bugcrowd

Why Crowdsourced Security?

Highlights

There is a fundamental imbalance between the creativity and motivations of cyber attackers, and those of enterprise security defenders.

Crowdsourced security eliminates this imbalance by harnessing white hat security researchers to find and eliminate vulnerabilities.

Crowdsourced security provides focused results to support rapid risk reduction, cost control, and lower operational overhead.

Crowdsourced security supports the most critical attack surfaces: web and APIs interfaces on server/cloud, mobile and IoT platforms.

Highly vetted, trusted security researchers and private programs diffuse concerns of risk associated with crowdsourced security.

Partnering with an established crowdsourced security platform largely eliminates administrative overhead and maximizes risk reduction.

Crowdsourced Security is a powerful tool - used by leading edge firms such as Google and Facebook - to decrease risk. However Crowdsourced Security is not yet well understood across the enterprise security community. This brief will define Crowdsourced Security and describe why it's a key element of any viable security architecture.

"Cybersecurity isn't a technology problem — it's a human one — and to compete against an army of adversaries we need an army of allies."

Casey Ellis, Founder, bugcrowd

The Attacker-Defender Mismatch

These days it's common knowledge that current approaches to enterprise security are not working. Usually the thinking is that the problem is a combination of complex yet ineffective technology and not enough staff. These points are certainly true, but the problem is more fundamental. The bigger issue is that there's a fundamental imbalance between the approach of the attacker and that of the defender. That imbalance is centered on human creativity, incentives and motivations.

Let's start with the area of human creativity. Cyber attacks are driven by teams of people, and their collective creativity drives their success. Black hat hackers are constantly at work, and finding vulnerabilities is something of a competition: Whoever solves the puzzle first "wins" and gets to decide how to exploit their discovery. Attackers share information and techniques, but combine that with their own creativity as they seek to overcome security defenses.

80% **of cyber-attacks are driven by organized crime rings, in which data, tools and expertise are widely shared.**

Enterprise security on the other hand focuses on technology, supported by a team that struggles to operationalize everything on their plate. Technology brings incredible leverage, but isn't as creative as humans, and cyber security is a fundamentally human problem. It's possible that someday technology will be able to mimic and predict human creativity, but that's not coming anytime soon. Technology is a lever, not a replacement for human input.

The second imbalance is that of motivations and incentives. The attacker is focused on results: they get nothing unless they're successful. Their motivation is centered on finding and exploiting a vulnerability, and they only have to succeed once to meet their objective. Since there's no "partial credit" for trying and failing, they will keep at it until they either succeed, or decide that there are other more promising targets. Of course this last point doesn't apply to the most sophisticated attackers that are after specific targets - they will simply keep at it as long as necessary.

The defender's situation is totally different. While the CISO's job might be tied to the lack of successful hacks, the security organization has very different incentives in place. The defenders are paid a fixed salary, and held to industry expectations for security architecture and operations. Their focus is on setting up credible defenses. They will leverage technology to set up a security architecture and incident response process, maybe do a little penetration testing using commonly known methodologies, and then hope for the best. Additionally, most enterprise security staff have no experience in actual hacking. Lastly, software engineering teams are strongly incentivized to develop code that supports the business as fast as possible, and security is an afterthought. In summary, enterprise security teams have the wrong incentives and goals in place, and they only have to miss one attack attempt to fail.

Crowdsourced Security in the Automotive Industry

FCA
FIAT CHRYSLER AUTOMOBILES

The automotive industry has rapidly embraced Crowdsourced Security in response to the risks of hacking. The public demonstration of successful vehicle hacking and subsequent recall of 1.4 million Fiat Chrysler vehicles in 2015 was perhaps the most well-publicized event, but there were many others. In one such demonstration the researcher was even able to subvert one of the vehicle electronic control units and use it to override the attempts of another ECU to maintain control of the car while under attack.

As a result, all the major automobile manufacturers now have public and private Crowdsourced Security programs, including Fiat Chrysler, demonstrating that Crowdsourced Security can suit most any type of business, not just technology firms.

Crowdsourced Security: A Human-Based Approach to Risk Reduction

Crowdsourced Security is a simple, powerful concept: Put the collective creativity of crowdsourced researchers on your side. In this model, you leverage a team of good-faith hackers (known as researchers) to harness the same creativity your adversaries are using to attack you. The concept is straightforward:

1. You define the attack surfaces you need to harden, for example web application front ends or a mobile application.

2. Depending on the type of program, you either publish the program broadly to the researcher community, or engage a more limited set of researchers in a private “invite only” program.

3. As vulnerabilities are found by the researchers, they are triaged to determine validity and severity.

4. You pay a reward (or grant public “kudos”) to the researcher for finding the problem, patch the vulnerability, and verify that the attack vector has been closed.

