Malicious OAuth applications abuse cloud email services to spread spam

By Microsoft Threat Intelligence

Published: 2022-09-22 · Archived: 2026-04-06 00:14:33 UTC

Microsoft researchers recently investigated an attack where malicious OAuth applications were deployed on compromised cloud tenants and then used to control Exchange Online settings and spread spam. The investigation revealed that the threat actor launched credential stuffing attacks against high-risk accounts that didn't have multi-factor authentication (MFA) enabled and leveraged the unsecured administrator accounts to gain initial access. The unauthorized access to the cloud tenant enabled the actor to create a malicious OAuth application that added a malicious inbound connector in the email server. The actor then used the malicious inbound connector to send spam emails that looked like they originated from the targets' domain. The spam emails were sent as part of a deceptive sweepstakes scheme meant to trick recipients into signing up for recurring paid subscriptions.

Microsoft has been monitoring the rising popularity of OAuth application abuse. One of the first observed malicious usage of OAuth applications in the wild is consent phishing. Consent phishing attacks aim to trick users into granting permissions to malicious OAuth apps to gain access to user's legitimate cloud services (mail servers, files storage, management APIs, etc.). In the past few years, Microsoft has observed that more and more threat actors, including nation-state actors, have been using OAuth applications for different malicious purposes – command-and-control (C2) communication, backdoors, phishing, redirections, and so on.

This recent attack involved a network of single-tenant applications installed in compromised organizations being used as the actor's identity platform to perform the attack. As soon as the network was revealed, all the related applications were taken down and notifications to customers were sent, including recommended remediation steps.

This blog presents the technical analysis of this attack vector and the succeeding spam campaign attempted by the threat actor. It also provides guidance for defenders on protecting organizations from this threat, and how Microsoft security technologies detect it.

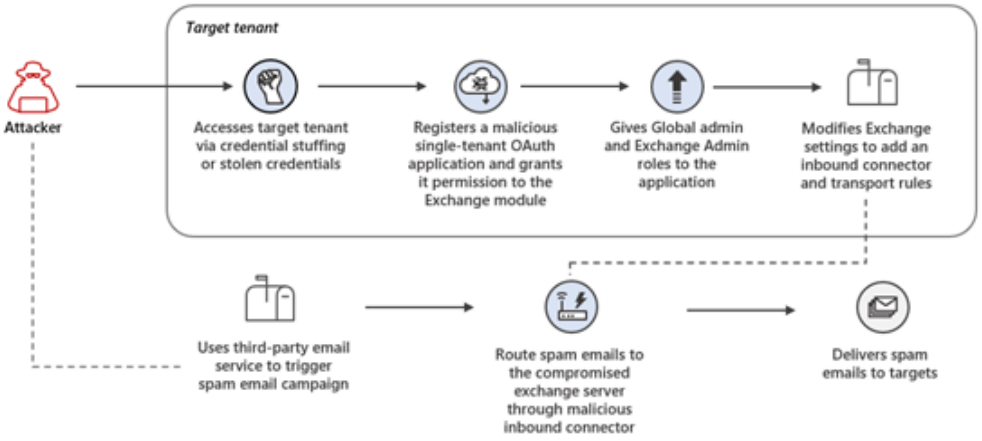

Figure 1. Overview of the attack chain. The time between application deployment and usage varied; there were cases where the actor took months before using the application.

Initial access

For the attack to succeed, the threat actor needed to compromise cloud tenant users with sufficient permissions that would allow the actor to create an application in the cloud environment and give it admin consent. The actor performed credential stuffing attacks against their targets, attempting to access users with the global admin role. The authentication attempts, which originated from a single IP address, were launched against the Azure Active Directory PowerShell application (app ID: 1b730954-1685-4b74-9bfd-dac224a7b894). The same application was later used to deploy the rest of the attack.

Based on the success ratio of the authentication attempts, it is inferred that the attacker used a dump of compromised credentials. The investigation also revealed that 86% of the compromised tenants had at least one admin with a real-time high risk score, which means they were flagged by Azure AD Identity Protection to be most likely compromised. It is also important to note that all the compromised admins didn't have MFA enabled, which could have stopped the attack. These observations amplify the importance of securing accounts and monitoring for high-risk users, especially those with high privileges.

