Effective Purple Teaming

Empower, Equip, And Enhance

Your Cybersecurity Operations.

PlexTrac

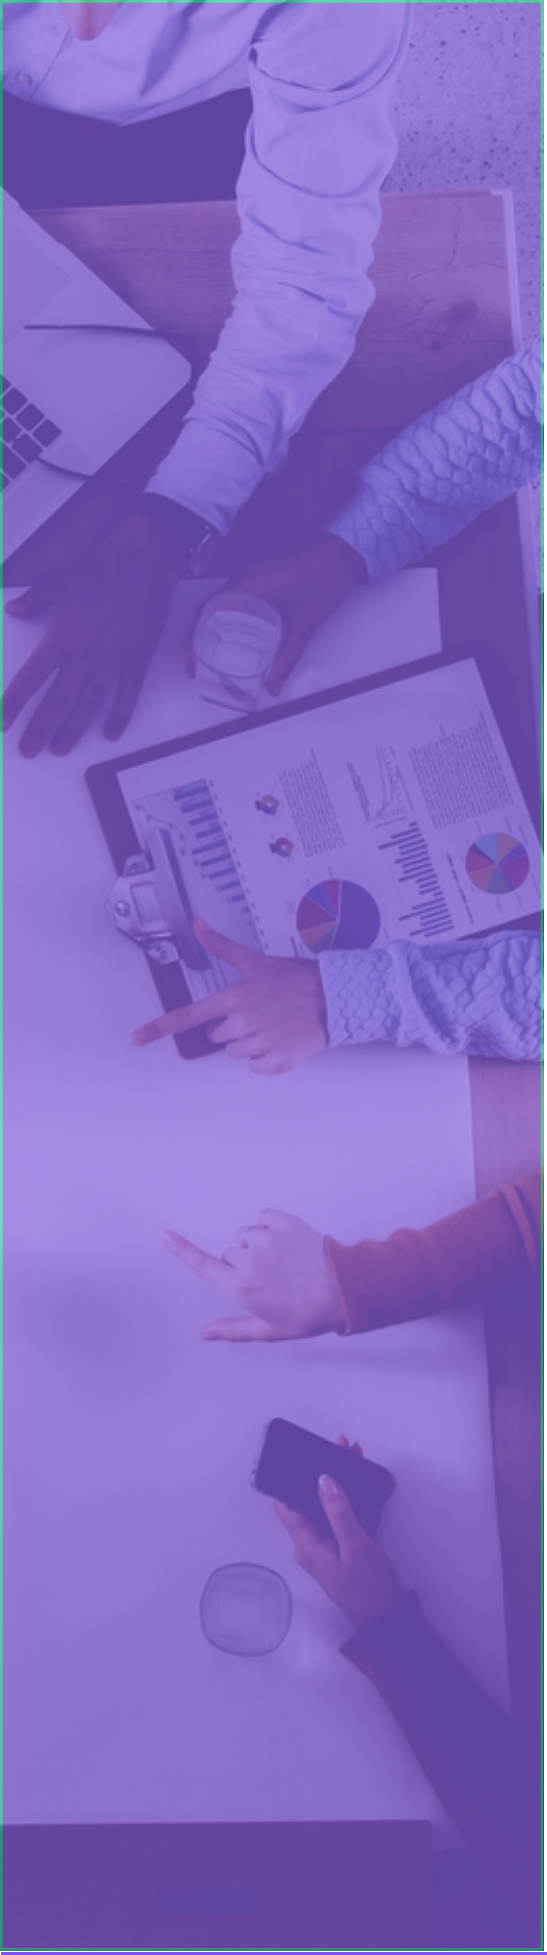

Table of Contents

Introduction

Cybersecurity is hard, presenting complex challenges to effectively managing enterprise risk. As the role of the Cybersecurity team has matured within organizations, the traditional roles of “Red Teams” and “Blue Teams” have been supplemented with the concept of “Purple Teams”. Security leaders have drawn an overwhelming consensus that purple team engagements, also known as purple teaming, provide immense value to rapid improvements in prevention, detection, and response techniques. Despite this consensus, little has been written to capture best practices for actually implementing Purple Team operations.

In this paper, we will review the problems with the status quo that have given rise to the purple teaming concept. We will discuss how purple teaming attempts to mitigate these problems at the theoretical level. Finally, we will offer guidance on how to put the theory into practice with concrete actions in your environment.

Definitions, Terminology, and Assumptions

Cybersecurity is a relatively new discipline, and as such even popular terms like “Red Team” and “Blue Team” may be interpreted and used differently by members of the community. So before leaping into discussion of the emerging concept of purple teaming, it is prudent to be clear about key terms.

Assessment - An assessment is any activity that is used to identify weaknesses or gaps in an organizations security controls or risk posture. This definition is purposely broad as it is intended to capture all proactive activities conducted towards an organization. Examples of assessments include penetration tests, vulnerability scans, risk assessments, compliance assessments, security questionnaires, etc.

Blue Team - Traditionally, the Blue Team refers to a subset of an organization’s technology team tasked with implementing the organization’s security controls and defending from cyber attacks. These key players are specifically tasked with the prevention, detection, triage and eradication of malicious cyber activity. As the defenders of the realm, they deploy a web of sensors that collect and aggregate data from across the environment using tools like Security Incident Event Management (SIEM) systems. They regularly build and exercise playbooks to guide their actions during the fog and friction of an actual incident. Increasingly, they are automating responses with tools like Security Orchestration and Response (SOAR) tools. Many IT professionals not specifically assigned to the security team routinely perform defensive functions such as patch management, hardening and ACL configuration. At PlexTrac, we take an inclusive view when referring to the Blue Team which incorporates all staff performing defensive activities. The Blue Team’s responsibilities are vast and often overwhelming, but in general the blue team is responsible for protecting the organization’s technology infrastructure from a breach.

Red Team - Red Teams exist to test the effectiveness of Blue Teams through proactive assessments. These professionals ideally use a defined methodology to thoroughly evaluate defenses, employing tools, tactics, techniques and procedures modeled after actual threat actors. These teams perform technical penetration testing, but may also use social engineering and counter-physical security skills to simulate adversary activities. Their tradecraft is often referred to as "offensive security," a nod to their role in supporting the overall security objectives. From the PlexTrac perspective, we also incorporate any form of assessment team as part of the Red Team. We broaden the definition to include teams such as Governance, Risk, and Compliance (GRC) or internal and external audit teams. Effectively we consider the Red Team to be any person or team charged with conducting assessments that result in actions to be taken by the Blue Team for remediation and risk reduction.

