

DEV-0537 criminal actor targeting organizations for data exfiltration and destruction

microsoft.com/security/blog/2022/03/22/dev-0537-criminal-actor-targeting-organizations-for-data-exfiltration-and-destruction/

March 22, 2022



March 24, 2022 update – As Microsoft continues to track DEV-0537’s activities, tactics, and tools, we’re sharing new detection, hunting, and mitigation information to give you additional insights on remaining vigilant against these attacks.

In recent weeks, Microsoft Security teams have been actively tracking a large-scale social engineering and extortion campaign against multiple organizations with some seeing evidence of destructive elements. As this campaign has accelerated, our teams have been focused on detection, customer notifications, threat intelligence briefings, and sharing with our industry collaboration partners to understand the actor’s tactics and targets. Over time, we have improved our ability to track this actor and helped customers minimize the impact of active intrusions and in some cases worked with impacted organizations to stop attacks prior to data theft or destructive actions. Microsoft is committed to providing visibility into the malicious activity we’ve observed and sharing insights and knowledge of actor tactics that

might be useful for other organizations to protect themselves. While our investigation into the most recent attacks is still in progress, we will continue to update this blog when we have more to share.

The activity we have observed has been attributed to a threat group that Microsoft tracks as DEV-0537, also known as LAPSUS\$. DEV-0537 is known for using a pure extortion and destruction model without deploying ransomware payloads. DEV-0537 started targeting organizations in the United Kingdom and South America but expanded to global targets, including organizations in government, technology, telecom, media, retail, and healthcare sectors. DEV-0537 is also known to take over individual user accounts at cryptocurrency exchanges to drain cryptocurrency holdings.

Unlike most activity groups that stay under the radar, DEV-0537 doesn't seem to cover its tracks. They go as far as announcing their attacks on social media or advertising their intent to buy credentials from employees of target organizations. DEV-0537 also uses several tactics that are less frequently used by other threat actors tracked by Microsoft. Their tactics include phone-based social engineering; SIM-swapping to facilitate account takeover; accessing personal email accounts of employees at target organizations; paying employees, suppliers, or business partners of target organizations for access to credentials and multifactor authentication (MFA) approval; and intruding in the ongoing crisis-communication calls of their targets.

The social engineering and identity-centric tactics leveraged by DEV-0537 require detection and response processes that are similar to insider risk programs—but also involve short response timeframes needed to deal with malicious external threats. In this blog, we compile the tactics, techniques, and procedures (TTPs) we've observed across multiple attacks and compromises. We also provide baseline risk mitigation strategies and recommendations to help organizations harden their organization's security against this unique blend of tradecraft.

Analysis

The actors behind DEV-0537 focused their social engineering efforts to gather knowledge about their target's business operations. Such information includes intimate knowledge about employees, team structures, help desks, crisis response workflows, and supply chain relationships. Examples of these social engineering tactics include spamming a target user with multifactor authentication (MFA) prompts and calling the organization's help desk to reset a target's credentials.

Microsoft Threat Intelligence Center (MSTIC) assesses that the objective of DEV-0537 is to gain elevated access through stolen credentials that enable data theft and destructive attacks against a targeted organization, often resulting in extortion. Tactics and objectives indicate this is a cybercriminal actor motivated by theft and destruction.

While this actor's TTPs and infrastructure are constantly changing and evolving, the following sections provide additional details on the very diverse set of TTPs we have observed that DEV-0537 is using.

Initial access

DEV-0537 uses a variety of methods that are typically focused on compromising user identities to gain initial access to an organization including:

- Deploying the malicious Redline password stealer to obtain passwords and session tokens
- Purchasing credentials and session tokens from criminal underground forums
- Paying employees at targeted organizations (or suppliers/business partners) for access to credentials and MFA approval
- Searching public code repositories for exposed credentials

Using the compromised credentials and/or session tokens, DEV-0537 accesses internet-facing systems and applications. These systems most commonly include virtual private network (VPN), remote desktop protocol (RDP), virtual desktop infrastructure (VDI) including Citrix, or identity providers (including Azure Active Directory, Okta). For organizations using MFA security, DEV-0537 used two main techniques to satisfy MFA requirements—session token replay and using stolen passwords to trigger simple-approval MFA prompts hoping that the legitimate user of the compromised account eventually consents to the prompts and grants the necessary approval.