Deploying malicious OAuth application

Once the threat actor gained access to privileged users, their next step was to set up the malicious application. Based on analysis of the event user agent (Swagger-Codegen/1.4.0.0/csharp) and how quickly the deployment of the application was done, it is likely that the actor ran a PowerShell script to perform the following Azure Active Directory (AAD) management activities in all targeted tenants:

- Register a new single–tenant application with the naming convention of *[domain name]_([a-zA-Z]){3}* (for example: Contoso_GhY)
- Add the legacy permission *Exchange.ManageAsApp* which can be used for app-only authentication of Exchange Online PowerShell module
- Grant admin consent to the above permission

- Give global admin and Exchange Online admin roles to the previously registered application
- Add application credentials (key/certificate/both)

The threat actor added their own credentials to the OAuth application, which enabled them to access the application even if the initially compromised global administrator changed their password.

The activities mentioned gave the threat actor control of a highly privileged application. It was observed that the threat actor did not always use the application right after it was deployed. In some cases, it took weeks or months before the application was utilized. Also, in organizations that didn't monitor for suspicious applications, the applications were deployed for months and used multiple times by the threat actor.

Modifying Exchange Online settings

The threat actor used the privileged application to authenticate the Exchange Online PowerShell module and modify the Exchange Online settings. There were two modifications which allowed them to perform the next step in the attack chain:

Create a new inbound connector

Connectors are a collection of instructions that customize the way email flows to and from organizations using Microsoft 365 or Office 365. Most organizations using Microsoft 365 and Office 365 don't need custom connectors for regular mail flow, but some use it when they need to process messages from another messaging system that's not running Exchange Online, or if they have a network appliance that performs policy checks and then routes messages to their Exchange Online service.

The threat actor set a new inbound connector with the naming convention *Ran_([a-zA-Z]){5}* (for example *Ran_xAFzd*). The purpose of the inbound connector was to allow mails from certain IPs (that are related to the attacker's infrastructure) to flow through the victim's Exchange Online service. This allowed the threat actor to send emails that looked like they originated from the compromised Exchange domain. The configuration information for the newly created connectors were seen in Exchange Online audit event *New-InboundConnector*. The following table shows the configuration parameters in the audit event related to this change:

Name	Value
"Name"	"Ran_jBelh"
"Enabled"	"True"
"CloudServicesMailEnabled"	"True"
"RestrictDomainsToCertificate"	"False"
"SenderDomains"	"smtp:*;1"
"SenderIPAddresses"	"170.75.174.97;170.75.172.8;170.75.170.69;170.75.174.95;54.39.94.145;149.56.200.36;158.69.21.185;66.70.201.131"

“RestrictDomainsToIPAddresses”	“True”
“ConnectorSource”	“HybridWizard”
“EFSkipIPs”	“170.75.174.97;170.75.172.8;170.75.170.69;170.75.174.95;54.39.94.145;149.56.200.36;158.69.21.185;66.70.201.131”
“TreatMessagesAsInternal”	“False”
“ConnectorType”	“OnPremises”
“RequireTls”	“False”

Create transport rules

Transport rules (also known as mail flow rules) are sets of actions that can be taken on any mail that flows in the organization. The threat actor utilized this feature to set 12 new transport rules with the naming convention of *Test01* to *Test012*. Each of these transport rules were responsible for deleting the following specific headers from every mail that flowed in the organization:

- X-MS-Exchange-ExternalOriginalInternetSender
- X-MS-Exchange-SkipListedInternetSender
- Received-SPF
- Received
- ARC-Authentication-Results
- ARC-Message-Signature
- DKIM-Signature
- ARC-Seal
- X-MS-Exchange-SenderADCheck
- X-MS-Exchange-Authentication-Results
- Authentication-Results
- X-MS-Exchange-AntiSpam-MessageData-ChunkCount

