Whitepaper: Red Teaming

Red Teaming – An option for all Enterprises

Date: 1 August 2016

Doc Ref: SOS-WP-RT-16081

Authors: Nathaniel Carew & Michael McKinnon

Table of Contents

Overview

This document aims to help security managers and business stakeholders understand what Red Teaming is, how it differs from other forms of security testing, and how it improves your enterprise's security posture.

We also provide insights to help you navigate your first Red Teaming engagement, to understand and communicate the right expectations for an assignment, and how to select a Red Teaming partner.

The information contained herein has been collected and compiled from the experience of many Red Teaming engagements performed by the Sense of Security Red Team over a number of years.

This paper does not intend to be a "how-to" guide on performing Red Teaming itself, and it does not cover the tools and confidential tactics used. However, it does explain the overall process to help demonstrate the skilful and detailed nature of Red Teaming.

What is Red Teaming

“Red Teams” and “Red Cells” first originated in the United States in the form of elite teams and military units designed to independently test the effectiveness of strategy, tactics and personnel. In typical military simulations and war games, a Red Team pretends to be an adversary, while a Blue Team defends patriotic interests.

Thus Red Teaming is broadly described as the act of simulating the actions of a real adversary, or playing “devil’s advocate”. According to the US Department of Defense, “Red teaming deepens the understanding of options available to adaptive adversaries and both complements and informs intelligence collection and analysis.”[1]

However, the use of Red Teaming is not exclusive to the military, and in the IT security industry the demand for Red Teaming exercises continues to grow – by challenging traditional organisational thinking and providing opportunities for an unbiased external view of real enterprise network defence and security.

In the private sector, Red Teams comprise highly-experienced security professionals who specialise in key disciplines, and collectively offer a diverse set of skills. The ability to simulate real-world attacks also requires a detailed knowledge of the latest techniques used by threat actors.

Red Teaming involves applying expert knowledge in a number of areas:

- **Strategic thinking** – the ability to think laterally and challenge long-standing assumptions
- **Tactical planning** – devising, coordinating and executing a plan to perfection
- **Industry knowledge** – expert level knowledge of all the latest threats and vulnerabilities
- **Social Engineering** – exploiting innate human vulnerabilities to leverage attacks
- **Physical attacks** – finding weaknesses in physical security systems and defenses
- **Virtual attacks** – using a cybersecurity arsenal to access confidential systems and data
- **Malware development** – using malicious payloads, and reverse engineering skills
- **Open source intelligence** – using publicly available information to gain an advantage

Along with these knowledge areas, the independent nature of Red Teaming brings with it many advantages, such as identifying likely real-world vulnerabilities, providing pragmatic opportunities that help build resilience and strengthen the security posture of your enterprise.

In summary, Red Teaming exercises are bespoke, intelligence-led security tests designed to replicate as closely as possible the evolving threat landscape as it relates to your business.

Effectiveness of Red Teaming

Red Teams are able to challenge aspects of your enterprise’s plans, programs and assumptions across multiple levels of the organisation simultaneously. It is the deliberately hostile adversarial action that distinguishes Red Teaming from other stand-alone management tools and makes it an important and effective undertaking.

[1] http://fas.org/irp/agency/dod/dsb/redteam.pdf

In contrast to traditional intelligence-based threat projections, Red Teaming is able to reveal a wider and deeper understanding of potential adversary options, including threat actor behaviors that may never have been previously considered.

Furthermore, relying solely on a combination of network scans, penetration tests and other similarly isolated tests may not yield the same results as when combined with a Red Teaming engagement. This is due to Red Teaming being able to leverage weaknesses across multiple domains at the same time.

Although individual penetration test reports may provide more technically detailed recommendations useful for compliance needs, Red Teaming reports will focus on complete end-to-end attack scenarios using the most critically discovered issues.

Even in Red Teaming assignments where an intrusion has not eventuated, in our experience useful recommendations are still aplenty, due to the process itself and the valuable observations made. Thus the effectiveness of Red Teaming is what arguably makes it one of the most important security engagement types, and underpins the growth in demand that we're seeing in this area.

