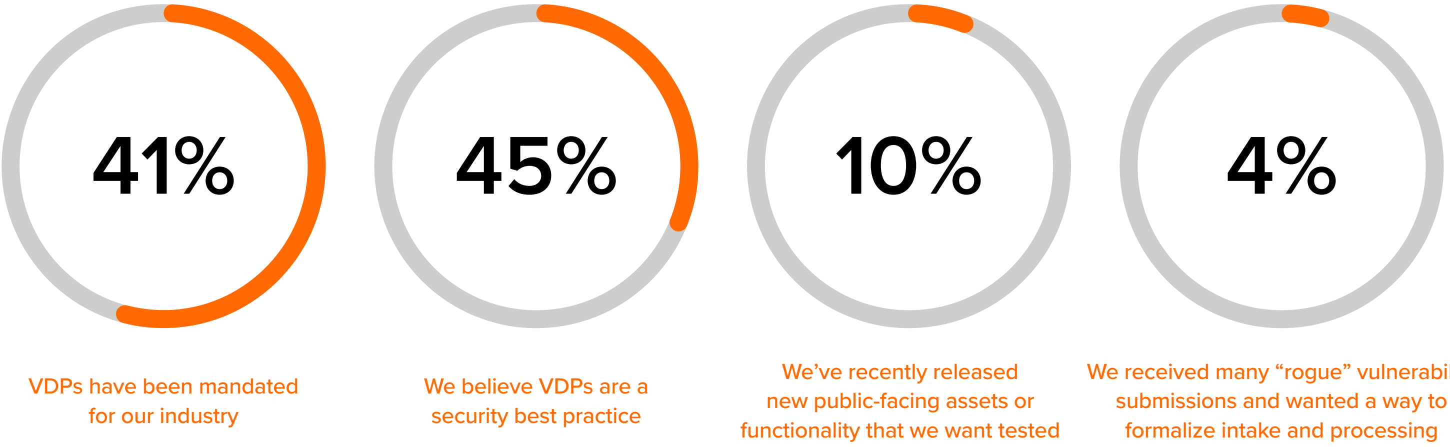

Of course, that's not all they're motivated by. The report goes on to show that researchers are also incentivized by **education, rewards, and recognition.** But unfortunately, 'recognition,' is all too often lumped in with 'reward.' Rewards and recognition are both gestures of appreciation, but are each rooted in different measures of value. VDP rewards may come in the form of kudos points, store credits, or, on occasion, payments. Recognition in a VDP program goes beyond the organization's acknowledgement of the researcher's contributions, and instead refers to the ability for the researcher to have their contributions recognized by the broader security community. It is *global* recognition, through disclosure.

79%
of organizations with a VDP have awarded monetary payments for exceptional findings

WHAT IS DISCLOSURE?

Sharing security vulnerabilities with the world enables similar organizations to get ahead of threats before they become larger problems. Communicating how and when these vulnerabilities were uncovered can **drastically reduce the frequency of their creation,** while improving the ability of security researchers to more readily spot related issues. Additionally, according to recent Bugcrowd research, **organizations that adopt disclosure terms see 30% more vulnerabilities than organizations that don't.**

Programs on the Bugcrowd platform that adopt disclosure terms see **30% more** vulnerabilities on average, versus organizations without

"Disclosure" has several meanings, referring both to the communication of a vulnerability to the organization within which it was discovered, and to external parties, usually in a public forum. While the first definition benefits the organization, and by extension, its direct customers, partners, and other stakeholders, the second, when done right, **benefits the entire digitally connected world.**

Coordinated disclosure terms set out our definition of good faith in the context of finding and reporting vulnerabilities;; they encourage rapid remediation while demonstrating commitment to, and appreciation of, the hacker community.

However, the term "disclosure" does carry an unfortunate and **misplaced stigma,** which is holding back security standards globally. A quick exploration of the varying types can help to clarify terms, and alleviate unfounded concerns.

THE SPECTRUM OF PUBLIC DISCLOSURE

Discretionary Disclosure

When organizations opt to enable coordinated disclosure, they signal **openness to consider public disclosure** of remediated vulnerabilities, in full or in redacted form, on a case-by-case basis. Ultimately, while disclosure may be requested by the finder, it remains at the sole discretion of the organization. Removing a vulnerability from consideration for coordinated disclosure is sometimes necessary when disclosing it creates significant risk to customers. This is the case with pacemakers, vehicles, and other IoT devices that are difficult to quickly recall or update remotely.

