A SANS Whitepaper

Red, Blue and Purple Teams: Combining Your Security Capabilities for the Best Outcome

Written by **Chris Dale**

October 2019

Sponsored by:

Core Security, a HelpSystems Company

Introduction

The terms Red Team and Blue Team are used to illustrate two different security teams within an organization working against each other. The goal of having adversarial teams is to further improve an organization's security posture. However, this approach is imperfect: It usually neglects fundamental flaws in how these teams work together. Instead of tight cooperation between the two teams, each side frequently misses out on opportunities which would benefit both.

With organizations unable to cover every base at defense, the offense continues to be as successful as ever at breaking in. If the offensive teams are *always* successful, it essentially means the goal of their job is not maximized—the goal being to make it harder for real attackers to break in. The same can be said about defense: If the team doesn't properly understand how attackers operate, it has few chances of architecting resilient infrastructure, which can hinder and limit attackers until they can be safely ejected from the networks.

The traditional adversarial relationship does not work; pitting the "Red Team" against the "Blue Team" is a thing of the past. We'll discuss how to intertwine the two units, creating a symbiotic relationship allowing for much better results for both teams, in what is often dubbed the "Purple Team" (see Figure 1).

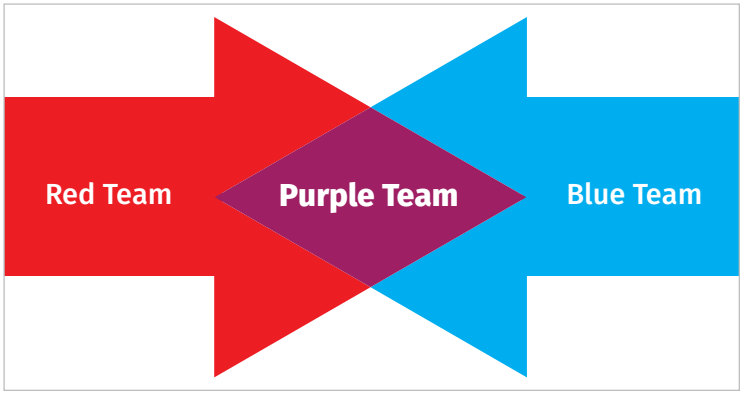

Figure 1. Red Team and Blue Team Come Together

Furthermore, automation plays a valuable role in today's security environments, with many Red Teams and Blue Teams often yielding improved results by accomplishing more with less. Because of this, we should strive to accomplish more through automation.

Red Team: Seeking Success

The Red Team is typically focused on breaking into the organization and demonstrating risk so the target organization can improve its security posture. If the team is successful in its efforts, it means the overall value and usefulness of the team is declining. It's important to realize the primary job of such a team is to make penetration testers' jobs harder, while all too often penetration testers pride themselves of breaking into organizations via trivial means.

This doesn't mean having a Red Team and penetration testing in general isn't a worthwhile investment. Penetration testing ranks high on the effectiveness of security controls and practices as noted in the SANS whitepaper "State of Application Security: Closing the Gap,"[1] however today's attacks warrant a fresh approach to improve not just the team's success in breaking in, but the corresponding defense side to make things harder in the long run—and essentially improve the organization's overall security, which has always been the intention of having teams.

Blue Team: Destined to Fail

The Blue Team is tasked with detecting adversaries and preventing them from breaking into the organization's infrastructure. Such a task is monumental: The boundaries of IT infrastructure for an organization are often undefined and ever-changing and the applications are tedious and hard to keep up-to-date—not to mention users who will always inadvertently provide attackers with a foothold into the organization. Another consideration: The organization faces the ever-imposing threat of zero-day exploits, which may be used against the infrastructure to provide an adversary access. Shadow IT, sometimes called unsanctioned IT, also adds for unmanageable infrastructure which may allow attackers ways into your infrastructure.[2] The Blue Team's challenge can be compared to being a goalie in a soccer match, except the goal post is ever changing and the criminals are using all kinds of balls to score goals.

Even with these challenges, the prevention stance is not futile. As more efforts are put into prevention, many attackers might be inclined to give up and move to easier targets. As we focus more attention into detecting compromises, hopefully we can thwart attackers before they can secure their objectives—causing the attackers to lose their investment of time and money while protecting the organization's objectives. This outcome means a win-win scenario, even if you were initially compromised.