Red Teaming vs. Penetration Testing

There are digital attack components to Red Teaming exercises that depend heavily on penetration testing skills (and tools), yet the objectives and outcomes of both testing activities are vastly different.

Penetration testing is designed to deliver an exhaustive battery of digital intrusion tests that provide recommendations ranging from "critical", all the way down to "informational". Thus while penetration testing is valuable for finding vulnerabilities, it is equally good at identifying compliance problems and opportunities.

In contrast, Red Teaming aims to exploit only the most effective vulnerabilities in order to achieve an end result – to capture a target and achieve a single-minded mission. Accordingly, Red Teaming is not a replacement for penetration testing as it provides nowhere near the same exhaustive review.

In our experience, Red Teaming engagements and penetration tests complement each other, with both distinct activities playing vital roles in securing your enterprise and ensuring it is compliant where required.

Red Teaming Scope and Methodology

While it is possible to develop and apply a general methodology suitable for most Red Teaming engagements, there is no one-size-fits-all approach. The scope of a Red Teaming exercise should ideally be as broad as possible, yet it must also be well defined – a contradiction that demonstrates the complexity of this type of undertaking.

Estimating the scope of a Red Teaming activity is critically important. Too narrow a scope and you might limit the effectiveness of the Red Team, and too broad a scope could have unintended consequences; for example, accidental damage to production systems or disruption to revenue flow.

A Red Teaming engagement must therefore be devised specific to your enterprise, including any seasonal variations and industry specific oddities. An engagement plan developed by your Red Teaming partner will help identify risks and reduce them to a set of test scenarios that mimic real-world threats.

These are commonly split into three areas.

1. digital attacks
2. social engineering
3. physical defence

In our experience, developing an engagement plan begins with identifying what your enterprise wishes to protect the most. This may involve consultation with many different stakeholders in the business.

One of the best starting points when identifying important assets, is to consider the different types of possible adversaries and their common motivations. Consider the following categories of threat actors:

- **Competitors/Organisations** – targeting competitive advantage knowledge and/or key Intellectual property, known to break in and steal schematics, source code, or sales and marketing information.
- **Organised Crime** – groups of criminals that intend to engage in illegal activity, most commonly for monetary gain through extorting money; commonly deploying banking malware or ransomware.
- **Nation-state sponsored** – a group employed by the government of a nation-state. The focus of this adversary can be pilfering data, intellectual property, and research and development data.
- **Hacktivists** – an actor that attacks for the purposes of drawing attention to a cause (such as free speech, environmental, legal or human rights), or alternatively hinder the support of a cause.
- **Cyber Terrorists** – carrying out an attack designed to cause alarm or panic with ideological or political goals and may target critical infrastructure or similar.
- **Security Professionals** – includes security researchers, computer scientists, anti-virus vendors, CERTs, threat intelligence (non-state-sponsored) groups, may have backdoors planted.
- **Law Enforcement** – including anyone involved in law enforcement (police, police cybercrime units, courts, judges) as well as attorneys and lawyers; in some countries they may target illegal businesses and undertake surveillance in order to bring them down.
- **Individual** – a person or group acting on their own, and not a member of any other threat category whose motives often include notoriety, curiosity, bragging rights, etc.
- **Unknown** – other actors such as trusted insiders, or any other party not listed above.

Understanding who your likely adversaries are will assist in formulating an effective Red Teaming engagement plan, including identification of the likely targets (assets) you need to protect.

Once the targets have been identified, along with indications as to the most “feared” type of threat actor, the Red Team will be able to formulate their plan of covert execution, within the scope of the engagement.

Before the execution of a Red Teaming exercise, however, there are many questions that need to be answered that will further contribute to defining the scope of the assignment.

Through a consultative process, your chosen Red Teaming partner should compile a comprehensive list of items that will shape the assignment to your specific needs, such as:

- **Operating hours** – e.g. allowed to perform out of business hours
- **Off-limit areas** – e.g. manufacturing facilities, hazardous or dangerous areas
- **Scheduling concerns** – e.g. peak production times or seasonal issues
- **Credentials provided** – e.g. for simulating insider threat scenarios
- **Digital scope** – e.g. any particular networks, equipment, IP addresses specifically excluded

- **Locations** – e.g. activity allowed at or near the homes of selected staff members
- **Privacy issues** – e.g. how to handle access to personally identifiable information
- **Allowed damage** – e.g. an agreed dollar amount of physical damage deemed acceptable

This is a short sample of hundreds of similar items that the Sense of Security Red Team has assembled during the course of previous Red Teaming engagements. It is vital that the Red Team considers as many factors ahead of the assignment as possible – this is one key difference that can set an experienced Red Teaming provider apart from one that has rarely performed this type of activity.