A key attribute of crowdsourced security is that motivations and incentives are aligned:

True Risk Reduction: Researcher compensation is tied to successful outcomes, that is, finding dangerous vulnerabilities that you need to know about

Speed: The first researcher to find a vulnerability is rewarded, encouraging researchers to work quickly.

Value: The more serious the vulnerability found, the bigger the reward to the researcher

This alignment drives the efficiency of the overall effort - everyone is aligned to the true goal of risk-reduction.

> “Bug bounty programs are going to become a baseline security control that you need to have.”
>
> **Alvaro Hoyos, CISO,** onelogin

In addition to the core elements of human creativity and incentives, Crowdsourced Security also offers a number of secondary but important advantages over traditional security approaches. First, Crowdsourced Security can easily be integrated into the software development lifecycle (SDLC). Finding vulnerabilities is a great start, but decreased risk is only achieved when vulnerabilities are actually eliminated. A solid Crowdsourced Security program integrates findings into software lifecycle tools such as Jira, making it efficient for development engineering to patch vulnerabilities.

Additionally, Crowdsourced Security is a perfect fit for agile, continuous development situations. As organizations move to agile CI/CD (continuous integration and development) software practices, security is struggling to develop risk mitigation strategies that can keep up. But ongoing Crowdsourced Security programs can easily respond to this challenge: As new code is rolled out, the researcher community can be immediately notified. This creates an incentive for them to quickly analyze the new code, since for them new code means new vulnerabilities to be found.

Crowdsourced Security supports today’s critical attack surfaces, on all key platforms.

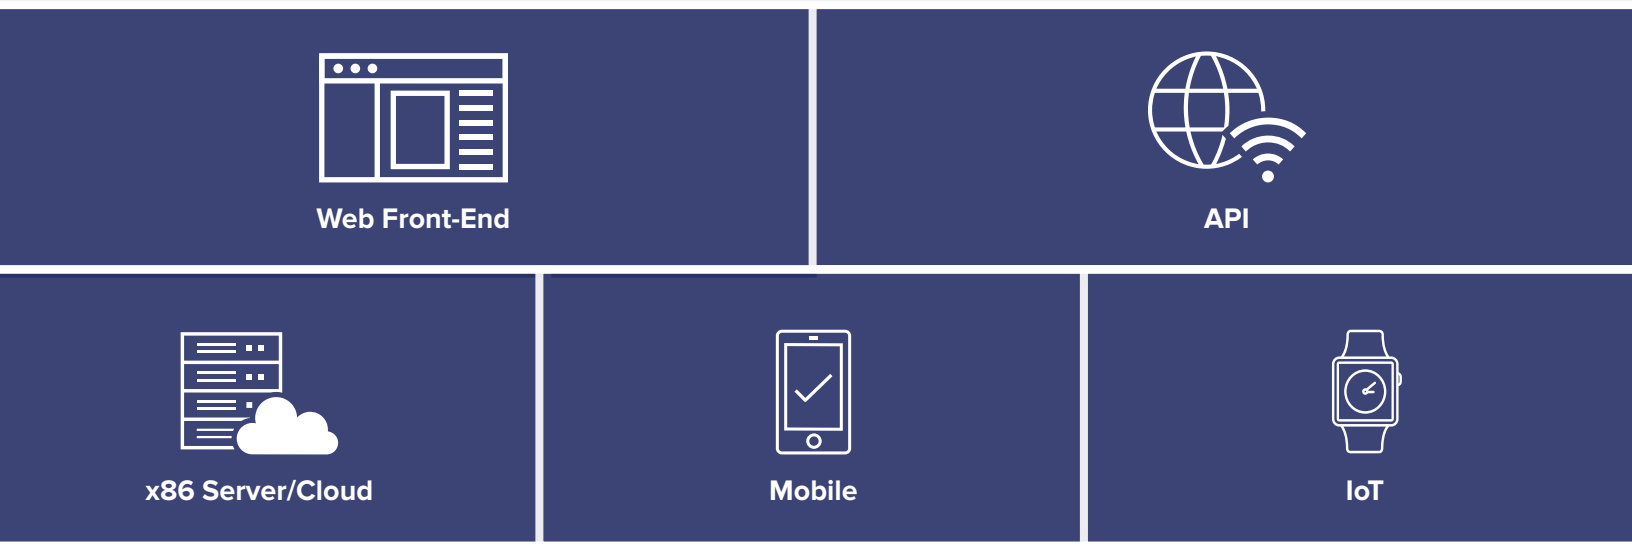

Crowdsourced Security supports today’s key attack surfaces, as well as “the unknown”. As organizations move to cloud architectures and applications, the biggest concerns are web application front ends and APIs, which may be deployed on IoT devices, mobile apps, or on-prem/cloud. All of these can be evaluated for risk by Crowdsourced Security. Furthermore, a public crowd program can uncover risk in areas unknown to the security organization, such as shadow IT applications or exposed perimeter interfaces.