In some cases, DEV-0537 first targeted and compromised an individual's personal or private (non-work-related) accounts giving them access to then look for additional credentials that could be used to gain access to corporate systems. Given that employees typically use these personal accounts or mobile phone numbers as their second-factor authentication or password recovery, the group would often use this access to reset passwords and complete account recovery actions.

Microsoft also found instances where the group successfully gained access to target organizations through recruited employees (or employees of their suppliers or business partners). DEV-0537 advertised that they wanted to buy credentials for their targets to entice employees or contractors to take part in its operation. For a fee, the willing accomplice must provide their credentials and approve the MFA prompt or have the user install AnyDesk or other remote management software on a corporate workstation allowing the actor to take control of an authenticated system. Such a tactic was just one of the ways DEV-0537 took advantage of the security access and business relationships their target organizations have with their service providers and supply chains.

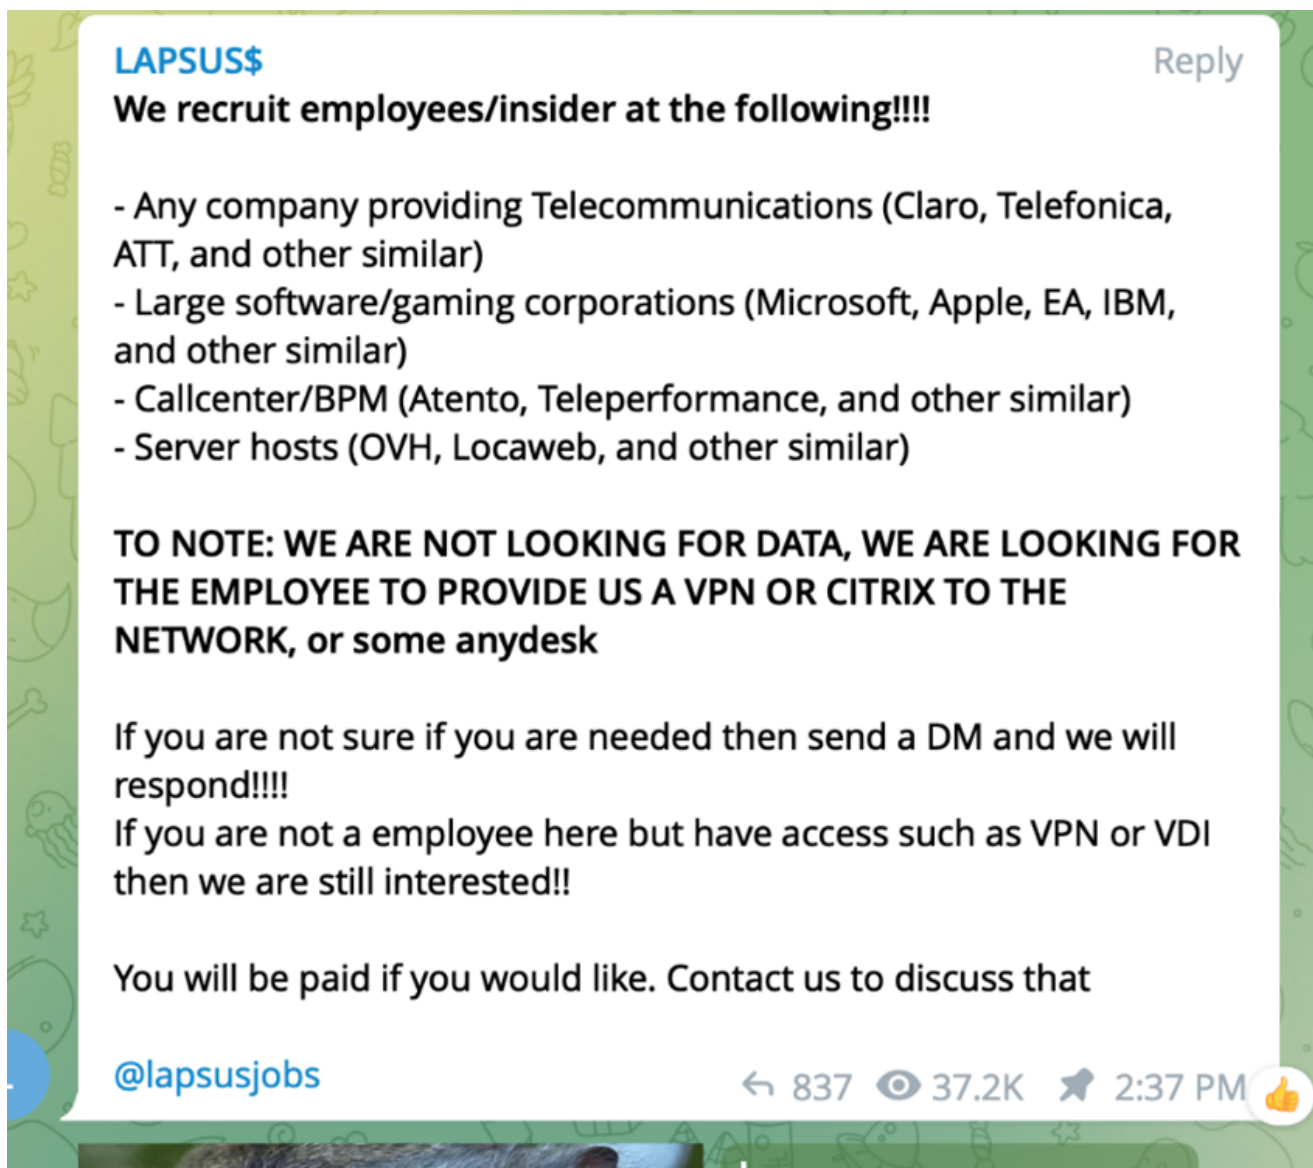


Figure 1. Screenshot of an ad recruiting employees to give out access to their employer's network

In other observed activity, DEV-0537 actors performed a SIM-swapping attack to access a user's phone number before signing into the corporate network. This method allows the actors to handle phone-based authentication prompts they need to gain access to a target.

Once standard user credentials or access was obtained, DEV-0537 typically connected a system to an organization's VPN. In some cases, to meet conditional access requirements, DEV-0537 registered or joined the system to the organization's Azure Active Directory (Azure AD).

Reconnaissance and privilege escalation

Once DEV-0537 obtained access to the target network using the compromised account, they used multiple tactics to discover additional credentials or intrusion points to extend their access including:

- Exploiting unpatched vulnerabilities on internally accessible servers including JIRA, Gitlab, and Confluence
- Searching code repositories and collaboration platforms for exposed credentials and secrets

They have been consistently observed to use AD Explorer, a publicly available tool, to enumerate all users and groups in the said network. This allows them to understand which accounts might have higher privileges. They then proceeded to search collaboration platforms like SharePoint or Confluence, issue-tracking solutions like JIRA, code repositories like GitLab and GitHub, and organization collaboration channels like Teams or Slack to discover further high-privilege account credentials to access other sensitive information.

DEV-0537 is also known to exploit vulnerabilities in Confluence, JIRA, and GitLab for privilege escalation. The group compromised the servers running these applications to get the credentials of a privileged account or run in the context of the said account and dump credentials from there. The group used DCSync attacks and Mimikatz to perform privilege escalation routines. Once domain administrator access or its equivalent has been obtained, the group used the built-in *ntdsutil* utility to extract the AD database.

In some cases, DEV-0537 even called the organization's help desk and attempted to convince the support personnel to reset a privileged account's credentials. The group used the previously gathered information (for example, profile pictures) and had a native-English-sounding caller speak with the help desk personnel to enhance their social engineering lure. Observed actions have included DEV-0537 answering common recovery prompts such as "first street you lived on" or "mother's maiden name" to convince help desk personnel of authenticity. Since many organizations outsource their help desk support, this tactic attempts to exploit those supply chain relationships, especially where organizations give their help desk personnel the ability to elevate privileges.

Exfiltration, destruction, and extortion

Based on our observation, DEV-0537 has dedicated infrastructure they operate in known virtual private server (VPS) providers and leverage NordVPN for its egress points. DEV-0537 is aware of detections such as impossible travel and thus picked VPN egress points that were geographically like their targets. DEV-0537 then downloaded sensitive data from the targeted organization for future extortion or public release to the system joined to the organization's VPN and/or Azure AD-joined system.

DEV-0537 has been observed leveraging access to cloud assets to create new virtual machines within the target's cloud environment, which they use as actor-controlled infrastructure to perform further attacks across the target organization.

If they successfully gain privileged access to an organization's cloud tenant (either AWS or Azure), DEV-0537 creates global admin accounts in the organization's cloud instances, sets an Office 365 tenant level mail transport rule to send all mail in and out of the organization to the newly created account, and then removes all other global admin accounts, so only the actor has sole control of the cloud resources, effectively locking the organization out of all access. After exfiltration, DEV-0537 often deletes the target's systems and resources. We've observed deletion of resources both on-premises (for example, VMware vSphere/ESXi) and in the cloud to trigger the organization's incident and crisis response process.

The actor has been observed then joining the organization's crisis communication calls and internal discussion boards (Slack, Teams, conference calls, and others) to understand the incident response workflow and their corresponding response. It is assessed this provides DEV-0537 insight into the victim's state of mind, their knowledge of the intrusion, and a venue to initiate extortion demands. Notably, DEV-0537 has been observed joining incident response bridges within targeted organizations responding to destructive actions. In some cases, DEV-0537 has extorted victims to prevent the release of stolen data, and in others, no extortion attempt was made and DEV-0537 publicly leaked the data they stole.

Impact

Early observed attacks by DEV-0537 targeted cryptocurrency accounts resulting in compromise and theft of wallets and funds. As they expanded their attacks, the actors began targeting telecommunication, higher education, and government organizations in South America. More recent campaigns have expanded to include organizations globally spanning a variety of sectors. Based on observed activity, this group understands the interconnected nature of identities and trust relationships in modern technology ecosystems and targets telecommunications, technology, IT services and support companies—to leverage their access from one organization to access the partner or supplier organizations. They have also been observed targeting government entities, manufacturing, higher education, energy, retailers, and healthcare.

Microsoft will continue to monitor DEV-0537 activity and implement protections for our customers. The current detections and advanced detections in place across our security products are detailed in the following sections.

Actor actions targeting Microsoft

This week, the actor made public claims that they had gained access to Microsoft and exfiltrated portions of source code. No customer code or data was involved in the observed activities. Our investigation has found a single account had been compromised, granting limited access. Our cybersecurity response teams quickly engaged to remediate the compromised account and prevent further activity. Microsoft does not rely on the secrecy of code as a security measure and viewing source code does not lead to elevation of risk. The

tactics DEV-0537 used in this intrusion reflect the tactics and techniques discussed in this blog. Our team was already investigating the compromised account based on threat intelligence when the actor publicly disclosed their intrusion. This public disclosure escalated our action allowing our team to intervene and interrupt the actor mid-operation, limiting broader impact.

Recommendations

Strengthen MFA implementation

Multifactor authentication (MFA) is one of the primary lines of defense against DEV-0537. While this group attempts to identify gaps in MFA, it remains a critical pillar in identity security for employees, vendors, and other personnel alike. See the following recommendations to implement MFA more securely:

Do:

- Require MFA for all users coming from all locations including perceived trusted environments, and all internet-facing infrastructure—even those coming from on-premises systems.
- Leverage more secure implementations such as FIDO Tokens, or Microsoft Authenticator with number matching. Avoid telephony-based MFA methods to avoid risks associated with SIM-jacking.
- Use Azure AD Password Protection to ensure that users aren't using easily guessed passwords. Our blog about password spray attacks outlines additional recommendations.
- Leverage passwordless authentication methods such as Windows Hello for Business, Microsoft Authenticator, or FIDO tokens to reduce risks and user experience issues associated with passwords.
- Implement user and sign-in risk-based policies that block high impact user actions like device enrollment and MFA registration.
- Break glass accounts should be stored offline and not be present in any sort of online password vaulting solution.
- Use automated reports and workbooks such as Azure Monitor workbooks for reports for detailed analysis on risk distribution, risk detection trends, and opportunities for risk remediation.
- Remind employees that enterprise or workplace credentials should not be stored in browsers or password vaults secured with personal credentials

Do NOT:

- Use weak MFA factors such as text messages (susceptible to SIM swapping), simple voice approvals, simple push (instead, use number matching), or secondary email addresses.
- Include location-based exclusions. MFA exclusions allow an actor with only one factor for a set of identities to bypass the MFA requirements if they can fully compromise a single identity.
- Allow credential or MFA factor sharing between users.

Require healthy and trusted endpoints

- Require trusted, compliant, and healthy devices for access to resources to prevent data theft.
- Turn on cloud-delivered protection in Microsoft Defender Antivirus to cover rapidly evolving attacker tools and techniques, block new and unknown malware variants, and enhance attack surface reduction rules and tamper protection.

Leverage modern authentication options for VPNs

VPN authentication should leverage modern authentication options such as OAuth or SAML connected to Azure AD to enable risk-based sign-in detection. Modern authentication enables blocking authentication attempts based on sign-in risk, requiring compliant devices for sign in, and tighter integration with your authentication stack to provide more accurate risk detections. Implementation of modern authentication and tight conditional access policies on VPN has been shown to be effective against DEV-0537's access tactics.

Strengthen and monitor your cloud security posture

DEV-0537 leverages legitimate credentials to perform malicious actions against customers. Since these credentials are legitimate, some activity performed might seem consistent with standard user behavior. Use the following recommendations to improve your cloud security posture:

- Review your Conditional Access user and session risk configurations:
 - Block or force password reset for high/medium user risk for all users
 - Block high sign-in risk logins for all users
 - Block medium sign-in risk logins for privileged users
 - Require MFA for medium sign-in risk logins for all other users
- Alerts should be configured to prompt a review on high-risk modification of tenant configuration, including but not limited to:
 - Modification of Azure AD roles and privileged users associated with those roles
 - Creation or modification of Exchange Online transport rules
 - Modification of tenant-wide security configurations

- Review risk detections in Azure AD Identity Protection
 - Risk detections highlight risky users and risky sign-ins
 - Administrators can review and confirm individual sign-ins listed here as compromised or safe
 - Read this article on how to investigate risk using Azure AD Identity Protection

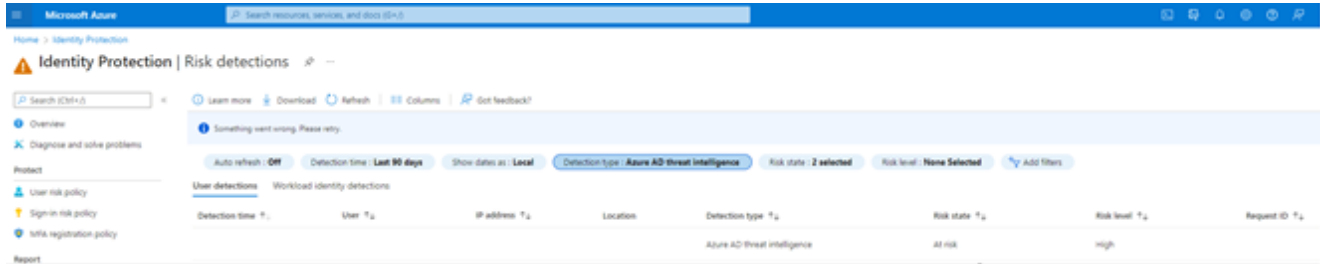


Figure 2. Using Azure AD Identity Protection to review risk detections

Improve awareness of social engineering attacks

Microsoft recommends raising and improving awareness of social engineering tactics to protect your organization. Educate members of your technical team to watch out for and report any unusual contacts with colleagues. IT help desks should be hypervigilant about suspicious users and ensure that they are tracked and reported immediately. We recommend reviewing help desk policies for password resets for highly privileged users and executives to take social engineering into consideration.

Embed a culture of security awareness in your organization by educating employees about help desk verification practices. Encourage them to report suspicious or unusual contacts from the help desk. Education is the number one defense against social engineering attacks such as this one and it is important to make sure that all employees are aware of the risks and known tactics.

Establish operational security processes in response to DEV-0537 intrusions

DEV-0537 is known to monitor and intrude in incident response communications. As such, these communication channels should be closely monitored for unauthorized attendees and verification of attendees should be performed visually or audibly.

We advise organizations to follow very tight operational security practices when responding to an intrusion believed to be DEV-0537. Organizations should develop an out-of-band communication plan for incident responders that is usable for multiple days while an investigation occurs. Documentation of this response plan should be closely held and not easily accessible.

Microsoft continues to track DEV-0537's activities, tactics, malware, and tools. We will communicate any additional insights and recommendations as we investigate their actions against our customers.

Detecting, hunting, and responding to DEV-0537 activities

Microsoft security products provide several detections that can help identify activities resembling DEV-0537 tactics. We're also sharing several Microsoft 365 Defender, Microsoft Defender for Cloud Apps, and Microsoft Sentinel hunting and detection queries that are linked in the following sections. We suggest reviewing the following detections and using the highlighted queries to enhance the investigation of potential activity in your environment.

Initial access

Microsoft Sentinel hunting queries

Sign-in from VPS providers – This query looks for successful sign-ins from known VPS provider network ranges with suspicious token-based sign-in patterns. This is not an exhaustive list of VPS provider ranges but covers some of the most prevalent providers observed.

Investigate unknown sign-in attempts from uncommon or unusual VPS providers.

Sign-in activity from NordVPN providers – This query looks for sign-in activity from NordVPN providers using the feed leveraging NordVPN API and is updated daily.

Investigate unknown sign-in attempts from VPN providers such as NordVPN unless it is commonly seen in your organization.

User sign-in IP address teleportation – This query looks at sign-in logs to identify user accounts that have signed in from two different countries or regions within a specified time window. By default, this is a 10-minute window either side of the previous sign-in.

Investigate the users signing in from multiple locations within a short span of time. It might detect users roaming onto VPNs. You can also exclude known VPN IP address ranges in the query.

Reconnaissance

Microsoft 365 Defender built-in detection: Multiple searches for sensitive data in SharePoint sites

This detection looks for instances where a user searched for sensitive data on SharePoint sites that an attacker can use as internal information to leverage in later attacks if the user's account is compromised.

Investigate the user account performing the queries to determine if it was compromised. Determine what, if any, sensitive information was accessed to assess the impact.

Note: Data used in this detection requires advanced audit to be enabled in Microsoft Defender 365 that includes the *SearchQueryInitiatedSharePoint* event type.

Privilege escalation

Microsoft 365 Defender built-in detection: Risky user created global admin

This detection will alert users based on the risk score proved by Azure AD Identity Protection when a new global admin was created by a user that had a risky sign-in. An attacker might have compromised the user account to perform lateral movement.

Investigate the new global admin account to determine if it was created legitimately and if the user account that performed the action was compromised.

Microsoft 365 Defender hunting queries

Multiple admin role removal operations done by a single user – This query looks for multiple users that had their administrator role removed by a single user within a certain period.

Investigate if the user account that removed the admin roles was compromised or if the actions were legitimate. If determined to be compromised, disable the account and reset the password. Restore access to affected accounts as needed.

'ElevateAccess' operation followed risky sign-in – This query looks for users who had a risky sign-in (based on Azure AD Identity Protection risk score) and then performed an 'ElevateAccess' action. 'ElevateAccess' operations can be used by global admins to obtain permissions over Azure resources.

Investigate the risky sign-ins and the following 'ElevateAccess' operation and disable the account if it was determined to be compromised.

Microsoft Sentinel hunting queries

User-assigned privileged role – This query identifies when a new privileged role is assigned to a user or when any account eligible for a role is given privileged access.

Investigate if the assignment of privileged access is unexpected or does not align to the role of the account holder. See Things to monitor in your security operations for privileged accounts for details.

User added to Azure AD privileged groups (near real-time (NRT) rule) – This query looks for instances when a user is added to any privileged groups.

Investigate any unusual additions to privileged groups, particularly administrator roles. For details, see Azure AD audit activity reference and administrator role permissions in Azure AD.

Multiple admin membership removals from newly created admin – This query detects when newly created global admin removes multiple existing global admins which can be an attempt by adversaries to lock down the organization and retain sole access.

Investigate reasoning and intention of multiple membership removal by new global admins and take necessary actions accordingly.

For Microsoft Sentinel customers who have onboarded Okta logs, the following queries can assist in investigating DEV-0537 activity across those logs:

Microsoft Sentinel + Okta logs hunting queries

Admin privilege granted (Okta) – This query searches for successful grant of administrator permissions to user/groups. Adversaries often attempt to assign administrator permission to users/group to maintain access as well as to elevate privileges.

Verify the behavior is known and filter out any expected activity and triage unknown. See Okta API event types for details.

Create API token (Okta) – This query searches for attempts to create new API token. Okta API tokens are used to authenticate requests to Okta APIs.

Investigate attempts to create new API token creation or authentication attempts. See Okta API event types for details.

Initiate impersonation session (Okta) – This query searches for impersonation events used in LAPSUS\$ activity. *User.session.impersonation* are rare events, normally triggered when an Okta Support person requests admin access for troubleshooting.

Review *user.session.impersonation* events and correlate that with legitimate opened Okta support tickets to determine if these are anomalous. See Okta API event types and Cloudflare's investigation of the January 2022 Okta compromise for details.

Rare MFA operations (Okta) – MFA helps prevent credential compromise. This query searches for rare MFA operations like deactivating, updating, resetting, and attempts to bypass MFA.

Adversaries often attempt these operations to compromise networks and high-value accounts.

Verify that the behavior is known and filter out anything that is expected. See Okta API event types for details.

Persistence

Microsoft 365 Defender hunting queries

Device registration after risky sign-in – This query looks for a new device registration in Azure AD preceded by a medium or high-risk sign-in session for the same user within a maximum of six hours.

Investigate the user account to determine if it is compromised. Disable user account, reset user password, and remove devices registered in Azure AD if compromised.

MFA method added after risky sign-in – This query looks for a new MFA method added to an account that was preceded by a medium or high-risk sign-in session for the same user within a maximum of six hours.

Investigate the user account to determine if it is compromised. If compromised, disable the user account, reset user password, and remove the MFA method added by threat actor.

Exfiltration, destruction, and extortion

Microsoft Defender for Cloud Apps built-in detection: Delete multiple VMs in a single session

This detection profiles your environment and triggers alerts when users delete multiple VMs in a single session, relative to the baseline in your organization. This might indicate an attempted breach.

Investigate the user account performing the deletion operations to determine if it was compromised or if the activities were performed legitimately and not part of a destructive attack.

Microsoft 365 Defender query

Upload multiple code repositories to external cloud domains – This query looks for accounts that uploaded multiple code repositories to external web domain.

Investigate if the accounts are compromised. If compromised, disable the accounts and reset the passwords. Assess the impact of what information was obtained, looking for any passwords, secrets, certificates, and others that the attacker might be able to leverage.

Note: This query uses 'FileUploadedToCloud' event which is only available for customers that enabled Microsoft Defender for Endpoint integration with Microsoft Defender for Cloud Apps. See [Integrate Microsoft Defender for Endpoint with Defender for Cloud Apps](#) for details)

Microsoft Sentinel hunting queries

Mass cloud resource deletions time series anomalies – This query generates baseline pattern of cloud resource deletions by a user and alert on an anomaly when any unusual spike is detected.

Investigate the anomalies from unusual or privileged users, they could be indication of a cloud infrastructure takedown by an adversary.

Mail redirect via ExO transport rule – This query identifies when Exchange Online transport rule configured to forward emails.

Investigate detections to determine if a malicious actor has configured a new mailbox to collect mail from multiple user accounts.

Time series anomaly for data size transferred to public internet – This query identifies anomalous or unusual data transfers to public networks. This detection identifies large deviations from a baseline pattern based on detection algorithms from the Sentinel-integrated Kusto Query Language (KQL) anomaly detection. The higher the score, the further it is from the baseline value. The output is aggregated to provide a summary view of unique source IP to destination IP address and port bytes sent traffic observed in the flagged anomaly hour. The source IP addresses which were sending less than *bytessentperhourthreshold* have been excluded, the value of which can be adjusted as needed. You might have to run queries for individual source IP addresses from *SourceIPlist* to determine if anything looks suspicious. Investigate any sudden increase in data transferred to unknown public networks as an indication of data exfiltration attempts.