DOCUMENTATION REPORT

Part of the BACHELOR DISSERTATION

Developing a Comprehensive Red Teaming Framework and Handbook

A Guide to Successful Engagement Strategies and Techniques

Bachelor	Applied Computer Science
Elective track Graduation	Cyber Security
Academic year	2023 - 2024
Student	Brendan Craven
Internal coach	Mattias De Wael
External promoter	Gert-Jan Wille

howest.be

Availability for consultation

The author gives permission to make this documentation report (part of the bachelor dissertation) available for consultation and to copy parts of this report for personal use. In all cases of other use, the copyright terms have to be respected, in particular with regard to the obligation to state explicitly the source when quoting results from this report.

3/06/2024

Foreword

For my mother Iva, my father Georgi, my friends, my lecturers, and everyone who worked so hard to support me and get me to where I am today.

Shoutout to the awesome folks whose work I've drawn from, particularly Jacob Oakley, Peter Kim, Marcus Carey, Jennifer Jin, and the University of Foreign Military and Cultural Studies – thanks for sharing your knowledge with the community!

Many thanks to all the industry professionals who made the time out of their extremely busy schedules to talk to me, to be interviewed, for verifying and scrutinizing all the things I said in this handbook.

Special thanks goes to Jayson E. Street, a red teaming veteran who made himself available on two separate occasions for two interviews and asked for absolutely nothing in return. Your unwavering, continuous contribution to this industry always astounds me.

Table of contents

1 Introduction

1.1 General

The field of red teaming has garnered increasing attention in both academic and practical contexts. Despite its growing prominence, there remains a noticeable gap in academic research regarding comprehensive methodologies and frameworks to guide red team operations effectively.

This thesis came about during my internship at Cyber3Lab in Howest, Bruges. My internship was mainly focused on the research for this thesis.

1.2 The problem

Current research in red teaming primarily focuses on practical methodologies, as well as exploring the integration of red teaming with broader cybersecurity strategies. While existing literature provides valuable insights into various aspects of red teaming, there is a gap in the availability of comprehensive guides and frameworks for conducting successful red teaming engagements. This gap shows that there is a need for research that advances the theoretical understanding of this field, as well as research that offers practical solutions to enhance the effectiveness of red teaming in real-world scenarios.

This thesis aims to address this gap by proposing a novel framework and handbook for conducting red teaming engagements. This contribution will attempt to synthesize existing research, best practices, and practical insights. The new framework will hopefully streamline the process of planning, executing, and analyzing red teaming work, particularly solving the major challenges identified in the research.

1.3 Research question

What is the main challenge that red teams face and is there a way to mitigate that challenge while maintaining efficiency in red teaming engagements?

1.4 Experiment

To answer this question, I will populate my thesis with two documents:

1. A ‘handbook’ of sorts, describing all the fundamental basics about red teaming. This will help me to identify what the challenges red teams face are in every single area of an engagement cycle.
    - Using this handbook, I will pick out what appears to be the biggest challenge that red teams deal with.
2. A new framework that attempts to mitigate the challenge identified, without sacrificing efficiency in engagements.

2 Experiment

The general approach to research for the thesis was:

1. **Literature review**: Conduct an extensive review of the existing literature on red teaming, cybersecurity testing methodologies, physical security testing, and related topics such as network security testing. The goal for this phase is to identify key concepts, methodologies, best practices, and gaps in the research.
2. **Case study analysis**: Select and analyse real-world case studies of successful and unsuccessful red teaming engagements. Evaluate the methods used, techniques, and outcomes to extract lessons learned.
3. **Expert interviews**: Conduct interviews with experienced red-teaming practitioners, cybersecurity professionals, and organizations. Gather insights on current practices, challenges faced, and areas for improvement, as well as to verify all the research done up to this point.
4. **Survey**: Develop a survey to gather quantitative data on red teaming practices, challenges, and organizational perceptions. Administer the survey to a diverse sample of cybersecurity professionals and organizations, whether or not they are engaged in red teaming activities.
5. **Framework development**: Based on all the information gathered from the literature review, case studies, interviews, and survey data, develop a comprehensive framework for structuring red teaming engagements. Define the key phases, activities, and deliverables of the framework. The important part is that this framework has to solve a problem/challenge that other frameworks struggle with.
6. **Handbook development**: Create a practical handbook based on this developed framework, providing step-by-step guidance and resources for each phase of the red teaming framework. This handbook *needs* to be user-friendly, accessible, and perhaps adaptable to the needs of different organizations.
7. **Validation and testing**: Validate the effectiveness of the proposed framework through pilot testing in a simulated red teaming engagement. Gather feedback from the simulation and iterate on the framework and handbook as needed.

A final note: The way this thesis is structured is exactly how the handbook document is structured. Two separate documents for both the handbook and the framework were created, and then merged into this one final document. The thesis will read similarly to the handbook, in which each section is a separate chapter of the actual book.

3 Disclaimer

Before diving into this thesis, a disclaimer, the disregard of which can lead to negative consequences: **Get permission to hack.**

This is a serious disclaimer, one must not go poking around in systems that aren't their own without the thumbs-up from the owner. Even if it's driven by curiosity, not malice, one can still land in hot water. The placement of this disclaimer at the outset is deliberate, as common sense does not always prevail, necessitating reminders akin to warning labels on shampoo bottles to deter risky behaviour.

Here are some cautionary tales of people getting in trouble for hacking without permission:

https://cyberscoop.com/dji-bug-bounty-drone-technology-sean-melia-kevin-finisterre/

https://www.forbes.com/sites/thomasbrewster/2015/12/17/facebook-instagram-security-research-threats/

If complete certainty about the legal aspects is not felt, it is advisable to consult a lawyer or reach out to the European Digital Rights (ERDi) network. They're a collection of NGOs set up to help navigate the fine line between legitimate research and unlawful behaviour.

Always stick to systems where authorization is granted to test, with this green light preferably taking the form of a signature on a contract. There are numerous legitimate bug bounty programs and sites with challenges and practice setups that exist to enhance hacking skills without risking jail time, and a standard practice lock goes for less than 20 Euro on Amazon.

One more thing: I'm not an expert in everything tech-related. The handbook section of this thesis is meant as supplementary material for answering the research question, so take everything my junior mind placed in here with a grain of salt.

4 Why this handbook exists

Ethical hacking is rapidly advancing, making it challenging to keep pace due to continuous industry developments. This fast evolution results from reciprocal pressure between offensive and defensive security tools: Improvements in defensive measures prompt the development of more sophisticated offensive tools, which in turn spur further advancements in defence.

In red team engagements, the primary challenges extend beyond identifying and exploiting vulnerabilities. They encompass navigating various hurdles throughout the engagement life cycle, including difficult clients, ambiguous rules of engagement, and scope inaccuracies. Addressing these challenges requires a combination of theoretical knowledge, the right mindset, strategic considerations, and practical skills.

This thesis is **not** a practical guide on hacking networks or covert entry and social engineering. There are courses for all of that online, certifications that exist to bolster existing skills, and a ton of practical material can be found for free from various sources on the Internet. I'm not the Lockpicking Lawyer, I'm not John Hammond, I'm far from the best person to be learning all that from. I do not, in fact, suffer from the Dunning-Kruger effect.

5 Red teaming explained

The term ‘red team’ is commonly used to describe offensive cybersecurity practices encompassing red teaming and both network and physical penetration testing. The origin of the term dates back to the Cold War era in the 1960s, to think tanks in the RAND corporation in Santa Monica, California, as well as the Pentagon. Simulations were ran where a ‘Red’ force symbolized the **adversary** in tests against organizations facing potential threats from Soviet forces. However, the concept of conducting simulated attacks to evaluate defenses and responses predates this historical period. [1]

A red team's objective is to simulate attacks on an organization to identify vulnerabilities in their information systems and facilities. This adversarial emulation aims to uncover weaknesses in detection, defense, and response to threat actors.

5.1 Advantages of red teaming

Red teaming uncovers an organization's security strengths and weaknesses, revealing how well they can defend against potential threats. The effectiveness of this assessment relies heavily on the skills and ethics of the hackers involved, along with the defined scope and rules of engagement. Generally, investing in red teaming proves more cost-effective than addressing security issues after they've been exploited by attackers [2].

Red teaming offers a distinct advantage by assessing an organization's readiness unlike other security measures. While conventional tools focus on enhancing defenses, detecting threats, and optimizing response, red teaming evaluates the effectiveness of these measures. It also pinpoints shortfalls in security posture such as overlooked threats, false alarms, or duplicative security measures, helping an organization streamline their resources. [3]

5.2 Disadvantages of red teaming

Implementing red teaming can be a complex endeavour, whether it's developed internally or outsourced. Outsourced red teaming services lack standards and transparency, which can make them unreliable. Additionally, red teaming can strain relationships within an organization and result in potentially damaging reports [4].

One of the main challenges in red teaming is the cost associated with setting up a team and acquiring necessary equipment. Skilled offensive security professionals are scarce and expensive, making it challenging to find trustworthy individuals with good judgment and communication skills. Many organizations would choose to opt for external red teaming services, known as ‘inorganic’ red teams, due to budget constraints, but even these can be costly and may result in short engagements that yield unreliable findings.

Organizations conducting in-house red teaming, known as ‘organic’ red teams, face their own set of challenges, especially when handling sensitive data governed by regulations like HIPAA and GDPR. Personnel conducting engagements must be well-versed in relevant laws and may require specific certifications.

Contractual obligations with external cloud infrastructure providers can also complicate red teaming efforts, as many agreements restrict testing activities. Special permissions may be necessary to conduct tests in these environments to avoid disrupting services or causing data loss.

Even after the engagement is complete, maintaining professionalism during the reporting phase is crucial. Adversarial behaviour from security staff members can hinder the red team's efforts, as they may attempt to limit the scope of testing or impede progress with defensive measures. Preserving a positive working relationship between red teams and security personnel is essential throughout all phases of the assessment.

Here's a hypothetical scenario: a red team is going to the client's building on a particular date to try to break into the server room, but that date gets leaked to internal staff. A security staff member can simply stand in front of the server room door only on that specific day, and deny access to anyone they do not recognize.

Strange behaviour, but expected when the work performed by red teams directly affects the outlook of the staff's own job performance, often in a negative manner.

5.3 Bigger challenges red teams face

Understanding the contemporary landscape of offensive security is crucial. Various challenges abound, yet a few key areas consistently pose significant hurdles. Red teams find themselves in a perpetual race against the very adversaries they seek to emulate. Moreover, industry standards often fail to accurately depict or facilitate thorough assessments. The inherently adversarial nature of the client-service provider dynamic complicates matters further, potentially leading to conflicts between assessors and clients. Even if these assessment-related challenges are surmountable, staffing issues within red teams persist.

5.3.1 APTs

Red teaming aims to replicate threats to an organization, particularly those that pose formidable challenges. One such example is Advanced Persistent Threats (APTs). Defining APTs proves elusive due to diverse interpretations, but for the purpose of this thesis, they are well-funded entities with defined objectives and organized operations [5]. These encompass nation-state actors like the NSA and CIA in the United States, as well as organized crime syndicates and terrorist organizations. APTs often employ sophisticated tactics, leveraging significant resources to procure advanced tools and personnel, including hackers of exceptional skill. They operate without legal constraints, acquiring tools and data through illicit means, enabling them to execute attacks with impunity.

APTs possess substantially greater resources than red teams, enabling them to invest millions in cutting-edge tools and personnel. Unlike red teams, they can procure illicit tools and personal data, bypassing legal restrictions. Moreover, APTs benefit from intangible resources such as national intelligence and corporate influences, facilitating complex attacks like blackmail, intellectual property theft, or industrial sabotage.

Additionally, APTs operate without temporal limitations. They persistently pursue their objectives, working tirelessly regardless of time constraints. They can probe a single target for months to years at a time, looking for cracks in security. In contrast, red team engagements are typically time-bound, limiting their ability to simulate the relentless nature of APT attacks.

While red teams operate within predefined scopes, APTs disregard such limitations, targeting individuals or systems beyond agreed-upon parameters. They leverage resources to coerce or compromise individuals, expanding the scope of their attacks with impunity. Unlike red teams, APTs exhibit no restraint, exploiting vulnerabilities wherever they arise.

Last but not least, red teams adhere to Rules of Engagement (ROE), guiding their ethical assessments within legal bounds. Conversely, APTs operate without regard for legality, utilizing any means necessary to achieve their objectives. This lack of constraint enables APTs to employ creative and aggressive tactics, amplifying their effectiveness compared to red teams.

5.3.2 In the field

Navigating the landscape of modern red teaming and physical pentesting presents a series of formidable hurdles. These hurdles range from grappling with regulatory standards to confronting a dearth of innovation, as well as contending with industry misconceptions surrounding these practices.

Adhering to regulatory standards poses a significant challenge for red teams and pentesters alike. The rigidity of these standards can impede engagement activities, particularly when those standards are either ambiguous or non-existent. Organizations harbouring sensitive data often impose stringent data usage policies, necessitating careful consideration and customization of engagement agreements. Additionally, compliance with industry and federal regulations, such as those outlined in HIPAA or GDPR, further complicates testing within these networks, demanding meticulous adherence to data protection protocols.

While innovation is inherent to the field of red teaming and pentesting, there exists a conspicuous absence of shared strategies and customized methodologies among organizations. This scarcity of information stems from the protective stance adopted by vendors and organic red teams, who view innovative approaches as proprietary or sensitive information. [4] Academic contributions to process improvement are limited, with a scarce amount of scholarly discourse on enhancing red team methodologies. Consequently, knowledge acquisition relies heavily on experienced practitioners or potentially risky assessments conducted by less seasoned staff.

