Inside the Mind of a CISO

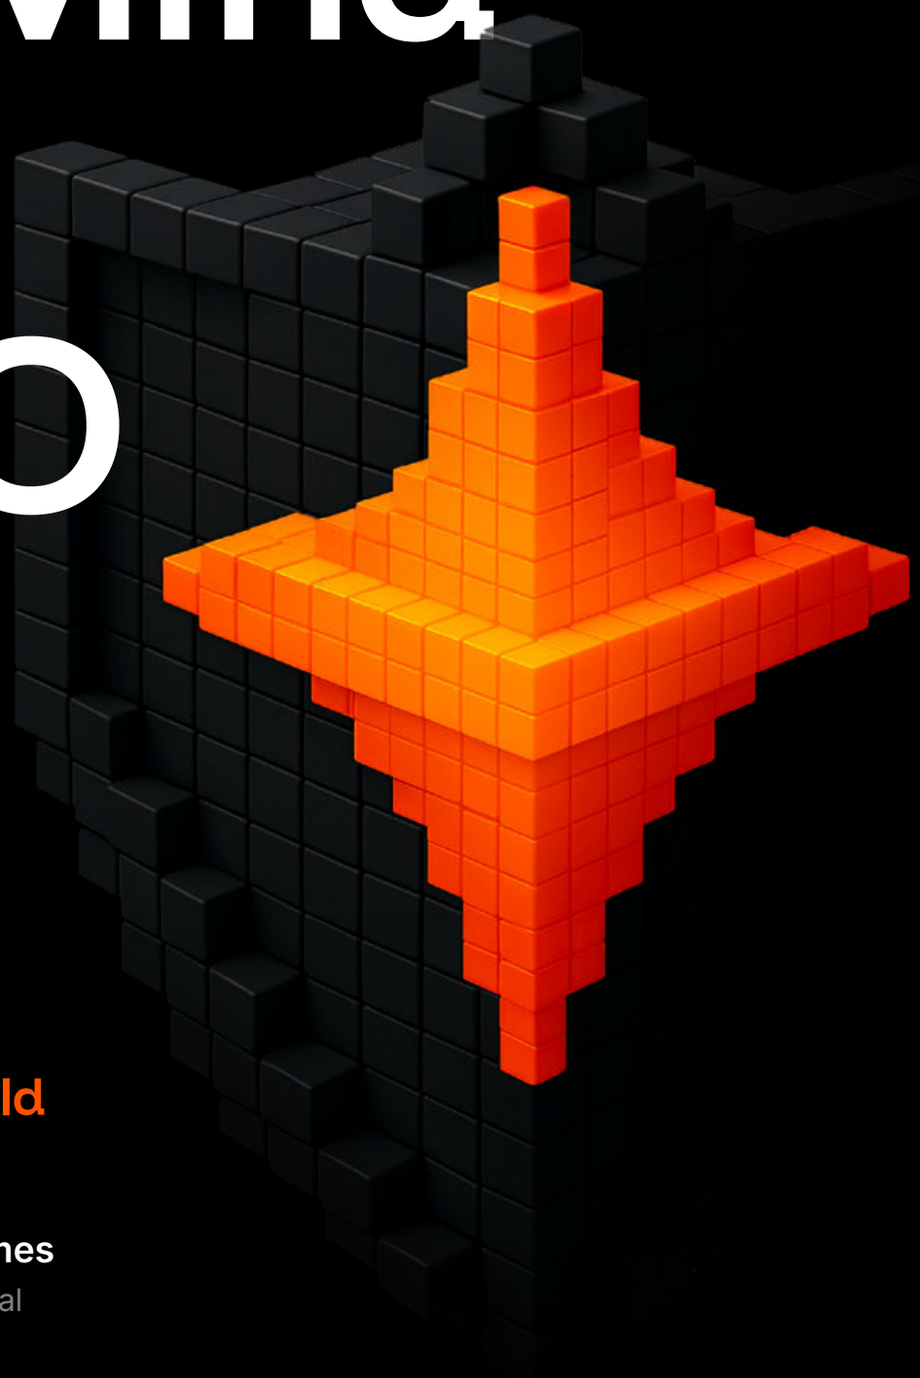

Resilience in an AI-accelerated world

The CISO's challenge: Measuring security outcomes
"Security without true adversarial testing is just an illusion"

The dos and don'ts for your next board deck

Ask a hacker
5 vulnerabilities to watch out for

bugcrowd

Table of Contents

An introduction from our CI&SO Nick McKenzie

As security leaders, we are in the thick of so much change. AI is everywhere, and frankly, many of us are sick of talking about it. How do we separate the wheat from the chaff?

We are in a high-stakes innovation race, but with every AI advance, the security landscape becomes exponentially more complex. Attackers are exploiting this complexity, but still targeting foundational layers like hardware and APIs. No single CISO can win this race alone. To thrive, we must move beyond isolated efforts and cultivate a collective resilience of collaboration—pooling our knowledge of the hacker community to outpace emerging threats together.

This community-driven approach is the only way to stay ahead—defeating attackers as one unified force. That's why, in this report, we are sharing a range of insights for CISOs, from vulnerability hot-spots to watch out for all the way through to strategies to confidently communicate with board members and justify investments.

The CISO position may be at the top of security leadership, but it's strengthened most by the collective intelligence of our community

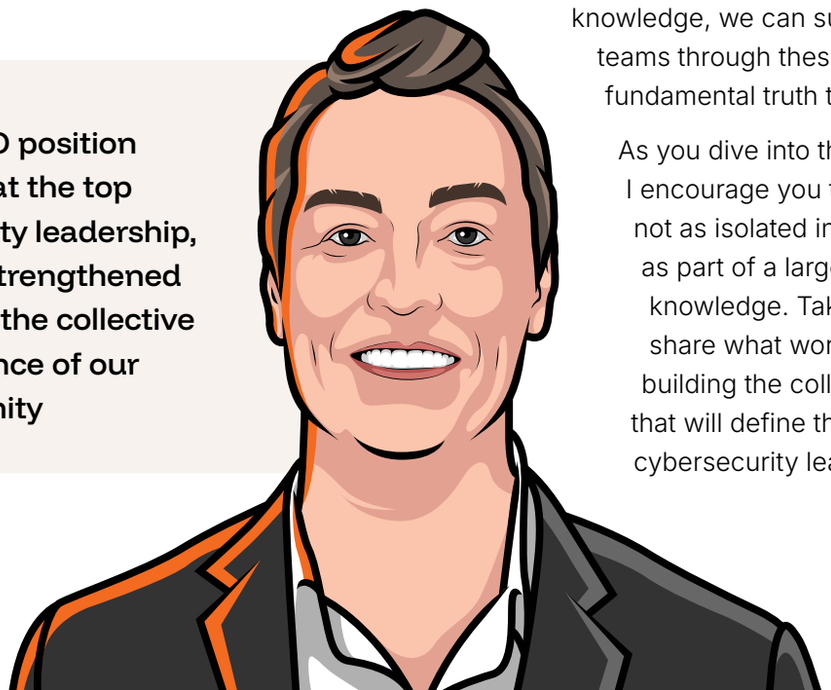

Let's look back on my original question about knowing where to anchor ourselves when it comes to AI. The reality is that we are at an inflection point. AI has fully taken over conversations about offensive security, so where do human testers and analysts end and where does AI begin? I guarantee that the moment you figure out how to answer that question, the goal posts will move again.

The key to success is understanding the role of humans, the role of AI, and the fact that the balance between the two will change over time. We can't get lost in the buzzwords. While others race to introduce flashy AI workflows and copilots, it is my belief that now is the time for sensible decision-making and adopting AI models where they make sense and provide true value.

Ultimately, CISO confidence in an AI-accelerated world comes from continuous, community-powered testing augmented by AI that translates risk for the board. This is what results in true security resilience.

The challenges we face are daunting, but they're not insurmountable when we work together. By tapping into a larger collection of knowledge, we can successfully lead our teams through these chaotic times—a fundamental truth this report highlights.

As you dive into this report, I encourage you to view each insight not as isolated information, but as part of a larger community of knowledge. Take what resonates, share what works, and continue building the collaborative spirit that will define the future of cybersecurity leadership. ■

Vulnerability trends

The trends, patterns, and themes we're seeing from hundreds of thousands of vulnerabilities submitted through the Bugcrowd Platform.

Recommendations

This digital magazine is made up of 10 articles, all examining different aspects of the CISO experience right now, whether you're a first-time CISO, a seasoned vet, or even an aspiring security leader. It's jam-packed with information, but **for those in a hurry,** here are a few highlights paired with actionable tips.

THE TOPIC	THE TL;DR	WHAT TO DO NEXT
The rise of API and hardware vulnerabilities	Last year, Bugcrowd saw an 88% increase in hardware vulnerabilities and a 10% increase in API vulnerabilities. 81% of researchers and hackers cite that they've encountered a new hardware vulnerability they had never seen before in the past 12 months. Prioritizing API and hardware testing ensures we proactively protect our systems and hardware so that CISOs can be more resilient and deliver secure experiences to users downstream.	→ Consider adding APIs and hardware to the scope of your offensive security testing programs. 
The need to operationalize attack surface intelligence	As the damage from cybercrimes increases rapidly, CISOs can't afford to wait weeks or months to act on their attack surface intelligence. To help CISOs truly reduce risk, security teams must integrate EASM intelligence into their offensive testing platforms so that there's a direct path from discovery to remediation.	→ Adopt an integrated approach to attack surface intelligence to demonstrate measurable improvements in security efficiency and faster remediation cycles. This enables you to prove the value and outcomes of a security program to external stakeholders.
The gift of objective feedback	Getting objective perspectives on where and why you are vulnerable is crucial for any CISO looking to build a stronger security program. The most mature organizations don't just value objective feedback, they prioritize it. CISOs must go beyond annual pen tests that only provide a snapshot of their security posture. They must invest in continuous testing that incentivizes expert feedback. A big part of this is fostering a culture where learning more about your attack surface, discovering the unknown, or being "beaten" by a red team is not seen as failure but as opportunity.	→ Leverage hackers, pentesters, and red teamers for offensive security testing to get a true understanding of where you're vulnerable. 

ARTICLE

The Vulnerability Intelligence Report

Every day, hackers in the Bugcrowd community submit hundreds of vulnerabilities via our Platform. These vulnerabilities range in criticality, target type, and submission category.

We analyzed hundreds of thousands of proprietary data points and vulnerabilities collected from across thousands of public and private engagements from January 1, 2024, to December 31, 2024.

Our goal is to provide security teams with the most up-to-date information on vulnerability trends to help them make educated decisions about their own risk and threat profiles.

Number of vulnerabilities

This graph shows the number of vulnerabilities over the past three years.