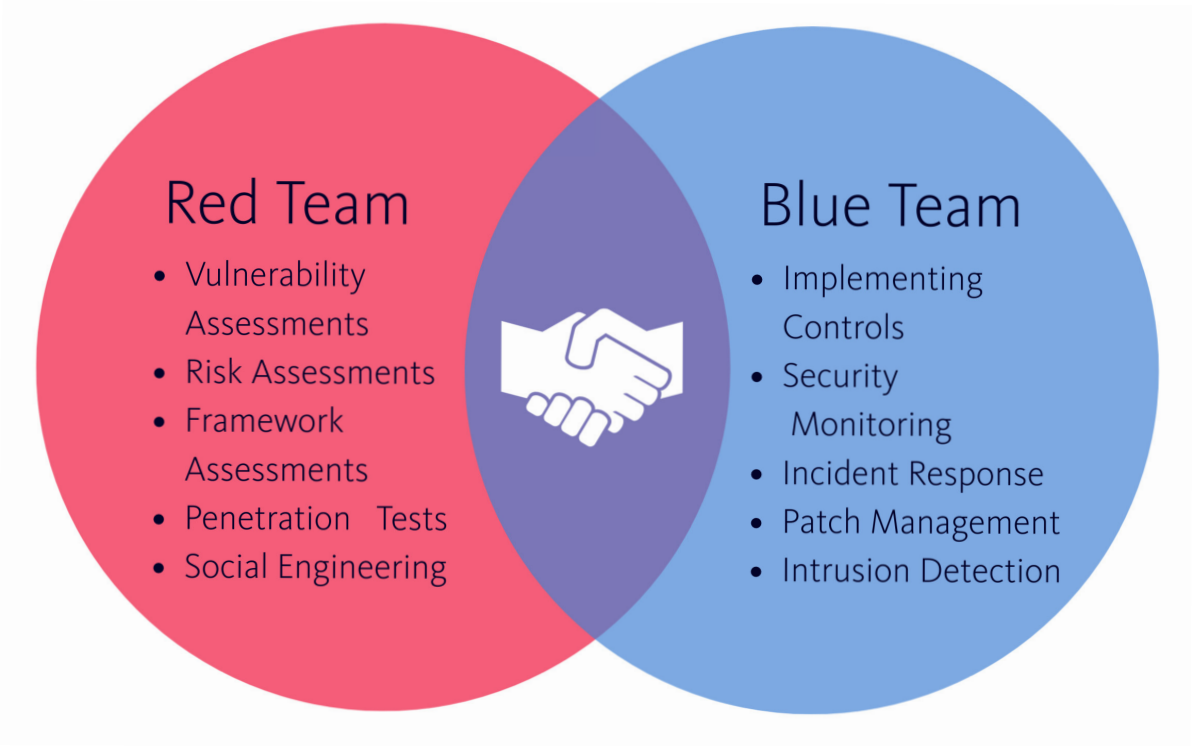

Purple Team - Traditionally, a Purple Team is considered a penetration testing team collaborating with a subset of the Blue Team and conducting a concrete, point-in-time assessment. The Red Team explains what attacks they are executing in real-time, with a goal of determining whether the Blue Team can either prevent or detect the attack in question. This paper will challenge this limited traditional role and discuss an broader, more effective use of the Purple Team concept. We argue that through a more expansive approach that leverages both technology and process, we can enhance the value from traditional purple operations and establish a new paradigm for achieving the mission of the cybersecurity program.

The Status Quo: Red vs. Blue

It is no secret that the challenges faced by cybersecurity teams are broad, including limitations in budget, time, and talent. Despite these challenges, the organization expects the security team to deliver a mature cybersecurity program that ensures the protection of the technical infrastructure and organization's most critical digital assets. Typically when building a security program, a team will start with the basic security controls and gradually add additional defenses in a ramp-up to a general assessment. This assessment could be a gap analysis for a specific framework, such as NIST 800-53, PCI-DSS, or CIS 20 or may include more general technical and non-technical efforts to identify key vulnerabilities. Once baseline controls are in place, the organization is ready for what may be considered the "ultimate" assessment - the penetration test. After one or more of these assessments, the Blue Team is tasked with fixing the identified issues while the Red Team moves onto another assessment or removes themselves from the picture altogether.

Thus, the current assessment paradigm involves multiple assessments by multiple teams (internal or external) where security issues and gaps get identified and then handed over to engineers or analysts responsible for investigating and ultimately remediating the risk. This is a perfectly logical approach - but logic doesn't always equate to efficiency. The time required to conduct an assessment, deliver the findings, remediate and then reassess the issues can take months (if not years). The existing paradigm also suffers from frictions which hinder collaboration and may even foster adversarial relationships. As mountains of findings pile up with limited resources for remediation, Blue Team begins to feel pummeled from multiple directions. Red Teams that lack exposure to the challenges of remediation may lose sight of the true goal - enhancement of an organization's cybersecurity posture. Additionally, a Red Team can often get comfortable in their current set of attack techniques because the Blue Team is slow to resolve known issues. This degrades the skillsets of Red Team operators who lack incentives to stay on the "bleeding edge" of real-world tactics, techniques and procedures.

