A CISO's guide to **Red Teaming**

bugcrowd

Read. Learn. Defend.

Table of Contents

Introduction

Every Chief Information Security Officer (CISO) knows that maintaining an organization's cyber defenses is a constant battle, requiring regular proactive measures to stay ahead of threat actors.

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 process) in a controlled but adversarial manner. The goal isn't mere compliance or checklist completion; it's to **preemptively expose weaknesses,** from unpatched systems to human errors, *before* a real attacker does. For a CISO, red teaming provides an unvarnished view of how an organization stands up to modern threats and where strategic reinforcements are needed.

For a CISO, red teaming provides an unvarnished view of how an organization stands up to modern threats and where strategic reinforcements are needed.

Red teaming can be used to test assumptions and incident response, validate detection capabilities, and inform risk reduction initiatives. Across industries and geographies, different CISOs will have different objectives based on the unique threats they face. The defensive control stack they deploy and the strategies they employ are all influenced by the real adversaries that an organization faces and the tools, tactics, techniques, and procedures (TTPs) those adversaries use. Red team engagements simulate these threats and translate the technical offensive journey of an attacker into **actionable insights,** supporting investment decisions, informing board-level reporting, sharpening security operations center (SOC) performance, and ultimately bolstering an organization's strategic resilience.

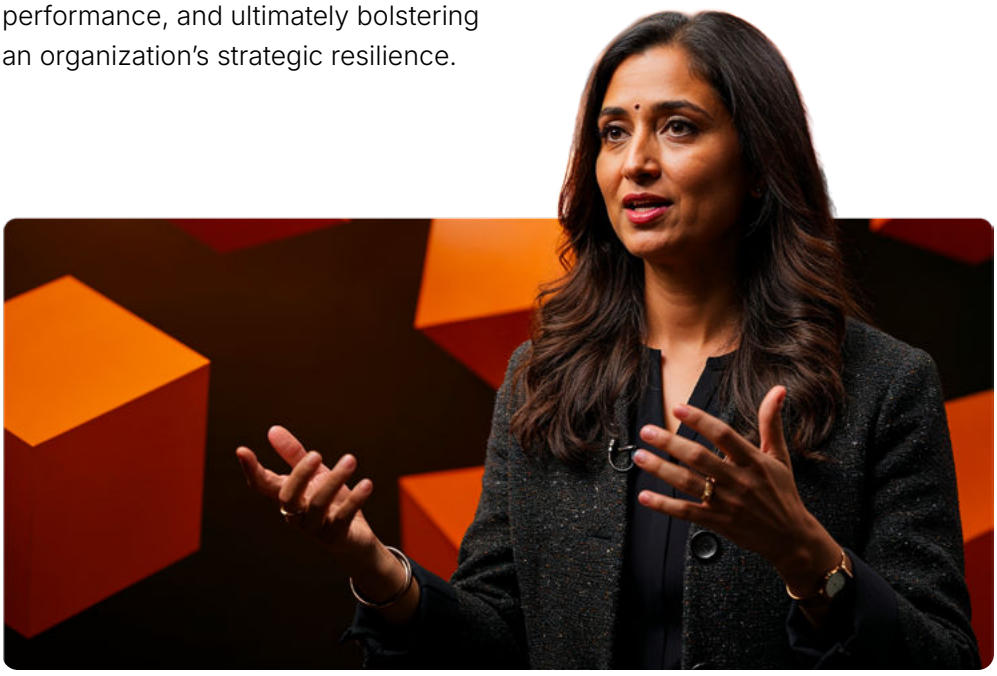

The role of red teaming in cybersecurity strategies

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. Therefore, red teaming serves **several critical functions** in a mature security program.

✔ Simulating real-world attacks to test defenses

A red team exercise gives organizations an opportunity to safely experience a full-spectrum cyberattack tailored to top threats facing their respective industries. By simulating and sometimes emulating the TTPs of real adversaries, red teaming provides the closest approximation of how prepared a company truly is for an actual incident. A red team will covertly attempt to achieve specific goals, such as access sensitive data and disrupt operations, without being detected, thereby stress-testing security controls under adversarial conditions. This **"live-fire" drill** often uncovers hidden vulnerabilities or attack paths that routine scans or compliance audits miss.

✔ Challenging assumptions and finding weak links

CISOs often have assumptions about what their security controls and staff can handle ("Our new email filter will catch all phishing" or "Our incident response playbook covers ransomware"). Red teaming provides a reality check by actively attempting to bypass controls and procedures. It forces an organization to confront the question: **"Are our defenses as effective as we think?"** For example, if multi-factor authentication (MFA) is assumed to stop all unauthorized access, a red team might test that assumption with tactics like MFA fatigue attacks or adversary-in-the-middle phishing to steal session tokens. Its success would reveal a gap in what was thought to be a strength. This process of assumption testing is invaluable; it illuminates systemic issues and blind spots in people, process, or technology that leadership might otherwise overlook.

Red teaming provides the closest approximation of how prepared a company truly is for an actual incident.

Assumption testing illuminates systemic issues and blind spots in people, processes, or technology.

✔ Validating detection and response (blue team effectiveness)

A core purpose of red teaming is to allow an organization to evaluate its ability to detect, respond to, and recover from sophisticated attacks (purple teams would actually evaluate PDR metrics systematically over a wide range of TTPs). Unlike a standard penetration test, which often stops at finding vulnerabilities and exploiting them, a red team will operate covertly and stealthily over days, weeks, or months. This allows an organization to see if the SOC or blue team notices the presence of the red team. This tests the **full incident response life cycle:** Was the attack detected? How long did it take (mean time to detect, MTTD)? Did the team appropriately contain and eradicate the threat (mean time to remediate, MTTR)? Were escalation and communication procedures followed?

A red team exercise effectively measures an organization's "immune response," allowing security leaders to analyze and assess the outcomes of their teams. (Again, this analysis can be done by the testers in a purple team engagement.) Success is not just in the red team achieving its objectives but also in what is learned about the defensive team's performance under pressure. Key metrics often emerge from these exercises. For instance, the results might reveal whether a breach was detected internally or by a third party and how long (in days) the red team "dwell time" was before detection. (Notably, Mandiant's M-Trends 2024 report shows progress in this area; in 2023, the global median attacker dwell time fell to 10 days, down from 16 days in 2022, reflecting improved detection capabilities.) A well-run red team engagement will produce **concrete data on detection gaps,** and a good internal control group can measure response times, which a CISO can use to drive improvements.