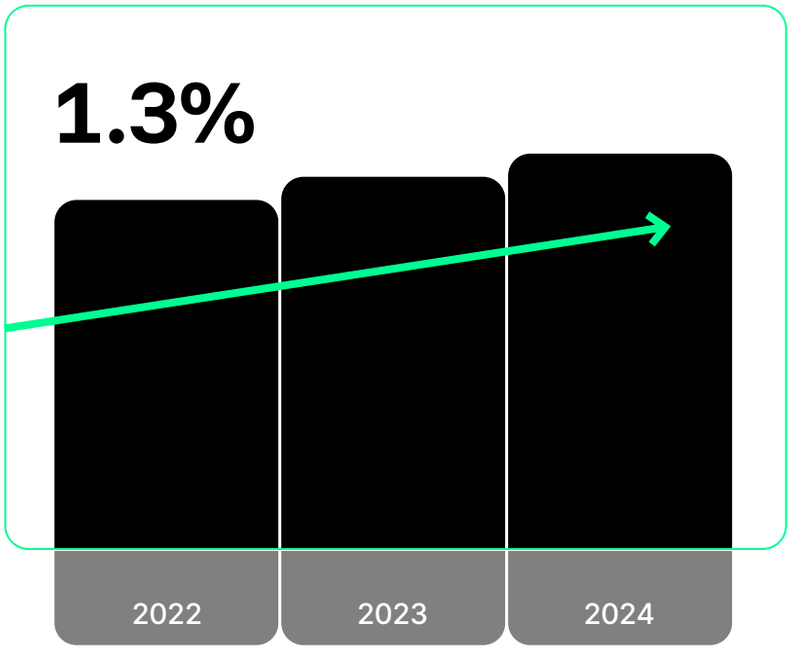

TRENDS

Over the past three years, the number of vulnerabilities found has stayed relatively consistent.

WHY?

The number of vulnerabilities is balanced by long-time engagements from more security mature organizations (which have a lower volume of vulnerabilities) and newer engagements that often have a higher volume of vulnerabilities.

Number of critical vulnerabilities

This graph shows the number of critical vulnerabilities over the past three years.

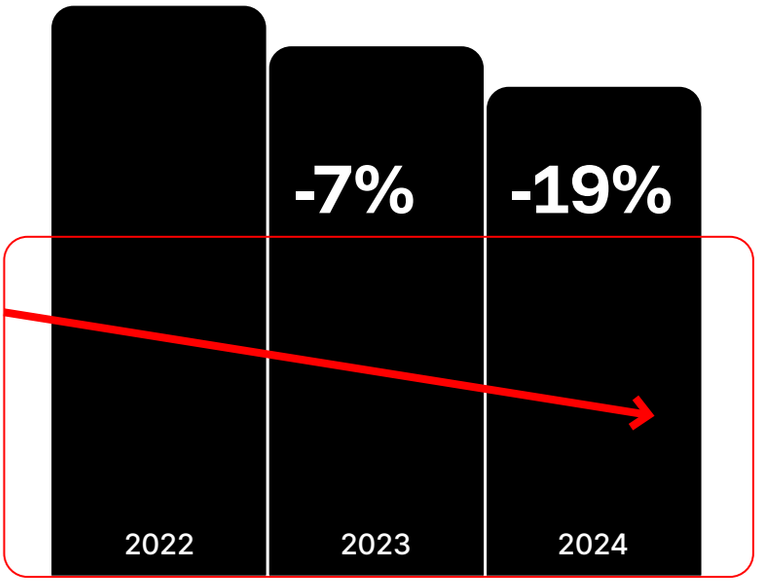

TRENDS

The number of critical vulnerabilities has gone down slightly year over year.

WHY?

Many new customers find that their first year of their engagement yields a high number of P1s. Over time, the number of P1s decreases—which isn't a bad thing! This is simply a signal of a program becoming more secure.

ASK A CISO

One CISO working for a long-time Bugcrowd customer shared his thoughts on this situation

> **It's a win-win situation—either the Crowd finds something we didn't see, in which case we can fix it, or they don't find anything, which validates our efforts.**

Number of critical vulnerabilities by target type

This graph shows the number of critical vulnerabilities (P1) by target type over the past three years.

API

↓ 25%

Website

↓ 30%

TRENDS

The number of critical vulnerabilities in API targets decreased by about 25%. Critical vulnerabilities in website targets decreased by 30%. There was a slight increase in critical vulnerabilities for Android, hardware, IoS, and network targets.

WHY?

The decrease in critical vulnerabilities in API and website targets is an incredibly encouraging stat. This points to customer API infrastructure becoming more secure. It tells us that developers are doing a good job at fixing bugs and securing their code. Keeping in mind that website target numbers indicate overall trends, we know that it is becoming harder to find P1s, giving customers assurance that they are becoming more secure.

Number of vulnerabilities by target type

This graph shows the number of vulnerabilities by target type over the past three years.

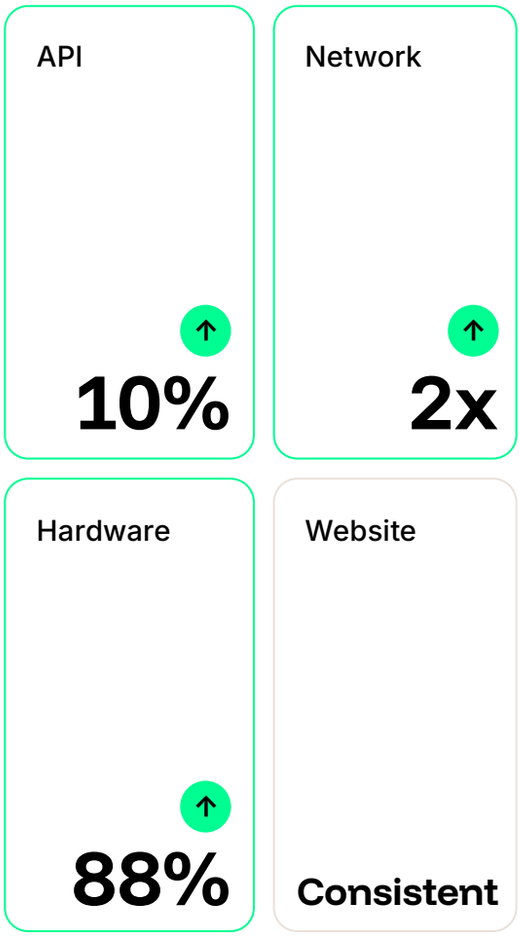

TRENDS

Over the past three years, API vulnerabilities increased by almost 10%. We also saw an 88% increase in hardware vulnerabilities. The number of network vulnerabilities doubled, and the number of website vulnerabilities stayed consistent.

WHY?

As a rule of thumb, it's helpful to use website target numbers as the ground truth when analyzing target data. This is because most engagements include websites in their scope.

Because website vulnerabilities stayed consistent, we can look at changes in other targets for additional insights. The increase in API and hardware vulnerabilities aligns with what we're seeing in the market—hardware is having a resurgence and API security is more important than ever. These numbers also tell us that organizations are diversifying their scope on their engagements. Many engagements start with a more limited scope, with website targets as the primary focus. As teams see the value in working with hackers, they will often expand their scope to include additional targets like IoT, network, and hardware.

ASK A CISO
API Testing

Dan Ford, CISO, ClassDojo

ClassDojo

APIs are the foundation of our platform, enabling key services and handling sensitive data. Because they directly expose business logic and functionality, they are a natural focus for attackers. Prioritizing API testing ensures that we are proactively protecting our systems and delivering secure experiences to our users.

APIs can expand the attack surface available to malicious actors, so securing them is critical. By combining internal testing with a comprehensive suite of unit tests, alongside live testing through our bug bounty program, we validate our defenses, catch subtle issues early, and maintain strong security as our platform evolves.

Partnering with Bugcrowd gives us access to diverse, skilled researchers who uncover vulnerabilities traditional testing might miss. Their insights strengthen our defenses, help us identify gaps, and continuously refine our internal security processes based on real-world attacker perspectives.

Payouts for vulnerabilities

This graph shows the average payouts for vulnerabilities over the past three years.

TRENDS

Over the past three years, the average, median, and 90th percentile of payouts have remained relatively consistent.

WHY?

Even in challenging times where budgets are being cut down and layoffs are common, security teams are maintaining their investments in crowdsourced security.

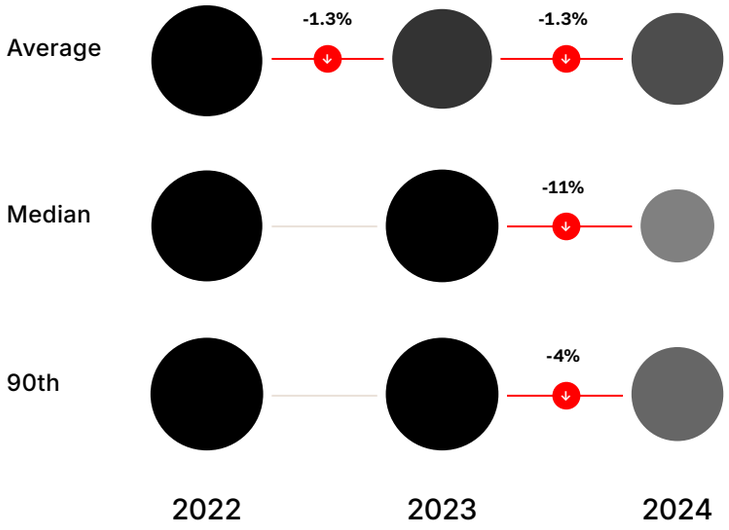

Payouts for critical vulnerabilities

This graph shows the average payouts for critical vulnerabilities (P1) over the past three years.

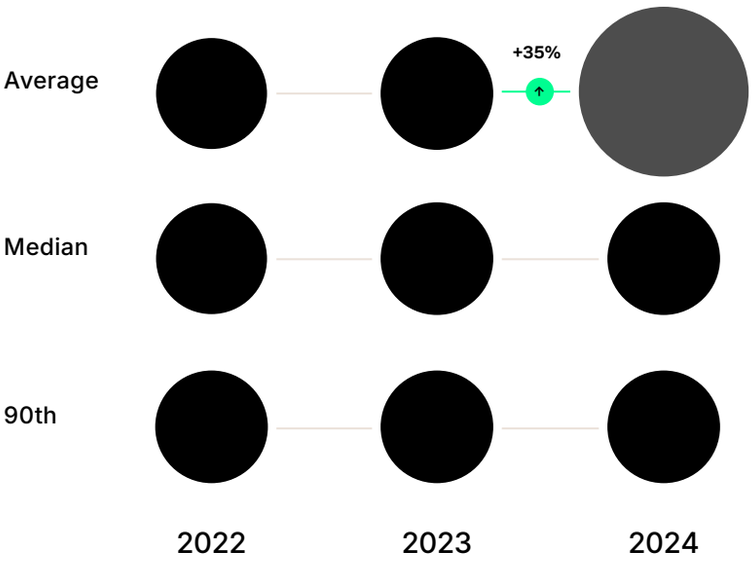

TRENDS

The average payouts for critical vulnerabilities increased by 32% in 2024. Median and 90th percentile critical vulnerability payouts remained the same.

WHY?

Notice how the graph above showed that overall vulnerability payouts remained relatively consistent while average critical vulnerability payouts increased each year? This points to organizations emphasizing critical vulnerability payouts. They are paying more for P1 vulnerabilities and balancing that by paying less for P3, P4, and P5 vulnerabilities.

Number of vulnerabilities by VRT category

This graph shows the top 10 most commonly reported VRT categories.

