Detecting and mitigating a multi-stage AiTM phishing and BEC campaign

microsoft.com/en-us/security/blog/2023/06/08/detecting-and-mitigating-a-multi-stage-aitm-phishing-and-bec-campaign/

June 8, 2023

ResearchThreat intelligenceMicrosoft DefenderBusiness email compromise12 min read

By

Microsoft Defender Experts uncovered a multi-stage adversary-in-the-middle (AiTM) phishing and business email compromise (BEC) attack against banking and financial services organizations. The attack originated from a compromised trusted vendor and transitioned into a series of AiTM attacks and follow-on BEC activity spanning multiple organizations. This attack shows the complexity of AiTM and BEC threats, which abuse trusted relationships between vendors, suppliers, and other partner organizations with the intent of financial fraud.

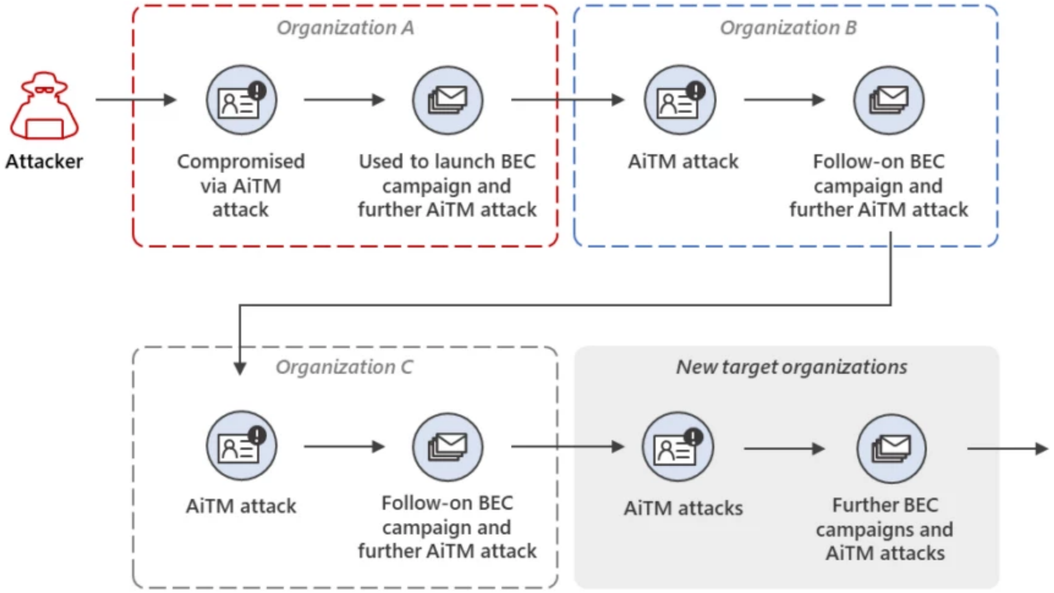
Attacker
Organization A
Compromised via AiTM attack
Used to launch BEC campaign and further AiTM attack
Organization B
AiTM attack
Follow-on BEC campaign and further AiTM attack
Organization C
AiTM attack
Follow-on BEC campaign and further AiTM attack
New target organizations
AiTM attacks
Further BEC campaigns and AiTM attacks

Figure 1. AiTM and BEC attacks spanning multiple suppliers and partner organizations

While the attack achieved the end goal of a typical AiTM phishing attack followed by business email compromise, notable aspects, such as the use of indirect proxy rather than the typical reverse proxy techniques, exemplify the continuous evolution of these threats. The use of indirect proxy in this campaign provided attackers control and flexibility in tailoring the phishing pages to their targets and further their goal of session cookie theft. After signing in with the stolen cookie through a session replay attack, the threat actors leveraged multifactor authentication (MFA) policies that have not been configured using security best practices in order to update MFA methods without an MFA challenge. A second-stage phishing campaign followed, with more than 16,000 emails sent to the target's contacts.

This attack highlights the complexity of AiTM attacks and the comprehensive defenses they necessitate. This sophisticated AiTM attack requires beyond the typical remediation measures for identity compromise such as a password reset. Affected organizations need to revoke session cookies and roll back MFA modifications made by the threat actor. The incident also highlights the importance of proactive threat hunting to discover new TTPs on previously known campaigns to surface and remediate these types of threats.

To launch this attack, the attackers used an AiTM phishing kit developed, maintained, and operated by a threat actor that Microsoft tracks as Storm-1167. As part of our threat actor tracking and naming taxonomy, Microsoft uses Storm-#### designations as a temporary name given to an unknown, emerging, or developing cluster of threat activity, allowing Microsoft to track it as a unique set of information until we reach high confidence about the origin or identity of the actor behind the activity.

AiTM with indirect proxy

Adversary-in-the-middle (T1557, T1111) is a type of attack that aims to intercept authentication between users and a legitimate authentication service for the purpose of compromising identities or performing other actions. The attackers position themselves between a user and the service to steal credentials and intercept MFA in order to capture the session cookie. The attackers can then replay the session with the stolen session cookie before the token expiration time and impersonate the user without user intervention or MFA. With this session, the attackers could access the affected user's resources and applications and perform business email compromise attacks and other malicious activities. More details about AiTM campaigns can be found on the blog Attackers use AiTM phishing sites as entry point to further financial fraud.

