BISHOPFOX

EBOOK

THE GUIDE TO GETTING RED TEAMING RIGHT

Table of Contents

Introduction

No matter your organization's security goals, challenges, and level of maturity - there is no single, more effective way of proving whether your organization is susceptible to today's threats or the effectiveness of your security controls than to emulate real-world attacks and see how your program stands up against an attacker's tactics and techniques. Whether the question is "Can an attacker deploy ransomware in my environment?" or "Can a rogue employee exfiltrate all of my customer data without detection?", the only definitive way to answer these types of questions is to simulate the attack and see how you stack up.

As attack surfaces continue to expand and threat actors conduct increasingly more daring and sophisticated attacks, Red Team testing enables organizations to strengthen security defenses by highlighting the state of security readiness (or lack thereof) and finding a path forward to combat current and emerging cyber threats while achieving meaningful risk reduction.

Red Teaming emulates a range of advanced adversaries and real-world attacks to challenge defenses against network-based, physical, and social exploits. Performed as realistically as possible, Red Team operations take the same steps an adversary would in achieving their objective, which often focuses on chaining vulnerabilities and weaknesses together to exploit and reach the overall attack objective.

In this eBook, we will explore:

- The strategy behind Red Teaming and its execution as part of a holistic cybersecurity program
- How simulations, purple teaming, and continuous testing can strengthen an organization's security posture

Red Teaming: An Overview

Red Teaming is a specialized form of objective-based security testing that simulates real-world threat scenarios and mimics the strategies and tactics conducted by a range of adversaries. Red Teaming begins with attack objectives or trophies and then attempts to define and produce many different pathways to achieving those attack objectives through not only network-based tactics but also social and physical paths as well.

What better way to test the capabilities of your Blue Team than by putting them to the ultimate test with a Red Team sparring partner executing real scenarios against established defenses to help you gain valuable insights? For security teams wondering if the expense and rigor of Red Teaming is worth the ROI, the answer is yes. You gain the ground truth on security gaps and vulnerability exposures, objectively and accurately evaluate existing defenses and reaction plans, and most importantly, strengthen overall security posture for the future.

Red Team operations are typically time and resource intensive efforts which can require weeks to months to perform. As such, it is vital to make optimal use of time and resources by ensuring that the threats being incorporated into your Red Teaming exercise are closely aligned with your organization's current threat landscape and unique business/technical landscape. To be successful, Red Team investments must have a clearly defined and documented purpose that is shared with key stakeholders from the beginning.

Transparency amongst participants and decision makers is crucial to building a successful foundation that supports Red Team operations. Ahead of any exercise, it is important to talk with executives and stakeholders about what Red Teaming entails and, of equal importance, what it doesn't entail. Red Teaming cannot exist in a vacuum, so it is useful to keep open lines of communication across the target organization and to get cross-functional support and involvement, including stakeholders from legal, physical security, risk, and other business lines. For example, physical security teams should receive advance notice of any activity that could result in law enforcement involvement.

Red Teaming vs. Pen Testing

Offensive security solutions come in many shapes and sizes, so it is important that Red Team programs are differentiated from other security solutions, particularly when communicating with stakeholders.

Specifically, distinguishing Red Team operations from penetration tests is key to ensuring that approaches and methodologies applied produce the output and value that the organization is anticipating from testing efforts. While penetration testing is a popular and very well-known offensive security solution, it aims to evaluate environments, products, websites, internal tools, and backend systems supporting a tool, and may be required to meet PCI and other compliance regulations.

As penetration tests do not typically test defender capabilities or require stealthy approaches, they are usually completed within shorter time frames, such as two to three weeks.

Red Teaming, on the other hand, mimics specific threat actors and associated tactics, techniques, and procedures (TTPs) across the organization's entire production environment. Red Team operations typically apply an open scope of potential attack paths across networks, physical access, and social media. Red Teaming also distinguishes itself from other offensive security solutions because it tests the defenders ability to detect and stop attacks, which is outside the scope of traditional pen tests.