1	**Broken access control**
2	**Cross site scripting XSS**
3	**Server security misconfiguration**
4	Sensitive data exposure
5	Broken authentication and session management
6	Other
7	Server-side injection
8	Unvalidated redirects and forwards
9	Cross-site request forgery
10	Application-level denial of service (DoS)

TRENDS

The VRT category that saw the largest increase in vulnerabilities was broken access control (40% increase). Other categories that saw increases over a smaller volume of submissions include cryptographic weakness and network security misconfiguration.

The only VRT category that saw a statistically significant decrease in vulnerabilities submitted was application-level denial of service (DoS). The categories that saw decreases over a smaller volume of submissions include automotive security misconfiguration, client-side injection, insecure data storage, and mobile security misconfiguration.

WHY?

It was no surprise to see that broken access control vulnerabilities increased so much in 2024. This is a common category that many hackers gravitate toward. Many hackers prefer finding a niche set of skills and going all in on the VRT categories that align with that skill set, and broken access control vulnerabilities are certainly popular with hackers who employ this style.

We saw about an 11% decrease in application-level denial of service (DoS) vulnerability submissions. This VRT category is often out of scope in engagements, so hackers are less likely to test these applications fearing legal consequences.

ASK A HACKER
The increase in broken access control vulnerabilities

DK999

When looking at this data, it's important to remember that apps are getting more and more complex. Given the increase in features and integrations, access controls are becoming harder to manage. Most broken access control issues are trivial to exploit, yet they carry a huge impact.

Any app goes through multiple development cycles with numerous code changes, keeping the attack surface dynamic. Developers are under pressure to release features quickly, meaning security often takes a backseat. Proper access control implementation can be time-consuming. With AI-assisted coding becoming common, we can expect the percentage of broken access control vulnerabilities to increase.

Between app complexity, rapid development, and the new AI adoption cycle, security is being neglected early on. I believe this is why we're seeing this vulnerability type surge.

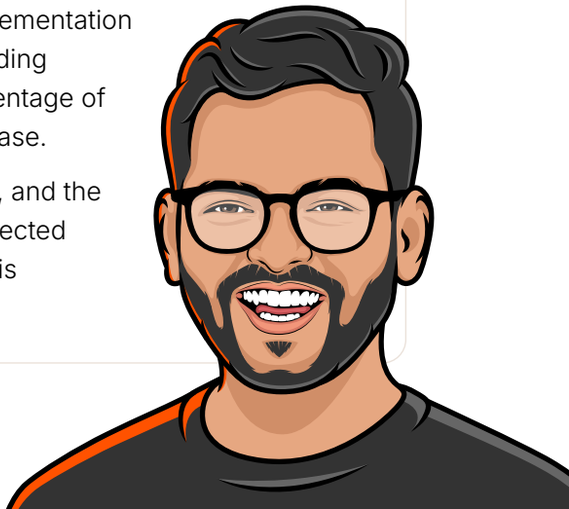

Number of critical vulnerabilities by VRT category

This graph shows the top 5 most commonly reported VRT category for P1s.

Sensitive data exposure P1s increased by 42%.

1. **Server security misconfiguration**
2. **Server-side injection**
3. **Broken access control**
4. Sensitive data exposure
5. Broken authentication and session management

TRENDS

The number of broken access control P1s increased by 36% in 2024. This category joined the top three in 2024. The top three categories of P1s rewarded in 2023 were broken authentication and session management, sensitive data exposure, and server-side injection.

WHY?

The increase in sensitive data exposure vulnerabilities is a key finding because it tells us that more personal identifiable information (PII) is being exposed to the world. PII includes items like names, addresses, emails, and social security numbers.

Unfortunately, in the process of code development, data exposure is still an afterthought. Given the number of P1s we're seeing in this category, we can assume that the PII that is being exposed is unencrypted. In the wrong hands, this type of data can lead to catastrophic consequences for customers and an organization's reputation.

Luckily, many hackers specialize in reconnaissance work in this specific category. There are thousands of hackers on the Bugcrowd Platform who consider sensitive data exposure vulnerabilities to be their bread and butter. They help organizations find these P1 vulnerabilities before threat actors do.

While many vulnerabilities are highly technical, complex, and mind-blowing, sensitive data exposure is much more, well, boring. But they can be cringe-inducing "how did I not know about this" issues with serious consequences.

From leaked credit card numbers to leaked employee tokens in GitHub repos, regulatory compliance compels organizations to take steps to secure sensitive data so that a data breach doesn't end up happening.

Attackers can end up living rent-free inside your systems, as data might be stolen for years before a breach is noticed. While you think everything is fine, attackers might be having a field day stealing intellectual property and PII on your customers or employees! And they don't just stop at breaching your systems; they'll extort you, dumping the data publicly when you refuse to pay.

ASK A HACKER
The increase in sensitive data exposure vulnerabilities

InsiderPhD

Ask a Hacker

Vulnerabilities to Watch out for...?

We asked expert hackers on the Bugcrowd Platform to break down the top five most commonly reported VRT categories for critical (P1) vulnerabilities last year.

These insights can help you understand the impact of some of the most common vulnerability types.

The potential impact of server security misconfiguration vulnerabilities

Server security misconfigurations remain one of the most common and dangerous weaknesses in modern environments. Misconfigured authentication, caching, or access controls can turn low-severity issues into critical breaches. Through my own work, I've located admin panels left exposed via default credentials, granting unrestricted system access. I've also uncovered a low-level rate limiting flaw on cached URLs containing sensitive documents protected by OTP codes. By combining predictable caching behavior with the rate limiting weakness, I was able to bypass the OTP requirement entirely and escalate the issue to a critical vulnerability. These cases show how small oversights can create significant risk when chained together.

Server security misconfiguration

Masonhck357

1

Vulnerabilities to Watch out for...?

Broken access control

2

DK999

The potential impact of broken access control vulnerabilities

Broken access control vulnerabilities should be a priority for every security team for three main reasons: ease of exploitation, likelihood of exploitation, and compliance requirements. Most of these vulnerabilities are easy to exploit, even for novice attackers. Threat actors are actively targeting these flaws to breach companies and leak data. Standards like the GDPR and HIPAA mandate strong access control. Failure to address these issues can result in significant fines and penalties.

From my experience, these vulnerabilities often leak critical information like PII, healthcare data, confidential system information, and internal documents. They are absolutely necessary to address.

Broken authentication and session management

Aituglo

3

The potential impact of broken authentication and session management vulnerabilities

Broken authentication and session management bugs are common vulnerabilities that often go unnoticed, with a very critical business impact. Even if you invest in a great firewall and EDR, keeping a completely clean dashboard, these bugs can run in the background, meaning someone can impersonate a legitimate user account without the security team getting any alerts.

This can be devastating from a business risk perspective. From a compliance and regulatory perspective, it can trigger GDPR or CCPA penalties because it commonly concerns sensitive customer data. There is also a downstream impact. Attackers can chain together these vulnerabilities, leading to more advanced attacks.

Vulnerabilities to Watch out for...?

The potential impact of sensitive data exposure vulnerabilities

The potential impact of sensitive data exposure can be a legal, financial, and reputational nightmare. Often, the sensitive data that is exposed—user names, addresses, IDs, and mobile numbers—are just part of an attack, and the attacker is pivoting their way deep into and across your network. By the time you find the breach, attackers have often already been working their way through your network for months. Meanwhile, they've already sold the data to the highest bidder, and now you and your organization are being targeted with phishing emails tweaked in just the right way to get you to engage.

These vulnerabilities are an obvious priority for security teams and CISOs. They must be identified and fixed quickly.

Sensitive data exposure

4

Brig

Server-side injection

Anon Hunter

5

The potential impact of server-side injection vulnerabilities

Whenever we type something into a website, a search box, a login form, or even a comment box, it sends that information to a server in the form of parameters and its value to process. Normally, a server should treat our input as plain, harmless text. But with server-side injections, attackers can send specially crafted text, called a payload, that the server will follow as if it were a legitimate command or instruction, considering my input as a part of its own coding.

This could lead to the attacker stealing all of your data, locking you out of your own system, demanding a ransom, or selling your stolen information on the dark web.

Hacking the NFL

Tomás Maldonado, National Football League (NFL)

Tomás Maldonado is a New York-based security leader and independent board director with over 25 years of experience across finance, media, manufacturing, and technology. He has been the CISO of the NFL for six years. As the largest and most popular sports league in North America, the NFL faces unique security challenges. The NFL is an organization of organizations—it is comprised of 32 clubs, each with unique operations, plus a league office, media properties, and global events. Tomás is in charge of **securing this entire ecosystem.**

We sat down with Tomás to learn more about his top priorities, his thoughts on AI governance, and his approach to proactive security.

What have been your top priorities since taking the helm as CISO of the NFL?

Our first priority has been to align security with the business's objectives and risk appetite. Cybersecurity cannot sit in isolation; it must support the NFL's mission and enable growth. We established a risk-based program mapped to standards like the NIST Cybersecurity Framework and made risk transparent to leadership so that they can make informed decisions about priorities and appetite.

The NFL isn't just one organization—it's an entire ecosystem. To secure it all, we built a unified framework that raises the baseline for every entity. Through consistent controls, shared playbooks, and regular assessments, we try to ensure no single point of weakness can impact the whole.

We have also invested heavily in culture and people. We don't see employees as the weakest link—we see them as potential security advocates. By equipping them with training and awareness, we've created an extended line of defense where everyone plays a part.

Finally, resilience has been central to our approach. We've strengthened threat detection, incident response, and data protection, but we haven't stopped there. We test these capabilities constantly through tabletop exercises and red team drills, ensuring that when the spotlight is on, security is seamless and the business can shine.

The issue of AI governance extends beyond tech into realms of compliance, operations, and brand reputation. How are you approaching and prioritizing AI governance?

AI governance can't live in a silo, so building an AI governance council that includes security, compliance, legal, and business leaders is necessary. Every AI use case should be reviewed for compliance, privacy, bias, and security concerns before it launches. We also monitor emerging regulations and translate those requirements into controls.

From an operational standpoint, I treat AI like any other critical system. This means securing data, testing models for manipulation, monitoring outputs for anomalies, and preparing incident response playbooks for AI-specific scenarios. I operate on a "security by design" principle, so innovation never outpaces safeguards.

The part I emphasize most is trust and brand integrity. We're entering an era where the line between real and fake is becoming increasingly blurred. Deepfakes and AI-generated content are a real risk to organizations. Companies investing in detection tools, validation processes for official communications, and crisis playbooks for AI-driven misinformation campaigns will be ahead of the curve. For me, AI governance is about protecting that fragile trust because once it's lost, it's incredibly difficult to win back.

In short, AI governance should be an extension of your security framework built on compliance, operational resilience, and brand protection; all of these elements must work in tandem.

How can CISOs effectively balance AI innovation and transformation with robust security and risk management?

For me, balance comes from embedding security from day one. Whenever a new AI initiative is proposed, my team runs risk assessments, applies guardrails, and ensures only the right data and systems are accessible. This way, we prevent later surprises.

But I don't view security as a roadblock. I often say, "If security is not enabling the business, then what are we doing?" Security should accelerate innovation, not stop it.

We celebrate when teams launch secure products and not just fast ones because this sets the tone that secure innovation is the standard.

Culturally, we work hard to make cybersecurity a partner to innovation. When business leaders understand why we're putting in guardrails, they become allies. Additionally, we highlight success stories where secure deployments allowed us to move faster or expand into new areas confidently.