Unlike campaigns we have previously reported, this attack did not use the reverse proxy method that AiTM kits like EvilProxy and NakedPages use, in which the attacker's server proxies the request from the application's legitimate sign-in page. Instead, the attack used AiTM attack with indirect proxy method, in which the attacker presented targets with a website that mimicked the sign-in page of the targeted application, as in traditional phishing attacks, hosted on a cloud service. The said sign-in page contained resources loaded from an attacker-controlled server, which initiated an authentication session with the authentication provider of the target application using the victim's credentials.

In this AiTM attack with indirect proxy method, since the phishing website is set up by the attackers, they have more control to modify the displayed content according to the scenario. In addition, since the phishing infrastructure is controlled by the attackers, they have the flexibility to create multiple servers to evade detections. Unlike typical AiTM attacks, there are no HTTP packets proxied between the target and the actual website.

When MFA is requested after successful password validation, the server displays a fake MFA page. Once the MFA is provided by the user, the attacker uses the same MFA token in the initiated session with the authentication provider. Following successful authentication, the session token is granted to the attacker, and victim is redirected to another page. The following diagram illustrates the AiTM attack observed in this scenario:

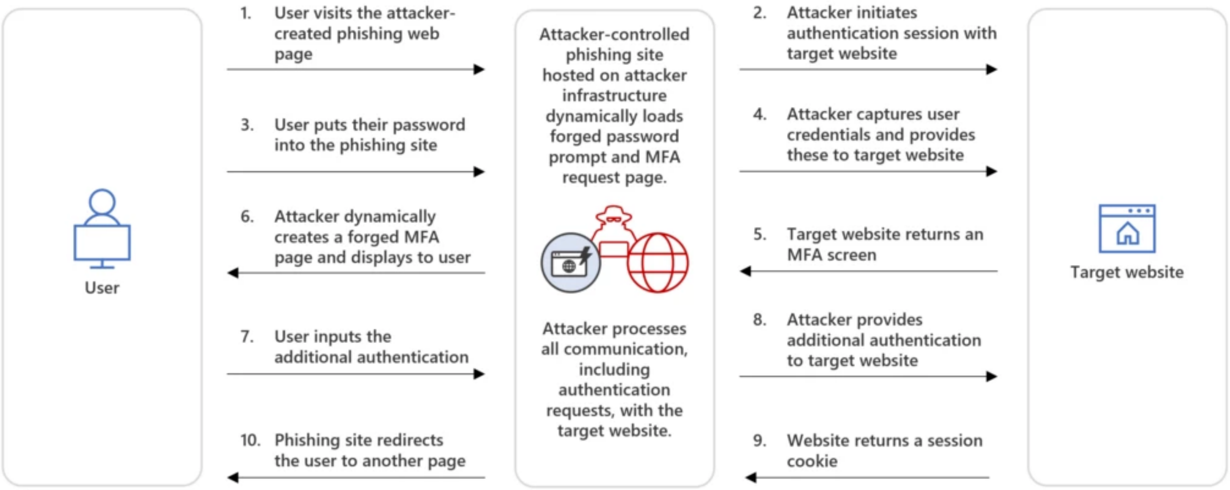
User
1. User visits the attacker-created phishing web page
3. User puts their password into the phishing site
6. Attacker dynamically creates a forged MFA page and displays to user
7. User inputs the additional authentication
10. Phishing site redirects the user to another page
Attacker-controlled phishing site hosted on attacker infrastructure dynamically loads forged password prompt and MFA request page.
Attacker processes all communication, including authentication requests, with the target website.
2. Attacker initiates authentication session with target website
4. Attacker captures user credentials and provides these to target website
5. Target website returns an MFA screen
8. Attacker provides additional authentication to target website
9. Website returns a session cookie
Target website

Figure 2. AiTM with indirect proxy

Attack chain: AiTM phishing attack leads to second-stage BEC

Our investigation into an AiTM phishing attack using the Storm-1167 AiTM kit uncovered details of a campaign that led to BEC activity. In the following sections, we present our in-depth analysis of the end-to-end attack chain.

