FEATURE

Red Teaming for Cybersecurity

A Practical Approach

Red teaming has been a buzzword in the cybersecurity industry for the past few years. This idea has gained even more traction in the financial sector as more and more central banks want to complement their audit-based supervision with a more hands-on and fact-driven mechanism. There is a practical approach toward red teaming that can be used by any chief information security officer (CISO) as an input to conceptualize a successful red teaming initiative.

Understanding Red Teaming

Red teaming is the process of providing a fact-driven adversary perspective as an input to solving or addressing a problem.[1] For instance, red teaming in the financial control space can be seen as an exercise in which yearly spending projections are challenged based on the costs accrued in the first two quarters of the year. In the cybersecurity context, red teaming has emerged as a best practice wherein the cyberresilience of an organization is challenged by an adversary's or a threat actor's perspective.

This is a powerful means of providing the CISO a fact-based assessment of an organization's security ecosystem. Such an assessment is performed by a specialized and carefully constituted team and covers people, process and technology areas. As a result, CISOs can get a clear understanding of how much of the organization's security budget is actually translated into a concrete cyberdefense and what areas need more attention. A practical approach on how to set up and benefit from a red team in an enterprise context is explored herein.

Why Invest in Red Teaming?

An organization invests in cybersecurity to keep its business safe from malicious threat agents. These threat agents find ways to get past the enterprise's security defense and achieve their goals. A successful attack of this sort is usually classified as a security incident, and damage or loss to an organization's information assets is classified as a security breach. While most security budgets of modern-day enterprises are focused on preventive and detective measures to manage incidents and avoid breaches, the effectiveness of such investments is not always clearly measured. Security governance translated into policies may or may not have the same intended effect on the organization's cybersecurity posture when practically implemented using operational people, process and technology means. In most large organizations, the personnel who lay down policies and standards are not the ones who bring them into effect using processes and technology. This contributes to an inherent gap between the intended baseline and the actual effect policies and standards have on the enterprise's security posture. Cyberthreats are constantly evolving, and threat agents are finding new ways to manifest new

Seemant Sehgal, CISA, CISM, CCNA, CEH, CIW-Security Analyst, ISO 27001 LI, ITIL SOA, PPO, PRINCE2, SABSA, SANS GSNA
Has been engaged with setting up vulnerability management functions for tech and financial sectors in Europe and the United States including leading a global red team for ING Group. His areas of expertise include cyberresilience, payment security (the Revised Payment Service Directive [PSD2], the Payment Card Industry Data Security Standard [PCI DSS]), International Organization for Standardization (ISO)/International Electrotechnical Commission (IEC) ISO/IEC 27001, cyberdefense and system and organization controls (SOC). He is an ardent promoter of the red team philosophy and is a key contributor/supporter of the Threat Intelligence Based Ethical Red teaming (TIBER) initiative.

security breaches. This dynamic clearly establishes that the threat agents are either exploiting a gap in the implementation of the enterprise's intended security baseline or taking advantage of the fact that the enterprise's intended security baseline itself is either outdated or ineffective. This leads to the question: How can one get the required level of assurance if the enterprise's security baseline insufficiently addresses the evolving threat landscape? Also, once addressed, are there any gaps in its practical implementation? This is where red teaming provides a CISO with fact-based assurance in the context of the active cyberthreat landscape in which they operate. Compared to the huge investments enterprises make in standard preventive and detective measures, a red team can help get more out of such investments with a fraction of the same budget spent on these assessments.

Setting Up a Red Team

All organizations are faced with two main choices when setting up a red team. One is to set up an in-house red team and the second is to outsource the red team to get an independent perspective on the enterprise's cyberresilience.

Both approaches have upsides and downsides. While an internal red team can stay more focused on improvements based on the known gaps, an independent team can bring a fresh perspective.

> "THE BEST APPROACH, HOWEVER, IS TO USE A COMBINATION OF BOTH INTERNAL AND EXTERNAL RESOURCES."

The best approach, however, is to use a combination of both internal and external resources. More important, it is critical to identify the skill sets that will be required to make an effective red team. The types of skills a red team should possess and details on where to source them for the organization follows.