By deleting these headers, the attacker tried defense evasion to prevent security products or email providers detecting or blocking their emails and increase the success rate of the spam campaign. The configuration information for the newly created transport rules were seen in Exchange Online audit event *New-TransportRule*. The following table shows the configuration parameters in the audit event related to this change:

Name	**Value**
“Name”	“Test06”
“SenderAddressLocation”	“Header”
“RemoveHeader”	“ARC-Message-Signature”

These modifications to the Exchange Online settings allowed the threat actor to perform their primary goal in the attack: sending out spam emails. After each spam campaign, the actor deleted the malicious inbound connector and transport rules to prevent detection, while the application remained deployed in the tenant until the next wave of the attack (in some cases, the app was dormant for months before it was reused by the threat actor).

```
Connect-ExchangeOnline -CertificateFilePath "<localpathtocert>" -CertificatePassword
(ConvertTo-SecureString -String "<password>" -AsPlainText -Force) -AppID "<appid>" -
Organization "<domain>.onmicrosoft.com"

New-InboundConnector -Name "Ran_GtrHs" -SenderDomains "smtp:*;1" -SenderIPAddresses
"x.x.x.x/26" -RestrictDomainsToIPAddresses $true -EFSkipIPs "x.x.x.x/26" -RequireTls
$false

New-TransportRule -Name "Test_skiplistedsender" -SenderAddressLocation "Header" -
RemoveHeader "X-MS-Exchange-SkipListedInternetSender"

New-TransportRule -Name "Test_removeoriginalsender" -SenderAddressLocation "Header" -
RemoveHeader "X-MS-Exchange-ExternalOriginalInternetSender"
```

Figure 2. An example of the PowerShell script used for setting new Exchange Online connector and transport rules

Spam email campaign sent through Exchange Online connector

The actor behind this attack has been actively running spam email campaigns for many years. Based on our research, this actor has sent high volumes of spam emails in short timeframes through other methods, such as connecting to mail servers from rogue IP addresses or sending directly from legitimate cloud-based bulk email sending infrastructure.

The actor’s motive was to propagate deceptive sweepstakes spam emails designed to trick recipients into providing credit card details and signing up for recurring subscriptions under the guise of winning a valuable prize. While the scheme possibly led to unwanted charges for targets, there was no evidence of overt security threats such as credential phishing or malware distribution.

The spam campaign carried the hallmarks of this actor: programmatically generated messages containing two visible hyperlinked images in the email body, as well as dynamic and randomized content injected within the HTML body of each mail message to evade spam filters.