Finally, we emphasize resilience. You can't block every threat; this is unrealistic. But you can prepare. We monitor AI systems, we scan for new vulnerabilities, and if something goes wrong, we respond quickly and learn from it. It's about embedding security into the DNA of innovation, so the organization can move forward safely and confidently.

How does proactive security and offensive security testing play a role in your overall security strategy?

Proactive testing is a cornerstone of our strategy. We don't believe in waiting for an incident to occur—we simulate attacks, run red team operations, and drill relentlessly. We do so many tabletop exercises that when a real incident happens, we have a plan. That preparation builds the confidence and speed we need when they matter most.

It's also about thinking like the adversary. I remind my team that unlike sports, cybersecurity has no rules—"We don't play fair with adversaries." This mindset drives us to simulate phishing, ransomware, and denial-of-service attacks against ourselves. If we can break our own defenses, we know where to shore them up.

For our key events, testing starts months in advance. We bring in partners to run scans, penetration tests, and tabletop drills. By event day, weaknesses we've found have been remediated, and security is invisible to customers and staff. The goal is to be boring from a cyber standpoint and exciting on the field.

Ultimately, proactive testing shapes what we do. It reinforces resilience because blocking every attack is impossible, but being prepared is.

It also helps validate our defenses, sharpen our responses, and keep our people vigilant. Offensive testing is how we stay one step ahead and ensure our defense is ready. ■

The CISO's Challenge: **Measuring Security Outcomes**

By Trey Ford

In Greek mythology, Sisyphus is punished to roll an immense boulder up a hill for the rest of eternity. As the boulder approaches the top, it immediately rolls back down.

From the ELT and board's perspective, CISOs can sometimes sound like Sisyphus when presenting our never-ending list of projects and asks. Every security program has a story full of milestones and gaps (from assessments, audits, best practices, customer requests, or some other source)—and there is always more to do and spend money on.

Budgetary constraints help us think critically and give us the opportunity to prioritize and innovate, though the roadmap and tradeoffs along the way are not always clear. When the board struggles to understand our vision, contextualize our risk investment strategy, or see how we measure success or failure, our boulder rolls back down the hill, requiring CISOs to start the process over again.

How we define "success" and "failure"

In reality, a security program can be a lot like our health and wellness journeys. Everyone is on their own path, and we are constantly having to navigate tradeoffs.

In my private life, I measure success in these areas by my ability to say yes to the things I care about—energy to say yes to family, capacity to be present and engage with friends, and ability to make time for sports and hobbies. Failure is when I don't have the energy to balance my work, travel, and the things that matter to me outside of work.

The difference between my personal goals and those of security programs is that the latter require that adversarial element to determine if we're executing at a level we're comfortable and confident in.

Furthermore, CISOs need to stretch a limited budget to balance people, process, and technology. The success of a program is measured in a handful of ways, but "an auditor approved" is the answer for so many.

However, can the lack of breach be considered a silent metric of success? (Reminder, we cannot prove a negative....) When we define success as a lack of incidents, justifying a constant increase in security spending to our boards is nearly impossible. In practicality, security without true adversarial testing is almost an illusion, leaning heavily on the "maturity" of best practices without pragmatic validation. This means that diversified research and testing clearly validates success, or identifies points of failure (opportunities for improvement),directly justifying our asks.

The culture we're building isn't about running from failure—it is aimed at continuous improvement and honest and objective feedback on what needs focus or prioritization.

Creating a safe environment for this level of objectivity is what changes our frame of reference from "failure" to a "growth mindset." This carries directly into program management and budgetary planning.

If the CISO community has learned anything through the zero-basis budget cycles over the last couple of years, it might be that the assumed nonnegotiable or brinksman position of "We need to be doing all of these things" doesn't easily stand up to scrutiny.

Adversarial testing: The path to objective measurement

NIST defines "resilience" as "the ability to maintain required capability in the face of adversity." So how do we measure this?

Adversarial testing evaluates our defenses by applying the tactics, techniques, and procedures of real-world attackers, highlighting deficiencies in our programs that rise above our agreed-upon risk profiles. Adversarial testers, like red teamers or ethical hackers, test resilience and provide actionable insights, highlighting high-priority gaps to address with a sense of purpose.

One way adversarial testing helps with objective measurement is it aids us in evaluating our technology investment stack. This area is notoriously difficult to be objective about—where are our people, process, and technology investments paying off or coming up short? We have a fear of asking how our technology investments are working, or even if they're working at all. Vendor evaluations are time-consuming, changes come with cost, and can be emotionally charged, so it's natural that there is an unwillingness to fire or rotate vendors/technologies. When leadership is confident in our objectivity in evaluating existing investments, we gain credibility.

When we engage in adversarial testing, we have the objective data to shine a light on our program to inform our decisions about what is and isn't working.

Why CISOs need adversarial testing to understand success and failure

Adversarial testing forces us to ask the hard questions and gives us an unparalleled view into the outcomes of our security spend. For most companies, this is almost like a Christmas card you send your customers and auditors—a once-a-year snapshot of your program. There's value, but moving beyond point-in-time assessments enables CISOs to confidently report program effectiveness.

By investing in adversarial testing, we quantify our security outcomes, identify gaps, and move beyond subjective assessments and maturity scores.

With the findings from adversarial testing, we can articulate and defend our asks to the risk committee and board, helping them make informed decisions about where we need to fund, where we need to defund, and what we need to adjust in the tech stack.

Bringing results to risk committees

From my perspective, the most successful, capable, and upwardly mobile CISOs operate in partnership with a risk committee. They regularly gather representatives from key leadership positions across an organization to sit down and evaluate the top risks to their business. These committees are an opportunity for businesses to look at their investments, assessments, audits, known technical deficiencies, and key concerns. In other words, CISOs use risk committees as an opportunity to align on difficult investment decisions associated with competing business risks.

In a time where zero-basis budgeting is becoming the norm, CISOs are constantly asked to defend every dollar and make difficult choices about what to cut.

Budget cuts affect every aspect of security planning, strategy, and operations—all of which are part of a complex tapestry woven across a business in alignment with the risk committee. Every time CISOs are asked to defund projects, they need fresh acceptance from the risk committee so that leadership can calibrate on the tradeoffs. CISOs can use the results of adversarial testing to justify these tradeoffs to the risk committee and make educated decisions to address risks and gaps.

A push toward resilience

When everything we ask for is "mission critical," we sound like Sisyphus, pushing our boulders up the hill over and over again. We must shift from incident prevention to measuring resilience. With the power of adversarial testing as a core component of our security programs, our asks are backed by evidence and we can tangibly demonstrate the value of our security investments.

Why does this resilience matter so much? Again, resilience is the ability to maintain required capability in the face of adversity. A strong security program means fewer disruptions to business, more effectively managed risk, and better processes to deal with incidents. We're building programs strong enough to protect what matters while letting teams focus on what they love outside of work.

Resilience isn't a destination but a series of daily choices and practices that become your way of operating.

When your security foundation is solid and continuously validated via adversarial testing, you're creating space for innovation, growth, and the kind of work–life balance that lets you say yes to what matters most. ■

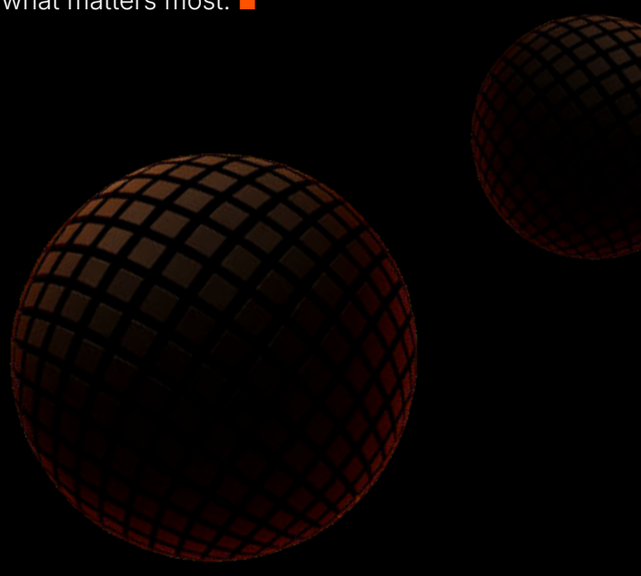

INFOGRAPHIC

The Dos and Don'ts of a Great Board Deck

Take your deck from 'pulse check' to a story the board will fund

Build upon previous decisions and discussions in a narrative approach.

Treat your board presentation as a status report instead of an ongoing story.

Build dashboards with meaningful, binary metrics that tell your story over time. Consistency is key!

Regularly change the metrics you report on or report on metrics that require deep security expertise.

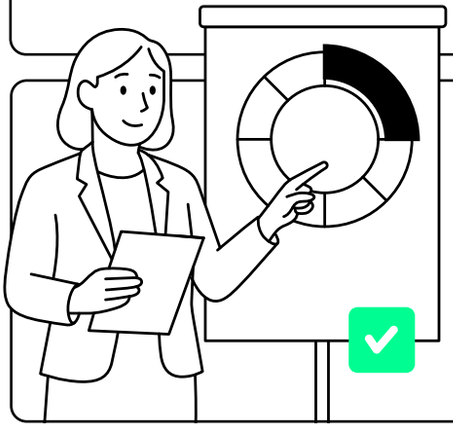

Include a maturity score based on a maturity model or risk management framework.

Only rely on maturity scores; combine with additional frameworks or metrics to show efficacy.

Use insights from adversarial testing to prove what is and isn't working.

Include asks with no justification.

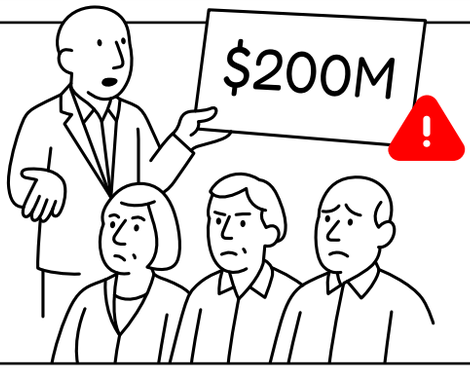

ARTICLE

Building a Board Deck

A Guide for CISOs

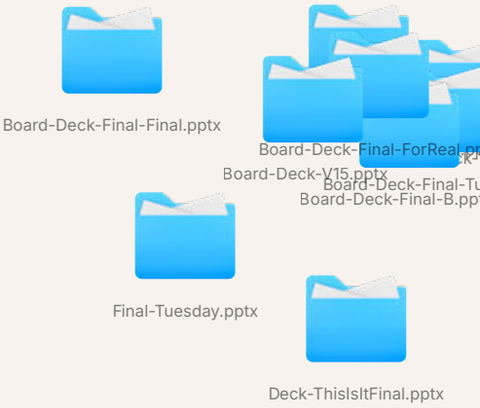

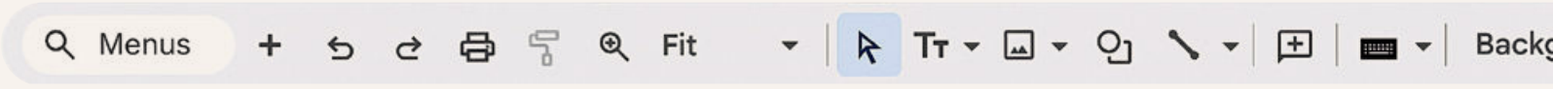

Boards routinely approve significant growth investments but freeze when CISOs ask for budget to fund security initiatives. This disconnect isn't a result of your presentation skills but the lack of context. Specifically, most board members lack the technical context to understand security risks (or tradeoffs) to evaluate your proposals against other initiatives, which makes it challenging to get their buy-in.

