bugcrowd

Ultimate Guide to

Red Teaming

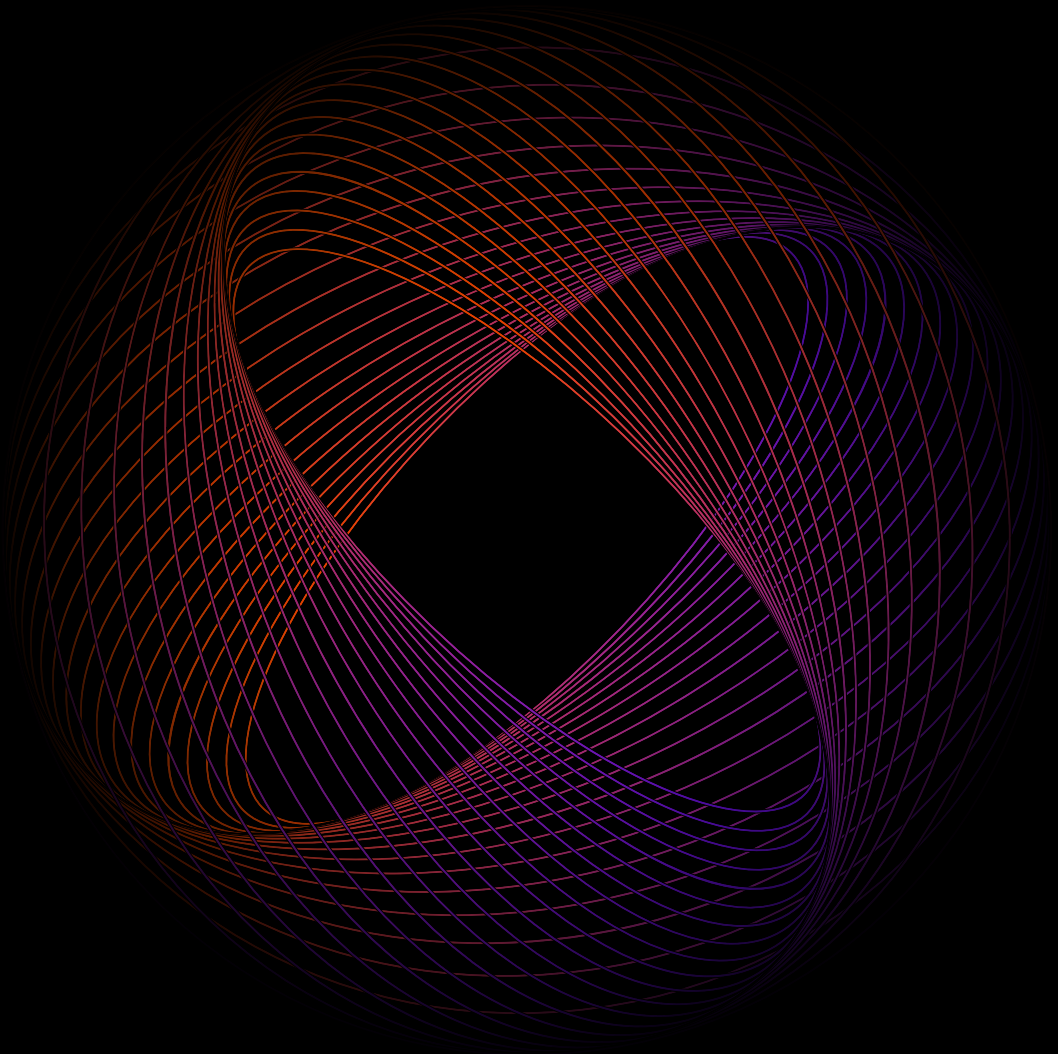

Table of Contents

Introduction

While the Cold War raged in the 1960s, the U.S. military invented a practice often used by cybersecurity teams today. The RAND Corporation, a think tank working with the U.S. military, simulated war games between the United States and the Soviet Union. On their simulated battlefield, Soviet units were red, and American units were blue. Out of this sprang the modern concept of red teaming.