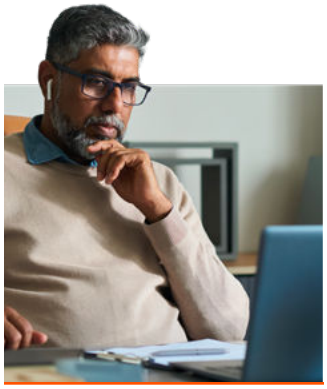

A red team exercise effectively measures an organization's "immune response," allowing security leaders to analyze and assess the outcomes of their teams.

✔ Identifying and prioritizing risks for reduction

Red teaming helps translate technical findings into **business risk terms.** By demonstrating the practical impact of certain vulnerabilities or process failures, it enables security leaders to prioritize what matters most. For example, a red team might show that a seemingly minor misconfiguration in the cloud could be chained with other weaknesses to cause a major data breach, turning an abstract risk into a vivid scenario for stakeholders. The exercise output is typically a report that classifies risks and recommends mitigations. This directly informs an organization's risk register and reduction efforts. Additionally, it ensures that remediation efforts and security investments are focused on the most dangerous attack paths *that attackers are likely to exploit,* rather than theoretical issues. In this way, red teaming acts as a **feedback loop** for strategic risk management, continuously aligning a security program with the evolving threat landscape.

✔ Strengthening security programs proactively

Overall, red teaming embodies a shift from reactive security (waiting for incidents to occur) to **proactive security.** It is a way to "train like you fight," exercising an organization's defenses regularly so that when a real incident happens, both technology and staff have essentially already survived similar scenarios. By uncovering weaknesses and prompting fixes, red teaming **drives continuous improvement.** It also has ancillary benefits: it raises security awareness among employees (they learn to be more vigilant if they know a phishing email could be a test). Mature organizations often institutionalize this by running frequent red team exercises and even continuous automated red teaming, often then leaning on collaborative purple team exercises to guide testing and collaboration and gather more insightful metrics.

One important distinction a CISO will note is the **difference between red teaming and traditional penetration testing.** While both are forms of ethical hacking, their scopes and objectives differ. Penetration testing is typically a narrowly scoped assessment of specific systems or applications for vulnerabilities, often performed with the knowledge of IT staff, to ensure those particular assets are secure. In contrast, red teaming is a broader, goal-oriented adversarial simulation: a red team mimics a real attacker by any means necessary, such as combining network exploits, social engineering, or physical intrusion, to achieve a goal, all while evading detection. Pen tests tend to be loud and announce every vulnerability found, whereas red teams are stealthy and look at the whole organization as an attacker would, often with only senior management aware it's a drill. The aim of a red team is not to enumerate every bug but to test defenders and defense-in-depth as a whole.

Both pen tests and red teaming are important in a security strategy. In fact, they complement each other.

Both pen tests and red teaming are important in a security strategy. In fact, they complement each other. Organizations with mature security programs employ both regular pen tests to harden specific assets and periodic red team exercises to evaluate holistic resilience. A CISO's challenge is to integrate these efforts such that findings from any test feed improvements across people, process, and technology.

In many sectors, the value of red teaming has become so recognized that it's mandated or strongly encouraged by **regulators and industry standards.** For instance, financial services in several regions must undergo intelligence-led red team exercises (e.g., the Bank of England's CBEST or the European Central Bank's TIBER-EU frameworks) to validate that critical assets like payment systems are well protected. These exercises are often conducted under strict oversight by regulators, aligning red team scenarios to real threat intelligence about adversaries targeting that sector.

This regulatory push underscores a key point: from the boardroom's perspective, red teaming is not just about finding holes—it's about assuring stakeholders (regulators, customers, and the board) that the institution's defenses work against high-end threats. Different industries face different threat profiles, and a CISO will tailor red team objectives to those mission-critical risks.

Threat landscape and red team objectives across industries

No two organizations are identical in their risk profiles; this is a fact every CISO understands well. An effective red team engagement must be contextualized to the industry and threat landscape of the target organization.

A bank, a hospital network, and a cloud software company each have different crown jewels and face adversaries with different motives. Having a tailored threat-informed approach ensures red team exercises simulate the attacks that matter most for that business. By understanding these sector-specific scenarios, a CISO can set red team goals that meaningfully challenge an organization's defenses and validate its readiness against the threats it is most likely to face in reality.

"How secure is this company we have just acquired?"

"How good or bad are our defenses?"

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

"How would a real threat target our company?"

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

"Does my security strategy reduce risk?"

Red teaming is effective in many different industries.
Let's look at three very different industries where CISOs leverage red teaming.

Finance and insurance

LEARN MORE

Banks, financial institutions, and insurers operate in a high-stakes threat environment. They are targeted by the full spectrum of adversaries, from nation-state APT groups to organized crime, due to the direct monetization potential and geopolitical impact of disrupting finance. Nation-state actors target banks to spy on or manipulate financial systems.

Given this threat landscape, a CISO in finance will direct red team exercises toward the scenarios that pose existential or significant business risk. Typical objectives include simulating attacks on high-value payment systems and transactions, such as trying to gain access to SWIFT terminals, core banking applications, trading systems, or internal fund transfer APIs. A red team might simulate an advanced attacker (e.g., an APT or well-funded criminal gang) whose goal is to fraudulently transfer funds or manipulate financial data.

Financial sector red teams also often test insider threat scenarios and physical security. For instance, one exercise might involve a red team member with covert physical access (maybe by tailgating into a data center or trading floor) to see if they can plug in a rogue device or access an unattended workstation, reflecting the risk of a malicious insider or intruder. This checks physical defenses and staff vigilance.

Healthcare and pharmaceuticals

Organizations in healthcare and pharmaceuticals face a dual-edged challenge: they hold extremely sensitive personal data and life-critical systems, yet often operate with legacy technology and tight resource constraints. Nation-state actors regularly target pharmaceutical companies and research institutions to steal intellectual property, including vaccine research, drug formulas, or clinical trial data.

Given these stakes, CISOs in healthcare and pharma design red team exercises to stress-test the resilience of critical healthcare operations and data protections. When it comes to healthcare providers, a red team might attempt to access medical devices or interfere with operations.

When it comes to a pharmaceutical company, a red team might emulate an APT by trying to exfiltrate drug research data from R&D labs or cloud storage. This tests the controls around sensitive data, including network egress monitoring, data loss prevention (DLP), and whether security teams notice unusually large transfers of data. It also highlights whether sensitive data is properly segmented and encrypted.