In summary, Red Teaming requires careful planning and attention to detail, and must be tailored to suit each individual enterprise for the best results. Stick to a well-used methodology from a trusted partner, identify your threat actors first, and understand what your targets for the Red Team will likely be.

Differences between real-world adversaries and Red Teams

There are important differences that exist between Red Teams and real-world adversaries. For reasons that are crucial on both ethical and legal grounds, Red Teaming is not intended to be a complete substitute for predicting or replicating an actual real-world attack.

Firstly, from an ethical point of view, all members of a commercially responsible Red Team must be police checked and their conduct should comply with a strict code-of-conduct. As security professionals, their conduct should already be in alignment with widely accepted ethical hacking industry standards.

In contrast, real-world adversaries are virtually guaranteed to operate without any ethical boundaries, and some threat actors likely have prior criminal convictions or are being actively pursued by law enforcement.

An enterprise needs to have confidence in the Red Teaming process at all levels, and this can only be achieved by the Red Teaming provider upholding the highest ethical standards available. Make sure your chosen Red Teaming partner fits the bill.

Interestingly, the upholding of ethical boundaries can provide challenges to a Red Teaming engagement resulting in the Red Team falling short of taking the “exact same actions” a real-world adversary would.

To compensate, an experienced Red Team should demonstrate a proof-of-concept that does not harm the enterprise, providing evidence of how a successful attack would play out. For example, an extreme destruction scenario that would otherwise threaten the existence of the enterprise if actually carried out.

The core legal and ethical differences are, we believe, vitally important to the integrity of Red Teaming and act to represent the trust and confidence that enterprises need to have in the IT security industry as a whole. For this reason, Red Teams must be distinctly separate from real-world adversaries at all times.

Moreover, as an advantage to Red Teaming, consider that real-world attacks are rarely attributable to a known threat actor, whereas a Red Team is commercially obligated to reveal themselves after an attack. For many Blue teams, it is a great relief when they eventually learn an incident was “only” caused by the Red Team.

Legal implications of Red Teaming

When it comes to legal issues, there are some instances where criminal law makes little distinction between the actions and the intent of certain alleged conduct. This can place Red Teams in predicaments during an assignment, particularly prevalent when performing physical intrusion tasks.

Different jurisdictions at state level, and also between countries, may apply during a Red Teaming assignment and must be carefully considered in advance, with appropriate legal advice sought.

Red Teams should always ensure they carry with them a letter of authority from the enterprise, in the event they are detained and questioned by law enforcement. This letter will normally include after-hours contact details of senior stakeholders from your enterprise.

It is not the role of a Red Team, ethically nor legally, to attempt to socially engineer their way out of situations involving law enforcement. But if confronted by private security staff or employees about their conduct, social engineering is certainly a tactic that may be used by the Red Team.

A Red Team "changing their story" can make matters worse, however, if they are subsequently handed off to local law enforcement officers and detained. In cases like this, even the validity of the letter of authority may come into question until further inquiries are made.

Potential physical harm to members of the Red Team must also be considered in Red Teaming engagements, particularly where armed guards or law enforcement are involved. Also, there are other physical risks such as electric shock, cuts from razor wire, falls from heights, or the dangers of interfering with mechanical systems such as elevators, as some examples.

Given the public liability challenges inherent in Red Teaming, it is the responsibility of the Red Teaming provider to ensure they're taking reasonable steps to reduce the operational risks, and be responsible for the welfare of their own staff. This point should serve as further evidence of the importance that experience plays in measuring the maturity of Red Teaming providers operating in this space.

Lastly, as Red Teaming aims to deliberately simulate hostile actions, it is not surprising that the Red Team may be treated like a real-world adversary by those not aware of the exercise. Managing as much of the legal risk as possible beforehand is paramount to the success of Red Teaming overall.