Red Teaming is also unique in that it relies heavily on full scope threat intelligence. While many security services and solutions utilize threat intelligence, it is primarily for information focused on malware families or vulnerabilities. Red Teaming, needs threat intelligence to understand the adversary's motivations and TTPs which are critical to performing emulation. How else would Red Teams know how to operate like an adversary?

For example, a typical Red Team engagement may act like a state-sponsored cyber espionage group that is attempting to compromise privately held source code or replicate a specific type of ransomware attack tied to a high-profile cybercriminal group. State-sponsored espionage campaigns in particular are highly complex and often have a long duration to support intelligence collection programs. Therefore, Red Team campaigns replicating this behavior will be conducted over several iterations.

PENETRATION TESTING VS. RED TEAMING

• Find and report on all exploitable vulnerabilities under controlled circumstances.	GOAL	• Provide a holistic and accurate picture of resilience to breach scenarios. Train and exercise defenders. Identify security gaps and test assumptions.
• Any/all vulnerabilities and misconfigurations within the scoped environment.	IDENTIFIES	• Vulnerabilities that relate to defined attacker objectives. Blind spots in security readiness and shortcomings of detective security controls.
• Specific application or aspect of attack surface (web application, mobile application, SDK, etc.).	SCOPE	• Broad attack surface, including people, places, and technologies (network/web) and the seams between them.
• What vulnerabilities are there?	ASKS	• How would we fare against a specific scenario? • How could an adversary achieve a particular objective?
• By request • Part of internal software development lifecycle, compliance, or other processes	SOURCE	• Self-driven, based on threat landscape • Tied to cyber threat intelligence and ITP strategic profiling • By request from leadership and lines of business

Types of Teams

We've briefly mentioned the concept of a Blue Team. But what does a Blue Team do, and how does it differ from a Red Team? And what's the purpose of a Purple Team?

Blue Teams are an organization's defenders and consists of individuals and groups who detect and respond to security incidents. Blue Teams identify when potentially malicious activities are occurring, investigate the behaviors, and act to contain them to limit their impact. Red Team successes can be valuable learning opportunities for Blue Teams and can gauge a Blue Team's current level of effectiveness along with key areas for improvement.

Red Teams are hackers authorized to emulate adversarial attacks and test organizations' defenses by identifying vulnerabilities and launching attacks in a controlled environment. Red Teams can mimic specific threat actors and their associated TTPs, providing the ultimate 'sanity check' for any attack scenario like ransomware, DDos attacks, or supply chain risks.

Purple Teams are a collaborative effort between both Red and Blue Teams, where both teams work closely together to strengthen an organization's ability to detect and respond to malicious behaviors. During a Purple Team engagement, communication and information sharing occurs regularly between Red and Blue Teams, in an effort to achieve a common goal of improving security monitoring and alerting."