LEARN MORE

Manufacturing and industrial (OT/ICS)

Manufacturers and industrial companies (including those with OT environments, such as factories, energy/utilities, chemical plants, and other industrial control system setups) face threats that span both traditional IT and industrial control realms. Nation-state military and intelligence units are a top concern, especially for critical infrastructure and defense industrial base manufacturers. These adversaries may seek to conduct cyber-sabotage or espionage.

In industrial settings, a CISO's red team engagement is usually designed to bridge the gap between IT and OT security and see if an adversary can traverse from a corporate network into operational domains. A classic objective is to simulate an IT-to-OT pivot attack.

Physical and social engineering tests are very relevant to manufacturing as well. A red team may attempt a physical intrusion into a plant or data center by impersonating maintenance staff or contractors.

LEARN MORE

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 the attacker.

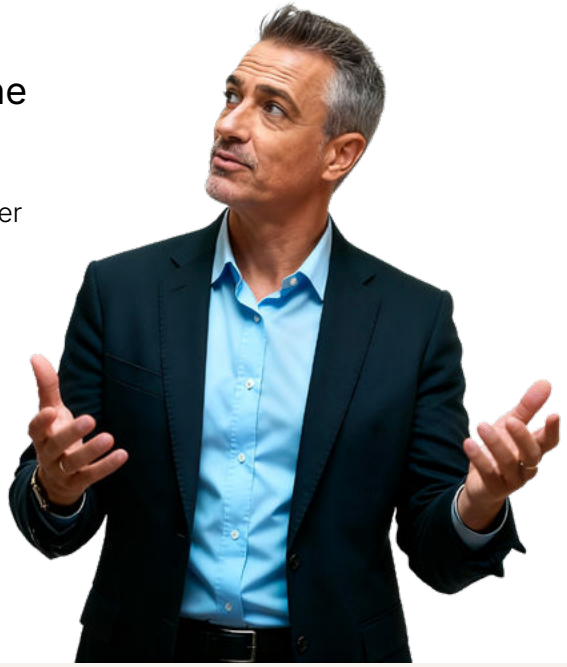

It's a little like a **"defensive onion"** because the more layers an attacker cuts through, the more likely they are to cry. However, one lesson red teaming continually reinforces is that adversaries are adept at finding ways around even well-crafted controls. It is crucial for security leaders to understand such **cat-and-mouse dynamics**: they reveal which controls are truly resilient and which may provide a false sense of security if not complemented by others. When it comes to security, as the saying goes, "Prevention is ideal, but detection is a must." This is because many preventive controls can and will be evaded by determined adversary efforts.

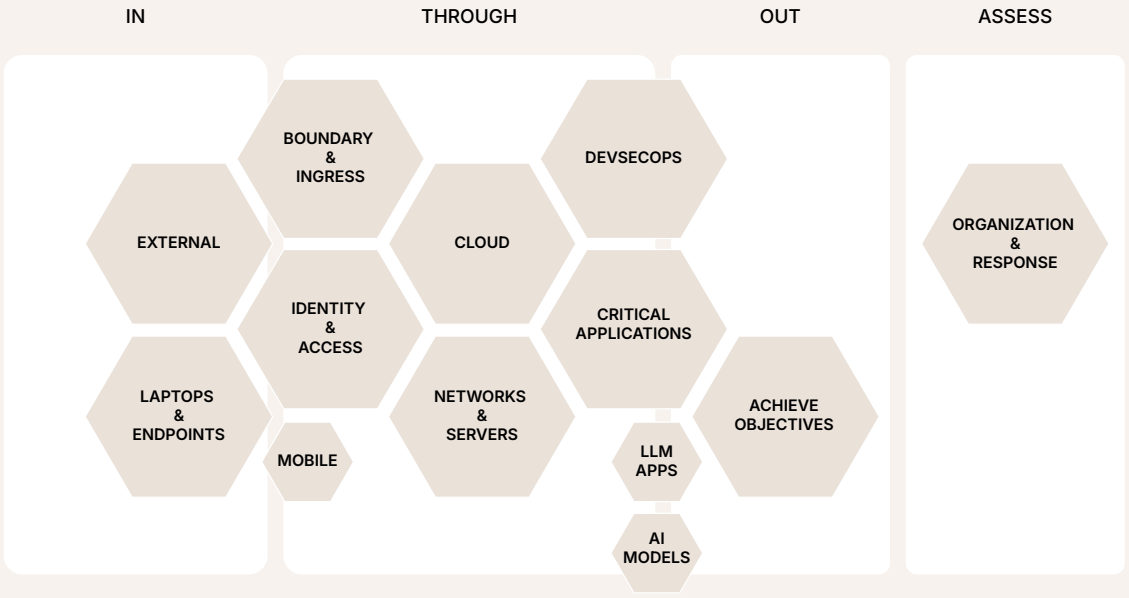

Technology, People & Processes

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

Virtually every company relies on authentication barriers, strong password policies, directory services (AD/Azure AD), single sign-on (SSO) for central identity, and MFA to add an extra layer beyond passwords. In high-security environments (finance and tech), phishing-resistant MFA (e.g., physical security keys or FIDO2 tokens) is deployed to thwart phishing and credential theft. Nevertheless, red teams routinely test the human and technical loopholes in these controls. One common evasion tactic is socially engineering users to unknowingly assist attackers. For example, the use of **MFA fatigue attacks** (or "prompt bombing") 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. Microsoft observed an explosion of these attacks; there were over 382,000 MFA fatigue attempts in 2022 alone. Indeed, this was how the notorious 2022 Uber breach began. Red teams replicate such tactics; they call up a user (as the Uber hacker did via WhatsApp) or send repeated prompts, often with a clever pretext ("IT is testing a new system; please approve the login"). Even a 1% success rate can be enough, but typically, red teamers have seen between 10% and 30% success rates with the exploitation of social engineering techniques. If a company only uses push-based MFA and lacks user education on this attack, a red team will likely succeed in gaining an initial foothold.

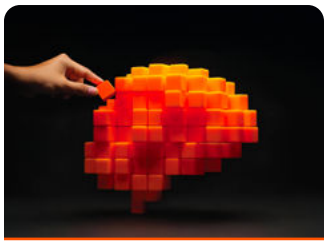

When it comes to social engineering, even a 1% success rate can be enough.