When to perform Red Teaming

Some policy frameworks place regimes such as red teaming and penetration testing at the very end of security control lists[2], suggesting they be performed only after a high level of enterprise maturity is achieved. And even from a traditional business process understanding, testing and quality assurance tasks are commonly *perceived* as being required only at the last stage of output in a production system.

Indeed, given the James Bond like buzz of Red Teaming – visions of darkly dressed operatives sneaking along fence-lines and cutting through razor-wire fences – it is no wonder this activity is seen as only being meant for an elite few.

[2] https://www.sans.org/critical-security-controls - CSC Critical Control No. 20

However, in direct contrast to the notion that Red Teaming should only be associated with high levels of process maturity, in our experience there are strong practical arguments for initiating it at any level.

A real-world adversary doesn't wait until their victim has a well-established security program – supporting the argument for introducing Red teaming at almost any stage of maturity.

In fact, Red teaming performed in the early stages of enterprise maturity can provide enormous value, such as identifying critical "low-hanging-fruit" issues that can then be used to set evidence-based priorities for fast-tracking security controls.

Repeat engagements / Red Team cycling

Red Teaming engagements should not be considered one-off activities; successive engagements can measure the effectiveness of changes to policy, readiness, awareness programs and security controls.

Choosing the best timeframe for a follow-up Red Teaming engagement will often depend on how many recommendations have been actioned from the previous assignment. This can vary significantly, but is usually when the enterprise stakeholders feel confident they have remediated or addressed most of the previous concerns.

In summary, the right-time for a first Red Teaming engagement is at any time, and arguably the sooner the better, followed by successive follow up engagements at intervals that have allowed for corrective controls to have been implemented.

Determining the length of a Red Teaming engagement

An adversary always has the upper hand when it comes to timing an attack with the element of surprise; as many incident handlers will declare, there's never a "good time" to experience a security incident.

Red Teams therefore need be allowed to maintain this same adversarial element of surprise to remain effective, and are also in the best position to estimate the most appropriate length of an engagement.

Time is also an important factor when considering the classic security maxim of "given enough time an adversary will always succeed". This may be true of a real adversary over the span of many years, but a Red Teaming engagement by contrast is always constrained by commercial realities that must prevail.

There are two time components associated with a typical engagement.

(i) the **total amount of effort** (measured in time) that the Red team is allowed to expend; and
(ii) a **window of opportunity timeframe** that provides the element of surprise advantage

If the Red team is not provided with enough time to plan and execute their attacks, there may not be any opportunity to provide meaningful recommendations to help strengthen the security of the enterprise.

Conversely, if the Red team is provided too much time, the project runs the risk of losing momentum, and the information obtained in reconnaissance phases may become stale before attacks are executed.

Therefore, the length of effort in an effective Red Teaming engagement should be estimated carefully and determined commensurate with the collective skills and experience of the Red Team members. This is a process best undertaken in consultation with the qualified Red Teaming practice.

Furthermore, Red Teaming providers should provide clear evidence of their established methodology when justifying the time required. In support of this, a US Army definition states that Red Teaming should be a “structured, iterative process executed by trained, educated and practiced team members” with the “capability to continuously challenge plans, operations, concepts, organizations and capabilities”[3].

The second time factor to be agreed is the window of opportunity timeframe designed to give the Red Team an appropriate element of surprise. This is when the Red Team is given explicit authority to mount their attacks, at any time of their own choosing.

This window of opportunity also provides comfort to the target enterprise, and in some cases may involve negotiation around high-risk seasonal times or events that could have an undue impact. This may also depend on how widely the Red Teaming activity has been disclosed inside the enterprise.

Calculating the best opportunity timeframe for a Red Teaming engagement involves making an estimation based on a multiple of the total work effort required for the engagement. For example, if a Red teaming engagement is estimated to require four weeks of total effort, then calculating a multiple of this timeframe usually by a factor of three, is generally a reasonable proposition.

In practice, an opportunity timeframe ranging anywhere from 3 – 6 months is generally enough time to retain a reasonable element of surprise for the Red Team.