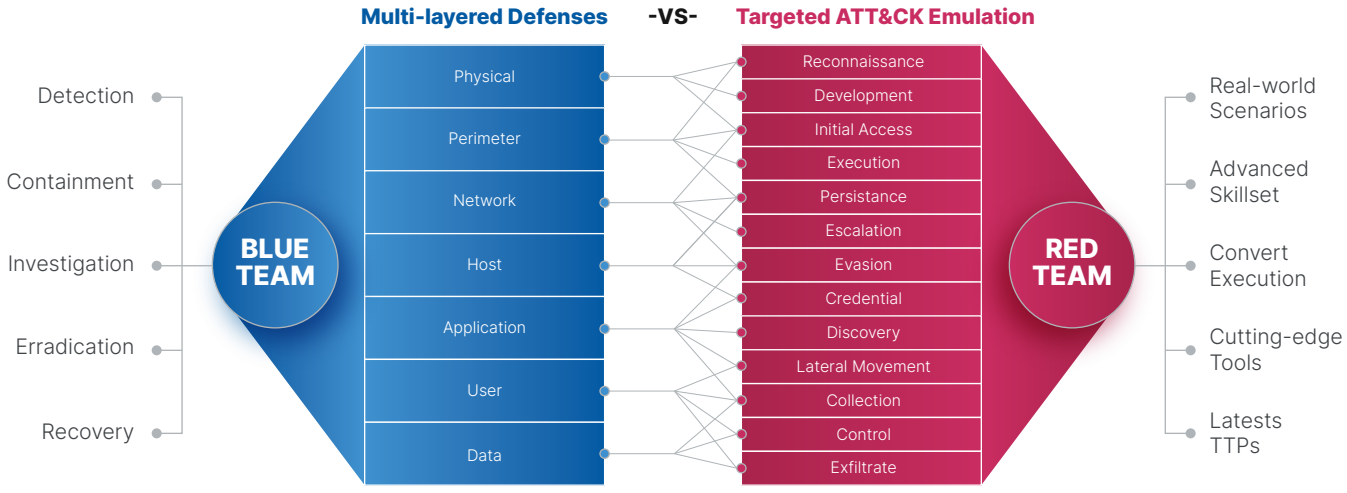

How Does Red Teaming Work?

Red Teaming is a cyberattack emulation operation in which the Red Team is allowed to leverage a variety of methods to capture a “trophy.” This trophy is previously identified by the organization engaging in a Red Team exercise, and it could be anything an organization deems of importance such as personally identifiable information (PII), persistent network access, or administrator credentials.

After one or more trophies are determined by the targeted organization, the Red Team plans and carries out various attack methods that an adversary would leverage to capture the trophy. It’s also important to note that organizations should determine in advance the level of knowledge the Red Team is provided about the target environment. It can be a zero-knowledge engagement (black box), or consultants may have some level of pre-existing knowledge of network and system architecture (gray box).

STAGES OF A RED TEAM ENGAGEMENT

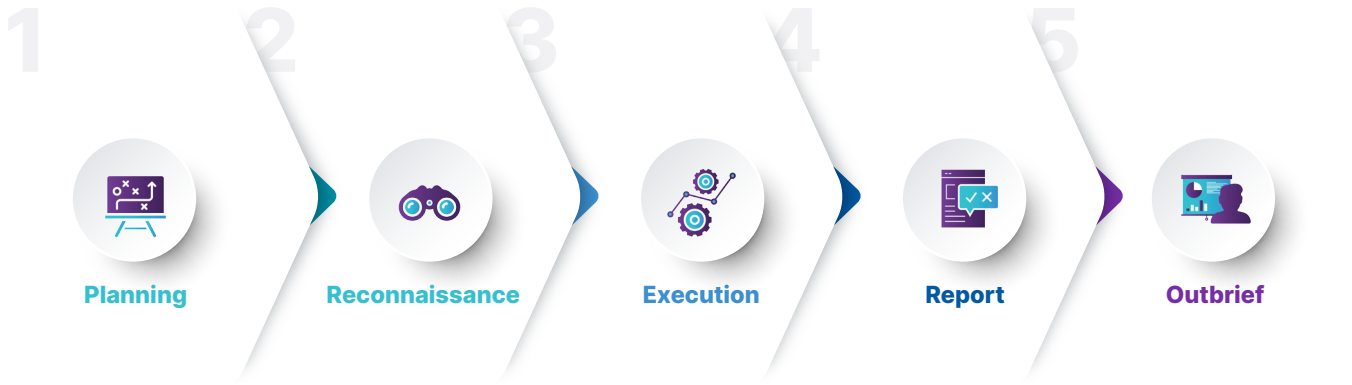

Planning

Meticulous planning of Red Team engagements ensures that exercises are performed in a way that provides maximum value while limiting various risks to an organization. A key output from planning should be clearly defined objectives and rules of engagement (particularly around processes for deconfliction throughout the exercise) to be followed throughout the operation. Consider limiting the number of individuals that have awareness of Red Team operations to those with a ‘need-to-know’ status in order to maintain a sense of realism throughout the engagement.

Planning is done to ensure that the exercise is set up to be successful and provides value. In addition to what is mentioned within this section, Red Team operation planning also includes determining objectives (flags), scope, roles and responsibilities, and key timelines for operation execution.