Another sophisticated ploy is **adversary-in-the-middle (AiTM) phishing,** which involves setting up a fake login page that proxies to the real one. When the target enters credentials and MFA, the attacker's site captures the session cookie. This essentially bypasses MFA, since the attacker never "breaks" it; they just ride the valid session token. This technique has been on the rise, with attackers deploying web proxy tools to steal session cookies and thereby "effectively bypassing MFA." Many companies that felt secure ("We have MFA on everything!") have been humbled when a red team demonstrates that stealing an authenticated session is as good as stealing the keys. Furthermore, such attacks often trigger no alerts because to the system, they look like a legitimate session (but launched from an unusual location). To counter this, leading organizations are now looking at phishing-resistant authentication, conditional access policies (device/user context checks), and monitoring for anomalous session patterns.

Endpoint security
(antivirus, EDR, and XDR)

Endpoints (servers, PCs, and laptops) are typically protected by antivirus or more advanced endpoint detection and response (EDR) agents that can detect malicious behavior patterns. Finance, tech, and other mature sectors often deploy full EDR/XDR coverage. EDR systems like CrowdStrike and Microsoft Defender are quite good at catching known malware and common attack behaviors, so how do red teams cope? They employ a variety of defense-evasion techniques. One approach is **obfuscation and custom tooling.** Red teamers make sure their malicious code does not match any known signatures and looks benign or unique. For example, a red team might write a custom implant or modify an open-source tool just enough to avoid AV heuristics. Recompiling malware with slight code changes can produce new hashes and evade some signature detection. Red teams also encrypt or encode their payloads so that static scanning can't easily recognize malicious patterns. They might only decrypt the payload in memory at runtime, giving traditional scanners nothing suspicious to latch onto on disk.

With enough skill, endpoint agents can be undermined. No single control is infallible.

Red team operators rely heavily on **"living off the land" techniques**, using legitimate system tools and binaries rather than custom malware to achieve their goals. For instance, instead of running an obvious malware executable, they might use PowerShell (a legitimate admin tool) to execute code from memory or abuse built-in tools like rundll32, wmic, or mshta that are present on Windows by default (these are often dubbed Living Off the Land Binaries, LOLBins). By doing so, they blend into normal system operations. **The best threat activity happens in plain sight;** malware that only uses native OS functions and doesn't escalate privileges immediately can operate without raising alarms. Real attackers do the same. Many advanced threats don't drop obvious binaries at all; they use known DLLs, scripts, and in-memory techniques to avoid leaving a footprint. In environments where EDR is configured to detect such things, red teams step up their game by using methods like disabling or bypassing EDR hooks. (Some known techniques involve unhooking userland EDR DLLs or using kernel vulnerabilities, though red teams do that carefully to avoid crashing systems.) There have even been cases where red teams deploy a Bring Your Own Vulnerable Driver (BYOVD) attack to disable endpoint protections (loading an old signed driver with a known flaw to get kernel access and kill security processes). These examples are extreme but illustrate that 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)

Perimeter security (like next-gen firewalls, intrusion prevention systems, and web application firewalls for web-facing apps) is a staple, especially in sectors like finance and telecom that historically had a strong network-centric security model. Additionally, many organizations rely on network segmentation to contain attacks (e.g., separating IT and OT networks or isolating sensitive data zones). Red teams often find ways to work around these **network barriers.** If segmentation is in place but certain bridges exist (like a jump server or a dual-homed system that connects to both networks), red teams will target those or target the users who have legitimate access to those. One company may say, "Our sensitive database is not reachable from the internet." A red team would find that a developer's machine with VPN access can reach it and render network segmentation moot by compromising that developer via phishing; they would get in through an authorized path, which is often easier than brute forcing the firewall.

With the shift to the cloud and remote work, traditional perimeters have become porous.

Moreover, with the shift to cloud and remote work, traditional perimeters have become porous. Red teams take advantage of this by **attacking cloud services directly** (bypassing on-prem controls) or by **abusing VPN and remote access solutions.** For example, if a VPN is protected by MFA but has a web portal vulnerable to something, they might exploit that to get in without needing to put in too much work. In some cases, red teams discovered legacy VPN credentials still active or default passwords on network gear, which is more common than one would hope, especially in telecom and manufacturing where legacy tech persists. Using those, they essentially walked right past the fortress walls through an old, unlocked gate. Such findings reinforce the need for rigorous network asset management and the retirement of legacy systems or at least compensating controls around them.

Another network evasion tactic involves encrypted channels and tunneling. Red teams often use encrypted command and control (C2) channels over common ports (HTTPS on 443, DNS tunneling on 53, etc.) to hide their traffic in normal flows. They might stand up an innocuous-looking cloud server or CDNs (shadow fronting) and have the compromised machine beacon out to it, blending with standard TLS traffic so that it's hard for an IDS to flag. If an organization relies solely on perimeter detections without deep packet inspection or anomaly detection, these communications can fly under the radar.

For CISOs, this is a reminder that rigorous external attack surface management and patching are still crucial. Red teams basically act as friendly opportunistic hackers here, finding the cracks before the bad guys do.

Email and endpoint hygiene vs. phishing

Since a huge portion of attacks begin with phishing, companies invest in email security filters, spam training, and phishing simulations. Red teams routinely test these by crafting very convincing phishing emails to see if they can slip past filters and lure employees. They might register lookalike domains or exploit trusted services (like using a Dropbox link, a Google form, or a calendar invite). One common scenario is using a **payload-less phish** with an email that just asks the user to click a link to a fake login (since many email filters focus on attachments). By not attaching malware and instead harvesting credentials or sessions via a link, the red team can evade attachment sandboxing controls. If the organization has good email filtering, the red team will adapt and maybe target users via other means (like SMS phishing or voice calls, aka "vishing"). This is why an increasing number of companies are adopting phishing-resistant MFA and user education. Still, all it takes is one clever email (perhaps impersonating the CEO or referencing a current event) sent at the right time to get a click. Red team reports often include statistics like:

CISOs must track valuable phishing metrics found in red team reports in order to build security awareness programs.

Out of 100 employees targeted with a tailored phishing scenario

- **X% clicked**
- **Y% even entered their credentials**

These are **valuable metrics** for the CISO to tailor security awareness programs.

Data protection and monitoring

Lastly, consider data-at-rest protections like encryption and DLP. 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 (e.g., sensitive files on a shared drive that is not actually encrypted). They might also reveal that once an attacker has admin rights on an active system, they can often get the data in plaintext from memory or via the application. DLP solutions that scan for sensitive data leaving a network can be bypassed by exfiltrating in chunks, using encryption, or even hiding data in images or other files. For instance, a red team might have exfiltrated a database by encoding it inside a series of DNS queries, which many DLP setups wouldn't catch if they aren't monitoring DNS traffic closely.