---

[1] "State of Application Security: Closing the Gap," May 2015, www.sans.org/reading-room/whitepapers/analyst/2015-state-application-security-closing-gap-35942, page 17, Table 5.

[2] "Unsanctioned Business Unit IT Cloud Adoption Increases Risk of Data Breaches and Financial Liabilities," www.gartner.com/smarterwithgartner/dont-let-shadow-it-put-your-business-at-risk/

Purple Team: Joint Efforts Yield Higher Returns on Investment

In our current security climate, the traditional adversarial relationship does not work. Pitting the Red Team against the Blue Team is a thing of the past. A better path forward is to intertwine the two units to create a "Purple Team" to allow the two teams to cooperate in much greater detail and enable much better results. We present advice on how both teams respectively can look at improving the merge of the two teams.

How the Red Team Can Contribute to the Purple Team

Again, the task of the Red Team is to make its own life harder. Thus, if the Red Team can break in via simple means, for example through leaked credentials, CVEs (common vulnerabilities and exposures) or other low hanging fruit, it means the team can likely put security investments elsewhere instead of on a penetration test. So how can the Red Team still contribute and not lose out on revenue or the sense of purpose? Here are some approaches to allow early proof-of-value and start the engagement among teams, with an angle towards making things harder.

Let's start with scoping. Scoping a penetration test is not a simple task. The Blue Team does not necessarily know its own scope, nor does it know how attackers operate. As a penetration tester, it is also hard to scope only via information provided by the target, usually via scoping meetings. This is not ideal by any means, as the reconnaissance, discovery and scanning processes of a penetration test will often discover new attack surfaces. Changing scope after concluding a scope is often a bittersweet experience, as the engagement might incur further expenses or may force the penetration testers to work on a limited scope, whereas real adversaries would not be impacted by such limitations.

Instead, we suggest a much more convenient, cost effective and efficient approach. The penetration test is split into two deliveries. First the engagement starts with an easy-to-scope reconnaissance, discovery and scanning phase directed towards the target. This engagement has a much lower cost than a full penetration test would, allowing the different teams to build up trust with one another. Or, in the case of in-house penetration testers, this engagement allows the team to properly focus on doing a thorough job on reconnaissance, scanning and discovery—phases which are often considered to be the most valuable to penetration testers. Furthermore, this makes it more enjoyable for the penetration testers, as they won't necessarily feel they're testing on the wrong targets, as they've already provided their opinions on what they feel should be considered in scope. Figure 2 details the abstract phases of a penetration test and how it can be divided into two deliveries, creating a win-win scenario for both parties involved.

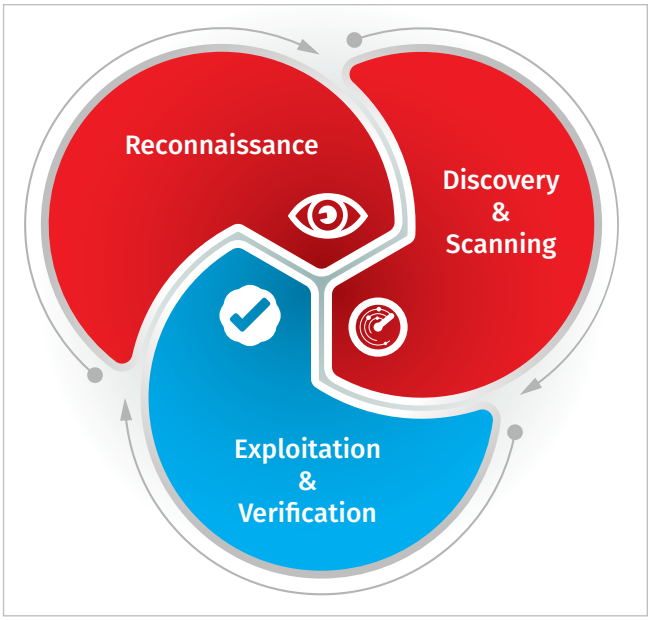

Figure 2. Cycles of the Penetration Test

The reconnaissance delivery can contain color-coded overviews of which assets the Red Team considers most important to include in the scope and can cover all available topics, which increases the attack surface. These include:

- Classic servers and applications, which are typical things a traditional penetration test would cover.
- Mobile applications and assets which are often overlooked by organizations.
- Users, preferably a listing of whom is in which business unit, for example, management, IT operations, developers, helpdesk and receptionists—Each one of these will have its own attack surface (i.e. management is often targeted via spear phishing fraud, IT is targeted via OSINT and helpdesk and receptionists might be vulnerable to password reset requests via phone calls).
- Leaked data, misconfigured servers and other vulnerabilities discovered without exploitation—These are low-hanging fruits, and ideally the target will fix these before the Red Team has its go against the infrastructure.

The team can propose a scope with a report detailing the identifiable attack surface the Red Team could discover. The Blue Team can then input any unidentified assets to the report, and the teams can both finally agree on a thorough and transparent scope containing the most important assets. In many cases, the reconnaissance report alone will contain enough details the Blue Team will want to address that this process will provide value.

The next thing the Red Team can do is revisit how it can broaden its penetration testing reports. These reports normally include an executive summary, an overview of findings and details regarding the different findings. When providing reports for the purpose of a Purple Team collaboration, the usual reports can be extended to increase more value for the Blue Team.

For example, the Red Team should strive to provide multiple suggestions for fixing the different vulnerabilities it found. These suggestions can be included in the executive summary so executives can see how risk can not only be removed but also mitigated in different ways. Additionally, with the Red Team's extensive knowledge of bypassing security, they can make further recommendations in the report on how to adequately detect and respond to the vulnerabilities identified. Ideally the report focuses not only on fixing vulnerabilities, but also on advising how the Blue Team can fix processes to more reliably and consistently manage such vulnerabilities in the future. Figure 3 displays how a traditional report can be extended and built upon to provide additional Blue Team value.

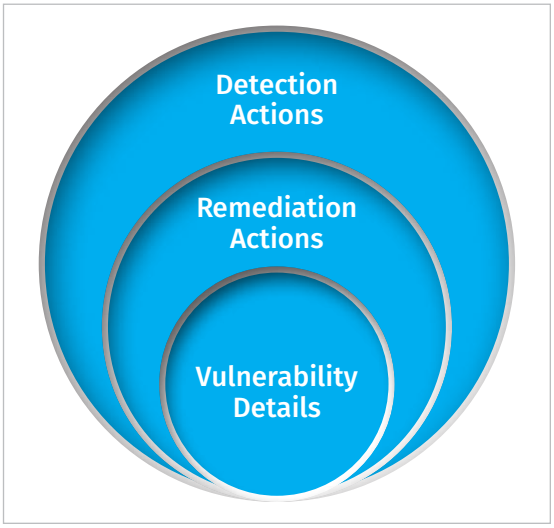

Figure 3. Layers of Red Team's Reports

Finally, the Red Team should consider different ways to present findings. Text editors are not necessarily the best format to use, as these do not promote ease of communications. While the Blue Team is doing triage, patching and verification, it ideally wants to have a way to follow up on individual items in a separate system. Furthermore, the Red Team wants to do continuous testing and assessments, and a portal to support this could allow for notifications and integration for the Blue Team.

How the Blue Team Can Contribute to the Purple Team

Collaboration is the key to building a successful Purple Team. Invite penetration testers to your collaboration platform and get the conversation started:

- Have an ongoing and easy discussion about the progress of the engagement.
- Use the platform to directly target individuals during a conversation thread, so each member of the team can interact quickly and effectively without having to stall out via emails and phone calls.
- Utilize Microsoft Teams, Slack and other communication platforms to easily invite third parties to directly interact with your team and vice versa.

Next, have the Red Team try to break individual pieces of code, such as libraries or the code for sanitization employed throughout the organization's applications. Instead of having the Red Team test on a wide scope in round one of the engagements, let it use its efforts on code which you believe should be safe and hardened. Furthermore, the Blue Team can focus its efforts on ensuring its software conforms to library requirements throughout the applications. Additionally, the Red Team can ensure IDEs (integrated development environments) report on unsafe programming logic which utilizes input from users or outputs data without the proper sanitization libraries.