Coordinated Disclosure

For more mature organizations, setting a "timer" on resolution and publishing for every vulnerability can further encourage more active discovery, though this structure often requires a dedicated team responsible for rapid remediation and communication. This approach is often taken by organizations who deem security to be a **strategic priority,** and need to invest in building the best possible relationship with the security community.

Coordinated disclosure is based on good faith, and is considered best practice for all parties involved as it **encourages rapid remediation** while demonstrating commitment to, and appreciation of, the hacker community.

Full Disclosure

Unlike the other approaches, full disclosure is not a program policy. Rather, it is an individual instance of public communication wherein the finder discloses a vulnerability before it has been fixed. Bruce Schneier defended the merits of full disclosure in 2007, suggesting that the threat of this act is sometimes necessary to force owners to fix vulnerabilities when they are unresponsive to hackers' well-intended communications. However, both parties often prefer to avoid this type of disclosure at all costs.

In fact, both non-disclosure and full disclosure are discouraged because of the **asymmetric cost** to only one party; either the finder is not given recognition for their effort to improve security, or the owner is not given an opportunity to fix a vulnerability before it becomes socialized in way that makes it more likely to be **maliciously exploited.** Disclosure should be undertaken in a way that protects the owner, rewards the finder, incentivizes further research, and enhances relationships between owners and the security community.

Non-Disclosure

When programs are marked as "non-disclosure," it is understood that the finder is not permitted to communicate any portion of a vulnerability beyond the confines of the organization itself, even after it has been resolved. For non-disclosure programs, no vulnerability, regardless of type or severity, can be shared. While these programs still receive submissions, they do not *encourage* them.

DOES YOUR ORGANIZATION ALLOW COORDINATED OR DISCRETIONARY DISCLOSURE?

WHEREIN THE FINDER OF A VULNERABILITY IS PERMITTED TO PUBLICLY DISCLOSE DETAILS OF THE FINDINGS AFTER IT HAS BEEN REMEDIATED?

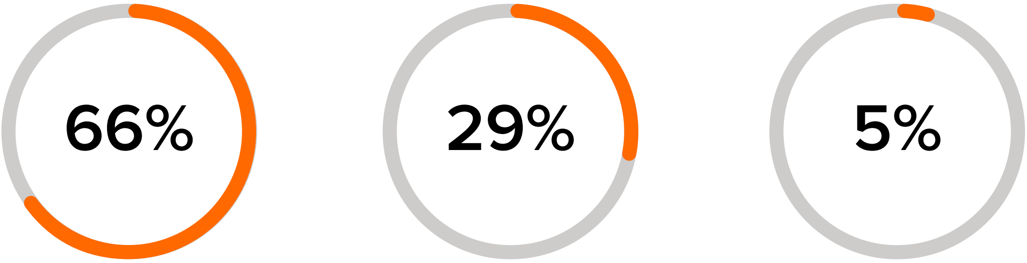

Yes, we allow virtually all vulnerabilities to be publicly disclosed

Yes, we allow some vulnerabilities to be publicly disclosed either fully, or with certain details redacted

No, we never allow vulnerabilities to be publicly disclosed

OBSTACLES TO DISCLOSURE

In addition to improving the security posture of other organizations, coordinated and discretionary disclosure policies **strengthen the relationship** between the organization and the security researcher community. Security researchers' reputations are their brands, and receiving acknowledgement for identifying an exceptionally complex vulnerability enhances the finder's reputation and increases their market value. Organizations that clearly state their willingness to collaborate on disclosing vulnerabilities in advance can expect better relationships with the security community, and often **greater program activity.** While the rationale seems straightforward enough for both parties, it's not quite that simple for many organizations.

Christian Toon, CISO at law firm Pinsent Masons, notes that perceived duties to stakeholders and the board can harm the outlook of certain owners when it comes to disclosure. "Many organizations see disclosure of a vulnerability to be an unnecessary admission of weakness that harms their reputation, but this is a short-term outlook" he states. "Embracing vulnerability disclosure creates a **security-first mentality,** builds your reputation within the security community and educates your board in the process. That way if there is ever a breach, the standard line 'we take our security seriously' will carry far more weight."