In summary, attackers play a **constant game of innovation against defenders.** For every defensive measure, there are known evasive techniques; EDR can be blinded, MFA can be sidestepped, users can be tricked, segmented networks have intersection points, and so on. The job of a red team is to demonstrate these in a controlled manner. For a CISO, seeing how these bypasses occur is immensely valuable: it directs them to strengthen defense in depth. It may lead to deploying complementary controls like user entity behavioral analytics (UEBA) to spot anomalies that evaded preventive blocks and ensuring comprehensive logging and monitoring to catch what isn't blocked. One clear example is the emphasis on detection and response. Given that even well-resourced companies like those in tech/finance get breached via novel means, having a capable detection program is vital. Many red team engagements conclude with recommendations on improving logging (so that those living-off-the-land techniques aren't invisible) and improving the correlation of signals. In effect, red teaming helps organizations not only plug holes but also refine their "radar" to spot the subtle signs of an intruder who found a way in.

Attackers are constantly innovating. For every defensive measure, there are known evasive techniques.

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. Addressing these findings is crucial for truly improving organizational resilience.

As noted earlier, the majority of breaches involve some human action or error, whether it's clicking a malicious link, using a weak password, or misconfiguring something. Red teams take advantage of this by incorporating social engineering attacks into their campaigns. This might involve phishing emails, phone calls (vishing), SMS (smishing), and even in-person deception (if it is in scope). For example, a red team member might call a company's IT helpdesk while posing as a frantic executive: "I'm travelling and can't log in. I've forgotten my token. Please, can you reset my password urgently? I have a board meeting in 10 minutes!" If the helpdesk has a weak authentication process for callers or something that can be bypassed with OSINT, it might just comply, effectively letting the attacker reset the exec's password and bypass MFA (maybe by also convincing them to disable MFA "just for now"). This tests **process adherence:** do employees follow security policy under pressure or do they bend the rules? The results can be illuminating. Many red teams find 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.

The majority of breaches involve some human action or error.

Red teaming also shines a light on **process failures and organizational silos.** For instance, an exercise might reveal that the incident response process looks good on paper but breaks down in practice. Perhaps the SOC analysts noticed some suspicious activity (say, the red team triggering a malware alert) but the incident never got properly escalated to management because the alert was dismissed as a false positive, was too low priority, was invisible on a "new" dashboard, or got stuck in an email queue. In a red team debrief, the timeline of "here's when we did X, here's when/if it was noticed, and here's how the staff responded" is incredibly valuable. It might show that the on-call process on a weekend is unclear, that the SOC is too understaffed to investigate every alert, or that they did respond but the communication to the broader team failed (e.g., they contained a server but didn't inform the app owners, resulting in confusion). These are systemic issues in incident response and crisis management that a red team helps identify without the cost of a real incident. Many organizations follow up a red team engagement with a tabletop or purple team exercise involving the same scenario to further drill response procedures, this time with key stakeholders aware.

CISOs should follow up a red team engagement with a tabletop or purple team exercise to further drill response procedures.

Another example of a systemic issue is poor joiner-mover-leaver (JML) processes for user accounts. A red team might find dormant accounts that still have access because HR and IT didn't disable them when someone left or because contractors were granted broad access but never audited. By exploiting such an account, a red team can reveal a breakdown in the identity governance process. This often prompts an organization to invest in identity management solutions or tighter HR–IT coordination. Similarly, change management processes can be tested: red teams have successfully deployed rogue devices or applications inside a network and watched if anyone notices an unapproved change. If they can operate for weeks off a server they set up and have no one question it, this indicates a gap in asset management and change detection.

A healthy organization, like a healthy body, should have reflexes; it should react when something is amiss.

Many red team engagements expose **issues of culture and coordination.** For example, in some companies, the security team might notice something odd but hesitate to raise an alarm for fear of false alarms or because the company has a blame culture. A healthy organization, like a healthy body, should have reflexes; it should react when something is amiss. If employees are afraid to report a lost badge or an unusual email because they think they'll be punished, that's a cultural issue that attackers exploit (they thrive in silence and fear). Red teams sometimes intentionally create obvious signs to see if employees report them. One metric could be the reporting rate. For example, Verizon noted in its 2024 DBIR that 20% of individuals reported phishing in simulated exercises. Only 11% of those who clicked still went on to report these instances, meaning some clicked and then realized their mistake. These numbers, while slowly improving, suggest a lot of attacks could still sail through without being reported. A CISO wants to improve such rates because human sensors (employees reporting incidents) are as important as technical sensors. Red team exercises boost these by raising awareness ("See how convincing a phish can be? Please report anything suspicious quickly!").

Importantly, red teams help find problems not just in IT systems but in the interplay of technology with people and process. For example, a red team may discover that backups are regularly performed (tech process) but the process to restore them is untested or takes too long—a huge gap if ransomware hits. Or they might reveal that while there is an incident response plan, the business continuity plans are not aligned (maybe the plan to continue operations during an IT outage is unrealistic). By simulating an incident, these disconnects become apparent. One could liken such simulations to a fire drill; you might have an evacuation plan on paper, but until you run a drill, you won't know that, say, one exit door is jammed or people congregate in the wrong place. Red teaming is a cyber fire drill that exposes the non-obvious issues.

Red teaming outcomes often highlight the need for organizational learning and adaptability. The most mature organizations treat each red team finding as a lesson, update their training and processes, and even share lessons across industries (anonymously, via ISACs, conferences, etc.). They foster a culture where being "beaten" by a red team is not failure but an opportunity to improve, akin to how regular exercise breaks muscle fibers only to rebuild them stronger.

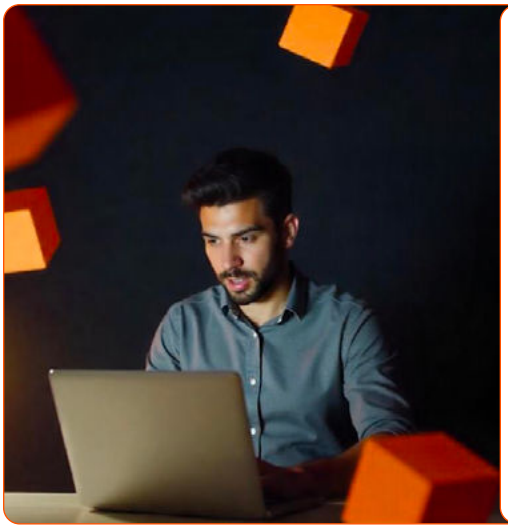

Consider the scenario of an advanced attacker who has partial access. The security team may detect and contain them on some machines. Does the IT team then quickly reimage those machines and remove persistence, or do bureaucratic delays allow the attacker (red team) to regain a foothold? Perhaps the red teamers observed that after being "caught" on one system, they still had valid credentials that nobody thought to revoke, letting them slip back in through a VPN. That's a process gap in remediation. Responding teams must not only remove malware but also invalidate credentials, sessions, etc. if they suspect compromise. Red team exercises can uncover if such follow-through is happening.