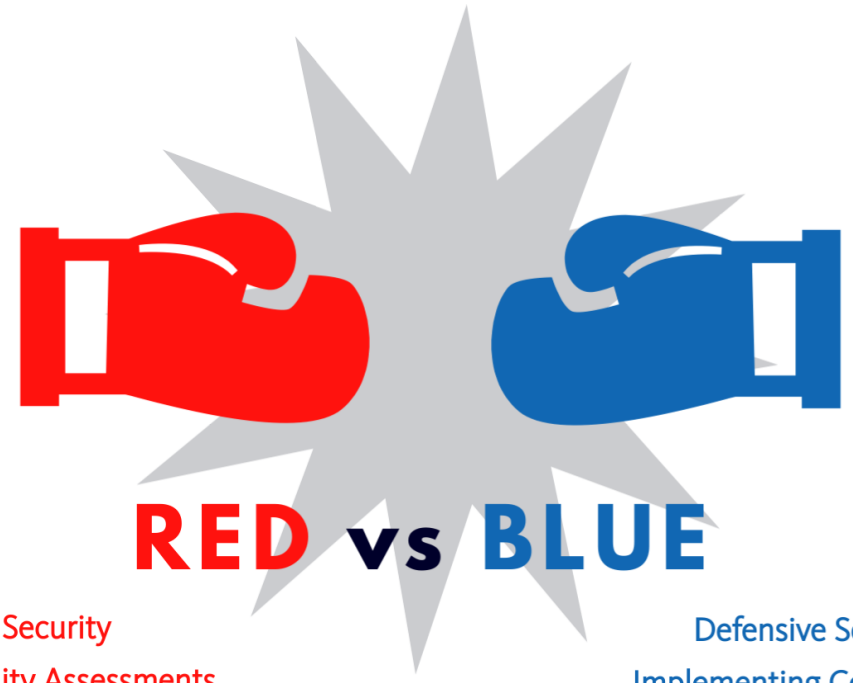

Offensive Security
Vulnerability Assessments
Risk Assessments
Framework Assessments
Penetration Tests
Social Engineering

Defensive Security
Implementing Controls
Security Monitoring
Incident Response
Patch Management
Intrusion Detection

The current paradigm for proactive security is heavily focused on periodic assessments with a defined start and stop to Red Team activities. This regimented engagement lifecycle, in which activities are performed separately and at separate times, contributes to the "us versus them" adversarial relationship that too often develops. Red Team activities are seldom communicated in a clear, consistent and timely fashion. Blue Team activities are not made visible to the Red Team, depriving them of the intelligence needed to refine the attack vectors to test blind spots. The lack of coordination extends to the methods we use for communication and coordination. Data generated, aggregated and enriched by both teams remains siloed in the tribes, spread out across multiple tools and platforms. Consolidation and analysis of progress data lags current activities, resulting in stale analytics for stakeholders and decision makers. Reports are abstracted into additional reports and presentations, depriving leaders with the real-time views of progress needed to support resource decisions. The inefficiencies inherent in the current paradigm combined with the constraints of time, talent, and budget ultimately result in security programs that are far too heavily reactive.

Challenges with the Status Quo

Today, most mature organizations build their information security program around the "Red/Blue" paradigm. Blue Teams self-assess to identify risk, implement continuous vulnerability management programs to mitigate risk, and (hopefully) detect and respond to incidents as they occur. Red Teams are often the "hired guns," brought in to occasionally test the defenses and identify previously-unknown gaps. Larger organizations may have permanent in-house Red Teams, but they typically are spread thin and operate similarly to consultant services with regards to frequency of engagement with any given business unit.

This is certainly an incredible improvement over the state of affairs at the turn of the 21st century, when dedicated defensive teams were a rarity and offensive security was in its infancy. This progress has been primarily driven by a recognition of the need to incorporate offensive security as a pillar of any information security program. Many of those curious folk with skills who were once disdained as unprofessional at best and criminal at worst have joined the mainstream security community as recognition of their value has gained acceptance.

While this progress is laudable, continuous improvement demands that we seek new ways to address the challenges we all face. Our adoption of offensive security is incomplete; we have opened the door to our hacking brethren, but tribalism persists which degrades our ability to maximize the value that offensive security principles can bring to our organizations.

Lack of Common Goals

If you ask a Blue Team member how they measure success, you are likely to hear a variation on one of the following themes:

- Absence of data breaches that impact the bottom line
- Number of malicious attacks thwarted at the perimeter
- Reduction in detection of incidents which trigger response actions

If you ask a Red Team member how they measure success, you may very well hear a different set of metrics:

- Level of access achieved on the target domain
- Evasion of detection when executing key functions
- Total number of vulnerabilities discovered

For both red and blue, these are commendable tactical objectives. But they are not necessarily strategic objectives. While "absence of data breaches impacting the bottom line" may seem like a strategic objective, it is actually a strategic outcome when another objective is met: Detection and response before the attacker achieves their objective. There is a common phrase oft repeated within the industry:

> "Prevention is ideal, but detection is a must."

Blue teams certainly want to stop the adversaries at the gates. Red teams want to uncover as many vulnerabilities as possible.

However, these are supporting objectives of what should be our primary concern in a world where the perimeter is rapidly dissolving: Detection and response.

All teams must share common strategic objectives to be successful, even if tactical objectives differ. In today's Red vs. Blue paradigm, it is not clear that the teams share common strategic objectives. We posit that the primary strategic objective should be detection and response, and our Purple Teams should be organized, trained and equipped to further this cause from a proactive perspective.

Information Security Skills Gap

We have all seen the numbers, and no matter which set you believe, the bottom line is that we don't have enough trained information security professionals and the problem is only going to get worse. Many organizations cannot find (or afford) in-house information security expertise. As a result, remediation of discovered vulnerabilities is often performed by IT staff. For simple remediation efforts like patching, this is perfectly fine. However, remediation of more complex vulnerabilities can be a challenge for personnel who lack an understanding of offensive security concepts. Furthermore, without a clear understanding of the root causes, those same personnel may be prone to repeat the processes that resulted in the vulnerability.

Continuing to silo our professionals into purely Red/Blue tribes will only exacerbate the existing

skills gap due to missed opportunities to pollinate wider audiences with offensive security principles. Red Teams exist to test, but ultimately the test is subordinate to a greater goal: to teach.

The current Red vs. Blue paradigm also ignores one of the foundational principles of teaching, one you have surely heard countless times before: Crawl, then walk, then run.

You can almost hear the sound of a starting gun at the commencement of an engagement. The Red Team begins stringing together all of their tactics, techniques and procedures to achieve their objectives. Unless the Blue Team is fortunate enough to detect each attack path in real time, they are missing key learning opportunities. They will never glean as much knowledge from a report as they can gain by observing and understanding as an activity is occurring. We know that a building block approach composed of partial-task training (PTT) events is a much more effective method for knowledge transfer. Nobody reads a book on JavaScript and then sits down to code a new web app. Everyone starts with a simple "Hello World" program, then progresses through a series of increasingly more complex exercises.

The information security skills gap is a daunting issue that few of us have the power to meaningfully impact at the macro level. However, remember the words of pioneering tennis legend Arthur Ashe: "Start where you are. Use what you have. Do what you can." You are in a position to change your organization. You have intelligent, savvy and eager employees on the Blue Team. You can effectively train them in offensive security principles and we will discuss how purple-teaming is the perfect medium for achieving this objective.

