Abnormal

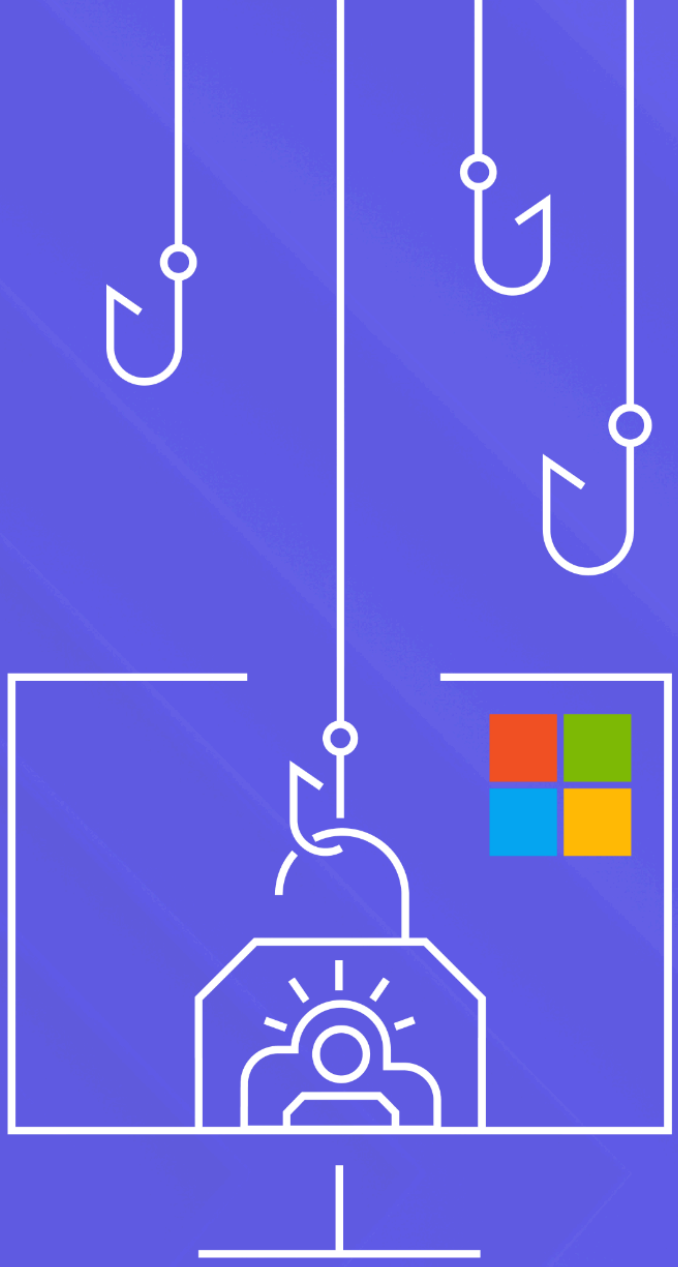

THREAT INTELLIGENCE REPORT

Targeting Microsoft ADFS

How Phishing Campaigns Bypass Multi-Factor Authentication to Enable Account Takeover

/ February 2025

Executive Summary

An ongoing phishing campaign is targeting organizations that rely on federated authentication systems, using spoofed Microsoft Active Directory Federation Services (ADFS) login pages to harvest credentials and bypass multi-factor authentication. ADFS is Microsoft's single sign-on solution, which allows users to authenticate across multiple applications and systems with a single set of credentials, providing secure access to both on-premises and cloud-based services.

What: Attack Progression

In this campaign, attackers exploit the trusted environment and familiar design of ADFS sign-in pages to trick users into submitting their credentials and second-factor authentication details. The success of these attacks is driven by highly convincing phishing techniques, including:

- **Spoofed sender addresses**: Emails appear as if they originate from trusted entities.
- **Legitimate branding**: Fraudulent login pages mimic the organization's official portal.
- **URL obfuscation**: Malicious links are crafted to mimic legitimate ADFS link structure.

Attackers further personalize phishing pages to match the organization's specific MFA setup. For example, targets using push notifications for second-factor authentication receive tailored instructions, such as approving an expected notification, which increases the likelihood of success.

This approach leverages nuanced psychological tactics to exploit human vulnerabilities and reinforce a false sense of legitimacy. By compromising ADFS accounts, attackers can gain access to the target's systems and data, enabling lateral phishing and financially-motivated attacks.

So What: Impact Analysis

TARGET SCOPE

150+

Organizations worldwide targeted by ADFS authentication attacks.

SECTORS

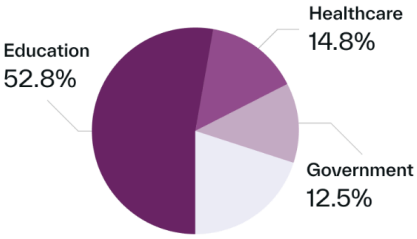

THREAT IMPACT

Exploiting ADFS through social engineering allows attackers to bypass MFA, leading to account takeovers, lateral phishing, and business email compromise attacks.

This campaign highlights how threat actors use social engineering to bypass MFA and exploit organizations—showcasing how important it is to migrate to modern systems, deploy AI-powered defenses, and update security awareness training to inform users of these threats.

Now What: Critical Actions for CISOs

Migrate from ADFS

Transition to Microsoft Entra for modern authentication

Deploy AI Defense

Implement behavioral AI to detect identity-based threats

Enhance Training

Update security awareness for AI-powered threats

Attack Overview

This phishing campaign leverages a combination of social engineering and technical exploitation to bypass security measures. The attack progresses through several stages, starting with the delivery of a phishing email that directs users to a spoofed ADFS login page, ultimately leading to account takeovers and the potential for lateral phishing and financially motivated attacks.

1. Phishing Email

The attack begins with an email, designed to appear as a notification from the organization's IT help desk. The message informs the recipient of an urgent/important update that requires their immediate attention and asks them to use the provided link to initiate the requested action: accept the revised policy, upgrade their system, etc. The theme of the emails varies, but the general motif is consistent across the campaign.

In the examples below, the messages claim the organization has recently updated its email system for enhanced security.