Misconceptions surrounding red teaming practices abound. One prevalent misconception revolves around the ambiguous distinction between penetration tests and red team assessments. [6] This blurring of definitions leads to misaligned expectations and compromises organizational security. Another misconception pertains to unrealistic expectations regarding assessment timelines and methodologies. Clients often demand short assessment windows without fully appreciating the time-intensive nature of penetration testing and red team engagements. Such unrealistic expectations perpetuate the notion that failure is intolerable, further exacerbating the pressure on red teams to deliver flawless results within constrained timeframes.

5.3.3 Adversarial clients

Red teaming professionals face significant challenges in successfully executing engagements. These challenges involve the need to outsmart, deceive, or uncover weaknesses within client organizations, while maintaining a professional relationship that ensures future collaboration. Clients typically consist of three groups: technical personnel responsible for security administration, managerial personnel overseeing organizational operations, and general users [4].

The technical personnel, including administrators and security staff, often present obstacles to conducting effective assessments. These challenges typically revolve around concerns of embarrassment or professional repercussions. For instance, some technical personnel may attempt to restrict the scope of assessment to avoid exposing vulnerabilities within their direct responsibilities. In other cases, they may actively interfere with the assessment process by targeting the assessors or undermining the significance of discovered vulnerabilities.

Similar to technical personnel, managerial staff can influence engagements, albeit for different reasons. They may seek to limit scope and duration to meet regulatory requirements cost-effectively or direct assessments towards specific areas to secure funding for remediation efforts. Additionally, some senior leaders may disregard assessment reports due to financial constraints or legal liabilities associated with documented vulnerabilities.

While average users may not directly impact testing, their reactions to assessment results can create adversarial dynamics. Users may feel embarrassed if they fall victim to social engineering tactics or face consequences for violating security policies.

Such actions taken by clients not only impede the assessment but also diminish the benefits of red team engagements. Adversarial clients can also strain relationships between the red team and their coworkers, further hindering the assessment process.

In interviews conducted with professional red teamers, one of the better answers provided for how to deal with adversarial clients was as follows:

“It’s easy to remain professional with people like that… you just tell them “I’m not here to embarrass you, I’m here to help you. I’m here to help you secure your systems by convincing your boss, or your boss’s boss, or your boss’s boss’s boss, to give you the money you need to do your job better.” And usually when I explain it that way to people they begin to relax, you can see the visible relief on their faces… just be professional about it all. It costs nothing to keep yourself grounded when things don’t go your way.”

5.4 Conclusion

Awareness of these challenges enables red team assessors to maintain professionalism and effectively communicate the significance of vulnerabilities to stakeholders. Helping clients understand how seemingly minor issues can lead to widespread compromise helps assessors overcome such obstacles and ensure the integrity of the engagement process.

6 How does one get started on red teaming?

Embarking on a journey into the realm of red teaming can seem like a daunting task for aspiring cybersecurity professionals. Contrary to popular belief, there is no singular path or shortcut to becoming a red teamer.

So the question then becomes: How is initiation into red teaming accomplished?

6.1 The myth of instant red teaming

That's the intriguing aspect: It isn't.

It's almost unheard of for people to simply 'get started' on being a red teamer. In most cases it takes years of work and experience to be able to start out in a position where entry into such a role becomes feasible.

According to Marcus Carey [7]:

"It is uncommon for people to start directly into red team jobs. The best way is to have or gain a skill such as internetworking, system administration, or software engineering and start out in a blue team role. Getting into a blue team role will allow you to gain cybersecurity experience and network with people in your dream role. You can network internally and externally from your organization at local events and regional cybersecurity conferences. There are a couple of certifications tailored to red teaming that can get you noticed by red teams looking to add some human resources."

6.2 Building a foundation

According to Marcus, building a strong foundation of skills and experience is essential to starting one's journey. He refers to a "blue team role". Blue teaming, as implied, is opposite to red teaming, dealing more with the defensive aspects of a system rather than investigating how that system can be broken or abused.

Networking is also important. Attending local events and cybersecurity conferences provides many opportunities to connect with like-minded, seasoned professionals, learn from their experiences, and even gain career opportunities in red teaming. Pursuing certifications tailored to red teaming, such as the CRTO (Certified Red Team Operator) or OSCP (Offensive Security Certified Professional) or CEH (Certified Ethical Hacker) can enhance one's visibility and credibility within the community, rendering prospective individuals more appealing to red teams scouting for new talent.

In an interview conducted for this thesis with Jayson Street, a hacker and security engagement professional, one of the most respected in the industry, he shared his backstory, recounting how he started his professional career in the police in a "gang task force" where he worked for years, gaining experience and a solid foothold in the security industry. Eventually, "…after being shot at one too many times…", he began getting jobs from clients for physical security engagements [8].

7 Red teaming mindset

Success in red teaming hinges not only on technical proficiency, but also on the ability to navigate complex human dynamics and decision-making processes. Social engineering is a commonly employed tactic in red teaming where deception is used to manipulate individuals into divulging confidential information. Dealing with people, understanding their motivations, is what makes someone good at manipulating others to draw out their weaknesses and vulnerabilities, but it also makes someone a good team player; someone who can work with others to achieve a common goal. This chapter explores the foundational mindset required for red teaming, and its pivotal role in enhancing security engagements.

7.1 Motivation

If one's motivations for red teaming are impure, they should not be a red teamer.

This is a point often misunderstood by many self-proclaimed 'industry professionals'. Many people have the wrong motivation to work in red teaming, and it causes problems for the entire industry.

The primary motivation for numerous red teamers is that they get to break things legally. They get to hack into systems or break into buildings with permission, and they love doing it because it makes them feel 'cool' or it gives them a sense of satisfaction to break into things. However, in this industry, motivations should ideally stem from a desire to fix clients' systems. The focus should not solely be to exploit vulnerabilities, but to help clients fix these weaknesses so nobody else can exploit them.

Here is an excerpt from the interview with Jayson Street:

"People don't understand, red teams don't work for themselves. You only exist as a red teamer to make the blue team better. And if your blue team isn't improving, if your client doesn't improve from one year to the next, that's because you suck. They're not paying you to break in, they're paying you for the report… You're not a rockstar, you're not a ninja… You're an auditor and tester that uses a different means of tools and a different set of methodology. You only exist because the blue teamers care about their defenses and they care about their networks and they want to secure their systems. And you're an advocate, not their adversary.

… We need to get over this whole idealization of red teamers because they get to go and commit crime or get to rob people. And I like to say I rob banks for a living because it sounds cool, you know? I won't lie about that… But no, in all actuality though, I am never going to give that kind of arrogance to a client or to anyone else in earnest, because you are doing a service. You are trying to humbly improve the defense in a corporation and protect them, and if you're not there to make them get better and make sure they're better protected and make sure they can be defending themselves against attacks then you're wrong and you're bad at your job.

It's never about "Oh I've got to find a vulnerability" or "I've gotta break in or they're not gonna think I did my job." That's not what they paid you for." [8]

7.2 Self-awareness

Through the cultivation of self-awareness, a better understanding can be gained of one's own tendencies, biases, and influences, which can help enable more objective evaluations and decisions.

The human factor plays a significant role in pentesting, especially physical pentesting and social engineering. Cultivating a healthy, empathetic, professional mindset is deemed important for maintaining good relationships with the people being worked with, despite the consistent adversarial actions taken against them.

7.2.1 Understanding undesired/unproductive tendencies

The journey to understanding undesired tendencies begins with a deep dive into self-awareness. It involves introspection, self-authorship, and reflection to uncover patterns of behaviour and emotional responses that may hinder one's effectiveness as a red teamer. By reflecting on one's own actions areas for personal and professional growth and improvement can be identified.

7.2.2 Understanding personal biases and assumptions

Biases are inherent to human cognition. These biases manifest in various forms, such as confirmation bias, where people unconsciously seek out information that confirms their preconceived beliefs, or anchoring bias, where they rely too heavily on initial impressions or information.

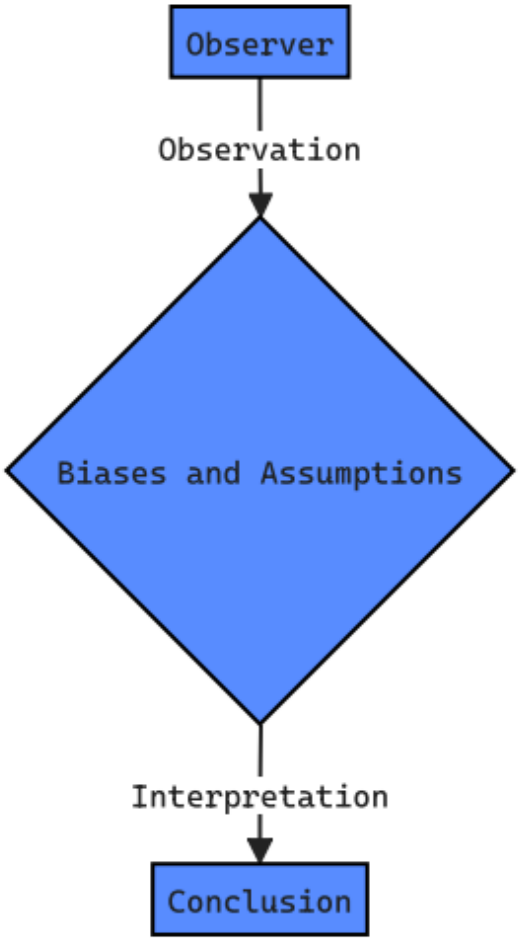

Figure 1 – Recognizing personal bias

CASE STUDY: BIAS

Let's look at an example inspired by a recent security engagement that one of the interviewees took part in, which shall remain undisclosed: The group was tasked with assessing the security of a large office building. During the reconnaissance phase, the interviewee noticed that the main entrance door appeared to be poorly secured, with a loose door frame and visible gaps around the edges. However, upon further investigation, him and his team discovered that the door was equipped with a state-of-the-art access control system, including biometric authentication and electronic locks which weren't visible from the pictures he gathered through social media.

Despite the evidence he had suggesting the door was adequately secured, the tester's initial observation of the door's physical condition had already planted a seed of doubt in their mind. This bias led them to focus disproportionately on finding vulnerabilities related to physical access, while overlooking other potential security risks in the building's digital infrastructure, or exploiting other simpler avenues like social engineering.

7.3 Self-authorship

Self-authorship is the ability to define and construct one's own identity, beliefs, values, and goals autonomously, rather than being solely influenced by external factors such as societal expectations or peer pressure. In the context of red teaming, it is an important concept for navigating the complexities of dealing with other humans.

Red teamers need a strong sense of self in order to make ethical decisions, challenge conventional thinking, overcome the throes of impostor syndrome, and innovate in the face of ever-evolving threats. They are tasked with exploring uncharted territory, often needing to devise clever strategies and adapt to dynamic environments, all of which require an understanding of their own values and limitations.

Robert Kegan introduced the concept of self-authorship, offering a comprehensive framework for enhancing self-awareness [9].

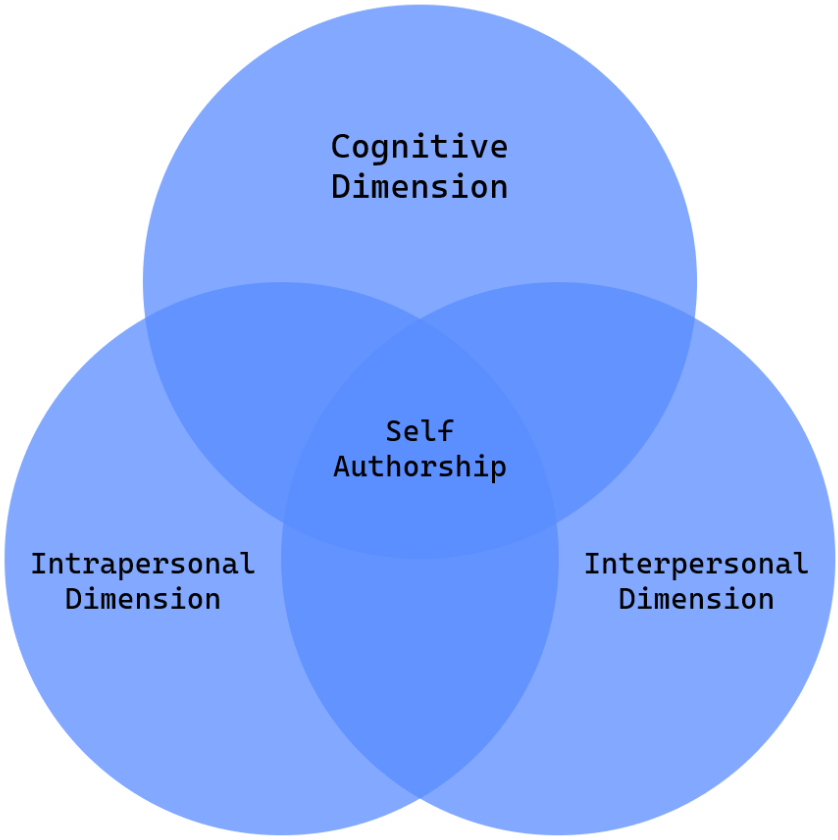

Figure 2 - Components of self-authorship

1. Cognitive dimension: Our knowledge is shaped by both social constructs and personal experiences.
2. Intrapersonal dimension: This encompasses our individual beliefs, values, and aspirations.
3. Interpersonal dimension: The perspectives of others.

These dimensions collectively encompass one's personal knowledge, their philosophies, and interpersonal interactions.