"The information security skills gap is a daunting issue."

Uninformed Threat Modeling

Which of these two types of test will provide the greatest value to your defensive efforts?

1. A penetration test where the testers use their best tricks to find any openings to whatever data they can compromise.
2. A penetration test where the Red Team analyzes historical attacks, understands what data is coveted by adversaries and uses this information to model realistic attack profiles.

In order to perform the latter, the blue team needs to share locally-generated threat intelligence with the Red Team during the planning phase. Unfortunately this rarely occurs. For example, when was the last time that a Red Team asked you to review incident response reports during the planning stage of an engagement?

Additionally, Red Teams should be informed about existing defenses prior to planning an engagement, or to take it further, the defenses you are truly focused on testing. Red Teams that choose attack vectors which are closely monitored are wasting resources. Red Team activities should guide the Blue Team to examine attack paths that are not well fortified to provide the greatest ROI. The Red Team should also be closely connected with the defenses the Blue Team wants to specifically test and augment the test plan accordingly. It is vital that the Blue Team knows the areas of focus have been addressed in addition to the new tactics and techniques exercised by the Red Team.

Lack of Post-Assessment Collaboration

In an ideal world, the process of knowledge transfer would be an ongoing process with collaboration occurring between teams independent of a defined test or assessment. An interim step towards that ideal world is simply post-assessment collaboration. No matter how detailed a finding or how many artifacts are included with the report, questions will inevitably arise.

Unfortunately most engagements with external assessment teams do not include post-assessment collaboration in the Statement of Work, and the blame is shared. Penetration testers enjoy hacking, and there is a cultural bias against becoming involved in the actual remediation work, so few even pitch the concept. Managers may not understand the value (but they do understand the high cost of billable hours), so do not ask for (or better yet, insist on) the consultant "sticking around" after the report is delivered.

This is a lose-lose paradigm and one that we must overcome as an industry. As the talent gap continues to grow, there will be an even greater reliance on external information security services. If we continue in a transactional mindset, where communication and collaboration starts and ends with a document-based report, we will never realize the efficiencies that are possible with outsourcing. Organizations will pay consultants to identify risk, but will lack the in-house expertise to address all the risks uncovered. Already today, we lament the number of reported findings that are never addressed; the image of the pen-test report "collecting dust in a desk drawer" has become an archetype in the industry. Granted, there are a variety of factors which contribute to unaddressed findings, to include budget, competing priorities and the rapid pace of technological change. But it is also true that remediation of many findings are beyond the skill set of in-house IT staff.

In this transaction-based approach, consultancies are also leaving money on the table. Billable hours are billable hours, but this isn't about milking the clients. If a block of additional hours enables more effective outcomes, then the value proposition for the customer is real.

Aside from the financial benefits, consultancies can also become more effective advisers to their clients when they have regular visibility on the true burdens of remediation. Some mature consultancies today provide a "burden of implementation" score with each recommendation, and some go so far as to provide an evaluation of the cost/benefit of a given remediation solution. This is extremely valuable information for the client but cannot be provided if the red team does not have relevant experience in actually implementing the solutions they prescribe.

In an enterprise environment with a dedicated internal Red Team, enlightened management can facilitate regular post-assessment collaboration between the Red and Blue teams. Unfortunately, this is not standard practice in most organizations. The information security workforce is a revolving door between internal testing teams and consultancies, and those same cultural biases against involvement in remediation work exist in internal teams. Blue Teams have their own biases as well - some rooted in real or perceived arrogance of Red-Teamers. Management is faced with two tribes that don't particularly want to work together, and lack an understanding of the value proposition in forcing collaboration.

Inadequate Tooling for Effective Collaboration

> "All teams face resource challenges regarding budget, time, and skillset."

Today the primary tool for communication between the red team and the blue team is the "final report," usually delivered in a PDF. This traditional format suffers from several inherent weaknesses.

First, the usability of the report for the consumer is inversely proportional to the length of the report. The more artifacts the testers include, the harder it is for blue teams to separate the signal from the noise. Testing teams collect reams of useful information, but make judgement calls about what to include in the report so as to avoid overwhelming the consumer.

Second, the process of transferring data from the report to the customer's workflow system is manual and laborious. Copy, paste, repeat. Inevitably, some data and artifacts never survive this process and the Blue Team is provided with a subset of the available information - which is itself a subset of the information that the testing team collected. Additionally, the current workflow system for tracking the findings from a report is predominantly a spreadsheet that doesn't trace back to previous reports or could get lost in the shuffle of personnel changes.

Finally, document-based delivery prevents the use of one of the most effective methods of communication: video. If you want to learn how to fix your washing machine, do you read the service manual? No, you go to Youtube. If a picture is worth a thousand words, a 30 second video clip can be worth 50 screenshots and a lot of unnecessary narrative.

With all of these challenges in the current paradigm, how do we continue to make progress in achieving the goal of avoiding compromise or detecting compromise as soon as possible within the lifecycle? The answer lies with a shift in the paradigm towards true purple teaming and effective collaboration.