To go with the health metaphor, small, controlled doses of stress (red team drills) build the resilience "muscle" of organizations. People become more astute (e.g., employees might start politely challenging that unknown person roaming the office or verifying unusual requests through a second channel). Processes get refined with each test (maybe the company institutes a stricter helpdesk verification protocol because the red team breached it). Over time, these adjustments lead to a security posture where both technology and humans are more prepared to prevent, detect, and respond to real threats in a concerted, agile manner.

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.

These then strengthen security strategy, justify investments, and inform stakeholders. CISOs report on cybersecurity readiness using red team metrics to drive improvements in SOC and IT systems and secure budget and resources to support high-level decision-making and long-term resilience.

Budgeting and investment

One of the most immediate impacts of a red team report is improvements in budgeting and project prioritization. This report is **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. For example, if a red team demonstrates that an attacker moved laterally with ease due to a lack of network segmentation, a CISO can justify spend on network micro-segmentation or network access control (NAC) upgrades. If the exercise shows that the SOC missed the attack because of limited log retention or coverage, it could justify expanding logging infrastructure or investing in a managed detection service. Often, these are things security teams have wanted to do, but now, they have the hard data to back the request. An executive might not loosen the purse strings just on hypotheticals, but telling them "During the simulation, our team failed to catch the intruder for 10 days because we lacked visibility in Cloud Region X, this $Y million investment will fix that" is compelling.

Metrics derived from a red team engagement play a starring role here. Organizations can coordinate with their blue teams to establish metrics like MTTD, MTTR, and a heat map of security control performance against specific tactics. A CISO might present to the board: "Our current average detection time for a breach (in this test) was 3 days. Our goal is to get it under 24 hours by end of year. To do that, we need to invest in XYZ technology or training." Another key metric can be the "eradication success rate"; can an organization be sure that it has truly eradicated an intruder? If a red team is able to regain access due to residual footholds, this highlights a need for **better clean-up tools or processes.** These metrics turn the nebulous concept of "security posture" into something measurable and improvable, which executives and boards appreciate.

Red team metrics turn concepts like "security posture" into something more measurable and improvable.

Findings Table

April 04, 2025 | SampleCorp - Intelligence Led Threat Simulation

Index	Title	Priority
RTO1	Lack of MFA Enforcement for Privileged Cloud Accounts	P1
RTO2	Cleartext Credentials Stored in Active Directory Description Fields	P2
RTO3	Excessive Privileges Assigned to Intern Account	P2
RTO4	Azure AD Connect Server Storing Extractable Sync Credentials	P2
RTO5	Unmonitored Federation Configuration Changes	P2
RTO6	Unrestricted RDP Access to Sensitive Infrastructure	P2
RTO7	Use of Unmonitored Contractor Workstation for Persistent Access	P2
RTO8	Alert for Suspicious LDAP Activity Ignored by SOC	P2
RTO9	Allowlisting of Cloudflare WARP Enabled Egress Bypass	P2
RT10	Phishing Protections Did Not Block OneNote WSF Payload	P3

Please refer to the Risk and Priority Key and Risk Matrix for additional details.

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 allocations toward more user-focused controls like advanced phishing training and new email filtering solutions or perhaps move more apps to SSO with phishing-resistant MFA. If findings consistently involve endpoint issues, maybe greater investment in EDR tuning or moving to a managed EDR service might be considered. Thus, red teams act as **feedback mechanisms** for whether previous investments are yielding results or if new ones are required. It's not uncommon for companies to conduct a red team engagement before a major security initiative and again after to validate that the needle has moved. For example, a Year 0 engagement might show many gaps. Year 2 could reveal far better detection and fewer paths, justifying the ROI of these improvements.

Board and executive reporting

Boards of directors today are acutely aware of cyber risks. Many ask management, "How do we know we're secure? Have we tested ourselves?" A red team exercise provides a narrative that the CISO can bring to the board to **answer these questions credibly.** A CISO might share a sanitized summary of the following scenario:

> **We conducted a simulated APT attack on our company. The red team, acting as attackers, attempted to breach our critical systems. Here's what they were able to do and here's where they were stopped.**

This storytelling is powerful. It avoids jargon and instead uses a plot, giving the board a clear picture of risk in context, not just as a theoretical. Crucially, it also **highlights improvements.** For example, a CISO might report, "Last year, our red team got to the crown jewels undetected. This year, we detected them halfway through the kill chain, an improvement, and next year, we aim to detect at initial ingress." Such framing shows progress and accountability.

How do we know we're secure? Have we tested ourselves?