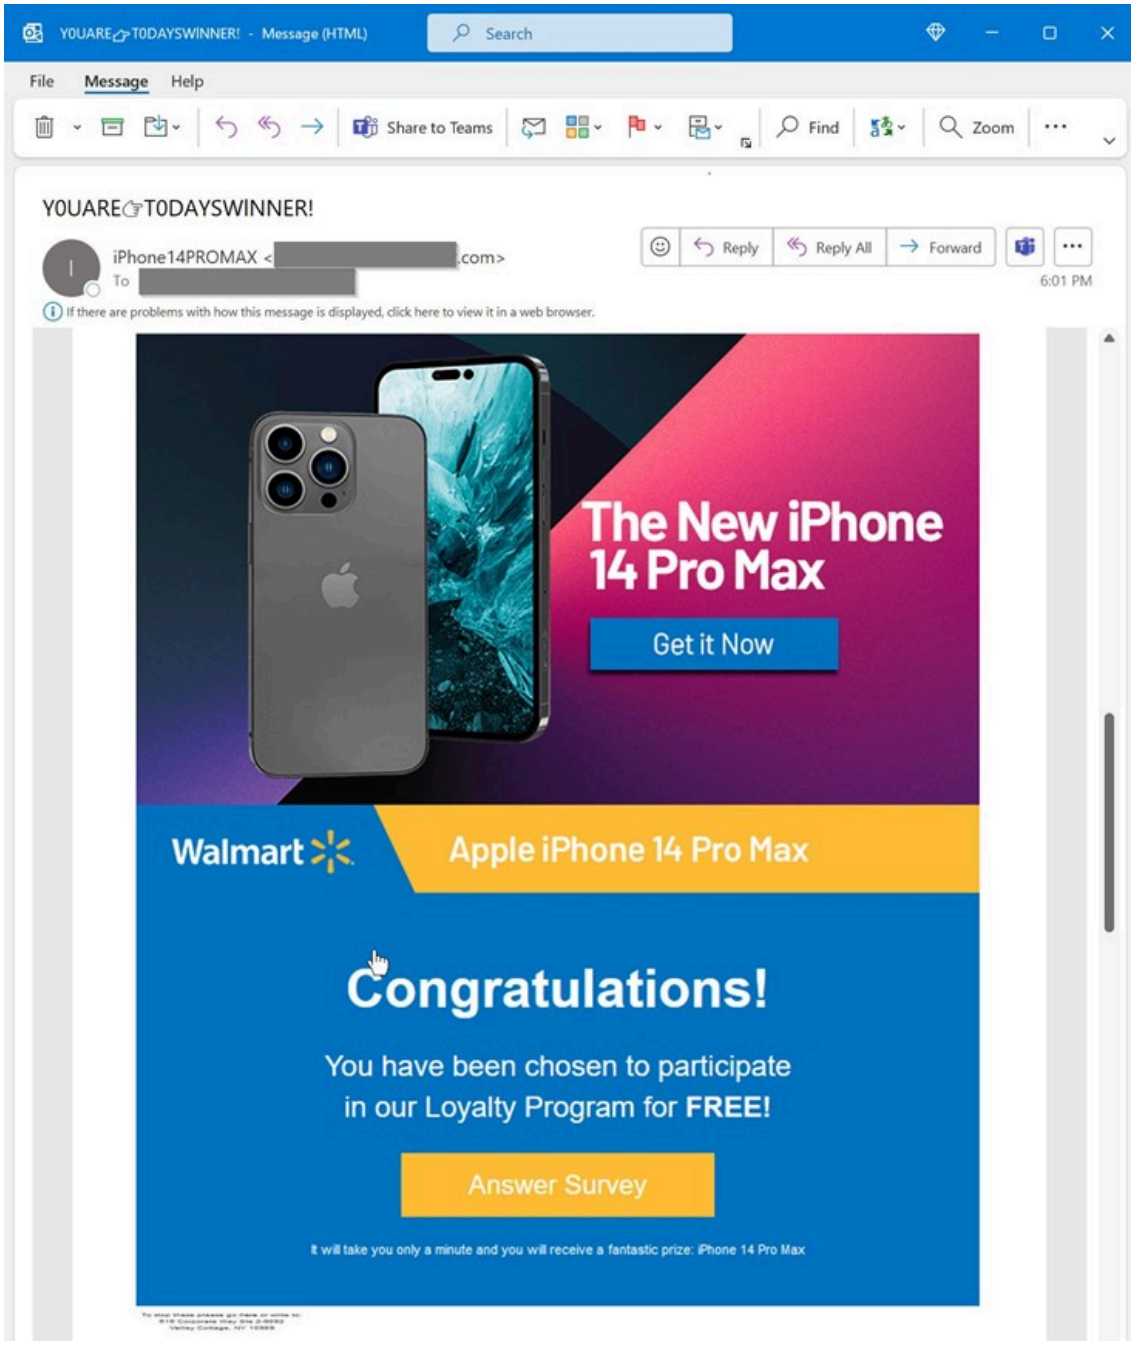

Figure 3. An example of spam email sent through the Exchange Online inbound connector

The hyperlinked images within the spam messages implied to recipients that they were eligible for a prize. Once clicked, the hyperlink directed recipients to a website where they were asked to complete a survey and provide credit card details to pay for the shipping of their prize. Familiar brand logos, names, and websites were used throughout the spam email, likely to give an illusion of legitimacy.