Some security activists will say the threat of full disclosure is necessary to keep owners honest and incentivize them to fix vulnerabilities. Many owners will say legal protections are necessary to prevent the threat of full disclosure becoming a vector for blackmail. A **solid legal framework** that recognizes the motivations of all parties is the best basis to facilitate vulnerability reporting and remediation.

NIST CYBERSECURITY FRAMEWORK AND LEGAL IMPLICATIONS

In the past year, the U.S. Federal Trade Commission (FTC) and Department of Justice (DOJ) have released guidance outlining the need for vulnerability disclosure programs (VDP). With support from major legislative bodies like the National Institute of Standards and Technology, widespread adoption of vulnerability disclosure programs is expected and necessary in the coming years.

WHAT IS THE NIST CYBERSECURITY FRAMEWORK?

The NIST Cybersecurity Framework is a set of policies meant to help the private sector in **strengthening their cybersecurity readiness and awareness.** The framework is published by the National Institute of Standards and Technology (NIST), under the US Department of Commerce.

Originally designed for critical infrastructure IT, it has since been adopted by private sector organizations as part of their risk management and cybersecurity practices. In fact, it's estimated that half of the organizations in the US use the framework. It has also been adopted by the information security agencies of other countries, including Italy, Israel, and Japan.

UPDATES TO THE FRAMEWORK

Since its inception in 2014, the framework has been updated several times to keep up with evolving threats. Version 1.1 was released in 2017, which included guidance on performing self-assessments, supply chain risk management, and vulnerability disclosure.

This revision is the result of a massive industry effort. During the spring of 2017 a number of organizations, including Rapid7, Duo Security, Cisco, Symantec (and yours truly, Bugcrowd) submitted a letter in response to NIST's call for public comment on the framework.

After the updates, it now includes the following:

RS.AN-5: Processes are established to receive, analyze and respond to vulnerabilities disclosed to the organization from internal and external sources (e.g. internal testing, security bulletins, or security researchers)

This language is very close to that suggested in the letter's primary recommendation: "Processes are established to receive, analyze, and respond to vulnerabilities disclosed to the organization from external sources."

The revised framework also mentions researchers in its Tier 1 implementation (pg. 10): This is an exciting addition and one that paves the way for the whitehat community to partner with organizations.

WHAT THE NIST UPDATE MEANS FOR VULNERABILITY ASSESSMENT

These updates mark an incredibly important move by the NIST. The news comes on the heels of another year of escalating cyberattacks and a growing focus from the federal government on vulnerability disclosure.

In the 2020 report by RiskBased Security, it was reported that **36 billion records were exposed by data breaches in the first half of 2020 alone.** Although the increased scope of cybersecurity threats is unfortunate, their sheer volume is causing policymakers to respond, and that's a positive thing.

Adding to the positive changes, the White House recently released the Federal IT Modernization Report. This report positions vulnerability disclosure as the best-practice approach to external security testing for the U.S. Government. This is another major step forward not only for the bug bounty model, but most importantly, for the security of everyone in the U.S.

Last year was undoubtedly another year of escalation in size, scope, and scale of cyberattacks. It goes without saying that most Americans have personally been impacted by a breach or know somebody who has been impacted.

> If the Department of Defense can have a VDP—*anyone can.*
>
> ETHAN DODGE • SECURITY ENGINEER, ATLASSIAN

VULNERABILITY DISCLOSURE LEGAL STATUS

Aligning the interests, incentives, and expectations for both hackers and host organizations primarily involves **frequent and clear communications,** but there is also a need to provide unambiguous **legal clarity and assurance.** Hacking involves testing, stressing, and sometimes even breaking software to rebuild and improve it. This creates problems in a legal system that defaults to ownership as a starting point, and presumes malice to be the motive for any party who uses and abuses software outside its supposed scope. As a result, the default legal status for vulnerability discovery and disclosure excludes good faith research.

The Computer Fraud and Abuse Act (CFAA) prohibits accessing a computer without authorization, or exceeding authorized access. This renders good faith testing of assets illegal where robust VDPs are not in place, and while the

number of hackers convicted for related offences is low, it has a chilling effect on the community; 60% of hackers do not submit vulnerabilities due to fear of legal retribution.

The Digital Millennium Copyright Act (DMCA) makes it illegal to circumvent controls that prevent access to copyrighted material, defined to include software. This applies even to legal owners of the products in question.