During the planning phase, the Red Team can utilize sources of threat intelligence such as The MITRE ATT&CK Navigator to strengthen operational realism and tailor attacks to current threat actors. Threat actor profiles should be incorporated into the operation proposal, so stakeholders understand the TTPs of the emulated threat group that the Red Team will mimic.

Reconnaissance & Execution

During the reconnaissance phase, Red Teams will enumerate the target's network architecture, domain and IP space, applications and technologies, email address format and other technical information. Additionally the Red Team will seek to uncover useful information about the target's processes and organizational culture to aid in tailoring attacks in subsequent phases. Reconnaissance is initially done as stealthily as possible, leveraging sources of information that will not alert internal Blue Teams, such as search engines, job postings, social networking sites and compromised credential sites.

Following the more passive and stealthy reconnaissance, the Red Team will begin more actively targeting in the execution phase and probe the target's applications, systems and networks to uncover more detailed information. This can include directly connecting to a target's systems and applications, interacting with them and even some limited network scanning and website spidering; although, these types of activities are typically done in a stealthy and careful manner to reduce the likelihood of detection by defenders.

Report & Outbrief

When Red Team operations conclude, a full scope report is prepared. The sections of the report may vary but should incorporate aspects like an executive summary, methodology/framework, planning, threat intelligence, attack narrative, analysis and response, and findings that include people, processes, and technology. The report should include both tactical and strategic recommendations.

Using standardized risk definitions and formally tracking findings in risk registers is vital for illustrating larger trends and patterns within the organization and for driving actual change from the work of Red Teams.

While the successful capture of the trophy is beneficial because it calls attention to an area of focus, your organization's real prerogative may be to test your defender's capabilities or determine how effective your security controls are when faced with a viable threat. The final deliverables illustrate the full value of the Red Team exercise.

Approaches & Methodologies

There are different varieties of Red Teaming security assessments, including breach simulations, Purple Teaming, and continuous testing. No matter what type of Red Team assessment methodology is being followed, the efforts should be based on a specifically-defined adversary and their known tactics, techniques, and procedures.

External Breach & Assumed Breach Simulations

Breach simulations are patterned after attack behaviors used by real-world adversaries in actual attacks observed in the wild, and testing is conducted in a safe but realistic environment. The Red Team provider should help determine which attack starting point (external or internal) will be most beneficial to test your security team.

EXTERNAL BREACH

An external breach places a Red Team fully outside a company's network to see if they can find a way in. Red Teams rely on their research of the targeted organization to find vulnerabilities in things like websites, assets, systems and applications which could enable an attacker to gain an initial foothold.

ASSUMED BREACH

An assumed breach, on the other hand, uses a scenario where the threat actor already has gained a foothold in the network. This is a more likely scenario to emulate if your organization is concerned about insider threats or social engineering schemes that capitalize on unwitting employees becoming the weak link giving an attacker an open invitation. An assumed breach will map the attack paths and exploit chains that allow an adversary to escalate privileges and steal data from your internal servers and databases.

In addition to the starting point, the level of knowledge and access granted to the Red Team can affect the accuracy of the engagement as well as its speed, efficiency, and coverage. As previously mentioned, there are three levels of knowledge to consider: black box (zero knowledge), white box (full knowledge), and gray box (a mix of both).

BLACK BOX

In a black-box engagement, the Red Team's role is that of a typical hacker, with no prior knowledge of the target system. The team isn't provided with any architecture diagrams or source code that is not publicly available. The limited knowledge makes black-box engagements the quickest to run, since the duration of the assignment largely depends on the team's ability to locate and exploit vulnerabilities in the target's outward-facing services.

WHITE BOX

White-box testing is the opposite of black-box testing, in which the Red Team is given full access to source code, architecture documentation, and so forth. The main challenge is sifting through the massive amount of data available to identify potential points of weakness, which can make white-box testing time-consuming in its own way. Red Teams with full knowledge can provide a comprehensive assessment of both internal and external vulnerabilities, making it the best choice for a holistic perspective of any security weaknesses. The close relationship between Red Teams and systems administrators/architects provides a high level of system knowledge but may affect the team's behaviors, since they operate based on knowledge not available to hackers.