While modern cybersecurity red teams focus on threat actors exploiting software and hardware vulnerabilities (and much less on invading armies), the core idea remains much the same. Modern red teaming simulates threat actors attacking a system. The resulting knowledge, revealing a wide range of vulnerabilities and impact points, can serve as an X-ray of an entire organization. With this information, security teams can mend weak spots by reducing vulnerabilities and forming better defensive controls.

Despite its value and applicability, red teaming is still underused. It's hard to hire a good red team. Even with a good red team, attack vectors change constantly, and new skills are required. It's even harder to update a red team while still keeping it small enough to be nimble.

At Bugcrowd, we're focused on making those challenges a thing of the past and on making red teams a part of every company's future.

Red teams are that powerful. That's why we made this guide: to give you a deeper understanding of red teaming and to suggest ways to incorporate it into your security strategy.

What is red teaming?

Red teaming is a security process where a team of operators (the “Red Team”) attacks your system using similar tactics, techniques, and procedures (TTPs) to threat actors. During the exercise, the “Blue Team” (unaware that it’s the Red Team, and not actual threat actors, attacking) will try to stop the attack using all tools at their disposal.

In an ideal world, both the Red and Blue Teams would write post-exercise reports on the methods they used. The entire security team would then use this report to patch the chain of vulnerabilities used in the attack, staying—for now—one step ahead of threat actors. In reality, teams may be too busy to write such reports, but they can still improve their defenses from the tacit knowledge gained.

Red teaming is always covert. The Blue Team doesn’t know the exercise is taking place.

So, they treat it like an actual security intrusion. In purple teaming, which is the non-covert version of red teaming, the Blue Team knows the exercise is happening. Purple Team operations can be collaborative, where the Blue Team communicates with the Red Team to check prevention, detection, and response protocols against the Red Team’s activities. Purple teaming can also be passive, where the Blue Team sees the activities of the Red Team but doesn’t communicate with them.

In its broadest sense, red teaming can be any adversarial simulation, not just security-focused ones. LLM red teaming involves researchers simulating adversarial behaviors and trying to get a model to say something harmful or unsafe. Researchers iteratively attack and then improve the model or guardrails to make it safer and less harmful. The techniques used are different, but the mindset is the same—play the role of a motivated adversary and simulate how they would break your system.

Red teaming vs. pen testing, VDPs, and bug bounties

People may wonder what the difference is between red teaming and the other security activities out there, namely pen tests, vulnerability disclosure programs (VDPs), and bug bounties.

In short, red teaming is open-scope, adversarial, deep, and covert. The other three activities don't fit this description. Pen testing usually focuses on a specific set of features and tests for common vulnerabilities. VDPs and bug bounties are open-scope as a best practice but are not adversarial because they don't fully assess impact. In contrast, red teaming involves teams using any means necessary (within the rules of engagement), including all the TTPs of threat actors, to get access to systems and exfiltrate relevant information to demonstrate impact. A red teaming exercise can also be done at a predictable cadence and is completely configurable.

All of these programs play a role in a strong security posture, but red teaming stands out for its ability to directly assess an organization's exploitability and resilience.

Red teaming use cases

Red team operations can be valuable in a variety of situations, from adhering to compliance mandates and securing new product launches to post-incident recovery. We'll cover why red teaming can be helpful in each of these scenarios.

Understanding current security posture

Red team operations provide a thorough exploration and understanding of your security. Red teams have an open scope, so they will try many more attack vectors than professionals using other security testing methods. After a red team test, a company can access much more information about the vulnerabilities in its attack surface and the gaps in its defenses. With this knowledge, the company can fill the gaps and improve its overall security posture.

Maximizing compliance

For companies and products that handle critical data, there are often mandatory security requirements. For example, companies handling payment data have strict security testing requirements as part of PCI-DSS compliance. The finance industry is governed by multiple frameworks such as CBEST, iCAST, CORIE, TIBER, and DORA. Red team assessments go above and beyond the usual pen testing that companies pursue as part of their compliance efforts. Not only do red team assessments provide security coverage and reveal gaps, but they also go deeper and test less-common attack vectors and determine the associated risk and root causes related to attack paths. As a result, companies can build stronger defenses while maintaining compliance. Within critical industries, red team frameworks exist to ease the path to security.