Finally, using Crowdsourced Security lowers security costs and operational overhead. There is no agent software on applications or clients, and no software instrumentation to support. There are no network devices or virtual appliances to install and manage. There is also little to no operational waste caused by false positives or low-priority events. And as has been noted, the reward payments to the researchers are completely based on actual risk identified. As security budgets come under increasing scrutiny, crowdsourcing becomes an obvious choice for simultaneously controlling costs while still aggressively protecting the business.

Crowdsourced Security & IoT: Aruba/HP Enterprise

aruba
a Hewlett Packard Enterprise company

As part of HP Enterprise, Aruba takes security very seriously and had a number of best practices in place to ensure their wireless products were secure. However, they still were concerned that they might be vulnerable, so they started a private bug-bounty engagement using Bugcrowd's platform. The results speak for themselves: They got their first high-severity vulnerability in the first 24 hours, and have received and accepted over 500 submissions in the first two years of the program. Based on that success, they have expanded their program to include more products and broader group of researchers.

Defusing the Risk Perception

Crowdsourced Security is clearly powerful, but common misconceptions remain. The primary concern is one of increased risk: the feeling that by exposing vulnerabilities to the crowd you are increasing the risk that those vulnerabilities will be exploited. To address this concern, it is important to first recognize that Crowdsourced Security programs can be private or public. Most programs are private, which means the details of the target are never shared publicly. Also in a private program, the researchers are limited to a hand-picked team based on multiple dimensions, in particular their expertise on the target and the level of trust they have earned.

The broader consideration is the level of trust in the researchers. Researcher trust is built up over time, based on actual behavior. In a well-run Crowdsourced Security program, you have many more insights about the researchers than what you gather during the vetting process for a typical employee or outsource firm. In very sensitive situations, you can require that researchers be background checked, geographically controlled, or even government security cleared.

Another point worth discussing is the level of exposure a target is getting outside of the Crowdsourced Security program. For targets such as Web or mobile applications, it's obvious the targets are already known to the world and exposed to attack.

Lastly, well-run programs require storing vulnerability submission details on a hardened platform that supports multi-factor authentificationn, data-at-rest encryption, and full security logging.

Starting a Crowdsourced Security program doesn't change that reality. In the case of enterprise internal applications, the network "perimeter" is very porous, and security best practices now assume the perimeter is breached. Therefore, it is ill-advised to assume that an internal application is unknown and will remain unknown to the attack community.

> "For us, the managed approach reduced our required time and effort by at least 80%, freeing up our security team to focus on other components of our security program."
>
> **Johnathan Hunt, VP, Information Security,** invision

The other common concern surrounds increased operational overhead, having to do with handling payments, processing high volumes of submissions, and even the hassle of understanding the (admittedly unique) culture and operating ethos of the hacker community. In the case of Bugcrowd as the program operator, this overhead is practically eliminated. At program outset, Bugcrowd creates a program brief, which sets out all the parameters, conditions, and payout levels that is communicated to the researcher community. Once the program starts, all researcher interactions are handled by Bugcrowd, including payouts and dispute resolution. All vulnerability submissions are verified, prioritized, and de-duplicated before they are forwarded to the customer for action. Bugcrowd also handles distribution of test credentials, monitors researcher engagement, and works with the client to recommend program changes such as researcher pool composition and payout levels as needed.

Summary

Current security approaches are largely ineffective because of the imbalance between attacker and defender with respect to creative approach and incentives. Crowdsourced Security eliminates this imbalance by providing a defensive approach that matches the resources and creativity used by the adversary. Additionally, Crowdsourced Security offers lower operational overhead and costs, supports today's key attack surfaces, and aligns well with both waterfall and agile/continuous development. Therefore, it should be a part of any organization's security architecture.

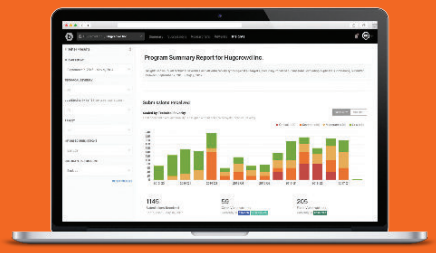

Getting Started

Want to learn more about how your organization can leverage vulnerability disclosure programs to start discovering and fixing high-value vulnerabilities missed by traditional security testing? Bugcrowd offers a full line of crowdsourced security solutions.

www.bugcrowd.com/get-started