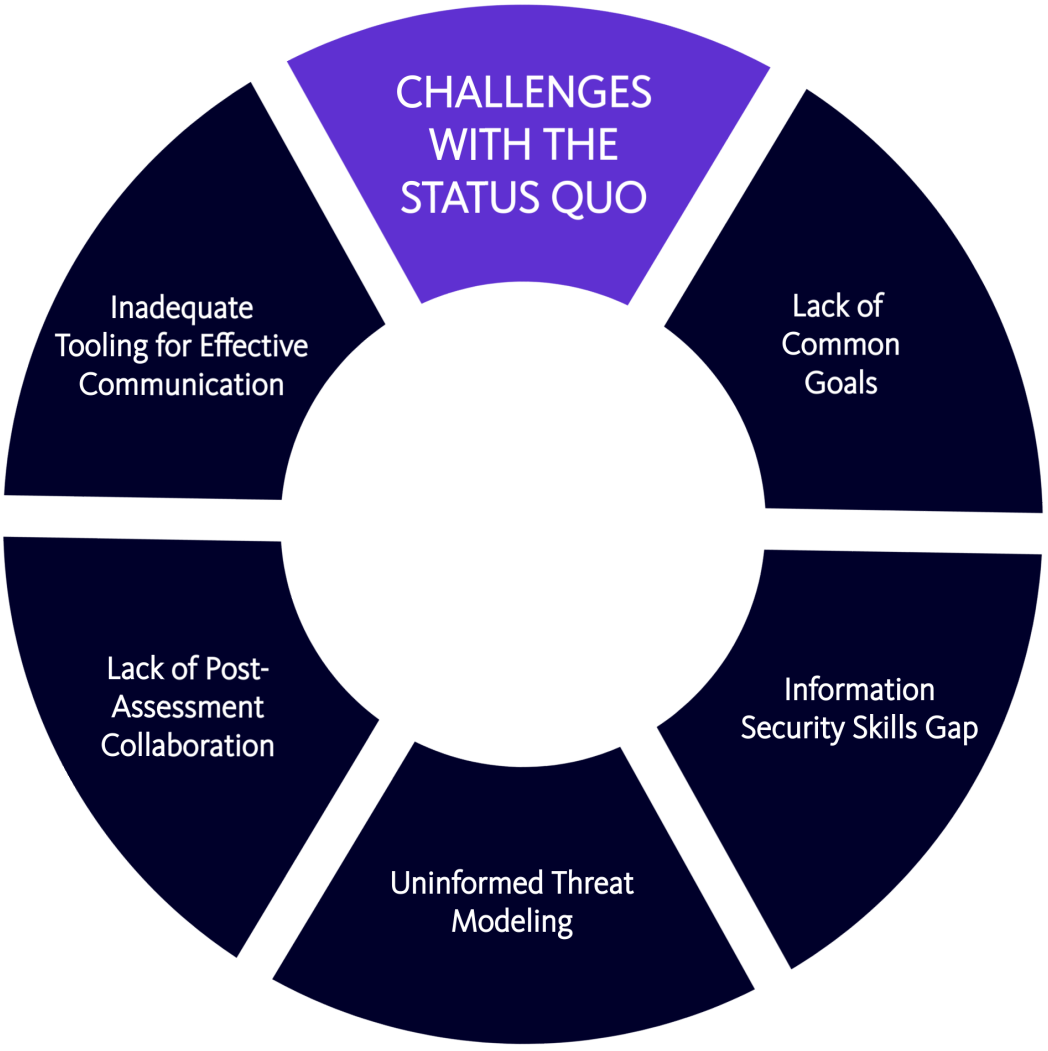

Purple Teaming to the Rescue:

Shifting the Paradigm

What is Purple Teaming?

Purple teaming is the collaborative function performed by Red Teams and Blue Teams to mitigate all of the pains discussed thus far. It's a new approach to collaborative testing and remediation that seeks to break down cultural barriers, improve communication and "level up" everyone's skills. It is also aimed at reducing the mean time to remediation for reported risks and vulnerabilities. Note that purple teaming is a role but not a job; there are no dedicated Purple Team members. A team member's function is either Red or Blue, but everyone's role is strictly purple with a common mission of detecting compromise as early as possible within the attack lifecycle. So what do this role look like? There is no canonical definition of purple teaming, but common tasks and objectives include:

- Design realistic tests based on shared priorities, informed by locally-derived threat intelligence and tailored to test the defenses' critical assets.
- Speed up the process of remediation through established channels for collaboration
- Prevent related future occurrences of issues through knowledge transfer of root causes
- Help foster an offensive security mindset across all members of the cybersecurity team

This all sounds wonderful but how does an organization build a well-functioning Purple Team? What activities are truly involved within purple teaming? And how do you know if you're succeeding? We'll break down the answers to help clarify the foundational elements of an effective Purple Team.

Team Organization

As discussed previously, Purple Teams are functions and not dedicated positions or job titles. However this does not mean that the relationships among team members should be unstructured. Supporting and supported roles should be clearly defined, to include:

Team Composition

Assignment of roles should be documented through internal policy documents or included in a Statement of Work / Master Services agreement. These roles should be well understood across both Red and Blue team functions. You should breakdown all current teams and activities within your security program and categorize them as either Red or Blue, and we encourage using the more expansive definitions of "red" and "blue" discussed earlier.

"Purple teaming is a role, not a job."

Team Member Functions

Roles and responsibilities need to be documented so that team members know who to go to and what deliverables to expect. In a world where everyone is always overtasked, if it isn't in writing, it's "not my job." Documentation of responsibilities demonstrates the organization's commitment to purple teaming and makes it easier for management to evaluate performance and hold members accountable.

Communications Plan

It is critical to understand what the communication lines are between Red Teams and Blue Teams as well as between the Purple Team and stakeholders. Depending on the scenario, it's possible that junior team members may be communicating directly with internal or external stakeholders or executives, thus it is important to have clear lines of communications established.

Activities and Cadence

With the team organized and clearly defined, the next phase is to establish the cadence with which Purple Team activities occur and the scope of those activities. There can be a lot of environment-driven license with these activities, but clear examples of best-practices and proven techniques are available. Purple teaming activities can be equated to that of executing a sprint within an agile workflow or scrum team. A Purple Team engagement should typically be a two or three week cycle that involves both the assessment and remediation efforts. This requires discipline on both the Red and Blue teams and also helps scope the planned activities to a reasonable and achievable set of objectives.

Let's assume you decide on a cadence of a two week time period for all activities to be conducted for a purple team engagement. The activities with the engagement include planning,

assessing, collaborating, remediating, and reporting. It is important to note that there is not a required order to these activities. Planning should initiate the engagement, but additional planning will occur throughout the engagement period.