Your role as CISO is to bridge this gap by helping the board and executive team calibrate risk tolerance and make informed tradeoffs that align with organizational goals. This requires translating risk to help them understand what level of risk they're comfortable accepting.

Let's break it down further.

What is a board looking for?

The first step is to understand what a board is looking for. Every board is looking for clarity on these three questions:

What do I need to know

The board wants a high-level understanding of the current state of the security system and risks that keep you up at night. This also includes any critical data points or trends that you and the team are monitoring.

What do you need from me

The board wants to know what you need them to do to prevent risks from materializing, whether it's greenlighting a funding request or getting executive alignment on a strategic direction.

Why do I care

The board needs to understand why the risks and trends matter for the business, whether it's a threat to operations or a regulatory/compliance need.

ASK A CISO

Dan Maslin, CISO, Monash University

What advice do you have for CISOs looking to effectively communicate risk to their board?

Once you've passed the commonly used risk rating matrix of "likelihood vs. consequence," you need to bring the reality to life. Go to the next level and develop threat-informed scenarios that are most likely to occur within the organization to make it real for your audience. Next, consider the key mitigations or controls for each scenario and rate the effectiveness of each.

For example, you might say, "Advanced threat actors leverage social engineering to manipulate staff into providing unauthorized access," and your top three controls are "staff training," "privileged access management," and "email protection," with an effectiveness rating for each.

Having a few of these scenarios—which are realistic because they are based on current intelligence—will bring risk to life and drive a good human-to-human discussion about what could happen and how risk can be mitigated.

Craft a narrative

Board members understand business stories better than security metrics. They want to see progression, learn from challenges, and understand how decisions play out over time. This is why the most effective CISO presentations are built around story arcs. Here's a rundown of how to begin crafting your narrative.

1. Make each meeting a new chapter

Think of your board presentation as part of an ongoing story, not a status report. You want to build on previous decisions and show how they're being addressed to create an ongoing story about the state of the organization's security. This requires you to translate technical risks into a compelling business narrative that helps the board understand why the risks matter for the business, which builds mutual understanding and trust.

For example, you could start an audit storyline with,

> **We've got the audit coming up next month, and we've expanded our scope. We'll likely see new action items because we've never thoroughly audited this.**

Then, in the next quarter, continue the arc:

> **Here's what was accomplished. Here's what we learned. Here's what it means for the business. This is what we're going to do about it.**

Even if you do your best to communicate your story, you might get conflicting feedback from the board. Don't panic—it's normal. Take what's valuable from the feedback and keep moving forward.

TIP Don't have all the data yet? Put red Xs in your presentation where those metrics would go. This builds transparency and trust—don't hide what you don't know. You can use this as an opportunity to ask for budget and resources to track them going forward.

2. Dashboards: A picture is worth a thousand words

To help your board buy your narrative, use dashboards to support your story with metrics. Focus on showing trend lines that demonstrate what's working, improving, or failing over time. It's best to use the same dashboard structure each quarter so that the board can quickly understand the data.

ASK A CISO

Tomás Maldonado, CISO, NFL

What advice do you have for CISOs looking to effectively communicate risk to their board?

When I engage with the board, my priority is transparency and reducing complexity in my messaging. I raise issues candidly, explain why they matter, and secure support for the solutions we need. This means being willing to deliver difficult news but with solutions and a go-forward approach. As CISOs, we can't sugarcoat; our responsibility is to escalate risks so leadership understands what's at stake.

Treat security as a business function; don't talk in terms of firewalls or CVEs. Instead, talk about how a risk could impact operations, revenue, or brand reputation. The reality is that no board member wants to see the organization in headlines for the wrong reasons, and they understand that cybersecurity protects both the business and the brand.

One of the most effective ways I get my message across is through storytelling. I'm a firm believer of *"never let a good incident go to waste."* When a high-profile incident hits the news, we map it back to our own business. This makes the risk tangible. For example, I have shown how a ransomware attack could disrupt operations and erode customer trust.

Finally, it's important to make board engagement routine. Don't just show up in a crisis; provide consistent updates on threat trends, resilience, and preparedness. This cadence builds trust and positions cybersecurity as a standing business priority, not a one-off conversation.

So, my advice is this:

- ✔ Be transparent, even with bad news—transparency builds trust.
- ✔ Speak the board's language—frame risks in language directors care about.
- ✔ Use real-world examples—make the risks relatable.
- ✔ Keep communication regular—there should be no surprises.

When you do this, the board begins to see cybersecurity as integral to the business, and they're far more willing to support the investments you need.

3. Choosing the right metrics

To build the best dashboard, you need the right metrics to tell your story. Focus on binary metrics: simple yes/no answers to questions like "do you have this coverage?" These work well because cyber insurance underwriters have learned they correlate to actual breach payouts. This can include:

- **End-of-life timelines and upgrade plans for software**
- **Coverage metrics (e.g., logging, EDR, system inventory completeness)**
- **SLA adherence by risk level (not total vulnerability count)**
- **Security baseline compliance**
- **Hygiene indicators (e.g., patch compliance rates, incident response training frequency, backup/recovery testing results)**

Just like taking your vitamins the day before going to the doctor doesn't improve your health, quick fixes just before the board meeting don't show real security health. Instead, focus on consistent trends over time.

4. Come prepared with options

Once you have your narrative, present options for your top risks to help the board understand how they can help. Highlight the cost, timeline, and resources for each priority to make sure the proposals are clear.

The finishing touches

You've built the narrative. Here are tips to ensure it lands effectively.

- **Know your fundamentals:** Make sure you have an in-depth understanding of your attack surface, data locations, and SLAs.
- **Align with your executive team:** Get consensus from leadership on your risk priorities and recommendations before your board meeting to present a united front.
- **Calibrate on the board's technical literacy:** Use this knowledge to decide the right context level for each topic.
- **Present with conviction:** State your confidence and conviction levels honestly

When it comes to boards, credibility is everything. If you're not believable, you're not safe.

The best way to build credibility is to create a clear, compelling narrative that your board can understand, changing them from security skeptics into advocates. ■

SUCCESS FORMULA

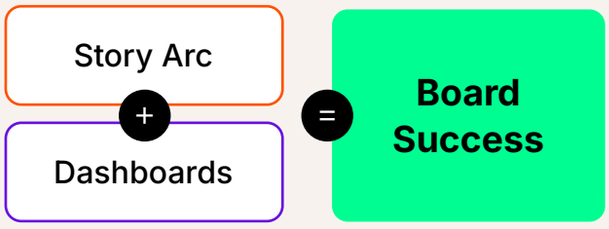

FROM SIMULATION TO STRENGTH

A CISO's Guide to Red Teaming

BY ALISTAIR G Director of Red Team Operations, Bugcrowd

I'm often struck by the parallels between maintaining personal health and an organization's cyber defenses. Regular checkups, stress tests, and immunizations help uncover hidden health issues before they become life-threatening—and in cybersecurity, red teaming plays a similar preventative role.

A red team exercise is a full-scope, real-world attack simulation that acts as the "diagnostic stress test" of an organization's security immune system. Conducted by ethical hackers, it probes a company's defenses (technology, people, and processes) in a controlled but adversarial manner. The goal isn't mere compliance or checklist completion; it's to proactively expose weaknesses, from unpatched systems to human errors, before a real attacker does. For a CISO, red teaming provides an unvarnished view of how their organization stands up to modern threats and where strategic reinforcements are needed.

I am new and I need budget. Can you show us our security holes?

How good or bad are our defenses?

Does my security strategy reduce risk?

Is my organization ready and able to respond to an attack?

How would a real threat target our company?

How secure is this company we have just acquired?

Ask anything

The role of red teaming in cybersecurity strategy

From a CISO's perspective, red teaming is not an isolated technical drill—it is a strategic tool that validates and strengthens an organization's security posture. CISOs often employ red team exercises to see how their enterprise detection and response mechanisms hold up under a simulated crisis. Red teaming serves several critical functions in a mature security program:

Simulating real-world attacks to test defenses

A red team can improve security resilience by simulating the TTPs used by threat actors that organizations would realistically face. This "live fire drill" often uncovers hidden vulnerabilities or attack paths that routine scans or compliance audits miss.

Identifying and prioritizing risks for reduction

Red teaming helps translate technical findings into business risk terms. Demonstrating the practical impact of certain vulnerabilities or process failures enables security leaders to prioritize what matters most.

Challenging assumptions and finding weak links

CISOs often have assumptions about what their security controls and staff can handle. Red teaming validates if existing security controls, policies, and procedures work as expected when under attack.

Strengthening security programs proactively

Overall, red teaming embodies a shift from reactive security (waiting for incidents to occur) to proactive security. By uncovering weaknesses and prompting fixes, red teaming drives continuous improvement.

Validating detection and response (blue team effectiveness)

Red teaming demonstrates how well blue teams can detect and respond to stealthy and evasive attacks. A well-run red team engagement will produce concrete data on detection gaps, and a good internal control group can measure response times, which the CISO can use to drive improvements.

In many sectors, the value of red teaming has become so recognized that it's mandated or strongly encouraged by regulators and industry standards. This regulatory push underscores a key point: from a boardroom's perspective, red teaming is not just about finding holes—it's about assuring stakeholders (regulators, customers, and the board) that an institution's defenses work against high-end threats.

Common defensive controls and red team evasion techniques

Across all these industries, organizations deploy a range of defensive controls to protect their assets. A CISO's mandate is to build a layered defense (people, process, and technology) such that if one layer fails, another will catch an attacker. I like to call this "the defensive onion" because the more layers an attacker cuts through, the more likely they are to cry. However, one lesson red teaming continually reinforces is that adversaries are adept at finding ways around even well-crafted controls. Understanding this cat-and-mouse dynamic is crucial for security leaders—it reveals which controls are truly resilient and which ones may provide a false sense of security if not complemented by others.

Identity and access controls (passwords, MFA, and SSO)

One common evasion tactic is socially engineering users to unknowingly assist attackers. For example, the use of MFA fatigue attacks has been widespread: an attacker uses a stolen password and keeps spamming a user's authenticator app with login approvals, hoping the user will eventually tap "allow" out of annoyance or confusion. Even a 1% success rate can be enough, but typically, I have seen successful exploitation between 10% and 30% of the time.

Data protection and monitoring

Many firms encrypt data on disk and rely on access controls, assuming that even if an attacker gets in, they can't easily access the most sensitive data without keys. Red teams sometimes reveal that encryption wasn't covering everything.

Email and endpoint hygiene vs. phishing

Red teams routinely craft convincing phishing emails, texts, and voice calls. They might register lookalike domains or exploit trusted services like calendar invites or Dropbox links. Even with increased user education, all it takes is one clever email at the right time to get a click.