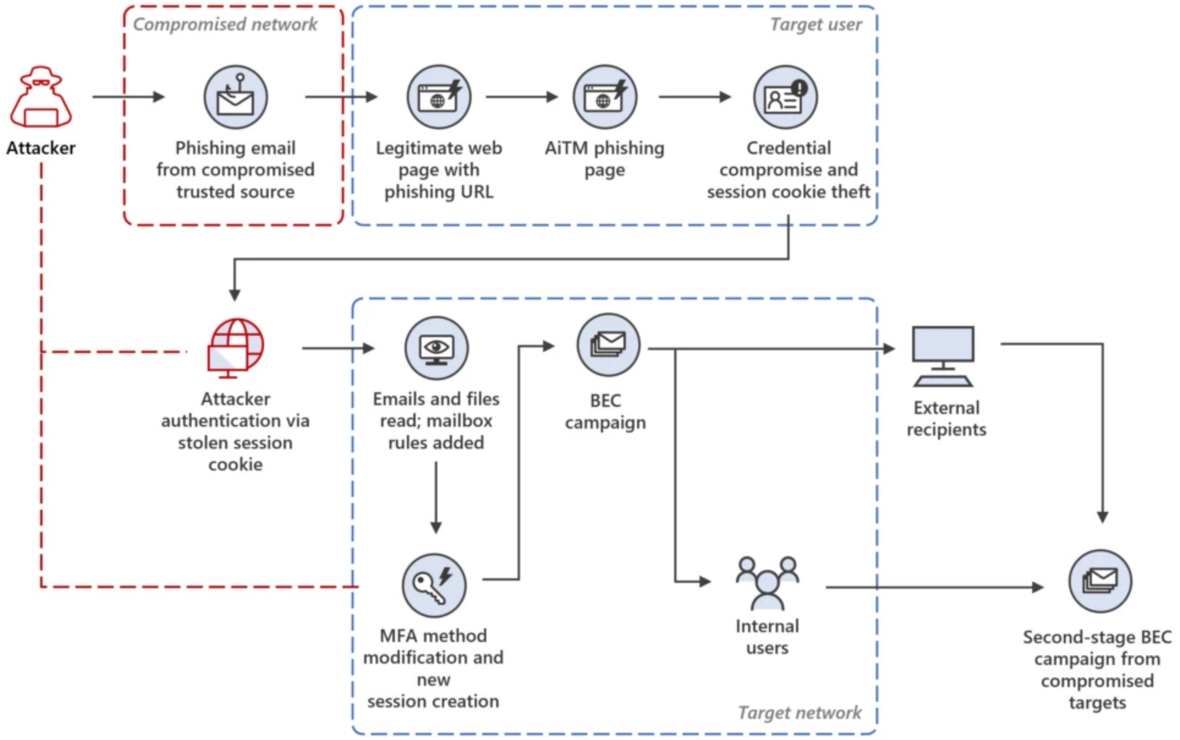
Compromised network
Target user
Attacker
Phishing email from compromised trusted source
Legitimate web page with phishing URL
AiTM phishing page
Credential compromise and session cookie theft
Attacker authentication via stolen session cookie
Emails and files read; mailbox rules added
BEC campaign
External recipients
MFA method modification and new session creation
Internal users
Second-stage BEC campaign from compromised targets
Target network

Figure 3. Attack chain from AiTM phishing attack to BEC

Stage 1: Initial access via trusted vendor compromise

The attack started with a phishing email from one of the target organizations' trusted vendors. The phishing email was sent with a seven-digit code as the subject. This code was unique for every target organization, which is likely a tracking mechanism for the attacker. The email body included a link to view or download a fax document. The link pointed to a malicious URL hosted on *canva[.]com*.

Sending phishing emails from a trusted vendor was one of the common behaviors that was observed for this threat actor across multiple targeted organizations. The intent of this behavior is to abuse the trusted vendor relationship and to blend with legitimate email traffic. A few of the target organizations had policies that automatically allow emails from trusted vendors, enabling the attacker to slip past detections.

Stage 2: Malicious URL click

Threat actors often abuse legitimate services and brands to avoid detection. In this scenario, we observed that the attacker leveraged the legitimate service Canva for the phishing campaign. Canva is a graphic design platform that allows users to create social media graphics, presentations, posters, and other visual content. Attackers abused the Canva platform to host a page that shows a fake OneDrive document preview and links to a phishing URL:

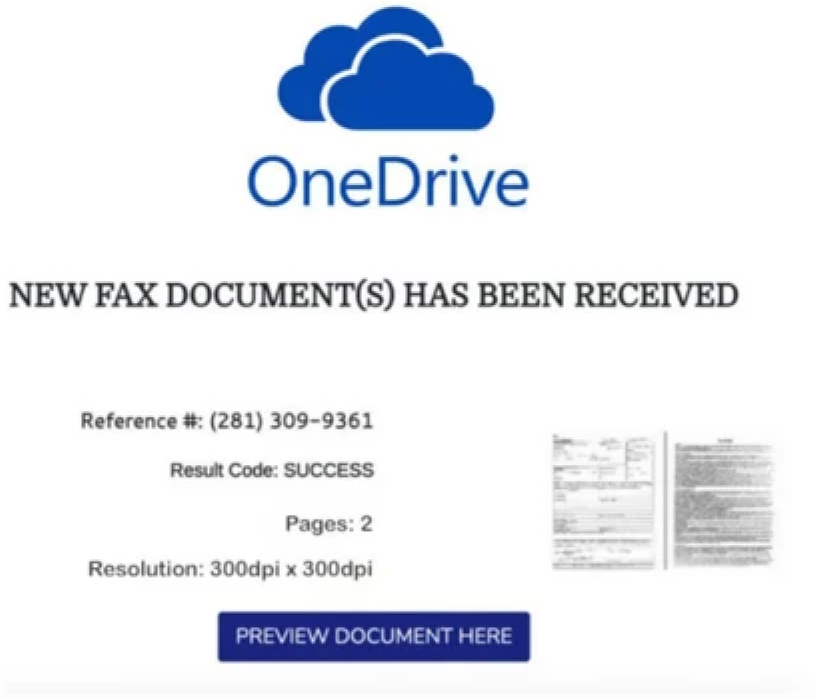

OneDrive

NEW FAX DOCUMENT(S) HAS BEEN RECEIVED

Reference #: (281) 309-9361

Result Code: SUCCESS

Pages: 2

Resolution: 300dpi x 300dpi

Figure 4. Screenshot of the intermediary page leading to AiTM landing page

Stage 3: AiTM attack

Accessing the URL redirected the user to a phishing page hosted on the Tencent cloud platform that spoofed a Microsoft sign-in page. The final URL was different for every user but showed the same spoofed sign-in page.

