

***THREAT INTEL
BASED ETHICAL
RED TEAMING***



SECURA

Vestdijk 59
5611 CA Eindhoven
Netherlands

Karspeldreef 8
1101 CJ Amsterdam
Netherlands

T +31 (0)88 888 3100
E info@secura.com
W www.secura.com

Follow us on   



Threat Intel Based Ethical Red Teaming

Introduction

Security is a wide and wild landscape. In a world with a large number of risks that belong to the category of ‘unknown unknowns’ and pushed along by sophisticated cybercriminals and nation state threat actors, financial institutions are combatting an ongoing flood of attacks. Coping with such events requires more than a dedicated SOC, it requires hands-on training, by doing. An increasingly popular way of testing and training in a controlled way is ‘Red Teaming’.

In this whitepaper, Secura will present its view on Red Teaming in the financial sector, especially focussing on the TIBER scheme (TIBER stands for Threat Intel Based Ethical Red Teaming) as promoted by the Dutch Central Bank, DNB (De Nederlandsche Bank). We will explain our vision on how to properly manage the processes surrounding successful Red Teaming (including how to manage operational risks, monitor execution and identify critical vulnerabilities), we take you through the detailed steps of TIBER and how we train the defenders (a.k.a. the Blue Team) to have a successful TIBER project execution.

Background

Originating in the military arena, Red Teaming is a security discipline that is gaining popularity in all sectors of critical national functions, especially in the financial sector.

By simulating full-spectrum information attacks, defenders get to practice their detection and mitigation skills in a managed and measurable way. If you want to measure how well your organisation detects and mitigates threats that are common in the threat landscape your organisation faces in the real world, then there really is only one way to know, and that is to test these processes by actually performing these attacks like a malicious attacker would. From spear phishing attacks to sophisticated cybercrime actors, and whether your detection capabilities are indeed seeing Advanced Persistent Threats (APTs). If you do not test this, you will never know your actual security posture.

Moreover, how people react and how well your incident response procedures are embedded in the organisation are of importance. The need for testing cyber resilience was first concretised in the financial sector in the UK, where in 2013, the Financial Policy Committee issued a recommendation to Her Majesty's Treasury. It is requested that regulators working with the core UK financial systems and the infrastructure providers that support them, have a programme to improve and test resilience to sophisticated cyber attacks. This became known as the 'CBEST'-program. The Dutch (and soon to be European) TIBER program is very similar to CBEST but differs in some aspects such as accreditation of Red Teaming providers and the precise role of Threat Intel providers.



Managing the Process

From the perspective of the financial institution being red-teamed, a simulated full-scope information attack needs to be managed in a very controlled way. The reasons for this are obvious: the risks on reputation damage and additional attack surface are non-negligible and must be tightly managed. From setting up communication channels to defining (de-) escalation paths and working with the ‘white team’: careful realism and preparation are key.



Figure 1: Key aspects for successful TIBER execution

Realism is also a requirement: cyber attacks follow patterns and scenarios. Working with Threat Intel (TI) means that the simulated attack tries to follow those patterns and scenarios as close as possible (although new and more sophisticated attacks are always being developed). Therefore preparing the scenarios whilst minimising the risks and staying within ethical and legal boundaries, is a large part of any successful Red Teaming exercise. Consider that a real attacker might choose to attack an organisation through a compromised home-PC of a system administrator working from home. This would not be permissible within RT because of legal restrictions. In the TIBER programme, we try to emulate these scenarios as much as possible, together with the white team who can provide a ‘leg-up’ in certain cases. However, managing the process starts with planning and preparation. This phase should be taken very seriously. A requirement is a dedicated project manager who works together with the Red Team lead and the White Team to create a schedule. Throughout the engagement, the schedule is followed, and adjusted where necessary. Risks and scenarios are assessed ongoing, and periodic white team feedback meetings are held. The project manager is the first point of contact for the white team.

Other important aspects of managing a TIBER engagement, are clear attribution, OPSEC and clean closure. Attribution is key to give fast and clear insight on whether the RT party is responsible for

reported incidents. Throughout the exercise, real attacks are performed by actual cybercriminals. The white team must be able to contact the red team 24/7 and find out, if it was actually the Red Team that performed a detected attack. If so, the white team can de-escalate and if not, then incident response procedures can be followed. This aspect of Red Teaming project management is only possible if all activities are meticulously logged and recorded by the red team.

Additionally, our own operational security (OPSEC) always has our attention: we are dealing with sensitive files and privileged access to our customers' core infrastructure, applications and data. Managing our own security during an engagement is therefore a generic critical success factor. Secura handles their own security in a variety of ways. First, we have our company security policies and practices, that are implemented using the ISO 27001 framework, for which we are certified. Second, we also provide our red team members with specialized methods for secure communications. Also, customer data is always stored on encrypted devices, and all copies (including backups) of any exfiltrated data are tracked and logged strictly. This means that in a moments' notice, we can provide the exact whereabouts and security status of any customer information we have gained access to during the course of the engagement. This ability has proven to be of immense value when dealing with extremely sensitive information.

Knowing what we have done exactly is also very important for the last step in any TIBER engagement: clean closure. Malware, fake accounts, manipulated devices, rogue network devices and copies of exfiltrated databases must all be cleaned up again. If Secura loses track of for instance a newly created domain administrator account, this could become a security exposure at a later time when it could potentially be abused by a real attacker.

Clean closure does not only mean managing the leftover digital remnants of the executed attacks. It also means providing the blue team with one or more evaluation sessions where the full timeline is replayed in a workshop. This gives the blue team the chance to see what they did detect, and what they did not. Don't forget, one goal of a TIBER engagement is to train the blue team. Without a detailed analysis together with the white, blue and red teams, that would be near impossible.

It should be clear that managing a TIBER engagement is in fact a significant portion of the actual work done during the project. This is why we use project tooling and create playbooks for all scenarios including backup scenarios that can be used if the preferred scenario does not play out the way we expect or want. Many security companies can send out phishing mails, plant malware or perform an internal penetration test. However, ensuring no new security vulnerabilities are introduced while





harvesting credentials via phishing, ensuring no other parties can be attacked using the used malware and providing near instant clarification on attribution is a different ball game. To bring those skills together in a managed and controlled TIBER engagement requires much more, and Secura has the experience and management processes in place to be able to do this in a repeatable and consistent manner.

Further measures must also be taken to minimise the risk to the customer and their reputation, and for this there will always be specific rules of engagement defined. Questions such as: is it allowed to use names of real people, third parties and vendors, when performing social engineering? Or are there any additional checks and authorisations to be performed before actually touching the 'crown jewels'. The rules of engagement will be drafted in the planning and preparation phase and reflect the risk appetite of the customer.

TIBER

In between the Planning & Preparation and Clean Closure steps as described above, lies the heart of any Red Teaming process: the 'kill chain'. These are the steps and activities that lead us to actual compromise, and from there, to the crown jewels that we are after. With near-military precision, we execute our playbooks, leverage any access we gain, moving sideways through the network while elevating privileges until we gain access to the desired goals, after which we perform the action that would be performed by an attacker in a controlled manner. An example of such an action can be the exfiltration of data or the execution of a specific financial transaction. The main steps in this kill chain deal with reconnaissance, exploitation of weaknesses, leveraging access, finding the crown jewels and exfiltration. In TIBER parlance these steps are:



The TIBER steps are presented in a serial manner, but in reality, it should be noted that (as in any real-life situation) any plans you make will rarely survive battlefield contact. This means that in the process, it is often necessary to move back some steps, for instance from 'elevate' back to 'gaining a foothold', because the attack was detected and mitigated by the blue team. Red teaming is like playing chess. You carefully put your pieces on the board, doing multiple and layered attacks, while protecting your king (hiding your presence). Below we will take you through each step.

Reconnaissance

An important objective in the TIBER scheme is to emulate realistic scenarios, using techniques and methods exactly like those used by real attackers. This is where the 'Threat Intel', or TI-part of TIBER comes into play. Within TIBER, the red team is provided with a general Threat Intel report for the financial sector of the country where the target is located. This TI report is built up by a TI-provider. Sometimes, this can be the same company as providing the red team. While the TI report delivers much information dealing with intelligence information on the threats and threat actors, for the red team the Target Intel is much more important. This is done through digital and physical reconnaissance.

If an attacker were to see your organisation as a target, and wanted to learn as much as possible about your internal processes, infrastructure and data, what could they find out? Any RT strategy starts with an information position, and the better the position, the better the strategy. Using mainly Open Source Intelligence (OSINT) a lot can be found out about a company and its employees. People divulge a lot of information on social media and websites such as LinkedIn, Twitter and Facebook, but also on technical forums such as StackOverflow, Tweakr and others. And don't forget the (meta-)data in all downloadable documents from the company website itself. Such (meta-)data will often reveal internal usernames, IP addresses, software versions used and other information that can make a difference to an attacker.

If, at a later stage, a scenario that includes physical access will be played out (for instance placing a rogue device in the network), then we will need to know how to get in to the offices of the target. Physical reconnaissance is therefore often required to find out things like what physical security measures are in place, where is the smokers-section (tailgating back into the office through the backdoor smokers' entrance can be a successful strategy) and if there are any publicly accessible network ports. Humans can also be used to obtain information (HUMINT). People who open specific email messages or click on links contained in them, disclose a lot of information about IP addresses, software versions, browser configuration and operating systems used. Everything combined, there is a wealth of information that can aid an attacker (and the red team) in their attacks.



Delivery

Delivering a malicious payload into the target network can take many forms. Deployment tactics that are often used in practice, are attacks that involve the human aspect, such as phishing attacks. They can be used to harvest credentials for core applications, but also deliver malware directly to workstations of users. The role of threat intel, target intel and reconnaissance is significant to perform this step successfully. To perform attacks such as these effectively, one will need to know email addresses, what spam filters are used, what email headers and footers are used et cetera. Besides this, the malware that will be delivered in this step will need to be tailored as much as possible to the workstations that are used by the customer to increase the chances of success and decrease the detectability of the attack. Also, the attack scenarios of threat actors common in the financial sector are relevant, in order to mimic those better.

Delivery of a payload can however also be done through a physical USB device, rogue network device, or compromised laptop. The individual scenario will be chosen based on metrics such as detectability and expected gain. In all cases, the delivery leads to exploitation of a vulnerability to gain a foothold.

Ultimately though, it is the detectability (by the blue team) of this step that is key. Early on in the engagement, we will deploy the stealthiest scenarios. This is because alerting the blue team cannot be undone: once they are actively hunting,

subsequent attacks have a higher chance of being detected. The noisier attacks with a higher detection probability are therefore only executed later on in the exercise. This build-up in attacks is often seen as plus in our approach to RT as it supports learning for the blue team efficiently.

Foothold

Gaining a foothold is achieved by successfully delivering an exploit, not being detected, and executing that exploit. This usually leads to a compromised system in the network of the target. The compromise itself in the case of TIBER, can take the form of an installation of our own piece of custom 'controlled malware', that then connects back to our own infrastructure and provides us access back in to the target network. The purpose of this malware in this stage is to obtain a stealthy foothold, that can act as a silent backdoor into the network of the customer.

Using this foothold, some low-key reconnaissance of the network can be performed without the risk of detection: DNS queries to map out internal host names, DHCP and WPAD requests to learn about routes and internet access, enumeration of local and domain windows users.

In order to connect back to the Secura Red Teaming C&C infrastructure, a number of hurdles must be taken by the 'malware'. First, it must (autonomously) find out how to get to the internet.

Also, since corporate proxies often terminate SSL connections, some form of data hiding (obfuscation) is necessary within the SSL tunnel. This makes such stealthy communication back to the RT slow, and prone to detection. However, once it is set up, it is possible (if used carefully) to retain such a foothold for a long time. Sophisticated real malware has been found in the wild, that uses the same tactics for evading detection.

Another method of gaining foothold that is also often used, but that does not have these specific drawbacks, is physical social engineering attacks. However, the main drawback of these attacks is that these must be physically delivered. Physical access opens a wide variety of different attacks that can be performed. From placing rogue devices in the network, to hiding keyloggers, to infecting workstations. Rogue devices, for instance based on small computers such as the Raspberry Pi, can be left in the network of the target, if physical access is obtained by the red team. The obvious (big) advantage is that this device could communicate out-of-band and therefore has a far lower risk of detection, while providing a foothold that can carry all kinds of tooling and interactive access for the red team.

Persist

After the foothold has been established, it is time to make sure that even if that one system is detected and removed, there are other ways to get back in. This step is highly dependent on what other vulnerable systems are found and further exploitation of the network and systems. This step could include, for instance, obtaining credentials that have a high level of privilege within the target organisation.

Another way of persistence can be achieved by having the remote-access malware remain dormant until it receives a trigger, for instance an email that is read on the system containing a link to a specific domain. Sending the email triggers the malware to become active again, after having been in hiding mode for some time.

There are a lot of variants of gaining persistence to be considered and actually they are only limited by the imagination and creativity of the attacker. In TIBER, when we try to emulate known *modus operandi* however, only a few often-used methods of persistence are relevant.

Move and Elevate

Moving through the network, more and more towards the crown jewels, is a very tricky part of any Red Teaming exercise. Within TIBER the steps move and elevate are separate, but in practice often these steps are combined, jumping from one server with certain access rights to another with more privileges.

It is highly dependent on the skill of the red team and blue team members, whether this stage will be detected. Obviously, the red team will start by





using methods that are not easily detectable. This includes dumping the RAM memory of a compromised workstation, looking for credentials that are valid on other systems in the network, and then using those to log in to other servers and workstations, harvesting credentials, password hashes and administrator account names as you go along. This will lead to the occasional failed login, but most detection mechanisms only trigger if the failed login count reaches a certain threshold.

The red team will focus on obtaining access to the previously defined goals of the Red Teaming exercise and it will not try to unnecessarily obtain access to other (confidential) information. In some situations the red team can directly elevate their privileges towards the crown jewels, and in some other cases it might be necessary to obtain Domain Administrator level privileges first. With such rights, the red team can install keyloggers and screen scrapers on any server or workstation, execute commands on all servers in the domain and add users with specific rights to many applications. While such usage of domain administrator rights is ultimately detectable, it is surprising how many organisations do not monitor the use of domain admin rights at all!

Exfiltrate

Once the crown jewels (or anything else interesting such as captured network traffic, the database of domain password hashes, or exchange email server database) have been reached, it is time to exfiltrate this data. Typically, this will be quite a bit of data, and moving it out of the network might be detected. This is where some creativity and patience comes in. Simply setting up an SSL session to a host on the internet and pumping data through might not be a good idea because the proxy server and firewall will both be on the lookout. Also, the SSL session might be terminated at the boundary for inspection by a Data Loss Prevention (DLP) solution.

If a rogue network device was planted with a separate 3G/4G connection, then that connection would be the most attractive way for an attacker to exfiltrate because it will not use the target's network: it will be out-of-band.

But also, if no rogue device is present it is possible to hide data in inconspicuous-looking requests. It is easy to hide data in variables of legitimate-looking web applications. However, this method will be extremely slow because hiding the data takes many bytes of unnecessary information. Patience is therefore required.

Success

When is a Red Teaming exercise a success? Some would say “when the crown jewels or flags have been reached without being detected by the blue team”. However, this definition also implies that the blue team will have learned little. On the other hand, it means that a plausible and realistic attack path has been exposed, that can now be closed or mitigated.

If you only consider the red team’s success, and not the blue team’s successes, you’re only seeing half the picture. We therefore strive for a balance. Yes, we try to successfully get to the flags, but we also look closely at what we think the blue team should be able to detect, and if we are successful, we will leave more and more indicators of compromise lying around so that in the end the detection level of the blue team is also assessed to a degree.

There is also another outcome that could be deemed a good lesson: we have not reached the flags but have also not been detected. This means that no undetectable path to the crown jewels has been identified within the given time frame, which is also worth noting.

Conclusion

From a Red Teaming perspective, the TIBER program is a framework that provides a common ground for red, white and blue teams to work with. However, the quality of the actual engagement is very much in the hands of the Red Teaming provider’s people and capabilities. Management of the engagement, project details, tooling, clean closure, risk management and reporting all hugely contribute to the success of the Red Teaming exercise. In some cases, these factors contribute more to success than the actual execution of the kill chain.

Secura’s experience in Red Teaming, combined with our capabilities, passion and TIBER-specific experience, provides our customers with the best possible basis for the clean, solid execution and management of TIBER engagements.

About Secura

Secura has worked in information security and privacy for nearly two decades. This is why we uniquely understand the challenges that you face like no one else and would be delighted to help you address your information security matters efficiently and thoroughly. We work in the areas of people, processes and technology. For our customers we offer a range of security testing services varying in depth and scope.

Secura has the mission to support organizations with up-to-date knowledge to work toward a bright and safe future.

Keep updated with the latest insights on digital security and subscribe to our periodical newsletter.

Contact us today at
info@secura.com or
visit secura.com for
more information.

SUBSCRIBE

TO OUR NEWSLETTER

