VECTRA
SECURITY THAT THINKS

2021 Q2 SPOTLIGHT REPORT

Vision and Visibility: Top 10 Threat Detections for Microsoft Azure AD and Office 365

TABLE OF CONTENTS

Vectra® protects business by detecting and stopping cyberattacks

Vectra® is the leader in threat detection and response—from cloud and data center workloads to user and IoT devices. Its Cognito® platform accelerates threat detection and investigation using AI to enrich network metadata it collects and stores with the right context to detect, hunt and investigate known and unknown threats in real time. Vectra offers four applications on the Cognito platform to address high-priority use cases. Cognito Stream™ sends security-enriched metadata to data lakes and SIEMs. Cognito Recall™ is a cloud-based application to store and investigate threats in enriched metadata. Cognito Detect™ uses AI to reveal and prioritize hidden and unknown attackers at speed. And Cognito Detect for Office365 and Azure AD™ finds and stops attacks in enterprise SaaS applications and the Microsoft 365 ecosystem. For more information, visit vectra.ai.

71% 

of users have suffered an account takeover of a legitimate user's account on average 7 times in the last year.*

HIGHLIGHTS

- Meaningful AI can provide constant analysis of how users access, use and configure cloud apps and can make all the difference in being able to detect and stop threats like account takeovers.
- The Top 10 Threat Detections seen across Microsoft Azure AD and Office 365 allow security teams to detect infrequent behavior that is abnormal or unsafe across their environments.
- Regardless of company size, Office 365 Risky Exchange Operation detection was at or near the top of the list of detections seen by Vectra customers.
- Common actions by actors in the Azure AD environment during a recent supply chain attack would map back to Vectra-defined detections and alert the security team about the threat.

* Securing Microsoft Office 365 in the new normal.

Detecting the "Out of the Ordinary"

The cloud continues to change everything we know about security, leaving the legacy approach to protecting assets obsolete. However, collecting the right data and having meaningful artificial intelligence (AI) can help pinpoint the ins and outs of attacks so security teams can focus on the threats that actually require attention, rather than spending valuable cycles on benign alerts. In this report, we'll discuss the top 10 threat detections seen across the Vectra customer base that help ratify attacks across Microsoft Azure AD and Office 365. All of the data presented represents real detection examples that security organizations receive from Vectra when something out of the ordinary happens.

Threat Detections Tell the Story

While this report focuses on the top 10 Azure AD and Office 365 detections our customers see by relative frequency, it's important to keep in mind that threat detection and response is easiest when adversaries take actions that are obviously malicious. The unfortunate reality for modern network defenders, however, is that adversaries increasingly find that such overt action is unnecessary when existing services and access used throughout an organization can simply be co-opted, misused and abused.

This makes it critical that modern network defenders understand the intersection that may exist between the types of actions an adversary would need to take to progress towards their objectives and the behaviors routinely taken by authorized users across the enterprise. In cases where these behaviors intersect, the key factors in distinguishing the adversary and insider threat from a benign user is intent, context and authorization. Having the insight and knowledge that meaningful AI can provide through constant analyzation of how users access, use and configure their cloud apps, while also knowing how your hosts, accounts and workloads are being accessed—can make all the difference.

The key factors in distinguishing the adversary and insider threat from a benign user is intent, context and authorization.

This difference can mean being able to see the clear contrast between something serious like a detection that notifies you about suspicious mail forwarding in Office 365 versus seeing a multitude of notifications that don't pose a threat, or maybe worse—not seeing anything. The security stakes are already at an all-time high with threats like account takeovers costing companies billions in annual losses, but the good news is that the behavior used in these criminal tactics are no longer a secret when you involve AI. Let's take a close look at the top 10 detections.

Top 10 Threat Detections

Many of the detections discussed in this report represent anomalous behavior and not all of these detections are due to malicious activity. Meanwhile, some may represent infrequent behavior that is abnormal for the environment, while others may represent behavior that is against policy. Regardless of the activity, these detections still represent a large attack surface that our customers need to manage. For more information about the science behind these detections, please see the white paper: The Data Science Behind Vectra AI Threat Detection Models.

> Regardless of the activity, these detections still represent a large attack surface that our customers need to manage.

Vision and Visibility: The Intersection Between Adversary and Defender

Vision:

Without a clear expectation of what is authorized, often in the form of prescriptive policy, security defenders will have difficultly doing anything but solving for the obvious threats. That is why organizations need to have a ***vision*** for what authorized use looks like when it comes to the cloud services they adopt, which is often expressed via policy.