GRAY BOX

A healthy mix of both zero knowledge and full knowledge can be found in a gray-box approach. While a black-box vantage examines a system from an outsider's perspective, a gray-box vantage provides access and knowledge of a user with some level of privilege to a system or network. The purpose is to provide a more focused and efficient assessment of a network's security than a black-box engagement. Using the design documentation for a network, Red Teams can focus their assessment efforts on the systems with the greatest risk and value from the start, rather than spending time determining this information on their own. An internal account on the system also allows testing of security inside the hardened perimeter and simulates an attacker with longer-term access to the network.

PRO-TIP

If you think you're up for the challenge, engage in one of these Red Team options. After performing a threat-based simulation, you'll have a more pragmatic sense of your preparedness for such a situation.

Purple Teaming

While many organizations may not have the budget and capacity to employ a full-time, internal Purple Team, outsourced Red Teams can work collaboratively with your Blue Team to perform Purple Teaming exercises. In this scenario, the Red Team attacks are coordinated with the Blue Team to determine what is and is not being detected and adjust defensive controls accordingly. As each aspect of the attack is performed, adjustments can be made on the fly by network defenders to modify and integrate defensive technologies to provide increased detection and prevention of offensive actions. Unlike Red Team scenarios, Purple Team operations are not trophy based. They also focus entirely on monitoring and detection; there is no aspect of looking for vulnerabilities for a Purple Team.

KEY SECURITY BENEFITS OF A PURPLE TEAMING EXERCISE

ENHANCED SECURITY KNOWLEDGE

Being able to observe and participate in attacks gives your Blue Team a better understanding of how attackers operate, enabling them to more effectively employ technologies to deceive actual attackers and study their TTPs.

BOOST PERFORMANCE WITHOUT BUDGET INCREASE

Combining defense and offense through Purple Teaming allows organizations to improve security monitoring functions faster and at less cost in comparison to other security solutions and/or hiring in-house talent.

STREAMLINING SECURITY IMPROVEMENTS

An alternative approach within the security industry is to view purple teaming as a conceptual framework that runs throughout an organization. This can nurture a collaborative culture that promotes continuous cybersecurity improvement.

GAIN CRITICAL INSIGHT

Purple Teaming gives your internal security team a critical understanding of gaps in your security posture and helps to identify areas for capability enhancement.

The outcome of this exercise is enhanced configuration and integration of existing defensive controls to maximize their potential to defend your network. Additionally, gaps may be identified where your existing technologies, people, or processes are not able to prevent or detect potentially dangerous techniques that are commonly used by adversaries. This empowers decision-makers to understand the need and benefit of additional defensive controls and make informed decisions when prioritizing budgetary requirements.

Continuous Red Teaming

Continuous Red Teaming allows organizations to establish a regular cadence of breach scenarios throughout the year, potentially targeting different parts of the attack surface. This gives teams regular practice and a more comprehensive, holistic look at their security posture.

The types of operations to be performed and their cadence can be designed to meet specific objectives. As an example, a continuous Red Team engagement could include a combination of biannual Red Team engagements and quarterly Purple Team assessments, or it could include quarterly Red Team breach scenarios and monthly Purple Team assessments.

Continuous testing with a third-party partner allows an organization to build a close and trusted partnership with an external firm who understands the environment and provides highly customized Red Team services that are most closely aligned to what an organization would have if they employed their own internal Red Team.

Tabletop Exercises

Tabletop exercises involve reviewing what your most immediate and dangerous threats are — and role-playing how you would react if those threats manifested into reality. All relevant stakeholders play a part in this exercise.

POPULAR SCENARIOS

- Your company is concerned about a ransomware attack wreaking havoc on your organization.
- You are worried that a potential breach of one of your third-party suppliers could lead to compromise of your own organization.
- You lose sleep over a denial-of-service attack bringing productivity to a halt.