Here, the target is asked to accept the latest policy updates using the provided link:

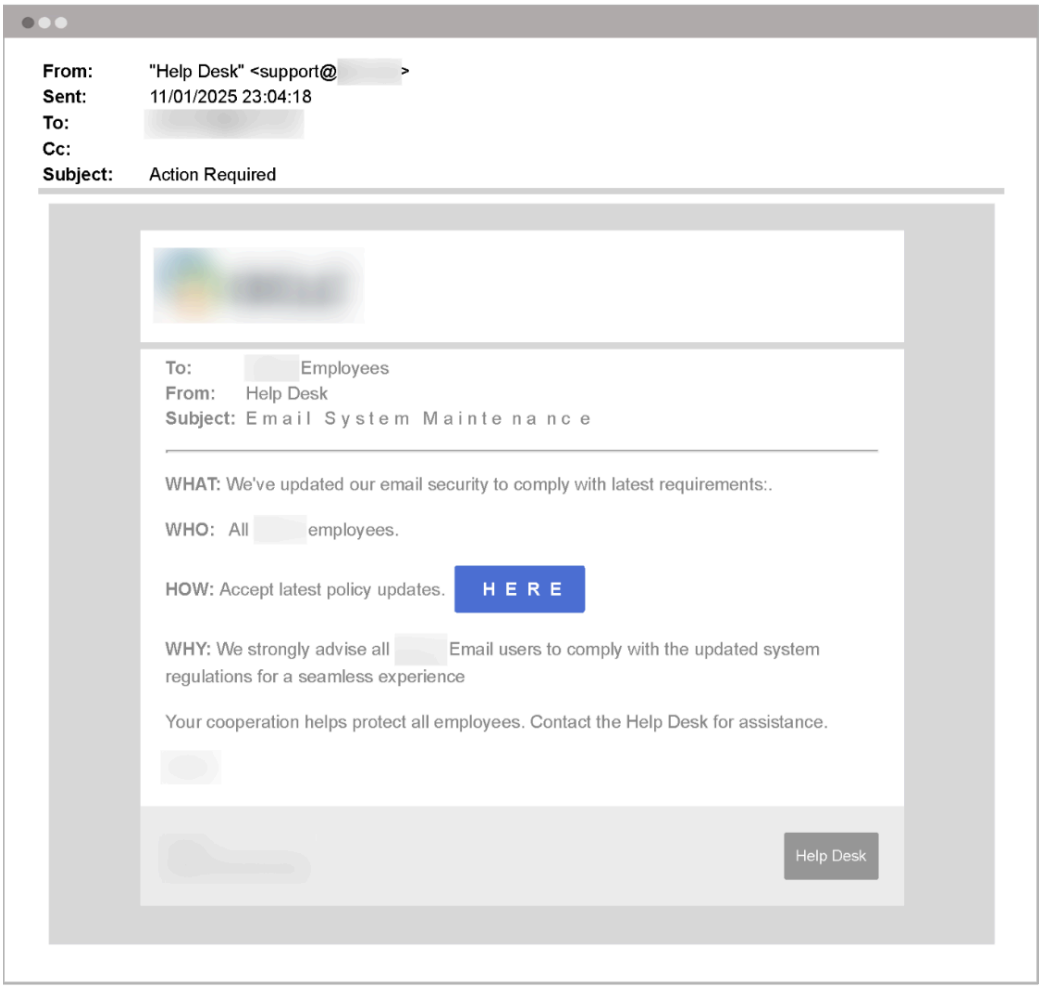

This email, sent to a different organization, requests that the target initiate an update to their email system via the embedded link:

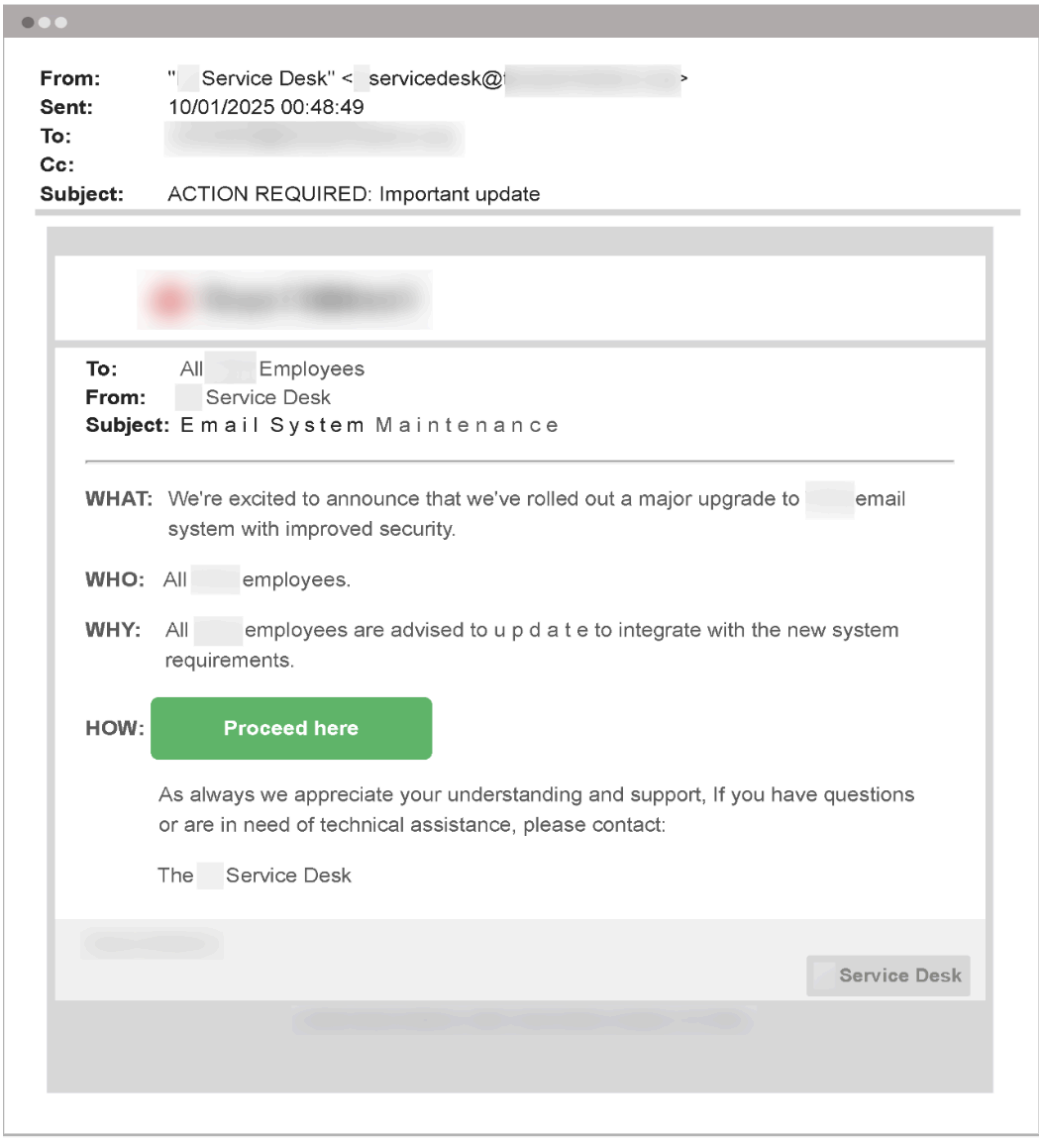

In every email, the embedded link(s) redirects the target to a phishing website mimicking the targeted organization's ADFS authentication page. The links include the following structure:

`http://subdomain.basedomain.com/targeted_organization.com/ls/adfs/ls/client-request-id=7c724&wa=wsignin10.html`

To increase the appearance of legitimacy, the attackers use spoofed sender addresses that incorporate the organization's name. They also include the targeted organization's logo in the body of the email, along with additional contact information, such as the company's real website or physical address.

In some of the emails, the attackers utilize a URL shortener to obfuscate the actual link destination. This can increase the likelihood of passing link verification checks and, as a result, decrease the chances of it being flagged as malicious.

Interestingly, while the emails are certainly designed to manufacture a sense of urgency, they don't fabricate as dire a situation as some attacks do. A common theme in malicious emails is claiming that if the target doesn't act immediately, they will lose access to something important within a matter of hours—their bank portal, Apple or Spotify account, or email account itself.

In this campaign, however, the attackers opt to use language that, though authoritative, does not imply that failure to perform the requested action promptly will have immediate, serious consequences. This tactic could be enough to not raise red flags in targets who know to be wary of emails with requests that prompt immediate action upon receipt.

While the attackers avoid an overly urgent tone, the goal remains the same: to entice the recipient into taking action. Should the target click on the embedded button, they are redirected to the phishing page.

2. Phishing Landing Page

The phishing landing page operates in three stages: credential collection, one-time password (OTP) collection, and confirmation. Each stage is designed to sequentially gather user credentials and OTPs, manipulate targets into approving MFA requests, and ultimately redirect them to the legitimate site to avoid suspicion.

Credential Collection

When users click on the link embedded in the phishing email, they are redirected to a fake ADFS login page mimicking the targeted organization's authentication portal. This phishing page features the targeted organization's logo, which is dynamically pulled from the legitimate website. It also incorporates specific branding elements, including relevant colors and imagery, that are consistent with Microsoft ADFS interfaces to enhance trust and credibility.

Below are screenshots of two spoofed ADFS portals used as part of the campaign.