Technical Skills

Highly skilled penetration testers who practice evolving attack vectors as a day job are best positioned in this part of the team. Scripting and development skills are utilized frequently during the execution phase, and experience in these areas, in combination with penetration testing skills, is highly effective. It is acceptable to source these skills from external vendors who specialize in areas such as penetration testing or security research. The main rationale to support this decision is twofold. First, it may not be the enterprise's core business to nurture hacking skills as it requires a very diverse set of hands-on skills. Second, if the enterprise wishes to raise the bar by testing resilience against specific threats, it is best to leave the door open for sourcing these skills externally based on the specific threat against which the enterprise wishes to test its resilience. As an example, in the banking industry, the enterprise may want to perform a red team exercise to test the ecosystem around automated teller machine (ATM) security, where a specialized resource with relevant experience would be needed. In another scenario, an enterprise may need to test its Software as a Service (SaaS) solution, where cloud security experience would be critical.

Tactical Skills

This part of the team requires professionals with penetration testing, incidence response and auditing skills. They are able to develop red team scenarios and communicate with the business to understand the business impact of a security incident. This part of the red team does not have to be too big, but it is crucial to have at least one knowledgeable resource made accountable for this area. Additional skills can be temporarily sourced based on the area of the attack surface on which the enterprise is focused. This is an area where the internal security team can be augmented.

Strategic Skills

Professionals with a deep and practical understanding of core security concepts, the ability to communicate with chief executive officers (CEOs) and the ability to translate vision into reality are best positioned to lead the red team. The lead role is either taken up by the CISO or someone reporting into the CISO. This role covers the end-to-end life cycle of the exercise. This includes getting

sponsorship; scoping; picking the resources; approving scenarios; liaising with legal and compliance teams; managing risk during execution; making go/no-go decisions while dealing with critical vulnerabilities; and making sure that other C-level executives understand the objective, process and results of the red team exercise. Lastly, this role also ensures that the findings are translated into a sustainable improvement in the organization's security posture. Although its best to augment this role from the internal security team, the breadth of skills required to effectively dispense such a role is extremely scarce.

Scoping the Red Team

The most critical aspect of scoping a red team is targeting an ecosystem and not an individual system. Hence, there is no predefined scope other than pursuing a goal. The goal here refers to the end objective, which, when achieved, would translate into a critical security breach for the organization. People, process and technology aspects are all covered as a part of this pursuit. How the scope will be approached is something the red team will work out in the scenario analysis phase. It is imperative that the board is aware of both the scope and anticipated impact.

> "THE MOST CRITICAL ASPECT OF SCOPING A RED TEAM IS TARGETING AN ECOSYSTEM AND NOT AN INDIVIDUAL SYSTEM."

At this stage, it is also advisable to give the project a code name so that the activities can stay classified while still being discussable. Agreeing on a small group who will know about this activity is a good practice. The intent here is not to inadvertently alert the blue team and ensure that the simulated threat is as close as possible to a real-life incident. The blue team includes all personnel that either directly or indirectly respond to a security incident or support an organization's security defenses. In the present cybersecurity context, all personnel of an organization are targets and, therefore, are also responsible for defending against threats. The secrecy around the upcoming red team exercise helps maintain the element of surprise and also tests the organization's capability to handle such surprises. Having said that, it is a good practice to include one or two blue team personnel in the red team to promote learning and sharing of knowledge on both sides.

Developing Red Team Scenarios

Simply put, this step is stimulating blue team colleagues to think like hackers. The quality of the scenarios will decide the direction the team will take during the execution. In other words, scenarios will allow the team to bring sanity into the chaotic backdrop of the simulated security breach attempt within the organization. It also clarifies how the team will get to the end goal and what resources the enterprise would need to get there. That said, there needs to be a delicate balance between the macro-level view and articulating the detailed steps that the team may need to undertake. In most cases, the scenario that was decided upon at the start is not the eventual scenario executed. This is a good sign and shows that the red team experienced real-time defense from the blue team's perspective and was also creative enough to find new avenues. This also shows that the threat the enterprise wants to simulate is close to reality and takes the existing defense into context.

While brainstorming to come up with the latest scenarios is highly encouraged, attack trees are also a good mechanism to structure both discussions and the outcome of the scenario analysis process. To do this, the team may draw inspiration from the methods that have been used in the last 10 publicly known security breaches in the enterprise's industry or beyond. **Figure 1** is an example attack tree that is inspired by the Carbanak malware, which was made public in 2015 and is allegedly one of the biggest security breaches in banking history.