Endpoint security (antivirus, EDR, and XDR)

Red teams employ custom tooling and obfuscation so that malicious code does not match any known signatures and looks benign or unique to slip through the cracks of EDR agents. With enough skill, endpoint agents can be undermined, highlighting to a CISO that no single control is infallible.

Network and perimeter defenses (firewalls, WAFs, and segmentation)

With the shift to cloud and remote work, traditional perimeters have become more porous. Red teams take advantage of this by attacking cloud services directly or by abusing VPN and remote access solutions, reminding CISOs that rigorous external attack surface management and patching are still crucial.

Beyond technical vulnerabilities: People and process

One of the most important insights a CISO gains from red teaming is that security is not just a technical problem—it's a human and organizational one. While vulnerability scanners and patch management address software flaws, red team exercises often reveal that the weakest links lie in human behavior and process deficiencies. A comprehensive red team doesn't just stop at hacking computers; it will probe the awareness and reactions of people, as well as the robustness of processes (incident response, change management, physical security procedures, etc.).

Red teaming goes beyond finding a misconfigured server or an open port and uncovers systemic issues such as employees being phished, IT support or helpdesk processes being tricked, or incident response playbooks failing under pressure. Many red teams find that they can gather a lot of information just by calling various departments and asking innocuous questions (pretexting as an auditor, new employee, etc.), a tactic known as elicitation. This might reveal internal lingo, names of key staff, or even details about what software or security measures are in place—all useful intel for further attacks.

Red teaming also shines a light on process failures and organizational silos. In a red team debrief, the timeline of "here's when we did X, here's when/if it was noticed, and this is how the staff responded" is incredibly valuable. It might show that the on-call process on weekends is unclear, the SOC is too understaffed to investigate every alert, or the SOC did respond but the communication to the broader team failed. These are systemic issues in incident response and crisis management that a red team helps identify without the cost of a real incident.

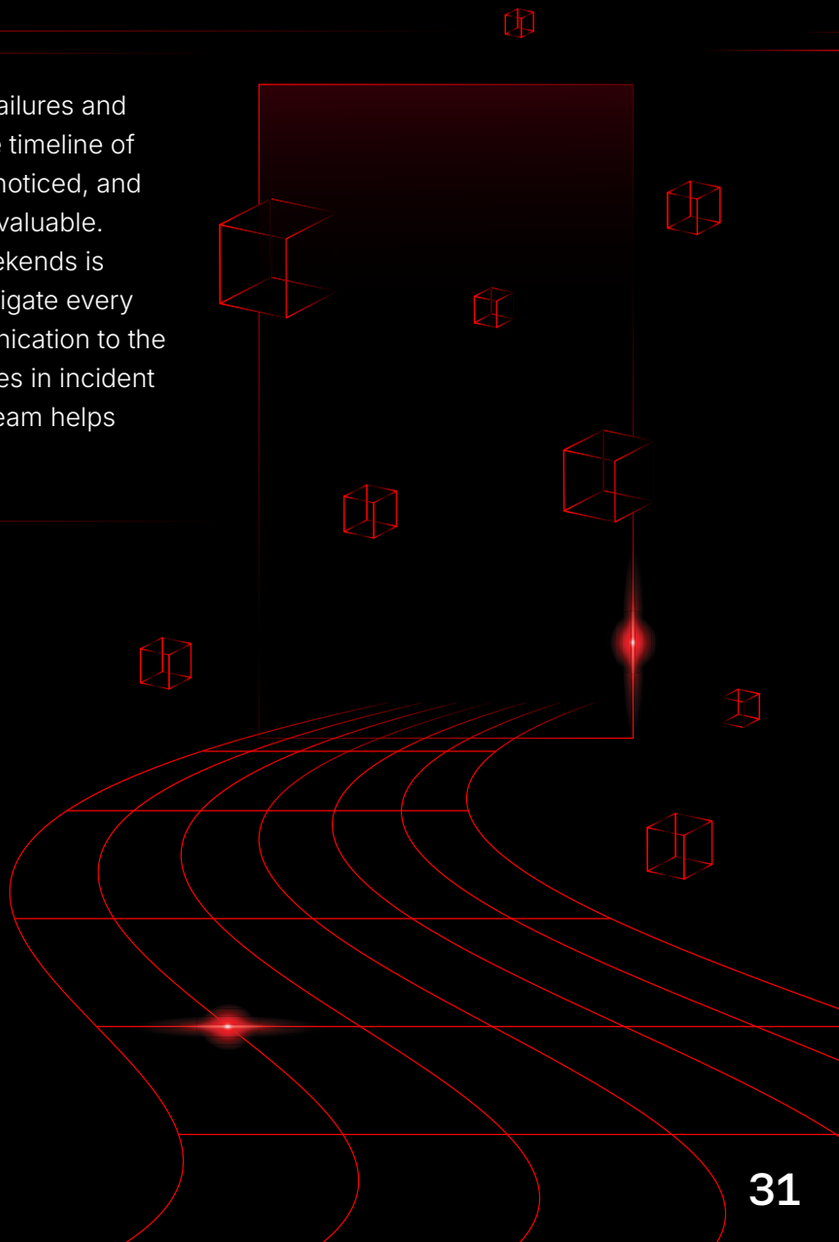

Red teaming outcomes often highlight the need for organizational learning and adaptability. The most mature organizations foster a culture where being "beaten" by the red team is not a failure but an opportunity to improve, akin to how regular exercise breaks down muscle fibers only to allow them to rebuild stronger. To go with the health metaphor, small, controlled doses of stress (red team drills) build the resilience "muscle" of an organization.

Leveraging red team outcomes for resilience and executive decision-making

A red team engagement is only as valuable as what an organization does with its results. For a CISO, the true deliverable of red teaming is not the successful "attack" itself but the actionable insights that emerge to strengthen security strategy, justify investments, and inform stakeholders.

Let's look at four areas where CISOs can benefit from the immediate impacts of red teaming:

Budgeting and investment

One of the most immediate impacts of a red team report is on budgeting and project prioritization. It provides concrete evidence of where an organization is exposed, often in a storytelling format ("We were able to steal the CEO's credentials and access sensitive M&A data because control X failed"). This can be incredibly persuasive when making the case for investments. Red team findings can also affect the strategic direction of security programs. For instance, if time and again red teams show that phishing is the entry point, a CISO might decide to shift budget into more user-focused controls like advanced phishing training, new email filtering solutions, or perhaps moving more apps to SSO with phishing-resistant MFA. Thus, a red team acts as a feedback mechanism for whether previous investments are yielding results or if new ones are required.

Board and executive reporting

Boards of directors today are acutely aware of cyber risk. Many ask management, "How do we know we're secure? Have we tested ourselves?" A red team exercise provides a narrative that a CISO can bring to their board to answer these questions credibly. This storytelling is powerful; it avoids jargon and instead uses a plot ("The attacker tried this, then this, we caught them here, but only after they had done that"). It gives the board a clear picture of risk in context, not just theoretically. Crucially, it also highlights improvements, which shows progress and accountability. Another board-level angle is using red team results to quantify potential impact reduction. Essentially, it's demonstrating cyber risk management in practice: find the problems, fix them, and reduce the likelihood or impact of a breach. Over time, repeated red team exercises can show a trend line, which can be translated into a risk reduction story for leadership

Driving SOC and blue team improvement

On a more operational level, red team findings are gold for the SOC and blue team. Every detection missed is an opportunity to create a new detection rule or refine an alert. Many SOCs will take the indicators of compromise (IoCs) from a red team activity (specific file hashes, command line strings, C2 domains, etc.) and retroactively check if their tools picked them up. If not, why? Perhaps the logs weren't there or thresholds were too high. They then improve those. Additionally, the exercise can be used to train a blue team in a "lessons learned" way. Some organizations even do replays or purple team sessions after the main covert red team is done. In effect, red teaming provides a continuous training loop for the defense team under realistic conditions.

Strategy

Regular red teaming fosters strategic cyber resilience. Resilience isn't just about preventing attacks; it's about ensuring that an organization can continue to operate and quickly recover even if an attack succeeds. Red team findings inform not just how to prevent breaches but how to limit damage and rebound from them. By incorporating red team scenarios into broader risk scenarios, leadership can develop a more robust risk management strategy. Another significant advantage is tracking improvement over time. A single red team exercise gives a snapshot; doing them regularly gives a trend. A CISO can set targets like "By next year's red team exercise, we aim to detect them at least at the data exfiltration stage, not after they have simulated customer data theft like this year." Achieving this goal would indicate improved resilience.

Conclusion

In the complex, ever-shifting cybersecurity landscape, CISO constantly ask: "Are we as prepared as we think we are?" Red teaming provides a profound and practical way to answer that question. Through the lens of simulated adversaries, it reveals the truth about an organization's defenses, the robust parts as well as the weak points, in a way no theoretical analysis can.
A CISO can leverage red teaming to test assumptions, sharpen detection and response, and ultimately drive down risk in alignment with real-world threats. These insights galvanize holistic fixes: better training, clearer processes, and more resilient architectures.

Red team outcomes give tangible metrics and stories that drive home the value of security initiatives. They help answer the tough questions from CEOs and boards like "How do we know our security investments are working?" by demonstrating improved detection times, fewer successful attack paths, and tested response procedures. In budgeting discussions, instead of relying on fear, uncertainty, and doubt, CISOs can point to red team exercises to say, "This is where we were, this is where we are now, and here's where we need to get to next."

For a CISO, red teaming is an indispensable tool for achieving and demonstrating cybersecurity excellence. With the insights gained from red teaming, and the resulting enhancements in strategy, controls, and culture, security leaders can sleep a bit more soundly at night, and assure their stakeholders that the organization's digital health is continuously monitored and improving. In the ongoing battle against cyber threats, red teaming ensures we are fighting fit and ready for whatever comes our way. ■

Securing a Leading AI Supercomputer

Dan Maslin, Monash University

Dan Maslin is an experienced technology executive based in Australia. For the past six years he has worked at Monash University, Australia's largest university with around 90,000 students and 20,000 staff, where he is Group Chief Information Security Officer and Head of Infrastructure Strategy.

In 2025, Monash University announced its investment in building and operating an advanced AI supercomputer to transform AI-driven research. This supercomputer is the first of its kind in Australia to utilize the NVIDIA GB200 NVL72 platform and is expected to deliver **unprecedented AI capability for research in areas from cancer detection to climate action.**

We sat down with Dan to learn more about this amazing project, AI governance, and his approach to proactive security.