7.3.1 Temperament

Temperament is a habitual, observable facet of personality. Decisions are made, communication occurs, and priorities are set through the lens of temperament [10]. This is known as the 'comfort zone'.

Temperament can be mapped onto a continuous scale, along a continuum between introversion and extroversion .

Introversion is characterized by a tendency to focus inward, reflecting on internal thoughts before expressing them outwardly. Introverts often prefer solitary activities and find quiet contemplation to be conducive to their processing.

On the other hand, extroversion involves directing attention outward, seeking stimulation from the external environment. Extroverts thrive on interaction with others, often thinking out loud and engaging in collaborative brainstorming sessions.

Within the realm of temperament, there are layers to consider:

1. The contextual self: How we behave in specific situations
2. The developed self: Learned behaviours and skills acquired over time
3. The core self: Innate predispositions inherited from birth

Temperament is important for red teaming because it shapes how individuals perceive challenges, make decisions, and react to threats. Understanding different temperaments helps red teams tailor their approaches and predict reactions to their actions.

CASE STUDY: TEMPERAMENT

Here are some notes taken out of another red teaming engagement the same interviewee was a part of:

His team was hired by a financial institution to conduct an exercise assessing its cybersecurity defenses, as well as its incident response processes. Their aim was to identify vulnerabilities in their digital networks, as well as to test the behaviour of their employees.

The team's aim was to simulate a data breach of sensitive customer information. Their main tactics were social engineering which included a mix of phishing and malware infiltration.

Here are some outtakes from their notes that could be considered good analyses of temperament:

1. The team observed how certain employees reacted to the phishing campaign. They then spoke to the employees afterwards, through individual questionnaires that asked about their security awareness. The team found that introverted employees hesitated to report suspicious activities out of fear of reprisal, while extroverted employees tended to overlook the subtle indicators of a phishing attempt.
2. The team found that employees who were aware of their security behaviours, or those who had undergone training in the past, exhibited more cautious behaviours regardless of their temperament.
3. The team discovered that understanding the mannerisms of company executives and IT personnel was also crucial. Leaders/management with different temperaments had to be told different things in the reports in order for them to fully grasp the importance of the findings. As an example, introverted leaders tended to prioritize robust technical defenses such as stringent access controls, while extroverted leaders emphasized cross-departmental communication and collaboration to mitigate risks.

7.4 Emotional intelligence

Emotional intelligence is the adept recognition and comprehension of emotions within oneself and others, and the utilization of this insight to navigate interactions and relationships. This encompasses four key dimensions:

1. Self-awareness: Emotions are recognized and understood as they arise, along with their influence on decision-making and outcomes during engagements. Additionally, recognition of how emotional responses impact rational thought processes.
2. Self-management: Emotional reactions are comprehended and used to shape behaviour and performance in a positive manner. This may involve the development of coping mechanisms to combat work stress and pressure from tight deadlines.
3. Social awareness: Perceiving and acknowledging the emotions of others, practicing cognitive empathy, empathizing with others' perspectives during an assessment. This includes active listening, asking pertinent questions, interpreting non-verbal cues to understand others' emotions, and more.
4. Relationship management: This focuses on accurately interpreting social dynamics and effectively interacting in different professional environments. It encompasses skills such as persuasion, negotiation, cooperation, and teamwork, all useful skills for a red teamer to have.

Good emotional intelligence plays a pivotal role in improving team dynamics, communication with clients, maintaining professionalism in an adversarial role, and effectiveness in navigating the complex human factors involved in activities like social engineering.

7.5 Interpersonal communication

Effective communication is not merely about conveying information; it's about engaging in a nuanced exchange that delves into ideas and insights. According to the *Red Team Handbook* [10], mastering interpersonal communication for red teaming goes beyond speaking; it entails three modes of listening:

1. Strategic listening: Seeking new information to facilitate choices or open a space for new ways of talking about a problem, using open and close-ended questions.
    a. Consider when to inject questions
    b. Ask clarifying questions and offer paraphrases
    c. Be on the lookout for discoveries
2. Empathic listening: Showing concern in support of emotions. Helping others feel safe and understood. Its absence can suggest impatience, disinterest, or even dismissal.
    a. Ask indirect questions to echo pieces of what someone says
    b. Don't interrupt
    c. Pay attention to the person's face (not just their mouth)
3. Active listening: Showing involvement and respect to foster social relationships. It is measured at the perceived quantity and quality of one's interest. Its absence may show a lack of concern or importance.
    a. Acknowledge what the other person is saying without interrupting
    b. Keep eye contact or pay attention to their face
    c. Expand on parts of what they are saying

7.5.1 Interpersonal conflict

Navigating conflict with clients necessitates active engagement from all parties involved, yielding mutual advantages. Taking a moment to reassess perspectives and ponder the differences and commonalities among team members' or clients' dispositions can unveil the underlying perceptions that sparked the conflict. This introspective process can then be utilized to formulate strategies aimed at bridging divides and de-escalating tensions.

7.6 Introspection

Self-reflection is a useful method for enhancing self-awareness. By turning focus inward and eliminating external distractions, the door is opened to a deeper understanding of one's own self, their thoughts, and behaviours. Daily journalling is a good tool for introspection.

Personal thoughts, revelations, inquiries, and notable occurrences should be written in this journal. The entries in this journal aren't just a mere recounting of the day's events; they demand a profound and deliberate examination of the day's themes.

CASE STUDY: CONFLICT RESOLUTION

In a high-profile security engagement, a red team was tasked with testing the security of a large law firm. The objective was to assess the susceptibility of employees to phishing attacks and other social engineering tactics. The engagement began smoothly, with the team employing various pretexting techniques to gather sensitive information on employees.

However, tensions arose within the team as conflicting personalities clashed over the future direction of the engagement. One team member, let's call him Alex, advocated for a more aggressive approach, pushing for rapid escalation and exploiting vulnerabilities as soon as possible, citing the approaching deadline for the assessment as a prime factor for his reasoning. Another team member, let's call her Sarah, favoured a more cautious approach, emphasizing the need for the team to blend into the network traffic carefully so as not to get caught.

As the engagement progressed, these differences in opinion led to heated debates and conflict within the team. Alex became increasingly frustrated with what he perceived as Sarah's overly conservative approach, while Sarah felt undermined and disrespected by Alex's domineering behaviour.

The situation reached a breaking point when a critical phishing email crafted by Alex inadvertently triggered a security alert, alerting the organization's IT security team to the pentesters' activities. Panic ensued as the red team realized they had been discovered, and blame was quickly assigned to Alex for his recklessness.

In the aftermath of the incident, conflict resolution techniques were needed to address the fractured dynamics within the team and salvage the engagement. The team leader, let's call him Harry, instead chose to ignore the issue and move on with the assessment in a vain attempt to salvage it. The team was not able to identify the root causes of their conflict, and a further breakdown in communication ensued, leading to a larger clash, the effects of which rippled through the rest of the engagement and caused the team to deliver underwhelming results. This caused Alex to resign from the team indefinitely.

What Harry should have done was call for a temporary halt to the engagement to allow for an open discussion about the underlying issue. During this session, which could be facilitated by an external

mediator, team members would be encouraged to express their concerns and grievances in a constructive manner. Active listening techniques could be used to ensure all perspectives are heard and understood. Through this process, the team could then identify the root cause of the issue and solve it.

That newfound understanding and collaboration would be a night-and-day difference compared to the underwhelming results they actually delivered.

7.7 Applied Critical Thinking

In our daily lives, our minds are constantly at work processing information and making decisions. Despite our experience in thinking thoughts, those thoughts can often lead people astray due to their own biases and assumptions. This can result in people making decisions that are unaligned with their objectives.

7.7.1 What is applied critical thinking?

Applied critical thinking (ACT) involves intentionally evaluating the strengths and weaknesses of thoughts and reasoning. Understanding and decision-making abilities are improved by analyzing our perceptions and interpretations of the world. This includes making implicit thoughts explicit, allowing assessment of their relevance and suitability to the situation.

7.7.2 Considering time

In situations where decisions are time-sensitive, often the case when pentesting an organization, incorporating ACT may initially seem impractical. However, by consistently applying ACT in our lives, reflexive methods to evaluate and adjust thinking over time are developed. Additionally, structured tools designed to fit within time constraints can support ACT. Whether facing tight deadlines, implicit timeframes, or self-imposed pressures, time is often the limiting factor that influences decision-making. Our minds tend to resort to shortcuts in these scenarios, leading to assumptions and settling for suboptimal solutions.

Red teamers can address these cognitive shortcuts by employing alternative approaches to understand different perspectives. Tools like the 5 W's (Who, What, When, Where, Why) can facilitate rapid examination of thought processes. Red teamers can learn to use ACT in their daily lives in order to recognize and adapt to different problems in their lives, leading to better solutions; here is another excerpt from the University of Foreign Military and Cultural Studies *Red Team Handbook*:

"One problem red teamers face often is 'cognitive autopilot', that is using known/easy solutions for problems faced before, using solution B for problem A. But because of the complexity of the world around us, multiple instances of problem A might not always have the same cause or characteristics. We fail to notice then that we are not actually dealing with the same problem and thus solution B will not work on it. By using ACT, we can identify differences between problems and alter our responses and apply more appropriate and effective solutions.

8 Scoping

Setting up a scope for an offensive security engagement means determining what will be assessed and when the assessment will happen.

These two attributes are tied closely to each other, and constraints on one affect the feasibility of the other. Shaping the “what” of the assessment scope is driven by the perceived or actual needs of the customer. The “when” refers to the schedule and window for assessment, and is affected by the availability of resources. Resource limitations impacting the “when” of the assessment are typically financial in nature from the customer.

8.1 Preventing incidents

Red team activities mimic real attackers and can be easily mistaken for real attacks. To minimize wasted response efforts on part of the customer, the assessment window needs to include not only the begin and end dates for the activity, but also those days during the week and the hours during the day which ethical hackers will be active. “Activity” can refer to human-involved attacks, enumeration, covert-entry attempts, automated functionality of red team tools, etc.

Aside from preventing wasted resources and a ton of frustration, a clear understanding of when scoped targets will be assessed is important to meet the needs of the customer as comprehensively as possible. What is being assessed may determine the schedule for assessment. The customer may tell the team they need a particular data center assessed and the team will take, for example, six weeks to assess and report on that target adequately, so six weeks is the desired amount of time for the engagement. Unfortunately scoping discussions almost never go like this. Typically it goes as follows: the organization has resources for four weeks of red team services and wants the same data center assessed. The limiter could be the resources or funds needed to pay for four weeks of services only, or the organic red team has a four-week window only available for the data center assessment because of other obligations.

8.2 Who

Incorrect scoping can both destroy any chance for successful engagements, as well as ruin working relationships between staff in the case of organic red teams or business relationships in the case of third-party red teams. It is important the right personnel be present or involved in scoping decisions so the assessment has the highest chances of meeting the needs of the customer, and that it is within operational capacity of assessment resources. An ideal scenario is that the red team and the customer have representatives from the technical and operational/managerial functional areas present during the scoping process.

8.3 When

The “when” of the scope is the easiest part of the scope to identify. The time periods involved in scoping are the start and end dates allotted for assessment and, if the client specifies, the time schedule for those dates. At a high level, the assessment window is simply the period of time when

the assessors test what has been determined to be approved as targetable. There is a need to be granular in identifying the assessment window.

8.4 What

The client organization is the pivotal factor in determining the scope of any engagement. However, aligning these needs with the actual requirements for assessment can be challenging. Involving technical and non-technical stakeholders in scope discussions leads to productive dialogue, ensuring an approach which will address the organization's needs while optimizing the use of their limited resources. To achieve this, the customer needs to answer key questions to guide scoping conversations:

- What is the motivation behind the assessment?
- Does the organization have any prior testing experience?
    - If so, what is the organization's prior experience with testing?
- What are the main outcomes the organization is looking to achieve through this test?
- What level of maturity characterizes their security infrastructure?
- What level of testing is the organization looking to conduct?
    - Assumed breach vs specific compromise vs holistic compromise, etc.

8.4.1 Assessment motivation

The organization's motivation for hacking assessments also impacts the scope of the engagement. These motivations stem from planned, scheduled, or unplanned events:

- Planned events: Orchestrated by the organization with specific objectives in mind, often tied to enhancing security posture. For example, a company might want to test a new, expanded physical security system that's been set up on their premises.
- Scheduled events: Externally mandated, driven by regulation or compliance requirements. Companies subject to HIPAA or NIS2 must undergo both physical and network security audits to ensure adherence to data handling requirements and physical security standards.
- Unplanned events: These present unique challenges. These events look like sudden audits, network disruptions due to a persistent attack, or security breaches like a break-in.
    - Unplanned events are the only kind of event that deal with post-incident measures, and thus is technically not a form of proactive security.

8.4.2 Prior testing experience

Determining prior testing experience could provide valuable information in understanding clients' needs. This relies heavily on their willingness to share such information, however.

Asking for detailed information, focusing on the who, what, when, where, and why of the prior tests, helps in gauging the effectiveness of those tests. If the assessment happened years ago, its relevance would probably be quite limited. But if it took place recently, focusing on overlooked systems from the prior assessments can maximize coverage of attack surfaces within limited timelines.

If the organization is willing to provide the reports from these tests, and assuming those reports were written well enough, it can streamline testing procedures.

It helps to know which organization wrote the report. Some organizations will run Nessus on a network, copy-paste the scan results, and present it as a 'penetration test'. A significant problem in the hacking space, and the IT field in general, is that many people skilled with a computer struggle with writing proper reports and documentation. Conversely, there are those who can write proper reports, but often have subpar hacking abilities.