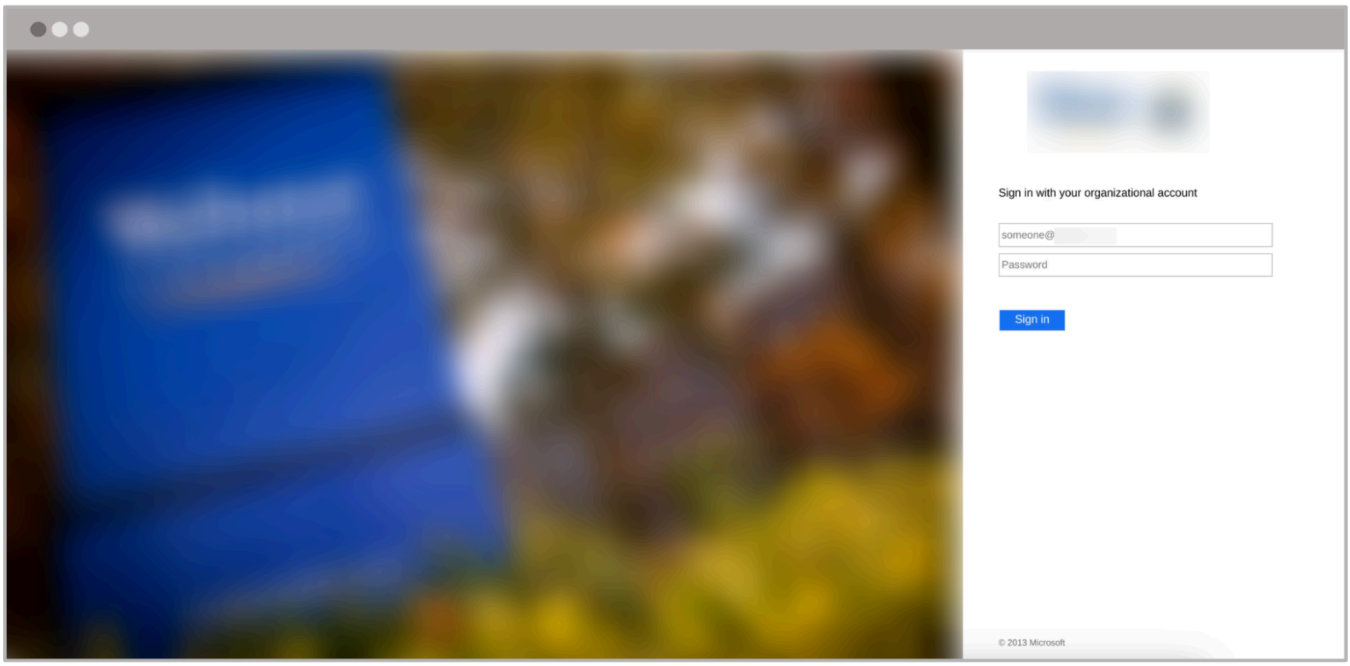

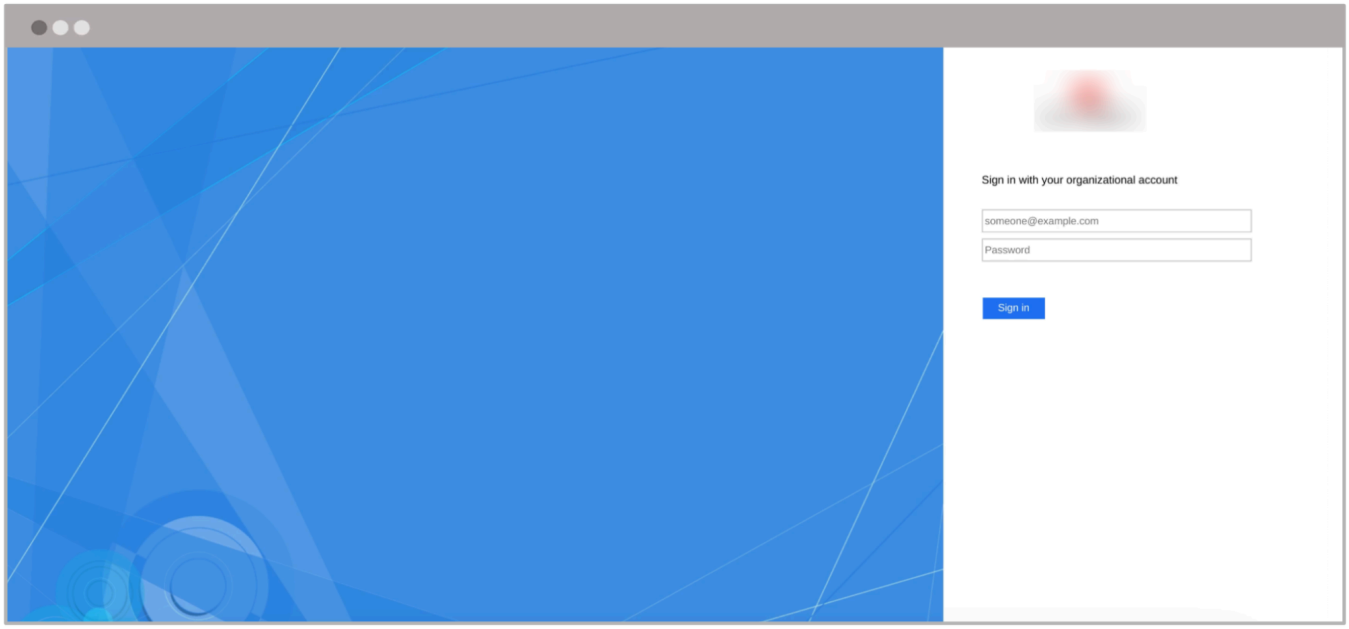

For reference, below is the example ADFS login portal from Microsoft's official learning platform, illustrating that the spoofed versions are indistinguishable from the real sign-in pages.

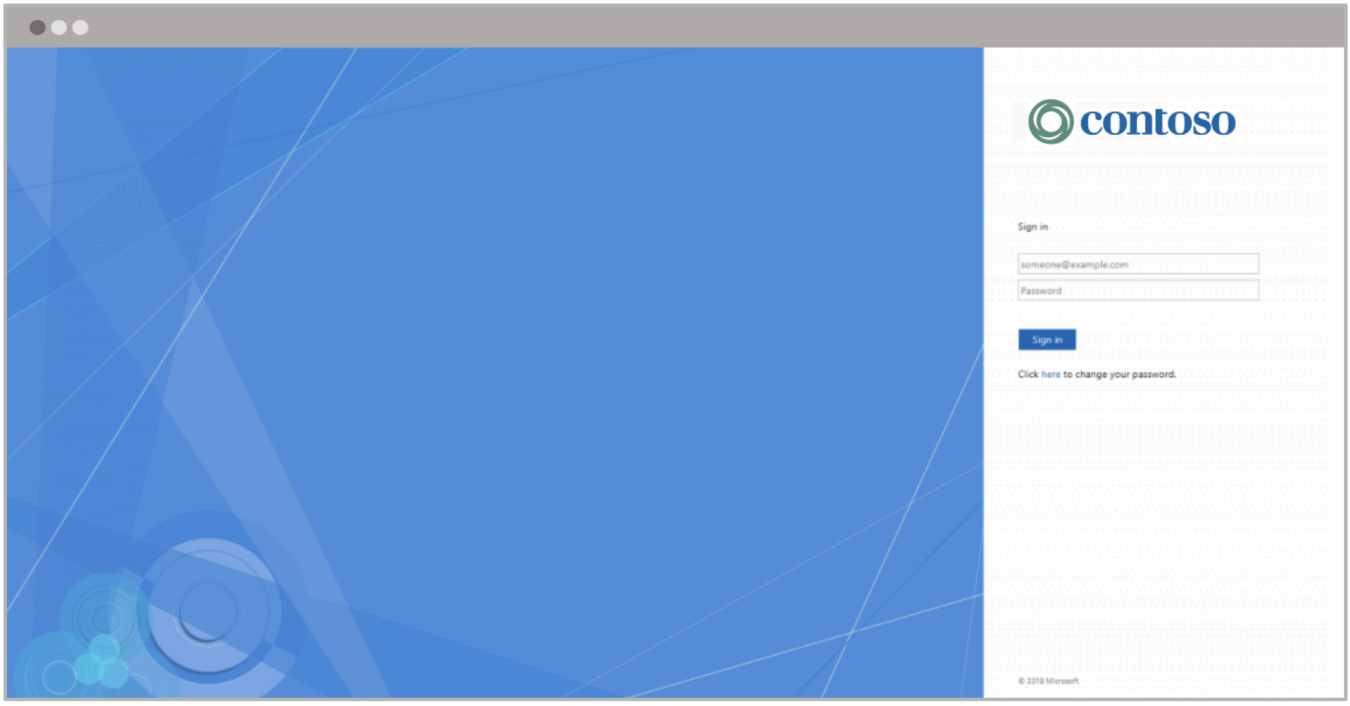

The interface on the spoofed ADFS sign-in page displays input fields for a username and password, giving users the impression they are logging into their organizational account.

While the page includes a simple client-side validation code to check the username format and ensure the password meets basic criteria—such as being non-empty and within a certain length—it does not verify the credentials against the targeted organization's systems. Thus, the form accepts invalid usernames and incorrect username-password combinations.

The code behind the form is displayed below:

```python
<form method="post" id="loginForm" action="radio.php">

    <div id="userNameArea">

        <input id="userNameInput" name="UserName" type="email"

               placeholder="someone@target_org.com" autocomplete="off">

    </div>

    <div id="passwordArea">

        <input id="passwordInput" name="Password" type="password"

               placeholder="Password" autocomplete="off">

    </div>

    <div id="submissionArea">

        <span id="submitButton" class="submit" tabindex="4"

              onclick="document.getElementById('loginForm').submit();">

            Sign in

        </span>

    </div>

                          <input      id="rl"      type="hidden"      name="rl"
value="http://outlook.com/target_org.com">

</form>
```

The login form includes a hidden field specifying the redirect location for post-authentication, directing users to the legitimate site (e.g., `http://outlook.com/target_org.com`) after completing the phishing workflow.

Second Factor Authentication

The phishing templates also include forms designed to capture the specific second factor required to authenticate the target's account, based on the organization's configured MFA settings. Abnormal observed templates targeting multiple commonly used MFA mechanisms, including Microsoft Authenticator, Duo Security, and SMS verification, examples of which are below.

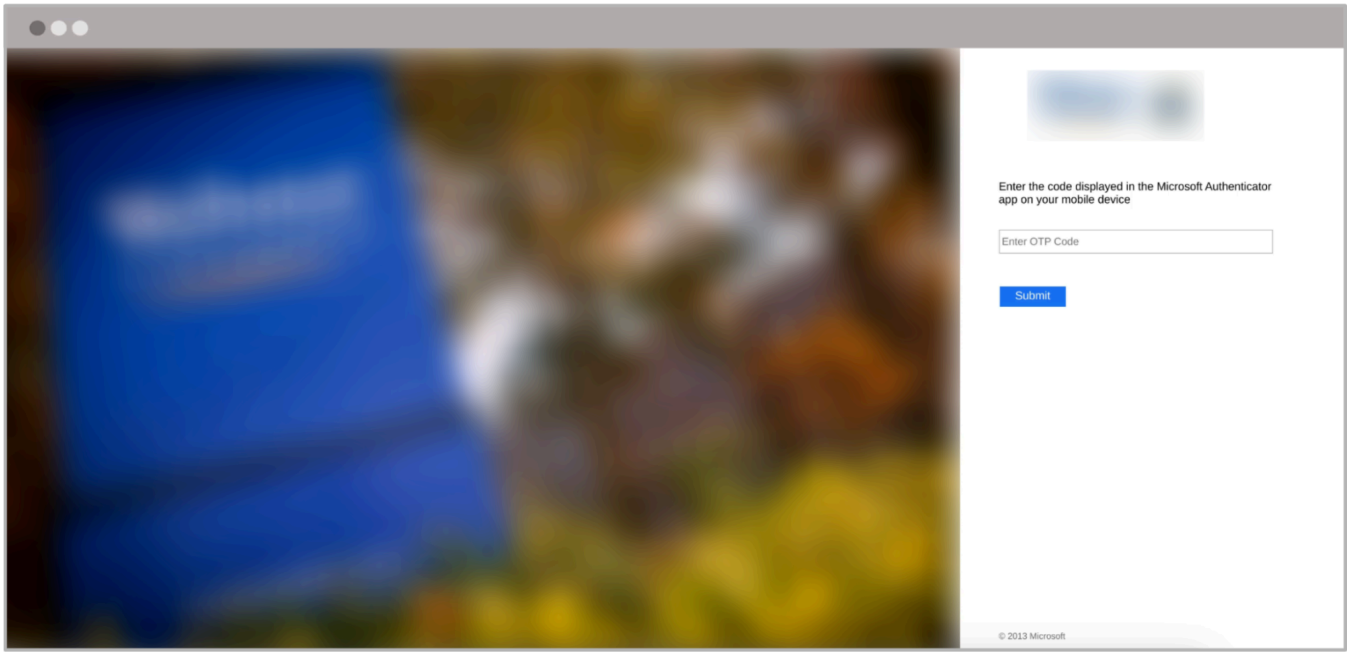

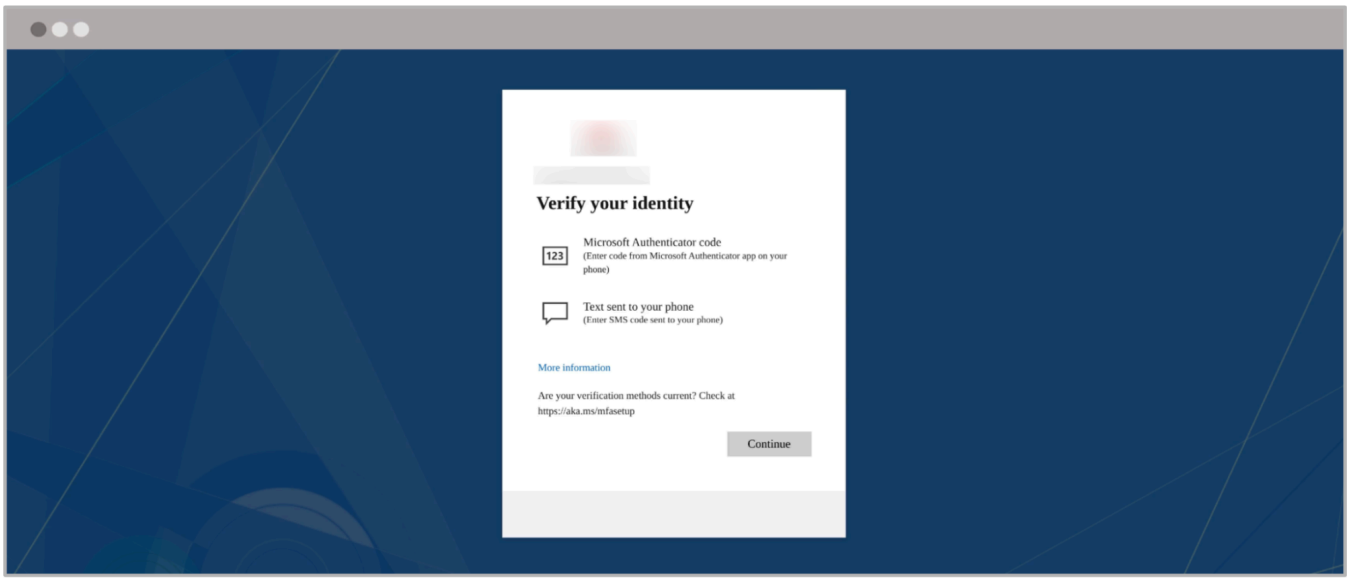

Similar to the login form that captures credentials, these forms do not validate the provided information. Instead, the collected MFA details, along with the previously acquired credentials, are submitted directly to the attacker's system.

Confirmation and Redirection

The penultimate step of the phishing workflow involves a message to the target informing them they may need to approve a push notification or respond to an automated call, as shown in the examples below:

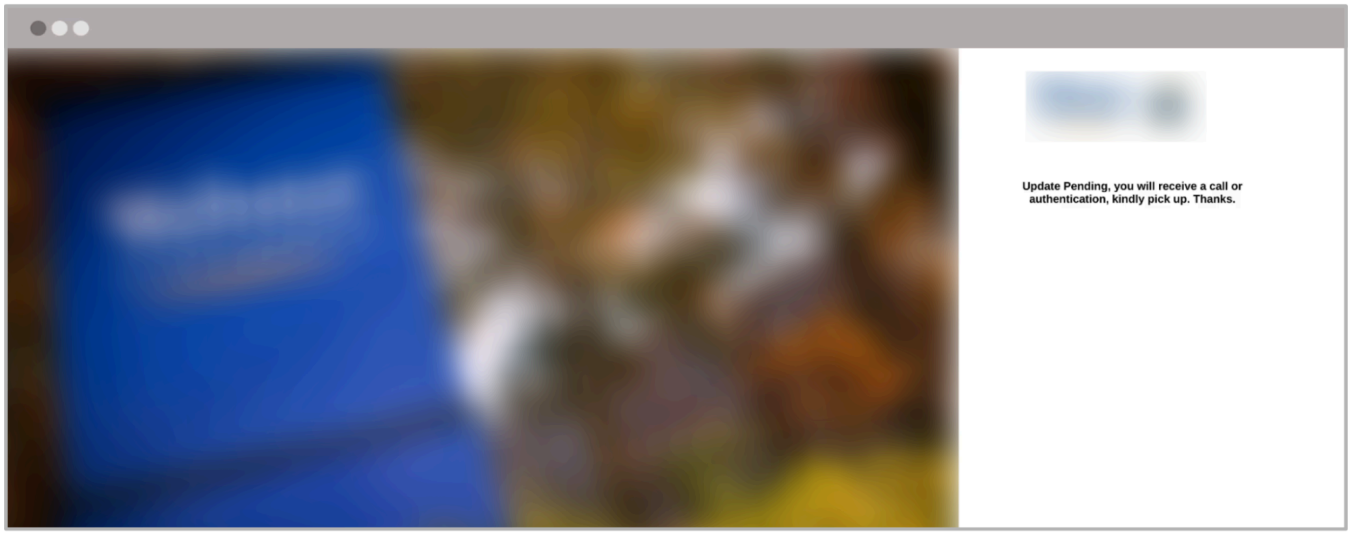

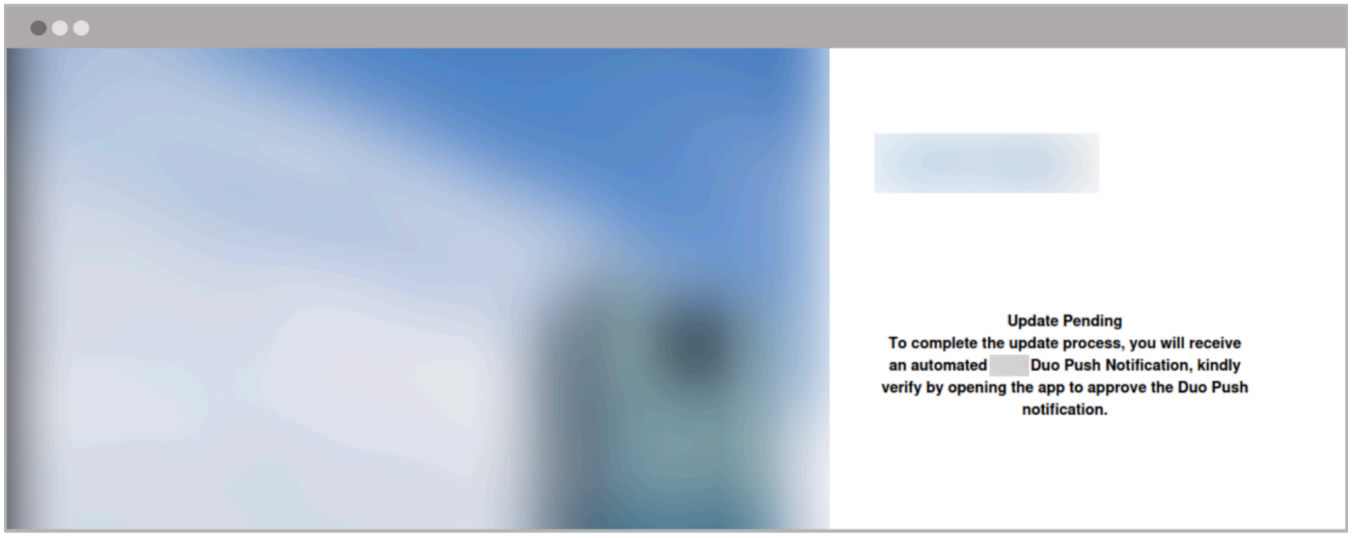

In the final step, the target is redirected to the legitimate sign-in page. This redirection serves multiple purposes:

- It reduces suspicion by reinforcing the illusion of legitimacy, creating the impression that the authentication process was successful.
- It ensures the target believes they have completed the requested update.

By relying on social engineering to manipulate the target into providing credentials, OTP, and approving the MFA request, the attacker is able to bypass the second-factor authentication and complete the account takeover.

3. Account Takeover (ATO)

If a target interacts with the phishing email, provides their credentials, and submits the second-factor authentication details, the threat actor proceeds with the next stage of the attack: account takeover.

Abnormal observed a pattern for the ATO cases, which included an initial failed login attempt, followed by a successful login once the MFA requirement was satisfied, illustrated in the activity timeline below. The sign-ins originated from Private Internet Access VPN, enabling attackers to source their connection from a range of IP addresses.

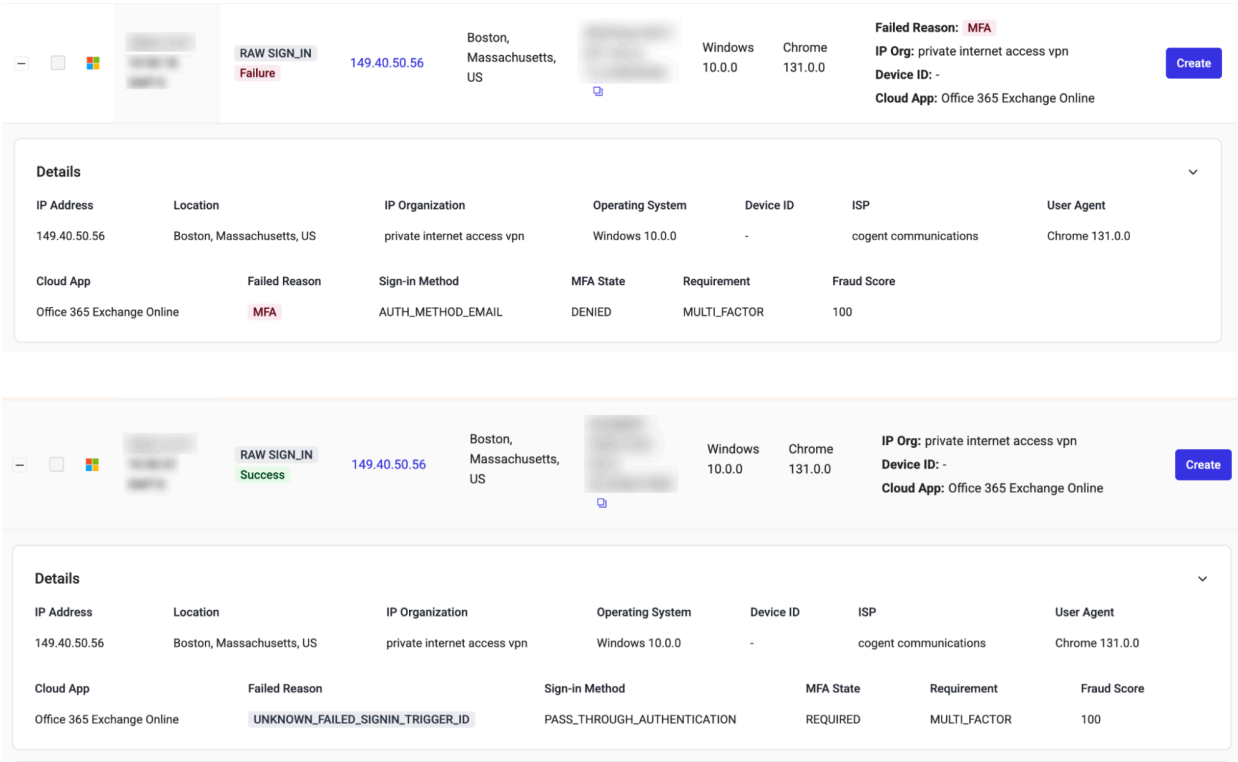

Post-Compromise Activity

The threat actors then leveraged access to the accounts, employing typical techniques for financially motivated email attacks, including reconnaissance, mail filter rule creation, and lateral phishing.

One observed tactic involved creating mail filters with deceptive naming conventions and keyword alterations. For instance, some mail filters were named using common words such as "recommended" to avoid suspicion. Others were crafted dynamically based on the target's name, derived by combining the first three characters of the target's last name with the first two characters of their first name.

To further evade detection, the filter conditions included obfuscated keywords, such as "Hish" for "Phish," "Elpdes" for "Help Desk," and "Otifica" for "Notification." By using non-obvious terms, the threat actors reduced the likelihood of security solutions or analysts identifying the filters as malicious. These tailored techniques ensured that any responses to lateral phishing emails were

intercepted and deleted, preventing the mailbox owner from noticing malicious activity or incoming replies.

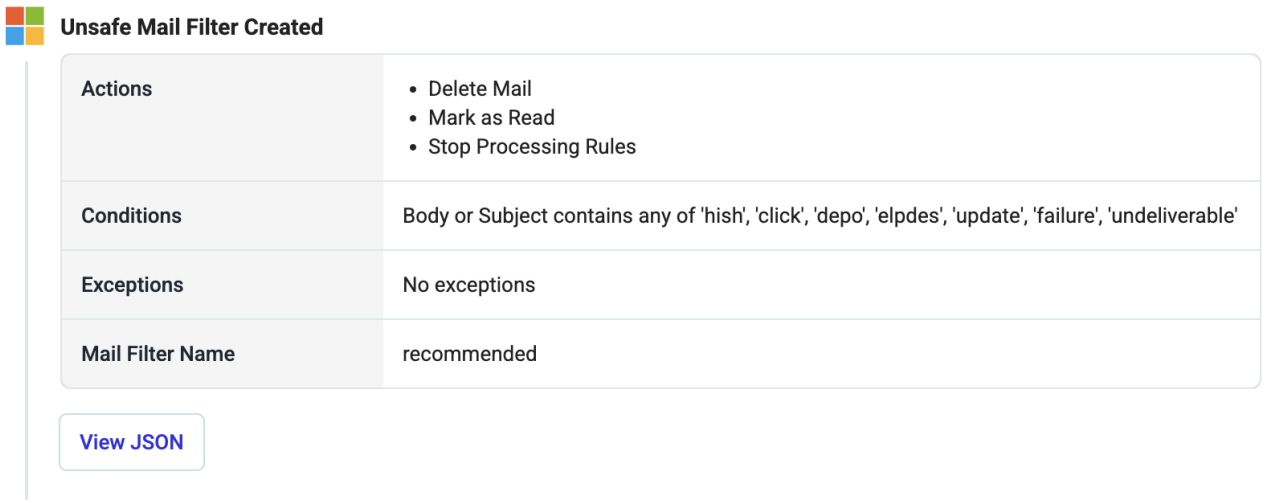

Another notable tactic employed was lateral phishing, where compromised accounts were leveraged to send phishing emails both laterally within the target's organization and externally to associated organizations. The use of compromised accounts enabled the attackers to exploit existing trust relationships, increasing the likelihood of recipients interacting with malicious emails.

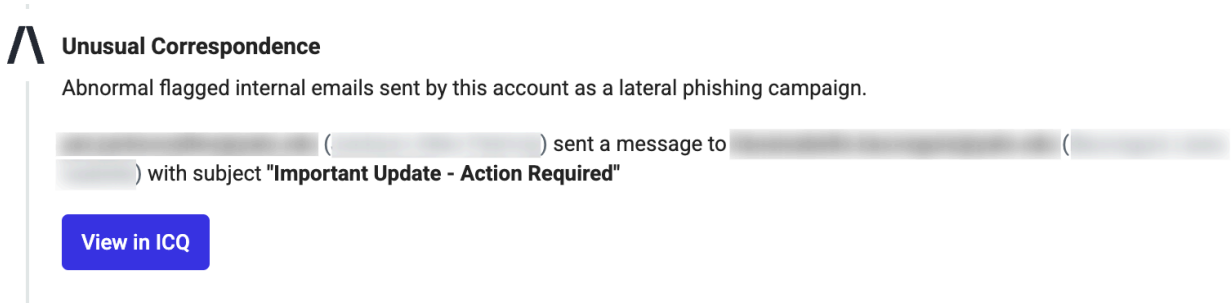

These campaigns were tailored to mimic legitimate communication patterns, often targeting colleagues, partners, or clients of the compromised individual. By leveraging trusted relationships within the organization, the attackers increase the likelihood of successful fraud and further data breaches.

Victimology

Abnormal researchers discovered that this phishing campaign has targeted over 150 organizations across a range of industries and geographic regions, specifically those relying on ADFS for authentication.

While Microsoft recommends organizations transition to its modern identity platform, Entra, attackers are exploiting the fact that many still depend on ADFS for authentication. This reliance is particularly prevalent in sectors with slower technology adoption cycles or legacy infrastructure dependencies—making them prime targets for credential harvesting and account takeovers.

Of the targeted organizations discovered by Abnormal researchers, the education sector has been the most heavily targeted, with over 50% of attacks affecting schools, universities, and other educational institutions. This highlights the attackers' preference for environments with high user volumes, legacy systems, fewer security personnel, and often less mature cybersecurity defenses.

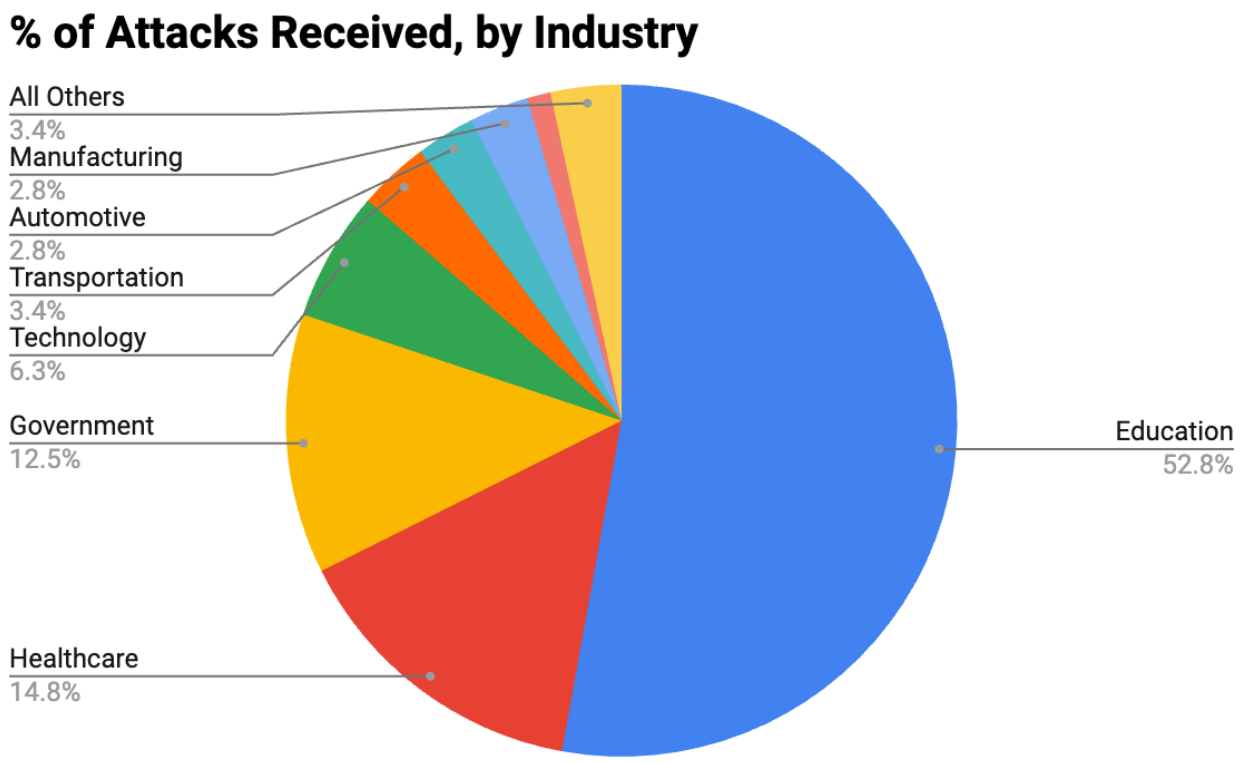

Other affected sectors include healthcare (14.8%), government (12.5%), technology (6.3%), and transportation (3.4%), reflecting a broader intent to exploit the continued reliance on ADFS technology rather than specifically targeting individual organizations or sectors. Geographically, most affected organizations were based in the United States, with notable incidents also found in Canada, Australia, and Europe.

Conclusion

This campaign demonstrates the deceptive tactics used by threat actors to bypass traditional security measures and highlights the risks associated with legacy authentication systems like ADFS.

It serves as a stark reminder that even advanced security measures like MFA can be bypassed when combined with human factors. Advanced phishing techniques and social engineering enable attackers to exploit these vulnerabilities, underscoring the critical need for organizations to adopt a multi-layered defense strategy:

1. **Transition to Modern Platforms**: Organizations should prioritize migrating to modern identity solutions, such as Microsoft Entra, to reduce reliance on legacy systems and strengthen defenses against credential-based attacks.
2. **Enhance User Awareness**: Educating users about phishing techniques and psychological tactics used by attackers is vital to reducing human vulnerabilities.
3. **Implement Robust Technical Defenses**: Deploy advanced email filtering, anomaly detection, and behavior monitoring technologies to identify and mitigate phishing attacks and detect compromised accounts early.

It's important to note that the attackers' focus on ADFS is less about targeting specific organizations and more about exploiting sectors that continue to rely on this technology. Organizations that are slower to transition to modern platforms, such as Microsoft Entra, are particularly susceptible to these attacks. However, by addressing these factors proactively, security leaders can reduce their exposure to deceptive phishing campaigns and better protect critical resources.

IOCs

The following Indicators of Compromise (IOCs) were identified during the investigation of this phishing campaign. These IOCs can help organizations detect and mitigate the attack by identifying phishing emails, malicious links, and other signs of compromise within their systems.

Email Subjects:

- Important Update <user@organization.com>
- Action Required
- Important Update
- Action Required: Important update
- Must Read
- Frontline Direct Deposit

Phishing Link URL Patterns:

- /&adfs/ls/client-request-id=7c724
- /adfs/ls/client-request-id=7c724

Mail Filters:

- Body or subject contains:
  - "elpdes"
  - "hish"
  - "click"
  - "depo"
  - "alert"
- Sender contains:
  - “otifica”
  - “elpdes”
  - “orkda”
  - “ayrol”
  - “norepl”

IP Addresses:

- 149.40.50.56

- 149.40.50.46
- 149.40.50.27
- 149.40.50.34
- 149.40.50.15
- 149.40.50.11
- 149.40.50.36
- 84.239.45.1
- 84.239.45.12
- 84.239.45.8