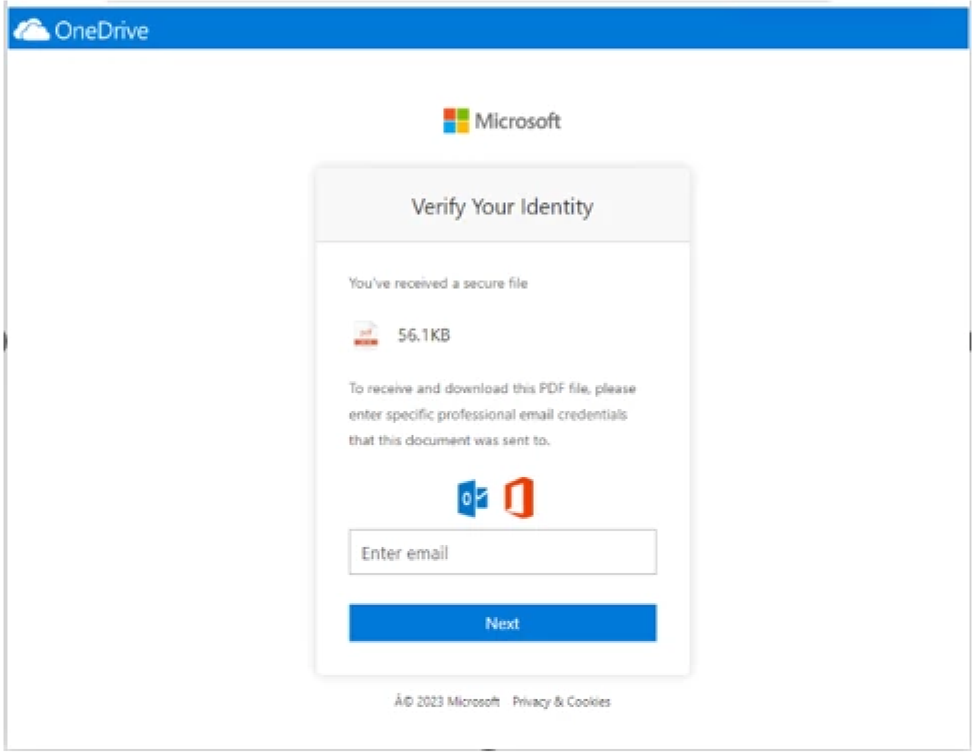

Verify Your Identity

You've received a secure file

56.1KB

To receive and download this PDF file, please enter specific professional email credentials that this document was sent to.

Enter email

Â© 2023 Microsoft Privacy & Cookies

Figure 5. Fake Microsoft sign-in page requesting the target's password

After the target provided the password on the phishing page, the attacker then used the credentials in an authentication session created on the target website. When the attacker is prompted with MFA in the authentication session, the attacker modified the phishing page into a forged MFA page (as seen below). Once the target completed the multifactor authentication, the session token was then captured by the attacker.

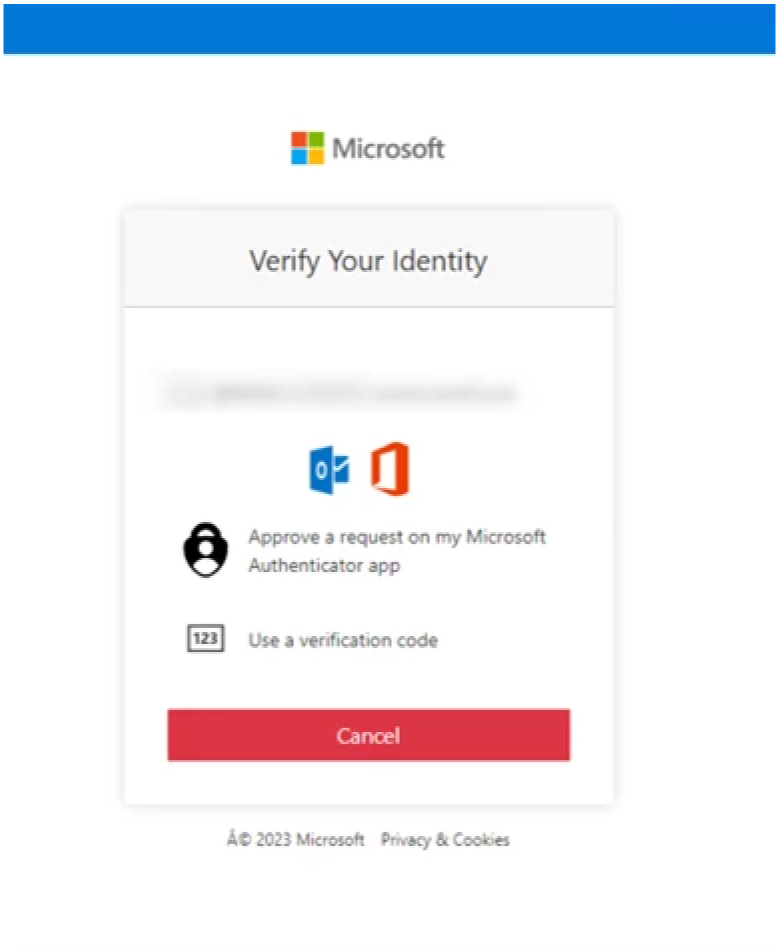
Microsoft

Verify Your Identity

[redacted]

Approve a request on my Microsoft Authenticator app

123 Use a verification code

Â© 2023 Microsoft Privacy & Cookies

Figure 6. Fake Microsoft MFA page requesting a verification code

The phishing pages for the AiTM attack were hosted on IP addresses located in Indonesia. The follow-on sign-ins described in the following sections were also observed from the same IP addresses.