8.4.3 Intended outcomes

This one is fairly simple: Understanding the goals an organization are looking to accomplish with a test will help to streamline the entire engagement by focusing on what the organization thinks is important.

Some organizations will state that a test is being conducted solely to check a regulatory box. Other organizations will mention that the test is meant as a response to an incident that already occurred. And a rare, golden few will convey that red teaming skills are needed because they genuinely want to understand how their existing, or new, security apparatus is holding up.

No matter what the motivation is, red teaming is still an arduous process when the motivations behind the exercise aren't made clear from the outset. Asking this question is vital to understanding where the scope of a project ought to lie.

8.4.4 Current security apparatus

When embarking on a red teaming engagement, team members need to gauge the level of security within the organization. This involves delving into existing security measures and procedures. It is good to know, then, what the existing security procedures and measures are. Is there a proper backup system in place for critical infrastructure? Who has access to the server room, and how does the organization control that access? Are there CCTV systems in place, and if not are there plans to set one up? What regulations or directives is the company subject to (HIPAA, GDPR, NIS2)? Who has access to critical systems? Does the company delegate some of their security responsibilities to third parties? What is their procedure for handling paper files, shredding documents, disposing hard drives, etc? Are employees allowed to bring their own devices into the building? Are employees able to bring USB drives home from work? What is the password policy?

And many, many more questions that should be asked in order to understand the security of the company and tailor the assessment properly. If the organization does not know the answers to these questions, then that should be mentioned in the findings during the reporting phase. Based off these answers, red teamers must also proceed with caution and adapt the approach they take during the engagement as necessary. For example, if it is made clear that the organization's password policy is not robust enough to defend against brute forcing, that should be an avenue of exploitation for the team to consider **if it is within scope**.

8.5 Further considerations

Another thing to consider with scoping is the overall subset of the security apparatus that will be tested. This is called the **footprint** of the scope. This determination is done by considering factors such as the **timeframe** and the **available targets**. While customers may, at first, request assessment of externally visible targets, relying solely on this criteria can lead to overlooking certain assets.

It's crucial to clarify whether internal hosts are included in the scope and whether pivoting between hosts is allowed. Many customers express discomfort with pivoting from initially compromised hosts.

This reluctance is ironic considering most compromises originate internally rather than from internet-based attacks. Red teams need to grasp the footprint of the scope thoroughly, even if customers cannot provide a detailed list of specific hosts to attack. Flexibility in adjusting this scope footprint may be needed to overcome constraints and ensure success.

External constraints, such as those imposed by network-based datacenters or cloud service providers, also affect the engagement scope. Cloud service agreements require notification and approval for testing against hosted systems, and some activities are strictly forbidden.

In other cases, organizations can have connections to third parties, governed by legal agreements delineating responsibilities and boundaries. For example, a public library or academic software provider or other academic-related organization could have dedicated connections with universities for collaborative purposes. Failure to consider these agreements and connections during scoping can lead to unintentional engagement with external entities or damage to unprepared systems. This is vital, as dealing with it early on will prevent illegal activity and ensure assessments target expected assets only.

8.5.1 The reverse scope

In many red teaming engagements, a lot of companies want to give red teams the freedom to use whatever tools and methods are at their disposal in order to most accurately simulate a real threat. Thus, some organizations would want to consider filling in a **reverse scope** instead of a normal scope.

A reverse scope is the practice of **excluding** the targets and methods on which the engagement is **not authorized** to do. Instead of have an extremely long list of all the methods the red team can use (brute force, phishing, tailgating, …) and all the systems that can be targeted, the company instead lists the few methods the red team cannot use during the engagement, common examples are denial of service, lockpicking, destruction of property, as well as listing the few systems they are not allowed to target.

9 Rules of engagement

The **ROE** delineates the **operational framework** for conducting assessments within the scope the red team defined after the initial planning phase. These rules establish the legitimacy and legality of all actions taken, safeguarding against legal statutes in place protecting people from hacking activities.

Both the customer and the red team need to acknowledge and sign a comprehensive ROE document prior to engagement. Legal advice should be sought during the drafting of this ROE to ensure its adequacy and compliance with relevant laws. A good recommendation for all red teamers is to have a template document set up with the help of a legal professional, that can then be revised and adjusted before engagements occur.

The ROE also fulfils more functions beyond legal significance. It **grants authorization** to red teamers to commence their activities without fear of legal repercussions, providing a safeguard against potential hostilities from the client. Copies of this document should be retained by both parties for reference and protection in case of disputes. Secondly, the ROE shields the service provider from unwarranted liability arising from damages incurred during the assessment, thereby mitigating financial risks. Thirdly, it outlines the responsibilities of the client organization in safeguarding its data and assets during and after the assessment.

Moreover, the document defines the parameters for addressing instances of gross negligence, encompassing actions that may result in harm to the client. This includes breaches of confidentiality, mishandling protected data, permanent damage to the premises as a result of covert entry tests, emotional or psychological damage suffered by employees as a result of negligent social engineering practices, and any actions deemed beyond the scope of ethical practices.

9.1 Activity types

Some activities within the ROE ought to be permitted. These activities can be categorized into various types of offensive security engagements:

- Physical infiltration
- Social engineering tactics
- Pentesting external networks
- Intrusion into internal networks
- Pivoting between systems
- Exploration of wireless network traffic

It is important to very clearly outline these activities and their parameters in both the scope and the ROE, so that both the client and the red team have an understanding of the nature of the engagement. This is done in order to ensure transparency and minimize unexpected outcomes for everyone involved.

9.1.1 Physical infiltration

Engaging in physical security assessments carries the largest risk for both the assessor and the client, often leading to its exclusion from engagements. Nevertheless it is an important area to cover, especially that it is now a required component of compliance with the NIS2 directive. While many

red teamers would desire to conduct physical break-ins, few possess the necessary expertise. The potential for physical damage and injury further complicates justification for including these activities.

Physical engagement falls into three categories:

- No-tech
    - Tactics like shoulder-surfing or tailgating which exploits human vulnerabilities without the use of technological aids. Leaving doors unlocked after hours also falls into this category, which can reveal weaknesses in policy adherence.
- Low-tech
    - This requires basic tools like lockpicks to breach physical security. These actions can pose lasting damage to systems, making them admittedly less appealing to clients who rely on legal and law-enforcement safeguards.
- High-tech
    - This makes use of electronic devices such as hardware keyloggers or network taps to gather information and facilitate cyber operations. These methods are comparatively low-risk and offer the most direct insights of a company's cybersecurity posture.

The more common ROE-defined activities will likely fall under either no-tech or high-tech, as they offer comprehensive approaches to evaluating security mechanisms, without the risk of lasting damage to those mechanisms.

9.1.2 Social engineering tactics

This category can also fall under physical security, depending on how it is done. If social engineering is done face-to-face, such as donning a worker's outfit and pretending to be performing maintenance on-site, then it falls under physical security as a no-tech activity.

No matter how social engineering is conducted, its overall aim is to manipulate targets into divulging information or performing actions that compromise an organization. Implementing such tactics during an assessment can be intricate. For example, tricking personnel via email, phone calls, texts, etc. can be tricky due to unforeseen complexities.

Consider the example here:

Approval is granted to attempt to convince employees to disclose their login credentials by posing as IT support personnel. This activity, commonly known as pretexting, evaluates the effectiveness of employee awareness training and authentication protocols. The team crafts a convincing pretext and starts calling employees, successfully obtaining login credentials from several individuals. However, as the day progresses, they begin to use these credentials and encounter a problem: one of the emails they were given was linked up to the employee's personal email address, and all the emails sent to their personal address are also sent to the work address, and are thus visible to the red team. What they've done is essentially gained access to the employee's personal emails, despite only desiring to compromise the work account. The team has at worst committed a crime and at best violated the terms of the ROE agreement.

Clear ROE is essential to avoid this, but getting organizations to accept liability for personal damages can be contentious. Beyond accidental breaches, phishing methods like spear-phishing and whaling pose additional risks. The latter targets high-ranking individuals, which can strain client relationships. Careful consideration is needed when drafting ROEs to avoid jeopardizing the entire assessment.

9.1.3 Pentesting external networks

External network testing is a common practice in red teaming. It involves conducting cyberattacks from outside the client's system boundaries. This targets external-facing assets. Assessors must specify the source addresses for attacks in their ROE document to distinguish red team activity from actual threats quickly. To avoid complications such as violating ISP agreements or cause blacklisting, a separate external infrastructure should be utilized.

Physical pentesting also falls under this category.

9.1.4 Intrusion into internal networks

Internal network testing refers to cyberattacks originating from within the target organization's network, aimed at internal assets. Convincing clients of the importance of internal network testing is a challenge, because there is a strange myth that systems lacking external access are invulnerable to attacks, and so are not valid targets. It's crucial to remind clients that a **significant portion** of cyber breaches stem from social engineering and insider threats, both of which will exploit internal network access. Opting for internal assessment, particularly within tight timeframes, yields cost-effective results. Most of these assessments begin with unprivileged access, often simulated through successful social engineering scenarios.

Assessors will initially operate with standard user-level access, progressively targeting other internal assets. Clear ROE is essential here, as assessments beginning within the security perimeter can rapidly spread to unexpected areas of the network.

9.1.5 Pivoting between systems

Pivoting plays a crucial role in the ROE. Pivoting encompasses two distinct actions. Firstly, it involves leveraging acquired access to penetrate deeper into the organization's network, targeting more devices.

Secondly, pivoting also encompasses exploiting one application to gain access to another and escalate privileges.

Red teamers have to keep in mind that many organizations will restrict or just outright deny such activities, preventing them from using access acquired in one device or application to attack another device or application.

9.1.6 Exploitation of wireless network traffic

Wireless attacks mirror physical infiltration techniques, requiring specialized expertise and tools. This activity introduces additional risk, necessitating clear guidelines within the ROE. Wireless attacks can be categorized as follows:

- Passive interception: Collecting traffic to crack encryption or identify credentials.
- Active exploitation: Targets vulnerabilities, such as insecure IoT devices, for data exfiltration or manipulation.
- Denial of Service: Disruption of, specifically, wireless service to favour the red team.

Incorporating wireless attacks into the ROE demands specifying the nature of operations and targetable technologies. Red team members must adhere to regulations governing wireless transmissions.

Considerations also arise regarding personal devices within the client organization's network. Care must be taken not to intercept or disrupt personal devices, or those devices outside the scope of the assessment. Wireshark, a popular network traffic capture tool, has options to specify a capture filter to limit the devices being intercepted.

9.2 Engagement type

In addition to defining activities in the ROE, it's important to categorize the type of engagement that will be done. Engagements typically fall into black-box, grey-box, or white-box testing.

Black-box testing involves minimal information on the target, sometimes only getting the target's name. While it can mimic a realistic attack scenario, it poses time constraints and risks exceeding legal boundaries.

Grey-box testing involves partial information, often including addresses and metadata, reducing some risks associated with black-box testing.

White-box testing offers comprehensive knowledge of the target organization. A list of information is provided to the red teamer, a blueprint of the organization's building is given, a network map is drafted for extra information on the host systems. Though common in red team scenarios, it's not inherently negative. Insider threats, for instance, would possess similar in-depth knowledge of the client's systems. White-box testing is as valuable as black-box testing in most cases.

9.3 Incident management

Establishing a chain of command is vital for handling incidents effectively. Incidents can be classified as one of two categories:

1. Discovery of illegal activities within the organization
2. Organization-specific illicit activities

Let's assume that evidence suggesting an employee in the organization is keeping encrypted CSAM files in their systems is uncovered. The ROE guidelines on reporting procedures must be followed. This is important because the client organization cannot pursue legal action against assessors if they report suspicions of illegal activities directly to the authorities instead of informing the organization first.

Consider this scenario: Evidence of drug trading by one of the company's executives is uncovered. The organization could attempt to prevent reporting by issuing a cease and desist order based on confidentiality agreements outlined in the ROE. However, a clause in the ROE regarding reporting illegal activities take precedence, ensuring assessors know their obligations upfront. Additionally, the ROE should detail the reporting expectations regarding potential illicit behaviour by organization personnel, such as sexual harassment or policy violations. All incidents, whether operational or security-related, should be reported through points of contact established for escalation purposes during engagements.

9.4 Tool usage

Defining the tools to be permitted for use is essential to prevent negligence or disruption. Both parties are protected this way.

As long as it's agreed upon that disruption is allowed in the ROE, any tool that accidentally causes disruption is still permitted. This ROE guideline also prevents the red team from resorting to risky tools out of frustration.

9.5 Certification requirements

There may be certain requirements, especially for red teams handling sensitive data, for assessors to hold certifications for handling said data. Data protected by GDPR may require someone certified with handling that data to operate on that system. Assessors would be approved on a case-by-case basis, with specific certified individuals being listed within the ROE if these requirements are necessary.

9.6 Personnel details

The ROE should include the identification and contact details of all personnel involved in the assessment from both parties.

10 Execution

When it comes to executing red team engagements, it is important to understand the professional methodology behind these activities, on top of knowing the actual activities themselves. This involves approaching assessments with professionalism; simply using sophisticated exploits and tools without professionalism yields little value.

Professionalism is especially important to keep in mind. The allure to stray from ethical or emotionally appropriate conduct is strong, but it can result in reputational damage, or lead to illegal actions that fall out of the scope of the assessment, which can nullify any potential benefits of the assessment. Immature actions undermine the credibility of the entire team. Especially when dealing with adversarial clients, the temptation to embarrass the security personnel must be resisted, with professionalism taking precedence. In addition to ethical and emotional awareness, adhering to best practices, employing effective tradecraft, and maintaining thorough operational documentation are essential for professional red teamers.