Incident recovery

Apart from helping companies patch up vulnerabilities, red team operations also help companies practice their post-incident recovery protocols. If a security team didn't spot any of the Red Team's attacks, this means the company in question needs to boost its detection practices. If the security team detected attacks but couldn't remove some of the red teamers from internal networks, then the team knows they'll need to work on their incident response. If the security team couldn't prevent a red team simulated ransomware attack, they will now better grasp the deficiencies in their backup and user alerting processes. The security team will essentially go through trials of incidents and recoveries, preparing them for real situations.

Red teams at work

At their core, red teams try to simulate threat actors attacking a system. Doing this well requires creativity and expertise, but the high-level process is similar for most red teams.

High-level process

Red team operations can be broken down into four phases: in, through, out, and assess. We'll cover each in the following.

IN PHASE

The in phase focuses on gaining initial access to a system or organization. During this phase, threat intelligence analysts provide threat intelligence about the target organization. In engagements with no threat intelligence, the red team performs the reconnaissance. This can include any of the following information:

Organization	Employee	Technology stack	Open source intelligence (OSINT)
Building locations, team structures, business units, phone numbers, emails, core products	Emails, phone numbers, LinkedIn accounts	Technology providers, email configurations, public IP addresses, open ports	Any publicly accessible information (e.g., previous compromises, leaked credentials)

Once they have collected and validated their threat intelligence, red teams devise attack scenarios. These are usually an initial access attack vector, threat profile, and set of objectives that a threat actor would target to try to gain initial access. An attack vector can be a combination of multiple vulnerabilities, misconfigurations, people, or legitimate tools and processes. The Red Team will then commence the attacks and execute on their attack vectors. As the operation proceeds, they will often modify the vectors or add new ones.

Once they gain initial access to a system, the next phase begins.

THROUGH PHASE

In the through phase, red team operators move laterally and escalate their access and privilege within the system, network, or organization. The goal of this phase is to find and execute any attack vectors that give privileged access to target data or systems. This can look like targeting employees within a company through phishing attacks or exploiting misconfigurations in cloud setups. Red team operators will chain attack vectors to gain the access they need; each vector will provide a small escalation, but the sum could get the Red Team right to their target.

OUT PHASE

With privileged access within an organization's systems, red team operators shift to simulating impact. A real threat actor may run a ransomware attack or leak critical user data. Red team operations won't harm a company, but they need to show potential impact. Experienced red team operators will weigh the potential impact simulations (e.g., deploying ransomware that encrypts only certain files or accessing the CEO's emails) by harm and adherence to objective.

Once a red team has chosen and executed an impact simulation, the active portion of the Red Team test has come to an end. The red team will then perform a clean up and attempt to remove indicators of compromise.

ASSESS PHASE

After a simulated attack ends, the Red Team writes an extensive report detailing all the attacks they tried (also called an attack narrative), the vulnerabilities or root cause issues they found, and the defenses they worked around. Some reports may even contain the full attack chains, which diagram the sequences of attack vectors in a graph format. By including as much detail as possible, the Red Team helps the security team understand exactly how the former were able to perform their attacks. The report also includes suggested recommendations for remediations for uncovered vulnerabilities or root causes. This report can be shared with the right stakeholders, with the goal of running root-cause analyses and fixing systems. Done well, the report can serve as a security roadmap for the entire organization.

Common frameworks

Frameworks provide a steady base for red team operations. They raise the floor by making it easier to communicate the results of an operation. Specific frameworks also exist for different domains (e.g., CBEST for finance), helping finance firms focus red team operations on regulation. Here are some of those frameworks:

MITRE ATT&CK

MITRE Adversarial Tactics, Techniques, and Common Knowledge (ATT&CK) is a repository of information on threat actor tactics and techniques. It splits threat actor behaviors based on 14 different tactics and lists specific techniques for each one. For example, within the tactic of lateral movement, the ATT&CK framework mentions techniques like internal spearphishing and remote service session hijacking. ATT&CK is often used after a red team operation to categorize attacks used in an exercise. Red team operators don't use ATT&CK during an operation since they need to be creative. Nevertheless, the ATT&CK framework maps many attacks (even creative ones) to a common set of understandable vectors, making it easier to communicate attacks in the debrief.