These laws were passed during a time when hacking was mostly done maliciously, before the advent of bug bounties, good faith hacking and a thriving community of professional security researchers. While the DMCA was amended in 2016 to allow security researchers to work on owned consumer devices in good faith, there are still legal gaps that need to be resolved before organizations can fully benefit from VDPs.

Organizations must draft terms of VDPs to allow and incentivize good faith testing and submission of vulnerabilities, in a way that keeps lawyers happy by ruling out backdoor entry points or loopholes for malicious actors. These agreements create legally robust **“safe harbor”** for well-intentioned researchers, which considerably increases the number and quality of vulnerabilities submitted.

One starting point to consider is Disclose.io, an open source standardization project that offers a boilerplate VDP framework instilling safe harbor and enabling good faith security research. This provides an **accessible legal agreement** for the research and disclosure of vulnerabilities, and uses standardized terms and policies to create a more welcoming space for hackers and researchers, many of whom do not speak English as a first language and have minimal legal knowledge. The safe harbor terms from disclose.io were adopted in 2020 by CISA DHS, the voting machine manufacturers, and a number of US States to encourage transparency and reporting of cybersecurity issues that could potentially impact elections.

VDP BEST PRACTICES

HOW TO START A VDP

Having a VDP is quickly becoming industry standard, and in fact is no longer optional for some. The Cybersecurity and Infrastructure Security Agency (CISA) issued a binding directive requiring all federal agencies to publish a VDP.

However, starting and managing one effectively can seem overwhelming. There are **five key steps** that every organization can follow to build a strong problem:

1. DECIDE ON SELF-MANAGED OR HOSTED

Organizations like Bugcrowd offer managed vulnerability disclosure programs to help **alleviate the time and effort** required to construct and run an effective disclosure program. Bugcrowd provides access to a cloud-hosted secure submission framework that enables individuals to submit security feedback from anywhere in the world. The fully-managed process includes design and management of email and website-embedded submission forms, validation, categorization, and prioritization of vulnerabilities, integration with the organization's software development tools for faster remediation, and researcher communication, points-based remuneration, and support. Additionally, leveraging Bugcrowd for program management with the option to have the program listed on Bugcrowd's researcher homepage, brings your program to the attention of registered hackers and researchers for increased likelihood of additional activity and submission volumes.

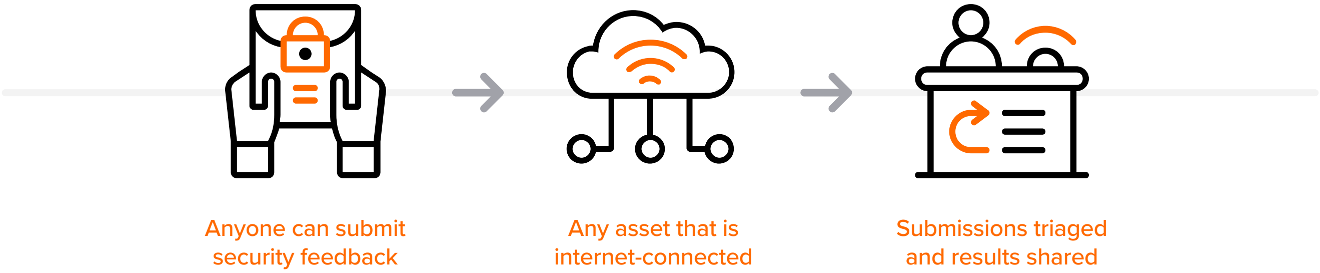
Anyone can submit security feedback
Any asset that is internet-connected
Submissions triaged and results shared

Companies with few internet-facing assets, limited resources, or still-maturing processes for accepting and remediating vulnerabilities may instead choose self-management, which usually equates to a more manageable flow of vulnerabilities. Of course, it's possible incoming submissions may outpace the ability of a thinly resourced team to respond in time, which can lead to tension between researcher and organization if communications are not prioritized. This tends to expedite the transition to a managed model, especially as evidence of urgency is usually quite easy to demonstrate to superiors.

2. CODIFY EXPECTATIONS

Organizations initiating a VDP should adhere to principles that **make the program scalable and robust.** This includes providing clear authorization of access to good faith researchers. This should include broad indications regarding acceptable conduct, as well as techniques that could be considered out of scope, DDoS or social engineering, for example. It should also determine the scope of assets covered by the policy, with restrictions for third party data or personal information, or a requirement that hackers use test accounts and dummy data when testing for vulnerabilities. Organizations with limited resources may also want to restrict the assets covered by the program to start with, to ensure they have the resources to deal with vulnerabilities submitted.