A tabletop exercise gives you the opportunity to map out one or more attack scenarios, as well as how you would handle such a situation. Plan for unexpected issues when conducting a tabletop exercise – you may have to consider panicking executives, the fatigue of overworked employees, or possibly news breaking about the attack (if one has already occurred).

Results & Reporting: What To Expect

Unlike a penetration test, the end report is the central deliverable of a Red Team exercise. Within that report the list of identified findings isn't the most critical piece; the most valuable information (and where you should focus attention) is the attack narrative that outlines each scenario attempted and how the target's defenses fared against it.

REPORT COMPONENTS

- **Detailed Attack Graphing**
  In-depth attack graphing to chart possible paths of attack, including analysis of architecture, vulnerable systems, and data at risk.
- **Severity Scoring**
  The potential impact of defensive gaps using a proprietary scoring method based on real-world observations and industry-standard methodologies such as OWASP and CVSS.
- **Attack Timeline and Execution Pathway Summaries**
  Timeframe of events with detailed breakdown of actions performed, defensive performance, and achievement against target objectives.
- **Detailed Findings Presentation and Reporting**
  A complete walkthrough of findings, with a live question and answer session, ensuring all stakeholders understand technical findings, risks, and recommendations.

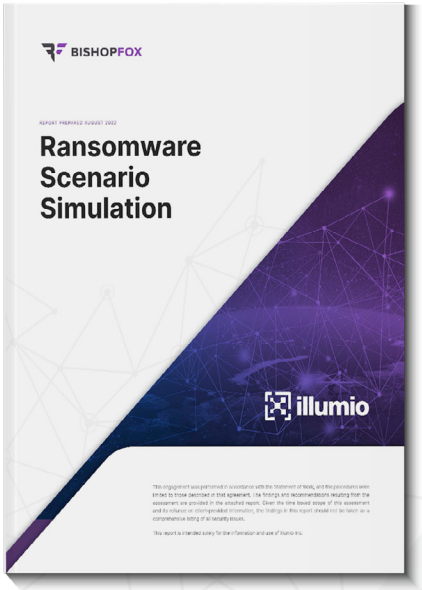

Are You Ready for Red Teaming?

If you are ready to embark on Red Teaming to proactively improve your security posture, you have several aspects to consider:

THREAT INTELLIGENCE

Invest time in understanding the current threat landscape that's impactful to your organization, industry, etc. While the Red Team that your organization hires will be deep in the weeds of threat intelligence, it will also help your own security teams to keep a pulse on the latest threat groups and tactics that are hitting your industry or geography.

INTERNAL SECURITY MATURITY

Red Teaming engagements are recommended for organizations with sophisticated cybersecurity programs. If you have already implemented ongoing employee training, social engineering activities, and penetration testing, then you are likely ready to challenge your security controls in a more life-like situation and are prepared for a Red Team engagement.

FINDING THE RIGHT PROVIDER

Red Teaming is a big commitment for any organization, so it is imperative to find the right partnership to reach your security goals. Partnering with an experienced third-party vendor is another option that provides you with a true outsider view of your infrastructure — most closely resembling an attacker's perspective. A third party will take the blinders off your security efforts and reveal weaknesses that may be missed by unintentionally biased insiders. Given the risks involved in this type of testing, it is important to select a vendor that is a good match for your company's specific threat landscape and risk tolerance.

Key Questions to Ask Red Team Vendors

01 "How will you customize the scenarios you're testing? How much control can we have over the scenarios?"

The vendor should be able to mimic the tools, techniques, and procedures (TTPs) and threat actors that you or your threat intelligence provider have been tracking. Additionally, confirm that the Red Team vendor's toolset can be customized to seed relevant indicators of compromise (IOCs) from the custom implants they use. Both will help to create as realistic an assessment of your organization's susceptibility to a cyberattack as possible.

02 "How do you plan to minimize the risk of production downtime?"

A Red Team's test plan should demonstrate a high degree of project management sophistication, including continuous communication of testing activities to "in the know" company stakeholders, advance notice of potentially damaging activities, and the demonstrated ability to pull the plug on activities that may have or are having noticeable negative impacts.