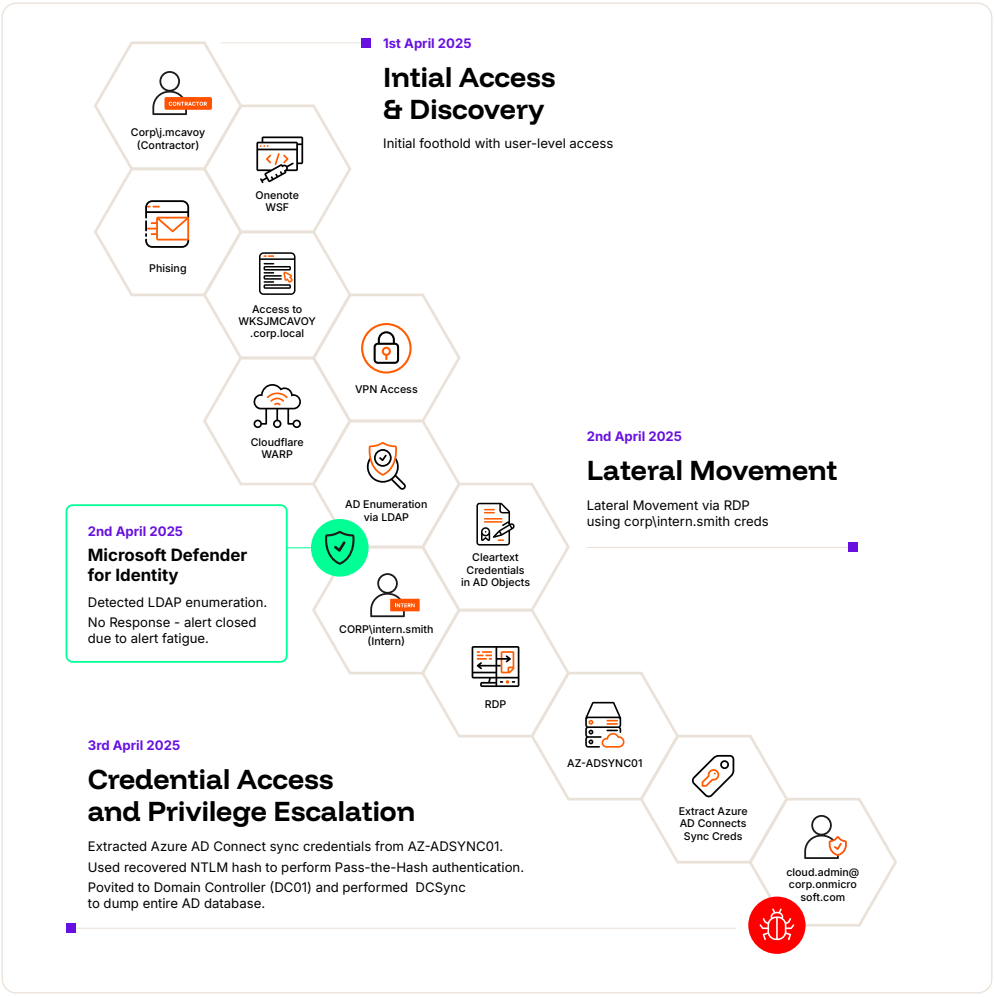

KEY

Significant Compromise
Domain Administrator or Enterprise Administrator

Defensive Action
Detection by security tooling

Boards also love **benchmarks and frameworks.** A CISO can map red team findings to frameworks like MITRE ATT&CK or NIST CSF to show coverage. For example, they might present a chart of the MITRE ATT&CK tactics where an organization has strong vs. weak detection coverage, as revealed by the red team. If a heat map shows green (good) on initial access and execution (maybe the red team caught phishing and malware execution) but red (needs work) on lateral movement and exfiltration (it missed data staging and exfil), this is a straightforward way to communicate priorities to a board. A more accurate picture can be gleaned from a purple team assessment where coverage is key and many red team TTPs are tested over the frameworks. This can provide better insight into real metrics, systemic issues, and control efficacy. The board doesn't need to know the technical nuance. They simply need to be shown "We're strong in these 5 areas, moderate in 3, and weak in 2, and plans are underway to address the weak spots." A red team exercise, and the subsequent organizational analysis, in essence provides evidence-based assurance. It's much more convincing than simply saying, "We think our security is good because we have X tools." Instead, it is much more effective to say, "We challenged our defenses with a real-world simulation and learned A, B, and C. Now we are fixing those and will retest." This approach resonates well with good corporate governance, akin to internal financial audits or business continuity drills that boards are familiar with.

Use red team results to quantify potential impact reduction and demonstrate cyber risk management.

Another board-level angle is using red team results to **quantify potential impact reduction.** If, for instance, a red team determined that had there been a real attack, the cost could have been a major data breach with millions in fines (this kind of data is sometimes hard to find, but Kaggle and industry reports like IBM can provide sufficient insight). In this case, the CISO can argue how the investments and changes post-red teaming are averting such costs. 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. Maybe the first one was able to compromise 5 high-impact assets, the next only 2, and so on, 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 (like specific file hashes, command line strings, and C2 domains) and retroactively check if their tools picked them up or not. If not, why? Perhaps the logs weren't there or thresholds were too high. They then improve those. For instance, if the red team used a tool that injected into dllhost.exe and it wasn't caught, the SOC might implement new behavior analytics around processes spawning unusual children or scanning memory. If the red team succeeded in moving laterally using WMI and that wasn't alerted, the SOC can tune their EDR to flag WMI execution originating from non-admin machines. This process of **detection engineering** is often accelerated by red teaming.

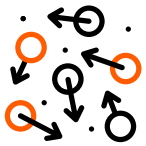

Red teaming provides a continuous training loop for the defense team under realistic conditions.

Additionally, red teaming can be used to train the blue team in a "lessons learned" way. Some organizations even do replays or purple team sessions after the main covert red team engagement is done. In a purple team model, the red and blue sit together; red shows "here's what we did, step by step" and blue verifies where they saw something or not. They might rerun parts of an attack and ensure detections fire. This collaboration **fosters knowledge transfer** so that the SOC is better prepared for real threats. Over time, these drills significantly sharpen defenders' skills. They start recognizing patterns ("This looks like that thing our red team did with DNS tunneling; investigate immediately"). In effect, red teaming provides a continuous training loop for the defense team under realistic conditions.

Strategy

At the highest level, regular red teaming cultivates what might be called **strategic cyber resilience.** Resilience isn't just about preventing attacks but also about ensuring that an organization can continue to operate and quickly recover even if an attack succeeds. Red team exercises often include objectives around testing DR plans or seeing if alternate systems stay intact. The findings thus inform not just how to prevent breaches but how to limit damage and rebound from them. If a red team found that a certain attack could take down a critical system and it would take days to rebuild, the CISO and IT can prioritize making that system more fault-tolerant or having hot standbys. In other words, red teaming can (in)validate the assumptions in business continuity plans.

By incorporating red team scenarios into broader risk scenarios, leadership can develop a **more robust risk management strategy.** For example, if multiple red teams show that third-party vendors are a consistent risk (perhaps they always manage to phish a contractor or use vendor credentials), an organization might decide to reduce that risk by changing how vendors connect. Perhaps they may implement stricter network segmentation for vendors or require hardware tokens for them. This creates alignment between technical findings and enterprise risk decisions, possibly even affecting contractual requirements for vendors.

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 engagement, we aim to detect the red team at least at the data exfiltration stage, not after it has simulated customer data theft as happened this year." This achievement would indicate improved resilience. Some advanced organizations have moved to a model of continuous red teaming, like Bugcrowd's Continuous RTaaS, one of our three red teaming models. The idea is to go beyond a once-a-year check to leverage the ongoing validation of controls in a DevOps-like continuous cycle. As a result, red team outcomes feed directly into the daily, weekly, and monthly tuning of defenses; regular red teaming ensures that defenses stay tuned. At the executive level, this becomes part of the organization's resilience story. The assumption changes from "if we get attacked" to "when attack attempts happen, we have confidence in our resilience and response, and here's the evidence of how we have tested and strengthened our internal processes."