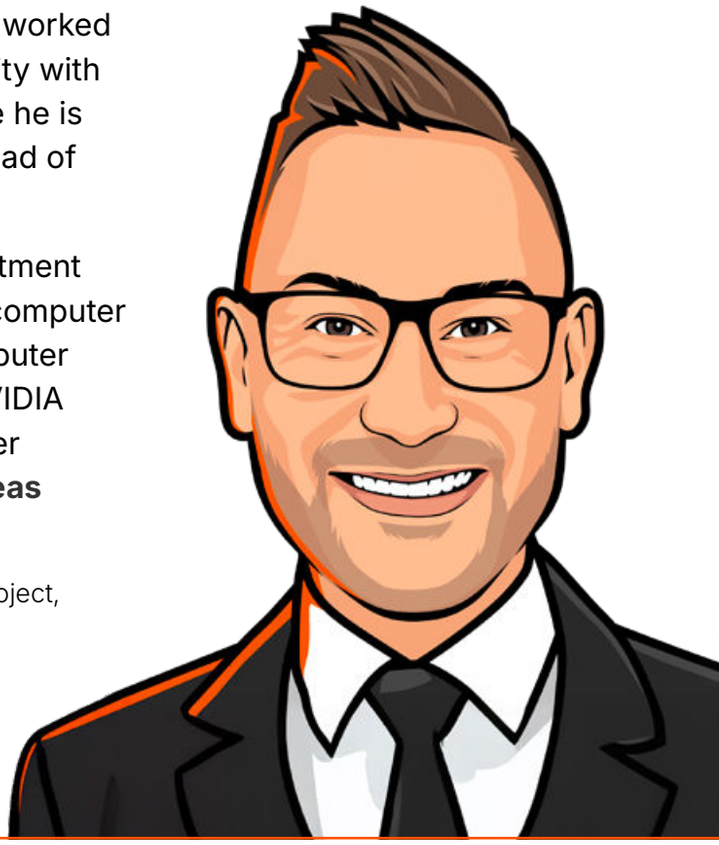

MONASH University

Can you tell us about how you're approaching security for this new AI supercomputer?

There are so many layers to this! To start, fortunately for me, the organization has a positive security culture and typically considers cyber, privacy, and sovereignty early on in projects. As CISO, I was brought into the project very early—more than 6 months before anything became public—and was on the evaluation panel for all parts of the project.

I needed to be comfortable on everything from the data center where we'd host it through to the supplier of the hardware. We landed on an arrangement with CDC as a data center and NVIDIA and Dell hardware.

I was able to query security considerations for every aspect—from physical security at the place of hosting to software and hardware supply chain assurance, the vetting of staff, and all parties' approach to vulnerability disclosure and inclusion in bug bounty programs.

Yes, that was a question they needed to respond to!

The issue of AI governance extends beyond tech into realms of compliance, operations, and brand reputation.

How are you approaching and prioritizing AI governance?

For Monash, AI governance runs even deeper. Aside from the usual corporate environment considerations around AI in operations, we also have to consider the impacts of AI on both research and education, both of which are likely to be heavily impacted in the coming years. In early 2024, Monash established an Artificial Intelligence Steering Committee, with more than a dozen members representing every corner of the university. Reporting directly to the Vice-Chancellor (the equivalent of the CEO in a corporation), the Committee exists to create a clear understanding of the risks and strategic benefits of using AI for education, research, and operations, both in the short and long term, and it oversees and informs decision-making on the use of AI across the Monash Group into the future.

Monash also has a publicly published AI Readiness Framework that is fairly comprehensive and considers the people, technology, and scaling aspects, and this is where governance is situated. It includes an organization-wide agreement on responsible use principles, internal policies, the risk management approach, and tracking of the evolving legal and regulatory landscape surrounding AI. So in short, AI governance is a product of organization-wide input, reporting into the most senior level of management.

How do proactive security and offensive security testing play a role in your overall security strategy?

Offensive security testing is absolutely at the core and one of the first principles we introduced when I joined five years ago. We can't scale to continuously proactively test our environment with our internal resources—we need a crowd.

We will never have the broad and expert skills internally to deeply test and provide effective assurance across everything, from mobile apps and building management systems to corporate IT and supercomputers; we need to leverage a variety of skills available within a crowd of ethical hackers to have confidence that we can know about a vulnerability first.

I've always said that we can't manage what we don't know about, so we're better off prioritizing the scalability and continuous visibility of our environment.

Can you highlight an initiative from your team over the past year that exemplifies excellence, innovation, and resilience?

Our team created and runs the Cyber Security Student Incubation Program, which was set up to do three things: build a reliable talent pipeline for the internal cyber security team, give students meaningful paid experience while they study, and help produce job-ready graduates who don't need to start from scratch in the industry. We recruit five students each year and give them part-time roles (usually 2–3 days a week for a year) paid at market rate and supported by structured training and mentoring.

This isn't unpaid work experience—they're treated as part of the team. We see it as win-win-win.

We win because we get access to new intelligent talent about to enter the market, the students win because they get real-life paid work experience for a year, and the industry wins because it gets a Monash graduate with a degree and a full year of hands-on real-life work experience. ■

Will AI Replace Security Research?

BY FRANCOIS GAUDREAULT aka P3t3r_R4bb1t

Hi. I'm Francois, also known as P3t3r_R4bb1t. I'm a cybersecurity leader with over 15 years of experience in information security, risk management, and ethical hacking.

I've served as the Senior Manager of Security and Enterprise Engineering at Wayfair, and I previously held key security roles at National Bank of Canada, Videotron, and GoSecure, where I led teams, managed multimillion-dollar budgets, and developed comprehensive security programs. As a top-ranked ethical hacker on Bugcrowd (#4 out of 100,000+ active hackers), I have identified over 1,700 valid vulnerabilities across public and private programs, including U.S. Federal Government systems, while also bringing my technical expertise and leadership skills to help organizations strengthen their cybersecurity posture through strategic risk management and offensive security initiatives.

Let's jump right into the topic of this article. AI agents and automated validators have gained traction recently in the hacking and cybersecurity space. Some self-proclaimed enterprise solutions are starting to leverage vulnerability disclosure programs (VDPs) or even private bug bounty programs to train and demonstrate full automation capabilities in AI agents.

Given my experience as both a hacker and a security leader, I'd like to share my thoughts on how AI will impact the hacking and security research space, as well as how CISOs should be approaching their offensive security testing in this new landscape.

All images in this article have been created with AI

Automation isn't new

The concept of leveraging scripts, workflows, and automation is not new in the bug bounty world. These approaches are likely as old as the concept of bug bounty itself. Of course, the landscape has evolved quite a bit over the last 5 years, now requiring less and less human interaction. Bugs are captured by the continuous scanning of assets and pushed to queues using webhooks. Findings are validated either manually or automatically and even pushed to platforms using prewritten and heavily templated reports.

So, what does AI automation actually bring to the table?

I would say it's simply the following:

- ✓ **Increased speed**
- ✓ **Larger asset coverage**
- ✓ **Drastically reduced complexity in tooling**
- ✓ **A basic level of thinking**

In its current state, I do not believe AI has the ability to provide additional depth (i.e., critical findings related directly to business-specific contexts) or the capability to efficiently circumvent proactive controls like a web application firewall (WAF) or bot detection technologies. For instance, how would an AI react if companies were to start implementing bot prevention at scale (or more simply, just denying traffic based on the AI traffic signature) to reduce the AI's reconnaissance capabilities? A human researcher can move around this limitation rather quickly.

AI is like a puppy

The other key reason why I believe AI will not replace human hunters in the short term, or perhaps even the longer term, is the need for AI to be trained. Currently, that training has to come from humans proficient in prompt engineering. Today, AI systems train on public data and complementary datasets. You don't know what you don't know, and the same applies to AI. In other words, an AI agent doesn't know what humans don't tell it. Thus, I strongly believe humans will continue to have an edge and maintain some control on that front.

Such training requirements may also trigger unwanted opacity in the future of vulnerability disclosure and research. Nobody wants their job to be replaced by AI. Therefore, in an AI-dominated world where companies fight for competitive advantage, will ethical security researchers continue to disclose their vulnerabilities publicly, or will they keep these techniques or findings to themselves for an extended amount of time? If we push this thinking slightly further, will researchers sell their research to AI companies instead? Similarly, will product manufacturers or companies disclose the vulnerabilities in their assets, or will they use incredibly vague statements (some businesses are already experts at this!) in their disclosures?

These are crucial questions to ask ourselves, and I myself am puzzled. On my end, I do see a potential case where an AI-dominated market may encourage additional secrecy, persuading bug bounty researchers or even AI companies to keep their edge in a highly competitive space.

It needs training

 All images in this article have been created with AI

Cost and architecture

Another interesting angle that could generate additional discussion and research is AI automation costs and architecture.

I discussed this before, but I personally tend to hunt manually. I limit myself to the bare-minimum tooling and automation. This strategy obviously can't scale to a larger scope and to multiple programs at the same time. This is an area where AI agents may drastically outpace researchers. But at what cost? And what do the architectures of these solutions look like?

While AI automation may revolutionize bug bounty research at scale, the economic reality reveals hidden costs that extend far beyond simple model usage fees. An AI system capable of meaningful vulnerability discovery across multiple programs requires sophisticated infrastructure orchestrating reconnaissance engines, specialized AI models, validation pipelines, evasion mechanisms, and continuous monitoring systems. Each component demands significant computational resources, storage capacity, and operational expertise to maintain effectiveness while avoiding detection by increasingly sophisticated bot-prevention systems. Architectural complexity grows exponentially when you take into account the need for distributed scanning, real-time data processing, model retraining, and compliance monitoring across diverse program requirements.

ITMOAC

Complementary rather than competitive

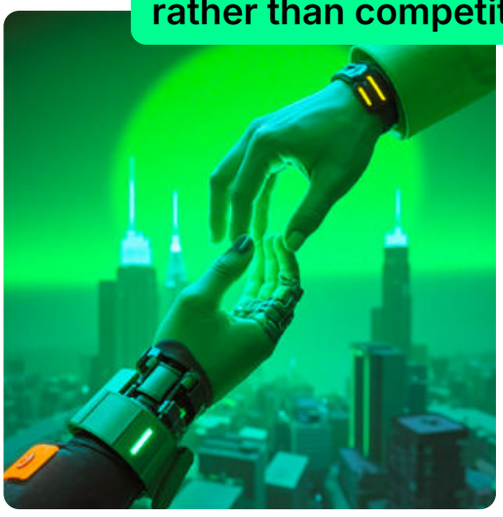

From a leading bug bounty researcher's point of view, AI-based automation should be able to drastically speed up bug hunting processes, help with reconnaissance on large scopes, highlight interesting aspects of a target, help pinpoint low-hanging fruit, and even submit issues to programs automatically. AI excels at processing vast amounts of data quickly, identifying patterns across extensive attack surfaces, and performing repetitive tasks that would consume significant human effort and time.

> **However, I personally see AI automation as far more relevant to enterprise attack surface monitoring solutions.**

These organizations have complex digital footprints that can benefit from AI systems that continuously scan, catalog, and assess their assets for potential vulnerabilities in real time.