Your vision for authorized use of cloud services should consider:

- Which services and behaviors are authorized?
- Under what context are they authorized?
- Are users authorized to leverage cloud storage and how should they interact with external entities?
- What operational parameters and safeguards are expected to accompany behaviors involving these cloud services?

Organizations need to have a ***vision*** for what authorized use looks like when it comes to the cloud services they adopt.

Visibility:

Even with a clear vision, organizations quickly get into trouble when there is a lack of ***visibility***. This is because they lack the ability to monitor and measure deviations from their vision. Solving for this challenge involves understanding the behaviors adversaries are motivated to take, and intentionally collecting and aggregating the data that uncovers these behaviors in a way that can be operationalized by security staff.

> Even with a clear vision, organizations quickly get into trouble when there is a lack of ***visibility***.

Taking Visibility to the Next Level

- **Services**: Are defenders able to detect malicious attacks carried out through enterprise cloud services? Things like detecting PowerAutomate abuse for command and control (C2) despite the protections deployed around those capabilities?
- **Management**: Are defenders able to identify the misuse and abuse of administrative and management functions, such as risky exchange operations that enable adversaries and insiders to escalate privileges or to collect and exfiltrate sensitive information?
- **Supply Chain**: Are defenders able to uncover instances where trusted suppliers and service providers have been compromised, giving adversaries a beachhead that bypasses an organization's preventative and protective controls?

Security leaders should be confident answering these questions as they require more than simple compliance checks or benchmarks—these questions are best pressure tested through active investigations, threat hunts and security testing exercises.

Relative Frequency Top 10 Threat Detections

This graph shows the top 10 threat detections in Vectra's customer base by relative frequency, plotted over time. Percentage of customers who triggered the detection per week.