> **A bank might realize from red team result patterns that insider threats are a big gap (since red teams acting as insiders had free rein). Strategically, it might establish an insider threat program, deploy user behavior analytics, or enforce stricter least privilege and monitoring on employees. These are all major initiatives that can stem from red team lessons.**

Red team results often feed into **compliance and external communication** as well. Many organizations include summaries of such testing in their annual reports or customer assurances to demonstrate diligence. It builds trust to be able to say, "We don't just meet compliance. We actively test ourselves regularly and improve." In sectors like finance and critical infrastructure, demonstrating this capability can even favorably influence regulators or insurance underwriters (cyber insurers are increasingly asking if organizations do red team exercises, as they indicate maturity).

Bugcrowd's Red Team as a Service (RTaaS)

Bugcrowd's Red Team as a Service (RTaaS) is a new delivery model where companies can purchase access to a crowdsourced red team, assembled to have the right skills to target their organization. During an operation, a red team can also scale up or down as they wish—crowdsourcing expertise in specific domains to give teams full flexibility.

Another dimension of flexibility is in frequency and collaboration. Companies can choose between assured, blended, or continuous red team operations. With assured red team operations, a bespoke, dedicated red team works with the company for a defined amount of time. This option focuses on reliability and deep client–team interactions, with one team running the entire red team assessment. With blended red team operations, crowdsourced operators conduct the initial red team phases (similar to a private, on-demand bug bounty). After the initial phases, selected crowd operators will focus on the latter phases of a red team assessment. This option combines traditional red teaming with scalable crowdsourced operations, giving companies flexibility and precision right where they need it. With continuous red team operations, both the in and through phases are modeled as private bug bounties. Small teams of vetted operators continually rotate through and attack systems, each bringing a different set of attack vectors. This option provides continuous coverage and replicates the most realistic attacks. With these options, companies can choose the model that works for their budget and security posture.

RTaaS offers CISOs flexible models that work with their budget and security posture.

RTaaS also fits neatly with our other products. Companies can engage red teams as complements to their bug bounties or VDPs. Crucially, the same collaboration tools that let us coordinate massive crowdsourced security tests also let us handle the collaboration within red teams.

Our Vulnerability Rating Taxonomy (VRT) and reporting tools also make it easy for red teams to create useful, informative posttest reports that companies can quickly implement. Customers can expect asset discovery, cyber threat intelligence, covert infrastructure building, attack approval chains, and event log-driven reporting on our Platform.

Our RTaaS program also makes life easier for red team operators, enhancing their effectiveness. Red team operators get access to Bugcrowd's External Attack Surface Management solution, which provides asset discovery tools to complement their own OSINT. We also provide operators with threat intelligence data from across our tools and crowdsourced teams, bolstering each team's knowledge. Operators can use our covert infrastructure tools to get started with their tests with minimal setup time. Attack approval chains let operators and red team managers quickly communicate and decide on the right attack vectors. Finally, our event log-driven reporting uses the covert infrastructure and attack approval chains to give companies continual visibility into red team activities and reports.

With this combination of flexible skill-based teams, easy collaboration, and detailed reporting, Bugcrowd's RTaaS removes the barriers to red team operations. Our goal is to help every company improve its security posture by identifying core issues with the help of a skilled red team that can think like a threat actor.

Our goal is to help every company improve its security posture by identifying core issues with the help of a skilled red team that can think like a threat actor.

Conclusion

In the complex, ever-shifting cybersecurity landscape, CISOs constantly ask, “Are we as prepared as we think we are?” Red teaming provides a profound and practical way to answer this question.

Through the lens of simulated adversaries, red teaming reveals the truth about an organization’s defenses—including the robust and 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. From finance to healthcare, and manufacturing to tech, each industry has its unique “attack surface anatomy” and threat motivations. A CISO’s red team program is tailored to these specifics. By incorporating detailed threat intelligence (the who and why of potential attackers) and focusing on mission-critical objectives, red teams ensure that organizations aren’t preparing for yesterday’s war but today’s and tomorrow’s. While technical vulnerabilities are important, often, the most impactful findings exist in the seams between technology, people, and process, whether it’s an employee who can be tricked by a cunning phish, a procedure not followed under duress, or a security control bypassed by an attacker’s ingenuity. These insights galvanize holistic fixes:

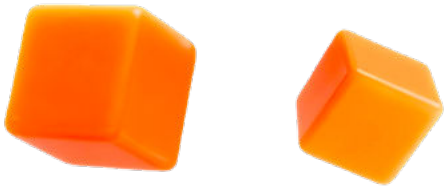

Better training **Clearer processes**

and more **Resilient architectures**

The benefits of red teaming accrue not only in shoring up defenses but also in strengthening an organization’s cyber culture and confidence. There is a noticeable shift in mindset that occurs when employees and leadership internalize red team lessons: security becomes a shared responsibility and an ongoing practice, not just a checkbox. For instance, after a red team engagement, it’s common to see employees more alert to suspicious emails or strangers; they’ve seen how easily a breach can start with a single lapse. IT and security staff often become more collaborative, and at the leadership level, the board and executives gain a clearer understanding of cyber risk in practical terms, enabling better governance and support.

Regular adversarial testing ensures that even as new technologies (like cloud, IoT, and AI) are adopted, an organization's risk blind spots are brought to light and can be addressed.

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." This shows a cycle of continuous improvement that management can appreciate. Red teaming allows organizations to track metrics for security posture much like a business might track and optimize KPIs for quality or uptime. It's not merely about preventing breaches; it's about ensuring that an organization can withstand and quickly recover from cyber shocks.

Red teaming is an indispensable tool for achieving and demonstrating cybersecurity excellence.

For a CISO, red teaming is an indispensable tool for achieving and demonstrating cybersecurity excellence. It brings together threat intelligence, creative hacking skills, and rigorous testing of defenses to provide a 360-degree view of security effectiveness. By addressing the findings from these exercises, organizations across all industries can significantly reduce their risk of catastrophic cyber incidents. Just as importantly, they can face inevitable attacks with confidence, knowing that they have drilled and hardened their responses. 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 their 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.

PLATFORM TOUR

See the Bugcrowd Platform in action

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

Unleash Human Creativity for Proactive Security

TRY BUGCROWD