Here are some measures that can be taken during the execution phase to ensure a successful engagement…

10.1 Keep the ROE in mind

At the outset of the engagement window, it's important to review the ROE document meticulously to confirm the approved activities and scope for the target set. This final check ensures the planned engagement is compliant with regulations and it also serves as a safety precaution.

Even before starting any activities, whether it be a physical pentest or a network pentest, keeping a signed copy of the ROE at hand is vital. That contract serves as the only safeguard. In the event things go wrong, such as getting caught breaking into the premises during a physical engagement, it's good to have the contract readily available to show so that the job doesn't result in immediate arrest.

10.2 Notify constantly

At the onset of operations, every day, the red team must notify the client of their engagement. This practice should also be repeated at the end of the day. Regular communication serves many purposes, mainly aiding the swift resolution of any conflicts arising from identified indicators, whether originating from the red team or actual malicious actors.

Consistently updating the client on daily work durations fosters transparent communication, ensuring awareness of the team's active status within the premises. Given how remote this kind of work can be, maintaining this perception of diligence minimizes doubts from the client.

Any post-operational impacts also need to be mentioned to the client, particularly script-based exploitations. This includes activities such as leaving USB drops in parking lots, where an employee working a late-night shift might decide to pick it up off the ground and plug it into their workstation after hours.

10.3 Prioritizing the team's activities

Effective red teaming relies on **mastering prioritization**. It's about knowing when to push boundaries without crossing the line into recklessness. Skilled red teamers can tread this line carefully.

For instance, consider the decision to run an exploit like Log4J, which carries the risk of system corruption if bad code is deployed, only when no alternatives exist, and after coordination with the client. On the other hand, hastily exploiting a SQL vulnerability without exploring less risky options that don't carry a risk of database corruption is imprudent.

Thus, prioritizing activities is crucial. Excessive caution impedes progress, overly aggressive tactics jeopardize the engagement. This is a problem even with seasoned assessors, who find themselves drawn into **rabbit holes** – time-consuming activities that detract from the overall objective of bolstering security.

10.4 Overcome fear of failure

Red teamers must learn to resist the fear of elitism and the **fear of detection**. Striving to remain undetected shouldn't overshadow the goal of enhanced security. Skilled red teamers will know how to blend into the network, mimicking normal user behaviour, and leveraging available protocols and credentials to minimize suspicion. Novice red teamers will learn this skill through time and experience, one should not expect to be a professional from the outset.

10.5 Keep notes! Constantly!

This is the only chapter title with exclamation marks, because it's something even seasoned red teamers get wrong constantly. Keeping a record of activities is invaluable to the entire engagement. The quality of the notes taken determine the quality of the report. Moreover, good notes also shield the team members from allegations of misconduct. Detailed notes will run through all the activities performed on the tested systems, with timestamps.

Thorough notes expedite analysis of system artifacts, crucial in team assessments. Standardized note-taking ensures everyone remains informed and aids in identifying potential issues caused by fellow assessors. Sharing notes with the client organization, if they request it, can also accelerate threat identification.

Operational notes encompass five main stages: enumeration, exploitation, post-access, system manipulation, and exit strategy.

Teams will pick and choose these depending on what their plan of attack or testing strategy is. Some teams even run through the same phases multiple times in one audit. For example, performing reconnaissance on a system, gaining access to that system, and then performing more reconnaissance as a post-access activity to identify new systems to pivot to.

Below is an example of notes taken from a physical pentest:

Enumeration

10:05 AM 11/4/2024: Conducted a thorough observation of the building exterior. Noted security cameras positioned at main entrance and side entrances, with a wide field of view covering potential entry points.

10:15 AM 11/4/2024: **Identified potential entry points: loading dock door and side emergency exit**. Dock door appears to have weak security measures, standard garage door, and no visible surveillance.

10:23 AM 11/4/2024: Observed employee badge usage at main entrance. **Noticed some employees holding doors for others**, indicates potential lapse in badge security enforcement.

Exploitation

11:10 AM 13/4/2024: Tailgated employee **through the side entrance**, asking them to hold the door open for us. Leveraged casual conversation to deflect suspicion.

11:18 AM 13/4/2024: Distracted security personnel stationed near server room by **posing as a delivery person with a package for the IT department**. Engaged in small talk to prolong the interaction and allow team members to access restricted areas.

11:36 AM 13/4/2024: Utilized pretexting, **posing as an IT contractor**, to gain unsupervised access to the server room. Exploited the assumption of authority to bypass verification procedures.

Post-access

12:03 AM 13/4/2024: Inside the server room, observed network switches and server racks, noting absence of physical controls on some equipment.

12:19 PM 13/4/2024: Connected to an **unsecured network port** using a preconfigured network tap to blend in with legitimate traffic. Conducted passive reconnaissance to identify critical systems and sensitive data repositories.

[list of systems identified]

[…]

12:32 PM 13/4/2024: Used gathered information to map out network topology and locate high-value targets such as file servers and database servers hosting sensitive information.

Server exploitation (post-exploitation system manipulation)

1:04 PM 13/4/2024: Exploited a known vulnerability in an **outdated version of Windows OS** running on a file server to gain **unauthorized access and escalate privileges**. Full activity log below:

1:00 PM 13/4/2024 from 192.168.23.6: nmap -O 192.168.23.2

Aggressive OS guesses: Microsoft Windows 10

1:02 PM 13/4/2024 from 192.168.23.6: nmap -Pn -p445 -script smb-vulnms17-010 -v 192.168.23.2

smb-vuln-ms17-010: VULNERABLE
…
…
…

1:04 PM 13/4/2024 from 192.168.23.6: msf exploit with port 5555 against endpoint 192.168.23.2 with a reverse_https meterpreter payload:
[+] 192.168.23.6:5555 – Exploit successful. Meterpreter session opened.

Exit strategy

1:17 PM 13/4/2024: Ensured all physical evidence of unauthorized access was removed. Closed server room door and reconnected any disconnected cables to maintain appearance of normalcy.

1:19 PM: Logged out of all compromised accounts and cleared command history on accessed systems to erase traces of intrusion.

1:24 PM: Ensured no noticeable changes to network configurations occurred.

1:28 PM: Exited the building, blending in with some employees leaving for lunch break to avoid raising suspicion.

Additional notes

- Maintained constant communication with team members to ensure smooth execution of operations
- Adhered to strict security measures to minimize risk of detection and attribution
- Identified specific vulnerabilities in physical security measures and employee practices.

11 Reporting

Regardless of how skilled the red team is in hacking, their efforts are futile if they cannot effectively convey their findings to the client. People who will read the reports drafted up may not have the same mindset or knowledge as the red team does. Moreover, those in leadership often need to understand issues in terms of financial impact, not just in terms of security.

Nonetheless, it is essential to convince the report's audience of the value in addressing the identified issues. Failure to do so undermines the entire engagement, and hampers the case for offensive security.

11.1 What to put in the report?

When compiling a report, it can either be a concise or a detailed document, and this can entirely depend on the client's wishes. Emphasizing findings in the report is crucial, but other aspects warrant attention.

The audience for the report most likely won't have been fully involved in the assessment process or communication chain. From technically adept security personnel to business-oriented senior leadership, the report needs to cater to all. Prior to detailing assessment findings, providing context on the assessment's participants, objectives, timeframe, and methodology is advisable. This ensures clarity, especially for those not privy to this information initially. Highlighting who conducted the assessment, what was assessed, and the duration of the assessment is vital.

Additionally, the report should advocate for future assessment activities. This is beneficial both for the company being tested, as well as the red team. After all, being hired again at some point leads to more work and thus greater income for the team.

Following this, providing a succinct overview of assessment activities is also beneficial. While detailed chronological information can be omitted, a high-level summary covering key enumeration, exploitation, pivot points, and other vulnerabilities is recommended. This narrative can be presented in bullet points or paragraph form. Even in cases of minimal successful exploitation, detailing enumeration is a sign of diligence.

Before presenting the findings, it's prudent to disclose any anomalies unearthed during the engagement. This can mean any non-security-related irregularities, such as unexpected devices in the network. Reiterating the identification of malicious activities detected during the assessment further reinforces the team's credibility. Addressing irregular security staff actions that hindered the assessment progress should also be considered. However, tact and diplomacy are essential, especially in the case where a third-party red team is used, to avoid creating adversarial dynamics.

Any odd behaviour that impeded progress can be acknowledged, while offering assistance to enhance the organization's overall security. The report needs to foster collaboration, not assign blame.

11.2 Categories of discovery

The most important part of any report is the discoveries themselves. It's crucial to recognize the various categories of discoveries, each with its nuances in presentation to the client.

Discoveries aren't always technical; they can encompass misconfigurations or the absence of configurations that facilitate successful attack manoeuvres. It can also take the form of deficiencies in policy that can compromise parts of the organization, an example of this is a weak password policy.

It's also valuable to disclose findings even if they weren't exploited successfully, in order to demonstrate their potential impact.

There are four main categories of discovery that can be included in a report:

1. Exploited vulnerabilities
    a. These refer to weaknesses actively used to breach the organization's defenses. Demonstrating the ability to compromise a system effectively communicates the severity of the vulnerability.
    b. However, the true impact lies in what unauthorized access or data exposure follows the initial compromise. Detailed documentation of the exploitation process provides valuable insights for security personnel in their attempts to mitigate the threat.
2. Non-exploited vulnerabilities
    a. Sometimes a vulnerability may be found, but not exploited. This can be for various reasons. Attempting the exploit may be overtly reckless, or it could present a danger to the system that is not allowed under the ROE.
    b. In such cases, selective targeting is important in order to demonstrate organizational risk effectively.
    c. Gaining access to a system may reveal additional vulnerabilities, obviating the need to demonstrate each one individually. Clients can also restrict access to certain systems deemed too sensitive for testing, prioritizing the safety of their infrastructure over testing.
3. Technical vulnerabilities
    a. These encompass weaknesses within software or hardware, often stemming from poor development practices.
    b. Distinguishing these vulnerabilities from others is important, since they may not be attributable to the organization's actions.
    c. Providing context, such as the recency of a vulnerability's disclosure, aids in understanding its significance for the organization. Addressing long-standing vulnerabilities may require effort beyond simply updating to the latest version of a software, or replacing the lock on a door.
4. Non-technical vulnerabilities
    a. These extend beyond software and hardware, encompassing abstract misconfigurations, such as a lack of policy, or bad adherence to said policies. These vulnerabilities are extremely dangerous for three reasons:
        i. They are often not reported on.
        ii. They are widespread. It's often humans at the source of non-technical vulnerabilities, specifically human behaviours.
        iii. They are difficult to identify because people often lie about their actions all the time. "No, my password is completely secure, I follow the policy set out by the company. I definitely do not share my workstation password with 3 other people

in my department". This is an example I've seen personally, in most companies I've tested.

b. Red team assessments can also uncover deficiencies in incident response, backup policies, policies regarding paper files, destruction of storage media, and general weaknesses in the organization's overall security strategy.

One thing to note for exploited vulnerabilities is that there are certain considerations when dealing with newly discovered vulnerabilities or those previously unexploited. When such vulnerabilities are discovered in third-party software, the organization may have limited influence over disclosure. In such cases, informed decisions on disclosure must be made, often requiring an NDA with the client while determining the best course of action. Conversely, if the vulnerable software is owned by the organization, they may wish to control the disclosure process, potentially necessitating NDAs with the red teamers.

11.3 Severity

Articulating findings effectively within the report involves informing clients about the severity of those findings. Without clarity on the relative importance of each issue, stakeholders will struggle to devise an appropriate remediation plan. Furthermore, the report plays a pivotal role in offering cost-benefit insights to the organization.

Preceding the detailed findings, many reports provide a synopsis of the assessment outcomes.

The way I used to do this, personally, was with a table. Something that looked like this:

	LOW RISK	MEDIUM RISK	HIGH RISK	TOTAL
NUMBER OF FINDINGS	2	4	1	7

This is a very simple and quick way to communicate findings, but there are some problems with this approach. This method fails to capture the full spectrum of risk. Severity ratings typically only focus on the threat posed to the specific system rather than its impact on the entire organization. A more effective approach involves conveying both the threat to the system and its potential ramifications for the organization as a whole. For instance, finding a simple padlock on the janitor's door would constitute a 'high risk' vulnerability with a 'low-impact' system, and wouldn't pose significant organizational risk.

To aid in prioritization, presenting a list of hosts and prompting the client to rank their potential risk can be invaluable. The table from before can be remade to reflect this new distribution of vulnerabilities:

	LOW RISK	MEDIUM RISK	HIGH RISK	TOTAL

LOW IMPACT	1	0	0	1
MEDIUM IMPACT	1	3	0	4
HIGH IMPACT	0	1	1	2
TOTAL	2	4	1	7

A matrix like this makes understanding the scope of the report easy to pick up. Finding out that a high-risk vulnerability exists on a high-impact system will help the organization prioritize their security remediation better, instead of focusing on the 3 medium-risk vulnerabilities on the medium-impact system first, as the first table would lead them to do.