A Red Team vendor should establish priorities and reasonable scope in terms of what constitutes as evidence that goals have been achieved to avoid unnecessary operational impacts. For example, instead of shutting down a critical system to prove they can, it would be sufficient for the Red Team to demonstrate that they have gained the access needed to do so. These types of safeguards can help to dramatically decrease the operational risk of Red Team testing.

03 "How will the communications of the command and control (C&C) implant be assured? How will information that may be exfiltrated via the implant be protected?"

Red Team vendors should be able to provide examples of detailed protective countermeasures they use, such as restricting access to C&C channels to client IP ranges or encrypting and signing all communications resulting from the exploit, including any exfiltrated data. A vendor should also utilize measures to prevent the implant from targeting people and environments that are out of scope, such as staff members or third-party vendors who happen to be on your network. They must also be able to explain how the implants will behave after the completion of testing. Implants that self-destruct are an effective protection measure for these scenarios.

Summary

Any Red Teaming method you choose will provide you with greater visibility into where your organization needs to devote more resources to improving security. Red Teaming is by far the most effective and holistic method of emulating how an attacker would likely compromise your organization, and it brings to light issues you wouldn't have otherwise considered and highlights tools and processes that are both effective and ineffective, enabling you to streamline investments. By defining where you are struggling most, your organization can enhance detection and response, and decrease the risk of an attacker finding a known vulnerability.

Over the past 30 years, Red Teaming has gone from being an esoteric military practice to a sophisticated and highly valuable security discipline.

We believe Red Teams can deliver even more value to customers by integrating Red Teaming, risk analysis, and threat modeling into a comprehensive program. These three are natural complements, and bringing them together can help organizations more clearly understand and manage their operational security risks across the full attack surface.

About Bishop Fox

Bishop Fox is recognized as the leading authority in offensive security, providing solutions ranging from continuous penetration testing, red teaming, and attack surface management to product, cloud, and application security assessments.

Over the past 16 years, we've worked with more than 25% of the Fortune 100, 8 of the top 10 global tech companies, and hundreds of other organizations to improve their security. Our award-winning Cosmos platform was named **Best Emerging Technology** in the 2021 SC Media Awards and our offerings are consistently ranked as "world class" in customer experience surveys.

Security isn't just a job to us. We do this because we love it — and because we're committed to the common good. We've been actively contributing to and supporting the security community for almost two decades and have published more than 16 open-source tools and 50 security advisories in the last five years.

Cosmos

Cosmos proactively defends dynamic attack surfaces by combining advanced technology, automation, and expert-driven testing to continuously identify and remediate high-risk exposures before attackers even know they exist.

Leveraging a proprietary asset discovery and exposure reconnaissance engine, Cosmos continuously discovers and maps your ever-changing attack surface and identifies dangerous vulnerabilities targeted by attackers.

Acting as an extension of your security team, our operators provide deep insights into findings, deliver real-time answers to pressing questions, and conduct on-demand retesting to validate remediation procedures and accelerate the closure of attack windows.

Consulting Services

Red Teaming

We utilize advanced offensive tools and tactics that mimic real-world adversaries to identify exploitable weaknesses in your organization while stress testing your incident responders and their playbooks for handling active, persistent attackers.

Ransomware Readiness

Putting your ransomware playbook to the ultimate test, Bishop Fox's security experts covertly execute ransomware attacks to measure the efficacy and communication of your security team (and other teams at large) to determine your true readiness against skilled, real-world hackers.

Network Testing

To safeguard your critical infrastructure, we locate vulnerabilities, attack paths, and exploit chains that internal and external threat actors could leverage to gain access to sensitive data and systems.

CONNECT WITH US

Get started today.

Are you ready to start "defending forward"? Get in touch with our offensive security experts today to explore solutions that meet your unique business needs.

Request a Meeting **Explore Cosmos**

8240 S. Kyrene Rd. • Tempe, AZ 85284
480.621.8967
hello@bishopfox.com • bishopfox.com