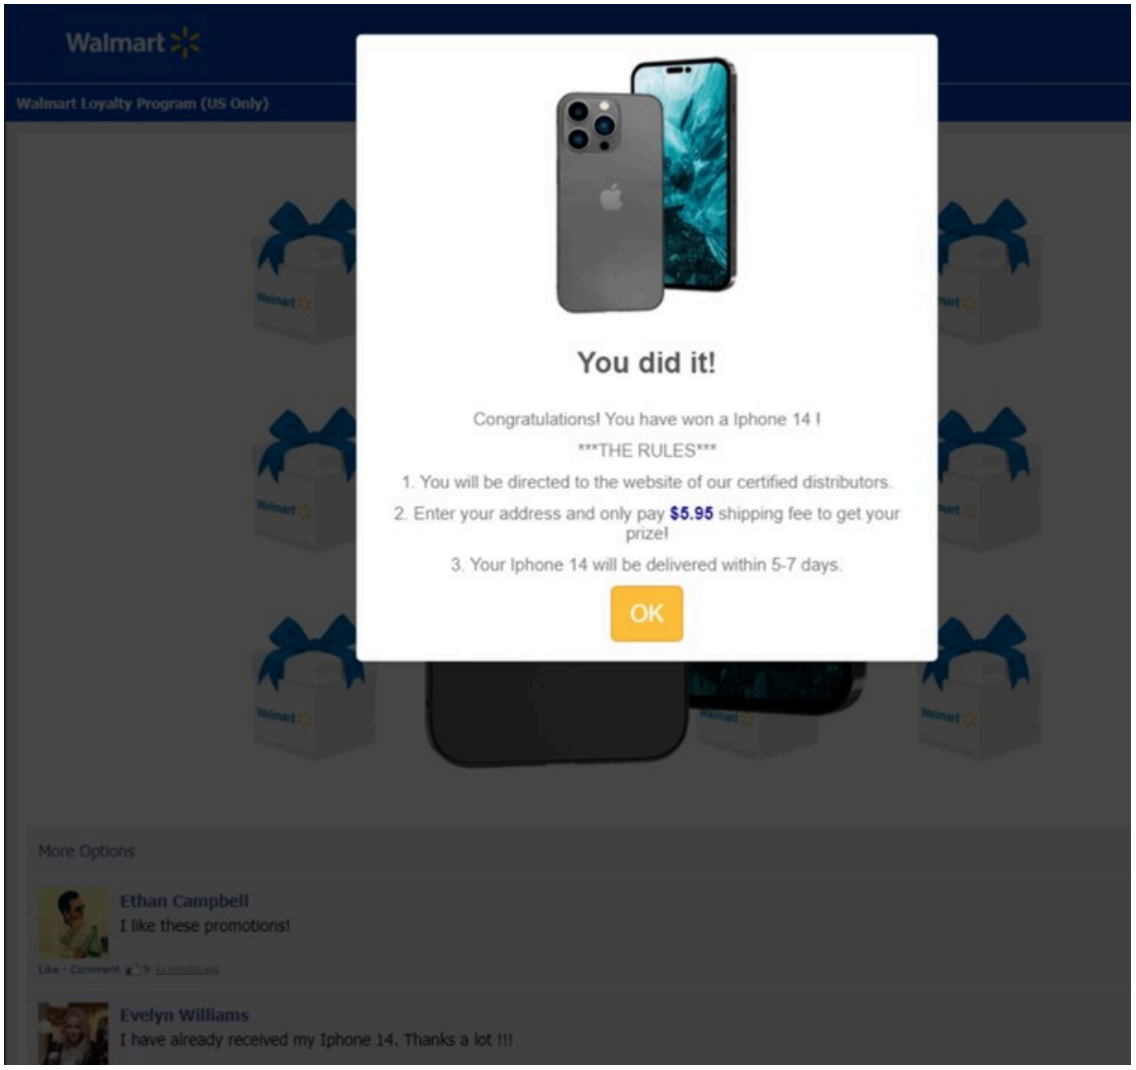

Figure 4. Clicking the image in the spam email leads to this website indicating a prize has been won

The fine print, visible only by scrolling to the very bottom of a subsequent page, revealed that in providing their credit card information, recipients were not paying a shipping fee for their prize, but were in fact agreeing to be charged fees for several paid subscription services in order to enter into a *sweepstakes* for the prize. It is likely the threat actor's main motive was potential financial gain from people who fell victim to this deceptive sweepstakes spam campaign.

