F-Secure

PURPLE TEAM

BUILDING RESILIENCE THROUGH COLLABORATION

An F-Secure Consulting whitepaper

THE F-SECURE GUIDE TO RAINBOW TEAMING

F-SECURE CONSULTING

F-Secure Consulting is a research-led cyber security consultancy, partnering with enterprises and early adopters worldwide. We exist to build resilience in an ever-changing digital world by providing evidence-based security advice. Our research drives service innovation, pushing the industry forward.

We're a multi-disciplinary team, equally intellectually curious and passionate about security. It's this that compels us to solve the world's most complex security challenges.

www.f-secure.com/consulting
Twitter: @FSecure_Consult
Linkedin: showcase/f-secure-consulting

CONTENTS

Driven by industry advancement in recent years, there is now a broader range of initiatives available to support the development of an organization's cyber security posture across the Predict, Prevent, Detect, and Respond (PPDR) model. Combined, these are colloquially referred to as a "Rainbow Team", delivering purple (collaborative), blue (defensive), red (offensive), and gold (crisis management) activities. When delivered sequentially and continuously, organizations gain the ability to utilize outputs from each development area and measure incremental improvement.

Each paper in this four-part series explores one such testing approach through the eyes of the teams – external and internal – leading and participating in the engagement. The aim: to demonstrate how the practical and technical delivery processes lead to real-world impact. For readers who have taken part in similar testing activities already, the series will help explain how to boost the benefits of that pre-existing investment.

The sequencing of rainbow teaming activities depends on the security testing and implementation your organization has carried out, and the experience of your security staff and senior security stakeholders.

THIS PAPER:

IDENTIFY HIGH-RISK ATTACK PATHS, AND USE THEM TO MEASURE AND IMPROVE PREVENTION AND DETECTION CAPABILITIES.

BACKGROUND

Purple teaming has traditionally been used to test and improve capability purely from a detection point of view, by utilizing the skills of red (offensive) and blue (defensive) teams in a collaborative simulation: sitting side by side, red attempting to attack or bypass controls in pursuit of an objective, with blue attempting to prevent them.

However, an alternative blend of exercises can be used to maximize efficiencies and increase the real-world relevance of purple teaming:

This two-pronged approach focuses on the likely paths an attacker would traverse in order to access an organization's critical assets, followed by an assessment of the offensive techniques the attacker could employ to remain undetected. The outcome is a comprehensive account of the organization's security posture. By enumerating multiple end-to-end attack paths and a vast array of attacker actions, this style of collaborative testing maximizes the value of the offensive and defensive exercises.

Organizations with a more mature security posture, those looking to prioritize the implementation of new security controls, and those that have inherited a new, unfamiliar environment are the most suitable candidates for these collaborative exercises. The learning would be valuable, for example, where a new CISO wishes to quickly gain an understanding of their organization's overall security posture and use those insights to build a roadmap for improvement. For those building their security from the ground-up, or for which security isn't a high-priority area, collaborative engagements will struggle to deliver valuable, actionable results.

The two exercises described are Attack Path Mapping (APM) and an Attack Detection Capability Assessment (ADCA). Together, they are used to assess organizations in a collaborative manner using defensive teams that contain individuals who are offensively trained.

APM and ADCA teams work together with business stakeholders in a continual cycle of discovery and feedback. In fact, collaborative engagements are less competitive exercises than they are a continuous dynamic. The buy-in and active involvement from these stakeholders is essential to the success of both exercises. Once completed, the results provide a baseline that can be cited to continually demonstrate detection improvement and ROI from people, process and technology.

ATTACK PATH MAPPING AND ATTACK DETECTION CAPABILITY ASSESSMENT

ATTACK PATH MAPPING

Collaborative engagements succeed when they are driven by a clear picture of the threats an organization is exposed to, as well as the risks that it wants to alleviate and manage. It's for this reason APM works so well. They highlight many paths to an organization's critical assets[1], demonstrating the TTPs a threat actor could successfully employ to achieve their goals. They also articulate the current posture of environments, both legacy and new, helping that organization prioritize the closing of paths and fixing of vulnerabilities.

Comparison of APM with other testing types:

Audit	Identifies the existence of controls but does not assess their effectiveness.
Penetration test	Tests the effectiveness of controls, but is narrower in scope and organizational context.
Red team exercise	Realistic and business-objective-led (esp. as a Targeted Attack Simulation (TAS)), but only evaluates one path, often that of least resistance. Also, red teaming is akin to black box activity, where consultants test with little or no prior insight into the client's security infrastructure.
APM	Determines the likely actionable paths an attacker would take to reach critical assets with the involvement of the client's SOC. Outcomes are qualified through recommendations.

[1] https://www.f-secure.com/en/consulting/our-thinking/what-is-attack-path-mapping

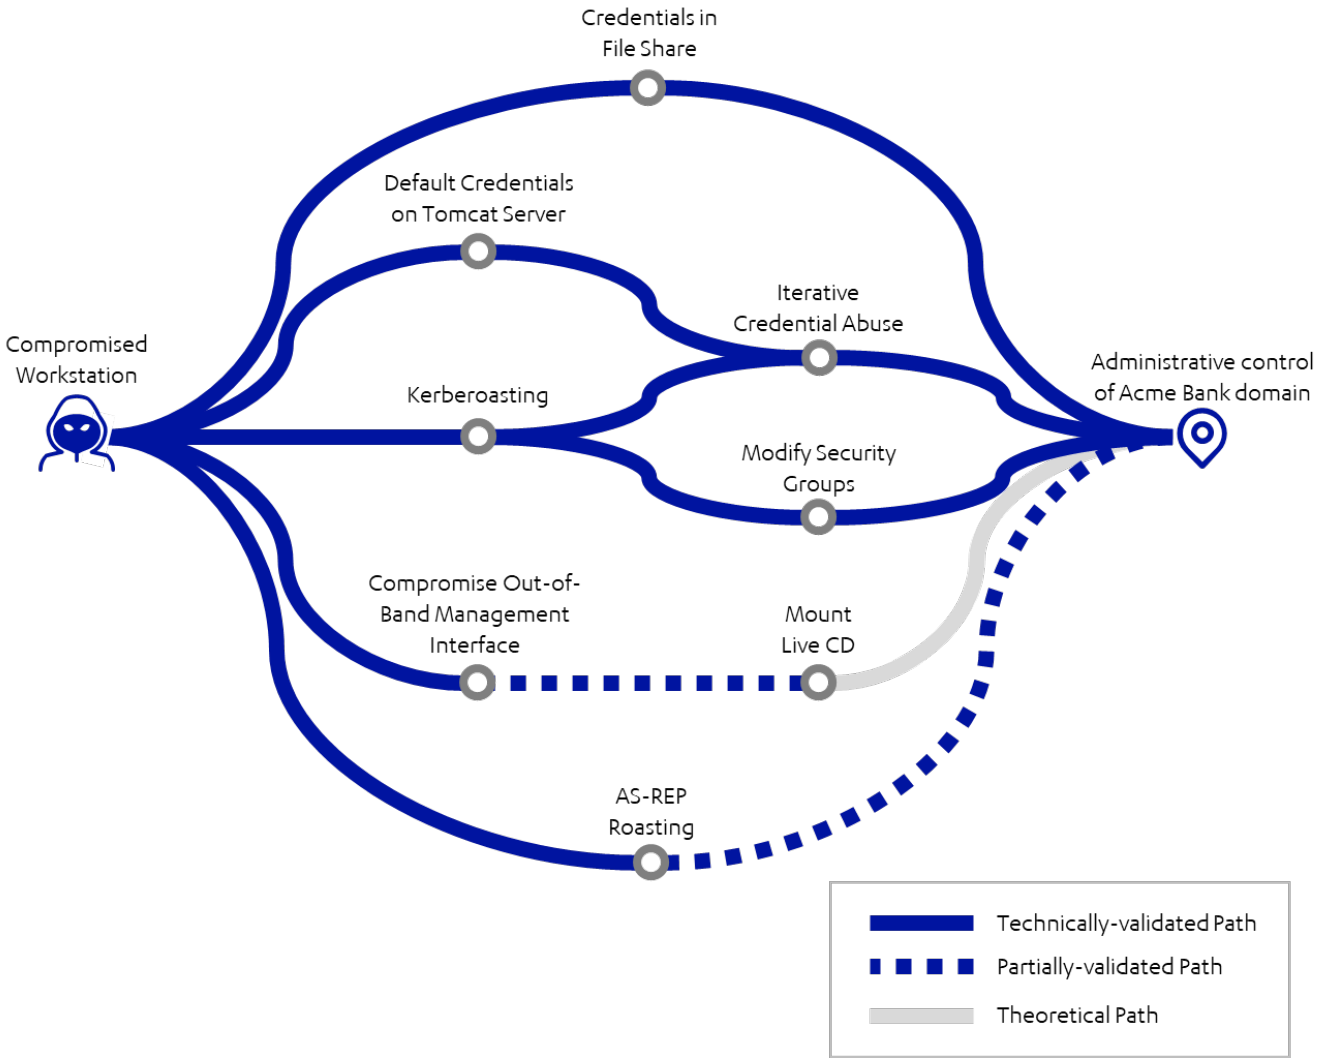

Fig. 1. Example APM graph showing steps to administrative control of corporate domain

ATTACK DETECTION CAPABILITY ASSESSMENT

An Attack Detection Capability Assessment (ADCA) works to answer the following questions:

- Is attack detection possible using the tooling and technology available?
- Do Security Operations Center (SOC) staff know how to identify and respond to malicious indicators?
- Are the necessary processes in place to ensure continuous improvement of detection capabilities?

If an APM has taken place, an ADCA may use the output of the exercise to test the organization's detection. This output comes in the form of an APM graph. When enumerated, opportunities for detective and/or preventative controls at the extremities of the graph are identified – highlighting the points where detection is vital. Though detection is the focus, a positive by-product of the engagement are insights into the business' preventative capability. This makes good business sense, as from the SOC's perspective, detection is vital. Even prevented attacker actions should be detected.

The role of the ADCA is to simulate attacks across identified paths before looking more generally at the spectrum of modern offensive techniques relevant to the organization. This aspect of the engagement is led by the mixed offensive and defensive team (purple team), and provides recommendations in answer to the following:

- What tactics, techniques, and procedures (TTPs) will threat actors likely use when specifically targeting the organization?
- What tools and knowledge does the SOC need to detect these attacks?
- Which of these attacks can be prevented by the organization?

In the short term, this enables the organization to determine the capabilities of its SOC – with regards to telemetry and actionable alerts. It is also valuable for the SOC to understand which attacks would be prevented within their environment, and to build its detective capability around techniques known to be effective within its environment. Longer term, it enables the development of a broader capability designed to detect an attack in its earliest stages, minimizing its overall impact.

WALKTHROUGH

This walkthrough provides an overview for the delivery of a collaborative engagement. For the purposes of this paper, we'll use a fictional client: Acme Bank. And to provide a true-to-life demonstration, we will base the walkthrough on the recent real-world engagements of our own purple team.

PHASE 0: PROJECT INITIATION

It's necessary for F-Secure to understand Acme Bank's perception of its security posture for true collaboration to take place. This includes what it considers to be its critical business assets, processes, and key risk events, as well as how they map to IT assets. For this reason, the first steps in preparation for the practical engagement are interviews and workshops with key technical and organizational stakeholders. Although these interviews are used to initiate activity, they won't be the last; interviews and collaborative exercises are scheduled throughout the engagement to ensure participant engagement and learning.

Acme Bank is a large financial enterprise with a global presence and over 30,000 employees, including a dedicated security team and 24/7 security operations. It uses many of the security processes and technologies common to the financial sector. In 2018, its SWIFT payment gateway was hit by a failed breach[2]. The subsequent post-breach investigation pointed to Carbanak, a cyber-criminal organization known best for their use of Cobalt Strike during targeted attacks[3]. The organization is keen to understand where existing preventative and detective controls failed, and in which areas of security additional investment would provide the greatest return. More specifically, it wishes to answer two questions:

- What are the likely attack paths towards the organization's most critical assets?
- How can I maximize its chances of detecting a threat traversing these attack paths?

An APM is selected to respond to the first; an APM will provide a series of recommendations aimed to prevent an attacker from traversing these paths in future. For the second, Acme Bank will need to understand what offensive techniques an attacker could employ, and which of these events will trigger alerts. Contrary to conventional wisdom, it is more important to detect every attacker, rather than every attacker action. As such, an ADCA is chosen to see where the SOC's detection can be improved.

[2] https://www.scmagazineuk.com/carbanak-active-latest-cyber-bank-heist-took-months-carry/article/1586531
[3] https://www.bitdefender.com/files/News/CaseStudies/study/262/Bitdefender-WhitePaper-An-APT-Blueprint-Gaining-New-Visibility-into-Financial-Threats-interactive.pdf

PHASE 1: APM

APM PLANNING

Planning for the APM begins by looking at Acme Bank's most critical assets. As with many financial services organizations, the applications and services listed as critical on their IT risk register fall into three categories:

- Trading platforms
- Settlement systems
- SWIFT network

The following objectives are then set:

- Gain a position to post a malicious payment and bypass the four-eyes check
- Obtain a foothold in the segregated SWIFT environment

One of Acme Bank's biggest concerns is their susceptibility to phishing campaigns due to their prevalence as an infiltration technique.[4] It is consequently decided that the APM engagement will begin from the position of an assumed compromise. The F-Secure team invites members of the SOC team to join and watch the exercise unfold and query decisions.

ATTACK POSITIONING

To simulate a successful phishing attack, the F-Secure team provides Acme Bank with a payload to execute on a workstation of their choice. The payload is then executed in the context of a new starter in the organization (their equivalent Active Directory (AD) privileges) thus demonstrating the potential to traverse any of the attack paths regardless of the compromised user.

F-Secure's payload generation tool, FoxDie, is used to create a C# project file containing the raw shell code for the Cobalt Strike payload (known as a beacon). When the .csproj file builds on the target system, the payload automatically decrypts itself and executes the malicious content, resulting in a foothold on that machine. This behavior of using built-in Windows functionality is often referred to as Living off the Land (or more colloquially, using LOLBAS[5]), and is a common technique employed by sophisticated threat actors[6].

As the .csproj file is encrypted, and therefore benign until it's decrypted and loaded into memory, endpoint solutions often find it difficult to statically signature. The team bypasses the endpoint anti-virus and the use cases configured by Acme Bank's SOC, achieving a foothold. This technique illustrates to the SOC the ways in which attackers may exploit native executables to obtain code execution.

[4] https://www.comtact.co.uk/blog/phishing-statistics-2019-the-shocking-truth
[5] https://lolbas-project.github.io/
[6] https://conference.apnic.net/48/assets/files/APIC778/Living-off-theLand-An-APT-case-study%20.pdf

Fig 2. Use of encrypted shellcode in .csproj file, executed with MSBuild to obtain a foothold

ATTACK PATH 1

The first attack path is designed to demonstrate lateral movement and privilege escalation to a domain administrator (DA) account.

Firstly, using Bloodhound[7], the team performs reconnaissance in AD, investigating opportunities for privilege escalation from their foothold position. A privileged service account is identified as the local administrator on eight servers – one of which has recently undergone maintenance from a domain administrator, signifying an easy opportunity to the purple team. A Kerberoasting[8] attack is used to obtain an encrypted Ticket-Granting Service (TGS) ticket, which, after a short period of offline brute-forcing, reveals the service account's plaintext password, "Sh@rePo1nt".

Leveraging the administrative privileges to one of the eight servers held by the compromised service account, PsExec from Windows SysInternals is used on one of the team's machines via the socket secure (SOCKS) proxy. This allows them to upload and execute the .csproj payload and establish internal Command-and-Control (C2) on the server. With this Cobalt Strike beacon, the team can attempt to retrieve the DA credentials in memory. Using the Mimikatz functionality built into Cobalt Strike, the DA's plaintext password is obtained without detection.

[7] https://github.com/BloodHoundAD/Bloodhound/wiki
[8] Abusing Windows implementation of the Kerberos protocol in an attempt to obtain the password of any account with an associated Service Principle Name (SPN). Only successful when the password is sufficiently weak.

ATTACK PATH 2

The second attack path focuses on exploiting misconfigurations in Acme Bank's network.

The team executes a port scan to search for other potential entry points across the estate, probing server IP ranges for commonly misconfigured services such as MQ, Tomcat, and Jenkins. They identify at least 10 Tomcat servers with default credentials[9], each allowing them to upload a web shell payload (shell.war) and obtain remote code execution.

Using a further LOLBAS technique[10], this web shell is used to download a payload with certutil.exe, and a foothold is achieved.

Tomcat is running as SYSTEM, the highest privilege level in the Windows user model. Hash dumping in Cobalt Strike is used to retrieve the NTLM hash of a local admin account.

Microsoft's Local Administrator Password Solution (LAPS) is partially deployed, and a pass the hash (PtH) attack against another server where a DA is logged in is successful; the user's credentials are extracted.

Using these attack paths, the team gains administrative control over the client's corporate AD environment.

[9] https://github.com/netbiosX/Default-Credentials/blob/master/Apache-Tomcat-Default-Passwords.mdown
[10] https://lolbas-project.github.io/

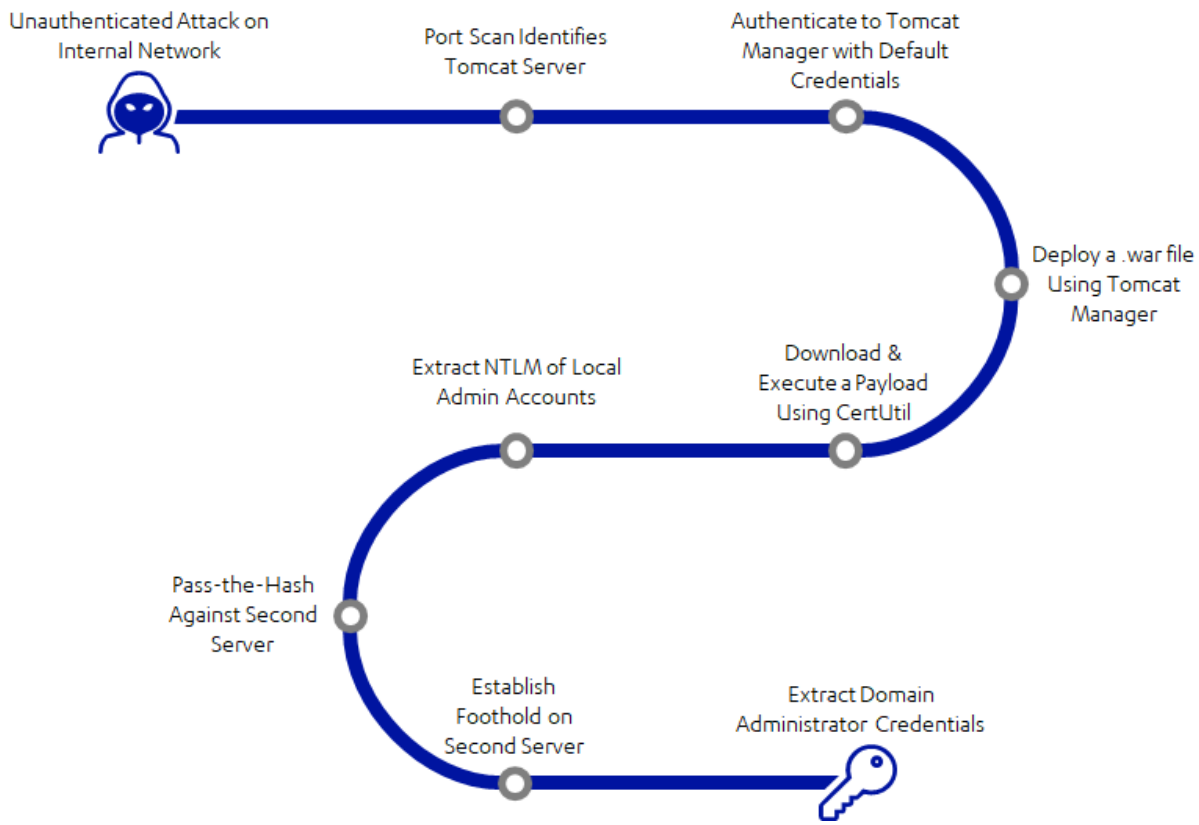

Fig. 3. Path of privilege escalation from instance of Tomcat

ACTIONS ON OBJECTIVES

With administrative control of Acme Bank's global corporate IT estate, the team may now leverage their privileges and progress towards the two objectives: gain a position to post a malicious payment and bypass the four-eyes check and obtain a foothold in the segregated SWIFT environment.

The team is pointed towards documentation available on Acme Bank's internal Confluence pages, providing instructions for onboarding SWIFT administrators to a jumpbox that straddles the Corporate and SWIFT networks. Using this information, it becomes apparent that access to the jumpbox is controlled using a different AD forest, which has no trust relationship with the previously compromised domain. Two potential attack paths are identified from the team's existing position to the compromise of the SWIFT network.

ATTACK PATH 1

Again, AD reconnaissance is carried out, and hosts are identified on the Acme Bank network where SWIFT administrators are logged in. This is done by searching for members of the SWIFT security groups within Bloodhound. By leveraging Impacket's wmiexec.py module[11], Windows Management Instrumentation (WMI) is used to move laterally to the workstation of a SWIFT administrator and deploy a Cobalt Strike beacon. A keystroke logger is set up through the beacon, so that the administrator's credentials can be captured next time they authenticate to the SWIFT jumpbox.

After waiting a short period of time, the target user logs into the jumpbox and their credentials are captured. The team waits until the user disconnects, then authenticates with the compromised credentials. They are now able to obtain a foothold in the SWIFT network, completing one of the assessment objectives.

[11] https://github.com/SecureAuthCorp/impacket

ATTACK PATH 2

Using DCSync, the team pulls all the usernames and NTLM password hashes from the compromised Corporate domain. They perform an offline brute force against them on their dedicated 6-GPU password cracking system, cracking 57% of all the passwords in 20 minutes.

The SWIFT usernames observed in Confluence appear to share the same format as the usernames in the compromised Corporate domain, enabling the team to identify accounts reusing passwords. A password spraying tool that targets Remote Desktop is used to attempt the username and password of all SWIFT administrators in the Corporate domain against the SWIFT jumpbox. Three accounts are successfully compromised, one of which is used to obtain a foothold on the SWIFT jumpbox, thereby validating another path towards the same objective as attack path 1.

The team meet with Acme Bank's SWIFT subject matter expert (SME) to discuss further attacks into the SWIFT environment. Armed with this information, multiple user accounts are accessed that each play a part in the legitimate payment workflow.

PHASE 2: ATTACK DETECTION CAPABILITY ASSESSMENT

ADCA PLANNING

In preparation for the ADCA, F-Secure's purple team works with Acme Bank's SOC to evaluate the APM findings; considering opportunities for the SOC to maximize its chances of detecting a threat traversing the attack paths. More widely, the ADCA hopes to achieve several objectives:

Provide a comprehensive understanding of the latest offensive TTPs, and those most likely to be employed against Acme Bank.

Identify where security analysts use appropriate detection measures to identify malicious indicators using available telemetry .

Identify the SOC's areas of strength, as well as where a lack of skills, tooling, or experience limit its ability to detect malicious activity.

ADCA BASELINING

The APM highlights the SOC's lack of actionable alerts needed for effective detection in the event of an attack. With this insight, the team is joined by other F-Secure consultants with detective specialisms.

This team will now work closely with the SOC to further explore its detection capability and deliver specific recommendations through the ADCA; specific focus is placed on building detection. The exercise ensures that the SOC is best equipped to detect malicious activity with the tools they have available. In contrast to the objective-led approach used throughout the APM, this ADCA is designed to provide Acme Bank with the best visibility of their current and potential detective capabilities.

The SOC is assessed by the purple team, in terms of its people, processes, and technology, with an overarching assessment of its maturity. Finally, recommendations are proposed based on what would have the greatest positive impact: be that the addition of a new log source to the SIEM, or offensive training for analysts to better understand and contextualize the attacks demonstrated through the APM.

The SOC is baselined across an adapted cyber kill-chain, which maps everything back to the MITRE ATT&CK framework. Where a typical red team might execute one or two techniques at each phase of the kill chain in order to achieve the objective, the purple team wishes to carry out as broad and as varied an array of attacks as possible, across multiple levels of perceived sophistication. This will show what the SOC is ultimately capable of.

Fig. 4. Adapted cyber kill chain

The purple team physically positions itself within the SOC, to familiarize itself with Acme Bank's tooling, as well as the people and processes at work. Analysts in the SOC team are also interviewed to enable the purple team to understand what their day-to-day activities look like.

It transpires that the SOC has an ELK-based detection platform, with Tanium deployed across the entire estate. This indicates an emphasis on endpoint telemetry, thus a relatively high level of visibility across many attacks. Given that most attacks target an organization's employees through phishing, spotting the initial compromise is vital – thus Acme Bank's strong endpoint visibility stands it in good stead.

To contextualize this capability with MITRE ATT&CK, the engagement team make use of their attack simulation tool, AttackSim, which automates attacker techniques. This frees up the team to focus on environment-specific and attack-path-specific techniques. The tool executes over 400 automated, common attacker techniques spanning the delivery, exploitation, persistence and internal reconnaissance phases of the cyber kill-chain - this includes sending emails containing malicious attachments and links to an inbox supplied by Acme Bank. The inbox is monitored to record which emails get blocked by the mail filters.

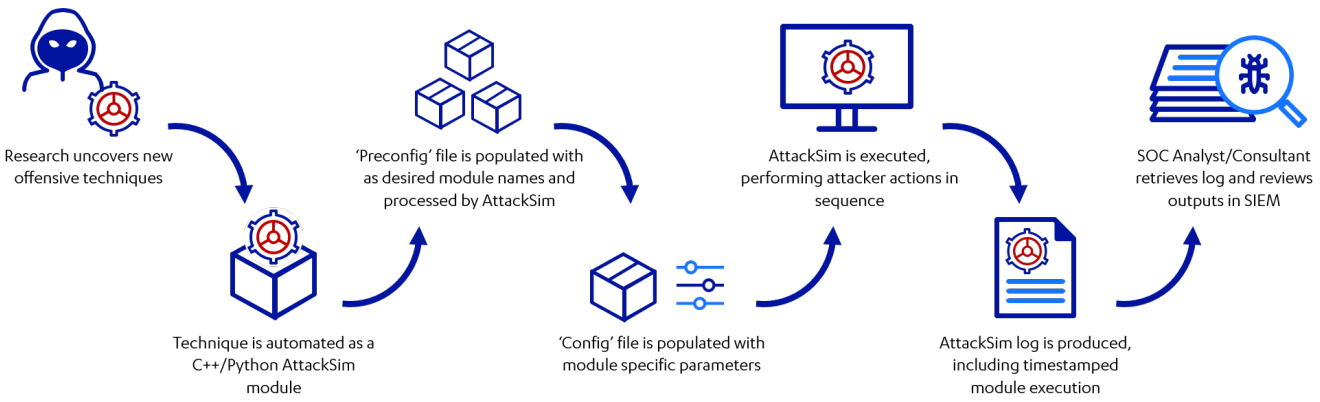

Fig. 5. Workflow of AttackSim module development and use

To complete the baselining technical testing, the team manually performs 50 additional attacker techniques. These are more advanced and sophisticated, and are determined using the outcomes of the APM report. For one test, the purple team carries out an RDP password guessing attack against the SWIFT jumpbox, replaying the actions of the APM. Working with the SOC, the team develops targeted detection for this attack, ensuring the SOC would be alerted should the same activity take place in the future.

The following information is logged for each attacker action:

- **Whether Acme Bank prevents the action.**
  For example, establish a C2 channel on a high port is blocked by the corporate firewall.

- **Whether Acme Bank is alerted to the activity.**
  For example, when using PowerShell to run a Bloodhound collector, an alert is generated by Tanium and raised to a SOC analyst.

- **Whether the telemetry exists for the given action.**
  For example, when a password spraying attack is performed against the domain, the SIEM records all the failed login attempts, but no use case had been configured to generate the alert.

THREAT INTEL

The engagement team now uses its internal threat intelligence platform to determine a list of threat actors that have historically targeted other financial institutions using SWIFT. Since it has previously attempted attacks on Acme Bank, Carbanak is identified as a threat actor of interest. This group uses the following key attacker actions:

- Persistence using local service creation
- Code execution using RunDll32
- Lateral movement using PsExec
- Credential dumping using Mimikatz functionality

With this knowledge, the team maps the activities of Carbanak to the data collected via the attack simulation process. It's determined that Acme Bank has the raw telemetry thanks to Tanium. However, because the organization does not have any use cases built for these specific actions, the above activities won't generate any alerts. Thus, should Acme Bank come under attack from Carbanak, it's likely that the threat would only be detected in its latter stages, if at all.

Other threat actors are modelled, providing estimates on the detective and preventative efforts that could be made against them. Importantly for Acme Bank, this data-driven approach means they are given an indication of where in the attack lifecycle they are most likely to detect each attacker.

As part of the exercise, the purple team decides to technically simulate one of these threat actors in a collaborative session with the SOC. The group FIN7 is used. The SOC analysts are asked to outline what they are detecting, telling the story of the simulated breach and threat actor. The whole SOC gains experience detecting an APT-level adversary in a safe environment.

Using the data collected, recommendations are presented in three key categories:

People

Hire more mid-level analysts to bridge the experience gap between the SOC manager and junior analysts.

Processes

Implement IR playbooks for all major categories of threat (ransomware, malware, phishing, etc).

Technology

Tune the mail filtering solution. During the delivery phase, this was identified as overly permissive, allowing executables and PowerShell scripts to be sent into the environment.

PHASE 3: ENGAGEMENT REPORTS

At its conclusion, each engagement provides Acme Bank with two reports:

APM report

This provides a high-level overview of Acme Bank's current security posture, and can be used to inform its overall security strategy moving forward, demonstrating where tactical and strategic improvements would be most beneficial. Changing default credentials on Tomcat servers and implementing MFA on jumpboxes, are two examples. The report also includes full technical details, including screenshots and commands executed.

ADCA report

The ADCA report gives Acme Bank a comprehensive overview of its existing detective capability, with a data-driven output that enables the organization to benchmark future capability improvements. Much like the APM, it provides recommendations relating to people, process, and technology, enabling Acme Bank to maximize the capability of its SOC.

The application of threat intelligence also adds more context to the findings of the ADCA, highlighting the likelihood of detecting Carbanak and other relevant actors, and the detection areas to focus on for greatest return on investment.

Considering the APM and ADCA reports together, Acme Bank is now aware of the potential routes that Carbanak might take (APM) and at what stages of the attack they would be best placed to detect them (ADCA).

SUMMARY OF OUTCOMES AND CONCLUSION

As a result of the APM and ADCA activities, Acme Bank has a clearer idea of which threat actors are likely to target them. Perhaps even more importantly, stakeholders now have a technical understanding of the steps these actors would take to do so.

Alongside this, Acme Bank's SOC analysts were able to spend time broadening their understanding of the latest offensive techniques and the ways to go about detecting them on the estate, all with the added context of which TTPs were effective from the APM.

The purple team carried out many activities via the ADCA – including simulating well-known threat actors Carbanak and FIN7. This gave Acme Bank the opportunity to see how its SOC would stand up against a targeted attack in a safe environment. Hundreds of other discrete attacker actions were also carried out to test its detection capabilities, and following that baseline testing, provided Acme Bank with a data set to demonstrate everything from control efficacy to an ongoing improvement in detective capability.

Overall, the key takeaways of the collaborative testing engagement were:

- An appreciation of the overall security posture of the organization
- An understanding of the performance of previous security investments
- The current and potential performance of a SOC in defending against the latest offensive techniques
- Actionable recommendations across its people, processes, and technology
- An understanding of the modern TTPs used by APT groups likely to attack, and which preventative measures to take in advance

REFERENCES

[1] **What is Attack Path Mapping?**
https://www.f-secure.com/en/consulting/our-thinking/what-is-attack-path-mapping

[2] **Carbanak still active, latest cyber-bank heist took months to carry out,**
https://www.scmagazineuk.com/carbanak-active-latest-cyber-bank-heist-took-months-carry/article/1586531

[3] **An APT Blueprint: Gaining New Visibility into Financial Threats,**
https://www.bitdefender.com/files/News/CaseStudies/study/262/Bitdefender-WhitePaper-An-APT-Blueprint-Gaining-New-Visibility-into-Financial-Threats-interactive.pdf

[4] **Phishing statistics 2019 - the shocking truth,**
https://www.comtact.co.uk/blog/phishing-statistics-2019-the-shocking-truth

[5][10] **Living Off The Land Binaries and Scripts (and also Libraries),**
https://lolbas-project.github.io/

[6] **Living off the Land: An APT case study,**
https://conference.apnic.net/48/assets/files/APIC778/Living-off-theLand-An-APT-case-study%20.pdf

[7] **Bloodhound,**
https://github.com/BloodHoundAD/Bloodhound/wiki

[9] **Apache Tomcat Default Credentials,**
https://github.com/netbiosX/Default-Credentials/blob/master/Apache-Tomcat-Default-Passwords.mdown

[11] **impacket,**
https://github.com/SecureAuthCorp/impacket

F-Secure

www.f-secure.com/consulting