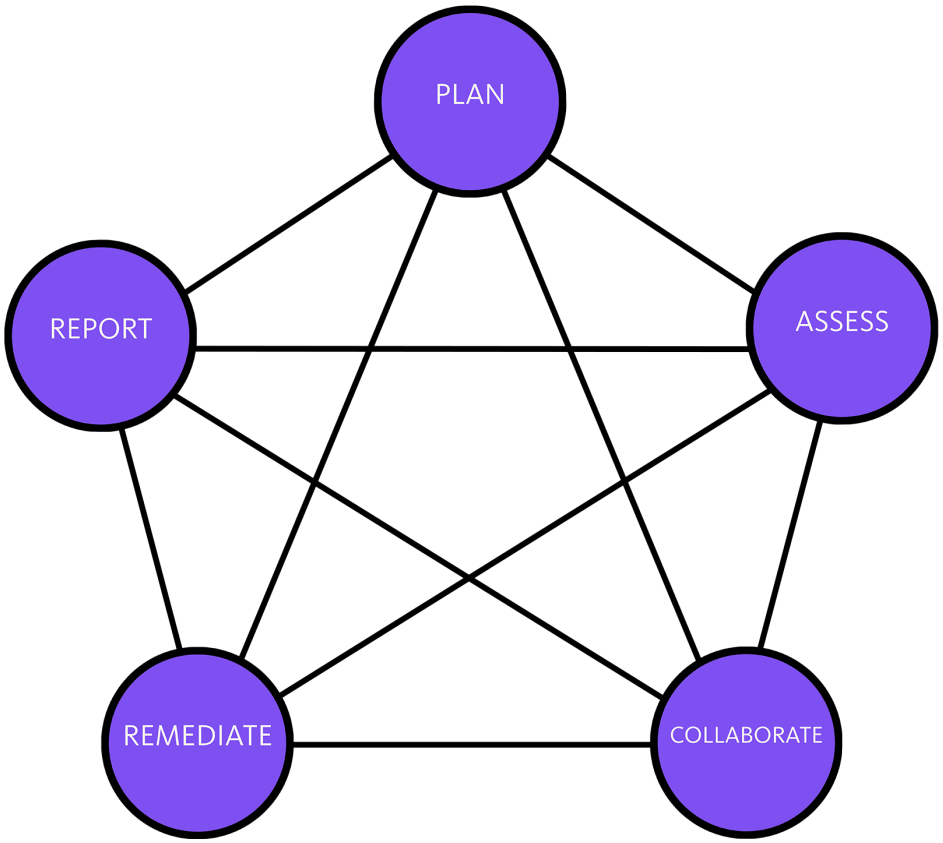

Plan

Within the planning phase, first the Red and Blue Teams must collaborate on what the outcome of the two week engagement period should be. This phase is imperative and is used to determine what gaps or issues need to be addressed from a program perspective. Planning concludes with a set of defined actions to perform, but it begins with establishment of objectives. What are the questions that the team expects to answer once the engagement is complete? Objectives should be clear, concise and unambiguous, and above all they must be achievable. Some examples of objectives for a purple team engagement:

- Determine the ability to detect data exfiltration via DNS tunnelling
- Evaluate effectiveness of lateral movement detection through analysis of netflow data

The objectives will drive the plan for testing activities, but how are the objectives determined? What questions do we need answers to from our activities? Locally-generated threat intelligence is an excellent resource, which is often gathered during the incident response process. If there are gaps in understanding how a previous malicious actor was able to perform their activities,

these can become the questions which drive our objectives.

With the objectives set, the respective teams can plan their activities. All planned activities should be measured against whether they will support completion of an objective. This may be outside the "comfort zone" for the Blue Team, which traditionally operates in a more reactive environment. The Blue Team cannot simply sit back and wait to see what is thrown at them. With a shared knowledge of the objectives, they must plan the activities that will provide the data to comparatively evaluate various detection and prevention mechanisms. Alternative theories should be tested to enable proper data collection among a range of potential solutions. Planning should be done for how and what data will be collected, with the end goal of supporting effective analysis and ultimately process refinement.

Assess

Though the plan should be robust, it is true that "No plan survives first contact with the enemy." This is not to say that plans should be abandoned, but they should be refined throughout the assessment process.

If a planned attack path is blocked, what alternative methods are available that still support the engagement objectives? Red team activities must be managed to meet tactical milestones, and this is best accomplished by dividing larger objectives into discrete and manageable tasks for which progress is more easily measured. Given that any engagement is time-constrained, it is vital to rapidly identify blockers that may degrade the ability to meet overall objectives. Artificially bypassing a defense isn't "cheating" if the result is meaningful data that can drive improvement. A door that is locked today may be open tomorrow, and there is much greater value in using the limited time available meaningfully.

The Blue Team must be continuously evaluating their performance against expectations. If a planned method of detection does not appear, how can it be refined? What data are we not collecting that might provide us with indicators of compromise? Do we need to refine our methods of parsing data to detect signal through noise? It is likely that blue teams will identify not only gaps in technical capabilities, but in knowledge of adversary techniques as well. These gaps should drive research, but also collaboration.

For assessment activities, the Red Team must be extremely targeted and specific with the focus on achieving their goal within a short period of time. Thus the Red Team should start with a big goal and break it down into small phases that can be accomplished in one to two week iterations. When assessment activities are occurring, the blue team must know what the current protection mechanisms are that exist with respect to the targeted attacks and how they plan to

detect the Red Team activities. Additionally, the Blue Team should also be conducting research on what additional techniques they might encounter by the red team and what additional controls they may need to have in place. This approach encourages a proactive mindset for Red and Blue Team activities.

Collaborate

It will be impossible for the Blue Team to accurately gauge whether they are detecting activities if they don't know what activities are occurring. The answer to the question, "Did you see it?" is much more nuanced than "yes" or "no." Consider these possible detection outcomes:

- A Red Team activity was not logged
- A Red Team activity was logged but the data did not generate an alert
- An alert was generated but not triaged properly
- An alert was acted upon, but the defensive response was ineffective
- A defensive response was effective in closing the vector, but not timely enough to prevent the attacker's objectives

For each of these potential outcomes, there are different lessons learned which will result in different process improvements. Each generates their own questions and threads to pull. But in the first three cases, the Blue Team cannot begin to ask these questions without an awareness that something happened. Thus to reap the benefits of purple teaming to their fullest, both sides need to have situational awareness of the totality of actions.

> "Real-time collaboration is required to achieve full return on investment."

While it is true that Blue Teams can perform forensic activities after an engagement to gather more data, the opportunities to refine and tune techniques is lost once Red Team activities are complete. This means that real-time collaboration is required to fully achieve the return-on-investment from the engagement. To ensure that this collaboration occurs, engagements should include regular checkpoints to provide each team an opportunity to confirm their understanding of the actions of the other. This can take the form of daily stand-up meetings, real-time collaboration through chat, or via embedded liaisons from the opposite team. Not every detail needs to be shared, but enough collaboration must occur to allow teams to refine their activities to test alternative responses.