Stage 4: Session cookie replay

In a stolen session cookie replay attack, the attacker uses the valid stolen cookie to impersonate the user, circumventing authentication mechanisms of passwords and MFA. In this campaign, we observed that the attacker signed in with the stolen cookie after a few hours from an IP address based in the United States. The attacker masqueraded as the target with this session replay attack and accessed email conversations and documents hosted in the cloud. In addition, the attacker generated a new access token, allowing them to persist longer in the environment.

Stage 5: MFA method modification

The attacker then proceeded to add a new MFA method for the target's account, which was through phone based one-time password (OTP), in order to sign in using the user's stolen credentials undetected. Adding a new MFA method, by default, does not require re-authentication. In this campaign, a common behavior that was observed was the attacker adding OneWaySMS, a phone-based OTP service, as a new MFA method in addition to the existing method used by the target. A phone number with the Iranian country code was observed added as the number used to receive the phone-based OTP.

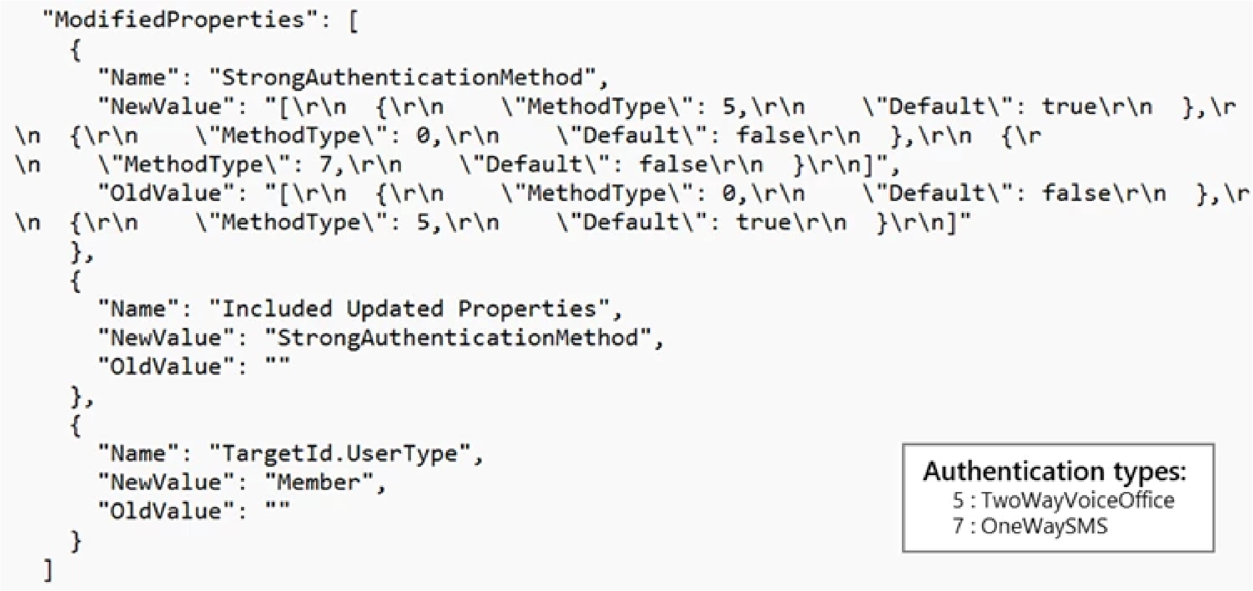

```
  "ModifiedProperties": [
    {
      "Name": "StrongAuthenticationMethod",
      "NewValue": "[\r\n  {\r\n    \"MethodType\": 5,\r\n    \"Default\": true\r\n  },\r
\n  {\r\n    \"MethodType\": 0,\r\n    \"Default\": false\r\n  },\r\n  {\r
\n    \"MethodType\": 7,\r\n    \"Default\": false\r\n  }\r\n]",
      "OldValue": "[\r\n  {\r\n    \"MethodType\": 0,\r\n    \"Default\": false\r\n  },\r
\n  {\r\n    \"MethodType\": 5,\r\n    \"Default\": true\r\n  }\r\n]"
    },
    {
      "Name": "Included Updated Properties",
      "NewValue": "StrongAuthenticationMethod",
      "OldValue": ""
    },
    {
      "Name": "TargetId.UserType",
      "NewValue": "Member",
      "OldValue": ""
    }
  ]
```

Authentication types:
5 : TwoWayVoiceOffice
7 : OneWaySMS

Figure 7. MFA configuration change from cloud application activity logs

Stage 6: Inbox rule creation

The attacker later signed in with the new session token and created an Inbox rule with parameters that moved all incoming emails on the user’s mailbox to the Archive folder and marked all the emails as read.