This can be further improved by providing a visual aid with the use of graphs:

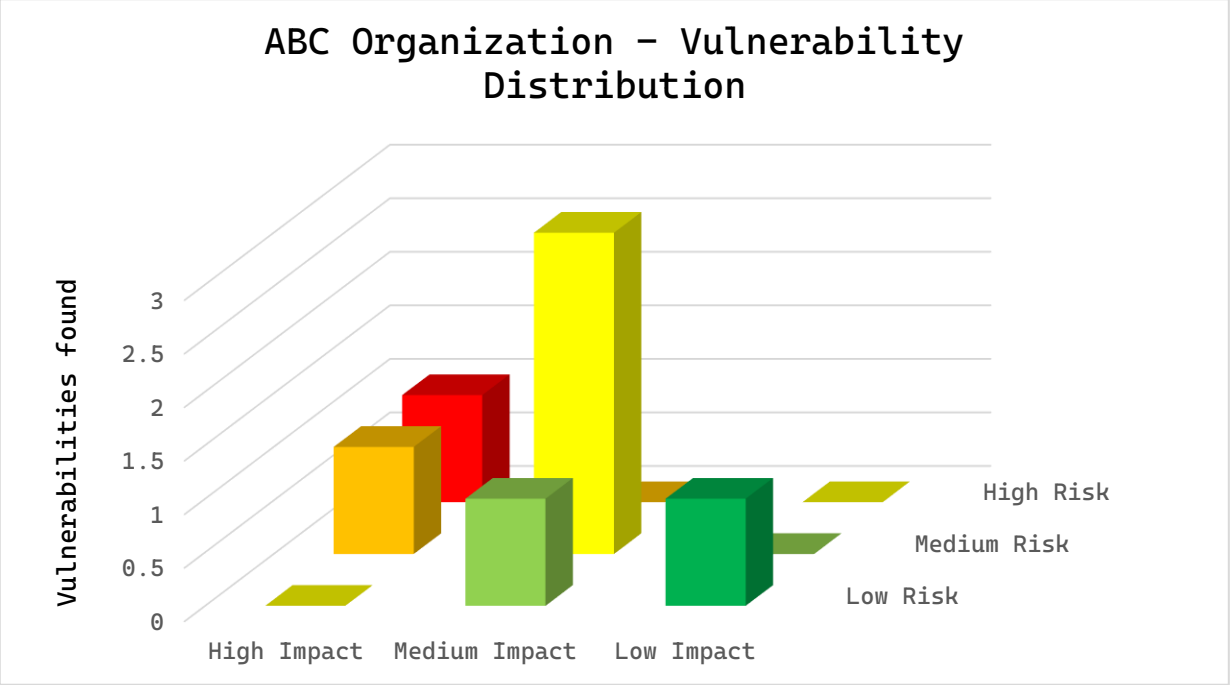

After showing this summary, it's good to then go through the findings themselves, in order of high-impact high-risk down to low-impact low-risk.

Here is a list of things to write down in the report for each finding:

- A short, descriptive title
- The severity rating
- The impact rating
- The category (exploited, non-exploited, technical, non-technical)
- The type (physical vulnerability, software vulnerability, misconfiguration, policy violation, etc.)
- CVE score (if applicable)
- What activity/activities lead to the discovery of this finding
- Systems where this issue was discovered (if applicable)
- A detailed description of the finding
- Mitigation/remediation steps

11.4 Delivery

Delivering the report is a step that is often done without much care put into it. The report is a vital part of the company's future security efforts, and it is prudent to ensure that it does not fall "into the wrong hands", so to speak.

Proper delivery ensures that the report adheres to confidentiality and integrity requirements.

1. Confidentiality: The report should, at minimum, be sent via email to the respective stakeholders via a password-locked archive file. Steps to decompress the file, as well as the password to the archive itself, should be sent via end-to-end-encrypted communication channels.
2. Integrity: To ensure the report is not tampered with en route to its destination, a hash value should be computed and sent along with the password in the previous step. By using this hash, the recipient of the report can ensure that its contents are as-is, and have not been tampered with by anyone else looking to undermine the engagement.

There are tools that can be used to automate delivery of reports, in a way that ensures both confidentiality and integrity, as well as availability, are met. An example of such a tool is PlexTrac.

11.5 Ending the engagement

A nice touch to any successful security engagement is the debriefing. While in-person debriefings aren't always possible, efforts should be made to facilitate a presentation by the assessing party post-report delivery. The purpose of the debrief is not to duplicate the report, to repeat it verbatim, but rather to provide supplementary information emphasizing the significance of both the findings and offensive security assessments in general.

During the debrief, the assessors narrates the story, transforming technical reports into a relatable presentation accessible to everyone, even those who have no idea of anything technical. The most important bit of this debrief is to ensure that it is demonstrated how seemingly minor findings cascade

into significant breaches. By doing this, the legitimacy of the red team and its activities are reinforced. This type of debrief enables red team exercises to simulate real-world threat scenarios.

I'd like to detail a story that Jayson told me, an incredible story he told me about the importance of the report:

"So everyone knows about the Beirut Bank Job, when they talk about you, when they talk about Jayson Street, they talk about the Beirut Bank Job. It's undoubtedly your most famous story ever. What I want to know is, is there another story, of another engagement you were a part of, apart from the Beirut Bank Job, that impacted you in a big way?"

"First off, I think what a lot of people don't understand is that the Beirut Bank Job was my biggest failure. It was my most massive failure, I didn't follow the scope, I entered the wrong building, I broke into someone else's systems without permission. I was extremely lucky not to get arrested from that, and I think why the story is so popular is that I managed to social engineer my way out of that situation, you know so that I didn't get arrested and thrown in prison… which could've happened very easily. So people need to know that this job, this bank job, was a massive failure.

My best, wonderful, success story that I've ever done that just makes me proud was in January of 2020, in the before-times… it was literally in 2019 I broke into this insurance company and they had big-four audit firms try to do pentesting and red teaming physically on their building. Never got in, they never succeeded in getting to the offices. On the first day of the assignment, when I was supposed to be doing mostly just recon, the guy who hired me, the point of contact, found me sitting at his chair, at his desk, after he got back from a meeting, with an actual employee's badge that I took off of their desk when I broke in. It didn't do much better after that, it was a very big thing.

But that wasn't the win! I'm rooting for my clients, I want my clients to do well. The success was this: I had made such an impact on them, and I did such a great job of **reporting what I found**, and making management understand **the importance of getting it fixed**, that within that year of 2019, the CEO had an all-hands meeting. In that meeting he only gets one hour (a year)… one hour to speak to all of his employees. So you know that it's got to be important, what he's got to say. He shares the vision, he shares what their numbers are and all that stuff. He spent 15 minutes of that one hour on security awareness and the responsibility of people to secure their workstations and secure their machines and to be wary of emails. So I come back in January and I legitimately get caught in every section. I didn't have to try to give anybody the win. I was successful in every section, whatever, it doesn't matter. Because it's not about getting in, you're gonna break in. I broke in, that was fine. It's about how quickly can you detect, how quickly can you react to a situation. That's what the new paradigm is. And these people, every single section there was at least one person that said: "No, I didn't get a memo, I didn't get a note, I don't know who you are, you can't touch my computer." I loved it. That was the best success I'd ever done… in every single section I was in, someone did the right thing. And that's all a company can ask for. A company can't ask for every employee to do the right thing but they need to have at least a few of them in every section to help everybody else, and that's the only company I'd seen that. That was such a wonderful thing to have."

11.6 What if nothing is found?

It's crucial to address assessments that yield no significant results. A lack of successful compromises does not equate to failure; rather, it underscores the primary goal of enhancing security posture. Rather than fixating on what was lacking, focus needs to shift to the thoroughness of the assessment. This can provide insights into potential gaps in security measures that require attention.

Furthermore, if constraints hindered the assessment's effectiveness, these limitations should be communicated. Reports for such assessments should guide customers towards refining future assessments. Suggestions may include adjusting the time given to tests, expanding the scope to cover more relevant attack surfaces, or allowing more engagements in the ROE to uncover security issues. While both client and red teamer optimize engagements for success, occasional low- or no-results engagements can still serve as valuable learning experiences for future engagements.

12 RAPTOR framework for anti-APT red teaming

The RAPTOR framework is a comprehensive approach to security engagements, designed from a combination of existing frameworks, and the knowledge gleaned in this project, to enhance the efficacy of red teaming, particularly against Advanced Persistent Threats (APTs).

Naturally, accurately mimicking the patterns of a real adversary brings some risk of system disruption or failure, and thus using this framework may be too bold for companies who would prefer to keep their systems unbothered. Organizations must calculate what amount of risk they are willing to tolerate when using this framework, and adapt accordingly.

12.1 Why is RAPTOR needed?

Counter-APT red teaming is a method of reverse red-teaming theorized by Jacob Oakley, evaluated during his doctoral research and in his dissertation. The red team is always struggling to accurately replicate and counter APTs and other advanced threats, due to inherent limitations already outlined in this thesis. Despite the skills of ethical hackers, the advantage usually lies with the actual threat, making effective emulation challenging.

As the demand for offensive security grows, organizations seek increasingly efficient solutions that minimize the time and resource impact. RAPTOR seeks to address these challenges by introducing modifications to the conventional red team process. By bending certain rules without violating legal/ethical boundaries, RAPTOR seeks to restore some balance to the tug-of-war against persistent threats.

Organizations frequently impose tight engagement windows, necessitating a methodology that delivers meaningful results within constraints. RAPTOR offers a streamlined approach to assessing security posture in such scenarios, ensuring impactful evaluations without compromising efficiency.

12.2 RAPTOR engagements

The goal of any RAPTOR engagement is to safeguard against APTs targeting an organization. The breaches resulting from APTs are termed ‘lethal compromises’ because they encompass scenarios that render an organization completely defunct, or even lead to loss of life. Although RAPTOR is primarily tailored for mitigating critical compromises, it can prove beneficial in various contexts within offensive security practices. Essentially, RAPTOR offers a streamlined and prioritized assessment of specific organizational subsets, aiding in combating APT threats, and conducting focused assessments on targeted assets.

Offensive security assessors should strive to outpace adversaries who invest significant time and effort in attacking entire organizations for valuable assets. Leveraging technical and operational resources, assessors must identify and prioritize assessment of critical items. Initiating campaigns with high-impact targets rather than navigating towards them ensures efficiency and operational advantage.

12.3 Blue team and red team integration

RAPTOR makes extensive use of both the blue and red teams, combining them into one cohesive unit in order to deliver fast, effective results.

Traditional red team assessments alone fall short in addressing the new challenges facing organizations today. The dynamic nature of the threat landscape renders vulnerability scans outdated shortly after an assessment. Moreover, conventional red teaming often prioritizes mimicking external attackers, neglecting internal threat vectors.

12.4 Unknown threats

RAPTOR helps deal with unknown threats, such as zero-day exploits, that capitalize on undiscovered vulnerabilities to bypass security. The process of scanning for vulnerabilities fails to discover zero-day exploits because, inherently, they have not been discovered yet. Nobody knows what to look for. This means that, even after a security assessment, threats may still exist.

The way RAPTOR helps defend against unknown vulnerabilities is by prioritizing the most critical assets first and then working to limit the potential avenues for access from those assets, and then spreading outward and assessing those other systems connected to it. This minimizes the risk that an APT, accessing a system from the inside or the outside, will be able to access critical areas.

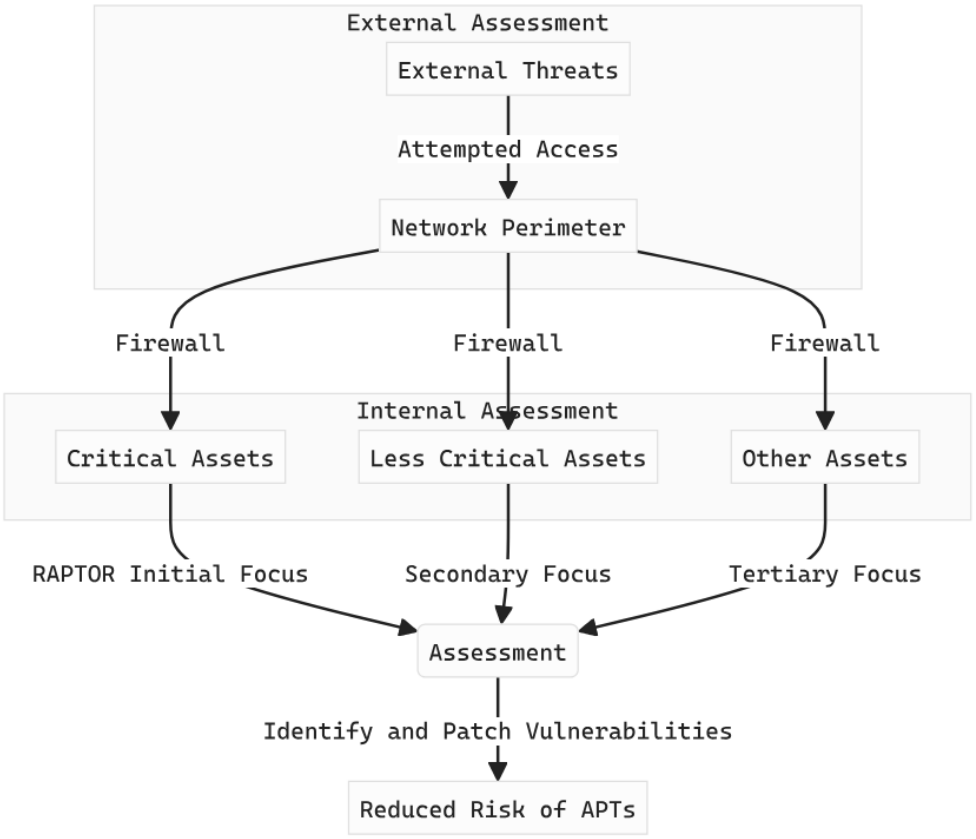

Figure 3 - Flowchart showing the process of a typical RAPTOR engagement

Here is a network diagram showing a comparison between a regular red teaming process and an example of a CAPTR/RAPTOR red teaming process:

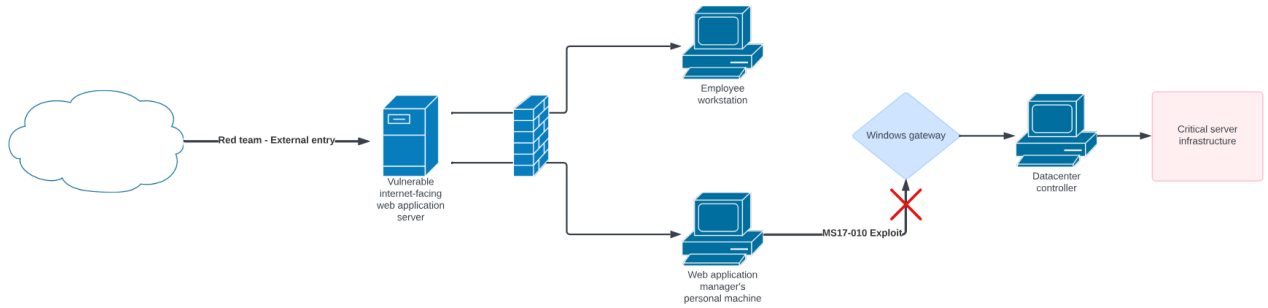

Figure 4 - Network diagram of a regular red teaming process

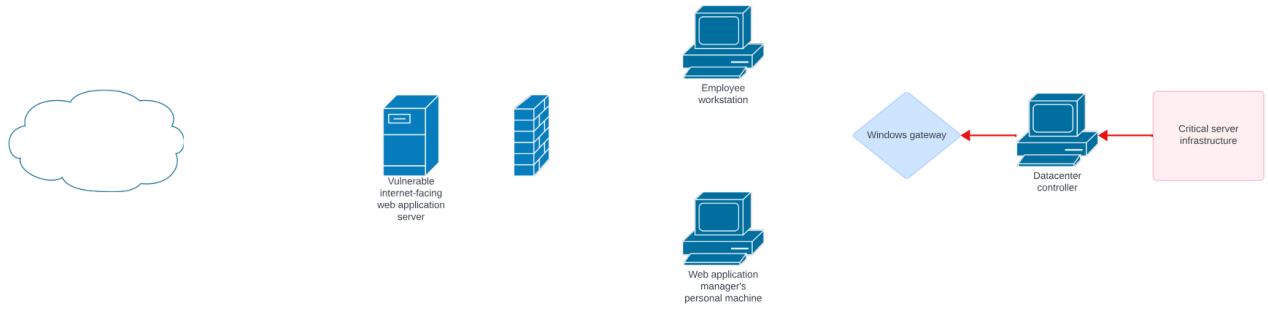

Figure 5 - Network diagram showing an example of a CAPTR/RAPTOR red teaming process

In this hypothetical scenario, the vulnerability MS17-010, which affects Windows machines, is a zero-day and thus has not been discovered yet.

In a typical engagement, the red team would enter the network, make their way through the network, and find themselves unable to bypass the windows gateway to access the critical server infrastructure. The red team would not be able to report on any of the vulnerabilities of that system because they have not managed to reach that point. But an APT, with considerable resources, would discover the zero-day and use it to bypass the gateway, gaining access to systems that have not been tested, and could thus cause considerable damage to the entire network.

Now take the CAPTR/RAPTOR engagement, where the focus is placed on the critical server infrastructure. Here, the red team starts from the core servers and moves outwards towards less critical peripheral systems. This way, even if the zero-day is exploited, the APT would find it more difficult to enter the critical infrastructure after bypassing the windows gateway because those systems were tested properly.

This provides an additional benefit during the auditing process: the risk to the critical system is reduced considerably, as it is not being bombarded by often unstable remote code execution originating from outside the critical system. Instead the testing takes place locally, inside the critical environment, and branches outward.

13 RAPTOR elements

1. Risk-oriented scoping
2. Assessment of vulnerabilities
3. Precision pentesting
4. Threat hunting
5. Operational resilience
6. Real-time response/remediation

13.1 Risk-oriented scoping

13.1.1 Objective

Scope out the most important assets to be assessed.

This approach optimizes assessment resources by targeting a subset of the organization that represents the 'worst-case scenario' if an APT gains access.

13.1.2 Deliverables

- Scope document that contains the most crucial assets in the organization.

13.1.3 Implementation

- Work with the security personnel, as well as any relevant staff and stakeholders to identify critical assets.
- Prioritize critical elements that have significant impact if compromised.

13.2 Assessment of vulnerabilities

13.2.1 Objective

This is the starting point of the actual engagement. This phase requires the team to directly assess the priority risk items identified in the scope.

13.2.2 Deliverables

- Plan of attack with potential attack scenarios, techniques, tools that will be used, or tactics that will be employed such as social engineering.

13.2.3 Implementation

- In this phase, scans and vulnerability enumeration occur. This can come either from the outside (external breach) from the Internet, or internally (assumed breach) from various points within the organization.
    - Where this phase is begun impacts several aspects of the assessment, including which attack surfaces are prioritized, and the vulnerabilities that get uncovered.
    - It's crucial to acknowledge that APTs can breach the perimeter and inner layers of the organization. Commencing from the outside can impede engagement progress and success.
- Once the team has a solid idea of what the systems look like, developing attack scenarios tailored to the target organization's industry comes next. The team needs to keep in mind the organization's tech stack and threat landscape.
- Attack techniques need to constantly evolve and change based on feedback gained and intelligence gathered during engagements.
- If the red team is to truly simulate a dangerous adversary, custom malware, zero-day exploits and multi-stage attack chains are a good pick to mimicking this behaviour. But keep in mind the scope, and the inherent risk that comes with using exploits.

13.3 Precision pentesting

13.3.1 Objective

Penetrate identified target systems with precision, exploiting the various vulnerabilities in both technical and human targets.

13.3.2 Deliverables

- Completed assessment report detailing identified vulnerabilities, their severity, and recommended remediation steps.

13.3.3 Implementation

- Exploit vulnerabilities, sticking to the attack plan drawn up in the previous phase.
- If possible, make use of social engineering to manipulate human targets and gain access to sensitive information or credentials.

13.4 Threat hunting

13.4.1 Objective

Hunt for signs of APT activity within the target environment and track adversary actions across the kill chain. This is where the red team starts to work a bit closer with the blue team in the organization.

13.4.2 Deliverables

- Document outlining indicators of compromise (IOCs).
- Incident response playbook detailing procedures for detecting, containing, and mitigating APT attacks. Most likely, it will be the blue team working on this document.

13.4.3 Implementation

- Deploy advanced monitoring and logging solutions to capture and analyse network traffic, system logs, and user behaviour.
- Utilize threat intelligence feeds and behavioural analysis techniques to identify anomalous activity and potential APT behaviour.
- Establish a centralized incident response team to coordinate detection, investigation, and response efforts.

13.5 Operational resilience

13.5.1 Objective

Build operational resilience to withstand APT attacks and minimize the impact of breaches through defense-in-depth strategies and preparedness measures. This is mostly blue team work, but the red team should work closely with them during this phase in order to identify shortfalls or gaps in security when considering what methods will be used for increasing resilience.

13.5.2 Deliverables

- Report outlining current security measures and incident response methods, gaps, and recommendations for improvement.
- Incident response plan detailing roles, responsibilities, and procedures for handling security incidents.

13.5.3 Implementation

- The blue team will work to implement defense-in-depth strategies, including network segmentation, least privilege access controls, and data encryption. They will then work to solve problems identified by the red team in their report. The red team will need to be open to answering questions and providing support post-assessment.
- The red team's role in this phase will be to conduct regular security assessments and penetration tests to proactively identify and address any vulnerabilities or shortfalls in the new resilience protocols.
- Another potential avenue is to establish communication channels and partnerships with external stakeholders, such as law enforcement agencies and industry peers, as needed.

13.6 Real-time response/remediation:

13.6.1 Objective

Respond to APT activity in real-time, with swift and decisive action to contain and mitigate the threat, followed by remediation measures to strengthen defenses.

13.6.2 Deliverables

- Incident response playbook detailing steps for detecting, analyzing, containing, and recovering from APT attacks.
- Post-incident analysis report documenting lessons learned and root cause analysis. This must be presented in conjunction with the **Threat tracking** and **Operational resilience** reports.

13.6.3 Implementation

- Establish incident response procedures and workflows to ensure timely detection and response to security incidents.
- Conduct tabletop exercises and simulation drills to test incident response capabilities and identify areas for improvement.
- Implement remediation measures, such as patching vulnerabilities, updating security policies, and enhancing employee training and awareness.

14 RAPTOR Risks

Conducting offensive security assessments carries inherent risks. These assessments often involve using potentially unstable exploits. These exploits can cause systems to crash. As engagements focus on critical assets, the risk escalates, akin to scenarios faced by APT hackers. Mitigating this risk involves careful scoping and establishing ROE prior to testing. However, even with precautions, there's still a chance of unforeseen consequences from exploitation techniques.

The RAPTOR approach minimizes risk to high-risk environments. By initiating assessments locally from identified critical assets, there's no threat of remote code execution vulnerabilities causing disruptions. Unlike traditional red teams relying on active scanning tools like Nessus scanner, the CAPTR team uses passive information gathering to identify pivot points. This approach reduces reliance on remote scanning tools and minimizes the risk posed to critical systems, enabling effective security assessments against APT-like threats.

15 RAPTOR Weaknesses/Disadvantages

In thoroughly examining the RAPTOR framework, it's important to address its limitations and scenarios where it might not be suitable. Challenges in launching RAPTOR engagements involve both inherent flaws in the approach and obstacles typical of any new concept competing with established norms. The methodology of RAPTOR assessments revolves around identifying vulnerabilities favoured by APT attackers to access critical assets (Precision Pentesting and Threat Tracking phases). Consequently, it's less effective against other threat types and divers entry points. This approach may overlook vulnerabilities accessible through internet-facing interfaces, leaving room for exploitation by less sophisticated attackers like automated tools or script kiddies. The primary hurdle lies in the initial stages of assessment, demanding a blend of security expertise and risk management acumen. Failure to align risk and security elements could compromise the entire evaluation process.

The reliance on data generated by the red team's scoping phase introduces a potential vulnerability. Initiating assessments from critical network points entails significant trust between organizations and testers, raising liability concerns. This heightened risk may deter organizations from opting for such assessments, impacting the service's market viability. Moreover, deeper network access necessitates closer collaboration with IT and security during testing, potentially influencing cost-benefit analyses. Ultimately, RAPTOR may not be suitable for organizations lacking critical assets or facing a diverse threat landscape. Their focus on APT threats and internal data security means they're not comprehensive solutions for all security assessment needs.

16 Comparative Analysis of RAPTOR and CAPTR Frameworks

This section provides a detailed analysis of the RAPTOR and CAPTR frameworks, discussing how RAPTOR can grant improvements to the capabilities of the existing CAPTR method of red teaming.

16.1 Overview of CAPTR Framework

The CAPTR (Counter-APT) framework, developed by Jacob Oakley, is designed to mitigate APTs by focusing on critical and high-impact assets within an organization [4]. CAPTR emphasizes efficient and effective security engagements through prioritizing evaluation of the most important assets and their vulnerabilities. Thus, RAPTOR is built to be a progeny of CAPTR, in a way.

16.2 Key components of CAPTR

1. Worst-case risk analysis and scoping
    a. Identifies high-risk items that can lead to catastrophic outcomes if compromised
2. Critical compromise initialization perspective
    a. Starts the assessment from the critical points in the network rather than from external points
    b. Focuses on identifying and exploiting vulnerabilities that lead to significant breaches
3. Vulnerability analysis and exploitation using reverse pivot chaining
    a. Utilizes passive information gathering to guide assessments towards high-risk items
    b. Reduces reliance on active scanning tools to minimize risk to critical systems

16.3 Advantages of CAPTR

- Efficiency: Concentrates on high-impact areas, ensuring resources are used effectively
- Reduced risk: Limits the use of potentially disruptive tools, thereby protecting critical assets from unintended damage during assessments
- Focused assessment: Tailors assessments to the most significant threats

All of these advantages also apply to RAPTOR.

16.4 Disadvantages of CAPTR

- Limited scope: May not identify all vulnerabilities
- High trust requirement: Requires significant trust between the organization and the assessors due to the sensitive nature of the assets being evaluated
- Complex coordination: Needs extensive coordination between IT and security teams, which can be resource-intensive

All of these disadvantages also apply to RAPTOR.

16.5 Advantages of RAPTOR

- Real-time adaptability: Keeps the assessment relevant by incorporating the latest threat intelligence
- Balanced risk approach: Combines impact and likelihood to prioritize risk effectively, while ensuring critical assets are always placed first.

16.6 Disadvantages of RAPTOR

- Resource intensive: Requires significant resources to maintain real-time threat intelligence and continuous assessment capabilities
- Complex implementation: More complex to implement and manage compared to focused frameworks like CAPTR

16.7 Comparative analysis

16.7.1 Scope and focus

- CAPTR: Focuses on high-impact, critical assets only, ensuring efficient use of resources but may leave peripheral vulnerabilities unchecked.
- RAPTOR: Attempts to balance efficiency with a holistic assessment, covering as many critical and peripheral assets as possible, ensuring comprehensive security but requiring more resources.