Red Teams benefit from this collaboration as well. Understanding the time required to detect and respond can directly influence the choice of tactics and techniques. If a detection has

occurred and a playbook initiated which will ultimately thwart a line of attack, continuing on that path will provide limited value. It is better to understand early that defenses were effective and be able to make an informed decision as to whether to allow the activities to continue or move to the next objective.

Remediate

Until this point, we have discussed purple engagements in the context of a discreet, time-constrained engagement of 2-3 weeks. Remediation activities can stretch for months (or longer), requiring resource approvals and procurement cycles. This does not mean that r emediation is not an integral aspect of any purple engagement. As issues are discovered, planning can begin on the necessary steps to eliminate the vulnerability and this plan can be generated collaboratively with the Red Team. Solutions that might appear adequate to the Blue Team can benefit from Red Team inspection, potentially preventing inadequate solutions that result in re-workv (or worse - unremediated vulnerabilities).

Red Teams benefit from remediation planning by gaining a more accurate understanding of the burdens associated with their recommendations. In almost any circumstance, there are multiple possible acceptable remediation solutions. Not all acceptable solutions may be equally "secure," but through collaboration during the remediation process, Red Teams gain greater understanding of what is feasible. A feasible solution that is acceptable is always preferable to a perfect solution that cannot be implemented due to resource constraints.

Even feasible solutions have a resource burden for implementation. Red Teams that are exposed to realistic cost/benefit analysis of their recommendations will benefit from learning to prioritize recommendations that prioritize return on investment. If I can remediate 10 "High" severity" findings for the cost of remediating one "Critical," is it the best course of action to prioritize remediation efforts based only on severity? Exposure to real resource decisions is crucial in assisting red teams make solid recommendations on the order of remediation. Red Teams that have no such exposure become philosophers on a hill.

Report

Reports that are informed by both offensive and defensive activities provide a more holistic assessment of the environment, but they also carry greater credibility. When both sides concur on the need to devote resources to remediate issues, those recommendations are more likely to carry the day with decision makers.

Recommendations that are informed and tempered by resource constraints are more readily implementable.

Recommendations that have the buy-in of those charged with implementation are more likely to be carried to fruition.

In short, the final page of any assessment should carry the signatures of both team leads. The report recommendations are the culmination of the engagement. We began with objectives to learn unknown aspects of our environment. Armed with newfound knowledge, a jointly-signed report communicates to leadership that a professional examination was administered which delivered recommendations free of tribal politics.

During the course of any engagement, threads will be uncovered which cannot be pulled under the scope of the current effort. The report should not neglect the consensus decision on the way forward. What have we uncovered that we should examine next? In addition to providing solid recommendations, a consensus on next steps will help garner the support to resource follow-on endeavors.

The most important thing to keep in mind regarding the reporting process of purple teaming is that it is iterative and dynamic, a result of each collaborative exercise.

Measurement and Progress Reporting

We have discussed purple teaming activities in the narrow context of defined 2-3 week engagements. In an ideal world there are no absolute timelines with purple-teaming. Engagements must be planned, and they may often be dual-purposed to meet compliance reporting or quarterly board report deadlines. But attackers don't have cycles or timelines, and thus neither should the purple team. As organizations mature in their purple-teaming processes, so can their engagements mature in complexity and scope. Organizations may reach a level of maturity where multiple purple team campaigns are occurring simultaneously; some "low and slow," some "smash-and-grab." Purple activities should ideally include a balance of these two, an approach that meets both security and business requirements.

Cybersecurity is now established as part of the duties expected of Board members, and many Board members are eager to help (if not technically skilled in the subject). As cybersecurity professionals, our duties include education to non-technical decision-makers. That requires regular reporting that provides meaningful metrics, presented in a format that doesn't require advanced technical skills. Recalling that the objective of purple teaming should be both prevention and a reduction in time to detection, our progress should include metrics such as:

- Mean time from vulnerability detection to remediation, based on severity
- Change in mean time from vulnerability detection to remediation

And if applicable:

- Mean time from initial compromise to detection
- Change in mean time from compromise to detection

A simple metric that can be very helpful (if measured properly) is the relation between new findings opened and findings that are closed. At PlexTrac, we refer to this as the "winning/losing" chart. It's an easily understandable metric - are you opening more findings than you are closing? If so, you either need to devote more resources to your problems or tune your reporting. Regardless of what metrics your organization deems important, consistency is paramount. Cybersecurity is a board-level issue, and board members are seldom cybersecurity experts. If we are to expect these leaders to make informed decisions, we owe them consistent metrics that enable them to detect and respond to change over time. Establish core metrics that become familiar to your leadership, and understand that changing these metrics will impact their decision-making functions.

Your organization can decide on the proper cadence for purple team engagements, but the paradigm is designed such that these activities are approached from an operational mindset in which they never stop. As such, you need to have a way to measure and show progress in real-time.

"Cybersecurity is a board-level issue."

Equipping Purple Teams

A tool is simply a solution that meets the intent of a validated requirement, and a requirement is simply a necessity for a team to meet its objectives. We posited earlier in this paper that the primary objectives for Purple Teams should be enhancement of the organization's detection and response capabilities. We further posited that the primary mechanism for the desired improvement is improved coordination, collaboration and joint training across the spectrum of cybersecurity activities.

Peeling the onion back one additional layer, purple teams ultimately need solutions that aggregate, normalize and present data from a plethora of disparate sources to empower the joint collaboration that purple teaming envisions. Ultimately, purple teams need solutions that differentiate the signal from the noise and empower everyone to be focusing on the most important tasks for securing the organization.

Aggregation, Normalization and Presentation

Security efforts are performed in parallel and in a wide variety of physical and logical environments. Scan results, tests, assessments and IR reports are generated in a wide variety of formats. Just as we need to "normalize" log data to separate signal from noise, we need to normalize the results of our information security efforts. If we want a Blue Team member to gather value from an application security scanner, we can't hand them a BurpSuite XML export. Similarly, we wouldn't expect a Red Team member to easily identify the "signal" if presented with a full NIST 800-53 assessment. Perhaps most importantly, we can't expect senior leaders in our organizations to gather meaningful insights from all the mechanisms for risk identification if we don't provide them with some common structure to present findings and overall risk.

Normalization does not, and should not, equate to filtering. Data should not be discarded in order to facilitate the understanding of secondary users. But because risk identification efforts do produce results in such a plethora of formats, our tools need to be flexible enough to capture as much data as desired while providing a structure that empowers secondary users to gain insights.