CBEST

CBEST is a framework for security testing, specifically for financial firms in the UK. It's designed to help financial firms secure their services in full compliance with regulations. CBEST splits the testing process into four phases: initiation, threat intelligence, penetration testing, and closure. The framework lists steps that must be taken in each phase along with rules on who can conduct testing. CBEST was the first framework of its kind, leading to the creation of others like iCAST, CORIE, and TIBER. In this way, CBEST paved the way for intelligence-led red team assessments and standardized them in critical industries.

Beyond frameworks

There's a whole universe of frameworks out there, both general (e.g., Lockheed Martin's Cyber Kill Chain) and domain-specific (e.g., I CS-Cert for industrial control systems). Red team operators often use these frameworks while learning the trade (often via pen testing) before developing custom tactics and techniques for the specific attack surfaces they will encounter in various organizations.

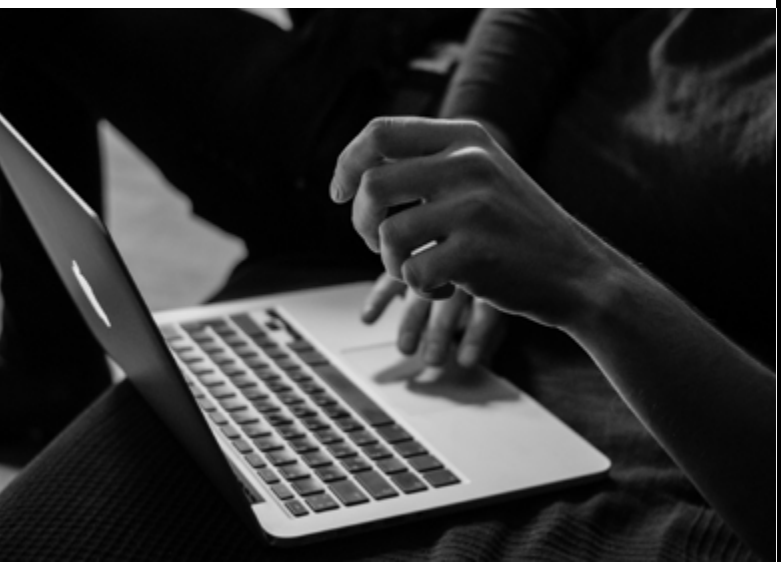

Red team engagements: Pros and cons

Red team assessments can provide a root-cause analysis of the risks in an organization. They probe deeper and more consistently than other security exercises, enabling continual testing instead of basic point-in-time measures. But setting up an adaptable red team is a challenge, even for well-resourced companies.

Benefits

By using a great red team, a company can simulate attacks from threat actors, patch up organization-wide root-cause issues, and stay one step ahead in the cybersecurity cat-and-mouse game. Red team assessments turn theory into reality by actually testing an organization's defenses to see where it is strong, weak, or exposed. The results of these assessments can inform a team's roadmap and help prioritize root-cause issues and risks. With effective iterative improvements or collaboration practices, red and blue (or purple) teams can level each other up, creating more advanced defenses and craftier attacks. Over time, this constant improvement in security posture reduces the risk of actual incidents.

Challenges

Companies struggle to find the right red team for their security posture. A common solution is to work with red team consultancies. The main problem with red team consulting is that it often relies on static red teams; they may not have the right skills for your specific attack surface. Traditional consultancies often lack the depth and breadth of skills needed for each company. Boutique firms can go deep in one area but can be expensive and slow. The consulting business model also means red team operators often work on projects back to back for years on end, leading to exhaustion or burnout. Lastly, external consulting teams can't always help companies fix their security holes after they've been discovered.

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.