3. EXPECTED TO ITERATE

Saketh Mailapur, Information Security Manager at consumer goods company Unilever, stresses the importance of **laying out a timeline and allowing time to build and review a data set.** “Starting a VDP can be overwhelming, so commit to a phased timeline that allows plenty of space for gathering data and making adjustments.” Mailapur points out that the scope should be revised in line with this data, “traffic on one site went up over 500% when we put our VDP in place, so we adjusted our policy to rule out automated scanning.”

No organization will land on their ideal scope, preferred disclosure policy, and most efficient communication process at their first attempt, so the best approach is to build iteratively. Toon says, “I advise those starting out with VDPs to fail fast and fix fast. Play around with parameters and approaches and gather plenty of data to inform yourself. As long as you don’t annoy or offend the security community or your board it will all be valuable.”

4. BE ACCESSIBLE

It is also important to give **clear guidance around communications,** within dedicated channels. This could be a security@companyname.com email address to begin with, but it is crucial to avoid single points of failure. Multiple channels, safeguards, and responsible parties can prevent an unchecked inbox or overactive spam filter from creating blind spots and associated risk.

5. FACTOR IN RESPECT

Finally, and perhaps most importantly, a VDP should **define clear disclosure standards based on good faith.** These define the strength of the relationship with hackers, and should align incentives to ensure that both parties benefit from interactions. Owners should receive as much detail of the vulnerability as possible, along with a good faith commitment from the hacker to stick to the agreed method of disclosure. Hackers should expect to get prompt replies to their submissions, and a **commitment to giving them appropriate recognition.**

78%

of organizations with a VDP are running more than one program

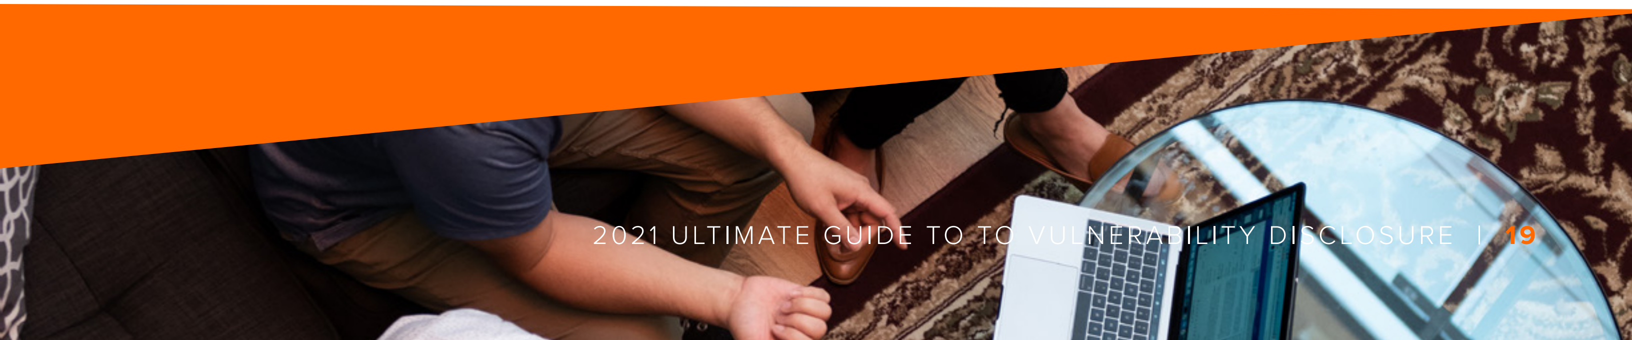

MANAGING A VDP

Those willing to implement best practice in vulnerability disclosure can both set a standard amongst peers while differentiating themselves against their competitors. Here are some steps that can make VDPs work best for organizations, partners and the security community.