Figure 1—Carbanak Attack Tree

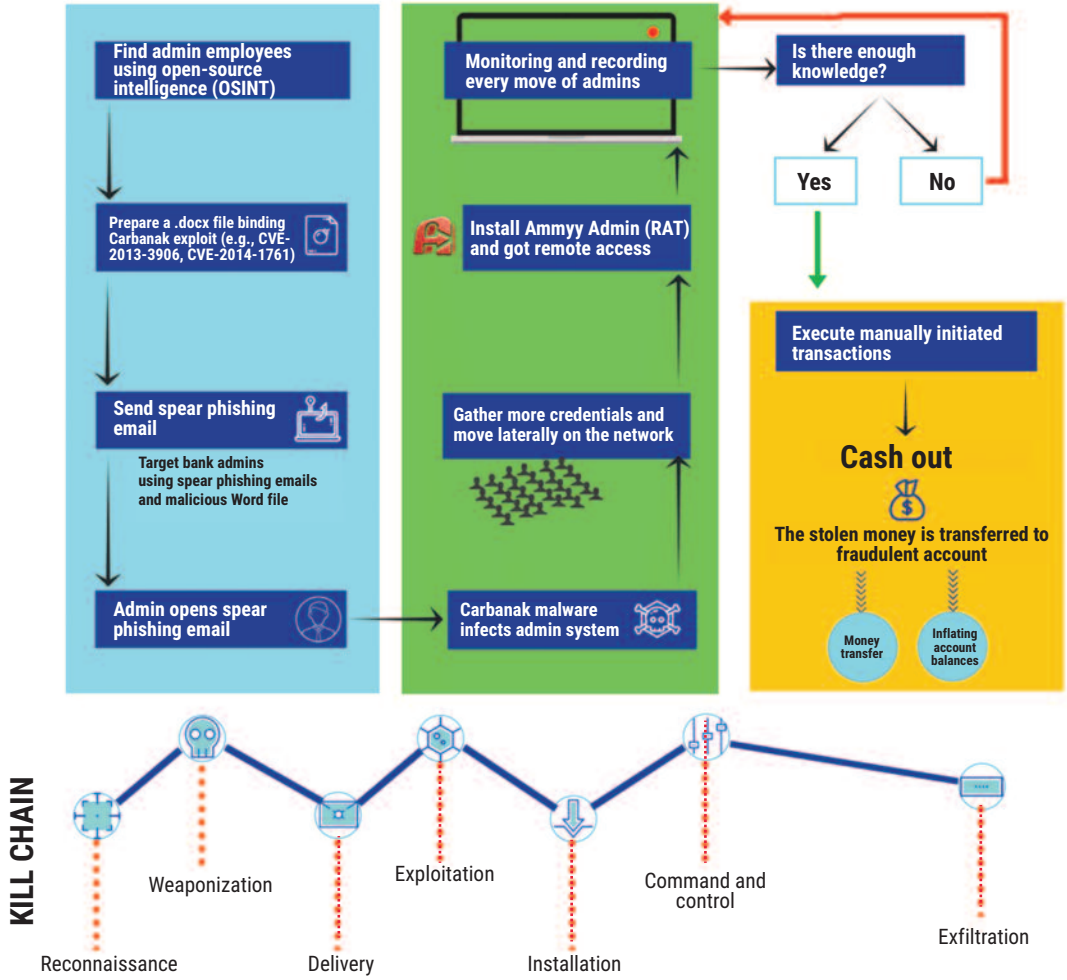

Execution Phase of a Red Team

This is perhaps the only phase that one cannot predict or prepare for in terms of events that will unfold once the team starts with the execution. By now, the enterprise has the required sponsorship, the target ecosystem is known, a team is set up, and the scenarios are defined and agreed upon. This is all the input that goes into the execution phase and, if the team did the steps leading up to execution correctly, it will be able to find its way through to the actual hack. The skill and experience of the people chosen for the team will decide how the surprises they encounter are navigated. Before the team begins, it is advisable that a "get out of jail card" is created for the testers. This artifact ensures the safety of the testers if encountered by resistance or legal prosecution by someone on the blue team. The get out of jail card is produced by the undercover attacker only as a last resort to prevent a counterproductive escalation.

Reporting Red Team Findings

Unlike a penetration test, the end report is not the central deliverable of a red team exercise. The report, which compiles the facts and evidence backing each fact, is certainly important; however, the storyline within which each fact is presented adds the required context to both the identified problem and suggested solution. A perfect way to find this balance would be to create three sets of reports.

The Storyline

The storyline describes how the scenarios played out. This includes the moments in time where the red team was stopped by an existing control, where an existing control was not effective and where the attacker had a free pass due to a nonexistent control. This is a highly visual document that shows the facts using pictures or videos so that executives are able to understand the context that would

otherwise be diluted in the text of a document. The visual approach to such storytelling can also be used to create additional scenarios as a demonstration (demo) that would not have made sense when testing the potentially adverse business impact. An example of such a demo would be the fact that a person is able to run a whoami command on a server and confirm that he or she has an elevated privilege level on a mission-critical server. However, it would create a much bigger impact on the board if the team can demonstrate a potential, but fake, visual where, instead of whoami, the team accesses the root directory and wipes out all data with one command. This will create a lasting impression on decision makers and shorten the time it takes to agree on an actual business impact of the finding.