All images in this article have been created with AI

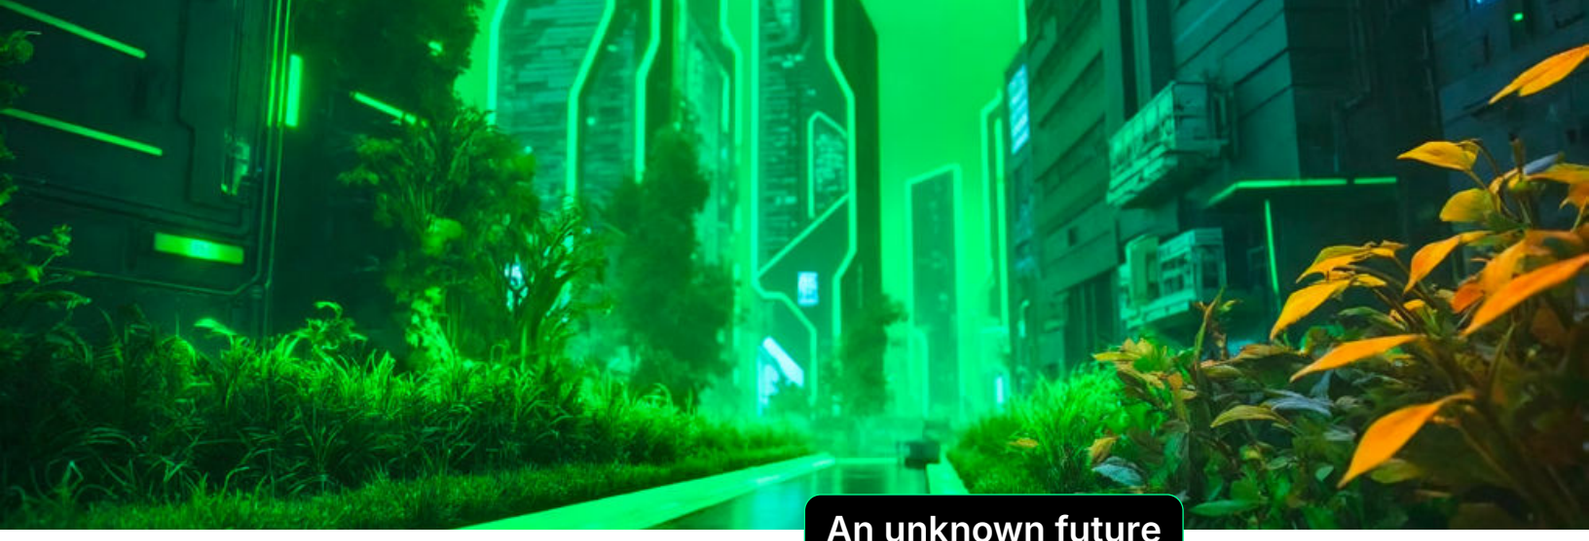

An unknown future

Using today's technology, I do not see this level of automation going as deep into a system as a human researcher would, and AI likely won't be able to find unique business-context vulnerabilities. Human researchers bring critical thinking, creativity, and contextual understanding that AI currently lacks. Researchers can identify logic flaws specific to business workflows, understand the nuanced implications of seemingly minor issues, and chain together multiple small vulnerabilities to yield significant security impacts. The most sophisticated vulnerabilities often require understanding not just technical implementation but also business logic, user behavior patterns, and organizational context.

> **Only human intuition and experience can provide this level of understanding.**

However, no one can really predict if a breakthrough will be made to significantly boost AI's capabilities. As AI becomes more and more sophisticated and capable of contextual reasoning, this gap might narrow.

With all that said, I remain confident that humans will continue to have a place of choice in the bug bounty (or even the cybersecurity) ecosystem, with the future likely showing a complementary relationship; AI will handle the breadth while humans will provide the depth and creative problem-solving that high-value, complex vulnerabilities demand. One thing is for sure: no one really knows how AI will effectively change the paradigm in the cybersecurity space. Only time will tell. ■

Why CISOs should care

CISOs should be concerned about the use of AI in cybersecurity primarily due to the significant increase in speed and efficiency it offers threat actors.

While the fundamental nature of cyber threats hasn't changed, AI's automation capabilities mean that vulnerabilities, especially low-hanging fruit on a perimeter, can be discovered and exploited far more rapidly than ever before. This acceleration could allow attackers to quickly pull in zero-day exploits through systematic testing. Although AI may struggle with complex business logic flaws or tricky injection attacks, its ability to quickly find and leverage simpler vulnerabilities still poses a substantial risk that security leaders cannot ignore. Ultimately, it's the unprecedented speed of both detection and exploitation that makes AI a critical concern for modern

From Assets to Action

Operationalizing Attack Surface Intelligence

Managing today's attack surface feels like a **never-ending game of whack-a-mole**—just as you get a handle on the current landscape, something changes, whether it's a new asset, attack vector, or vulnerability. As a result, security teams find themselves constantly reacting rather than staying ahead, which creates blind spots that attackers can exploit.

To proactively safeguard their assets, many organizations turn to external attack surface management (EASM) to improve visibility. However, these tools operate in isolation from offensive testing workflows, which usually have different logins and reporting structures. The result? Critical intelligence sits idly in the EASM tool, disconnected from remediation efforts.

To help CISOs truly reduce risk, security teams must integrate EASM intelligence into their offensive testing platforms so there's a direct path from discovery to remediation.

PRESS START

The disconnected state of security tooling

As organizations scale, their attack surfaces become increasingly complex to manage. Development teams are constantly deploying new infrastructure, like cloud services, APIs, and proprietary LLMs, creating a dynamic environment that's nearly impossible to track in real time. This is further exacerbated by the rise of third-party integrations and shadow IT, which expand attack surfaces unpredictably.

But visibility is just one part of the equation—**CISOs must also be able to prioritize assets based on business risk**, ensuring resources are focused where they matter most.

This means having accurate, up-to-date intelligence on each asset: exposure status, environment, criticality, and any validated vulnerabilities.

Most security teams try to fill this gap themselves, using EASMs with some combination of spreadsheets, open-source tools, and internal systems. Each solution has its own login, workflow, and data model, creating a patchwork approach that leads to stale data, duplicated effort, and inconsistent context across tools—slowing down remediation and increasing exposure risk.

The case for integrating EASMs with offensive testing

To bridge the gap between discovery and action, security teams should integrate their EASMs with offensive testing workflows. This creates an automated pipeline where newly discovered intelligence is immediately prioritized and validated through offensive testing methods like bug bounties, red team engagements, and pen testing. The result: teams respond to threats as quickly as they emerge and continue to stay one step ahead of attackers.

For example, when an EASM identifies a new subdomain with an exposed admin panel, teams can immediately scope targeted testing through an integrated platform to determine if it's exploitable and what data is at risk—fully leveraging their attack intelligence for swift remediation.

At a more strategic level, integrating these workflows fundamentally shifts security operations from reactive firefighting to intelligence-led decision-making. Security teams not only become more efficient in their daily tasks but can proactively and confidently prioritize vulnerabilities based on real-time, actionable insights—leading to smarter, faster, and more informed security strategies.

The bottom line

As the damage from cybercrimes increases rapidly, CISOs can't afford to wait weeks or months to act on their attack surface intelligence.

By adopting this integrated approach, CISOs demonstrate measurable improvements in security efficiency and faster remediation cycles, which enable them to prove the value and outcomes of a security program to external stakeholders.

Conclusion

If you're a CISO, think back to what some of your early jobs in security looked like. Chances are that the space is now unrecognizable. Perhaps you remember receiving your patches in a folder filled with floppy discs, and phrases like "artificial intelligence" felt like they belonged in The Matrix, not the office.

And now here we are. We're not just at the precipice of change in this new AI landscape; we've jumped. The question is, do you have a parachute that you can dependably deploy, allowing you to land safely?

At Bugcrowd, we're doing a lot with AI, but we don't believe it's the silver bullet that can solve every CISO problem. As a leader in the offensive security testing space, it's our responsibility to use critical judgment, embrace AI with caution, and most importantly, share our knowledge with the community.

In this edition of *Inside the Mind of a CISO*, we covered some of the biggest priorities and pain points for security leaders. As we wrap up, let's look at three ways Bugcrowd can help CISOs achieve greater security resilience.

Three ways Bugcrowd can help

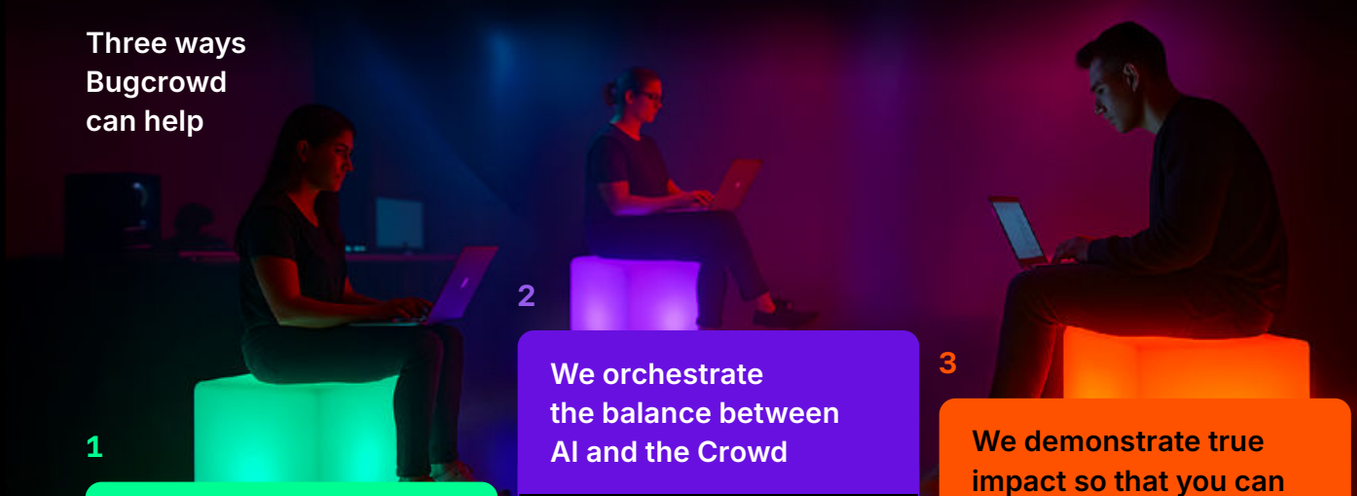

1

We give you the gift of objective feedback

You've likely heard the question, "What keeps you up at night as a CISO?" The answer is simple—it's the unknown. Ultimately, we all need a way to objectively measure security outcomes. If you're vulnerable, you would want to know about it.

By partnering with Bugcrowd, CISOs can lean on the expertise of a global hacking community to help them find and fix vulnerabilities faster.

The Crowd offers continuous testing from experts with a massive range of specialties and skill sets. When CISOs tap into the Crowd for their insights, they're not just accessing increased security resiliency; they're accessing peace of mind.

2

We orchestrate the balance between AI and the Crowd

CISOs shouldn't be expected to keep up with every nuance of where the Crowd ends and AI begins—at this point, the goal posts are moving too quickly. Bugcrowd is here to cut through the complexity. Using our security expertise, we make sensible decisions about where the adoption of AI models makes sense and where human ingenuity is still king.

For over a decade, Bugcrowd has helped organizations know the right levers to pull in their security programs at the right time to find and fix unknown vulnerabilities faster. AI is simply another powerful lever we pull for our customers, bringing the best outcomes possible.

3

We demonstrate true impact so that you can take informed actions

CISOs are in the business of putting out fires all day, every day. The noise is constant, and the "what ifs" never end.

Bugcrowd can provide clear visibility into your attack surface, simplifying prioritization so you know where to focus first. We also give you the ability to take action right in the Platform.

For those ready to take their security testing to the next level, they can kick off world-class red teaming engagements with Bugcrowd.

Red teaming measures the true impact of a potential breach. For CISOs, red teaming provides an unvarnished view of how their organization stands up to modern threats and where strategic reinforcements are needed. ■

INSIDE THE MIND OF A CISO

bugcrowd