A winner will be selected from all eligible entrants. The iPhone 14 Pro winner will be directly contacted by email.

One day trial subscription to Please Prize Club for usd $49.99, if you continue your subscription after the trial period we will automatically bill your card for the monthly fee of $99.97 until you cancel your subscription. You will also get subscribed to our health app for which you will be charged a total of $189.94 today. You will also get charged $128.94 after two days and then every month to retain access to our health app. You will have access to the services as long as your subscription is active. 24/7 customer support is available by phone and email. Just call or contact us at .com. iPhone 14 Pro

iPhone 14 Pro Sweepstakes Entry SWEEPSTAKES ABBREVIATED RULES: NO PURCHASE NECESSARY... Void where prohibited and outside US. Open to legal residents of the 50 US (and DC) who are 18 yrs+ (19 yrs+ in AL and NE, 21 yrs+ in MS) at entry. Sweepstakes begins 12:00 AM PST on September 1 and ends 11:59 PM PST on September 31, 2023. 2 ways to enter: (a) sign up for a one week trial subscription of the Fitness App and pay $49.99 (if you do not cancel, you will be billed $99.97/month in addition to

Figure 5. The fine print shows the chance to win a prize is only through a sweepstakes after paying fees

Likely to achieve scale and further maximize the chances of successful email delivery, the actor triggered this spam campaign from cloud-based outbound email infrastructure outside of Microsoft, mainly Amazon SES and Mail Chimp. These email platforms enable sending of mass bulk email, normally for marketing and other legitimate purposes.

The campaign also utilized techniques to generate unique dynamic URLs underpinning the hyperlinked images within each spam email message, as shown in the examples below.

```
http://asas4as.s564as.[redacted].com/Yz5czUViOnnhAFA2KBS4qnt2iXt3zyPijE
http://asas4as.s564as.[redacted].com/7EIPUbiOaMcqFWlOGSmpMHLiAuZnVF2P
http://asas4as.s564as.[redacted].com/dwPTJU4igwJcYF1mbeFdzMH5iKtVeDhAm
```

Figure 6. Examples of dynamic URL generation (domain obfuscated)

This spam campaign exclusively targeted consumer email accounts. In the case of spam messages sent to Microsoft-hosted consumer email accounts (outlook.com), the spam emails were moved into customers' junk folders before they could be viewed and clicked.

Mitigations

While the follow-on spam campaign targets consumer email accounts, this attack targets enterprise tenants to use as infrastructure for this campaign. This attack thus exposes security weaknesses that could be used by other threat actors in attacks that could directly impact affected enterprises.

As the main initial access vector of the attack was to obtain the admin's credentials, we recommend organizations take the following steps to reduce their attack surface:

Mitigate credential guessing attacks risks

A key step in reducing the attack surface is securing the identity infrastructure. The most common initial access vector observed in this attack was account compromise through credential stuffing, and all the compromised administrator accounts did not have MFA enabled. Implementing security practices that strengthen account credentials such as enabling MFA raises the cost of an attack.

Enable conditional access policies

Conditional access policies are evaluated and enforced every time the user attempts to sign in. Organizations can protect themselves from attacks that leverage stolen credentials by enabling policies such as device compliance or trusted IP address requirements.

Enable continuous access evaluation

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.

Enable security defaults

While some of the features mentioned above require paid subscriptions, the security defaults in Azure AD, which is mainly for organizations using the free tier of Azure Active Directory licensing, are sufficient to better protect the organizational identity platform, as they provide preconfigured security settings such as MFA, protection for privileged activities, and others.