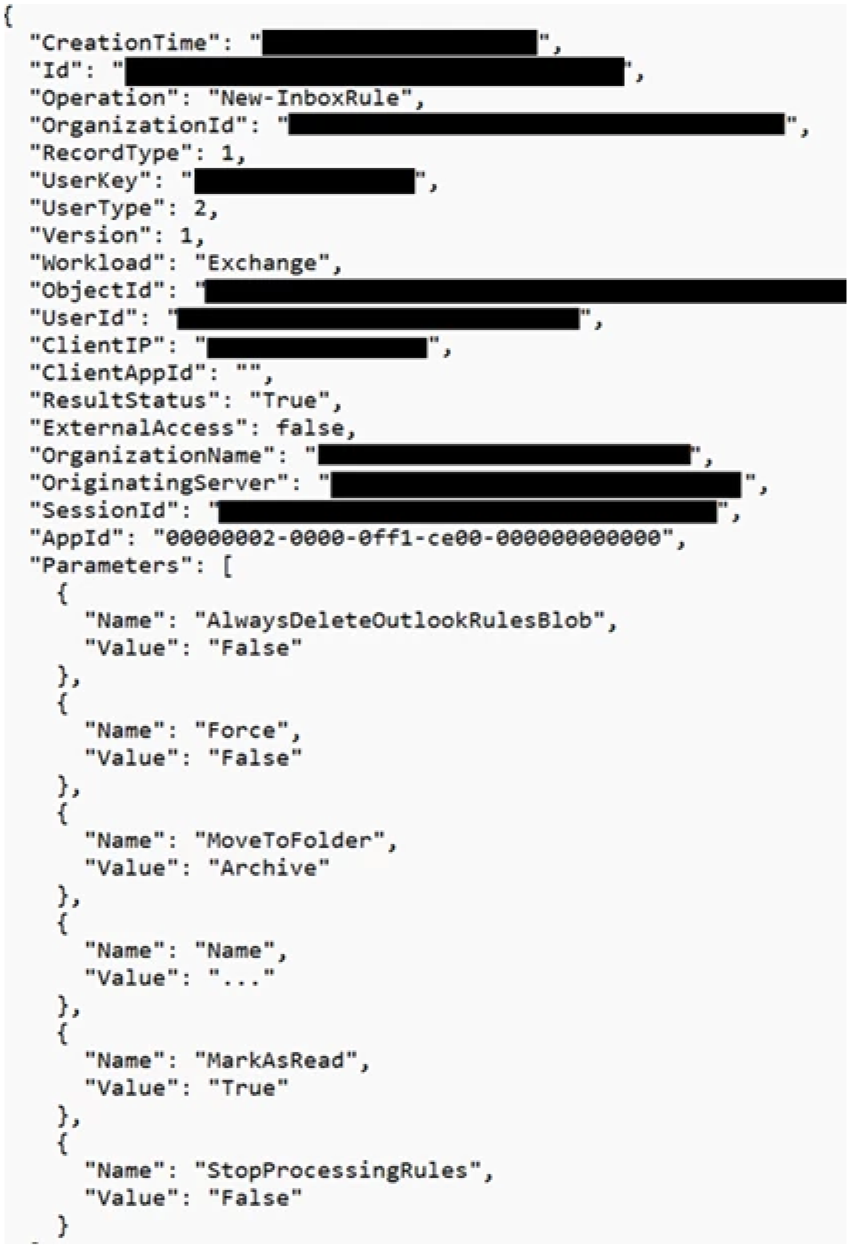

```
{
  "CreationTime": "[redacted]",
  "Id": "[redacted]",
  "Operation": "New-InboxRule",
  "OrganizationId": "[redacted]",
  "RecordType": 1,
  "UserKey": "[redacted]",
  "UserType": 2,
  "Version": 1,
  "Workload": "Exchange",
  "ObjectId": "[redacted]
  "UserId": "[redacted]",
  "ClientIP": "[redacted]",
  "ClientAppId": "",
  "ResultStatus": "True",
  "ExternalAccess": false,
  "OrganizationName": "[redacted]",
  "OriginatingServer": "[redacted]",
  "SessionId": "[redacted]",
  "AppId": "00000002-0000-0ff1-ce00-000000000000",
  "Parameters": [
    {
      "Name": "AlwaysDeleteOutlookRulesBlob",
      "Value": "False"
    },
    {
      "Name": "Force",
      "Value": "False"
    },
    {
      "Name": "MoveToFolder",
      "Value": "Archive"
    },
    {
      "Name": "Name",
      "Value": "..."
    },
    {
      "Name": "MarkAsRead",
      "Value": "True"
    },
    {
      "Name": "StopProcessingRules",
      "Value": "False"
    }
```

Figure 8. Inbox rule creation

Stage 7: Phishing campaign

Followed by Inbox rule creation, the attacker initiated a large-scale phishing campaign involving more than 16,000 emails with a slightly modified Canva URL. The emails were sent to the compromised user's contacts, both within and outside of the organization, as well as distribution lists. The recipients were identified based on the recent email threads in the compromised user's inbox. The subject of the emails contained a unique seven-digit code, possibly a tactic by the attacker to keep track of the organizations and email chains.

Stage 8: BEC tactics

The attacker then monitored the victim user's mailbox for undelivered and out of office emails and deleted them from the Archive folder. The attacker read the emails from the recipients who raised questions regarding the authenticity of the phishing email and responded, possibly to falsely confirm that the email is legitimate. The emails and responses were then deleted from the mailbox. These techniques are common in any BEC attacks and are intended to keep the victim unaware of the attacker's operations, thus helping in persistence.