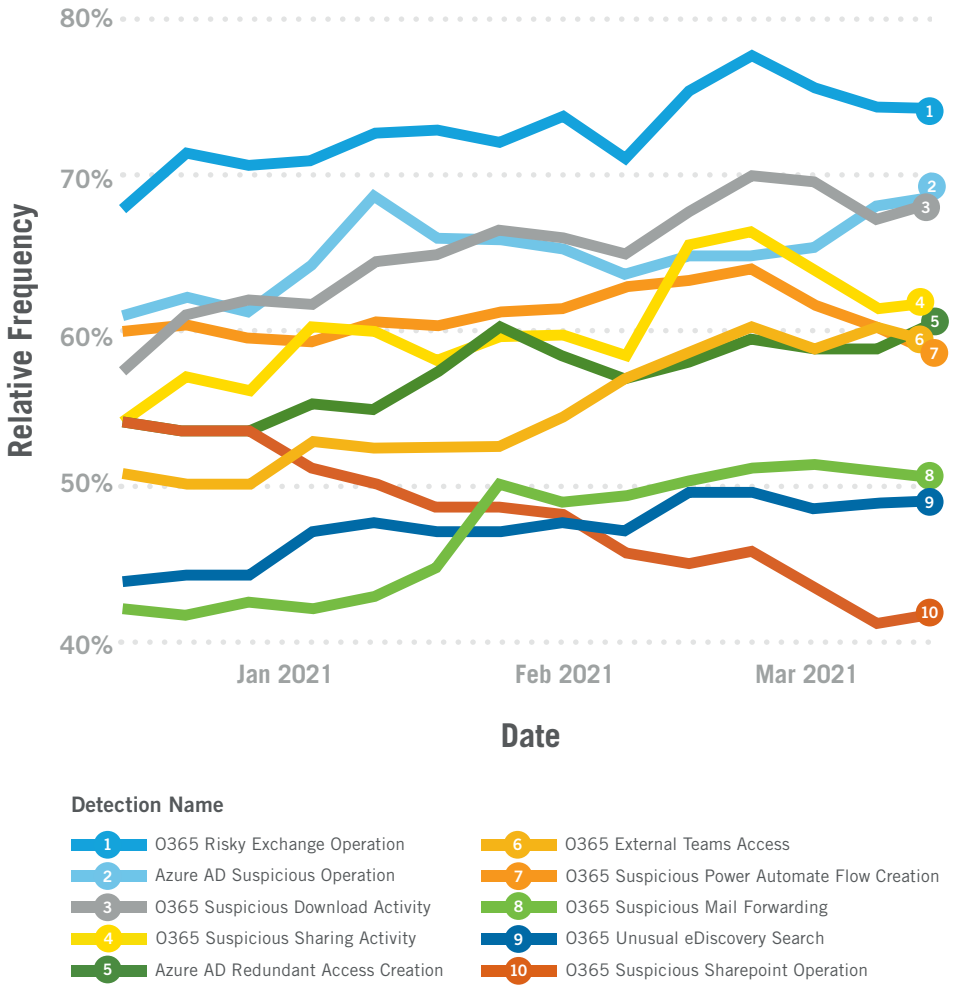

Top 10 Threat Detections

1 O365 Risky Exchange Operation
Abnormal Exchange operations have been detected that may indicate an attacker is manipulating Exchange to gain access to specific data or further attack progression.

2 Azure AD Suspicious Operation
Abnormal Azure AD operations have been detected that may indicate attackers are escalating privileges and performing admin-level operations after regular account takeover.

3 O365 Suspicious Download Activity
An account was seen downloading an unusual number of objects which may indicate an attacker is using SharePoint or OneDrive download functions to exfiltrate data.

4 O365 Suspicious Sharing Activity
An account was seen sharing files and/or folders at a volume that is higher than normal which may indicate an attacker is utilizing SharePoint to exfiltrate data or maintain access after initial access has been remediated.

5 Azure AD Redundant Access Creation
Administrative privileges have been assigned to an entity which may indicate redundant access is being created by the attacker to guard against remediation.

6 O365 External Teams Access
An external account has been added to a team in O365 Teams which may indicate an adversary has added an account under their control.

7 O365 Suspicious Power Automate Flow Creation
An abnormal Power Automate Flow creation has been observed which may indicate an attacker is configuring a persistence mechanism.

8 O365 Suspicious Mail Forwarding
Mail forwarding which may be used as a collection or exfiltration channel without the need to maintain persistence.

9 O365 Unusual eDiscovery Search
A user is creating or updating an eDiscovery search which may indicate an attacker has gained access to eDiscovery capabilities and is now performing reconnaissance.

10 O365 Suspicious Sharepoint Operation
Abnormal administrative SharePoint operations that may be associated with malicious activities.

Not just convenient collaboration

A number of these threat detections represent activities that provide ease of use, collaboration with external parties and provisioning of administrative access to the Azure AD environment. Being able to easily share documents from OneDrive or SharePoint facilitates the sharing of information with external parties, but it also provides a means for an attacker to gain access to information already in place or staged in these services. In this instance, most of the detections are Office 365 Suspicious Download Activity, Office 365 Suspicious SharePoint Operation and Office 365 Suspicious Sharing Activity.

Some of the detections could be triggered from communicating and collaborating with external users via Microsoft Teams, which is certainly convenient for legitimate users. However, it can also be a convenient means for attackers to find useful information or obtain documents and information. It's not uncommon to receive detections that notify about Office 365 External Teams access.

A number of these threat detections represent activities that provide ease of use, collaboration with external parties and provisioning of administrative access to the Azure AD environment.

How do the Threat Detections Measure up Based on Company Size?

Top 10 Threat Detections – Small Company

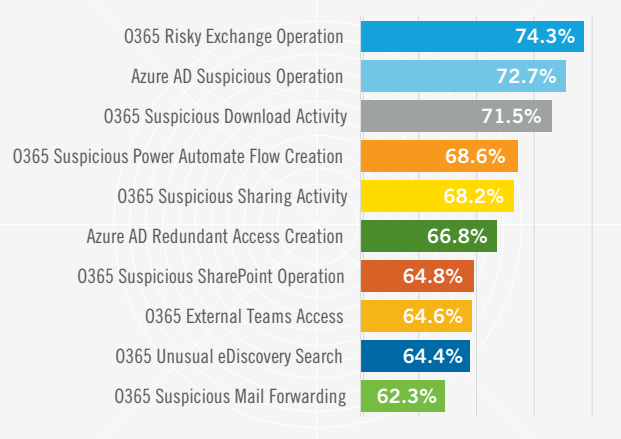

Top 10 Threat Detections – Medium Company

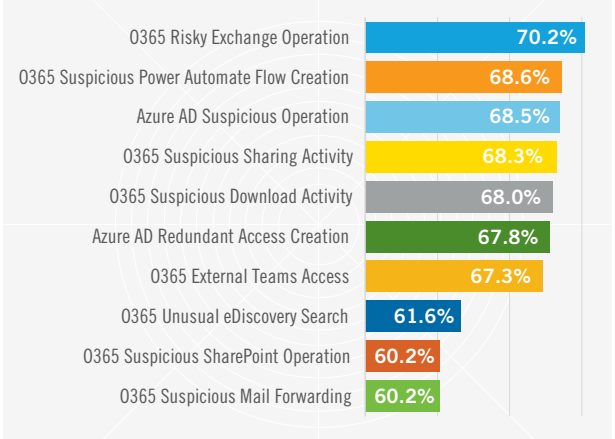

Top 10 Threat Detections – Large Company

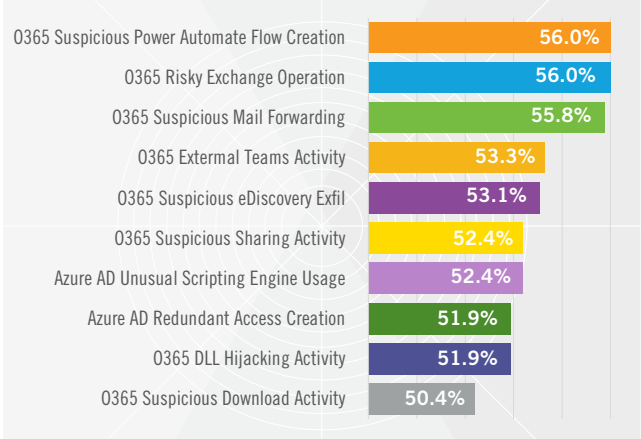

Vectra calculated the relative frequency of threat detections that were triggered during a three-month span based on customer size. Looking at the top 10 threat detections per customer size (small, medium and large)—the larger the customer, the smaller the percentage of detections that trigger on a per detection type basis. This general trend of larger companies triggering fewer detections when compared to smaller companies tells us that the larger companies' users and administrators may perform Azure AD and Office 365 activity more consistently compared to smaller organizations. Here's what we found:

Larger companies' users and administrators may perform Azure AD and Office 365 activity more consistently compared to smaller organizations.

Detection similarities between medium and small companies

Result: Looking at the breakdown of which detection types were in the top 10 for each customer size, we see medium and small companies have the same top 10 threat detections, albeit with different detection type rankings. Large customers' top 10 breakdown shows us that Office 365 DLL Hijacking, Office 365 Unusual Scripting Engine and Office 365 Suspicious eDiscovery Exfil were in the top 10, however those did not make the top 10 for medium and small companies. Medium and small companies included Office 365 Suspicious SharePoint Operation, Office 365 Suspicious eDiscovery Search and Azure AD Suspicious Operation in their top 10, while large companies did not.

> Storage of applications in the cloud facilitates actors being able to replace or insert malicious executables and DLLs into commonly accessed shares.

Analysis: This tells us large companies have a higher prevalence of users accessing applications stored in OneDrive or SharePoint. Storage of applications in the cloud facilitates actors being able to replace or insert malicious executables and DLLs into commonly accessed shares, making the access of these file types more frequent while being less noticeable and a more effective technique for actors.

Detecting the compromise

The fact that Suspicious Office 365 eDiscovery Search did not make large companies' top 10 but Office 365 Suspicious eDiscovery Exfil did, tells us that the users who perform these types of searches at large companies generally perform eDiscovery more consistently, and pull-down data from

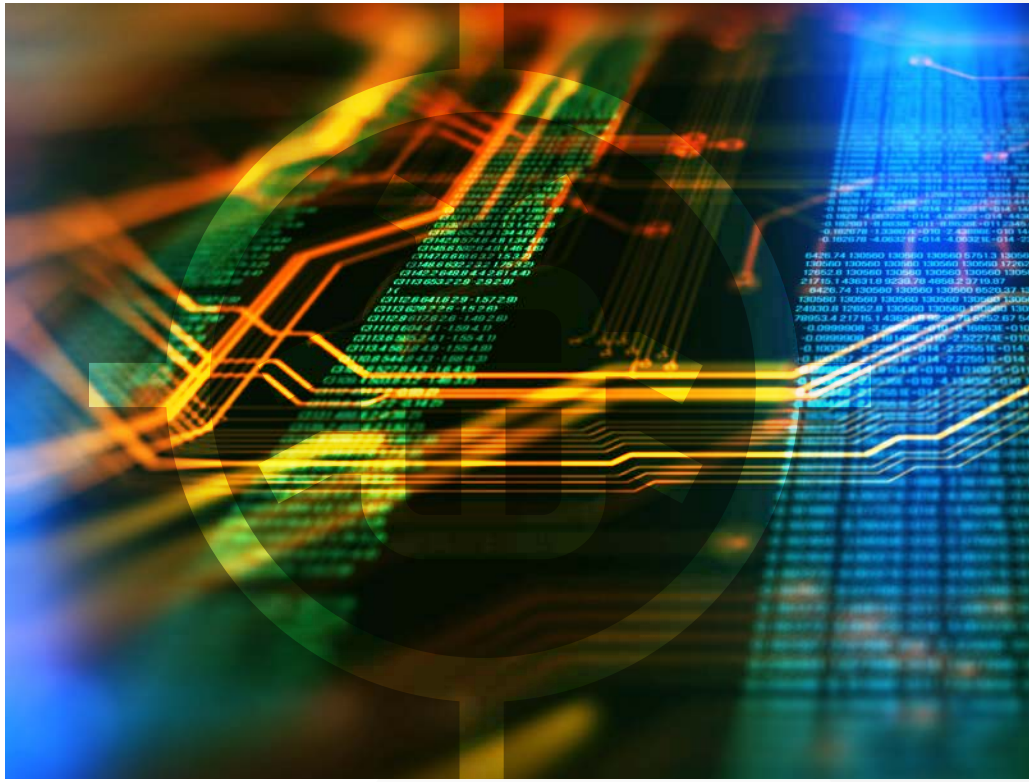

those searches more frequently compared to smaller companies. Office 365 eDiscovery Exfil detection will trigger on every occurrence of an account previewing or exporting data from an eDiscovery search. Access to eDiscovery provides almost complete unfettered access across all components of Office 365 to an actor, allowing them to easily search for and obtain information. Vectra identified one instance of eDiscovery being utilized by an actor to monitor a SOC team's response to the actor's presence in the internal network. Without Vectra, the organization would have been blind to the fact that an Office 365 account was compromised allowing the actor to re-establish a presence after the initial compromise was remediated.

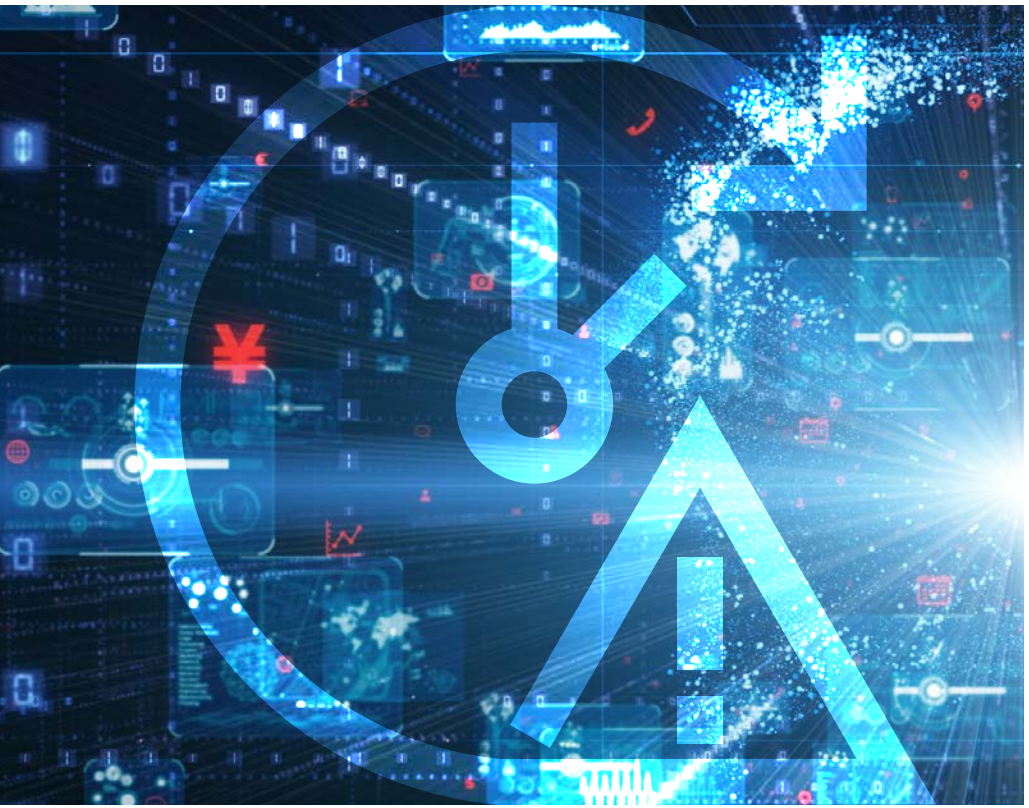

Medium and small companies have Azure AD Suspicious Operation as their top three and top two detections respectively.

Azure AD suspicious operation: A top three threat detection for medium and small companies

Result: Medium and small companies have Azure AD Suspicious Operation as their top three and top two threat detections respectively. Azure AD Suspicious Operation identifies changes to the environment that may be related to an actor taking over an account and elevating privileges of that account.

Analysis: This threat detection did not make the top 10 list for large organizations which tells us that in large organizations, more frequent administration of Azure AD may be required compared to smaller organizations. This means smaller organizations should be vigilant in reviewing alerts to make sure that abnormal behavior does not represent malicious administrative action after an account takeover and privilege escalation.

O365 risky exchange operation topping the list

Result: Office 365 Risky Exchange Operation was at or near the top of the list for each company size segment.

Analysis: This particular threat detection identifies activity that is concerning and could range from the collection and exfiltration of sensitive information, the running of scripts or providing a backdoor. As we saw with a recent major supply chain attack, actors are routinely going after an organization's email to obtain access to sensitive information. The presence of this threat detection denotes that it is abnormal for the account to perform the detected activity. This is an indication that an administrator is performing work of this type infrequently, or an actor is manipulating Exchange in the tenant to obtain access to data or further their attack progression.

Mapping Attack Surfaces to a Supply Chain Attack

Let's explore a recent supply chain attack and detail how reported common actions by the actors in the Azure AD environment would map to the Vectra-defined threat detections. In this attack example, the actor utilized access to a widely deployed vendor's product to gain access to many organizations' local area networks.

Supply chain attacks have shown significant effort and skill in evading preventive controls that involve network sandboxes, endpoint and multifactor authentication (MFA). Including:

- Extensive checks to ensure that it was not in a sandbox or other malware analysis environment.
- Use of code signing and legitimate processes to evade common endpoint controls.
- Novel in-memory dropper to evade file-based analysis in distributing the C2 beacon.
- MFA bypass using stolen security assertion markup language (SAML) session signing keys.

The level of skill and focus required to cleanly bypass endpoint controls is a tribute to recent advances in endpoint detection and response (EDR). However, it is also a reminder that a determined and sophisticated adversary will always be able to bypass prevention and endpoint controls.

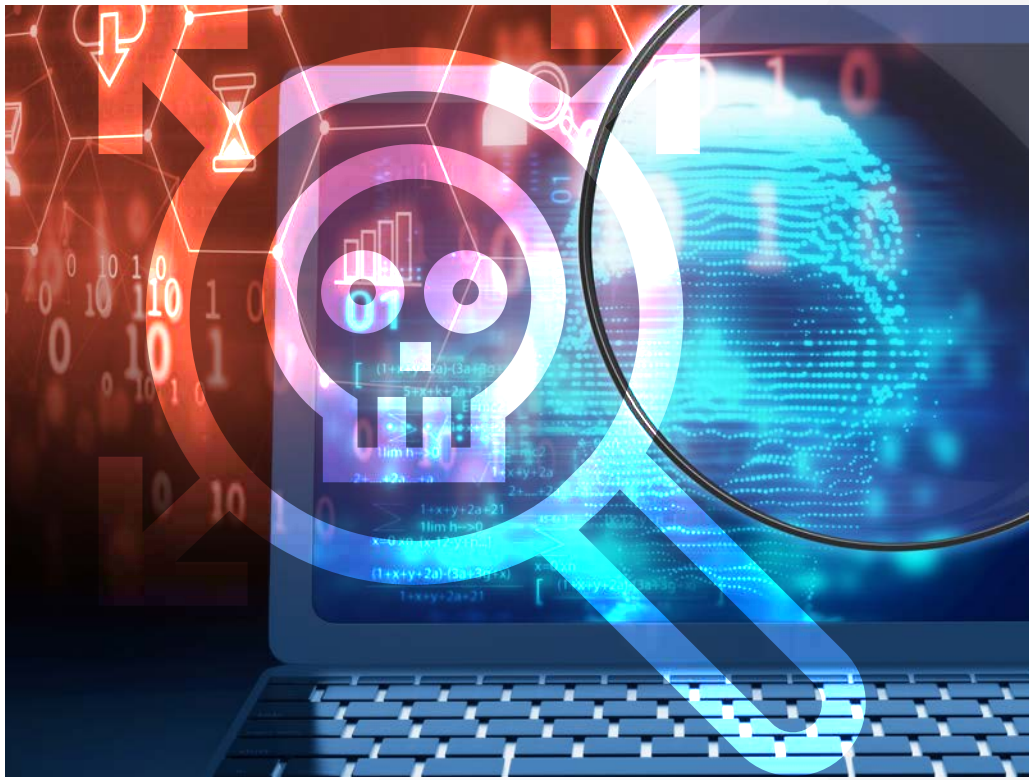

Supply Chain Attacks have shown significant effort and skill in evading preventive controls that involve network sandboxes, endpoint and multifactor authentication (MFA).

For additional information about Supply Chain Attacks, this video details a recent compromise. Watch Now!

Let's find out how

This particular supply chain attack highlights an attack surface that organizations are not well equipped to defend.

Supply Chain Attack / SolarFlare's Path from Network to Cloud

1

Leverage Solarwinds backdoor (Sunburst)

- Activate initial C2 with DNS response
- Download and establish full C2 (Cobalt Strike)

2

Gain Domain Admin

- Network recon to find domain credentials
- Credential based lateral movement from SolarWinds

3

Steal SAML Signing Certificate

- Leverage domain admin credentials to get to ADFS

4

Access Azure AD & Microsoft Office 365

- Admin access to Azure AD with forged ticket
- Add persistence with changes to configurations and permissions
- Access email with native tools

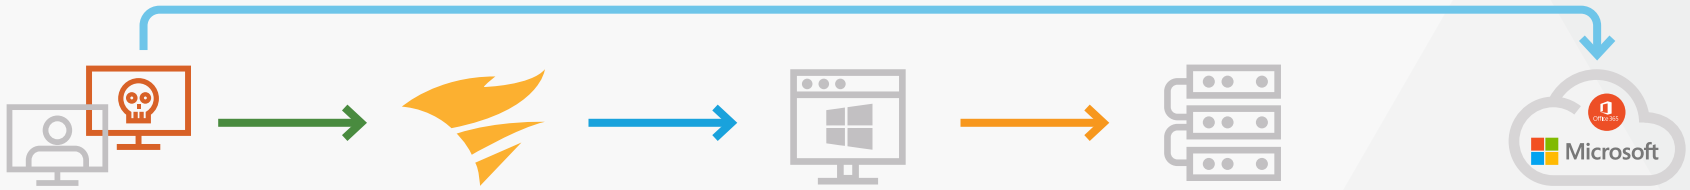

Analysis of step 3: Forging tokens to gain access to Microsoft Azure AD and Office 365

In many instances the actor utilized this access to infiltrate the organization's cloud tenant with a primary goal of accessing business email and documents. The movement from on-premise to cloud is in the opposite direction we typically expect an actor to move based on historical observation. Similar to an attacker operating in the local area network, the actor performed multiple actions that would be considered to fall within one of the following categories: Command and Control, Lateral Movement or Exfiltration.

One of the first motions this actor took was to compromise or modify the authentication infrastructure. This allowed the actor to forge SAML tokens and gain access to the Azure AD and Office 365 environment, bypassing MFA and other security posture validations. The utilization of the forged tokens to sign in would trigger a Vectra **Azure AD Suspicious Sign-on Detection**, which is in the C2 category. This threat detection alerts an organization when there is an anomalous login compared to what is known to be normal for that account. In our analysis, this detection was not in the top 10 overall.

This indicates that abnormal access to the environment is not observed frequently with most customers and should be scrutinized even more when this detection is triggered. Lastly, this highlights that visibility isn't just a question of having data without a way of also inferring context. You have to understand behaviors that are both suspicious and desirable for adversaries, which requires both an understanding of what is expected, and the threat model of an unfolding attack.

Alternative method to step 3: Modifying trust

Another action that the actor may have taken related to the authentication infrastructure was to modify the trust relationship. Modifying the trust relationship would require compromising an account with the required rights or adding those rights to an account already compromised. In the latter scenario, adding the required administrative rights to an account would trigger the **Azure AD Redundant Access Creation Detection**, which is also in the C2 category. Modification of the trust relationship configuration with the privileged account would trigger the **Azure AD Suspicious Operation Detection** in the Lateral Movement category. Our analysis shows that this threat detection is in the top 10 list for each customer based on size.

> The actor's reported actions indicate that a common goal was to obtain access to information contained in email and documents.

Analysis of step 4: Gaining access to information contained in emails

The actor's reported actions indicate that a common goal was to obtain access to information contained in email and documents. Once access was established, the actor would need to perform additional actions that can be considered lateral movement. The actor is reported to have commonly added credentials of a compromised account to existing applications or service principles. This action would trigger the **Azure AD Suspicious Detection** previously discussed.

In some cases, the actor would add new applications or service principals, which would then add permissions to the desired applications. This latter case would trigger the **Azure AD Suspicious OAuth Application Detection**. These last activities highlight two additional attack surfaces that should be on customers' radar to monitor closely so they have clear vision for both modifications that are expected and authorized—and for those that are not.

The Difference: Truly Knowing "Your" Account Behavior

All too often the above scenario has become a reality for companies that fail to recognize the difference between attacker behavior and privileged account usage. That's because the distinction between the two isn't obvious and a large reason why 30% of organizations suffer account takeovers every month. As we discussed, this is also a good reminder that determined adversaries will find a way to bypass prevention and endpoint controls. It all goes back to having the required vision and visibility—knowing what authorized use looks like and understanding the behaviors adversaries are willing to take. The good news is—this doesn't have to be a difficult task to accomplish.

Meaningful AI can help close the gap in your Azure AD and Office 365 accounts, so you have the right data to detect when something is out of the ordinary. Just think back to the top 10 threat detections discussed and what they told us—are you aware when a suspicious attachment gets downloaded or forwarded, or perhaps when a risky or suspicious operation occurs in your cloud? If not, it probably wouldn't hurt to know.

To detect what's out of the ordinary in your Azure AD or Office 365 environment, see more about Cognito Detect for Office 365.

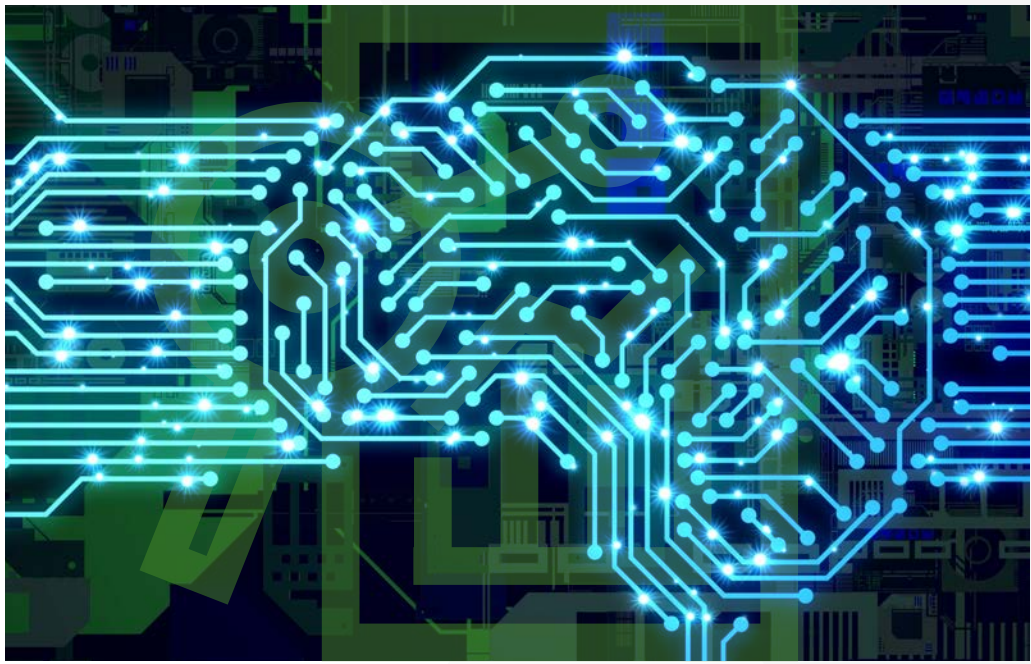

Meaningful AI can help close the gap in your Azure AD and Office 365 accounts, so you have the right data to detect when something is out of the ordinary.

Email info@vectra.ai vectra.ai