- **Align expectations** – Researchers should feel legally protected and know exactly how to report a bug and what to expect throughout the process. Don't be afraid to over-communicate.
- **Provide clear legal guidance** – Use standardized terms and clear examples to encourage good faith interaction, and authorize conduct under CFAA by providing explicit consent to access systems.
- **Ground interactions in good faith** – Allow for the accidental overreach of scope by hackers done in good faith. Ensure your policy prioritizes relationships and industry norms over strict interpretations of the guidelines.
- **Remediate efficiently** – Prioritize your end users and the vulnerability finder by getting to work resolving the bug and validating the fix quickly.
- **Start a dialogue** – VDPs are a two-way street and there are long term benefits to working on your end of the relationship with hackers through clear communication and appropriate incentives.
- **Troubleshoot the process** – Remove single points of failure in communications channels, seek feedback from researchers and commit to flexibility in your VDP philosophy and operations.
- **Take an integrated approach** – VDPs are just one in a number of overlapping tools and procedures that make up your security posture. Ensure all processes and products are configured to move in the same direction.
- **Know your limits** – Depending on your current security posture, VDPs can be overwhelming. Work with your team and/or VDP provider to configure a manageable solution.

COMBINING VDPS WITH BUG BOUNTY OR PEN TESTING

Bug bounty programs allow organizations to direct targeted, rigorous testing at business-critical assets. Similarly, pen test programs enable organizations to focus on compliance-related assets, or those where a structured methodology would improve how security posture is communicated to partners, investors, and customers. Vulnerabilities found through these programs qualify for financial rewards, so most organizations limit scope for budgetary reasons, and may also impose limited testing windows. While economical, this creates gaps in coverage, and wrongfully assumes that all potential vulnerabilities can and will be surfaced through an exclusive (often private) crowd of researchers.

NIST 800-53 r5 codified the idea that a public bug bounty program is actually a subset of a VDP, and is specifically a VDP where monetary rewards are optionally offered as thanks to the finder.

Each program has its strengths and limitations. Toon at Pinsent Masons notes, "pentesting has been recognized and accepted by the audit community, which makes it useful for assets where compliance is of particular importance. But the limits in scope and partners involved means it can become rigid and less effective over time." VDP programs add a much needed, yet economical, catchment for vulnerabilities surfaced by anyone, anywhere. But when is the right time to implement?

99% of organizations already run or would consider running a pay-per-finding bug bounty alongside their VDP for targeted testing on priority assets

The market has tied itself in knots trying to create a linear maturity model for when and how to "progress" between a VDP, Bug Bounty, and/or Pen Test. However, each should be viewed as providing complementary benefits, with adoption driven by individual goals and resources rather than maturity. Atlassian's Adeloye considers a VDP to be the first building block in external testing-- "a superset that can include a bug bounty program," though it's equally common for VDPs to be the final addition to a comprehensive crowdsourced approach. While an agreed upon sequence might make for tidier budgeting, it also goes against the organic, adaptive, and sometimes unruly nature of security. Every organization is different.

BUGCROWD
VULNERABILITY DISCLOSURE PROGRAMS

Let's go back to the common security goals that we mentioned earlier in the report. While organizations deal with a poor signal-to-noise ratio, difficulty prioritizing what matters, and rapid release cycles that leave vulnerabilities unaddressed, Bugcrowd Vulnerability Disclosure Programs help reduce risk, improve security ROI, accelerate digital transformation, and help organizations drive better decisions through contextual intelligence. Bugcrowd provides a framework to securely accept, triage, and rapidly remediate vulnerabilities submitted from the global security community. These can range from web applications, to APIs, to mobile apps, IOT devices, and more. Bugcrowd VDP integrates with your security/development stack to quickly fix vulnerabilities. This all happens from our all-in-one SaaS platform.

The platform combined with our in-house team of security experts escalates high-priority issues within hours and averaging triage completion within one business day. This means that all the vulnerabilities coming through the researcher crowd are first vetted by Bugcrowd, according to objective rating standards like our own Vulnerability Rating Taxonomy (VRT).

In short, Bugcrowd removes virtually all overhead for your security team so they can focus on reducing risk by remediating the vulnerabilities identified—making your security team happy and efficient. Now more than ever, vulnerability disclosure programs help organizations thrive in the digital era.

> Bugcrowd's Vulnerability Disclosure Program is one of the best value for money services that we have. The annual cost of the program is the same cost of one traditional penetration test and the VDP has given us around a 100-fold increase in actionable intelligence.
>
> **DAN MASLIN**
>
> CISO, MONASH UNIVERSITY

Curious to learn how your organization can leverage a Vulnerability Disclosure Program? Contact us now!

www.bugcrowd.com/contact-us