16.7.2 Threat intelligence integration

- CAPTR: Primarily focuses on internal vulnerabilities with less emphasis on real-time threat intelligence
- RAPTOR: Integrates real-time threat intelligence, ensuring the assessment is relevant and up-to-date

16.8 Enhancing CAPTR with RAPTOR principles

Integrating elements of RAPTOR into CAPTR could address some of the latter's limitations:

- Incorporating real-time threat intelligence can enhance the relevancy of assessments conducted, and allows red teams to adapt their focus based on emerging threats.
- Balancing risk assessment with priority assets can help ensure both high-impact and high-likelihood vulnerabilities are addressed.
    - More holistic scoping provides a more comprehensive security assessment without compromising too much on resource efficiency.
- Using RAPTOR's comprehensive reporting strategies, keeping in mind the deliverables of each phase of the process, helps the blue team come up with actionable mitigation strategies.

17 Conclusion

This thesis aimed to address the critical gaps in current research with regard to methodologies and frameworks guiding red team operations. This paper sought to enhance effectiveness and efficiency of red teaming engagements, especially against APTs.

Here are some of the main highlights in the research to consider:

1. Literature review and case studies: The review of real-world case studies provided a foundation for identifying current challenges red teams face, and also helped identify best practices.
2. Interviews: Insights gathered from professionals in the industry, such as Jacob Oakley and Jayson Street, as well as regular cybersecurity practitioners, helped shape the framework and the content of the handbook.
3. RAPTOR framework: This framework was introduced as a targeted approach to mitigate APT-like threats. RAPTOR works to minimize risks to critical assets while ensuring the organization still maintains compliance with security audits.
    a. There are limitations of the framework. It's less suitable for environments where the primary threat comes from less advanced attackers, or more diverse attack vectors. Before using this framework, organizations have to make informed decisions about their risk tolerance and security strategies.
4. Implementation: This thesis offers practical guidelines for planning, executing, and analyzing red teaming engagements, all to offer solutions to issues and challenges red teams face.

In conclusion, this paper contributes significantly to the field of cybersecurity by offering a novel framework and theoretical, as well as practical, tools for red teaming. The framework could represent a new step forward in addressing the challenges posed by APTs. Future research will build on this thesis by refining the framework and exploring its applicability in more diverse threat landscapes.

AI Engineering Prompts

1. Give me a list of questions for interviewing a professional red teamer. The interview needs to yield answers that could be valuable to a bachelor's thesis, keep that in mind.
2. Should I include a chapter on cultural analysis for my red teaming handbook? How could that be relevant? What uses would cultural analysis have for red teaming?
3. Give some examples of applying this cultural analysis.
4. What about a chapter on mindset, such as self-awareness?
5. Help me write supplementary material for a chapter for my red teaming handbook based on the following text:
   The journey to understanding your own undesired or unproductive tendencies (and acting to overcome them) starts with self-awareness. Objective evaluations and decisions can only be made by self-aware individuals who understand the characteristics of themselves that influence the end result of such evaluations.

Bibliography

[1] M. Zenko, Red Team - How to Succeed By Thinking Like the Enemy, Basic Books, 2015.

[2] “Red Teaming: Discussing the Basics,” IBA Group, 4 April 2024. [Online]. Available: https://ibagroupit.com/insights/red-teaming-discussing-the-basics/. [Accessed 10 April 2024].

[3] “Understanding Pentesting vs. Red Teaming,” Cyderes, 14 December 2020. [Online]. Available: https://www.cyderes.com/blog/penetration-testing-vs-red-teaming. [Accessed 13 April 2024].

[4] J. G. Oakley, Professional Red Teaming - Conducting Successful Cybersecurity Engagements, New York: Apress, 2019.

[5] “Glossary - Advanced Persistent Threat,” NIST, [Online]. Available: https://csrc.nist.gov/glossary/term/advanced_persistent_threat. [Accessed 19 April 2024].

[6] “Pentest vs. Red Team Assessment,” turingpoint, 24 March 2020. [Online]. Available: https://turingpoint.de/en/blog/pentest-vs-red-team-assessment/. [Accessed 12 May 2024].

[7] J. Jin, Tribe of Hackers Red Team: Tribal Knowledge from the Best in Offensive Cybersecurity, Indianapolis: Wiley, 2019.

[8] B. Craven, O. Mokhamed Makhmud and J. E. Street, Composers, *Red Teaming with Jayson Street.* [Sound Recording]. Howest. 2024.

[9] R. Kegan, In Over Our Heads: The Mental Demands of Modern Life, Harvard University Press, 1998.

[10] University of Foreign Military and Cultural Studies, Red Team Handbook, 2012.

Appendices

Figure 1: Recognizing personal bias

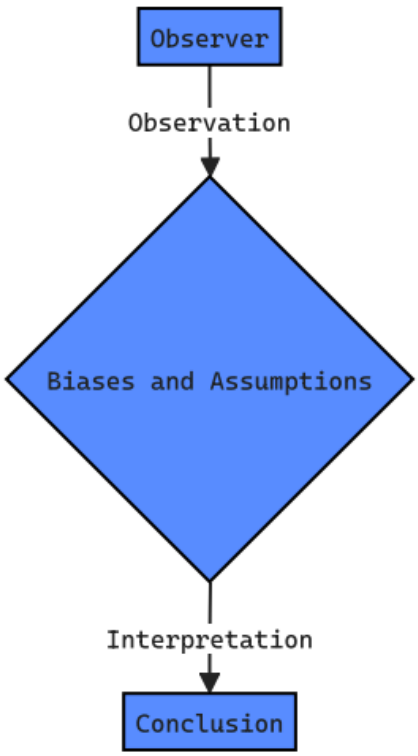

Figure 2: Components of self-authorship

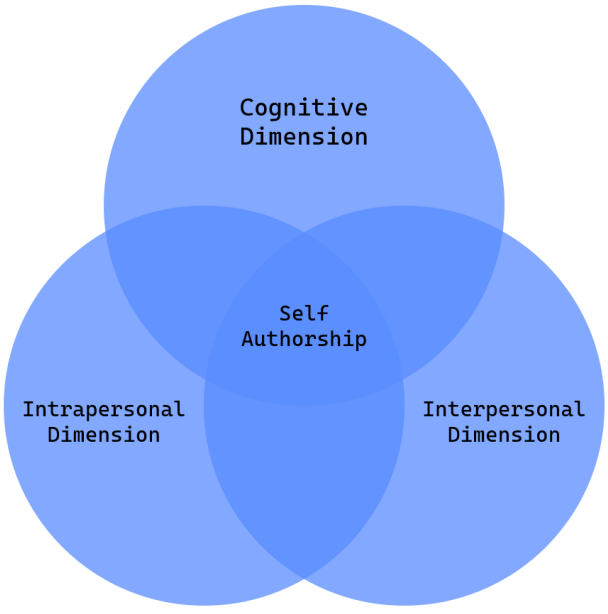

Figure 3 – Flowchart showing the process of a typical RAPTOR engagement

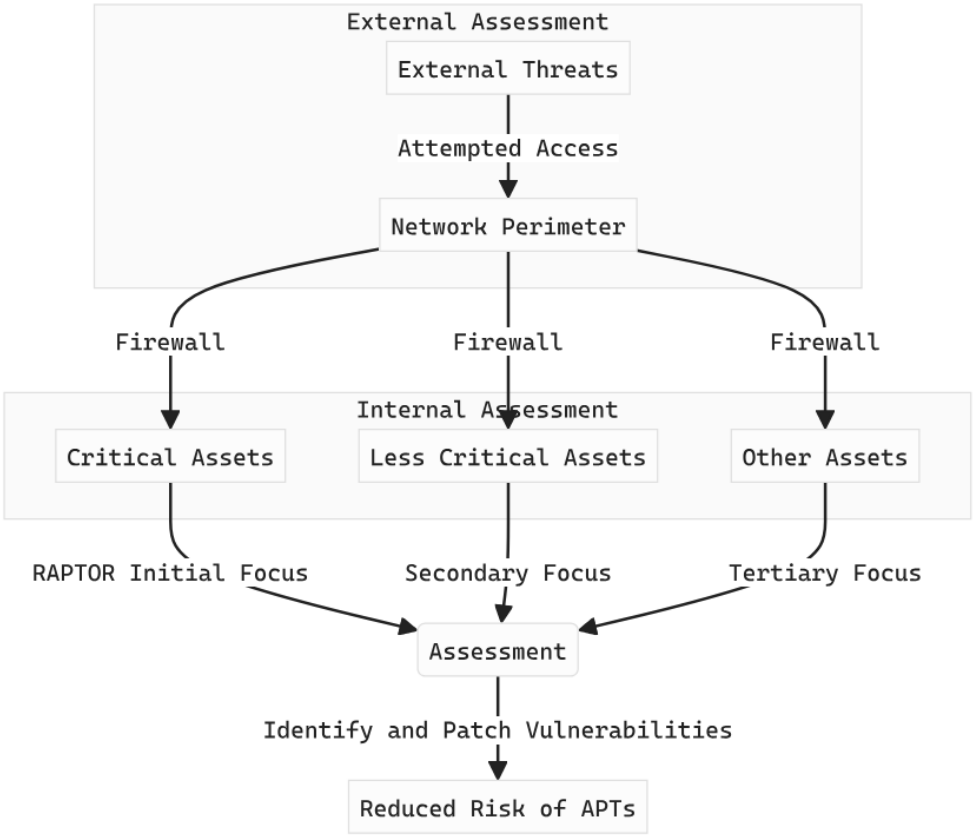

Figure 4 – Network diagram of a regular red teaming process

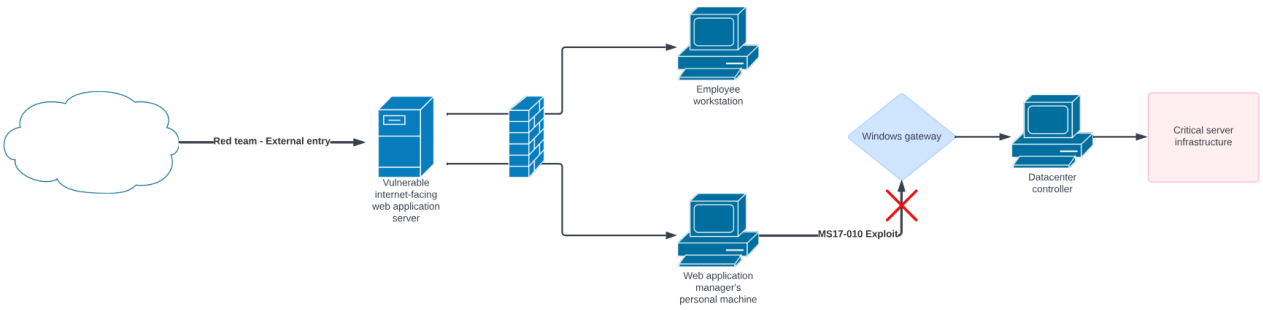

Figure 5 – Network diagram showing an example of a CAPTR/RAPTOR red teaming process:

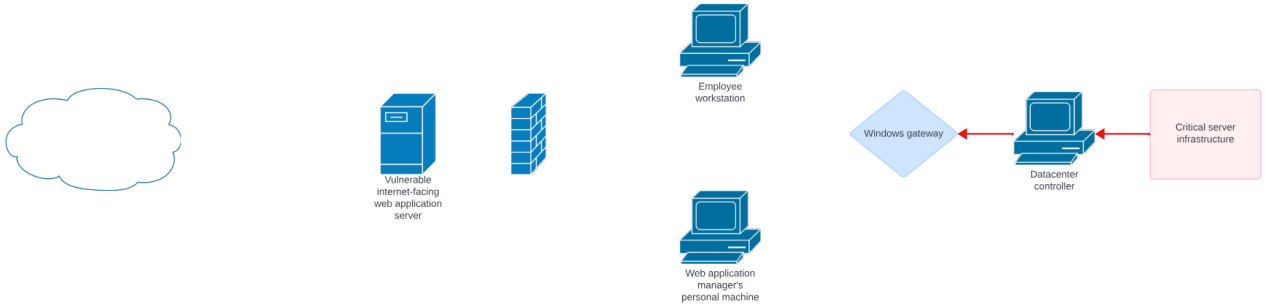

Plagiarism Check

https://www.duplichecker.com/

Duplichecker

Dupli Checker Plagiarism Report

Originality Report

6.44%	93.56%	16187	115534
Plagiarized Content	Unique Content	Words	Characters

Primary Sources

1	https://github.com/lenvdv/auto-dj... ... dissertation available for consultation and to copy parts of this master dissertation for personal use. In the case of any other use, the copyright terms ...	0.08%
2	https://www.diva-portal.org/smas... by E Meurs · 2023 — In all cases of other use, the copyright terms have to be respected, in particular with regard to the obligation to state explicitly the ...	0.08%
3	https://upr-info-database.uwazi.io... ... respected, in particular with regard to the right to freedom of expression. Austria. 5 - Specific action. Noted. Freedom of opinion and expression ...	0.08%
4	https://www.ugent.be/ea/en/educ... regard to the obligation to state explicitly the source when quoting results from this master dissertation." (gevolgd door datum) / (followed by the date). 4 ...	0.08%
5	https://medium.com/@craiglawso... Thanks for sharing your knowledge with the community. Sign up to discover human stories that deepen your understanding of the world. Free.	0.08%
6	https://books.google.com/books?i... The Ten Commandments: Interpretation: Resources for the Use ...	0.08%
7	https://hmsom.edu/en/education/... ... best practices, and gaps in the research for further work. As part of this review, they were able to identify several key components found in high-quality ...	0.08%

Report was generated on Sun, May 12, 2024 Page 1 of 27

This report was generated on a draft version of the thesis, Sunday May 12 2024 using the professional version of Duplichecker. The full report can be found on LEHO uploaded alongside this document.