Standard Report With Findings and Recommendations

The second report is a standard report very similar to a penetration testing report that records the findings, risk and recommendations in a structured format. This report is built for internal auditors, risk managers and colleagues who will be directly engaged in mitigating the identified findings.

Purple Teaming Report With Event and Technical Logs

The third report is the one that records all technical logs and event logs that can be used to reconstruct the attack pattern as it manifested. This report is a great input for a purple teaming exercise. Purple teaming is the process in which both the red team and blue team go through the sequence of events as they happened and try to document how both parties viewed the attack. This is a great opportunity to improve skills on both sides and also improve the cyberdefense of the organization.

Learn, Improve and Red Team Again

To learn and improve, it is important that both detection and response are measured from the blue team. Once that is done, a clear distinction between what is nonexistent and what needs to be improved further can be observed. This matrix can be used as a reference for future red teaming exercises to assess how the cyberresilience of the organization is improving. As an example, a matrix can be captured that measures the time it took for an employee to report a spear-phishing attack or the time taken by the computer emergency response team (CERT) to seize the asset from the user, establish the actual impact, contain the threat and execute all mitigating actions. These matrices can then be used to prove if the enterprise's investments in certain areas are paying off better than others based on the scores in subsequent red team exercises. **Figure 2** can be used as a quick reference card to visualize all phases and key activities of a red team.

Figure 2—Red Team Overview

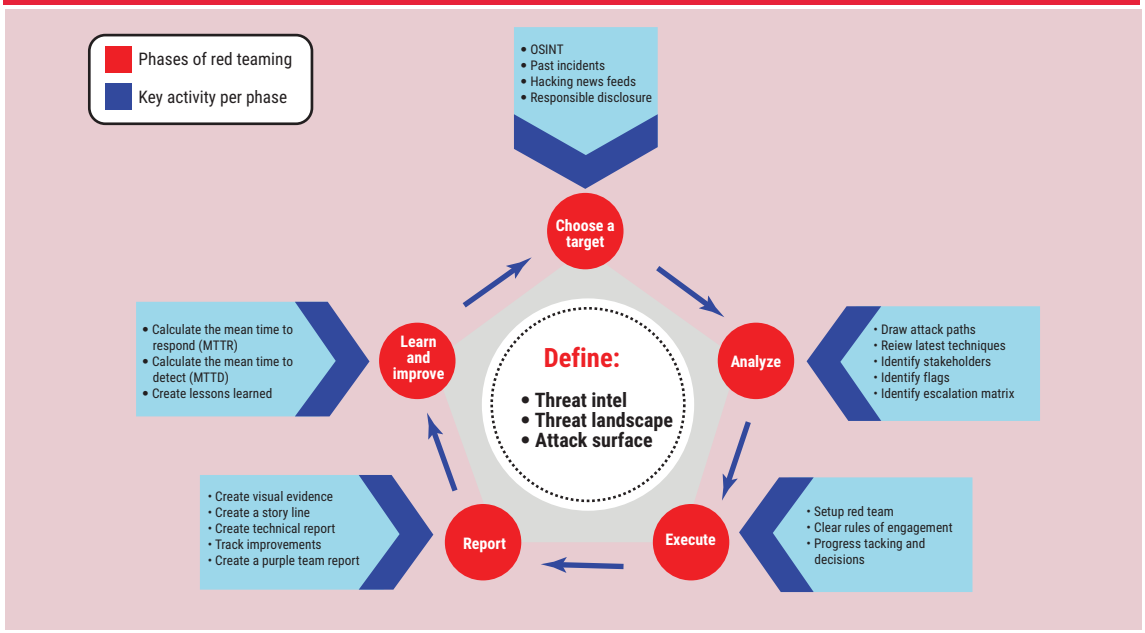

Conclusion

A crucial element in the setup of a red team is the overall framework that will be used to ensure a controlled execution with a focus on the agreed objective. The importance of a clear split and mix of skill sets that constitute a red team operation cannot be stressed enough. This covers strategic, tactical and technical execution. When used with the right sponsorship from the executive board and CISO of an enterprise, red teaming can be an extremely effective tool that can help constantly refresh cyberdefense priorities with a long-term strategy as a backdrop.

Endnotes

1 *Red Team Journal*, "Red Teaming and Alternative Analysis," *https://redteamjournal.com/red-teaming-and-alternative-analysis/*