"Purple teams need solutions that differentiate the signal from the noise."

Even once normalized, not all consumers will need or want to be presented with all data. Thus our solutions should provide the same degree of flexibility in presentation as provided during aggregation. Today, many cybersecurity leaders struggle with workarounds to parse and present data such as Jira dashboards. While these workarounds can be effective, they represent an unnecessary investment of time by people that should be doing security - not analytics. Purple Team solutions should make the tailoring of data effortless and be capable of providing it in various mediums.

PlexTrac was designed from the ground up to be the aggregation, normalization and presentation platform that purple teams need to facilitate collaboration and coordination.

Aggregation: PlexTrac supports collection of all sources used to identify information security risk, to include:

- Importation of all leading network and application scanner results
- Manual findings from penetration tests, to include an infinitely-customizable field set
- Questionnaire / Framework-based assessments such as PCI, NIST, CIS, COBIT, ISO, etc.

Normalization: By organizing all risks into a common data structure, PlexTrac facilitates rapid understanding by secondary consumers. Ultimately, the method by which risk is identified is irrelevant; critical risks uncovered during a compliance inspection are no less valuable than those discovered by a crack pentest team. Normalizing risk data provides the conceptual framework for leaders to understand risk from all sources and make informed resource decisions.

Presentation: PlexTrac's advanced analytics allow leaders to view the data the need with a few easy clicks in a convenient web-based platform, without the need for creation of custom dashboards or excel macros. Because remediation efforts are tracked within the Platform, the data is never stale. Custom export templates allow for standardized and professional traditional document-based presentation, tailored to exactly the data you intend to present.

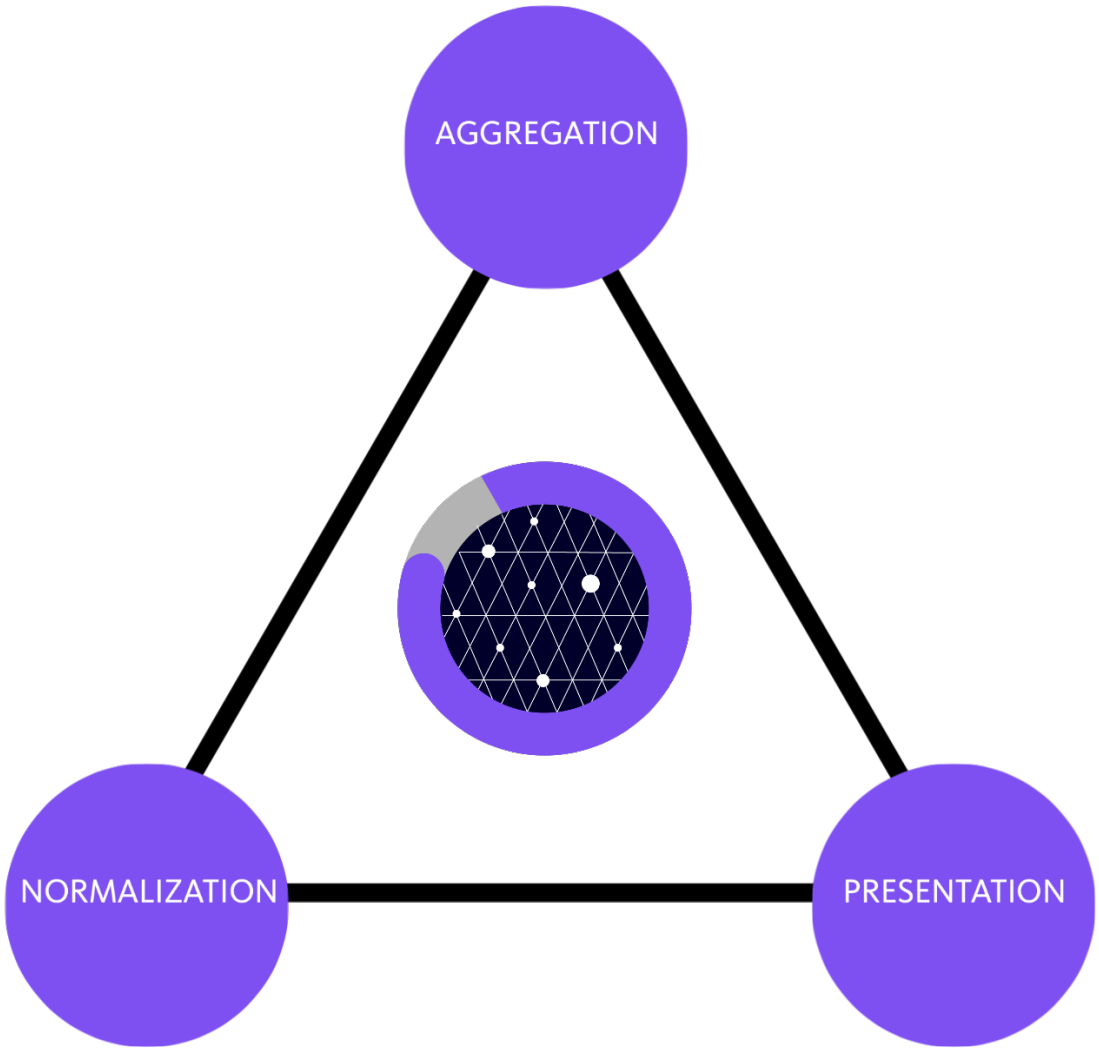

Conclusion

Purple teaming is in its infancy as an operational concept. The goals and terminology have rapidly gained acceptance in the information security community, however little guidance has been offered on how to actually implement joint Red / Blue operations. We do not claim that the guidance offered here is the best possible; we simply want to help move the conversation from beyond the notional to the practical.

At PlexTrac we are fortunate to have close relationships with some of the best teams in the industry, and they are kind enough to share their experiences with us. What is needed to further elevate the community is case studies from those who are implementing purple team operations in their environment. The data set available is currently too small. Only through continued sharing of successful (and unsuccessful) efforts will we establish accepted "best practices." Are you ready to help? Contact us at support@plextrac.com to discuss how we might incorporate your lessons-learned into future versions of this paper.

Interested in exploring PlexTrac as a tool to equip your purple team? Contact us at sales@plextrac.com to initiate the conversation.