It is important to note that a window of opportunity that is too short may result, in some cases, with a Blue Team (enterprise defenders) lying in wait for the Red Team to strike; poised in a “red alert” status that would otherwise have worn-off after a longer period of time.

Conversely, having an opportunity timeframe that is too long is arguably going to leave the enterprise potentially vulnerable for a longer period of time; remember that the key advantage to Red Teaming is to provide feedback and recommendations to the enterprise to help improve their security.

Choosing a Red Teaming partner

Selecting a suitable Red Teaming partner, we believe, should rely on making comparisons based on a four-point criteria of: knowledge, experience, certification, and qualification.

Firstly, a suitable provider should be able to clearly demonstrate their knowledge in Red Teaming. This should be focussed on highlighting areas where risk to your enterprise can be minimised – such as understanding the legal and ethical challenges, and how their process and methodology will get results.

As stated in this document, a great deal of knowledge in this area can only be obtained through the experience of being actively engaged in Red Teaming activities. Mature and capable Red Teaming providers are generally those that have conducted multiple assignments already; those that have learned first-hand the issues involved, and captured that knowledge for your benefit.

Certifications and appropriate licenses are also important factors to consider, and this point doesn’t just refer to the commonly accepted IT security certifications either.

[3] https://www.army.mil/standto/archive_2014-07-22

Your prospective Red Teaming partner should be a registered local security practitioner with staff that has been police-checked and cleared. For example, Sense of Security is the holder of a Master Security License at an organisational level in New South Wales, and all staff have been police-checked.

When it comes to IT security credibility, your Red Teaming selection criteria should compare providers based on their teams having achieved some of the more technically-challenging certifications such as CREST[4] or OSCP[5]. Other certifications that involve generalist knowledge may not provide as much direct benefit to Red Teaming engagements, but still demonstrate a partner's overall dedication to security.

Moreover, your Red Teaming partner should be suitably qualified to perform the many tasks and activities that make such an undertaking successful. Established consulting practices with long-term commitment to the development of their employees, through ongoing education, participating in conferences, and delivering workshops, should therefore be first on the list of preferred providers.

The future of Red Teaming

As the popularity and adoption of Red Teaming grows, we believe it will likely continue to evolve and adapt to suit the needs of different types of enterprises.

One emerging concept is that of "Purple Teaming" whereby more collaboration between the Red and Blue teams is encouraged (hence resulting in the colour Purple) and this may have some interesting advantages.

Purple Teaming potentially gives the Red Team a much greater insider knowledge of the enterprise before mounting an attack. This insight may speed up Red Team engagement planning and help shorten the amount of effort required; since scenarios that are most likely to work can be quickly identified.

However, collusion with the Blue Team may inadvertently dilute the element of surprise – weakening the real-world nature of Red Teaming that would otherwise provide the benefit of independence.

Whatever the future of Red Teaming, we're excited by the prospect of continuing to learn and adapt, as we continue to provide this highly valuable engagement to enterprises of all kinds.

Summary

Red Teaming exercises are not reserved for the Military, or for enterprises that wait until they consider themselves to be mature – while their adversaries don't.

Although Red Teaming is not an insignificant undertaking, the immense value it provides to improve security resilience cannot be understated.

While it is natural for a lot of emphasis to be placed on the highly talented individual members of a Red Team, results are best achieved when combined with an entire process and methodology behind it.

Ultimately, a successful Red Teaming engagement requires flawless planning (70%) and execution (30%), delivered against a road-tested and proven process, combined with the experience of a talented team.

[4] https://www.crestaustralia.org/
[5] https://www.offensive-security.com/information-security-certifications/

About Sense of Security

Sense of Security Pty Limited is an Australian based information security and risk management consulting practice delivering industry leading services and research to organisations throughout Australia and abroad.

Our strategic approach to security provides our clients with a capability to understand the security risks relevant to their organisation and knowledge to protect their information assets. We provide expertise in governance & compliance, strategy & architecture through to risk assessment, assurance & technical security testing.

For more information, please contact us:

Web: **www.senseofsecurity.com.au**

Email: **info@senseofsecurity.com.au**

Phone: **1300 922 923**