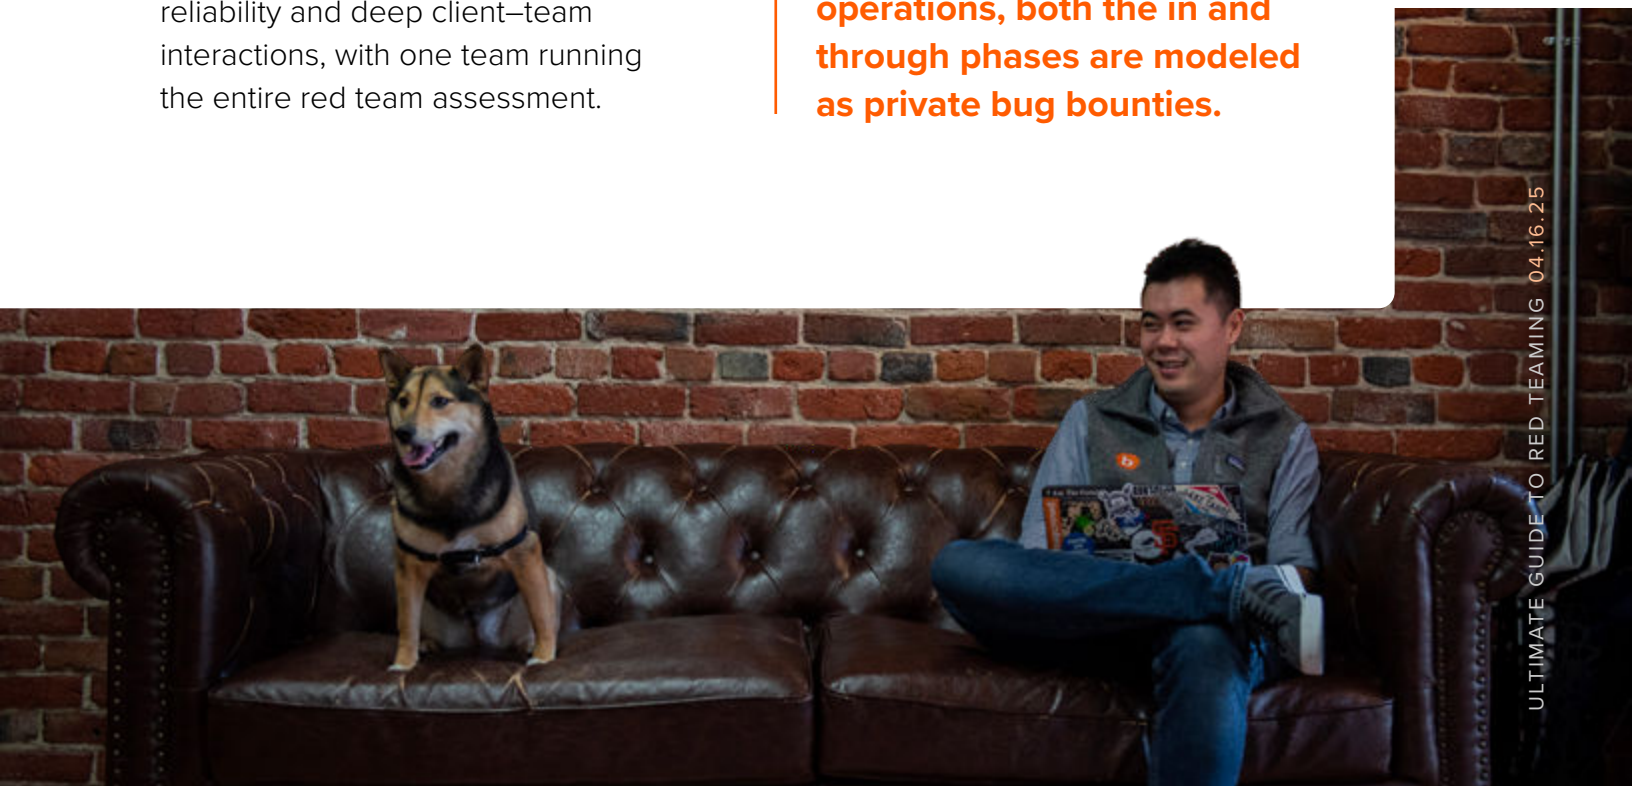

Small teams of vetted operators continually rotate through and attack systems, each bringing a different set of attack vectors. This option provides continuous coverage and replicates the most realistic attacks. With these options, companies can choose the model that works for their budget and security posture.

RTaaS 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.

bugcrowd

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