Stage 9: Accounts compromise

The recipients of the phishing emails from within the organization who clicked on the malicious URL were also targeted by another AiTM attack. Microsoft Defender Experts identified all compromised users based on the landing IP and the sign-in IP patterns.

Stage 10: Second-stage BEC

The attacker was observed initiating another phishing campaign from the mailbox of one of the users who was compromised by the second AiTM attack. Microsoft revoked the compromised user's session cookie, intervening with the second-stage attack.

Microsoft Defender Experts: Extending security and threat defense

This AiTM attack's use of indirect proxy is an example of the threat's increasingly complex and evolving TTPs to evade and even challenge conventional solutions and best practices. Proactively hunting for and quickly responding to threats thus becomes an even more important aspect in securing organization networks because it provides an added layer to other security remediations and can help address areas of defense evasion.

Microsoft Defender Experts is part of Microsoft's global network of more than 8,000 security analysts and researchers who, through our managed services like Microsoft Defender Experts for Hunting, help extend organizations' ability to defend their environment, manage security, and even augment SOC teams. Our experts also enrich our vast cross-domain signals and let us deliver coordinated threat defense in our security products and solutions.

In this incident, because our experts actively research for new AiTM and BEC techniques, they were able to create advanced hunting detections for the Defender Experts service. These detections, combined with our experts' own analyses of the anomalous emails and user behavior, enabled them to uncover the attack at its early stages, analyze the entire attack chain, and identify and promptly reach out to affected and targeted customers through Defender Experts Notifications. They then continuously monitored the attack for any additional compromised users or changes in the phishing patterns as it rapidly unfolded into a large-scale campaign.

Defender Experts also initiated rapid response with Microsoft 365 Defender to contain the attack including:

- Automatically disrupting the AiTM attack on behalf of the impacted users based on the signals observed in the campaign
- Initiating zero-hour auto purge (ZAP) in Microsoft Defender for Office 365 to find and take automated actions on the emails that are a part of the phishing campaign

Defender Experts further worked with customers to remediate compromised identities through the following recommendations:

- Revoking session cookies in addition to resetting passwords
- Revoking the MFA setting changes made by the attacker on the compromised user's accounts
- Require re-challenging MFA for MFA updates as the default

Mitigation and protection guidance

Microsoft 365 Defender detects suspicious activities related to AiTM phishing attacks and their follow-on activities, such as session cookie theft and attempts to use the stolen cookie to sign into Exchange Online. To further protect themselves from similar attacks, organizations should also consider complementing MFA with conditional access policies, where sign-in requests are evaluated using additional identity-driven signals like user or group membership, IP location information, and device status, among others.

Mitigating AiTM phishing attacks

The general remediation measure for any identity compromise is to reset the password for the compromised user. However, in AiTM attacks, since the sign-in session is compromised, password reset is not an effective solution. Additionally, even if the compromised user's password is reset and sessions are revoked, the attacker can set up persistence methods to sign-in in a controlled manner by tampering with MFA. For instance, the attacker can add a new MFA policy to sign in with a one-time password (OTP) sent to attacker registered mobile number. With this persistence mechanisms in place, the attacker can have control over the victim's account despite conventional remediation measures.

While AiTM phishing attempts to circumvent MFA, implementation of MFA still remains an essential pillar in identity security and highly effective at stopping a wide variety of threats. MFA is the reason that threat actors developed the AiTM session cookie theft technique in the first place. Organizations are advised to work with their identity provider to ensure security controls like MFA are in place. Microsoft customers can implement through various methods, such as using the Microsoft Authenticator, FIDO2 security keys, and certificate-based authentication.

Defenders can also complement MFA with the following solutions and best practices to further protect their organizations from such attacks:

- **Use security defaults** as a baseline set of policies to improve identity security posture. For more granular control, **enable conditional access policies, especially risk-based access policies.** Conditional access policies evaluate sign-in requests using additional identity-driven signals like user or group membership, IP location information, and device status, among others, and are enforced for suspicious sign-ins. Organizations can protect themselves from attacks that leverage stolen credentials by enabling policies such as compliant devices, trusted IP address requirements, or risk-based policies with proper access control.
- **Implement continuous access evaluation.**
- **Invest in advanced anti-phishing solutions** that monitor and scan incoming emails and visited websites. For example, organizations can leverage web browsers that automatically identify and block malicious websites, including those used in this phishing campaign, and solutions that detect and block malicious emails, links, and files.
- **Continuously monitor suspicious or anomalous activities.** Hunt for sign-in attempts with suspicious characteristics (for example, location, ISP, user agent, and use of anonymizer services).

Detections

Because AiTM phishing attacks are complex threats, they require solutions that leverage signals from multiple sources. Microsoft 365 Defender uses its cross-domain visibility to detect malicious activities related to AiTM, such as session cookie theft and attempts to use stolen cookies for signing in.