Most Blue Teams are already running vulnerability scans, however Red Teams often break in via simple means. It's a good idea to either allow the Red Team to issue a preliminary vulnerability scan and provide an executive summary of this report or to feed the Red Team an existing scan and get its take on it before running further penetration testing activities. Often, we will see such early interactions provide higher value penetration tests and allow the Blue Team to better understand its weaknesses before committing to a test.

Likewise, before committing to a scope of testing, it's a good idea to consider giving the Red Team credentials to high risk assets which could have serious consequences when compromised. Often, applications will have more vulnerabilities post authentication. This exercise is especially valuable if the application is high risk, since it would identify potential threats in advance of the scenario where a user's credentials are compromised. If multi-factor authentication is enabled for the application, the risk decreases and one might consider focusing the testing scope on just the façade and other external assets. Figure 4 illustrates the integral nature of communication and collaboration to the evolution of the Purple team.

Figure 4. Communication and Collaboration in the Purple Team

Automation: A Necessary Focus for All Teams

If dev-ops has taught us anything, it is that automation is key to successfully battling threats. The Red Team can do breach simulations to improve its penetration testing processes, allowing the team to focus on important aspects of testing instead of focusing on tedious and repetitious tasks. Likewise, the Blue Team can take advantage of similar measures so it can focus on detection tactics, techniques and procedures (TTPs) used by the Red Team. JPCert has an overview of the logs created by Windows upon executing the tools most likely used by a network infiltrator,[3] and Thailand Cert shares an overview of the many different threat actors and the tools it uses.[4] Ahead of time, the Blue Team should ensure its networks are capable of preventing and detecting the execution of the Red Team TTPs, which can be done through automation. The Blue Team should also consider investing in efforts to automatically shun and isolate suspected threats in the network. Instead of simply disconnecting a host from the network, have the host become part of a private-VLAN while it's triaged and controlled by the Blue Team. The team can also capture and replay attacks across environments, ensuring conformity throughout the organization.

Be Cautious with Automation

When implementing automation, proceed with care. Consider how the automation process logs into devices, for example to automatically triage a compromised host. Also, be careful to avoid inadvertently sharing more information—such as credentials—to the attackers, potentially giving them further control of the environment. For example, a vulnerability scanner operating with domain administrator credentials might decide to scan an attacker-controlled host and thus the credentials might be sniffed and either cracked or the hashes passed.

Successful Integration Between Teams

In an example of a successful Purple Team engagement, the Blue Team used a collaboration platform to provide the Red Team with a screen sharing session to one of the potentially vulnerable systems. Within 10 minutes, the Red Team concluded a false positive which would've otherwise taken hours to figure out. Furthermore, the Blue Team discovered crashes in the organization's payment platform, unsure if it was caused by Red Team or others. By providing logs and discussions, together the teams were able to uncover a vulnerability against the availability of the platform. The platform was properly patched, fixing an issue which could've caused damages for years to come.

Conclusions

For increased success in defeating advanced adversaries and increasing the overall security posture of our organizations, the terms Blue and Red must merge under the umbrella Purple Team. The teams must stop working as only adversaries, and instead start collaborating and working in unison in the future. The potential is tremendous and has a very low bar to get started. Information Security has come far, but it is essential to keep developing and looking at new ways to further allow us to defend.

[3] "Tool Analysis Results Sheet," https://jpcertcc.github.io/ToolAnalysisResultSheet/

[4] "Threat Group Cards: A Threat Actor Encyclopedia," www.thaicert.or.th/downloads/files/A_Threat_Actor_Encyclopedia.pdf

About the Author

Chris Dale

SANS instructor Chris Dale teaches SANS SEC504: Hacker Tools, Techniques, Exploits, and Incident Handling. As head of the penetration testing and incident handling groups at Netsecurity, a company based in Norway, Chris brings significant security expertise and a background in system development, IT operations and security management. He is passionate about security, and he regularly gives presentations and teaches at conferences and workshops. Chris holds the GCIH, GPEN, GSLC and GMOB certifications and participates in panel debates and government-related working groups to recommend and improve security in the Norwegian private and public sectors.

Sponsor

SANS would like to thank this paper's sponsor: