

RED TEAM

BUILDING RESILIENCE THROUGH TARGETED ATTACK SIMULATIONS

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.

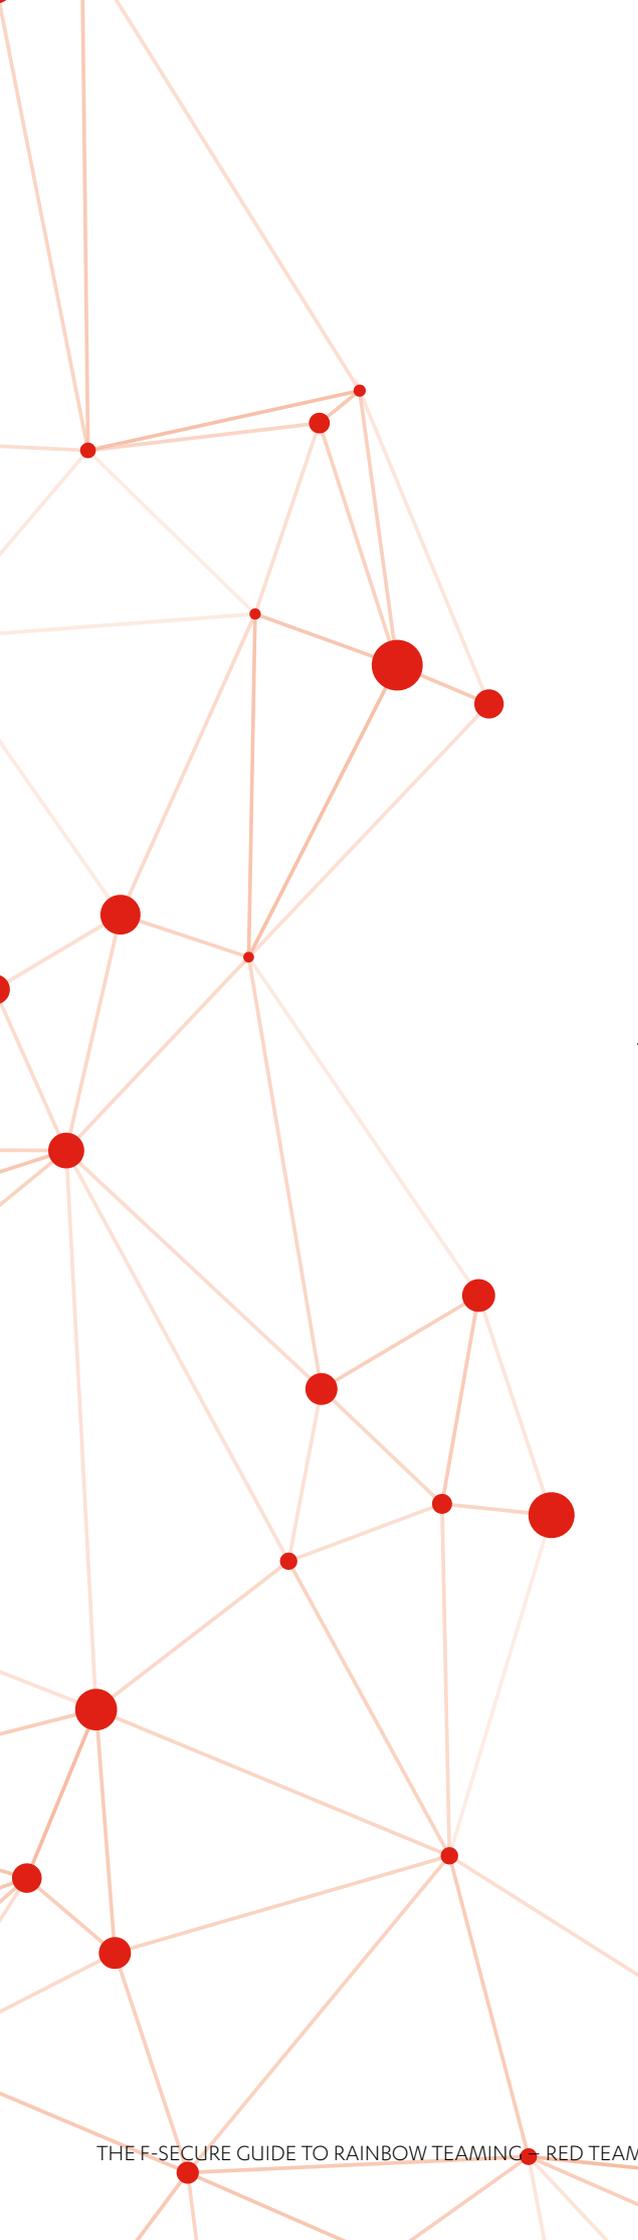
www.f-secure.com/consulting

Twitter: @FSecure_Consult

LinkedIn: [showcase/f-secure-consulting](https://www.linkedin.com/showcase/f-secure-consulting)

CONTENTS

INTRODUCTION TO RAINBOW TEAMING	01
BACKGROUND	03
WALKTHROUGH	04
PHASE 0: PROJECT INITIATION	04
PHASE 1: ATTACK POSITIONING	05
INITIAL RECONNAISSANCE	05
FOOTHOLD	06
INTERNAL RECONNAISSANCE	07
LATERAL MOVEMENT AND PRIVILEGE ESCALATION	07
FULL CONTROL OF ACTIVE DIRECTORY	09
PHASE 2: ATTACK EXECUTION	11
PHASE 3: BREACH NOTIFICATION AND RESPONSE COLLABORATION	11
SUMMARY OF OUTCOMES AND CONCLUSION	12
REFERENCES	15



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:

A network diagram background consisting of red circular nodes connected by thin orange lines, forming a complex web-like structure across the page.

**ASSESS
DETECTION
AND RESPONSE
PERFORMANCE
THROUGH AN
AUTHENTIC
TARGETED
ADVERSARY
SIMULATION
EXERCISE.**

BACKGROUND

Achieving total security remains an impossible task. Any organization can be compromised by a motivated, persistent, and capable attacker – the variable is how much time and effort the organization's defensive posture demands of the attacker. The organizations most resilient to attack are those with the capabilities to effectively detect and respond.

Targeted Attack Simulations (TAS) provide an opportunity for organizations to experience a realistic attack with specific objectives that would undermine its ability to operate. As a result, the target's defensive capabilities are exercised and enhanced by enabling processes to be rehearsed under the pressure of a major incident.

The key aspects of a successful TAS are:

Attack objectives that align with business risk

The potential outcome of a successful attack must resonate with senior executives. Without a broader focus on the organization and its business processes, security problems will continue to be perceived as an IT problem, as opposed to a critical business risk.

Collaborative learning opportunities designed to upskill defensive teams

A TAS should provide an opportunity to experience, experiment, and learn. It is not a pass or fail measure of performance that requires 'controlling' and 'damage limitation'. Instead, the attack activities must be used as part of a cooperative training exercise that allows defensive teams to observe how an attack might proceed in their environment and to question the attackers who undertook it.

Realistic attack tactics and techniques

The attack should progress in a manner that is not limited to isolated systems or attack stages that prevent the chaining of realistic attacker actions that lead to a realistic goal. A TAS must also be informed by the offensive tactics and techniques observed in real-world incidents as well as cutting-edge research from the offensive security industry. This will provide a learning experience to best prepare defensive teams for a real threat actor.

WALKTHROUGH

The following walkthrough depicts a TAS delivery. 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 TAS team.

PHASE 0: PROJECT INITIATION

Acme Bank has engaged F-Secure to perform a simulated attack on the organization. The TAS team technical lead and delivery manager from F-Secure meet with the key stakeholders from Acme Bank to define attack objectives and project management scope.

Acme Bank is a large financial enterprise, with a global presence and over 30,000 employees. This includes a dedicated security team and 24/7 Security Operations Center (SOC). It has invested in standard security hygiene, as well as improving detection and response capabilities against well-known attack techniques via collaborative, purple teaming activity executed 12 months prior.

The bank now wishes to undergo a simulated attack in order to provide its security team with a realistic challenge and learning experience. Present at the kick-off meeting are the CISO, head of cyber security, and an internal project manager.

The meeting is used to define the following aspects of the scheduled TAS:

Attack objectives

The core objective is to compromise payment systems and demonstrate the plausibility of an unauthorized transfer.

Circle of trust

The group of people within Acme Bank that knows about the project is defined. These individuals – also known as the “white team” under schemes like TIBER – will receive regular updates on attack progress.

Secure communications

Mechanisms for secure delivery of documentation and progress updates are defined. This is necessary to allow the TAS team to provide technical details of how the organization was compromised, without exposing it to any further risk.

Risk management

Rules are defined to state how the TAS team should proceed if a critical risk is identified that requires immediate remediation, or if the investigation of a real attack could be hindered by their activities.

Incident escalation

The standard incident escalation procedure is discussed to ensure that, in the event the TAS team's activities are discovered, the escalation and recovery process does not unintentionally harm Acme Bank or the team itself – for example, through being reported to regulatory bodies or the police.

This information is captured in the project initiation document and approved by both parties before the attack activities are launched.

PHASE 1: ATTACK POSITIONING

INITIAL RECONNAISSANCE

The attack is initiated with a reconnaissance phase. The TAS team reviews Acme Bank's Internet-facing estate, enumerates potential phishing targets via social media, and identifies security controls that could hinder compromise attempts and post-compromise actions.

As part of the estate review, certificate transparency logs maintained by Google are used to identify subdomains for which Acme Bank has recently registered SSL certificates. A certificate for "interns.acmebank.com" is found among the subdomains. This page is found to host a series of programming challenges for technology graduates as part of Acme Bank's annual internship program within publicly accessible Git repositories.

The commit history for these repositories is searched for sensitive data, and the TAS team happens upon a set of credentials for a document-sharing platform used by Acme Bank. Though it appears to be used for sharing files with clients, vendors, and interns, the credentials only provide access to a handful of documents relating to the organization's internship program.

Still following the trail of its internship program, a post on social media from Acme Bank's Head of Tech reveals the company is currently recruiting for its cyber security detection engineering team. The post tellingly states that the ideal candidate has experience with feeds from a popular Endpoint Detection and Response (EDR) agent and SaaS mail filtering and sandboxing solution.

The TAS team has experienced this same email filtering solution in previous engagements, where it prevented numerous phishing attacks by capturing and subsequently blocking malicious attachments. After conceding that phishing will not be the most efficient route to establish a foothold, the team's strategy is altered.

Returning to the compromised document-sharing platform, they discover in-built functionality that "invites" target users to view documents via the email address Securedoc@AcmeBank.com. The TAS team takes the opportunity to abuse the trust which Acme Bank employees potentially place in this system. A malicious macro-enabled document is uploaded to the platform within the internship folder, and target users are invited to view the document.

The attack team attempts to increase its chances of bypassing Acme Bank's EDR agent, the macro embedded within the malicious document leverages the parent process ID (PPID) spoofing technique. PPID spoofing alters the hierarchy of processes to evade detection by allowing the execution of a malicious process from any legitimate parent process. This masks any suspicious parent-child process relationships – for example, a parent Microsoft Office process creating a child rundll32 process designed to further execute a library of code and procedures.

Closely monitored commands, such as "regsvr32.exe" and "wmic.exe", are also avoided. These are typically used by threat actors and recognized as such because of their ability to provide attackers with legitimate functionality to execute code. Instead, the team relies on memory injection to migrate their malicious code out of the Office application and into another process.

FOOTHOLD

The macro-enabled document is opened by the second employee targeted by the team. A “low and slow” HTTPS command and control (C2) profile is used to establish communications with the TAS team’s attack infrastructure. The requests and responses are disguised to look like benign web traffic and avoid scrutiny by intercepting proxies.

A foothold has been successfully established within the environment. The TAS team’s priority is now to obtain some initial situational awareness and attempt to spread access to other systems before performing “noisier” post-compromise activities more likely to trigger detection. The attackers wish to solidify their presence in the network as soon as possible. By doing so, even if these activities are caught further into the compromise, it will be more difficult for Acme Bank to fully contain or eradicate their presence.

A document store is identified mapped to the local machine. After a quick investigation, this is found to contain financial models in the form of macro-enabled spreadsheets. The compromised user has editing permissions for the files, and it becomes evident to the TAS team – due to the multitude of files and “last modified” timestamps – that they are in frequent use by many individuals. Exploiting their prevalence, the team adds extra macro code containing a malicious payload to one of the documents.

Using the C3 framework¹ to establish communications over legitimate Office 365 applications, the TAS team enables a backdoor when the document is opened – all whilst blending in with legitimate network traffic. Where many organizations rely on factors such as domain age and categorization when building a strategy to detect anomalous network traffic, the C3 framework bypasses these detection strategies. It does this by piggybacking on trusted and well-known services that are used for legitimate business purposes. This use of Office 365 emulates advanced real-world threat actors who are increasingly leveraging legitimate services for C2². C3 provides the TAS team with a robust and difficult-to-detect foothold within Acme Bank’s network.

^[1] <https://labs.f-secure.com/blog/hunting-for-c3>

^[2] <https://attack.mitre.org/techniques/T1102>

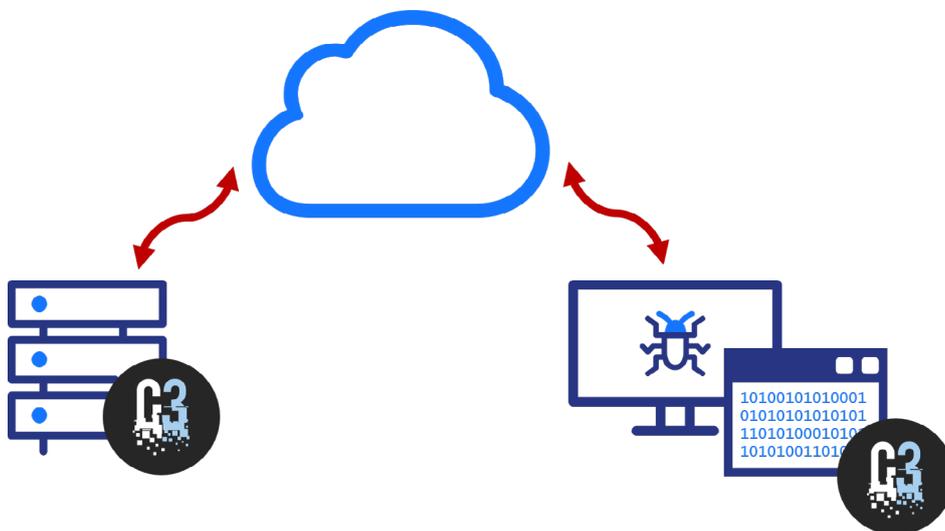


Fig. 1. The C3 framework allows attackers to funnel traffic through legitimate applications, such as Office 365

INTERNAL RECONNAISSANCE

With abundant access to the network, the team begins to carry out more extensive internal reconnaissance and credential-gathering activities, including:

Taking a copy of Active Directory (AD) information for offline analysis

This is taken by querying for data from LDAP. The information will help identify control paths that can be abused to compromise target users and systems.

Performing a Kerberoasting attack on service tickets

Once subjected to offline password cracking, service account passwords are recovered. A password for a service account associated with several TEST systems is found.

Searching file shares and intranet portals for passwords

A spreadsheet is discovered that contains SSH credentials for Linux servers in the TEST environment.

BLUE TEAM EVENT

A high-priority alert is raised when a Kerberos service ticket is requested for a honeypot service account. This account and its associated detection rule were created in response to findings from the organization's past collaborative engagements. The SOC uses the EDR agent and other log sources for remote investigation of the source workstation.

Other indicators of compromise are identified based on the team's understanding of modern attack techniques:

Outbound traffic is observed proceeding towards a domain that has never been visited by any other employee. This activity only began a few days prior to the alert. The SOC decides that this pattern of traffic indicates C2 behavior.

An anomalous scheduled task is identified running a Dynamic-link library (DLL) every morning. The SOC determines this to be the persistence mechanism installed by their attackers.

No further attack activities are detected relating to this incident, so it is marked as "eradicated" in the SOC's incident tracker.

A review of the network traffic originating from the compromised machine finds no evidence of direct lateral movement to other workstations or servers. No other hosts on the network have communicated with the C2 channel, and a decision is made to quarantine the host to contain the attack.

The affected user is supplied with a new laptop, and their account is marked for additional monitoring over the next week in case the attackers attempt to regain access to the network.

LATERAL MOVEMENT AND PRIVILEGE ESCALATION

Before the team can traverse further systems using the stolen credentials, the initial compromised workstation goes offline. Access doesn't appear the following day, and detection is assumed. Other workstations that were compromised with the backdoored document continue to operate, however. This enables the TAS team to gain access across three more systems in the TEST environment: one Windows and two Linux servers. They believe that despite the initial foothold being quarantined, their movements were quick enough to obtain multiple redundant access points in lightly monitored areas of the network.

Mimikatz is used to gather credentials from memory on the Windows server. This includes the plaintext passwords for several TEST domain accounts. There is no direct trust relationship between the TEST and PROD domains that can be abused, but a TEST user's corresponding account in the PROD domain is found to have the same password.

The data collected by the team's copy of AD shows that this account has administrative access to a number of production servers, including one with permissions for unconstrained delegation.

Delegation is a legitimate feature within Windows that allows a server or service account to impersonate other users. A typical example of delegation would be a web application legitimately impersonating users when interacting with a database.

Unconstrained delegation – the most powerful form of Windows delegation – can be leveraged to impersonate any user that requests access to a service on a compromised host through Kerberos through the following process:

1. User makes a ticket request to a service with unconstrained delegation
2. A copy of the Ticket Granting Ticket (TGT) gets added to the ticket
3. The TGT becomes available to the service when the client presents the ticket during authentication
4. Client gains ability to impersonate user

The exploitation of unconstrained delegation relies on two key prerequisites:

- Administrative control over a server that's configured for unconstrained delegation
- The ability to coerce a user to authenticate over Kerberos to a service on the compromised host.

Should these requirements be met, the following process occurs (see fig. 2.):

1. User performs a Kerberos authentication process with the domain controller (DC) in order to get a ticket for the target service on the compromised server
2. The domain controller stores a copy of the user's TGT in the service ticket
3. The user presents the service ticket to the server
4. The server decrypts the service ticket using the service account password hash
5. The user's TGT is extracted from the service ticket and stored in memory on the server
6. The attacker, with administrative control over the server, extracts the TGT from memory and uses it to authenticate to the domain controller as the user

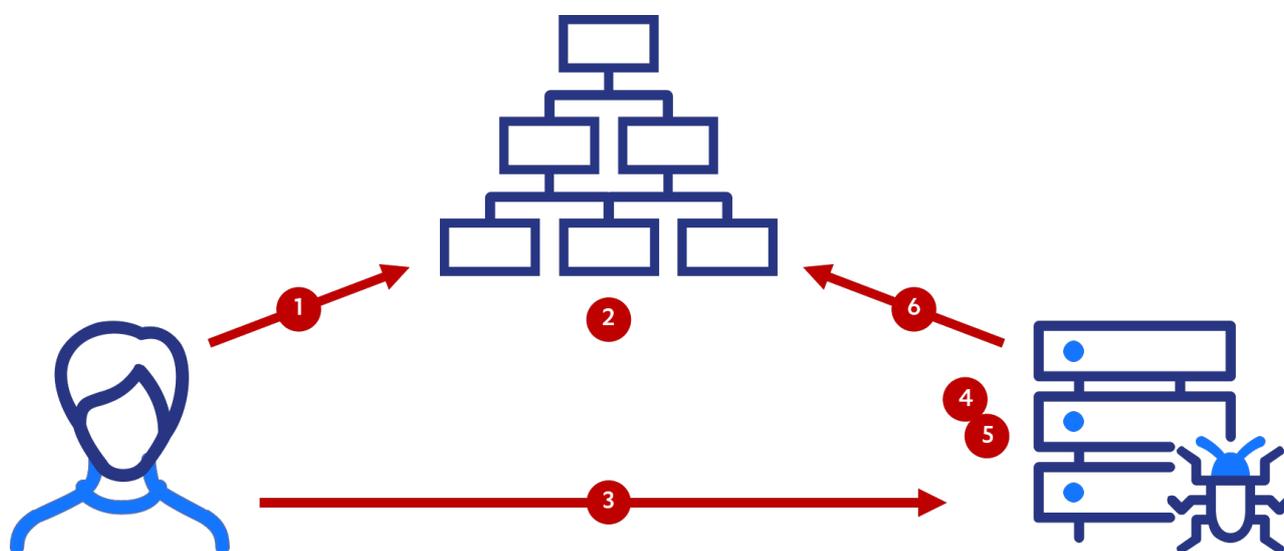


Fig. 2. Exploitation of unconstrained delegation

FULL CONTROL OF ACTIVE DIRECTORY

A Microsoft Remote Procedure Call (RPC) is made to a domain controller, triggering it to authenticate to the compromised server with unconstrained delegation. As a result, the DC sends a copy of its Kerberos TGT to the compromised server. The TAS team extracts the TGT from memory, consequently performing a DCSync attack³. This attack allows the team to impersonate a DC, giving them the privilege to request password hashes for any user account in the PROD domain.

These hashes can subsequently be abused to impersonate any user within the compromised domain through Pass-the-Hash (PtH) attacks. Full control over the forest can also be achieved through a golden ticket attack, which leverages the “krbtgt” account password hash to impersonate privileged user accounts with arbitrary group membership.

^[3] <https://attack.mitre.org/techniques/T1003>

PHASE 2: ATTACK EXECUTION

With the ability to impersonate any account, the TAS team is now positioned to carry out actions on objectives. The circle of trust is notified of their progress in order to enable subject matter experts (SMEs) to validate the team's attack plans on critical systems.

Close to the target, the TAS team now focuses on developing a detailed understanding of Acme Bank's payment architecture and processes. It's essential at this late stage to carry out the objective in a controlled and targeted way. The team begins by inspecting the target systems, payment flows, and relevant security controls using the following sequence of actions:

- Gathering documentation from internal information stores, including SharePoint and Confluence
- Compromising one of the PROD payment application servers by impersonating a PROD support administrator
- Inspecting server and application logs in search of constituent payment message processing components

The TAS team discovers that the payment server runs a Java application, which reads and writes SWIFT payment messages to an IBM Message Queue (MQ). These messages are processed downstream by additional payment systems and transmitted via the SWIFT network.

Application configuration files are retrieved from the server. It's possible to extract the credentials used by the application to access payment MQ channels. As a result, the attack team can connect to the MQ service with the same level of privileges as the application itself. MQ network traffic is pivoted through the compromised server to demonstrate access to the target message queues. Since the TAS team is interacting directly with the backend message queues, it is believed that the four-eyes approval controls enforced by the application itself can be bypassed.

Using its practiced knowledge of the application and SWIFT messaging, the team develops a proof-of-concept payment instruction to transfer a small sum between two accounts owned by Acme Bank. This message and the proposed attack are validated with Acme Bank SMEs before being sent to the target payments queue. The message is processed successfully, thus demonstrating an end-to-end attack path to make arbitrary money transfers.

PHASE 3: BREACH NOTIFICATION AND RESPONSE COLLABORATION

BLUE TEAM EVENT

A tier one analyst begins triaging a security alert for execution of a suspicious PowerShell command. The detection rule has highlighted the use of PowerShell with an encoded command:

```
powershell.exe -nop -w hidden -enc JAB-
ZAD0ATgB1AHCALQBAGIAagB1AGMAdAagAEKAT-
wAuAE0AZQBtAG8AcgB5AFMAdABYAGUAYQBtAC-
gALABbAEMAbwBuAHYAZQByAHQAXQA6ADoARgB-
yAG8AbQBCAGEAcwB1ADYANABTAHQAcgBpAG4A-
ZwAoACIASAA0AHMASQBBAEEAQBBAAEEAQBBAAEEA-
QQBLADEAVwBhADIALwBpAFMAaABMADkAbgBQAHCAS-
wBmAHgAZwBKAFUAQ . . .
```

The analyst notices that the process was initiated by an administrative service account associated with a centralized backup solution. A brief search with the EDR solution finds no other similar PowerShell usage by this account. Furthermore, the affected server is identified as a jump host used by database administrators. They escalate the incident for further investigation because these facts are indicative of the presence of a high level of compromise within the network.

The SOC begins to hunt for evidence to understand:

- The purpose of the PowerShell command in order to confirm the activity is malicious and show what the attackers might be targeting
- Which other servers may have been compromised using the backup service account
- How the attackers obtained access to this account and to the network initially

The SOC begins by restricting internet access from the affected server, then logging in to the server in order to retrieve more detailed PowerShell event logs.

The suspicious command detected by the SOC has been run by the TAS team in order to intentionally trigger the incident response process. As the attack objectives near completion, the breach notification process commences to ensure that the defensive teams have an opportunity to attempt their investigation and response to a widespread compromise within the network. This allows them to exercise their attack detection and response playbooks, as well as the performance of their tools during a real incident.

An administrative service account has been impersonated via a golden ticket attack in order to compromise multiple jump hosts used for sensitive administrative operations. These include production support for central database clusters and the core payment application. A PowerShell command matching a common pattern of malicious execution has been run on the database operation's jump host. Had this not been detected, noisier activities would have been performed until an actual breach notification document was issued to point the SOC towards a specific indicator of compromise.

The jump hosts were being controlled with an internal C2 channel based on SMB communications. Therefore, partially quarantining the server by restricting internet access did not affect the attack team's ability to interact with the server. When the incident responder logged into the server using remote desktop, their credentials were stolen from memory using Mimikatz. The TAS team use these credentials to compromise additional systems.

BLUE TEAM EVENT

The SOC's analysis of centrally-collected event logs for Windows authentications finds that the compromised service account has been used for suspicious access to three administrative jump hosts. The access in these cases is identified as anomalous due to the source of the logon (a user workstation) and the time at which it occurred – the normal backup process usually occurred out of hours.

The lead investigator identifies a jump host associated with a production payment application. The decision is made to declare a major incident and escalate to the business for potential regulatory reporting requirements.

At this point, the incident response team is informed by the circle of trust that this is part of a simulated attack being performed by a red team. The incident will not be escalated any further in order to avoid any business disruptions. However, the investigation should be continued as part of an exercise in understanding the effectiveness of internal incident response processes.

A member of the F-Secure incident response team is briefed on the full extent of the attack activities before joining the bank's defensive teams to observe and support their investigation. This allows F-Secure to closely observe Acme Bank's detection and response through interactions with its people, processes, and tools, allowing for the engagement reports to include detailed analyses and recommendations. Furthermore, it provides an opportunity for the SOC to obtain 'on the job' incident response training on best practices and approaches to dealing with a live incident.

At the conclusion of the technical delivery phases, Acme Bank's defensive teams have been able to identify many of the hosts compromised by the TAS team and block associated C2 channels. Some aspects of the attack remained undetected, due to the scale of compromise and use of esoteric C2 channels via the C3 framework.

Following this, the TAS delivery team provides the SOC with a full debrief of the attack path and their observations from the attackers' perspective. This also provides them with an opportunity to question the attackers face-to-face and understand why certain actions were, or were not, taken. The discussions between the two groups are used as further input towards the final deliverables.

SUMMARY OF OUTCOMES AND CONCLUSION

A key objective of the TAS was to provide Acme Bank with the opportunity to experience, experiment, and learn from a real-world targeted attack. This adds to the return on defensive investment, areas of improvement, and an understanding of their exposure to advanced adversaries.

During the engagement, for example, the TAS team performed a kerberoasting attack and successfully compromised a set of credentials. This action also triggered an alert in Acme Bank's SOC, due to a ticket request being raised for a fake service account set up as a honeypot. The SOC demonstrated their ability to respond to this alert and effectively quarantine some of the affected hosts. Acme Bank was thus able to determine the efficacy of this particular defensive countermeasure.

Through the collaboration between Acme Bank's SOC and the TAS team, an awareness of the full attack path and its associated Indicators of Compromise (IoCs) was developed. This activity provided a significant improvement in Acme Bank's hunt capability, giving the SOC analysts the opportunity to: develop new playbooks, increase the efficiency of detection and response, and enhance the coverage of their tools. These takeaways are drawn directly from the real-world experience of responding to active attackers operating within Acme Bank's network who themselves react to the SOC's attempts to evict them.

Overall, the key takeaways from the TAS included:

- Significant improvement in hunt capability, tool coverage, process awareness, and efficiency
- The need to develop new playbooks that match evolving APTs, such as for file shares and C2
- Identifying and plugging gaps in detection and monitoring
- Practicing identification and eviction of threats
- Having rehearsed crisis management within a real-world scenario

The results of the offensive exercise go much further than finding an end-to-end attack path, by also demonstrating the benefit of learnings from past collaborative activities, and the success of defensive measures in hindering the threat as it advanced.

REFERENCES

^[1] **Hunting for C3,**

<https://labs.f-secure.com/blog/hunting-for-c3/>

^[2] **Web Service,**

<https://attack.mitre.org/techniques/T1102/>

^[3] **Credential Dumping,**

<https://attack.mitre.org/techniques/T1003/>



F-Secure®

www.f-secure.com/consulting

© F-Secure Consulting 2020