Using Microsoft Defender for Cloud Apps connectors, Microsoft 365 Defender raises AiTM-related alerts in multiple scenarios. For Azure AD customers using Microsoft Edge, attempts by attackers to replay session cookies to access cloud applications are detected by Defender for Cloud Apps connectors for Office 365 and Azure. In such scenarios, Microsoft 365 Defender raises the following alert:

> Stolen session cookie was used

In addition, signals from these Defender for Cloud Apps connectors, combined with data from the Defender for Endpoint network protection capabilities, also triggers the following Microsoft 365 Defender alert on Azure AD environments:

> Possible AiTM phishing attempt

A specific Defender for Cloud Apps connector for Okta, together with Defender for Endpoint, also helps detect AiTM attacks on Okta accounts using the following alert:

> Possible AiTM phishing attempt in Okta

Other detections that show potentially related activity are the following:

Microsoft Defender for Office 365

- Email messages containing malicious file removed after delivery
- Email messages from a campaign removed after delivery
- A potentially malicious URL click was detected
- A user clicked through to a potentially malicious URL
- Suspicious email sending patterns detected

Microsoft Defender for Cloud Apps

- Suspicious inbox manipulation rule
- Impossible travel activity
- Activity from infrequent country
- Suspicious email deletion activity

Azure AD Identity Protection

- Anomalous Token
- Unfamiliar sign-in properties
- Unfamiliar sign-in properties for session cookies

Microsoft 365 Defender

- BEC-related credential harvesting attack
- Suspicious phishing emails sent by BEC-related user

Hunting queries

Microsoft Sentinel

Microsoft Sentinel customers can use the following analytic templates to find BEC related activities similar to those described in this post:

In addition to the analytic templates listed above Microsoft Sentinel customers can use the following hunting content to perform Hunts for BEC related activities:

Further reading

For the latest security research from the Microsoft Threat Intelligence community, check out the Microsoft Threat Intelligence Blog: https://aka.ms/threatintelblog.

To get notified about new publications and to join discussions on social media, follow us on Twitter at https://twitter.com/MsftSecIntel.

Related Posts

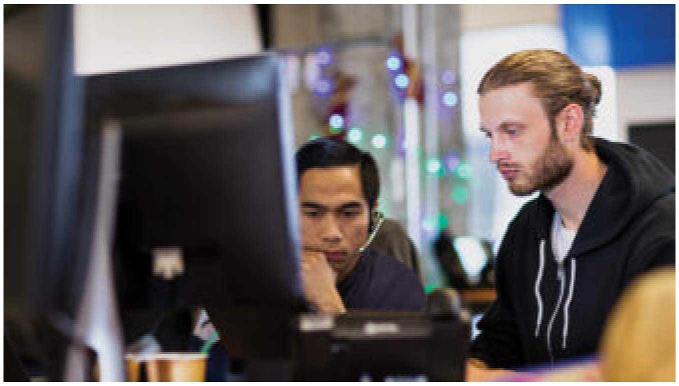

Microsoft Incident Response lessons on preventing cloud identity compromise

In real-world customer engagements, Microsoft IR sees combinations of issues and misconfigurations that could lead to attacker access to customers' Microsoft Entra ID tenants. Reducing risk and exposure of your most privileged accounts plays a critical role in preventing or detecting attempts at tenant-wide compromise.

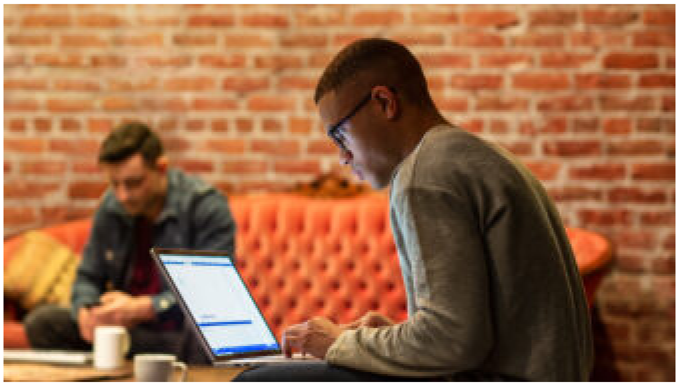

Octo Tempest crosses boundaries to facilitate extortion, encryption, and destruction

Microsoft has been tracking activity related to the financially motivated threat actor Octo Tempest, whose evolving campaigns represent a growing concern for many organizations across multiple industries.

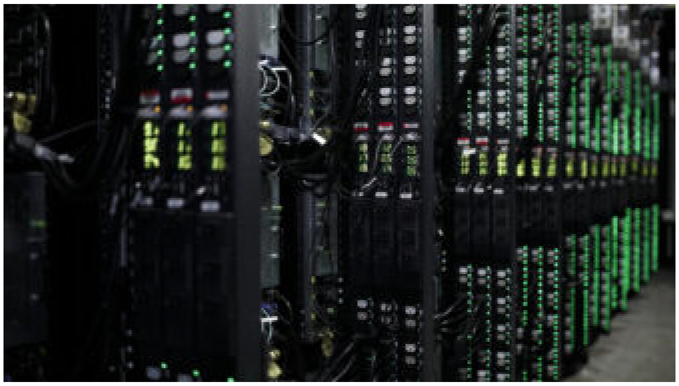

Defending new vectors: Threat actors attempt SQL Server to cloud lateral movement

Microsoft security researchers recently identified an attack where attackers attempted to move laterally to a cloud environment through a SQL Server instance. The attackers initially exploited a SQL injection vulnerability in an application within the target’s environment to gain access and elevated permissions to a Microsoft SQL Server instance deployed in an Azure Virtual Machine (VM). The attackers then used the acquired elevated permission to attempt to move laterally to additional cloud resources by abusing the server’s cloud identity.