Leveraging its cross-signal capabilities, Microsoft 365 Defender alerts customers using Microsoft Defender for Office 365, Microsoft Defender for Cloud Apps, Application governance add-on, and Azure Active Directory Identity Protection to detect the techniques covered in the attack through the attack chain. Each product can provide a different aspect for protection to cover the techniques observed in this attack:

Microsoft 365 Defender

Suspicious email-sending pattern from new Exchange inbound connector – This alert is generated when a suspicious email-sending pattern originating from a new Exchange inbound connector is detected. This behavior might suggest that an attacker set a malicious inbound connector to allow anonymous relay through the organization's Exchange Online service.

Recommended reading to response for malicious connector incidents:

- Respond to a compromised connector in Microsoft 365
- Alert grading for malicious exchange connectors

New transport rule removing antispam header – This alert is generated when a new transport rule to remove antispam header is detected.

Suspicious inbound connector and transport rule created to remove sender email headers – This alert is generated when a suspicious inbound connector and transport rule is created to remove headers that identify the true source address of sender. This might indicate a spam campaign is ongoing from the organization's mailbox.

Suspicious Azure AD app creation – This alert is generated when a user account creates Azure Active Directory OAuth application with suspicious characteristics, as observed in this campaign.

Azure AD app registration by risky user – This alert is generated when a user account with high risk score as calculated by AAD Identity Protection is creating a new OAuth application and grants admin consent to it.

Microsoft Defender for Office 365

Microsoft Defender for Office 365 detects threat activity associated with this spamming 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 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 zero-hour auto purge (ZAP) if this event occurs.

Microsoft Defender for Cloud Apps

Activity from suspicious IP addresses. This alert is generated when there is activity from an IP address that has been identified as risky by Microsoft Threat Intelligence or by the organization. These IP addresses were identified as being involved in malicious activities, such as performing password spray, botnet C2, and may indicate a compromised account.

Activity from a password-spray associated IP address. This alert is generated when a successful sign-in from an IP address that had been identified as participating in password spray was observed.

App governance

App governance is an add-on to Microsoft Defender for Cloud Apps, which can detect malicious OAuth applications that make sensitive Exchange Online Administrative activities along with other threat detection alerts. Activity related to this campaign will trigger the following alert:

OAuth app with suspicious metadata has exchange permission. This alert is generated when a line of business OAuth app with suspicious metadata has privilege to manage permission over Exchange Online, which can lead an OAuth App to perform data collection or exfiltration activities or attempts to access and retrieve sensitive information.

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.This alert flags anomalies in the token claims, token age, and other authentication attributes.
- **Anonymous IP address.** This alert flag sign-in attempts from anonymous IP addresses (for example, Tor browser or anonymous VPN).

Hunting queries

To locate related activity, Microsoft 365 Defender customers can run the following advanced hunting queries:

Applications given the "Exchange.ManageAsApp" permission:

```
CloudAppEvents
| where Timestamp > ago(30d)
| where ActionType == "Add app role assignment to service principal."
| where RawEventData.ResultStatus == "Success"
| where RawEventData has "dc50a0fb-09a3-484d-be87-e023b12c6440" //Exchange.ManageAsApp Role Id
| project Timestamp, AccountObjectId,AccountDisplayName, AppId =
RawEventData.ModifiedProperties[0].NewValue, AppName = RawEventData.ModifiedProperties[6].NewValue
```

New transport rules that remove anti-spam headers from emails:

```
CloudAppEvents
| where ActionType == "New-TransportRule"
| mvexpand Param = RawEventData.Parameters
| where Param.Name == "RemoveHeader" and Param.Value contains "X-MS-Exchange-AntiSpam-MessageData"
| project Timestamp, AccountObjectId, AccountDisplayName, ServicePrincipalId =
tostring(RawEventData.AppId)
```

> **Update 09/23/2022** – Updated instances 'Exchange server' to 'Exchange Online" to clarify that attackers were not able to compromise on-premises Exchange servers, but that the permissions they gained during their compromise of cloud tenants allowed the attackers to modify Exchange Online settings.

Source: https://www.microsoft.com/en-us/security/blog/2022/09/22/malicious-OAuth-applications-used-to-compromise-email-servers-and-spread-spam/